



Learning and selfconfirming equilibria in network games [☆]

Pierpaolo Battigalli ^{a,d,e}, Fabrizio Panebianco ^{b,f}, Paolo Pin ^{c,e,*}

^a Department of Decision Sciences, Università Bocconi, Milan, Italy

^b Department of Economics and Finance, Università Cattolica, Milan, Italy

^c Department of Economics and Statistics, Università di Siena, Italy

^d IGIER, Università Bocconi, Milan, Italy

^e BIDS, Università Bocconi, Milan, Italy

^f CLE, Università Cattolica, Milan, Italy

Received 31 December 2020; final version received 27 June 2023; accepted 7 July 2023

Available online 22 July 2023

Abstract

Consider a set of agents who play a network game repeatedly. Agents may not know the network. They may even be unaware that they are interacting with other agents in a network. Possibly, they just understand that their optimal action depends on an unknown state that is, actually, an aggregate of the actions of their neighbors. In each period, every agent chooses an action that maximizes her instantaneous subjective expected payoff and then updates her beliefs according to what she observes. In particular, we assume that each agent only observes her realized payoff. A steady state of the resulting dynamic is a **selfconfirming equilibrium** given the assumed feedback. We identify conditions on the network externalities, agents' beliefs, and learning dynamics that make agents more or less active (or even inactive) in steady state compared to Nash equilibrium.

[☆] We thank Editor Tilman Börgers, an anonymous Associate Editor, and two anonymous referees. We also thank Nicola Bariletto, Federico Bobbio, Davide Bordoli, Yann Bramoullé, Ben Golub, Sebastiano Della Lena, Nicodemo De Vito, Nicolò Generoso, Julien Manili, Anna Merotto, Paola Moscariello, Alessandro Pavan, Giulio Principi, Sergey Savchenko, Nicolas Sourisseau, Yves Zenou, and seminar participants at 24th CTN Workshop in Aix-en-Provence, Bergamo, Bocconi, CISEI in Capri, Milano Bicocca, Milano Cattolica, 7th European Meeting on Networks in Cambridge, Marseille, Nazarbayev University, NSE at Indiana University, Pompeu Fabra, Siena, UTS Sydney, Venice. Pierpaolo Battigalli, Fabrizio Panebianco, and Paolo Pin gratefully acknowledge funding from, respectively, the European Research Council (ERC) grant 324219, the Spanish Ministry of Economía y Competitividad project ECO2017-87245-R, and the Italian Ministry of Education Progetti di Rilevante Interesse Nazionale (PRIN) grant 2017ELHNNJ.

* Corresponding author.

E-mail address: paolo.pin@unisi.it (P. Pin).

<https://doi.org/10.1016/j.jet.2023.105700>

0022-0531/© 2023 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

© 2023 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

JEL classification: C72; D83; D85

Keywords: Learning; Selfconfirming equilibrium; Network games; Observability by active players; Shallow conjectures

1. Introduction

Social networks can be quite complex. Think about friendship networks or networks of people interacting online (such as Twitter, Facebook, Instagram, and so on). These networks often consist of thousands (or millions) of interacting agents or firms, and agents rarely know how the network is shaped.¹ In this paper, we provide a novel approach to analyze how incomplete information about the network affects behavior and learning processes. We propose a framework in which agents may not know how the network affects their payoffs, how the network is shaped, or even that they are interacting in a network.

The standard solution concept used to study the behavior of agents in network games is Nash equilibrium, with the motivation that learning and adaptation converge to a profile of actions in which every player best responds to the actions of the other players. Nash equilibrium action profiles are limit outcomes of learning paths when agents have perfect feedback about the payoff-relevant aspects of others' behavior. Yet, as we shall argue, such perfect feedback hypothesis may be too strong for some social networks applications and, if learning is based on imperfect feedback, non-Nash action profiles may result as the steady-state limits of learning paths. Indeed, such limits under (possibly) imperfect feedback are characterized by the selfconfirming equilibrium concept. With this, we analyze the effects of milder conditions on information feedback.

In our analysis we assume that the only feedback agents receive is their realized payoff. This implies that they do not always identify the payoff-relevant aspects of the actions of others, represented by a payoff state.

We analyze how agents use the feedback they receive to update their conjectures about the payoff state and best respond to them, and we characterize their limit behavior under different settings of local and global externalities. We study conditions under which agents are more or less active (or even inactive) in steady state compared to Nash equilibrium. These conditions are based on the network structure and on the type of externalities, on the conjectures that agents have, and on the rules that they use to update their conjectures. Thus, in some applications, knowing these conditions, a social planner or the owner of the network can try to change the beliefs of people to induce them to increase their activity levels.

1.1. Preview of the model and results

To be more specific about our modelling approach, let us introduce an example that will guide us through the whole discussion. Consider an online social network with many users, like Twitter,

¹ For example, Breza et al. (2018) provide evidence from Indian rural villages on the fact that people have limited knowledge about the social networks of personal relations in which they are embedded, at odds with many of the existing theoretical models of strategic interactions in networks. Actually, even if the decision makers in our model play a game, we often call them 'agents' instead of 'players' when we want to emphasize that they need not reason strategically to choose their actions.

and a simultaneous-moves game in which each user i decides her level of activity $a_i \geq 0$ in the social network. The payoff that agents get from their activity depends on the social interaction. We start considering the case in which only local externalities are at play, and then extend the model to the case in which there are also global externalities. In particular, active user i receives idiosyncratic externalities – that can be positive or negative – from the other users with whom she is in contact in the social network. The externality from user j to user i is proportional to the time that they both spend on the social network, a_i and a_j . Sticking to a quadratic specification, which yields linear best replies, we assume that the payoff function of i is²

$$u_i(a_i, \mathbf{a}_{-i}) = \alpha_i a_i - \frac{1}{2} a_i^2 + a_i \sum_{j \in I \setminus \{i\}} z_{ij} a_j. \tag{1}$$

In equation (1), I is the set of agents, or individuals, in the social network, a_i is the activity level of $i \in I$, \mathbf{a}_{-i} is the profile of activities of all the other users in I , and $\alpha_i > 0$ represents the individual pleasure of i from being active on the social network in isolation, which results in the *bliss point* of activity in autarchy. For each $j \in I \setminus \{i\}$, parameter z_{ij} represents the intensity (absolute value) and type (sign) of the externality from j to i . We say that j affects i , or that j is a **peer** (or a **neighbor**) of i , if $z_{ij} \neq 0$.

The network described by the matrix \mathbf{Z} of all the z_{ij} 's is assumed to be *exogenous*. As a first approximation, this fits a *directed* online social network like Twitter or Instagram, where users do not have full control on who follows them.³

Under this interpretation, i receives positive or negative externalities from those who follow her proportional to her activity. We do not assume that player i knows all the z_{ij} 's. She may not know them either because she cannot observe who is following her,⁴ or because she knows her followers but she does not know the sign or intensity of their externality. The payoff of i represents both the pleasure that i gets from participating in the platform and what i can indirectly observe about her own popularity. We consider that i cannot choose the style of what she writes, since she just follows her exogenous nature. In this interpretation, a_i represents both the amount of time that i spends on the platform and the amount of posts that i writes, and this can make her more or less appreciated, according to how her style combines with the (typically unobserved) tastes of each of her followers. In our setting, player i may also set $a_i = 0$. Indeed, we interpret \mathbf{Z} as a network of *opportunities* of interaction, with players deciding endogenously whether they want to be active or inactive. When they are inactive, not only the network becomes irrelevant for them, but they also become irrelevant for the payoffs of other players.

² This is the class of linear–quadratic network games originally analyzed by Ballester et al. (2006), as we discuss in the next section. We use **boldface** symbols to denote vectors (in this case, action profiles) and matrices.

³ An endogenous directed network in which player i decides who to follow (the z_{ji} entries of matrix \mathbf{Z}) but not who is following her (the z_{ij} entries of matrix \mathbf{Z}) seems to us in line with our assumption of exogenous network. That is because, in this modification of our model, a player affects the payoff of those that she follows but her payoff is not affected by their choices, including, if the network were endogenous, who they follow. So, endogenizing \mathbf{Z} would mean to endogenize choices (link choices) that are payoff-irrelevant for the players.

⁴ In many platforms, for users with many followers it is not practical to keep track of the list of followers and, even if possible, the effective *interactions* are driven by many opaque algorithmic decisions, like whom the algorithm decides to show the messages. There are online social networks, like **Reddit**, which actually do not provide this information at all to their users. **Reddit**, in particular, provides a measure to each user, called *karma*, which is apparently based on how many other people follow – and how much they like – what that user posts. However, the algorithm on which this measure is based is not public.

To analyze learning dynamics and their steady states, we have to specify what agents observe after their choices, at the end of any given period. Continuing the running example of activity on Twitter, user i , typically, does not observe the sign of the externalities and the activity of others. However, she gets indirect measures of her level of appreciation that come, for instance, from her conversations and experiences in the real world, where her activity on Twitter affects her social and professional real life. If the players are small firms using Twitter for advertising, they will observe their actual profits. Players of this game may have wrong beliefs about the details of the game they are playing (e.g., the structure of the network, or the value of the parameters) and about the actions of other players. Consequently, they update their beliefs in response to the feedback they receive, which is assumed to be their realized payoff, and maximize their instantaneous expected payoff given such updated beliefs. This updating process yields learning paths that do not necessarily converge to a Nash equilibrium of the game.

Next we also consider an extra **global** term in the payoff function:

$$u_i(a_i, \mathbf{a}_{-i}) = \alpha_i a_i - \frac{1}{2} a_i^2 + a_i \sum_{j \in I \setminus \{i\}} z_{ij} a_j + \gamma \sum_{k \in I \setminus \{i\}} a_k. \tag{2}$$

We can interpret this extra term $\gamma \sum_{k \in I \setminus \{i\}} a_k$ as an additional utility that i gets, regardless of being active or inactive. So, while local network effects impact an individual’s marginal utility from activity in the network, global network effects do not impact this marginal utility. Moreover, in this case, what agents can learn radically changes with respect to the previous case without global externalities, because the presence of the global term makes it harder to identify the impact of neighbors on one’s own realized payoff.

Although we let agents be largely unaware of the nature and extent of network externalities, we rely on the following minimal *maintained assumption*: each agent knows how her payoff (utility) and information feedback depend on her action and on a **payoff state**, which in turn depends on other agents’ actions in the given network (but the agent may ignore the latter dependence). With this, each agent best responds to her conjecture about the payoff state, observes her realized payoff, and – in equilibrium – her conjecture must be consistent with the feedback received, that is, **confirmed**. Note that conjectures may be confirmed without being correct. A profile of actions and conjectures satisfying these requirements forms a **selfconfirming equilibrium** (SCE), whereby agents best respond to conjectures that can be wrong, but are nonetheless believed to be true, as they are consistent with the available evidence.

In our analysis, we assume that *agents observe only their realized payoff*. Given the assumed properties of the payoff functions, it follows that there exists a discontinuity at activity level 0 in what agents learn from their feedback. In particular, we show that if *externalities are only local* (i.e., positive or negative peer effects) as in equation (1), an active player i is always able to exactly infer from her feedback the realized payoff state $x_i = \sum_{j \in I \setminus \{i\}} z_{ij} a_j$ (e.g., how good it is for her to have a Twitter account), even if she may have a wrong conjecture about how many neighbors she has or what her neighbors chose. Indeed, we say that in a selfconfirming equilibrium *active* agents have correct **shallow** conjectures about the payoff state, but possibly wrong **deep** conjectures about the parameters and the actions of others. Actually, agents may even be unaware that the payoff state is determined by others within an interactive network structure; in this case, they do not hold deep conjectures.

Conversely, an *inactive* agent receives uninformative feedback, because inactivity makes payoff independent of the state. If – given her conjecture – she finds it subjectively optimal to be

inactive, such lack of information about the payoff state creates an “inactivity trap,” allowing her possibly wrong conjecture to persist. Furthermore, if there are only non-negative local externalities (hence, strategic complementarities), the inactivity of some agents also induces the remaining agents to be (weakly) less active than what they would be in Nash equilibrium. Actually, for the application to online social networks, such inactivity trap seems to be perceived by the platforms, to the point that many of them, after some period of inactivity of agents, start sending emails about what is happening on the online social network to provide a positive signal and make agents more prone to be active again.⁵

Since network games without global externalities are easier to analyze and relevant in their own right, we first study this special case and then extend the analysis to games with both local and global externalities. When agents observe only their realized payoff, the presence of *global externalities* impacts the way in which conjectures are confirmed or revised. Recall that in our setting *a game is not solely characterized by the best-reply functions, but also by the structure of the payoff/feedback functions*. This implies that additional SCE action profiles are possible compared to the case with only local externalities. Indeed, we show that the SCE action profiles studied for the latter special case correspond to the equilibria of games with local *and* global externalities, in which agents have correct conjectures about the global aggregate. But there are other SCEs in which conjectures about global aggregates are wrong. For the sake of simplicity, we focus on the case of *positive* local and global externalities, in which being inactive is dominated. Even in this simple case, agents may have a continuum of confirmed conjectures about the relative size of the two externalities. Indeed, there are multiple SCEs because, even if they are active, players may have false but confirmed conjectures making them choose actions that are not objective best replies. In detail, we find that active agents are not able to perfectly infer the size of the local externality due to the confound induced by the global externality: the realized payoff, a one-dimensional feedback, does not allow to retrieve a two-dimensional (local-global) externality. In particular, since we assume positive externalities, we show that agents’ perception of their role in the network determines whether in an SCE they are more or less active than predicted by Nash equilibrium. Thus, overall activity and (possibly) welfare are higher if agents think that (externalities are positive and that) they are more linked than in reality.

The paper is structured as follows. In Section 2 we discuss the related literature. Section 3 presents the basic framework and equilibrium concept. In Section 4 we analyze network games with only local externalities, whereas in Section 5 we analyze a more general model that accounts for global externalities. Section 6 concludes.

We devote appendices to proofs and technical results. Appendix A analyzes properties of feedback and selfconfirming equilibria in a class of games including as a special case the linear-quadratic network games that we consider in the main text. Appendix B reports existing and novel results in linear algebra, that we use to find sufficient conditions for unique and interior Nash equilibria in network games. Appendix C contains the proofs of the results presented in the main text.

⁵ For example, in November 2019, Twitter sent emails to all its inactive users, under the justification of permanently remove inactive accounts. Source: [The Verge](#).

2. Related literature

We model interactions through *linear–quadratic network games*. We focus on this class of games because it has well-known properties, and it has been used for modelling a variety of environments where strategic interaction is local and can be described by a network structure, as surveyed by Zenou (2016) and Bramoullé and Kranton (2016). Moreover, these games belong to the larger class of *nice games* (Moulin, 1984), for which we provide in Appendix A some general results. Bramoullé et al. (2014) show that other payoff functions lead to the same best-reply functions, hence, to the same Nash equilibria of linear–quadratic network games. However, we focus on *selfconfirming equilibria* (SCE), and, since realized payoffs affect feedback, the entire payoff function is relevant, not just the corresponding best-reply function. Thus, we rely in our analysis on the specific original payoff function of network games, as introduced in the economic literature by Ballester et al. (2006).

We call “selfconfirming equilibria” the steady states of learning processes when static or dynamic games are played recurrently, independently of the specific assumptions about feedback (monitoring) at the end of each one-period play. This concept encompasses what used to be called “conjectural equilibrium” as well as the original “selfconfirming equilibrium” of Fudenberg and Levine (1993). In an SCE, agents best respond to confirmed conjectures that may be inconsistent with sophisticated strategic reasoning. The latter has been added to SCE relating it to rationalizability. See Section IV of Battigalli et al. (2015) and the relevant references therein for a more detailed discussion of different versions of these concepts. Here we focus on SCE, while we analyze SCE with rationalizable conjectures in the appendix of our working paper version. Lipnowski and Sadler (2019) apply a concept akin to rationalizable SCE for games where feedback about the behavior of others is described by a network topology: agents have correct conjectures about the strategies of their peers (neighbors), but their payoff may depend on the whole strategy profile and it is not observed *ex post*. We instead assume that agents observe (only) their realized payoff and that the network describes how the payoff of each agent is affected by the actions of her neighbors (with global externalities, there is also an influence of other players on own payoffs not mediated by the network structure). We interpret the recent model of Bochet et al. (2020) as another interesting application of the SCE concept to a network game where agents observe, besides their realized payoff, the behavior of their neighbors. In their game, agents play a Tullock contest with incomplete information about the structure of externalities. The equilibrium concept that they use is, actually, a refinement of SCE whereby agents wrongly believe that they compete for a local rather than a global resource.

McBride (2006) applies SCE to games of network formation with asymmetric information. In his model, agents observe (only) the private information of other agents they link to, and possibly of agents to whom they are indirectly linked. We instead assume that the network is exogenous and actions are activity levels. We allow for information incompleteness, but – with the partial exception of Section 5 – we do not assume that agents are necessarily aware of the states of nature (e.g., the possible network structures), hence we do not assume that agents necessarily reason about them.⁶ Frick et al. (2022) apply a refinement of rationalizable SCE to analyze a model with asymmetric information and assortative matching. The refinement is obtained by assuming that

⁶ De Marti and Zenou (2015) consider network formation games where players do not know the externalities in the network, which are random, but their analysis concerns Bayesian-Nash equilibria, and players have correct *ex-ante* beliefs.

agents neglect the assortativity of matching when they make inferences from feedback. Foerster et al. (2021) share elements of Lipnowski and Sadler (2019) and of McBride (2006). As in the former, agents observe the behavior of those with whom they are linked; furthermore, they also observe public links. As in the latter, theirs is a model of network formation. They assume that beliefs satisfy a kind of rationalizable SCE condition. Unlike those papers, however, Foerster et al. (2021) do not explicitly analyze the equilibria of a non-cooperative game, but rather adopt a reduced-form notion of stability akin to Jackson and Wolinsky (1996).

