Unbeatable Imitation*

Peter Duersch† Jörg Oechssler ‡ Burkhard C. Schipper§

November 10, 2011
– 2nd revised version –

Abstract

We show that for many classes of symmetric two-player games, the simple decision rule "imitate-if-better" can hardly be beaten by any strategy. We provide necessary and sufficient conditions for imitation to be unbeatable in the sense that there is no strategy that can exploit imitation as a money pump. In particular, imitation is subject to a money pump if and only if the relative payoff function of the game is of the rock-scissors-paper variety. We also show that a sufficient condition for imitation not being subject to a money pump is that the relative payoff game is a generalized ordinal potential game or a quasiconcave game. Our results apply to many interesting examples of symmetric games including 2×2 games, Cournot duopoly, price competition, public goods games, common pool resource games, and minimum effort coordination games.

Keywords: Imitate-the-best, learning, symmetric games, relative payoffs, zero-sum games, rock-paper-scissors, finite population ESS, generalized ordinal potential games, quasiconcave games.

JEL-Classifications: C72, C73, D43.

*We thank the associate editor and two anonymous referees for their suggestion. Moreover, 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, Wien, 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

"Whoever wants to set a good example must add a grain of foolishness to his virtue: then others can imitate and yet at the same time surpass the one they imitate - which human beings love to do." Friedrich Nietzsche

1 Introduction

Psychologists and behavioral economists stress the role of simple heuristics or rules for human decision making under limited computational capabilities (see Gigerenzer and Selten, 2002). While such heuristics lead to successful decisions in some particular tasks, they may be suboptimal in others. It is plausible that decision makers may cease to adopt heuristics that do worse than others in relevant situations. If various heuristics are pitted against each other in a contest, then in the long run the heuristic with the highest payoff should survive.

The competing heuristics could be anything from very simple to rational, omniscient, and forward looking ones. Even if a specific rule is not currently among the contestants, there can always be a "mutation", i.e., an invention of a new rule, that enters the pool of rules. A heuristic that does very badly against other rules will not be around for long as it will not belong to the top performers. Consequently, we would like to raise the following question: Is there a simple adaptive heuristic that, in large classes of economically relevant situations, cannot be beaten by any strategy including even those of a rational, omniscient and forward looking opponents?

Arguably, a necessary condition for a heuristic to be evolutionarily viable is that it cannot be exploited without bounds by opponents. We call such an unbounded exploitation opportunity a "money pump". More precisely, we say that a heuristic is subject to a money pump if the sum of payoff differences an opponent can achieve against this heuristic is unbounded. Or equivalently, if there is a cyclic strategy of the opponent, in which the opponent constantly earns more than the heuristic.

In this paper, we show that for many classes of symmetric two-player games, the simple decision rule "imitate-if-better" is unbeatable in the sense that there cannot be a money pump.[1] Since our results hold for all possible strategies of the imitator's opponent,

[1] In Duersch, Oechssler, and Schipper (2011) we consider a stricter notion of unbeatable and provide sufficient conditions for imitation to satisfy this condition. There, we call imitation "essentially unbeatable" if in the infinitely repeated game there exists no strategy of the opponent with which she can obtain, in total, over an infinite number of periods, a payoff difference that is more than the maximal payoff difference for the one–period game.

they also apply to strategies by truly sophisticated opponents. In particular, the opponent may be infinitely patient, forward looking, and free of mistakes. More importantly, the opponent can be aware of the fact that she is matched against an imitator. That is, she may know exactly what her opponent, the imitator, would do at all times, including the imitator's starting value. Finally, the opponent may be able to commit to any strategy including any closed-loop strategy. This is a property of imitation that is not shared by any other standard learning rule we are aware of.

The idea for this paper emerged from a prior observation in experimental data. In Duersch, Kolb, Oechssler, and Schipper (2010), subjects played against computers that were programmed according to various learning algorithms in a Cournot duopoly. On average, human subjects easily won against all of their computer opponents with one exception: the computer following the rule "imitate-if-better", the rule that simply prescribes to mimic the action of another player if and only if the other player received a higher payoff in the previous period. This suggested to us that imitation may be hard to beat by other strategies including strategies by forward–looking players.

In order to determine the classes of game to which our results apply, we are able to provide *necessary and sufficient conditions* for imitation to be subject to a money pump. The paradigmatic example for a money pump is playing repeatedly the game rock–paper–scissors, in which, obviously, an imitator can be exploited without bounds. The main result of this paper is that imitation is subject to a money pump *if and only if* the *relative payoff* game in question contains a generalized rock–paper-scissors submatrix.

