

Why Learning Doesn't Add Up: Equilibrium
Selection with a Composition of Learning Rules

Russell Golman*

April 8, 2010

*Department of Applied and Interdisciplinary Mathematics, University of Michigan, 2082 East Hall, 530 Church

Street, Ann Arbor, MI 48109, USA. Telephone: 734 764-7436. Fax: 734 763-0937. E-mail: rgolman@umich.edu

Abstract

In this paper, we investigate the aggregate behavior of populations of learning agents. We compare the outcomes in homogenous populations learning in accordance with imitate the best dynamics and with replicator dynamics to outcomes in populations that mix these two learning rules. New outcomes can emerge. In certain games, a linear combination of the two rules almost always attains an equilibrium that homogenous learners almost never locate. Moreover, even when almost all weight is placed on one learning rule, the outcome can differ from homogenous use of that rule. Thus, allowing even an arbitrarily small chance of using an alternative learning style can shift a population to select a different equilibrium.

KEYWORDS: Adjustment dynamics, basins of attraction, best response dynamics, equilibrium selection, evolutionary game, imitate the best dynamics, learning, replicator dynamics.

JEL classification: C72, C73

1 Introduction

Models of individual learning in games often produce the same dynamical systems as population-based evolutionary adjustment processes. For example, an appropriate form of reinforcement learning leads to the replicator dynamics (Borgers and Sarin 1997), originally proposed by theoretical biologists to describe the growth of a haploid population (Taylor and Jonker 1978). Similarly, the best response dynamics (Gilboa and Matsui 1991) can arise as the continuous time limit of a belief-based learning rule like fictitious play or in a population of myopic best responders. The models that lead to the replicator dynamics and those that lead to the best response dynamics make mutually incompatible assumptions about how strategies are updated, offering no middle ground.

In this paper, we adopt a population-based framework and propose combining an imitative variant of best response dynamics (Hofbauer, 1995) with the replicator dynamics to form a tractable composite rule. We find that a population is sensitive to the presence of more than one learning style, as well as to how the learning rules are put together, in that composite learning dynamics may have basins of attraction that share little overlap with each other or with the homogeneous learning dynamics. We focus in particular on a straightforward linear combination of imitate-the-best and replicator learning rules, which we show to select an equilibrium not attained by homogenous use of either pure rule in certain matrix games. Introducing even arbitrarily small weight on an alternative learning rule may cause the resulting linear

combination of rules to locate a different equilibrium than the pre-existing pure rule. A feature of the class of games we consider is the existence of a temporary initial best response – an action that looks good at first, but that cannot survive against itself. The temptation to try such an action instead of an equilibrium action creates complex dynamics, which allow the effects of subtle differences in learning styles to build up over time.

Learning in games, at the population level, can be modeled by dynamic adjustment processes (Swinkels 1993, Fudenberg and Levine 1998). In imitate the best dynamics (as with best response dynamics), players slowly adopt myopic best responses, ignoring the possibility that the population mixed strategy may soon change. For a matrix game, the resulting flows are piecewise linear in the direction of the best response to the current population state. In replicator dynamics, players copy others' strategies, so more successful actions spread in the population.¹ These population dynamics can be derived from the revision protocols that describe how the individual agents learn.

In this paper, we combine revision protocols that lead to imitate the best and replicator dynamics respectively, forming a rule in which either learning protocol may be tried with some positive probability. We compare equilibrium selection following from the composition of learning styles and from pure imitate the best and replicator dynamics, and we find they may be markedly different from each other in their long-run behavior. For a specific class of coordination games, we show that introducing even a

¹Actions initially not present in the population will never be tried with these imitative dynamics, so we always consider initial points in the interior of the strategy space.

small presence of a second revision protocol into a population changes which equilibrium the population locates. These results underscore the importance of accurately determining exactly how agents learn.

Our results build off of existing comparisons of basins of attraction and equilibrium selection under the best response dynamics and replicator dynamics (Golman and Page 2010a, Golman and Page 2010b). The contribution of this paper is to explore the behavior of a composition of dynamics. We find that heterogeneity across learning styles may bring about new outcomes not foreseeable from analysis of the homogenous dynamics induced by the component learning styles.

Our treatment of equilibrium selection relies on analysis of the deterministic dynamical system. The learning rules define paths through the strategy space that attain equilibria in the games we consider. The equilibrium that is selected is thus a function of the initial population mixed strategy point and the learning dynamic. Nevertheless, we obtain results that hold throughout the interior of the strategy space by focusing on cases in which the basin of attraction of a particular equilibrium approaches the entire space. Our approach does not assume random shocks that allow the population to move from one equilibrium to another, as in the stochastic stability literature (Foster and Young 1990, Young 1993, Kandori et al. 1993, Ellison 2000, Binmore and Samuelson 1999). But, our findings complement that sort of ultralong run analysis because an equilibrium whose basin of attraction approaches the entire strategy space will be uniquely stochastically stable.

The rest of the paper is organized as follows. The next section defines the revision

protocols and the learning dynamics derived from them. In Section 3, we present our result comparing equilibrium selection of these rules. Section 4 concludes with a discussion of the importance of models that capture heterogeneity of learning styles.

