Games and Economic Behavior 123 (2020) 182–206

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Fictitious play in networks ☆

Christian Ewerhart [a,*], Kremena Valkanova [b]

[a] Department of Economics, University of Zurich, Schönberggasse 1, CH-8001 Zurich, Switzerland
[b] Zurich Graduate School of Economics, University of Zurich, Blümlisalpstrasse 10, CH-8006 Zurich, Switzerland

ARTICLE INFO

Article history:
Received 6 December 2016
Available online 3 July 2020

JEL classification:
C72
D83
D85

Keywords:
Fictitious play
Networks
Zero-sum games
Conflicts
Potential games
Miyasawa's theorem
Robinson's theorem

ABSTRACT

This paper studies fictitious play in networks of noncooperative two-person games. We show that continuous-time fictitious play converges to the set of Nash equilibria if the overall n-person game is zero-sum. Moreover, the rate of convergence is $1/\tau$, regardless of the size of the network. In contrast, arbitrary n-person zero-sum games with bilinear payoff functions do not possess the continuous-time fictitious-play property. As extensions, we consider networks in which each bilateral game is either strategically zero-sum, a weighted potential game, or a two-by-two game. In those cases, convergence requires a condition on bilateral payoffs or, alternatively, that the network is acyclic. Our results hold also for the discrete-time variant of fictitious play, which implies, in particular, a generalization of Robinson's theorem to arbitrary zero-sum networks. Applications include security games, conflict networks, and decentralized wireless channel selection.

1. Introduction

Fictitious play (Brown, 1949, 1951; Robinson, 1951) refers to a class of simple and intuitive models of learning in games. The common element of such models is that a player is assumed to respond optimally to an evolving belief on the behavior of her opponents, where the player's belief at any point in time is formed on the basis of the empirical frequencies of strategy choices made by her opponents up to that point. Understanding the conditions under which fictitious play converges to Nash equilibrium is important because such results help to clarify the intuition that equilibrium play may be reached even if players are not perfectly rational.[1] While variants of fictitious play are known to converge in large classes of two-person games, the case of n-person games has been explored to a somewhat lesser extent.[2]

☆ Comments by an Advisory Editor and two anonymous reviewers helped us to substantially improve the exposition of the paper. Josef Hofbauer kindly provided detailed comments on an earlier version. For useful discussions and comments, we would like to thank Ulrich Berger, Sergiu Hart, Dov Monderer, Stephen Morris, Hervé Moulin, Dov Samet, and Aner Sela. Kremena Valkanova would like to thank MIT and Drew Fudenberg for their generous hospitality during her stay in the study year 2019-2020, gratefully acknowledging funding from the Swiss National Science Foundation (grant No. P1ZHP1_184169).

* Corresponding author.
E-mail addresses: christian.ewerhart@econ.uzh.ch (C. Ewerhart), kremena.valkanova@econ.uzh.ch (K. Valkanova).

[1] The literature on fictitious play is too large to be surveyed here. For an introduction to the theory of learning in games, see Fudenberg and Levine (1998). The literature on learning in social networks has been surveyed by Acemoglu and Ozdaglar (2011). For a concise discussion of epistemic vs. dynamic foundations of Nash equilibrium, see Krishna and Sjöström (1997).

[2] Positive convergence results allowing for more than two players have been established, in particular, for games solvable by iterated dominance (Milgrom and Roberts, 1991), games with identical interests (Monderer and Shapley, 1996b; Harris, 1998), and certain classes of star-shaped network games (Sela, 1999). Shapley (1964) pointed out that fictitious play need not converge to Nash equilibrium in two-person non-zero-sum games, and he also noted that

https://doi.org/10.1016/j.geb.2020.06.006
0899-8256/

Inspired by recent developments in the literature (Daskalakis and Papadimitriou, 2009; Cai and Daskalakis, 2011; Cai et al., 2016), the present paper studies the dynamics of fictitious play in general classes of network games. We start by considering what we call *zero-sum networks* (Bregman and Fokin, 1987, 1998; Daskalakis and Papadimitriou, 2009; Cai and Daskalakis, 2011; Cai et al., 2016). These are n-person zero-sum games that can be represented as a network of two-person games. This class of games is actually quite large and includes practically relevant examples of resource allocation games such as generalized Blotto and security games. Our first main result says that any continuous-time fictitious-play (CTFP) path in a zero-sum network converges in payoffs to zero at rate $1/\tau$ (where τ denotes time), regardless of the size of the network. In particular, CTFP converges to the set of Nash equilibria in any zero-sum network. However, as we also show with an example, arbitrary n-person zero-sum games with bilinear payoff functions need not possess the CTFP property, i.e., the network assumption is crucial for our conclusions.

To prove our continuous-time convergence result, we employ the standard Lyapunov approach (Brown, 1951; Hofbauer, 1995; Harris, 1998; Berger, 2006; Hofbauer and Sorin, 2006; Hofbauer and Sandholm, 2009). Thus, we consider a Lyapunov function that aggregates, across all players in the network, the maximum payoff that could be obtained by optimizing against the empirical frequency distribution of prior play. It is then shown that the Lyapunov function diminishes along the CTFP path as $\tau \to \infty$. In slight departure from the literature, however, this property is established by putting an upper bound on a Dini derivative instead of calculating the usual (directional) derivate. As we argue in an Appendix, this may be seen as a certain simplification vis-à-vis existing proofs.[3]

To gauge the role of the zero-sum assumption, we consider three additional classes of network games. First, we look at networks of strategically zero-sum games, or *conflict networks*.[4] Thus, each bilateral game in the network is assumed to be best-response equivalent in mixed strategies to a zero-sum game. Moulin and Vial (1978) noted that fictitious play converges in this class of two-person games. In conflict networks, CTFP converges as well, provided that valuations in the bilateral games satisfy a condition that we call *pairwise homogeneity of valuations*. This assumption is satisfied, for example, in transfer networks considered by Franke and Öztürk (2015). When valuations are heterogeneous, however, convergence need not hold in general. Intuitively, the aggregation of bilateral payoffs does not commute with the equivalence relation, because payoff transformations that turn two bilateral games into zero-sum games need not be identical. We illustrate this possibility with an example of a network conflict that does not settle down under CTFP even though the unique Nash equilibrium is peaceful. But convergence can still be obtained with heterogeneous valuations when the underlying network is *acyclic*, i.e., when it is a disjoint union of trees.[5] Second, we assume that bilateral games are weighted potential games. Applications include channel selection problems in wireless communication networks, and the spreading of ideas and technologies over social networks, for instance. Extending the analysis of Cai and Daskalakis (2011), we show that CTFP converges to equilibrium in any network of exact potential games. However, as we illustrate with still another example, fictitious play need not converge in general networks of weighted potential games. Instead, in analogy to the previously considered case, the convergence result holds for weighted potential games under the condition that the underlying network structure is acyclic. Third, by combining our findings for conflict networks with pairwise homogeneous valuations and for networks of exact potential games, we obtain a generalization of Miyasawa's (1961) theorem to network games on arbitrary graphs.

As an extension, the paper looks at discrete-time fictitious play (DTFP), generalizing Robinson's (1951) famous result for two-person zero-sum games to arbitrary n-person zero-sum networks, and showing that any DTFP is belief affirming, which extends a result of Monderer et al. (1997). Finally, we discuss the possibility of correlated beliefs which is a relevant aspect in multiplayer games.

Related literature. The first paper studying fictitious play in an environment similar to ours is Sela (1999). His observation was that some of the convergence results for two-person games generalize quite easily to n-person games with a "one-against-all" structure. In that setting, one player located in the center of the star-shaped network chooses a *compound* strategy that is the same in every bilateral interaction. Then, the network game can be transformed into a two-person game in which the choices of the peripheral players are orchestrated by a single agent that maximizes the sum of the payoffs of those players. Under a specific tie-breaking rule, the DTFP process in the reduced game turns out to be identical, for any given initial condition, to the DTFP process in the "one-against-all" game. Thereby, DTFP and CTFP properties in the network game can be established as a corollary of results for two-person games, provided that the bilateral games are all of the same type, like zero-sum, with identical payoffs, or generic two-by-two. It is, however, not obvious how this approach could

this observation extends to games with more than two players. A three-person non-zero-sum counterexample has been constructed by Jordan (1993). See also Gaunersdorfer and Hofbauer (1995).

[3] Driesen (2009) pursues an alternative route to simplification, yet at the cost of assuming pure choices, which might interfere with existence (cf. Harris, 1998, p. 242). See also Shamma and Arslan (2004) who unify existing Lyapunov arguments in a set-up with a "soft-max" best response.

[4] Recent papers that look at conflict networks include Bozbay and Vesperoni (2018), Dziubiński et al. (2016a), Franke and Öztürk (2015), Huremović (2014), Jackson and Nei (2015), König et al. (2017), Kovenock et al. (2015), Kovenock and Roberson (2018), Matros and Rietzke (2018), and Xu et al. (2019), among others. For a survey, see Dziubiński et al. (2016b).

[5] E.g., any star-shaped network considered by Sela (1999) is acyclic, but the network shown in Fig. 1 below is not acyclic.

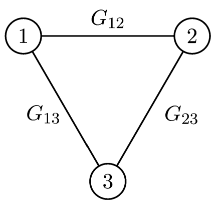

Fig. 1. A network game.

be applied to networks that are not star-shaped.[6] The present paper extends the results of Sela (1999) to general network structures. We also drop the tie-breaking rule and deal more explicitly with the case of CTFP.

An interesting recent strand of literature, related to the interdisciplinary field of algorithmic game theory, has taken up the study of networks of two-person games.[7] Cai et al. (2016) have clarified how far results traditionally known only for two-person zero-sum games (such as solvability by a linear program, existence of a value, equivalence of max-min and equilibrium strategies, exchangeability of Nash equilibria, and the relationship to coarse correlated equilibrium) can be extended to zero-sum networks. Moreover, in that class of games, discrete-time no-regret learning algorithms converge to the set of Nash equilibria (Daskalakis and Papadimitriou, 2009; Cai and Daskalakis, 2011).[8] As will be discussed, those results do not allow any immediate conclusions regarding the convergence of fictitious play in zero-sum networks. However, our convergence result for discrete-time fictitious play (Proposition 5) certainly has a similar flair as the existing results regarding no-regret behavior.

Another natural class of network games is defined by the requirement that players possess identical payoff functions in each bilateral game (Cai and Daskalakis, 2011).[9] Under the condition that pairwise interactions are games with identical payoffs, the network game is shown to possess an exact potential, which implies that, in this class of games, certain learning algorithms converge to equilibrium. In particular, the discrete dynamics of pure best responses converges to the set of Nash equilibria. However, Cai and Daskalakis (2011) do not discuss fictitious play.

The remainder of this paper is structured as follows. Section 2 contains preliminaries. Convergence of CTFP in zero-sum networks is established in Section 3. Section 4 deals with additional classes of games. DTFP is considered in Section 5. Section 6 discusses correlated beliefs. Section 7 concludes. Appendices provide auxiliary results from the literature, details on two of our examples, and a discussion of the case of two-person zero-sum games.

2. Preliminaries

2.1. Network games

There is a finite set $V = \{1, \ldots, n\}$ of players (countries, firms, consumers, political institutions, etc).[10] Let $E \subseteq V \times V$ be a set of bilateral relationships. Any two players $i, j \in V$ are either in interaction (i.e., $(i, j) \in E$) or not (i.e., $(i, j) \notin E$). Thus, the pair (V, E) is a graph, and we assume that it is (i) undirected (i.e., $\forall i, j : (i, j) \in E \Leftrightarrow (j, i) \in E$) and (ii) irreflexive (i.e., $\forall i : (i, i) \notin E$).[11] Each edge $(i, j) \in E$ represents a finite two-person game G_{ij} between players i and j, where we assume that G_{ij} and G_{ji} refer to the same game, yet in the first (second) case from i's (from j's) perspective. The respective sets of *bilateral* strategies for players i and j in G_{ij} will be denoted by S_{ij} and S_{ji}. Similarly, payoff functions for players i and j in G_{ij} will be denoted by $u_{ij} : S_{ij} \times S_{ji} \to \mathbb{R}$ and $u_{ji} : S_{ji} \times S_{ij} \to \mathbb{R}$, respectively. Note that the *first* argument in a bilateral payoff function u_{ij} always refers to player i's strategy. E.g., in the expression $u_{21}(s_{21}, s_{12})$, strategy s_{21} is 2's bilateral strategy vis-a-vis player 1, and s_{12} is 1's bilateral strategy vis-a-vis player 2.

Let $N(i) = \{j : (i, j) \in E\}$ denote the set of player i's *neighbors*. Fig. 1 shows a network with three players, each of them interacting with two neighbors. As usual, the number of neighbors of a player corresponds to the player's degree as a node

---

6 Daskalakis and Papadimitriou (2009) make a related point in their discussion of the complexity of computing Nash equilibrium in networks of zero-sum games.

7 Network games have, of course, a long tradition in game theory. See, e.g., the recent survey by Bramoullé and Kranton (2016) and references given therein.

8 Under *no-regret behavior*, a player's expected payoff against past play is asymptotically weakly lower than her average payoff experience. An example is the multiplicative-weights adaptive learning algorithm (Freund and Schapire, 1999). See Cesa-Bianchi and Lugosi (2006). For useful discussions of the relationship between no-regret learning and fictitious play, see Hart and Mas-Colell (2001) and Viossat and Zapechelnyuk (2013).

9 This class includes, e.g., binary coordination games, as considered by Bramoullé and Kranton (2016, Prop. 3). See also Bramoullé et al. (2014) and Bourlès et al. (2017). For early uses of potential methods in network models, see Blume (1993) and Young (1993).

10 Finiteness of the network seems essential for the convergence of fictitious play. However, we have not looked specifically into this issue. Blume (1993) studies the strategic interaction of players that are located on an infinite lattice. Morris (2000) considers best-response dynamics in locally finite networks of coordination games.

11 While the network structure (V, E) is exogenous, our set-up is consistent with the assumption that players strategically choose a subset of their neighbors as (potential) partners (see, e.g., Jackson, 2005). Along these lines, the dynamic evolution of a network is subsumed as well.

in the network of interactions. In the example, the network is complete, but this is not assumed.[12] Let $\varnothing \neq X_i \subseteq \times_{j\in N(i)} S_{ij}$ denote the set of *multilateral* strategies of player i (defense policies, trade quotas, promotional strategies, prices, invitation or acceptance of friendship, etc.).

For a given multilateral strategy $x_i \in X_i$ of player i, we denote by $\sigma_{ij}(x_i) = s_{ij} \in S_{ij}$ the corresponding bilateral strategy vis-a-vis player j. Going over all neighbors of player i, it becomes clear that any multilateral strategy $x_i \in X_i$ may be considered as a vector of bilateral strategies $x_i = \{s_{ij}\}_{j\in N(i)} = \{\sigma_{ij}(x_i)\}_{j\in N(i)}$. It is important to note that we allow for $X_i \subsetneq \times_{j\in N(i)} S_{ij}$, which is the interesting case in most applications. For example, there could be budget constraints, limited resources (e.g., planes in a military conflict), or constraints regarding price coherence across platforms.

An important special case arises if bilateral and multilateral strategy spaces coincide, i.e., if σ_{ij} defines a one-to-one correspondence $X_i \simeq S_{ij}$ for any $i \in V$ and any $j \in N(i)$. In that case, player i's multilateral strategy $x_i \in X_i$ implements the same *compound* strategy $x_i \simeq s_{ij}$ in each bilateral game with neighbor $j \in N(i)$, so that X_i corresponds to the diagonal in $\times_{j\in N(i)} S_{ij}$. Settings along these lines have been considered, in particular, by Sela (1999) and Cai et al. (2016), and will also be used in our examples. Clearly, our set-up is no less general than those settings.

In an n-player *network game* G, each player $i \in \{1,\dots,n\}$ chooses a multilateral strategy $x_i \in X_i$, and receives payoff

$$u_i(x_i, x_{-i}) \equiv u_i(x_i, x_{N(i)}) = \sum_{j\in N(i)} u_{ij}(\sigma_{ij}(x_i), \sigma_{ji}(x_j)), \tag{1}$$

where $x_{-i} \in X_{-i} = \times_{j\neq i} X_j$ specifies a multilateral strategy for every player $j \neq i$ different from player i, while $x_{N(i)} = \{x_j\}_{j\in N(i)} \in X_{N(i)} = \times_{j\in N(i)} X_j$ specifies a multilateral strategy for every neighbor $j \in N(i)$. Reflecting the local nature of interaction and payoffs, it will be assumed below that each player i forms a belief about $x_{N(i)}$ only, rather than about x_{-i}. While this assumption is not required for our results, it may be considered somewhat more plausible in a learning context. Finally, let $X = \times_{i=1}^n X_i$ denote the set of multilateral strategy profiles in G.

