Convergence in games with continua of equilibria

Sebastian Bervoets, Mathieu Faure

Convergence in Games with Continua of Equilibria

Sebastian Bervoets*and Mathieu Faure†

May 12, 2020

Abstract

In game theory, the question of convergence of dynamical systems to the set of Nash equilibria has often been tackled. When the game admits a continuum of Nash equilibria, however, a natural and challenging question is whether convergence to the set of Nash equilibria implies convergence to a Nash equilibrium. In this paper we introduce a technique developed in Bhat and Bernstein (2003) as a useful way to answer this question. We illustrate it with the best-response dynamics in the local public good game played on a network, where continua of Nash equilibria often appear.

Keywords: Convergence; Continua of Nash Equilibria; Best-Response Dynamics.

JEL Codes: C62, C65

1 Introduction

The question of convergence of dynamical systems to some Nash equilibrium has often been explored in economics. Usually, convergence is discussed in contexts where the set of Nash equilibria is finite (see, for instance, the series of papers about convergence to the Cournot solution - Theocharis (1960), Fisher (1961), Hahn (1962), Seade (1980) among many others. In other games, see i.a. Arrow and Hurwicz (1960) or Rosen (1965)). Yet continua of Nash equilibria may appear in several economic situations. As Seade (1980) points out when discussing convergence of dynamical systems, "Things would

*Corresponding author. Aix Marseille Univ, CNRS, EHESS, Centrale Marseille, AMSE, Marseille, France. Email address: sebastian.bervoets@univ-amu.fr. Postal address: Aix Marseille School of Economics - 5 Boulevard Maurice Bourdet - 13205 Marseille Cedex 01 - France.

†Aix Marseille Univ, CNRS, EHESS, Centrale Marseille, AMSE, Marseille, France. Email address: mathieu.faure@univ-amu.fr.

get trickier (...) if equilibria happened not to be regular, that is not even locally unique, isolated. This, one can dismiss as a non-generic, 'unlikely' occurrence, although that is often a risky stand to take." Proving convergence in that case becomes problematic. In fact, to the best of our knowledge, no paper in economics addresses this issue.

When Nash equilibria are isolated, proving convergence amounts to showing that the distance between any solution curve and the set of Nash equilibria goes to zero. This is also necessary, but no longer sufficient, when equilibria are not isolated. Actually, the solution curve could very well approach the set of Nash equilibria, without ever converging to one specific element of that set. Convergence to an equilibrium when continua of equilibria exist has been explored in the dynamical systems literature (see for instance the book by Aulbach (2006) devoted to this problem). However, these techniques generally require strong regularity assumptions which fail to hold in most economic situations. In particular, they assume that the state space is an open set, while economic variables (such as prices, time allocation, efforts, quantities.) are typically defined on non-open sets[1]. This makes these convergence results inapplicable[2].

Two other methods to prove pointwise convergence of a dynamical system in presence of continua of equilibria have been developed, in Bhat and Bernstein (2010) and in Panageas and Piliouras (2016). The first is based on arclength, where the idea is to prove that every orbit has finite arc length. The authors then exploit a well-chosen Lyapunov function in order to prove their result. In the second, the authors prove convergence of the replicator dynamics in potential population games in presence of continua of equilibria. Their method consists in constructing a local Lyapunov function in the neighborhood of a given omega-limit point of the replicator dynamics. In both cases, the technique crucially relies on finding an appropriate Lyapunov function. This technique might be difficult to generalize or even to adapt to other settings, since the Lyapunov functions are specific to each problem and there is no systematic way of finding one.

In this paper we present and adapt the *non-tangency technique* (introduced in Bhat and Bernstein (2003)), which does not rely on Lyapunov functions nor does it require regularity assumptions. We illustrate how it works by analyzing a standard dynamical system - continuous-time best-response dynamics- in the local public good game introduced in Bramoullé and Kranton (2007). This game has received considerable attention

[1] Usually, agents' actions would be defined on $[0, +\infty[$ or on some compact subset.

[2] For instance, the techniques in Aulbach (2006) rely on the analysis of the Jacobian matrix of the dynamical system at every point of the manifold of equilibria. Obviously, when the state space is not an open set, this Jacobian matrix is not defined everywhere.

over recent years.[3]

In this game, players are placed on a network and interact only with their neighbors. The game has linear best responses and strategic substitutes, where individuals' payoffs depend on the sum of their neighbors' actions. It has been extensively studied in the recent literature, both for its great simplicity and for its rich structure of Nash equilibria. Of course the structure of the set of equilibria critically depends on the structure of the network, and in fact, Bervoets and Faure (2019) show in a companion paper that a substantial fraction of networks have continua of Nash equilibria. These ingredients combine to make this game the perfect candidate for our analysis of convergence.

The dynamical system that we consider is continuous-time best-response dynamics. We choose to focus on this specific dynamical system for various reasons. First, it is widely used in economics (see for instance the papers mentioned above about convergence to the Cournot solution). Second, it is related to many other dynamical system, in the sense that proving asymptotic properties of the solutions of this dynamical systems can be helpful for studying other dynamical systems (for instance replicator dynamics, see Hofbauer et al. (2009)). Also many adaptive dynamics can be studied through a careful analysis of the best-response dynamics (for instance fictitious play, see Benaïm et al. (2005) or similar algorithms, see Leslie and Collins (2006)). And third, it is simple enough to allow for a clear exposition of the non-tangency technique, when a more complex system would necessarily interfere with the understanding of the proof.