2 The Learning Dynamics

We assume a symmetric game, with agents recurrently randomly matched from a single population. The set of actions is finite, $\mathcal{A} = \{1, \dots, n\}$. We let x_i refer to the fraction of the society choosing action $i \in \mathcal{A}$. The population mixed strategy $\mathbf{x} = (x_1, \dots, x_n)$ is an element of Δ^{n-1} , the $(n-1)$ -dimensional simplex where $x_i \geq 0$ for all i and $\sum_i x_i = 1$. Denote by $\pi_i(\mathbf{x})$ the payoff to action i . Naturally, payoffs are a function of the population mixed strategy, $\pi_i : \Delta^{n-1} \rightarrow \mathfrak{R}$. Denote the vector of these payoffs by $\vec{\pi}(\mathbf{x}) = (\pi_1(\mathbf{x}), \dots, \pi_n(\mathbf{x}))$. Let $\text{BR}(\mathbf{x})$ be the set of best replies to \mathbf{x} ,

$$\text{BR}(\mathbf{x}) = \arg \max_{\mathbf{v} \in \Delta^{n-1}} \mathbf{v} \cdot \vec{\pi}(\mathbf{x}).$$

Best response dynamics can be written as

$$\dot{\mathbf{x}} \in \text{BR}(\mathbf{x}) - \mathbf{x}. \tag{1}$$

The *imitate the best dynamics* are

$$\dot{x}_i = x_i \left(I(i|\mathbf{x}) - \sum_j I(j|\mathbf{x}) x_j \right) \tag{2}$$

where the best response indicator function $I(i|\mathbf{x}) = 1$ if $\pi_i(\mathbf{x}) = \pi_{\max}(\mathbf{x})$ and 0 otherwise. These differential equations have points of discontinuity, so a (Filippov) solution

(i.e., an absolutely continuous function that coincides with the unique solution whenever the differential equations are continuous) must satisfy (2) only for almost all t (Smirnov, 2002; Aubin and Cellina, 1984; Benaim, et al., 2005). When there is a unique best response – say action b – (which, for almost any payoff matrix, there is at almost every point in the simplex) equation (2) simplifies to $\dot{\mathbf{x}} = x_b (\text{BR}(\mathbf{x}) - \mathbf{x})$. The additional factor of x_b here, as compared with equation (1), arises because the learning process underlying these dynamics relies on imitation, and best responses are only adopted at a rate proportional to how prevalent they are. In the interior of the strategy space, the trajectory of the imitate the best dynamics matches (one of the solutions of) the best response dynamics, though with different speed.²

The *replicator dynamics* are

$$\dot{x}_i = x_i(\pi_i(\mathbf{x}) - \bar{\pi}(\mathbf{x})) \tag{3}$$

where $\bar{\pi}(\mathbf{x}) = \mathbf{x} \cdot \vec{\pi}(\mathbf{x})$ is the average payoff.

Both imitate the best and replicator dynamics can arise from Poisson alarm clock models (see Sandholm (2009) and Schlag (1998)) in which a revision protocol

²How the imitate the best dynamics relate to the best response dynamics at the boundary of the strategy space is an open question. Clearly, though, if the best response is absent from the population, the dynamics stop altogether. An additional property of the imitate the best dynamics that sets it apart from best response dynamics is that every Nash Equilibrium (even in mixed strategies) must be a steady state. Imitate the best dynamics are also related to the imitative logit dynamics as the limiting case when the logit rationality parameter approaches infinity (Weibull, 1995).

determines the probability that an individual agent switches to another action (as a function of population state and current payoffs) when one's clock rings. The expected motion of this stochastic process defines the learning dynamic.

The paths of the replicator and the imitate the best (and similarly the best response) dynamics are invariant under positive affine transformations of the payoff matrix.³ Payoff transformations do, however, affect the speed of the flow of replicator dynamics. Flows are considered identical if their paths are the same, regardless of speed, because the time parameter has no intrinsic scale.

When trying to combine imitate the best and replicator dynamics, the time parameter becomes relevant. To put both learning dynamics on the same footing, we adopt the framework of a generalized imitation revision protocol. When an agent's alarm goes off, the agent picks an opponent at random and observes this opponent's action. (The selection of an opponent to observe is entirely independent of the random matching of opponents to play against.⁴) The agent observes action j with probability x_j . The agent then switches to the opponent's strategy with probability q that depends on the opponent's payoff. If $q(\pi_j(\mathbf{x})) = \frac{\pi_j(\mathbf{x}) - \pi_{\min}(\mathbf{x})}{\pi_{\max}(\mathbf{x}) - \pi_{\min}(\mathbf{x})}$, the revision protocol generates the replicator dynamics, with a continual rescaling of time. For this *replicator protocol*, the probability of imitation increases linearly in the observed

³Transformations of payoffs that do not affect the dynamics can still determine which of multiple equilibria is payoff dominant. Thus, this traditional equilibrium selection criterion is unrelated to the dynamical equilibrium selection we presently consider.

⁴See Boylan (1992) and Gilboa and Matsui (1992) for justification of random matching in a countably infinite population and Alos-Ferrer (1999) for the case of a continuum of agents.

payoff. Alternatively, if $q(\pi_j(\mathbf{x})) = 1$ when $\pi_j(\mathbf{x}) = \pi_{\max}(\mathbf{x})$ and 0 otherwise, the revision protocol generates the imitate the best dynamics. We term this the *imitate the best protocol*.

One natural way to combine imitate the best and replicator dynamics is to suppose that agents use the replicator protocol with probability p and the imitate the best protocol with probability $1 - p$.⁵ A *linear combination of imitate the best and replicator dynamics* emerges:

$$\dot{x}_i = x_i \left[p \frac{\pi_i(\mathbf{x}) - \bar{\pi}(\mathbf{x})}{\pi_{\max}(\mathbf{x}) - \pi_{\min}(\mathbf{x})} + (1 - p) \left(I(i|\mathbf{x}) - \sum_j I(j|\mathbf{x}) x_j \right) \right]. \quad (4)$$

We can recover homogeneous use of a single learning protocol by setting $p = 1$ or $p = 0$. With $p = 1$, equation (4) becomes

$$\dot{x}_i = x_i \left(\frac{\pi_i(\mathbf{x}) - \bar{\pi}(\mathbf{x})}{\pi_{\max}(\mathbf{x}) - \pi_{\min}(\mathbf{x})} \right). \quad (5)$$

The factor $\frac{1}{\pi_{\max}(\mathbf{x}) - \pi_{\min}(\mathbf{x})}$ comes from the replicator revision protocol, where it is necessary to ensure that we can interpret q as a probability, i.e., that $q \leq 1$. With the appropriate rescaling of time, it would drop out of equation (5), producing the replicator dynamics of equation (3). With $p = 0$, the linear combination dynamics (4) reduces to the imitate the best dynamics (2).

