

Infection and immunization: a new class of evolutionary game dynamics

Samuel Rota Bulò^a, Immanuel M. Bomze^{b,*}

^a*DSI, University of Venice, Italy*

^b*ISDS, University of Vienna, Austria*

Abstract

Building upon a central paradigm of evolutionary game theory, namely the invasion barrier, we propose the new Infection and Immunization Dynamics (INFIMMDYN), modelling a plausible adaptation process in a large population. For general games, this yields a novel refinement of the Nash equilibrium concept based on dynamical arguments, close in spirit to Nash's original "mass action" idea in his Ph.D. thesis. For partnership games, INFIMMDYN exhibits a better asymptotic behavior compared to other popular procedures like Fictitious Play and Replicator Dynamics. We establish even support separation of INFIMMDYN in finite time, which can never be achieved by any interior-point method like those mentioned above. In fact, this property has not yet been established for any other evolutionary game dynamics.

Keywords: revision protocol, evolutionary stability, equilibrium selection, fictitious play, replicator dynamics, best-response, learning population games

1. Introduction

Since John Maynard Smith's seminal contribution [21], a main goal of evolutionary game theory has been to predict or motivate animal behavior using the principles and tools of non-cooperative game theory. Within this context, animal conflicts are modeled in terms of a game involving (typically) a small number of pure strategies, thereby allowing one to obtain analytical solutions, whenever possible. Recently, however, a new trend is emerging which poses novel, challenging questions pertaining to the computational aspects related to evolutionary game theoretic notions. Indeed, in many application fields one is often interested in finding equilibria of evolutionary games which involve thousands or even millions of (pure) strategies. These include applications in economics [7] but also, for example, problems in combinatorial optimization [3, 26], medical imaging and computer vision [19, 34], statistics and machine learning [25], network design [35] and security [27]. Here, standard approaches to equilibrium selection (such as the replicator dynamics) pose serious efficiency as well as efficacy problems. Further, we note that

*Corresponding author

Email addresses: srotabul@dsi.unive.it (Samuel Rota Bulò), immanuel.bomze@univie.ac.at (Immanuel M. Bomze)

Preprint submitted to Games and Economic Behaviour

June 24, 2010

recently some interest has been directed towards studying the complexity of computing equilibria [12, 13].

Motivated by the above considerations, we present a new class of evolutionary dynamics, called Infection and Immunization Dynamics (INFIMMDYN). Our dynamics is built upon a central paradigm of evolutionary game theory, namely the invasion barrier. While in part of the literature, this notion has mostly served the purpose of static evolutionary stability analysis, its real significance is revealed in the dynamical context, determining when and how a population is driven to equilibrium [8, 36, 5, 1]. In this work we will also follow this approach.

Indeed, for any population \mathbf{x} not in equilibrium, there exists at least one mixed strategy \mathbf{y} that is a better response to \mathbf{x} than \mathbf{x} to itself and against which \mathbf{x} has no invasion barrier. Once such a strategy (called “infective” strategy) is given, it stands to reason that by introducing a small share of mutant agents playing strategy \mathbf{y} in the current population (we call this process “infection”), they will spread until the invasion barrier against them becomes positive, i.e., until the new population turns out to be “immune” against \mathbf{y} . This process resembles the function of a vaccine, by which a small share of virus is introduced in the body, in order to stimulate its immune system and prevent the effects of future infections. By reiterating this process of infection and immunization, we aim at obtaining a population for which no infective strategy can be found anymore, because in that case we have reached a Nash equilibrium. Indeed, we show that fixed points of INFIMMDYN are Nash equilibria and vice versa, independently from the way we select infective strategies at each iteration. This is thus a property of the whole class of infection and immunization dynamics, and yields a new characterization of Nash equilibrium strategies, allowing for a whole variety of new refinement concepts by qualifying the dynamic type (asymptotical or Lyapunov stable, for instance) of such a fixed point. Note that Nash’s original idea of “mass action” in his unpublished Ph.D. thesis [23], see also [18, 2] is very close in spirit to this approach.

Within the class INFIMMDYN, we focus our attention on a special dynamics called Pure INFIMMDYN, which allows us to obtain a computationally efficient approach for finding equilibria in evolutionary games. Intuitively, the evolutionary process can be interpreted as follows: as time passes by, an advertisement on the basis of the aggregate behavior of the population tells the agents that a certain pure strategy is *trendy* or is *out-of-fashion*. A strategy is trendy, if it is performing best in terms of payoff in the population, whereas it is out-of-fashion, if it is the worst still alive performing strategy in the population. Both variants will be taken into account: in contrast to the best-reply approach, a trendy strategy is chosen only if its *absolute* deviation from the average payoff is largest among all absolute deviations. Otherwise, the largest absolute deviation is provided by an out-of-fashion strategy, and we move straight away from it by help of its *co-strategy* to be defined below. In terms of the general formulation, this phase encodes a particular selection function for infective strategies, which basically increases (decreases) the share of agents playing the trendy (out-of-fashion) strategy, as long as there is no barrier to the invasion. Hence, assuming that agents can gather information only about the announced strategy, they will be inclined to switch to the trendy strategy, or abandon the one out-of-fashion.

Most of the convergence results presented here concern properties of PURE INFIMMDYN in partnership games. We also establish the property of support separation in finite time, which has not yet been established for any other evolutionary game dynamics.

This property allows us to determine the support of Nash equilibria after finitely many iterations, which is valuable from a computational point of view.

It is worth noting that we will focus throughout the paper on dynamics in discrete time rather in continuous time, as the latter would anyway require some auxiliary discretization policy when implemented. We discuss such policies in Section 3.5 for continuous-time adjustment dynamics based upon our new class of dynamics.

The paper is organized as follows. Section 2 recalls theory and notations pertaining to evolutionary game dynamics. Here, we also provide a formalization of discrete-time revision protocols, which, to our knowledge, has not been done before. In Section 3 we develop the theoretical motivation behind our new class of evolutionary dynamics, starting from the basic notion of invasion barrier. We also introduce and discuss PURE INFIMM-DYN, a special instance from our class admitting a straightforward interpretation from a parsimonious microfoundation point-of-view. Section 4 provides general results holding for INFIMM-DYN, as well as results holding for PURE INFIMM-DYN, especially in the case of partnership games. Finally, the last section provides some experimental results.

2. Evolutionary game dynamics

In this section we will review some principles of evolutionary dynamics. For a complete overview of the topic we refer to [36, 16, 31]. Here we adapt some concepts from a continuous-time to a discrete-time setting, which also may not have been done before in the literature.

2.1. Preliminaries and notations

In this paper, we will focus on two-player symmetric games, where each player has the same payoff function. Let $S = \{1, 2, \dots, n\}$ be the finite set of *pure strategies* available to the players and $A = (a_{ij})$ the $n \times n$ payoff or utility matrix [36], where a_{ij} is the payoff that a player gains when playing the strategy i against an opponent playing strategy j . Denote by $\mathbf{e} \in \mathbb{R}^n$ the vector in n -dimensional Euclidean space \mathbb{R}^n with unity entries. Then the set of *mixed strategies* is the standard simplex

$$\Delta = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{e}^\top \mathbf{x} = 1, \quad \mathbf{x} \geq 0\}.$$

The expected payoff that a player obtains by playing the strategy i against an opponent playing a mixed strategy \mathbf{x} is

$$\pi(\mathbf{e}^i | \mathbf{x}) = (A\mathbf{x})_i = \sum_{j \in S} a_{ij} x_j,$$

where $\mathbf{e}^i \in \mathbb{R}^n$ is the i -th column of the $n \times n$ identity matrix which represents the i -th vertex of Δ , and sometimes this is identified with the pure strategy (index) i .

Hence, the expected payoff received by adopting a mixed strategy \mathbf{y} is

$$\pi(\mathbf{y} | \mathbf{x}) = \sum_{i \in S} y_i \pi(\mathbf{e}^i | \mathbf{x}) = \mathbf{y}^\top A\mathbf{x}.$$

The expected payoff of the entire population is given by

$$\pi(\mathbf{x}) = \pi(\mathbf{x} | \mathbf{x}) = \mathbf{x}^\top A\mathbf{x}.$$

In the sequel, we will also use the following notations:

$$\pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) = \pi(\mathbf{y}|\mathbf{x}) - \pi(\mathbf{x}),$$

and

$$\pi(\mathbf{y} - \mathbf{x}) = \pi(\mathbf{y} - \mathbf{x}|\mathbf{y}) - \pi(\mathbf{y} - \mathbf{x}|\mathbf{x}).$$

Given any (pure or) mixed strategy \mathbf{x} , we denote by

$$\begin{aligned} \sigma(\mathbf{x}) &= \{i \in S : x_i > 0\} && \text{the support of } \mathbf{x}, \\ \tau_-(\mathbf{x}) &= \{i \in S : \pi(\mathbf{x} - \mathbf{e}^i|\mathbf{x}) > 0\} && \text{the losing pure strategies at } \mathbf{x}, \\ \tau_+(\mathbf{x}) &= \{i \in S : \pi(\mathbf{x} - \mathbf{e}^i|\mathbf{x}) < 0\} && \text{the winning pure strategies at } \mathbf{x} \text{ and} \\ \tau_0(\mathbf{x}) &= \{i \in S : \pi(\mathbf{x} - \mathbf{e}^i|\mathbf{x}) = 0\} && \text{the balancing pure strategies at } \mathbf{x}. \end{aligned} \quad (1)$$

The terms “losing” and “winning” above always refer to a performance below or above the average performance, not in an absolute sense. The support $\sigma(\mathbf{x})$ defines the set of pure strategies used with positive probability in state \mathbf{x} .

The *best replies* $\beta(\mathbf{x})$ against a mixed strategy \mathbf{x} is the set of strategies that maximize the expected payoff when played against \mathbf{x} , namely

$$\beta(\mathbf{x}) = \arg \max_{\mathbf{z} \in \Delta} \pi(\mathbf{z}|\mathbf{x}).$$

A strategy profile (pair) $(\mathbf{x}, \mathbf{y}) \in \Delta^2$ is a Nash equilibrium point if $\mathbf{x} \in \beta(\mathbf{y})$ and $\mathbf{y} \in \beta(\mathbf{x})$. Since we consider symmetric games, where the two players are indistinguishable, only symmetric pairs (\mathbf{x}, \mathbf{x}) of strategies are of interest. Therefore, we call a mixed strategy \mathbf{x} a (*symmetric*) *Nash (equilibrium) strategy* if \mathbf{x} is a best reply to itself, i.e., $\mathbf{x} \in \beta(\mathbf{x})$ or

$$\pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) \leq 0 \quad \text{for all } \mathbf{y} \in \Delta.$$

This implies that $\tau_+(\mathbf{x}) = \emptyset$, which in turn implies that for all $\mathbf{x} \in \Delta$, $\sigma(\mathbf{x}) \subseteq \tau_0(\mathbf{x})$. For this reason we call $\tau_0(\mathbf{x})$ the *extended support of a Nash strategy*. So every pure strategy in the support of a Nash strategy \mathbf{x} is balancing, while all strategies outside $\sigma(\mathbf{x})$ earn a payoff that is less than or equal to $\pi(\mathbf{x})$, i.e., they are either losing or balancing.

2.2. Evolutionary and Neutral stability

A key concept in evolutionary game theory is that of evolutionarily stable strategy [20, 21, 22]. Consider a scenario where individuals are repeatedly drawn from a large population to play a two-player symmetric game. Each player is not supposed to behave rationally or have a complete knowledge of the details of the game, but he acts according to a pre-programmed pure or mixed strategy.

Let $\mathbf{x} \in \Delta$ be the *incumbent* strategy. and consider a different *mutant* strategy $\mathbf{y} \in \Delta$. Let the share of mutants in the (*postentry*) population be ε , where $\varepsilon \in (0, 1)$. The *postentry payoff* of the incumbent strategy is thus $\pi(\mathbf{x}|\mathbf{w})$ and that of the mutant strategy $\pi(\mathbf{y}|\mathbf{w})$ where $\mathbf{w} = \varepsilon\mathbf{y} + (1 - \varepsilon)\mathbf{x} \in \Delta$. Biological intuition suggests that evolutionary forces select against the mutant strategy if and only if its postentry payoff is lower than that of the incumbent strategy,

$$\pi(\mathbf{x}|\mathbf{w}) > \pi(\mathbf{y}|\mathbf{w}).$$

A strategy $\mathbf{x} \in \Delta$ is said to be an *evolutionarily stable strategy* (ESS) if this inequality holds for any mutant strategy $\mathbf{y} \neq \mathbf{x}$, granted the population share of mutants is sufficiently small [20, 22]. Specifically, this is true if and only if \mathbf{x} is a Nash equilibrium (*equilibrium condition for ESS*) and for any best reply \mathbf{y} to \mathbf{x} different from \mathbf{x} , i.e., such that $\pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) = 0$ and $\mathbf{y} \neq \mathbf{x}$, we have $\pi(\mathbf{x} - \mathbf{y}|\mathbf{y}) > 0$ (*stability condition for ESS*).

A less stringent concept than ESS is that of *neutral stability*. Relaxing the above conditions for ESS, we say that a Nash equilibrium strategy \mathbf{x} is *neutrally stable* if any best reply \mathbf{y} to \mathbf{x} satisfies $\pi(\mathbf{x} - \mathbf{y}|\mathbf{y}) \geq 0$.

2.3. Population games and evolutionary game dynamics

Population games provide a general model of strategic interactions among large numbers of anonymous agents. In this article, we will focus on games played by a single population, in which agents are not differentiated by roles. Thus, each agent from a unit-mass population chooses a strategy from S .

The *population state* at a given moment t is fully described by the distribution of strategy choices $\mathbf{x}^{(t)} \in \Delta$. The payoff to strategy i , denoted by $F_i : \Delta \rightarrow \mathbb{R}$, is a continuous function of the population state. If the set of strategies S is fixed, the *population game* is fully described by $F : \Delta \rightarrow \mathbb{R}^n$.

