N-person games where imitation always hits a better reply

Ana B. Ania*

Department of Economics, University of Vienna, Hohenstaufeng. 9, A1010 Vienna, Austria.
Phone: +43-1-4277 37433 Fax: +43-1-4277 9374.
E-mail: ana-begona.ania-martinez@univie.ac.at

This paper identifies some symmetric n-person games where the behavioral rule imitate the best always hits a better reply. In particular, that is the case in Bertrand games with homogeneous product, in minimum-effort games, and in games with positive network externalities.*Journal of Economic Literature* Classification Numbers: C72, D83.

Key Words: learning, imitation, Bertrand oligopoly, network externalities, minimum-effort games.

1. INTRODUCTION

In recent years there has been a growing interest on learning rules based on imitation. Many papers in the evolutionary literature postulate imitative behavior as for example Alós-Ferrer *et al.*[1], Eshel *et al.*[4], and Vega-Redondo [11] to name only some. Imitation is also one of the possible underlying behavioral principles that yield the well-established replicator dynamics' equations (see e.g. Weibull [12]).

The experimental literature provides some account for the use of imitation in economic decision making. Under specific information structures, namely when the actions chosen by others as well as the success of those actions (payoffs in a game) are observed, imitation appears to be widely used in complex situations. Learning by imitation can be regarded as a way of saving decision-making costs. Even when it may lead to suboptimal decisions, agents seem to use it because it requires very restricted information and cognitive resources. See e.g. Pingle and Day [8] for a discussion and experimental evidence of imitation as a mode of economizing behavior. Huck *et al.*[5] also provide experimental evidence for the use of imitation in the context of an oligopolistic market.

A theoretical justification for the use of imitation is provided by Schlag [10], who characterizes a class of imitative rules – proportional imitation – which have certain optimality properties in *any* possible multi-armed bandit problem, where agents have to choose among actions with uncertain bounded payoffs. These results, though, cannot be extended to strategic situations. Ania [2] shows that there are no nontrivial behavioral rules that have Schlag's optimality properties for every possible game.

The present paper focuses on a different property of imitative rules in the framework of n-player symmetric games. In particular, some classes of games are identified, where a certain imitation rule – *imitate the best* – always selects a better reply, namely Bertrand games with homogeneous product, minimum-effort games, and games with positive network externalities. In this sense, we take a first step in the direction of identifying a more general class of games, where imitate the best always hits a better reply.

The paper is organized as follows. In Section 2 we introduce the notation and basic definitions. Sections 3, 4, and 5 consider respectively Bertrand games, minimum-effort games, and games with network externalities. Section 6 comments on the relation between imitation and better-reply dynamics. Section 7 concludes.

2. BEHAVIORAL RULES: IMITATE THE BEST

Consider the symmetric game $\Gamma = \langle N, S, \pi \rangle$, where $N = \{1, \ldots, n\}$ is the set of players, S is the set of strategies for each player, and $\pi : S^n \to \mathbb{R}^n$ is the symmetric payoff function that associates to each strategy profile a payoff for each player. By symmetric we mean that $\pi_i(s_1, \ldots, s_n) = \pi_{\sigma(i)}(s_{\sigma(1)}, \ldots, s_{\sigma(n)})$ for all $\sigma \in \Sigma_n$, and all $i \in N$, where Σ_n denotes the symmetric group of permutations of n elements.

A *behavioral rule* $F : S^n \times \mathbb{R}^n \to [\Delta(S)]^n$ maps, for each player, strategies and payoffs into the set of probability distributions over strategies. Given any strategy profile $\mathbf{s} = (s_1, \ldots, s_n)$ and any vector of payoffs $\tilde{\pi} = (\pi_1, \ldots, \pi_n)$, $F^i_s(\mathbf{s}, \tilde{\pi})$ is the probability that player i chooses strategy $s \in S$ after observing profile $\mathbf{s}$ and payoffs $\tilde{\pi}$. Obviously, for each player i, $\sum_{s \in S} F^i_s(\mathbf{s}, \tilde{\pi}) = 1$.

