Schipper, Burkhard C.

Working Paper

Imitators and optimizers in Cournot oligopoly

Working Paper, No. 05-37

Provided in Cooperation with:
University of California Davis, Department of Economics

Suggested Citation: Schipper, Burkhard C. (2005) : Imitators and optimizers in Cournot oligopoly, Working Paper, No. 05-37, University of California, Department of Economics, Davis, CA

This Version is available at:
https://hdl.handle.net/10419/31362

Imitators and optimizers in Cournot oligopoly

Burkhard C. Schipper
University of California, Davis

December 14, 2005

Paper # 05-37

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 la Vega-Redondo (1997). Optimizers play a myopic best response to the opponents' previous output. Firms are allowed to make mistakes and deviate from their 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).

Imitators and Optimizers in Cournot Oligopoly*

Burkhard C. Schipper†
Department of Economics
University of California, Davis

December 14, 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 their 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).

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

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

---

*I thank among others Carlos Alós-Ferrer, Itzhak Gilboa, Jörg Oechssler, and Yossi Spiegel as well as audiences of various seminars and conferences for very helpful comments. Financial support from the Deutsche Forschungsgemeinschaft through SFB/TR 15 and the DAAD is gratefully acknowledged.

†University of California, Davis, Department of Economics, One Shields Avenue, Davis, CA 95616, USA, Email: bcschipper@ucdavis.edu

"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 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

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

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 response. 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 the market index or similar means of relative comparison.

In the proofs of our results, we rely on two main concepts, quasisubmodularity of payoff functions and stochastic stability analysis. Quasisubmodularity (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 quasisubmodularity 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

[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 Professor Selten for pointing me to the "majorant game".

a lower total market quantity. The Cournot oligopoly satisfies this property by definition (see Lemma 1). A similar version of this property, in which the total market quantity is replaced by the opponents' quantity, 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 the class of aggregative quasisubmodular 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 of 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 (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.

The paper is organized as follows: Section 2 introduces the model and the decision rules. It is followed in section 3 by a 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 a function of the individual quantity and the total quantity,

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

We restrict our analysis to a symmetric oligopoly since imitation is more reasonable if firms face similar conditions of production.[3]

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) it 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 $p(Q') < p(Q)$.*

[3]Our model could be extended to asymmetric settings, in which each firm may imitate opponents that are most similar to her in terms of the payoff function.

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

If Assumption 1 is modified such that p is weakly decreasing, then π_i is submodular in (q_i, Q) on $\Gamma \times \{0, \delta, 2\delta, ..., n\nu\delta\}$.

Remark 1 (Quasisubmodularity) *Strict submodularity of π_i implies that π_i is strictly quasisubmodular in (q_i, Q) on $\Gamma \times \{0, \delta, 2\delta, ..., n\nu\delta\}$ (but not vice versa), i.e., for all $q_i'' > q_i'$, $Q'' > Q'$,*

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

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

The intuition for quasisubmodularity is that if a firm with a higher quantity has a higher profit than a firm with a lower quantity for a given market quantity, then the first firm has also a higher profit for a lower market quantity. That is, in some cases the property allows us to compare profits of firms operating with different quantities in the same market. The observation that the payoff function is quasisubmodular 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 by

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

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*

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

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

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

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

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

Definition 3 means that an optimizer sets an output level that is a best response to the opponents' total output in the previous period. In the last section we discuss how our results generalize to more sophisticated optimizers.

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

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.

At every output revision opportunity t, each firm follows her decision rule with probability $(1 - \varepsilon)$, $\varepsilon \in (0, a]$, $a < 1$, and with probability ε she randomizes with full support Γ. As a matter of convention, we call a firm *mutating* at t if she 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. $P(\varepsilon)$ 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) for all $\omega \in A$ and for all $\omega' \notin A$, $p_{\omega\omega'} = 0$; and

(ii) for all $\omega, \omega' \in A$, there exists a finite $m \in \mathbb{N}$ s.t. $p^{(m)}_{\omega\omega'} > 0$, $p^{(m)}_{\omega\omega'}$ 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* $i \in N$,

$$q_i^* p(Q^*) - c(q_i^*) \geq q_i p(Q^*) - c(q_i), \text{ for all } q_i \in \Gamma, \text{ with } Q^* := \sum_{i \in N} q_i^*. \tag{9}$$

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 response. 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* $i \in N$,

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

It is known that in Cournot oligopoly with a homogeneous population of optimizers, the sequential best response process converges under certain conditions to Cournot Nash equilibrium. 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 set their best response quantities anyway. Thus the symmetric Cournot Nash equilibrium is an absorbing state. However, is it the unique absorbing state? Consider the following state:[4]

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.*

$$q^S p(\theta n q^S + (1-\theta) n q^D) - c(q^S) > q p(\theta n q^S + (1-\theta) n q^D) - c(q), \text{ for all } q \neq q^S, \quad (11)$$

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

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

In the Pseudo-Stackelberg solution all imitators set identical outputs. This output maximizes profits of imitators given that they do not perceive any influence on the price and the optimizers set identical best responses. Clearly, this outcome has features of the 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 all $q \in \Gamma$, $q \neq q^S$, hence also for $q^D \neq q^S$. I.e., it follows that if $q^D \neq q^S$ then[5]

$$\pi_i(q^S, q^D, n, \theta) > \pi_j(q^S, q^D, n, \theta), \text{ for all } i \in I \text{ and for all } j \in N \backslash I.$$

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

[5]For 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 $i \in N \backslash I$.