The most common example of population game is a *pairwise contest* where individuals are randomly matched to play a symmetric game with payoff matrix A . As already seen in Section 2.1, if \mathbf{x} is the current population state, then $\pi(\mathbf{e}^i|\mathbf{x}) = (A\mathbf{x})_i$ is the average payoff of strategy i in population. Thus, in this case, $F(\mathbf{x}) = A\mathbf{x}$. In general, the payoff function may not take a linear form, e.g., if more than two individuals participate in a contest [9].

In a considerable part of the literature, continuous-time evolutionary dynamics are preferred, which take the form [31]

$$\dot{\mathbf{x}}^{(t)} = \mathbf{f}^F(\mathbf{x}^{(t)}), \quad (2)$$

where \mathbf{f}^F is a suitable system function involving F . A possible time discretization can be obtained from (2) by putting

$$V^F(\mathbf{x}^{(t)}) = \mathbf{x}^{(t)} + \kappa(\mathbf{x}^{(t)})\mathbf{f}^F(\mathbf{x}^{(t)}) \quad (3)$$

where $\kappa(\mathbf{x}^{(t)})$ is a positive scaling factor, which may depend on $\mathbf{x}^{(t)}$ and ensures that $V^F(\mathbf{x}^{(t)}) \in \Delta$. This yields the following discrete-time dynamics over Δ :

$$\mathbf{x}^{(t+1)} = V^F(\mathbf{x}^{(t)}). \quad (4)$$

If V^F takes the form

$$x_i^{(t+1)} = x_i^{(t)} g_i(\mathbf{x}),$$

where the functions g_i satisfy $\sum_i x_i g_i = 1$ on Δ , we speak of *imitation dynamics*. The simplex and its faces are invariant under imitation dynamics. This observation has two implications: fixed points under imitation dynamics may not be Nash equilibria, and trajectories never reach the boundary of the face from which it started in finite time.

Imitation dynamics are said to be *payoff monotonic* [36] if

$$g_i(\mathbf{x}) > g_j(\mathbf{x}) \iff F_i(\mathbf{x}) > F_j(\mathbf{x}).$$

For payoff monotonic imitation dynamics the *folk theorem* holds: Nash equilibria are rest points, strict Nash equilibria are asymptotically stable, and rest points that are stable or limit points of interior orbits are Nash equilibria. The converse implications in general do not hold.

A particular case of payoff monotonic dynamics constitutes the most popular evolutionary game dynamics: the Replicator Dynamics (RD). This RD was originally introduced to model a discrete-time selection process, and later adapted by [33] for continuous-time game dynamics. One of the first systematic accounts on the discrete-time version for pairwise contests can be found in [21, Appendix]:

$$x_i^{(t+1)} = x_i^{(t)} \frac{\pi(\mathbf{e}^i | \mathbf{x}) + C}{\pi(\mathbf{x}) + C}, \quad (5)$$

where C is a constant needed to render the denominator positive and can be considered as a “background payoff”. It follows that under RD we have a growth in the fraction of i -strategists if $i \in \tau_+(\mathbf{x})$, whereas if $i \in \tau_-(\mathbf{x})$ then the share of i -strategists will decrease. Finally, we have a fixed point if $\sigma(\mathbf{x}) \subseteq \tau_0(\mathbf{x})$.

2.4. Discrete-time revision protocols

Instead of defining evolutionary dynamics directly, we derive them from explicit models of myopic individual choice. This can be accomplished by using the notion of *revision protocol* $\rho_{ij} : \mathbb{R}^n \times \Delta \rightarrow \mathbb{R}_+$ proposed in [17]. Here, seemingly for the first time in literature, we adapt the theory to the discrete-time dynamic setting.

Given a population state \mathbf{x} of a population game F , the revision protocol specifies for each pair of strategies i and j a nonnegative number $\rho_{ij}(F(\mathbf{x}), \mathbf{x})$ representing the share of current i -strategists that switch to strategy j (we abbreviate $\rho_{ij}(F(\mathbf{x}), \mathbf{x})$ with ρ_{ij} wherever possible). Moreover, we take only proper switches into account and therefore $\rho_{ii} = 0$ for all $i \in S$. Hence, we necessarily have that $\sum_{j \in S} \rho_{ij} \leq 1$ for all $i \in S$ and that $0 \leq \rho_{ij} \leq 1$ for all $i, j \in S$. Clearly, the share of i -strategists that decide not to switch is given by $1 - \sum_{j \in S} \rho_{ij}$.

An interpretation of the revision protocol consists in thinking that, as time passes, agents are randomly drawn from the population and granted the opportunity to switch strategies. When an i -strategist receives such an opportunity, he switches to strategy j with probability ρ_{ij} . For future reference, we observe that the sum $\sum_{j \in S} \rho_{ij}$ describes the probability at which i -strategists actively revise their choice of strategy. For example, if the revision protocol is of the imitative form $\rho_{ij} = x_j \hat{\rho}_{ij}$ for some well-behaved $\hat{\rho}_{ij}$, we can suppose that the revising agent selects his candidate strategy j not by drawing a strategy at random, but by drawing an opponent at random and observing the opponent’s strategy.

A population game F and a revision protocol ρ generate together an evolutionary dynamics called *mean dynamics* for F and ρ :

$$x_i^{(t+1)} = V_i^F(\mathbf{x}^{(t)}) = \sum_{j \in S} x_j^{(t)} \rho_{ji}(F(\mathbf{x}^{(t)}), \mathbf{x}^{(t)}) - x_i^{(t)} \sum_{j \in S} \rho_{ij}(F(\mathbf{x}^{(t)}), \mathbf{x}^{(t)}) + x_i^{(t)}.$$

Here the first term describes the inflow into strategy i from other strategies, whereas the second term provides the outflow from strategy i to other strategies. The difference

between these two terms is the net share of change in the use of strategy i , that added to the original share provides us with the new share of i -strategists. For example, the revision protocol of RD is

$$\rho_{ij} = x_j \frac{[\pi(\mathbf{e}^j - \mathbf{e}^i | \mathbf{x})]_+}{\pi(\mathbf{x}) + C},$$

where $[x]_+ = \max\{0, x\}$. The probability to switch from i to j increases if the average payoff gained by j -strategists is higher than that of the i -strategists.

3. Infection and immunization processes

Now we introduce a new class of evolutionary dynamics, inspired by infection and immunization processes. But before let us shortly recall the key concepts of barriers in evolutionary game theory, and prove an immunization result.

3.1. Score functions and barriers

Let $\mathbf{x} \in \Delta$ be the *incumbent* population state, \mathbf{y} be the *mutant* population invading \mathbf{x} and let $\mathbf{z} = (1 - \varepsilon)\mathbf{x} + \varepsilon\mathbf{y}$ be the population state obtained by injecting into \mathbf{x} a small share of \mathbf{y} -strategists. Then the *score function* of \mathbf{y} versus \mathbf{x} (introduced by [8]) is given by

$$h_{\mathbf{x}}(\mathbf{y}, \varepsilon) = \pi(\mathbf{y} - \mathbf{x} | \mathbf{z}) = \varepsilon\pi(\mathbf{y} - \mathbf{x}) + \pi(\mathbf{y} - \mathbf{x} | \mathbf{x}).$$

Following [9], we define the (*neutral*) *invasion barrier* $b_{\mathbf{x}}(\mathbf{y})$ of $\mathbf{x} \in \Delta$ against any mutant strategy \mathbf{y} as the largest population share of \mathbf{y} -strategists such that for all smaller positive population shares ε , \mathbf{x} earns a higher or equal payoff than \mathbf{y} in the post-entry population \mathbf{z} . Formally

$$b_{\mathbf{x}}(\mathbf{y}) = \inf(\{\varepsilon \in (0, 1) : h_{\mathbf{x}}(\mathbf{y}, \varepsilon) > 0\} \cup \{1\}).$$

Now \mathbf{x} is neutrally stable if and only if it is protected by a positive invasion barrier [9]:

$$\mathbf{x} \text{ is neutrally stable if and only if } b_{\mathbf{x}}(\mathbf{y}) > 0 \text{ for all } \mathbf{y} \in \Delta.$$

Given populations $\mathbf{x}, \mathbf{y} \in \Delta$, we say that \mathbf{x} is *immune* against \mathbf{y} if $b_{\mathbf{x}}(\mathbf{y}) > 0$. Trivially, a population is always immune against itself. Note that, \mathbf{x} is immune against \mathbf{y} if and only if either $\pi(\mathbf{y} - \mathbf{x} | \mathbf{x}) < 0$ or $\pi(\mathbf{y} - \mathbf{x} | \mathbf{x}) = 0$ and $\pi(\mathbf{y} - \mathbf{x}) \leq 0$. If $\pi(\mathbf{y} - \mathbf{x} | \mathbf{x}) > 0$ we say that \mathbf{y} is *infective* for \mathbf{x} . Hence, the set of infective strategies for \mathbf{x} is given by

$$\Upsilon(\mathbf{x}) = \{\mathbf{y} \in \Delta : \pi(\mathbf{y} - \mathbf{x} | \mathbf{x}) > 0\}.$$

Consider $\mathbf{y} \in \Upsilon(\mathbf{x})$; clearly, this implies $b_{\mathbf{x}}(\mathbf{y}) = 0$. If we allow for invasion of a share ε of \mathbf{y} -strategists as long as the score function of \mathbf{y} versus \mathbf{x} is positive, at the end we will have a share of $\delta_{\mathbf{y}}(\mathbf{x})$ mutants in the post-entry population, where

$$\delta_{\mathbf{y}}(\mathbf{x}) = \inf(\{\varepsilon \in (0, 1) : h_{\mathbf{x}}(\mathbf{y}, \varepsilon) \leq 0\} \cup \{1\}).$$

Note that if \mathbf{y} is infective for \mathbf{x} , then $\delta_{\mathbf{y}}(\mathbf{x}) > 0$, whereas if \mathbf{x} is immune against \mathbf{y} , then $\delta_{\mathbf{y}}(\mathbf{x}) = 0$. Further note that all the above concepts can be straightforwardly extended to contests with more than two participants and/or correlated individual behavior, where the score functions may be nonlinear in ε ; see, e.g. [9] and references therein. In our

two-person context, score functions are affine-linear, so that there is a simpler expression for $\delta_{\mathbf{y}}(\mathbf{x})$ if $\mathbf{y} \in \Upsilon(\mathbf{x})$:

$$\delta_{\mathbf{y}}(\mathbf{x}) = \begin{cases} \min \left\{ \frac{\pi(\mathbf{x}-\mathbf{y}|\mathbf{x})}{\pi(\mathbf{y}-\mathbf{x})}, 1 \right\}, & \text{if } \pi(\mathbf{y}-\mathbf{x}) < 0 \\ 1, & \text{otherwise.} \end{cases} \quad (6)$$

It is easy to see that $\delta_{\mathbf{y}}(\mathbf{x})$ is scale- and column-wise shift-invariant with respect to the payoff matrix A , i.e., it yields the same value for any payoff matrix of the form $sA + \mathbf{e}\mathbf{v}^\top$, where $s > 0$ and $\mathbf{v} \in \mathbb{R}^n$.

The next result shows that if we allow for invasion of a population in state \mathbf{x} by an infective strategy \mathbf{y} , and the extent of this infection is $\delta_{\mathbf{y}}(\mathbf{x})$, then the post-entry population will become immune against \mathbf{y} .

Proposition 1. *Let $\mathbf{y} \in \Upsilon(\mathbf{x})$ and let*

$$\mathbf{z} = [1 - \delta_{\mathbf{y}}(\mathbf{x})]\mathbf{x} + \delta_{\mathbf{y}}(\mathbf{x})\mathbf{y}.$$

Then \mathbf{z} is immune against \mathbf{y} .