3. Framework

3.1. Network games

Consider a finite set of agents (or players) I , with cardinality $n = |I|$ and generic element i . Agents are located in a network $\mathbf{Z} \in \mathbb{R}^{I \times I}$, here expressed as an adjacency matrix, with $z_{ii} = 0$ for each i in I . Each agent $i \in I$ chooses an action a_i from a compact interval $A_i = [0, \bar{a}_i]$.⁷ For each $i \in I$, $\mathbf{A}_{-i} := \times_{j \neq i} A_j$ denotes the set of feasible action profiles $\mathbf{a}_{-i} = (a_j)_{j \in I \setminus \{i\}}$ for players different from i . For each $i \in I$, we posit two compact intervals $X_i := [\underline{x}_i, \bar{x}_i] \subset \mathbb{R}$ and $Y_i := [0, \bar{y}_i] \subset \mathbb{R}_+$ of **payoff states for i** , with the interpretation that i 's payoff is determined by her action a_i , the interaction between a_i and state x_i , and the additive term y_i according to the quadratic utility function

$$v_i: \begin{aligned} A_i \times X_i \times Y_i &\rightarrow \mathbb{R}, \\ (a_i, x_i, y_i) &\mapsto \alpha_i a_i - \frac{1}{2} a_i^2 + a_i x_i + y_i. \end{aligned} \tag{3}$$

Payoff state x_i is determined by the actions of i 's neighbors – the agents with non-zero weight in adjacency matrix \mathbf{Z} – according to the linear **aggregator**⁸

$$\ell_i: \begin{aligned} \mathbf{A}_{-i} &\rightarrow X_i, \\ \mathbf{a}_{-i} &\mapsto \sum_{j \neq i} z_{ij} a_j. \end{aligned} \tag{4}$$

Since the codomain of ℓ_i is $[\underline{x}_i, \bar{x}_i]$, we are effectively assuming that

$$\underline{x}_i \leq \sum_{j \in N_i^-} z_{ij} \bar{a}_j, \bar{x}_i \geq \sum_{j \in N_i^+} z_{ij} \bar{a}_j,$$

where $N_i^- := \{j \in I : z_{ij} < 0\}$ denotes the set of neighbors of player i that have a negative effect on the payoff state of i , and $N_i^+ := \{j \in I : z_{ij} > 0\}$ denotes the set of neighbors of player i that have a positive effect on the payoff state of i .

We also consider a non-strategic global externality, that is, a payoff state y_i determined by all the co-players' actions according to the proportional aggregator:

$$g_i: \begin{aligned} \mathbf{A}_{-i} &\rightarrow Y_i \\ \mathbf{a}_{-i} &\mapsto \gamma \sum_{j \neq i} a_j, \end{aligned} \tag{5}$$

⁷ Note that in the network literature it is common to assume $A_i = \mathbb{R}_+$. For the case of local externalities with complementarities, we consider constraints on the parameters so that assuming an upper bound on actions is without loss of generality for the analysis of Nash equilibria and of selfconfirming equilibria without global externalities. When externalities are global the upper bound may become binding, and we discuss this issue below in the paper.

⁸ In principle, we can allow for non-linear aggregators as in Feri and Pin (2020). However, in this paper, we focus on the linear case. In Appendix A we provide results for the non-linear case.

where $\gamma \geq 0$. Since the codomain of g_i is $[0, \bar{y}_i]$, we are assuming that $\gamma \sum_{j \neq i} \bar{a}_j \leq \bar{y}_i$. The special case of **no global externalities** obtains if $\gamma = 0$ and every player i knows it, or at least knows that $y_i = 0$.

With this, we derive the **payoff function**

$$u_i: A_i \times \mathbf{A}_{-i} \rightarrow \mathbb{R},$$

$$(a_i, \mathbf{a}_{-i}) \mapsto v_i(a_i, \ell_i(\mathbf{a}_{-i}), g_i(\mathbf{a}_{-i})).$$

Since y_i does not interact with a_i , $x_i = \ell_i(\mathbf{a}_{-i})$ is the payoff-relevant state that i has to guess in order to choose a subjectively optimal action. We let

$$r_i(x_i) := \begin{cases} 0, & \text{if } x_i \leq -\alpha_i, \\ \alpha_i + x_i, & \text{if } -\alpha_i < x_i < \bar{a}_i - \alpha_i, \\ \bar{a}_i, & \text{if } x_i \geq \bar{a}_i - \alpha_i. \end{cases} \tag{6}$$

denote the continuous and piecewise linear **best-reply function** of player $i \in I$. Note that, since $\alpha_i > 0$, we may have $r_i(x_i) = 0$ only if $\underline{x}_i < 0$.

We assume that *the game is repeatedly played by agents maximizing their instantaneous payoff*. Each agent i knows her utility function $v_i : A_i \times X_i \times Y_i \rightarrow \mathbb{R}$ as specified in eq. (3), hence also its domain $A_i \times X_i \times Y_i = [0, \bar{a}_i] \times [\underline{x}_i, \bar{x}_i] \times [0, \bar{y}_i]$ and the “stand-alone” parameter α_i , but we do not assume that the aggregators parameters (\mathbf{Z}, γ) are known.⁹ Actually, for most of our analysis it does not even matter that agents understand that payoff states aggregate the actions of others according to eq.s (4) and (5). After each play, agents get an imperfect feedback about the payoff states. Specifically, we assume that *each agent observes only her realized utility/payoff*. What agent i learns in a given period after choosing action a_i and observing her realized payoff \hat{v}_i is that $(x_i, y_i) \in \{(x'_i, y'_i) : v_i(a_i, x'_i, y'_i) = \hat{v}_i\}$, that is,

$$(x_i, y_i) \in \begin{cases} \{(x'_i, y'_i) : y'_i = \hat{v}_i\}, & \text{if } a_i = 0, \\ \{(x'_i, y'_i) : \alpha_i a_i - \frac{1}{2}a_i^2 + a_i x'_i + y'_i = \hat{v}_i\}, & \text{if } a_i > 0. \end{cases}$$

In words, if i is inactive she can infer y_i but has no clue about x_i , if she is active she obtains joint information about y_i and x_i that she cannot disentangle.

If there are *no global externalities*, that is, if each i knows that $y_i = 0$, then being inactive reveals nothing, because $v_i(0, x_i, 0) = 0$ independently of x_i , while being active reveals that

$$x_i = \frac{\hat{v}_i - \alpha_i a_i + \frac{1}{2}a_i^2}{a_i} = \frac{\hat{v}_i}{a_i} - \alpha_i + \frac{1}{2}a_i.$$

With the aforementioned assumptions about feedback, the interactive situation is represented by the mathematical structure

$$NG = \langle I, (A_i, X_i, Y_i, v_i, \ell_i, g_i)_{i \in I} \rangle,$$

determined by eq.s (3), (4), and (5), which we call **linear-quadratic network game with just observable payoffs**, or simply **network game**.

To choose an action, a subjectively rational agent i must have some deterministic or probabilistic conjecture about the payoff state x_i . Yet, her post-feedback update about x_i depends on what she thinks about y_i , because she gets imperfect joint feedback about both. Therefore,

⁹ Except for the case of no global externalities, see above.

we model how i forms conjectures about x_i and y_i . We refer to conjectures about the states x_i and y_i as **shallow conjectures**, as opposed to **deep conjectures**, which concern the specific network topology \mathbf{Z} , the global externality parameter γ (when positive), and the actions of other players \mathbf{a}_{-i} . In our equilibrium analysis it is sufficient to focus on *deterministic shallow conjectures*.¹⁰ Indeed, for each $i \in I$ and every probabilistic conjecture $\mu_i \in \Delta(X_i \times Y_i)$, there exists a corresponding deterministic conjecture $(\hat{x}_i, \hat{y}_i) \in X_i \times Y_i$ that justifies the same action a_i^* as the unique best reply.¹¹ Deep conjectures are relevant for the analysis of strategic thinking, e.g., reasoning based on common belief in rationality, but our equilibrium concept does not rely on strategic thinking (with the partial exception of Section 5.1).

3.2. Selfconfirming equilibrium

We analyze a notion of equilibrium that characterizes the steady states of learning dynamics and therefore relaxes the mutual-best-reply condition of the Nash equilibrium concept. Recall that our approach allows for the possibility of agents being unaware of many aspects of the game. In equilibrium, agents best respond to (deterministic) shallow conjectures consistent with the feedback that they receive given the true parameter values (\mathbf{Z}, γ) .

Definition 1. A profile $(a_i^*, \hat{x}_i, \hat{y}_i)_{i \in I} \in \times_{i \in I} (A_i \times X_i \times Y_i)$ of actions and (shallow) deterministic conjectures is a **selfconfirming equilibrium (SCE) at (\mathbf{Z}, γ)** if, for each $i \in I$,

1. (*subjective rationality*) $a_i^* = r_i(\hat{x}_i)$,
2. (*confirmed conjecture*) $v_i(a_i^*, \hat{x}_i, \hat{y}_i) = v_i(a_i^*, \ell_i(\mathbf{a}_{-i}^*), g_i(\mathbf{a}_{-i}^*))$.

The two conditions require that: 1) each agent best responds to her own conjecture; 2) the conjecture in equilibrium must belong to the ex post information set, so that the expected payoff (feedback) coincides with the realized payoff (feedback) given a_i^* , $x_i = \ell_i(\mathbf{a}_{-i}^*)$, and $y_i = g_i(\mathbf{a}_{-i}^*)$, where the aggregators ℓ_i and g_i are determined by \mathbf{Z} and γ as in (4) and (5) respectively. We say that $\mathbf{a}^* = (a_i^*)_{i \in I}$ is a **selfconfirming action profile at (\mathbf{Z}, γ)** if there exists a corresponding profile of conjectures $(\hat{x}_i, \hat{y}_i)_{i \in I}$ such that $(a_i^*, \hat{x}_i, \hat{y}_i)_{i \in I}$ is a selfconfirming equilibrium, and we let $\mathbf{A}_{\mathbf{Z}, \gamma}^{SCE}$ denote the set of these action profiles. Also, we denote by $\mathbf{A}_{\mathbf{Z}}^{NE}$ the set of (pure) Nash equilibria of the game (hence neglecting the non-strategic global externalities), that is,

$$\mathbf{A}_{\mathbf{Z}}^{NE} := \{ \mathbf{a}^* \in \times_{i \in I} A_i : \forall i \in I, a_i^* = r_i(\ell_i(\mathbf{a}_{-i}^*)) \}.$$

Since, the joint best-reply function $\mathbf{a}^* \mapsto (r_i(\ell_i(\mathbf{a}_{-i}^*)))_{i \in I}$ is a continuous self-map on the compact and convex subset $\times_{i \in I} [0, \bar{a}_i] \subseteq \mathbb{R}^I$, Brouwer Fixed Point Theorem implies that a Nash equilibrium exists. Hence, we obtain the existence of selfconfirming equilibria. Indeed, a Nash equilibrium \mathbf{a}^* corresponds to a selfconfirming equilibrium with correct conjectures $(a_i^*, \hat{x}_i, \hat{y}_i)_{i \in I} = (a_i^*, \ell_i(\mathbf{a}_{-i}^*), g_i(\mathbf{a}_{-i}^*))_{i \in I}$. To summarize:

Remark 1. For every \mathbf{Z} and γ , there is at least one Nash equilibrium, and every Nash equilibrium is a selfconfirming action profile at (\mathbf{Z}, γ) :

¹⁰ Assuming that each player i knows that $y_i = 0$, if there are no global externalities..

¹¹ See the analysis in Appendix A.1.

$$\emptyset \neq \mathbf{A}_{\mathbf{Z}}^{NE} \subseteq \mathbf{A}_{\mathbf{Z},\gamma}^{SCE}.$$

In the following sections we study selfconfirming equilibria and learning, first when there are only local externalities, and then when also global externalities are considered.

4. Local externalities

In this section, we analyze the set of selfconfirming equilibria and the learning paths in linear-quadratic network games with just observable payoffs and *without global externalities*, that is, when $\gamma = 0$ and each player knows it, or at least knows that $y_i = 0$. Several proofs are derived from the results in Appendix A, which refers to the case of generic network games with feedback, and from the results in Appendix B. The proofs themselves are collected in Appendix C. In subsection 4.1 we characterize the set of selfconfirming equilibria at $(\mathbf{Z}, 0)$, $\mathbf{A}_{\mathbf{Z},0}^{SCE}$, relating them to the Nash equilibria of auxiliary reduced games and we classify equilibria according to the set of active agents. In subsection 4.2 we provide properties of \mathbf{Z} that imply uniqueness of active agents' equilibrium actions. In subsection 4.3 we analyze learning paths.

4.1. Nash equilibrium and structure of the SCE set

Let I_0 denote the **set of players for whom being inactive is justifiable** (that is, undominated):¹²

$$I_0 := \{i \in I : \exists x_i \in X_i, r_i(x_i) = 0\} = \{i \in I : \alpha_i + \underline{x}_i \leq 0\}.$$

Also, for each non-empty subset of players $J \subseteq I$, let $\mathbf{A}_{\mathbf{Z}}^{NE,J}$ denote the set of Nash equilibria of the auxiliary game with player set J obtained by imposing $a_i = 0$ for each $i \in I \setminus J$, that is,

$$\mathbf{A}_{\mathbf{Z}}^{NE,J} = \left\{ \mathbf{a}_J^* \in \times_{j \in J} A_j : \forall j \in J, a_j^* = r_j \left(\ell_j \left(\mathbf{a}_{J \setminus \{j\}}^*, \mathbf{0}_{I \setminus J} \right) \right) \right\},$$

where $\mathbf{0}_{I \setminus J} \in \mathbb{R}^{I \setminus J}$ is the profile that assigns 0 to each $i \in I \setminus J$. If $J = \emptyset$, let $\mathbf{A}_{\mathbf{Z}}^{NE,J} = \{\emptyset\}$ by convention, where \emptyset is the pseudo-action profile such that $(\emptyset, \mathbf{0}_I) = \mathbf{0}_I$.¹³ We relate the set of selfconfirming equilibria to the sets of Nash equilibria of such auxiliary games.

Proposition 1. *In a linear-quadratic network game with just observable payoffs, and without global externalities, the set of selfconfirming action profiles at $(\mathbf{Z}, 0)$ is*

$$\mathbf{A}_{\mathbf{Z},0}^{SCE} = \bigcup_{J: I \setminus J \subseteq I_0} \mathbf{A}_{\mathbf{Z}}^{NE,J} \times \{\mathbf{0}_{I \setminus J}\},$$

that is, in each selfconfirming action profile \mathbf{a}^* , a subset $I \setminus J$ of players for whom being inactive is justifiable choose 0, and every other player chooses the best reply to the actions of her co-players. Therefore, in each selfconfirming action profile \mathbf{a}^* and for each player $i \in I$,

¹² This definition is motivated by Lemma 1 in Appendix A, in which we analyze also the more general case of probabilistic conjectures and we explain why restricting attention to deterministic conjectures is without loss of generality.

¹³ As we do in set theory with the empty set, when we consider functions whose domain is a subset J of some index set I , it is convenient to have a symbol for the pseudo-function with empty domain. For example, if $J \subseteq I = \mathbb{N}$, such functions are (finite and countably infinite) sequences and \emptyset denotes the empty sequence.

$$\begin{aligned}
 a_i^* = 0 &\Rightarrow x_i \leq -\alpha_i, \\
 a_i^* > 0 &\Rightarrow \left(\alpha_i + \sum_{j \in I} z_{ij} a_j^* > 0 \wedge a_i^* = \min \left\{ \bar{a}_i, \alpha_i + \sum_{j \in I} z_{ij} a_j^* \right\} \right). \tag{7}
 \end{aligned}$$

Note that, if being inactive is justifiable for every agent ($I_0 = I$), then $\mathbf{0}_I \in \mathbf{A}_{\mathbf{Z},0}^{SCE}$. In the polar opposite case, being inactive is unjustifiable for every agent ($I_0 = \emptyset$) and the SCE action profiles coincide with Nash equilibrium (NE) profiles. Thus, the SCE set can be characterized by means of the NE of the auxiliary games in which only active agents are considered. If, for example, for every given set $J \subseteq I$ there is a unique NE of the corresponding auxiliary game (Proposition 2 provides sufficient conditions), then $|\mathbf{A}_{\mathbf{Z},0}^{SCE}| = 2^{|I_0|}$, because for each J with $I \setminus J \subseteq I_0$ there is exactly one SCE where the set of active agents is J . Since each auxiliary game has at least one NE (see Remark 1), $2^{|I_0|}$ is a lower bound on the number of SCE's. If we assume strategic substitutes, then the Nash equilibria for each auxiliary game in which only agents in $J \subseteq I$ may be active, can be characterized as in Bramoullé et al. (2014). Note that in this case, some of the agents in J can be active and some inactive. It is also interesting to note that, if the local externalities are all non-negative (strategic complementarities), agents will always be (weakly) less active in any SCE compared to the NE.¹⁴ Appendix A.3 discusses the equilibrium characterization for the general case of non linear-quadratic network games.

4.2. Relative uniqueness

We now list and briefly discuss some properties of the weighted adjacency matrix \mathbf{Z} that will be used throughout the text but are *not* maintained assumptions.¹⁵ In what follows, we will assume some of these properties to retrieve sufficient conditions for the existence and stability of selfconfirming equilibria. In particular, they imply the uniqueness of SCE actions relative to any given set J of active players. We refer to Appendix B for a deeper discussion of these assumptions and their implications.

Assumption 1. Matrix \mathbf{Z} of size n has bounded values, i.e., for each $i, j \in I$, $|z_{ij}| < \frac{1}{n}$.

This assumption simply states that there is no agent who has a link with an excessive weight compared to the size of the network.

Assumption 2. Matrix \mathbf{Z} has the same sign property, i.e., for each $i, j \in I$, $sign(z_{ij}) = sign(z_{ji})$, where the sign function can have values $-1, 0$ or 1 .

This assumption requires a sort of symmetry in the way two agents influence each other. Namely, the local externalities they impose on each other must be both positive or both negative.¹⁶ The next assumption, instead, requires that all local externalities are strictly negative.

¹⁴ We are also implicitly assuming high enough upper bounds on the actions, so that the NE is indeed unique. On this, see previous footnote 7.

¹⁵ That is, they appear explicitly among the hypotheses of some of the subsequent propositions.

¹⁶ This sign condition is used in Bervoets et al. (2019) to prove convergence to Nash equilibria in network games, under a particular form of learning.

Assumption 3. Matrix \mathbf{Z} is negative, i.e., for each $i, j \in I, z_{ij} \leq 0$.

The following assumption provides some conditions on the spectral properties of the network. Recall here the spectral radius $\rho(\mathbf{Z})$ of \mathbf{Z} is the largest absolute value of its eigenvalues.

Assumption 4. Matrix \mathbf{Z} is limited, i.e., $\rho(\mathbf{Z}) < 1$.

This assumption, which is quite common in the network literature, states that the network should not be excessively tight or densely interconnected. In fact, $\rho(\mathbf{Z})$ serves as a measure of this connectedness, as it falls between the minimum and maximum sums of the entries in each row. Note that Assumption 1 implies Assumption 4.

