Learning and Selfconfirming Equilibria in Network Games

P. Battigalli F. Panebianco P. Pin

ASSA Meetings, 4 January 2020

Motivation

State of the Art

Analysis of interaction in networks, main ingredients:

- Strategic interaction (**explicit**): Individual incentives to act, local externalities (with neighbors), global externalities.
- Information (**Implicit**):
  - The network is usually assumed to be common knowledge.
  - When this is not the case: common knowledge of random network statistics, and random draw of a new network every period (mean-field analysis).

Motivation

State of the Art

Based on the assumption that the network is common knowledge, with very few exceptions, the literature uses Nash equilibrium (or some form of Bayesian Nash) to predict behavior.

HOWEVER

We are not aware of any discussion in this literature of whether Nash equilibrium (NE) is the appropriate solution concept. Indeed:

1. Even under the assumption of common knowledge of the network, the set of **rationalizable** actions (those *consistent with rationality and common belief in rationality*) may be larger than the NE set.
2. In most applications it is not reasonable to assume that agents know the network, hence it is even less reasonable to assume that it is common knowledge.
3. In some cases, NE has a learning foundation. But we found little discussion of such cases, i.e., of how agents may learn to play an NE in network games.

Aim of the Paper

This paper provides the first systematic discussion of what happens when agents:

- play a network game maximizing instantaneous payoff;
- **may** ignore (i) the type of strategic interaction on links; (ii) how the network is shaped; (iii) even that they are part of a network.
- they understand how their payoffs depend on an unknown "payoff state", which is (actually) a function of the actions of their neighbors;
- they have some conjecture about this payoff-relevant state;
- receive some feedback after playing and update conjectures;
- they *may have wrong conjectures* about the strategic environment and still keep them because they are "confirmed" (not falsified).

Our Tool: Selfconfirming Equilibrium

The selfconfirming equilibrium (SCE) concept requires that:

- **Subjective Rationality:** Agents *best reply to* their subjective *conjectures*.
- **Confirmation:** These conjectures may be wrong, yet they are *consistent with available evidence (confirmed).*

WHY DO WE USE IT?

- SCE has a clear learning foundation (appropriate for games with incomplete information).
- Let a temporal sequence of action profiles be *consistent with adaptive learning*, *if* it converges, *it converges to an SCE*, but not necessarily to an NE.
- Under some conditions about feedback and payoff functions, SCE is equivalent to NE, which—indirectly—provides a learning foundation for NE.

Summary of Results

1. We characterize the structure of the SCE set in a class of "standard" network games, adding specific assumptions about feedback and payoff functions, and we relate SCE to NE.
2. We provide conditions under which the NE concept has a learning foundation, despite the fact that agents may have incomplete information.
3. We show that being active/inactive is crucial for agents to make correct inferences about the payoff state.
4. We show when learning converges to a non-NE SCE.
5. We study the case of local and global externalities and show how *perceived centrality* determines the long-run outcome (SCE).
6. We also study what happens when (some feature of) the network is common knowledge and agents play a SCE with rationalizable conjectures.

An Explicative Example

Consider the case of a linear-quadratic network game:

$$\begin{array}{rccc} u_i: & A_i \times A_{-i} \times \mathcal{Z} & \rightarrow & \mathbb{R} \\ & (a_i, a_{-i}, \mathbf{Z}) & \longmapsto & \alpha_i a_i - \frac{1}{2} a_i^2 + a_i \sum_{j \in I} z_{ij} a_j \end{array}$$

(with parameters $\alpha_i > 0$, $\mathbf{Z} \in \mathcal{Z} \subseteq \mathbb{R}^{I \times I}$ =adjacency matrix, $z_{ii} = 0$). At the end of each period agents *observe their payoff realization*.

1. **Active players** can perfectly infer payoff state $x_i = \sum_{j \in I \setminus \{i\}} z_{ij} a_j$ backing it out of realized payoff (each i understands how u_i depends on a_i and x_i, possibly not how x_i depends on a_{-i}).
2. **Inactive players** may hold any conjecture about the payoff state, which cannot be inferred.

$\Longrightarrow$ **Observability by active players**

An Explicative Example

Consider the case of a linear-quadratic network game with global externalities

$$\begin{array}{llcl} u_i: & A_i \times A_{-i} \times \mathcal{Z} \times \mathbb{R} & \longrightarrow & \mathbb{R} \\ & (a_i, \mathbf{a}_{-i}, \mathbf{Z}, \gamma_i) & \longmapsto & \alpha_i a_i - \frac{1}{2} a_i^2 + a_i \sum_{j \in I} z_{ij} a_j + \gamma_i \sum_{j \neq i} a_j \end{array}$$

(with $\alpha_i > 0$, $z_{ii} = 0$). At the end of each period agents observe their payoff realization.