In the next section we present the local public good game and describe the structure of the set of Nash equilibria. In section 3 we define the continuous-time best-response dynamics and state our main result about convergence. We also provide an intuitive sketch of the proof, while the formal proof is in the appendix.

2 The local public good game

Consider a game $\mathcal{G} = (\mathcal{N}, X, u)$, where $\mathcal{N} = \{1, \ldots, N\}$ is the set of players, $X = \times_{i=1,\ldots,N} X_i$ where $X_i = [0, +\infty[$ is the action space, and $u = (u_i)_{i=1,\ldots,N}$ is the vector of payoff functions.

Agents are placed on a network represented by an undirected graph $\mathbf{G}$. By convention, we also denote by $\mathbf{G}$ the adjacency matrix of the graph, where $\mathbf{G}_{ij} = \mathbf{G}_{ji} = 1$ if players i and j are linked, and $\mathbf{G}_{ij} = 0$ otherwise. The set of neighbors of player i is $N_i(\mathbf{G}) := \{j \in \mathcal{N}, \mathbf{G}_{ij} = 1\}$. The game we focus on is the local public good game

[3] See Allouch and King (2019) for a recent survey on public good games on networks.

introduced in Bramoullé and Kranton (2007), where the payoff function is

$$u_i(x) = b\left(x_i + \sum_{j \in N_i(\mathbf{G})} x_j\right) - cx_i \tag{1}$$

where $c > 0$ is the marginal cost of effort and $b(.)$ is a differentiable, strictly increasing concave function. In order to rule out trivial cases, we assume that $\lim_{x\to+\infty} b'(x) < c < b'(0)$ and normalize so that $b'(1) = c$.

As can be seen by equation (1), a disconnected player i (i.e. $\mathbf{G}_{ij} = 0$ for all $j \in \mathcal{N}$) will choose to play $x_i^* = 1$. When connected, either the neighbors of i provide less than 1 and i aims to fill the gap to reach 1, or else neighbors provide more than 1 and i enjoys the benefits without exerting any effort.

Thus agents have a unique best response, given by:

$$\forall i \in \mathcal{N},\ \ Br_i(x_{-i}) = \max\left\{1 - \sum_{j \in N_i(\mathbf{G})} x_j, 0\right\}. \tag{2}$$

The set of Nash equilibria is therefore the set of all profiles x^* such that every player for which $x_i^* > 0$ is such that $x_i^* = 1 - \sum_{j\in N_i(\mathbf{G})} x_j^*$, while every player for which $x_i^* = 0$ is such that $\sum_{j\in N_i(\mathbf{G})} x_j^* \geq 1$. It should be noted that what matters for our analysis is the form of the best responses, not of the payoff functions (1). Thus our results extend to any payoff function resulting in best responses of the form (2).

The set of Nash equilibria can take complex forms. Importantly for our purposes here, this game often admits continua of equilibria. Let us illustrate this on the simplest possible network, the pair (Figure 1).

The profile $(1, 0)$ where $x_1 = 1$ and $x_2 = 0$ is a Nash equilibrium. The profile $(0, 1)$ is also a Nash equilibrium, together with, for instance, $(\frac{1}{3}, \frac{2}{3})$. In fact, any profile of the form $(\alpha, 1-\alpha)$, with $0 \leq \alpha \leq 1$, is a Nash equilibrium. We say that there is a *continuum* of equilibria[4], represented by the *connected component* $\Lambda = \{(\alpha, 1-\alpha) : \alpha \in [0, 1]\}$.

Definition 1 *Let $NE(\mathbf{G})$ denote the set of all Nash equilibria of the public good game played on network $\mathbf{G}$. Then, Λ is a connected component of $NE(\mathbf{G})$ if and only if:*

1 - Λ is connected

[4]The reader can find a detailed description of the possible structures of this set in Bervoets and Faure (2019). In that paper, we also conjecture that the fraction of networks on which the game admits at least one continuum of Nash equilibria goes to 1 as N grows, and we stress that continua are generally complex objects of potentially high dimensions.

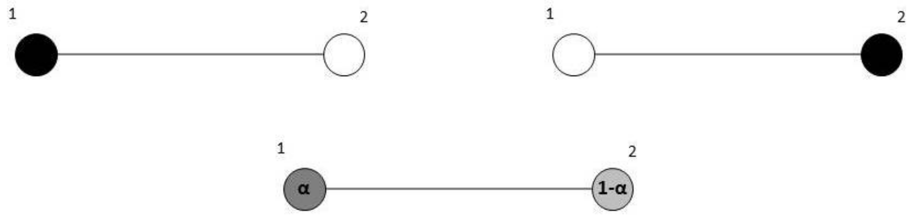

Figure 1: The pair network. Upper left and right panels: circles in black represent agents playing action 1 while white circles represent agents playing action 0; Lower panel: an action profile where agent 1 plays action α and agent 2 plays action $1-\alpha$. These three profiles belong to a continuum of Nash equilibria.

