Robustness of perfectly competitive equilibria to memory in imitative learning

Carlos Alós-Ferrer* Ana B. Ania†

April 3, 2009

Abstract

We extend the analysis in Alós-Ferrer and Ania (2005) to allow for memory in the imitative learning process. Adding memory allows to detect and correct unprofitable deviations and makes it harder to disturb Nash equilibria with spiteful deviations that achieve a relative advantage. We show for the case of general aggregative submodular games that perfectly competitive outcomes are still stochastically stable of the imitative learning process after the introduction of memory.

Keywords: imitation, learning, finite-population ESS, aggregative games

JEL Classification: C73

1 Introduction

The concept of finite-population evolutionarily stable strategy (ESS) introduced by Schaffer (1988) is defined in the same spirit as its continuum-population counterpart (Maynard Smith and Price, 1973; Maynard Smith, 1974, 1982) – a strategy that, once adopted by the entire population, cannot be invaded by any small fraction of deviators. Invasion is possible only when deviators perform at least as well as non-deviators after deviation and

*Department of Economics, University of Konstanz. Email: carlos.alos-ferrer@uni-konstanz.de

†Department of Economics, University of Vienna and Department of Economics, University of Munich. Email: ana.ania@lrz.uni-muenchen.de

is, thus, based on simultaneous payoff comparisons. Instead, a Nash equilibrium of the game is a strategy profile at which no individual can strictly increase payoffs by unilaterally deviating. The underlying motive for deviation is conceptually different in both definitions of equilibrium. Yet, in a continuum population, deviations have a negligible impact on opponents' payoffs and, thus, gaining a relative advantage can only be achieved through an increase of individual payoffs. It is well known that, as a consequence, ESS in a continuum population are always Nash equilibrium strategies of the underlying game. Instead, in a finite population, a relative advantage may be obtained by deviating from a Nash equilibrium, if the harm imposed on others is larger than the harm imposed on oneself. For this reason, ESS in a finite population may correspond to *spiteful* behavior in the sense of Hamilton (1971) and not necessarily to Nash equilibrium behavior.[1]

Alós-Ferrer and Ania (2005) show two important properties of the concept of finite-population ESS for the case of symmetric aggregative games — games for which payoffs can be expressed as a function of own strategy and some monotonic aggregate of players' strategies. In such games, a generalization of a perfectly competitive equilibrium can be defined as a fixed point of the problem in which players maximize individual payoffs disregarding their impact on the aggregate. If the game exhibits strategic substitutability between individual strategies and the aggregate, then strategies played in a perfectly competitive equilibrium are always ESS. Moreover, in that case, ESS are robust to the introduction of any fraction of deviators; this corresponds to the concept of global stability in Schaffer (1988, p. 473).[2] These properties guarantee that a perfectly competitive equilibrium is stochastically stable in a perturbed imitation dynamics in which players copy any strategy that obtained highest payoffs in the last round, occasionally experimenting

[1]On the comparison between continuum and finite-population ESS see e.g. Vega-Redondo (1996, pp. 13-35). Schaffer (1989, p. 38) gives an interpretation of a finite-population ESS as a Nash equilibrium of the transformed game where players attempt to maximize payoffs relative to the average of opponents' payoffs. Neill (2004) shows that the difference between Nash equilibrium and ESS may persist even when the population size is large but finite and unknown. Ania (2008) and Guse et al. (2008) show equivalence between finite-population ESS and Nash equilibrium in some classes of games.

[2]Possajennikov (2003) was the first to establish the connection between perfectly competitive equilibrium and finite-population ESS. Leininger (2006) gives an alternative proof by induction of the global stability result.

with any available strategy (Alós-Ferrer and Ania, 2005, Proposition 3 and Corollary 1).[3] The relationship to perfectly competitive equilibrium and the property of stochastic stability for imitative learning indicate that finite-population ESS is an important equilibrium concept in economic models of strategic interaction in which decision makers lack the information to act fully rationally and observed successful behavior strongly influences individuals' decisions.

An important critique, however, to the definition of ESS and to this simple imitation dynamics is precisely that it is based on simultaneous payoff comparisons only. In particular, a Nash equilibrium can be abandoned after a deviation resulting in a relative payoff advantage, even though such deviations always result in a reduction of the deviator's payoff. A slight increase of sophistication would immediately preclude the success of such deviations — if individuals had at least one period of memory, they would recognize and be able to correct payoff-reducing deviations, Nash equilibria would not be easily abandoned. Indeed, in an imitative process with memory, deviations only succeed if they obtain both an absolute and a relative payoff advantage.

In the present paper we extend the analysis of Alós-Ferrer and Ania (2005) along those lines. We consider a perturbed imitation dynamics with memory in general aggregative submodular games and show that perfectly competitive equilibria are still stochastically stable and, thus, observed with positive probability in the long run. The paper also generalizes some of the main results in Alós-Ferrer (2004), which considered this kind of dynamics for the particular case of a Cournot oligopoly and provided conditions under which all output levels ranging from Cournot to Walrasian output are stochastically stable.

The paper is organized as follows. Section 2 reviews the definition of ESS and the main result on stability under perturbed imitation. Section 3 reviews the relationship between perfectly competitive equilibrium and ESS in generalized submodular aggregative games. Sections 4 and 5 are the core of the present paper, extending the analysis of the previous sections to imitative learning with memory. All proofs are relegated to the Appendix.

[3] See also Schipper (2004).

2 Stability of finite-population ESS

In the present section we review the definition of finite-population ESS and the stochastic stability result for a perturbed imitation dynamics (Alós-Ferrer and Ania, 2005, Proposition 4).

Consider the game defined by the set of players $I = \{1, \ldots, n\}$ with common and finite strategy set S. A strategy profile is denoted $\mathbf{s} = (s_1, \ldots, s_n)$. Payoffs to player i are given by a function $\pi_i : S^n \to \mathbb{R}$ which is assumed to be symmetric; i.e. invariant to any permutation of opponents' strategies. Formally,

$$\pi_i(\mathbf{s}) = \pi(s_i|\mathbf{s}_{-i}) = \pi(s_i|\mathbf{s}'_{-i})$$

where s_i denotes i's strategy in $\mathbf{s}$, $\mathbf{s}_{-i} = (s_j)_{j\neq i}$, and $\mathbf{s}'_{-i}$ is any permutation of $\mathbf{s}_{-i}$.

ESS and globally stable ESS

A strategy $s \in S$ is a finite-population *evolutionarily stable strategy* (ESS) if, for all $s' \in S$,

$$\pi(s|s', s, \overset{n-2}{\ldots}, s) \geq \pi(s'|s, \overset{n-1}{\ldots}, s).$$

We say that $s \in S$ is *m-stable* if, for all $s' \in S$,

$$\pi(s|s', \overset{m}{\ldots}, s', s, \overset{n-1-m}{\ldots}, s) \geq \pi(s'|s', \overset{m-1}{\ldots}, s', s, \overset{n-m}{\ldots}, s).$$

We say that $s \in S$ is *globally stable* if it is m-stable for all $m \in \{1, \ldots, n-1\}$. Finally, we say that s is *strict* ESS, *strictly* m-stable, or *strictly* globally stable, if the inequalities above are fulfilled strictly.