⁵Another possibility, instead of assuming that each agent sometimes uses each protocol, is that there could be some agents who always use the imitate the best protocol and other agents who always use the replicator protocol. The theoretical analysis is more complicated in this case, but by capturing heterogeneity of learning styles across the population, the model appears more realistic. This formulation is presented in the appendix.

Note that the linear combination dynamics, as described by (4), is invariant under positive affine transformations of the payoff matrix, and the relative growth rate of any action is bounded between -1 and 1 (and a priori cannot be bounded any more tightly). This would not be the case for the more straightforward combination

$$\dot{x}_i = x_i \left[p (\pi_i(\mathbf{x}) - \bar{\pi}(\mathbf{x})) + (1 - p) \left(I(i|\mathbf{x}) - \sum_j I(j|\mathbf{x}) x_j \right) \right].$$

This straightforward combination dynamic could behave quite differently from (4) because the factor $\pi_{\max}(\mathbf{x}) - \pi_{\min}(\mathbf{x})$ could get bigger or smaller over time and thus shift more or less weight to the replicator dynamics. Nevertheless, Theorem 1 would hold with this dynamic in place of (4) if we were to divide all the entries in our payoff matrix by N^2 .⁶

3 Basins of Attraction

Strict Nash equilibria are asymptotically stable for the pure replicator dynamics, the imitate the best dynamics, and the linear combination dynamics.⁷ However, agreement of local stability properties does not ensure similar long-run behavior. In

⁶Dividing our payoff matrix by N^2 would have the effect of making $\pi_{\max}(\mathbf{x}) - \pi_{\min}(\mathbf{x}) \approx 1$. (We cannot rescale payoffs to make $\pi_{\max}(\mathbf{x}) - \pi_{\min}(\mathbf{x}) = 1$ because it varies with time.) Establishing an analogue of Theorem 1 is not a trivial corollary because this small variation in the learning dynamic could potentially have a big impact on equilibrium selection (that is the point of our result, after all), but in this case, the logic of our proof would still go through.

⁷See Hofbauer, et. al. (1979), Hofbauer (2000), Hofbauer and Sigmund (2003), and Hofbauer, et. al. (2009) for more stability results concerning the best response and replicator dynamics.

games with multiple equilibria, we must consider basins of attraction.

Imitate the best dynamics and replicator dynamics often produce very similar basins of attraction for their equilibria. However, Golman and Page (2010a) provide a class of three-by-three coordination games in which replicator dynamics has basins of attraction arbitrarily different from those of best response dynamics (and hence from imitate the best dynamics as well). Games with such divergence between the behavior of the replicator dynamics and the imitate the best dynamics lend themselves to the analysis of our composite dynamics. Intuition might lead one to believe that a combination of imitate the best and replicator dynamics would always share some overlap in basins of attraction with each component rule, but this is not so. The linear combination dynamics can be fundamentally different from its parts. Here, we identify a class of games in which imitate the best dynamics and replicator dynamics and the linear combination of the two all have vanishing overlap in their basins of attraction at the same time. That is, by combining an imitate the best protocol and a replicator protocol, the population can locate an equilibrium that neither rule could find on its own.

Consider a class of symmetric games with five actions and the following payoff

matrix:

$$A = \begin{pmatrix} 0 & 0 & 0 & 0 & N \\ \ln N & \ln N & \ln N - N & \ln N - N & \ln N \\ 0 & 0 & 1 & 1 - \frac{1}{N} & 0 \\ 1 - \frac{1}{N} & 1 - \frac{1}{N} & 1 - \frac{1}{N} & 1 - \frac{1}{N} + \frac{1}{N^2} & 1 - \frac{1}{N} \\ -N^2 & -N^2 & -N^2 & -N^2 & -N^2 \end{pmatrix} \quad (6)$$

where the parameter N is large.⁸ Actions 2, 3, and 4 are all strict equilibrium actions.

The payoffs described by (6) may appear complicated, but they permit a simple interpretation of the game. We should think of it as a coordination game, but with a couple of extra options. Action 5 is just a bad choice. So bad, in fact, that all other actions look good in comparison to it if an agent is using the replicator learning protocol. While such strongly dominated strategies are often ignored when modeling games, it's hard to argue that they don't exist, and there is both theoretical and experimental evidence that they can affect equilibrium selection (Ochs 1995, Basov 2004). Action 1, on the other hand, looks tempting because it can exploit action 5. In fact, action 1 is almost always the initial best response, but it's only a temporary best response because its payoff shrinks as agents learn to avoid action 5. Action 2 is high risk, high reward; agents want to coordinate on it if and only if they almost never see actions 3 or 4 played. Actions 3 and 4 are safer, with moderate payoffs. Of the pair, action 3 would yield a slightly higher equilibrium payoff, whereas action 4

⁸Our result holds for more general classes of games as long as the terms in the payoff matrix keep the same relative orders of magnitude (cf. Golman, 2009).

has the highest worst-case payoff.

In Theorem 1, we show that the imitate the best dynamics lead to everyone playing action 2, the replicator dynamics to everyone playing action 3, and the linear combination dynamics to everyone taking action 4. The imitate the best dynamics selects the action that is a best response to the temporary best response; the replicator dynamics selects the action that remains better than average as long as it spreads; and the linear combination dynamics selects something else entirely.