2 - $\Lambda \subset NE(\mathbf{G})$ *and*

3 - there exists an open neighborhood U *of* Λ *such that* $U \cap NE(\mathbf{G}) = \Lambda$.

An important characteristic of the set of Nash equilibria in this class of games is the following: given any network $\mathbf{G}$, the set of Nash equilibria $NE(\mathbf{G})$ can be described as a union of connected components, i.e. $NE(\mathbf{G}) = \cup_{i=1}^{L} \Lambda_i$, where every Λ_i is a connected component (see Bervoets and Faure (2019) for a formal proof).

3 Convergence of the Best-Response Dynamics

As seen in (2), best responses are unique. Let

$$Br : X \to X,\ x \mapsto Br(x) := (Br_1(x_{-1}), ..., Br_n(x_{-n})).$$

The continuous-time best-response dynamics is given by the following dynamical system:

$$\dot{x}(t) = -x(t) + Br(x(t)) \tag{3}$$

The map $Br(.)$ being Lipschitz, the ordinary differential equation (3) has a unique solution curve for any initial condition in $\mathbb{R}^N$. Because we restrict attention to $X = \mathbb{R}_+^N$ instead of $\mathbb{R}^N$ and to positive times ($t \geq 0$), we consider the semi-flow

$$\varphi : (x, t) \in X \times \mathbb{R}_+ \to \varphi(x, t) \in X, \tag{4}$$

where, for $t \geq 0$, $\varphi(x, t)$ is the unique solution of (3) at time t, with initial condition $x \in \mathbb{R}_+^N$. Then, for any $x_0 \in X$ and any $r > 1$, $\varphi(x_0, t) \in [0, r]^N$ for t large enough.

Indeed, if $x_i \geq r$ then $\dot{x}_i \leq 1 - r < 0$. This means that, for any $x_0 \in X$, $\omega(x_0)$ is nonempty and contained in $[0, 1]^N$.

We say that the system (3), starting from a point $x \in X$, *converges to a set* $S \subset X$ if

$$\lim_{t \to +\infty} d(\varphi(x, t), S) = 0$$

Equivalently, given $x \in X$, we call y an omega limit point of x for (3) if there exists a non-negative sequence $t_n \uparrow_n +\infty$ such that $\varphi(x, t_n) \to y$. The set of omega limit points of x is called the *omega limit set of* x and denoted $\omega(x)$. Starting from a point $x \in X$, saying that (3) converges to a set S is equivalent to saying that $\omega(x) \subset S$, as long as the trajectories of system (3) are bounded. If, for any initial condition x, the set $\omega(x)$ is a singleton we will say that the system is *pointwise convergent*.

We are now ready to state our result.

Theorem 1 *The continuous-time best-response dynamics defined in (3) is pointwise convergent to the set of Nash equilibria.*

In other words, for any initial conditions x_0, the solution curve starting from x_0 converges to a Nash equilibrium as t increases to infinity:

$$\forall x_0 \in X, \ \exists x^* \in NE(\mathbf{G}) \text{ such that } \ lim_{+\infty} \varphi(x_0, t) = x^*.$$

While it is straightforward to show the convergence of system (3) to the set of Nash equilibria, establishing this stronger result is much more laborious. One ingredient of the proof is adapting the non-tangency method developed by Bhat and Bernstein (2003) to our settings. We now sketch out the proof of the theorem avoiding technical details, and illustrate how it works on a very simple example (see the appendix for a completely detailed proof).

Let us denote by f the vector field of the dynamical system: $f(x) := -x + Br(x)$. We know that for any initial condition x_0 we have $\omega(x_0) \subset \Lambda$, where Λ is a connected component of Nash equilibria. Assume x is one point in the continuum of equilibria Λ, and $x \in \omega(x_0)$. We want to show that $\omega(x_0) = \{x\}$. One way of doing so is to look at how the vector field f behaves when the system approaches x. The way it behaves is captured by the possible directions the system can take when it approaches x. This is called the *direction cone* at point x, i.e. the convex cone generated by $f(U)$ for arbitrarily small open neighborhoods U of x.

Now we need to compare this set of possible directions with the set of directions that the system should take in order to stay in component Λ when close to x. These directions are given by the *tangent cone* at x. If the tangent cone and the direction cone have a nonempty intersection, then the system could keep moving along Λ, thus satisfying $\omega(x_0) \subset \Lambda$, without ever converging. If, on the contrary, the intersection between the tangent cone and the direction cone is the null vector, then we are guaranteed that the vector field f will be non-tangent to Λ at point x.

We illustrate how the proof works on a very simple case: consider the pair. In this case, as already stated above, there is one connected component of Nash equilibria $\Lambda = \{(\alpha, 1-\alpha), \alpha \in [0,1]\}$. Assume that we start at x_0 with $x_0^1 + x_0^2 < 1$ and that there exists some $\hat{x} = (\hat{x}_1, \hat{x}_2) \in \omega(x_0)$, that is interior (i.e. $1 > \hat{x}_1, \hat{x}_2 > 0$). By definition, $\hat{x}$ is a Nash equilibrium so that $\hat{x}_1 + \hat{x}_2 = 1$. At some point, the solution curve will enter a small neighborhood of $\hat{x}$. In that neighborhood, along the trajectory of the BRD, we have $\dot{x}_1 = \dot{x}_2 = (1 - x_1 - x_2)$, and the direction cone is given by all vectors proportional to $(1,1)$.