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. Consider first the imitators: all imitators set identical outputs and each of them is strictly better off than is any optimizer. Hence an imitator has no reason to deviate from her output. Optimizers do not deviate too from their output since each of them plays the best response. Thus the Pseudo-Stackelberg solution is an absorbing state.

Existence of Pseudo-Stackelberg solution is analogous to existence of competitive solution in Vega-Redondo (1997) and the existence of Cournot Nash equilibrium. Standard assumptions on costs, i.e., strictly increasing marginal costs and small fixed costs, and Assumption 3 in the next section suffice. By the strict Inequality (11), quasisubmodularity, and Assumption 3 in the next section, the Pseudo-Stackelberg solution must be unique if it exists (see Lemma 2 (vi) in the appendix).

Example 1 (Pseudo-Stackelberg Solution) *Suppose $n = 4$ and $\theta = \frac{1}{2}$ such that we have two imitators and two optimizer. Inverse demand is given by $p(q_1+q_2+q_3+q_4) = 10-q_1-q_2-q_3-q_4$, where q_1, q_2 are the outputs of the imitators and q_3, q_4 are the outputs of the optimizers. Costs are strictly convex and identically for all firms $c(q_i) = \frac{1}{2}q_i^2, i \in \{1,2,3,4\}$. Each optimizer maximizes individually profits according to the Cournot conjecture as follows (since the setting is symmetric q^I denotes the identical output of either imitator): $(10-2q^I-q_3-q_4)q_3-\frac{1}{2}q_3^2 \longrightarrow \max_{q_3}$, which leads to $b_3(2q^I+q_4) = \frac{10-2q^I-q_4}{3}$. Substituting optimizer 4's best response into optimizer 3's best response, we get an optimizer's symmetric equilibrium quantity given the any imitators' symmetric quantity, $q^D = \frac{5-q^I}{2}$. Each imitator maximizes profits taking the price $\bar{p}$ fixed (as indicated by $\bar{p}$), i.e., for $i = 1,2$, $\bar{p}q_i - \frac{1}{2}q_i^2 \longrightarrow \max_{q_i}$, which leads to price equals marginal cost, $\bar{p} = q_i =: q^I$. Substituting inverse demand for $\bar{p}$ taking into account the optimizers' equilibrium quantities leads to $q^S = \frac{5}{2}$. It follows that $q^D = \frac{5}{4}$, $\pi_i = \frac{100}{32}$ for $i = 1,2$ and $\pi_i = \frac{75}{32}$ for $i = 3,4$. Note that each imitator sets a higher output than each optimizer. Moreover, each imitator is strictly better off than each optimizer.*

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* $i \in N \backslash I$, $q^D := b(\theta n q^I + ((1-\theta)n - 1)q^D)$,

(iii) $\pi_i(q^I, q^D, n, \theta) \geq \pi_j(q^I, q^D, n, \theta)$, *for all* $i \in I$ *and all* $j \in N \setminus I$,

(iv) $\pi_i(q^I, q^D, n, \theta) = \pi_j(q^I, q^D, n, \theta)$ *for all* $i \in I$ *and all* $j \in N \setminus I$, *iff* $q^I = q^D$.

Each Pseudo-Stackelberg state is an absorbing state. 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 response 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 $\pi_i(q^D, q^I, n, \theta) = \pi_j(q^D, q^I, n, \theta)$ for all $i \in I$ and all $j \in N \setminus I$. To see the purpose of the other direction note that if $\pi_i(q^D, q^I, n, \theta) = \pi_j(q^D, q^I, n, \theta)$ for some $i \in I$ and some $j \in N \setminus I$, and $q^I \neq q^D$ then some 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 the competitive solution is not a Pseudo-Stackelberg state since optimizers do not set a best response in the competitive solution (unless $n \to \infty$ or $\theta = 1$). Finally, if marginal costs are strictly increasing then the Pseudo-Stackelberg solution is a Pseudo-Stackelberg state since $q^S > q^D$ are such that $\pi_i(q^S, q^D, \theta, n) > \pi_j(q^S, q^D, \theta, n)$ for all $i \in I$ and all $j \in N \setminus I$. Thus properties (i) to (iv) of Definition 7 of Pseudo-Stackelberg states are satisfied.

[6]Again, we assume here that the best response is unique. The uniqueness condition later in Assumption 3 ensures that the best response to the opponents' output is indeed a singleton (see Lemma 2 in the appendix).

Can there be other absorbing sets like cycles? Following example shows that there are cases where we can answer this question in the affirmative.

Example 2 (Cycle) *Consider three firms, $n = 3$, two imitators and one optimizer, $\theta = \frac{2}{3}$, linear demand $p(Q) = 10 - Q$, and fixed cost $c(q_i) = \frac{85}{2}$, and let $\Gamma = \{0, 5\}$. It is straight forward to compute that the monopoly output, $q^M = b(0) = 5$. Moreover, $b(2 \cdot 5) = 0$. The profit from setting the monopoly output when one other firm sets the monopoly output too is strictly negative, $\pi(q^M, 2q^M) = -\frac{5}{2}$. The profit of a monopolist is strictly positive, $\pi(q^M, q^M) = \frac{5}{2}$. Consider the state in which imitators play zero output and the optimizer the monopoly output, $\omega = (0, 0, 5)$. Clearly, while the optimizer plays the best response, the imitators are worse off. Hence, ω is not a Pseudo-Stackelberg state. Imitators have an incentive to adjust. In particular, they will mimic the optimizer and subsequently the optimizer will adjust to that. Thus we will reach a state $\omega' = (5, 5, 0)$. Again, the optimizer plays the best response while the imitators are worse off. It is not a Pseudo-Stackelberg state and imitators have an incentive to adjust quantities by mimicking the optimizer. After the optimizer adjusts, we reach again ω. Note that inertia does not destroy such cycle. That is, it is still true that from ω we reach after a couple of periods ω' and vice versa. In fact, this game has no Pseudo-Stackelberg state, Pseudo-Stackelberg solution nor symmetric pure strategy Cournot Nash equilibrium.*