Theorem 1 *Consider the class of games with payoff matrix given by equation (6).*

As the payoff parameter N grows arbitrarily large, the imitate the best dynamics, the replicator dynamics, and the linear combination dynamics (with any positive weight on both components, i.e., any $p \in (0, 1)$) from an arbitrarily large portion of the strategy space flow to a single strict equilibrium featuring action 2, 3, or 4 respectively. Notably, all three dynamics share with each other vanishing overlap in their basins of attraction.

Proof See appendix.

For an arbitrarily large portion of the state space, all three dynamics start out heading towards the initial, temporary best response (action 1). But, there are subtle differences between the flows of each dynamic that will greatly affect which action spreads when there's hardly anybody left to exploit with action 1.

With imitate the best dynamics, the relative payoffs of the other actions are irrelevant for their decay rate; agents switch away from actions 3 and 4 just as quickly

as they switch from action 5. This mitigates the downside of action 2, and with its residual high reward, it becomes the next best response. By the improvement principle for best response / imitate the best dynamics (Monderer and Sela, 1997; Sela, 2000; Viossat, 2008), if a strict equilibrium action is ever the unique best response to the population state, it remains so thereafter. Thus, action 2 eventually spreads under the imitate the best dynamics.

With replicator dynamics, the flow towards the initial best response comes at the expense of the bad choice (action 5). As the payoff to action 5 is made arbitrarily bad relative to all other payoffs, all other actions initially spread along with the initial best response. The spread of actions 3 and 4 inhibits action 2, and the success of action 1 is temporary, so we effectively have actions 3 and 4 racing to achieve critical mass in a coordination game. Although action 4 starts out in the lead because it profits more against action 1, action 3 generates for itself more positive feedback and thus has a much lower critical mass threshold (the relevant comparison in (6) is $a_{33} - a_{34} \gg a_{44} - a_{43}$). Thus, the replicator dynamics yields eventual coordination on action 3.

For the linear combination dynamics, actions 3 and 4 neither keep pace with the spread of the initial best response nor decay as fast as the bad choice. They do survive enough to still inhibit action 2. But the initial lead in payoffs that action 4 has relative to action 3, and the fact that the linear combination dynamics favors a best response more than the replicator dynamics does, means that action 4 now has the advantage. Thus, the linear combination dynamics allows action 4 to become

entrenched.

Not only can the linear combination dynamics differ from its component parts, but this deviation in long-run behavior can even occur with almost all weight on one piece of the linear combination. Thus, the mere possibility of using the replicator protocol, however unlikely, is enough to completely shift the long-term behavior of a population of agents using primarily the imitate the best protocol, and vice versa.

4 Discussion

In a large population of agents, we should expect heterogeneity of behavior (Hommes 2006, Kirman 2006).⁹ Camerer and Ho (1999) correctly recognized that individuals may combine diverse basic learning rules. Our result demonstrates that the presence of multiple learning styles in the same population, and the specifics of how they all come into play, affects equilibrium selection. We find that for a specific class of coordination games, the outcome of a learning dynamic that is a convex combination of imitate the best dynamics and replicator dynamics is completely different than a convex combination of the outcomes produced by these individual rules. The simple assumption that the combination of two learning dynamics behaves like one of its component pieces is wrong. The explanation is that in games with many actions and multiple equilibria, the trajectory of the combination dynamics can enter parts of the

⁹Heterogeneity can take many forms. Ely and Sandholm (2005) consider agents who all use a common learning rule, but who have diverse preferences, i.e., heterogeneous payoff functions.

strategy space that would not be approached by either of the component dynamics. Thus, the basins of attraction can be entirely reshaped when agents use multiple learning rules.

Hofbauer et al. (2009) have shown that the time-average of the replicator dynamics is a vanishingly perturbed solution of the best-reply dynamics. Our result, along with the findings of Golman and Page (2010a; 2010b), illustrates that this does not imply that these dynamics will necessarily select the same equilibria. Even a small perturbation in the learning process can have a strong effect on long run behavior.

Appendix

The Two-Subpopulation Dynamics

Consider a partition of the population into subpopulations of best-responders and replicators, \mathbf{y} and \mathbf{z} respectively.¹⁰ For all i , $x_i = y_i + z_i$. Learning styles are fixed for all time, so $\sum_i z_i = \rho$ is the fraction of agents using the replicator protocol, and $\sum_i y_i = 1 - \rho$ is the fraction using the imitate the best protocol.¹¹ We assume the two subpopulations are well-mixed, both for the purposes of playing the game and for finding an agent to possibly imitate. Thus, the dynamics of these two subpopulations

¹⁰There is reason to think that optimizers and imitators coexist in the same population (Conlisk, 1980).

¹¹In a biological setting, we might want to allow the relative sizes of the subpopulations to vary with the success of the learning rule in use. The possibility of the learning rule itself evolving or being learned would be an interesting extension.

are coupled together. For the subpopulation learning with the replicator protocol, the following dynamics emerges:

$$\begin{aligned} \dot{z}_i = & \sum_j z_j z_i \left(\frac{\pi_i(\mathbf{x}) - \pi_{\min}(\mathbf{x})}{\pi_{\max}(\mathbf{x}) - \pi_{\min}(\mathbf{x})} \right) - z_i \sum_j z_j \left(\frac{\pi_j(\mathbf{x}) - \pi_{\min}(\mathbf{x})}{\pi_{\max}(\mathbf{x}) - \pi_{\min}(\mathbf{x})} \right) + \\ & \sum_j z_j y_i \left(\frac{\pi_i(\mathbf{x}) - \pi_{\min}(\mathbf{x})}{\pi_{\max}(\mathbf{x}) - \pi_{\min}(\mathbf{x})} \right) - z_i \sum_j y_j \left(\frac{\pi_j(\mathbf{x}) - \pi_{\min}(\mathbf{x})}{\pi_{\max}(\mathbf{x}) - \pi_{\min}(\mathbf{x})} \right). \end{aligned} \quad (7)$$