Now the tangent cone is defined by all the directions compatible with staying in Λ. So if x is in the relative interior, the admissible deviations are of the form $(+\alpha, -\alpha)$. The tangent cone is thus given by all vectors proportional to $(1,-1)$. Obviously, the intersection of the tangent cone and the direction cone can only be the null vector. As a consequence, the vector field is non-tangent to x, which is what we need to prove convergence. This example is relatively simple to deal with, especially since we use the simplest graph of interactions. In general however, the set of equilibria will possess corners, and finding a precise and convenient expression of the tangent cone is not straightforward. We provide all the technical details in the appendix.

4 Comments

Our method with respect to other dynamics. Our proof method nicely works for the best-response dynamics because it is piece-wise linear and therefore exhibits simple directional cones close to these continua of Nash equilibria. It also induces individuals who have more than 1 around them to simultaneously reduce their efforts, while those having less than 1 simultaneously increase their efforts. Thus, out of equilibrium both move in the same direction. On the contrary, once at equilibrium, changes in actions which allow individuals to stay in Λ are necessarily in opposite directions, since the sum of their efforts needs to remain constant.

Other dynamical systems might lead individuals to move in opposite directions even when out of equilibrium, forbidding us to use this method. Consider for instance the

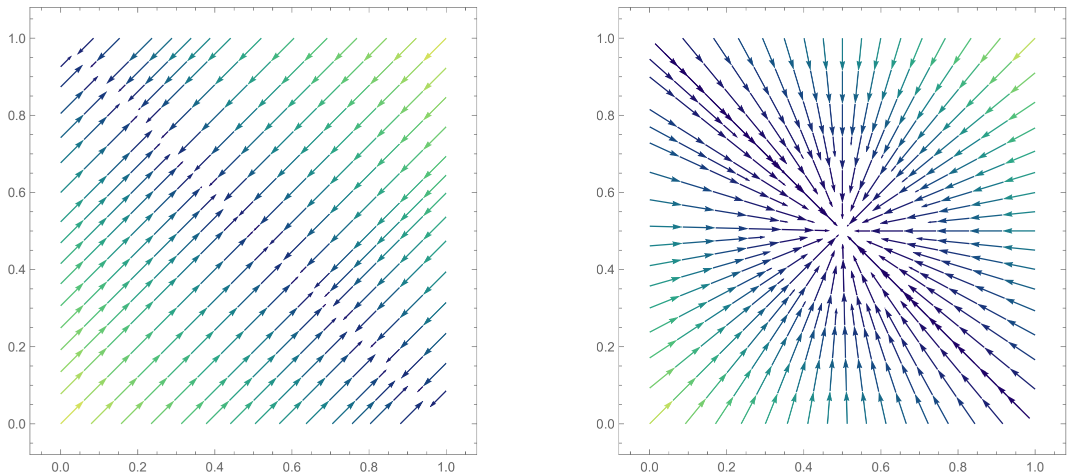

(a) Best-response dynamics

(b) Modified dynamics

Figure 2: Two different dynamical systems converging point-wise. On the left panel (a), the best-response dynamics, for which trajectories arrive non-tangentially to $\Lambda = \{(x, 1-x) : \ x \in [0,1]\}$. On the right panel (b), a point-wise convergent dynamical system, where every trajectory starting away from Λ converges to $(\frac{1}{2}, \frac{1}{2})$.

following dynamical system:

$$\begin{cases} \dot{x_1} &= (\frac{1}{2} - x_1)| - x_1 + Br_1(x)| \\ \\ \dot{x_2} &= (\frac{1}{2} - x_2)| - x_2 + Br_2(x)| \end{cases}$$

The behavioral difference of the dynamics is illustrated on Figure (2).

The set of equilibria of this modified dynamics is again $\Lambda = \{(\alpha, 1-\alpha), \alpha \in [0,1]\}$, the tangent cone is the same as for the best-response dynamics and as can be seen on Figure (2), the system is point-wise convergent. However, the non-tangency method cannot be used to prove this because the direction cone at any $\hat{x} \in \Lambda$ contains $(1, -1)$ (resp. $(-1, 1)$) if $\hat{x}_1 \leq \hat{x}_2$ (resp. $\hat{x}_1 \geq \hat{x}_2$). Therefore the intersection of the tangent cone with the direction cone now contains non-zero vectors.