In some cases, we can write $\mathbf{Z} = \mathbf{W}\mathbf{Z}_0$, where \mathbf{W} is a diagonal matrix, and $\mathbf{Z}_0 \in \{0, 1\}^{I \times I}$ is the basic underlying topology of the network. Whenever this is the case, matrix \mathbf{Z} represents a basic network combined with an additional idiosyncratic effect by which every agent i weights the effects of others on her. These effects are modeled by the parameter w_i (the i^{th} entry in the diagonal of \mathbf{W}).¹⁷ The next assumption adds a symmetry condition on \mathbf{Z}_0 .

Assumption 5. Matrix \mathbf{Z} is symmetrizable, i.e., it can be written as $\mathbf{Z} = \mathbf{W}\mathbf{Z}_0$, with \mathbf{W} diagonal and \mathbf{Z}_0 symmetric. Moreover, \mathbf{W} has all strictly positive entries in the diagonal.

Note that if \mathbf{Z} is symmetrizable then all its eigenvalues are real. Moreover, since \mathbf{W} has all strictly positive entries in the diagonal, Assumption 5 implies that the sign condition (Assumption 2) holds.

Our final assumption is discussed in Bramoullé et al. (2014) and combines Assumptions 4 and 5 above.

Assumption 6. Matrix $\mathbf{Z} = \mathbf{W}\mathbf{Z}_0$ is symmetrizable-limited, i.e., \mathbf{Z} is symmetrizable and the matrix $\bar{\mathbf{Z}}$, whose entries are defined, for each $i, j \in I$, as $\bar{z}_{ij} = z_{0,ij} \sqrt{w_i w_j}$, is limited.

Our previous results about the characterization of selfconfirming equilibria state that we can choose any subset $J \subseteq I_0$ of agents and have them inactive in an SCE. However, we cannot ensure that the other agents are active, because their best response in the reduced game could be to stay inactive, since the Nash equilibrium of the reduced game in which only agents in $I \setminus J$ are considered may have both active and inactive agents. The next result goes in the direction of specifying under what sufficient conditions this does not happen. Given the matrix \mathbf{Z} , and given $J \subseteq I$, we call \mathbf{Z}_J the submatrix which has only rows and columns corresponding to the elements of J .

Proposition 2. Consider a linear-quadratic network game and a subset of players $J \subseteq I$ such that $I \setminus J \subseteq I_0$ (that is, $\alpha_i + \underline{x}_i \leq 0$ for each $i \notin J$). Suppose that \mathbf{Z}_J satisfies at least one of the three conditions below:

¹⁷ Then the payoff of $i \in I$ at a given profile \mathbf{a} of the original game is

$$u_i(\mathbf{a}) = \alpha_i a_i - \frac{1}{2} a_i^2 + a_i w_i \sum_{j \in I} z_{0,ij} a_j = \alpha_i a_i - \frac{1}{2} a_i^2 + a_i \sum_{j \in I} z_{ij} a_j .$$

Table 1

Selfconfirming equilibria of the network from Fig. 1, with positive (resp., negative) externalities of intensity 0.2 (resp., -0.2). Columns correspond to subsets of active players. The unique Nash Equilibrium is in bold.

	All	{1, 2, 3}	{1, 2, 4}	{1, 3, 4}	{2, 3, 4}	{1, 2}	{1, 3}	{1, 4}	{2, 3}	...	\emptyset
a_1	0.1257	0.1	0.125	0.128	0	0.1	0.1	0.125	0		0
a_2	0.1603	0.1346	0.15	0	0.144	0.12	0	0	0.1154		0
a_3	0.0412	0.731	0	0.720	0.1	0	0.1	0	0.0729		0
a_4	0.1336	0	0.125	0.14	0.12	0	0	0.125	0		0

1. it has bounded values (Assumption 1);
2. it is negative and limited (Assumptions 3 and 4);
3. it is symmetrizable-limited (Assumption 6).

Then, the following statements hold:

- the auxiliary game with player set J has a unique and strictly positive Nash equilibrium: $\mathbf{A}_J^{NE} = \{\mathbf{a}_j^{NE}\}$ with $a_j^{NE} > 0$ for all $j \in J$;
- $(\mathbf{a}_J^{NE}, \mathbf{0}_{I \setminus J})$ is a selfconfirming equilibrium at $(\mathbf{Z}, 0)$.

Proposition 2 provides sufficient conditions to have sets of active and inactive players in a selfconfirming equilibrium. In particular, if any of the three conditions is satisfied for every subset of I , and if being inactive is justifiable for all the players ($I_0 = I$), then the set of SCE's has the same cardinality as the power set 2^I , that is 2^n . The first sufficient condition about (sub)matrix \mathbf{Z}_J is novel, while the other two were obtained respectively by Ballester et al. (2006) and Stańczak et al. (2006), and by Bramoullé et al. (2014).

We provide below an example with mixed externalities.

Example 1. Proposition 2 provides alternative *sufficient* conditions for an interior Nash Equilibrium (NE) in the auxiliary game with player set J . Fig. 1 provides an example of game that does not satisfy any of them, but still has a unique interior NE. We set $\alpha_i = 0.1$ for each player i . Every blue arrow represents a positive externality of intensity 0.2. The two red arrows represent negative externalities of intensity -0.2. This network game has a unique NE, and 16 SCE's. Table 1 shows them all (redundant doubletons and singletons are omitted).

4.3. Learning paths

Definition 1 of selfconfirming equilibrium and the characterization stated in Proposition 1 identify steady states: if agents' conjectures are confirmed (not contradicted) by the feedback they receive, these conjectures will not change in the next interactions. However, we may wonder how agents get to play SCE action profiles and if these profiles are stable.¹⁸

We first point out that SCE has solid learning foundations.¹⁹ The following result is specifically relevant for this paper (see Gilli, 1999 and Chapter 7 of Battigalli et al., 2023). Consider a

¹⁸ Throughout all our analysis, players perform adaptive learning given an exogenously fixed (but possibly unknown) network. For models in which players adaptively change also their links, with a quadratic payoff function analogous to ours, and the overall network evolve endogenously, see König and Testone (2011) and König et al. (2014).

¹⁹ See, for example, Battigalli et al. (2019), Fudenberg and Kreps (1995), and the references therein.




Fig. 1. A network with 4 nodes. Blue (resp., red) arrows represent positive (resp., negative) externalities.

temporal sequence (path) of action profiles $(\mathbf{a}_t)_{t=0}^\infty$. Then, if $(\mathbf{a}_t)_{t=0}^\infty$ is consistent with adaptive learning²⁰ and $\mathbf{a}_t \rightarrow \mathbf{a}^*$, it follows that \mathbf{a}^* must be a selfconfirming action profile.

To ease the analysis, we consider conjectural best-reply paths for shallow conjectures. For each network \mathbf{Z} , each period $t \in \mathbb{N}_0$, and each agent $i \in I$, $a_{i,t} = r_i(\hat{x}_{i,t})$ is the best reply to $\hat{x}_{i,t}$. After actions are chosen, given the feedback received, agents update their conjectures. If her previous conjecture is confirmed, then agent i keeps it, otherwise she updates it using as new conjecture the one that would have been correct in the previous period. Thus,

$$\hat{x}_{i,t+1} = \begin{cases} \hat{x}_{i,t} & \text{if } a_{i,t} = 0, \\ \ell_i(\mathbf{a}_{-i,t}) & \text{if } a_{i,t} > 0, \end{cases} \tag{8}$$

and, from (6) we obtain

$$a_{i,t+1} = r_i(\hat{x}_{i,t+1}) = \begin{cases} 0, & \text{if } \hat{x}_{i,t} \leq -\alpha_i, \\ \bar{a}_i, & \text{if } \hat{x}_{i,t+1} \geq \bar{a}_i - \alpha_i, \\ \alpha_i + \hat{x}_{i,t+1}, & \text{otherwise.} \end{cases}$$

We will consider the possibility that the upper bound \bar{a}_i is reached only in the analysis of diverging dynamics. Given our assumptions about feedback, *being inactive is an absorbing state*: if an agent is inactive at time t she will remain so also at time $t + 1$. If instead the agent is active ($a_{i,t} > 0$), feedback is such that the agent can perfectly infer the payoff state $x_{i,t} = \ell_i(\mathbf{a}_{-i,t})$, and so she updates conjectures according to (8), which becomes the updated conjecture. This is a conjectural best-reply path. The result cited above implies that if the path described above converges, then it must converge to a selfconfirming equilibrium, i.e., a rest point where players keep repeating their choices.

In this subsection, we analyze the local stability of such rest points (cf. Bramoullé and Kranton, 2007).

Definition 2 (*Conjectural best-reply paths*). A sequence of profiles of actions and shallow deterministic conjectures $(\mathbf{a}_t, \hat{\mathbf{x}}_t)_{t \in \mathbb{N}_0}$ is a **conjectural best-reply path** if it has the following features:

1. Each player $i \in I$ starts at time 0 with a belief, and beliefs are represented by a profile of shallow deterministic conjectures $\hat{\mathbf{x}}_0 = (\hat{x}_{i,0})_{i \in I}$.

²⁰ In a *finite* game, a path of play $(\mathbf{a}_t)_{t=0}^\infty$ is consistent with adaptive learning if for every \hat{t} , there exists some T such that, for every $t > \hat{t} + T$ and $i \in I$, $a_{i,t}$ is a best reply to some *deep* conjecture μ_i that assigns probability 1 to the set of action profiles \mathbf{a}_{-i} consistent with the feedback received from \hat{t} through $t - 1$. The definition for compact-continuous games is a bit more complex (see Milgrom and Roberts, 1991, who assume perfect feedback).

2. In each period t , players best reply to their conjectures: for each $i \in I$, $a_{i,t} = \min\{\max\{\alpha_i + \hat{x}_{i,t}, 0\}, \bar{a}_i\}$.
3. At the beginning of each period $t + 1$, each player i keeps her period- t shallow conjecture if she was inactive, and updates her conjecture to period- t revealed payoff state if she was active, that is, $\hat{x}_{i,t+1} = \frac{u_i(\mathbf{a}_t)}{a_{i,t}} - \alpha_i + \frac{1}{2}a_{i,t}$.

Observe that the system is deterministic and the initial conditions completely determine the paths. From conditions (7) and (8), the system is not linear because, for each $i \in I$ and $t \in \mathbb{N}_0$,

$$\hat{x}_{i,t+1} = \begin{cases} \hat{x}_{i,t} & \text{if } \hat{x}_{i,t} \leq -\alpha_i, \\ \sum_{j \in I} z_{ij} a_{j,t} & \text{if } \hat{x}_{i,t} > -\alpha_i. \end{cases}$$

Clearly an SCE of the game is always a rest point of these learning paths. Indeed, every SCE $(\mathbf{a}^*, \hat{\mathbf{x}})$ is – trivially – the limit of the constant conjectural best-reply path starting at $(\mathbf{a}_0, \hat{\mathbf{x}}_0) = (\mathbf{a}^*, \hat{\mathbf{x}})$. Furthermore, the set of inactive agents in a conjectural best-reply path can only increase:

$$I_0(\hat{\mathbf{x}}_t) \subseteq I_0(\hat{\mathbf{x}}_{t+1}),$$

where $I_0(\hat{\mathbf{x}})$ denotes the set of inactive agents given profile of conjectures $\hat{\mathbf{x}} = (\hat{x}_i)_{i \in I}$.

We now consider the stability of such rest points.

Definition 3. A profile $\mathbf{a}^* \in \mathbf{A}_{\mathbf{Z},0}^{SCE}$ is locally stable if there exists a profile of conjectures $\hat{\mathbf{x}}$ such that $(\mathbf{a}^*, \hat{\mathbf{x}})$ is a selfconfirming equilibrium, and if there exists an $\epsilon > 0$ such that, for each $\hat{\mathbf{x}}_0$ with $\|\hat{\mathbf{x}}_0 - \hat{\mathbf{x}}\| < \epsilon$ (where $\|\cdot\|$ is the Euclidean norm), the conjectural best-reply path, starting at $\hat{\mathbf{x}}_0$, has a limit and it is such that $\lim_{t \rightarrow \infty} \mathbf{a}_t = \mathbf{a}^*$.

Since $(\mathbf{a}_t, \hat{\mathbf{x}}_t)_{t \in \mathbb{N}_0}$ is determined by the initial conjectures $\hat{\mathbf{x}}_0$, we analyze stability with respect to perturbations of $\hat{\mathbf{x}}_0$. Our notion of stability with respect to conjectures relates to the standard notion of stability with respect to actions in the following way. First of all, since played actions are justified by some conjectures, the only reason for these actions to change is a perturbation of the justifying conjectures, but this is not a sufficient condition. If all agents are active, the two definitions have the same consequences in terms of stability, since a perturbation with respect to actions happens if and only if every agent’s conjecture is perturbed. Indeed, each active agent i has perfect feedback about x_t , and always chooses the best reply to neighbors’ actions in previous period. However, consider an SCE with inactive agents, who choose the null action as a corner solution, that is, whose subjective expected marginal utility for increasing activity is strictly negative. For such agents a small perturbation of their conjectures would not change their null subjective best reply. This is so because inactive agents have imperfect feedback and cannot infer the value of the local externality aggregator. This implies that if an action profile is locally stable with respect to action perturbations, then it is also locally stable under conjectures perturbations, but the converse does not hold. Specifically, forcing inactive agents to be active may lead some of them to be active forever. The two definitions would be equivalent under perfect feedback for all agents. Note finally that a temporary perturbation of shallow conjectures $\hat{\mathbf{x}}_0$ has the same effect of a temporary shock in the parameter vector $\{\alpha_i\}_{i \in I}$. By looking at the first-order conditions, they both induce the same effect on agents’ best reply and on payoffs.

Each SCE is characterized by a set of active agents. So, given an action profile $\mathbf{a} = (a_i)_{i \in I}$, let $I_{\mathbf{a}} := \{i \in I : a_i > 0\}$ denote the set of active players at profile \mathbf{a} . Also let $I_0^* := \{i \in I : \alpha_i + \underline{x}_i < 0\}$ (a subset of I_0) denote the set of agents for whom being inactive is a “corner

solution” for a set of conjectures with nonempty interior. For each action profile \mathbf{a} , $\mathbf{Z}_{I_{\mathbf{a}}}$ denotes the sub-matrix with rows and columns corresponding to players who are active in \mathbf{a} . The following result provides sufficient conditions for an SCE to be locally stable.

Proposition 3. *The action profile in a selfconfirming equilibrium $(\mathbf{a}^*, \hat{\mathbf{x}})$, such that $\hat{x}_i \neq -\alpha_i$ for each $i \in I$, is locally stable if*

- Assumption 4 holds for matrix $\mathbf{Z}_{I_{\mathbf{a}^*}}$;
- $I \setminus I_{\mathbf{a}^*} \subseteq I_0^*$.

Intuitively, consider a sufficiently small perturbation of players’ conjectures. The first condition ensures that active players keep being active and their actions converge back to the unique Nash equilibrium of the auxiliary game with player set $I_{\mathbf{a}^*}$. The second condition ensures that inactive players keep being inactive. Next, we provide alternative sufficient conditions that allow to find the subsets of active agents associated to SCE’s.

Proposition 4. *Consider the action profile \mathbf{a}^* in a selfconfirming equilibrium $(\mathbf{a}^*, \hat{\mathbf{x}})$ such that $I \setminus I_{\mathbf{a}^*} \subseteq I_0^*$ and $\hat{x}_i \neq -\alpha_i$ for each $i \in I$. If $\mathbf{Z}_{I_{\mathbf{a}^*}}$ satisfies at least one of the three conditions below:*

1. *it has bounded values (Assumption 1),*
2. *it is negative and limited (Assumptions 3 and 4),*
3. *it is limited and symmetrizable (Assumptions 4 and 5),*

then \mathbf{a}^ is locally stable. Moreover, for every $J \subseteq I_{\mathbf{a}^*}$ such that $I \setminus J \subseteq I_0^*$, $\mathbf{a}^{**} = (\mathbf{a}_J^{NE}, \mathbf{0}_{I \setminus J})$ is a locally stable SCE action profile, where \mathbf{a}_J^{NE} is the unique and strictly positive Nash equilibrium action profile of the auxiliary game restricted to player set J .*

The proof is based on results from linear algebra. In fact, if an adjacency matrix satisfies one of the conditions from Proposition 4, then also every submatrix of that matrix satisfies that property.

We know that there may be SCE’s that are not Nash equilibria, because some agents are inactive even if inactivity is not a best response to the actions of others. Proposition 4 provides an additional observation. Under the stated conditions, for any given SCE action profile \mathbf{a}^* with set of active agents $I_{\mathbf{a}^*}$, any subset $J \subseteq I_{\mathbf{a}^*}$ such that $I \setminus J \subseteq I_0^*$ is associated to a stable SCE where all agents in J are active, and the other agents are inactive.

The following example shows that we can reach SCE’s that are not NE’s also if the initial beliefs induce strictly positive actions for all agents at the beginning of the learning paths.

Example 2. Consider the case of 4 players with the network matrix $\mathbf{Z} \in \{-0.2, 0, 0.2\}^{I \times I}$ shown in Fig. 1, and, for every i , $\alpha_i = 0.1$. This is a case of externalities that can be positive or negative. Fig. 2 shows the learning paths of actions that start from different initial conditions. In one case (left panel) the path converges to the unique Nash equilibrium of this game (the dotted lines), in the other (right panel) the path makes a player inactive after two rounds and converges to a selfconfirming equilibrium which is not Nash.




Fig. 2. Positive and negative externalities. Starting from different conjectures, given the same network (from Figure 1), the learning process may converge to the unique Nash equilibrium (left panel – dotted lines are the Nash equilibrium) or to an SCE which is not a Nash equilibrium (right panel). For active players, actions are just an upward shift of conjectures by amount α_i . In the right panel, for the inactive player 3 the action is 0 from step 2 on.

5. Local and global externalities

In many applications the feedback that active players receive is not enough to find out the objectively optimal response. Users of online platforms may not understand *ex post* the objective best response to others’ activity. In our context, this means that perfect feedback may *not* hold even for active players. In particular, this is the case if players just observe their realized payoffs, but there may be global externalities, which introduce a confound. This implies there may be other equilibria besides those analyzed above. Assuming that local externalities are positive, the following analysis yields two important observations. First, players may be more active if they think that they are more linked in the network than they actually are, and this can be welfare improving for the whole society. Second, agents with excessive perceived connectedness may prevent convergence of best reply paths to interior equilibria. Recall Definition 1 (of self-confirming equilibrium), based on general linear-quadratic network games with just observable payoff (see equations (3)-(5)). We can characterize the set of SCE’s as follows:

Proposition 5. A profile of actions and conjectures $(a_i^*, \hat{x}_i, \hat{y}_i)_{i \in I} \in \times_{i \in I} (A_i \times X_i \times Y_i)$ in a linear-quadratic network game with just observable payoffs and local and global externalities is a selfconfirming equilibrium at (\mathbf{Z}, γ) if and only if, for every $i \in I$,

1. $a_i^* = 0$ implies $\hat{x}_i \in [x_i, -\alpha_i]$ and $\hat{y}_i = \gamma \sum_{j \neq i} a_j^*$;
2. $a_i^* > 0$ implies $a_i^* = \min\{\alpha_i + \hat{x}_i, \bar{a}_i\}$ and $\hat{y}_i = \gamma \sum_{k \neq i} a_k^* + a_i^* \left(\sum_{j \neq i} z_{ij} a_j^* - \hat{x}_i \right)$.