Since the existence of a rock–paper–scissors submatrix may be cumbersome to check in some instances, we also provide sufficient conditions for imitation not to be subject to a money pump that are based on more familiar concepts of quasiconcavity and generalized ordinal potentials. These sufficient conditions allow to cover examples such as all symmetric 2×2 games, Cournot duopoly, Bertrand duopoly, rent seeking, public goods games, common pool resource games, minimum effort coordination games, Diamond's search, Nash demand bargaining, etc.

To gain some intuition for why imitation is hard to beat, consider the game of "chicken" presented in the following payoff matrix.

$$\begin{array}{cc} & \begin{array}{cc} \text{swerve} & \text{straight} \end{array} \\ \begin{array}{c} \text{swerve} \\ \text{straight} \end{array} & \left(\begin{array}{cc} 3,3 & 1,4 \\ 4,1 & 0,0 \end{array} \right) \end{array}$$

Suppose that initially the imitator starts out with playing "swerve". What should a forward looking opponent do? If she 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 an heuristic.

The behavior of learning heuristics has previously been studied mostly for the case when all players use the same heuristic. For the case of imitate-the-best,[2] 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 forward looking, very rational and patient player. Apart from experimental evidence in Duersch, Kolb, Oechssler, and Schipper (2010) we are not aware of any work that deals with this issue.[3]

A recent paper by Feldman, Kalai, and Tennenholtz (2010) has a similar but complementary objective to ours. They study whether a strategy which they call "copycat" can be beaten in a symmetric two-player game by an arbitrary opponent who may have full knowledge of the game and may play any history dependent strategy. The copycat strategy, on the other hand, can only observe past actions of both players. Remarkably, the copycat strategy can nearly match the average payoff of the opponent. The strategy used by copycat is to equalize the occurrences of action profiles (x, y) and (y, x) for any $x, y \in X$. To achieve this, an auxiliary two-player zero-sum game is introduced whose

[2] For the two-player case, imitate-the-best and imitate-if-better are almost equivalent, the difference being that the latter specifically prescribes a tie-breaking rule (for the case of both players having equal payoffs in the previous round). Since we use imitate-if-better only in the two-player case, we do not need to specify what happens if more than one other player is observed.

[3] 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. Juang (2002) considered a similar question for coordination games.

payoff at (x, y) is the difference of frequencies of (x, y) and (y, x) played so far. Thus, similar to our approach, the authors use the idea of imitation and auxiliary zero-sum games. Yet, their copycat rule is far more sophisticated than our imitation rule as it entails finding a (possibly mixed) minmax strategy in the auxiliary zero–sum game in each round. On the other hand, the possible opponents are less omniscient than the opponents in our setting. In particular, the opponents in their setting cannot perfectly predict the imitator's action in the next round, which explains why Feldman et al.'s result applies even to rock–paper–scissors.

The article is organized as follows. In the next section, we present the model, provide a formal definition for being unbeatable, and our main result, which provides a necessary and sufficient condition for a money pump. Section 3 provides sufficient conditions for imitation not being subject to a money pump. We conclude with Section 4, 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 (finite or infinite) set of pure actions X. For each player, the bounded payoff function is denoted by $\pi : X \times X \longrightarrow \mathbb{R}$, 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 $\Delta : X \times X \longrightarrow \mathbb{R}$ 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 $\Delta(x, y) = -\Delta(y, x)$.

The *imitator* follows the simple rule "imitate-if-better". To be precise, the imitator adopts the opponent's action if and only if in the previous round the opponent's payoff was strictly higher than that of the imitator. Formally, the action of the imitator y_t in period t given the action of the other player from the previous period x_{t-1} is

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

for some initial action $y_0 \in X$.

Our aim is to determine whether there exists a strategy of the imitator's opponent that obtains 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.

Our notion of unbeatable is that an opponent can only take a finite advantage of an imitator. One can argue that the finite disadvantage should not play a role in the long run as time goes to infinity.

Definition 2 (No money pump) *We say that* imitation is not subject to a money pump *if there exists a finite bound M such that for any initial action of the imitator y_0 and any sequence $\{x_t\}$ of actions of the opponent*

$$\sum_{t=0}^{T} \Delta(x_t, y_t) \leq M, \quad \text{for all } T \geq 0, \tag{2}$$

where y_t is given by equation (1).

The name of this condition is motivated by the observation that in a finite game, imitation is not subject to a money pump if the opponent cannot create a cycle of actions that strictly improve her relative payoff at every step. This is reminiscent of "no money pumps" in economics. The following definitions make this precise.