Proof. Let us write δ for $\delta_{\mathbf{y}}(\mathbf{x})$. Note that $\delta > 0$ since $\mathbf{y} \in \Upsilon(\mathbf{x})$. Now if $\delta = 1$ then $\mathbf{z} = \mathbf{y}$ and, therefore, \mathbf{z} is trivially immune against \mathbf{y} . Hence, we can assume $\delta < 1$ in the sequel. By linearity of $\mathbf{y}^\top A \mathbf{z}$ in \mathbf{y} we can express the score function of any convex combination representing post-entry mixtures, $\mathbf{y}^\lambda = (1 - \lambda)\mathbf{x} + \lambda\mathbf{y}$ of \mathbf{x} and \mathbf{y} , versus any other such convex combination \mathbf{y}^μ as follows:

$$h_{\mathbf{y}^\mu}(\mathbf{y}^\lambda, \varepsilon) = (\lambda - \mu)h_{\mathbf{x}}(\mathbf{y}, \varepsilon\lambda + (1 - \varepsilon)\mu).$$

By setting $\lambda = 1$ and $\mu = \delta$, we obtain

$$h_{\mathbf{z}}(\mathbf{y}, \varepsilon) = \varepsilon(1 - \delta)h_{\mathbf{x}}(\mathbf{y}, \varepsilon + (1 - \varepsilon)\delta).$$

Note that $h_{\mathbf{z}}(\mathbf{y}, \varepsilon) \leq 0$ for all $\varepsilon \in (0, 1)$, by definition of δ and since

$$\varepsilon + (1 - \varepsilon)\delta = \delta + \varepsilon(1 - \delta) > \delta.$$

Hence we have $b_{\mathbf{z}}(\mathbf{y}) = 1$, and \mathbf{z} is immune against \mathbf{y} . □

3.2. Infection and immunization dynamics (INFIMMDYN)

The core idea of our method consists in selecting a strategy \mathbf{y} which is infective for the current population \mathbf{x} . By allowing for invasion, according to Proposition 1, we obtain a new population \mathbf{z} which is immune to \mathbf{y} . This idea suggests the following class of new dynamics which for evident reasons is called *Infection and Immunization Dynamics* (INFIMMDYN):

$$\mathbf{x}^{(t+1)} = \delta_{\mathcal{S}(\mathbf{x}^{(t)})}(\mathbf{x}^{(t)})[\mathcal{S}(\mathbf{x}^{(t)}) - \mathbf{x}^{(t)}] + \mathbf{x}^{(t)}, \quad (7)$$

where $\mathcal{S} : \Delta \rightarrow \Delta$ is a *strategy selection* function, which returns an infective strategy for \mathbf{x} if it exists, or \mathbf{x} otherwise:

$$\mathcal{S}(\mathbf{x}) = \begin{cases} \mathbf{y} & \text{for some } \mathbf{y} \in \Upsilon(\mathbf{x}) \text{ if } \Upsilon(\mathbf{x}) \neq \emptyset, \\ \mathbf{x} & \text{otherwise.} \end{cases} \quad (8)$$

By reiterating this process of *immunization* we aim at reaching a population state \mathbf{x} that cannot be infected by any other strategy. If this is the case then \mathbf{x} is a fixed point under dynamics (7), but also a Nash strategy:

Theorem 1. *Let $\mathbf{x} \in \Delta$ be a strategy. Then the following statements are equivalent:*

- (a) $\Upsilon(\mathbf{x}) = \emptyset$: *there is no infective strategy for \mathbf{x} ;*
- (b) \mathbf{x} *is a Nash strategy;*
- (c) \mathbf{x} *is a fixed point under dynamics (7).*

Proof. A strategy \mathbf{x} is a Nash strategy if and only if $\pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) \leq 0$ for all $\mathbf{y} \in \Delta$. This is true if and only if $\Upsilon(\mathbf{x}) = \emptyset$. Further, $\delta = 0$ implies $\mathcal{S}(\mathbf{x}) = \mathbf{x}$. Conversely, if $\mathcal{S}(\mathbf{x})$ returns \mathbf{x} , then we are in a fixed point. By construction of $\mathcal{S}(\mathbf{x})$ this happens only if there is no infective strategy for \mathbf{x} . \square

If \mathbf{x} is not fixed under (7), i.e., is not a Nash strategy, straightforward intuition renders selection of an infective strategy in a way easier than it could seem at first glance. Let \mathbf{x} be the current population and let \mathbf{y} be a strategy. The *co-strategy* of \mathbf{y} with respect to \mathbf{x} is given by

$$\bar{\mathbf{y}}_{\mathbf{x}} = (1 + \bar{\varepsilon})\mathbf{x} - \bar{\varepsilon}\mathbf{y},$$

where

$$\bar{\varepsilon} = \max\{\varepsilon \in \mathbb{R} : (1 + \varepsilon)\mathbf{x} - \varepsilon\mathbf{y} \in \Delta\} \geq 0.$$

For any strategy \mathbf{y} , if both $\pi(\mathbf{y} - \mathbf{x}|\mathbf{x})$ and $\bar{\varepsilon}$ are not zero, then either $\mathbf{y} \in \Upsilon(\mathbf{x})$ or $\bar{\mathbf{y}}_{\mathbf{x}} \in \Upsilon(\mathbf{x})$ in an exclusive sense. If $\mathbf{y} = \mathbf{e}^i$, then $\bar{\mathbf{y}}_{\mathbf{x}}$ is the mixed strategy that maintains the strategies other than i in their original proportions in \mathbf{x} while scaling their total weight to 1. In this case, $\bar{\varepsilon} = x_i/(1 - x_i)$. Geometrically, $\bar{\mathbf{y}}_{\mathbf{x}}$ is the intersection between the simplex boundary and the half line originated in \mathbf{e}^i and passing through \mathbf{x} ; see Figure 1.

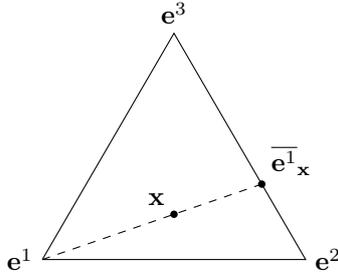


Figure 1: Example of co-strategy of the pure strategy \mathbf{e}^1 with respect to \mathbf{x} .

The next section introduces a particular instance of our new class of dynamics, where the strategy selection function returns only infective pure strategies or their relative co-strategies.

3.3. A pure strategy selection function

Depending on how we choose the function $\mathcal{S}(\mathbf{x})$ in (7), we may obtain different dynamics. One in particular, which is simple and leads to nice properties, consists in allowing only infective pure strategies or their respective co-strategies. This way, our equilibrium selection process closely resembles a vertex-pivoting method, as opposed to interior-point approaches like replicator dynamics or best-response dynamics [16].

Consider the strategy selection function $\mathcal{S}_{\text{Pure}}(\mathbf{x})$, which finds a pure strategy i maximizing $|\pi(\mathbf{e}^i - \mathbf{x}|\mathbf{x})|$, and returns \mathbf{e}^i , $\overline{\mathbf{e}}^i_{\mathbf{x}}$ or \mathbf{x} according to whether $i \in \tau_+(\mathbf{x})$, $i \in \tau_-(\mathbf{x}) \cap \sigma(\mathbf{x})$ or $i \in \tau_0(\mathbf{x})$: Let $\mathcal{M}(\mathbf{x})$ be a (randomly or otherwise selected) pure strategy such that

$$\mathcal{M}(\mathbf{x}) \in \arg \max \{ \pi(\mathbf{e}^i - \mathbf{x}|\mathbf{x}) : i \in \tau_+(\mathbf{x}) \} \cup \{ \pi(\mathbf{x} - \mathbf{e}^i|\mathbf{x}) : i \in \tau_-(\mathbf{x}) \cap \sigma(\mathbf{x}) \}.$$

Then $\mathcal{S}_{\text{Pure}}(\mathbf{x})$ can be written as

$$\mathcal{S}_{\text{Pure}}(\mathbf{x}) = \begin{cases} \mathbf{e}^i & \text{if } i = \mathcal{M}(\mathbf{x}) \in \tau_+(\mathbf{x}) \\ \overline{\mathbf{e}}^i_{\mathbf{x}} & \text{if } i = \mathcal{M}(\mathbf{x}) \in \tau_-(\mathbf{x}) \cap \sigma(\mathbf{x}) \\ \mathbf{x} & \text{otherwise.} \end{cases}$$

For obvious reasons, we refer to INFIMMDYN with selection function $\mathcal{S}_{\text{Pure}}(\mathbf{x})$ as PURE INFIMMDYN.

Note that the search space for an infective strategy is reduced from Δ to a finite set. Therefore, it is not obvious that $\mathcal{S}_{\text{Pure}}(\mathbf{x})$ is a well-defined selection function, i.e., it satisfies (8). The next result shows that this is indeed the case.

Proposition 2. *There exists an infective strategy for \mathbf{x} if and only if $\mathcal{S}_{\text{Pure}}(\mathbf{x})$ is infective for \mathbf{x} .*

Proof. By construction of $\mathcal{S}_{\text{Pure}}$, if $\mathcal{S}_{\text{Pure}}(\mathbf{x}) = \mathbf{x}$ then for all $i \in S$, $i \notin \tau_+(\mathbf{x})$, which in turn implies that there exists no infective strategy for \mathbf{x} . Otherwise, if $\mathcal{S}_{\text{Pure}}(\mathbf{x}) = \mathbf{e}^i$ then \mathbf{e}^i is infective, since $\pi(\mathbf{e}^i|\mathbf{x}) > \pi(\mathbf{x})$, whereas if $\mathcal{S}_{\text{Pure}} = \overline{\mathbf{e}}^i_{\mathbf{x}}$, then we know that $\pi(\mathbf{e}^i|\mathbf{x}) < \pi(\mathbf{x})$. But then we have

$$\pi(\mathbf{x} - \overline{\mathbf{e}}^i_{\mathbf{x}}|\mathbf{x}) = \bar{\varepsilon} \pi(\mathbf{e}^i - \mathbf{x}|\mathbf{x}) = \frac{x_i}{1 - x_i} [\pi(\mathbf{e}^i|\mathbf{x}) - \pi(\mathbf{x})] < 0,$$

which proves that $\overline{\mathbf{e}}^i_{\mathbf{x}}$ is infective. Hence, $\mathcal{S}_{\text{Pure}}$ always returns an infective strategy or \mathbf{x} . To prove the first implication, let $\mathbf{y} \in \Delta$ be infective for \mathbf{x} . Then

$$0 < \pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) = \sum_{i=1}^n y_i \pi(\mathbf{e}^i - \mathbf{x}|\mathbf{x}).$$

It follows that there exists at least one infective pure strategy for \mathbf{x} , which implies that $\mathcal{S}_{\text{Pure}}(\mathbf{x})$ is infective. The converse trivially holds. \square

3.4. PURE INFIMMDYN: fashions and the revision protocol

A possible interpretation of PURE INFIMMDYN is as follows: as time passes by, an advertisement on the basis of the aggregate behavior of the population tells the agents that a certain pure strategy is trendy or is out-of-fashion. A strategy is trendy, if it is the best performing one in terms of payoff in the population, whereas it is out-of-fashion, if it is the worst performing strategy still alive in the population. The choice among them depends on which strategy deviates most from the average payoff. Note that if \mathbf{x} is the current population, k is trendy if and only if $\mathcal{S}_{\text{Pure}}(\mathbf{x}) = \mathbf{e}^k$, whereas k is out-of-fashion if and only if $\mathcal{S}_{\text{Pure}}(\mathbf{x}) = \overline{\mathbf{e}^k_{\mathbf{x}}}$; for obvious reasons we only treat the nontrivial case $\mathcal{S}_{\text{Pure}}(\mathbf{x}) \neq \mathbf{x}$.

Consider a scenario where agents can gather informations only about the announced strategy k which in turn keeps its trendy (out-of-fashion) status as long as its score function with respect to the population is positive (negative). As long as a strategy remains trendy, agents playing other strategies will switch to it when they receive the possibility of re-evaluating their strategy (changing this way the population state). On the other hand, as long as a strategy is out-of-fashion, agents playing that strategy will switch to another strategy with a probability proportional to its current popularity. Once a strategy loses its status, a new advertisement will be done on the basis of the current population aggregate behaviour.

In terms of a discrete-time revision protocol, if k is advertised as a trendy pure strategy, then the revision protocol of the agents becomes

$$\rho_{ij} = \begin{cases} \delta_{\mathbf{e}^k}(\mathbf{x}) & \text{if } j = k \text{ and } i \neq k \\ 0 & \text{otherwise,} \end{cases}$$

where individuals playing a non trendy strategy switch to the trendy one with probability $\delta_{\mathbf{e}^k}(\mathbf{x})$. While if k is announced as out-of-fashion then the revision protocol of the agents becomes

$$\rho_{ij} = \begin{cases} \frac{x_j}{1-x_k} \delta_{\overline{\mathbf{e}^k_{\mathbf{x}}}}(\mathbf{x}) & \text{if } i = k \text{ and } j \neq k \\ 0 & \text{otherwise,} \end{cases}$$

where individuals playing the out-of-fashion strategy abandon it and switch to a different one with a probability proportional to its frequency (current popularity).

3.5. Other strategy selection functions; discretization of continuous time

There are many other possibilities to define the selection function. Before we discuss this in a general context, let us go back in time and consider one of the first game dynamics ever studied [10], namely *fictitious play*: consider a finite population of agents, which is iteratively incremented by injecting a new individual. This new player chooses (in a rational way) the “preprogrammed” strategy that he will play for the rest of the game from the set of best replies to the strategy mix in the current population. At time $t + 1$, the new entrant chooses a strategy $\mathbf{r}^{(t+1)} \in \beta(\mathbf{x}^{(t)})$, where $\mathbf{x}^{(t)} = \sum_{k=1}^t \mathbf{r}^{(k)}/t$. The resulting discrete-time dynamics that reflects the evolution of the distribution of strategies in the population is given by

$$\mathbf{x}^{(t+1)} = \frac{\mathbf{r}^{(t+1)} + \mathbf{x}^{(t)}}{t+1}, \quad \mathbf{r}^{(t+1)} \in \beta(\mathbf{x}^{(t)}). \quad (9)$$

Going from discrete generations to a continuous-time dynamics, the *best-response dynamics* can be derived [14, 15], which takes the form of a differential inclusion:

$$\dot{\mathbf{x}} \in \beta(\mathbf{x}) - \mathbf{x}. \quad (10)$$

An alternative interpretation of (10) is that in a very large population, a small fraction of individuals revise their strategy, choosing best replies to the current mean population strategy \mathbf{x} . Such players behave rationally, but myopically.