We discuss how the presence of the global externality term in the utility function changes the characterization of selfconfirming equilibria. Although we maintain the assumption of just observable payoffs, with global externalities it is not anymore the case that active players have perfect feedback about the payoff state. Indeed, for all $i \in I$ and for all pairs of realized externalities (x_i, y_i) , $v_i(0, x_i, y_i) = y_i$. Thus, on the one hand inactive players have correct conjectures about the global externality, but may have incorrect conjectures about the local externality. On the other hand, active players are not able to determine the relative magnitude of the local effects with respect to the global effects. Given any strictly positive action a_i^* , the confirmed conjectures condition yields $(\hat{y}_i - y_i) = a_i^* (x_i - \hat{x}_i)$. Then, in equilibrium, if agent i overestimates (underestimates) the local externality, she must compensate this error by underestimating (overestimating) the global externality. Compared to the case of only local externalities, we have that:

(i) active agents may have a wrong conjecture about the payoff state; thus, (ii) it is not possible to completely characterize the set of SCE's by means of Nash equilibria of the auxiliary games restricted to the active players.

Yet, the analysis of Section 4 allows to identify a subset of selfconfirming equilibria, those where agents have correct (shallow) conjectures about the global payoff state.

Remark 2. The set of SCE action profiles of the network game with only local externalities is included in the set of SCE action profiles of the game with local and global externalities, that is, $\mathbf{A}_{\mathbf{Z},0}^{SCE} \subseteq \mathbf{A}_{\mathbf{Z},\gamma}^{SCE}$. Specifically, if $(a_i^*, \hat{x}_i)_{i \in I}$ is an SCE of the game with only local externalities, then $(a_i^*, \hat{x}_i, \hat{y}_i)_{i \in I}$ with $\hat{y}_i = \gamma \sum_{k \neq i} a_k^*$ for each $i \in I$ is an SCE of the game with local and global externalities.

Indeed, by Proposition 1, in profile $(a_i^*, \hat{x}_i)_{i \in I}$ of the game with only local externalities, each inactive player has a (trivially) confirmed conjecture that makes her choose 0, and each active player must have a correct conjecture about the local externality. In profile $(a_i^*, \hat{x}_i, \hat{y}_i)_{i \in I}$ of the game with global externalities, conjectures $(\hat{y}_i)_{i \in I}$ about these externalities are correct by assumption. Thus, by Proposition 5, $(a_i^*, \hat{x}_i, \hat{y}_i)_{i \in I}$ is an SCE.

Proposition 5 characterizes the set of SCEs considering only shallow conjectures. Therefore, it allows for the possibility that agents know very little about the situation of strategic interaction they are facing. As we mentioned in the Introduction, an agent i may even be unaware that she is interacting with other agents and that the external states x_i and y_i aggregate the actions of others. Suppose instead that each i knows that she is interacting with others and knows the parametric form of her payoff function, but has incomplete information about the parameter values. *When is an SCE action profile supportable by deep conjectures?* Assume that each player i only knows that co-players' actions and global externalities are non-negative, while – as far as they know – some links z_{ij} might be negative. We can show that, for each SCE with shallow conjectures $(a_i^*, \hat{x}_i, \hat{y}_i)_{i \in I}$, there is at least one corresponding profile of deep (deterministic) conjectures $(\hat{\mathbf{a}}_{-i}^i, \hat{\mathbf{Z}}^i, \hat{\gamma}^i)_{i \in I}$ inducing the shallow conjectures $(\hat{x}_i, \hat{y}_i)_{i \in I}$. For example, in SCEs with at least 2 active players, we can consider the case where each i believes that all co-players j are equally active at level $\hat{a}_j^i = \hat{a} > 0$. With this, we can derive corresponding conjectures about parameters $(\hat{z}_{ij}^i)_{j \neq i}$ (specifying the rest of $\hat{\mathbf{Z}}^i$ arbitrarily) and $\hat{\gamma}^i$ that induce \hat{x}_i and \hat{y}_i . It follows that $(a_i^*, \hat{\mathbf{a}}_{-i}^i, \hat{\mathbf{Z}}^i, \hat{\gamma}^i)_{i \in I}$ is an SCE with deep conjectures. In particular, if i is inactive it must be the case that $\hat{x}_i < 0$ and i must believe that some links z_{ij} are negative.

To ease the following analysis, in the remainder of this whole section, we assume that (i) each agent i has the same stand-alone parameter $\alpha > 0$ and upper bound $\bar{\alpha}$, and (ii) $\gamma > 0$ (that is, global externalities are real, not just a possibility in the mind of players). We assume also that (iii) matrix \mathbf{Z} is non-negative, and (iv) either condition 1. or 3. of Proposition 2 is satisfied, so that there exists a unique NE. Finally, (v) we assume that the admissible range of possible best replies of each player contains the upper bound $\bar{\alpha}$.

Understanding how conjectures are shaped in an SCE also allows us to shed some light on the efficiency properties of the SCE's. First of all note that the problem of finding a maximizer of the sum of the utilities is a concave quadratic problem and there exists a bliss point. The presence of positive externalities makes the unique Nash equilibrium be Pareto-dominated by other actions

profiles. Moreover, the presence of a bliss point makes an arbitrary increase of agents' actions not always welfare improving. Let us analyze these issues in detail.

Given the presence of global externalities, it is straightforward to see that the Nash equilibrium is inefficient. Now consider an SCE action profile \mathbf{a}^{SCE} (possibly \mathbf{a}^{NE}). This action profile is justified by some profile of confirmed conjectures $(\hat{x}_i, \hat{y}_i)_{i \in I}$. Then, we can find another SCE, $\mathbf{a}'^{SCE} \geq \mathbf{a}^{SCE}$, such that \mathbf{a}'^{SCE} yields a higher aggregate payoff than \mathbf{a}^{SCE} . A possible way to find such an equilibrium is to decrease, for each $i \in I$, the global externality (shallow) conjecture \hat{y}_i . To keep the confirmation condition, it is necessary to increase the local (shallow) conjectures $(\hat{x}_i)_{i \in I}$, and thus to increase the best-reply actions. This, in turn, makes the local and global externality states $(x_i, y_i)_{i \in I}$ increase. However, this makes it necessary that the local conjectures are further increased, which induces another increase in actions, and so on. The following proposition imposes a condition for the existence of an interior SCE.

Proposition 6. Fix a profile of local conjectures $\hat{\mathbf{x}}$. If for every pair of agents (i, j) the following inequality is satisfied

$$\sum_{k \in I \setminus \{i, j\}} z_{ik} (\alpha + \hat{x}_k) - z_{ij} \sum_{h \in I \setminus \{i, j\}} (\alpha + \hat{x}_h) \alpha \geq 0, \tag{9}$$

then, for every profile of global conjectures $\hat{\mathbf{y}}$ with $\hat{y}_i < \bar{a}(\alpha \sum_{k \in I \setminus \{i\}} z_{ij} + \gamma n)$ for every i , there exists a unique SCE with local conjectures $\hat{\mathbf{x}}$ and action profile \mathbf{a}^* , with $a_i^* < \bar{a}$.

The condition of the proposition imposes concavity on some fixed point equations derived from the best replies functions, and then ensures existence and uniqueness of this fixed point. Note that such condition is always satisfied if $\alpha \leq 1$ and $\mathbf{Z} = w\mathbf{Z}_0$, with $w > 0$, and $\mathbf{Z}_0 \in \{0, 1\}^{I \times I}$ is the unweighted network. That is: every strictly positive z_{ij} has the same value for each pair of agents i and j in I . Otherwise, the larger the number of agents, the more likely it is that the condition is violated for some pair (i, j) for which z_{ij} is high. If the network is composed of just two agents, this condition is always satisfied.

To better understand the structure of the equilibrium set, we introduce additional assumptions about what agents know or think they know about the strategic environment. This is a way to restrict their conjectures. We provide some insights along two different dimensions: *i)* What happens if agents know something about the magnitude of the externalities? *ii)* What happens if agents have definite beliefs about the relative size of local with respect to global externality?

5.1. Knowledge of externalities parameters

We assume that $\mathbf{Z} = w\mathbf{Z}_0$, where $w > 0$. This means that there is a homogeneous positive externality w between all connected players, so that equation (2) becomes:

$$u_i(a_i, \mathbf{a}_{-i}) = \alpha_i a_i - \frac{1}{2} a_i^2 + a_i w \sum_{j \in I \setminus \{i\}} z_{0,ij} a_j + \gamma \sum_{k \in I \setminus \{i\}} a_k . \tag{10}$$

We do not impose any further restriction over the network structure \mathbf{Z}_0 , but we assume that players understand that they interact in a network and know w and γ . Given these assumptions, we need to slightly modify our definitions of aggregators and conjectures. In detail, aggregators

about local and global externalities do not internalize w and γ , respectively, and the conjectures concern the aggregate actions of the neighbors (local) and of all other players (global).

Consider the case in which $\mathbf{Z} = w\mathbf{Z}_0^c$, where \mathbf{Z}_0^c is the matrix of the complete basic network (i.e., $z_{0,i,j} = 1$ for all non-diagonal entries). Note that if all players are certain that the network is a complete one, then, for each $i \in I$, $\hat{x}_i = \hat{y}_i$, and this ensures uniqueness of the SCE with such complete-network conjectures. Then the SCE can just be indexed by the conjecture about the local externality.²¹ Given (w, γ) , let $(a_i^c(w, \gamma), \hat{x}_i^c(w, \gamma))_{i \in I}$ denote the unique SCE in which, for each $i \in I$, $\hat{x}_i^c(w, \gamma)$ is the (confirmed) shallow conjecture induced by $\hat{\mu}_i^c \in \{\mathbf{Z}_0^c\} \times \mathbf{A}_{-i}$, that is, a (confirmed) deep conjecture in which i thinks she belongs to a complete network.

Proposition 7. Consider a linear quadratic network game with global externalities, with $0 < w < \frac{1}{n-1}$, and where all agents know w and γ . Let $\mathbf{a}_{\mathbf{Z}_0}^{NE}$ and $\mathbf{a}_{\mathbf{Z}_0^c}^{NE}$ be the unique Nash equilibria of the game played on $(w\mathbf{Z}_0, \gamma)$ and $(w\mathbf{Z}_0^c, \gamma)$, respectively. Assume also that the profile of upper bounds is above $\mathbf{a}_{\mathbf{Z}_0}^{NE}$. Then, (1) for each $i \in I$, $a_i^c(w, \gamma)$ is increasing in the ratio $\frac{\gamma}{w}$; (2) $\lim_{\frac{\gamma}{w} \rightarrow 0} \mathbf{a}^c(w, \gamma) = \mathbf{a}_{\mathbf{Z}_0}^{NE}$; and (3) $\lim_{\frac{\gamma}{w} \rightarrow \infty} \mathbf{a}^c(w, \gamma) = \mathbf{a}_{\mathbf{Z}_0^c}^{NE}$.

So, independently of the basic network \mathbf{Z}_0 , if all players believe to be more linked than they actually are and $\frac{\gamma}{w}$ is large, then the action profile approaches what they would choose in the NE of the game played on the complete network, where every player is linked to every other player.

As it will be clear from Section 5.2, this result implies that the learning paths are self-reinforcing. Players maintain wrong conjectures about the network structure and they infer $\ell_i(\mathbf{a}_{-i})$ from the payoff that they receive as feedback, using (10). This implies that, in a conjectural best-reply path, as they increase their own action they infer a higher $\ell_i(\mathbf{a}_{-i})$ and a lower $g_i(\mathbf{a}_{-i})$, to which they will respond with an even higher action. Nevertheless, this process does not diverge to hit the upper bounds of the action profiles, and it reaches the Nash equilibrium on the complete network.

Proposition 7 is a limiting result. However, for some networks where NE's and SCE's can be easily computed analytically, we can show that the SCE actions converge rapidly to the actions of the NE for the complete network as γ/w becomes large. Fig. 3 shows how this happens when every player has the same number of links (regular network) and when there is a central player and every other player is linked only to her (star network).

In the Introduction we discussed the possible application of our model to online social networks, where the provider may have the possibility to affect the beliefs of the users. The previous result applies to the case where users know the value of the parameters w and γ , and their overall number n . If we further assume that the profits of the provider are positively correlated with the overall activity on the platform, the provider may have an incentive to make people feel more connected than they actually are. So, if $\frac{\gamma}{w}$ is large (which means, in our interpretation, that most of the payoff for the agents is obtained from using the platform *per se*, and not from actual interaction), and if these parameters are known to the users, companies make more profit by letting players think that they have a lot of followers. With this application in mind, in the end of this section we will extend the discussion about the implications of biased beliefs on aggregate welfare.

Proposition 7 is based on the assumption that players know the values of γ and w . However, if they have wrong beliefs about γ , overestimating it, their actions would even exceed those of

²¹ The discussion below about conjectured ratios will make this point clear.

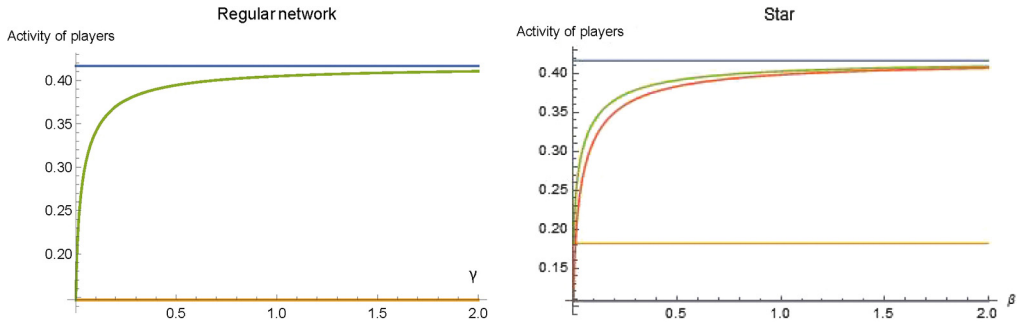


Fig. 3. The panels show the SCE common activity level as a function of parameter γ when each agent thinks she is connected to every other agent. Both cases have parameters $\alpha = 0.1$, $w = 0.04$ and $n = 20$. The left panel is for the regular network with common degree 8: in blue we have the action that would be played in the NE of the complete network; in yellow the NE of the regular network; in green the SCE action. The right panel is for the star network: in blue we have the action that would be played in the NE of the complete network; in yellow and purple the NE action profiles for the center and the spokes, respectively, in the star network; in green and red the SCE action profile for the center and the spokes, respectively.

Table 2

Simulations for the case of $\alpha = 0.1$, $w = 0.2$, and $\gamma = 1$. Columns refer to 1) NE of the line network; 2) Nash equilibrium of complete network; 3) SCE in the line network in which each $i \in I$ believes that $\ell_i(\mathbf{a}_{-i}^*) = \frac{\gamma}{w} g_i(\mathbf{a}_{-i}^*)$.

	Line NE	Complete Network NE	SCE
a_1	0.130	0.167	1.569
a_2	0.152	0.167	1.679
a_3	0.130	0.167	1.569

the NE of the complete network. This is shown in the next example, where agents do not know the true value of γ and, overestimating the ratio between local and global externalities, they play actions that are much above the action that they would play in the NE of the complete network.

Example 3. Consider three agents in a star network (i.e., a line). Let agent 2 be the center. Then, for every SCE, $\ell_2(\mathbf{a}_{-2}^*)$ is proportional to $g_2(\mathbf{a}_{-2}^*)$, always with the same ratio $\frac{\gamma}{w}$, while this is not true for agents 1 and 3. We assume that each agent thinks that the network is complete, so every $i \in I$ thinks that $\ell_i(\mathbf{a}_{-i}^*)$ is proportional to $g_i(\mathbf{a}_{-i}^*)$. In this case agents 1 and 3 believe to be more linked than they actually are. Table 2 provides the Nash equilibria for the actual network and for the complete network, and the selfconfirming equilibrium actions for some specification of the parameters.

This numerical exercise shows that, when agents overestimate the impact of local externalities, we get a *multiplier* effect that makes SCE actions increase at a level even larger than what would be predicted in a complete network by Nash equilibrium. This follows from how agents misinterpret their feedback. In particular, thinking to be in a complete network makes agents 1 and 3 overestimate local externalities. Take for instance agent 1. Given any \mathbf{a}_{-1} , she chooses a subjective best reply higher than the objective best reply since she overestimates the local externality. This high action has the effect of increasing the global externality term for agent 3. Agent 3, by overestimating the local externality, partly attributes this higher global externality to the local externality term, and chooses an action larger than predicted by Nash equilibrium.

The choice of agent 3 increases in turn the global externality perceived by agent 1, and so on. At the same time agent 2, as neighbors choose higher actions, increases her own action level. This effect goes on and gives rise to a multiplier effect. The limit of such a conjectural best reply path is a selfconfirming equilibrium in which actions are almost ten times larger than the complete network NE actions

We call $c_i := \frac{\hat{x}_i}{\hat{y}_i}$ the **conjectured ratio** of player i with respect to local and global externalities. Then, given a profile $(c_i)_{i \in I}$, one can rewrite the SCE conditions as a non-linear system of n equations in n unknowns solved either for $(\hat{x}_i)_{i \in I}$ or $(\hat{y}_i)_{i \in I}$, and characterize the set of SCE's given the imposed restrictions. This is what we will use in the next section when studying the learning paths.