1. **Active players** cannot infer $x_i = \sum_{i \in I} z_{ij} a_j$.
2. **Inactive players** can infer the global externality term.

$\Longrightarrow$ Lack of observability by active players.

The Model

Primitives

- Set I of agents, with $|I| = n$.
- A network (adjacency matrix) $\mathbf{Z} \in \mathcal{Z} \subseteq \mathbb{R}^{I \times I}$, generic entry z_{ij}, with $z_{ii} = 0$.
- Neighborhood $N_i := \{j \in I : z_{ij} \neq 0\}$.
- Parameter space $\mathcal{Z} \subseteq \mathbb{R}^{I \times I}$ (compact).
- Actions: $a_i \in A_i = [0, \bar{a}]$.

The Model

Aggregator - Payoff relevant state

Let $X_i = [\underline{x}_i, \overline{x}_i]$ include all possible **payoff states** x_i for player/node i. Network parameters $\mathbf{Z}$ determine a local **aggregator** ℓ_i of co-players' actions, representing *local externalities:*

$$\begin{array}{rccc} \ell_i : & A_{-i} \times \mathcal{Z} & \longrightarrow & X_i \\ & (\mathbf{a}_{-i}, \mathbf{Z}) & \longmapsto & \sum_{j \in I} z_{ij} a_j \end{array}$$

such that $\ell_i \left(A_{-i} \times \mathcal{Z} \right)$ is connected (an interval).
Note: The model can be extended to non linear ℓ_i.

The **payoff-relevant state** for i is $x_i = \ell_i(\mathbf{a}_{-i}, \mathbf{Z})$.

The Model

Payoffs

Parameterized Payoff function

Agent i understands that his utility depends on (a_i, x_i) as follows:

$$\begin{array}{rccl} v_i : & A_i \times X_i & \longrightarrow & \mathbb{R} \\ & (a_i, x_i) & \longmapsto & \alpha_i a_i - \frac{1}{2} a_i^2 + a_i x_i \end{array}$$

Given $x_i = \ell_i(\mathbf{a}_{-i}, \mathbf{Z})$, we obtain i's **parameterized payoff function** in the network game: for each a_i, $u_{i,a_i} = v_{i,a_i} \circ \ell_i$, that is,

$$\begin{array}{rccl} u_i : & A_i \times A_{-i} \times \mathcal{Z} & \longrightarrow & \mathbb{R} \\ & (a_i, \mathbf{Z}) & \longmapsto & \alpha_i a_i - \frac{1}{2} a_i^2 + a_i \ell_i(\mathbf{a}_{-i}, \mathbf{Z}) \end{array}$$

Note: We can extend to general $v_i(a_i, \ell_i(\mathbf{a}_{-i}, \mathbf{Z}))$ *strictly quasi-concave* in a_i and *continuous* ($\Rightarrow$ technically, **nice game**).

The Model

Conjectures

Shallow Conjectures

- Agent i may be unaware of the variables/parameters $(\mathbf{a}_{-i}, \mathbf{Z})$ and the aggregator ℓ_i.
- He just cares about the aggregate x_i [without necessarily knowing that $x_i = \ell_i(\mathbf{a}_{-i}, \mathbf{Z})$].
- Each i has a *deterministic* conjecture $\widehat{x}_i \in X_i$ (*w.l.o.g. in nice games*).

Deep conjectures If i knows $u_i : A_i \times A_{-i} \times \mathcal{Z} \to \mathbb{R}$, then he forms conjectures $(\hat{\mathbf{a}}_{-i}, \hat{\mathbf{Z}})$ (relevant for strategic thinking).

Agents *maximize* their *instantaneous expected payoff*. Let

$$r_i(\hat{x}_i) := \arg\max_{a_i} v_i(a_i, \hat{x}_i)$$

(r_i is a *function*). Simple *linear-quadratic* model (assuming $\bar{a}$ large):

$$a_i = \max\{\alpha_i + \hat{x}_i, 0\}.$$

Action a_i^* is **justifiable** if $\exists \hat{x}_i \in X_i$, $a_i^* = r_i(\hat{x}_i)$ (in nice games, equivalent to a_i^* undominated).

The Model

Feedback

At the end of the period each i gets a **feedback** message according to a *known* function

$$f_i : \quad A_i \times X_i \quad \longrightarrow \quad M$$

Parameterized **network game with feedback:**

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

Here we assume that *payoff and feedback coincide*: $f_i = v_i$ for every $i \in I$. Hence, *realized payoffs* are *observable*.