Intuitively, s is an ESS if it gives higher payoff than any deviating strategy after deviation from the symmetric profile. As a first difference to its continuum-population counterpart, when only the smallest fraction of deviators is allowed ($1/n$), the deviator in a finite population does not face other deviators. A strategy s is m-stable if it gives higher payoff than the deviating strategy after m simultaneous deviations to the same s'. The concept of m stability captures robustness to the introduction of a larger fraction of deviators. Finally, s is globally stable if it is robust to any fraction of deviators.

Note that s is an ESS if and only if

$$s \in \arg\max_{s'} \left[\pi(s'|s, \overset{n-1}{\ldots}, s) - \pi(s|s', s, \overset{n-2}{\ldots}, s) \right]$$

A finite-population ESS maximizes the payoff *difference* between the deviator and non-deviators after deviation. Therefore, it need not be a Nash equilibrium of the underlying game, defined by absolute payoff maximization given opponents' strategies. This discrepancy comes about in finite populations, where single deviations have an impact on opponents' payoffs and deviating from a Nash equilibrium may result in a larger damage on opponents' payoffs and, thus, in a relative payoff advantage.

Schaffer (1989, pp. 31-34) illustrates this for the case of a Cournot duopoly with linear demand and linear costs. Starting at the Cournot-Nash equilibrium, if a firm sets output equal to the Walrasian output (the one that would equate price to marginal cost), market price goes down and both firms attain lower profits. The deviating firm, however, with a larger market share, has higher profits than the non-deviating firm and attains a relative advantage. Starting at a Walrasian equilibrium, instead, no deviation can result in a relative advantage. Reducing output (e.g. by choosing a best response) would increase payoffs to the non-deviator more than the deviator's payoff, resulting in a relative disadvantage. Alternatively, increasing output would lead both firms to losses, which are higher for the deviator with a higher market share.

Imitation dynamics

Suppose the game described above is played recurrently in discrete time. Players in I know their set of strategies S, but they do not necessarily know the payoff function π and, thus, may not be able to compute a best response, nor to infer how successful strategies really are. Instead, every $t = 1, 2, \ldots$ players observe the strategies currently chosen, $\mathbf{s}(t)$, and the vector of payoffs obtained, $(\pi_1(\mathbf{s}(t)), \ldots, \pi_n(\mathbf{s}(t)))$. At the end of each t, each player has a probability $\lambda \in (0, 1)$ to get a revision draw. In the event of strategy revision, players choose one of the strategies observed that gave highest payoffs in the current round of play. Thus, if called to revise their strategies, players *imitate the best* performing strategies observed. Before choices are made and independently of the outcome of imitation, each player gets an experimentation draw with probability $\varepsilon \in [0, 1)$.

In the event of experimentation, players choose any strategy in S with positive probability. Imitation and experimentation opportunities are drawn independently across players and across time periods.

The resulting stochastic process is a Markov chain with state space S^n. For each $\varepsilon > 0$, the process is irreducible and aperiodic and, thus, has a unique invariant distribution, μ^ε. Our focus is on the so called *limit invariant distribution*, $\mu = \lim_{\varepsilon \to 0} \mu^\varepsilon$. The states in the support of the limit invariant distribution can be characterized using well known techniques proposed by Foster and Young (1990), Young (1993), and Kandori et al. (1993) for the analysis of this kind of evolutionary learning models. These states, called *stochastically stable states* or *long-run equilibria*, are the ones observed most of the time along any sample path of the process. The next lemma states that, if the game has a strictly globally stable ESS as defined above, then this is essentially the only strategy we observe in the long run of the perturbed imitation dynamics just described.

Lemma 1. If s^E is a strictly globally stable ESS, then $\mathbf{s}^E = (s^E, \ldots, s^E)$ is the unique stochastically stable state of the perturbed imitation dynamics.

We provide here an intuitive sketch of proof using Ellison (2000, Theorem 1). We refer to Alós-Ferrer and Ania (2005) for further details. The uniqueness result that follows from strict ESS is shown in Proposition 2(i) in Section 4 below. Intuitively, if there were two different strictly globally stable strategies, then one of them would violate the definition of ESS. We now turn to the issue of stochastic stability.

The unperturbed imitation process (with $\varepsilon = 0$) quickly leads to symmetric profiles. To see this, note that at any given period, there is positive probability that all players revise their strategies simultaneously and, in doing so, that they all choose the same strategy out of the set of best performing strategies. Thus, from any state there is positive probability to reach a symmetric profile. Once a symmetric profile has been reached, it cannot be abandoned by imitation alone.

For $\varepsilon > 0$ it is now enough to restrict attention to symmetric profiles and *count* the number of mutations needed to exit and to reach any symmetric profile. Suppose the process starts at the state $\mathbf{s}^E$, where all players choose the strictly globally stable ESS. The fact that s^E is a strict ESS implies that a single deviator choosing any other strategy attains strictly lower payoffs than non deviators and will not be followed. It therefore takes more than one experimenting deviator to exit $\mathbf{s}^E$. On the other hand, the property of strict global stability, implies that, from any other symmetric profile, a single deviator experimenting with strategy s^E attains strictly higher payoffs than non deviators and will be followed with positive probability. Thus, a single experimenting deviator is enough to take the system from any state to $\mathbf{s}^E$. In Ellison's terminology, $\mathbf{s}^E$ has coradius 1 (can be reached with at most one deviation) and radius strictly higher than 1 (requires more than one simultaneous deviation to be abandoned). This makes it much more likely that the process moves to state $\mathbf{s}^E$ than it moves out of it and the state where the strictly globally stable ESS is played is, thus, the unique long-run equilibrium.

3 Perfectly competitive equilibrium

In the present section we turn to generalized aggregative games, restricting attention to symmetric games. For this common class of games we define the concept of aggregate-taking strategy (ATS), a substantial generalization of the concept of perfectly competitive equilibrium. We then review the relationship between ATS and ESS for generalized submodular aggregative games obtained in Alós-Ferrer and Ania (2005, Proposition 3).

Generalized submodular aggregative games

A game has an aggregative structure if the payoff to any player can be alternatively written as (a) a function of the strategy profile, (b) a function of the player's own strategy and an aggregate of *opponents'* strategies, and (c) a function of the player's own strategy and an aggregate of *all* players' strategies. Formally, a symmetric game with common strategy set $S \subseteq \mathbb{R}$ is a generalized aggregative game if payoffs can be alternatively written using one of the three following payoff functions[4]

$$\pi(s_i|\mathbf{s}_{-i}) = \widetilde{\pi}(s_i, g^n(\mathbf{s})) = \widehat{\pi}(s_i, g^{n-1}(\mathbf{s}_{-i})),$$

where $g^k : S^k \to \mathbb{R}$ defines a family of symmetric and strongly increasing aggregate func-

[4]It is not essential that strategies are real numbers. Bellow we only need that S is an ordered set.

tions,[5] which can be recursively constructed as $g^{k+1}(s_1, \dots, s_{k+1}) = g(s_{k+1}, g^k(s_1, \dots, s_k))$; the function $\widetilde{\pi} : S \times \mathbb{R} \to \mathbb{R}$ expresses payoffs depending on the player's own strategy and an aggregate of all strategies; respectively, $\widehat{\pi} : S \times \mathbb{R} \to \mathbb{R}$ expresses payoffs depending on the player's own strategy and opponents' strategies.