5.2. Learning with global externalities

We now study conjectural best reply paths with global externalities. To simplify the analysis, we assume a *fixed conjectured ratio* for each agent. Differently from Section 5.1, we do not assume agents to know anything about the parameters characterizing the strategic environment. In each period, there are infinitely many profiles of feasible pairs $(\hat{x}_{i,t}, \hat{y}_{i,t})_{i \in I}$ consistent with agents' feedback. For each $i \in I$, and each period $t \in \mathbb{N}$, let $v_{i,t} = u_i(a_{i,t}, \mathbf{a}_{-i,t})$ be the realized payoff that agent i observes. Then, given $v_{i,t-1}$, and considering that agents perfectly recall their past actions, $\hat{y}_{i,t}$ is uniquely determined as a function of $\hat{x}_{i,t}$. In particular, if at each time period t agent i 's conjectures $\hat{x}_{i,t}$ and $\hat{y}_{i,t}$ are consistent with the feedback received at the previous period, we obtain

$$\hat{y}_{i,t+1} = v_{i,t} - \alpha_i a_{i,t} + \frac{1}{2} (a_{i,t})^2 - a_{i,t} \hat{x}_{i,t+1}.$$

Then, we can focus on the path of $\hat{x}_{i,t}$, given by

$$\hat{x}_{i,t+1} = \frac{v_{i,t} - \hat{y}_{i,t+1}}{a_{i,t}} - \alpha_i + \frac{1}{2} a_{i,t}. \tag{11}$$

In this case, active agents do not have perfect feedback, because players' conjectures are bi-dimensional, but feedback (the realized payoff) is one-dimensional. This brings also indeterminacy to the updating rule that players use. To avoid bifurcations at each time period t , we need to use simplifying assumptions on conjectures. We define, for each $i \in I$ and each $t \in \mathbb{N}_0$,²²

$$c_{i,t} := \frac{\hat{x}_{i,t}}{\hat{y}_{i,t}}, \tag{12}$$

and in the following we assume that this *conjectured ratio* is constant along paths of learning dynamics for each player i .

Assumption 7. For each $i \in I$ and for each $t \in \mathbb{N}$, $c_{i,t} = c_{i,t+1} = c_i$.

From equation (11) we get the following learning path, for each agent at each time period:

$$\hat{x}_{i,t+1} = x_{i,t} + \frac{y_{i,t}}{a_{i,t}} - \frac{\hat{y}_{i,t+1}}{a_{i,t}}, \tag{13}$$

²² In doing so, we implicitly assume that players think there are active co-players. This is a reasonable assumption, because under positive externalities any best response should be at least α .

where $x_{i,t}$ and $y_{i,t}$ are the true realized values of the payoff states. Plugging in $c_i = \frac{\hat{x}_{i,t}}{\hat{y}_{i,t}}$ we get, for each t and i ,

$$\hat{x}_{i,t+1} = \frac{c_i}{1 + c_i a_{i,t}} (a_{i,t} x_{i,t} + y_{i,t}). \tag{14}$$

Note that the true ratio of player i at time t is

$$c'_{i,t} := \frac{x_{i,t}}{y_{i,t}},$$

with $c'_{i,t} \in \left[0, \frac{\sum_{j \neq i} z_{ij}}{\gamma}\right]$. For this reason, we also assume that the conjectured ratio of each player i is such that $c_i \in \left(0, \frac{\sum_{j \neq i} z_{ij}}{\gamma}\right]$, and this specifies the set of all admissible conjectured ratios.

With this, the learning dynamic from (13) can be written as

$$\hat{x}_{i,t+1} = c_i y_{i,t} \frac{a_{i,t}^* c'_{i,t} + 1}{a_{i,t}^* c_i + 1}, \tag{15}$$

which implies that the conjecture $\hat{x}_{i,t+1}$ is correct only if $c_i = c'_{i,t}$.

Assuming non-binding upper bounds, we look at best responses $a_{i,t+1} = \alpha_i + \hat{x}_{i,t+1}$, and study the existence and characterization of the steady state of this learning process. Recall that $y_{i,t} = \gamma \sum_{j \neq i} a_{j,t}$. To find a fixed point we look at the system of n equations, one for each i ,

$$H_i(\mathbf{a}^*, \mathbf{c}) := \alpha_i + c_i \left(\gamma \sum_{j \neq i} a_j^* \right) \frac{a_i^* c'_i + 1}{a_i^* c_i + 1} - a_i^* = 0. \tag{16}$$

For comparison, we also study the system of equations that yield the Nash equilibrium of this network game, that is, for each i :

$$F_i(\mathbf{a}^*) := \alpha_i + \sum_{j \neq i} z_{ij} a_j^* - a_i^* = 0. \tag{17}$$

Let $\mathcal{A} \subset [\alpha, \infty)^J$ denote the set of the solutions of system (16). We have the following result.

Proposition 8. *If the system defined by (17) admits a solution \mathbf{a}^* with non-negative entries, then for each profile \mathbf{c} of conjectured ratios also the system defined by (16) admits a solution. Moreover, there is a homeomorphism Φ between the set of all profiles \mathbf{c} and \mathcal{A} . The homeomorphism Φ is strictly monotone with respect to the lattice order on the domain of all profiles \mathbf{c} and the codomain \mathcal{A} .*

The assumption of non-negative solutions implies a unique NE of the game, and we refer to Proposition 2 for sufficient conditions for uniqueness. This result provides information only on the steady states of our learning paths. It is important because it establishes a one-to-one function between profiles of conjectured ratios and SCEs: there is one and only one SCE strategy profile for each profile \mathbf{c} but there may SCEs that do not result from the hypothesized learning paths. The homeomorphism also provides continuity in the initial conjectures, as a marginal change in

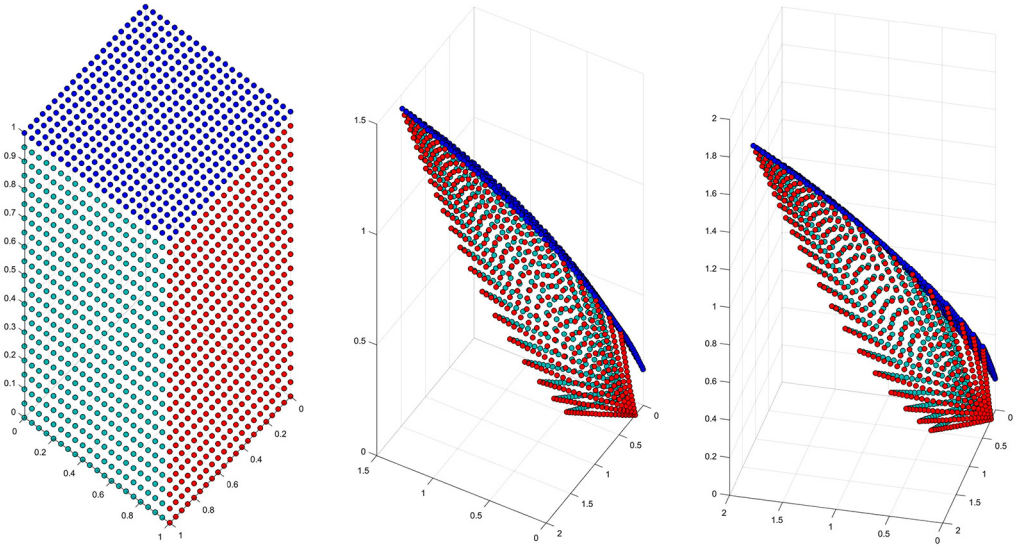


Fig. 4. Simulations showing the homeomorphism of Proposition 9 for the case of 3 nodes, as discussed in Example 4. The left panel shows vectors of conjectured ratios. The central panel shows the corresponding SCE conjecture profile \hat{x} when the network is a line (the node that has conjectured ratio 1 in the red dots is the central node). The right panel shows the corresponding SCE conjecture profile \hat{x} when the network is a complete triangle.

the conjectured ratios will result in a marginal change in the resulting SCE, even if this function may be highly non-linear, as shown in the example below.

Example 4. Under the conditions of Proposition 9, we use equation (14) to express learning paths converging to the SCE implicitly defined by (16). This allows us to provide a graphical illustration of Proposition 8, for the case of three nodes. We do this for the case of a line network (where each of the two links is bidirectional), and for the case of a complete network. We consider equation (10), with $\gamma = 1$ and $w = 0.2$. Fig. 4 shows the results. We can start from any pattern of conjectured ratios for the three nodes. The left panel shows the profile of conjectured ratios when at least one node has maximal conjectured ratio (the three faces of the cube have different colors, according to which node has the maximal centrality). The central panel shows the corresponding SCE conjecture profile \hat{x} when the network is a line (the node that has conjectured ratio 1 in the red dots is the central node). The right panel shows the corresponding SCE conjecture profile \hat{x} when the network is a complete triangle. The figure suggests that homeomorphism Φ (from Proposition 8) is highly non-linear, because of the self-reinforcement process in beliefs that we discussed in Example 3. The figure also shows that, as stated by Proposition 8, homeomorphism Φ respects the lattice order on the two sets.

Monotonicity implies that increasing the conjectured ratio of one player will have a weakly monotonic effect on the action of that player and those of other players in the corresponding SCE. A final *caveat* to keep in mind is that the homeomorphism is implied by the particular learning path that we are assuming, which is based on constant conjectured ratios. Considering the paths in this special case, in the following proposition we show that if local and global externalities are not too large, the learning paths always converge.

Proposition 9. *If, for each player $i \in I$, $0 < c_i \gamma (n - 1) < \sum_{j \neq i} z_{ij} < 2$, then the paths defined by (15) always converge to the unique solution of (16), which is locally stable.²³*

It should be noted that if $\gamma = 0$, i.e., global externalities are objectively absent, the assumptions of Proposition 9 are more general than assuming $|\sum_{j \neq i} z_{ij}| < 1$, which in turn implies that Assumption 4 holds and hence that the learning paths converge. That is because we are focusing on a precise learning path in which players act as if global externalities were present. Moreover, in a game with $\gamma > 0$, if for some players the conjectured ratios are too high, the learning paths defined by (16) may not converge to an interior solution, but rather hit the upper bounds of the feasible action profiles.

Proposition 8 tells us that a non-negative shift in each conjectured ratio will always result in a non-negative shift of each agent’s action in the resulting SCE. However, Proposition 9 gives an implicit warning. Too high conjectured ratios may imply that the sufficient conditions for stability are lost, and convergence to the virtual SCE that we would have without upper boundaries may not occur. Note also that, summing up equation (2) for all the players, the aggregate welfare is maximized if \mathbf{a}^* solves the following linear system of equalities

$$\forall i \in I, a_i^* = \alpha_i + (n - 1)\gamma + \sum_{j \in I \setminus \{i\}} (z_{ij} + z_{ji})a_j^* .$$

To better understand this aspect, consider the online social networks application we often referred to. The results of this last subsection apply to the case where agents do not know the parameters of the model and their own total number, but have only a conjecture about the ratio of the benefits from just using the platform, and from the actual strategic interaction on the platform. Social platforms like Facebook and Twitter often provide information to users about the activity of their peers. The social platform Reddit does not show to users their followers, but only a measure of popularity called *karma*. A rationale for this marketing strategy may be that these companies want to change the beliefs of players, making them feel more important (i.e., more followed) in the social network. Even a benevolent social planner may want to set the conjectured ratios to the level for which the social optimum is achieved. However, according to our model, if conjectured ratios are too high, the learning paths may diverge. For example, in the context of the model and from the assumptions of Proposition 9 a conjectured ratio is *too high* if $c_i \geq \frac{\sum_{j \neq i} z_{ij}}{\gamma(n-1)}$, because in this case learning can lead to SCE where the activity of some player i hits her upper bound a_i and the strategy profile is inefficiently high for the players.

This is shown in the following example.

Example 5. We replicate the same exercise that we did in Example 4, but only for the case of the complete triangle. However we do it for a wider range of conjectured ratios. Fig. 5 shows that in this case there may be combinations of conjectured ratios that prevent convergence of the learning paths to interior equilibria.

²³ Definition 3 of local stability extends naturally to the case of learning with global externalities with paths of the form $(\mathbf{a}_t, \hat{\mathbf{x}}_t, \hat{\mathbf{y}}_t)_{t \in \mathbb{N}_0}$.

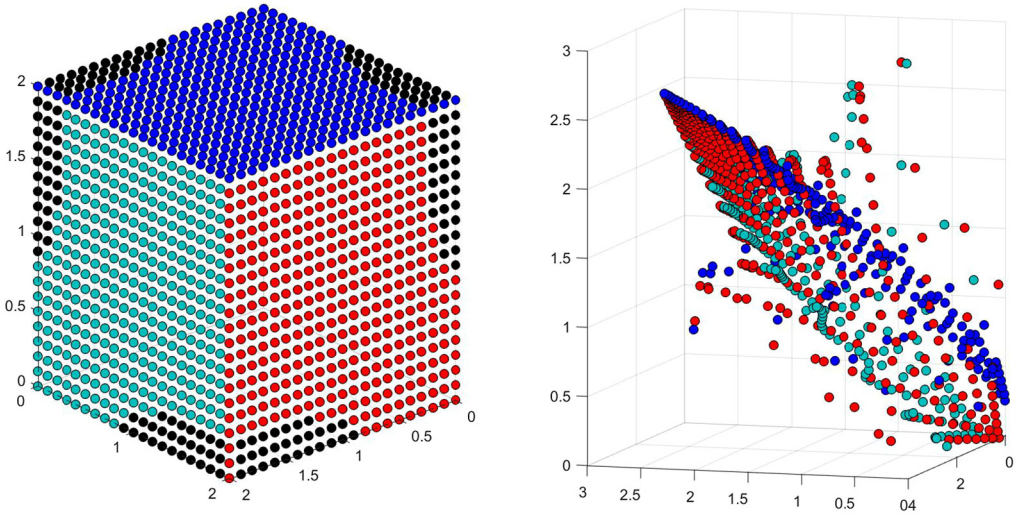


Fig. 5. Simulations showing the homeomorphism of Proposition 9 for the case of 3 nodes, as discussed in Example 5. The left panel shows vectors of conjectured ratios. With respect to Figure 4, we allow for higher values of conjectured ratios. Black dots represent cases for which the learning paths diverge. The right panel shows the corresponding SCE conjecture profile \hat{x} when the network is a complete triangle, and when the learning paths are converging.

6. Conclusion

In this paper we offer a novel approach to network games. A key application of network games is in modelling large societies with millions of agents–nodes and non regular distributions of connections. It is natural to assume that players may ignore the complete structure of the network; this prevents them from performing sophisticated strategic reasoning possibly leading to a Nash equilibrium. Instead, they just best respond to some subjective beliefs affected by the information feedback they receive. We analyze simple conjectural best-reply paths and show that in some cases they converge to stable Nash equilibria. However, we also characterize those situations in which stable action profiles are not Nash equilibria, but rather selfconfirming equilibrium action profiles in which some (if not *all*) players have wrong beliefs and yet the feedback they receive is consistent with such beliefs.

One natural application of this approach is to directed online social platforms like Twitter and Instagram, where links need not to be reciprocated. Using a linear–quadratic structure for the payoff function we have also laid the ground for a tractable welfare analysis of the model. However, policy implications are not straightforward if we want to consider the long run benefits of connections and not only the instantaneous payoffs of the users of those platforms.

Declaration of competing interest

Declarations of interest: none.

Data availability

No data was used for the research described in the article.

Appendix A. Selfconfirming equilibria in parameterized nice games with aggregators

In this section we develop a more general analysis of selfconfirming equilibria in a class of games that contains the linear-quadratic network games with just observable payoffs studied in the main text. To ease reading, we make this section self-contained, repeating some definitions from the main text. We write this section focusing on *local externalities*, because the analysis that follows mainly concerns best-replies, that are not affected by the presence of global externalities. Thus, all the considerations about best-replies in this section also apply to the case of games with both local and global externalities.

A **parameterized nice game with aggregators and feedback** is a structure

$$G = \langle I, \mathcal{Z}, (A_i, \ell_i, v_i, f_i)_{i \in I} \rangle$$

where

- I is the finite **players set**, with cardinality $n = |I|$ and generic element i .
- $\mathcal{Z} \subseteq \mathbb{R}^m$ is a **compact parameter space**.
- $A_i = [0, \bar{a}_i] \subseteq \mathbb{R}_+$, a **compact interval**, is the **action space** of player i with generic element $a_i \in A_i$.
- $X_i = [\underline{x}_i, \bar{x}_i] \subseteq \mathbb{R}$, a **compact interval**, is the **space of payoff states** for i .
- $\ell_i : \mathbf{A}_{-i} \times \mathcal{Z} \rightarrow X_i$ (where $\mathbf{A}_{-i} = \times_{j \in I \setminus \{i\}} A_j$) is a **continuous parameterized aggregator** of the actions of i 's co-players such that its **range** $\ell_i(\mathbf{A}_{-i} \times \mathcal{Z})$ is **connected**.²⁴
- $v_i : A_i \times X_i \rightarrow \mathbb{R}$ is the **utility function** of player i , which is **strictly quasi-concave** in a_i and **continuous**,²⁵ and from which we derive the **parameterized payoff function**

$$\begin{aligned} u_i : A_i \times \mathbf{A}_{-i} \times \mathcal{Z} &\rightarrow \mathbb{R}, \\ (a_i, \mathbf{a}_{-i}, \mathbf{Z}) &\mapsto v_i(a_i, \ell_i(\mathbf{a}_{-i}, \mathbf{Z})). \end{aligned}$$

Thus, $x_i = \ell_i(\mathbf{a}_{-i}, \mathbf{Z})$ is the payoff relevant state that i has to guess in order to choose a subjectively optimal action. With this, for each $\mathbf{Z} \in \mathcal{Z}$, $\langle I, (A_i, u_{i,\mathbf{Z}})_{i \in I} \rangle$ is a nice game (cf. Moulin 1984), and $\langle I, \mathcal{Z}, (A_i, u_i)_{i \in I} \rangle$ is a **parameterized nice game**. We let

$$\begin{aligned} r_i : X_i &\rightarrow A_i \\ x_i &\mapsto \arg \max_{a_i \in A_i} v_i(a_i, x_i) \end{aligned}$$

denote the **best-reply function** of player i . The Maximum theorem implies that r_i is continuous.

- Let $M \subseteq \mathbb{R}$ be a set of “messages,” $f_i : A_i \times X_i \rightarrow M$ is a **continuous feedback function** that describes what i observes (a “message,” e.g., a monetary outcome) after taking any action a_i given any payoff state x_i .

On top of the formal assumptions stated above, we maintain the following *minimal informal assumption* about players’ knowledge of the game:

²⁴ Since the range of each section $\ell_{i,\mathbf{Z}}$ must be a compact interval, we require that the union of the compact intervals $\ell_{i,\mathbf{Z}}(\mathbf{A}_{-i})$ ($\mathbf{Z} \in \mathcal{Z}$) is also an interval, which must be compact because \mathcal{Z} is compact and ℓ_i continuous.

²⁵ That is, v_i is jointly continuous in (a_i, x_i) and, for each $x_i \in [\underline{x}_i, \bar{x}_i]$, the section $v_{i,x_i} : [0, \bar{a}_i] \rightarrow \mathbb{R}$ has a unique maximizer a_i^* (that typically depends on x_i), it is strictly increasing on $[0, a_i^*]$, and it is strictly decreasing on $[a_i^*, \bar{a}_i]$. Of course, the monotonicity requirement holds vacuously when the relevant sub-interval is a singleton.

