Adaptive Learning in Weighted Network Games

Citation for published version (APA):

Bayer, P., Herings, P. J.-J., Peeters, R., & Thuijsman, F. (2017). *Adaptive Learning in Weighted Network Games*. Maastricht University, Graduate School of Business and Economics. GSBE Research Memoranda No. 025 <https://doi.org/10.26481/umagsb.2017025>

Document status and date:

Published: 24/10/2017

DOI:

[10.26481/umagsb.2017025](https://doi.org/10.26481/umagsb.2017025)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Péter Bayer,
P. Jean-Jacques Herings,
Ronald Peeters, Frank Thuijsman

**Adaptive Learning in
Weighted Network Games**

RM/17/025

GSBE

Maastricht University School of Business and Economics
Graduate School of Business and Economics

P.O. Box 616
NL- 6200 MD Maastricht
The Netherlands

Adaptive Learning in Weighted Network Games*

Péter Bayer[†] P. Jean-Jacques Herings[‡]

Ronald Peeters[§] Frank Thuijsman[¶]

October 12, 2017

Abstract

This paper studies adaptive learning in the class of weighted network games. This class of games includes applications like research and development within interlinked firms, crime within social networks, the economics of pollution, and defense expenditures within allied nations. We show that for every weighted network game, the set of pure Nash equilibria is non-empty and, generically, finite. Pairs of players are shown to have jointly profitable deviations from interior Nash equilibria. If all interaction weights are either non-negative or non-positive, then Nash equilibria are Pareto inefficient. We show that quite general learning processes converge to a Nash equilibrium of a weighted network game if every player updates with some regularity.

Keywords: Networks, Learning, Public goods, Potential games.

JEL classification: C72, D74, D83, D85, H41.

*This project is financed by the Centre for Human Enhancement and Learning.

[†]Department of Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands.
E-mail: p.bayer@maastrichtuniversity.nl

[‡]Department of Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands.
E-mail: p.herings@maastrichtuniversity.nl

[§]Department of Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands.
E-mail: r.peeters@maastrichtuniversity.nl

[¶]Department of Data Science and Knowledge Engineering, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands. E-mail: f.thuijsman@maastrichtuniversity.nl

1 Introduction

The theory of learning is of fundamental importance in game theory. With most of the focus in the non-cooperative game theory literature being devoted to the study of equilibria – various concepts, characterization of the equilibrium set, properties, refinements – it is critical to understand how equilibrium is reached. However, the main concepts of equilibrium theory, and in particular, the concept of Nash equilibrium, have proven difficult to validate, especially in one-shot games. To quote Fudenberg and Levine (1998): “One traditional explanation of equilibrium is that it results from analysis and introspection by the players in a situation where the rules of the game, the rationality of the players, and the players’ payoff functions are common knowledge. Both conceptually and empirically, these theories have many problems.” One of the main goals of learning in game theory is to provide such a motivation. For one-shot games this is typically achieved by interpreting the equilibrium points as results of a series of updates by the players acting in a recurrent setting of that game. These updates are made in response to observed moves by their opponents, with various assumptions on rationality. Ideally, as the players discover more about the game and about their opponents, their collective decisions should, in time, resemble equilibrium play. As such, the learning literature focuses mainly on the stability and convergence properties of various learning processes.

The class of games in which we frame our analysis is the class of weighted network games. This class of games corresponds to the private provision of local public goods games introduced by Bramoullé and Kranton (2007) for undirected graphs and generalized by Bramoullé et al. (2014) for weighted networks. For a comprehensive overview of related models see sections 3 and 4 of Jackson and Zenou (2014). The main practical reason this class of games is worth studying, is its wide range of applications in various subfields of economic theory, including R&D within interlinked firms (König et al., 2014), crime within a social network (Ballester et al., 2006), and peer effects with spatial interactions (Blume et al., 2010). Further applications include pollution models as in Leontief (1970) as well as defense expenditures within an international community as studied by Sandler and Hartley (1995) and Sandler and Hartley (2007). Networks offer a simple way to model complex interactions between many decision makers. The simplest network models are undirected graphs, in which a link between a pair of players indicates a direct interaction. Since players may be indirectly affected by the neighbors of their neighbors, and so on, each interaction may be relevant for each player, resulting in a profoundly rich model. In weighted networks, interaction weights with arbitrary values, either positive or negative, are used to characterize the way that pairs of interacting players influence each other.

The parameters of a weighted network game are the weighted network itself, describing the interactions between the players, a vector of targets that describes the players' needs, and a vector of upper bounds representing the players' highest possible activity levels. Each player has a concave benefit function of the weighted aggregate activity and a linear cost function of his own activity. We show that the set of Nash equilibria of weighted network games is non-empty and generically finite. Additionally, we show that under quite general conditions pairs of players can jointly improve their payoffs, so Nash equilibria are not strong. We also give conditions such that they are not Pareto efficient.

We study a class of learning processes with the following features. The players update their decisions at discrete points in time, maximizing their payoffs for a single period. The updates determine the status quo of the next period. At any given period, only one player is allowed to update, the actions of every other player remain the same. This class of learning processes includes e.g. the improvement paths of Monderer and Shapley (1996).

Weighted network games are generalized aggregative games (Dubey et al., 2006) as well as best-response potential games (Voorneveld, 2000), but may not belong to the class of ordinal potential games (Monderer and Shapley, 1996). Since weighted network games generally do not have an ordinal potential, better-response dynamics may not converge, and we show the possibility of non-convergence by an example.

Our main results concern the properties of adaptive learning processes centered around the best responses. We find that convergence to the set of Nash equilibria requires two conditions: (1) each update has to take the player closer to his contemporary best response, and (2) with some regularity, every player must have the possibility to update. Furthermore, we show that such processes converge to a Nash equilibrium point, if (3) the set of Nash equilibria is finite. The first and second conditions concern the players, and may be interpreted as assumptions of cautiousness and activity, respectively. The third condition concerns the parameters of the weighted network game, and is generically satisfied, as mentioned above. The main significance of our results is in the fact that convergence of the learning process to a Nash equilibrium can be achieved with relatively weak assumptions on the behavior of the players.

To our knowledge, our paper is the first to consider discrete-time learning processes in the setting of weighted network games. Bramoullé et al. (2014) and Bervoets and Faure (2016) study best-response dynamics in continuous time. Bervoets et al. (2016) considers a two-stage stochastic learning process with experimenting players that converges with probability one. Eskin et al. (2012) considers a similar game of incomplete information played on a graph.

The paper proceeds as follows. Section 2 introduces weighted network games. Section 3

contains the characterization of the set of Nash equilibria and its welfare properties. In Section 4 we explore the convergence properties of learning processes centered around the best response. Section 5 concludes.

2 Weighted network games

Let $I = \{1, \dots, n\}$ denote the set of players with $n \geq 2$. The action set of player $i \in I$ is $X_i = [0, \bar{x}_i]$ for $\bar{x}_i > 0$. Let $x_i \in X_i$ denote player i 's action. The action profile of all players is denoted by $x = (x_j)_{j \in I}$ and the action profile of all players except i by $x_{-i} = (x_j)_{j \neq i}$. Similarly, $X = \prod_{i \in I} X_i$ denotes the set of action profiles and $X_{-i} = \prod_{j \in I \setminus \{i\}} X_j$ the set of action profiles for all players other than i .

Definition 2.1. The tuple $G = (I, X, (\pi_i)_{i \in I})$ is called a *weighted network game* if for every $i \in I$ the payoff function $\pi_i: X \rightarrow \mathbb{R}$ is given by:

$$\pi_i(x) = f_i\left(\sum_{j \in I} w_{ij}x_j\right) - c_i x_i,$$

with cost parameters $c_i > 0$, interaction parameters $w_{ij} \in \mathbb{R}$, and benefit functions $f_i: \mathbb{R} \rightarrow \mathbb{R}$.

Assumption 2.2. For every $i \in I$, $w_{ii} = 1$, and for every $i, j \in I$, $w_{ij} = w_{ji}$. Furthermore, for every $i \in I$, the benefit function f_i is twice continuously differentiable and satisfies the following properties: (1) $f'_i > 0$, (2) $f''_i < 0$, and (3) there exists $t_i \in \mathbb{R}$ such that $f'_i(t_i) - c_i = 0$.

The interpretation is the following. Each player $i \in I$ produces a specialized good using a linear production technology. The costs of producing one unit of the good are equal to c_i . The production of player i is denoted by x_i . Each player consumes his own good as well as those of his neighbors. The total amount of consumption of player i is $\sum_{j \in I} w_{ij}x_j$, the benefit of consumption is $f_i(\sum_{j \in I} w_{ij}x_j)$, and the desired amount of consumption, called target value, is t_i . For player i , the parameter w_{ij} captures the substitutability of one unit of player j 's good to his own. If $w_{ij} > 0$, then player j 's production generates positive externalities for player i . Specifically, if $w_{ij} = 1$, then player i 's enjoyment of player j 's good equals that of his own good. If $w_{ij} \in (0, 1)$, then player i derives less benefits from player j 's good than from his own. If $w_{ij} \in (1, \infty)$, then player i enjoys the good of player j more than that of his own. Negative values of w_{ij} indicate that player j 's production has negative external effects on player i 's benefits, with $w_{ij} \in (0, -1)$, $w_{ij} = -1$, and $w_{ij} \in (-1, -\infty)$ indicating that the negative effects are smaller, equal, or greater in magnitude than the positive effects of equal amounts of the own good. The assumption $w_{ii} = 1$ is a normalization. The symmetry assumption $w_{ij} = w_{ji}$

for $i, j \in I$ is also made in previous studies like Dubey et al. (2006), Bramoullé and Kranton (2007), and Bramoullé et al. (2014). The asymmetric case $w_{ij} \neq w_{ji}$ is relatively unexplored in the local public good setting, and hence an interesting direction for future research (see Bourlès et al. (2017) for a model of transfers with asymmetric interactions).

The consumption benefit is strictly increasing in x_i and concave by properties (1) and (2) of the benefit function. The marginal benefits equal the marginal costs of production at the target value t_i by property (3). Note that since the target value t_i can be above or below values achievable by using action profiles in X , property (3) of the benefit function is without loss of generality.

Example 2.3. Let $I = \{1, 2\}$ be a set of two countries that have to decide on the level of their defense expenditures. We take $X_i = [0, Z_i]$, where Z_i denotes the GDP of Country $i \in I$, and

$$W = \begin{pmatrix} 1 & w_{12} \\ w_{12} & 1 \end{pmatrix}.$$

The increasing, concave functions f_1 and f_2 indicate the countries' benefits from defense. Let $t_i = 0.01Z_i$, indicating that both countries have a target value for defense expenditure of 1% of their GDP. This is the amount they would spend on defense if the other nation spends nothing.