Given a symmetric two-player game (X, π), a path in the action space $X \times X$ is a sequence of action profiles $(x_0, y_0), (x_1, y_1), ...$ A path is constant if $(x_t, y_t) = (x_{t+1}, y_{t+1})$ for all $t = 0, 1, ...$ Otherwise, the path is called non–constant. A non–constant finite path $(x_0, y_0), ..., (x_n, y_n)$ is a *cycle* if $(x_0, y_0) = (x_n, y_n)$ for some $n > 1$. Let us call a cycle an *imitation cycle* if for all (x_t, y_t) and (x_{t+1}, y_{t+1}) on the path of the cycle $\Delta(x_t, y_t) > 0$ and $y_{t+1} = x_t$. An imitation cycle is thus a particular cycle along which one player always obtains a strictly positive relative payoff and the other player mimics the action of the first player in the previous round. Thus, an imitation cycle never contains an action profile on the main-diagonal of the payoff matrix.

Lemma 1 *For any finite symmetric game (X, π), imitation is subject to a money pump if and only if there exists an imitation cycle.*

PROOF. Consider a finite symmetric game (X, π) and its relative payoff game (X, Δ). We show that if imitation is subject to a money pump, then there is a imitation cycle. The converse is trivial.

Since the game is finite, there can not be infinitely many strictly positive relative payoff improvements unless there is a cycle. To show that such a cycle implies an imitation cycle, suppose by contradiction that there exists a period t such that $\Delta(x_t, y_t) \leq 0$. W.l.o.g. assume that $\Delta(x_{t+1}, y_{t+1}) > 0$. This is w.l.o.g. because we assumed a money pump. By equation (1) the imitator will not imitate in $t+1$ the previous period's action of the opponent, i.e., $y_{t+1} = y_t$. But then, there must be a cycle with $x_t = x_{t+1}$, $x_{t+1} = x_{t+2}$, ... By applying this argument to any period t for which $\Delta(x_t, y_t) \leq 0$, we can construct a cycle with $\Delta(x_t, y_t) > 0$ for all t. The decision rule of the imitator then requires that $y_{t+1} = x_t$ for all t, which proves that such a cycle is an imitation cycle. □

2.1 A Necessary and Sufficient Condition

The game rock–paper-scissors is the paradigmatic example for how an imitator can be exploited without bounds by a clever opponent. In our terminology, imitation is subject to a money pump.

Example 1 (Rock-Paper-Scissors) Consider the well known rock-paper-scissors game.[4]

$$\begin{array}{c} \\ R \\ P \\ S \end{array} \begin{array}{c} \begin{array}{ccc} R & P & S \end{array} \\ \left(\begin{array}{ccc} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{array} \right) \end{array}$$

If the imitator starts for instance with R, then the opponent can play the cycle P-S-R... In this way, the opponent could win in every period and the imitator would lose in every period. Over time, the payoff difference would grow without bound in favor of the opponent.

We can generalize Example 1 by noting that the crucial feature of the example is that a money pump is created by the fact that for each action of the imitator there is an action of the opponent which yields her a strictly positive relative payoff and which yields the imitator a strictly negative relative payoff.

[4] In the following, we will often represent symmetric payoff matrices by the matrix of the row player's payoffs only.

Definition 3 (gRPS Matrix) *A symmetric zero-sum game* (X, π) *is called a generalized rock-paper-scissors (gRPS) matrix if for each column there exists a row with a strictly positive payoff to the row player, i.e., if for all* $y \in X$ *there exists a* $x \in X$ *such that* $\pi(x, y) > 0$*.*

It should be fairly obvious that if a zero–sum game contains somewhere on a main-diagonal a submatrix that is a generalized rock-paper-scissors matrix, then this is sufficient for a money pump as the opponent can make sure that the process cycles forever in this submatrix. What is probably less obvious is that the existence of such a submatrix is also necessary for a money pump.

Definition 4 (gRPS Game) *A symmetric zero-sum game* (X, π) *is called a generalized rock-paper-scissors (gRPS) game if it contains a submatrix* $(\bar{X}, \bar{\pi})$ *with* $\bar{X} \subseteq X$ *and* $\bar{\pi}(x, y) = \pi(x, y)$ *for all* $x, y \in \bar{X}$*, and* $(\bar{X}, \bar{\pi})$ *is a gRPS matrix.*

This leads us to our main result.

Theorem 1 *Imitation is subject to a money pump in the finite symmetric game* (X, π) *if and only if its relative payoff game* (X, Δ) *is a gRPS game.*

The proof follows from Lemma 1 and the following lemma.