- Each player i knows v_i and f_i .

Unless we explicitly say otherwise, we instead do not assume that i necessarily knows \mathbf{Z} , or function ℓ_i , or even that i understands that her payoff is affected by the actions of other players. However, since i knows the feedback function $f_i : A_i \times X_i \rightarrow M$ and the action she takes, what i infers about the payoff state x_i after she has taken action a_i and observed message m is that

$$x_i \in f_{i,a_i}^{-1}(m) := \{x'_i : f_i(a_i, x'_i) = m\}.$$

A.1. Conjectures

If player i only knows the feedback function f_i , but does not know how the payoff state x_i is determined, then she just forms a conjecture about x_i . If instead i knows that x_i is determined by the actions of others given parameter \mathbf{Z} through the aggregator ℓ_i , then i forms a conjecture about $(\mathbf{a}_{-i}, \mathbf{Z})$.

Definition 4. A **shallow conjecture** for $i \in I$ is a probability measure $\mu_i \in \Delta(X_i)$. A **deep conjecture** for i is a probability measure $\bar{\mu}_i \in \Delta(\mathbf{A}_{-i} \times \mathcal{Z})$. An action a_i^* is **justifiable** if there exists a shallow conjecture μ_i such that

$$a_i^* \in \operatorname{argmax}_{a_i \in A_i} \int_{X_i} v_i(a_i, x_i) \mu_i(dx_i);$$

in this case we say that μ_i **justifies** a_i^* . Similarly, we say that deep conjecture $\bar{\mu}_i \in \Delta(\mathbf{A}_{-i} \times \mathcal{Z})$ **justifies** a_i^* if the shallow conjecture induced by $\bar{\mu}_i$ ($\mu_i = \bar{\mu}_i \circ \ell_i^{-1} \in \Delta(X_i)$) justifies a_i^* .

The following lemma summarizes well known results about nice games (see, e.g., Battigalli et al., 2023) and some straightforward consequences for the more structured class of nice games with aggregators considered here. We include the proof to make the exposition self-contained.

Lemma 1. *The best-reply function $r_i : X_i \rightarrow A_i$ is continuous, hence its range $r_i(X_i)$ is a compact interval, just like X_i . Furthermore, for each $a_i^* \in A_i$, the following are equivalent:*

- a_i^* is justifiable,
- $a_i^* \in r_i(X_i)$ (that is, a_i^* is justified by a deterministic shallow conjecture),
- there is no a_i such that $v_i(a_i^*, x_i) < v_i(a_i, x_i)$ for all $x_i \in X_i$ (that is, a_i^* is not dominated by any other pure action).

Proof. With a slight abuse of notation, we let $r_i(\mu_i)$ denote the set of best replies to (shallow) conjecture μ_i :

$$r_i(\mu_i) := \operatorname{argmax}_{a_i \in A_i} \int_{X_i} v_i(a_i, x_i) \mu_i(dx_i).$$

By the Maximum theorem $\mu_i \mapsto r_i(\mu_i)$ has a closed graph, which – under the stated assumptions – is equivalent to upper hemi-continuity. By strict quasi-concavity, the restriction of the best-reply correspondence to the domain X_i of deterministic conjectures is single-valued; hence, it must be a continuous function.

Fix any closed (hence, compact) sub-interval $C \subseteq X_i$. Let $ND_{i,p}(C)$ denote the set of **actions that are not strictly dominated by other pure actions**. By inspection of the definitions, it holds that

$$r_i(C) \subseteq r_i(\Delta(C)) \subseteq ND_{i,p}(C).$$

We prove that $ND_{i,p}(C) \subseteq r_i(C)$, that is, $A_i \setminus r_i(C) \subseteq A_i \setminus ND_{i,p}(C)$, which therefore implies the thesis. Since r_i is a continuous function on $X_i \supseteq C$ and C is compact and connected, $r_i(C)$ is compact and connected as well, hence, it is a compact interval. Therefore, it is enough to show that all the actions below $\min r_i(C)$ or above $\max r_i(C)$ are dominated. Fix any $a_i < \min r_i(C)$, by strict quasi-concavity,

$$\forall x_i \in C, v_i(a_i, x_i) < v_i(\min r_i(C), x_i) \leq v_i(r_i(x_i), x_i).$$

Therefore, every $a_i < \min r_i(C)$ is strictly dominated by $\min r_i(C)$. A similar argument shows that every $a_i > \max r_i(C)$ is strictly dominated by $\max r_i(C)$. Since there are no other actions outside $r_i(C)$, this concludes the proof. \square

Corollary 1. *Suppose that the aggregator ℓ_i is onto. Then, an action of player i is justifiable if and only if it is justified by a deterministic (Dirac) deep conjecture.*

Proof. The “if” part is trivial. For the “only if” part, fix a justifiable action a_i^* arbitrarily. By Lemma 1, there is some $x_i \in X_i$ such that $a_i^* = r_i(x_i)$. Since the aggregator ℓ_i is onto, there is some $(\mathbf{a}_{-i}, \mathbf{Z}) \in \ell_i^{-1}(x_i)$ such that

$$a_i^* \in \arg \max_{a_i \in A_i} u_i(a_i, \mathbf{a}_{-i}, \mathbf{Z}).$$

Hence a_i^* is justified by the deep conjecture $\delta_{(\mathbf{a}_{-i}, \mathbf{Z})}$, that is, the Dirac measure supported by $(\mathbf{a}_{-i}, \mathbf{Z})$. \square

With this, from now on we mostly restrict our attention to (shallow, or deep) *deterministic conjectures*.

A.2. Feedback properties

Definition 5. Feedback f_i satisfies **observable payoffs** (OP) relative to v_i if there is a function $\bar{v}_i : A_i \times M \rightarrow \mathbb{R}$ such that

$$v_i(a_i, x_i) = \bar{v}_i(a_i, f_i(a_i, x_i))$$

for all $(a_i, x_i) \in A_i \times X_i$; if the section \bar{v}_{i,a_i} is injective for each $a_i \in A_i$, then we say that f_i satisfies **just observable payoffs** (JOP) relative to v_i . Game G satisfies (just) observable payoffs if, for each player $i \in I$, feedback f_i satisfies (J)OP relative to v_i .

If f_i satisfies JOP, we may assume without loss of generality that $f_i = v_i$, because, for each action a_i , the partitions of X_i induced by the preimages of v_{i,a_i} and f_{i,a_i} coincide:

Remark 3. Feedback f_i satisfies JOP relative to v_i if and only if

$$\forall a_i \in A_i, \left\{ v_{i,a_i}^{-1}(u) \right\}_{u \in v_{i,a_i}(X_i)} = \left\{ f_{i,a_i}^{-1}(m) \right\}_{m \in f_{i,a_i}(X_i)}. \tag{18}$$

Proof. (Only if) Fix $a_i \in A_i$. Since f_i satisfies JOP relative to v_i , $v_{i,a_i}(X_i) = (\bar{v}_{i,a_i} \circ f_{i,a_i})(X_i)$ (by OP), for each $u \in v_{i,a_i}(X_i)$ there is a unique message $m_{a_i,u} = \bar{v}_{i,a_i}^{-1}(u)$ (by injectivity of \bar{v}_{i,a_i}), and

$$\begin{aligned} v_{i,a_i}^{-1}(u) &= \{x_i \in X_i : v_i(a_i, x_i) = u\} \\ &= \{x_i \in X_i : \bar{v}_i(a_i, f_i(a_i, x_i)) = u\} \\ &= \{x_i \in X_i : f_i(a_i, x_i) = m_{a_i,u}\} = f_{i,a_i}^{-1}(m_{a_i,u}), \end{aligned}$$

which implies eq. (18).

(If) Suppose that eq. (18) holds. For every $a_i \in A_i$ and $m \in f_{i,a_i}(X_i)$ select some $\xi_i(a_i, m) \in f_{i,a_i}^{-1}(m)$. Let

$$D := \bigcup_{a_i \in A_i} \{a_i\} \times f_{i,a_i}(X_i).$$

With this,

$$\xi_i : D \rightarrow X_i$$

is a well defined function. Domain D is the set of action-message pairs for which the definition of \bar{v}_i matters. Define \bar{v}_i as follows:

$$\bar{v}_i(a_i, m) = \begin{cases} v_i(a_i, \xi_i(a_i, m)) & \text{if } (a_i, m) \in D, \\ 0 & \text{otherwise.} \end{cases}$$

By construction, eq. (18) implies that

$$\forall (a_i, x_i) \in A_i \times X_i, \bar{v}_i(a_i, f_i(a_i, x_i)) = v_i(a_i, x_i).$$

Hence, OP holds. Furthermore, for all $a_i \in A_i, m', m'' \in f_{a_i}(X_i)$,

$$\begin{aligned} m' \neq m'' &\Rightarrow \xi_i(a_i, m') \neq \xi_i(a_i, m'') \\ &\Rightarrow v_i(a_i, \xi_i(a_i, m')) \neq v_i(a_i, \xi_i(a_i, m'')) \\ &\Rightarrow \bar{v}_i(a_i, m') \neq \bar{v}_i(a_i, m'') \end{aligned}$$

where the first and the second implications follow from eq. (18) ($\xi_i(a_i, m')$ and $\xi_i(a_i, m'')$ belong to different cells of the coincident partitions, hence yield different utilities), and the third holds by construction. Therefore, \bar{v}_{i,a_i} is injective for every a_i , which means that JOP holds. \square

Definition 6. Feedback f_i satisfies **observability if and only if i is active** (OiffA) if section f_{i,a_i} is injective for each $a_i > 0$ and constant for $a_i = 0$. Game G satisfies **observability by active players** if OiffA holds for each i .

Remark 4. If a network game is *linear-quadratic* and satisfies *just observable payoffs*, then it satisfies observability by active players.

Proof. By Remark 3 JOP implies that, for each $a_i \in A_i$,

$$\left\{ v_{i,a_i}^{-1}(u) \right\}_{u \in v_{i,a_i}(X_i)} = \left\{ f_{i,a_i}^{-1}(m) \right\}_{m \in f_{i,a_i}(X_i)}.$$

The linear-quadratic form of v_i implies that, for every $x_i \in X_i$,

$$v_{i,0}^{-1}(v_{i,0}(x_i)) = X_i, \\ \forall a_i > 0, v_{i,a_i}^{-1}(v_{i,a_i}(x_i)) = \{x_i\}.$$

These equalities imply that $f_{i,0}$ is constant and f_{i,a_i} is injective for $a_i > 0$, that is, NG satisfies observability by active players. \square

Definition 7. Function f_i satisfies **own-action independence** (OAI) of feedback about the state if, for all justifiable actions a_i^*, a_i^o and all payoff states \hat{x}_i, x_i ,

$$f_i(a_i^*, \hat{x}_i) = f_i(a_i^*, x_i) \Rightarrow f_i(a_i^o, \hat{x}_i) = f_i(a_i^o, x_i).$$

Game G satisfies own-action independence of feedback about the state if, for each player $i \in I$, feedback f_i satisfies OAI.

In other words, OAI says that if player i cannot distinguish between two payoff states \hat{x}_i and x_i when she chooses some given justifiable action a_i^* , then she cannot distinguish between these two states when he chooses any other justifiable action a_i^o . This is equivalent to requiring that the partitions of X_i of the form $\left\{ f_{i,a_i}^{-1}(m) \right\}_{m \in f_{i,a_i}(X_i)}$ coincide across justifiable actions, i.e. across actions $a_i \in r_i(X_i)$ (see Lemma 1).

The following lemma – which holds for any game, not just nice games – states that, under payoff observability and own-action independence, an action is justified by a confirmed conjecture if and only if it is a best reply to the actual payoff state:

Lemma 2. If f_i satisfies observable payoffs relative to v_i and own-action independence of feedback about the state, then for all $(a_i^*, x_i) \in A_i \times X_i$ the following are equivalent:

1. there is some $\hat{x}_i \in X_i$ such that $a_i^* \in \arg \max_{a_i \in A_i} v_i(a_i, \hat{x}_i)$ and $f_i(a_i^*, \hat{x}_i) = f_i(a_i^*, x_i)$,
2. $a_i^* \in \arg \max_{a_i \in A_i} v_i(a_i, x_i)$.

Proof. (Cf. Battigalli et al., 2015) It is obvious that 2 implies 1 independently of the properties of f_i . To prove that 1 implies 2 under the stated assumptions, suppose that f_i satisfies OP-OAI and let \hat{x}_i be such that 1 holds. Let a_i^o be a best reply to the actual state x_i . We must show that also a_i^* is a best reply to x_i . Note that both a_i^* and a_i^o are justifiable; hence, by OAI, $f_i(a_i^*, \hat{x}_i) = f_i(a_i^*, x_i)$ implies $f_i(a_i^o, \hat{x}_i) = f_i(a_i^o, x_i)$. Using OP, condition 1, and OAI as shown in the following chain of equalities and inequalities, we obtain

$$v_i(a_i^*, x_i) \stackrel{\text{(OP)}}{=} \bar{v}_i(a_i^*, f_i(a_i^*, x_i)) \stackrel{\text{(1)}}{=} \bar{v}_i(a_i^*, f_i(a_i^*, \hat{x}_i)) \stackrel{\text{(OP)}}{=} v_i(a_i^*, \hat{x}_i) \stackrel{\text{(1)}}{\geq} \\ v_i(a_i^o, \hat{x}_i) \stackrel{\text{(OP)}}{=} \bar{v}_i(a_i^o, f_i(a_i^o, \hat{x}_i)) \stackrel{\text{(1,OAI)}}{=} \bar{v}_i(a_i^o, f_i(a_i^o, x_i)) \stackrel{\text{(OP)}}{=} v_i(a_i^o, x_i).$$

Since a^o is a best reply to x_i and $v_i(a_i^*, x_i) \geq v_i(a_i^o, x_i)$, it must be the case that also a_i^* is a best reply to x_i . \square

In the main text we defined SCE for the special case in which the feedback and utility function of each player coincide. More generally, a profile $(a_i^*, \hat{x}_i)_{i \in I}$ of actions and shallow conjectures is a **selfconfirming equilibrium at $\mathbf{Z} \in \mathcal{Z}$** of the parameterized nice game with aggregators and feedback G if $a_i^* = r_i(\hat{x}_i)$ (best reply) and $\hat{x}_i = f_i(a_i^*, \ell_i(\mathbf{a}_{-i}^*, \mathbf{Z}))$ (confirmed conjecture) for every $i \in I$. As in the main text (but neglecting global externalities), $\mathbf{A}_{\mathbf{Z}}^{\text{SCE}}$ and $\mathbf{A}_{\mathbf{Z}}^{\text{NE}}$ respectively denote the set of SCE and NE action profile of the game with parameters \mathbf{Z} .

Corollary 2. *Suppose that the parameterized nice game with aggregators and feedback G satisfies observable payoffs and own-action independence of feedback about the state. Then, for each $\mathbf{Z} \in \mathcal{Z}$, the sets of selfconfirming action profiles and Nash equilibrium action profiles coincide.*

Proof. By Remark 1, we only have to show that $\mathbf{A}^{SCE} \subseteq \mathbf{A}^{NE}$. Fix any $\mathbf{a}^* = (a_i^*)_{i \in I} \in \mathbf{A}^{SCE}$ and any player i . By definition of SCE and by Lemma 1, there is some $\hat{x}_i \in X_i$ such that $a_i^* \in r_i(\hat{x}_i)$ and $f_i(a_i^*, \hat{x}_i) = f_i(a_i^*, \ell_i(\mathbf{a}_{-i}^*))$. By Lemma 2, $a_i^* \in r_i(\ell_i(\mathbf{a}_{-i}^*))$. This holds for each i , hence $\mathbf{a}^* \in \mathbf{A}^{NE}$. \square

Corollary 2 provides sufficient conditions for the equivalence between SCE and NE action profiles. Next, we give sufficient conditions that allow a characterization of \mathbf{A}^{SCE} by means of Nash equilibria of auxiliary games.

A.3. Equilibrium characterization

If $a_i \in [0, \bar{a}_i]$ is interpreted as an activity level (e.g., effort) by player i , then it makes sense to say that i is **active** if $a_i > 0$ and **inactive** otherwise. Let I_0 denote the **set of players for whom being inactive is justifiable**. Note that, by Lemma 1,

$$I_0 = \{i \in I : \min r_i(X_i) = 0\}.$$

Also, for each \mathbf{Z} and non-empty subset of players $J \subseteq I$, let $\mathbf{A}_{\mathbf{Z}}^{NE,J}$ denote the set of Nash equilibria of the auxiliary game with players set J obtained by letting $a_i = 0$ for each $i \in I \setminus J$, that is,

$$\mathbf{A}_{\mathbf{Z}}^{NE,J} := \left\{ \mathbf{a}_J^* \in \times_{j \in J} A_j : \forall j \in J, a_j^* = r_j \left(\ell_j \left(\mathbf{a}_{J \setminus \{j\}}^*, \mathbf{0}_{I \setminus J}, \mathbf{Z} \right) \right) \right\},$$

where $\mathbf{0}_{I \setminus J} \in \mathbb{R}^{I \setminus J}$ is the profile that assigns 0 to each $i \in I \setminus J$. If $J = \emptyset$, let $\mathbf{A}_{\mathbf{Z},J}^{NE} = \{\emptyset\}$ by convention, where \emptyset is the pseudo-action profile such that $(\emptyset, \mathbf{0}_I) = \mathbf{0}_I$.

Since here we focus on games without global externalities, we ease notation and let $\mathbf{A}_{\mathbf{Z}}^{SCE}$ (instead of $\mathbf{A}_{\mathbf{Z},0}^{SCE}$) denote the set of selfconfirming action profiles given \mathbf{Z} .