Denote the *carrier* of a strategy profile $\mathbf{s}$ by

$$C(\mathbf{s}) := \{s \in S \ : \ s = s_i, \ i \in N\}.$$

We say that a behavioral rule F is *imitating* if $F^i_s(\mathbf{s}, \tilde{\pi}) = 0$ for all $s \notin C(\mathbf{s})$ and all $i \in N$; i. e. if it never prescribes to choose any strategy not observed in $\mathbf{s}$.

Given any $\mathbf{s}$, let

$$I(\mathbf{s}) := \{s \in S \ : \ s = s_i \text{ for some } i \in N \text{ and } \pi_i(\mathbf{s}) \geq \pi_k(\mathbf{s}) \ \forall k \in N\}.$$

The set $I(\mathbf{s}) \subseteq C(\mathbf{s})$ contains the strategies that gave highest payoffs in the profile $\mathbf{s}$. A rule F corresponds to *imitate the best* if, for any strategy profile $\mathbf{s}$, (symmetric) payoff function π, and player i, $F^i_s(\mathbf{s}, \pi(\mathbf{s})) = 0$ for all $s \notin I(\mathbf{s})$. Note this defines a class of behavioral rules. If $I(\mathbf{s})$ is not a singleton, the only constraint placed on the distribution $F^i(\mathbf{s}, \pi(\mathbf{s}))$ is that its support must be contained in $I(\mathbf{s})$.

Denote by $(\mathbf{s} \setminus s_i; s)$ the strategy profile where all players but i choose strategies according to $\mathbf{s}$ and player i chooses $s \in S$. Given $\mathbf{s}$, let

$$B_i(\mathbf{s}) := \{s \in S \ : \ \pi_i(\mathbf{s} \setminus s_i; s) \geq \pi_i(\mathbf{s})\}$$

that is, the set of *better reply* strategies for player i. Of course, $s_i \in B_i(\mathbf{s})$ for all $\mathbf{s}$.

We say that a behavioral rule F *always hits a better reply* in game Γ if for all $\mathbf{s} \in S^n$, $i \in N$, and $s \in S$ such that $F^i_s(\mathbf{s}, \pi(\mathbf{s})) > 0$, we have that $s \in B_i(\mathbf{s})$. That is, starting at any $\mathbf{s}$, if only one player revises her strategy at a time, following F will (weakly) improve the revising player's payoffs.

The following definitions extend the optimality properties defined by Schlag [10] to the case in which each player's behavior depends on the strategies and payoffs of all players.

Suppose that an external agent were in the position to enter at any $\mathbf{s}$ to randomly, uniformly replace any player $i \in N$. The entering player's expected payoff of following rule F when entering at $\mathbf{s}$ is given by

$$\pi(F, \mathbf{s}) = \frac{1}{n} \sum_{i=1}^{n} \sum_{s \in S} F^i_s(\mathbf{s}, \pi(\mathbf{s})) \cdot \pi_i(\mathbf{s} \setminus s_i; s) \tag{1}$$

Let $\overline{\pi}(\mathbf{s})$ denote the average payoff at $\mathbf{s}$, which would correspond to the entering player's expected payoff if she did not change strategy with respect to her predecessor's, that is if she followed a *never switch* rule. A behavioral rule F is called *improving* if $EIP_F(\mathbf{s}, \pi) := \pi(F, \mathbf{s}) - \overline{\pi}(\mathbf{s}) \geq 0$ for all $\mathbf{s} \in S^n$, and all (symmetric) payoff functions $\pi : S^n \mapsto \mathbb{R}^n$, where EIP_F stands for the expected improvement of rule F. The rationale behind this concept is that it would be myopically rational for the entering player to adopt an improving rule.