Lemma 2 *Consider a finite symmetric game* (X, π) *with its relative payoff game* (X, Δ)*.* (X, Δ) *is a gRPS game if and only if there exists an imitation cycle.*

Proof. "$\Leftarrow$": If there exists an imitation cycle in (X, Δ), let $\bar{X}$ be the orbit of the cycle, i.e., all actions of X that are played along the imitation cycle. For each action (i.e., column) $y \in \bar{X}$, there exists an action (i.e., row) $x \in \bar{X}$ such that $\Delta(x, y) > 0$. Hence, $(\bar{X}, \bar{\Delta})$, where $\bar{\Delta}$ is defined by $\bar{\Delta}(x, y) = \Delta(x, y)$ for all $x, y \in \bar{X}$, is a gRPS submatrix. Thus, (X, Δ) is a gRPS game.

"$\Rightarrow$": If the relative payoff game (X, Δ) is a gRPS game, then it contains a gRPS submatrix $(\bar{X}, \bar{\Delta})$. That is, for each column of the matrix game $(\bar{X}, \bar{\Delta})$ there exists a row with a strictly positive relative payoff to player 1. Let the initial action of the imitator y be contained in $\bar{X}$. If the opponent selects such a row $x \in \bar{X}$ for which she earns a strict positive relative payoff, i.e., $\Delta(x, y) > 0$, then she will be imitated by the imitator in the next period. Yet, at the next period, when the imitator plays x, the opponent has another action $x' \in \bar{X}$ with a strictly positive relative payoff, i.e., $\Delta(x', x) > 0$. Thus the imitator will imitate her in the following period. More generally, for each action $y \in \bar{X}$ of the imitator, there is another action $x \in \bar{X}$, $x \neq y$ of the opponent that earns the latter a strictly positive relative payoff. Since $\bar{X}$ is finite, such a sequence of actions must contain a cycle. Moreover, we just argued that $\Delta(x_t, y_t) > 0$ and $y_{t+1} = x_t$ for all t. Thus, it is an imitation cycle. □

Our discussion of the "chicken" example in the introduction extends to all symmetric 2×2 games. Note that the relative payoff game of any symmetric 2×2 game cannot be a generalized rock–paper–scissors matrix since the latter must be a symmetric zero–sum game. If one of the row player's off-diagonal relative payoffs is $a > 0$, then the other must be $-a$ violating the definition of a gRPS matrix. Thus Theorem 1 implies that for any symmetric 2×2 game imitation is not subject to a money pump.

Remark 1 *In any symmetric* 2×2 *game, imitation is not subject to a money pump.*

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

2.2 The Relation to fESS

Theorem 1 is used to obtain an interesting necessary condition for imitation being not subject to a money pump.

Proposition 1 *Let* (X, π) *be a finite symmetric game with its relative payoff game* (X, Δ). *If* (X, Δ) *has no pure equilibrium, then imitation is subject to a money pump.*

Proof. By Duersch, Oechssler, and Schipper (2012, Observation 4), (X, Δ) has no symmetric pure equilibrium if and only if it is a gRPS matrix. Thus, if (X, Δ) has no symmetric pure equilibrium, then it is a gRPS game. Hence, by Theorem 1 imitation is subject to a money pump. □

In previous studies of imitation (see e.g. Alós-Ferrer and Ania, 2005; Schipper, 2003; Vega-Redondo, 1997), the concept of a finite population evolutionary stable strategy (Schaffer, 1988, 1989) played a prominent role in the analysis.

Definition 5 (fESS) *An action* $x^* \in X$ *is a* finite population evolutionary stable strategy (fESS) *of the game* (X, π) *if*

$$\pi(x^*, x) \geq \pi(x, x^*) \text{ for all } x \in X. \tag{3}$$

In terms of the relative payoff game, inequality (3) is equivalent to

$$\Delta(x^*, x) \geq 0 \text{ for all } x \in X.$$

Already Schaffer (1988, 1989) observed that the fESS of the game (X, π) and the symmetric pure Nash equilibria of the relative payoff game (X, Δ) coincide. Thus Proposition 1 implies:

Corollary 1 *If the finite symmetric game (X, π) has no fESS, then imitation is subject to a money pump.*

In other words, the existence of a fESS is a necessary condition for imitation not being subject to a money pump. The reason for the existence of a fESS not being sufficient is that there could be a gRPS submatrix of the game ("disjoint" from the fESS profile) that gives rise to an imitation cycle.

Since the relative payoff game of a symmetric zero-sum game is a gRPS game if and only if the underlying symmetric zero-sum game is a gRPS game, we obtain from Theorem 1 the following corollary.

Corollary 2 *Imitation is subject to a money pump in the finite symmetric zero-sum game (X, π) if and only if (X, π) is a gRPS game.*

3 Sufficient Conditions

The existence of a gRPS submatrix may be cumbersome to check in some instances. Therefore, we provide below two sufficient conditions for imitation not to be subject to a money pump that are based on the more familiar concepts of quasiconcavity and generalized ordinal potentials. We impose these properties on the relative payoff games rather than on the underlying games.