Asymmetric networks. The proof of our main result uses the symmetry of the interaction network. However the relationship between network symmetry and pointwise convergence is not very clear. We could have asymmetric networks for which the dynamical system converges to the set of Nash equilibria, but for which our proof cannot be used to prove pointwise convergence (or to prove non-pointwise convergence). In Bayer et al. (2020), the authors shows that the public good game played on networks which are *rescalable into a symmetric one*, or exhibit *weak externalities* are best-response potential games (for definitions, see Bayer et al. (2020)). As we show in the first lines of the appendix, the existence of a best-response potential guarantees that the continuous-time best-response will converge to the set of Nash equilibria. However, this does not imply pointwise convergence when the equilibria are not isolated. More importantly, we could have asymmetric networks for which the system cycles. Consider the following asymmetric network with 10 agents.

$$G = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 1 & 1 & 1 & 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 1 \\ 0 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 1 & 0 \end{bmatrix}.$$

The game generated by this network has a Nash equilibrium at

$$x^* = (1/3, 1/3, 1/8, 1/8, 1/8, 1/8, 1/12, 1/12, 1/12, 1/12)$$

which is part of the continuum of equilibria

$$\Lambda = \left\{ x^* + (0, 0, \alpha, -\alpha, \beta, -\beta, 0, 0, 0, 0) : \ \alpha, \beta \in \left[-\frac{1}{8}, \frac{1}{8} \right] \right\}$$

By denoting $(a)_+ := Max\{a, 0\}$, the best-response dynamics is given by the system

$$\dot{x}_i = -x_i + \left(1 - \sum_j \mathbf{G}_{ij} x_j \right)_+ , \quad i = 1, ..., 10. \tag{5}$$

Network G is neither rescalable into a symmetric one nor does it exhibit weak externalities. In fact, we show that system (5) exhibits cycles. We construct these cycles by studying a related three-dimensional linear system, for which existence of periodic orbits can be easily shown:

$$\begin{bmatrix} \dot{X} \\ \dot{Y} \\ \dot{Z} \end{bmatrix} = -2(A + I) \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} + 2 \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \tag{6}$$

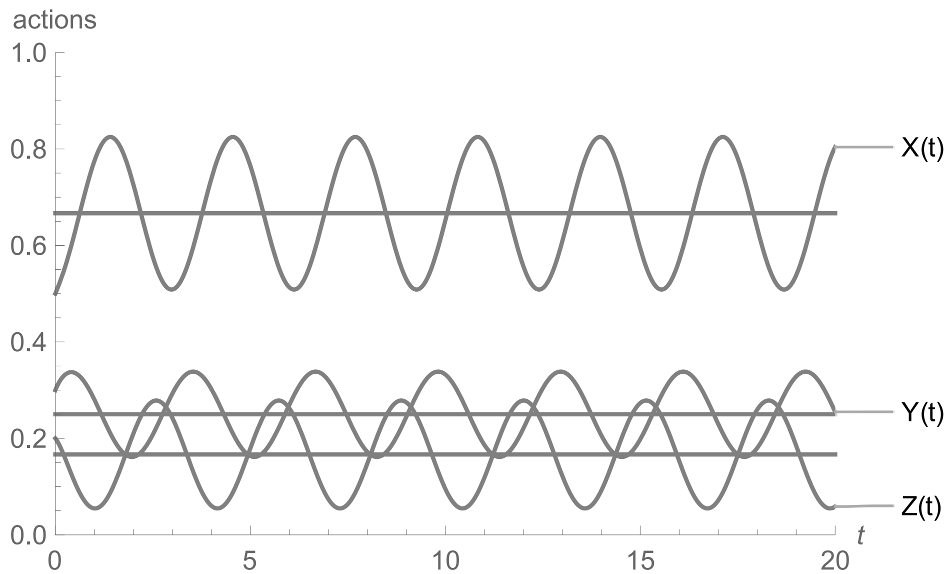

Figure 3: The solution of the dynamical system (6), starting at $X_0 = 0.5, Y_0 = 0.3, Z_0 = 0.2$. The constant lines represent the equilibrium of the game with interaction matrix A $((X^*, Y^*, Z^*) = (\frac{2}{3}, \frac{1}{4}, \frac{1}{6}))$ around which the solutions cycle.

where

$$A = \begin{bmatrix} 0 & 0 & 2 \\ 1 & 0 & 1/2 \\ 1/2 & 2 & 0 \end{bmatrix}$$

This system has an equilibrium in $(X^*, Y^*, Z^*) = (2/3, 1/4, 1/6)$. The eigenvalues of matrix $(A+I)$ are $\{3, i, -i\}$, hence system (6) exhibits a continuum of periodic orbits around (X^*, Y^*, Z^*), with as small amplitude as one wants. Note that generic initial conditions near the equilibrium generate asymptotically cycling trajectories (and not exact cycles). We illustrate this by drawing an asymptotically cycling trajectory in Figure (4).

Consider a non-trivial periodic orbit $(X(t), Y(t), Z(t))$ of system (6), with initial conditions close enough from the equilibrium, so that $X(t)+4Z(t) \leq 2,\ 2X(t)+Y(t)+Z(t) \leq 2$ and $X(t)+4Y(t)+Z(t) \leq 2$, for any $t \geq 0$. Then $(X(t), Y(t), Z(t))_{t \geq 0}$ is a solution of the system

$$\begin{aligned}\dot{X} &= -X + (2 - X - 4Z)_+ \\ \dot{Y} &= -Y + (2 - 2X - Y - Z)_+ \\ \dot{Z} &= -Z + (2 - X - 4Y - Z)_+\end{aligned}$$

Now let $x(t) := \frac{1}{2}\left(X(t), X(t), Y(t), Y(t), Y(t), Y(t), Z(t), Z(t), Z(t), Z(t)\right)$ for $t \geq 0$. Then $(x(t))_{t\geq 0}$ is a non-trivial periodic orbit of system (5). To see this, observe for instance that $\dot{x}_1(t) = \frac{1}{2}\dot{X}(t) = -X(t) + 1 - 2Z(t)$. Since $x_1(t) = x_2(t) = \frac{1}{2}X(t)$ and $x_7(t) = x_8(t) = x_9(t) = x_{10}(t) = \frac{1}{2}Z(t)$, we get

$$\dot{x}_1(t) = -x_1(t) + 1 - x_2(t) - (x_7(t) + x_8(t) + x_9(t) + x_{10}(t))$$

which corresponds to the first equation of system (5).