We denote by $\Delta(X_i)$ the set of player i's mixed multilateral strategies. Thus, any $\mu_i \in \Delta(X_i)$ is a (column) vector representing a probability distribution on X_i. Payoff functions extend to mixed multilateral strategies in the usual way. Specifically, if $\mu_{-i} \in \times_{j\neq i} \Delta(X_j)$ is a profile of mixed strategies for all players except player i, and if $\mu_{N(i)} \in \times_{j\in N(i)} \Delta(X_j)$ denotes the resulting profile of mixed strategies for the neighbors of player i, then player i's expected payoff from the mixed strategy $\mu_i \in \Delta(X_i)$ is given by $u_i(\mu_i, \mu_{-i}) = E[u_i(x_i, x_{-i})] = u_i(\mu_i, \mu_{N(i)}) = E[u_i(x_i, x_{N(i)})]$, where the expectations are taken with respect to (μ_i, μ_{-i}) and $(\mu_i, \mu_{N(i)})$, respectively. Player i's mixed best-response correspondence MBR_i assigns to any profile $\mu_{N(i)} \in \times_{j\in N(i)} \Delta(X_j)$ the set of mixed strategies $\mu_i^* \in \Delta(X_i)$ such that $u_i(\mu_i^*, \mu_{N(i)}) = \max_{\mu_i\in\Delta(X_i)} u_i(\mu_i, \mu_{N(i)})$. Further, the mixed best-response correspondence MBR of the network game G assigns to any profile $\mu = (\mu_1,\dots,\mu_n) \in \times_{i=1}^n \Delta(X_i)$ the Cartesian product $\mathrm{MBR}(\mu) = \times_{i=1}^n \mathrm{MBR}_i(\mu_{N(i)})$. Since strategy spaces are finite, a mixed-strategy Nash equilibrium $\mu^* = (\mu_1^*,\dots,\mu_n^*)$, i.e., a fixed point of MBR, exists by Nash's theorem.

2.2. Continuous-time fictitious play with independent beliefs

We start by considering the learning process in *continuous time* (Brown, 1951; Rosenmüller, 1971). Moreover, we will initially assume that beliefs formed by a player are *independent* across neighbors, so that empirical frequencies are accounted for as marginal distributions only. Note that this latter assumption is consistent with a common interpretation of fictitious play, according to which players take it as given that their opponents adhere to some independently chosen mixed strategy. Both assumptions will be relaxed in later sections of this paper.

Let $m : [0,\infty) \to \Delta(X_1) \times \dots \times \Delta(X_n)$ be a measurable *path* specifying, for any point in time $\tau \geq 0$ and for any player $i \in \{1,\dots,n\}$, a mixed strategy $m_i(\tau) \in \Delta(X_i)$. As time is continuous, independent averaging over time amounts to integrating the components of m over a non-degenerate interval. Consequently, the *(continuous-time) independent average* $\alpha : (0,\infty) \to \Delta(X_1) \times \dots \times \Delta(X_n)$ of the path m is defined, for any time $\tau > 0$ and for any player $i \in \{1,\dots,n\}$, by the belief

$$\alpha_i(\tau) \equiv \alpha_i(\tau, m) = \frac{1}{\tau} \int_0^\tau m_i(\tau')d\tau' \in \Delta(X_i). \tag{2}$$

We will use the following definition of continuous-time fictitious play.

Definition 1 *(CTFP).* A continuous-time fictitious play (with independent beliefs) is a measurable mapping $m : [0,\infty) \to \times_{i=1}^n \Delta(X_i)$ such that $m(\tau) \in \mathrm{MBR}(\alpha(\tau))$ for all $\tau \geq 1$.

[12] We shall use standard terminology of graph theory (see, e.g., Bollobás, 2013). Thus, a network (V, E) is *complete* if $N(i) = V\backslash\{i\}$ for all $i \in V$; it is *acyclic* when there is no finite sequence of pairwise distinct players $i_1,\dots,i_\kappa \in V$ with $\kappa \geq 3$ such that $(i_1,i_2) \in E$, $(i_2,i_3) \in E$, ..., $(i_{\kappa-1},i_\kappa) \in E$, and $(i_\kappa,i_1) \in E$; a network is *star-shaped* if there is a player $i \in V$ such that $N(i) = V\backslash\{i\}$ and $N(j) = \{i\}$ for any $j \in V\backslash\{i\}$; finally, a network is *connected* if, for any $i, j \in V$ with $i \neq j$, there is a finite sequence $i_1, i_2,\dots,i_\kappa \in V$ with $i_1 = i$ and $i_\kappa = j$ such that $(i_1,i_2) \in E$, $(i_2,i_3) \in E$, ..., and $(i_{\kappa-1},i_\kappa) \in E$.

Thus, Definition 1 requires optimality at *any* point in time $\tau \geq 1$, whereas the closely related notion of CTFP1 used by Harris (1998) demands optimality at *almost any* point in time.[13]

The following lemma assures us of the existence of a CTFP learning process.

Lemma 1. *A CTFP exists.*

Proof. By Lemma A.1 in the Appendix, a measurable path $\widehat{m} : [0, \infty) \to \times_{i=1}^{n} \Delta(X_i)$ exists such that $\widehat{m}(\tau) \in \text{MBR}(\alpha(\tau, \widehat{m}))$ for all $\tau \in [1, \infty)\backslash\mathcal{N}$, where $\mathcal{N} \subseteq \mathbb{R}$ is a set of measure zero. We construct a modified path $m : [0, \infty) \to \times_{i=1}^{n} \Delta(X_i)$ by letting $m(\tau) = \widehat{m}(\tau)$ for any $\tau \in [0, \infty)\backslash\mathcal{N}$, and by choosing a mixed best response $m(\tau) \in \text{MBR}(\alpha(\tau, \widehat{m}))$ for any $\tau \in \mathcal{N}$. Then, clearly, $\alpha(\tau, m) = \alpha(\tau, \widehat{m})$ for any $\tau \geq 1$, so that $m(\tau) \in \text{MBR}(\alpha(\tau, m))$ for any $\tau \geq 1$. The claim follows. □

The proof is based upon an existence result of Harris (1998) for CTFP1 paths.[14] The CTFP1 path is modified by replacing any suboptimal mixed best response by an optimal mixed best response. Since the original path is changed on a set of measure zero, averages stay the same, and the resulting path is a continuous-time fictitious play according to Definition 1.

Next, we define convergence of CTFP in a given network game G. Recall that, for an arbitrary measurable path $m : [0, \infty) \to \Delta(X_1) \times \cdots \times \Delta(X_n)$, the independent average $\tau \mapsto \alpha(\tau) \equiv \alpha(\tau, m)$ is a continuous path in the space of mixed strategy profiles, $\Delta(X_1) \times \cdots \times \Delta(X_n)$. Denote by $\mathcal{A}(m)$ the set of all accumulation points of $\alpha(., m)$.[15] Convergence in continuous time is then defined by the requirement that $\mathcal{A}(m)$ is a subset of the set of Nash equilibria of G.

Definition 2. A path m is said to converge to the set of Nash equilibria if every limit distribution $\mu^* \in \mathcal{A}(m)$ is a Nash equilibrium in G.

The reader is cautioned that convergence is to a set where all points are Nash equilibria, meaning that the trajectory of the CTFP dynamics might not converge to a specific Nash equilibrium. Only if the Nash equilibrium is unique, convergence implies that the trajectory will converge to a single point. We will say that G has the *continuous-time fictitious-play property* if any CTFP in G converges to the set of Nash equilibria.

It should be noted at this point that a network player that optimizes simultaneously against several opponents behaves differently from an unconstrained player in a bilateral game. Indeed, as pointed out by Sela (1999, Ex. 4 & 6), the fictitious-play property in the bilateral games (regardless of whether it holds in continuous or discrete time) is not generally informative about the corresponding property in the network game. Thus, even if the bilateral games have the fictitious-play property, this need not be the case for the network game. Conversely, if neither of the bilateral games possesses the fictitious-play property, the network game may still possess the fictitious-play property.

3. Zero-sum networks

By a *zero-sum network*, we mean a network game G that is zero-sum as an n-person game, i.e., a network game in which $u_1 + \cdots + u_n \equiv 0$. Clearly, if each bilateral game G_{ij} in a given network is two-person zero-sum, i.e., if $u_{ij} + u_{ji} \equiv 0$ for any $i, j \in \{1, \ldots, n\}$ with $j \neq i$, then the network game G is an n-person zero-sum game. However, the converse is not generally true. Moreover, even if payoffs are given in the n-person normal form, there are efficient ways to check if the game is a zero-sum network.[16]

The following result is the first main result of the present paper.

Proposition 1. *Any zero-sum network G has the CTFP property.*

Proof. For an arbitrary profile of mixed strategies $\mu = (\mu_1, \ldots, \mu_n) \in \Delta(X_1) \times \cdots \times \Delta(X_n)$, we define the Lyapunov function

$$\mathcal{L}(\mu) = \sum_{i=1}^{n} \max_{x_i \in X_i} \left\{ u_i(x_i, \mu_{N(i)}) - u_i(\mu_i, \mu_{N(i)}) \right\}, \tag{3}$$

where $\mu_{N(i)} = \{\mu_j\}_{j \in N(i)}$ denotes the restriction of μ to the neighbors of player i. As a direct consequence of the definition, $\mathcal{L}(\mu) \geq 0$. Further, given that G is an n-person zero-sum game, we may rewrite the Lyapunov function as

[13] A formal definition of CTFP1 is provided in Appendix A. Which definition one is using is, to our understanding, a matter of taste. E.g., Ostrovski and van Strien (2014) use CTFP for $n = 2$ players. In any case, all of our results remain valid when CTFP is replaced by CTFP1.

[14] For a formal statement, see Appendix A.

[15] Thus, $\mathcal{A}(m)$ consists of all strategy profiles that are limit points of some convergent sequence of independent averages, $\{\alpha(\tau_q, m)\}_{q=1}^{\infty}$, where $\{\tau_q\}_{q=1}^{\infty}$ is any sequence in $[1, \infty)$ such that $\lim_{q \to \infty} \tau_q = \infty$. Because the set of mixed strategy profiles $\Delta(X_1) \times \cdots \times \Delta(X_n)$ is compact, there will be at least one such limit point, i.e., $\mathcal{A}(m) \neq \varnothing$.

[16] In the literature, zero-sum networks are known as *zero-sum polymatrix games* or *separable zero-sum multiplayer games*. For additional background and discussion, see Bregman and Fokin (1987, 1998), Daskalakis and Papadimitriou (2009), Cai and Daskalakis (2011), and Cai et al. (2016).

$$\mathcal{L}(\mu) = \sum_{i=1}^{n} \left\{ \left(\max_{x_i \in X_i} u_i(x_i, \mu_{N(i)}) \right) - u_i(\mu_i, \mu_{N(i)}) \right\} \tag{4}$$

$$= \left(\sum_{i=1}^{n} \max_{x_i \in X_i} u_i(x_i, \mu_{N(i)}) \right) - \underbrace{\sum_{i=1}^{n} u_i(\mu_i, \mu_{N(i)})}_{=0} \tag{5}$$

$$= \sum_{i=1}^{n} \max_{x_i \in X_i} u_i(x_i, \mu_{N(i)}). \tag{6}$$

Take a CTFP path $m : [0, \infty) \to \Delta(X_1) \times \cdots \times \Delta(X_n)$, and let $\alpha(\tau) \equiv \alpha(\tau, m)$ denote the independent average at some point in time $\tau > 1$. Then, because $m_i(\tau)$ is a mixed best response to $\alpha_{N(i)}(\tau)$ for $i \in \{1, \ldots, n\}$, and because interactions are bilateral,

$$\mathcal{L}(\alpha(\tau)) = \sum_{i=1}^{n} u_i(m_i(\tau), \alpha_{N(i)}(\tau)) \tag{7}$$

$$= \sum_{i=1}^{n} \sum_{j \in N(i)} u_{ij}(m_i(\tau), \alpha_j(\tau)) \tag{8}$$

$$= \sum_{i=1}^{n} \sum_{j \in N(i)} m_i(\tau) \cdot A_{ij} \alpha_j(\tau), \tag{9}$$

where A_{ij} is the matrix of player i's payoffs in the bilateral game G_{ij}, i.e.,[17]

$$\mu_i \cdot A_{ij} \mu_j \equiv u_{ij}(\mu_i, \mu_j). \tag{10}$$

Multiplying the expression for $\mathcal{L}(\alpha(\tau))$ found in (9) by τ, and subsequently using the definition of $\alpha_j(\tau)$, yields

$$\tau \mathcal{L}(\alpha(\tau)) = \sum_{i=1}^{n} \sum_{j \in N(i)} m_i(\tau) \cdot A_{ij} \int_0^{\tau} m_j(\tau') d\tau'. \tag{11}$$

Consider now a player $i \in \{1, \ldots, n\}$, and some $\widehat{\tau} \in (1, \tau)$. Then, given that $m_i(\widehat{\tau})$ is a mixed best response to $\alpha_{N(i)}(\widehat{\tau})$, we have

$$\sum_{j \in N(i)} m_i(\widehat{\tau}) \cdot A_{ij} \alpha_j(\widehat{\tau}) \geq \sum_{j \in N(i)} m_i(\tau) \cdot A_{ij} \alpha_j(\widehat{\tau}). \tag{12}$$

Adding up across players, and subsequently multiplying through with $\widehat{\tau}$, one obtains

$$\widehat{\tau} \mathcal{L}(\alpha(\widehat{\tau})) \geq \sum_{i=1}^{n} \sum_{j \in N(i)} m_i(\tau) \cdot A_{ij} \int_0^{\widehat{\tau}} m_j(\tau') d\tau'. \tag{13}$$

Subtracting inequality (13) from equation (11), one arrives at

$$\tau \mathcal{L}(\alpha(\tau)) - \widehat{\tau} \mathcal{L}(\alpha(\widehat{\tau})) \leq \sum_{i=1}^{n} \sum_{j \in N(i)} m_i(\tau) \cdot A_{ij} \int_{\widehat{\tau}}^{\tau} m_j(\tau') d\tau'. \tag{14}$$

We divide this inequality by $\tau - \widehat{\tau} > 0$ and consider the limit for $\widehat{\tau} \to \tau$. Then, by the fundamental theorem of calculus, the limit on the right-hand side exists for almost any $\tau > 1$, in which case

$$\limsup_{\widehat{\tau} \to \tau, \widehat{\tau} < \tau} \frac{\tau \mathcal{L}(\alpha(\tau)) - \widehat{\tau} \mathcal{L}(\alpha(\widehat{\tau}))}{\tau - \widehat{\tau}} \leq \sum_{i=1}^{n} \sum_{j \in N(i)} m_i(\tau) \cdot A_{ij} m_j(\tau) = 0. \tag{15}$$

Thus, the upper-left Dini derivative of $\tau \mathcal{L}(\alpha(\tau))$ is a.e. weakly negative. Since, in addition, the mapping $\tau \mapsto \tau \mathcal{L}(\alpha(\tau))$ is continuous, this implies that $\tau \mathcal{L}(\alpha(\tau))$ is monotone decreasing (cf. Royden, 1988, p. 99). Hence, there is a constant $C \geq 0$

[17] Here and in the sequel, the dot $\cdot$ denotes the scalar product between two vectors.

such that $\mathcal{L}(\alpha(\tau)) \leq C/\tau$ for any $\tau \geq 1$. Noting that the individual terms of the Lyapunov function (3) are all positive, we therefore obtain for any player $i \in \{1, \ldots, n\}$ and for any pure strategy $x_i \in X_i$ that

$$u_i(x_i, \alpha_{N(i)}(\tau)) - u_i(\alpha_i(\tau), \alpha_{N(i)}(\tau)) \leq \frac{C}{\tau} \qquad (\tau \geq 1). \tag{16}$$

Take now any accumulation point $\mu^* \in \mathcal{A}(m)$ of the path $\alpha(.)$. Then, there exists a sequence $\{\tau_q\}_{q=1}^{\infty}$ in $[1, \infty)$ such that $\lim_{q\to\infty} \tau_q = \infty$ and $\lim_{q\to\infty} \alpha(\tau_q) = \mu^*$. Evaluating (16) at $\tau = \tau_q$, and considering the limit for $q \to \infty$, shows that

$$u_i(x_i, \mu^*_{N(i)}) - u_i(\mu^*_i, \mu^*_{N(i)}) \leq 0, \tag{17}$$

i.e., x_i is not a profitable deviation for player i. Since the inequality holds for any $i \in \{1, \ldots, n\}$ and any $x_i \in X_i$, the strategy profile μ^* is necessarily a Nash equilibrium. This completes the proof. □

Thus, the convergence result for two-person zero-sum games extends to zero-sum networks in a rather straightforward way. It may be noted that Proposition 1, in particular, provides a proof of existence of a Nash equilibrium in the considered class of n-person games.[18]

Apart from a minor modification (see Appendix C for details), the proof presented above follows the literature by combining the Lyapunov method with an envelope argument (Brown, 1951; Hofbauer, 1995; Harris, 1998; Berger, 2006; Hofbauer and Sorin, 2006; Hofbauer and Sandholm, 2009). Intuitively, $\mathcal{L}(\mu)$ measures each player's scope for individual improvement relative to μ, and aggregates the result across all n players in the network. Then, as equation (11) shows, the product $\tau\mathcal{L}(\alpha(\tau))$ corresponds to the total (across all players in the network) of the maxima over "cumulative payoffs." However, the network game is zero-sum, so that the sum of instantaneous payoffs vanishes. Therefore, $\tau\mathcal{L}(\alpha(\tau))$ cannot increase in τ.[19]

Regarding the rate of convergence, we mention that a minor refinement of the proof shows that, as in the case of two-person zero-sum games considered by Harris (1998), the rate of convergence in payoffs, i.e., the rate by which the Lyapunov function approaches zero, is precisely $1/\tau$.