Selfconfirming equilibrium

Definition

A profile $(a_i^*, \widehat{x}_i)_{i \in I} \in \times_{i \in I}(A_i \times X_i)$ of actions and (shallow, deterministic) conjectures is a **selfconfirming equilibrium at $\mathbf{Z}$** of G if, for each $i \in I$,

1. *(subjective rationality)* $a_i^* = r_i(\widehat{x}_i)$
2. *(confirmed conjecture)* $f_i(a_i^*, \widehat{x}_i) = f_i(a_i^*, \ell_i(\mathbf{a}_{-i}^*, \mathbf{Z}))$.

Let $\mathbf{A}_{\mathbf{Z}}^{SCE}$ (resp., $\mathbf{A}_{\mathbf{Z}}^{NE}$) denote the set of SCEs (resp., NEs) action profiles at $\mathbf{Z}$.

Equilibrium Characterization

Only Active Players

Consider the case of *strictly positive justifiable actions* (in linear-quadratic model: $\underline{x}_i > -\alpha_i$). This happens when there are local strategic complementarities or mild substitution (e.g., Ballester et al., 2006).
Linear-quadratic models satisfy *observability by active players* (if $a_i > 0$, i can infer x_i from the realized payoff=feedback).

Theorem

Consider a nice network game with feedback G satisfying observability by active players. Assume that, for all $i \in I$ and $\widehat{x}_i \in X_i$, $r_i(\widehat{x}_i) > 0$. Then, for each $\mathbf{Z}$, $\mathbf{A}_{\mathbf{Z}}^{SCE} = \mathbf{A}_{\mathbf{Z}}^{NE}$.

To get the intuition consider the linear-quadratic case and recall that the payoff is given by $\alpha_i a_i - \frac{1}{2}a_i^2 + a_i x_i$, in every SCE $a_i^* > 0$ (i is active) hence i infers x_i and best respond to a correct conjecture.

Equilibrium Characterization

Unrestricted Set of payoff states

We study what happens if $\min r_i(X_i) \leq 0$ for some $i \in I$.
Let $I^0 = \{i \in I : \min r_i(X_i) \leq 0\}$ =set of players for whom being inactive is justifiable.

Theorem

Fix $\mathbf{Z} \in \mathcal{Z}$. Assume observability by active players. A profile $(a_i^*, \widehat{x}^i)_{i\in I}$ of actions and conjectures is a self-confirming equilibrium if there exists $J \subseteq I$ such that

1. $\forall i \in I \diagdown J \subseteq I^0$, $a_i^* = r_i\left(\widehat{x}^i\right) = 0$.
2. $\forall i \in J$, $a_i^* = r_i(\hat{x}_i) > 0$ *and* $\widehat{x}^i = \ell^i(\mathbf{a}_{-i}^*, \mathbf{Z}^*)$.

1. Inactive agents with possibly wrong conjectures;
2. Active agents with correct shallow conjectures.

Let $\mathbf{A}_{\mathbf{Z},J}^{NE}$ denote the **NE** set **at Z** with **restricted player set** $J \subseteq I$. In *linear-quadratic games*,

$$\mathbf{A}_{\mathbf{Z}}^{SCE} = \cup_{I \setminus J \subseteq I^0} \mathbf{A}_{\mathbf{Z},J}^{NE} \times \left\{\mathbf{0}_{I \diagdown J}\right\}.$$

Example

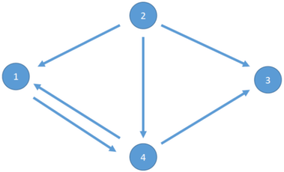

Figure: A network between 4 nodes. Every arrow is for an externality of equal magnitude and sign.

	All	$\{1,2,3\}$	$\{1,2,4\}$	$\{1,3,4\}$	$\{2,3,4\}$	$\{1,2\}$	...	$\emptyset$
a_1	0.1292	0.1	0.125	0.1292	0	0.1		0
a_2	0.1750	0.14	0.15	0	0.144	0.12		0
a_3	0.1	0.1	0	0.1	0.1	0		0
a_4	0.1458	0	0.125	0.1458	0.12	0		0

Table: Self confirming equilibria of the network, with all positive externalities of 0.2. The unique Nash Equilibrium is in bold.

Learning

One can adapt Milgrom & Robert's (GEB, 1991) "consistency with adaptive learning" to our framework with imperfect feedback (cf. Gilli RED, 1999):

Theorem (Convergence to SCE)

If $(\mathbf{a}_t)_{t=0}^{\infty}$ is consistent with adaptive learning and $\mathbf{a}_t \to \mathbf{a}^*$, then $\mathbf{a}^*$ is an SCE action profile.

Consider the *linear-quadratic* model.

Definition (Learning Process)

Players start at time 0 with a vector of conjectures $\widehat{\mathbf{x}}_0$. In each period t players best respond to their conjectures: $\forall i \in I$, $a^*_{i,t} = \max\{\alpha_i + \widehat{x}_{i,t}, 0\}$.
At the end of each period t players update their conjectures so that, if $a^*_{i,t} = 0$, then $\widehat{x}_{i,t+1} = \widehat{x}_{i,t}$; if instead $a^*_{i,t} > 0$, then $\widehat{x}_{i,t+1} = \frac{u_i(\mathbf{a}^*)}{a^*_{i,t}} - \alpha_i + \frac{1}{2}a^*_{i,t}$.

