



DEPARTMENT OF ECONOMICS

Working Paper Series

When is Tit-For-Tat unbeatable?

Peter Duersch
University of Heidelberg
Joerg Oechssler
University of Heidelberg
Burkhard Schipper
University of California, Davis

January 22, 2013

Paper # 13-1

We characterize the class of symmetric two-player games in which tit-for-tat cannot be beaten even by very sophisticated opponents in a repeated game. It turns out to be the class of exact potential games. More generally, there is a class of simple imitation rules that includes tit-for-tat but also imitate-the-best and imitate-if-better. Every decision rule in this class is essentially unbeatable in exact potential games. Our results apply to many interesting games including all symmetric 2x2 games, and standard examples of Cournot duopoly, price competition, public goods games, common pool resource games, and minimum effort coordination games.

Department of Economics
One Shields Avenue
Davis, CA 95616
(530)752-0741

http://www.econ.ucdavis.edu/working_search.cfm

When is tit-for-tat unbeatable?*

Peter Duersch[†] Jörg Oechssler[‡] Burkhard C. Schipper[§]
– revised version –

January 22, 2013

Abstract

We characterize the class of symmetric two-player games in which tit-for-tat cannot be beaten even by very sophisticated opponents in a repeated game. It turns out to be the class of exact potential games. More generally, there is a class of simple imitation rules that includes tit-for-tat but also imitate-the-best and imitate-if-better. Every decision rule in this class is essentially unbeatable in exact potential games. Our results apply to many interesting games including all symmetric 2x2 games, and standard examples of Cournot duopoly, price competition, public goods games, common pool resource games, and minimum effort coordination games.

Keywords: Imitation, tit-for-tat, decision rules, learning, exact potential games, symmetric games, repeated games, relative payoffs, zero-sum games.

JEL-Classifications: C72, C73, D43.

*We are extremely thankful to two anonymous referees who provided very valuable input. This paper subsumes the earlier working paper “Once beaten, never again: Imitation in two-player potential games”. Some of the material was previously circulated in a companion paper “Unbeatable Imitation”. We thank Carlos Alós-Ferrer, Chen Bo, Drew Fudenberg, Alexander Matros, Klaus Ritzberger, Karl Schlag, and John Stachurski for interesting discussions. Seminar audiences at Australian National University, Melbourne University, Monash University, UC Davis, UC San Diego, the Universities of Heidelberg, Konstanz, Vienna, and Zürich, the University of Queensland, the University of Oregon, Calpoly, at the International Conference on Game Theory in Stony Brook, 2009, the Midwestern Economic Theory Conference in Evanston 2010, and at the Econometric Society World Congress 2010 in Shanghai contributed helpful comments.

[†]Department of Economics, University of Heidelberg, Email: peter.duersch@awi.uni-heidelberg.de

[‡]Department of Economics, University of Heidelberg, Email: oechssler@uni-hd.de

[§]Department of Economics, University of California, Davis, Email: bcschipper@ucdavis.edu

1 Introduction

In a repeated two-player game, tit-for-tat refers to the strategy in which a player always chooses the opponent's action from the previous round (Rapoport and Chammah, 1965). Axelrod (1980a, 1980b) observed in his famous tournaments that tit-for-tat when started with “cooperation” is surprisingly successful in sustaining cooperation in the prisoners' dilemma. Part of the success is due to the fact that it can hardly be exploited by an opponent who may follow various more complex decision rules, a fact already noted in Axelrod (1980a). Axelrod and Hamilton (1981) suggest that tit-for-tat is successful from an evolutionary point of view (for a critique, see Selten and Hammerstein, 1984, or Nowak and Sigmund, 1993).