3.1 Relative Payoff Games with Generalized Ordinal Potentials

Potential functions are often useful for obtaining results on convergence of learning algorithms to equilibrium, existence of pure equilibrium, and equilibrium selection.[5] The

[5]For some of the classes of games considered here there exist convergence results for various learning processes although convergence results for imitation are rare (see Alós-Ferrer and Ania, 2005, Schipper, 2003, and Vega-Redondo, 1997). Note, however, that our results do not follow from any results in the literature since we do not consider a pair of imitators but rather one imitator against an arbitrary decision rule of the opponent.

following notion was introduced by Monderer and Shapley (1996).

Definition 6 (Generalized ordinal potential games) *The symmetric game* (X, π) *is a generalized ordinal potential game if there exists a generalized ordinal potential function* $P : X \times X \longrightarrow \mathbb{R}$ *such that for all* $y \in X$ *and all* $x, x' \in X$,[6]

$$\pi(x, y) - \pi(x', y) > 0 \quad \textit{implies} \quad P(x, y) - P(x', y) > 0,$$
$$\pi(x, y) - \pi(x', y) > 0 \quad \textit{implies} \quad P(y, x) - P(y, x') > 0.$$

Note that every exact potential game is a weighted potential game, every weighted potential game is an ordinal potential game, and every ordinal potential game is a generalized ordinal potential game. Monderer and Shapley (1996, Lemma 2.5 and the first paragraph on p. 129) show that any finite strategic game admitting a generalized ordinal potential possesses a pure Nash equilibrium. Thus, if (X, π) is a finite symmetric game with relative payoff game (X, Δ) and the latter is a generalized ordinal potential game, then (X, π) possesses a fESS.

A sequential path in the action space $X \times X$ is a sequence $(x_0, y_0), (x_1, y_1), ...$ of profiles $(x_t, y_t) \in X \times X$ such that for all $t = 0, 1, ...$, the action profiles (x_t, y_t) and (x_{t+1}, y_{t+1}) differ in exactly one player's action. A sequential path is a strict improvement path if for each $t = 0, 1, ...$, the player who switches her action at t strictly improves her payoff. A finite sequential path $(x_0, y_0), ..., (x_m, y_m)$ is a *strict improvement cycle* if it is a strict improvement path and $(x_0, y_0) = (x_m, y_m)$.

Lemma 3 *If* (X, Δ) *does not contain a strict improvement cycle, then it does not contain an imitation cycle.*[7]

Proof. We prove the contrapositive. I.e., if (X, Δ) contains an imitation cycle, then it contains a strict improvement cycle. Let $(x_0, y_0), ..., (x_m, y_m)$ be an imitation cycle. From this imitation cycle, we construct a strict improvement cycle as follows: For $t = 0, ..., m - 1$, we add the element (x_t, y_{t+1}) as successor to (x_t, y_t) and predecessor

[6] Given the symmetry of (X, π), the second line below plays the role usually played by the quantifier "for all players" in the definition of potential games.

[7] Ania (2008, Proposition 3) presents a similar result according to which if all players are imitators and imitation is payoff improving, then the fESS is a Nash equilibrium action. This is different from Lemma 3 as we consider an imitator against an opponent who herself may not imitate and focus on the relationship between relative payoff games that possess a generalized ordinal potential and imitation cycles.

to (x_{t+1}, y_{t+1}). That is, instead of simultaneous adjustments of actions at each period as in an imitation cycle, we let players adjust actions sequentially by taking turns. The imitator adjusts from (x_t, y_t) to (x_t, y_{t+1}) and the opponent from (x_t, y_{t+1}) to (x_{t+1}, y_{t+1}) for $t = 0, ..., m-1$. This construction yields a sequential path.

We now show that it is a strict improvement cycle. First, for the imitator, whenever he adjusts in $t = 0, ..., m-1$, we claim $\Delta(y_t, x_t) < \Delta(y_{t+1}, x_t) = 0$. Note that by symmetric zero-sum, $\Delta(y_t, x_t) = -\Delta(x_t, y_t) < 0$ because (x_t, y_t) is an element of an imitation cycle, i.e., $\Delta(x_t, y_t) > 0$. $\Delta(y_{t+1}, x_t) = 0$ because the imitator mimics the action of the opponent, $y_{t+1} = x_t$. Thus $\Delta(y_{t+1}, x_t) = \Delta(x_t, x_t) = 0$ by symmetric zero-sum.