Note: We consider *stability with respect to perturbations of conjectures.*

Stability

From now on $\mathbf{Z}$ is the network $\mathbf{Z} \in \mathbb{R}^I \times \mathbb{R}^I$ (intensities included). For a given $\mathbf{a}^*$, let $I_{\mathbf{a}^*} := \left\{ i \in I : a_i^* > 0 \right\}$ (active players).

Theorem

Consider a selfconfirming strategy profile $\mathbf{a}^* \in \mathbf{A}_{\mathbf{Z}}^{SCE}$. If $\mathbf{Z}_{I_{\mathbf{a}^*}}$ satisfies at least one of the three conditions below:

1. *it has bounded values ($\left| z_{ij} \right| < \frac{1}{n}$ $i, j \in I_{\mathbf{a}^*}$),*
2. *it is negative and limited (spectral value $\rho \left(\mathbf{Z}_{I_{\mathbf{a}^*}} \right) < 1$),*
3. *or it is symmetrizable-limited,*

then $\mathbf{a}^*$ is locally stable and, for every $J \subseteq I_{\mathbf{a}^*}$, the following is another locally stable selfconfirming equilibrium: $a^{**} = \left(\mathbf{a}_J^{NE}, \mathbf{0}_{I \setminus J} \right)$, where $\mathbf{0}_J < \mathbf{a}_J^{NE}$ and $\mathbf{A}_{J,\mathbf{Z}}^{NE} = \left\{ \mathbf{a}_J^{NE} \right\}$.

Global Externalities

$$u_i(\mathbf{a},\mathbf{Z}) = \alpha_i a_i - \frac{1}{2}a_i^2 + a_i \underbrace{\sum_{j\neq i} z_{ij}a_j}_{\text{local: } \ell_i(\mathbf{Z},\mathbf{a}_{-i})\in\mathbb{R}} + \underbrace{\gamma\sum_{j\neq i} a_j}_{\text{global: } g_i(\mathbf{a}_{-i},\gamma)\in\mathbb{R}}$$

All the results on NE (and rationalizability with complete information, see below) are exactly the same as before, since $g_i(\mathbf{a}_{-i},\gamma)$ does not affect best responses. But, global ext. $\Rightarrow$confound in feedback$\Rightarrow$lack of observability by active players, different relationship SCE/NE, different dynamics.

	Line NE	Complete 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

Table: Simulations for the case of $\alpha = 0.1$ and $\gamma = 0.2$. Columns refer to 1) Nash Equilibrium of the line network; 2) Nash equilibrium of complete network; 3) selfconfirming equilibrium in the star network in which agent believe that $\ell_i\left(\mathbf{a}^*_{-i},\mathbf{Z}\right) = g_i\left(\mathbf{a}^*_{-i},\gamma\right)$.

Knowledge of the Network and Rationalizability (preliminary)

Do (common) knowledge of the network and its use in iterated strategic reasoning help in the selection of some SCEs? We use the concept of **selfconfirming equilibrium in rationalizable conjectures**, for which a kind of learning foundation is available (unlike its more demanding cousin, the rationalizable-SCE).

1. We assume symmetric and common knowledge of the uncertainty space $\mathcal{Z}$ (for simplicity, no private information). We study 2 cases: i) the network is common knowledge $\mathcal{Z} = \{\mathbf{Z}\}$; ii) Only signs and bounds are commonly known.
2. Agents form **deep conjectures** $\mu_i \in \Delta(\mathbf{A}_{-i} \times \mathcal{Z})$. Since we have a nice game, w.l.o.g. we can consider *deterministic deep conjectures*, $\mu_i \in \mathbf{A}_{-i} \times \mathcal{Z}$.
   1. Under common knowledge of the network, essentially $\mu^i = \hat{\mathbf{a}}^i_{-i} \in \mathbf{A}_{-i}$;
   2. Under symmetric ignorance (only sign+bounds knowledge), $\mu^i = \left(\hat{\mathbf{a}}^i_{-i}, \hat{\mathbf{Z}}^i\right) \in \mathbf{A}_{-i} \times \mathcal{Z}$.
3. Do iterated elimination of non-best replies to deep (deterministic) conjectures. Results depend on what is commonly known and complements *vs* substitutes.

Join us at

THE 12TH WORLD CONGRESS OF THE ECONOMETRIC SOCIETY

2020 World Congress of the Econometric Society

August 17, 2020-August 21, 2020

Bocconi University

Milan, Italy

THANK YOU