Remark 2. 1. If a behavioral rule F always hits a better reply in a game $\Gamma = \langle N, S, \pi \rangle$, then F is improving in Γ, i. e. $EIP_F(\mathbf{s}, \pi) \geq 0$ for all $\mathbf{s} \in S^n$. The fact that F always hits a better reply in Γ implies that for all $i \in N$ and all $s \in S$ such that $F^i_s(\mathbf{s}, \pi(\mathbf{s})) > 0$, we have that $\pi_i(\mathbf{s} \setminus s_i; s) \geq \pi_i(\mathbf{s})$. Therefore, using (1) we have

$$EIP_F(\mathbf{s}, \pi) = \frac{1}{n} \sum_{i=1}^{n} \left(\sum_{s \in S} F^i_s(\mathbf{s}, \pi(\mathbf{s})) \cdot \pi_i(\mathbf{s} \setminus s_i; s) - \pi_i(\mathbf{s}) \right) \geq 0$$

In the following sections we will provide examples of games where *imitate the best* always hits a better reply; that is, for all $i \in N$ and all $\mathbf{s} \in S^n$, $I(\mathbf{s}) \subseteq B_i(\mathbf{s})$. Remark 2.1 then shows that in those games imitate the best will also be improving.

Suppose now that all players would follow rule F simultaneously. For any $\mathbf{s}$ let $\overline{\pi}(F, \mathbf{s})$ be the corresponding expected average payoff.

$$\overline{\pi}(F, \mathbf{s}) = \frac{1}{n} \sum_{\mathbf{s}' \in S^n} \left[\left(\prod_{i=1}^{n} F^i_{s'_i}(\mathbf{s}, \pi(\mathbf{s})) \right) \cdot \sum_{i=1}^{n} \pi_i(\mathbf{s}') \right] \tag{2}$$

A behavioral rule F is called *payoff increasing* in game Γ if $\overline{\pi}(F, \mathbf{s}) - \overline{\pi}(\mathbf{s}) \geq 0$ for all $\mathbf{s} \in S^n$, i. e. if it always increases the expected average payoff from any strategy profile. The rationale behind this property is that a benevolent planner would prescribe a payoff increasing rule to all players.[1]

3. BERTRAND GAMES WITH HOMOGENEOUS PRODUCT

In this section we provide a first example of a class of n-person games where imitate the best always hits a better reply, namely the class of Bertrand games where all firms face the same decreasing demand function and increasing and convex cost function. Consider an industry where identical firms $N = \{1, \ldots, n\}$ set prices. All firms face identical cost function $C(q)$ that depends on the produced quantity q and is increasing and convex, with $C(0) = 0$ for simplicity. Suppose customers buy only from the firm with minimum price. Let $D(p)$ be a positive, decreasing demand function. In case of ties, demand splits equally. Given $\mathbf{p} = (p_1, \ldots, p_n)$, call $p = \min\{p_1, \ldots, p_n\}$. The profits to firm i are given by

$$\pi_i(\mathbf{p}) = \begin{cases} p_i \frac{D(p_i)}{m} - C(\frac{D(p_i)}{m}) & \text{if} \quad p_i = p \\ 0 & \text{if} \quad p_i \neq p \end{cases}$$

where $m = |\{i \in N | p_i = p\}|$ is the number of firms with minimum price p.

In this class of games, imitate the best prescribes to mimic with probability one the price charged by the firm with highest profits. If several firms obtain the same maximum profits, then one of them will be imitated according to a pre-specified probability distribution (not necessarily with full support).

The next result proves that imitate the best always hits a better reply in symmetric Bertrand games.

PROPOSITION 3.1. *In a symmetric Bertrand game with decreasing demand, increasing and convex costs, and equal splitting in case of ties,* $I(\mathbf{p}) \subseteq B_i(\mathbf{p})$ *for all* $\mathbf{p} \in \mathbb{R}^n$ *and all* $i \in N$.

