General Revision Protocols in Best Response Algorithms for Potential Games

Pierre Coucheney, Stéphane Durand, Bruno Gaujal, Corinne Touati

General Revision Protocols in Best Response Algorithms for Potential Games

Pierre Coucheney*, Stéphane Durand†, Bruno Gaujal¶‡§ and Corinne Touati¶‡§

* Univ. of Versailles, PRISM, France

† ENS Lyon, France

‡ Univ. Grenoble Alpes, LIG, F-38000 Grenoble, France

§ CNRS (French National Center for Scientific Research), LIG, F-38000 Grenoble, France

¶ Inria

***Abstract*—In this paper, we characterize the revision sets in different variants of the best response algorithm that guarantee convergence to pure Nash Equilibria in potential games. We prove that if the revision protocol is separable (to be defined in the paper), then the greedy version as well as smoothed versions of the algorithm converge to pure Nash equilibria. If the revision protocol is not separable, then convergence to Nash Equilibria may fail in both cases. For smoothed best response, we further show convergence to Nash Equilibria with optimal potential when players can only play one by one. Again this may fail as soon as simultaneous play is allowed, unless the number of players is two. We also provide several examples/counter-examples testing the domain of validity of these results.**

Keywords*—*Potential Games; Best Response; Logit Dynamics.

I. Introduction

Potential games have been introduced in [1] and have proven very usefuf, especially in the context of routing games, first mentioned in [2] and exhaustively studied ever since, in the transportation as well as computer science litterature, see for example [3]–[5] and for distributed optimization (see for example [6]).

While it is well-known that the Best Response Algorithm (BRA) converges to pure Nash equilibria (NE) in potential games [7], the robustness of this result w.r.t. the underlying assumptions has attracted surprisingly little attention.

In this paper we investigate the following question: under which revision protocol does BRA converge to NE? The classical BRA assumes that the revision protocol is *asynchronous*: At each round, a single player is given a chance to revise its strategy.

It can be shown (see Section II) that violating this assumption and allowing for *simultaneous* revision protocols can compromize the convergence to any Nash Equilibrium. So why should one consider simultaneous revisions anyway? Because such revisions may happen in practical cases. Indeed, partially simultaneous revisions can be implemented in a fully distributed way (for example, each player decides to play (or not) at each round, independently of the others). On the other hand, an asynchronous revision requires a central controller, or an election protocol between the players, to select the next player to play. In practical cases, this is either cumbersome or even impossible when the game involves many players with no coordination mechanism between them.

This paper is organized as follows. In Section III, we give a necessary and sufficient condition on the revision protocol to converge to NE in all potential games. This condition is a structural condition on the support of the revision law. It is rather easy to check and can be enforced in fully distributed games. We also provide a game-dependent version of this condition that is based on the neighboring graph between players induced by the game. This also guarantees convergence to a Nash equilibrium of the game. Section IV extends our result to a smoothed version of BRA, that provides, in contrast with BRA, a guarantee of convergence to one Nash Equilibrium with maximal potential for special cases of separable revision laws (more precisely when players can only play one by one). In the last section (§V), several examples show that the conditions given in the theorems cannot be improved.

II. Potential Games, Best Response Algorithm and the Revision Protocol

We consider a finite game $\mathfrak{G} \stackrel{\text{def}}{=} (\mathcal{N}, \mathcal{A}, u)$ consisting of

- a finite set of *players* $\mathcal{N} = \{1, \dots, N\}$;
- a finite set $\mathcal{A}_k$ of *actions* (or *pure strategies*) for each player $k \in \mathcal{N}$; The set of (action) *profiles* or *states* of the game is $\mathcal{A} \stackrel{\text{def}}{=} \prod_k \mathcal{A}_k$;
- the players' *payoff functions* $u_k : \mathcal{A} \to \mathbb{R}$ that players seek to maximize.

We define the classical *best response correspondence* $\text{BR}_k(x)$ as the set of all actions that maximizes the payoff for player k under profile x:

$$\text{BR}_k(x) \stackrel{\text{def}}{=} \left\{ \underset{\alpha \in \mathcal{A}_k}{\text{argmax}}\ u_k(\alpha; x_{-k}) \right\}. \tag{1}$$

A *Nash equilibrium* (NE) is a fixed point of the correspondence, i.e. a profile x^* such that $x^*_k \in \text{BR}_k(x^*)$ for every player k.

Iteratively playing a best response may not converge in general. We consider here the specific class of potential games for which convergence is ensured.