The payoff function $\widetilde{\pi}$ (alternatively $\widehat{\pi}$) is said to be (quasi)-*submodular* (in a weak sense) if it satisfies the *dual single crossing property* in $(s, x) \in S \times \mathbb{R}$; i.e. if for all $s'' > s'$ and $x'' > x'$ both conditions below hold

$$\widetilde{\pi}(s'', x') \leq \widetilde{\pi}(s', x') \Longrightarrow \widetilde{\pi}(s'', x'') \leq \widetilde{\pi}(s', x'')$$
$$\widetilde{\pi}(s'', x') < \widetilde{\pi}(s', x') \Longrightarrow \widetilde{\pi}(s'', x'') < \widetilde{\pi}(s', x'')$$

We use here the ordinal version of submodularity which is weaker than decreasing differences (see e.g. Topkis, 1998, p. 59); in words, if a given increase in the value of a player's strategy from s' to s'' does not increase payoffs (alternatively, reduces payoffs) given an aggregate value of x', the same increase in the strategy will not increase payoffs (alternatively, reduce payoffs) for any higher value of the aggregate x''.

In the sequel, we refer to *generalized submodular aggregative games* as games that can be rewritten in the three different forms just presented, using some aggregate function, and for which either the function $\widetilde{\pi}(s, x)$, or the function $\widehat{\pi}(s, x)$, or both satisfy the dual single-crossing property. It should be apparent from the context which of this submodularity assumptions we are imposing.

To understand intuitively the extent to which these different requirements impose different restrictions on the game, it is worth considering the simple example of a Cournot oligopoly (Alós-Ferrer and Ania, 2005, Example 1). Profits in a Cournot oligopoly can be written using total output as an aggregate. It is easy to check that a decreasing demand function is enough to guarantee submodularity of $\widetilde{\pi}$. Submodularity of $\widehat{\pi}$ requires the additional, though common assumption of decreasing marginal revenues.

[5]By symmetric we mean invariant to strategy permutations. A function is strongly increasing if it is strictly increasing in every coordinate. In our results, the latter can sometimes be relaxed to allow for aggregates like the minimum function (cf. proof of Lemma 2).

ESS and perfectly competitive equilibrium

We can now define an aggregate-taking strategy (ATS) as a strategy $s^* \in S$ that solves the following fix-point problem

$$s^* \in \arg\max_s \widetilde{\pi}(s, g^n(s^*, \ldots, s^*)).$$

We say that s^* is a strict ATS if it is a strict maximizer of $\widetilde{\pi}$.

The concept of ATS is usually interpreted as a generalization of a perfectly competitive equilibrium. It is a payoff maximizer at the aggregate that results when all players choose s^*. Choosing s^* implicity requires that players disregard not only the externality they exert on others (as in a Nash equilibrium), but also their own effect on the aggregate. In applications, however, it does not necessarily coincide with what we would call a competitive equilibrium in the economic sense. For example, in Dixit's model of oligopolistic competition with product differentiation and linear demand, payoffs can be written in terms of an aggregate of strategies. The corresponding ATS, however, does not correspond to price equal marginal cost. See Takana (2000) for an evolutionary analysis of this model.

The next proposition considers a particular case in Alós-Ferrer and Ania (2005, Proposition 3). The proof is included in the Appendix for the sake of completeness.

Proposition 1. If s^* is a strict ATS, then it is always a strictly globally stable ESS.

It follows immediately from Proposition 1 and Lemma 1 that a strict ATS must be the long-run outcome of the perturbed imitation dynamics described in Section 2. This result is contained in the next Corollary.

Corollary 1. If s^* is a strict ATS, then it is the unique stochastically stable state of the perturbed imitation dynamics.

If learning is mainly driven by imitation of strategies that are currently observed to be the most successful, with some infrequent experimentation, then in the long run essentially all we observe is an ATS. Following an imitative behavioral rule, individuals turn out to behave as if they maximized payoffs given some aggregate. In most cases, this will differ from Nash equilibrium behavior; i.e. imitation will lead away from Nash equilibrium. By definition, deviating from a Nash equilibrium reduces the deviator's payoffs. Yet, this deviation can be successfully followed by imitators if non-deviators' payoffs are reduced even more (effect of *spite*). It is in this sense that imitation is a extremely naïve behavioral rule — recent payoff information is disregarded in favor of immediate inter-player payoff comparisons.

4 Imitation dynamics with memory

In the present section we extend the analysis to imitative behavioral rules with some periods of memory. The length of memory will turn out to be inconsequential for our results. As we will see, however, introducing memory changes the nature of the dynamics considerably. Memory allows, in particular, intertemporal payoff comparisons that enable detection and correction of mistakes or harmful deviations, like the ones resulting when deviating from a Nash equilibrium.

Description of the dynamics

Let us consider the following variation of the dynamics presented in Section 2. Every $t = 1, 2, \ldots$ players observe the strategies chosen and the payoffs obtained in the last $k > 1$ periods. The state of the process is given by the vector of strategy profiles in the last k periods, $\mathbf{s}_t^k = (\mathbf{s}(t-k), \ldots, \mathbf{s}(t))$. Observed payoffs are given by a vector $(\pi_{t-k}, \ldots, \pi_t)$ where $\pi_\ell = (\pi_1(\mathbf{s}(\ell)), \ldots, \pi_n(\mathbf{s}(\ell)))$ with $t - k \leq \ell \leq t$. Let

$$O(\mathbf{s}_t^k) = \{s \in S \mid s = s_i(\ell) \text{ for some } i \in I \text{ and some } \ell \in \{t-k, \ldots, t\}\}$$

be the set of strategies observed in the last k periods. Define

$$B(\mathbf{s}_t^k) = \left\{s \in O(\mathbf{s}_t^k) \mid \pi_i(\mathbf{s}(\ell)) \geq \pi_j(\mathbf{s}(\ell')) \text{ for all } j \in I \text{ and all } \ell' \in \{t-k, \ldots, t\}\right\}.$$

The set $B(\mathbf{s}_t^k)$ contains the strategies that gave highest payoffs in the last k periods.

At the end of each t, each player has a probability $\lambda \in (0, 1)$ to get a revision draw. In the event of strategy revision, players choose any of the strategies in $B(\mathbf{s}_t^k)$ with positive probability; i.e. players adopt one of the best performing strategies of the last k periods.

Before choices are made and independently of the outcome of imitation, each player gets an experimentation draw with probability $\varepsilon \in [0, 1)$. In the event of experimentation, players choose any strategy in S with positive probability. Imitation and experimentation opportunities are drawn independently across players and across time periods.

The resulting stochastic process is a Markov chain analogous to the one analyzed in Section 2. The state space is now S^{kn}. For each $\varepsilon > 0$, the process is still irreducible and aperiodic. Notice there is always positive probability that all individuals experiment every period with any strategy and that no revision or experimentation takes place in any given period. This implies that there is positive probability to reach any state from any other state in at most k periods, which gives irreducibility. Similarly, there is positive probability that the process stays at any give state for one period, which gives aperiodicity. Therefore, the Markov chain has a unique invariant distribution, μ^{ε}. Again we focus on the *limit invariant distribution*, $\mu = \lim_{\varepsilon \to 0} \mu^{\varepsilon}$. If $\mu(\mathbf{s}) > 0$, we say that $\mathbf{s}$ is *stochastically stable* – the state will be observed in the long run when the probability of experimentation is very small.

It is important to highlight the difference between this and the dynamics presented in Section 2 and, in particular, clarify the effect that the introduction of memory will have in our arguments. Payoffs to any observed strategy are compared here to all payoffs observed in the last k periods; in particular, any deviating player compares current and past payoffs. This allows to detect the reduction of payoffs due to a deviation from a Nash equilibrium profile. Deviations can only be followed successfully by imitation now, if they achieve an absolute as well as a relative payoff advantage. In particular, a simple individual deviation from Nash equilibrium cannot succeed.

