Pairwise Interact-and-Imitate Dynamics

Ennio Bilancini
IMT School for Advanced Studies Lucca, Italy
Leonardo Boncinelli (leonardo.boncinelli@unifi.it)
University of Florence
Nicola Campigotto
University of Pisa

Research Article

Keywords: evolutionary dynamics, pairwise interact-and-imitate dynamics

Posted Date: March 11th, 2021

DOI: https://doi.org/10.21203/rs.3.rs-244622/v1

Pairwise Interact-and-Imitate Dynamics

Ennio Bilancini[1], **Leonardo Boncinelli**[2,*], **and Nicola Campigotto**[3]

[1]IMT School for Advanced Studies Lucca, Laboratory for the Analysis of compleX Economic Systems, 55100 Lucca, Italy

[2]University of Florence, Department of Economics and Management, 50127 Florence, Italy

[3]University of Pisa, Department of Economics and Management, 56124 Pisa, Italy

[*]Corresponding author: leonardo.boncinelli@unifi.it

ABSTRACT

This paper introduces and studies a class of evolutionary dynamics — Pairwise Interact-and-Imitate Dynamics (PIID) — in which agents are matched in pairs, engage in a symmetric game, and imitate the opponent with a probability that depends on the difference in their payoffs. We provide a condition on the underlying game, named supremacy, and show that the population state in which all agents play the supreme strategy is globally asymptotically stable. We extend the framework to allow for payoff uncertainty, and check the robustness of our results to the introduction of some heterogeneity in the revision protocol followed by agents. Finally, we show that PIID can allow the survival of strictly dominated strategies, leads to the emergence of inefficient conventions in social dilemmas, and makes assortment ineffective in promoting cooperation.

Introduction

In evolutionary game-theoretic models, it is standard practice to assume that agents make decisions according to short-sighted adaptive rules. These include avoidance of strategies which performed poorly in the past, best response to the empirical distribution of opponents' strategies, and imitation of successful peers[1,2]. The last has been shown to be common in both human and animal groups, and is generally recognized as a cognitively parsimonious social heuristic[3–5].

Distinctions among imitative rules, in turn, can be made as to what drives behavior, who is imitated, and how much information is needed for choices to be made. For example, rules of the kind 'copy the first person you see' make actions only depend on their popularity, whereas other rules make decisions a function of observed payoffs. The target of comparisons may consist of either a single agent or a (possibly large) group of individuals, and information requirements can range from rather minimal to extremely high levels, yielding a taxonomy of imitative revision protocols[6–8]. The less demanding classes require agents to know the payoff of a single individual, either oneself or someone else, as in the case where agents change strategy depending on the difference between their payoff and a certain aspiration level[9]. Next come those revision protocols which require individuals to observe both their own payoff and the payoff of another agent; this class includes the family of rules considered in this paper. Finally, at the opposite end of the taxonomy are those protocols which require agents to know either the average payoff of all strategies in the population or the information necessary to compute average payoffs[10].

An important, but easily overlooked, aspect of imitative dynamics is the relationship between the structure of interactions and agents' reference groups. The interaction structure specifies how agents are matched, e.g., in a uniform random manner or according to other criteria. For example, in games played on graphs, interactions depend on the underlying topology of relations[11,12]; non-uniform random matching rules include type- and action-assortative rules, which reflect the tendency of individuals to interact with people who are similar to themselves in terms of observed characteristics and behavior[13–16]. An agent's reference group consists of those individuals whom that agent observes and takes as a reference for comparison purposes. Interaction structure and reference groups are frequently treated as separate entities: whenever an agent receives a revision opportunity, she randomly selects another individual, observes this individual's strategy, and switches to it with a probability that depends on relative payoffs[17–19]. A possible form of separation between interaction structure and reference groups is obtained in case of global random interactions and local comparisons based on observation of a fixed set of individuals[20,21]. This is most reasonable in the case of games against nature or when agents cannot observe their opponents' payoffs. However, cases also exist in which the decoupling of interaction structure and reference groups does not hold. Often people can only observe, and act upon, the behavior of those with whom they interact. This idea is recurrent in the literature on games on networks (even if there are exceptions[22]) where typically agents play with and imitate their immediate neighbors[23–25], or indirect neighbors[26,27]. The interplay between the structure of interactions and agents' reference groups may help understand the result that local interactions can favor the evolution of cooperation in the case of death-birth and imitative processes but not in the case of birth-death processes, which make defection emerge in the long run[28]. This finding can be seen as originating from differences in the degree of overlap between interaction structure and reference groups. Indeed, while in all cases agents only interact with their nearest neighbors, the target of comparisons differs from case to case — in birth-death processes, the set of nearest neighbors; in death-birth processes, the set of next-nearest neighbors; in imitative processes, a mixture of nearest and next-nearest neighbors. This can also be interpreted in terms of the inclusive fitness theory perspective, focusing on the interplay between genetic relatedness (a measure of assortment) and local competition brought about by limited dispersal (local imitation), in the presence of overlapping generations.

Building on this insight, this paper introduces and studies a class of evolutionary dynamics in which interaction structure and reference groups overlap, that is, where those whom one interacts with are also those with whom she compares herself. When given a revision opportunity, an agent playing strategy i against an opponent playing strategy j will switch to j with positive probability whenever the payoff from j against i is greater than the payoff from i against j. We name this protocol *Pairwise Interact-and-Imitate*.