Proof. For all strategy profiles $\mathbf{p}(n) = (p, p, \overset{n}{\ldots}, p)$, all firms share the market and obtain the same profits. By following imitate the best, none of them will

[1] Schlag [10] shows that in games against nature a rule is improving if and only if it is payoff increasing in all possible games. Ania [2] shows that these two properties are not equivalent in n-person games.

change price and profits will not change after any price revision; $p \in B_i(\mathbf{p}(n))$ for all $i \in N$.

Consider states of the form $\mathbf{p}(m) = (p_1, p_2, \ldots, p_n)$ where w.l.o.g. $p_1 = p_2 = \cdots = p_m = p$ with $1 \leq m < n$ and $p < p_j,\ j = m+1, \ldots, n$. The profits of the firms with higher than minimum price is $\pi_j(\mathbf{p}(m)) = 0$ while those of the firms with minimum price are given by

$$\pi_i(\mathbf{p}(m)) = p\frac{D(p)}{m} - C\left(\frac{D(p)}{m}\right) = \frac{D(p)}{m}\left[p - AC\left(\frac{D(p)}{m}\right)\right],\ i = 1, \ldots, m$$

where $AC(q) = C(q)/q$ denotes average cost and the last equality is understood to hold only if $D(p) > 0$. To show that from all these strategy profiles imitate the best always hits a better reply, we have to consider the following cases separately.

First, if $\pi_i(\mathbf{p}(m)) > 0$, then $I(\mathbf{p}(m)) = \{p\}$, since $\pi_j(\mathbf{p}(m)) = 0$. Notice $AC(q)$ is increasing in q.[2] Thus after revision, for any $j = m+1, \ldots, n$

$$\begin{aligned}\pi_j(\mathbf{p}(m+1)) = \frac{D(p)}{m+1}\left[p - AC\left(\frac{D(p)}{m+1}\right)\right] &\geq \\ \frac{D(p)}{m+1}\left[p - AC\left(\frac{D(p)}{m}\right)\right] &> 0\end{aligned} \tag{3}$$

This implies that $p \in B_i(\mathbf{p}(m))$ for all $i \in N$.

Second, if $\pi_i(\mathbf{p}(m)) < 0$, then $I(\mathbf{p}(m)) = \{p\ :\ p = p_j,\ j = m+1, \ldots, n\}$, since highest profits are $\pi_j(\mathbf{p}(m)) = 0$. Obviously, if $m > 1$, then $I(\mathbf{p}(m)) \subset B_i(\mathbf{p}(m))$ for all $i \in N$. As long as there is more than one firm setting minimum price p, any unilateral deviation to a price in $I(\mathbf{p}(m))$ will yield the deviating firm zero profits. For firms $i = 1, \ldots, m$ this allows to avoid losses, while for firms $j = m+1, \ldots, n$ this does not change profits. Consider now the case $m = 1$ with $\pi_1(\mathbf{p}(1)) = D(p)[p - AC(D(p))] < 0$. Call $p' = \min\{p_2, \ldots, p_n\}$. If the firm setting p imitates any $p_j \in I(\mathbf{p}(1)),\ p_j \neq p'$, it will face no demand and avoid losses. If it imitates p', its profits after imitation will be

$$\frac{D(p')}{m'}\left[p' - AC\left(\frac{D(p')}{m'}\right)\right] \tag{4}$$

[2] Consider $q' \geq q > 0$, by convexity of $C(q)$,

$$C(q) \leq \frac{q}{q'}C(q') + \left(1 - \frac{q}{q'}\right)C(0) = \frac{q}{q'}C(q')$$

Thus, $AC(q') \geq AC(q)$. If $C(0) > 0$ and costs are convex, then $AC(q)$ is U-shaped, but average variable costs are increasing. Actually, this is all we need here, although for simplicity we assumed zero fixed costs.

where $m' \geq 2$ is the number of firms setting price p' after imitation. If the expression in square brackets in equation (4) is positive, then after imitation profits instead of losses are achieved. If it is still negative, note that, since $p' > p$, the demand faced after imitation is smaller ($D(p')/m' \leq D(p') \leq D(p)$), thus also $AC(D(p')/m') \leq AC(D(p))$; that is, less is sold to a lower loss per unit which yields lower total losses. In any case, profits of firm 1 after imitation will be higher than $\pi_1(\mathbf{p}(1))$. For firms $j = 2, \ldots, n$ imitation will not change profits. Thus $I(\mathbf{p}(1)) \subseteq B_i(\mathbf{p}(1))$ for all $i \in N$.