Bounds to stochastic stability

We now proceed to characterize the set of long-run equilibria of this modified version of imitation. In a first step, we proceed to limit the set of states that are candidates for stochastic stability. The next results show that monotonicity of the payoff function with respect to the aggregate and submodularity of payoffs are sufficient to guarantee that strategies in a long-run equilibrium are contained in an interval bounded between a Nash equilibrium and a strict ATS (respectively, by Proposition 1 a strictly globally stable ESS).

Lemma 2. Consider a generalized submodular aggregative game with $\widetilde{\pi}(s, x)$ strictly decreasing in x. If s^E is a globally stable ESS and s^N is a symmetric Nash equilibrium, then $s^E \geq s^N$. If at least one of them is strict, then $s^E > s^N$.

Lemma 2 reproduces what we know for the particular case of a Cournot oligopoly; total output in a Walrasian equilibrium is larger than output in a Cournot equilibrium. As the lemma shows, this comparison can be extended to any globally stable ESS and Nash equilibrium of the game.

Proposition 2. Consider a generalized submodular aggregative game with $\widetilde{\pi}(s, x)$ strictly decreasing in x. Let s^E be a strict ESS. Then:

(i) s^E is strictly globally stable and the unique ESS;

(ii) if $\mathbf{s} = (s, \ldots, s)$ is stochastically stable, then $s \leq s^E$.

Provided the game has a strict ESS, Proposition 2 shows that this will be the upper bound to the set of strategies that can be observed in a stochastically stable state. The intuition of the proof is as follows. The first part of the result, strict global stability and uniqueness, is a simple technical consequence of monotonicity of the aggregate and submodularity. The second part of the result can be argued as follows. A strict ESS cannot be disturbed by a single experimenting player, since any deviation would result in a strict payoff disadvantage. However, starting at a symmetric profile with a strategy $s > s^E$, it is easy to see that a single deviation to s^E attains both an absolute, and a relative payoff advantage. To see this, notice that a deviation to s^E would result in a reduction of the aggregate. All payoffs would increase, leading to an absolute advantage. At the same time, strict global stability also guarantees that s^E attains a relative advantage.

Proposition 3 below is then the corresponding result that gives the lower bound of the set, corresponding to a Nash equilibrium. Recall we used two different ways of rewriting the game in the definition of an aggregative game. For the next proposition, the payoff function $\widehat{\pi}$, expressing payoffs as a function of the player's own strategy and an aggregate of opponents' strategies, becomes important. Otherwise, the result and its proof is analogous to that of Proposition 2.

Proposition 3. Consider a generalized submodular aggregative game with $\widetilde{\pi}(s, x)$ strictly decreasing in x. Let $\mathbf{s}^N = (s^N, \ldots, s^N)$ be a strict symmetric Nash equilibrium. Then:

(i) $\mathbf{s}^N$ is the unique symmetric Nash equilibrium of the game;

(ii) if $\mathbf{s} = (s, \ldots, s)$ is stochastically stable, then $s \geq s^N$.

It is worth pointing out that the proof of Proposition 3(i) shows and exploits the following strong property of a Nash equilibrium strategy in this case. For all $m \in \{1, \ldots, n-1\}$ and all $s \in S$ we have

$$\widehat{\pi}(s^N, g^{n-1}(s, \overset{m}{\ldots}, s, s^N, \overset{n-1-m}{\ldots}, s^N)) \geq \widehat{\pi}(s, g^{n-1}(s, \overset{m}{\ldots}, s, s^N, \overset{n-1-m}{\ldots}, s^N))$$

In words, a Nash equilibrium strategy s^N is always a better reply against an opponents' profile of the type $(s, \ldots, s, s^N, \ldots, s^N)$ than s itself. This property is analogous to the global stability property of ESS and allows to use the same kind of arguments as in Proposition 2.

Lemma 1, Propositions 2 and 3 together now imply the next theorem which establishes that strategies in a long-run equilibrium are always contained in the interval between the Nash equilibrium and the ESS.

Theorem 1. Consider a generalized submodular aggregative game, with g^n symmetric and strongly increasing and both $\widetilde{\pi}(s, x)$ and $\widehat{\pi}(s, x)$ satisfying (ordinal) submodularity. Assume $\widetilde{\pi}(s, x)$ is strictly decreasing in x. Let s^E be a strict ESS and s^N be a strict symmetric Nash equilibrium of the game. Then, stochastically stable states of the perturbed imitation dynamics for any memory size $k \geq 1$ are symmetric profiles $\mathbf{s} = (s, \ldots, s)$ with $s^N \leq s \leq s^E$.

Alós-Ferrer (2004, Theorem 1) is a direct corollary of this result. It considers a symmetric Cournot oligopoly with decreasing marginal revenues. In that setup our Proposition 1 implies that the individual output level corresponding to a strict Walrasian equilibrium is always a strict ESS. It then follows from our Theorem 1 that mimicking any of the output levels that led to highest profits in the last k rounds, together with occasional experimentation, would lead the market to some intermediate outcomes between Cournot and Walrasian equilibrium.

5 Stability of ATS with memory

The analysis in Section 4 shows how adding memory to the imitation process makes Nash equilibria harder to destabilize. The question addressed in the present section ist whether the result on stochastic stability of ATS (cf. Corollary 1) is still valid after the introduction of memory. Surprisingly, this question is answered in the affirmative. To obtain our main result contained in Theorem 2 below, we will need to add some structure to the aggregate function and two additional technical assumptions to the payoff function, upper-hemicontinuity and quasiconcavity of $\widetilde{\pi}$.

In the present section we assume that $S \subseteq \mathbb{R}$ is a regular grid of the form

$$S = \{\underline{s}, \underline{s} + \delta, \underline{s} + 2\delta, \ldots, \overline{s} - 2\delta, \overline{s} - \delta, \overline{s}\}$$

with δ sufficiently low so that $s^N, s^* \in S$.

Definition 1. We say that an aggregate g^n is *decreasing in variance* if, for all $s \in S$ and all $\delta > 0$,

$$g^n(s - \delta, s + \delta, s, \ldots, s) \leq g^n(s, \ldots, s)$$

The sum is obviously decreasing in variance. We offer here a practical sufficient condition for an aggregate to be decreasing in variance.

Proposition 4. Suppose

$$g^n(s_1, \ldots, s_n) = f\left(\sum_{i=1}^{n} h(s_i)\right)$$

where $n \geq 2$, $h : S \mapsto \mathbb{R}$ is a concave function and $f : \mathbb{R} \mapsto \mathbb{R}$ is increasing. Then, g^n is decreasing in variance.

Notice that our inductive definition of aggregate forces $f(h(s_i)) = s_i$, i.e. f is a left-inverse of h.

As an immediate application, rent-seeking games (where $S \subseteq \mathbb{R}_+$) have aggregates which are decreasing in variance for any parameter $0 < r < 1$ (decreasing returns to scale). Indeed, the aggregate

$$g^n(s_1, \ldots, s_n) = \left[\sum_{i=1}^{n} (s_i)^r\right]^{1/r}$$

is in the form given in the last lemma with $f(x) = x^{1/r}$ (obviously increasing for positive x) and $h(s) = s^r$, which is concave for $0 < r < 1$.