In this paper we extend the analysis to larger strategy sets and ask what is the class of symmetric two-player games in which tit-for-tat cannot be beaten by *any* other decision rule? To be precise, suppose you play a repeated symmetric game against one opponent. Suppose further that you know that this opponent uses tit-for-tat. Since the rule is deterministic, you therefore know exactly what your opponent will do in all future periods and how he will react to your actions. The question we pose in this paper is whether you can use this knowledge to exploit the tit-for-tat player in the sense that you achieve a higher payoff than the tit-for-tat player. We call tit-for-tat *essentially unbeatable* if in the repeated game there exists no strategy of the tit-for-tat player's opponent with which the opponent can obtain, in total, over a possibly infinite number of periods, a payoff difference that is more than the maximal payoff difference between outcomes in the one-period game. It turns out that tit-for-tat is essentially unbeatable if and only if the game is an exact potential game. Exact potential games have been introduced by Monderer and Shapley (1996) to show existence of and convergence to pure Nash equilibria. They are studied widely in the literature on learning in games (see Sandholm, 2010). The class of symmetric two-player exact potential games includes many meaningful games such as all symmetric 2x2 games, and standard examples of Cournot duopoly, price competition, public goods games, common pool resource games, and minimum effort coordination games.

More generally, we show that there is a class of decision rules such that any rule in this class is essentially unbeatable if the game is a two-player symmetric exact potential game. This class of decision rules features variants of imitation. It includes tit-for-tat but also imitate-the-best (Vega-Redondo, 1997) and imitate-if-better (Duersch, Oechssler, and Schipper, 2012a), where the latter prescribes to mimic the opponent's action from the

previous period if and only if the opponent received a higher payoff in the previous period.

To gain some intuition for our results, consider the game of “chicken” presented in the following payoff matrix.

	swerve	straight
swerve	3, 3	1, 4
straight	4, 1	0, 0

What should a forward looking opponent do if she knows that she is facing a tit-for-tat player?¹ To beat the tit-for-tat player she can play “straight” against “swerve” once. But if she wants to play this again, she also has to play once “swerve” against “straight”, which equalizes the score. Thus, the maximum payoff difference the opponent can obtain against a tit-for-tat player is 3, the maximal one-period payoff differential.

Suppose now that the imitator uses the rule “imitate if and only if the other player obtained a higher payoff in the previous round” and starts out with playing “swerve”. If the opponent decides to play “straight”, she will earn more than the imitator today but will be copied by the imitator tomorrow. From then on, the imitator will stay with “straight” forever. If she decides to play “swerve” today, then she will earn the same as the imitator and the imitator will stay with “swerve” as long as the opponent stays with “swerve”. Suppose the opponent is a dynamic *relative* payoff maximizer. In that case, the dynamic relative payoff maximizer can beat the imitator at most by the maximal one-period payoff differential of 3. Now suppose the opponent maximizes the sum of her *absolute* payoffs. The best an absolute payoff maximizer can do is to play swerve forever. In this case the imitator cannot be beaten at all as he receives the same payoff as his opponent. In either case, imitation comes very close to the top-performing heuristics and there is no evolutionary pressure against such a heuristic.

These results extend our recent paper Duersch, Oechssler, and Schipper (2012a), in which we show that imitation is subject to a money pump (i.e., can be exploited without bounds) if and only if the relative payoff game is of the rock-paper-scissors variety. The current results are stronger because they show that imitation can only be exploited with a bound that is equal to the payoff difference in the one-period game. Furthermore, the class of imitation rules considered here is significantly broader. In particular, it includes unconditional imitation rules like tit-for-tat. However, this comes at the cost of having to restrict ourselves to the class of symmetric two-player exact potential games. But as mentioned above, many economically relevant games satisfy this property.

¹In this paper we shall call the tit-for-tat player, or more generally, the imitator “he” and all possible opponents “she”.

The behavior of learning heuristics such as imitation has previously been studied mostly for the case when all players use the same heuristic. For the case of imitate-the-best, Vega-Redondo (1997) showed that in a symmetric Cournot oligopoly with imitators, the long run outcome converges to the competitive output if small mistakes are allowed. This result has been generalized to aggregative quasisubmodular games by Schipper (2003) and Alós-Ferrer and Ania (2005). Huck, Normann, and Oechssler (1999), Offerman, Potters, and Sonnemans (2002), and Apesteguia et al. (2007, 2010) provide some experimental evidence in favor of imitative behavior. In contrast to the above cited literature, the current paper deals with the interaction of an imitator and a possibly forward looking, very rational and patient player. Apart from experimental evidence in Duersch, Kolb, Oechssler, and Schipper (2010) and our own paper Duersch, Oechssler, and Schipper (2012a) we are not aware of any work that deals with this issue. For a Cournot oligopoly with imitators and myopic best reply players, Schipper (2009) showed that the imitators' long run average payoffs are strictly higher than the best reply players' average payoffs.