If $w_{12} = 0$, then neither country benefits from the other's defense expenditure, nor are they threatened by it. This may indicate neutrality or a significant geographical distance. If $w_{12} > 0$, the two nations are allies, and the game becomes a game of strategic substitutes. In this case both nations benefit from the other's defense spending and therefore, national defense expenditures are likely to be lower than 1% of GDP. If $w_{12} < 0$, the two nations are hostile to each other, and the game is a game of strategic complements. In this case the nations are hurt or threatened by the other's defense spending and hence, defense expenditures will likely exceed 1% of GDP.

In case $w_{12} = 1$, Example 2.3 results in the 2-player pure public good model of defense expenditure between allies, while $0 < w_{12} < 1$ gives the symmetric version of the limited substitutability public good model of defense expenditure between allies, developed by Sandler and Hartley (2001). They do not consider the case $w_{12} < 0$. In the subsequent section, we discuss how the set of Nash equilibria of this particular game depends on the parameters t_1, t_2 , and w_{12} in more detail.

Our setup allows the modeling of more intricate relationships between players, as illustrated by the following example.

Example 2.4. Let $I = \{1, 2, 3\}$ be a set of three countries deciding on the level of their defense expenditures, $X_i = [0, Z_i]$, and

$$W = \begin{pmatrix} 1 & -1 & 1 \\ -1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix},$$

and $t = (0.03Z_1, 0.03Z_2, 0.01Z_3)$. In this example, Countries 1 and 2 are rivals, but both of them are friendly to Country 3. An example of this type of relationship may be that of Israel, Saudi Arabia, and the U.S. since the 2000s.

Example 2.4 and similar examples that feature intransitive relationships between countries cannot be modeled in the spirit of Sandler and Hartley (2001). Our setup therefore reflects more closely the possible intricacies of diplomatic relationships, and can be used to model any system of alliances and threats, provided that the relationship between any two nations is symmetric.

Games with strategic substitutes and complements are of great relevance in the economic literature. The game of Example 2.3 can be interpreted as a game where two firms choose their output to maximize their profits, with the parameter w_{12} deciding whether their products are net substitutes or net complements. An interaction matrix similar to Example 2.4's may describe the relationship between two competitor firms producing substitute goods, e.g. plane manufacturers Airbus and Boeing, and a third firm producing a complementary good, e.g. a kerosene supplier Exxon Mobil. Other such examples include gaming consoles, Xbox and Playstation with a game developer EA Sports, or tea companies Lipton and Twinings with a sugar company Südzucker. Weighted network games provide a framework to model any type of relationship structure with any number of companies.

We denote the set of weighted network games satisfying Assumption 2.2 by $\mathcal{G}$. Weighted network games in $\mathcal{G}$ are described by (i) the vector of upper bounds $\bar{x} \in \mathbb{R}_{++}^n$ characterizing the players' action spaces, (ii) a weighted network, which is a matrix $W = (w_{ij})_{i,j \in I}$ capturing the nature of the interactions between the players, and (iii) a vector of targets $t \in \mathbb{R}^n$ describing levels where marginal benefits are equal to marginal costs of the players. Since $w_{ii} = 1$ for every $i \in I$ and $w_{ij} = w_{ji}$ for every $i, j \in I$, the number of free parameters in W is $n(n-1)/2$. Let $w \in \mathbb{R}^{n(n-1)/2}$ denote the column vector of the upper triangular elements of W . We define the set of parameters $P = \mathbb{R}_{++}^n \times \mathbb{R}^{n(n-1)/2} \times \mathbb{R}^n$. Then, for $(\bar{x}, w, t) \in P$, let $\mathcal{G}(\bar{x}, w, t)$ be the set of weighted network games in $\mathcal{G}$ with upper bounds $\bar{x}$, interaction parameters w , and targets t . A weighted network game in such a set is characterized by the benefit functions f_i and cost parameters c_i .

For the benefit functions f_i , the properties in Assumption 2.2 imply that for every $x_{-i} \in X_{-i}$, $\pi_i(x_i, x_{-i})$ has a unique global maximizer in X_i . For player $i \in I$, let $b_i: X \rightarrow X_i$ denote his

best-response function, i.e. $b_i(x) = \operatorname{argmax}_{x_i \in X_i} \pi_i(x_i, x_{-i})$ for every $x \in X$. We now show that for a fixed configuration $(\bar{x}, w, t) \in P$, all games in $\mathcal{G}(\bar{x}, w, t)$ are best-response equivalent, which implies that all games in $\mathcal{G}(\bar{x}, w, t)$ have the same set of Nash equilibria.

Lemma 2.5. *Let $(\bar{x}, w, t) \in P$ and let $G \in \mathcal{G}(\bar{x}, w, t)$ be a weighted network game. For every $i \in I$ and $x \in X$ it holds that*

$$b_i(x) = \begin{cases} 0 & \text{if } t_i - \sum_{j \neq i} w_{ij} x_j < 0, \\ t_i - \sum_{j \neq i} w_{ij} x_j & \text{if } t_i - \sum_{j \neq i} w_{ij} x_j \in [0, \bar{x}_i], \\ \bar{x}_i & \text{if } t_i - \sum_{j \neq i} w_{ij} x_j > \bar{x}_i. \end{cases} \quad (1)$$

Proof. By differentiation of the payoff function we get

$$\frac{\partial \pi_i(x_i, x_{-i})}{\partial x_i} = f'_i\left(\sum_{j \in I} w_{ij} x_j\right) - c_i.$$

The first order condition of unconstrained maximization is satisfied if $f'_i(\sum_{j \in I} w_{ij} x_j) - c_i = 0$. Using property (3) of f_i this is satisfied if $\sum_{j \in I} w_{ij} x_j = t_i$.

If $t_i - \sum_{j \neq i} w_{ij} x_j \in [0, \bar{x}_i]$, then it follows that $b_i(x) = t_i - \sum_{j \neq i} w_{ij} x_j$. Note that the second order condition of maximization is satisfied due to the concavity of f_i , and, therefore, of π_i .

If $t_i - \sum_{j \neq i} w_{ij} x_j < 0$, then for every $x_i \in X_i$ it holds that $t_i < \sum_{j \in I} w_{ij} x_j$. Invoking properties (2) and (3) of f_i , for every $x_i \in X_i$ we have $f'_i(\sum_{j \in I} w_{ij} x_j) < c_i$, meaning that $\partial \pi_i(x_i, x_{-i}) / \partial x_i$ is uniformly negative. Therefore, π_i is maximized for the lowest possible value of x_i , so $b_i(x) = 0$.

Similarly, if $t_i - \sum_{j \neq i} w_{ij} x_j > \bar{x}_i$, then for every $x_i \in X_i$ it holds that $t_i > \sum_{j \in I} w_{ij} x_j$. Properties (2) and (3) of f_i guarantee that for every $x_i \in X_i$ we have $f'_i(\sum_{j \in I} w_{ij} x_j) > c_i$, and that $\partial \pi_i(x_i, x_{-i}) / \partial x_i$ is uniformly positive. Therefore, π_i is maximized for the highest possible value of x_i , so $b_i(x) = \bar{x}_i$. This concludes the proof. $\blacksquare$

It is useful to define a player's unconstrained best response, the contribution level they would choose if instead of $[0, \bar{x}_i]$, their set of available actions were equal to $\mathbb{R}$. For player $i \in I$ and action profile $x \in X$, let $\hat{b}_i(x) = t_i - \sum_{j \neq i} w_{ij} x_j$ denote this value. Clearly, $\hat{b}_i(x) \neq b_i(x)$ implies that $t_i - \sum_{j \neq i} w_{ij} x_j \notin [0, \bar{x}_i]$ and therefore the actual best response is on the boundary: $b_i(x) \in \{0, \bar{x}_i\}$.

For $i \in I$, we define the numbers $\underline{b}_i$ and $\bar{b}_i$ by $\underline{b}_i = \min_{x_{-i} \in X_{-i}} (t_i - \sum_{j \neq i} w_{ij} x_j)$ and $\bar{b}_i = \max_{x_{-i} \in X_{-i}} (t_i - \sum_{j \neq i} w_{ij} x_j)$. Since the set X_{-i} is compact, both $\underline{b}_i$ and $\bar{b}_i$ are well-defined. It is easily seen that the unconstrained best response of player i always belongs to the interval $[\underline{b}_i, \bar{b}_i]$.

Lemma 2.5 shows that for every player $i \in I$ and every action profile $x \in X$ such that $b_i(x) \in (0, \bar{x}_i)$, ceteris paribus changing a player j 's action by Δx_j changes player i 's best response by $-w_{ij}\Delta x_j$.

3 Nash equilibria

Since the paper's main focus is on the convergence of adaptive learning processes to the set of Nash equilibria, as a precursor we characterize the relevant properties of this set.

We first show that a weighted network game from Definition 2.1 satisfying Assumption 2.2 is a best-response potential game (Voorneveld, 2000). A game with set of players I , action space X , and payoff functions $(\pi_i)_{i \in I}$ is a best-response potential game if there exists a function $\phi: X \rightarrow \mathbb{R}$ such that for every $i \in I$ and every $x \in X$ it holds that

$$\operatorname{argmax}_{x_i \in X_i} \pi_i(x_i, x_{-i}) = \operatorname{argmax}_{x_i \in X_i} \phi(x_i, x_{-i}). \quad (2)$$

We call ϕ the best-response potential of game (I, X, π) .

Proposition 3.1. *Every weighted network game $G \in \mathcal{G}(\bar{x}, w, t)$ with $(\bar{x}, w, t) \in P$ is a best-response potential game with potential function $\phi: X \rightarrow \mathbb{R}$ defined by*

$$\phi(x) = x^\top t - \frac{1}{2} x^\top W x, \quad x \in X.$$

Proof. We show that for every $x \in X$ and every $i \in I$ it holds that

$$\operatorname{argmax}_{x_i \in X_i} \pi_i(x_i, x_{-i}) = \operatorname{argmax}_{x_i \in X_i} \phi(x_i, x_{-i}).$$

The left-hand side of the equality above equals $b_i(x)$. For the right-hand side, notice that $x^\top W$ is the row vector of consumption levels of each player, $x^\top W = (\sum_{j \in I} w_{1j} x_j, \dots, \sum_{j \in I} w_{nj} x_j)$, using the fact that W is symmetric. Multiplication by x gives

$$x^\top W x = x_1 \sum_{j \in I} w_{1j} x_j + \dots + x_n \sum_{j \in I} w_{nj} x_j.$$

Differentiating $\phi(x_i, x_{-i})$ by x_i leads to

$$\frac{\partial \phi(x_i, x_{-i})}{\partial x_i} = t_i - \frac{1}{2} \left(2x_i + \sum_{j \neq i} w_{ij} x_j + \sum_{j \neq i} w_{ji} x_j \right) = t_i - \sum_{i \in I} w_{ij} x_j,$$

where the last equality uses the symmetry of W .