Second, for the opponent, whenever she adjusts in $t = 1, ..., m$, $\Delta(x_t, y_t) > \Delta(x_{t-1}, y_t) = 0$ because (x_t, y_t) is an element of an imitation cycle, so $\Delta(x_t, y_t) > 0$. Moreover, the imitator mimics the action of the opponent, i.e., $y_t = x_{t-1}$, and thus $\Delta(x_{t-1}, y_t) = \Delta(x_{t-1}, x_{t-1}) = 0$. Hence $(x_0, y_0), (x_0, y_1), (x_1, y_1), ..., (x_{m-1}, y_m), (x_m, y_m)$ is indeed a strict improvement cycle. □

The converse is not true as the following counter-example shows.

Example 2 Consider the following relative payoff game.[8]

$$\Delta = \begin{array}{c} \\ a \\ b \\ c \end{array} \begin{array}{c} \begin{array}{ccc} a & b & c \end{array} \\ \begin{pmatrix} 0 & 0 & -1 \\ 0 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix} \end{array}$$

Clearly, this game is not a gRPS game. Thus, by Lemma 1 it does not possess an imitation cycle. However, we can construct a strict improvement cycle (b, a), (c, a), (c, c), (b, c) and (b, a).

Proposition 2 *Let (X, π) be a finite symmetric game with its relative payoff game (X, Δ). If (X, Δ) is a generalized ordinal potential game, then imitation is not subject to a money pump.*

Proof. Monderer and Shapley (1996, Lemma 2.5) show that a finite strategic game has no strict improvement cycle (what they call the finite improvement property) if and

[8] This example appears also in Ania (2008, Example 2), where it is used to demonstrate that the class of games where imitation is payoff improving (when all players are imitators) is not a subclass of generalized ordinal potential games.

only if it is a generalized ordinal potential game. Since this result holds for any finite strategic game, it holds also for any finite symmetric zero-sum game (X, Δ).

Lemma 3 shows that if (X, Δ) does not contain a strict improvement cycle, then it does not contain an imitation cycle. Thus Lemma 1 implies that imitation is not subject to a money pump. $\square$

If the converse were true, then the class of generalized ordinal potential relative payoff games and relative payoff games that are not gRPS games would coincide. Yet, the converse is not true. This follows again from Example 2. It is not a gRPS game but due to the existence of a strict improvement cycle by Monderer and Shapley (1996, Lemma 2.5) it does not possess a generalized ordinal potential.

For an example of a game whose relative payoff game is a generalized ordinal potential game consider the following payoff function π of a coordination game with an outside option. A generalized ordinal potential function of the relative payoff game is given by P.

$$\pi = \begin{array}{c} \\ A \\ B \\ C \end{array} \begin{array}{c} \begin{array}{ccc} A & B & C \end{array} \\ \begin{pmatrix} 4 & -1 & 0 \\ 2 & 3 & 0 \\ 0 & 0 & 0 \end{pmatrix} \end{array} \quad P = \begin{array}{c} \\ A \\ B \\ C \end{array} \begin{array}{c} \begin{array}{ccc} A & B & C \end{array} \\ \begin{pmatrix} -2 & -1 & -2 \\ -1 & 0 & 0 \\ -2 & 0 & 0 \end{pmatrix} \end{array} .$$

3.2 Quasiconcave Relative Payoff Games

Here we show that imitation is not subject to a money pump if the relative payoff game is quasiconcave.

Definition 7 (Quasiconcave) *A symmetric two-player game (X, π) is quasiconcave (or single-peaked) if there exists a total order $<$ on X such that for each $x, x', x'', y \in X$ and $x' < x < x''$, we have that $\pi(x, y) \geq \min \{\pi(x', y), \pi(x'', y)\}$.*

For a matrix game this definition implies that the row player's payoff has in each column a single peak. In Duersch, Oechssler, and Schipper (2012, Theorem 7), we show that if X is finite and Δ is quasiconcave, then an equilibrium of (X, Δ) and therefore a fESS of (X, π) exists.

Proposition 3 *Let (X, π) be a finite symmetric game with relative payoff game (X, Δ). If (X, Δ) is quasiconcave, then imitation is not subject to a money pump.*

Proof. Suppose (X, Δ) is a finite quasiconcave game. Consider a symmetric submatrix (X', Δ') where $X' \subset X$ and Δ' is the restriction of Δ to X'. It follows directly from Definition 7 that (X', Δ') is also a finite quasiconcave game. Lemma 8 in Duersch, Oechssler, and Schipper (2012), then implies that (X', Δ') is not a gRPS matrix. Since we picked an arbitrary $X' \subset X$, (X, Δ) is not a gRPS game. Thus, by Theorem 1, imitation is not subject to a money pump. □

The following corollary may be useful for applications. Let $X \subset \mathbb{R}^m$ be a finite subset of a finite dimensional Euclidean space. A function $f : X \longrightarrow \mathbb{R}$ is *convex (resp. concave)* if for any $x, x' \in X$ and for any $\lambda \in [0, 1]$ such that $\lambda x + (1 - \lambda)x' \in X$, $f(\lambda x + (1 - \lambda)x') \leq (\geq)\lambda f(x) + (1 - \lambda)f(x')$.