Lemma 3. *Suppose that the parameterized nice game with aggregators and feedback G satisfies observability by active players. Then, the set of selfconfirming action profiles is*

$$\mathbf{A}_{\mathbf{Z}}^{SCE} = \bigcup_{J: I \setminus J \subseteq I_0} \mathbf{A}_{\mathbf{Z}}^{NE,J} \times \{\mathbf{0}_{I \setminus J}\}.$$

Proof. Fix \mathbf{a}^* and let J be the set of players i such that $\bar{a}_i^* > 0$. Fix \mathbf{Z} arbitrarily. Suppose that $\mathbf{a}^* \in \mathbf{A}_{\mathbf{Z}}^{SCE}$ and fix any $i \in I$. If $a_i^* = 0$, then 0 is justifiable for i , that is $i \in I_0$. If $a_i^* > 0$, observability by active players implies that f_{i,a_i^*} is injective, that is, action a_i^* reveals the payoff state, which implies that the (shallow) conjecture justifying a_i^* is correct: $a_i^* = r_i(\ell_i(\mathbf{a}_{-i}^*))$. Hence, $\mathbf{a}_J^* \in \mathbf{A}_{\mathbf{Z}}^{NE,J}$. Thus, $\mathbf{a}^* = (\mathbf{a}_J^*, \mathbf{a}_{I \setminus J}^*)$ is such that $a_i^* = 0$ for each $i \in I \setminus J \subseteq I_0$, and $a_j^* = r_j(\ell_j(\mathbf{a}_{J \setminus \{j\}}^*, \mathbf{0}_{I \setminus J})) > 0$ for each $j \in J$. Hence,

$$\mathbf{a}^* = (\mathbf{a}_J^*, \mathbf{a}_{I \setminus J}^*) \in \mathbf{A}_{\mathbf{Z}}^{NE,J} \times \{\mathbf{0}_{I \setminus J}\} \text{ with } I \setminus J \subseteq I_0.$$

Let $I \setminus J \subseteq I_0$ and $(\mathbf{a}_J^*, \mathbf{a}_{I \setminus J}^*) \in \mathbf{A}_Z^{NE,J} \times \{\mathbf{0}_{I \setminus J}\}$. Since G satisfies observability by active players, for each $i \in I \setminus J$, any conjecture justifying $a_i^* = 0$ (any $\hat{x}_i \in r_i^{-1}(0)$) is trivially confirmed. For each $j \in J$, $a_j^* > 0$ is by assumption the best reply to the correct, hence confirmed, shallow conjecture $\hat{x}_j = \ell_j(\mathbf{a}_{J \setminus \{j\}}^*, \mathbf{0}_{I \setminus J})$. Hence, $(\mathbf{a}_J^*, \mathbf{a}_{I \setminus J}^*) = (\mathbf{a}_J^*, \mathbf{0}_{I \setminus J}) \in \mathbf{A}_Z^{SCE}$. \square

Appendix B. Interior Nash equilibria

Proposition 1 shows that, given our maintained assumptions about the network game with feedback, selfconfirming action profiles can be characterized as Nash equilibria of auxiliary games with a restricted set of players, which must include all those for whom being inactive is unjustifiable (dominated), but may leave out any player for whom inactivity is justifiable (undominated). We now provide some results about existence of these SCE's that will be useful in proving Proposition 2. We first present sufficient conditions that are present in the literature for the existence and uniqueness of interior Nash equilibria, then we provide some original results.

In this appendix we formulate the problem with the approach of linear algebra. We consider a square matrix $\mathbf{Z} \in \mathbb{R}^{n \times n}$ such that $z_{ii} = 0$ for all $i \in \{1, \dots, n\}$. We denote by \mathbf{I} the identity matrix, $\lambda_{max}(\mathbf{Z})$ the maximal eigenvalue of \mathbf{Z} , $\rho(\mathbf{Z})$ the spectral radius of \mathbf{Z} (i.e., the largest absolute value of its eigenvalues), $\mathbf{1}$ the vector of all 1's, $\mathbf{0}$ the vector of all 0's, and \gg the strict partial ordering between vectors (meaning that all the entries in the first vector are coordinatewise strictly greater than the entries in the second vector). With this notation, the condition for the existence of a unique Nash equilibrium which is also interior is $(\mathbf{I} - \mathbf{Z})^{-1} \cdot \mathbf{1} \gg \mathbf{0}$.

Proposition 10. Consider a square matrix $\mathbf{Z} \in \mathbb{R}^{n \times n}$ such that (i) $\rho(\mathbf{Z}) < 1$, (ii) for each $i \in I$, $z_{ii} = 0$, and (iii) for each $j \neq i$, $z_{ij} \leq 0$. Then $(\mathbf{I} - \mathbf{Z})^{-1} \cdot \mathbf{1} \gg \mathbf{0}$.²⁶

Some results can be provided also when the sign of the externalities are mixed. Recall that matrix \mathbf{Z} is symmetrizable if there exists a diagonal matrix \mathbf{W} and a symmetric matrix \mathbf{Z}_0 such that $\mathbf{Z} = \mathbf{W}\mathbf{Z}_0$. Note that, if \mathbf{Z} is symmetrizable, then all its eigenvalues are real. If for all i , $z_{ii} = 0$, and \mathbf{Z} is symmetrizable, we define the symmetric matrix $\tilde{\mathbf{Z}}$ to be such that $\tilde{z}_{ij} = z_{ij} \sqrt{w_i w_j}$.

Proposition 11. Consider a square matrix $\mathbf{Z} \in \mathbb{R}^{n \times n}$ such that (i) for each $i \in I$, $z_{ii} = 0$, (ii) \mathbf{Z} is symmetrizable, and (iii) $|\lambda_{max}(\tilde{\mathbf{Z}})| < 1$. Then $(\mathbf{I} - \mathbf{Z})^{-1} \cdot \mathbf{1} \gg \mathbf{0}$.²⁷

Finally, we provide below a novel alternative condition.

Proposition 12. Consider a square matrix $\mathbf{Z} \in \mathbb{R}^{n \times n}$ such that (i) for each $i \in I$, $z_{ii} = 0$ and (ii) for each $i \neq j$, $|z_{ij}| < \frac{1}{n}$. Then $(\mathbf{I} - \mathbf{Z})^{-1} \cdot \mathbf{1} \gg \mathbf{0}$.

Proof. Let $\mathbf{B} := (\mathbf{I} - \mathbf{Z})$. First of all, by Gershgorin circle theorem, \mathbf{B} has all eigenvalues, possibly complex, with real part strictly between 0 and 2, so $\det(\mathbf{B}) \neq 0$.

²⁶ This is Theorem 1 in Ballester et al. (2006). The same result is in Appendix A in Stańczak et al. (2006).

²⁷ See Section VI of Bramoullé et al. (2014), generalizing Proposition 2 therein. Note that in their payoff specification externalities have a minus sign, while in (3) we have a plus sign: this is why we have a condition on the maximal eigenvalue and not on the minimal eigenvalue.

Consider the n vectors $\mathbf{b}^1, \dots, \mathbf{b}^n$ given by the n rows of \mathbf{B} , and take the hyperplane in \mathbb{R}^n passing by those n points:

$$H := \{\mathbf{h} \in \mathbb{R}^n : \exists \alpha \in \mathbb{R}^n, \alpha' \cdot \mathbf{1} = 1 \wedge \mathbf{h} = \mathbf{B}'\alpha\}.$$

Now, consider the following vector

$$\mathbf{v} := \mathbf{B}^{-1}\mathbf{1}.$$

Note that each v_i is exactly the sum of the entries in i^{th} row of \mathbf{B}^{-1} . However, \mathbf{v} is also a vector perpendicular to H . This is because for each $\mathbf{h} \in H$, there exists $\alpha \in \mathbb{R}^n$ such that

$$\begin{aligned} \mathbf{h} \cdot \mathbf{v} &= (\mathbf{B}'\alpha)' \cdot \mathbf{B}^{-1}\mathbf{1} \\ &= \alpha' \mathbf{1} \\ &= \sum_{i=1}^n \alpha_i = 1, \end{aligned}$$

which is a constant.

Now, we want to show that H does not pass through the convex region of vectors with all negative elements: $H \cap (-\infty, 0]^n = \emptyset$. In fact, it is impossible to find $\mathbf{w} \in \mathbb{R}^n$, such that $\mathbf{w}' \cdot \mathbf{1} = 1$ and $\mathbf{B}'\mathbf{w} \ll \mathbf{0}$. Suppose, by way of contradiction, that such vector \mathbf{w} exists. Let $k := \arg \max_{i \in \{1, \dots, n\}} \{w_i\}$ (note that $w_k > 0$ because $\sum_{i=1}^n w_i = 1$), then, calling \mathbf{b}_k the k^{th} row of matrix \mathbf{B} , we get

$$\mathbf{b}_k \cdot \mathbf{w} = w_k + \sum_{j \neq k} w_j b_{jk} > w_k - \sum_{j \neq k} |w_j| |z_{jk}| > w_k \left(1 - \sum_{j \neq k} |z_{jk}| \right) > 0,$$

which is a contradiction.

Finally, we show that if a hyperplane H satisfies $H \cap (-\infty, 0]^n = \emptyset$, then its perpendicular vector from the origin has all strictly positive entries, and this concludes the proof. We do so by induction on n .

1. **$n = 2$:** This is easy to show graphically. In the Cartesian plane the hyperplane is a line. Since this line does not intersect the negative hortalant $(-\infty, 0]^2$, it must cross both axes in their strictly positive part: call these intersection points A and B . So, the segment that from the origin crosses this line perpendicularly will cross it in a point C that lies on the line between A and B .
2. **Induction hypothesis:** Suppose it is true for $n - 1$.
3. **Inductive step:** a hyperplane $H \subset \mathbb{R}^n$ that satisfies $H \cap (-\infty, 0]^n = \emptyset$ does not pass through the origin. So, it has an orthogonal vector \mathbf{v} such that $\mathbf{v} \in H$. By assumption on H , \mathbf{v} cannot have all elements non-strictly positive. So, there exists $i \in \{1, \dots, n\}$ such that $v_i > 0$. Let us take $P_{-i} = \{\mathbf{p} \in \mathbb{R}^n : p_i = 0\}$. Call H_{-i} the intersection of H with P_{-i} . Take the vector \mathbf{v}_{-i} that is the projection of \mathbf{v} on P_{-i} . This vector has all entries equal to \mathbf{v} , except for entry i which is null. Also, \mathbf{v}_{-i} is perpendicular to H_{-i} .
By assumption on H , $H_{-i} \cap (-\infty, 0]^{n-1} = \emptyset$. Moreover, by the induction hypothesis, \mathbf{v}_{-i} has all strictly positive entries, except from entry i . Finally, since also $v_i > 0$, the result follows.

Notice that, if \mathbf{Z} satisfies the conditions of Proposition 12, then it must also hold that $|\lambda_{max}(\mathbf{Z})| < 1$, because of Gershgorin circle theorem. However, the condition that $|\lambda_{max}(\mathbf{Z})| < 1$ is in general not sufficient to guarantee that $(\mathbf{I} - \mathbf{Z})^{-1} \mathbf{1} \gg \mathbf{0}$. \square

Appendix C. Proofs of propositions

Proposition 1

Proof. By Remark 4, NG satisfies observability by active players. Hence, Lemma 3 in Appendix A and the best-reply equation yield the result. \square

Proposition 2

Proof. Conditions 1, 2, and 3 correspond, respectively, to the conditions in Propositions 12, 10, and 11 from Appendix B. \square

Proposition 3

Proof. Let us consider separately the two sets $I \setminus I_{\mathbf{a}^*}$ and $I_{\mathbf{a}^*}$ of inactive and active agents.

For every $i \in I \setminus I_{\mathbf{a}^*}$, $\alpha_i + \underline{x}_i < 0$; thus, $a_i^* = 0$ is a best reply to every conjecture $\hat{x}_i \in (\underline{x}_i, -\alpha_i)$ and a sufficiently small perturbation of \hat{x}_i does not make i become active.

Now, let us focus on the subset $I_{\mathbf{a}^*}$ of active agents. For each $i \in I_{\mathbf{a}^*}$, a perturbation in \hat{x}_i induces a change in the corresponding best reply. Let us focus on perturbations that are small enough so that all actions of agents in $I_{\mathbf{a}^*}$ remain strictly positive. Since $\rho(\mathbf{Z}) < 1$ is a strict inequality, Assumption 4 guarantees that the limiting points of the discrete path system defined for actions by (7) and (8) are locally stable, because the non-null eigenvalues and eigenvectors of the Jacobian of this system are the same eigenvalues and eigenvectors of $\mathbf{Z}_{I_{\mathbf{a}^*}}$.

Thus, there is $\epsilon > 0$ such that the perturbation of beliefs given by any \mathbf{x}_0 with $\|\mathbf{x}_0 - \hat{\mathbf{x}}\| < \epsilon$ is small enough so that inactive agents keep being inactive and all actions of active agents in $I_{\mathbf{a}^*}$ remain strictly positive.

In this way, the discrete system defined for actions by (7) and (8) converges back to \mathbf{a}^* . \square

Proposition 4

Proof. For all the action profiles considered in the proposition the inactive players are choosing a best response for an open set of conjectures; thus, being inactive is robust to small perturbations of justifying non-falsified conjectures. With this, we can focus on the active agents. Note that if we take an active agent i from $I_{\mathbf{a}^*}$ and we make him inactive, then the new matrix $\mathbf{Z}_{I_{\mathbf{a}^*} \setminus \{i\}}$ for active players is a sub-matrix of $\mathbf{Z}_{I_{\mathbf{a}^*}}$ obtained deleting the row and the column corresponding to agent i . This process can be repeated removing more active agents, which means that if we remove a subset $J \subset I_{\mathbf{a}^*}$ of the active agents, then the new matrix $\mathbf{Z}_{I_{\mathbf{a}^*} \setminus J}$ is a sub-matrix of $\mathbf{Z}_{I_{\mathbf{a}^*}}$ obtained deleting all the rows and the columns corresponding to every agent $j \in J$.

So, given the results from Propositions 2 and 3, to prove the statement, we need to prove that if an adjacency matrix satisfies one of the three conditions, then also every sub-matrix of that matrix, which is obtained deleting one row and one column with the same index, satisfies that condition. By induction this will be true for every sub-matrix of that matrix, which is obtained deleting any subset of rows and columns with the same indices.

For Point 1 the result is clear, because a property that holds for all the elements of a matrix will hold also for all the elements of a sub-matrix of that matrix.

Point 2 is based on two assumptions. Assumption 3 is still valid if we remove one column and one row of a matrix because it is a property of all the elements of that matrix. To check for Assumption 4, let us consider the following implications of the Perron–Frobenius theorem (see, e.g., Savchenko, 2003): (i) for a matrix with all positive entries, there exists a real eigenvalue (often called the **Perron root**) which is equal to its spectral radius; (ii) the Perron root of any principal submatrix of such a matrix does not exceed that of the original matrix. In our case, Assumption 3 implies that our matrix can be seen as a matrix with all positive elements with a minus sign in front, and this proves the statement.

Point 3 holds because of a generalization of the Cauchy interlace theorem applied to symmetrizable matrices (see Kouachi, 2016 and McKee and Smyth, 2020). We know that the magnitude of the eigenvalues of the sub-matrix of a symmetrizable matrix, obtained deleting one row and one column with the same index, are between the magnitudes of the minimal and the maximal eigenvalues of the old matrix. So, the sub-matrix of a limited matrix, which is obtained deleting one row and one column with the same index, is limited. The resulting sub-matrix is also symmetrizable. That is because the original matrix was obtained as the product of a diagonal and a symmetric matrix, and to obtain the sub-matrix we can delete the corresponding rows and columns in those diagonal and symmetric matrices: the two matrices will maintain their properties and the result will be our sub-matrix. □

Proposition 5

Proof. Fix a an SCE $(a_i^*, \hat{x}_i, \hat{y}_i)_{i \in I}$. For each i in I , rationality implies

$$a_i^* = \min\{\max\{0, \alpha_i + \hat{x}_i\}, \bar{a}_i\}.$$

Agent i then thinks that

$$v_i^* = \alpha_i a_i^* - \frac{1}{2} (a_i^*)^2 + a_i^* \hat{x}_i + \hat{y}_i,$$

where v_i^* denotes the realized and observed payoff, so that

$$\hat{y}_i = v_i^* - \alpha_i a_i^* + \frac{1}{2} (a_i^*)^2 - a_i^* \hat{x}_i. \tag{19}$$

Substituting the expression of the true actual payoff

$$v_i^* = \alpha_i a_i^* - \frac{1}{2} (a_i^*)^2 + a_i^* x_i + y_i$$

into (19), we get the dependence between \hat{y}_i and \hat{x}_i :

$$\hat{y}_i = y_i + a_i^* (x_i - \hat{x}_i).$$

The first and second items in the proposition are derived, respectively, if $a_i^* = 0$ or if $a_i^* > 0$. □

Proposition 6

Proof. By substituting, for each $i \in I$, the subjectively rational choice into the confirmed conjecture condition, we get the following:

$$(\alpha + \hat{x}_i) \left(\hat{x}_i - \sum_{j \in I \setminus \{i\}} z_{ij} (\alpha + \hat{x}_j) \right) = \left(\gamma \sum_{k \in I \setminus \{i\}} (\alpha + \hat{x}_k) - \hat{y}_i \right). \tag{20}$$

This condition holds for each $i \in I$, so that we have a non-linear system of n equations and $2n$ unknowns. Still, from (20) we can provide useful insights to understand how conjectures are shaped in an SCE.

First of all, note that (20) is linear in \hat{y}_i . Thus, given any profile $(\hat{x}_i)_{i \in I}$, there exists a unique profile $(\hat{y}_i)_{i \in I}$ consistent with the confirmed conjectures condition. Moreover, we can also compute a bound for each \hat{y}_i . Indeed, for each $i \in I$, $\hat{x}_i > 0$. Then, since $a_i = \alpha + \hat{x}_i \leq \bar{a}$, for each $i \in I$, and given other agents' conjectures, it must be that $y_i \leq \alpha \sum_{j \in I \setminus \{i\}} z_{ij} a_j + \gamma \sum_{k \in I \setminus \{i\}} a_k \leq$

$$\bar{a} \left(\alpha \sum_{k \in I \setminus \{i\}} z_{ij} + \gamma n \right).$$

Given a profile $(\hat{y}_i)_{i \in I}$, condition (20), also allows us to characterize the corresponding SCE profile $(\hat{x}_i)_{i \in I}$. Solving the second-order polynomial, we get that the only positive solution for each \hat{x}_i is given by

$$\hat{x}_i = \frac{1}{2} \left(\sum_{j \in I \setminus \{i\}} z_{ij} (\alpha + \hat{x}_j) - \alpha + \sqrt{\left(\sum_{j \in I \setminus \{i\}} z_{ij} (\alpha + \hat{x}_j) + \alpha \right)^2 + 4\gamma \sum_{k \in I \setminus \{i\}} (\alpha + \hat{x}_k) - 4\hat{y}_i} \right). \tag{21}$$

Note that, at an SCE, each \hat{x}_i is increasing in others' conjectures about local externality, and decreasing in own \hat{y}_i . Indeed, given \hat{y}_i , an increase in any \hat{x}_j increases j 's action and thus it increases the global externality. Given that only positive externalities are considered, if \hat{y}_i is kept fixed, at SCE i has no other option than having a higher \hat{x}_i . On the contrary, if \hat{y}_i increases keeping fixed $(\hat{x}_j)_{j \in I \setminus \{i\}}$, then actual local and global externalities for i are unchanged. However, if i thinks y_i to be higher, she necessarily needs to decrease \hat{x}_i . Given that equilibrium \hat{x}_i is monotonically decreasing in \hat{y}_i , we can also easily compute an upper bound for \hat{x}_i by simply letting $\hat{y}_i = 0$ in (21).