The article is organized as follows. In the next section, we present the model and provide a formal definition for being unbeatable. Section 3 introduces exact potential games. Necessary and sufficient conditions for tit-for-tat to be essentially unbeatable are given in Section 4 followed by a number of potential applications. We finish with Section 5, where we summarize and discuss the results.

2 Model

We consider a symmetric two-player game (X, π) , in which both players are endowed with the same compact set of pure actions X . For each player, the continuous payoff function is denoted by $\pi : X \times X \rightarrow \mathbb{R}$, where $\pi(x, y)$ denotes the payoff to the player choosing the first argument when his opponent chooses the second argument. We will frequently make use of the following definition.

Definition 1 (Relative payoff game) *Given a symmetric two-player game (X, π) , the relative payoff game is (X, Δ) , where the relative payoff function $\Delta : X \times X \rightarrow \mathbb{R}$ is defined by*

$$\Delta(x, y) = \pi(x, y) - \pi(y, x).$$

Note that, by construction, every relative payoff game is a symmetric zero-sum game since $\Delta(x, y) = -\Delta(y, x)$.

Next, we specify a number of different imitation rules.

Definition 2 (Tit-for-tat) *A player plays strategy tit-for-tat if he plays in each period $t \geq 1$ whatever his opponent did in the preceding period $t - 1$, i.e., $y_t = x_{t-1}$. Moreover, we allow y_0 to be arbitrary.²*

While tit-for-tat is an *unconditional* imitation rule, in economics imitation is often thought to be payoff dependent. The following decision rule was studied by Vega-Redondo (1997). A player follows *imitate-the-best* if for $t \geq 1$ his action is given by

$$y_t \in \begin{cases} \{x_{t-1}\} & \text{if } \Delta(x_{t-1}, y_{t-1}) > 0 \\ \{x_{t-1}, y_{t-1}\} & \text{if } \Delta(x_{t-1}, y_{t-1}) = 0 \\ \{y_{t-1}\} & \text{if } \Delta(x_{t-1}, y_{t-1}) < 0 \end{cases}$$

and arbitrary $y_0 \in X$. Duersch, Oechssler, and Schipper (2012a) use a version of it. A player follows *imitate-if-better* if for $t \geq 1$ his action is given by

$$y_t = \begin{cases} x_{t-1} & \text{if } \Delta(x_{t-1}, y_{t-1}) > 0 \\ y_{t-1} & \text{otherwise} \end{cases}$$

and arbitrary $y_0 \in X$.

The following class of imitation rules includes all three of the above rules, tit-for-tat, imitate-the-best, and imitate-if-better.³

Definition 3 (Imitation) *We call a player an imitator if for $t \geq 1$, his action is given by*

$$y_t \in \begin{cases} \{x_{t-1}\} & \text{if } \Delta(x_{t-1}, y_{t-1}) > 0 \\ \{x_{t-1}, y_{t-1}\} & \text{otherwise} \end{cases}$$

and arbitrary $y_0 \in X$.

That is, the imitator always adopts the opponent's action if in the previous round the opponent's payoff was strictly higher than that of the imitator. In other words, if the imitator decides to stick to his action, the other player must have had a weakly lower payoff.

²In Axelrod's tournament on the prisoners' dilemma, tit-for-tat submitted by Anatol Rapoport also prescribed "cooperate" as initial action (see Axelrod, 1980a, 1980b). While the prisoners' dilemma has a well-defined cooperative action, not all games possess such an action. We therefore consider a definition of tit-for-tat without restrictions on the initial action.

³In the context of the prisoner's dilemma, the class includes strategies like two-tits-for-tat, "always D", and grimm trigger but not tit-for-two-tats or Pavlov (Nowak and Sigmund, 1993).