We are now ready to show stochastic stability of ATS. The next theorem is a more restrictive version of Theorem 1.

Theorem 2. Consider a generalized submodular aggregative game with g^n decreasing in variance, and $\widetilde{\pi}(s, x)$ strictly decreasing in x and strictly quasiconcave and upper-semicontinuous in s. Let s^* be an ATS and s^N be a strict, symmetric Nash equilibrium. Then, stochastically stable states of the perturbed imitation dynamics for any memory size $k \geq 1$ are symmetric profiles $\mathbf{s} = (s, \ldots, s)$ with $s^N \leq s \leq s^*$. Further, $\mathbf{s}^* = (s^*, \ldots, s^*)$ is always stochastically stable.

Note that strict quasiconcavity of the payoff function implies that an ATS is always a strict maximizer of $\widetilde{\pi}$. Recall also from Proposition 1 that a strict ATS is always a strictly globally stable ESS. Then, Theorem 1 bounds the set of stochastically stable states for strategies in the interval $[s^N, s^*]$.

The proof of stochastic stability of $\mathbf{s}^*$ is based on two observations. First, symmetric profiles with strategy in the interval $[s^N, s^*]$ are robust to single deviations. Essentially, a single deviation upwards may attain a relative but not an absolute advantage. Respectively, a single deviation downwards may attain an absolute but not a relative advantage. However, two simultaneous deviations are enough to exit those states. The main idea is that any two deviations from such a symmetric profile that result in a lower value of the aggregate will increase all payoffs. Moreover, submodularity implies that payoffs to the highest strategy will increase most (cf. Lemma 3). Therefore, starting at any $\mathbf{s} = (s, \ldots, s)$ with $s \in [s^N, s^*]$, it is always possible to move the process *upwards* with two simultaneous deviations of the form $s - \delta$ and $s + \delta$. Strategy $s + \delta$ attains higher absolute and relative payoffs and will be followed by imitation. This allows to move the process from $\mathbf{s}$ to $\mathbf{s}^*$ and even further up to higher values of s with chains of two simultaneous deviations. Recalling from the proof of Theorem 1 that, from symmetric profiles with strategies higher than s^*, it is always possible to move back to $\mathbf{s}^*$ with a single deviation, we have that $\mathbf{s}^*$ is always stochastically stable.

Theorem 2 does not exclude that other strategies in the interval $[s^N, s^*]$ are observed in the long run. For the case of a Cournot oligopoly Alós-Ferrer (2004, Theorem 2) gives conditions such that all strategies between the Cournot-Nash equilibrium output and the Walrasian (competitive) output are observed in the long run. In the next section we generalize this result.

6 Stability of Nash equilibrium

In the present section we provide sufficient conditions for the stochastic stability of all symmetric states with strategies in the interval $[s^N, s^*]$. In the proof of Theorem 2 we showed that below s^* it is always possible to move the process to symmetric profiles with higher values of s using two deviations, one to a slightly lower strategy, $s-\delta$, and one to a slightly higher strategy, $s+\delta$. Assuming that the value of the aggregate does not increase after two such deviations, submodularity implies that $s+\delta < s^*$ attains higher profits and will be followed by imitation. Notice that the same argument would go through if we considered deviations to $s+\delta$ and any $s' < s-\delta$, due to monotonicity of the aggregate; that it, below s^* transitions upwards are always possible with two deviations that do not increase the value of the aggregate.

To show stability of all strategies in the interval $[s^N, s^*]$, we proceed as follows. First, we show that there exist strategy $\widehat{s} > s^*$, such that from state $\widehat{\mathbf{s}} = (\widehat{s}, \ldots, \widehat{s})$ the Nash equilibrium state $\mathbf{s}^N$ can be reached after a single deviation. Then we show that state $\widehat{\mathbf{s}}$ can always be reached with two deviations from $\mathbf{s}^*$, one to a sufficiently low strategy, which we denote s^0, and one to $\widehat{s}$. In summary, what we show is that the process can move with a chain of one or at most two deviations up to $\widehat{\mathbf{s}}$; from there one deviation is enough to move the process back down to state $\mathbf{s}^N$.

Consider states of the type $\mathbf{s} = (s^N, s^*, \overset{n-2}{\ldots}, s^*, s)$ with $s > s^*$. These can be interpreted as the profiles resulting after two simultaneous deviations from $\mathbf{s}^*$, one to the Nash strategy, s^N, and one to a larger strategy s. Denote $y(s) = g^n(s^N, s^*, \overset{n-2}{\ldots}, s^*, s)$, the corresponding aggregate after deviation. Let $\widehat{s} \in S$ be such that $\widehat{s} > s^*$ and

$$\widetilde{\pi}(s^N, y(\widehat{s})) \geq \widetilde{\pi}(\widehat{s}, y(\widehat{s})).$$

By submodularity of $\widetilde{\pi}$, if $\widetilde{\pi}(\widehat{s}, y(\widehat{s})) \leq \widetilde{\pi}(s^N, y(\widehat{s}))$, then $\widetilde{\pi}(\widehat{s}, y) \leq \widetilde{\pi}(s^N, y)$ for all $y > y(\widehat{s})$. In particular, for $y = g^n(s^N, \widehat{s}, \ldots, \widehat{s})$ this implies that it is possible to reach $\mathbf{s}^N$ from $\widehat{\mathbf{s}}$ with a single deviation to s^N. However, at $y(\widehat{s})$ we may have that

$$\widetilde{\pi}(s^*, y(\widehat{s})) > \widetilde{\pi}(s^N, y(\widehat{s})) \geq \widetilde{\pi}(\widehat{s}, y(\widehat{s})).$$

Thus, it may not possible to reach $\widehat{\mathbf{s}}$ from $\mathbf{s}^*$ with the two deviations considered above. Yet, we can find s^0 sufficiently lower than s^N so as to make that transition possible.

To see this, consider states of the type $\mathbf{s} = (s, s^*, \overset{n-2}{\ldots}, s^*, \widehat{s})$ with $s < s^N$ and $\widehat{s}$ as above. These can be interpreted as the states resulting after two deviations from $\mathbf{s}^*$, one to $\widehat{s}$ and one to a lower strategy s. Denote $x(s) = g^n(s, s^*, \overset{n-2}{\ldots}, s^*, \widehat{s})$, the corresponding aggregate after deviation. Let $s^0 \in S$ be such that $s^0 < s^N$ such that

$$\widetilde{\pi}(s^*, x(s^0)) \leq \widetilde{\pi}(\widehat{s}, x(s^0)).$$

Provided s^0 and $\widehat{s}$ exist, we can now show stability of Nash equilibrium as follows. From all symmetric profiles with $s < s^*$, it is possible to reach $\mathbf{s}^*$ with a chain of one or two deviations. From $\mathbf{s}^*$ two deviations to s^0 and $\widehat{s}$ will move the process to the state $\widehat{\mathbf{s}}$. From $\widehat{\mathbf{s}}$, one deviation to s^N is enough to take the process back to $\mathbf{s}^N$. The result is summarized in the following theorem. The formal proof is relegated to the Appendix.

Theorem 3. Consider a generalized submodular aggregative game with g^n continuous and decreasing in variance, and $\widetilde{\pi}(s, x)$ strictly decreasing in x and strictly concave and upper-semicontinuous in s. Let s^* be an ATS and s^N be a strict, symmetric Nash equilibrium. Then, all symmetric profiles $\mathbf{s} = (s, \ldots, s)$ with $s^N \leq s \leq s^*$ are stochastically stable states of the perturbed imitation dynamics for any memory size $k \geq 1$.