Last, if $\pi_i(\mathbf{p}(m)) = 0$, then $I(\mathbf{p}(m)) = \{p \,:\, p = p_i,\ i = 1, \ldots, n\} = C(\mathbf{p})$, since all firms active or idle make zero profits. That is, at these strategy profiles, imitate the best prescribes to imitate any of the observed prices. Again we distinguish the cases $m > 1$ and $m = 1$. Suppose $m > 1$, then for any $i = 1, \ldots, m$ imitate the best will not change profits, and for any $j = m+1, \ldots, n$ that imitates p the new profits will be as in equation (3) positive. Notice that, if $D(p) = 0$, $\pi_j(\mathbf{p}(m+1)) = \pi_j(\mathbf{p}(m)) = 0$. Suppose $m = 1$, then for all $j = m+1, \ldots, n$ everything is analogous to the case $m > 1$, and for firm 1 everything is analogous to the case $\pi_1(\mathbf{p}(1)) < 0$ considered above. ▌

Remark 3. 1. By Proposition 3.1, all imitate-the-best rules always hit a better reply in Bertrand games with homogeneous product. It follows from Remark 2.1 that these rules are improving in such games. We conclude this section by showing that not all imitate-the-best rules are payoff increasing in that class of games. This demonstrates that hitting a better reply is not sufficient for a behavioral rule to be payoff increasing. We should stress the fact, though, that there are some imitate-the-best rules which are indeed payoff increasing in Bertrand games.

Following the notation introduced previously in this section, at any strategy profile $\mathbf{p}(m)$,

$$\overline{\pi}(\mathbf{p}(m)) = \frac{m}{n}\left[p\frac{D(p)}{m} - C\left(\frac{D(p)}{m}\right)\right] = \frac{D(p)}{n}\left[p - AC\left(\frac{D(p)}{m}\right)\right]$$

Note first that, if $\pi_i(\mathbf{p}(m)) > 0$, then $I(\mathbf{p}(m)) = \{p\}$. Thus, if all firms follow imitate the best simultaneously, since $AC(\cdot)$ is increasing we have

$$\overline{\pi}(F, \mathbf{p}(m)) - \overline{\pi}(\mathbf{p}(m)) = \frac{D(p)}{n}\left[AC\left(\frac{D(p)}{m}\right) - AC\left(\frac{D(p)}{n}\right)\right] \geq 0$$

That is, from any strategy profile where all firms in the industry obtain profits, imitate the best will increase average expected payoff in the industry. Therefore, if all payoffs were positive, the class of rules would be payoff increasing.

To see that for some imitate-the-best rules average expected payoff may decrease consider the following minimal example. Take $N = \{1, \ldots, 4\}$ and consider the strategy profile $\mathbf{p}(2) = (p_1, \ldots, p_4)$ with $p = p_1 = p_2 < p_3 = p + \epsilon < p_4$ for small $\epsilon > 0$. Assume the demand and cost functions are such that[3]

$$\pi_i(\mathbf{p}(2)) = \frac{D(p)}{2}\left[p - AC\left(\frac{D(p)}{2}\right)\right] < 0; \qquad i = 1, 2$$

Then $\overline{\pi}(\mathbf{p}(2)) = \frac{1}{2}\pi_i(\mathbf{p}(2)) < 0$

Now consider the extreme imitating rule that prescribes firms with maximum profits not to change price and firms with lower than maximum profits to mimic the highest price observed. In our example, this would imply that firms 1 and 2 should imitate p_4 and firms 3 and 4 would not change price. The resulting strategy profile after imitation would be $\mathbf{p_3}(1) = (p_4, p_4, p_3, p_4)$ with probability one. The profit of the firm setting minimum price after imitation will be given by

$$\pi_3(\mathbf{p_3}(1)) = D(p+\epsilon)\left[p + \epsilon - AC\left(D(p+\epsilon)\right)\right].$$