As long as $\beta(\mathbf{x}(t)) = \{\mathbf{e}^i\}$ for all times $t \in [0, T]$ along a trajectory of dynamics (10), we may integrate the system to

$$\mathbf{x}(t) = (1 - e^{-t})\mathbf{e}^i + e^{-t}\mathbf{x}(0), \quad t \in [0, T]. \quad (11)$$

Hence the movement is along a straight ray emanating from $\mathbf{x}(0)$ directed towards \mathbf{e}^i , and we could take, e.g., the point $\mathbf{x}(T)$ as the successive iterate of $\mathbf{x}(0)$ in a discrete dynamics. However, it is unclear whether $\mathbf{x}(T)$ is still infective by \mathbf{e}^i , and this can be overcome by simply defining

$$\mathcal{S}_\beta(\mathbf{x}) = \begin{cases} \mathbf{e}^i, & \text{if } i \in \tau_+(\mathbf{x}) \text{ and } \mathbf{e}^i \in \beta(\mathbf{x}), \\ \mathbf{x}, & \text{else.} \end{cases} \quad (12)$$

More generally, let us start from a continuous-time evolutionary *adjustment* dynamics (2) in the sense of [16, p.98], i.e., $\mathbf{x}^\top A^\top \mathbf{f}^F(\mathbf{x}) = [\mathbf{f}^F(\mathbf{x})]^\top A \mathbf{x} > 0$ unless $\mathbf{f}^F(\mathbf{x}) = \mathbf{o}$, and build a discrete-time dynamics according to (4). Then define

$$\mathcal{S}_F(\mathbf{x}) = \begin{cases} V^F(\mathbf{x}), & \text{if } \mathbf{f}^F(\mathbf{x}) \neq \mathbf{o}, \\ \mathbf{x}, & \text{else.} \end{cases} \quad (13)$$

Here, $\mathcal{S}_F(\mathbf{x})$ satisfies (8) if and only if it holds that

$$\mathbf{f}^F(\mathbf{x}) = \mathbf{o} \iff \mathbf{x} \text{ is a Nash strategy.}$$

This is true for example in the cases of continuous-time best-response dynamics (here by taking a selection $\mathbf{f}^{BR}(\mathbf{x}) \in \beta(\mathbf{x})$, which may be multivalued), the Brown-von Neumann-Nash dynamics [11, 24], the Maynard Smith-dynamics [32, 30], and the projection dynamics [17]. By contrast, the RD does not satisfy the property. In fact, if started from a face of the simplex, it may end up with a fixed point which is not a Nash equilibrium. For a closely related discussion see [29, Section 5]. A way of getting rid of the problem with RD is to add a “recovery” condition, yielding a strategy selection function like the following one:

$$\mathcal{S}_{RD}(\mathbf{x}) = \begin{cases} \mathbf{e}^i, & \text{if } i \in \tau_+(\mathbf{x}) \setminus \sigma(\mathbf{x}), \\ RD(\mathbf{x}), & \text{if } \tau_+(\mathbf{x}) \setminus \sigma(\mathbf{x}) = \emptyset, \end{cases} \quad (14)$$

where $RD(\mathbf{x})$ can be considered a discretization of the continuous-time replicator equations according to (3), or simply a step of the discrete-time replicator dynamics (5). It is straightforward to see that the strategy selection function in (14) satisfies (8). Whenever there exist pure strategies not in the support of \mathbf{x} , which are indeed infective for \mathbf{x} , these strategies will appear in the population after the subsequent INFIMMDYN steps. Indeed,

the replicator step will be taken only if $\tau_+(\mathbf{x}) \setminus \sigma(\mathbf{x})$ is empty. Note that if we start INFIMMODYN with the \mathcal{S}_{RD} strategy selection function from the interior of the simplex, the “recovery” condition of (14) will never be met, as $\sigma(\mathbf{x}^{(t)}) = S$ for all $t \geq 0$. If we define $RD(\mathbf{x})$ as (5), we obtain an interesting variant of the standard discrete-time RD, where either we take a complete or partial replicator step, according to the value of $\delta_{RD(\mathbf{x})}(\mathbf{x})$.

Note that any possibly occurring ties can be broken by an arbitrary rule (e.g., randomly or by taking always the strategy with the smallest label). As a final remark, we stress that by definition of INFIMMODYN, any choice of strategy selection function renders forward-invariance of the standard simplex, because the updated strategy is a convex combination of the previous strategy and the one selected by \mathcal{S} .

4. INFIMMODYN for partnership games

This section is devoted to providing both general theoretical results characterizing all possible infection and immunization dynamics, as well as specific properties holding for PURE INFIMMODYN, in the case of partnership games.

Beforehand let us note the general property for changes in average payoff along trajectories; for convenient notation we abbreviate \mathbf{x} for $\mathbf{x}^{(t)}$ and δ for $\delta_{\mathcal{S}(\mathbf{x})}(\mathbf{x})$ on the right-hand side below. Then we have

$$\pi(\mathbf{x}^{(t+1)}) - \pi(\mathbf{x}^{(t)}) = \delta [h_{\mathbf{x}}(\mathcal{S}(\mathbf{x}), \delta) + \pi(\mathbf{x}|\mathcal{S}(\mathbf{x})) - \pi(\mathbf{x})]. \quad (15)$$

Partnership games include coordination and anti-coordination games, and sometimes are called doubly symmetric games, because they are symmetric games with symmetric payoff matrices $A = A^\top$.

4.1. Population payoff increases under INFIMMODYN

The following result can be seen as a variant of the celebrated Theorem of Natural Selection, see, e.g. [16]: we show that average payoff is strictly increasing along any non-constant trajectory of the dynamics (7), provided that the payoff matrix is symmetric.

Theorem 2. *Let $\{\mathbf{x}^{(t)}\}_{t \geq 0}$ be a trajectory of (7) with $A = A^\top$. Then for all $t \geq 0$,*

$$\pi(\mathbf{x}^{(t+1)}) \geq \pi(\mathbf{x}^{(t)}),$$

with equality if and only if $\mathbf{x}^{(t)} = \mathbf{x}^{(t+1)}$.

Proof. Again, let us write \mathbf{x} for $\mathbf{x}^{(t)}$ and δ for $\delta_{\mathcal{S}(\mathbf{x})}(\mathbf{x})$. By double symmetry of the game (i.e., due to symmetry of the payoff matrix A), we get for any two strategies $\pi(\mathbf{y}|\mathbf{x}) = \pi(\mathbf{x}|\mathbf{y})$. For $\mathbf{y} = \mathcal{S}(\mathbf{x})$ the construction of $\mathcal{S}(\mathbf{x})$ implies $\pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) \geq 0$. Further, (15) yields now

$$\pi(\mathbf{x}^{(t+1)}) - \pi(\mathbf{x}^{(t)}) = \delta [h_{\mathbf{x}}(\mathbf{y}, \delta) + \pi(\mathbf{y} - \mathbf{x}|\mathbf{x})].$$

If $\mathbf{x}^{(t+1)} \neq \mathbf{x}^{(t)}$, then \mathbf{x} is no Nash strategy, and $\mathbf{y} = \mathcal{S}(\mathbf{x})$ returns by Theorem 1 an infective strategy. Hence $\delta > 0$ and

$$h_{\mathbf{x}}(\mathbf{y}, \delta) + \pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) \geq \pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) > 0$$

(in fact, if $\delta < 1$, then even $h_{\mathbf{x}}(\mathbf{y}, \delta) = 0$), so that we obtain a strict increase of the population payoff. On the other hand, if $\pi(\mathbf{x}^{(t+1)}) = \pi(\mathbf{x}^{(t)})$, then the above equation implies $\delta = 0$ or

$$h_{\mathbf{x}}(\mathbf{x}, \delta) = \pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) = 0,$$

due to nonnegativity of both quantities above. In particular, we have $\delta = 0$ or $\pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) = 0$. In both cases, $\mathbf{y} = \mathcal{S}(\mathbf{x})$ cannot be infective for \mathbf{x} . Thus $\Upsilon(\mathbf{x}) = \emptyset$ and \mathbf{x} must be a fixed point, according to Theorem 1. This establishes the last assertion of the theorem. \square

So INFIMMDYN generates a *growth transformation* on Δ for the population payoff $\pi(\mathbf{x})$. Despite the fact that this is a quadratic form in \mathbf{x} , we cannot invoke general convergence principles for continuous growth transformations since, unlike many familiar evolutionary game dynamics, the transformation given by INFIMMDYN may exhibit discontinuities. In this respect, INFIMMDYN resembles the best-response paradigm, but the former behaves more nicely, as will be detailed below.

Given a strictly increasing sequence $s = \{s_j\}_{j=0}^{\infty}$ on \mathbb{N} and a point $\mathbf{x}_0 \in \Delta$, we denote with $\xi(\mathbf{x}_0, s)$ the subsequence of the trajectory of (7) starting from point \mathbf{x}_0 , which is given by

$$\xi(\mathbf{x}_0, s) = \{\mathbf{x}^{(s_j)}\}_{j=0}^{\infty}.$$

In the sequel, we will omit the starting point where unnecessary, writing just $\xi(s)$, and we will refer to $\xi(s)$ as a subsequence of (7).

Further, as usual, denote by

$$\omega(\mathbf{x}^{(0)}) = \{\mathbf{q} \in \Delta : \mathbf{x}^{(s_j)} \rightarrow \mathbf{q} \text{ as } j \rightarrow \infty \text{ for some subsequence } \xi(\mathbf{x}^{(0)}, s)\}$$

the set of all limit points (the *omega-limit*) of the trajectory starting at $\mathbf{x}^{(0)}$.

Lemma 1. *Let $\{\mathbf{x}^{(t)}\}_{t \geq 0}$ be a trajectory of (7) with $A = A^\top$. Then there exists γ such that the following properties hold*

1. $\pi(\mathbf{x}^{(t)}) \leq \gamma$ for all $t \geq 0$;
2. $\lim_{t \rightarrow \infty} \pi(\mathbf{x}^{(t)}) = \gamma$;
3. $\pi(\mathbf{q}) = \gamma$, for all $\mathbf{q} \in \omega(\mathbf{x}^{(0)})$;
4. if $\pi(\mathbf{x}^{(T)}) = \gamma$ for some T , then $\mathbf{x}^{(T)}$ is a fixed point.

Proof. See appendix. \square

4.2. The role of Nash strategies and ESSs under INFIMMDYN

We now will establish that any limit point of a trajectory will be a Nash strategy, and that evolutionarily stable strategies are asymptotically stable under INFIMMDYN, given the base game is of partnership type, i.e., $A^\top = A$. It should be noted that the results of the previous Lemma 1 as well as the following Lemma 2 and Theorem 3 are well known for continuous-time dynamics satisfying the myopic adjustment condition, and where fixed points are Nash strategies and vice versa, see [28]; some of the arguments can already be found in [16]. However, our treatment differs in two aspects: time is discrete here, and the dynamical system lacks continuity; like in the best reply dynamics, the latter property carries over to continuous time, too.

Lemma 2. Let $\{\mathbf{x}^{(t)}\}_{t \geq 0}$ be a trajectory of (7) with $A = A^\top$ and abbreviate $\delta^{(t)} = \delta_{\mathcal{S}(\mathbf{x}^{(t)})}(\mathbf{x}^{(t)})$. Then

$$\lim_{t \rightarrow \infty} \pi(\mathcal{S}(\mathbf{x}^{(t)}) - \mathbf{x}^{(t)} | \mathbf{x}^{(t)}) = 0. \quad (16)$$

Furthermore, we have

$$\lim_{t \rightarrow \infty} \left[\delta^{(t)} \right]^2 \pi(\mathcal{S}(\mathbf{x}^{(t)}) - \mathbf{x}^{(t)}) = 0. \quad (17)$$

Proof. See appendix. \square

Theorem 3. Limit points of the PURE INFIMMDYN are Nash strategies, provided $A = A^\top$.

Proof. Let us write \mathcal{S} for $\mathcal{S}_{\text{Pure}}$. Assume that there exists a convergent subsequence $\xi(s)$ of (7) with limit point \mathbf{q} a non Nash strategy. Since \mathbf{q} is not a Nash strategy, by Theorem 1, $\mathcal{S}(\mathbf{x}^{(s_j)}) \in \Upsilon(\mathbf{x}^{(s_j)})$ for any $j \geq 0$. Moreover, there exists $i \in S$ such that $\pi(\mathbf{e}^i - \mathbf{q} | \mathbf{q}) > 0$. Hence, by construction of $\mathcal{S}_{\text{Pure}}$, for all $j \geq 0$,

$$\pi(\mathcal{S}(\mathbf{x}^{(s_j)}) - \mathbf{x}^{(s_j)} | \mathbf{x}^{(s_j)}) \geq \pi(\mathbf{e}^i - \mathbf{x}^{(s_j)} | \mathbf{x}^{(s_j)}).$$

By taking limits we obtain

$$\liminf_{j \rightarrow \infty} \pi(\mathcal{S}(\mathbf{x}^{(s_j)}) - \mathbf{x}^{(s_j)} | \mathbf{x}^{(s_j)}) \geq \liminf_{j \rightarrow \infty} \pi(\mathbf{e}^i - \mathbf{x}^{(s_j)} | \mathbf{x}^{(s_j)}) = \pi(\mathbf{e}^i - \mathbf{q} | \mathbf{q}) > 0,$$

which contradicts (16). This concludes the proof. \square

Theorem 4. A state \mathbf{q} is asymptotically stable for PURE INFIMMDYN if and only if \mathbf{q} is an ESS, provided $A = A^\top$.