Definition 1 (Potential games and its variants). A game is an (exact) *potential game* [7] if it admits a function (called the potential) $F : \mathcal{A} \to \mathbb{R}$ such that for any player k and any *unilateral* deviation of k from action profile x to x'

$$u_k(x) - u_k(x') = F(x) - F(x'). \tag{2}$$

A game is a *generalized ordinal potential game* [7] (or G-potential game for short) if there is $F : \mathcal{A} \to \mathbb{R}$ such that, for any player k and any *strictly profitable unilateral* deviation of k from action profile x to x', $F(x') > F(x)$.

A game is a *best-response potential game* [8] (or BR-potential game for short) if there is $F : \mathcal{A} \to \mathbb{R}$ such that for any player k and action profile x

$$\mathrm{BR}_k(x) = \left\{ \operatorname*{argmax}_{\alpha \in \mathcal{A}_k} F(\alpha, x_{-k}) \right\}. \tag{3}$$

As shown in [8], BR-potential games are characterized by the fact that any sequence of profiles generated by unilateral best response, and containing at least one strict improvement, is not a cycle. In particular, it can be seen that exact potential games are BR-potential games, but there exist G-potential games that are not BR-potential games. Yet, by imposing that the best response correspondence is univaluated (our next assumption), it becomes true that any G-potential game is also a BR-potential game.

To avoid ties, and unless otherwise mentioned, we make the following assumption in the rest of the paper.

Assumption 1 (Uniqueness of Best Response). We assume that the Best Response correspondence is a univaluated function; for any player k and profile x:

$$\mathrm{BR}_k(x) = \operatorname*{argmax}_{\alpha \in \mathcal{A}_k} u_k(\alpha; x_{-k}). \tag{4}$$

For example, this assumption holds when players face different payoffs at each state:

$$u_k(\alpha, x_{-k}) \neq u_k(\beta, x_{-k}), \tag{5}$$

whenever $\alpha \neq \beta$, for all k and all x. This can be imposed by perturbing the payoffs. Also, for general payoffs, breaking ties can be done by ranking players and actions using an arbitrary fixed order.

As mentioned earlier, a G-potential game that satisfies this assumption is necessarily a BR-potential game. Hence all the results will be stated for BR-potential games in the following.

Let us now focus on the *revision protocol*, that is the sequence of players that can revise their strategy . We first consider a version of the asynchronous *Best Response Algorithm* (asyncBRA) where the next player is selected according to a revision protocol driven by a random process. Its probability distribution ρ is called the *revision law* over the players:

$$\forall k \in \mathcal{N}, x \in \mathcal{A},\ \mathbb{P}(\text{selected player} = k \,|\, \text{profile } x) = \rho_x(k).$$

We assume that the revision law is *communicating* meaning that the probability of choosing any player is strictly positive over time. Without this assumption, the algorithm may clearly not converge to a NE in general.

Theorem 2 (asyncBRA converges to NE [7]). *For any BR or G-potential game $\mathfrak{G}$, Algorithm 1 converges in finite time, almost surely, to a Nash Equilibrium of $\mathfrak{G}$.*

Proof: This is a well known result, we only provide a sketch of the proof.

Algorithm 1: Best Response Algorithm (asyncBRA) with random selections

foreach *player* $k \in K$ **do**
  $stop_k := false$
repeat
  Pick player $k \in \mathcal{N}$ using law ρ_x
  Select new action $\alpha_k := \mathrm{BR}_k(x)$
  $stop_k := \mathbf{1}_{\{\alpha_k = x_k\}}$;
  $x_k := \alpha_k$;
until $stop_1 \wedge stop_2 \wedge \cdots \wedge stop_N$;

Clearly, when a NE is reached, the algorithm stops when each player have had the opportunity to revise its strategy.

Otherwise, a player has incentive to change its strategy, which will lead to a strict improvement of the potential. By characterization of BR-potential games, this state will never be visited again. Since the number of states is finite, the algorithm will reach a NE in finite time, almost surely. ∎

Let us now consider the algorithm BRA under a general revision protocol ρ that allows several players to change their strategy simultaneously. In that case, the revision law ρ is a distribution over sets of players, that, in full generality, depends on the current profile x:

$$\forall K \subset \mathcal{N}, x \in \mathcal{A},$$
$$\mathbb{P}(\text{set of selected players} = K \,|\, \text{profile } x) = \rho_x(K).$$

The sets whose probability is positive define the *support of* ρ_x, denoted by $\mathcal{S}(\rho_x)$. In the following, we will assume that the support does not depend on the profile, and that each player has a chance to revise its strategy.