Setting the derivative of ϕ with respect to x_i equal to zero gives the extreme point $x_i = t_i - \sum_{j \neq i} w_{ij} x_j$, and as long as $t_i - \sum_{j \neq i} w_{ij} x_j \in [0, \bar{x}_i]$, it is the unique maximum, since the second derivative is $-w_{ii} = -1$.

If $t_i - \sum_{j \neq i} w_{ij} x_j < 0$, then, since $x_i \geq 0$, the first derivative of ϕ with respect to x_i is uniformly negative on $[0, \bar{x}_i]$ hence the unique maximum is achieved for the minimal contribution, $x_i = 0$.

Similarly, if $t_i - \sum_{j \neq i} w_{ij} x_j > \bar{x}_i$, then the first derivative is uniformly positive, meaning that, in $[0, \bar{x}_i]$, the unique maximum is achieved for the maximal contribution, $x_i = \bar{x}_i$. ■

For games with continuous action sets, the existence of a best-response potential is of particular interest. Jensen (2010) and Ewerhart (2017) study this question and its implications in the class of aggregative games, and contest games, respectively. It is a pivotal step in our analysis as well, for the best-response potential guarantees that the set of Nash equilibria is non-empty and has strong implications on the convergence properties of best-response dynamics.

In addition, for network games on unweighted graphs, Bervoets and Faure (2016) show the following property:

$$\text{sgn} \left(\frac{\partial \pi_i(x_i, x_{-i})}{\partial x_i} \right) = \text{sgn} \left(\frac{\partial \phi(x_i, x_{-i})}{\partial x_i} \right).$$

This property can be easily generalized for weighted network games. Games that satisfy this property are called locally ordinal potential games. In Section 4 we show that weighted network games do not belong to the class of ordinal potential games (Monderer and Shapley, 1996) as they admit better-response cycles (Example 4.5).

We now turn to the characterization of Nash equilibria. We first show existence. For $(\bar{x}, w, t) \in P$, let $X^*(\bar{x}, w, t)$ denote the set of Nash equilibria of a game in $\mathcal{G}(\bar{x}, w, t)$.

Proposition 3.2. *For every $(\bar{x}, w, t) \in P$, it holds that $X^*(\bar{x}, w, t) \neq \emptyset$.*

Proof. Let $b : X \rightarrow X$ be the function such that its component $i \in I$ is equal to b_i , the best-response function of player i . Since X is non-empty, compact, and convex, and b is continuous, the existence of x^* such that $b(x^*) = x^*$ is guaranteed by Brouwer's fixed-point theorem. ■

Proposition 3.2 states that each weighted network game has a Nash equilibrium. Note that this does not hold in case the action space is unbounded, as negative interaction parameters may cause an infinite increase of best replies. This is illustrated in the following example.

Example 3.3. Consider the 2-player weighted network game of Example 2.3 with $w_{12} = -1$. Then, the unconstrained best responses are

$$\hat{b}_1(x) = x_2 + t_1$$

and

$$\widehat{b}_2(x) = x_1 + t_2.$$

In this example, the only Nash equilibrium is $x^* = \bar{x}$. If the strategy sets were unbounded, no Nash equilibrium would exist.

Since the games we study are best-response potential games (Proposition 3.1), we can characterize the set of Nash equilibria as the solution set of the constrained optimization of the best-response potential.

Proposition 3.4. *Let some $(\bar{x}, w, t) \in P$ be given. It holds that $x^* \in X^*(\bar{x}, w, t)$ if and only if x^* satisfies the Karush-Kuhn-Tucker (KKT) conditions, i.e. for every $i \in I$, there exist $\lambda_i, \mu_i \in \mathbb{R}_+$ such that*

$$\begin{aligned} t_i - \sum_{j \in I} w_{ij} x_j^* + \lambda_i - \mu_i &= 0, \\ x_i^* &\geq 0, \quad \bar{x}_i \geq x_i^*, \\ \lambda_i x_i^* &= 0, \quad \mu_i (\bar{x}_i - x_i^*) = 0. \end{aligned}$$

Proof. Since every game in $\mathcal{G}(\bar{x}, w, t)$ is a best-response potential game with potential function ϕ , every Nash equilibrium satisfies the stated KKT conditions.

Since, for every $i \in I$, for every $x \in X$, it holds that

$$\frac{\partial^2 \phi(x)}{\partial x_i^2} = -1 < 0,$$

every point satisfying the KKT conditions yields a Nash equilibrium. ■

Our version of this result aligns with Lemma 1 of Bramoullé et al. (2014), formulated for networks with all weights equal to 0 or 1. In section VI they mention the generalization for weighted networks and refer to the best-response potential as the exact potential of the linear-quadratic model of Ballester et al. (2006).

In what follows we derive conditions for the interaction matrix and the vector of targets that guarantee the equilibrium set to be finite. For $(\bar{x}, w, t) \in P$, let $\Xi(\bar{x}, w, t)$ denote the set of solutions (x, λ, μ) to the KKT conditions of Proposition 3.4. For $H \subseteq I$, let

$$\Xi_H(\bar{x}, w, t) = \{(x, \lambda, \mu) \in \Xi(\bar{x}, w, t) : \forall i \in H, \lambda_i = \mu_i = 0, \text{ and } \forall i \in I \setminus H, \max\{\lambda_i, \mu_i\} > 0\}.$$

In words, $\Xi_H(\bar{x}, w, t)$ denotes the set of solutions to the KKT conditions of Proposition 3.4 such that for every player in H neither complementarity condition is binding and for every player outside H exactly one complementarity condition is binding. Note that both complementarity conditions cannot be binding simultaneously.

The set $X_H^*(\bar{x}, w, t)$ is obtained by taking the projection of $\Xi_H(\bar{x}, w, t)$ to the set of action profiles X , $X_H^*(\bar{x}, w, t) = \text{proj}_X \Xi_H(\bar{x}, w, t)$, where proj_X is the projection mapping to X . If $x^* \in X_H^*(\bar{x}, w, t)$, then $x^* \in X^*(\bar{x}, w, t)$ and for every $i \in I \setminus H$ we have $x_i^* \in \{0, \bar{x}_i\}$. It follows that $X^*(\bar{x}, w, t) = \bigcup_{H \subseteq I} X_H^*(\bar{x}, w, t)$. Then, clearly, the set of Nash equilibria $X^*(\bar{x}, w, t)$ is finite if and only if for every $H \subseteq I$, $X_H^*(\bar{x}, w, t)$ is finite.

Let bounds $\bar{x} \in \mathbb{R}_{++}^n$, interaction parameters $w \in \mathbb{R}^{n(n-1)}$, and a set $H \subseteq I$ be given. The set of target vectors for which $X_H^*(\bar{x}, w, t)$ is infinite is denoted by

$$T_H = \{t \in \mathbb{R}^n : |X_H^*(\bar{x}, w, t)| = \infty\}.$$

The set $\bar{T}_H$ denotes its closure. Further, let

$$T = \{t \in \mathbb{R}^n : |X^*(\bar{x}, w, t)| = \infty\}$$

denote the set of target vectors that yield infinitely many Nash equilibria and let $\bar{T}$ denote the closure of T .

Lemma 3.5. *For every $\bar{x} \in \mathbb{R}_{++}^n$, for every $w \in \mathbb{R}^{n(n-1)/2}$, for every $H \subseteq I$, the set $\bar{T}_H$ has Lebesgue measure zero.*

Proof. Let some $\bar{x} \in \mathbb{R}_{++}^n$, some $w \in \mathbb{R}^{n(n-1)/2}$, and some $H \subseteq I$ be given.

First consider the case $H = \emptyset$. Then, for every $t \in \mathbb{R}^n$ it holds that $X_\emptyset^*(\bar{x}, w, t) \subseteq \prod_{i \in I} \{0, \bar{x}_i\}$, meaning that for every $t \in \mathbb{R}^n$ we have $|X_\emptyset^*(\bar{x}, w, t)| < \infty$. It follows that $T_\emptyset = \bar{T}_\emptyset = \emptyset$.

Now consider the case $H \neq \emptyset$. We show that there exists a set $U_H \subset \mathbb{R}^n$ of Lebesgue measure zero such that $\bar{T}_H \subseteq U_H$.

For every $t \in \mathbb{R}^n$, for every $x^* \in X_H^*(\bar{x}, w, t)$, we have that

$$\begin{aligned} x_i^* &\in \{0, \bar{x}_i\}, & i &\in I \setminus H, \\ t_i - \sum_{j \in H} w_{ij} x_j^* - \sum_{j \in I \setminus H} w_{ij} x_j^* &= 0, & i &\in H. \end{aligned}$$

Let $W_H = (w_{ij})_{i,j \in H}$ denote the submatrix of W that we obtain by removing every row and every column whose index is not contained in H . Further, let $W_{H,-H} = (w_{ij})_{i \in H, j \in I \setminus H}$, $t_H = (t_i)_{i \in H}$, $x_H^* = (x_i^*)_{i \in H}$, and $x_{-H}^* = (x_i^*)_{i \in I \setminus H}$. Now, the previous system of equations can be written in matrix form as

$$W_H x_H^* = t_H - W_{H,-H} x_{-H}^*. \quad (3)$$

Therefore, by the Rouché-Capelli theorem, $|X_H^*(\bar{x}, w, t)| = \infty$ implies $\text{rank}(W_H) < |H|$. So $T_H = \bar{T}_H = \emptyset$ whenever $\text{rank}(W_H) = |H|$. Consider the case where $\text{rank}(W_H) < |H|$. For $y \in \prod_{i \in I \setminus H} \{0, \bar{x}_i\}$, let U_H^y be the set of target vectors t such that $t_H - W_{H,-H} y$ belongs to the

span of W_H . Notice that U_H^y is an $(n - |H| + \text{rank}(W_H))$ -dimensional vector space and therefore a closed set of Lebesgue measure zero. Let $U_H = \bigcup_{y \in \prod_{i \in I} \{0, \bar{x}_i\}} U_H^y$. Since U_H is a union of finitely many closed sets of Lebesgue measure zero, it is also closed and is of Lebesgue measure zero. Notice that $t \in \mathbb{R}^n \setminus U_H$ implies that $t \in \mathbb{R}^n \setminus T_H$, since for every $t \in \mathbb{R}^n \setminus U_H$ the system

$$W_H x_H^* = t_H - W_{H,-H} x_{-H}^*$$

has no solutions in x_H^* . It follows that $T_H \subseteq U_H$. Furthermore, since U_H is closed, we also have $\bar{T}_H \subseteq U_H$. ■

We prove Lemma 3.5 for the closure of the set T_H , which implies that the set of target vectors with infinitely many Nash equilibria is not only small in a measure theoretic sense, but also in a topological sense.