Pairwise Interact-and-Imitate protocol is closely related to the local replicator dynamics[29], where individuals are uniformly matched at random in groups of size N to play pairwise games and imitate each other depending on payoff differences: for $N = 2$ a Pairwise Interact-and-Imitate is obtained for the special case of uniform random matching (while we consider more general interaction structures). Earlier works that relate most closely to our own include the pairwise comparative model of traffic dynamics, where changes from one route to another occur at a frequency that depends upon differences in traveling costs[30]. Another related revision protocol is one where agents imitate a randomly selected individual with a probability proportional to how much better the latter performed than them[17]. It is known in the literature that the set of rest points of the resulting imitative dynamics comprises the Nash equilibria of the underlying game and the states that would be a Nash equilibrium if the strategies unused at those states were removed from the game[7]. A possible variant concerns protocols that are direct instead of imitative, meaning that revising agents sample candidate strategies directly rather than by observing someone else's behavior[6,7]. The dynamics resulting from pairwise interactions with more informationally demanding imitative rules are also studied, for instance in a setup in which agents know others' payoffs, the probabilities with which individuals can meet, and the distribution of strategies in the population[31]. Another possibility is when agents evaluate alternative strategies by computing their average payoffs in a randomly sampled reference group, and then choose whichever strategy yielded the highest average payoff[32]. The imitate-the-best protocol, which prescribes copying the strategy chosen by the most successful individual, has been shown to be nearly unbeatable in a wide variety of symmetric two-player games[33]. The implications of the imitate-the-best rule have also been discussed in oligopoly settings[34]. Both imitate-the-best and imitate-the-best-average behaviors have been studied experimentally in the Cournot game[35].

The purpose of our paper is twofold. First, we introduce the Pairwise Interact-and-Imitate revision protocol and study the resulting dynamics in symmetric games. We give a condition on the stage game, named supremacy, and show that the population state in which all agents choose the supreme strategy is globally asymptotically stable. Roughly speaking, a strategy is supreme if it always yields a payoff higher than the payoff received by an opponent playing a different strategy. We then generalize the framework to allow for payoff uncertainty, we check the robustness of our results to the introduction of some heterogeneity in revision protocols, and we show that PIID can allow the survival of strictly dominated strategies. Second, we apply the revision protocol to social dilemmas, showing that PIID causes the emergence of inefficient conventions and makes assortment ineffective in facilitating cooperation.

Results

The model

Consider a unit-mass population of agents who repeatedly interact in pairs to play a symmetric stage game. The set of strategies available to each agent is finite and denoted by $S \equiv \{1,\ldots,n\}$. A population state is a vector $x \in X \equiv \{x \in \mathbb{R}^n_+ : \sum_{i\in S} x_i = 1\}$, with x_i the fraction of the population playing strategy $i \in S$. Payoffs are described by a function $F : S \times S \to \mathbb{R}$, where $F(i, j)$ is the payoff received by an agent playing strategy i when the opponent plays strategy j. As a shorthand, we refer to an undirected pair of individuals, one playing i and the other playing j, as an ij pair. The set of all possible pairs is denoted by $\mathscr{P}$.

The interaction structure is modeled as a function $p : X \times \mathscr{P} \to [0,1/2]$ subject to $\sum_{ij\in\mathscr{P}} p_{ij}(x) = 1/2$ (since the mass of pairs is half the mass of agents), with $p_{ij}(x)$ indicating the mass of ij pairs formed in state x. Note that the mass of ij pairs can never exceed $\min\{x_i,x_j\}$, that is, $p_{ij}(x) \leq \min\{x_i,x_j\}$ for all x. We assume that p is continuous in X and that if $x_i > 0$ and $x_j > 0$, then $p_{ij}(x) > 0$ — meaning that if in a certain state strategies i and j are played by someone, then the probability of an ij pair being formed is strictly positive. In the case of uniform random matching, $p_{ii} = x_i^2/2$ and $p_{ij} = x_i x_j$ for any i and $j \neq i$.

The revision protocol is modeled as a function $\phi : X \times \mathscr{P} \to [-1,1]$, where $\phi_{ij}(x) \in [-1,1]$ is the probability that an ij pair will turn into an ii pair minus the probability that it will turn into a jj pair, conditional on the population state being x and an ij pair being formed. We assume that ϕ is continuous in X, and note that $\phi_{ii} = 0$ for all $i \in S$. Our main assumption on the revision protocol is the following, which is met, among others, by pairwise proportional imitative and imitate-if-better rules[17].

Assumption 1. *For every $x \in X$, $\phi_{ij}(x) > 0$ if $F(i, j) > F(j,i)$.*

In what follows we consider a dynamical system in continuous time with state space X, characterized by the following equation of motion.

Definition 1 (Pairwise Interact-and-Imitate Dynamics – PIID). *For every $x \in X$ and every $i \in S$:*

$$\dot{x}_i = \sum_{j\in S} p_{ij}(x)\phi_{ij}(x). \tag{1}$$

Main findings

Global asymptotic convergence

In any purely imitative dynamics, if $x_i(t) = 0$, then $x_i(t') = 0$ for every $t' > t$. This implies that we cannot hope for global asymptotic convergence in a strict sense. Thus, to assess convergence towards a certain state x in a meaningful way, we restrict our attention to those states where all strategies that have positive frequency in x have positive frequency as well. We denote by X_x the set of states whose support contains the support of x.