Acknowledgments

This work was supported by French National Research Agency Grants ANR-17-EURE-0020. Mathieu Faure gratefully acknowledges the support of the French National Research Agency, under grant ANR CIGNE (ANR-15-CE38-0007-01).

Appendix - Proof of Theorem 1

We start the proof by showing that omega-limit sets of the best-response dynamics are contained in the set of Nash equilibria. Let the network $\mathbf{G}$ be fixed and define $P : X \to \mathbb{R}$ as follows:

$$P(x) = \sum_i x_i \left(1 - \frac{1}{2}x_i - \frac{1}{2}\sum_j \mathbf{G}_{ij}x_j\right) \tag{7}$$

defined in (7) is a strict Lyapunov function for $\dot{x} = -x(t) + Br(x(t))$, that is:[5]

- for $x \in NE(\mathbf{G})$ the map $t \mapsto P(\varphi(x,t))$ is constant;
- for $x \notin NE(\mathbf{G})$ the map $t \mapsto P(\varphi(x,t))$ is strictly increasing.

This property implies that only Nash equilibria can be local maxima of the Lyapunov function P. More importantly it For any $x \in \omega(x_0)$, we have $P(x) = \lim_{t\to+\infty} P(\varphi(x_0,t))$. Now $\omega(x_0)$ being invariant directly implies that $P(\varphi(x,t)) = P(x)$ for any t, which means that $x \in NE$ because $x \mapsto P(\varphi(x,t))$ is strictly increasing for $x \notin NE(\mathbf{G})$. We can actually go farther: recall that the set NE can be written as $NE = \cup_{i=1}^{L} \Lambda_i$. Then, since $\omega(x_0)$ is connected, it must necessarily be contained in a connected component of NE, i.e. $\omega(x_0) \subset \Lambda_l$ for some Λ_l.

However, convergence to the set of Nash equilibria does not imply convergence to a Nash equilibrium. In order to prove convergence to an equilibrium point from any initial condition i.e. that $|\omega(x_0)| = 1$ for any $x_0 \in X$, we show that the vector field of the dynamical system (3) is non-tangent to the set of Nash equilibria.

In order to proceed, we will use a proposition which determines some characteristics of the components of Nash equilibria. Let x be a profile of actions. We define the following subsets of agents:

$$I(x) := \{i \in \mathcal{N}; x_i = 0 \text{ and } \sum_{j\in N_i(\mathbf{G})} x_j = 1\}$$

$$SI(x) := \{i \in \mathcal{N}; x_i = 0 \text{ and } \sum_{j\in N_i(\mathbf{G})} x_j > 1\}$$

$$A(x) := \{i \in \mathcal{N}; x_i > 0\}$$

$I(x)$ is the set of inactive players, $SI(x)$ is the set of strictly inactive players, and $A(x)$ is the set of active players. Note that, if $\hat{x}$ is a Nash equilibrium, then $(I(\hat{x}), SI(\hat{x}), A(\hat{x}))$ form a partition of $\mathcal{N}$. We have the following:

[5] See for instance Bervoets and Faure (2019) for a proof of this claim.

Lemma 1 *Let $x_0 \in X$. There exists $\hat{x} \in \omega(x_0)$ and $\epsilon > 0$ such that*

$$\forall y \in B(\hat{x}, \epsilon) \cap \omega(x_0), I(y) = I(\hat{x}).$$

Proof Choose some $\hat{x} \in \omega(x_0)$ with minimal set of inactive players, i.e. for any $x \in \omega(x_0)$, $I(x)$ is not strictly contained in $I(\hat{x})$. Since $\hat{x}_i > 0$ for all $i \in A(\hat{x})$ and $\hat{x}_i + \sum_j \mathbf{G}_{ij}\hat{x}_j > 1$ for all $i \in SI(\hat{x})$, there exist $\epsilon > 0$ such that, for any $y \in B(\hat{x}, \epsilon)$, by continuity we have $y_i > 0$ for all $i \in A(\hat{x})$ and $y_i + \sum_j \mathbf{G}_{ij} y_j > 1$ for all $i \in SI(\hat{x})$. As a consequence $I(y) \subset I(\hat{x})$ $\forall y \in B(\hat{x}, \epsilon)$, and since $\hat{x}$ is minimal, it directly implies that $I(y) = I(\hat{x})$. ■

One implication of this lemma is that it is always possible to find a Nash equilibrium $\hat{x}$ in $\omega(x_0)$ for which the partition of agents into $I(\hat{x}), SI(\hat{x}), A(\hat{x})$ remains the same in a neighborhood.