Corollary 1. *The rate of convergence of CTFP in any zero-sum network is* $1/\tau$.

Proof. It suffices to note that, while inequality (14) holds likewise for any $\widehat{\tau} \in (\tau, \infty)$, the direction of the inequality is reversed when dividing by $\tau - \widehat{\tau} < 0$. Considering then the lower-right Dini derivative, the mapping $\tau \mapsto \tau\mathcal{L}(\alpha(\tau))$ is seen to be not only monotone decreasing, but also monotone increasing, hence constant. □

The observation that the rate of convergence does not depend on the size of the network may be surprising. However, it should be noted that convergence is measured here on the aggregate level. Thus, for a large network, the value of the Lyapunov function provides little information about the scope of improvement that is feasible for an *individual* player. Therefore, it might indeed take longer in a larger network to reach, say, an ε-equilibrium.

One might conjecture that the zero-sum property alone is sufficient to guarantee convergence of CTFP also in n-person games with $n \geq 3$. However, as a straightforward adaption of Shapley's (1964) example illustrates, this is not the case. I.e., there are multiplayer zero-sum games in which CTFP need not converge.

Example 1 *(Three-person zero-sum game).* Consider the game G^1 between three players depicted in Fig. 2.[20] In G^1, player 1 and player 2 each have three pure strategies, whereas player 3 has just one pure strategy. Therefore, to see what happens in equilibrium or under fictitious play, player 3 may be safely ignored. But with player 3 eliminated from the game, the two-person game between players 1 and 2 is of the Shapley (1964) type, so that nonconvergence obtains.[21]

Example 1 shows that the network assumption in Proposition 1 cannot be easily dropped. Actually, the example shows a bit more, namely that CTFP need not converge even in a three-person zero-sum game with bilinear payoffs. To see the point, note that the payoff functions in the mixed extension of G^1 are indeed bilinear. E.g., player 3's payoff reads

$$u_3 = (-1) \times \begin{pmatrix} \mathrm{pr}\{x_1 = \mathrm{T}\}\mathrm{pr}\{x_2 = \mathrm{L}\} + \mathrm{pr}\{x_1 = \mathrm{M}\}\mathrm{pr}\{x_2 = \mathrm{L}\} \\ + \mathrm{pr}\{x_1 = \mathrm{M}\}\mathrm{pr}\{x_2 = \mathrm{M}\} + \mathrm{pr}\{x_1 = \mathrm{B}\}\mathrm{pr}\{x_2 = \mathrm{M}\} \\ + \mathrm{pr}\{x_1 = \mathrm{T}\}\mathrm{pr}\{x_2 = \mathrm{R}\} + \mathrm{pr}\{x_1 = \mathrm{B}\}\mathrm{pr}\{x_2 = \mathrm{R}\} \end{pmatrix}. \tag{18}$$

[18] However, in contrast to the case of two-person zero-sum games, this observation does not imply a minmax theorem for zero-sum networks. For a generalization of the minmax theorem to zero-sum networks, the reader is referred to Cai et al. (2016).

[19] An alternative way to establish Proposition 1 would involve showing that any zero-sum network (considered as a population game with unit-sized populations) is *stable* in the sense of Hofbauer and Sandholm (2009). We are grateful for Josef Hofbauer for pointing this connection out to us.

[20] Here and elsewhere in the paper, payoff vectors are arranged diagonally in each box, starting with player 1's payoff in the respective upper-left corner.

[21] That conclusion does not depend on the fact that player 3 has only one strategy. In fact, we have constructed (details omitted) an example of a $2 \times 2 \times 2$ zero-sum game with a Shapley hexagon, similar to the three-person non-zero-sum example of Jordan (1993). On a related note, we conjecture that Example 1 as well as the later examples of the present paper could be made robust by introducing additional strategies for all players.

	$x_3 = O$		
	$x_2 = L$	$x_2 = M$	$x_2 = R$
$x_1 = T$	1, 0, -1	0, 0, 0	0, 1, -1
$x_1 = M$	0, 1, -1	1, 0, -1	0, 0, 0
$x_1 = B$	0, 0, 0	0, 1, -1	1, 0, -1

Fig. 2. The game G^1.

Thus, to obtain the conclusion of Proposition 1, it does not in general suffice to assume that payoffs may be represented as a sum of bilinear terms.[22]

While the network assumption in Proposition 1 cannot be easily dropped, the zero-sum assumption may be relaxed to a certain extent for convergence in network games, as will be discussed in the next section.

4. Additional classes of network games

4.1. Conflicts

A bilateral game G_{ij} will be called a *conflict* if there exist valuations $v_{ij} > 0$ and $v_{ji} > 0$, success functions $p_{ij} : S_{ij} \times S_{ji} \to [0, 1]$ and $p_{ji} : S_{ji} \times S_{ij} \to [0, 1]$, as well as cost functions $c_{ij} : S_{ij} \to \mathbb{R}$ and $c_{ji} : S_{ji} \to \mathbb{R}$ such that

$$u_{ij}(s_{ij}, s_{ji}) = p_{ij}(s_{ij}, s_{ji})v_{ij} - c_{ij}(s_{ij}), \tag{19}$$

$$u_{ji}(s_{ji}, s_{ij}) = p_{ji}(s_{ji}, s_{ij})v_{ji} - c_{ji}(s_{ji}), \tag{20}$$

and

$$p_{ij}(s_{ij}, s_{ji}) + p_{ji}(s_{ji}, s_{ij}) = 1 \tag{21}$$

hold for any $s_{ij} \in S_{ij}$ and $s_{ji} \in S_{ji}$. Examples include discrete variants of probabilistic contests (Tullock, 1980; Hirshleifer, 1989; Lazear and Rosen, 1981), the first-price all-pay auction (Baye et al., 1996), and Colonel Blotto games (Roberson, 2006).[23]

A network game G will be called a *conflict network* if the bilateral game G_{ij} is a conflict for each pair $(i, j) \in E$. We will say that a conflict network has *pairwise homogeneous valuations* if valuations of two players coincide in any pairwise conflict, i.e., if $v_{ij} = v_{ji}$ for any $(i, j) \in E$. An example is Franke and Öztürk's (2015) transfer network where the net valuations of winning a conflict are assumed identical across players. If valuations are not pairwise homogeneous, we will (somewhat loosely) say that the conflict network exhibits heterogeneous valuations.

Proposition 1 can be extended to these additional classes of games as follows.

Proposition 2. *Let G be either (i) a conflict network with pairwise homogeneous valuations, or (ii) an acyclic conflict network. Then, G has the CTFP property.*

Proof. *(i)* Starting from the conflict network G, we construct another network game $\widetilde{G}$ on the same graph and with identical strategy sets by letting payoffs in the bilateral game $\widetilde{G}_{ij}$ be given by

$$\widetilde{u}_{ij}(s_{ij}, s_{ji}) = u_{ij}(s_{ij}, s_{ji}) - \frac{v_{ij}}{2} + c_{ji}(s_{ji}). \tag{22}$$

[22] To resolve the apparent contradiction, it suffices to note that in a network game, any bilinear term arising in u_3 would have to multiply a probability of a pure strategy for player 3 with a probability of a pure strategy for another player. In relationship (18), however, u_3 contains bilinear terms that combine a probability of a pure strategy for player 1 with a probability of a pure strategy for player 2, which is inconsistent with the definition of a network game. We are grateful to one of the anonymous reviewers for hinting towards this distinction.

[23] The class of conflicts defined above corresponds precisely to the class of *strategically zero-sum games* (Moulin and Vial, 1978). However, we will use the terminology of conflict because it is more suggestive and because it allows making an important distinction (homogeneous vs. heterogeneous valuations) that is absent from the theory of strategically zero-sum games but crucially needed below.

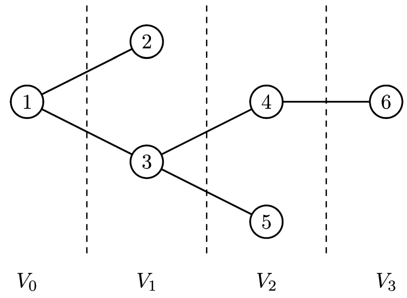

Fig. 3. Partitioning the set of nodes in an acyclic network.

Then, clearly, player i's payoff function in $\widetilde{G}$ reads

$$\widetilde{u}_i(x_i, x_{N(i)}) = \sum_{j \in N(i)} \widetilde{u}_{ij}(\sigma_{ij}(x_i), \sigma_{ji}(x_j)) \tag{23}$$

$$= u_i(x_i, x_{N(i)}) - \sum_{j \in N(i)} \left\{ \frac{v_{ij}}{2} - c_{ji}(\sigma_{ji}(x_j)) \right\}, \tag{24}$$

which shows that $\widetilde{G}$ is best-response equivalent in mixed strategies to G.[24] We claim that each bilateral game $\widetilde{G}_{ij}$ is two-person zero-sum. Indeed, for any pair $(i, j) \in E$, using (19)-(21) and $v_{ij} = v_{ji}$, we have

$$\widetilde{u}_{ij}(s_{ij}, s_{ji}) + \widetilde{u}_{ji}(s_{ji}, s_{ij})$$

$$= u_{ij}(s_{ij}, s_{ji}) - \frac{v_{ij}}{2} + c_{ji}(s_{ji}) + u_{ji}(s_{ji}, s_{ij}) - \frac{v_{ji}}{2} + c_{ij}(s_{ij}) \tag{25}$$

$$= p_{ij}(s_{ij}, s_{ji})v_{ij} - c_{ij}(s_{ij}) - \frac{v_{ij}}{2} + c_{ji}(s_{ji})$$

$$+ p_{ji}(s_{ji}, s_{ij})v_{ji} - c_{ji}(s_{ji}) - \frac{v_{ji}}{2} + c_{ij}(s_{ij}) \tag{26}$$

$$= 0, \tag{27}$$

which proves the claim. As a result, $\widetilde{G}$ is a zero-sum network. By Proposition 1, CTFP converges in $\widetilde{G}$. Using the best-response equivalence in mixed strategies, both the set of continuous-time fictitious plays and the set of Nash equilibria are the same for G and $\widetilde{G}$. Hence, CTFP converges also in G.

(ii) Without loss of generality, the network may be assumed to be connected. The set of players can then be partitioned into a finite number of subsets $V_0, V_1, \ldots, V_L$, where $i \in V_\ell$ for $\ell \in \{0, \ldots, L\}$ if and only if the graph-theoretic distance between players 1 and i equals ℓ. See Fig. 3 for illustration.

Thus,

$$V_0 = \{1\}, \tag{28}$$

$$V_1 = N(1), \text{ and} \tag{29}$$

$$V_\ell = \{\textstyle\bigcup_{i \in V_{\ell-1}} N(i)\} \backslash V_{\ell-2} \qquad (\ell = 2, \ldots, L). \tag{30}$$

We will describe an iterative construction that transforms the conflict network with bilateral payoff function u_{ij} and *heterogeneous* valuations v_{ij} into a conflict network with bilateral payoff functions $\widehat{u}_{ij}$ and *homogeneous* valuations $\widehat{v}_{ij}$. For this, we initialize the iteration by letting $\widehat{u}_{1j} = u_{1j}$ and $\widehat{v}_{1j} = v_{1j}$ for any neighbor $j \in N(1)$. We start now with $\ell = 1$ and consider some player $j \in V_\ell$. Since the network is acyclic, there is precisely one player $i \in V_{\ell-1}$ such that $j \in N(i)$. We rescale player j's payoff function u_{jk} in her relationship with any neighbor $k \in N(j)$ by letting

$$\widehat{u}_{jk}(s_{jk}, s_{kj}) = \frac{\widehat{v}_{ij}}{v_{ji}} \cdot u_{jk}(s_{ki}, s_{ik}), \tag{31}$$

so that player j's valuation becomes

[24] We call two network games G and $\widetilde{G}$ *best-response equivalent in mixed strategies* (Monderer and Shapley, 1996b; Morris and Ui, 2004) if for any $i \in \{1, \ldots, n\}$ and any $\mu_{N(i)} \in \Delta(X_{N(i)})$, we have $\arg\max_{\mu_i \in \Delta(X_i)} u_i(\mu_i, \mu_{N(i)}) = \arg\max_{\mu_i \in \Delta(X_i)} \widetilde{u}_i(\mu_i, \mu_{N(i)})$. Because of the network structure, it actually suffices to check the condition for any $\mu_{N(i)} \in \bigtimes_{j \in N(i)} \Delta(X_j)$, rather than for any $\mu_{N(i)} \in \Delta(X_{N(i)})$ (cf. Section 6).

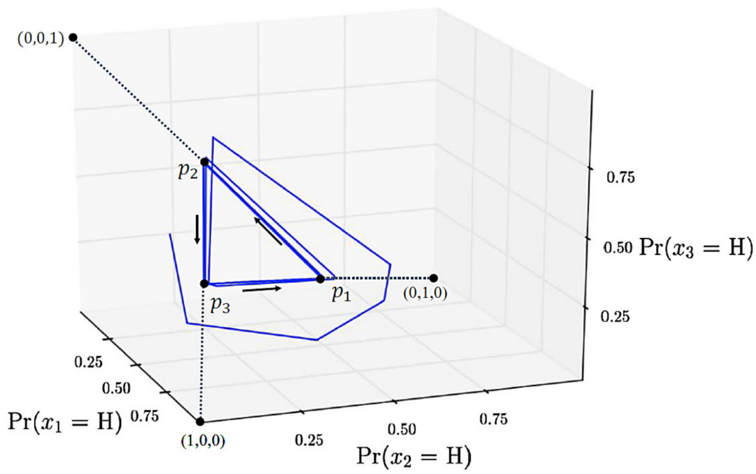

Fig. 4. Fictitious play need not converge in a network of conflicts.

$$\widehat{v}_{jk} = \frac{\widehat{v}_{ij}}{v_{ji}} \cdot v_{jk}. \tag{32}$$

Since the factor $(\widehat{v}_{ij}/v_{ji})$ does not depend on the neighbor $k \in N(j)$, such rescaling does not affect player j's multilateral best-response correspondence. Moreover, in the special case $k = i$, the resulting bilateral game $\widehat{G}_{ij}$ (with payoff functions $\widehat{u}_{ij}$ for player i and $\widehat{u}_{ji}$ for player j) is a conflict with homogeneous valuations, since $\widehat{v}_{ji} = \widehat{v}_{ij}$ by equation (32). Once this is accomplished for any $j \in V_\ell$, the running index ℓ is incremented, and the rescaling procedure repeated. After the iteration has reached $\ell = L$, we end up with a conflict network $\widehat{G}$ with pairwise homogeneous valuations that is best-response equivalent in mixed strategies to G. Convergence of CTFP follows, therefore, from part (i). □

The intuition is as follows. If valuations are homogeneous, then each bilateral game is essentially a constant-sum game with costly strategies.[25] Suppose that the cost of a player in any bilateral conflict is not lost, but reaches the other player as a subsidy. Then, as the size of the subsidy does not depend on the player's choice of strategy, her best-response correspondence remains unaffected. However, if the subsidy is implemented in any bilateral conflict, the network game becomes constant-sum, and we are done. If valuations are heterogeneous, and the underlying network structure is acyclic, then the conflict network can be transformed into a conflict network with pairwise homogeneous valuations by a simple iteration that starts at player 1 and works its way through the tree, where in each step, valuations and cost functions of a new set of players are rescaled so as to render the backward-looking conflict homogeneous.