Definition 2 (Supremacy). *Strategy $i \in S$ is supreme if $F(i, j) > F(j,i)$ for every $j \in S\setminus\{i\}$.*

Equivalently, a strategy i is supreme if the relative payoff from playing i against any other strategy[33] is positive, i.e., if $F(i,j) - F(j,i) > 0$ for all $j \neq i$. We note that, under PIID, the notion of supremacy is closely related to that of asymmetry[36,37], in that $F(i,j) > F(j,i)$ implies that agents can only switch from strategy j to strategy i.

Proposition 1. *If $i \in S$ is a supreme strategy, then state $x^* \equiv \{x \in X : x_i = 1\}$ is globally asymptotically stable for the dynamical system with state space X_{x^*} and PIID as equation of motion.*

Relation to replicator dynamics

To further characterize the dynamics induced by the pairwise interact-and-imitate protocol, we make two additional assumptions. First, matching is uniformly random, meaning that everyone in the population has the same probability of interacting with everyone else. Second, the probability that an agent has to imitate the opponent is proportional to the difference in their payoffs if the opponent's payoff exceeds her own, and is zero otherwise.

Let

$F(i,x) := \sum_j x_j F(i,j)$,

$F(x,i) := \sum_j x_j F(j,i)$, and

$F(x,x) := \sum_i \sum_j x_i x_j F(i,j)$.

Under these assumptions, at any point in time, the motion of x_i is described by:

$$\dot{x}_i = \sum_{j \neq i} x_j x_i \left[F(i,j) - F(j,i)\right] = x_i \sum_j x_j \left[F(i,j) - F(j,i)\right]$$
$$= x_i \left[F(i,x) - F(x,i)\right], \tag{2}$$

which is a modified replicator equation. According to (2), for every strategy i chosen by one or more agents in the population, the rate of growth of the fraction of i-players, $\dot{x}_i/x_i$, equals the difference between the expected payoff from playing i in state x and the average payoff received by those who are matched against an agent playing i. In contrast, under the standard replicator dynamics[38], the fraction of agents playing i varies depending on the excess payoff of i with respect to the current average payoff in the whole population, i.e., $\dot{x}_i = x_i \left[F(i,x) - F(x,x)\right]$.

A noteworthy feature of the replicator dynamics is that it is always payoff monotone: for any $i, j \in S$, the proportions of agents playing i and j grow at rates that are ordered in the same way as the expected payoffs from the two strategies[39]. In the case of PIID, this result fails. To see this, simply consider any symmetric 2×2 game where $F(i,j) > F(j,i)$ but $F(j,x) > F(i,x)$ for some $x \in X$, meaning that i is the supreme strategy but j yields a higher expected payoff in state x. If strategies are updated according to the interact-and-imitate protocol, this game only admits switches from i to j, therefore violating payoff monotonicity. This can have important consequences, including the survival of pure strategies that are strictly dominated.

Survival of strictly dominated strategies

An important topic in evolutionary game theory is to what extent does support exist for the idea that strictly dominated strategies will not be played. It has been shown that if a strategy i does not survive the iterated elimination of pure strategies strictly dominated by other pure strategies, then the fraction of the population playing i converges to zero in all payoff monotone dynamics[40,41]. This result does not apply in our case, as PIID is not payoff monotone.

More precisely, under PIID, a strictly dominated strategy may be supreme and, therefore, not only survive but even end up being adopted by the whole population. This suggests that, from an evolutionary perspective, support for the elimination of dominated strategies may be weaker than is often thought. Our result contributes to the literature studying conditions under which evolutionary dynamics fail to eliminate strictly dominated strategies in some games, considering a case which was not previously studied by the literature[42].

To see that a strictly dominated strategy may be supreme, consider the simple example shown in Figure 1. Here each agent has a strictly dominant strategy to play A; however, since the payoff from playing B against A exceeds that from playing A against B, strategy B is supreme. Thus, by Proposition 1, the population state in which all agents choose B is globally asymptotically stable.

	A	B
A	4	2
B	3	1

Figure 1. A game where the supreme strategy is strictly dominated.

Applications

Having obtained general results for the class of finite symmetric games, we now restrict the discussion to the evolution of behavior in social dilemmas. We show that, under the conditions of Proposition 1, inefficient conventions emerge in the Prisoner's Dilemma, Stag Hunt, Minimum Effort, and Hawk-Dove games. Furthermore, these results hold both without and with the assumption that agents interact assortatively.

Ineffectiveness of assortment

Consider the 2×2 game represented in Figure 2. If $c > a > d > b$, then mutual cooperation is Pareto superior to mutual defection, but agents have a dominant strategy to defect. The resulting stage game is the Prisoner's Dilemma, whose unique Nash equilibrium is (B,B). Moreover, since $F(B,A) > F(A,B)$, B is the supreme strategy and the population state in which all agents defect is globally asymptotically stable.

We stress that defection emerges in the long run for every matching rule satisfying our assumptions, and therefore also in the case of assortative interactions. Assortment reflects the tendency of similar people to clump together, and can play an important role in the evolution of cooperation[43–46]. Intuitively, when agents meet assortatively, the risk of cooperating in a social dilemma may be offset by a higher probability of playing against other cooperators. However, under PIID, this is not the case; the decision whether to adopt a strategy or not is independent of expected payoffs, and like-with-like interactions have no effect other than reducing the frequency of switches from A to B.

	A	B
A	a	b
B	c	d

Figure 2. A 2×2 stage game.