Example 2 appears to be contrived by the restricted action space. But even with a finer grid there are less trivial cycles. In order to rule out any cycles, we impose following assumption:

Assumption 2 *Let q^S denote the imitators output in the Pseudo-Stackelberg solution and let $q^{(1-\theta)n}$ be the unique symmetric $(1-\theta)n$-firm Cournot Nash equilibrium output, i.e., $q^{(1-\theta)n} = b(((1-\theta)n - 1)q^{(1-\theta)n})$. We assume that $q^S > q^{(1-\theta)n}$.*

This assumption ensures that when imitators mimic identical best responses of optimizers, they will eventually play a Pseudo-Stackelberg state quantity, and thus no cycles can arise. It is a strong assumption, but it is for instance satisfied in the standard Cournot quadropoly with linear demand and convex cost of Example 1.[7] However, in Example 4 we show that this assumption is not necessary for our results.

[7]This would be enough for our aim of showing that there is an interesting class of counter-examples to the proposition that optimizers are better off than non-optimizers.

4 Results

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

Assumption 3 *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 assumption states that the slopes of the best response 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 response 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 response correspondence are strictly larger than -1 is made in order to obtain an unique best response. 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 response correspondence is in fact a function (see Vives, 2000, p. 43). In Lemma 2 we show that by Assumptions 3, total output is increasing in imitators output, and that if the Pseudo-Stackelberg solution exists, it must be unique.

The next assumption is very useful to compare in some cases profits of players with different output facing the same market output. Together with quasisubmodularity it allows us also to compare in some cases profits of players with different output when we change the market output.

Assumption 4 (Quasiconcavity) *Let π_i be quasiconcave in q_i on Γ for all $Q \in \{0, \delta, 2\delta, ..., n\nu\delta\}$, i.e., for any $q_i, q_i', q_i'' \in \Gamma$ with $q_i' \leq q_i \leq q_i''$,*

$$\pi_i(q_i, Q) \geq \min\{\pi_i(q_i', Q), \pi(q_i'', Q)\}, \text{ for all } Q \in \{0, \delta, 2\delta, ..., n\nu\delta\}. \tag{14}$$

In Lemma 3 in the appendix we show, that this assumption together with quasisubmodularity and Assumption 3 implies that for any quantity between the Cournot Nash equilibrium and the Pseudo-Stackelberg solution, there is a Pseudo-Stackelberg state in which imitators play symmetrically this quantity.

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).

Theorem 1 *Let $\varepsilon = 0$ and suppose previous assumptions hold. Then $Z = \{\{\omega\} : \omega \in H\}$.*

In previous section, we have argued already that each Pseudo-Stackelberg state must be an absorbing set. To show that every absorbing set is a Pseudo-Stackelberg state, we use Assumptions 2 and 3 to show that there aren't any cycles and any adjustment process converges to a Pseudo-Stackelberg state. Such adjustment process involves profit comparisons by imitators for which we make use quasisubmodularity and quasiconcavity.

Let S denote the support of the long run distribution φ^*. By standard results, Theorem 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 2 *If $\theta \in (0, 1]$, then under previous assumptions it is never true that $S = \{\omega^\circ\}$.*

This result follows from Theorem 1 above and Lemmata 4 and 5 in the appendix. In particular, we show in Lemma 5 that the Pseudo-Stackelberg solution can be reached from any Pseudo-Stackelberg state by just one suitable mutation followed by the unperturbed adjustment process. This is so because an imitator when switching to the output of the Pseudo-Stackelberg solution

may decrease her own payoff but decreases the payoffs of any other players even more. By Lemma 4, the instability against a single mutation is sufficient to conclude that the Cournot Nash equilibrium can never be 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 *Under previous assumptions, 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 2 does not exclude any absorbing states from the support of the long run distribution. Since in Lemma 5, which is used to prove Theorem 2, the Pseudo-Stackelberg solution figures prominently, one may conjecture that the Pseudo-Stackelberg solution is the unique long run outcome. This would be also analogous to Vega-Redondo's (1997) competitive solution as long run outcome in a homogeneous population of imitators. It turns out that this conjecture is false. An imitator deviating from the Pseudo-Stackelberg solution may become worse off than others and consequently should switch back. However, in the meantime, some optimizers may have already adjusted their best responses. This adjustment may lead temporarily to higher profits of those optimizers. Subsequently, imitators may mimic those optimizers such that the adjustment process may lead to another Pseudo-Stackelberg state. In fact, we prove the following result:

Example 3 *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). To prove the example, we show first that we can find a sequence of single mutations by which we can move step-wise through the set of Pseudo-Stackelberg states, starting from the Pseudo-Stackelberg state with the lowest (resp. largest) output of imitators up (resp. down) to Pseudo-Stackelberg solution (Lemma 6). 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 Pseudo-Stackelberg state with the largest or lowest output of imitators (Lemma 8 and 9). 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 can conclude by Lemma 4 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 Theorem 1 implies that the Cournot Nash equilibrium is the unique absorbing set. Mutations do 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 5 in the appendix implies Vega-Redondo's (1997) result.[8] I.e., the competitive solution is the unique long run outcome.

Following simple example illustrates a process with mutations and adjustments.