Corollary 3 *Let $(\mathbb{R}^m, \pi)$ be a symmetric two-player game for which $\pi(\cdot, \cdot)$ is concave in its first argument and convex in its second argument. If the players' actions are restricted to a finite subset X of the finite dimensional Euclidian space $\mathbb{R}^m$, then imitation is not subject to a money pump.*

Bargaining is an economically relevant situation involving two players. Our results imply that imitation is not subject to a money pump in bargaining as modeled in the Nash Demand game.

Example 3 (Nash Demand Game) Consider the following version of the Nash Demand game (see Nash, 1953). Two players simultaneously demand an amount in $\mathbb{R}_+$. If the sum is within a feasible set, i.e., $x + y \leq s$ for $s > 0$, then player 1 receives the payoff $\pi(x, y) = x$. Otherwise $\pi(x, y) = 0$ (analogously for player 2). The relative payoff function is quasiconcave. If the players' demands are restricted to a finite set, then Proposition 3 implies that imitation is not subject to a money pump.

Finally, we would like to remark that Example 2 is an instance of a quasiconcave relative payoff game but due to the strict improvement cycle it does not posses a generalized ordinal potential. Moreover, in Duersch, Oechssler, and Schipper (2012, Example 9) we show that there are relative payoff games that are neither gRPS games nor quasiconcave.

4 Discussion

We have shown in this paper that imitation is a behavioral rule that is surprisingly robust to exploitation by *any* strategy.[9] This includes strategies by truly sophisticated opponents. The only class of symmetric games in which imitation can really be beaten is the class of games whose relative payoff function is a generalized rock–paper–scissors game. According to Lemma 2 this is also the class of games in which there is an imitation cycle, i.e., a cycle in which the opponent always jumps to a new action which in turn is imitated by the imitator in the next round. One can show that many symmetric two-player games such as 2×2 games, Cournot duopoly, price competition, public goods games, common pool resource games, and minimum effort coordination games are not generalized rock-paper-scissors games (see Duersch, Oechssler, and Schipper, 2011, 2012). Thus, it seems fair to say that imitation is very hard to beat in large and generic classes of economically relevant games.

One may suspect that the property of not being subject to a money pump is shared by many other reasonable learning rules.[10] This impression, however, is false. Consider the following symmetric 3×3 game, which has the structure of a Cournot game with 3 actions, where A corresponds to the collusive or Stackelberg follower action, B to the Cournot-Nash action, and C to the Stackelberg leader action.

$$\begin{array}{c} \\ A \\ B \\ C \end{array} \begin{array}{c} \begin{array}{ccc} A & B & C \end{array} \\ \left(\begin{array}{ccc} 15,15 & 12,20 & 7,15 \\ 20,12 & 13,13 & 6,9 \\ 15,7 & 9,6 & 0,0 \end{array} \right) \end{array}$$

This game is not a generalized rock-paper-scissors game. Thus, imitation is not subject to a money pump. Yet, all belief based learning processes we are aware of (e.g. best response dynamics, fictitious play or any of its variants) would be subject to a money pump. A forward looking opponent would simply play C forever. Any belief based learning process (but also reinforcement type learning processes) would then eventually best respond to C by playing the "Stackelberg follower" position A, which creates a money pump. The same holds more generally for all games in which a Stackelberg leader achieves a higher payoff than the follower (as for instance in Cournot games).

We do not know of any rule that shares with imitate-if-better the property of being

[9] Since our results allow for *any* strategy of the opponent, they also apply to opponents that dynamically maximize a (discounted) sum of (absolute or relative) payoffs.

[10] We thank an anonymous referee for encouraging this discussion.

unbeatable in a large class of games.[11] For example, there are important differences between imitate-if-better and *unconditional* imitation, when behavior is imitated regardless of its success. A well known example of the latter is tit–for–tat. To see the difference, consider again the above game. Tit–for–tat is subject to a money pump by following a cycle ($A \to B \to C \to A \dots$), whereas imitate-if-better is not. The reason for this difference is that an 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.