Emergence of the maximin convention

If $a > c > b$, $a > d$, and $d > b$, then the game in Figure 2 becomes a Stag Hunt game, which contrasts risky cooperation and safe individualism. The payoffs are such that both (A,A) and (B,B) are strict Nash equilibria, that (A,A) is Pareto superior to (B,B), and that B is the maximin strategy, i.e., the strategy which maximizes the minimum payoff an agent could possibly receive. We also assume that $a + c \neq c + d$, so that one of A and B is risk dominant[47]. If $a + b > c + d$, then A (Stag) is both payoff and risk dominant. When the opposite inequality holds, the risk dominant strategy is B (Hare).

Since $F(B,A) > F(A,B)$, B is supreme independently of whether or not it is risk dominant to cooperate. This can result in large inefficiencies because, in the long run, the process will converge to the state in which all agents play the riskless strategy regardless of how rewarding social coordination is. As in the case of the Prisoner's Dilemma, this holds for all matching rules satisfying our assumptions.

Evolution of effort exertion

In a minimum effort stage game, agents simultaneously choose a strategy i, usually interpreted as a costly effort level, from a finite subset S of $\mathbb{R}$. An agent's payoff depends on her own effort and on the minimum effort in the pair:

$$F(i,j) = \alpha \min\{i,j\} - \beta i,$$

where $\beta > 0$ and $\alpha > \beta$ are the cost of and the benefit from effort, respectively. The best response to a choice of j by the opponent is to choose j as well, and coordinating on any common effort level gives a Nash equilibrium. Nash outcomes can be Pareto-ranked, with the highest-effort equilibrium being the best possible outcome for all agents. Thus, choosing a high i is rationalizable and potentially rewarding but may also result in a waste of effort.

Under PIID, any $i > j$ implies $\phi_{ij} < 0$ by Assumption 1, meaning that agents will tend to imitate the opponent when the opponent's effort is lower than their own. The supreme strategy is therefore to exert as little effort as possible, and the population state in which all agents choose the minimum effort level is the unique globally asymptotically stable state.

Emergence of aggressive behaviors

Consider again the payoff matrix shown in Figure 2. If $c > a > b > d$, then the stage game is a Hawk-Dove game, which is often used to model the evolution of aggressive and sharing behaviors. Interactions can be framed as fights over a contested resource. When two Doves (who play A) meet, they share the resource equally, whereas two Hawks (who play B) engage in a brawl and suffer a cost. Also, when a Dove meets a Hawk, the latter takes the entire prize. Again we have that $F(A,B) < F(B,A)$, implying that B is the supreme strategy and that All Hawk is the sole asymptotically stable state.

The inefficiency that characterizes the (B,B) equilibrium in the Hawk-Dove game arises from the cost that Hawks impose on one another. This can be viewed as stemming from the fact that neither agent owns the resource prior to the interaction or cares about property. A way to overcome this problem may be to introduce a strategy associated with respect for ownership rights, the Bourgeois, who behaves as a Dove or Hawk depending on whether or not the opponent owns the resource[48]. If we make the standard assumption that each member of a pair has a probability of $1/2$ to be an owner, then in all interactions where a Bourgeois is involved there is a 50 percent chance that she will behave hawkishly (i.e., fight for control over the resource) and a 50 percent chance that she will act as a Dove.

Let R and C denote the agent chosen as row and column player, respectively, and let ω_R and ω_C be the states of the world in which R and C owns the resource. The payoffs for the resulting Hawk-Dove-Bourgeois game are shown in Figure 3. If agents behave as expected payoff maximizers, then All Bourgeois can be singled out as the unique asymptotically stable state. Under PIID, this is not so; depending on who owns the resource, an agent playing C against an opponent playing B may either fight or avoid conflict and let the opponent have the prize. It is easy to see that $F(C,B \mid \omega_R) = F(B,C \mid \omega_C) = d$, meaning that the payoff from playing C against B, conditional on owning the resource, equals the payoff from playing B against C conditional on not being an owner. In contrast, the payoff from playing C against B, conditional on not owning the resource, is always worse than that of the opponent, i.e., $F(C,B \mid \omega_C) = b < c = F(B,C \mid \omega_R)$. Thus, in every state of the world, B (Hawk) yields a payoff that is greater or equal to that from C (Bourgeois).

Moreover, since $F(B,A) > F(A,B)$ in both states of the world, strategy B is weakly supreme by Definition 4, and play unfolds as an escalation of hawkishness and fights.

	A	B	C
A	a	b	a
B	c	d	c
C	c	d	c

(a) Row player owns the resource.

	A	B	C
A	a	b	b
B	c	d	d
C	a	b	b

(b) Column player owns the resource.

Figure 3. The Hawk-Dove-Bourgeois game.

Discussion

We have studied a novel class of evolutionary dynamics, named Pairwise Interact-and-Imitate Dynamics, in which agents choose whether or not to change strategy by comparing their payoff with that of their opponent. Our main result is that, under PIID, if there exists a supreme strategy (that is, a strategy that always yields a payoff higher than the payoff received by an opponent playing a different strategy), then the state in which the whole population chooses the supreme strategy is globally asymptotically stable. Importantly, the supreme strategy may be a dominated strategy, and the profile of strategies played in the asymptotically stable state may not be a Nash equilibrium.

Under PIID, externalities have an important role. Whenever a strategy, say i, generates an increase (decrease) in the payoff of an opponent playing a different strategy, say j, it is more (less) likely that i will be updated in favor of j. Moreover, this is so regardless of the payoff received when using the same strategy as the opponent. Thus, ceteris paribus, strategies that impose negative (positive) externalities are more (less) likely to be selected by evolution, possibly leading to inefficient outcomes.