Our aim is to determine whether there exists a strategy with which the imitator's opponent can obtain substantially higher payoffs than the imitator. We allow for any strategy of the opponent, including very sophisticated ones. In particular, the opponent may be infinitely patient and forward looking, and may never make mistakes. More importantly, she may know exactly what her opponent, the imitator, will do at all times, including the imitator's starting value. She may also commit to any closed loop strategy.

Consider now a situation in which the imitator starts out with a very unfavorable initial action. A clever opponent who knows this initial action can take advantage of it. Suppose that from then on the opponent has no strategy that makes her better off than the imitator. Arguably, the disadvantage in the initial period should not play a role in the long run. This motivates the following definition.

Definition 4 (Essentially unbeatable) *We say that imitation is essentially unbeatable if for any strategy of the opponent, the imitator can be beaten in total by at most the maximal one-period payoff differential, i.e., if for any sequence of actions by the opponent (x_0, x_1, \dots) and any initial action y_0 ,*

$$\sum_{t=0}^T \Delta(x_t, y_t) \leq \max_{x, y} \Delta(x, y), \text{ for all } T \geq 0 \quad (1)$$

where y_t is defined by a specific imitation rule.

3 Exact potential games

The question of unbeatability of imitation is closely linked to the class of exact potential games (Monderer and Shapley, 1996).

Definition 5 (Exact potential games) *The symmetric game (X, π) is an exact potential game if there exists an exact potential function $P : X \times X \rightarrow \mathbb{R}$ such that for all $y \in X$ and all $x, x' \in X$,*⁴

$$\begin{aligned} \pi(x, y) - \pi(x', y) &= P(x, y) - P(x', y), \\ \pi(x, y) - \pi(x', y) &= P(y, x) - P(y, x'). \end{aligned}$$

⁴Given the symmetry of (X, π) , the second equation plays the role usually played by the quantifier “for all players” in the definition of potential games.

Duersch, Oechssler, and Schipper (2012b) show that a symmetric two-player game is an exact potential game if and only if its relative payoff game is also an exact potential game. Furthermore, in symmetric two-player games, exact potential games have a relative payoff function that is additively separable, and has increasing and decreasing differences. All these definitions may appear to be restrictive. However, we will show below that there is a fairly large number of important examples that fall into this class.

Definition 6 (Additively separable) *A relative payoff function Δ is additively separable if $\Delta(x, y) = f(x) + g(y)$ for some functions $f, g : X \rightarrow \mathbb{R}$.*

Additive separable games have been studied in Balder (1997) and Peleg (1998).

Definition 7 (Increasing/decreasing differences) *A (relative) payoff function Δ has decreasing (resp. increasing) differences on $X \times X$ if there exists a total order $>$ on X such that for all $x'', x', y'', y' \in X$ with $x'' > x'$ and $y'' > y'$,*

$$\Delta(x'', y'') - \Delta(x', y'') \leq (\geq) \Delta(x'', y') - \Delta(x', y'). \quad (2)$$

Δ is a valuation if it has both decreasing and increasing differences.

Games with increasing differences have been introduced by Topkis (1998).

It is interesting that in our context all of the above properties define the same class of games.

Remark 1 *For symmetric two-player games the following conditions are equivalent: (i) (X, π) is an exact potential game. (ii) (X, Δ) is an exact potential game. (iii) Δ has increasing differences. (iv) Δ has decreasing differences. (v) Δ is additively separable.*

Duersch, Oechssler, and Schipper (2012b, Theorem 20) show that (i) and (ii) are equivalent. There, we also show that (iii) and (iv) are equivalent for all symmetric two-player zero-sum games (Proposition 13). Hence, (iii) or (iv) imply that Δ is a valuation. Brânzei, Mallozzi, and Tijs (2003, Theorem 1) show that (ii) is equivalent to Δ being a valuation for zero-sum games. Finally, Topkis (1998, Theorem 2.6.4.) shows equivalence of (v) and Δ being a valuation for zero-sum games.

4 Results

We are ready to state our main result.

Theorem 1 *Let (X, π) be a symmetric two-player game, where X is compact and π is continuous. Tit-for-tat is essentially unbeatable if and only if (X, π) is an exact potential game.*

PROOF. We prove here only “ \Rightarrow ”. The converse follows directly from Proposition 1 below.