Proof. If the payoff matrix is symmetric, every limit point of PURE INFIMMDYN is a Nash strategy by Theorem 3. Moreover ESSs are strict local maximizers of $\pi(\mathbf{x})$ over Δ and vice versa, see, e.g., [6]. If \mathbf{q} is asymptotically stable, then there exists a neighborhood U of \mathbf{q} in Δ such that any trajectory starting in U converges to \mathbf{q} . By Theorem 2 this implies that $\pi(\mathbf{q}) > \pi(\mathbf{x})$ for all $\mathbf{x} \in U \setminus \{\mathbf{q}\}$. Hence, \mathbf{q} is a strict local maximizer of $\pi(\mathbf{x})$ and therefore \mathbf{q} is an ESS. Conversely, if \mathbf{q} is an ESS then \mathbf{q} is a strict local maximizer of $\pi(\mathbf{x})$ and an isolated Nash strategy. Hence, there exists a neighborhood U of \mathbf{q} in Δ which contains no other Nash strategy. Choose $\rho > 0$ so small that $B = \{\mathbf{x} \in \Delta : \|\mathbf{x} - \mathbf{q}\| \leq \rho\} \subseteq U$. Then the set $C = \{\mathbf{x} \in \Delta : \|\mathbf{x} - \mathbf{q}\| = \rho\}$ is a compact subset of $U \setminus \{\mathbf{q}\}$, which means that $\eta = \pi(\mathbf{q}) - \max_{\mathbf{x} \in C} \pi(\mathbf{x}) > 0$. Next define

$$V = \{\mathbf{x} \in \Delta : \|\mathbf{x} - \mathbf{q}\| < \rho \text{ and } \pi(\mathbf{x}) > \pi(\mathbf{q}) - \eta\},$$

which is a (relatively to Δ) open neighborhood of \mathbf{q} contained in U . We claim that V is forward invariant under any INFIMMDYN; indeed, by construction we know

$$\pi\left((1 - \lambda)\mathbf{x}^{(t)} + \lambda\mathbf{x}^{(t+1)}\right) \geq \pi(\mathbf{x}^{(t)}) > \pi(\mathbf{q}) - \eta \quad \text{if } \mathbf{x}^{(t)} \in V \quad (18)$$

holds for all $\lambda \in [0, 1]$, which implies $\mathbf{x}^{(t+1)} \in V$, since otherwise we had $\|\mathbf{x}^{(t+1)} - \mathbf{q}\| \geq \rho$ (note (18), case $\lambda = 1$), which entailed $(1 - \lambda)\mathbf{x}^{(t)} + \lambda\mathbf{x}^{(t+1)} \in C$ for some $\lambda \in [0, 1]$, in contradiction to the definition of η in conjunction with (18). So by forward invariance of V , any limit point of the iterated sequence must belong to B and therefore to U . But then there is only one limit point, namely \mathbf{q} , which is the only Nash strategy in U . So any sequence starting in V converges to \mathbf{q} , and hence \mathbf{q} is asymptotically stable. \square

Note a strict local maximizer \mathbf{q} of a general function g need not necessarily be isolated in the set of stationary (Karush-Kuhn-Tucker) points of g . However, the dynamical properties of ESSs and Nash strategies under the continuous-time replicator dynamics (see, e.g. [6, Theorem 10]), imply exactly this: any ESS is an isolated Nash strategy; see also [36].

We close this subsection with an investigation of the relation between Lyapunov stability under PURE INFIMMDYN and neutral stability.

Theorem 5. *Suppose that a state \mathbf{q} is a Lyapunov stable fixed point for PURE INFIMMDYN. Then \mathbf{q} is neutrally stable, provided $A = A^\top$.*

Proof. If \mathbf{q} is a fixed point, we know by Theorem 1 that \mathbf{q} is a Nash strategy, so that $\tau_+(\mathbf{q}) = \emptyset$ and $\sigma(\mathbf{q}) \subseteq \tau_0(\mathbf{q})$. Further, by continuity, we may choose a closed neighbourhood V of \mathbf{q} in Δ such that

$$\sigma(\mathbf{x}) \supseteq \sigma(\mathbf{q}) \quad \text{and} \quad \tau_-(\mathbf{x}) \supseteq \tau_-(\mathbf{q}) \quad \text{for all } \mathbf{x} \in V.$$

Since \mathbf{q} is Lyapunov stable, we know there is a neighbourhood U of \mathbf{q} in Δ such that any trajectory $\mathbf{x}^{(t)}$ under PURE INFIMMDYN with $\mathbf{x}^{(0)} \in U$ satisfies $\mathbf{x}^{(t)} \in V$ for all $t \geq 0$. Since V is closed, also any limit point $\hat{\mathbf{x}} \in \omega(\mathbf{x}^{(0)})$ belongs to V . Further, Theorem 3 implies that $\hat{\mathbf{x}}$ is also a Nash strategy, so that we obtain

$$\tau_-(\mathbf{q}) \subseteq \tau_-(\hat{\mathbf{x}}) = S \setminus \tau_0(\hat{\mathbf{x}})$$

and

$$\sigma(\hat{\mathbf{x}}) \subseteq \tau_0(\hat{\mathbf{x}}) \subseteq S \setminus \tau_-(\mathbf{q}) = \tau_0(\mathbf{q}),$$

and also, on the other hand by choice of V ,

$$\sigma(\mathbf{q}) \subseteq \sigma(\hat{\mathbf{x}}) \subseteq \tau_0(\hat{\mathbf{x}}).$$

Hence, by virtue of $A^\top = A$,

$$\pi(\mathbf{x}^{(0)}) \leq \pi(\hat{\mathbf{x}}) = \pi(\mathbf{q}|\hat{\mathbf{x}}) = \pi(\hat{\mathbf{x}}|\mathbf{q}) = \pi(\mathbf{q}),$$

so that \mathbf{q} is a local maximizer of π over Δ . We now show that this inequality holds on a large set:

$$\pi(\mathbf{x}) \leq \pi(\mathbf{q}) \quad \text{for all } \mathbf{x} \in \Delta \text{ with } \sigma(\mathbf{x}) \subseteq \tau_0(\mathbf{q}). \quad (19)$$

Indeed, choose $s \in (0, 1)$ so small that $\mathbf{x}_s = (1-s)\mathbf{q} + s\mathbf{x} \in U$ where U is the above neighbourhood. Of course, also $\sigma(\mathbf{x}_s) \subseteq \tau_0(\mathbf{q})$ for all these s , so that $\pi(\mathbf{x}_s - \mathbf{q}|\mathbf{q}) = 0$ and

$$0 \geq \pi(\mathbf{x}_s) - \pi(\mathbf{q}) = \pi(\mathbf{x}_s - \mathbf{q}) = s^2\pi(\mathbf{x} - \mathbf{q}) = s^2[\pi(\mathbf{x}) - \pi(\mathbf{q})].$$

Since the latter expression never changes sign for any positive s , we arrive at (19). This property in turn is easily seen to be equivalent to the definition of neutral stability. Hence the assertion. \square

Remark: Unfortunately we are not aware of any proof of the result converse to Theorem 5, which would establish equivalence of Lyapunov stability under PURE INFIMMDYN and neutral stability. An analogous equivalence result for the replicator dynamics has been established in [6]. Note that the proof method of Theorem 5 can easily be adapted to the case of replicator dynamics, yielding an argument much simpler than the retrograde perturbation method in [6].

4.3. Convergence of PURE INFIMMDYN for partnership games

We will now establish conditions on a single trajectory (i.e., the starting point $\mathbf{x}^{(0)}$) which guarantee convergence of this trajectory. Note that these trajectory-specific conditions are weaker than the usual global regularity conditions on the payoff matrix A , namely that there are only finitely many Nash strategies in the game. Again, the latter property of A is generic, but as argued in [6], sometimes models have to be treated which in a larger context are degenerate. Examples are normal-form formulations of role-dependent contests or those with genetical (i.e., exogenous) constraints [16].

If the transition given by PURE INFIMMDYN were continuous, we could have inferred in a quite general way that $\omega(\mathbf{x}^{(0)})$ is a *connected* subset of Nash strategies, but this is not the case. We only know from continuity of π that

$$\omega(\mathbf{x}^{(0)}) \subseteq \{\mathbf{q} \in \Delta : \pi(\mathbf{q}) = \gamma \text{ and } \mathbf{q} \text{ is a Nash strategy}\} \quad (20)$$

from Lemma 1 and Theorem 3.

We need the following auxiliary result.

Lemma 3. *Let $\{\mathbf{x}^{(t)}\}_{t \geq 0}$ be a trajectory under PURE INFIMMDYN based on $A = A^\top$, and suppose*

$$\lim_{t \rightarrow \infty} \pi(\mathbf{x}^{(t)}) = \pi(\mathbf{e}^i).$$

If $\mathcal{S}_{\text{Pure}}(\mathbf{x}^{(T)}) = \mathbf{e}^i$ at some $T \geq 0$, then $\mathbf{x}^{(T+1)}$ is a fixed point.

If $\mathcal{S}_{\text{Pure}}(\mathbf{x}^{(T)}) = \overline{\mathbf{e}}^i_{\mathbf{x}^{(T)}}$ at some $T \geq 0$, then $\mathbf{x}^{(T+1)} = \overline{\mathbf{e}}^i_{\mathbf{x}^{(T)}}$.

Proof. See appendix. □

Theorem 6. *Consider a trajectory under PURE INFIMMDYN with $A = A^\top$ starting at $\mathbf{x}^{(0)}$ such that $\omega(\mathbf{x}^{(0)})$ is finite. Then the trajectory converges, i.e., $\omega(\mathbf{x}^{(0)})$ contains exactly one Nash strategy.*

Proof. See appendix. □

Corollary 1. *If the partnership game with payoff matrix $A = A^\top$ admits only finitely many symmetric Nash equilibrium strategies, then every trajectory under PURE INFIMMDYN converges.*

Proof. The assertion follows directly from Theorem 6 and (20). □

The example of Phase Portrait 47 in [4] shows that the assumption of Theorem 6 is strictly weaker than that of Corollary 1.

4.4. Support separation in finite time

In this section, we establish a property of PURE INFIMMDYN that seems to be a distinctive feature among all evolutionary game dynamics, namely support separation in finite time (we are not aware of any similar result for any other evolutionary game dynamics). Recall that the $\sigma(\mathbf{q})$ of a Nash strategy \mathbf{q} , includes all pure strategies really used by \mathbf{q} . Suppose that this set $\sigma(\mathbf{q})$ is known or, to be more precise, pick any set τ satisfying

$$\sigma(\mathbf{q}) \subseteq \tau \subseteq \tau_0(\mathbf{q}), \quad (21)$$

so that τ comprises only pure strategies balancing at \mathbf{q} . Because of the Nash property $\sigma(\mathbf{q}) \subseteq \tau_0(\mathbf{q})$, such a set τ always exists, but generally the computational and combinatorial complexity of finding Nash strategies mainly consists in determining such a τ ; recall there are *a priori* $2^n - 1$ possibilities. Given τ , it is a mere matter of simple linear algebra to determine the Nash strategy \mathbf{q} .

Indeed, denote by A_τ the submatrix of A obtained by extracting the rows and columns corresponding to τ , and similarly with \mathbf{x}_τ and \mathbf{u}_τ , a vector with unity entries only. Further, put $\lambda = 1/(\mathbf{u}_\tau^\top A_\tau^{-1} \mathbf{u}_\tau)$ and solve $A_\tau \mathbf{x}_\tau = \lambda \mathbf{u}_\tau$. Then the positive coordinates of \mathbf{q} are those of \mathbf{x}_τ while necessarily all $q_i = 0$ for $i \notin \tau$, due to (21). Note that *regular* Nash strategies \mathbf{q} are characterized by the fact that any balancing pure strategy is actually used with positive probability, i.e., $\sigma(\mathbf{q}) = \tau_0(\mathbf{q})$, and that regularity of Nash strategies, as well as nonsingularity of the matrices A_τ , is generic, cf., e.g. [6].

Geometrically speaking, the property of support separation in finite time allows us to find the face of the simplex where a Nash equilibrium is lying, after finitely many steps of the dynamics. So any iterative equilibrium selection process starting with a large support $\sigma(\mathbf{x}^{(0)})$, e.g., in the interior of Δ , may be considered finite if the (extended) support of a Nash strategy is separated in finite time T , i.e., if $\sigma(\mathbf{x}^{(t)}) \subseteq \tau_0(\lim_{t \rightarrow \infty} \mathbf{x}^{(t)})$ for all $t \geq T$. Of course, we already employ Theorem 6 to ensure convergence. The point is that we need not wait for the coordinates $x_i(t)$ to stabilize, but rather know after T steps the support and therefore can calculate $\mathbf{q} = \lim_{t \rightarrow \infty} \mathbf{x}^{(t)}$ by choosing $\tau = \sigma(\mathbf{x}^{(T)})$ and simply solving a linear equation system as detailed above.

Theorem 7. *Let $\{\mathbf{x}^{(t)}\}_{t \geq 0}$ be a convergent (cf. Theorem 6) trajectory of PURE INFIMMDYN and suppose that $A = A^\top$. There exists $T \geq 0$ such that*

$$\sigma(\mathbf{q}) \subseteq \sigma(\mathbf{x}^{(t)}) \subseteq \tau_0(\mathbf{q}) \quad \text{for all } t \geq T,$$

where $\mathbf{q} = \lim_{t \rightarrow \infty} \mathbf{x}^{(t)}$.

Proof. See appendix. □

5. Computational evidence and illustration

5.1. Effectiveness for complex partnership games

In order to provide an example of the effectiveness of PURE INFIMMDYN, we compare it against RD and FP on two partnership games, identified by the following payoff matrices:

$$A_1 = \begin{pmatrix} 0 & 2 & 1 \\ 2 & 0 & 1 \\ 1 & 1 & 1 \end{pmatrix} \quad \text{and} \quad A_2 = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 1 & 2 \\ 0 & 2 & 1 \end{pmatrix}.$$

Both games, A_1 and A_2 , have infinitely many Nash strategies and exhibit a somehow more complex dynamic behavior than in a typical partnership game with a finite set of Nash equilibria. Specifically, in game A_1 the set of Nash equilibria is given by $\{\alpha \mathbf{e}^3 + (1 - \alpha)(\mathbf{e}^1 + \mathbf{e}^2)/2 : \alpha \in [0, 1]\}$, whereas the set of Nash strategies in game A_2 is given by $\{\alpha \mathbf{e}^1 + (1 - \alpha)\mathbf{e}^2 : \alpha \in [0, 1]\} \cup \{(\mathbf{e}^2 + \mathbf{e}^3)/2\}$.

For each game, we selected a mixed strategy $\mathbf{x}^{(0)} \in \Delta$ as starting point for the three dynamics and we let them run to convergence. For RD we chose the standard setting

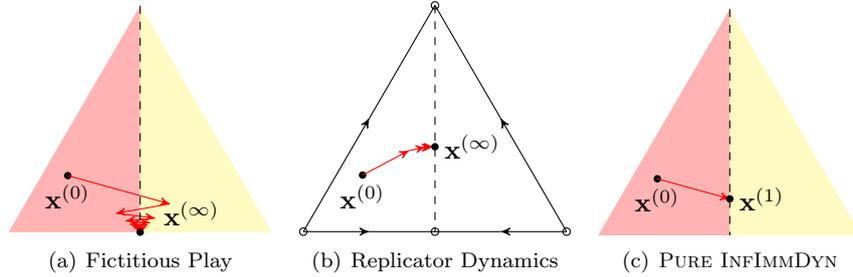


Figure 2: Example of trajectories of FP, RD and PURE INFIMMDYN on the partnership game with payoff matrix A_1 . See text for color codes in the simplex partitions.