The intuition behind Lemma 3.5 is that for a fixed subset of players, the set of interior Nash equilibria corresponds to the solution set of a linear system which, generically, has only one solution. The case of infinitely many solutions, and hence, the possibility of infinitely many Nash equilibria obtains only if the rank of the interaction matrix is not full and the target vector belongs to a vector space parallel to the span of the interaction matrix, which is only the case for a set of target vectors of Lebesgue measure zero.

As an illustration of Lemma 3.5 for the case where all targets are very large or very small, i.e. for every $i \in I$ we have $t_i > \max_{x \in X} \sum_{j \in I} w_{ij} x_j$ or $t_i < \min_{x \in X} \sum_{j \in I} w_{ij} x_j$, the set of Nash equilibria is a subset of the corners of the strategy space X , i.e. $X^*(\bar{x}, w, t) \subseteq \prod_{i \in I} \{0, \bar{x}_i\}$, and is therefore finite.

Lemma 3.6. *For every $\bar{x} \in \mathbb{R}_{++}^n$, for every $w \in \mathbb{R}^{n(n-1)/2}$, the set $\bar{T}$ has Lebesgue measure zero.*

Proof. We show that there exists a set $U \subseteq \mathbb{R}^n$ of Lebesgue measure zero such that $\bar{T} \subseteq U$.

Let $U = \bigcup_{H \subseteq I} U_H$. Since U is a union of finitely many sets of Lebesgue measure zero, it has Lebesgue measure zero. Since $T = \bigcup_{H \subseteq I} T_H$, and $T_H \subseteq U_H$ for every $H \subseteq I$, it also holds that $T \subseteq U$. Once again, since U is closed, we have $\bar{T} \subseteq U$. ■

Corollary 3.7. *For every $\bar{x} \in \mathbb{R}_{++}^n$, for every $w \in \mathbb{R}^{n(n-1)/2}$, for almost every $t \in \mathbb{R}^n$, the weighted network game $G \in \mathcal{G}(\bar{x}, w, t)$ has a finite number of Nash equilibria.*

The generic finiteness of the set of Nash equilibria is illustrated in the following example.

Example 3.8. Fix parameters $\bar{x}$ and w_{12} in a weighted network game with two players. We have shown in Corollary 3.7 that for almost every $t \in \mathbb{R}^2$ the set $X^*(\bar{x}, w_{12}, t)$ is finite. For every

$t \in \mathbb{R}^2$, the set $X_\emptyset^*(\bar{x}, w_{12}, t)$ is trivially finite, and it is easy to see that the sets $X_{\{1\}}^*(\bar{x}, w_{12}, t)$ and $X_{\{2\}}^*(\bar{x}, w_{12}, t)$ are finite in the case of two players.

We therefore only check the interior solutions to the KKT problem of this game as defined in Proposition 3.4, i.e. where all Lagrange parameters λ_i, μ_i are zero. In an interior solution $x^* \in X_{\{1,2\}}^*(\bar{x}, w_{12}, t)$, we have $\widehat{b}_1(x^*) = x_1^*$ and $\widehat{b}_2(x^*) = x_2^*$, therefore

$$\begin{aligned} x_1^* &= t_1 - w_{12}x_2^*, \\ x_2^* &= t_2 - w_{12}x_1^*. \end{aligned}$$

In case w_{12} is not equal to 1 or -1 , rearranging yields

$$\begin{aligned} x_1^* &= \frac{t_1 - w_{12}t_2}{1 - (w_{12})^2}, \\ x_2^* &= \frac{t_2 - w_{12}t_1}{1 - (w_{12})^2}. \end{aligned}$$

Therefore, for every $w_{12} \in \mathbb{R} \setminus \{-1, 1\}$, we have $|X_{\{1,2\}}^*(\bar{x}, w_{12}, t)| \leq 1$. Whether or not the set of interior equilibria is empty depends on whether x^* is an element of X .

If $w_{12} = -1$, it is easy to check that $t_1 = t_2 = 0$ yields infinitely many Nash equilibria x^* with $x_1^* = x_2^*$. There can also be infinitely many Nash equilibria when $t_1 + t_2 = 0$. There are no interior Nash equilibria for different values of t as then the system of best responses is inconsistent. Similarly, if $w_{12} = 1$ then there can only be infinitely many Nash equilibria if $t_1 = t_2$. Indeed, if $\bar{x}_1 + \bar{x}_2 > t_1$, then there are infinitely many Nash equilibria x^* with $x_1^* + x_2^* = t_1$, if $\bar{x}_1 + \bar{x}_2 = t_1$, then there is a unique interior Nash equilibrium, and if $\bar{x}_1 + \bar{x}_2 < t_1$, then there are no interior Nash equilibria.

Our result makes use of the generic uniqueness of interior equilibria, but the network structure allows the existence of a finite number of corner equilibria. Ballester and Calvó-Armengol (2010), Belhaj et al. (2014) and Allouch (2015) provide results for the uniqueness of Nash equilibrium.

We conclude this section by discussing efficiency properties of the equilibria. We consider efficiency in the Pareto sense. Other models consider efficiency in terms of minimizing total efforts/production (Bramoullé and Kranton, 2007; Goyal, 2012), or maximizing total welfare Helsley (2014). We first show that a pair of players with a non-zero interaction parameter can always jointly deviate from an interior equilibrium to a better action profile. For a subset of players $H \subseteq I$ and $\delta \in \mathbb{R}$, let $\delta^H \in \mathbb{R}^n$ denote the vector such that $\delta_i^H = \delta$ for $i \in H$ and $\delta_i^H = 0$ for $i \in I \setminus H$.

Proposition 3.9. *Let $(\bar{x}, w, t) \in P$ be such that for some $i, j \in I$ with $i \neq j$ it holds that $w_{ij} \neq 0$. Let $x^* \in X^*(\bar{x}, w, t)$ be a Nash equilibrium of a game $G = (I, X, \pi) \in \mathcal{G}(\bar{x}, w, t)$ such that $x_i^* \in (0, \bar{x}_i)$ and $x_j^* \in (0, \bar{x}_j)$. Then there exists $\delta \in (0, \min\{x_i^*, x_j^*, \bar{x}_i - x_i^*, \bar{x}_j - x_j^*\})$ such that $\pi_i(x^*) < \pi_i(x^* + \text{sgn}(w_{ij})\delta^{i,j})$ and $\pi_j(x^*) < \pi_j(x^* + \text{sgn}(w_{ij})\delta^{i,j})$.*

Proof. We first discuss the case $w_{ij} > 0$.

Since $x_i^* \in (0, \bar{x}_i)$ and $x_j^* \in (0, \bar{x}_j)$, it holds that $\sum_{k \in I} w_{ik} x_k = t_i$ and $\sum_{k \in I} w_{jk} x_k = t_j$. Hence, for $\delta > 0$ such that $x_i^* + \delta \leq \bar{x}_i$ and $x_j^* + \delta \leq \bar{x}_j$, we have

$$\frac{\pi_i(x^* + \delta^{i,j}) - \pi_i(x^*)}{\delta} = \frac{f_i(t_i + \delta(1 + w_{ij})) - f_i(t_i)}{\delta} - c_i.$$

Since f_i is concave, we have

$$f_i(t_i + \delta) \leq f_i(t_i + \delta(1 + w_{ij})) - \delta w_{ij} f'_i(t_i + \delta(1 + w_{ij})).$$

Therefore, we can write

$$\frac{\pi_i(x^* + \delta^{i,j}) - \pi_i(x^*)}{\delta} \geq \frac{f_i(t_i + \delta) - f_i(t_i)}{\delta} - c_i + w_{ij} f'_i(t_i + \delta(1 + w_{ij})).$$

Let $\varepsilon_i = \min_{x \in X} f'_i(\sum_{k \in I} w_{ik} x_k)$. Since f'_i is a continuous function, its minimum over the compact set X is well-defined. Notice that Assumption 2.2 guarantees that $\varepsilon_i > 0$. Thus,

$$\frac{\pi_i(x^* + \delta^{i,j}) - \pi_i(x^*)}{\delta} \geq \frac{f_i(t_i + \delta) - f_i(t_i)}{\delta} - c_i + w_{ij} \varepsilon_i.$$

Also due to the continuity of f'_i , the term $(f_i(t_i + \delta) - f_i(t_i))/\delta - c_i$ converges to zero as δ goes to zero. Hence, for sufficiently small positive δ , we have $\pi_i(x^* + \delta^{i,j}) - \pi_i(x^*) > 0$. The same argument applies to agent j .

The case $w_{ij} < 0$ follows from very similar arguments. ■

Proposition 3.9 implies that interior Nash equilibria are not strong Nash equilibria since there are profitable deviations by coalitions of two linked players. It is then easy to derive the next proposition, stating that interior Nash equilibria are not Pareto efficient provided that the interaction parameters are either all non-negative or all non-positive. See Elliott and Golub (2013) for a characterization of efficient Nash equilibria in the non-negative case.

Proposition 3.10. *Let $(\bar{x}, w, t) \in P$ be such that $w \geq 0$ or $w \leq 0$ and, for some $i, j \in I$ with $i \neq j$, it holds that $w_{ij} \neq 0$. Let $x^* \in X^*(\bar{x}, w, t)$ be a Nash equilibrium of a game $G = (I, X, \pi) \in \mathcal{G}(\bar{x}, w, t)$ such that $x_i^* \in (0, \bar{x}_i)$ and $x_j^* \in (0, \bar{x}_j)$. Then there exists $\delta \in (0, \min\{x_i^*, x_j^*, \bar{x}_i - x_i^*, \bar{x}_j - x_j^*\})$ such that the action profile $x^* + \text{sgn}(w_{ij})\delta^{i,j}$ is a Pareto improvement over x^* .*

Proof. We first consider the case where $w_{ij} > 0$. As per Proposition 3.9, there exists $\delta \in (0, \min\{x_i^*, x_j^*, \bar{x}_i - x_i^*, \bar{x}_j - x_j^*\})$ such that $\pi_i(x^* + \delta^{i,j}) > \pi_i(x^*)$ and $\pi_j(x^* + \delta^{i,j}) > \pi_j(x^*)$. Since $w \geq 0$ it follows that for every other player $h \in I \setminus \{i, j\}$ we have $f_h(\sum_{k \in I} w_{hk} x_k^* + \delta^{i,j}) \geq f_h(\sum_{k \in I} w_{hk} x_k^*)$, while his own action did not change, and therefore $\pi_h(x^* + \delta^{i,j}) \geq \pi_h(x^*)$, meaning that players i and j increasing their action by δ yields a Pareto improvement.