For $\epsilon > 0$ small enough,

$$\overline{\pi}(F, \mathbf{p}(2)) - \overline{\pi}(\mathbf{p}(2)) = \frac{1}{4}\pi_3(\mathbf{p_3}(1)) - \frac{1}{2}\pi_i(\mathbf{p}(2)) < 0$$

This imitating rule prescribes to mimic some of the prices that gave maximum profits with probability zero. There are, however, imitating rules with full support over the set of prices that gave maximum payoffs, in the example $I(\mathbf{p}(2)) = \{p_3, p_4\}$, such that the probability of ending up in the strategy profile $\mathbf{p_3}(1)$ after profile $\mathbf{p}(2)$ is close to one. For such rules average expected payoff of the industry will also decrease in that case.

Before we conclude it is worthwhile mentioning that there are indeed imitate-the-best rules that are payoff increasing in Bertrand games with homogeneous product. E. g. it is easy to see that this is the case for the rule that prescribes to mimic with probability one the lowest price among those that gave maximum payoffs.

4. MINIMUM-EFFORT GAMES

Consider the game where each player in $N = \{1, \ldots, n\}$ chooses an effort level $e_i \in \mathbb{R}_+$. Player i's payoff is given by

$$\pi_i(e_1, \ldots, e_n) = a \cdot \min\{e_1, \ldots, e_n\} - b \cdot e_i + c,$$

where $a,\ b$, and c are constants with $a > b \geq 0$. This is called a minimum-effort coordination game or a Stag Hunt game (see e. g. Crawford [3]).

[3]Take for example $D(p) = 12 - 2p,\ C(q) = q^2$, and $\mathbf{p}(2) = (2, 2, 2.5, 3)$.

PROPOSITION 4.1. *In a minimum-effort game, imitate the best always hits a better reply.*

Proof. Given $\mathbf{e} = (e_1, \ldots, e_n)$, $I(\mathbf{e}) = \arg\min\{e_1, \ldots, e_n\}$. Take any j with $e_j \notin I(\mathbf{e})$. If j imitates, it will adopt the only effort level in $I(\mathbf{e})$ and obtain payoff

$$(a-b)\cdot\min\{e_1,\ldots,e_n\} + c \geq a\cdot\min\{e_1,\ldots,e_n\} - b\cdot e_j + c$$

∎

Remark 4. 1. Remark 2.1 implies that imitate the best is improving in minimum-effort games. Note that in this case imitate the best is also payoff increasing. For every time a player imitates the minimum effort level, her payoff increases and those of all other players do not change, increasing average payoff.

5. GAMES WITH POSITIVE NETWORK EXTERNALITIES

In the present section we consider a subclass of the so-called congestion games, first defined by Rosenthal [9]. Congestion games model situations in which players have to decide on a set of facilities. These may refer for example to a selection of primary factors in a production process, to a route of roads to travel between two points, or to a software package. The feature that characterizes congestion games is that the payoffs to each player depend on the number of players choosing each facility. This dependence is negative if the cost of using a facility increases with the number of users, and hence the name congestion games. If, on the other hand, facilities become more attractive (respectively less costly) when more users choose them, payoffs increase with the number of players choosing the same facilities, and the game displays *positive network externalities*. We now recall the definition of a congestion game in general.