However, it is worth noting that PIID does not necessarily lead to inferior outcomes with respect to other evolutionary dynamics. The simple example of Figure 4 shows this. For instance, under Pairwise Proportional Imitation[7,17], if the fraction of agents playing A is sufficiently large, then the system moves towards the state where the whole population plays A — which, however, is Pareto dominated by everyone playing B.

Conversely, under PIID, the system always moves towards the state where everyone plays B (since B is supreme). This result holds in general, i.e., even without the additional assumptions required to represent evolution by means of the modified replicator equation.

	A	B
A	2	0
B	1	3

Figure 4. A game where PIID selects an efficient outcome.

Overall, our findings provide a case for why individual behaviors may direct evolution towards outcomes that do not meet Pareto efficiency and strategy dominance criteria. Rather, our dynamics depend on which strategy, if any, is supreme, i.e., it systematically outperforms other strategies when these are chosen by one's opponents. This implies that the outcome of evolution can be either very undesirable or very desirable, depending on how large the payoff from the supreme strategy is when this strategy is chosen by everyone in the population. Moreover, since the structure of interactions plays no role in determining which strategy is supreme, the long-run equilibrium selected by PIID is not affected by the presence of institutions and other factors that influence how agents interact, such as those generating assortment.

We have shown that when applied to study the evolution of behavior, PIID leads to clear-cut and sometimes surprising results in a large variety of games. However, not all classes of games are suited to our revision protocol; in this paper we have only considered symmetric games, leaving aside those cases where agents can choose among different strategies or have different payoff functions. When agents' strategy sets differ from one another, it does not seem very reasonable to assume that choices are updated according to a pairwise imitative rule based on payoff differentials. However, we believe that an imitative protocol like ours may still be applied in a meaningful way to those cases in which agents have the same strategy set but differ in some other respect. For instance, in a setting where agents differ in wealth, a poor individual may be driven to imitate the strategy chosen by a rich individual earning a high payoff, even if this is due to differences in wealth rather than in strategy.

An extension of the model developed here would be to consider the case of a finite population of agents. This would make comparisons with some of the literature easier[29], but would come at the cost of hindering the analysis when introducing payoff uncertainty and studying how PIID relates to the replicator dynamics. Another extension would concern moving from two-player to n-player symmetric games, which would require defining the class of Groupwise Interact-and-Imitate Dynamics, beyond adjusting the notion of supremacy to consider the relative performance of a strategy towards profiles of others' strategies.

Finally, a question that may be worthy of further investigation is how the dynamics will behave when no supreme strategy exists. To answer, one may define a binary relation $\succcurlyeq$ such that $i \succcurlyeq j$ if and only if $F(i,j) \geq F(j,i)$. One may then define $\succcurlyeq^*$ as the transitive closure of $\succcurlyeq$, and let $\mathscr{S} := \{i \in S : i \succcurlyeq^* j \; \forall \; j \in S\}$ be set of supremal strategies. Our conjecture is that, under PIID, all strategies that do not belong to $\mathscr{S}$ will die out independently of the structure of interactions; however, the precise characterization of limit sets may depend on details of the payoff structure, the interaction structure, and the revision protocol.

Methods

Lyapunov's method

To prove Proposition 1 we use Lyapunov's second method for global stability. We want to show that $f : X_{x^*} \to [0,1)$, with

$$f(x) = 1 - x_i,$$

is a strict Lyapunov function. It is easy to see that f is of class C^1, that $f(x) > 0$ for every $x \in X_{x^*} \setminus \{x^*\}$, and that $f(x^*) = 0$. We are left to show that (i) $\dot{f}(x^*) = 0$ and (ii) $\dot{f}(x) < 0$ for every $x \in X_{x^*} \setminus \{x^*\}$. Taking the time derivative of f and using Definition 1, we can write:

$$\dot{f}(x) = \frac{\partial f(x)}{\partial x_i}\dot{x}_i = -\sum_{j \in S} p_{ij}(x)\phi_{ij}(x).$$

We observe that $p_{ij}(x^*) = 0$ for every $j \in S$, which implies $\dot{f}(x^*) = 0$. For every $x \in X_{x^*} \setminus \{x^*\}$, there exists $k \neq i$ such that $x_k > 0$. Since $x_i > 0$ and $x_k > 0$, we therefore have that $p_{ik}(x) > 0$. Moreover, since $F(i,j) > F(j,i)$ for every $j \in S \setminus \{i\}$, we have that $\phi_{ik}(x) > 0$ by Assumption 1. To see that $\dot{f}(x) < 0$ for every $x \in X_{x^*} \setminus \{x^*\}$, we finally note that $p_{ij}(x)$ is always non-negative and that $\phi_{ij}(x)$ is positive for all $j \in S \setminus \{i\}$ by the definition of supremacy and Assumption 1.

Supremacy under uncertainty

So far we have only considered games in which payoffs are not subject to any uncertainty in their realization. Here we extend our analysis by allowing for this possibility, which is relevant for a variety of applications (e.g., the Hawk-Dove and Hawk-Dove-Bourgeois games).

Let Ω be a finite set of states of the world and let q_ω be the probability of state $\omega \in \Omega$ occurring. We write $F(i,j|\omega)$ to denote the payoff received by an agent playing strategy i against an opponent playing j when state of the world ω occurs. For every $x \in X$ and every $\omega \in \Omega$, let $\phi(x|\omega)$ be the $n \times n$-matrix with typical element $\phi_{ij}(x|\omega)$, the latter being the net inflow from j to i when the population state is x, the state of the world is ω, and an ij pair is formed.