The case $w \leq 0$ follows from similar arguments. ■

Together, Propositions 3.9 and 3.10 characterize some of the most important properties of interior equilibria: they are neither strong nor efficient.

4 Convergence of learning processes

Within the framework of weighted network games, we consider learning processes where players update their strategies sequentially. That is, given some initial action profile, one player changes his action, while that of every other player remains the same. Then, another player makes a change under similar circumstances, and so on. First, we are interested in the question whether action profiles chosen by best-responding and better-responding players may cycle. Generally, the non-existence of best-response cycles is a necessary but not sufficient condition of the convergence of best-response dynamics (Kukushkin, 2015) and therefore this is an important first step towards our convergence results.

Let $\mathbb{N} = \{1, 2, \dots\}$ denote the set of positive integers and let $\mathcal{K} = \{\{1, 2\}, \{1, 2, 3\}, \dots\} \cup \mathbb{N}$ be a collection of index sets. For $K \in \mathcal{K}$, we denote by K^- the set that results from K by leaving out its highest element. Notice that K^- is equal to K if $K = \mathbb{N}$.

Definition 4.1. Let some $G \in \mathcal{G}$ and $K \in \mathcal{K}$ be given. A sequence of action profiles $(x^k)_{k \in K}$ is a *path* in the game G if

1. for each $k \in K^-$ there exists a player i^k such that $x_{-i^k}^{k+1} = x_{-i^k}^k$,
2. there is at least one $k \in K^-$ such that $x^{k+1} \neq x^k$.

If $x_{-i^k}^{k+1} = x_{-i^k}^k$, and $x_{i^k}^{k+1} \neq x_{i^k}^k$, then we call i^k the updating player at period k .

As per Definition 4.1, a path is a sequence where at most one player has changed his contribution between any two successive action profiles, while there are at least two different action profiles in the sequence.

Definition 4.2. Let some $G \in \mathcal{G}$ and $K \in \mathcal{K}$ be given. A path $(x^k)_{k \in K}$ is *best-response compatible* in game G if for every $k \in K^-$ it holds that

1. if $x^{k+1} = x^k$, then there exists $i^k \in I$ such that $x_{i^k}^{k+1} = x_{i^k}^k = b_{i^k}(x^k)$,
2. if $x_{i^k}^{k+1} \neq x_{i^k}^k$, then $x_{i^k}^{k+1} = b_{i^k}(x^k)$.

Definition 4.3. Let some $G \in \mathcal{G}$ and $K \in \mathcal{K}$ be given. A path $(x^k)_{k \in K}$ is *better-response compatible* in game G if for every $k \in K^-$ it holds that

1. if $x^{k+1} = x^k$, then there exists $i^k \in I$ such that $x_{i^k}^{k+1} = x_{i^k}^k = b_{i^k}(x^k)$,
2. if $x_{i^k}^{k+1} \neq x_{i^k}^k$, then $\pi_{i^k}(x^{k+1}) > \pi_{i^k}(x^k)$.

Definitions 4.2 and 4.3 capture two of the simplest and best-known learning processes. In case of a best-response compatible path, each updating player moves to his best available option. In case of a better-response compatible path, updating players are only required to strictly improve their payoffs. Clearly, a best-response compatible path is also a better-response compatible path.

Definition 4.4. Let some $G \in \mathcal{G}$ and $K = \{1, \dots, m\} \in \mathcal{K}$ be given. A finite path $(x^k)_{k \in K}$ in the game G is a *cycle* if $x^1 = x^m$.

It is well known that best-response dynamics do not produce cycles in best-response potential games (Voorneveld, 2000), which includes weighted network games by Proposition 3.1. Better-response dynamics do not generate cycles in ordinal potential games (Monderer and Shapley, 1996). The following example shows that better-response cycles can occur within weighted network games.

Example 4.5. Let $I = \{1, 2\}$, $X_1 = X_2 = [0, 4]$, and $t_1 = t_2 = 1$. Moreover, let the payoff functions be given by

$$\pi_1(x_1, x_2) = 2\sqrt{x_1 + 0.6x_2} - x_1$$

and

$$\pi_2(x_1, x_2) = 2\sqrt{x_2 + 0.6x_1} - x_2.$$

It is easy to check that π_1 and π_2 satisfy the properties laid down in Definition 2.1 and Assumption 2.2 with $w_{12} = 0.6$, $f_1(z) = f_2(z) = 2\sqrt{z}$, and $c_1 = c_2 = 1$.

Table 1 presents a sequence of action profiles that constitutes a better-response cycle for this example.

k	x_1^k	x_2^k	$\pi_1(x^k)$	$\pi_2(x^k)$
1	0	0.1	0.49	0.53
2	3	0.1	0.50	2.66
3	3	0	0.46	2.68
4	0.1	0	0.53	0.49
5	0.1	3	2.66	0.50
6	0	3	2.68	0.46
7	0	0.1	0.49	0.53

Table 1: Actions played and payoffs in the better-response cycle in Example 4.5.

Note that the changes in player 1's choice of actions between periods 1 and 2 and between periods 3 and 4, as well as those for player 2 between periods 4 and 5 and between periods 6 and 7 are quite large, given the action space. Columns 5 and 6 of Table 2 present for each period the distance between the current action and both the best response and the action chosen of the player updating his action. Notice that in periods 1 and 4, the actions chosen are more than twice as far away from the current action than the best response is, meaning that the updating player, despite the increase in payoffs, has moved farther from his optimal decision than he originally was. We refer to this as extreme overshooting beyond the best response. Our main interest for the remainder of this section is to show how the extent of overshooting in a process affects its convergence properties.

k	x_1^k	x_2^k	$b_{i^k}(x^k)$	$ b_{i^k}(x^k) - x_{i^k}^k $	$ x_{i^k}^{k+1} - x_{i^k}^k $	α_k
1	0	0.1	0.94	0.94	3	-2.19
2	3	0.1	0	0.1	0.1	0
3	3	0	1	2	2.9	-0.45
4	0.1	0	0.94	0.94	3	-2.19
5	0.1	3	0	0.1	0.1	0
6	0	3	1	2	2.9	-0.45
7	0	0.1				

Table 2: The size of action changes in the better-response cycle of Table 1.

As before, for a path $(x^k)_{k \in K}$, let $(i^k)_{k \in K^-}$ denote the updating player in period k if there was a change in the action profile and let it denote any other player if there was not. Furthermore, for $k \in K^-$, let the overshooting coefficient $\alpha_k \in \mathbb{R} \cup \{-\infty, \infty\}$ be defined as

$$\alpha_k = \frac{x_{i^k}^{k+1} - b_{i^k}(x^k)}{x_{i^k}^k - b_{i^k}(x^k)},$$

where we take the convention that in case the denominator is 0, $\alpha_k = -\infty$ if the numerator is negative, $\alpha_k = 0$ if the numerator is 0, and $\alpha_k = +\infty$ if the numerator is positive. Column 7 of Table 2 shows the values of α_k in the better-response cycle of Example 4.5.

The coefficient α_k determines the extent of overshooting of the updating player beyond the best response. If there is overshooting, then α_k is negative. There is no overshooting if α_k is positive. If $\alpha_k = 0$ then the updating player moved to the best response. If $\alpha_k \in \{-\infty, \infty\}$ then $x_{i^k}^k = b_{i^k}(x^k)$, so the player moved away from a best response. If $\alpha_k < -1$, then the new action is farther from the best response relative to the action before the update.

Values of α_k in $(0, 1)$ correspond to a better response, while in case $\alpha_k > 1$ the payoff of the updating player is lower than before. For negative values of α_k , the threshold between better

and worse replies depends on the payoff function. Naturally, the possible values that α_k may take depend on $\bar{x}$.

As suggested by Example 4.5, sequences of action profiles that feature extreme overshooting beyond the best response may cycle. We therefore characterize sequences by their extent of overshooting.

Definition 4.6. A path $(x^k)_{k \in K}$ in a game $G \in \mathcal{G}$ is α -centered for some $\alpha > 0$ if for every $k \in K^-$ it holds that $|\alpha_k| < \alpha$.

A best-response compatible path is α -centered for every $\alpha > 0$. Furthermore, for every $\alpha > 0$ there exist paths that are better-response compatible and α -centered, but are not best-response compatible. For instance, it is easy to see that every sequence $(\alpha_k)_{k \in K^-}$ such that for every $k \in K^-$, $\alpha_k \in [0, \min\{\frac{\alpha}{2}, \frac{1}{2}\}]$, is both α -centered and better-response compatible. The restriction of being α -centered on a better-response dynamic captures a form of cautiousness by the players, as they do not engage in updates that take them very far from their optimal choice. For finite values of α , players do not change their action in an α -centered path if they are at their best response, as that would imply $|\alpha_k| = \infty$.

For the remainder of this paper we mainly consider α -centered paths with $\alpha \in (0, 1)$. In these paths, every updating player moves closer to his current best response.

We define the overshooting coefficient $\hat{\alpha}_k$ similar to α_k , replacing the best-response function b with the unconstrained best-response function $\hat{b}$. For a path of action profiles $(x^k)_{k \in K}$, we define

$$\hat{\alpha}_k = \frac{x_{i^k}^{k+1} - \hat{b}_{i^k}(x^k)}{x_{i^k}^k - \hat{b}_{i^k}(x^k)}, \quad k \in K^-.$$

The relationship between α_k and $\hat{\alpha}_k$ is summarized in the following lemma.

Lemma 4.7. Let $(x^k)_{k \in K}$ be a path of action profiles in a game $G \in \mathcal{G}$. The following statements hold for every $k \in K^-$:

- (i) $\alpha_k \neq \hat{\alpha}_k$ implies $b_{i^k}(x^k) \in \{0, \bar{x}_{i^k}\}$.
- (ii) $\alpha_k \in (0, 1)$ implies $\hat{\alpha}_k \in (0, 1)$.
- (iii) $\alpha_k \in (-1, 0)$ implies $\alpha_k = \hat{\alpha}_k$.
- (iv) $\alpha_k = 0$ implies $0 \leq \hat{\alpha}_k \leq 1$.
- (v) $\hat{\alpha}_k = 1$ implies $x^{k+1} = x^k$.

Proof. (i). If $b_{ik}(x^k) \in (0, \bar{x}_{ik})$, then it holds that $b_{ik}(x^k) = \widehat{b}_{ik}(x^k)$ and thus $\alpha_k = \widehat{\alpha}_k$.