For tit-for-tat to be essentially unbeatable, for any $T > 0$ there must not be a limit cycle $(x_t, y_t)_{t=0}^T$, with $\sum_{t=0}^T \Delta(x_t, y_t) > 0$ in which the tit-for-tat-player plays $y_t = x_{t-1}$ for $t \geq 1$ and $y_0 = x_T$.

This implies in particular that for every 3-cycle $(x_0, x_2), (x_1, x_0), (x_2, x_1), (x_0, x_2) \dots$ it must hold that

$$\Delta(x_0, x_2) + \Delta(x_1, x_0) + \Delta(x_2, x_1) \leq 0. \quad (3)$$

Since this must hold for every 3-cycle, it must also hold for the reverse cycle $(x_2, x_0), (x_1, x_2), (x_0, x_1), (x_2, x_0), \dots$, yielding

$$\Delta(x_2, x_0) + \Delta(x_1, x_2) + \Delta(x_0, x_1) \leq 0. \quad (4)$$

Since (X, Δ) is a symmetric zero-sum game, inequalities (3) and (4) imply

$$\Delta(x_0, x_2) + \Delta(x_1, x_0) + \Delta(x_2, x_1) = 0.$$

Hence,

$$\Delta(x_0, x_2) + \Delta(x_2, x_1) = -\Delta(x_1, x_0) = \Delta(x_0, x_1)$$

and thus (since $\Delta(x, x) = 0$ for any $x \in X$),

$$\Delta(x_0, x_2) - \Delta(x_2, x_2) = \Delta(x_0, x_1) - \Delta(x_2, x_1).$$

Since this equation holds for any 3-cycle, we have that Δ is a valuation. I.e., for all $x'', x', x \in X$,

$$\Delta(x'', x) - \Delta(x', x) = \Delta(x'', x') - \Delta(x', x').$$

By Remark 1 it implies that (X, π) is an exact potential game. \square

One direction of Theorem 1 actually holds for the more general class of imitation rules given in Definition 3.

Proposition 1 *Let (X, π) be a symmetric two-player game, where X is compact and π is continuous. If (X, π) is an exact potential game, then any imitation rule in the class given in Definition 3 is essentially unbeatable.*

PROOF. By Remark 1, if (X, π) is a symmetric two-player exact potential game, then Δ is additively separable. I.e., $\Delta(x, y) = f(x) + g(y)$ for some functions $f, g : X \rightarrow \mathbb{R}$. Since the relative payoff game is a symmetric zero-sum game, we have that

$$\Delta(x, x) = f(x) + g(x) = 0,$$

and hence

$$\Delta(x, y) = f(x) - f(y).$$

Let (x_0, x_1, \dots) be a sequence of actions generated by the opponent's strategy, and let $\{\Delta(x_t, y_t)\}_{t=0,1,\dots}$ be her associated sequence of relative payoffs when the imitator follows an imitation rule satisfying Definition 3. We claim that

$$f(x_t) - f(y_{t+1}) \leq 0, \text{ for all } t \geq 0. \quad (5)$$

This follows because either the imitator imitates, i.e., $y_{t+1} = x_t$ or he does not, $y_{t+1} = y_t$, in which case it must have been the case that $\Delta(x_t, y_t) = f(x_t) - f(y_t) \leq 0$.

Given (5), the sum of relative payoffs satisfies for any $T \geq 0$,

$$\sum_{t=0}^T \Delta(x_t, y_t) = \sum_{t=0}^T (f(x_t) - f(y_t)) \leq f(x_T) - f(y_0) = \Delta(x_T, y_0) \leq \max_{x,y} \Delta(x, y),$$

where $\max_{x,y} \Delta(x, y)$ exists because π is continuous and X is compact. \square

In contrast to tit-for-tat, the existence of an exact potential function is *not* a necessary condition for being essentially unbeatable for imitate-the-best and imitate-if-better. That is, going from tit-for-tat to the more general class of imitation rules comes at the cost of losing the converse of Theorem 1. To see this, consider the following game, which is not an exact potential game.