The first two terms of (7) represent learning through imitation of other agents who also use the replicator protocol while the last two terms represent learning from members of the other subpopulation. Let $\bar{\pi}_\mu(\mathbf{x})$ be the average payoff in subpopulation μ , meaning that $\mathbf{z} \cdot \vec{\pi}(\mathbf{x}) = \rho \bar{\pi}_z(\mathbf{x})$ and $\mathbf{y} \cdot \vec{\pi}(\mathbf{x}) = (1 - \rho) \bar{\pi}_y(\mathbf{x})$. The dynamics (7) simplifies to

$$\dot{z}_i = \rho z_i \left(\frac{\pi_i(\mathbf{x}) - \bar{\pi}_z(\mathbf{x})}{\pi_{\max}(\mathbf{x}) - \pi_{\min}(\mathbf{x})} \right) + \rho y_i \left(\frac{\pi_i(\mathbf{x}) - \pi_{\min}(\mathbf{x})}{\pi_{\max}(\mathbf{x}) - \pi_{\min}(\mathbf{x})} \right) - (1 - \rho) z_i \left(\frac{\bar{\pi}_y(\mathbf{x}) - \pi_{\min}(\mathbf{x})}{\pi_{\max}(\mathbf{x}) - \pi_{\min}(\mathbf{x})} \right). \quad (8)$$

For the subpopulation learning with the imitate the best protocol, the dynamics are:

$$\dot{y}_i = I(i|\mathbf{x})(1 - \rho)(z_i + y_i) - y_i \sum_j I(j|\mathbf{x})(z_j + y_j). \quad (9)$$

We refer to the system defined by (8) and (9) as the two-subpopulation dynamics.¹²

With twice the variables, the two-subpopulation dynamics is much more complicated than the linear combination dynamics, but strict Nash equilibria are still asymptotically stable, and for some games, at least, we can show that the two dynamics behave alike (Golman, 2009).

¹²We can recover homogeneous learning dynamics, as in equation (5) or (2), from the two-subpopulation dynamics by setting $\rho = 1$ or $\rho = 0$ respectively.

Formal Statement and Proof of Theorem 1.

Let m be the Lebesgue measure on the strategy space Δ^{n-1} . Given the parameter N , let $G(N)$ be a class of symmetric normal form games with payoff matrix A depending on those parameters. That is, the payoff to action i against action j , a_{ij} , is a function of N . Then, the expected payoff to action i given random matching with population mixed strategy $\mathbf{x} \in \Delta^{n-1}$ is $\pi_i(\mathbf{x}) = \sum_j a_{ij} x_j$. Henceforth, we omit the payoff function's argument for ease of notation, writing π_i in place of $\pi_i(\mathbf{x})$. The payoff vector can be represented succinctly with matrix multiplication, $\vec{\pi} = A\mathbf{x}$. Given a learning rule \mathcal{R} and an equilibrium action i of the game $G(N)$, let $B(\mathcal{R}, i, N)$ denote the basin of attraction of $(x_i = 1, x_{-i} = 0)$. Let \mathbf{R} denote the replicator dynamics, \mathbf{B} the imitate the best dynamics, and $\mathbf{L}(p)$ the linear combination using the replicator protocol with probability p .

Theorem 1 *Consider the class of games with payoff matrix given by equation (6).*

For any $\mathcal{R} \in \{\mathbf{B}, \mathbf{R}, \mathbf{L}(p)\}$, with any $p \in (0, 1)$,

$$\lim_{N \rightarrow \infty} m(B(\mathcal{R}, i, N)) = m(\Delta^4)$$

for some equilibrium action i , i.e., the imitate the best dynamics, the replicator dynamics, and the linear combination dynamics with any positive weight on both components each flow to a single strict equilibrium from an arbitrarily large portion of the strategy space. Still,

$$\lim_{N \rightarrow \infty} \sum_{i \in \mathcal{A}} m(B(\mathbf{B}, i, N) \cap B(\mathbf{R}, i, N)) = 0$$

and

$$\lim_{N \rightarrow \infty} \sum_{i \in \mathcal{A}} m(B(\mathbf{L}(\mathbf{p}), i, N) \cap B(\mathbf{B}, i, N)) = 0$$

and

$$\lim_{N \rightarrow \infty} \sum_{i \in \mathcal{A}} m(B(\mathbf{L}(\mathbf{p}), i, N) \cap B(\mathbf{R}, i, N)) = 0,$$

i.e., all three dynamics share with each other vanishing overlap in their basins of attraction. (In particular, imitate the best dynamics leads to the spread of action 2, replicator dynamics action 3, and the linear combination dynamics action 4.)

Proof Take N to be large. We first show that from almost any starting point in the strategy space, the imitate the best dynamics locates the equilibrium in which only action 2 is played. Observe that if action 2 is ever the best response, it remains so forever, by the improvement principle.

Observe that $\pi_2 \leq \ln N$ as well as $\pi_3 \leq 1$ and $\pi_4 \leq 1$ everywhere. Consider starting points that satisfy

$$x_5(0) > \frac{1}{\ln N}, \tag{10}$$

a condition that can be met anywhere in the interior of the strategy space by taking N to be sufficiently large. Initially the best response must be action 1, because inequality (10) implies that $\pi_1 > \ln N$. So all x_i , $i \neq 1$, have the same relative decay rate. While $x_i \geq \frac{\ln N}{N} x_i(0)$,

$$\pi_1 \geq \ln N x_5(0) > 1.$$

If x_i is to reach this lower bound, at the same time

$$\pi_2 \geq \ln N (1 - x_3(0) - x_4(0)) > 1.$$

Thus, π_2 grows above π_3 and π_4 before π_1 falls to that level. Eventually, action 2 must become the best response, and the dynamics heads towards it thereafter.