However, rescaling does not work in general conflict networks. Indeed, as the following example illustrates, fictitious play need not converge in cyclic conflict networks with heterogeneous valuations.

Example 2 *(Network of conflicts).* Suppose there are three players $i = 1, 2, 3$, where $X_1 = X_2 = X_3 = \{\mathrm{L}, \mathrm{H}\}$. The network structure is a triangle. In the network game G^2, each bilateral game is a conflict with heterogeneous valuations. Specifically, one assumes

$$v_{i,i+1} = 6,\ v_{i,i-1} = 3, \tag{33}$$

$$c_{i,i+1}(\mathrm{L}) = c_{i,i-1}(\mathrm{L}) = 0,\ c_{i,i+1}(\mathrm{H}) = c_{i,i-1}(\mathrm{H}) = 1 \tag{34}$$

$$p_{i,i+1}(\mathrm{H},\mathrm{L}) = \tfrac{3}{4},\ p_{i,i+1}(\mathrm{L},\mathrm{H}) = \tfrac{1}{3},\ p_{i,i+1}(\mathrm{L},\mathrm{L}) = p_{i,i+1}(\mathrm{H},\mathrm{H}) = \tfrac{1}{2}, \tag{35}$$

where $i+1$ refers to player 1 if $i = 3$, and similarly, $i-1$ refers to player 3 if $i = 1$. Denote by $r_i = \mathrm{pr}\{x_i = \mathrm{H}\}$ the probability that player i uses strategy H. There is a unique Nash equilibrium $(r_1^*, r_2^*, r_3^*) = (0, 0, 0)$.[26] Moreover, fictitious play need not converge to equilibrium, but may follow a triangle-shaped path that runs in a round-robin fashion through the points

$$\cdots \to p_1 = (\tfrac{2}{7}, \tfrac{4}{7}, \tfrac{1}{7}) \to p_2 = (\tfrac{1}{7}, \tfrac{2}{7}, \tfrac{4}{7}) \to p_3 = (\tfrac{4}{7}, \tfrac{1}{7}, \tfrac{2}{7}) \to \ldots, \tag{36}$$

as illustrated in Fig. 4. On the linear segment from p_1 to p_2, for instance, players 1 and 2 each choose L, while player 3 chooses H, so that the process moves in the direction of the corner point $(0, 0, 1)$. A similar logic applies to the other two

[25] This useful interpretation is borrowed from Ben-Sasson et al. (2007).

[26] So all players would choose the effort level L. For proofs, see Appendix B.

segments of the triangle. Thus, the dynamic conflict does not settle down, which shows that the assumption of homogeneous valuations cannot be dropped in general.

4.2. Potential games

We introduce notions of increasing flexibility first for bilateral games and then for network games. A bilateral game G_{ij} is said to possess *identical payoff functions* if $u_{ij}(s_{ij}, s_{ji}) = u_{ji}(s_{ji}, s_{ij})$ for all $s_{ij} \in S_{ij}$ and $s_{ji} \in S_{ji}$. A bilateral game G_{ij} is an *exact potential game* (Monderer and Shapley, 1996a) if there exists a potential function $P_{ij} : S_{ij} \times S_{ji} \to \mathbb{R}$ such that

$$u_{ij}(s_{ij}, s_{ji}) - u_{ij}(\widehat{s}_{ij}, s_{ji}) = P_{ij}(s_{ij}, s_{ji}) - P_{ij}(\widehat{s}_{ij}, s_{ji}), \tag{37}$$

$$u_{ji}(s_{ji}, s_{ij}) - u_{ji}(\widehat{s}_{ji}, s_{ij}) = P_{ij}(s_{ij}, s_{ji}) - P_{ij}(s_{ij}, \widehat{s}_{ji}), \tag{38}$$

for all $s_{ij}, \widehat{s}_{ij} \in S_{ij}$ and $s_{ji}, \widehat{s}_{ji} \in S_{ji}$. Next, a bilateral game G_{ij} is a *weighted potential game* (Monderer and Shapley, 1996a) if there exists a potential function $P_{ij} : S_{ij} \times S_{ji} \to \mathbb{R}$ as well as weights $w_{ij} > 0$ and $w_{ji} > 0$ such that

$$u_{ij}(s_{ij}, s_{ji}) - u_{ij}(\widehat{s}_{ij}, s_{ji}) = w_{ij}\left\{P_{ij}(s_{ij}, s_{ji}) - P_{ij}(\widehat{s}_{ij}, s_{ji})\right\}, \tag{39}$$

$$u_{ji}(s_{ji}, s_{ij}) - u_{ji}(\widehat{s}_{ji}, s_{ij}) = w_{ji}\left\{P_{ij}(s_{ij}, s_{ji}) - P_{ij}(s_{ij}, \widehat{s}_{ji})\right\}, \tag{40}$$

for all $s_{ij}, \widehat{s}_{ij} \in S_{ij}$ and $s_{ji}, \widehat{s}_{ji} \in S_{ji}$.[27] Potential games and weighted potential games belong to the class of games *with identical interests* (Monderer and Shapley, 1996b), i.e., they are best-response equivalent in mixed strategies to a game with identical payoff functions.[28] A network game G will be said to be an *exact potential network* if all bilateral games G_{ij} are exact potential games. Finally, a network game G will be referred to as a *weighted potential network* if all bilateral games G_{ij} are weighted potential games.

Cai and Daskalakis (2011) have shown that if all bilateral games in a network have identical payoff functions then the network game admits an exact potential (the welfare function). The following result applies their reasoning, but also offers some extensions because bilateral games may here be exact potential games or even weighted potential games. We also use a slightly different proof. Specifically, we construct the potential of the network game as the sum over all potentials rather than as the sum over all payoff functions. Moreover, in the case of weighted potential games, we employ a similar induction argument as in the proof of Proposition 2(ii).

Proposition 3. *Let G be either (i) an exact potential network, or (ii) a weighted potential network on an acyclic graph. Then, G has the CTFP property.*

Proof. *(i)* Suppose that G is an exact potential network. Since any bilateral game G_{ij} is an exact potential game, there exists a potential function $P_{ij} : S_{ij} \times S_{ji} \to \mathbb{R}$ such that equations (37) and (38) hold for any $s_{ij}, \widehat{s}_{ij} \in S_{ij}$ and $s_{ji}, \widehat{s}_{ji} \in S_{ji}$. In particular, by exchanging the roles of players i and j in equation (38) and comparing with (37), we obtain

$$P_{ij}(s_{ij}, s_{ji}) - P_{ij}(\widehat{s}_{ij}, s_{ji}) = P_{ji}(s_{ji}, s_{ij}) - P_{ji}(s_{ji}, \widehat{s}_{ij}). \tag{41}$$

Consider now the aggregate potential $\mathcal{P} : X \to \mathbb{R}$ defined through

$$\mathcal{P}(x) = \frac{1}{2} \sum_{(i,j)\in E} P_{ij}(s_{ij}, s_{ji}), \tag{42}$$

where $s_{ij} = \sigma_{ij}(x_i) \in S_{ij}$ and $s_{ji} = \sigma_{ji}(x_j) \in S_{ji}$. It is claimed that $\mathcal{P}$ is an exact potential for G. For this, fix a player $i \in \{1, \ldots, n\}$. Then, for any $x_i \in X_i$, $\widehat{x}_i \in X_i$, and $x_{-i} \in X_{-i}$, writing $\widehat{s}_{ij} = \sigma_{ij}(\widehat{x}_i) \in S_{ij}$, it is straightforward to check that

$$u_i(x_i, x_{-i}) - u_i(\widehat{x}_i, x_{-i})$$

$$\underset{(1)}{=} \sum_{j\in N(i)} \left\{u_{ij}(s_{ij}, s_{ji}) - u_{ij}(\widehat{s}_{ij}, s_{ji})\right\} \tag{43}$$

$$\underset{(37)}{=} \sum_{j\in N(i)} \left\{P_{ij}(s_{ij}, s_{ji}) - P_{ij}(\widehat{s}_{ij}, s_{ji})\right\} \tag{44}$$

[27] The following observation shows that the notation need not lead to confusion: Let G_{ij} be a weighted potential game with potential P_{ij} and weight w_{ij} for player i and weight w_{ji} for player j. Define a potential $P_{ji} : S_{ji} \times S_{ij} \to \mathbb{R}$ by $P_{ji}(s_{ji}, s_{ij}) = P_{ij}(s_{ij}, s_{ji})$. Then, the bilateral game G_{ji} (i.e., the game G_{ij} with the roles of players i and j exchanged) is a weighted potential game with potential P_{ji} and weight w_{ji} for player j and weight w_{ij} for player i.

[28] Thus, games with identical interests relate to games with identical payoff functions in a similar way as strategically zero-sum games relate to zero-sum games.

	$x_2 = H$	$x_2 = L$
$x_1 = H$	1, 5	0, 0
$x_1 = L$	0, 0	1, 5

	$x_3 = H$	$x_3 = L$
$x_2 = H$	1, 5	0, 0
$x_2 = L$	0, 0	1, 5

	$x_1 = H$	$x_1 = L$
$x_3 = H$	0, 0	1, 5
$x_3 = L$	1, 5	0, 0

Fig. 5. A network of weighted potential games.

$$\underset{(41)}{=} \frac{1}{2}\left\{\sum_{(i,j)\in E}\{P_{ij}(s_{ij},s_{ji}) - P_{ij}(\widehat{s}_{ij},s_{ji})\} + \sum_{(j,i)\in E}\{P_{ji}(s_{ji},s_{ij}) - P_{ji}(s_{ji},\widehat{s}_{ij})\}\right\} \tag{45}$$

$$= \frac{1}{2}\left\{\sum_{(i,j)\in E} P_{ij}(s_{ij},s_{ji}) + \sum_{(j,i)\in E} P_{ji}(s_{ji},s_{ij}) + \sum_{\substack{(j,k)\in E\\ j\neq i\neq k}} P_{jk}(s_{jk},s_{kj})\right\}$$

$$- \frac{1}{2}\left\{\sum_{(i,j)\in E} P_{ij}(\widehat{s}_{ij},s_{ji}) + \sum_{(j,i)\in E} P_{ji}(s_{ji},\widehat{s}_{ij}) + \sum_{\substack{(j,k)\in E\\ j\neq i\neq k}} P_{jk}(s_{jk},s_{kj})\right\} \tag{46}$$

$$\underset{(42)}{=} \mathcal{P}(x_i,x_{-i}) - \mathcal{P}(\widehat{x}_i,x_{-i}). \tag{47}$$

Hence, $\mathcal{P}$ is indeed an exact potential for the n-person game G. Therefore, from Lemma A.2 in the Appendix, any CTFP1 converges to the set of Nash equilibria. Since every CTFP is, in particular, a CTFP1, also any CTFP converges to the set of Nash equilibria.

(ii) From the weighted potential network G, an exact potential network $\widehat{G}$ is constructed as follows. We start by assigning players to subsets $V_0, V_1, \ldots, V_L$ according to their distance ℓ from player 1, as in the proof of Proposition 2. Then, we initiate an iteration by letting $\ell = 1$. Consider any player $j \in V_\ell$, and recall that there is precisely one player $i \in V_{\ell-1}$ such that $j \in N(i)$. By assumption, the bilateral game G_{ij} is a weighted potential game. Therefore, there exists a potential function $P_{ij} : S_{ij} \times S_{ji} \to \mathbb{R}$, as well as weights $w_{ij} > 0$ and $w_{ji} > 0$ such that (39) and (40) hold for any $s_{ij}, \widehat{s}_{ij} \in S_{ij}$ and $s_{ji}, \widehat{s}_{ji} \in S_{ji}$. Note that P_{ij} may be chosen such that $w_{ij} = 1$. Given this normalization, we rescale all bilateral payoff functions of player j by letting $\widehat{u}_{jk} = u_{jk}/w_{ji}$, for any $k \in N(j)$. Then, the bilateral game $\widehat{G}_{ij}$ with payoff functions $\widehat{u}_{ij} = u_{ij}$ for player i and $\widehat{u}_{ji} = u_{ji}/w_{ji}$ for player j is easily seen to admit the exact potential P_{ij}. Moreover, any bilateral game G_{jk} with $k \neq i$ remains a weighted potential game when u_{jk} is replaced by $\widehat{u}_{jk}$. All players $j \in V_\ell$ are dealt with in this fashion. Then, the index ℓ is incremented, and the iteration continued. When the iteration ends at $\ell = L$, the bilateral games $\widehat{G}_{ij}$ form an exact potential network $\widehat{G}$. Moreover, since each player's payoff has merely been rescaled, $\widehat{G}$ is best-response equivalent in mixed strategies to the weighted potential network we started from. Hence, invoking part (i), CTFP converges to the set of Nash equilibria. □

The intuition of the first part is simple. Because bilateral games possess exact potentials, the aggregate potential reflects incentives precisely as the network game.[29] The intuition of the second part is very similar to Proposition 2(ii).

Sela (1999, Prop. 12) has shown that a star-shaped network of generic weighted potential two-by-two games is, when reduced to a two-person game, best-response equivalent in mixed strategies to a game with identical payoff functions. This implies the CTFP property. Proposition 3 shows that the "one-against-all" assumption, the assumption on the number of strategies, and the genericity of payoffs may be dropped without weakening the conclusion.

The following example shows that a general network consisting of weighted potential games need not have the fictitious-play property.

Example 3 *(Network of weighted potential games).* Consider the following game G^3 between three players $i = 1, 2, 3$, each of them having two compound strategies, i.e., $X_1 = X_2 = X_3 = \{H, L\}$. Bilateral payoffs are specified in Fig. 5.

Thus, each player is involved in two coordination games, where for player i, coordination with player $i - 1$, is more valuable than coordination with player $i + 1$.[30] Moreover, there is a twist in the coordination between players 1 and 3.

[29] Extending the proof, one can easily convince oneself that arbitrary networks (i.e., hypergraphs) of multiplayer exact potential games admit an exact potential.

[30] We use the same notational conventions as in the previous example. Further, all proofs related to Example 3 can be found in Appendix B.

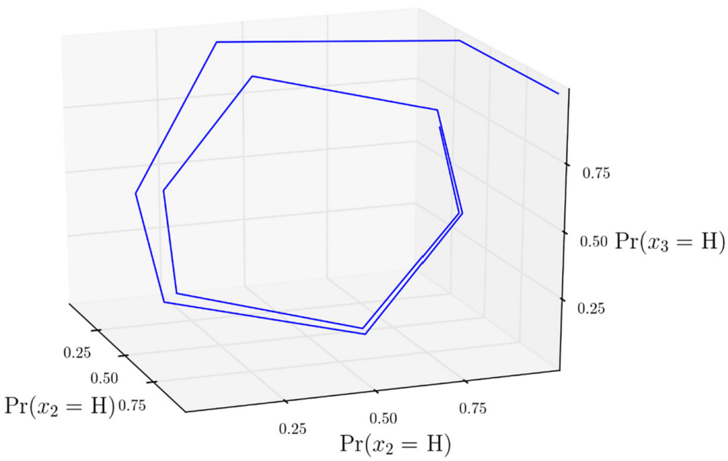

Fig. 6. A network of weighted potential games without the CTFP property.

The game G^3 admits a unique Nash equilibrium in which each player randomizes with equal probability over her two alternatives. However, CTFP runs indefinitely through the hexagonal cycle

$$\begin{aligned} \ldots &\to p_1 = (a, b, c) \to p_2 = (1-c, a, b) \to p_3 = (1-b, 1-c, a) \to \\ &\to p_4 = (1-a, 1-b, 1-c) \to p_5 = (c, 1-a, 1-b) \to p_6 = (b, c, 1-a) \to \ldots, \end{aligned} \tag{48}$$

where $(a, b, c) = (\frac{4}{9}, \frac{8}{9}, \frac{7}{9})$, and the entries correspond to the respective probabilities of choosing H. Fig. 6 shows a numerical fictitious-play path approaching the hexagon.