$$\pi = \begin{array}{c|ccc} & A & B & C \\ \hline A & 0, 0 & 0, -1 & -1, 0 \\ B & -1, 0 & 0, 0 & 0, 10 \\ C & 0, -1 & 10, 0 & 0, 0 \end{array} \quad \Delta = \begin{array}{c|ccc} & A & B & C \\ \hline A & 0 & 1 & -1 \\ B & -1 & 0 & -10 \\ C & 1 & 10 & 0 \end{array}$$

It is easy to see that imitate-the-best and imitate-if-better are essentially unbeatable for this game. However, tit-for-tat could be exploited without any bound by following

a cycle of actions ($A \rightarrow B \rightarrow C \rightarrow A \dots$). The reason for this difference is that an imitate-the-best or imitate-if-better player would never leave action C whereas a tit-for-tat player can be induced to follow the opponent from C to A .

In the chicken game discussed in the Introduction, imitation was essentially unbeatable since the maximal payoff difference was 3. Axelrod (1980a, b) observed that tit-for-tat was unexploitable by other decision rules in the prisoners' dilemma. More generally, since every symmetric 2×2 is an exact potential game, Proposition 1 implies the following corollary.

Corollary 1 *In any symmetric 2×2 game, imitation is essentially unbeatable.*

It is also easy to see why the corollary is true without making use of the notion of exact potential. Let $X = \{x, x'\}$. Consider a period t in which the opponent achieves a strictly positive relative payoff, $\Delta(x, x') > 0$. (If no such period t in which the opponent achieves a strictly positive relative payoff exists, then trivially imitation is essentially unbeatable.) The relative payoff game is symmetric zero-sum and hence $\Delta(x', x) = -\Delta(x, x')$ and $\Delta(x, x) = \Delta(x', x') = 0$. Therefore, the action combination (x, x') is the only one that makes imitation worse off. However, in this case, imitation will immediately imitate and play x in the next round. The only way to move imitation back to play x' is for the opponent to playing x' first, leading to the payoff $\Delta(x', x) = -\Delta(x, x') < 0$. Thus, every period with a positive relative profit for the opponent must be preceded by a period with a negative relative profit of the same absolute value. Depending on the exact imitation rule, the imitator might not follow in the first period where the opponent plays x' , leading to even worse relative profits, but it will always follow to x right away.

Note that “Matching pennies” is not a counter-example since it is not symmetric.

A sufficient condition for the additive separability of relative payoffs and thus the existence of an exact potential for the symmetric two-player game (X, π) is provided in the next result.

Corollary 2 *Consider a game (X, π) with a compact action set X and a payoff function that can be written as $\pi(x, y) = f(x) + g(y) + a(x, y)$ for some continuous functions $f, g : X \rightarrow \mathbb{R}$ and a symmetric function $a : X \times X \rightarrow \mathbb{R}$ (i.e., $a(x, y) = a(y, x)$ for all $x, y \in X$). Then imitation is essentially unbeatable.*

The following examples demonstrate that the assumption of additively separable relative payoffs is not as restrictive as may be thought at first glance.

Example 1 (Cournot Duopoly with Linear Demand) Consider a (quasi) Cournot duopoly given by the symmetric payoff function $\pi(x, y) = x(b - x - y) - c(x)$ with $b > 0$. Since $\pi(x, y)$ can be written as $\pi(x, y) = bx - x^2 - c(x) - xy$, Corollary 2 applies, and imitation is essentially unbeatable.

Example 2 (Bertrand Duopoly with Product Differentiation) Consider a differentiated duopoly with constant marginal costs, in which firms 1 and 2 set prices x and y , respectively. Firm 1's profit function is given by $\pi(x, y) = (x - c)(a + by - \frac{1}{2}x)$, for $a > 0$, $b \in [0, 1/2)$. Since $\pi(x, y)$ can be written as $\pi(x, y) = ax - ac + \frac{1}{2}cx - \frac{1}{2}x^2 - bcy + bxy$, Corollary 2 applies, and imitation is essentially unbeatable.