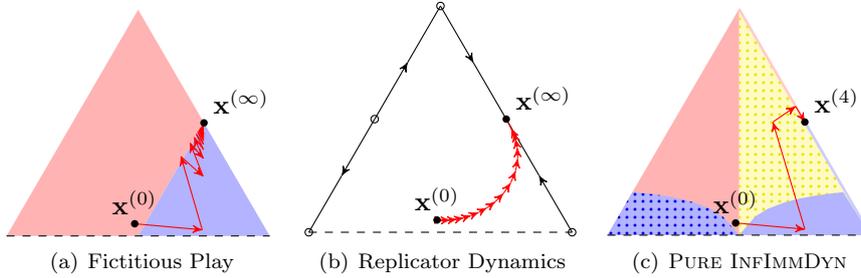


Figure 3: Example of trajectories of FP, RD and PURE INFIMMDYN on the partnership game with payoff matrix A_2 . See text for color codes in the simplex partitions.

$C = 0$. Figures 2 and 3 show the trajectories that have been generated for the partnership games A_1 and A_2 , respectively. In the illustrations for PURE INFIMMDYN and FP, the simplices have been partitioned into colored regions according to the values of $\mathcal{S}_{\text{Pure}}(\mathbf{x})$ and $\beta(\mathbf{x})$, respectively. Yellow, red, and blue regions correspond to \mathbf{e}^1 , \mathbf{e}^2 , and \mathbf{e}^3 , while dotted yellow and blue ones correspond to co-strategies $\bar{\mathbf{e}}^1_{\mathbf{x}}$ and $\bar{\mathbf{e}}^3_{\mathbf{x}}$. Further, since RD is face-invariant, we pictured on the simplex' boundary the direction of evolution along the face.

Although there are infinitely many Nash equilibria, it can be proved that for both games the ω -limits of every trajectory are finite sets and therefore, by Theorem 6, PURE INFIMMDYN converges. In particular, PURE INFIMMDYN is always able to converge in finitely many steps, namely in *one* (!) step for game A_1 and in at most four steps for game A_2 . We chose this worst case among all possible trajectories for illustration. Note that in contrast to our dynamics, RD and FP need infinitely many steps to converge starting from any non-fixed point. It should be remarked, however, that there are special cases (where a pure strategy is a Nash equilibrium) for which FP may converge in finite time, too. As RD is an interior-point method, this never happens with RD, unless one by chance already starts in a fixed point.

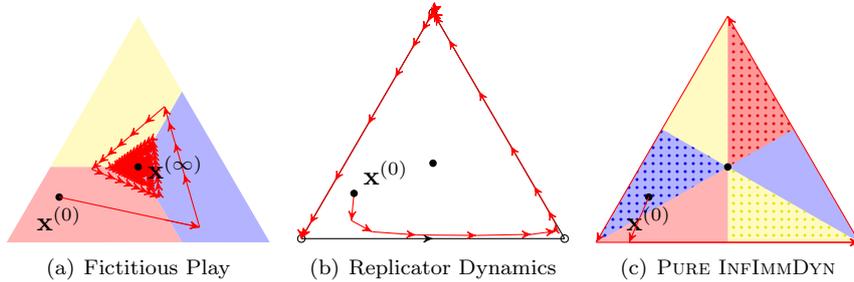


Figure 4: Example of trajectories of FP, RD and PURE INFIMMODYN on the standard RPS game $A_3(1)$, with color code for simplex partitions according to the values of $\beta(\mathbf{x})$ and $\mathcal{S}_{\text{PURE}}(\mathbf{x})$, respectively: yellow for \mathbf{e}^1 ; red for \mathbf{e}^2 ; blue for \mathbf{e}^3 ; dotted yellow for $\mathbf{e}^1_{\mathbf{x}}$; dotted red for $\mathbf{e}^2_{\mathbf{x}}$; and dotted blue for $\mathbf{e}^3_{\mathbf{x}}$.

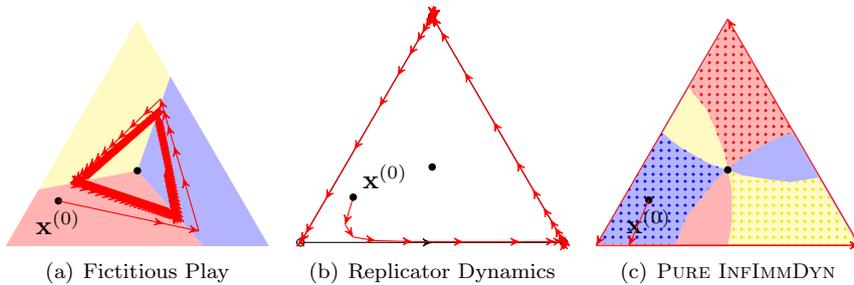


Figure 5: Example of trajectories of FP, RD and PURE INFIMMODYN on a bad RPS game $A_3(0.5)$. Color code as in Figure 4.

5.2. Acceleration for non-partnership games

We also report results on some non-partnership games. The first one is the well-known Rock-Paper-Scissor (RPS) game, which is identified by the following class of payoff matrices:

$$A_3(w) = \begin{pmatrix} 1 & 0 & w+1 \\ w+1 & 1 & 0 \\ 0 & w+1 & 1 \end{pmatrix}$$

where $w \geq 0$. Following Sandholm's classification and terminology [31], we obtain the *standard RPS* game for $w = 1$; a *good RPS* game for $w > 1$; and a *bad RPS* game if $w < 1$. Note that $w = 1$ is equivalent to a zero-sum game with ± 1 on the off-diagonal entries for FP and PURE INFIMMODYN. By contrast, the discrete-time RD, unlike its continuous-time version, is not invariant by adding a constant value to every entry.

In all three cases the only Nash equilibrium is located in the barycenter of the simplex. This Nash equilibrium is an ESS if $w > 1$ but not asymptotically stable (neither under the discrete-time RD nor under PURE INFIMMODYN), which shows that Theorem 4 cannot hold for every base game. By contrast, in good RPS games ($w > 1$) the continuous-time RD converges to equilibrium. As it turns out from Figures 4, 5 and 6, in the case of standard and good RPS games, FP is the only dynamics which converges to

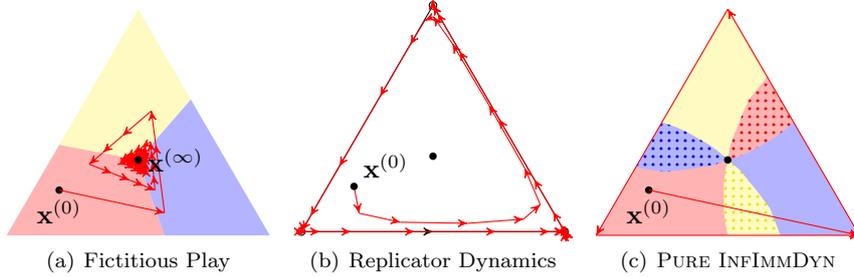


Figure 6: Example of trajectories of FP, RD and PURE INFIMMODYN on a good RPS game $A_3(2)$. Color code as in Figure 4.

Nash equilibrium. In all cases, PURE INFIMMODYN cycles over the three pure strategies, which numerically makes us aware immediately of the cycling behavior. By contrast, RD asymptotically approaches the simplex' boundary, slowing down close to the vertices, which may in unfavorable numerical circumstances, wrongly suggest convergence to a pure strategy.

To enforce convergence of PURE INFIMMODYN and RD to Nash equilibrium in the case of good RPS games, we have to slow them down. We can slow down RD by increasing the constant C in (5), while for PURE INFIMMODYN, we may scale down δ in (7) or define a different selection function. However, by slowing down PURE INFIMMODYN, we lose the feature of being a non-interior-point method.

There are non-partnership games where PURE INFIMMODYN converges quickly. This is illustrated by the following games:

$$A_4 = \begin{pmatrix} 1 & 2 & 2 \\ 0 & 1 & 4 \\ 2 & 2 & 1 \end{pmatrix} \quad \text{and} \quad A_5 = \begin{pmatrix} 1 & 0 & 2 \\ 3 & 1 & 0 \\ 3 & 2 & 1 \end{pmatrix} .$$

Each game has a unique Nash equilibrium strategy, namely the barycenter $(\mathbf{e}^1 + \mathbf{e}^2 + \mathbf{e}^3)/3$ for A_4 and the boundary point $(\mathbf{e}^1 + 2\mathbf{e}^3)/3$ for A_5 .

In Figures 7 and 8 we can see sample trajectories of the three dynamics for games A_4 and A_5 , respectively. Like in the case of partnership games, PURE INFIMMODYN converges to Nash equilibrium within very few steps, as opposed to RD and FP, which require again infinitely many steps.

6. Conclusions and outlook

This paper proposes a new class of dynamics which directly builds upon a central paradigm of evolutionary game theory, namely the invasion barrier. Recall that \mathbf{p} is an evolutionarily stable strategy (ESS) if and only if \mathbf{p} is protected by a positive invasion barrier against any possible invading strategy $\mathbf{x} \neq \mathbf{p}$. A similar paradigm holds for neutral stability.

Based upon these concepts, we propose the new Infection and Immunization Dynamics (INFIMMODYN), modelling a plausible adaptation process in a large population

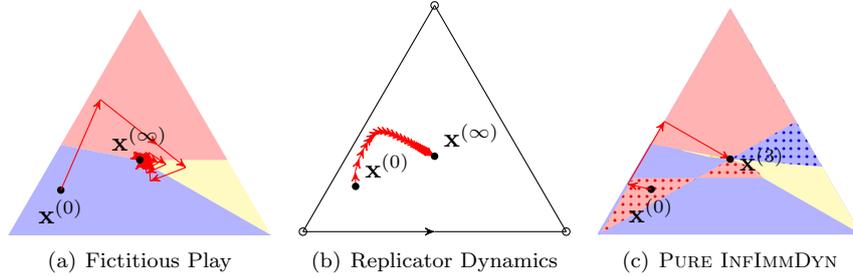


Figure 7: Example of trajectories of FP, RD and PURE INFIMMODYN on game A_4 . Color code as in Figure 4.

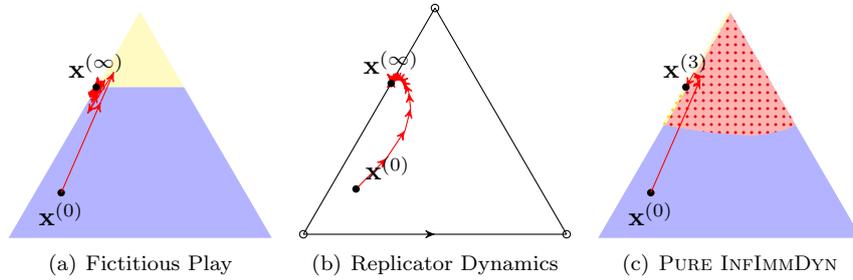


Figure 8: Example of trajectories of FP, RD and PURE INFIMMODYN on game A_5 . Color code as in Figure 4.

without imposing too sophisticated arguments: for instance, every step requires just twice as many simple payoff comparisons as there are pure strategies. Based upon this process, we derived a new characterization of symmetric Nash equilibrium strategies as fixed points, thereby opening the road to several refinement of this concepts gained by qualifying these fixed points, e.g. requiring asymptotic or Lyapunov stability. Further, we adapted the notion of revision protocols to arrive at a possible microfoundation of adaptive dynamics also in discrete time, and applied this to the most immediate variant of INFIMMODYN based on pure strategy selection. We also discussed discretization policies of continuous-time adaptive dynamics in this context.

Unlike many popular evolutionary game dynamics, INFIMMODYN lacks continuity, in a way similar to the best reply dynamics. Likewise it exhibits a favorable asymptotic behavior, e.g., convergence for partnership games. In addition, we here have equivalence of evolutionary with asymptotic dynamical stability. So INFIMMODYN combines positive features of different approaches, with the additional advantage that in many cases, the asymptotics is improved compared to other familiar procedures like Fictitious Play (FP) and Replicator Dynamics (RD).

We also established even support separation of INFIMMODYN in finite time, which can never be achieved by any interior-point method like RD (and in most cases, neither by FP). In fact, this property – which becomes a central issue in view of the rapidly increasing importance of game theoretic methods in high-performance algorithms for

applications like artificial intelligence and/or combinatorial optimization – has not yet been established for any other evolutionary game dynamics.

Finally, we illustrated by several examples of RSP and other non-partnership games, that INFIMMDYN may have computational advantages also in case of asymmetric payoff matrices, even if no convergence takes place. While further theoretical investigations, e.g., of general zero-sum games or alternatives to the pure selection function, are beyond the scope of this paper, they constitute an interesting research topic for the future.

Acknowledgements. By their thoughtful comments and diligent suggestions, an anonymous Associate Editor and two anonymous referees contributed significantly to improvements of earlier versions of this paper. The authors also are indebted to Reinhard Ullrich for his valuable remarks. Finally, we acknowledge financial support from the FET programme within the EU FP7, under the SIMBAD project (contract 213250).

7. Appendix: Proofs of some results

Here we present the more involved proofs of some results.