(ii). We only need to consider the case $\alpha_k \neq \widehat{\alpha}_k$. By (i) we have $b_{ik}(x^k) \in \{0, \bar{x}_{ik}\}$. Take the case $b_{ik}(x^k) = 0$. Then it holds that $\widehat{b}_{ik}(x^k) < 0$, so $x_{ik}^{k+1} < x_{ik}^k$ due to $0 < \alpha_k < 1$, and thus

$$0 < \alpha_k = \frac{x_{ik}^{k+1}}{x_{ik}^k} < \frac{x_{ik}^{k+1} - \widehat{b}_{ik}(x^k)}{x_{ik}^k - \widehat{b}_{ik}(x^k)} = \widehat{\alpha}_k < 1.$$

The case $b_{ik}(x^k) = \bar{x}_{ik}$ follows from similar arguments.

(iii). Since $\alpha_k \in (-1, 0)$, we have

$$\text{sgn}(x_{ik}^k - b_{ik}(x^k)) = -\text{sgn}(x_{ik}^{k+1} - b_{ik}(x^k)) \neq 0.$$

Therefore, it must hold that $b_{ik}(x^k) \in (0, \bar{x}_{ik})$, otherwise x_{ik}^{k+1} would not be in X_{ik} . It follows that $b_{ik}(x^k) = \widehat{b}_{ik}(x^k)$.

(iv). Once again, we only need to discuss the case $\alpha_k \neq \widehat{\alpha}_k$, so $b_{ik}(x^k) \in \{0, \bar{x}_{ik}\}$. Consider the case $b_{ik}(x^k) = 0$. We have that $\widehat{b}_{ik}(x^k) < 0$ and $x_{ik}^{k+1} = 0$ since $\alpha_k = 0$. It holds that

$$0 = \alpha_k \leq \frac{-\widehat{b}_{ik}(x^k)}{x_{ik}^k - \widehat{b}_{ik}(x^k)} = \widehat{\alpha}_k \leq 1.$$

The case $b_{ik}(x^k) = \bar{x}_{ik}$ follows from similar arguments.

(v). In case $\widehat{\alpha}_k = 1$, we have

$$x_{ik}^{k+1} - \widehat{b}_{ik}(x^k) = x_{ik}^k - \widehat{b}_{ik}(x^k),$$

so $x_{ik}^{k+1} = x_{ik}^k$. ■

In the following proposition we show the relation between the value of $\widehat{\alpha}_k$ and changes in the value of the potential function as defined in Proposition 3.1. This relationship will prove crucial in our convergence analysis.

Proposition 4.8. *Let $(x^k)_{k \in K}$ be a path of action profiles in a game $G \in \mathcal{G}$ such that, for every $k \in K^-$, $\widehat{\alpha}_k \in \mathbb{R}$. Then it holds that*

$$\phi(x^{k+1}) - \phi(x^k) = \frac{1}{2}(1 - \widehat{\alpha}_k)(1 + \widehat{\alpha}_k)(\widehat{b}_{ik}(x^k) - x_{ik}^k)^2, \quad k \in K^-.$$

Proof. Using the definition of ϕ gives

$$\phi(x^{k+1}) - \phi(x^k) = \sum_{i \in I} (x_i^{k+1} - x_i^k) t_i - \frac{1}{2} \sum_{i \in I} x_i^{k+1} \left(\sum_{j \in I} w_{ij} x_j^{k+1} \right) + \frac{1}{2} \sum_{i \in I} x_i^k \left(\sum_{j \in I} w_{ij} x_j^k \right).$$

Using the symmetry of the interaction matrix W and taking advantage of the fact that x^{k+1} is the successor of x^k in a path, we substitute $x_{-i^k}^k = x_{-i^k}^{k+1}$ to get

$$\phi(x^{k+1}) - \phi(x^k) = (x_{i^k}^{k+1} - x_{i^k}^k)t_i - \frac{1}{2}((x_{i^k}^{k+1})^2 - (x_{i^k}^k)^2) - (x_{i^k}^{k+1} - x_{i^k}^k) \sum_{j \neq i^k} w_{i^k j} x_j^k.$$

Factoring out $x_{i^k}^{k+1} - x_{i^k}^k$ yields

$$\phi(x^{k+1}) - \phi(x^k) = (x_{i^k}^{k+1} - x_{i^k}^k) \left[(t_i - \frac{1}{2}(x_{i^k}^{k+1} + x_{i^k}^k)) - \sum_{j \neq i^k} w_{i^k j} x_j^k \right].$$

Substituting $\widehat{b}_{i^k}(x^k) = t_i - \sum_{j \neq i^k} w_{i^k j} x_j^k$ gives

$$\begin{aligned} \phi(x^{k+1}) - \phi(x^k) &= (x_{i^k}^{k+1} - x_{i^k}^k) (\widehat{b}_{i^k}(x^k) - \frac{1}{2}(x_{i^k}^{k+1} + x_{i^k}^k)) \\ &= \frac{1}{2}(x_{i^k}^{k+1} - x_{i^k}^k) [\widehat{b}_{i^k}(x^k) - x_{i^k}^{k+1} + \widehat{b}_{i^k}(x^k) - x_{i^k}^k]. \end{aligned}$$

Finally, substituting $x_{i^k}^{k+1} = (1 - \widehat{\alpha}_k) \widehat{b}_{i^k}(x^k) + \widehat{\alpha}_k x_{i^k}^k$ and $\widehat{b}_{i^k}(x^k) - x_{i^k}^{k+1} = \widehat{\alpha}_k (\widehat{b}_{i^k}(x^k) - x_{i^k}^k)$ gives

$$\phi(x^{k+1}) - \phi(x^k) = \frac{1}{2}(1 - \widehat{\alpha}_k)(1 + \widehat{\alpha}_k)(\widehat{b}_{i^k}(x^k) - x_{i^k}^k)^2.$$

■

Proposition 4.8 says that in a path of action profiles, the change of the potential is only determined by the magnitude of $\widehat{\alpha}_k$. Each time the updating player gets closer to his unconstrained best response by his update, the value of the potential function increases, and each time he gets further from the unconstrained best response, the value of the potential function decreases. This property can be exploited to assess the possibility of cycles in the set of α -centered paths.

Proposition 4.9. *A game $G \in \mathcal{G}$ has no 1-centered cycle.*

Proof. Suppose that $(x^k)_{k \in K}$ is a 1-centered cycle. For every $k \in K^-$ it holds by Lemma 4.7 that $\widehat{\alpha}_k \in (-1, 1]$ and therefore, by Proposition 4.8, that $\phi(x^{k+1}) - \phi(x^k) \geq 0$.

By Definition 4.1, each path has at least one pair of successive action profiles that are different. Let $k' \in K^-$ be such that $x^{k'+1} \neq x^{k'}$. Since the path $(x^k)_{k \in K}$ is 1-centered, Lemma 4.7 implies $|\widehat{\alpha}_{k'}| < 1$, and therefore, by Proposition 4.8, $\phi(x^{k'+1}) - \phi(x^{k'}) > 0$. Together with the fact that $\phi(x^{k+1}) - \phi(x^k) \geq 0$ for every $k \in K^-$, we obtain a contradiction to $(x^k)_{k \in K}$ being a 1-centered cycle. ■

By Proposition 4.9, if every update moves the updating player closer to his best response, then better-response cycles cannot exist. This implies the non-existence of best-response cycles. Notice that the cycle in Example 4.5 is not 1-centered, hence Proposition 4.9 is not applicable.

Furthermore, notice that for $\alpha > 1$, cycling is possible in an α -centered path. For example, if we have $i^1 = i^2$ and $\alpha_1 = \alpha_2 = -1$, then (x^1, x^2, x^3) constitutes a cycle. This means that $\alpha \leq 1$ is a necessary and sufficient condition for the non-existence of α -centered cycles.

Even in 1-centered cycles, a player may get farther away from his best response as a result of updates by other players, meaning that subsequent updates for any given player are not necessarily smaller in magnitude than previous ones. Nevertheless, we can show that the distance between consecutive elements of any α -centered path with $\alpha < 1$ converges to zero.

Proposition 4.10. *Let $(x^k)_{k \in \mathbb{N}}$ be an α -centered path in a game $G \in \mathcal{G}$ such that $\alpha < 1$. Then it holds that $\lim_{k \rightarrow \infty} \|x^{k+1} - x^k\|_2 = 0$.*

Proof. We use the fact that $x_{i^k}^{k+1} = (1 - \hat{\alpha}_k) \hat{b}_{i^k}(x^k) + \hat{\alpha}_k x_{i^k}^k$ to obtain

$$\|x^{k+1} - x^k\|_2^2 = (x_{i^k}^{k+1} - x_{i^k}^k)^2 = (1 - \hat{\alpha}_k)^2 (\hat{b}_{i^k}(x^k) - x_{i^k}^k)^2.$$

Applying Proposition 4.8 gives

$$\|x^{k+1} - x^k\|_2^2 = 2 \frac{1 - \hat{\alpha}_k}{1 + \hat{\alpha}_k} (\phi(x^{k+1}) - \phi(x^k)).$$

Since the path is α -centered with $\alpha < 1$, by Lemma 4.7 we have $-\alpha < \hat{\alpha}_k \leq 1$. It follows that

$$\|x^{k+1} - x^k\|_2^2 \leq 2 \frac{1 + \alpha}{1 - \alpha} (\phi(x^{k+1}) - \phi(x^k)). \quad (4)$$

By Proposition 4.8 we have that the sequence $(\phi(x^k))_{k \in \mathbb{N}}$ is monotonically increasing. Furthermore, since ϕ is continuous and the set X is compact, the sequence $(\phi(x^k))_{k \in \mathbb{N}}$ is also bounded, and hence it is convergent, so $\phi(x^{k+1}) - \phi(x^k) \rightarrow 0$ as $k \rightarrow \infty$. Since the right-hand side of (4) converges to zero, it follows that $\|x^{k+1} - x^k\|_2^2 \rightarrow 0$ as $k \rightarrow \infty$. This implies our result. $\blacksquare$

Proposition 4.10 follows from the monotonicity, and therefore, the convergence of the values of the potential function along an α -centered path, by applying Proposition 4.8 to translate differences in the value of the potential function to distances between action profiles.

With similar tools we can show that the distance between the current action and the best response to it approaches zero for an updating player.

Proposition 4.11. *Let $(x^k)_{k \in \mathbb{N}}$ be an α -centered path in a game $G \in \mathcal{G}$ such that $\alpha < 1$. Then it holds that $\lim_{k \rightarrow \infty} |b_{i^k}(x^k) - x_{i^k}^k| = 0$.*

Proof. Suppose it does not hold that $\lim_{k \rightarrow \infty} |b_{i^k}(x^k) - x_{i^k}^k| = 0$. Then the sequence $(i^k, x^k)_{k \in \mathbb{N}}$ has a converging subsequence (i^{k^ℓ}, x^{k^ℓ}) with limit (i, x) such that $|b_i(x) - x_i| = \varepsilon > 0$. We distinguish three cases: (a) $\hat{b}_i(x) \in (0, \bar{x}_i)$, (b) $\hat{b}_i(x) \leq 0$, and (c) $\hat{b}_i(x) \geq \bar{x}_i$.