We now show that the replicator dynamics leads to the equilibrium in which action 3 is always played. Under the replicator dynamics,

$$\begin{aligned} \dot{x}_3 &= x_3 \left[x_3(1 - x_3) + \left(1 - \frac{1}{N}\right)x_4(1 - x_3) - \left(1 - \frac{1}{N}\right)x_4 - \frac{1}{N^2}x_4^2 + \right. \\ &\qquad \qquad \qquad \left. (Nx_3 + Nx_4 - \ln N)x_2 + N^2x_5 - Nx_5x_1 \right] \\ &= x_3 \left[x_3(x_1 + x_2 + x_5) + \frac{1}{N}x_3x_4 - \frac{1}{N^2}x_4^2 + (Nx_3 + Nx_4 - \ln N)x_2 + N^2x_5 - Nx_5x_1 \right] \end{aligned}$$

and we claim this is positive. We need to restrict to starting points that satisfy

$$x_3(0) > \frac{\ln N}{N}, \tag{11}$$

another inequality that can be met anywhere in the interior of the strategy space by taking N sufficiently large. Inequality (11) implies that $\frac{1}{N}x_3x_4 - \frac{1}{N^2}x_4^2 > 0$ and $Nx_3 + Nx_4 - \ln N > 0$. We have $N^2x_5 - Nx_5x_1 > 0$ easily. Thus, $\dot{x}_3 > 0$ under the replicator dynamics.

Lastly, we show that the linear combination dynamics finds the equilibrium in which everyone takes action 4. Inspection of the payoff matrix reveals that π_5 is always π_{\min} . For sufficiently large N , initial points (except on the boundary of the strategy space) will satisfy

$$x_4(0) > N^{1-p} \frac{\ln N}{N}. \tag{12}$$

Among other things, this means that $\pi_2 < 0$ initially. Assuming additionally that $x_5(0) > \frac{1}{N}$, as will also hold for sufficiently large N , then $\pi_1 = \pi_{\max}$ initially. While

this remains the case, the linear combination dynamics yield

$$\dot{x}_1 = x_1 \left(p \frac{Nx_5 - \bar{\pi}}{N^2 + Nx_5} + (1-p)(1-x_1) \right) \quad (13)$$

$$\dot{x}_3 = x_3 \left(p \frac{x_3 + (1 - \frac{1}{N})x_4 - \bar{\pi}}{N^2 + Nx_5} - (1-p)x_1 \right) \quad (14)$$

$$\dot{x}_4 = x_4 \left(p \frac{1 - \frac{1}{N} + \frac{1}{N^2}x_4 - \bar{\pi}}{N^2 + Nx_5} - (1-p)x_1 \right) \quad (15)$$

$$\dot{x}_5 = x_5 \left(p \frac{-N^2 - \bar{\pi}}{N^2 + Nx_5} - (1-p)x_1 \right). \quad (16)$$

While $x_5 > \frac{1}{N^2 - N}$ and $x_4 > \frac{\ln N}{N}$, straightforward but tedious calculation reveals that the average payoff $\bar{\pi} < 0$, and we have the following bounds:

$$\dot{x}_1 > (1-p)x_1(1-x_1) \quad (17)$$

$$\dot{x}_3 > -(1-p)x_1x_3 \quad (18)$$

$$\dot{x}_4 > -(1-p)x_1x_4 \quad (19)$$

$$\dot{x}_5 > -x_5. \quad (20)$$

Also note that

$$\begin{aligned} \frac{-N^2 - \bar{\pi}}{N^2 + Nx_5} &= \frac{-N^2 - (-N^2x_5 + Nx_5x_1 - (Nx_3 + Nx_4)x_2 + \ln Nx_2 + \pi_3x_3 + \pi_4x_4)}{N^2 + Nx_5} \\ &< \frac{-N^2(1-x_5) - Nx_5x_1 + (Nx_3 + Nx_4)x_2}{N^2 + Nx_5} \\ &= \frac{-N^2(x_1) - Nx_5x_1 + (Nx_3 + Nx_4)x_2 - N^2(x_2 + x_3 + x_4)}{N^2 + Nx_5} \\ &< \frac{(-N^2 - Nx_5)x_1}{N^2 + Nx_5} = -x_1. \end{aligned}$$

So we can bound the left hand side of equation (16),

$$\dot{x}_5 < -x_1x_5. \quad (21)$$

So

$$x_1(t) > \frac{1}{1 + \frac{1-x_1(0)}{x_1(0)}e^{-(1-p)t}} \quad (22)$$

$$x_3(t) > x_3(0)e^{-(1-p)\int_0^t x_1(s)ds} \quad (23)$$

$$x_4(t) > x_4(0)e^{-(1-p)\int_0^t x_1(s)ds} \quad (24)$$

$$x_5(t) < x_5(0)e^{-\int_0^t x_1(s)ds} \quad (25)$$

$$x_5(t) > x_5(0)e^{-t}. \quad (26)$$

For sufficiently large N , initial points in the interior can be made to satisfy $x_1(0) > \frac{1}{N-1}$. Observe that $\pi_4 > \pi_3 + (1 - \frac{1}{N})x_1 - \frac{1}{N}x_3$. As long as the dynamics as described by equation (13) apply, $x_1(t) \geq x_1(0)$, and so $\pi_4(t) > \pi_3(t)$.