We replace Assumption 1 with the following.

Assumption 2. *For every $x \in X$ and every $\omega \in \Omega$, $\phi_{ij}(x|\omega) > 0$ if $F(i,j|\omega) > F(j,i|\omega)$.*

We also assume that, at any point in time, the net population inflow from j to i is obtained by averaging ϕ_{ij} over the set of states of the world, i.e.:

$$\phi_{ij}(x) = \sum_{\omega \in \Omega} q_\omega \phi_{ij}(x|\omega). \tag{3}$$

The convergence result of Proposition 1 can be extended to the case of uncertainty if we define supremacy as follows.

Definition 3 (Supremacy under uncertainty)**.** *In the presence of uncertainty, strategy $i \in S$ is supreme if $F(i,j|\omega) > F(j,i|\omega)$ for every $\omega \in \Omega$ and every $j \in S \setminus \{i\}$.*

To see that the result holds, note that if strategy i is supreme by Definition 3 and Assumption 2 holds, then $\phi_{ij}(x) > 0$ for every $j \in S \setminus \{i\}$ and every $x \in X$. Moreover, ϕ is continuous in X if $\phi(\cdot,\omega)$ is continuous in X for every $\omega \in \Omega$. By a reasoning analogous to that used in the proof of Proposition 1, we therefore have that state $\{x \in X : x_i = 1\}$ is globally asymptotically stable for the dynamical system with state space X_{x^*} and PIID as equation of motion.

A less restrictive definition of supremacy under uncertainty is given below.

Definition 4 (Weak supremacy under uncertainty)**.** *In the presence of uncertainty, strategy $i \in S$ is weakly supreme if $F(i,j|\omega) \geq F(j,i|\omega)$ for every $\omega \in \Omega$ and every $j \in S \setminus \{i\}$, and if for every $j \in S \setminus \{i\}$ there exists $\hat{\omega} \in \Omega$ such that $F(i,j|\hat{\omega}) > F(j,i|\hat{\omega})$.*

Under the conditions of Definition 4, our convergence result holds if we strengthen Assumption 2 as follows.

Assumption 3. *For every $x \in X$ and every $\omega \in \Omega$, $\phi_{ij}(x|\omega) > 0$ if and only if $F(i,j|\omega) > F(j,i|\omega)$.*

Here the 'only if' is required to deal with those states of the world where $F(i,j|\omega) = F(j,i|\omega)$, which are not ruled out by Definition 4.

An even weaker definition of supremacy can be given when focusing on a specific ϕ. For instance, consider the case where the probability that agents have to imitate the opponent is proportional to the difference in their payoffs if the opponent's payoff exceeds their own, and is zero otherwise. Under this protocol, letting the expected payoff from playing i against j be $\mathbb{E}[F(i,j)] = \sum_{\omega \in \Omega} q_\omega F(i,j|\omega)$, we can define the following.

Definition 5 (Supremacy in expectation under uncertainty)**.** *In the presence of uncertainty, strategy $i \in S$ is supreme in expectation if $\mathbb{E}[F(i,j)] > \mathbb{E}[F(j,i)]$ for every $j \in S \setminus \{i\}$.*

It can now be seen that:

$$\phi_{ij}(x) > 0 \Leftrightarrow \sum_{\omega \in \Omega} q_w [F(i,j \mid \omega) - F(j,i \mid \omega)] > 0 \Leftrightarrow \mathbb{E}[F(i,j)] > \mathbb{E}[F(j,i)],$$

meaning that the net population inflow from j to i is positive if and only if i is supreme in expectation. This suffices to replicate the result of Proposition 1.

Heterogeneous revision protocols

Although we believe the interact-and-imitate protocol to be reasonable in many circumstances, it may well be the case that agents also rely on other revision protocols occasionally. If our results crucially hinged on the assumption that agents always follow the interact-and-imitate rule, they would be of little interest.

Let the set of possible states of the world be $\Omega = \{\omega_1, \omega_2\}$. Suppose that agents follow the pairwise interact-and-imitate protocol in state ω_1 and a different revision protocol in state ω_2. Now let us define a continuous function $\rho : X \times \mathscr{P} \to [-1,1]$, where $\rho_{ij}(x) \in [-1,1]$ is the probability that an ij pair will turn into an ii pair minus the probability that it will turn into a jj pair, conditional on the population state being x, an ij pair being formed, and the state of the world being ω_2. Note that ρ may also reflect the fact that members of a pair interact after, rather than before, having updated their strategies, in which case the interact-and-imitate protocol cannot be applied.

We define:

$$\hat{\phi}_{ij}(x) := (1-\varepsilon)\phi_{ij}(x) + \varepsilon\rho_{ij}(x), \tag{4}$$

with $\varepsilon \in (0,1)$. The equation of motion for this dynamical system is the following.

Definition 6 (Quasi pairwise interact-and-imitate dynamics – QPIID)**.** *For every $x \in X$ and every $i \in S$:*

$$\dot{x}_i = \sum_{j \in S} p_{ij}(x)\hat{\phi}_{ij}(x). \tag{5}$$

Note that Assumption 1 concerns ϕ only, and that we have no analogous assumption for ρ. This notwithstanding, a convergence result in the spirit of Proposition 1 can be obtained if agents follow the ρ protocol rarely enough.