Recall that f denotes the vector field of the system (3): $f(x) := -x + Br(x)$. The next definition formally introduces the direction cone associated with f.

Definition 2 *Given $x \in X$, the* direction cone *$\mathcal{F}_x$ of f at x is defined by*

$$\mathcal{F}_x := \cap_{\epsilon > 0} coco(f(B(x, \epsilon)) \setminus \{0\}),$$

where $coco(A)$ is the cone generated by the convex hull of A.

In practice, $v \in \mathcal{F}_x$ if and only if for any $\epsilon > 0$ there exists $\lambda > 0$ and u in the convex hull of $f(B(x, \epsilon) \setminus \{0\})$ such that $v = \lambda u$.

Definition 3 *Let $A \subset X$ and $x \in A$. The* tangent cone *to A in x is the set of directions $v \in \mathbb{R}^N$ such that there exists a sequence $(x_n)_n$ in A converging to x and $h_n \downarrow_n 0$ with the following property:*

$$v = \lim_{n \to +\infty} \frac{x_n - x}{h_n}.$$

We denote this set $T_x A$.

We now state the result of Bhat and Bernstein (2003), adapted to our context.

Proposition 1 (Proposition 5.2 in Bhat and Bernstein (2003)) *Let $x_0 \in X$ and $\hat{x} \in \omega(x_0)$. Then if $T_{\hat{x}}\omega(x_0) \cap \mathcal{F}_{\hat{x}} \subset \{0\}$ we have $\omega(x_0) = \{\hat{x}\}$.*

Let us determine the tangent cone. Let $\hat{x}$ be a Nash equilibrium and Λ be the connected component containing $\hat{x}$. Then $v \in T_{\hat{x}}\Lambda$ if and only if

$$
\begin{aligned}
v_i &\geq 0 \ \forall i \in I(\hat{x}) && (8)\\
v_i &= 0 \ \forall i \in SI(\hat{x}) && (9)\\
v_i + \sum_j \mathbf{G}_{ij} v_j &= 0 \ \forall i \in A(\hat{x}) && (10)\\
v_i + \sum_j \mathbf{G}_{ij} v_j &= 0 \ \forall i \in I(\hat{x}) \text{ s.t. } v_i > 0, && (11)\\
v_i + \sum_j \mathbf{G}_{ij} v_j &\geq 0 \ \forall i \in I(\hat{x}) && (12)
\end{aligned}
$$

These conditions are the "marginal movements" which allow every agent to stay in Λ. Inactive agents can either remain inactive, become strictly inactive or become active. Condition (6) says that their efforts cannot become negative. Condition (9) says that if they become active, then the sum of efforts around them, own effort included, should remain constant. Condition (10) says that the sum of efforts around them, own effort included, can only increase. Condition (7) says that strictly inactive agents can only remain strictly inactive. Condition (8) says that active agents can move in any direction such that the sum of efforts around them, own effort included, remains constant. These conditions combined characterize the tangent cone at $\hat{x}$.
One useful corollary of Lemma 1 is the following:

Corollary 1 *Let $\hat{x}$ be as in Lemma 1. Then for any $v \in T_{\hat{x}}\omega(x_0)$, $v_i = 0$ and $v_i + \sum_j \mathbf{G}_{ij} v_j = 0 \quad \forall i \in I(\hat{x})$.*

Proof of Corollary 1. Let $v \in T_{\hat{x}}\omega(x_0)$ and $i \in I(\hat{x})$. There exists some sequence x^n converging to $\hat{x}$ and h_n going to zero such that

$$v_i = \lim_n \frac{x_i^n - \hat{x}_i}{h_n}.$$

For large enough n we then have $x_i^n = 0$, which implies that $v_i = 0$. Also

$$v_i + \sum_j \mathbf{G}_{ij} v_j = \lim_n \frac{1}{h_n} \left(x_i^n + \sum_j \mathbf{G}_{ij} x_j^n - (\hat{x}_i + \sum_j \mathbf{G}_{ij} \hat{x}_j) \right) = 0$$

because $i \in I(x^n)$ for large enough n. ■

We can now prove our main result.

Proof of Theorem 1. We must prove that $|\omega(x_0)| = 1$ for any $x_0 \in X$. Let $x_0 \in X$ and $\hat{x} \in \omega(x_0)$ be as in Lemma 1. We assume without loss of generality that $T_{\hat{x}}\omega(x_0) \cap \mathcal{F}_{\hat{x}}$ is nonempty (otherwise there is nothing to prove). Pick $v \in T_{\hat{x}}\omega(x_0) \cap \mathcal{F}_{\hat{x}}$. We must prove that $v = 0$. According to the last corollary, we have $v_i = 0$ and $v_i + \sum_j \mathbf{G}_{ij} v_j = 0$ $\forall i \in I(\hat{x})$.