Case (a). $\widehat{b}_i(x) \in (0, \bar{x}_i)$.

There is $\ell' \in \mathbb{N}$ such that, for every $\ell \geq \ell'$, $i^{k^\ell} = i$, $b_i(x^{k^\ell}) \in (0, \bar{x}_i)$, and $|b_i(x^{k^\ell}) - x_i^{k^\ell}| \geq \varepsilon/2$. It follows that $\widehat{\alpha}_{k^\ell} = \alpha_{k^\ell}$, so by Proposition 4.8,

$$\phi(x^{k^{\ell+1}}) - \phi(x^{k^\ell}) \geq \frac{1}{2}(1 - \alpha_{k^\ell})(1 + \alpha_{k^\ell})\frac{1}{4}\varepsilon^2 > \frac{1}{2}(1 - \alpha)(1 + \alpha)\frac{1}{4}\varepsilon^2, \quad \ell \geq \ell'. \quad (5)$$

By Proposition 4.8, we have that the sequence $(\phi(x^k))_{k \in \mathbb{N}}$ is monotonically increasing, so the subsequence $(\phi(x^{k^\ell}))_{\ell \in \mathbb{N}}$ is monotonically increasing, and by (5) it tends to infinity. This contradicts the fact that the continuous function ϕ has a maximum on the compact set X .

Case (b). $\widehat{b}_i(x) \leq 0$.

We have that $b_i(x) = 0$ and $x_i = \varepsilon$. There is $\ell' \in \mathbb{N}$ such that, for every $\ell \geq \ell'$, $i^{k^\ell} = i$, $b_i(x^{k^\ell}) \leq x_i^{k^\ell}$, and $|b_i(x^{k^\ell}) - x_i^{k^\ell}| \geq \varepsilon/2$. If $\widehat{b}_i(x^{k^\ell}) \geq 0$, then $\widehat{\alpha}_{k^\ell} = \alpha_{k^\ell}$. Otherwise, we have $\widehat{b}_i(x^{k^\ell}) < 0$, so $b_i(x^{k^\ell}) = 0$, and

$$0 \leq \widehat{\alpha}_{k^\ell} = \frac{x_i^{k^{\ell+1}} - \widehat{b}_i(x^{k^\ell})}{x_i^{k^\ell} - \widehat{b}_i(x^{k^\ell})} \leq \frac{x_i^{k^{\ell+1}} - \underline{b}_i}{x_i^{k^\ell} - \underline{b}_i} = \frac{\alpha_{k^\ell} x_i^{k^\ell} - \underline{b}_i}{x_i^{k^\ell} - \underline{b}_i} \leq \frac{\frac{1}{2}\alpha_{k^\ell}\varepsilon - \underline{b}_i}{\frac{1}{2}\varepsilon - \underline{b}_i} < \frac{\frac{1}{2}\alpha\varepsilon - \underline{b}_i}{\frac{1}{2}\varepsilon - \underline{b}_i}. \quad (6)$$

The right-hand side of (6), denoted by β , belongs to $(\alpha, 1)$, so it holds that

$$-\alpha < \widehat{\alpha}_{k^\ell} \leq \beta, \quad \ell \geq \ell'.$$

By Proposition 4.8, we have that

$$\phi(x^{k^{\ell+1}}) - \phi(x^{k^\ell}) \geq \frac{1}{2}(1 - \alpha_{k^\ell})(1 + \alpha_{k^\ell})\frac{1}{4}\varepsilon^2 > \frac{1}{2}(1 - \beta)(1 + \beta)\frac{1}{4}\varepsilon^2, \quad \ell \geq \ell'. \quad (7)$$

By Proposition 4.8, we have that the sequence $(\phi(x^k))_{k \in \mathbb{N}}$ is monotonically increasing, so the subsequence $(\phi(x^{k^\ell}))_{\ell \in \mathbb{N}}$ is monotonically increasing, and by (7) it tends to infinity. This contradicts the fact that the continuous function ϕ has a maximum on the compact set X .

Case (c). $\widehat{b}_i(x) \geq \bar{x}_i$.

We can derive a contradiction along similar lines as in Case (b).

Since all three cases lead to a contradiction, we conclude that $\lim_{k \rightarrow \infty} |b_{i^k}(x^k) - x_{i^k}^k| = 0$. ■

Proposition 4.11 shows convergence to the best response for all updating players. In order to achieve convergence to a Nash equilibrium, we need convergence to the best response for all players. This can only be achieved if all players update regularly, otherwise nothing guarantees convergence for a player who, for instance, never updates. We therefore define the notion of updating in every ℓ periods, which is going to be the final condition for our main result.

Definition 4.12. Player $i \in I$ updates in every ℓ periods in a path of action profiles $(x^k)_{k \in \mathbb{N}}$ in a game $G \in \mathcal{G}$ if for every $k \in \mathbb{N}$ there exists $k' \in \{k, \dots, k + \ell - 1\}$ such that either $[x_i^{k'} \neq x_i^{k'+1}]$ or $[x^{k'} = x^{k'+1} \text{ and } x_i^{k'} = x_i^{k'+1} = b_i(x^{k'})]$.

A player satisfies Definition 4.12 if in every length ℓ segment of the path there is an action profile at which he updated or there exists a pair of successive action profiles that are identical and the player is at his best response.

We state a technical lemma.

Lemma 4.13. *Let $(x^k)_{k \in \mathbb{N}}$ be a path of action profiles in a game $G \in \mathcal{G}$ such that $\lim_{k \rightarrow \infty} \|x^{k+1} - x^k\|_2 = 0$. For every $\varepsilon > 0$, for every $\ell \in \mathbb{N}$, there exists $M \in \mathbb{N}$ such that, for every $m > M$, for every $k \in \{m, \dots, m + \ell - 1\}$, we have $\|x^k - x^m\|_2 < \varepsilon$.*

Proof. Let some $\varepsilon > 0$ and some $\ell \in \mathbb{N}$ be given. The statement obviously holds for $\ell = 1$, so consider the case $\ell > 1$.

For every $\delta > 0$ there exists $M_\delta \in \mathbb{N}$ such that for every $m > M_\delta$ we have $\|x^{m+1} - x^m\|_2 < \delta$, since $\lim_{k \rightarrow \infty} \|x^{k+1} - x^k\|_2 = 0$. We take $\delta = \varepsilon/(\ell - 1)$ and consider an arbitrary $m > M_\delta$.

Then, by the triangle inequality, for every $k \in \{m, \dots, m + \ell - 1\}$ we can write

$$\|x^k - x^m\|_2 \leq \|x^k - x^{k-1}\|_2 + \dots + \|x^{m+1} - x^m\|_2 < \frac{k-m}{\ell-1} \varepsilon \leq \varepsilon.$$

Therefore, M_δ is a suitable candidate for M . ■

We are ready to present our main results. First we show that an α -centered path with $\alpha < 1$ in which every player updates regularly gets arbitrarily close to the set of Nash equilibria.

Theorem 4.14. *Let $(\bar{x}, w, t) \in P$ and let $(x^k)_{k \in \mathbb{N}}$ be an α -centered path in a game $G \in \mathcal{G}(\bar{x}, w, t)$ such that $\alpha < 1$. If there is $\ell \in \mathbb{N}$ such that every player updates in every ℓ periods, then every cluster point of $(x^k)_{k \in \mathbb{N}}$ belongs to $X^*(\bar{x}, w, t)$.*

Proof. Since every linear function is Lipschitz continuous, the function $b_i: X \rightarrow X_i$ is Lipschitz continuous for every $i \in I$. Denote the Lipschitz constant of b_i by L_i .

Let x be a cluster point of $(x^k)_{k \in \mathbb{N}}$. We prove the result by showing that, for every $i \in I$, for every $\varepsilon > 0$, $|b_i(x) - x_i| < \varepsilon$.

Let $i \in I$ and $\varepsilon > 0$ be given.

Let $M_1 \in \mathbb{N}$ be such that, for every $m > M_1$, for every $k \in \{m, \dots, m + \ell - 1\}$, we have $\|x^k - x^m\|_2 < \varepsilon/(3 + 2L_i)$. Lemma 4.13 guarantees the existence of such an M_1 .

Let $M_2 \in \mathbb{N}$ be such that for every $m > M_2$ it holds that $|b_{i^m}(x^m) - x_{i^m}^m| < \varepsilon/(3 + 2L_i)$. Proposition 4.11 guarantees the existence of such an M_2 .

Let $m > \max\{M_1, M_2\}$ be such that $\|x^m - x\|_2 < \varepsilon/(3 + 2L_i)$. Such an m must exist, since x is a cluster point of the sequence $(x^k)_{k \in \mathbb{N}}$.

If player i updates in every ℓ periods, then there exists $k' \in \{m, \dots, m + \ell - 1\}$ such that $|b_i(x^{k'}) - x_i^{k'}| < \varepsilon/(3 + 2L_i)$, where we use that $m > M_2$. Since $m > M_1$ as well, it holds that $\|x^{k'} - x^m\|_2 < \varepsilon/(3 + 2L_i)$, and by the choice of m we have $\|x - x^m\|_2 < \varepsilon/(3 + 2L_i)$. In particular, it follows that $|x_i^{k'} - x_i^m| < \varepsilon/(3 + 2L_i)$ and $|x_i^m - x_i| < \varepsilon/(3 + 2L_i)$. By the triangle inequality we get

$$|b_i(x^{k'}) - x_i| \leq |b_i(x^{k'}) - x_i^{k'}| + |x_i^{k'} - x_i^m| + |x_i^m - x_i| < \frac{3\varepsilon}{3 + 2L_i}.$$

Also, $\|x - x^m\|_2 < \varepsilon/(3 + 2L_i)$ and $\|x^m - x^{k'}\|_2 < \varepsilon/(3 + 2L_i)$ imply that $\|x - x^{k'}\|_2 < 2\varepsilon/(3 + 2L_i)$. Using the Lipschitz continuity of b_i , we get

$$|b_i(x) - b_i(x^{k'})| < \frac{2L_i\varepsilon}{3 + 2L_i}.$$

Summing up, we have

$$|b_i(x) - x_i| \leq |b_i(x) - b_i(x^{k'})| + |b_i(x^{k'}) - x_i| < \frac{2L_i\varepsilon}{3 + 2L_i} + \frac{3\varepsilon}{3 + 2L_i} = \varepsilon.$$

This concludes the proof. ■

Theorem 4.14 combines the results in Propositions 4.8, 4.10, and 4.11. It states that when players update their actions regularly, an α -centered path with $\alpha < 1$ will cluster around the set of Nash equilibria. The strength of Theorem 4.14 is that it holds for every weighted network game.