In contrast to the case of zero-sum networks, we make no claims whatsoever regarding the rate of convergence in networks of weighted potential games.[31]

4.3. Two-by-two games

Finally, we consider two-by-two games, i.e., two-person games in which each player has just two strategies.[32] It will be recalled (e.g., Krishna and Sjöström, 1997, Prop. 3) that any two-by-two game without weakly dominated or identical strategies is best-response equivalent in mixed strategies to either a zero-sum game or to a game with identical payoff functions. Miyasawa's theorem is therefore customarily presented as a corollary of the corresponding results for those classes of games. The approach will be similar here.

A network game G will be called a *network of strategically similar two-by-two games* if (i) X_i has precisely two elements, for any $i \in \{1, \ldots, n\}$, and (ii) either all bilateral games G_{ij} are conflicts with pairwise homogeneous valuations, or all bilateral games G_{ij} admit an exact potential.

Proposition 4. *Let G be a network of strategically similar two-by-two games. Then, G has the CTFP property.*

Proof. Immediate from Propositions 2 and 3. □

Proposition 4 extends Miyasawa's theorem to arbitrary network structures. As seen above, the assumptions on the bilateral games can be relaxed if the underlying network structure is acyclic. This implies, in particular, a related result by Sela (1999, Cor. 13) for star-shaped networks. However, Examples 2 and 3 show that it is not possible to generalize Proposition 4 to arbitrary networks of strategically zero-sum games, nor to arbitrary networks of weighted potential games.[33]

---

[31] However, a recent paper by Swenson and Kar (2017) might throw some light on this difficult question.

[32] Miyasawa (1961) has shown that (discrete-time) fictitious play converges in every two-by-two game. That result was later seen to depend on a particular tie-breaking rule (Monderer and Sela, 1996). With generic payoffs, however, the theorem holds. For further discussion of this important special case, see Metrick and Polak (1994), Monderer and Sela (1997), and Sela (1999).

[33] Similarly, it does not seem possible to obtain a general convergence result in mixed networks, i.e., in networks where some bilateral games are zero-sum, while others reflect identical interests.

5. Discrete-time fictitious play

While the continuous-time variant of fictitious play considered above is analytically more convenient, there are reasons to be interested also in the discrete-time variant. For instance, the first major result in the literature by Robinson (1951) concerned the discrete-time process in two-person zero-sum games. It has often been suggested that the two processes should behave similarly. This is also intuitive because the incremental changes in the discrete-time process become smaller and smaller over time. Harris (1998) has developed a very useful approach that, indeed, allows transferring results for the continuous-time case to the discrete-time case. Below, we will use his approach to extend some of our conclusions to the case of DTFP.[34]

In contrast to the analysis so far, time progresses now in stages. At any given stage $t \in \mathbb{N}_0 = \{0, 1, 2, \ldots\}$, each player $i \in \{1, \ldots, n\}$ is assumed to choose a pure multilateral strategy $y_i(t) \in X_i$.[35] We denote by $\check{y}_i(t) \in \Delta(X_i)$ the Dirac measure on X_i that places all probability weight on $y_i(t)$. Then, the empirical frequencies of pure-strategy choices made by player i before stage $t \in \mathbb{N} = \{1, 2, 3, \ldots\}$ may be summarized in the *discrete-time independent average*

$$\alpha_i^d(t) \equiv \alpha_i^d(t, y(.)) = \frac{1}{t} \sum_{t'=0}^{t-1} \check{y}_i(t'). \tag{49}$$

Let $\alpha^d(t) \in \Delta(X_1) \times \cdots \times \Delta(X_n)$ denote the (column) vector whose i-th entry is $\alpha_i^d(t)$. For any profile $\mu_{N(i)} \in \bigtimes_{j \in N(i)} \Delta(X_j)$, we will write $\mathrm{BR}_i(\mu_{N(i)})$ for the set of pure strategies $y_i \in X_i$ such that the corresponding Dirac measure $\check{y}_i \in \Delta(X_i)$ satisfies $\check{y}_i \in \mathrm{MBR}_i(\mu_{N(i)})$. Further, for a profile of mixed strategies $\mu = (\mu_1, \ldots, \mu_n) \in \bigtimes_{i=1}^n \Delta(X_i)$, let $\mathrm{BR}(\mu) = \bigtimes_{i=1}^n \mathrm{BR}_i(\mu_{N(i)})$.

Definition 3 *(DTFP).* A discrete-time fictitious play in the network game G is a sequence of multilateral pure-strategy profiles $y(.) = \{y(t)\}_{t=0}^\infty$ such that, for some $t^* \in \mathbb{N}$, it holds that $y(t) \in \mathrm{BR}(\alpha^d(t, y(.)))$ for any $t \geq t^*$.

Thus, in a DTFP $y(.)$, there is a stage $t^* \geq 1$ from which onwards players optimize against historical averages of their neighbors' past behavior. However, there is no restriction on choices before stage t^*.

For a given DTFP $y(.) = \{y(t)\}_{t=0}^\infty$, the corresponding sequence of independent averages $\{\alpha^d(t)\}_{t=1}^\infty$ defined in (49) is a sequence in $\Delta(X_1) \times \cdots \times \Delta(X_n)$. We denote by $\mathcal{A}^d(y(.))$ the set of all accumulation points of $\{\alpha^d(t)\}_{t=1}^\infty$.

Definition 4. A sequence of multilateral pure-strategy profiles $y(.) = \{y(t)\}_{t=0}^\infty$ is said to converge to the set of Nash equilibria if every limit distribution $\mu^* \in \mathcal{A}^d(y(.))$ is a Nash equilibrium in G.

In analogy to the continuous case, we will say that G has the *discrete-time fictitious-play property* if any DTFP in G converges to the set of Nash equilibria. The following result generalizes Robinson's (1951) theorem to n-person zero-sum networks.

Proposition 5. *Any zero-sum network G has the DTFP property.*

Proof. The proof has three parts. It is shown first that, for any DTFP, the corresponding process of discrete-time independent averages solves a particular difference inclusion. Subsequently, we verify that the set of Nash equilibria in G is a global uniform attractor[36] of the best-response population dynamics (Gilboa and Matsui, 1991; Matsui, 1992). Finally, we apply an approximation result of Hofbauer and Sorin (2006) to conclude that any accumulation point of the process of discrete-time independent averages is a Nash equilibrium.

(i) Take a zero-sum network G and a DTFP process $y(.) = \{y(t)\}_{t=0}^\infty$ in G. Thus, $y(.) = \{y(t)\}_{t=0}^\infty$ is a sequence in X such that, for some $t^* \in \mathbb{N}$, it holds that $y(t) \in \mathrm{BR}(\alpha^d(t, y(.)))$ for any $t \geq t^*$. Fix some player $i \in \{1, \ldots, n\}$ and some $t \geq t^*$. Then, by simple algebraic manipulation,

$$\alpha_i^d(t+1) = \frac{1}{t+1} \sum_{t'=0}^{t} \check{y}_i(t') \tag{50}$$

[34] The idea is to change the time scale, such that the differential inclusion defining the continuous-time process becomes *autonomous* and, in fact, equivalent to the best-response population dynamics, as in Gilboa and Matsui (1991) and Matsui (1992). Thereby, as detailed in Harris (1998) and Hofbauer and Sorin (2006), it is feasible to exploit the near-convergence of a sufficiently delayed discrete-time process. That method of proof extends to arbitrary finite n-person games in which fictitious play converges uniformly across initial conditions (cf. Hofbauer, 1995, p. 23). It follows from the proof of Proposition 1 that the class of zero-sum networks satisfies this condition.

[35] It can be checked that our results hold likewise, with essentially unchanged proofs, when DTFP is defined in terms of mixed strategies.

[36] See Appendix A for definitions and additional background.

$$= \frac{1}{t+1}\check{y}_i(t) + \frac{1}{t+1}\sum_{t'=0}^{t-1}\check{y}_i(t') \tag{51}$$

$$= \frac{1}{t+1}\check{y}_i(t) + \frac{t}{t+1}\alpha_i^d(t). \tag{52}$$

Hence, recalling that $\check{y}_i(t) \in \mathrm{MBR}_i(\alpha^d_{N(i)}(t))$, the discrete-time process $\{\alpha^d(t)\}_{t=1}^{\infty}$ of independent averages is seen to solve the difference inclusion

$$\alpha^d(t+1) \in \frac{1}{t+1}\mathrm{MBR}(\alpha^d(t)) + \frac{t}{t+1}\alpha^d(t), \tag{53}$$

for any $t \geq t^*$.

(ii) Let $Z_G \subseteq \times_{i=1}^n \Delta(X_i)$ denote the set of Nash equilibria in G. We claim that Z_G is a global uniform attractor of the differential inclusion

$$\dot{z} \in \mathrm{MBR}(z) - z. \tag{54}$$

To see this, let $\varepsilon > 0$, and take some solution z of (54). Then, $z : [0,\infty) \to \times_{i=1}^n \Delta(X_i)$ is an absolutely continuous mapping satisfying

$$\frac{\partial z(\widetilde{\tau})}{\partial \widetilde{\tau}} \in \mathrm{MBR}(z(\widetilde{\tau})) - z(\widetilde{\tau}) \tag{55}$$

at all points in time $\widetilde{\tau} \in [0,\infty)$ at which z is differentiable. Let now $\tau = \exp(\widetilde{\tau})$, and define the *rescaled solution* $\zeta : [1,\infty) \to \times_{i=1}^n \Delta(X_i)$ associated with z by the relationship $\zeta(\tau) = z(\widetilde{\tau})$. Then, clearly, $\partial z(\widetilde{\tau})/\partial\widetilde{\tau} = \tau \cdot (\partial\zeta(\tau)/\partial\tau)$. Using (55), we get

$$\frac{\partial(\tau\zeta(\tau))}{\partial\tau} = \tau \cdot \frac{\partial\zeta(\tau)}{\partial\tau} + \zeta(\tau) \in \mathrm{MBR}(\zeta(\tau)) \tag{56}$$

at all $\tau \geq 1$ where ζ is differentiable. Since ζ is absolutely continuous, the set $\mathcal{N}_\zeta$ of points where ζ is not differentiable has measure zero. Choose now a path $m : [0,\infty) \to \times_{i=1}^n \Delta(X_i)$ satisfying $m(\tau) = \zeta(1)$ for $\tau \in [0,1)$, $m(\tau) = \partial(\tau\zeta(\tau))/\partial\tau$ for $\tau \in [1,\infty)\backslash\mathcal{N}_\zeta$, and $m(\tau) \in \mathrm{MBR}(\zeta(\tau))$ for $\tau \in \mathcal{N}_\zeta$. Then, from equation (2), for any $i \in \{1,\ldots,n\}$ and any $\tau \geq 1$,

$$\alpha_i(\tau, m) = \frac{1}{\tau}\left\{\zeta_i(1) + \int_1^\tau m_i(\tau')d\tau'\right\} = \zeta_i(\tau). \tag{57}$$

Hence, from (56), we see that $m(\tau) \in \mathrm{MBR}(\alpha(\tau,m))$ for any $\tau \geq 1$. Thus, m is a CTFP. By Corollary 1, $\mathcal{L}(\alpha(\tau,m))$ declines at rate $1/\tau$. Moreover, since $\mathcal{L}$ is bounded, there exists a constant C, *independent* of m, such that

$$\mathcal{L}(\alpha(\tau,m)) \leq \frac{C}{\tau} \qquad (\tau \geq 1). \tag{58}$$

To provoke a contradiction, suppose there is a sequence $\{\widetilde{\tau}_\kappa\}_{\kappa=1}^\infty$ in $[0,\infty)$ with $\lim_{\kappa\to\infty}\widetilde{\tau}_\kappa = \infty$, and a sequence $\{z^{(\kappa)}\}_{\kappa=1}^\infty$ of solutions of (54), such that

$$\underline{d}(z^{(\kappa)}(\widetilde{\tau}_\kappa), Z_G) > \varepsilon \qquad (\kappa \in \mathbb{N}), \tag{59}$$

where $\underline{d}(.,.)$ denotes the Hausdorff distance (see Appendix A for a definition). Let $\tau_\kappa = \exp(\widetilde{\tau}_\kappa)$, and denote by $\zeta^{(\kappa)}$ the rescaled solution associated with $z^{(\kappa)}$. Then, from (57) and (58),

$$\mathcal{L}(\zeta^{(\kappa)}(\tau_\kappa)) \leq \frac{C}{\tau_\kappa} \qquad (\kappa \in \mathbb{N}). \tag{60}$$

By the compactness of $\Delta(X_1) \times \cdots \times \Delta(X_n)$, the sequence $\{\zeta^{(\kappa)}(\tau_\kappa)\}_{\kappa=1}^\infty = \{z^{(\kappa)}(\widetilde{\tau}_\kappa)\}_{\kappa=1}^\infty$ has a converging subsequence. Denote the limit by μ^*. By (60), and the continuity of $\mathcal{L}$, we have $\mathcal{L}(\mu^*) = 0$. Hence, $\mu^* \in Z_G$ and, consequently, $\underline{d}(\mu^*, Z_G) = 0$. However, for *some* κ large enough, $d(z^{(\kappa)}(\widetilde{\tau}_\kappa), \mu^*) < \varepsilon$, so that via (59),

$$\underline{d}(\mu^*, Z_G) \geq \underline{d}(z^{(\kappa)}(\widetilde{\tau}_\kappa), Z_G) - d(z^{(\kappa)}(\widetilde{\tau}_\kappa), \mu^*) > 0. \tag{61}$$

The contradiction shows that, indeed, Z_G is a global uniform attractor of (54).

(iii) We have to show that $y(.) = \{y(t)\}_{t=0}^\infty$ converges to the set of Nash equilibria. For this, let $\mu^* \in \mathcal{A}^d(y(.))$ be an accumulation point of $\{\alpha^d(t, y(.))\}_{t=1}^\infty$. By part (i) of the proof, the difference inclusion

$$P_{t+1} \in \beta_t\Phi(P_t) + (1-\beta_t)P_t \qquad (t \in \mathbb{N} = \{1,2,\ldots\}), \tag{62}$$

with $\beta_t = 1/(t^* + t)$ for $t \in \mathbb{N}$, admits the solution

$$P_t = \alpha^d(t^* + t - 1) \qquad (t \in \mathbb{N}). \tag{63}$$

Moreover, by part (ii), Z_G is a global uniform attractor of (54). Applying Lemma A.3 from the Appendix, it follows that, for any $\varepsilon > 0$, there exists $t^\#(\varepsilon) \in \mathbb{N}$ such that, for any $t \geq t^\#(\varepsilon)$, we have $\underline{d}(P_t, Z_G) \leq \varepsilon$. By assumption, μ^* is an accumulation point of $\{\alpha^d(t)\}_{t=1}^\infty$, and consequently, also of the subsequence $\{P_t\}_{t=1}^\infty$. Hence, $d(\mu^*, P_t) < \varepsilon$ for infinitely many t. In particular, there exists $t_0 \geq t^\#(\varepsilon)$ such that $d(\mu^*, P_{t_0}) < \varepsilon$. Moreover, since $\underline{d}(P_{t_0}, Z_G) \leq \varepsilon$ and Z_G is compact, we find $\mu^{**} \in Z_G$ of G such that $d(P_{t_0}, \mu^{**}) \leq \varepsilon$. Combining these observations, the triangle inequality implies

$$\underline{d}(\mu^*, Z_G) \leq d(\mu^*, \mu^{**}) \tag{64}$$
$$\leq d(\mu^*, P_{t_0}) + d(P_{t_0}, \mu^{**}) \tag{65}$$
$$< 2\varepsilon. \tag{66}$$

Since this holds for any $\varepsilon > 0$, it follows that $\underline{d}(\mu^*, Z_G) = 0$. But Z_G is compact. Hence, $\mu^* \in Z_G$, which proves the proposition. □

The result above extends Sela's (1999, Prop. 7) result for star-shaped networks in three ways. First, the restriction regarding the network structure is dropped. Second, Proposition 5 imposes the zero-sum assumption only on the network, rather than on each of the bilateral games. Finally, no assumptions regarding tie-breaking are used here. In sum, it is feasible to address additional applications such as security games (Cai et al., 2016) and, as has been seen, conflict networks.