Appendix

Proof of Proposition 1. Let s^* be a strict ATS. From the symmetric profile where all players choose s^*, consider m simultaneous deviations to any $s \neq s^*$ with $1 \leq m \leq n-1$ and denote $x^* = g^n(s^*, \overset{n}{\ldots}, s^*)$ and $y^m = g^n(s, \overset{m}{\ldots}, s, s^*, \overset{n-m}{\ldots}, s^*)$. If $s > s^*$, then $y^m > x^*$ by strong monotonicity of the aggregate. Since s^* is a strict ATS, we have $\widetilde{\pi}(s^*, x^*) > \widetilde{\pi}(s, x^*)$ for all $s \neq s^*$. Submodularity of $\widetilde{\pi}$ implies $\widetilde{\pi}(s^*, y^m) > \widetilde{\pi}(s, y^m)$ for all m. Alternatively, if $s < s^*$, $y^m < x^*$. By contradiction, suppose the strict global stability property is not fulfilled; i.e. for some s and some m we have $\widetilde{\pi}(s^*, y^m) \leq \widetilde{\pi}(s, y^m)$. Submodularity implies $\widetilde{\pi}(s^*, x^*) \leq \widetilde{\pi}(s, x^*)$, contradicting the definition of strict ATS. ■

Proof of Lemma 2. Let s^E be a globally stable ESS and s^N a Nash equilibrium of the game. Denote the corresponding aggregates before and after deviation from a Nash equilibrium as $x^N = g^n(s^N, \ldots, s^N)$ and $y = g^n(s^E, s^N, \ldots, s^N)$. We must have that

$$\widetilde{\pi}(s^N, x^N) \geq \widetilde{\pi}(s^E, y) \geq \widetilde{\pi}(s^N, y).$$

The first inequality follows from the definition of Nash equilibrium (deviating from a Nash equilibrium reduces payoffs); the second inequality follows from the definition of ESS (after deviation, the ESS has a relative advantage at the new aggregate y). Suppose by contradiction that $s^N > s^E$. Then $x^N > y$ and, by strict monotonicity, we have that $\widetilde{\pi}(s^N, y) > \widetilde{\pi}(s^N, x^N)$, a contradiction.

Alternatively, if either s^N or s^E is strict, one of the inequalities above becomes strict. By contradiction, if now $s^N \geq s^E$, then $x^N \geq y$ and $\widetilde{\pi}(s^N, y) \geq \widetilde{\pi}(s^N, x^N)$, which gives the corresponding contradiction.[6] ■

Proof of Proposition 2. To show (i) let s^E be a strict ESS and, starting at the state $\mathbf{s}^E$, consider m deviations to some $s \neq s^E$. Denote $y^m = g^n(s, \overset{m}{\ldots}, s, s^E, \overset{n-m}{\ldots}, s^E)$ with $1 \leq m \leq n-1$. Suppose, first, that $s > s^E$, then $y^{m'} > y^m$ for all $m' > m$.

[6] It is worth pointing out that the assumption of strong monotonicity of the aggregate included in our definition of aggregative game in Section 3 could be relaxed here if either s^N or s^E is strict. This follows immediately from this proof and it would allow to accommodate aggregates like the minimum function.

By definition of strict ESS, we have $\widetilde{\pi}(s^E, y^1) > \widetilde{\pi}(s, y^1)$. Submodularity then implies $\widetilde{\pi}(s^E, y^m) > \widetilde{\pi}(s, y^m)$ for all m. Respectively, for deviations to lower strategies, suppose by contradiction that there exit m and $s < s^E$ sucht that $\widetilde{\pi}(s, y^m) \geq \widetilde{\pi}(s^E, y^m)$. Since now $y^{m'} < y^m$ for all $m' > m$, by submodularity, this would imply $\widetilde{\pi}(s, y^1) \geq \widetilde{\pi}(s^E, y^1)$, contradicting the definition of strict ESS. This shows that strict ESS implies strictly globally stable ESS. To show uniqueness notice strict global stability implies in particular that $\widetilde{\pi}(s, y^{n-1}) < \widetilde{\pi}(s^E, y^{n-1})$ for all $s \neq s^E$, which in turn implies that s cannot be ESS.

To show (ii) let the process start at some $\mathbf{s} = (s, \ldots, s)$ with $s > s^E$, then a single deviation to s^E would result in an absolute as well as a relative payoff advantage. To see this, denote $x = g^n(s, \ldots, s)$ and, as before, $y^m = g^n(s, \overset{m}{\ldots}, s, s^E, \overset{n-m}{\ldots}, s^E)$. Starting at $\mathbf{s} = (s, \ldots, s)$, consider a deviation from s to s^E resulting in state $\mathbf{s}' = (s, \ldots, s, s^E)$ with aggregate $y^{n-1} < x$. Since payoffs are strictly decreasing in the aggregate, we have $\widetilde{\pi}(s, y^{n-1}) > \widetilde{\pi}(s, x)$ – payoffs increase for non-deviators. Moreover, since s^E is strictly globally stable $\widetilde{\pi}(s^E, y^{n-1}) > \widetilde{\pi}(s, y^{n-1})$ – at the new aggregate the deviator choosing s^E has strictly higher payoffs. In summary, after deviation, the deviator choosing s^E has higher payoffs than non-deviators and payoffs to non-deviators have increased; both together imply that the deviator choosing s^E obtains a relative as well as an absolute payoff advantage from deviation and will be followed at the next revision opportunity. This shows that a single experiment is enough to move the process from $\mathbf{s}$ to $\mathbf{s}^E = (s^E, \ldots, s^E)$. Since s^E is a strict ESS, however, a single deviation is not enough to abandon $\mathbf{s}^E$. It is thus more likely that the process will move from $\mathbf{s}$ to $\mathbf{s}^E$ than the other way around and, thus, states $\mathbf{s}$ with $s > s^E$ cannot be stochastically stable.[7] ∎

Proof of Proposition 3. The proof is analogous to that of Proposition 2 using the function $\widehat{\pi}$ to express payoffs, instead of the function $\widetilde{\pi}$. To show (i) consider strategy

[7] Technically, if we construct $\mathbf{s}$-trees of minimum cost we get that connecting states $\mathbf{s}$ with $s > s^E$ to an $\mathbf{s}^E$-tree can be done at cost 1, but connecting $\mathbf{s}^E$ to any other $\mathbf{s}$-tree would imply cost strictly higher than 1. Thus $\mathbf{s}^E$-trees always have lower cost than $\mathbf{s}$-trees with $s > s^E$ and the latter cannot be stochastically stable (see e.g. Young, 1993, Theorem 2). Alternatively, we can partition the state space into $S^- = \{\mathbf{s} \mid s \leq s^E\}$ and $S^+ = \{\mathbf{s} \mid s > s^E\}$. By Ellison (2000, Theorem 1), $CR(S^-) = 1$ and $R(S^-) > 1$; since $R(S^-) > CR(S^-)$, stochastically stable states must be contained in S^-.