As our final result of this paper, we show that in the generic case of a finite set of Nash equilibria, every infinite α -centered path with $\alpha < 1$ converges to a Nash equilibrium, provided that the players update regularly.

Theorem 4.15. *Let $(\bar{x}, w, t) \in P$ and let $(x^k)_{k \in \mathbb{N}}$ be an α -centered path with $\alpha < 1$ in a game $G \in \mathcal{G}(\bar{x}, w, t)$ such that $|X^*(\bar{x}, w, t)| < \infty$. If there is $\ell \in \mathbb{N}$ such that every player updates in every ℓ periods, then there exists $x^* \in X^*(\bar{x}, w, t)$ such that $\lim_{k \rightarrow \infty} x^k = x^*$.*

Proof. Let Y denote the non-empty set of cluster points of $(x^k)_{k \in \mathbb{N}}$. Theorem 4.14 implies that every element of Y is a Nash equilibrium. We therefore only have to show that Y is a singleton. We know that the set Y is finite, since $Y \subseteq X^*(\bar{x}, w, t)$ and the set $X^*(\bar{x}, w, t)$ is finite by assumption.

Let some $y \in Y$ be given. Since the set Y is finite, there exists $\varepsilon > 0$ such that for every $x \in X \setminus \{y\}$ with $\|x - y\|_2 \leq \varepsilon$ it holds that $\phi(x) - \phi(y) < 0$. Take $\varepsilon > 0$ sufficiently small such that the set

$$D(y) = \{x \in X : \frac{\varepsilon}{2} \leq \|x - y\|_2 \leq \varepsilon\}$$

is non-empty. Since $D(y)$ is also compact, the number $\underline{\phi} = \max_{x \in D(y)} \phi(x)$ is well-defined. Note that $\phi(y) > \underline{\phi}$.

Since ϕ is continuous and $y \in Y$, we have that $\lim_{k \rightarrow \infty} \phi(x^k) = \phi(y)$. So there exists $M_1 \in \mathbb{N}$ such that for every $k > M_1$ it holds that $\phi(x^k) > \underline{\phi}$. Furthermore, since $\lim_{k \rightarrow \infty} \|x^{k+1} - x^k\|_2 = 0$ by Proposition 4.10, it holds that there exists $M_2 \in \mathbb{N}$ such that for every $k > M_2$ we have $\|x^{k+1} - x^k\|_2 < \varepsilon/2$.

Let $m > \max\{M_1, M_2\}$ be such that $\|x^m - y\|_2 < \varepsilon/2$. Such an m must exist due to the fact that $y \in Y$. We argue that for every $k > m$ we have $\|x^k - y\|_2 < \varepsilon/2$. Suppose to the contrary that there exists $k > m$ with $\|x^k - y\|_2 \geq \varepsilon/2$ and let k be the smallest such number. Since $k > M_2$ and $\|x^{k-1} - y\|_2 < \varepsilon/2$, we have $\|x^k - y\|_2 < \varepsilon$, hence $x^k \in D(y)$ and $\phi(x^k) \leq \underline{\phi} < \phi(x^m)$, contradicting the fact that the sequence $(\phi(x^k))_{k \in \mathbb{N}}$ is non-decreasing.

We have shown that for every $\varepsilon > 0$ sufficiently small, there exists $m \in \mathbb{N}$ such that for every $k > m$ it holds that $\|x^k - y\|_2 < \varepsilon/2$. It follows that y is the only cluster point of $(x^k)_{k \in \mathbb{N}}$. ■

Theorem 4.15 provides sufficient conditions for an α -centered path to converge to a Nash equilibrium. The conditions are as follows: every player must update regularly, α must be less than one, and the set of Nash equilibria has to be finite. The latter condition holds generically as stated in Corollary 3.7.

We conclude this section by an example illustrating the tightness of these sufficient conditions.

Example 4.16. Consider the case with no strategic interaction, $w = 0$, and interior target values, for every $i \in I$, $0 < t_i < \bar{x}_i$. Then a game $G \in \mathcal{G}(\bar{x}, w, t)$ has a single Nash equilibrium, $x^* = t$. Since the Nash equilibrium set is finite, Theorem 4.14 applies. It is easy to see that to achieve convergence to the Nash equilibrium, the $\alpha < 1$ condition cannot be weakened even in this simple case.

Letting $\alpha = 1$ allows for $\lim_{k \rightarrow \infty} \alpha_k = 1$, which means that Proposition 4.10 no longer holds. In this case, the distance to the best reply, which is equal to the target value t_i , is no longer converging to zero. We have no convergence to the Nash equilibrium.

5 Conclusion

In this paper we consider weighted network games, a class of games with a very wide range of applications, where direct, pairwise player interactions are described by a matrix of weights. We show that this class of games is a subset of the class of best-response potential games and that its set of Nash equilibria is generically finite.

Pairs of linked players can always benefit from jointly deviating in an interior equilibrium. Two players whose contributions are strategic substitutes of each other can jointly increase their actions to increase their payoffs, while players whose contributions are strategic complements can jointly decrease their actions to improve their payoffs. In case all players' actions are strategic substitutes or all players' actions are strategic complements, such deviations lead to Pareto improvements. Therefore, in general, equilibria are neither strong nor efficient.

We study a large class of better-response learning processes. The convergence properties of these processes are determined by their centering parameter, which indicates to what extent players can overshoot their best responses. If players move closer to the best response at each update, as is the case for best-response dynamics and better-response dynamics with a centering parameter of one, then the players get arbitrarily close to the set of Nash equilibria and converge to a single Nash equilibrium whenever the set of Nash equilibria is finite, which is generically the case. This is due to the fact that the best-response potential function is in every variable symmetric around the best response of the players, hence moving closer to the best response increases the value of the potential. In the case of better-response dynamics with unrestricted overshooting, it is shown that cycles may arise.

The restrictions on overshooting that guarantee convergence to a Nash equilibrium in the general case are the same as in a trivial game with no strategic interaction. The reason for this is that the best-response potential function can be shown to increase whenever an updating player moves closer to his best response, irrespective of the values of the interaction parameters. Our results hence identify a rich class of learning processes that produce Nash equilibria, including cautious better-reply dynamics.

Topics that are left unexplored in this paper include asymmetric interaction parameters. In this case the existence of a best-response potential is no longer guaranteed, and hence, best-response cycles may occur. Asymmetric interaction is a more general framework that allows the modeling of a wider range of decision making scenarios. Examples include pollution between two neighboring cities where one of the cities is upriver, hence it is affected far less by the pollution of its neighbor than vice versa. Another topic is the issue of inefficiency of equilibria. These inefficiencies may disappear in different – possibly more centralized – classes of learning processes. Finally, the beliefs that shape the updates themselves are left unmodeled and unexplored in this paper. These topics are open for future research.

References

- Allouch, N., 2015. On the private provision of public goods on networks. *Journal of Economic Theory*, 157, 527-552.
- Ballester, C., Calvó-Armengol, A., and Zenou, Y., 2006. Who's who in networks. Wanted: The key player. *Econometrica*, 74(5), 1403-1417.
- Ballester, C. and Calvó-Armengol, A., 2010. Interactions with hidden complementarities. *Regional Science and Urban Economics*, 40(6), 397-406.
- Belhaj, M., Bramoullé, Y., and Deroïan, F., 2014. Network games under strategic complementarities. *Games and Economic Behavior*, 88, 310-319.
- Bervoets, S., Bravo, M., and Faure, M., 2016. Learning and convergence to Nash in games with continuous action sets. Working paper.
- Bervoets, S., Faure, M., 2016. Best-response in pure public good games on networks. Working paper.
- Blume, L.E., Brock, W.A., Durlauf, S.N., and Ioannides, Y.M., 2010. Identification of social interactions. In *Handbook of Social Economics*, edited by Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, 853-964. Amsterdam: North Holland.
- Bourlès, R., Bramoullé, Y., and Perez-Richet, E., 2017. Altruism in networks. *Econometrica*, 85(2), 675-689.
- Bramoullé, Y., and Kranton, R., 2007. Public goods in networks. *Journal of Economic Theory* 135(1), 478-494.
- Bramoullé, Y., Kranton, R., and D'Amours, M., 2014. Strategic interaction and networks. *American Economic Review* 104(3), 898-930.
- Dubey, P., Haimanko, O., and Zapechelnyuk, A., 2006. Strategic complements and substitutes, and potential games. *Games and Economic Behavior* 54(1), 77-94.
- Eksin, C., Molavi, P., Ribeiro, A., and Jadbabaie, A., 2012. Learning in linear games over networks. In *Communication, Control, and Computing (Allerton)*, 50th Annual Allerton Conference (434-440). IEEE.

- Ewerhart, C., 2017. The lottery contest is a best-response potential game. *Economics Letters* 155, 168-171.
- Elliott, M., and Golub, B., 2013, June. A network approach to public goods. In *Proceedings of the fourteenth ACM conference on Electronic commerce* (377-378). ACM.
- Fudenberg, D., and Levine, D.K., 1998. *The Theory of Learning in Games*. MIT press.
- Goyal, S., 2012. *Connections: An Introduction to the Economics of Networks*. Princeton University Press.
- Helsley, R.W., and Zenou, Y., 2014. Social networks and interactions in cities. *Journal of Economic Theory*, 150, 426-466.
- Jensen, M.K., 2010. Aggregative games and best-reply potentials. *Economic Theory* 43(1), 45-66.
- Jackson, M.O., and Zenou, Y., 2014. Games on networks. *Handbook of Game Theory*, Vol. 4, Peyton Young and Shmuel Zamir, eds., Elsevier Science.
- König, M., Liu, X. and Zenou, Y., 2014. R&D networks: Theory, empirics and policy implications. Unpublished manuscript.
- Kukushkin, N.S., 2015. Cournot tatonnement and potentials. *Journal of Mathematical Economics*, 59, 117-127.
- Leontief, W., 1970. Environmental repercussions and the economic structure: An input-output approach. *The Review of Economics and Statistics* 52(3), 262-271.
- Monderer, D., and Shapley, L., 1996. Potential games. *Games and Economic Behavior* 14(1), 124-143.
- Sandler, T., and Hartley, K., 1995. *The Economics of Defense*. Cambridge University Press.
- Sandler, T., and Hartley, K., 2001. Economics of alliances: The lessons for collective action. *Journal of Economic Literature* 39(3), 869-896.
- Sandler, T., and Hartley, K., 2007. *Handbook of Defense Economics: Defense in a globalized world*. Elsevier.
- Voorneveld, M., 2000. Best-response potential games. *Economic Letters* 66, 289-295.