Example 4 *Suppose $n = 3$ and $\theta = \frac{2}{3}$ such that we have two imitators and one optimizer. Inverse demand is given by $p(q_0 + q_1 + q_2) = 10 - q_0 - q_1 - q_2$ where q_0 is the output of the optimizer and q_1, q_2 are the outputs of the first and second imitator respectively. Costs are strictly convex and identically for all firms, $c(q_i) = \frac{1}{2}q_i^2, i \in \{0, 1, 2\}$. The Pseudo-Stackelberg solution can be computed as $q^S = \frac{20}{7}$, $q^D = \frac{10}{7}$, $\pi_i = \frac{200}{49}$ for $i = 1, 2$ and $\pi_0 = \frac{150}{49}$. Note that the monopoly output $q^M = \frac{10}{3} > q^S$. Hence, Assumption 2 is violated. Nevertheless, there*

[8] 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

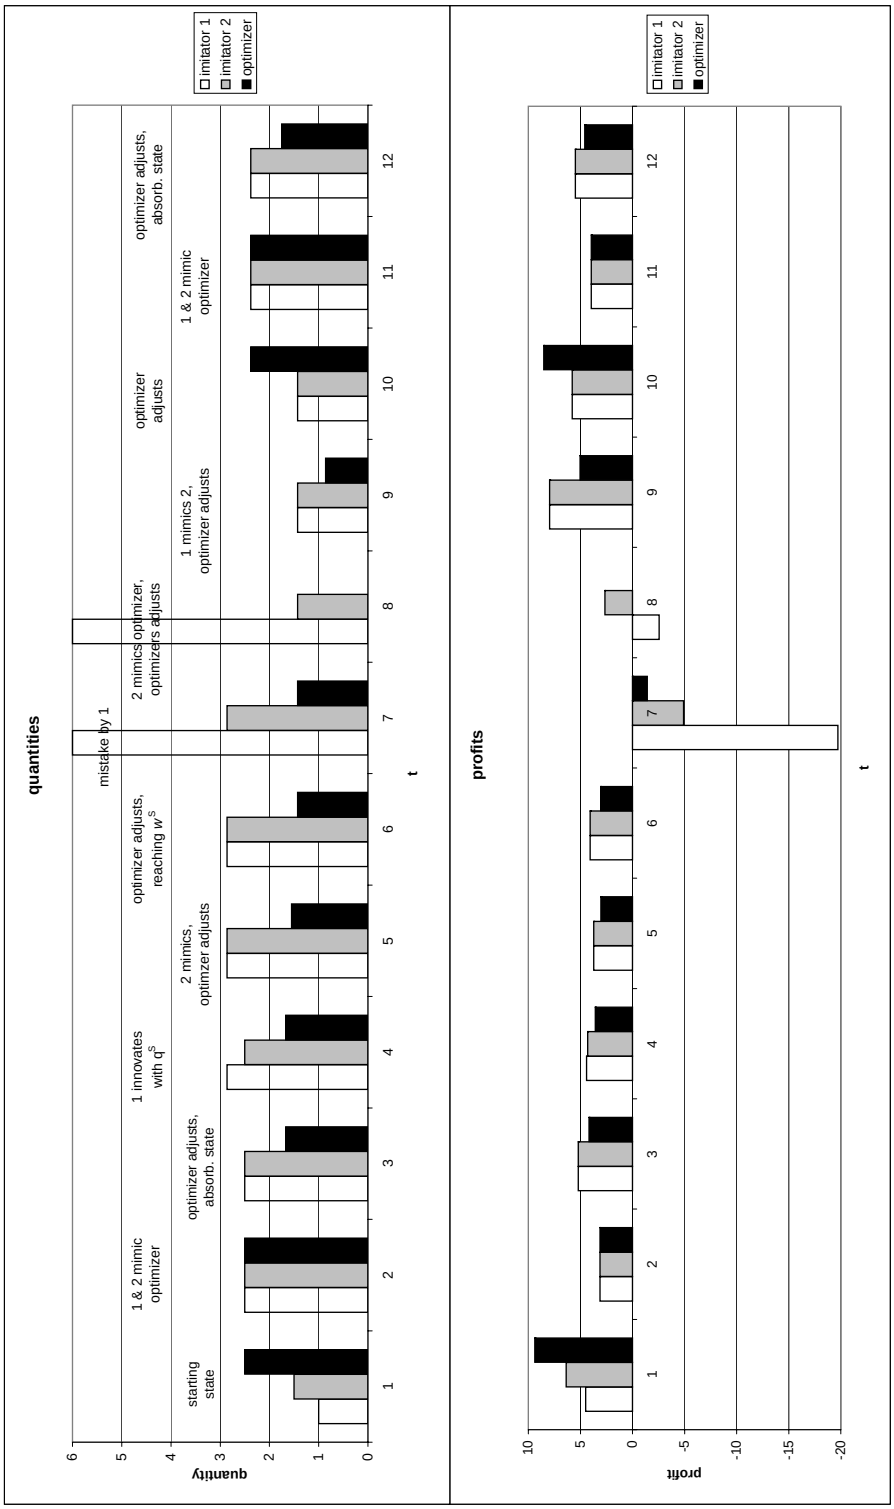

exists no cycle. To see this note that any best response by the optimizer to some output of imitators is below the Pseudo-Stackelberg state with the largest symmetric output of imitators. Moreover, observe that the symmetric quantities of an imitator in any Pseudo-Stackelberg states are bounded above by the monopoly output and below by the Cournot Nash equilibrium output. Suppose that the optimizer is better off then any imitator at some state along the cycle. Then imitators mimic the optimizer. If the optimizer's quantity was between the largest Pseudo-Stackelberg state and the Cournot Nash equilibrium state, then we reach a Pseudo-Stackelberg state after the optimizer adjusts too. If the optimizer's quantity was below the Cournot Nash equilibrium, then the optimizer will subsequently adjust to a output above the Cournot Nash equilibrium output but below the monopoly output. Since this is not a Pseudo-Stackelberg state, imitators must be worse off. Consequently they will mimic the optimizer. Thus they reach a Pseudo-Stackelberg state quantity. After the optimizer adjusts in the next step, we must have reached Pseudo-Stackelberg state. That is, if the process starts outside the Pseudo-Stackelberg state it leads to a Pseudo-Stackelberg state. This example shows that Assumption 2 is not necessary.