Now let $i \in A(\hat{x})$. There exists $\epsilon > 0$ such that $f_i(x) = 1 - \left(x_i + \sum_j \mathbf{G}_{ij} x_j\right) = (\hat{x} - x)_i + \sum_j \mathbf{G}_{ij}(\hat{x} - x)_j$ for any $x \in U := B(\hat{x}, \epsilon)$. Also $f_i(x) = -x_i$ for $i \in SI(\hat{x})$. As a consequence $f(U) \subset C(U)$, where

$$C(U) := \{w \in \mathbb{R}^N : \exists x \in U \text{ s.t. } w_i = (\hat{x}-x)_i + \sum_j \mathbf{G}_{ij}(\hat{x}-x)_j \ \ \forall i \in A(\hat{x}),\ \ w_i = -x_i \ \forall i \in SI(\hat{x})\},$$

which is convex. By definition of $co(f(U))$, we therefore have $co(f(U)) \subset C(U)$. Now since $v \in coco(f(U))$, we have $v = \lambda w$, with $\lambda > 0$ (if $\lambda = 0$ there is nothing to prove) and $w \in C(U)$ associated with some $x \in U$. As $v \in T_{\hat{x}}\omega(x_0)$, we have $x_i = -w_i = -\frac{v_i}{\lambda} = 0$ for $i \in SI(\hat{x})$.

Let us prove that $v = 0$. Note that $v_i = \lambda \left((\hat{x} - x)_i + \sum_j \mathbf{G}_{ij}(\hat{x} - x)_j\right)$ for $i \in A(\hat{x})$. We also have $v_i = 0$ for $i \notin A(\hat{x})$. Thus,

$$\begin{aligned}
\frac{1}{\lambda}\|v\|^2 &= \sum_{i \in N} v_i \left((\hat{x} - x)_i + \sum_{j \in N} \mathbf{G}_{ij}(\hat{x} - x)_j\right) && (13)\\
&= \sum_{i \in N} v_i(\hat{x} - x)_i + \sum_{j \in N}\sum_{i \in N} v_i \mathbf{G}_{ij}(\hat{x} - x)_j && (14)\\
&= \sum_{i \in N} v_i(\hat{x} - x)_i + \sum_{j \in N}(\hat{x} - x)_j \sum_{i \in N} \mathbf{G}_{ji} v_i && (15)\\
&= \sum_{i \in N} v_i(\hat{x} - x)_i - \sum_{j \in N}(\hat{x} - x)_j v_j && (16)\\
&= 0 && (17)
\end{aligned}$$

where we used the fact that $\mathbf{I_d} + \mathbf{G}$ is symmetric to obtain (15). FInally (16) comes from the fact that $(\hat{x} - x)_j = 0$ for $j \in SI(\hat{x})$, and that $\sum_{i \in N} \mathbf{G}_{ji} v_i = -v_j$ for any $j \notin SI(\hat{x})$ (by (11) and (12)). The proof is complete. ■

References

Allouch, N. and King, M. (2019). Constrained public goods in networks. *Journal of Public Economic Theory*, 21(5):895–902.

Arrow, K. J. and Hurwicz, L. (1960). Stability of the gradient process in n-person games. *Journal of the Society for Industrial and Applied Mathematics*, 8(2):280–294.

Aulbach, B. (2006). *Continuous and discrete dynamics near manifolds of equilibria*, volume 1058. Springer.

Bayer, P., Kozics, G., Szőke, N. G., et al. (2020). Best-response dynamics in directed network games. Technical report, Department of Economics, Central European University.

Benaïm, M., Hofbauer, J., and Sorin, S. (2005). Stochastic approximations and differential inclusions. *I. SIAM Journal on Optimization and Control*, 44:328–348.

Bervoets, S. and Faure, M. (2019). Stability in games with continua of equilibria. *Journal of Economic Theory*, 179:131–162.

Bhat, S. P. and Bernstein, D. S. (2003). Nontangency-based lyapunov tests for convergence and stability in systems having a continuum of equilibria. *SIAM Journal on Control and Optimization*, 42(5):1745–1775.

Bhat, S. P. and Bernstein, D. S. (2010). Arc-length-based lyapunov tests for convergence and stability with applications to systems having a continuum of equilibria. *Mathematics of Control, Signals, and Systems*, 22(2):155–184.

Bramoullé, Y. and Kranton, R. (2007). Public goods in networks. *Journal of Economic Theory*, 135(1):478–494.

Fisher, F. M. (1961). The stability of the cournot oligopoly solution: The effects of speeds of adjustment and increasing marginal costs. *The Review of Economic Studies*, 28(2):125–135.

Hahn, F. H. (1962). The stability of the cournot oligopoly solution. *The Review of Economic Studies*, 29(4):329–331.

Hofbauer, J., Sorin, S., and Viossat, Y. (2009). Time average replicator and best-reply dynamics. *Mathematics of Operations Research*, 34(2):263–269.

Leslie, D. and Collins, E. (2006). Generalised weakened fictitious play. *Games and Economic Behavior*, 56(2):285–298.

Panageas, I. and Piliouras, G. (2016). Average case performance of replicator dynamics in potential games via computing regions of attraction. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 703–720. ACM.

Rosen, J. B. (1965). Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica:*, pages 520–534.

Seade, J. (1980). The stability of cournot revisited. *Journal of Economic Theory*, 23(1):15–27.

Theocharis, R. D. (1960). On the stability of the cournot solution on the oligopoly problem. *The review of economic studies*, 27(2):133–134.