As in the case of two-person zero-sum games, the transformation of the CTFP process into a discrete-time process comes at a cost, which is the slower rate of convergence. More specifically, the discrete-time process is known to "overshoot," which makes it generally hard to nail down its rate of convergence. For two-person zero-sum games, Robinson's proof allows deriving an upper bound on the rate of convergence in payoffs (Shapiro, 1958). The resulting estimate is of order $O(t^{-1/(\nu_1+\nu_2-2)})$, where ν_i is the number of strategies for player $i = 1, 2$.[37] Given the lack of a direct extension of Robinson's proof to zero-sum networks, however, that upper bound is not easily generalized to zero-sum networks. Improving on Shapiro's upper bound, Karlin's strong conjecture says (or more precisely, said) that DTFP converges in payoffs at rate $O(t^{-1/2})$ in two-person zero-sum games, regardless of the number of strategies. Daskalakis and Pan (2014) have recently disproved that conjecture, showing that the rate of convergence in an asymmetric two-person zero-sum game in which both players have the same number of strategies ν may be as low as $O(t^{-1/\nu})$. That lower bound holds, obviously, also for zero-sum networks.[38]

The "Rosy Theorem" of Monderer et al. (1997, Th. A) says that a player's expected payoff in DTFP at any given stage is weakly higher than her average payoff experience. Using Robinson's theorem, this result implies that every DTFP in a two-person zero-sum game is *belief-affirming*, which means that, in the limit, the gap between expected payoffs at any given stage and the average payoff experience vanishes for each player. This is just the first part of Monderer et al. (1997, Th. B). Using Proposition 5, we can derive the following extension.

Corollary 2. *Let G be a zero-sum network. Then, any DTFP process $\{y(t)\}_{t=0}^\infty$ in G is belief-affirming.*

Proof. Let $\alpha^d(.)$ denote, as before, the independent discrete-time average of $\{y(t)\}_{t=0}^\infty$. As has been shown in the proof of Proposition 5, $\lim_{t\to\infty} \underline{d}(\alpha^d(t), Z_G) = 0$. Therefore, $\lim_{t\to\infty} \mathcal{L}(\alpha^d(t)) = 0$. Using (4)-(6), this implies

$$\lim_{t\to\infty} \left\{ \sum_{i=1}^{n} \max_{x_i \in X_i} u_i(x_i, \alpha^d_{N(i)}(t)) \right\} = 0. \tag{67}$$

By the "Rosy Theorem,"

$$\max_{x_i \in X_i} u_i(x_i, \alpha^d_{N(i)}(t)) \geq \frac{1}{t} \sum_{t'=0}^{t-1} u_i(y(t')), \tag{68}$$

for any $t \in \mathbb{N}$ and any player $i \in \{1, \ldots, n\}$. Hence,

$$\limsup_{t\to\infty} \left\{ \left\{ \max_{x_i \in X_i} u_i(x_i, \alpha^d_{N(i)}(t)) \right\} - \frac{1}{t} \sum_{t'=0}^{t-1} u_i(y(t')) \right\} \geq 0, \tag{69}$$

[37] If the game is symmetric, and consequently both players have the same number of strategies $\nu_1 = \nu_2 \equiv \nu$, then the upper bound may be sharpened to $O(t^{-1/(\nu-1)})$.

[38] For related work, see Gjerstad (1996), Conitzer (2009), and Brandt et al. (2013).

for any $i \in \{1, \ldots, n\}$. Suppose now that $\{y(t)\}_{t=0}^{\infty}$ is not belief-affirming. Then, for *some* player, inequality (69) holds strictly. Adding up across players, and exploiting the zero-sum property, this yields

$$\limsup_{t\to\infty} \left\{ \sum_{i=1}^{n} \max_{x_i \in X_i} u_i(x_i, \alpha_{N(i)}^d(t)) \right\} > 0, \tag{70}$$

in conflict with relationship (67). This proves the claim. □

However, there is an important difference to the case of two-person zero-sum games, which is that Nash equilibrium payoffs are not necessarily unique in zero-sum networks (cf. Cai et al., 2016). Hence, the second part of Monderer et al. (1997, Th. B), which states an equality between a player's long-run payoff experience and the value of the game, cannot be easily generalized to zero-sum networks.

Cai and Daskalakis (2011) have shown that, if every node in a network game plays a *no-regret* sequence of mixed strategies over sufficiently many stages, then the resulting frequency distributions over pure strategies form an ε-equilibrium. By definition, no-regret is a property of DTFP that is less stringent than being belief-affirming. Specifically, under no-regret behavior, a player's expected payoff against past play is asymptotically *weakly below* her average payoff experience, while in a belief-affirming DTFP process, a player's expected payoff is asymptotically *equal to* her average payoff experience. Therefore, Monderer et al. (1997, Th. B) implies that DTFP in two-person zero-sum games has no-regret. Given Corollary 2, the same conclusion holds for DTFP in zero-sum networks.[39]

The following should now be immediate.

Corollary 3. *Let G be a network game that satisfies the assumptions of any of the Propositions 1 through 4. Then, any DTFP process in G converges to the set of Nash equilibria.*

Proof. For zero-sum networks, the claim follows directly from Proposition 5. Since conflict networks with pairwise homogeneous valuations and acyclic conflict networks are best-response equivalent in mixed strategies to zero-sum networks, the claim holds also for these classes of games. Next, it was shown above that networks of exact potential games admit an exact potential. In this case, therefore, the claim follows from Monderer and Shapley (1996b, Th. A).[40] Finally, to deal with acyclic networks of weighted potential games, it suffices to recall that any such network game is best-response equivalent in mixed strategies to a network of exact potential games. □

6. The case of joint beliefs

So far, we assumed that players' beliefs are independent across neighbors. In a general network game, however, a player might observe correlations between the behavior of her neighbors.[41] In this section, we will explore the implications of assuming that players take account of such correlations.

Let $\widetilde{\mu}_{N(i)} \in \Delta(X_{N(i)})$ be player i's *joint belief* over strategies chosen by i's neighbors, where the tilde indicates that correlation is feasible. Player i's expected payoff from a mixed strategy $\mu_i \in \Delta(X_i)$ is written as $u_i(\mu_i, \widetilde{\mu}_{N(i)}) = E[u_i(x_i, x_{N(i)})]$, where the expectation is taken with respect to $(\mu_i, \widetilde{\mu}_{N(i)})$. Player i's mixed best-response correspondence MBR_i extends as usual to joint beliefs, $\widetilde{\mu}_{N(i)} \in \Delta(X_{N(i)})$, in the sense that $\text{MBR}_i(\widetilde{\mu}_{N(i)})$ is the set of mixed strategies $\mu_i^* \in \Delta(X_i)$ such that $u_i(\mu_i^*, \widetilde{\mu}_{N(i)}) = \max_{\mu_i \in \Delta(X_i)} u_i(\mu_i, \widetilde{\mu}_{N(i)})$. Similarly, the mixed best-response correspondence MBR of the network game G extends to arbitrary probability distributions $\widetilde{\mu} \in \Delta(X)$ by letting $\text{MBR}(\widetilde{\mu}) = \times_{i=1}^{n} \text{MBR}_i(\widetilde{\mu}_{N(i)})$, where $\widetilde{\mu}_{N(i)}$ denotes the marginal of $\widetilde{\mu}$ on $X_{N(i)}$. Let $m : [0, \infty) \to \Delta(X_1) \times \cdots \times \Delta(X_n)$ be a measurable path specifying each player i's mixed strategy at any point in time $\tau \geq 0$. Then, the *continuous-time joint average* $\widetilde{\alpha} : (0, \infty) \to \Delta(X)$ of the path m is defined at time $\tau > 0$ as

$$\widetilde{\alpha}(\tau) \equiv \widetilde{\alpha}(\tau, m) = \frac{1}{\tau} \int_0^{\tau} m(\tau')d\tau', \tag{71}$$

where the integral is now taken in $\Delta(X)$ rather than for each player separately. Similarly, if $\{y(t)\}_{t=0}^{\infty}$ is a sequence in X, we define the *discrete-time joint average* $\widetilde{\alpha}^d$ as

[39] One may even go one step further. Using the "Rosy Theorem," we see that DTFP exhibits no-regret if and only if it is belief-affirming. This observation might help to see that, despite the obvious similarity in the conclusion, there is no simple way of deducing Proposition 5 from existing results on the convergence of no-regret behavior.

[40] To see this, note that our definition of convergence to equilibrium is equivalent to the one used by Monderer and Shapley (1996b, pp. 260-261). Moreover, while their analysis formally restricts attention to the case $t^* = 1$, they mention that their convergence result holds more generally. Indeed, their proof extends in a straightforward way to $t^* \geq 1$.

[41] Indeed, correlation of the limit profile is a well-documented possibility (Young, 1993; Fudenberg and Kreps, 1993; Jordan, 1993; Fudenberg and Levine, 1998).

$$\widetilde{\alpha}^d(t) \equiv \widetilde{\alpha}^d(t, y(.)) = \frac{1}{t}\sum_{t'=0}^{t-1} \check{y}(t') \qquad (t = 1, 2, 3, \ldots), \tag{72}$$

where $\check{y}(t')$ denotes the Dirac measure in $\Delta(X)$ that assigns all probability weight to $y(t')$. The following definition of joint fictitious play should contain no surprises.

Definition 5 *($\widetilde{CTFP}$, $\widetilde{DTFP}$).* A continuous-time fictitious play with joint beliefs is a measurable mapping $m : [0, \infty) \to \times_{i=1}^n \Delta(X_i)$ such that $m(\tau) \in \mathrm{MBR}(\widetilde{\alpha}(\tau, m))$ for any $\tau \geq 1$.[42] Similarly, a discrete-time fictitious play with joint beliefs is a sequence $\{y(t)\}_{t=0}^{\infty}$ in $\times_{i=1}^n X_i$ such that, for some $t^* \in \mathbb{N}$, it holds that $y(t) \in \mathrm{BR}(\widetilde{\alpha}^d(t, y(.)))$ for any $t \geq t^*$.

Denote by $\widetilde{\mathcal{A}}(m)$ and $\widetilde{\mathcal{A}}^d(y(.))$, respectively, the set of all accumulation points of $\widetilde{\alpha}(.)$ and $\widetilde{\alpha}^d(.)$. The definition of convergence is adapted as follows.

Definition 6. We will say that the joint probability distribution $\widetilde{\mu}^* \in \Delta(X)$ is an observational equilibrium when each player i's marginal distribution $\widetilde{\mu}_i^* \in \Delta(X_i)$ is a mixed best response to the marginal $\widetilde{\mu}_{N(i)}^* \in \Delta(X_{N(i)})$, i.e., when $u_i(\widetilde{\mu}_i^*, \widetilde{\mu}_{N(i)}^*) \geq u_i(\mu_i, \widetilde{\mu}_{N(i)}^*)$ for any $i \in \{1, \ldots, n\}$ and any $\mu_i \in \Delta(X_i)$. We will further say that a path m, or a sequence $y(.)$, converges observationally to Nash if any $\widetilde{\mu}^* \in \widetilde{\mathcal{A}}(m)$, or $\widetilde{\mu}^* \in \widetilde{\mathcal{A}}^d(y(.))$, is an observational equilibrium.

Thus, in an observational equilibrium, each player's marginal distribution over pure strategies optimizes against the joint strategy profile. For example, any coarse correlated equilibrium in a zero-sum network is an observational equilibrium, as follows from the analysis of Cai et al. (2016). The corresponding notions of observational convergence require that any accumulation point of the joint average of the fictitious-play path or sequence is an observational equilibrium.

In the proof of the following result, we generalize an insight due to Sela (1999, Lemma 1) to arbitrary networks, and apply it to the present situation.

Proposition 6. *Let G be an arbitrary network game satisfying assumptions of any of the Propositions 1 through 4. Then any $\widetilde{CTFP}$, and likewise any $\widetilde{DTFP}$, converges observationally to Nash.*

Proof. Let $\widetilde{\mu} \in \Delta(X)$ be a probability distribution over pure strategy profiles in G, and let $\mu = (\mu_1, \ldots, \mu_n) \in \Delta(X_1) \times \cdots \times \Delta(X_n)$ denote the corresponding profile composed of marginal distributions. Then clearly, because interactions are bilateral, and because expectations ignore correlations, expected payoffs satisfy

$$u_i(\mu_i, \widetilde{\mu}_{N(i)}) = \sum_{j \in N(i)} u_{ij}(\mu_i, \mu_j), \tag{73}$$

for any $i \in \{1, \ldots, n\}$. Therefore, $\mathrm{MBR}(\widetilde{\mu}) = \mathrm{MBR}(\mu)$. The claim follows. □

Intuitively, correlation is irrelevant for the best-response correspondence in a network game G because all interactions are bilateral. As a consequence, a path m is a CTFP if and only if it is a $\widetilde{\text{CTFP}}$, and a sequence $y(.) = \{y(t)\}_{t=0}^{\infty}$ is a DTFP if and only if it is a $\widetilde{\text{DTFP}}$. Therefore, in the considered classes of network games, fictitious play with joint beliefs converges to the set of potentially correlated profiles in which each player's marginal distribution is a mixed best response both to her neighbors' jointly taken marginal distribution and to the product of her neighbors' independently taken marginal distributions.

7. Concluding remarks

In this paper, we have identified new classes of network games in which fictitious play processes converge to the set of Nash equilibria. In particular, the analysis has led to simple conditions on bilateral payoffs and the network structure that are sufficient to guarantee convergence of continuous-time fictitious play even when a player's decisions across bilateral games are interdependent. We have also constructed examples of multiplayer games that show that these conditions cannot be easily relaxed.

Applications are manifold and include security games, conflict networks, and decentralized wireless channel selection, for instance. Moreover, the findings confirm the intuition that equilibrium behavior in important types of social interaction can be reached without assuming strong forms of economic rationality.

[42] Adapting the proof of Harris (1998), existence of $\widetilde{\text{CTFP}}$ can be verified for any n-person game. For network games, however, existence follows more easily from the proof of Proposition 6 below.

Regarding the discrete-time variant of fictitious play, our results entail an extension of Robinson's (1951) classic result. This might serve as a basis for further analysis and simulation exercises. But the derivation also provides an additional illustration of the intimate relationship between the continuous-time and the discrete-time processes that has been suggested in many studies.

Obviously, we did not address all open issues related to fictitious play in network games. For instance, we did not examine stochastic fictitious play. It should be noted, however, that the convergence of stochastic fictitious play follows directly from our results for networks of exact potential games and, similarly, for acyclic networks of weighted potential games. Moreover, it may be conjectured that the techniques developed by Hofbauer and Sandholm (2002) apply also to the zero-sum networks and conflict networks discussed in the present paper.

Last but not least, there is some recent work that digs deeply into the differential topology and projective geometry of fictitious-play paths in two-person zero-sum games (van Strien, 2011; Berger, 2012). Exploring the potentially interesting implications of such approaches for zero-sum networks remains, however, beyond the scope of the present study.[43]

Appendix A. Results from the literature used in the proofs

A.1. Existence of continuous-time fictitious play

Harris (1998, Sec. 3) defines fictitious play in continuous-time in a way that is somewhat more flexible than our definition. Specifically, he assumes optimizing behavior at *almost* every point in time $\tau \geq 1$, while our Definition 1 assumes optimizing behavior at *every* point in time $\tau \geq 1$.

Definition A.1. (Harris, 1998) A continuous-time fictitious play of the first kind (CTFP1) is a measurable path $\widehat{m} : [0, \infty) \to \bigtimes_{i=1}^{n} \Delta(X_i)$ with the property that $\widehat{m}(\tau) \in \mathrm{MBR}(\alpha(\tau, \widehat{m}))$ for all $\tau \in [1, \infty)\backslash\mathcal{N}$, where $\mathcal{N} \subseteq \mathbb{R}$ is a set of measure zero.

It is obvious from the definition that any CTFP, as defined in the body of the paper, is in particular a CTFP1. As noted by Harris (1998), general results from the theory of differential inclusions imply that a continuous-time fictitious play of the first kind exists for any finite normal-form game.

Lemma A.1. *(Harris, 1998) CTFP1 exists.*

Proof. See Harris (1998, p. 244, paragraph following Prop. 7). □

Lemma A.1 is used in the proof of Lemma 1.