profiles of the form $\mathbf{s} = (s, \overset{n-m}{\ldots}, s, s^N, \overset{m}{\ldots}, s^N)$. Let $y^m = g^{n-1}(s, \overset{n-1-m}{\ldots}, s, s^N, \overset{m}{\ldots}, s^N)$ be the aggregate when $0 \leq m \leq n-1$ opponents choose s^N. We want to show first that, if $\mathbf{s}^N$ is a strict Nash equilibrium, then s^N is always a better reply against an opponents' profile with aggregate y^m than s itself. By definition $\widehat{\pi}(s^N, y^{n-1}) > \widehat{\pi}(s, y^{n-1})$ for all $s \neq s^N$. Let $s > s^N$, then $y^m > y^{n-1}$ for all $m < n-1$ and submodularity implies $\widehat{\pi}(s^N, y^m) > \widehat{\pi}(s, y^m)$. Suppose now by contradiction that there exist $s < s^N$ and $m < n-1$ such that $\widehat{\pi}(s^N, y^m) \leq \widehat{\pi}(s, y^m)$; since $y^m < y^{n-1}$, by submodularity $\widehat{\pi}(s^N, y^{n-1}) \leq \widehat{\pi}(s, y^{n-1})$, contradicting the definition of strict Nash equilibrium. It follows that

$$\widehat{\pi}(s^N, y^m) > \widehat{\pi}(s, y^m)$$

for all $s \neq s^N$ and all $m \in \{0, \ldots, n-1\}$. In particular, for $m = 0$ this implies that no other s can be played in a symmetric Nash equilibrium.

To show (ii) suppose the process starts at $\mathbf{s} = (s, \ldots, s)$ with $s < s^N$ and consider a deviation from s to s^N resulting in state $\mathbf{s}' = (s, \overset{n-1}{\ldots}, s, s^N)$. By (i) above this deviation must result in a payoff improvement for the deviator, $\widehat{\pi}(s^N, y^0) > \widehat{\pi}(s, y^0)$. Rewriting this inequality in terms of $\widetilde{\pi}$, we must have $\widetilde{\pi}(s^N, g^n(s^N, y^0)) > \widetilde{\pi}(s, g^n(s, y^0))$. Moreover, since $\widetilde{\pi}$ is decreasing in the aggregate, we also have that $\widetilde{\pi}(s, g^n(s, y^0)) > \widetilde{\pi}(s, g^n(s^N, y^0))$. These two inequalities together imply that the deviator choosing s^N attains an absolute as well as a relative advantage and will be followed by imitation. Thus, one experiment is enough to move the process from $\mathbf{s}$ to $\mathbf{s}^N$. However, the converse is not true, since a deviation from the strict Nash equilibrium yields strictly lower payoffs. By the same arguments used at the end of the proof of Proposition 2(ii), states $\mathbf{s}$ with $s < s^N$ cannot be stochastically stable. ■

Lemma 3. Consider a generalized aggregative game with compact S. Assume that $\widetilde{\pi}(s, x)$ is upper-hemicontinuous, strictly quasiconcave in s, and submodular. Then, the following exist and are decreasing:

$$z(x) = \arg\max_{s' \in S} \widetilde{\pi}(s', x)$$
$$\overline{z}(s) = z(g^n(s, \ldots, s))$$

Proof. Existence of a maximum is guaranteed by upper-hemicontinuity and compactness of S. Strict quasiconcavity implies that the maximum is also unique and strict. Submodularity of $\widetilde{\pi}$ is used to show the maximum is decreasing. By contradiction, suppose $x' < x''$ and $z(x') < z(x'')$. By definition of $z(x')$ we must have that $\widetilde{\pi}(z(x''), x') < \widetilde{\pi}(z(x'), x')$. Submodularity implies that $\widetilde{\pi}(z(x''), x'') < \widetilde{\pi}(z(x'), x'')$ which contradicts the definition of $z(x'')$. Since g^n is increasing, it follows that $\overline{z}(x)$ is also decreasing. ■

Proof of Proposition 4. Let $s \in S$, $\delta > 0$ with $s \geq \delta$. First, notice that, since h is strictly concave,

$$h(s) = h\left(\frac{1}{2}(s+\delta) + \frac{1}{2}(s-\delta)\right) > \frac{1}{2}h(s+\delta) + \frac{1}{2}h(s-\delta)$$

and thus $h(s-\delta) + h(s+\delta) + (n-2)h(s) < nh(s)$. The conclusion follows from the fact that f is increasing. ■

Proof of Theorem 2. If s^* is an ATS, it is strict by strict quasiconcavity. By Proposition 1, s^* is a strictly globally stable ESS; by Proposition 2(i), s^* is also the unique ESS. Theorem 1 implies that all stochastically stable states of the imitation dynamics for any memory size $k \geq 1$ are symmetric profiles $\mathbf{s} = (s, \ldots, s)$ with strategy s in the interval $[s^N, s^*]$. Further, the proofs of Propositions 2 and 3 imply that symmetric profiles with strategies *out* of this interval are destabilized with a single deviation. In particular, any symmetric profile with $s > s^*$ can be connected to $\mathbf{s}^*$ with a single deviation, and any symmetric profile with $s < s^N$ can be connected to $\mathbf{s}^N$ with a single deviation.

Note also that, by definition of ATS, $s^* = \overline{z}(s^*)$ and recall by Lemma 3 that $z(x)$ and $\overline{z}(s)$ are both decreasing. In particular, this implies that at any $\mathbf{s} = (s, \ldots, s)$ with $s \in [s^N, s^*]$ and denoting $x = g^n(s, \ldots, s)$ we have $z(x) \geq s^*$. Moreover, any deviation, or any set of deviations, from $\mathbf{s}$ resulting in aggregate $x' \leq x$ leads to $z(x') \geq s^*$ and $\widetilde{\pi}(s, x')$ is increasing in the interval $[s^N, s^*]$. This observation immediately implies that single or multiple deviations with lower strategies, which result in lower aggregate values, can never attain a relative advantage and, thus, they cannot be successful.

We will now proceed to show that combinations of two deviations involving at least one higher strategy but a lower value of the aggregate will render the highest strategy as the most successful of the last k rounds. In particular, we will show that any state $\mathbf{s} = (s, \ldots, s)$ with $s^N \leq s \leq s^* - \delta$ can be connected to the state $\mathbf{s}' = (s + \delta, \ldots, s + \delta)$ with two simultaneous deviations to $s - \delta$ and $s + \delta$.

To see this, denote $x = g^n(s, \ldots, s)$ and $y = g^n(s - \delta, s + \delta, s, \ldots, s)$, respectively the aggregate before and after these two deviations. Note that $y \leq x$ since the aggregate is decreasing in variance. Hence, $z(y) \geq z(x) \geq s^* \geq s+\delta$. Thus, $\widetilde{\pi}(\cdot, y)$ is strictly increasing below s^*, implying that

$$\widetilde{\pi}(s - \delta, y) < \widetilde{\pi}(s, y) < \widetilde{\pi}(s + \delta, y).$$

After deviation, strategy $s + \delta$ has a relative advantage. But, since $\widetilde{\pi}$ is decreasing in the aggregate, $\widetilde{\pi}(s, y) \geq \widetilde{\pi}(s, x)$ for all s; it follows that $s + \delta$ also earns larger payoffs than s in previous periods. As a result, $s + \delta$ will be adopted at any revision opportunity.

It may even be possible to make the same kind of argument with a single deviation to a larger strategy. Indeed, it may be possible to disturb $\mathbf{s}$ with a single deviation to some $s' > s$ as long as $y' = g^n(s', s, \ldots, s) > x$ is such that $\widetilde{\pi}(\cdot, y')$ is still increasing in $[s, s']$ (of course, now $z(y') \leq z(x)$).