Figure 1 below illustrates an example of a 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 Pseudo-Stackelberg 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 Pseudo-Stackelberg state in period 12.

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 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.[9] 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 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.[10] 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 responses against their beliefs leading them to play a smaller output with smaller profits than imitators. Even if they could temporarily "low-ball" the

[9]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 will emerge. This is followed up recently by Thijssen (2005).

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

imitators with some quantities, imitators would erase any profit advantage by mimicking those quantities such that they 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 response 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 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. 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. Hence, long memory imitators may not survive when facing single period memory imitators (as in this article) in Cournot games.

What does it take to make optimizers better off than imitators in our setting? A possible answer to this question is pursued by a follow up paper by Hehenkamp and Kaarbøe (2005). They show in a special example of a two-player game with strategic substitutes that if parameters of the game (like the demand) change as fast as players can make decisions, then it is possible that optimizers are better off than are imitators if the strategic interdependencies are sufficiently weak. It is clear from our analysis that if the game changes less rapidly such that long run outcomes can emerge sufficiently often in between, then imitators will be still better off than optimizers in the long run. Moreover, one can envision an evolutionary stable speed of changing decisions in which sufficiently quick imitators (relative to the speed of change of the environment) would win against optimizer of any speed.

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 2, 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 payoff functions of a Cournot game are quasisubmodular in the individual and the total quantity. This is closely related to strategic substitutes. It allows us in some cases to compare profits of firms with different quantities facing a change of the market quantity. Our results are likely to be generalized to 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[11] and suppose that the imitator demands a share larger than 50% of the pie. What can an optimizer do? She 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.").[12] Our result is likely to be generalized to a class of aggregative quasisubmodular 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 response 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 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. However, Dürsch et al. (2005) show already in a recent experiment on a repeated Cournot duopoly that subjects earn on average higher profits than the opponent when playing against computers programmed to best response, fictitious play, reinforcement learning and trail & error learning but do much worse against computers programmed to imitation.

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

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

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 two lemmata and remark:

Lemma 2 *If Assumptions 1 and 3 hold then we conclude the following:*

(i) For any fixed quantities by imitators, the game in which the set of optimizers is the set of players has a Nash equilibrium.

(ii) Sequential best response converges to such Nash equilibrium in finite time.

(iii) Any optimizer's best response is unique, and in above Nash equilibrium best responses are symmetric.[13]

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

(v) If ω^S exists, then it is unique.

Proof. For any fixed quantities by imitators, the game in which the set of players is the set of optimizers satisfies Assumptions 1 and 3. Hence, (i) and (ii) are implied Dubey, Haimanko, and Zapechelnyuk (2005), and (iii) is implied by Vives (2000, p. 43).

(iv) For $\theta = 1$, the result follows trivially. Consider now $\theta \in (0, 1)$. Suppose to the contrary that $x' < x$ and $x' + (1-\theta)nq^{D'} \geq x + (1-\theta)nq^{D}$. Last inequality is equivalent to $x' - x \geq (1-\theta)n(q^D - q^{D'})$. Since by assumption $0 > x' - x$ we conclude that $0 > (1-\theta)n(q^D - q^{D'})$. Since $\theta \in (0, 1)$ and $n \geq 1$, 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 $q'_{-i} \geq q_{-i}$, then by Assumption 3 (strictly decreasing best responses) $q^{D'} \leq q^D$ (with equality if $q'_{-i} = q_{-i}$), a contradiction to above. Suppose now that we have $q'_{-i} < q_{-i}$.

[13] This implies that the Cournot Nash equilibrium is unique and symmetric.

Then by Assumption 3, $q'_{-i} + q^{D'} < q_{-i} + q^D$, a contradiction to $x' + (1-\theta)nq^{D'} \geq x + (1-\theta)nq^D$ above.

(v) 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 $Q^{S''} > Q^{S'}$ by (iv). By Assumption 1 (Lemma 2 and Remark 1, 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)). □

Recall that by q^I we denote the symmetric output of imitators and by q° a firm output in the symmetric Cournot Nash equilibrium.

Lemma 3 *If Assumptions 1, 3 and 4 hold, then for any $q \in \{q^\circ, q^\circ + \delta, ..., q^S\}$, there exists a state $\omega \in H$ such that any imitator's output $q^I_\omega = q$.*

Proof. By definition, $\omega^\circ, \omega^S \in H$. Consider any q^I_ω s.t. $q^\circ < q^I_\omega < q^S$. We claim that $q^D_\omega < q^\circ$, where $q^D_\omega = b(\theta nq^I_\omega + ((1-\theta)n-1)q^D_\omega)$. Suppose to the contrary that $q^D_\omega \geq q^\circ$. Then $\theta nq^\circ + ((1-\theta)n-1)q^\circ < \theta nq^I_\omega + ((1-\theta)n-1)q^D_\omega$. We conclude by Assumption 3, $q^D_\omega < q^\circ$, a contradiction. Hence $q^D_\omega < q^\circ$. By transitivity, $q^D_\omega < q^I_\omega$.