Example 3 (Public Goods) Consider the class of symmetric public good games defined by $\pi(x, y) = g(x, y) - c(x)$ where $g(x, y)$ is some symmetric monotone increasing benefit function and $c(x)$ is an increasing cost function. Usually, it is assumed that g is an increasing function of the sum of provisions, $x + y$. Various assumptions on g have been studied in the literature such as increasing or decreasing returns. In any case, Corollary 2 applies, and imitation is essentially unbeatable.

Example 4 (Common Pool Resources) Consider a common pool resource game with two appropriators. Each appropriator has an endowment $e > 0$ that can be invested in an outside activity with marginal payoff $c > 0$ or into the common pool resource. Let $x \in X \subseteq [0, e]$ denote the opponent's investment into the common pool resource (likewise y denotes the imitator's investment). The return from investment into the common pool resource is $\frac{x}{x+y}(a(x+y) - b(x+y)^2)$, with $a, b > 0$. So the symmetric payoff function is given by $\pi(x, y) = c(e - x) + \frac{x}{x+y}(a(x+y) - b(x+y)^2)$ if $x, y > 0$ and ce otherwise (see Walker, Gardner, and Ostrom, 1990). Since $\Delta(x, y) = (c(e - x) + ax - bx^2) - (c(e - y) + ay - by^2)$, Proposition 1 implies that imitation is essentially unbeatable.

Example 5 (Minimum Effort Coordination) Consider the class of minimum effort games given by the symmetric payoff function $\pi(x, y) = \min\{x, y\} - c(x)$ for some cost function $c(\cdot)$ (see Bryant, 1983, and Van Huyck, Battalio, and Beil, 1990). Corollary 2 implies that imitation is essentially unbeatable.

Example 6 (Synergistic Relationship) Consider a synergistic relationship among two individuals. If both devote more effort to the relationship, then they are both better off,

but for any given effort of the opponent, the return of the player's effort first increases and then decreases. The symmetric payoff function is given by $\pi(x, y) = x(c + y - x)$ with $c > 0$ and $x, y \in X \subset \mathbb{R}_+$ with X compact (see Osborne, 2004, p.39). Corollary 2 implies that imitation is essentially unbeatable.

Example 7 (Diamond's Search) Consider two players who exert effort searching for a trading partner. Any trader's probability of finding another particular trader is proportional to his own effort and the effort by the other. The payoff function is given by $\pi(x, y) = \alpha xy - c(x)$ for $\alpha > 0$ and c increasing (see Milgrom and Roberts, 1990, p. 1270). The relative payoff game of this two-player game is additively separable. By Proposition 1 imitation is essentially unbeatable.

5 Discussion

We have shown in this paper that there is a class of imitation rules that is surprisingly robust to exploitation by any strategy in symmetric two-player exact potential games. This includes strategies by truly sophisticated opponents. The property that imitation is unbeatable in these games seems to be unique among commonly used learning rules. We are not aware of other rules outside our class of imitation rules share this property. For example, many commonly used belief learning rules, such as best response learning or fictitious play, can easily be exploited in all games in which a Stackelberg leader achieves a higher payoff than the follower (as e.g. in Cournot games). Against such rules, the opponent can simply stubbornly choose the Stackelberg leader action knowing that the belief learning player will eventually converge to the Stackelberg follower action. Thus, belief learning rules can be beaten without bounds in such games. Yet, it remains an open question for future research whether there are other behavioral rules that perform equally well as imitation.

In Duersch, Oechssler, and Schipper (2012a) we show some extensions to more general classes of two-player games including relative payoff games with a generalized ordinal potential that come at cost of weakening the criterion of essentially unbeatable and focusing on the rule "imitate-if-better". Extensions to n -player games must be left for further research.