Consider the set of players $N = \{1, 2, \ldots, n\}$ that choose among a set of facilities $M = \{1, 2, \ldots, m\}$. Each player's strategy set is given by $S \subseteq \mathcal{P}(M)$.[4] A strategy of player i will be any $\mathtt{s}_i \subseteq S$, so players choose collections of facilities. E. g. if facilities model roads, a strategy represents a route of roads; if facilities model software (email, browser, word processor, etc.), a strategy represents a software package. Given a strategy profile $\mathbf{s} = \{\mathtt{s}_1, \mathtt{s}_2, \ldots, \mathtt{s}_n\}$, call $x_j(\mathbf{s}) = |\{i \in N : \ j \in \mathtt{s}_i\}|$ for $j \in M$; i. e. $x_j(\mathbf{s})$ is the number of players that make use of facility j at $\mathbf{s}$. Let $c_j(x_j(\mathbf{s}))$ denote the cost to each player of using facility j, which depends on the number of players that use j simultaneously.

[4]In principle, each player could have a different strategy set $S^i \subseteq \mathcal{P}(M)$. Here we concentrate on symmetric congestion games where all players have access to all subsets of facilities.

Given $\mathbf{s}$, the payoff to each player i is

$$\pi_i(\mathbf{s}) = -\sum_{j \in \mathtt{S}_i} c_j(x_j(\mathbf{s})).$$

We refer to a congestion game with decreasing cost functions c_j as a game with *positive network externalities*. The choice of a software package or communication system would fall into this category. Actually in this case the name congestion game does not seem appropriate. The next proposition shows that in a game with positive network externalities of this kind imitation always hits a better reply.

PROPOSITION 5.1. *For every symmetric game $\Gamma = (N, M, S, \pi)$ with positive network externalities, imitate the best always hits a better reply.*

Proof. Recall that, given any strategy profile $\mathbf{s}$,

$$I(\mathbf{s}) = \{\mathtt{S} \in S \ : \ \mathtt{S} = \mathtt{S}_i \text{ for some } i \in N \text{ and } \pi_i(\mathbf{s}) \geq \pi_k(\mathbf{s}) \, \forall k \in N\}.$$

Starting at $\mathbf{s}$, suppose any player, say l, updates her strategy following imitate the best. Call $\mathbf{s}' = (\mathbf{s} \setminus \mathtt{S}_l; \mathtt{S_i})$ the resulting strategy profile with $\mathtt{S}_i \in I(\mathbf{s})$. We have

$$\begin{aligned} x_j(\mathbf{s}') &= \quad x_j(\mathbf{s}) \quad\ \ \forall\, j \in (\mathtt{S}_i \cap \mathtt{S}_l) \\ x_j(\mathbf{s}') &= \ x_j(\mathbf{s}) + 1 \ \ \forall\, j \in (\mathtt{S}_i \setminus \mathtt{S}_l) \end{aligned}$$

Since $\mathtt{S}_i \in I(\mathbf{s})$ and c_j is decreasing, this implies that

$$\pi_l(\mathbf{s}') = -\sum_{j \in \mathtt{S}_i} c_j(x_j(\mathbf{s}')) \geq -\sum_{j \in \mathtt{S}_i} c_j(x_j(\mathbf{s})) \geq -\sum_{j \in \mathtt{S}_l} c_j(x_j(\mathbf{s})) = \pi_l(\mathbf{s}).$$

Therefore, $\mathtt{S}_i \in B_l(\mathbf{s})$. ▮

Remark 5. 1. Remark 2.1 implies that imitate the best is improving in games with positive network externalities. Note that in this case imitate the best is also payoff increasing. For every time a player imitates a set of facilities that give maximum payoffs, it decreases the cost of using those facilities for every player, thus increasing the payoff of all users.

6. FINAL REMARKS

In the games discussed above, if only one agent revises her action at a time, following imitate the best results in single-player (weak) improvement paths, i. e.

paths where the only revising agent achieves a (weak) payoff improvement after revision. These paths are a subset, or particular case, of those that would result from a better-reply dynamics defined as follows. Each period exactly one agent is selected at random and given revision opportunity. The agent then selects one action out of her set of better replies according to a probability distribution with full support.