By definition of ω^S (Definition 6),

$$\pi_i(q^S, \theta nq^S + (1-\theta)nq^{D_S}) > \pi_i(q^D_\omega, \theta nq^S + (1-\theta)nq^{D_S}) \tag{15}$$

with $q^{D_S} = b(\theta nq^S + ((1-\theta)n-1)q^{D_S})$. By Assumption 4 (quasiconcavity),

$$\pi_i(q^I_\omega, \theta nq^S + (1-\theta)nq^{D_S}) \geq \pi_i(q^D_\omega, \theta nq^S + (1-\theta)nq^{D_S}). \tag{16}$$

By Lemma 2 (iv), $\theta nq^I_\omega + (1-\theta)nq^D_\omega < \theta nq^S + (1-\theta)nq^{D_S}$. Hence by Assumption 1, Lemma 1, and Remark 1 (strict quasisubmodularity),

$$\pi_i(q^I_\omega, \theta nq^I_\omega + (1-\theta)nq^D_\omega) > \pi_i(q^D_\omega, \theta nq^I_\omega + (1-\theta)nq^D_\omega). \tag{17}$$

It follows that $\omega \in H$. □

Proof of Theorem 1.

Recall that Z is the collection of all absorbing sets of the unperturbed decision dynamics when $\varepsilon = 0$. We need to show that previous assumptions imply $Z = \{\{\omega\} : \omega \in H\}$.

First, we show that every $\{\omega\}$ with $\omega \in H$ is an absorbing state. 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 deviate from its best response 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. Hence $\{\omega\}$ is an absorbing state.

Second, we show that there are no other absorbing sets other than $\{\{\omega\} : \omega \in H\}$. Consider any state $\omega' \notin H$. At least one condition of (i) to (iv) of Definition 7 is violated. From any $\omega' \notin H$ there is an unperturbed adjustment path based on the decision rules leading in the subsequent periods to a state $\omega'' = (q^I, \overset{\theta n}{\ldots}, q^I, q^D, \overset{(1-\theta)n}{\ldots}, q^D)$ in which (a) all imitators play identical quantities q^I and (b) all optimizers play identical best responses $q^D = b(\theta n q^I + ((1-\theta)n - 1)q^D)$. (a) follows from the imitators' decision rule (Definition 2) since by inertia there is positive probably that in the next period all imitators adjust to the same quantity. (b) follows by Assumption 3 (Lemma 2), since each optimizer can reach in finite time the symmetric best response given that imitators play $\theta n q^I$. Note that ω'' satisfies condition (i) and (ii).

Suppose now that (iii) or (iv) are not satisfied in ω''. Then

$$\pi_i(q^I, \theta n q^I + (1-\theta) n q^D) \leq \pi_i(q^D, \theta n q^I + (1-\theta) n q^D)$$

and $q^I \neq q^D$. To show that there exists an unperturbed adjustment path based on the decision rules leading in the subsequent periods to a state $\omega \in H$, it is sufficient to show that

$$\pi_i(q^D, \theta n q^D + (1-\theta) n q^{D'}) \geq \pi_i(q^{D'}, \theta n q^D + (1-\theta) n q^{D'})$$

with strict inequality if and only if $q^D \neq q^{D'}$, and $q^{D'} = b(\theta n q^D + ((1-\theta)n - 1)q^{D'})$. I.e., it is sufficient to show that if imitators mimic the optimizers' outputs in ω'' and optimizers play their symmetric best response to it, the resulting state is a Pseudo-Stackelberg state.

Consider the case $q^D > q^{D'}$: By definition of ω^S (Definition 6), we have

$$\pi_i(q^S, \theta n q^S + (1-\theta) n q^{D_S}) > \pi_i(q^{D'}, \theta n q^S + (1-\theta) n q^{D_S})$$

with $q^{D_S} = b(\theta n q^S + ((1-\theta)n - 1)q^{D_S})$. We claim that $q^S > q^D$ for $q^D = b(\theta n q^I + ((1-\theta)n - 1)q^D)$ for any $q^I \in \Gamma$. By Assumption 2, $q^S > q^{(1-\theta)n} = b(((1-\theta)n - 1)q^{(1-\theta)n}) = b(\theta n \cdot 0 + ((1-\theta)n - 1)q^{(1-\theta)n})$. By Lemma 2 (iv), $\theta n q^I + ((1-\theta)n - 1)q^D > \theta n \cdot 0 + ((1-\theta)n - 1)q^{(1-\theta)n}$ for $q^I \geq 0$ and equality for $q^I = 0$. We need to show that $q^D \leq q^{(1-\theta)n}$: Suppose to the contrary that $q^D > q^{(1-\theta)n}$. Since $q^I \geq 0$, $\theta n q^I + ((1-\theta)n - 1)q^D \geq \theta n \cdot 0 + ((1-\theta)n - 1)q^{(1-\theta)n}$ (with equality if $q^I = 0$). By Assumption 3, it implies $q^{(1-\theta)n} \geq q^D$ (with equality if $q^I = 0$), a contradiction. Hence, $q^S > q^D > q^{D'}$. By Assumption 4 (quasiconcavity),

$$\pi_i(q^D, \theta n q^S + (1-\theta) n q^{D_S}) \geq \min\{\pi_i(q^S, \theta n q^S + (1-\theta) n q^{D_S}), \pi_i(q^{D'}, \theta n q^S + (1-\theta) n q^{D_S})\}.$$