Proposition 2. *If $i \in S$ is supreme, then there exists $\bar{\varepsilon} > 0$ such that, for every $\varepsilon \in (0,\bar{\varepsilon})$, state $x^* \equiv \{x \in X : x_i = 1\}$ is globally asymptotically stable for the dynamical system with state space X_{x^*} and QPIID as equation of motion.*

We want to show that if $F(i,j) > F(j,i)$ for every $j \in S \setminus \{i\}$, then there exists $\bar{\varepsilon} > 0$ such that, for every $\varepsilon \in (0,\bar{\varepsilon})$, $\hat{\phi}(x) > 0$ for every $x \in X$. Moreover, since both ϕ and ρ are continuous in X, $\hat{\phi}$ is continuous in X as well. As a consequence, the statement in Proposition 2 can be proven by replicating the argument used in the proof of Proposition 1.

We consider the worst case to have $\hat{\phi}(x)$ positive. For every $j \in S \setminus \{i\}$, we have that $\bar{\phi}_{ij} := \min_{x \in X} \phi_{ij}(x)$ exists, since ϕ_{ij} is a continuous function on a compact set. Moreover, by continuity of ϕ_{ij} and noting that $\phi_{ij}(x) > 0$ due to $F(i,j) > F(j,i)$ and Assumption 1, we also have that $\bar{\phi}_{ij} > 0$. We define $\bar{\phi}_i := \min_{j \neq i} \bar{\phi}_{ij}$, and note that $\bar{\phi}_i > 0$. Also, we define:

$$\bar{\varepsilon} := \frac{\bar{\phi}_i}{1+\bar{\phi}_i}. \tag{6}$$

Since ρ_{ij} cannot be smaller than -1, if $\varepsilon < \bar{\varepsilon}$ then we have that $\hat{\phi}(x) > 0$ for every $x \in X$, as can be easily checked from (4).

We stress that the bound on ε given in (6) is independent of which ρ protocol is being considered. A more precise value for the bound can be obtained by making specific assumptions about the strategy revision process. For example, suppose that in state of the world ω_1, agents imitate the opponent with unit probability 1 whenever the latter receives a higher payoff than they do (so that $\phi_{ij} = 1$ if $F(i,j) > F(j,i)$). In this case, we can have the maximum possible amount of heterogeneity in revision protocols (i.e., $\bar{\varepsilon} = 1/2$) and still have global asymptotic stability of the state in which the whole population chooses the supreme strategy.

References

1. Fudenberg, D. & Levine, D. K. *The theory of learning in games* (MIT press, 1998).

2. Young, H. P. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions* (Princeton University Press, 1998).

3. Pingle, M. & Day, R. H. Modes of economizing behavior: Experimental evidence. *J. Econ. Behav. & Organ.* **29**, 191–209 (1996).

4. Rendell, L. *et al.* Why copy others? insights from the social learning strategies tournament. *Science* **328**, 208–213 (2010).

5. Gigerenzer, G. & Gaissmaier, W. Heuristic decision making. *Annu. Rev. Psychol.* **62**, 451–482 (2011).

6. Sandholm, W. Pairwise comparison dynamics and evolutionary foundations for nash equilibrium. *Games* **1**, 3–17 (2010).

7. Sandholm, W. H. *Population Games and Evolutionary Dynamics* (MIT press, 2010).

8. Newton, J. Evolutionary game theory: a renaissance. *Games* **9**, 31 (2018).

9. Bjrönerstedt, J. & Weibull, J. W. Nash equilibrium and evolution by imitation. In Arrow, K., Colombatto, E., Perlman, M. & Schmidt, C. (eds.) *The Rational Foundations of Economic Behavior* (Macmillan, 1996).

10. Robson, A. J. & Vega-Redondo, F. Efficient equilibrium selection in evolutionary games with random matching. *J. Econ. Theory* **70**, 65–92 (1996).

11. Szabó, G. & Fáth, G. Evolutionary games on graphs. *Phys. Reports* **446**, 97–216 (2007).

12. Jackson, M. O. & Zenou, Y. Games on networks. In Young, H. P. & Zamir, S. (eds.) *Handbook of Game Theory, Volume 4*, 95–164 (North Holland, 2014).

13. Alger, I. & Weibull, J. W. Homo moralis—preference evolution under incomplete information and assortative matching. *Econometrica* **81**, 2269–2302 (2013).

14. Alger, I. & Weibull, J. W. Evolution and kantian morality. *Games Econ. Behav.* **98**, 56–67 (2016).

15. Bilancini, E., Boncinelli, L. & Wu, J. The interplay of cultural intolerance and action-assortativity for the emergence of cooperation and homophily. *Eur. Econ. Rev.* **102**, 1–18 (2018).

16. Newton, J. The preferences of homo moralis are unstable under evolving assortativity. *Int. J. Game Theory* **46**, 583–589 (2017).

17. Schlag, K. H. Why imitate, and if so, how? a boundedly rational approach to multi-armed bandits. *J. Econ. Theory* **78**, 130–156 (1998).

18. Alós-Ferrer, C. & Schlag, K. H. Imitation and learning. In Anand, P., Pattanaik, P. & Puppe, C. (eds.) *Handbook of Rational and Social Choice* (Oxford University Press, 2009).

19. Izquierdo, S. S. & Izquierdo, L. R. Stochastic approximation to understand simple simulation models. *J. Stat. Phys.* **151**, 254–276 (2013).

20. Khan, A. Coordination under global random interaction and local imitation. *Int. J. Game Theory* **43**, 721–745 (2014).

21. Cui, Z. & Wang, R. Collaboration in networks with randomly chosen agents. *J. Econ. Behav. & Organ.* **129**, 129–141 (2016).