Monderer and Shapley [7] show that minimum-effort games, and games with positive network externalities of the type discussed in this paper are potential games. They also show that finite potential games fulfill the property that every (strict) improvement path is finite – called the *finite improvement property* (FIP). In finite games, the FIP implies that the one-sided better reply dynamics mentioned above converges to the set of Nash equilibria. [5]

Consider now an imitation dynamics where each period one agent is selected at random to revise her strategy according to imitate the best with full support. For minimum-effort games and games with positive network externalities and except for degenerate payoff functions,[6] one can easily see that from any non Nash state there exist nontrivial strict improvement paths of positive probability under the imitation dynamics. This implies that the imitation dynamics converges to Nash equilibria. [7] However, these are very particular games where all states where all agents choose the same action – monomorphic states – correspond to Nash equilibria. If this were not the case, e. g. in the Bertrand game, it is obvious that imitation cannot converge to Nash equilibrium in general, since all monomorphic states are absorbing.

7. CONCLUSIONS

Behavioral rules based on imitation have been increasingly postulated in learning and evolutionary economic models in later years. The results obtained with such a stylized behavioral principle as imitation can be of great simplicity and at the same time surprisingly powerful. The incentives to use imitation in complex, strategic frameworks remain unexplored, although in experimental work it has been shown that imitation is frequently used by subjects who are provided with the adequate information.

In the present work we have focused on the ability of imitation to find better replies to the current strategy profile. The motivation to explore such a property

[5] Friedman and Mezzetti [6] show that the better-reply dynamics converges to pure-strategy Nash equilibrium in finite games satisfying the property that from any action profile there exists a finite sequence of single-player improvements leading to a Nash equilibrium – called the *weak finite improvement property* (FIP). Such games include games with the FIP.

[6] Specifically $b > 0$ for minimum-effort games and c_j strictly decreasing for games with positive network externalities.

[7] Since imitation reduces the number of observed strategies, it is easy to argue that convergence does not depend on the games being finite.

is that we would expect imitation to persist in strategic situations where its use always allows to improve the player's payoff.

The paper is only a first step into a general characterization of the class of games with that property. We have seen that imitate the best always hits a better reply in Bertrand games with homogeneous product, in minimum-effort games, and in games with positive network externalities.

REFERENCES

1. C. Alós-Ferrer, A.B. Ania, and K.R. Schenk-Hoppé, An Evolutionary Model of Bertrand Oligopoly, *Games Econ. Behav.* **33** (2000), 1-19.
2. A.B. Ania, Learning by Imitation when Playing The Field, Working Paper Nr. 0005, Department of Economics, University of Vienna (2000).
3. V.P. Crawford, An Evolutionary Interpretation of Van Huyck, Battalio, and Beil's Experimental Results on Coordination, *Games Econ. Behav.* **3** (1991), 25-59.
4. I. Eshel, L. Samuelson, and A. Shaked, Altruist, Egoist, and Hooligans in a Local Interaction Model, *Amer. Econ. Rev.* **88** (1998), 157-179.
5. S. Huck, H.-T. Norman, and J. Oechssler, Learning in Cournot oligopoly – an experiment, *Econ. J.* **109** (1999), C80-C95.
6. J.W. Friedman and C. Mezzetti, Learning in Games by Random Sampling, *J. Econ. Theory* **98** (2001), 55-84.
7. D. Monderer and L.S. Shapley, Potential Games, *Games Econ. Behav.* **14** (1996), 124-143.
8. M. Pingle and R.H. Day, Modes of economizing behavior: Experimental evidence, *J. Econ. Behav. Organ.* **29** (1996), 191-209.
9. R.W. Rosenthal, A Class of Games Possessing Pure-Strategy Nash Equilibria, *Int. J. Game Theory* **2** (1973), 65-67.
10. K. Schlag, Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits, *J. Econ. Theory* **78** (1998), 130-156.
11. F. Vega-Redondo, The Evolution of Walrasian Behavior, *Econometrica* **65** (1997), 375-384.
12. J.W. Weibull, *Evolutionary Game Theory,* Cambridge: MIT Press (1995).