A.2. Convergence of continuous-time fictitious play in weighted potential games

In a later section of his paper, Harris (1998, Sec. 8) considers convergence of fictitious play in continuous-time for the class of finite weighted potential games (with n players), and establishes the following result.

Lemma A.2. *(Harris, 1998) Any CTFP1 in a finite weighted potential game converges to the set of Nash equilibria.*

Proof. See Harris (1998, Th. 24). □

As any CTFP is, a fortiori, a CTFP1, Lemma A.2 implies that any finite weighted potential game has the CTFP property as defined in the present paper. This observation is used in the proof of Proposition 3.

A.3. Differential inclusions

Harris (1998, Sec. 7) derives convergence in discrete time from uniform convergence in continuous time. For the reader's convenience, this theory is reviewed below, where we closely follow the exposition in Hofbauer and Sorin (2006, pp. 221-222).

[43] Further, one might want to seek conditions that ensure that the results of the present paper continue to hold if all bilateral games are dominance solvable (Milgrom and Roberts, 1991) or exhibit strategic complementarities and diminishing returns (Krishna, 1992; Berger, 2008). However, as discussed in Sela (1999), this last route seems less promising because interdependencies between choices in bilateral games undermine such structural properties of the bilateral games.

	x_{i+1} = H	x_{i+1} = L
x_i = H	2, .5	3.5, .75
x_i = L	2, 1	3, 1.5

	x_3 = H		x_3 = L	
	x_2 = H	x_2 = L	x_2 = H	x_2 = L
x_1 = H	2.5, 2.5, 2.5	4, 2.75, 3	3, 4, 2.75	4.5, 3.75, 3.5
x_1 = L	2.75, 3, 4	3.75, 3.5, 4.5	3.5, 4.5, 3.75	4.5, 4.5, 4.5

Fig. 7. Bilateral payoffs (left panel), and the normal form of G^2 (right panel).

Let Z be a nonempty, compact, and convex subset of some Euclidean space $\mathbb{R}^M$, where $M \geq 1$, and let Φ be an upper semi-continuous,[44] compact-valued, and convex-valued correspondence from Z to itself.[45] We consider the formal relationship

$$\dot{z} \in \Phi(z) - z, \tag{74}$$

which is commonly referred to as a *differential inclusion*. By a *solution* of (74), we mean any absolutely continuous mapping $z : [0, \infty) \to Z$ satisfying

$$\frac{\partial z(\tau)}{\partial \tau} \in \Phi(z(\tau)) - z(\tau) \tag{75}$$

at all points of differentiability of z.

Let $d(.,.)$ denote the Euclidean distance in $\mathbb{R}^M$. Then, for any point $z_1 \in \mathbb{R}^M$ and any nonempty subset $Z_0 \subseteq \mathbb{R}^M$, we may define the *Hausdorff distance* by $\underline{d}(z_1, Z_0) = \inf_{z_0 \in Z_0} d(z_1, z_0)$.

Definition A.2. A subset $Z_0 \subseteq Z$ is called a global uniform attractor of the differential inclusion (74) if, for any $\varepsilon > 0$, there exists $\tau^{\#}(\varepsilon) \geq 0$ such that, for any solution z of (74) with $z(0) \in Z$, and for any $\tau \geq \tau^{\#}(\varepsilon)$, it holds that $\underline{d}(z(\tau), Z_0) \leq \varepsilon$.

Let $\{\beta_t\}_{t=1}^{\infty}$ be a sequence of parameters in $[0, 1]$, strictly decreasing to zero as $t \to \infty$, and such that $\Sigma_{t=1}^{\infty} \beta_t = \infty$. Then, a discrete-time counterpart to differential inclusion (74) is given by

$$P_{t+1} \in \beta_t \Phi(P_t) + (1 - \beta_t) P_t \qquad (t \in \mathbb{N} = \{1, 2, \ldots\}), \tag{76}$$

where $\{P_t\}_{t=1}^{\infty}$ is a sequence in Z.

Lemma A.3. *(Hofbauer and Sorin, 2006) Assume that $Z_0 \subseteq Z$ is a global uniform attractor of the differential inclusion (74). Then, for any $\varepsilon > 0$, there exists $t^{\#}(\varepsilon) \in \mathbb{N}$ such that for any solution $\{P_t\}_{t=1}^{\infty}$ of (76), and any $t \geq t^{\#}(\varepsilon)$, we have $\underline{d}(P_t, Z_0) \leq \varepsilon$.*

Proof. See Hofbauer and Sorin (2006, Prop. 7). □

We make use of Lemma A.3 in the proof of Proposition 5.

Appendix B. Details on Examples 2 and 3

Details on Example 2. Payoffs in the bilateral conflict $G_{i,i+1}$, where $i \in \{1, 2, 3\}$, are shown in the left panel of Fig. 7. One can check that G^2 has the payoffs shown in the right panel of Fig. 7.

Lemma B.1. *In G^2, player $i \in \{1, 2, 3\}$ optimally chooses H if $r_{i-1} - 2r_{i+1} \geq 0$ (and L if $r_{i-1} - 2r_{i+1} \leq 0$).*

Proof. Using the bilateral payoff functions, one finds

$$u_i(\mathrm{H}, r_{i+1}, r_{i-1}) - u_i(\mathrm{L}, r_{i+1}, r_{i-1})$$

$$= u_{i,i+1}(\mathrm{H}, r_{i+1}) - u_{i,i+1}(\mathrm{L}, r_{i+1}) + u_{i,i-1}(\mathrm{H}, r_{i-1}) - u_{i,i-1}(\mathrm{L}, r_{i-1}) \tag{77}$$

$$= r_{i+1} \cdot 0 + (1 - r_{i+1}) \cdot \tfrac{1}{2} + r_{i-1} \cdot (-\tfrac{1}{4}) + (1 - r_{i-1}) \cdot (-\tfrac{1}{2}) \tag{78}$$

$$= \tfrac{1}{4}(r_{i-1} - 2r_{i+1}). \tag{79}$$

[44] The correspondence Φ is called *upper semi-continuous* (Aubin and Cellina, 1984, p. 41) if at any $z_0 \in Z$ and for any open set $\mathcal{O}$ containing $\Phi(z_0)$, there exists an open neighborhood $\mathcal{M}$ of z_0 such that $\Phi(\mathcal{M}) \subseteq \mathcal{O}$.

[45] As is well-known, these assumptions hold, in particular, for the correspondence MBR on $\Delta(X_1) \times \cdots \times \Delta(X_n)$.

	x_3 = H			x_3 = L	
	x_2 = H	x_2 = L		x_2 = H	x_2 = L
x_1 = H	1, 6, 5	0, 0, 0	x_1 = H	6, 5, 1	5, 1, 6
x_1 = L	5, 1, 6	6, 5, 1	x_1 = L	0, 0, 0	1, 6, 5

Fig. 8. The game G^3.

The claim follows. □

Lemma B.2. *G^2 has a unique Nash equilibrium, given by $(r_1^*, r_2^*, r_3^*) = (0, 0, 0)$.*

Proof. Clearly, (L, L, L) is a pure-strategy Nash equilibrium in G^2. Next, we show that the equilibrium is unique. Suppose first that there is a completely mixed-strategy equilibrium (r_1, r_2, r_3). Then, it follows from Lemma B.1 that $r_3 - 2r_2 = 0$, $r_1 - 2r_3 = 0$, and $r_2 - 2r_1 = 0$. But the sole solution of this system is $(r_1, r_2, r_3) = (0, 0, 0)$. Hence, G^2 does not admit a completely mixed equilibrium. Further, if another equilibrium exists, at least one player must choose a pure strategy. Without loss of generality, suppose this is player 3. Assume first that $r_3 = 1$, so that player 3 chooses H. Then the iterated elimination of strictly dominated strategies implies that player 2 chooses L and that player 1 chooses H, but then player 3 would want to deviate to $x_3 =$ L. Next, assume that player 3 chooses L. In this case, if $r_1 > 0$, then player 2 chooses H and consequently $r_1 = 0$, which is impossible. If, however, player 1 chooses L with probability one, then player 3 would want to deviate unless also player 2 chooses L with probability one. But that only brings us back to $(r_1, r_2, r_3) = (0, 0, 0)$. This proves uniqueness, and hence, the lemma. □

Next, recall the coordinates of the points at which the process changes its direction:

$$\cdots \to p_1 = (\tfrac{2}{7}, \tfrac{4}{7}, \tfrac{1}{7}) \to p_2 = (\tfrac{1}{7}, \tfrac{2}{7}, \tfrac{4}{7}) \to p_3 = (\tfrac{4}{7}, \tfrac{1}{7}, \tfrac{2}{7}) \to \ldots . \tag{80}$$

The following lemma shows that this path is indeed a stable cycle.

Lemma B.3. *(i) At point p_i, with $i \in \{1, 2, 3\}$, player i optimally chooses L, while players $i - 1$ and $i + 1$ are both indifferent. (ii) There is $\lambda > 1$ such that*

$$\lambda(p_1 - (0, 1, 0)) = p_3 - (0, 1, 0), \tag{81}$$

$$\lambda(p_2 - (0, 0, 1)) = p_1 - (0, 0, 1), \tag{82}$$

$$\lambda(p_3 - (1, 0, 0)) = p_2 - (1, 0, 0). \tag{83}$$

Proof. *(i)* From Lemma B.1, player $i = 1$ optimally chooses L if $r_3 - 2r_2 \leq 0$. But at $p_1 = (\frac{2}{7}, \frac{4}{7}, \frac{1}{7})$, we even have $r_3 - 2r_2 = -1 < 0$. Moreover, players 2 and 3 are indifferent at p_1 because $r_1 - 2r_3 = 0$ and $r_2 - 2r_1 = 0$. The points p_2 and p_3 can now be dealt with by straightforward symmetry considerations. This proves the first claim. *(ii)* Note that, with $\lambda = 2 > 1$,

$$\lambda(p_1 - (0, 1, 0)) = 2 \cdot (\tfrac{2}{7}, -\tfrac{3}{7}, \tfrac{1}{7}) = (\tfrac{4}{7}, -\tfrac{6}{7}, \tfrac{2}{7}) = p_3 - (0, 1, 0). \tag{84}$$

The other two equations follow, again, by symmetry. This proves the second claim, and hence, the lemma. □

To see how Lemma B.3 implies the existence of a stable path, note that, on the linear segment from p_1 to p_2, players 1 and 2 indeed optimally choose L, while player 3 is indifferent and can be assumed to choose H. Thus, consistent with the definition of a CTFP, the path moves from p_1 in the direction of the vector (0, 0, 1). However, the collinearity condition (82) ensures that p_2 is eventually reached.[46] Similar arguments can be used to deal with the remaining segments.

Details on Example 3. The payoff matrix of G^3 is shown in Fig. 8.

Lemma B.4. *In G^3, player 1 optimally chooses H if $5r_3 - r_2 \leq 2$ (and L if $5r_3 - r_2 \geq 2$); player 2 optimally chooses H if $5r_1 + r_3 \geq 3$ (and L if $5r_1 + r_3 \leq 3$); player 3 optimally chooses H if $5r_2 - r_1 \geq 2$ (and L if $5r_2 - r_1 \leq 2$).*

[46] Indeed, equation (82) obviously implies that p_2 may be written as a strict convex combination of p_1 and (0, 0, 1). See Fig. 4 for illustration.

Proof. As for player 1, one notes that

$$\begin{aligned} & u_1(\mathrm{H}, r_2, r_3) - u_1(\mathrm{L}, r_2, r_3) \\ & = r_2 r_3 \cdot (-4) + (1-r_2) r_3 \cdot (-6) + r_2(1-r_3) \cdot 6 + (1-r_2)(1-r_3) \cdot 4 \quad (85) \\ & = 4 + 2r_2 - 10r_3. \quad (86) \end{aligned}$$

This proves the claim regarding player 1. Similarly,

$$\begin{aligned} & u_2(r_1, \mathrm{H}, r_3) - u_2(r_1, \mathrm{L}, r_3) \\ & = r_1 r_3 \cdot 6 + (1-r_1) r_3 \cdot (-4) + r_1(1-r_3) \cdot 4 + (1-r_1)(1-r_3) \cdot (-6) \quad (87) \\ & = -6 + 10r_1 + 2r_3, \quad (88) \end{aligned}$$

proving the claim regarding player 2. Finally,

$$\begin{aligned} & u_3(r_1, r_2, \mathrm{H}) - u_3(r_1, r_2, \mathrm{L}) \\ & = r_1 r_2 \cdot 4 + (1-r_1) r_2 \cdot 6 + r_1(1-r_2) \cdot (-6) + (1-r_1)(1-r_2) \cdot (-4) \quad (89) \\ & = 4 - 2r_1 + 10r_2. \quad (90) \end{aligned}$$

This proves the final claim, and hence, the lemma. □

Lemma B.5. *G^3 has a unique Nash equilibrium, given by $(r_1^*, r_2^*, r_3^*) = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$.*

Proof. Suppose first that player 1 chooses H with probability one. Then, by iterated elimination of strictly dominated strategies, player 2 chooses H, and so does player 3. But then, player 1 would deviate to L. Suppose next that player 1 chooses L with probability one. Then, by the iterated elimination of strictly dominated strategies, player 2 chooses L, and so does player 3. But then player 1 would deviate to H. Thus, there is no equilibrium in which player 1 plays a pure strategy. Thus, $r_1 \in (0, 1)$, and by Lemma B.4, $5r_3 - r_2 = 2$. Clearly, this precludes $r_3 = 0$ and $r_3 = 1$, so that also player 3 must randomize. By Lemma B.4, $5r_2 - r_1 = 2$. But this excludes $r_2 = 0$ and $r_2 = 1$. Hence, any equilibrium is necessarily completely mixed. There is a unique completely mixed equilibrium, because the system of indifference relationships $-r_2 + 5r_3 = 2$, $5r_1 + r_3 = 3$, and $-r_1 + 5r_2 = 2$ has the unique solution $(r_1, r_2, r_3) = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$. □

Recall the cycle

$$\begin{aligned} & \ldots \to p_1 = (a, b, c) \to p_2 = (1-c, a, b) \to p_3 = (1-b, 1-c, a) \to \quad (91) \\ & \quad \to p_4 = (1-a, 1-b, 1-c) \to p_5 = (c, 1-a, 1-b) \to p_6 = (b, c, 1-a) \to \ldots, \end{aligned}$$

where $(a, b, c) = (\frac{4}{9}, \frac{8}{9}, \frac{7}{9})$ correspond to the respective probabilities to play H. The following two lemmas show that the cyclic process is indeed a CTFP.