Proof of Lemma 1:

Proof. Note that as a consequence of Theorem 2, for all $\mathbf{y} \in \omega(\mathbf{x}^{(0)})$, we have that $\pi(\mathbf{y})$ is constant. In fact, if \mathbf{q} and \mathbf{p} are two limit points of the convergent subsequences $\xi(s)$ and $\xi(r)$, respectively, which are chosen in a way as to have $r_j < s_j < r_{j+1}$ for all $j \geq 0$, we obtain

$$\pi(\mathbf{x}^{(r_j)}) \leq \pi(\mathbf{x}^{(s_j)}) \leq \pi(\mathbf{x}^{(r_{j+1})}),$$

and by taking the limit we have

$$\pi(\mathbf{p}) \leq \pi(\mathbf{q}) \leq \pi(\mathbf{p}),$$

which is true only if $\pi(\mathbf{q}) = \pi(\mathbf{p})$. Hence, the image through π of all limit points gives a constant γ . Therefore, the sequence $\{\pi(\mathbf{x}^{(t)})\}_{t \geq 0}$ converges to γ , and, by Theorem 2, it is non-decreasing and upper bounded by γ . From this, it also follows that if for some T , $\pi(\mathbf{x}^{(T)}) = \gamma$ then for all $t \geq T$, $\pi(\mathbf{x}^{(t)}) = \gamma$, which, by Theorem 2, implies that $\mathbf{x}^{(T)}$ is a fixed point of INFIMMDYN. \square

Proof of Lemma 2:

Proof. By Lemma 1, we know $\pi(\mathbf{x}^{(t+1)}) - \pi(\mathbf{x}^{(t)}) \rightarrow 0$ as $j \rightarrow \infty$. Due to symmetry of the payoff matrix we have $\pi(\mathbf{x}|\mathcal{S}(\mathbf{x})) = \pi(\mathcal{S}(\mathbf{x})|\mathbf{x})$, so that the above entails via (15) the relation

$$\lim_{t \rightarrow \infty} \delta^{(t)} \left[h_{\mathbf{x}}(\mathcal{S}(\mathbf{x}), \delta^{(t)}) + \pi(\mathcal{S}(\mathbf{x}^{(t)}) - \mathbf{x}^{(t)}|\mathbf{x}^{(t)}) \right] = 0. \quad (22)$$

To establish (16), consider a partition of the trajectory $\{\mathbf{x}^{(t)}\}$ into two (without loss of generality infinite) subsequences $\xi(s)$ and $\xi(r)$ where for all $j \geq 0$ we have that $\delta^{(s_j)} = 1$ and $\delta^{(r_j)} < 1$ (if either of $\xi(s)$ or $\xi(r)$ is finite, we need only consider the other one for establishing the assertion). By (22) it follows that

$$\begin{aligned} 0 &= \lim_{j \rightarrow \infty} \delta^{(s_j)} \left[h_{\mathbf{x}^{(s_j)}}(\mathcal{S}(\mathbf{x}^{(s_j)}), \delta^{(s_j)}) + \pi(\mathcal{S}(\mathbf{x}^{(s_j)}) - \mathbf{x}^{(s_j)}|\mathbf{x}^{(s_j)}) \right] \\ &= \lim_{j \rightarrow \infty} \left[h_{\mathbf{x}^{(s_j)}}(\mathcal{S}(\mathbf{x}^{(s_j)}), 1) + \pi(\mathcal{S}(\mathbf{x}^{(s_j)}) - \mathbf{x}^{(s_j)}|\mathbf{x}^{(s_j)}) \right]. \end{aligned} \quad (23)$$

Since the last two terms above are always nonnegative, it holds in particular that

$$\pi(\mathcal{S}(\mathbf{x}^{(s_j)}) - \mathbf{x}^{(s_j)} | \mathbf{x}^{(s_j)}) \rightarrow 0 \quad \text{as } j \rightarrow \infty.$$

Considering the other subsequence, it follows that $h_{\mathbf{x}^{(r_j)}}(\mathcal{S}(\mathbf{x}^{(r_j)}), \delta^{(r_j)}) = 0$, by definition of $\delta^{(r_j)}$, and therefore, using (6) for $\delta^{(r_j)} < 1$ and (22),

$$\begin{aligned} -\frac{\pi(\mathcal{S}(\mathbf{x}^{(r_j)}) - \mathbf{x}^{(r_j)} | \mathbf{x}^{(r_j)})^2}{\pi(\mathcal{S}(\mathbf{x}^{(r_j)}) - \mathbf{x}^{(r_j)})} &= \delta^{(r_j)} \pi(\mathcal{S}(\mathbf{x}^{(r_j)}) - \mathbf{x}^{(r_j)} | \mathbf{x}^{(r_j)}) = \\ &= \delta^{(r_j)} \left[h_{\mathbf{x}^{(r_j)}}(\mathcal{S}(\mathbf{x}^{(r_j)}), \delta^{(r_j)}) + \pi(\mathcal{S}(\mathbf{x}^{(r_j)}) - \mathbf{x}^{(s_j)} | \mathbf{x}^{(r_j)}) \right] \rightarrow 0 \quad \text{as } j \rightarrow \infty, \end{aligned}$$

which in turn implies that $\pi(\mathcal{S}(\mathbf{x}^{(r_j)}) - \mathbf{x}^{(r_j)} | \mathbf{x}^{(r_j)}) \rightarrow 0$. Hence (16) follows.

To establish (17), note that $\delta^{(t)}$ is bounded, so (16) implies

$$\delta^{(t)} \pi(\mathcal{S}(\mathbf{x}^{(t)}) - \mathbf{x}^{(t)} | \mathbf{x}^{(t)}) \rightarrow 0 \quad \text{as } t \rightarrow \infty.$$

Hence (22) yields $\delta^{(t)} h_{\mathbf{x}^{(t)}}(\mathcal{S}(\mathbf{x}^{(t)}), \delta^{(t)}) \rightarrow 0$. But

$$h_{\mathbf{x}^{(t)}}(\mathcal{S}(\mathbf{x}^{(t)}), \delta^{(t)}) = \delta^{(t)} \pi(\mathcal{S}(\mathbf{x}^{(t)}) - \mathbf{x}^{(t)}) + \pi(\mathcal{S}(\mathbf{x}^{(t)}) - \mathbf{x}^{(t)} | \mathbf{x}^{(t)}),$$

and the last term goes to zero using (16) once more, so the limit for (17) follows. \square

Proof of Lemma 3:

Proof. Put $\mathbf{y} = \mathcal{S}_{\text{pure}}(\mathbf{x}^{(T)})$ in both cases, and abbreviate by $\delta = \delta_{\mathbf{y}}(\mathbf{x}^{(T)})$. First we deal with the case $\mathbf{y} = \mathbf{e}^i$. If $\mathbf{x}^{(T)} = \mathbf{e}^i$ then trivially, $\mathbf{x}^{(T+1)}$ is a fixed point. Hence, we can assume in the sequel that $\mathbf{x}^{(T)} \neq \mathbf{e}^i$, which indeed implies that $\mathbf{e}^i \in \Upsilon(\mathbf{x}^{(T)})$ by definition of $\mathcal{S}_{\text{pure}}$. Now, if $\delta = 1$, then $\pi(\mathbf{x}^{(T+1)}) = \pi(\mathbf{e}^i) = \gamma$, and, by Lemma 1, it follows that $\mathbf{x}^{(T+1)} = \mathbf{e}^i$ is a fixed point. On the other hand, if $\delta < 1$, then by definition of δ , we have $h_{\mathbf{e}^i}(\mathbf{x}^{(T)}, \delta) = 0$. Note that, according to (6), this happens only if

$$0 < \delta = -\frac{\pi(\mathbf{e}^i - \mathbf{x}^{(T)} | \mathbf{x}^{(T)})}{\pi(\mathbf{e}^i - \mathbf{x}^{(T)})} < 1,$$

where $\pi(\mathbf{e}^i - \mathbf{x}^{(T)}) < 0$. By exploiting this inequality, we obtain

$$\pi(\mathbf{e}^i - \mathbf{x}^{(T)}) < -\pi(\mathbf{e}^i - \mathbf{x}^{(T)} | \mathbf{x}^{(T)}) < 0,$$

which implies that

$$\pi(\mathbf{e}^i - \mathbf{x}^{(T)} | \mathbf{e}^i) < 0.$$

But then,

$$\pi(\mathbf{x}^{(T+1)}) - \gamma = \delta \pi(\mathbf{e}^i - \mathbf{x}^{(T)} | \mathbf{x}^{(T)}) + \pi(\mathbf{x}^{(T)}) - \pi(\mathbf{e}^i) = -\frac{\pi(\mathbf{e}^i - \mathbf{x}^{(T)} | \mathbf{e}^i)^2}{\pi(\mathbf{e}^i - \mathbf{x}^{(T)})} > 0,$$

which contradicts Lemma 1. Hence, the only admissible case is when $\delta = 1$, which proves that $\mathbf{x}^{(T)}$ is a fixed point.

Next we deal with the co-strategy case $\mathbf{y} = \overline{\mathbf{e}}^i_{\mathbf{x}^{(T)}}$. We show that $\pi(\mathbf{y} - \mathbf{x}^{(T)}) \geq 0$,

which entails by (6) that $\delta = 1$. Now since the co-strategy of \mathbf{e}^i is selected, we know $i \in \tau_-(\mathbf{x}^{(T)}) \cap \sigma(\mathbf{x}^{(T)})$, hence $x_i^{(T)} > 0$ and $\pi(\mathbf{x}^{(T)} - \mathbf{e}^i | \mathbf{x}^{(T)}) > 0$, so that

$$\pi(\mathbf{y} - \mathbf{x}^{(T)}) = \left(\frac{x_i^{(T)}}{x_i^{(T)} - 1} \right)^2 \pi(\mathbf{e}^i - \mathbf{x}^{(T)}) = \quad (24)$$

$$= \left(\frac{x_i^{(T)}}{x_i^{(T)} - 1} \right)^2 \left[\pi(\mathbf{e}^i) - \pi(\mathbf{x}^{(T)}) + 2\pi(\mathbf{x}^{(T)} - \mathbf{e}^i | \mathbf{x}^{(T)}) \right] > 0, \quad (25)$$

where the latter inequality follows from the above and from Lemma 1. So indeed $\delta = 1$, and $\mathbf{x}^{(T+1)} = \mathbf{y}$ follows. \square

Proof of Theorem 6:

Proof. Suppose $\omega(\mathbf{x}^{(0)})$ contains more than one point. Since there are finitely many pairs of limit points by hypothesis, there exists a pair $\{\mathbf{p}, \mathbf{q}\} \subseteq \omega(\mathbf{x}^{(0)})$ and a subsequence s_j such that $\mathbf{x}^{(s_j)} \rightarrow \mathbf{p}$ and its successors $\mathbf{x}^{(s_j+1)} \rightarrow \mathbf{q} \neq \mathbf{p}$ as $j \rightarrow \infty$. Abbreviate by $\mathbf{y} = \mathcal{S}_{\text{pure}}(\mathbf{x}^{(s_j)})$ and by $\delta_j = \delta_{\mathbf{y}}(\mathbf{x}^{(s_j)})$. By thinning this subsequence (s_j) appropriately if necessary, we can assume that either

$$\mathbf{y} = \mathbf{e}^i \quad \text{for all } j \geq 0, \quad (26)$$

or

$$\mathbf{y} = \overline{\mathbf{e}^i}_{\mathbf{x}^{(s_j)}} \quad \text{for all } j \geq 0. \quad (27)$$

Now we use Lemma 2 again: in case (26), we use (17) to arrive at

$$\pi(\mathbf{e}^i - \mathbf{p}) \lim_{j \rightarrow \infty} \delta_j^2 = \lim_{j \rightarrow \infty} \delta_j^2 \pi(\mathbf{e}^i - \mathbf{x}^{(s_j)}) = 0,$$

and conclude that $\pi(\mathbf{e}^i - \mathbf{p}) = 0$ or $\delta_j \rightarrow 0$. In the first case, we use

$$\pi(\mathbf{e}^i) - \pi(\mathbf{p}) = \pi(\mathbf{e}^i - \mathbf{p}) - 2\pi(\mathbf{e}^i - \mathbf{p} | \mathbf{p}) = -2 \lim_{j \rightarrow \infty} \pi(\mathbf{e}^i - \mathbf{x}^{(s_j)} | \mathbf{x}^{(s_j)}) = 0,$$

so that $\pi(\mathbf{x}^{(s_j)}) \rightarrow \pi(\mathbf{e}^i)$ holds, which entails by Lemma 3 that even $\mathbf{x}^{(s_1)}$ is a fixed point by (26), which contradicts the assumption. Therefore we know that the stepsize

$$\|\mathbf{x}^{(s_j+1)} - \mathbf{x}^{(s_j)}\| \leq 2\delta_j \rightarrow 0 \quad \text{as } j \rightarrow \infty,$$

which means $\mathbf{p} = \mathbf{q}$ which is also absurd. Hence (27) must obtain: along s_j , always the co-strategy is infective. The stepsize now is equal to

$$\|\mathbf{x}^{(s_j+1)} - \mathbf{x}^{(s_j)}\| = \delta_j \|\mathbf{y} - \mathbf{x}^{(s_j)}\| = \delta_j \frac{x_i^{(s_j)}}{1 - x_i^{(s_j)}} \|\mathbf{e}^i - \mathbf{x}^{(s_j)}\|. \quad (28)$$