22. Ohtsuki, H., Pacheco, J. M. & Nowak, M. A. Evolutionary graph theory: Breaking the symmetry between interaction and replacement. *J. Theor. Biol.* **246**, 681–694 (2007).

23. Fosco, C. & Mengel, F. Cooperation through imitation and exclusion in networks. *J. Econ. Dyn. Control.* **35**, 641–658 (2011).

24. Tsakas, N. Imitating the most successful neighbor in social networks. *Rev. Netw. Econ.* **12**, 403–435 (2014).

25. Cimini, G. Evolutionary network games: Equilibria from imitation and best response dynamics. *Complexity* **2017**, 1–14 (2017).

26. Alós-Ferrer, C. & Weidenholzer, S. Contagion and efficiency. *J. Econ. Theory* **143**, 251–274 (2008).

27. Alós-Ferrer, C. & Weidenholzer, S. Imitation and the role of information in overcoming coordination failures. *Games Econ. Behav.* **87**, 397–411 (2014).

28. Ohtsuki, H., Hauert, C., Lieberman, E. & Nowak, M. A. A simple rule for the evolution of cooperation on graphs and social networks. *Nature* **328**, 502–505 (2006).

29. Hilbe, C. Local replicator dynamics: a simple link between deterministic and stochastic models of evolutionary game theory. *Bull. Math. Biol.* **73**, 2068–2087 (2011).

30. Smith, M. J. The stability of a dynamic model of traffic assignment—an application of a method of lyapunov. *Transp. Sci.* **18**, 245–252 (1984).

31. Helbing, D. A mathematical model for behavioral changes by pair interactions. In Haag, G., Mueller, U. & Troitzsch, K. G. (eds.) *Economic Evolution and Demographic Change: Formal Models in Social Sciences*, 330–348 (Springer, 1992).

32. Ellison, G. & Fudenberg, D. Word-of-mouth communication and social learning. *Q. J. Econ.* **110**, 93–125 (1995).

33. Duersch, P., Oechssler, J. & Schipper, B. C. Unbeatable imitation. *Games Econ. Behav.* **76**, 88–96 (2012).

34. Vega-Redondo, F. The evolution of walrasian behavior. *Econometrica* **65**, 375–384 (1997).

35. Apesteguia, J., Huck, S. & Oechssler, J. Imitation—theory and experimental evidence. *J. Econ. Theory* **136**, 217–235 (2007).

36. Peski, M. Generalized risk-dominance and asymmetric dynamics. *J. Econ. Theory* **145**, 216–248 (2010).

37. Newton, J. Conventions under heterogeneous choice rules (2019). Forthcoming in *Rev. Econ. Stud*. Available at SSRN: https://ssrn.com/abstract=3383471.

38. Taylor, P. D. & Jonker, L. B. Evolutionary stable strategies and game dynamics. *Math. Biosci.* **40**, 145–156 (1978).

39. Weibull, J. W. *Evolutionary Game Theory* (MIT Press, 1995).

40. Nachbar, J. H. "evolutionary" selection dynamics in games: Convergence and limit properties. *Int. J. Game Theory* **19**, 59–89 (1990).

41. Samuelson, L. & Zhang, J. Evolutionary stability in asymmetric games. *J. Econ. Theory* **57**, 363–391 (1992).

42. Hofbauer, J. & Sandholm, W. H. Survival of dominated strategies under evolutionary dynamics. *Theor. Econ.* **6**, 341–377 (2011).

43. Eshel, I. & Cavalli-Sforza, L. L. Assortment of encounters and evolution of cooperativeness. *Proc. Natl. Acad. Sci.* **79**, 1331–1335 (1982).

44. Bergstrom, T. C. The algebra of assortative encounters and the evolution of cooperation. *Int. Game Theory Rev.* **5**, 211–228 (2003).

45. Bergstrom, T. C. Measures of assortativity. *Biol. Theory* **8**, 133–141 (2013).

46. Allen, B. & Nowak, M. A. Games among relatives revisited. *J. Theor. Biol.* **378**, 103–116 (2015).

47. Harsanyi, J. C. & Selten, R. *A General Theory of Equilibrium Selection in Games* (MIT Press, 1988).

48. Maynard Smith, J. *Evolution and the Theory of Games* (Cambridge University Press, 1982).

Acknowledgements

The authors gratefully acknowledge financial support from the Italian Ministry of Education, University and Research (MIUR) through the PRIN project Co.S.Mo.Pro.Be. "Cognition, Social Motives and Prosocial Behavior" (grant n. 20178293XT) and from the IMT School for Advanced Studies Lucca through the PAI project Pro.Co.P.E. "Prosociality, Cognition, and Peer Effects".

Author contributions statement

E.B., L.B., and N.C. contributed equally to setting the model, developing theoretical results and finding applications, as well as to writing the paper.

Additional information

The authors declare no competing interests.

Figures

	A	B
A	4	2
B	3	1

Figure 1

A game where the supreme strategy is strictly dominated.

	A	B
A	a	b
B	c	d

Figure 2

A 2x2 stage game.

	A	B	C
A	a	b	a
B	c	d	c
C	c	d	c

(a) Row player owns the resource.

	A	B	C
A	a	b	b
B	c	d	d
C	a	b	b

(b) Column player owns the resource.

Figure 3

The Hawk-Dove-Bourgeois game.

	A	B
A	2	0
B	1	3

Figure 4

A game where PIID selects an efficient outcome. 248