By taking the second derivative of the right hand side of (21), with respect to \hat{x}_j , we obtain

$$\frac{\partial^2 \hat{x}_i}{\partial \hat{x}_j^2} = -\frac{2\gamma}{\sqrt{\Gamma(\hat{x}_j)^{3/2}}} \left(z_{ij} \left(\sum_{k \in I \setminus \{i, j\}} z_{ik} (\alpha + \hat{x}_k) - z_{ij} \sum_{h \in I \setminus \{i, j\}} (\alpha + \hat{x}_h) + \alpha \right) + \gamma \right),$$

where $\Gamma(\hat{x}_j)$ is an always positive quadratic expression of \hat{x}_j . If, for every couple of agents i and j in I , the inequality

$$\sum_{k \in I \setminus \{i, j\}} z_{ik} (\alpha + \hat{x}_k) - z_{ij} \sum_{h \in I \setminus \{i, j\}} (\alpha + \hat{x}_h) + \alpha \geq 0, \tag{22}$$

is satisfied, then \hat{x}_i is concave in each \hat{x}_j . So, there is always a unique finite solution to the system where each player has the higher possible belief about \hat{x}_j . In this solution, as we assume that either condition 1. or 3. of Proposition 2 is satisfied, we derive a unique $(a_i^*)_{i \in I}$ with $a_i^* < \bar{a}$ for each i . If, \hat{x}_i is convex in some \hat{x}_j , then the process may self-reinforce and it is possible that a corner solution is reached. \square

Proposition 7

Proof. Before proving the result we need to consider a slight modification of aggregator and conjectures.

Let

$$\begin{aligned} \tilde{\ell}_i: \mathbf{A}_{-i} &\rightarrow \tilde{X}_i, \\ \mathbf{a}_{-i} &\mapsto \sum_{j \neq i} z_{0,ij} a_j \end{aligned} \tag{23}$$

and

$$\begin{aligned} \tilde{g}_i: \mathbf{A}_{-i} &\rightarrow \tilde{Y}_i \\ \mathbf{a}_{-i} &\mapsto \sum_{j \neq i} a_j \end{aligned} \tag{24}$$

be the equivalent of ℓ_i and g_i , when we do not incorporate the parameters on which there is mutual knowledge. Similarly, let \hat{x}_i and \hat{y}_i be the shallow conjectures about \tilde{x}_i and \tilde{y} , respectively. Then, we need to provide a definition of selfconfirming equilibrium consistent with the hypotheses about the knowledge of the agents. \square

Definition 8. A profile $(a_i^*, \hat{x}_i, \hat{y}_i)_{i \in I} \in \times_{i \in I} (A_i \times \tilde{X}_i \times \tilde{Y}_i)$ of actions and (shallow) deterministic conjectures is a **selfconfirming equilibrium** at $(\mathbf{Z}_0, \omega, \gamma)$ of a network game with global externalities with mutual knowledge of (ω, γ) if, for each $i \in I$,

1. (subjective rationality) $a_i^* = r_i(\hat{x}_i)$;
2. (confirmed conjecture) $f_i(a_i^*, \hat{x}_i, \hat{y}_i; \omega, \gamma) = f_i(a_i^*, \tilde{\ell}_i(\mathbf{a}_{-i}^*, \mathbf{Z}_0), \tilde{g}_i(\mathbf{a}_{-i}^*); w, \gamma)$.

We are now ready to prove the result.

Consider first the Nash equilibrium of the game with payoff function (10) played on a complete network. For each $i \in I$, $a_{\mathbf{Z}_c, i}^{NE} = r_i(w \sum_{k \in I \setminus \{i\}} a_{\mathbf{Z}_c, k}^{NE})$. Because of symmetry, for each $i \in I$, $a_{\mathbf{Z}_c, i}^{NE} = \frac{\alpha_i}{1 - (n-1)w}$.

Given a selfconfirming equilibrium action profile \mathbf{a}^c , each player i , by perfect recall of her own action, can correctly infer that

$$a_i^c w \tilde{x}_i + \gamma \tilde{y}_i = a_i w \sum_{j \in I \setminus \{i\}} z_{0,ij} a_j + \gamma \sum_{k \in I \setminus \{i\}} a_k, \tag{25}$$

so that, her shallow conjectures must be such that

$$a_i^c w \hat{x}_i + \gamma \hat{y}_i = a_i w \sum_{j \in I \setminus \{i\}} z_{0,ij} a_j + \gamma \sum_{k \in I \setminus \{i\}} a_k. \tag{26}$$

At the same time, by deep conjecture $\bar{\mu}_i^c$ each player i thinks to be linked with all the other players. Then $\hat{x}_i = \hat{y} = \hat{x}_i^c$, and her shallow conjectures are such that

$$a_i^c w \hat{x}_i + \gamma \hat{y}_i = (a_i w + \gamma) \hat{x}_i^c. \tag{27}$$

So, by (26)-(27) we have

$$\hat{x}_i^c = \frac{a_i w \sum_{j \in I \setminus \{i\}} z_{0,ij} a_j + \gamma \sum_{k \in I \setminus \{i\}} a_k}{a_i w + \gamma}.$$

As externalities are positive and $a_i > 0$, γ and $a_i w$ are just weights in a weighted average. If $\frac{\gamma}{w} = 0$, then $\hat{x}_i^c = \sum_{j \in I \setminus \{i\}} z_{0,ij} a_j^c$, i.e., conjecture \hat{x}_i^c is correct, so that $\mathbf{a}^c = \mathbf{a}_{\mathbf{Z}_0}^{NE}$. Finally, $\lim_{\frac{\gamma}{w} \rightarrow \infty} \hat{x}_i^c = \sum_{k \in I \setminus \{i\}} a_k^c$ so that, since upper bounds are not binding, at this limit we have $\mathbf{a}^c = \mathbf{a}_{\mathbf{Z}_c}^{NE}$.

Proposition 8

Proof. First, we derive some properties. Recall that we assumed a common bliss point in isolation: $\alpha_i = \alpha$ for each $i \in I$, and that c_i is the conjectured ratio of i . Each equation in the system given by (16) can be written as an upward parabola $b_1 a_i^2 + b_2 a_i + b_3 = 0$, in the following way:

$$H_i(\mathbf{a}, \mathbf{c}, \gamma, \mathbf{Z}) = \underbrace{c_i}_{:=b_1} a_i^2 + \underbrace{\left(1 - \alpha c_i - c_i \left(\sum_{j \in I} z_{ij} a_{j,t}\right)\right)}_{:=b_2} a_i - \underbrace{\left(1 + c_i \left(\gamma \sum_{j \neq i} a_{j,t}\right)\right)}_{:=b_3} = 0. \tag{28}$$

So, for each $i \in I$, the solution a_i^* is such that $H_i(\mathbf{a}, \mathbf{c}, \gamma, \mathbf{Z}) = 0$ lays in the right-arm of this upward parabola, where $\frac{dH_i}{da_i} \Big|_{a_i=a_i^*} > 0$. Each $H_i(\mathbf{a}, \mathbf{c}, \gamma, \mathbf{Z})$ is linear in c_i .

Equation (28) holds at the unique positive solution (because $b_3 > 0$):

$$a_i^* = \frac{-b_2 + \sqrt{b_2^2 + 4b_1 b_3}}{2b_1}, \tag{29}$$

so that a_i^* can be seen as a continuous function of b_1 , b_2 and b_3 . Considering that a_i^* is increasing in b_1 (which is bounded by 1), decreasing in b_2 and increasing in b_3 , it follows that each a_i^* increases in each a_j , with $j \neq i$. Moreover, each a_i^* increases in c_i , so that

$$\frac{da_i}{dc_i} \Big|_{a_i=a_i^*} > 0.$$

If b_2 is bounded (from below), then a_i^* is bounded above by

$$\lim_{b_1 \rightarrow 1} \frac{-b_2 + \sqrt{b_2^2 + 4b_1 b_3}}{2b_1} = \frac{-b_2 + \sqrt{b_2^2 + 4b_3}}{2},$$

which is in turn bounded above by $\sqrt{b_3}$ (because if a and b are positive, $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$).

Second, we show that there is a homeomorphism. There is a continuous function that assigns to each $\mathbf{c} \in [0, 1]^n$ an element $\mathbf{a}^* \in \mathcal{A}$, because

- either $c_i = 0$ and then $a_i^* = \alpha$, since (from (29)):

$$\lim_{c_i \rightarrow 0} a_i^* = \alpha;$$

- or $c_i > 0$ and then each a_i^* is continuously increasing in each a_j with $j \neq i$. Also, b_2 is bounded (from below), because the system defined by (17) admits a solution, and then also any linear transformation of this system will admit a finite solution, which means that b_2 is limited.

Since b_2 is bounded (from below), then a_i^* is bounded above by

$$\sqrt{1 + c_i \left(\gamma \sum_{j \neq i} a_{j,t} \right)}.$$

But this upper limit is sub-linear, and then the system defined by (16) admits a finite solution.

So, applying system (16), for each $\mathbf{c} \in [0, 1]^n$, we obtain a unique profile $\mathbf{a}^* \in \mathcal{A}$, and this function is continuous because (29) is continuous.

To analyze the relation between \mathbf{a}^* and \mathbf{c} , we already know that each a_i^* is increasing in c_i and in all the other a_j^* , with $j \neq i$, which in turn are increasing in c_j . This shows that a_i^* is strictly monotone with respect to the lattice order of the domain of all profiles $\mathbf{c} \in [0, 1]^n$.

Strict monotonicity and continuity imply that the function from $\mathbf{a} \in \mathcal{A}$ to $\mathbf{c} \in [0, 1]^n$ is invertible. \square

Proposition 9

Proof. To obtain the rest points of the paths defined by (15), we consider the system derived from (16) for each i :

$$H_i(\mathbf{a}, \mathbf{c}, \gamma, \mathbf{Z}) = \alpha + c_i \left(\gamma \sum_{j \neq i} a_{j,t} \right) \frac{a_{i,t} c'_{i,t} + 1}{a_i c_i + 1} - a_i = 0,$$

with $c'_{i,t} = \frac{\sum_{j \in I} z_{ij} a_{j,t}}{\gamma \sum_{j \neq i} a_{j,t}}$. We can compute its Jacobian, with respect to \mathbf{a} . We know from the proof of Proposition 8 that each entry of this Jacobian is strictly positive. If we prove that each row of this Jacobian sums to less than 1, by the Gershgorin circle theorem we will have that the Jacobian is limited (as defined in Assumption 4), so that the process is always a contraction and the rest points are stable (see, e.g., Galor, 2007). The Jacobian J is such that, for each $i, j \in I$:

$$\begin{cases} J_{ij} = \frac{c_i}{a_i c_i + 1} (\gamma + a_i z_{ij}) & , \text{ for } j \neq i \\ J_{ii} = c_i \left(\gamma \sum_{j \neq i} a_j \right) \left(\frac{c'_i}{a_i c_i + 1} - c_i \frac{a_i c'_i + 1}{(a_i c_i + 1)^2} \right) - 1 & , \text{ otherwise.} \end{cases}$$

The sum of each row of the Jacobian is

$$\sum_{j \in I} J_{ij} = \frac{c_i}{a_i c_i + 1} \left(\gamma \left(\sum_{j \neq i} a_j \right) \left(c'_i - c_i \frac{a_i c'_i + 1}{a_i c_i + 1} \right) + a_i \left(\sum_{j \neq i} z_{i,j} \right) + \gamma(n - 1) \right) - 1. \tag{30}$$

Let us analyze expression (30) with respect to a_i , for any $a_i \geq 0$.

First note that

$$\lim_{a_i \rightarrow \infty} \sum_{j \in I} J_{ij} = \sum_{j \neq i} z_{ij} - 1, \tag{31}$$

whose absolute value is less than one by assumption.

Moreover,

$$\lim_{a_i \rightarrow 0} \sum_{j \in I} J_{ij} = c_i \gamma \left(\left(\sum_{j \neq i} a_j \right) (c'_i - c_i) + (n - 1) \right) - 1. \tag{32}$$

An interior maximum or minimum of the numerical expression (30), with respect to a_i , must satisfy first order condition

$$\begin{aligned} & - \left(\frac{c_i}{a_i c_i + 1} \right)^2 \left(\gamma \left(\sum_{j \neq i} a_j \right) \left(c'_i - c_i \frac{a_i c'_i + 1}{a_i c_i + 1} \right) + a_i \left(\sum_{j \neq i} z_{ij} \right) + \gamma(n - 1) \right) \\ & + \frac{c_i}{a_i c_i + 1} \left(\gamma \left(\sum_{j \neq i} a_j \right) \left(\frac{c_i}{a_i c_i + 1} \right) \left(c'_i - c_i \frac{a_i c'_i + 1}{a_i c_i + 1} \right) + \left(\sum_{j \neq i} z_{ij} \right) \right) = 0. \end{aligned}$$

The last expression can be simplified and results in

$$c_i \gamma (n - 1) = \sum_{j \neq i} z_{ij},$$

which is independent of a_i . So, the only candidates for being minima or maxima for expression (30) are its values in the extrema, namely (31) and (32).

Also, the sign of the first derivative of (30) with respect to a_i is equal to the sign of $\sum_{j \neq i} z_{ij} - c_i \gamma (n - 1)$. So, if $c_i \gamma (n - 1) < \sum_{j \neq i} z_{ij}$ we have that (30) is strictly increasing in a_i , and then (31) is strictly greater than (32).

The value of (31) is between -1 and 1 , by assumption, because $0 < \sum_{j \neq i} z_{ij} < 2$.

The quantity in (32) is minimized by $c_i \rightarrow 0$; and $c'_i \rightarrow 0$. In this case (32) goes to -1 from the right, and for every $c_i > 0$ it will be greater than -1 . This completes the proof, because we have shown that any row of the Jacobian J sums to a number between -1 and 1 . \square

References

Ballester, C., Calvó-Armengol, A., Zenou, Y., 2006. Who’s who in networks. Wanted: the key player. *Econometrica* 74, 1403–1417.
 Battigalli, P., Catonini, E., de Vito, N., 2023. *Game Theory: Analysis of Strategic Thinking*. Book in preparation.
 Battigalli, P., Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., 2015. Self-confirming equilibrium and model uncertainty. *Am. Econ. Rev.* 105, 646–677.
 Battigalli, P., Francetich, A., Lanzani, G., Marinacci, M., 2019. Learning and self-confirming long-run biases. *J. Econ. Theory* 183, 740–785.
 Bervoets, S., Bravo, M., Faure, M., 2019. Stability in games with continua of equilibria. *J. Econ. Theory* 179, 131–162.
 Bochet, O., Faure, M., Long, Y., Zenou, Y., 2020. Perceived Competition in Networks. CEPR Discussion Papers.

- Bramoullé, Y., Kranton, R., 2007. Public goods in networks. *J. Econ. Theory* 135, 478–494.
- Bramoullé, Y., Kranton, R., 2016. Games played on networks. In: Galeotti, Andrea, Rogers, Brian, Bramoullé, Yann (Eds.), *Oxford Handbook on the Economics of Networks*.
- Bramoullé, Y., Kranton, R., D'Amours, M., 2014. Strategic interaction and networks. *Am. Econ. Rev.* 104, 898–930.
- Breza, E., Chandrasekhar, A.G., Tahbaz-Salehi, A., 2018. Seeing the forest for the trees? An investigation of network knowledge. NBER working papers.
- De Martí, J., Zenou, Y., 2015. Network games with incomplete information. *J. Math. Econ.* 61, 221–240.
- Feri, F., Pin, P., 2020. Externalities aggregation in network games. *Int. Econ. Rev.* 61, 1635–1658.
- Foerster, M., Mauleon, A., Vannetelbosch, V.J., 2021. Shadow links. *J. Econ. Theory* 197, 105325.
- Frick, M., Iijima, R., Ishii, Y., 2022. Dispersed behavior and perceptions in assortative societies. *Am. Econ. Rev.* 112 (9), 3063–3105.
- Fudenberg, D., Kreps, D., 1995. Learning in extensive-form games I. Self-confirming equilibria. *Games Econ. Behav.* 8, 20–55.
- Fudenberg, D., Levine, D.K., 1993. Self-confirming equilibrium. *Econometrica*, 523–545.
- Galer, O., 2007. *Discrete Dynamical Systems*. Springer Science & Business Media.
- Gilli, M., 1999. Adaptive learning in imperfect monitoring games. *Rev. Econ. Dyn.* 2, 472–485.
- Jackson, M.O., Wolinsky, A., 1996. A strategic model of social and economic networks. *J. Econ. Theory* 71, 44–74.
- König, M.D., Tessone, C.J., 2011. Network evolution based on centrality. *Phys. Rev. E* 84, 056108.
- König, M.D., Tessone, C.J., Zenou, Y., 2014. Nestedness in networks: a theoretical model and some applications. *Theor. Econ.* 9, 695–752.
- Kouachi, S., 2016. The Cauchy interlace theorem for symmetrizable matrices. [arXiv:1603.04151](https://arxiv.org/abs/1603.04151).
- Lipnowski, E., Sadler, E., 2019. Peer-confirming equilibrium. *Econometrica* 87, 567–591.
- McBride, M., 2006. Imperfect monitoring in communication networks. *J. Econ. Theory* 126, 97–119.
- McKee, J., Smyth, C., 2020. Symmetrizable integer matrices having all their eigenvalues in the interval $[-2, 2]$. *Algebraic Combin.* 3, 775–789.
- Milgrom, P., Roberts, J., 1991. Adaptive and sophisticated learning in normal form games. *Games Econ. Behav.* 3 (1), 82–100.
- Moulin, H., 1984. Dominance solvability and Cournot stability. *Math. Soc. Sci.* 7, 83–102.
- Savchenko, S., 2003. On the Perron roots of principal submatrices of co-order one of irreducible nonnegative matrices. *Linear Algebra Appl.* 361, 257–277.
- Stańczak, S., Wiczanowski, M., Boche, H., 2006. *Resource Allocation in Wireless Networks: Theory and Algorithms*. Springer Science & Business Media.
- Zenou, Y., 2016. Key players. In: Galeotti, Andrea, Rogers, Brian, Bramoullé, Yann (Eds.), *Oxford Handbook on the Economics of Networks*.