We conclude by definition of ω^S (Definition 6),

$$\pi_i(q^D, \theta nq^S + (1-\theta)nq^{D_S}) \geq \pi_i(q^{D'}, \theta nq^S + (1-\theta)nq^{D_S}).$$

By Assumption 1, Lemma 1 and Remark 1, strict quasisubmodularity of π_i implies

$$\pi_i(q^D, Q) > \pi_i(q^{D'}, Q) \text{ for all } Q < \theta nq^S + (1-\theta)nq^{D_S}.$$

Since $q^S > q^D$ we have by Lemma 2 (iv), $\theta nq^S + (1-\theta)nq^{D_S} > \theta nq^D + (1-\theta)nq^{D'}$. Hence

$$\pi_i(q^D, \theta nq^D + (1-\theta)nq^{D'}) > \pi_i(q^{D'}, \theta nq^D + (1-\theta)nq^{D'}).$$

Consider now the case $q^D < q^{D'}$. We show that there exists an unperturbed adjustment path based on the decision rules leading in the subsequent periods to a state satisfying the previous case. If $q^D < q^{D'}$ and $\omega'' \notin H$, there is by inertia and the imitators' decision rule (Definition 2) positive probability that in the next period all imitators adjust to the same quantity $q^{D'}$. By Assumption 3 (Lemma 2) and inertia, let all optimizers reach in finite time the symmetric best response $q^{D''} = b(\theta nq^{D'} + ((1-\theta)n-1)q^{D''})$ given that imitators adjusted previously to identical outcomes $q^{D'}$. If the resulting state ω''' is in H, then ω' and ω'' were not absorbing. Suppose now that $\omega''' \notin H$. We claim that $q^{D'} > q^{D''}$, hence we are in previous case. Suppose to the contrary that $q^{D''} \geq q^{D'}$. Together with $q^{D'} > q^D$ we conclude $n\theta q^{D'} + ((1-\theta)n-1)q^{D''} > n\theta q^D + ((1-\theta)n-1)q^{D'}$. By Assumption 3, it implies $q^{D''} < q^{D'}$, a contradiction. This completes the proof of Theorem 1. □

In order to characterize the support of the unique limiting invariant distribution, we consider small perturbations, $\varepsilon > 0$. 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 $B(A) = \{\omega \in \Gamma^n | \exists m \in \mathbb{N}, \exists \omega' \in A \text{ s.t. } p^{(m)}_{\omega\omega'} > 0\}$. 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 there exists an $\omega \in A$, $M(\omega) \cap B(A') \neq \emptyset$. 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 4 by Nöldeke and Samuelson (1993, 1997) and Samuelson (1994).

Lemma 4 (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 4 is contained in Samuelson (1997), Lemma 7.1 and Proposition 7.7., proof pp. 236-238.

Proof of Theorem 2.

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

Lemma 5 *Given previous assumptions, if* $\omega^S \in H$ *then* $M(\omega) \cap B(\{\omega^S\}) \neq \emptyset$, *for all* $\omega \in H \setminus \{\omega^S\}$.[14]

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

$$q^S p((\theta n - k)q + kq^S + (1-\theta)nq^D) - c(q^S) > qp((\theta n - k)q + kq^S + (1-\theta)nq^D) - c(q), \tag{18}$$

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_i is strictly quasisubmodular (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 (iv), we conclude that $Q'' > Q'$. If $q'' < q'$ then $\theta n q'' < (\theta n - k)q' + kq''$. By Lemma 2 (iv), we conclude that $Q'' < Q'$. 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 (18). 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 (18). Finally, set $k = 1$ to see that one suitable mutation only is required to connect every $\omega \in H$ to $\omega^S \in H$. □

Since Lemma 5 holds for any absorbing state except the Pseudo-Stackelberg solution, it holds also for the Cournot Nash equilibrium ω°. By Lemma 4, this implies Theorem 2. □

Proof of Example 3.

[14] This result can be easily extended to all states in which imitators play identical quantities and optimizers play identical best responses, i.e., all states satisfying conditions (i) and (ii) of Definition 7. Such generalization may be of interest if one does not want to rule out cycles and gives up Assumption 2.

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.

Denote by $\bar{\omega}$ (resp. $\underline{\omega}$) the Pseudo-Stackelberg state with the largest (resp. smallest) symmetric output of the imitators. I.e., if $\bar{q}^I$ is the symmetric output of any imitator in $\bar{\omega}$, there should not exist $\omega \in H$ s.t. $q^I_\omega > \bar{q}^I$ (analogously for $\underline{\omega}$).

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

Lemma 6 *Under previous assumptions we conclude:*

(i) If $\omega^S \in H$, then there exists an increasing sequence $\omega_1, ..., \omega_m \in H$ with $\omega_1 = \underline{\omega}$ 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 $0 < k < \theta n$

$$\begin{aligned} q^I_{\omega_{j+1}} p((\theta n - k) q^I_{\omega_j} + k q^I_{\omega_{j+1}} + (1-\theta) n q^D) - c(q^I_{\omega_{j+1}}) > \\ q^I_{\omega_j} p((\theta n - k) q^I_{\omega_j} + k q^I_{\omega_{j+1}} + (1-\theta) n q^D) - c(q^I_{\omega_j}), \end{aligned} \tag{19}$$

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

By Assumption 1, Lemma 1 and Remark 1, π_i is strictly quasisubmodular. Set $q'' = q^I_{\omega_{j+1}}$, $q' = q^I_{\omega_j}$, $Q' = (\theta n - k) q^I_{\omega_j} + k q^I_{\omega_{j+1}} + (1-\theta) n q^{D'}$ and $Q'' = \theta n q^I_{\omega_{j+1}} + (1-\theta) n q^{D''}$, 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.