There are other modifications that may cause the imitate-if-better rule to lose the property of being unbeatable. For instance, we assumed that an imitator sticks to his action in case of a tie in payoffs. To see what goes wrong with an alternative tie-braking rule consider a homogenous Bertrand duopoly with constant marginal costs. Suppose the imitator starts with a price equal to marginal cost. If the opponent chooses a price strictly above marginal cost, her profit is also zero. If nevertheless, the opponent were imitated, she could start the money pump by undercutting the imitator until they reach again price equal to marginal cost and then start the cycle again. To sum up, the property that imitate-if-better is unbeatable in a large class of games seems to be unique among commonly used learning rules. Yet, it remains an open question for future research whether there are other behavioral rules that perform equally well as imitate-if-better.

The restriction of our analysis to two–player games is certainly a limitation. While a full treatment of the n–player case is beyond the scope of the current paper, we provide here an example that shows how imitation can be beaten in a standard Cournot game when there are three players. Let the inverse demand function be $p(Q) = 100 - Q$ and the cost function be $c(q_i) = 10q_i$. Now consider the case of two relative payoff maximizers and one imitator. Writing a vector of quantities as (q_I, q_M, q_M), it is easy to check that the following sequence of action profiles $(0, 22.5, 22.5)$, $(22.5, 0, 68)$, $(0, 22.5, 22.5)$, $(22.5, 68, 0)$, $(0, 22.5, 22.5)$ ... is an imitation cycle. The two maximizers take turns in inducing the imitator to reduce his quantity to zero by increasing quantity so much that price is below marginal cost. Since the other maximizer has zero losses, she is imitated in the next period, which yields half of the monopoly profit for both maximizers. Clearly, this requires coordination among the two maximizers but this can be achieved in an infinitely repeated game by the use of a trigger strategy. Thus, imitation is subject to a money pump. Recall, however, that we pitted imitation against truly sophisticated opponents in a particular game. Whether imitation can be beaten also by less sophis-

[11] Of course there may be close variants of imitate-if-better that share this property, e.g. rules that imitate only with a certain probability like Schlag's (1998) proportional imitation rule.

ticated (e.g. human) opponents in a wider class of games remains to be seen in future experiments and in theoretical work on n-player games.

References

[1] Alós-Ferrer, C. and A.B. Ania (2005). The evolutionary stability of perfectly competitive behavior, *Economic Theory* **26**, 497–516.

[2] Ania, A. (2008). Evolutionary stability and Nash equilibrium in finite populations, with an application to price competition, *Journal of Economic Behavior and Organization* **65**, 472–488.

[3] Apesteguia, J., Huck, S., and J. Oechssler (2007). Imitation - Theory and experimental evidence, *Journal of Economic Theory* **136**, 217–235.

[4] 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.

[5] 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.

[6] Duersch, P., Oechssler, J., and B.C. Schipper (2012). Pure strategy equilibria in symmetric two-player zero-sum games, *International Journal of Game Theory*, forthcoming, DOI: 10.1007/s00182-011-0302-x.

[7] Duersch, P., Oechssler, J., and B.C. Schipper (2011). Once Beaten, Never Again: Imitation in Two-Player Potential Games, mimeo.

[8] Feldman, M., Kalai, A., and M. Tennenholtz (2010). Playing Games without Observing Payoffs, Proceedings of Innovations in Computer Science, 106-110.

[9] Gigerenzer, G. and R. Selten (2002). *Bounded rationality. The adaptive toolbox*, Cambridge, MA: The MIT Press.

[10] Hehenkamp, B., Leininger, W., and A. Possajennikov (2004). Evolutionary equilibrium in Tullock contests: Spite and overdissipation, *European Journal of Political Economy* **20**, 1045–1057.

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

[12] Juang, W.-T. (2002). Rule Evolution and Equilibrium Selection, *Games and Economic Behavior* **39**, 71-90.

[13] Monderer, D. and L.S. Shapley (1996). Potential games, *Games and Economic Behavior* **14**, 124–143.

[14] Offerman, T., Potters, J., and J. Sonnemans (2002). Imitation and belief learning in an oligopoly experiment, *Review of Economic Studies* **69**, 973–997.

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

[16] Schaffer, M.E. (1988). Evolutionarily stable strategies for a finite population and a variable contest size, *Journal of Theoretical Biology* **132**, 469–478.

[17] Schipper, B.C. (2009). Imitators and optimizers in Cournot oligopoly, *Journal of Economic Dynamics and Control* **33**, 1981–1990.

[18] Schipper, B.C. (2003). Submodularity and the evolution of Walrasian behavior, *International Journal of Game Theory* **32**, 471–477.

[19] Schlag, K. (1998). Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits, *Journal of Economic Theory* **78**, 130-56.

[20] Vega-Redondo, F. (1997). The evolution of Walrasian behavior, *Econometrica* **65**, 375–384.