While $x_5(t) > \frac{1-\frac{1}{N}}{N}$, inequality (25) implies $e^{-\int_0^t x_1(s)ds} > \frac{1-\frac{1}{N}}{Nx_5(0)}$. So, using inequality (24),

$$x_4(t) > x_4(0) \left(\frac{1 - \frac{1}{N}}{Nx_5(0)} \right)^{1-p}. \quad (27)$$

We can ensure $1 - \frac{1}{N} > x_5(0)$ for initial points in the interior of the strategy space by taking N sufficiently large, so inequality (27) simplifies to $x_4(t) > x_4(0) \left(\frac{1}{N}\right)^{1-p}$, and then by inequality (12), we get $x_4(t) > \frac{\ln N}{N}$. So $\pi_2(t) < 0$ as long as $x_5(t) > \frac{1-\frac{1}{N}}{N}$.

Let $t_c = t : \pi_1(t) = \pi_4(t)$. Thus, dividing both sides by N and dropping a positive term from the right side, we have $x_5(t_c) > \frac{1-\frac{1}{N}}{N}$. This means $\pi_4(t) > \pi_2(t)$ and $\pi_4(t) > \pi_3(t)$ for all $t \leq t_c$. The dynamics described by equations (17) through (21) hold until t_c , at which point π_4 becomes π_{\max} .

Because $\pi_4(t) > \pi_3(t)$ for all $t \leq t_c$, we can also conclude that

$$\frac{x_4(t_c)}{x_4(0)} > \frac{x_3(t_c)}{x_3(0)}. \quad (28)$$

We know that $x_5(t_c) < \frac{1}{N}$, too. Inequality (26) then implies $e^{-t_c} < \frac{1}{Nx_5(0)}$, and using inequality (22), we come to

$$x_1(t_c) > \frac{1}{1 + \frac{1-x_1(0)}{x_1(0)} \left(\frac{1}{Nx_5(0)}\right)^{1-p}}. \quad (29)$$

After t_c , while π_4 remains π_{\max} , the linear combination dynamics yield

$$\begin{aligned} \dot{x}_1 &= x_1 \left(p \frac{\pi_1 - \bar{\pi}}{N^2 + \pi_4} - (1-p)x_4 \right) \\ \dot{x}_3 &= x_3 \left(p \frac{\pi_3 - \bar{\pi}}{N^2 + \pi_4} - (1-p)x_4 \right) \\ \dot{x}_4 &= x_4 \left(p \frac{\pi_4 - \bar{\pi}}{N^2 + \pi_4} + (1-p)(1-x_4) \right). \end{aligned}$$

Let

$$f(t) = p \frac{\pi_3 - \bar{\pi}}{N^2 + \pi_4} - (1-p)x_4,$$

so that $\dot{x}_3 = x_3 f(t)$. For $t \geq t_c$,

$$x_3(t) = x_3(t_c) e^{\int_{t_c}^t f(s) ds}. \quad (30)$$

Because $\pi_3 \leq 1$ and $\pi_1 \geq 0$, we know that $\dot{x}_1 \geq x_1 (f(t) - p \frac{1}{N^2})$. So

$$x_1(t) \geq x_1(t_c) e^{\int_{t_c}^t f(s) ds} \cdot e^{-\frac{p}{N^2}(t-t_c)}. \quad (31)$$

As $\pi_4 > \pi_3$, we know $\dot{x}_4 > x_4 (f(t) + 1 - p)$, and

$$x_4(t) \geq x_4(t_c) e^{\int_{t_c}^t f(s) ds} \cdot e^{(1-p)(t-t_c)}. \quad (32)$$

Recall that $\pi_4 > \pi_3$ if $(1 - \frac{1}{N})x_1 > \frac{1}{N}x_3$. Putting together equations (30) and (31), this holds as long as $(1 - \frac{1}{N})x_1(t_c)e^{-\frac{p}{N^2}(t-t_c)} > \frac{1}{N}(1 - x_1(t_c))$. Plugging in inequality (29) and arranging terms, we find that $\pi_4 > \pi_3$ if

$$e^{-(t-t_c)} > \left[\frac{1}{N-1} \frac{1-x_1(0)}{x_1(0)} \left(\frac{1}{Nx_5(0)} \right)^{1-p} \right]^{\frac{N^2}{p}}. \quad (33)$$

Note that we also obtain $\pi_4 > \pi_3$ if $\frac{1}{N^2}x_4 > \frac{1}{N}x_3$. Putting together equations (30), (32) and (28), this holds when $\frac{1}{N^2}x_4(0)e^{(1-p)(t-t_c)} > \frac{1}{N}x_3(0)$. Arranging terms, we obtain $\pi_4 > \pi_3$ when

$$e^{-(t-t_c)} < \left(\frac{\frac{1}{N^2}x_4(0)}{\frac{1}{N}x_3(0)} \right)^{\frac{1}{1-p}}. \quad (34)$$

By combining equations (33) and (34), we can be sure $\pi_4 > \pi_3$ for all t if

$$\left(\frac{x_4(0)}{Nx_3(0)} \right)^{\frac{1}{1-p}} > \left[\frac{1}{N-1} \frac{1-x_1(0)}{x_1(0)} \left(\frac{1}{Nx_5(0)} \right)^{1-p} \right]^{\frac{N^2}{p}}. \quad (35)$$

Inequality (35) holds for any interior initial point when N is sufficiently large because the right hand side of the inequality decays exponentially in N .

While $\pi_4 = \pi_{\max}$, obviously x_4 is increasing, and so π_2 remains negative for all t after t_c . Thus, action 4 remains the best response forever, and the linear combination dynamics locates the equilibrium in which only action 4 is played. ■

References

- [1] Alos-Ferrer, C., (1999). Dynamical Systems with a Continuum of Randomly Matched Agents. *J. Econ. Theory* 86, 245-267.