Note that each $q^I_{\omega_{j+1}} \in [\underline{q}^I, q^S] \cap \Gamma$. We claim that $\pi_i(q^I_{\omega_j}, Q'') = \min\{\pi_i(q^S, Q''), \pi_i(q^I_{\omega_j}, Q'')\}$. We know that $\pi_i(q^S, Q^S) > \pi_i(q^I_{\omega_j}, Q^S)$ by definition. Since $q^I_{\omega_j} < q^S$ for $j = 1, ..., m-1$, we know by Lemma 2 that $Q^S > Q''$. By strict quasisubmodularity (upper Formula (3)) $\pi_i(q^S, Q'') > \pi_i(q^I_{\omega_j}, Q'')$ and the claim follows. By quasiconcavity $\pi_i(q^I_{\omega_{j+1}}, Q'') \geq \pi_i(q^I_{\omega_j}, Q'')$. It implies Inequality (19) by strict quasisubmodularity (upper Formula (3)). (i) follows from setting $k = 1$ in Inequality (19).

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

This completes the proof of Lemma 6. □

Lemma 7 *Let $q^{\circ 2} \in \Gamma$ be a firm's Cournot duopoly equilibrium output.*[15] *There exist p, c, θ, δ, ν, and finite n such that*

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

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

$$\pi(q, \nu\delta) < 0, \forall q > 0, \tag{22}$$

$$q^{\circ 2} = \bar{q}^I, \tag{23}$$

$$\underline{q}^I = q^\circ, \tag{24}$$

π_i is quasiconcave in q_i, $\omega^S \in \Gamma^n$ exists, and all previous assumptions are satisfied.

Proof of Lemma. Consider p, c, θ, δ, ν, and n in Example 3, 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 (20) to (24) as well as all previous assumptions hold, and that ω^S exists. Moreover, since π_i is strictly concave in q_i, it is quasiconcave. □

Lemma 8 *Let p, c, θ, δ, ν, and n be such that the properties of Lemma 7 hold. Then $M(\omega) \cap B(\{\underline{\omega}\}) \neq \emptyset$, for all $\omega \in H$.*

Proof of Lemma. Suppose in t any arbitrary state $\omega(t) \in H$. By Lemma 2 such state exists. W.l.o.g. suppose 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 7, Inequality (20), $\pi_j(t+1) < 0$, for all $j \in N$. W.l.o.g. assume that 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 $q_k(t+2) = 0$. $Q(t+2) \geq (2n-3)q^\circ$. By Lemma 7, Inequality (20), $\pi_j(t+1) < 0$, for all $j \in N \backslash \{k\}$. Since $D_O(t+2) = D_I(t+2) = 0$ assume w.l.o.g. that 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 7, Inequality (21), $\pi_i(t+3) \leq 0$. Since $D_O(t+3) = b((n-1)q^\circ) = q^\circ$ assume w.l.o.g. that another firm $k \in N \backslash I$ has the opportunity to adjust output in $t+4$. Since $\pi_k(t+4) > \pi_j(t+4)$, $j \in N \backslash \{k\}$ and $D_I(t+4) = q^\circ$ we can assume w.l.o.g. that all $j \in I$ adjust output. Let all remaining optimizers adjust output in the subsequent periods such that with positive probability ω° is reached in finite time (by Lemma 2). By Lemma 7, $\omega^\circ = \underline{\omega}$. 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 $M(A) \cap B(\{\underline{\omega}\}) \neq \emptyset$, for all $A \in Z$. □

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

Lemma 9 *Let p, c, θ, δ, ν, and n be such that the properties of Lemma 7 hold. Then $M(\omega) \cap B(\{\bar{\omega}\}) \neq \emptyset$, for all $\omega \in H$.*

Proof of Lemma. Suppose in t any arbitrary state $\omega(t) \in H$. W.l.o.g. assume that in $t+1$ a mutation by an imitator $i \in I$ occurs setting a large quantity $\nu\delta$ such that by Lemma 7, Inequality (22), $\pi_j(t+1) < 0$, for all $j \in N$. W.l.o.g. assume that 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 (22), we have $\pi^I(t+2) < \pi^D(t+2)$. W.l.o.g. assume that 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 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 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 subsequent periods such that by 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 7, $\bar{q}^I = q^{\circ 2}$, we have that $\omega^{\circ 2} = \bar{\omega}$. □

In Lemma 6 we showed that we can connect the Pseudo-Stackelberg states by an increasing (starting with Pseudo-Stackelberg state with the smallest output of imitators) 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. In Lemma 8 and 9 we showed that we can connect by a single suitable mutation followed by the decision dynamics any Pseudo-Stackelberg state to the Pseudo-Stackelberg state with the largest (resp. smallest) output of imitators if the properties of Lemma 7 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 4 it follows that $S = H$. This completes the proof of Example 3. □

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, Games and Economic Behavior, forthcoming.

[12] Dürsch, P., Kolb, A., Oechssler, J. and Schipper, B. C., 2005. Rage against the machines: How subjects learn to play against computers, University of California, Davis.

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

[14] 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.

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

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

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

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

[19] Hehenkamp, B. and O. Kaarbøe, 2005. Imitators and optimizers in a changing environment, University of Dortmund.

[20] 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.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

[44] Thijssen, J., 2005. Nearly-complete decomposability and stochastic stability with an application to Cournot oligopoly, Trinity Economic Paper 2005/6.

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

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

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

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

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