References

- [1] Alós-Ferrer, C. and A.B. Ania (2005). The evolutionary stability of perfectly competitive behavior, *Economic Theory* **26**, 497–516.
- [2] Apesteguia, J., Huck, S., and J. Oechssler (2007). Imitation - Theory and experimental evidence, *Journal of Economic Theory* **136**, 217–235.
- [3] Apesteguia, J., Huck, S., Oechssler, J., and Weidenholzer, S. (2010). Imitation and the evolution of Walrasian behavior: Theoretically fragile but behaviorally robust, *Journal of Economic Theory*, **145**, 1603–1617.
- [4] Axelrod, R. (1980a). Effective choice in prisoner’s dilemma, *Journal of Conflict Resolution* **24**, 3–25.
- [5] Axelrod, R. (1980b). More effective choice in prisoner’s dilemma, *Journal of Conflict Resolution* **24**, 379–403.
- [6] Axelrod, R. and Hamilton, W.D. (1981). The evolution of cooperation, *Science* **211**, 1390–1396.
- [7] Balder, E. (1997). Remarks on Nash equilibria for games with additively coupled payoffs, *Economic Theory* **9**, 161–167.
- [8] Brânzei, R., Mallozzi, L., and S. Tijs (2003). Supermodular games and potential games, *Journal of Mathematical Economics* **39**, 39–49.
- [9] Bryant, J. (1983). A simple rational expectations Keynes-type coordination model, *Quarterly Journal of Economics* **98**, 525–528.
- [10] Duersch, P., Kolb, A., Oechssler, J., and B.C. Schipper (2010). Rage against the machines: How subjects play against learning algorithms, *Economic Theory* **43**, 407–430.
- [11] Duersch, P., Oechssler, J., and B.C. Schipper (2012a). Unbeatable imitation, *Games and Economic Behavior* **76**(1), 88–96.
- [12] Duersch, P., Oechssler, J., and B.C. Schipper (2012b). Pure strategy equilibria in symmetric two-player zero-sum games, *International Journal of Game Theory*, 41(3), 553–564.

- [13] Huck, S., Normann, H.-T., and J. Oechssler (1999). Learning in Cournot oligopoly - An experiment, *Economic Journal* **109**, C80–C95.
- [14] Milgrom, P. and J. Roberts (1990). Rationalizability, learning, and equilibrium in games with strategic complementarities, *Econometrica* **58**, 1255–1277.
- [15] Monderer, D. and L.S. Shapley (1996). Potential games, *Games and Economic Behavior* **14**, 124–143.
- [16] Nowak, M. and K. Sigmund (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner’s Dilemma game. *Nature* **364**, 56–58.
- [17] Peleg, B. (1998). Almost all equilibria in dominant strategies are coalition-proof, *Economics Letters* **60**, 157–162.
- [18] Offerman, T., Potters, J., and J. Sonnemans (2002). Imitation and belief learning in an oligopoly experiment, *Review of Economic Studies* **69**, 973–997.
- [19] Osborne, M. (2004). *An introduction to game theory*, Oxford University Press.
- [20] Rapoport, A. and Chammah, A.M. (1965). *Prisoner’s dilemma: A study of conflict and cooperation*, Michigan University Press.
- [21] Sandholm, W. (2010). *Population games and evolutionary dynamics*, MIT Press.
- [22] Schipper, B.C. (2009). Imitators and optimizers in Cournot oligopoly, *Journal of Economic Dynamics and Control* **33**, 1981–1990.
- [23] Schipper, B.C. (2003). Submodularity and the evolution of Walrasian behavior, *International Journal of Game Theory* **32**, 471–477.
- [24] Schlag, K. (1998). Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits, *Journal of Economic Theory* **78**, 130–56.
- [25] Selten, R. and Hammerstein, P. (1984). Gaps in Harley’s argument on evolutionarily stable learning rules and in the logic of “tit for tat”, *Behavioral and Brain Sciences* **7**, 115–116.
- [26] Topkis, D. M. (1998). *Supermodularity and complementarity*, Princeton, New Jersey: Princeton University Press.
- [27] Van Huyck, J., Battalio, R., and R. Beil (1990). Tacit coordination games, strategic uncertainty and coordination failure, *American Economic Review* **80**, 234–248.

- [28] Vega-Redondo, F. (1997). The evolution of Walrasian behavior, *Econometrica* **65**, 375–384.
- [29] Walker, J.M., Gardner, R., and E. Ostrom (1990). Rent dissipation in a limited-access Common-Pool resource: Experimental evidence, *Journal of Environmental Economics and Management* **19**, 203–211.