A $\mathbf{s}^*$-tree of minimum cost can be constructed as follows. Connect any $\mathbf{s}$ with $s < s^N$ to state $\mathbf{s}^N$ with 1 deviation as in Proposition 3. Connect any $\mathbf{s}$ with $s > s^*$ to state $\mathbf{s}^*$ with 1 deviation as in Proposition 2. For any $\mathbf{s}$ with $s \in [s^N, s^* - \delta]$, there are two possibilities. If $\mathbf{s}$ can be destabilized with one mutation, it must be to a larger strategy $s' > s$. Hence, connect $\mathbf{s}$ to $\mathbf{s}' = (s', \ldots, s')$, with either $s' = s^*$, or still $s' \in [s^N, s^* - \delta]$. Otherwise, connect $\mathbf{s}$ to $\mathbf{s}' = (s + \delta, \ldots, s + \delta)$ with two deviations. Since all transitions can be made at cost at most 2 and at least two deviations are needed to exit s^*, this is a minimal-cost tree, which proves the Theorem. ■

Proof of Theorem 3. To complete the proof we only need to show existence of s^0 and $\widehat{s}$. Denote $x^* = g^n(s^*, \ldots, s^*)$ and $y(s) = g^n(s^N, s^*, \ldots, s^*, s)$. Let $s^1 \in \mathbb{R}$ be such that $\widetilde{\pi}(s^N, x^*) = \widetilde{\pi}(s^1, x^*)$. By strict concavity of $\widetilde{\pi}$ we have that s^1 exists and is strictly greater than s^*. Recall $s^* = z(x^*)$ and, thus, $\widetilde{\pi}(s^*, x^*) > \widetilde{\pi}(s^1, x^*)$. Let $s^2 \in \mathbb{R}$ be such that $y(s^2) = x^*$. By continuity of g^n, s^2 exists and, by strong monotonicity, it is unique and strictly greater than s^*. Note that either $x^* \leq y(s^1)$ and then $s^2 \leq s^1$, or $x^* > y(s^1)$ and then $s^2 > s^1 > s^*$ and strict concavity implies that $\widetilde{\pi}(s^2, x^*) < \widetilde{\pi}(s^1, x^*) = \widetilde{\pi}(s^N, x^*)$. Define $\widehat{s} = \max\{s^1, s^2\}$, or take the next larger strategy in S, if $\max\{s^1, s^2\}$ is not a point in the grid. By submodularity, since $y(\widehat{s}) \geq x^*$, we must have

$$\widetilde{\pi}(\widehat{s}, y(\widehat{s})) \leq \widetilde{\pi}(s^N, y(\widehat{s})).$$

Moreover, since $y(\widehat{s}) \geq x^*$, by Lemma 3, $z(y(\widehat{s})) \leq s^*$. Since $\widehat{s} > s^*$, strict concavity of $\widetilde{\pi}$ implies that

$$\widetilde{\pi}(\widehat{s}, y(\widehat{s})) < \widetilde{\pi}(s^*, y(\widehat{s})).$$

In particular, this implies that, if the process starts at the ATS state $\mathbf{s}^*$, two deviations to s^N and $\widehat{s}$ respectively would not be enough to move the process to state $\widehat{\mathbf{s}}$.

Suppose, however, the process starts at the symmetric profile $\widehat{\mathbf{s}} = (\widehat{s}, \ldots, \widehat{s})$, with $\widehat{s}$ as constructed above. Then a single deviation to s^N would result in a state of the type $(s^N, \widehat{s}, \overset{n-1}{\ldots}, \widehat{s})$ with aggregate $\widehat{y} := g^n(s^N, \widehat{s}, \overset{n-1}{\ldots}, \widehat{s}) > y(\widehat{s})$. Submodularity then implies

$$\widetilde{\pi}(\widehat{s}, \widehat{y}) \leq \widetilde{\pi}(s^N, \widehat{y})$$

and the deviation would be followed by imitation. I.e. it is possible to reach state $\mathbf{s}^N$ from state $\widehat{\mathbf{s}}$ with a single deviation.

To show existence of s^0, notice that there exists x^0 such that $\widehat{s} = \arg\max_s \widetilde{\pi}(s, x^0)$. To see this note that, by Lemma 3, $z(x)$ is unique and thus strictly decreasing. Set for example s^0 to be such that $x(s^0) = x^0$. ■

References

Alós-Ferrer, C.: Cournot versus Walras in dynamic oligopolies with memory. *Int. J. Ind. Organ.* **22**, 193-217 (2004).

Alós-Ferrer, C., Ania, A.B.: The evolutionary stability of perfectly competitive behavior. *Econ. Theory* **26**, 497-516 (2005).

Ania, A.B.: Evolutionary stability and Nash equilibrium in finite populations, with an application to price competition. *J. Econ. Behav. Organ.* **65**, 472-488 (2008).

Ellison, G.: Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *Rev. Econ. Stud.* **67**, 17-45 (2000).

Foster, D., Young, P.: Stochastic evolutionary game dynamics. *Theor. Popul. Biol.* **38**, 219-232 (1990).

Guse, T., Hehenkamp, B., Possajennikov, A.: On the equivalence of Nash and evolutionary equilibrium in finite populations. CeDEx Discussion Paper No. 200806, University of Nottingham (2008).

Hamilton, W.D.: Selection of selfish and altruistic behavior in some extreme models. In: Eisenberg, J.F., Dillon, W.S. (eds.) *Man and Beast: Comparative Social Behavior.* Smithsonian Institution, Washington D.C. (1971).

Kandori, M., Mailath, G.J., Rob, R.: Learning, mutation, and long-run equilibria in games. *Econometrica* **61**, 29-56 (1993).

Leininger, W.: Fending off one means fending off all: evolutionary stability in quasi-submodular aggregative games. *Econ. Theory* **29**, 713-719 (2006).

Maynard Smith, J.: The theory of games and the evolution of animal conflict. *J. theor. Biol.* **47**, 209-221 (1974).

Maynard Smith, J.: *Evolution and the theory of games.* Cambridge University Press, Cambridge (1982).

Maynard Smith, J., Price, G.R.: The logic of animal conflict. *Nature* **246**, 15-18 (1973).

Neill, D.B.: Evolutionary stability for large populations. *J. theor. Biol.* **227**, 397-401 (2004).

Possajennikov, A.: Evolutionary foundations of aggregate-taking behavior. *Econ. Theory* **21**, 921-928 (2003).

Schaffer, M.: Evolutionarily stable strategies for a finite population and a variable contest size. *J. theor. Biol.* **132**, 469-478 (1988).

Schaffer, M.: Are profit-maximizers the best survivors? A Darwinian model of economic natural selection. *J. Econ. Behav. Organ.* **12**, 29-45 (1989).

Schipper, B.: Submodularity and the evolution of Walrasian behavior. *Int. J. Game Theory* **32**, 471-477 (2004).

Tanaka, Y.: Stochastically stable states in an oligopoly with differentiated goods: equivalence of price and quantity strategies. *J. Math. Econ.* **34**, 235-253 (2000).

Topkis, D.M.: Supermodularity and complementariy. Princeton: Princeton University Press (1998).

Vega-Redondo, F.: *Evolution, Games, and Economic Behaviour.* Oxford University Press, Oxford (1996).

Young, P.: The evolution of conventions. *Econometrica* **61**, 57-84 (1993).