- [2] Aubin, J., Cellina, A., (1984). *Differential inclusions: set-valued maps and viability theory*. Springer.
- [3] Basov, S., (2004). Equilibrium Selection in Coordination Games: Why do dominated strategies matter? *Econ. Bull.* 3, 1-3.
- [4] Benaim, M., Hofbauer, J., Sorin, S., (2005). Stochastic Approximations and Differential Inclusions. *SIAM J. Control Optim.* 44, 3283-48.
- [5] Binmore, K., Samuelson, L., (1999). Evolutionary Drift and Equilibrium Selection. *Rev. Econ. Stud.* 66, 363-394.
- [6] Borgers, T., Sarin, R., (1997). Learning Through Reinforcement and Replicator Dynamics. *J. Econ. Theory* 77, 1-14.
- [7] Boylan, R., (1992). Laws of Large Numbers for Dynamical Systems with Randomly Matched Individuals. *J. Econ. Theory* 57, 473-504.
- [8] Camerer, C., Ho, T.H., (1999). Experience Weighted Attraction Learning in Normal Form Games. *Econometrica* 67, 827-874.
- [9] Conlisk, J., (1980). Costly Optimizers Versus Cheap Imitators. *J. Econ. Behav. Organ.* 1 (3), 275-293.
- [10] Ellison, G., (2000). Basins of Attraction, Long-Run Stochastic Stability, and the Speed of Step-by-Step Evolution. *Rev. Econ. Stud.* 67 (1), 17-45.

- [11] Ely, J., Sandholm, W.H., (2005). Evolution in Bayesian Games I: Theory. *Games Econ. Behav.* 53, 83-109.
- [12] Foster, D., Young, H.P., (1990). Stochastic Evolutionary Game Dynamics. *Theoretical Population Biology* 38, 219-232.
- [13] Fudenberg, D., Levine, D.K., (1998). *The Theory of Learning in Games*. Cambridge: MIT Press.
- [14] Gilboa, I., Matsui, A., (1991). Social Stability and Equilibrium. *Econometrica* 59, 859-867.
- [15] Gilboa, I., Matsui, A., (1992). A Model of Random Matching. *J. Math. Econ.* 21, 185-197.
- [16] Golman, R. (2009). Essays on Population Learning Dynamics and Boundedly Rational Behavior. Ph.D. dissertation, University of Michigan.
- [17] Golman, R., Page, S., (2010a). Basins of Attraction and Equilibrium Selection Under Different Learning Rules. *J. Evol. Econ.* 20, 49-75.
- [18] Golman, R., Page, S., (2010b). Individual and Cultural Learning in Stag Hunt Games With Multiple Actions. *J. Econ. Behav. Organ.* 73, 359-376.
- [19] Hofbauer, J., (1995). Imitation Dynamics for Games, mimeo.
- [20] Hofbauer, J., (2000). From Nash and Brown to Maynard Smith: Equilibria, Dynamics and ESS. *Selection* 1, 81-88.

- [21] Hofbauer, J., Schuster, P., Sigmund, K., (1979). A Note on Evolutionary Stable Strategies and Game Dynamics. *J. Theor. Biol.* 81, 609-612.
- [22] Hofbauer, J., Sigmund, K., (2003). Evolutionary Game Dynamics. *Bull. Am. Math. Soc.* 40, 479-519.
- [23] Hofbauer, J., Sorin, S., Viossat, Y., (2009). Time Average Replicator and Best-Reply Dynamics. *Math. Oper. Res.* forthcoming.
- [24] Hommes, C., (2006). Heterogeneous Agent Models in Economics and Finance. In: Judd, K., Tesfatsion, L. (eds) Handbook of Computational Economics, vol 2, Agent-Based Computational Economics. North Holland, Amsterdam.
- [25] Kandori, M., Mailath, G., Rob, R., (1993). Learning, Mutation and Long Run Equilibria in Games. *Econometrica* 61, 29-56.
- [26] Kirman, A., (2006). Heterogeneity in Economics. *J. Econ. Interact. Coord.* 1, 89-117.
- [27] Monderer, D., Sela, A., (1997). Fictitious Play and No-Cycling Conditions. Sfb504 Publications 97-12, University of Mannheim.
- [28] Ochs, J., (1995). Coordination Problems. *Handbook of Experimental Economics*. Kagel, J., Roth, A. (eds) Princeton: Princeton University Press.

- [29] Sandholm, W. H., (2009). Pairwise Comparison Dynamics and Evolutionary Foundations for Nash Equilibrium. Working paper, University of Wisconsin. <http://www.ssc.wisc.edu/~whs>.
- [30] Schlag, K. H., (1998). Why Imitate, and If So, How? A Boundedly Rational Approach to Multi-armed Bandits. *J. Econ. Theory* 78, 130-156.
- [31] Sela, A., (2000). Fictitious Play in 23 Games. *Games Econ. Behav.* 31, 152-162.
- [32] Smirnov, G., (2002). *Introduction to the theory of differential inclusions*. Providence: AMS Bookstore.
- [33] Swinkels, J. M., (1993). Adjustment Dynamics and Rational Play in Games. *Games Econ. Behav.* 5, 455-484.
- [34] Taylor, P.D., Jonker, L., (1978). Evolutionarily Stable Strategies and Game Dynamics. *Math. Biosci.* 40, 145-156.
- [35] Viossat, Y., (2008). Evolutionary Dynamics May Eliminate All Strategies Used in Correlated Equilibria. *Mathematical Social Sciences* 56, 27-43.
- [36] Weibull, J., (1995). *Evolutionary Game Theory*. Cambridge, MA: The M.I.T. Press.
- [37] Young, H.P., (1993). The Evolution of Conventions. *Econometrica* 61, 57-84.