Now we take limits as $j \rightarrow \infty$. To avoid again the absurd $\mathbf{p} = \mathbf{q}$, we deduce that, unless $\mathbf{p} = \mathbf{e}^i$, now $\delta_j \geq \bar{\delta} > 0$ holds for all $j \geq 0$, along with $p_i > 0$, so that (17) entails

$$\left(\frac{p_i}{p_i - 1} \right)^2 \pi(\mathbf{e}^i - \mathbf{p}) = \lim_{j \rightarrow \infty} \pi(\overline{\mathbf{e}^i}_{\mathbf{x}^{(s_j)}} - \mathbf{x}^{(s_j)}) = 0$$

and therefore $\pi(\mathbf{e}^i) - \pi(\mathbf{p}) = \pi(\mathbf{e}^i - \mathbf{p}) - 2\pi(\mathbf{e}^i - \mathbf{p} | \mathbf{p}) = 0 - 0 = 0$ because $i \in \sigma(\mathbf{p}) \subseteq \tau_0(\mathbf{p})$. Note that $\pi(\mathbf{p}) = \pi(\mathbf{e}^i)$ and $p_i > 0$ trivially also holds in case $\mathbf{p} = \mathbf{e}^i$, where taking limits in (28) could be problematic. Next we invoke Lemma 3 again, wherefrom we know $\mathbf{x}^{(s_j+1)} = \mathbf{y} = \overline{\mathbf{e}^i}_{\mathbf{x}^{(s_j)}}$ for all $j \geq 0$. In particular, $x_i^{(s_j+1)} = 0$ holds. Now the only way to increase $x_i^{(s_j+1)}$ from zero along the original sequence is the selection $\mathcal{S}(\mathbf{x}^{(t)}) = \mathbf{e}^i$ for some $t \geq s_j + 1$. But then, as $\pi(\mathbf{x}^{(t)}) \rightarrow \gamma = \pi(\mathbf{p}) = \pi(\mathbf{e}^i)$, we know from Lemma 3 that $\mathbf{x}^{(t+1)}$ must be fixed, again a contradiction. This means that $x_i^{(t)} = 0$ for all $t \geq s_j + 1$, in particular for all $t = s_k > s_j + 1$ for $k > j$, which would entail $p_i = 0$, which we ruled out above. So we have shown that the assumption of more than one element in $\omega(\mathbf{x}^{(0)})$ is absurd, and the theorem is proved. \square

Proof of Theorem 7:

Proof. By continuity and definition (1), we can specify a $T \geq 0$ such that $\tau_-(\mathbf{q}) \subseteq \tau_-(\mathbf{x}^{(t)})$ and $\sigma(\mathbf{q}) \subseteq \sigma(\mathbf{x}^{(t)})$ for all $t \geq T$. Further, \mathbf{q} is a Nash strategy by Theorem 3, so that $\sigma(\mathbf{q}) \subseteq \tau_0(\mathbf{q}) = S \setminus \tau_-(\mathbf{q})$. Hence we know

$$\tau_+(\mathbf{x}^{(t)}) \subseteq S \setminus \tau_-(\mathbf{x}^{(t)}) \subseteq S \setminus \tau_-(\mathbf{q}) = \tau_0(\mathbf{q}) \quad \text{for all } t \geq T. \quad (29)$$

This inclusion implies that once $\sigma(\mathbf{x}^{(t_0)}) \subseteq \tau_0(\mathbf{q})$ for some $t_0 \geq T$, then this inclusion of supports is guaranteed for all later $t \geq t_0$. Indeed, by definition of PURE INFIMMDYN we either have $\sigma(\mathbf{x}^{(t_0+1)}) = \sigma(\mathbf{x}^{(t_0)}) \cup \{i\}$ where $i = \mathcal{M}(\mathbf{x}^{(t_0)}) \in \tau_+(\mathbf{x}^{(t_0)}) \subseteq \tau_0(\mathbf{q})$ by (29) and thus also $\sigma(\mathbf{x}^{(t_0+1)}) \subseteq \tau_0(\mathbf{q})$, or the update rule uses the co-strategy $\overline{\mathbf{e}^j}_{\mathbf{x}^{(t_0)}}$ for some $j \in \sigma(\mathbf{x}^{(t_0)}) \cap \tau_-(\mathbf{x}^{(t_0)})$. In this case, generally, the support either shrinks or stays the same in the following iteration, i.e., $\sigma(\mathbf{x}^{(t_0+1)}) \subseteq \sigma(\mathbf{x}^{(t_0)})$, so that again $\sigma(\mathbf{x}^{(t_0+1)}) \subseteq \tau_0(\mathbf{q})$ holds. Next we argue by contradiction. In light of the above reasoning, the statement can be wrong only if for all $t \geq T$ there is an index $j \in \sigma(\mathbf{x}^{(t)}) \setminus \tau_0(\mathbf{q})$. Now, by increasing T if necessary, we can quantify the continuity arguments in that there is a real number $\gamma > 0$ such that for all $t \geq T$ we have

$$\begin{aligned} |\pi(\mathbf{e}^i - \mathbf{x}^{(t)} | \mathbf{x}^{(t)})| &< \gamma & \text{if } i \in \tau_0(\mathbf{q}), \text{ and} \\ \pi(\mathbf{x}^{(t)} - \mathbf{e}^j | \mathbf{x}^{(t)}) &> \gamma & \text{if } j \in \sigma(\mathbf{x}^{(t)}) \setminus \tau_0(\mathbf{q}) \subseteq \tau_-(\mathbf{q}). \end{aligned} \quad (30)$$

Hence, by the preceding argument, $\sigma(\mathbf{x}^{(t)}) \setminus \tau_0(\mathbf{q})$ is never empty and equals $\sigma(\mathbf{x}^{(t)}) \cap \tau_-(\mathbf{x}^{(t)})$. This implies, by construction of $\mathcal{S}_{\text{Pure}}$ and relation (30), that for all $t \geq T$,

$$\mathcal{S}_{\text{Pure}}(\mathbf{x}^{(t)}) = \overline{\mathbf{e}^j}_{\mathbf{x}^{(t)}} \quad \text{for some } j \in \sigma(\mathbf{x}^{(t)}) \setminus \tau_0(\mathbf{q}).$$

We already saw that this yields shrinking or constant supports, and we will now show that indeed in every iteration one index is removed from the support, which gives a contradiction after at most $|\tau_-(\mathbf{q})|$ steps. To this end, let

$$T' = \inf \left\{ t \geq T : \delta_{\overline{\mathbf{e}^j}_{\mathbf{x}^{(t)}}}(\mathbf{x}^{(t)}) = 1 \quad \text{for all } j \in \sigma(\mathbf{x}^{(t)}) \setminus \tau_0(\mathbf{q}) \right\}.$$

Note that $T' < \infty$ because for all $j \in \sigma(\mathbf{x}^{(t)}) \setminus \tau_0(\mathbf{q})$, either $\pi(\mathbf{e}^j - \mathbf{x}^{(t)}) \geq 0$ and then $\delta_{\overline{\mathbf{e}^j}_{\mathbf{x}^{(t)}}}(\mathbf{x}^{(t)}) = 1$, see (6); or $\pi(\mathbf{e}^j - \mathbf{x}^{(t)}) < 0$, in which case we derive from (30)

$$\frac{x_j^{(t)} - 1}{x_j^{(t)}} \frac{\pi(\mathbf{x}^{(t)} - \mathbf{e}^j | \mathbf{x}^{(t)})}{\pi(\mathbf{e}^j - \mathbf{x}^{(t)})} > \gamma \frac{x_j^{(t)} - 1}{x_j^{(t)} \pi(\mathbf{e}^j - \mathbf{x}^{(t)})} \rightarrow +\infty \quad \text{as } t \rightarrow \infty$$

because $x_j^{(t)} \rightarrow q_j = 0$ and $\pi(\mathbf{e}^j - \mathbf{x}^{(t)})$ is bounded, so that, again by (6),

$$\delta_{\mathbf{e}^j_{\mathbf{x}^{(t)}}}(\mathbf{x}^{(t)}) = \min \left\{ 1, \frac{x_j^{(t)} - 1}{x_j^{(t)}} \frac{\pi(\mathbf{x}^{(t)} - \mathbf{e}^j | \mathbf{x}^{(t)})}{\pi(\mathbf{e}^j - \mathbf{x}^{(t)})} \right\}$$

is eventually equal to one in any case. Hence after finitely many steps, for all $j \in \sigma(\mathbf{x}^{(t)}) \setminus \tau_0(\mathbf{q})$ we have $\delta_{\mathbf{e}^j_{\mathbf{x}^{(t)}}}(\mathbf{x}^{(t)}) = 1$ and therefore $T' < \infty$. This implies that after at most $|\tau_-(\mathbf{q})|$ steps and for all following ones, the support of $\mathbf{x}^{(t)}$ will be a subset of $\tau_0(\mathbf{q})$. \square

References

- [1] Balkenborg, D., Schlag, K., 2001. Evolutionarily stable sets. *Int. J. Game Theory* 29 (4), 571–595.
- [2] Bjoernerstedt, J., Weibull, J., 1996. The rational foundations of economic behaviour. MacMillan; New York, Ch. Nash equilibrium and evolution by imitation, pp. 151–171.
- [3] Bomze, I., Pelillo, M., Stix, V., 2000. Approximating the maximum weight clique using replicator dynamics. *IEEE Trans. Neural Networks* 11, 1228–1241.
- [4] Bomze, I. M., 1995. Lotka-Volterra equation and replicator dynamics: new issues in classification. *Biological Cybernetics* 72 (5), 447–453.
- [5] Bomze, I. M., 1998. *Game Theory, Experience, Rationality*. Kluwer, Dordrecht, Ch. Uniform barriers and evolutionarily stable sets, pp. 225–244.
- [6] Bomze, I. M., 2002. Regularity vs. degeneracy in dynamics, games, and optimization: a unified approach to different aspects. *SIAM Review* 44, 394–414.
- [7] Bomze, I. M., 2005. Portfolio selection via replicator dynamics and projections of indefinite estimated covariances. *Dyn. of Cont., Discr. and Impulsive Systems B* 12, 527–564.
- [8] Bomze, I. M., Pötscher, B. M., 1989. *Game Theoretical Foundations of Evolutionary Stability*. Springer.
- [9] Bomze, I. M., Weibull, J. W., 1995. Does neutral stability imply Lyapunov stability? *Games and Econ. Behaviour* 11, 173–192.
- [10] Brown, G. W., 1951. Activity analysis of production and allocation. Wiley, Ch. Iterative solution of games by fictitious play, pp. 374–376.
- [11] Brown, G. W., von Neumann, J., 1950. Contributions to the theory of games. Princeton University Press, Ch. Solutions of games by differential equations.
- [12] Conitzer, V., Sandholm, T., 2008. New complexity results about Nash equilibria. *Games and Econ. Behaviour* 63 (2), 621–641.
- [13] Etesami, K., Lochbihler, A., 2007. The computational complexity of evolutionarily stable strategies. *Int. J. Game Theory* 37 (1), 93–113.
- [14] Gilboa, I., Matsui, A., 1991. Social stability and equilibrium. *Econometrica* 59, 859–867.
- [15] Hofbauer, J., 1995. Stability for the best response dynamics, University of Vienna, unpublished manuscript.
- [16] Hofbauer, J., Sigmund, K., 1998. *Evolutionary games and population dynamics*. Cambridge University Press.
- [17] Lahkar, R., Sandholm, W. H., 2008. The projection dynamic and the geometry of population games. *Games and Econ. Behaviour* 64, 565–590.
- [18] Leonard, R. J., 1995. From parlour games to social science: von Neumann, Morgenstern, and the creation of game theory. *J. Econ. Literature* 33, 1928–1944.
- [19] Lohmann, G., Bohn, S., 2002. Using replicator dynamics for analyzing fMRI data of the human brain. *IEEE Trans. Medical Imaging* 21 (5), 485–492.
- [20] Maynard Smith, J., 1974. The theory of games and the evolution of animal conflicts. *J. Theor. Biology* 47, 209–221.
- [21] Maynard Smith, J., 1982. *Evolution and the theory of games*. Cambridge University Press.
- [22] Maynard Smith, J., Price, G. R., 1973. The logic of animal behaviour. *Nature* 16, 15–18.
- [23] Nash, J. F., 1950. Noncooperative games. Ph.D. thesis, Math. Dept., University of Princeton.
- [24] Nash, J. F., 1951. Non-cooperative games. *Annals of Math.* 54, 286–295.

- [25] Pavan, M., Pelillo, M., 2007. Dominant sets and pairwise clustering. *IEEE Trans. Pattern Anal. Machine Intell.* 29 (1), 167–172.
- [26] Pelillo, M., Torsello, A., 2006. Payoff-monotonic game dynamics and the maximum clique problem. *Neural Computation* 18 (5), 1215–1258.
- [27] Pouget, F., Dacier, M., Zimmerman, J., Clark, A., Mohay, G., 2006. Internet attack knowledge discovery via clusters and cliques of attack traces. *J. Inform. Assurance and Security* 1, 21–32.
- [28] Sandholm, W. H., 2001. Potential games with continuous player sets. *J. Econ. Theory* 97, 81–108.
- [29] Sandholm, W. H., 2005. Excess payoff dynamics and other well-behaved evolutionary dynamics. *J. Econ. Theory* 124, 149–170.
- [30] Sandholm, W. H., 2010. Pairwise comparison dynamics and evolutionary foundations for Nash equilibrium. *Games* 1, 3–17.
- [31] Sandholm, W. H., 2011. *Population games and evolutionary dynamics*. MIT Press, to appear Jan. 2011.
- [32] Smith, M. J., 1984. The stability of a dynamic model of traffic assignment – an application of a method of Lyapunov. *Transportation Science* 18, 245–252.
- [33] Taylor, P. D., Jonker, L., 1978. Evolutionary stable strategies and game dynamics. *Math. Biosciences* 40, 145–156.
- [34] Torsello, A., Rota Bulò, S., Pelillo, M., 2006. Grouping with asymmetric affinities: a game-theoretic perspective. In: *IEEE Conf. Computer Vision and Patt. Recogn. (CVPR)*. pp. 292–299.
- [35] Vasilakos, A., Anastasopoulos, M., 2007. *Advances in evolutionary computing for system design*. Vol. 66. Springer Berlin, Ch. Application of evolutionary game theory to wireless mesh networks, pp. 249–267.
- [36] Weibull, J. W., 1995. *Evolutionary game theory*. Cambridge University Press.