Lemma B.6. *(i) At p_1, player 1 optimally chooses L, player 2 is indifferent, and player 3 optimally chooses H. (ii) At p_2, players 1 and 2 optimally choose L, whereas player 3 is indifferent. (iii) At p_3, player 1 is indifferent, while players 2 and 3 optimally choose L. (iv) At p_4, player 1 optimally chooses H, player 2 is indifferent, and player 3 optimally chooses L. (v) At p_5, players 1 and 2 optimally choose H, whereas player 3 is indifferent. (vi) At p_6, player 1 is indifferent, while players 2 and 3 optimally choose H.*

Proof. *(i)* By Lemma B.4, player 1 optimally chooses L if $5r_3 - r_2 \geq 2$. At $p_1 = (a, b, c) = (\frac{4}{9}, \frac{8}{9}, \frac{7}{9})$, we even have $5r_3 - r_2 = 3 > 2$. Next, player 2 is indifferent if $5r_1 + r_3 = 3$. This is true at p_1. Finally, player 3 optimally chooses H if $5r_2 - r_1 \geq 2$. At p_1, we even have $5r_2 - r_1 = 4 > 2$. *(ii)* Player 1 optimally chooses L if $5r_3 - r_2 \geq 2$. At $p_2 = (1-c, a, b) = (\frac{2}{9}, \frac{4}{9}, \frac{8}{9})$, we even have $5r_3 - r_2 = 4 > 2$. Next, player 2 optimally chooses L if $5r_1 + r_3 \leq 3$. At p_2, we even have $5r_1 + r_3 = 2 < 3$. Finally, player 3 is indifferent if $5r_2 - r_1 = 2$. At p_2, we indeed have $5r_2 - r_1 = 2$. *(iii)* Player 1 is indifferent if $5r_3 - r_2 = 2$. And indeed, at $p_3 = (1-b, 1-c, a) = (\frac{1}{9}, \frac{2}{9}, \frac{4}{9})$, we have $5r_3 - r_2 = 2$. Player 2 optimally chooses L if $5r_1 + r_3 \leq 3$. At p_3, we even have $5r_1 + r_3 = 1 < 3$. Finally, player 3 optimally chooses L if $5r_2 - r_1 \leq 2$. At p_3, we even have $5r_2 - r_1 = 1 < 2$. *(iv)* Player 1 optimally chooses H if $5r_3 - r_2 \leq 2$. At $p_4 = (1-a, 1-b, 1-c) = (\frac{5}{9}, \frac{1}{9}, \frac{2}{9})$, we even have $5r_3 - r_2 = 1 < 2$. Next, player 2 is indifferent if $5r_1 + r_3 = 3$. This is true at p_4. Finally, player 3 optimally chooses L if $5r_2 - r_1 \leq 2$. At p_4, we even have $5r_2 - r_1 = 0 < 2$. *(v)* Player 1 optimally chooses H if $5r_3 - r_2 \leq 2$. At $p_5 = (c, 1-a, 1-b) = (\frac{7}{9}, \frac{5}{9}, \frac{1}{9})$, we even have $5r_3 - r_2 = 0 < 2$. Next, player 2 optimally chooses H if $5r_1 + r_3 \geq 3$. At p_5, we even have $5r_1 + r_3 = 4 > 3$. Finally, player 3 is indifferent if $5r_2 - r_1 = 2$. And at p_5, we indeed have this. *(vi)* Player 1 is indifferent if $5r_3 - r_2 = 2$. And indeed, at $p_6 = (b, c, 1-a) = (\frac{8}{9}, \frac{7}{9}, \frac{5}{9})$, we have this. Player 2 optimally chooses H if $5r_1 + r_3 \geq 3$. At p_6, we even have $5r_1 + r_3 = 5 > 3$. Finally, player 3 optimally chooses H if $5r_2 - r_1 \geq 2$. At p_6, we even have $5r_2 - r_1 = 3 > 2$. This proves the last claim and therefore the lemma. □

The final lemma shows that process (91) is a stable cycle.

Lemma B.7. *There is* $\lambda > 1$ *such that*

$$\lambda(p_2 - (0,0,1)) = p_1 - (0,0,1) \tag{92}$$

$$\lambda p_3 = p_2 \tag{93}$$

$$\lambda(p_4 - (1,0,0)) = p_3 - (1,0,0) \tag{94}$$

$$\lambda(p_5 - (1,1,0)) = p_4 - (1,1,0) \tag{95}$$

$$\lambda(p_6 - (1,1,1)) = p_5 - (1,1,1) \tag{96}$$

$$\lambda(p_1 - (0,1,1)) = p_6 - (0,1,1). \tag{97}$$

Proof. One can easily verify that equations (92)-(97) are satisfied for $\lambda = 2$. □

Appendix C. The case of a two-person zero-sum game

This section presents a self-contained proof of convergence of CTFP for finite two-person zero-sum games. Needless to say, the sole departure from existing proofs (Hofbauer, 1995; Harris, 1998; Berger, 2006; Hofbauer and Sorin, 2006) consists in the use of the Dini derivative.

Lemma C.1. *(Brown, 1951; Hofbauer, 1995; Harris, 1998) Any finite two-person zero-sum game has the continuous-time fictitious-play property.*

Proof. We closely follow the steps of the proof of Proposition 1. Denote by X_i player i's strategy space, for $i \in \{1,2\}$, and by A player 1's payoff matrix. For $\mu = (\mu_1, \mu_2) \in \Delta(X_1) \times \Delta(X_2)$, let

$$\mathcal{L}(\mu) = \max_{\xi_1 \in \Delta(X_1)} \xi_1 \cdot A\mu_2 - \min_{\xi_2 \in \Delta(X_2)} \mu_1 \cdot A\xi_2. \tag{98}$$

Take any CTFP m, with corresponding process of independent averages α. Then,

$$\mathcal{L}(\alpha(\tau)) = m_1(\tau) \cdot A\alpha_2(\tau) - \alpha_1(\tau) \cdot Am_2(\tau), \tag{99}$$

for any $\tau \geq 1$. Moreover, for any $\widehat{\tau} \in (0, \tau)$,

$$\mathcal{L}(\alpha(\widehat{\tau})) \geq m_1(\tau) \cdot A\alpha_2(\widehat{\tau}) - \alpha_1(\widehat{\tau}) \cdot Am_2(\tau). \tag{100}$$

Hence,

$$\tau\mathcal{L}(\alpha(\tau)) - \widehat{\tau}\mathcal{L}(\alpha(\widehat{\tau})) \leq m_1(\tau) \cdot A \int_{\widehat{\tau}}^{\tau} m_2(\tau')d\tau' - \int_{\widehat{\tau}}^{\tau} m_1(\tau')d\tau' \cdot Am_2(\tau). \tag{101}$$

Dividing by $\tau - \widehat{\tau}$, and taking the limit $\widehat{\tau} \to \tau$ shows that

$$\limsup_{\widehat{\tau} \to \tau, \widehat{\tau} < \tau} \frac{\tau\mathcal{L}(\alpha(\tau)) - \widehat{\tau}\mathcal{L}(\alpha(\widehat{\tau}))}{\tau - \widehat{\tau}} \leq m_1(\tau) \cdot Am_2(\tau) - m_1(\tau) \cdot Am_2(\tau) = 0 \tag{102}$$

for a.e. $\tau \geq 1$. Given that $\tau\mathcal{L}(\alpha(\tau))$ is continuous in τ, this implies that $\tau\mathcal{L}(\alpha(\tau))$ is declining (Royden, 1988, Prop. 2, p. 99). But $\mathcal{L} \geq 0$, so that $\mathcal{L}(\alpha(\tau)) \to 0$ as $\tau \to \infty$. Thus, any accumulation point of α is indeed a Nash equilibrium. □

References

Acemoglu, D., Ozdaglar, A., 2011. Opinion dynamics and learning in social networks. Dyn. Games Appl. 1, 3–49.
Aubin, J.-P., Cellina, A., 1984. Differential Inclusions. Springer.
Baye, M.R., Kovenock, D., de Vries, C.G., 1996. The all-pay auction with complete information. Econ. Theory 8, 291–305.
Ben-Sasson, E., Kalai, A.T., Kalai, E., 2007. An approach to bounded rationality. In: Advances in Neural Information Processing Systems, pp. 145–152.
Berger, U., 2006. A generalized model of best response adaption. Int. Game Theory Rev. 8, 45–66.
Berger, U., 2008. Learning in games with strategic complementarities revisited. J. Econ. Theory 143, 292–301.
Berger, U., 2012. Non-algebraic convergence proofs for continuous-time fictitious play. Dyn. Games Appl. 2, 4–17.
Blume, L.E., 1993. The statistical mechanics of strategic interaction. Games Econ. Behav. 5, 387–424.
Bollobás, B., 2013. Modern Graph Theory. Springer.
Bourlès, R., Bramoullé, Y., Perez-Richet, E., 2017. Altruism in networks. Econometrica 85, 675–689.

Bozbay, I., Vesperoni, A., 2018. A contest success function for networks. J. Econ. Behav. Organ. 150, 404–422.
Bramoullé, Y., Kranton, R., 2016. Games played on networks. In: Bramoullé, Y., Galeotti, A., Rogers, B. (Eds.), Oxford Handbook on the Economics of Networks. Oxford University Press.
Bramoullé, Y., Kranton, R., D'Amours, M., 2014. Strategic interaction and networks. Am. Econ. Rev. 104, 898–930.
Brandt, F., Fischer, F., Harrenstein, P., 2013. On the rate of convergence of fictitious play. Theory Comput. Syst. 53, 41–52.
Bregman, L.M., Fokin, I.N., 1987. Methods of determining equilibrium situations in zero-sum polymatrix games. Optimizatsia 40, 70–82 (in Russian).
Bregman, L.M., Fokin, I.N., 1998. On separable non-cooperative zero-sum games. Optimization 44, 69–84.
Brown, G.W., 1949. Some notes on computation of games solutions. Report P-78. Rand Corporation.
Brown, G.W., 1951. Iterative solution of games by fictitious play. In: Koopmans, T.C. (Ed.), Activity Analysis of Production and Allocation. Wiley, New York, pp. 374–376.
Cai, Y., Candogan, O., Daskalakis, C., Papadimitriou, C.H., 2016. Zero-sum polymatrix games: a generalization of minmax. Math. Oper. Res. 41, 648–655.
Cai, Y., Daskalakis, C., 2011. On minmax theorems for multiplayer games. In: Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 217–234.
Cesa-Bianchi, N., Lugosi, G., 2006. Prediction, Learning, and Games. Cambridge University Press, New York.
Conitzer, V., 2009. Approximation guarantees for fictitious play. In: Communication, Control, and Computing, 47th Annual Allerton Conference on IEEE.
Daskalakis, C., Pan, Q., 2014. A counter-example to Karlin's strong conjecture for fictitious play. In: Foundations of Computer Science (FOCS), IEEE 55th Annual Symposium.
Daskalakis, C., Papadimitriou, C.H., 2009. On a network generalization of the minmax theorem. In: International Colloquium on Automata, Languages, and Programming. Springer, Berlin, pp. 423–434.
Driesen, B., 2009. Continuous fictitious play in zero-sum games. Research Memorandum RM/09/049. University of Maastricht.
Dziubiński, M., Goyal, S., Minarsch, D., 2016a. Dynamic conflict on a network. In: Proceedings of the 2016 ACM Conference on Economics and Computation, pp. 655–656.
Dziubiński, M., Goyal, S., Vigier, A., 2016b. Conflict and networks. In: The Oxford Handbook of the Economics of Networks. Oxford University Press.
Franke, J., Öztürk, T., 2015. Conflict networks. J. Public Econ. 126, 104–113.
Freund, Y., Schapire, R.E., 1999. Adaptive game playing using multiplicative weights. Games Econ. Behav. 29, 79–103.
Fudenberg, D., Kreps, D.M., 1993. Learning mixed equilibria. Games Econ. Behav. 5, 320–367.
Fudenberg, D., Levine, D.K., 1998. The Theory of Learning in Games, vol. 2. MIT Press, Cambridge.
Gaunersdorfer, A., Hofbauer, J., 1995. Fictitious play, Shapley polygons, and the replicator equation. Games Econ. Behav. 11, 279–303.
Gilboa, I., Matsui, A., 1991. Social stability and equilibrium. Econometrica 59, 859–867.
Gjerstad, S., 1996. The rate of convergence of continuous fictitious play. Econ. Theory 7, 161–177.
Harris, C., 1998. On the rate of convergence of continuous-time fictitious play. Games Econ. Behav. 22, 238–259.
Hart, S., Mas-Colell, A., 2001. A general class of adaptive strategies. J. Econ. Theory 98, 26–54.
Hirshleifer, J., 1989. Conflict and rent-seeking success functions: ratio vs. difference models of relative success. Public Choice 63, 101–112.
Hofbauer, J., 1995. Stability for the best response dynamics. Preprint.
Hofbauer, J., Sandholm, W.H., 2002. On the global convergence of stochastic fictitious play. Econometrica 70, 2265–2294.
Hofbauer, J., Sandholm, W.H., 2009. Stable games and their dynamics. J. Econ. Theory 144, 1665–1693.
Hofbauer, J., Sorin, S., 2006. Best response dynamics for continuous zero-sum games. Discrete Contin. Dyn. Syst., Ser. B 6, 215–224.
Huremović, K., 2014. Rent seeking and power hierarchies: a noncooperative model of network formation with antagonistic links. Working Papers 2014.45, Fondazione Eni Enrico Mattei.
Jackson, M.O., 2005. A survey on network formation models: stability and efficiency. In: Group Formation in Economics: Networks, Clubs, and Coalitions, pp. 11–49.
Jackson, M.O., Nei, S., 2015. Networks of military alliances, wars, and international trade. Proc. Natl. Acad. Sci. 112, 15277–15284.
Jordan, J., 1993. Three problems in learning mixed-strategy equilibria. Games Econ. Behav. 5, 368–386.
König, M., Rohner, D., Thoenig, M., Zilibotti, F., 2017. Networks in conflict: theory and evidence from the great war of Africa. Econometrica 85, 1093–1132.
Kovenock, D., Roberson, B., 2018. The optimal defense of networks of targets. Econ. Inq. 56, 2195–2211.
Kovenock, D., Sarangi, S., Wiser, M., 2015. All-pay 2×2 hex: a multibattle contest with complementarities. Int. J. Game Theory 44, 571–597.
Krishna, V., 1992. Learning in games with strategic complementarities. HBS Working Paper 92-073. Harvard University.
Krishna, V., Sjöström, T., 1997. Learning in games: fictitious play dynamics. In: Hart, S., Mas-Colell, A. (Eds.), Cooperation: Game-Theoretic Approaches, vol. 155. Springer, pp. 257–273.
Lazear, E.P., Rosen, S., 1981. Rank-order tournaments as optimum labor contracts. J. Polit. Econ. 89, 841–864.
Matros, A., Rietzke, D., 2018. Contests on Networks. Mimeo.
Matsui, A., 1992. Best response dynamics and socially stable strategies. J. Econ. Theory 57, 343–362.
Metrick, A., Polak, B., 1994. Fictitious play in 2 × 2 games: a geometric proof of convergence. Econ. Theory 4, 923–933.
Milgrom, P., Roberts, J., 1991. Adaptive and sophisticated learning in normal form games. Games Econ. Behav. 3, 82–100.
Miyasawa, K., 1961. On the convergence of the learning process in a 2 × 2 non-zero-sum two-person game. Research Memorandum No. 33. Princeton University.
Monderer, D., Samet, D., Sela, A., 1997. Belief affirming in learning processes. J. Econ. Theory 73, 438–452.
Monderer, D., Sela, A., 1996. A 2×2 game without the fictitious play property. Games Econ. Behav. 14, 144–148.
Monderer, D., Sela, A., 1997. Fictitious play and no-cycling conditions. Working Paper. University of Mannheim.
Monderer, D., Shapley, L.S., 1996a. Potential games. Games Econ. Behav. 14, 124–143.
Monderer, D., Shapley, L.S., 1996b. Fictitious play property for games with identical interests. J. Econ. Theory 68, 258–265.
Morris, S., 2000. Contagion. Rev. Econ. Stud. 67, 57–78.
Morris, S., Ui, T., 2004. Best response equivalence. Games Econ. Behav. 49, 260–287.
Moulin, H., Vial, J.-P., 1978. Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. Int. J. Game Theory 7, 201–221.
Ostrovski, G., van Strien, S., 2014. Payoff performance of fictitious play. J. Dyn. Games 1, 621–638.
Roberson, B., 2006. The Colonel Blotto game. Econ. Theory 29, 1–24.
Robinson, J., 1951. An iterative method of solving a game. Ann. Math. 54, 296–301.
Rosenmüller, J., 1971. Über Periodizitätseigenschaften spieltheoretischer Lernprozesse. Probab. Theory Relat. Fields 17, 259–308 (in German).
Royden, H.L., 1988. Real Analysis, 3rd edition. Macmillan, New York.
Sela, A., 1999. Fictitious play in 'one-against-all' multi-player games. Econ. Theory 14, 635–651.
Shamma, J.S., Arslan, G., 2004. Unified convergence proofs of continuous-time fictitious play. IEEE Trans. Autom. Control 49, 1137–1142.
Shapiro, H.N., 1958. Note on a computation method in the theory of games. Commun. Pure Appl. Math. 11, 587–593.
Shapley, L.S., 1964. Some topics in two-person games. Adv. Game Theory 52, 1–29.

Swenson, B., Kar, S., 2017. On the exponential rate of convergence of fictitious play in potential games. In: 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton). IEEE, pp. 275–279.
Tullock, G., 1980. Efficient rent-seeking. In: Buchanan, J.M., Tollison, R., Tullock, G. (Eds.), Toward a Theory of the Rent-Seeking Society. Texas A&M University Press, pp. 267–292.
van Strien, S., 2011. Hamiltonian flows with random-walk behaviour originating from zero-sum games and fictitious play. Nonlinearity 24, 1715–1742.
Viossat, Y., Zapechelnyuk, A., 2013. No-regret dynamics and fictitious play. J. Econ. Theory 148, 825–842.
Xu, J., Zenou, Y., Zhou, J., 2019. Networks in Conflict: A Variational Inequality Approach.
Young, P.H., 1993. The evolution of conventions. Econometrica 61, 57–84.