Assumption 2 (Constant support). The support of the revision law ρ,

(i) is constant wrt the profile x (hence is just denoted $\mathcal{S}(\rho)$):

$$\forall x, y \in \mathcal{A}, \mathcal{S}(\rho_x) = \mathcal{S}(\rho_y); \tag{6}$$

(ii) covers all players:

$$\forall x \in \mathcal{A}, \cup_{K \in \mathcal{S}(\rho_x)} K = \mathcal{N} \tag{7}$$

Algorithm 2: Best Response Algorithm BRA(ρ) with general revision protocol ρ

foreach *player* $k \in \mathcal{N}$ **do**
  $stop_k := false$
repeat
  Pick a set of players $K \subset \mathcal{N}$ using law ρ_x
  foreach *player* $k \in K$ *simultaneously* **do**
    Select new action $\alpha_k := \mathrm{BR}_k(x)$;
    $stop_k := \mathbf{1}_{\{\alpha_k = x_k\}}$;
  $x := \alpha$;
until $stop_1 \wedge stop_2 \wedge \cdots \wedge stop_N$;

To keep notations simple, for every set K of players, we will denote by $\mathrm{BR}_K(x)$ the action profile obtained by simultaneous best responses for each player in K, under x. Hence,

$x' = \text{BR}_K(x)$ means that, for all $k \in K$, $x'_k = \text{BR}_k(x)$ and for all $j \notin K$, $x'_j = x_j$.

When several players move simultaneously, the potential may not be increasing and then the convergence of BRA to a NE of the game is not guaranteed as shown in the following example.

Example 1 (**No convergence to NE for simultaneous revision**). Let us consider a 2-player 2-action (a and b) potential game with the following (exact) potential:

$$F = \begin{array}{|c||c|c|} \hline 1\backslash 2 & a & b \\ \hline a & 0 & 2 \\ b & 3 & 1 \\ \hline \end{array}$$

If the revision protocol always makes the two players play simultaneously ($\rho(\{1,2\}) = 1$) and if Algorithm 2 starts with action profile (a, a), then, during the run, both players keep changing their strategy simultaneously from (a, a) to (b, b) and back, and they never reach the two NEs (a, b) and (b, a).

III. Separability of the revision protocol

This section is dedicated to the description of a necessary and sufficient condition on the revision protocol ρ that guarantees convergence to a NE for Algorithm BRA(ρ) in *all* BR-potential games. It should be clear that this condition only depends on the support of ρ (which sets have positive probability under ρ) rather than on the actual values of the probabilities. Indeed, since players are selected at each round independently of the previous rounds, Borel-Cantelli lemma implies that any sequence of revision sets with positive probability will occur infinitely often regardless of the probability values.

To state a necessary and sufficient condition for convergence to a NE, we need to introduce several definitions. A family $\mathcal{F}$ of sets of player is a *set cover* of $\mathcal{N}$ if $\cup_{K \in \mathcal{F}} K = \mathcal{N}$.

Definition 3 (**Separable family**). Let $\mathcal{F}$ be a set cover, consider the following iterative elimination process: as long as there is a singleton (say $\{k\}$) in $\mathcal{F}$, remove player k from all sets in $\mathcal{F}$. Then, $\mathcal{F}$ is *separable* if the elimination process reduces $\mathcal{F}$ to the empty set. A revision law ρ is *separable* if its support is separable.

Another way to state the separability property is the following: $\mathcal{F}$ is separable if and only if there is a permutation $(k_1, k_2, \dots k_N)$ of $\mathcal{N}$ such that the following sets all belong to $\mathcal{F}$:

$$\begin{aligned} K_1 &= \{k_1\}, \\ K_{i+1} &= \{k_{i+1}\} \cup L_i, \ \forall i \in \{1, \dots, N-1\} \end{aligned} \tag{8}$$

where L_i is any (possibly empty) set included in $\{k_1, \dots k_i\}$.

Yet another characterization, similar to the previous one, but more compact, is that $\mathcal{F}$ contains N sets $K_1, \dots K_N$ such that, for all i,

$$K_i \setminus \cup_{j<i} K_j \text{ is a singleton.} \tag{9}$$

Obviously, if a family $\mathcal{F}$ is separable, then adding any set to $\mathcal{F}$ preserves separability. It should also be clear that if $\mathcal{F}$ is separable, then the family $\mathcal{F}'$, obtained by removing one player from all the sets in $\mathcal{F}$, is separable over the remaining $N-1$ players.

Example 2 (**Separable families and laws**).

- If $\mathcal{F}$ contains all the singletons, $\{k_1\}, \{k_2\}, \cdots, \{k_N\}$ then $\mathcal{F}$ is separable.
- If the family $\mathcal{F}$ contains the sets $\{k_1\}, \{k_1, k_2\}, \cdots, \{k_1, k_2, \dots, k_N\}$ then $\mathcal{F}$ is separable.
- The following revision law ρ is separable: Each player k_i chooses to play independently of the others, with some positive probability p_i. This revision law is separable because each singleton has a positive probability to be played. For all i,
$$\rho(\{k_i\}) = p_i \prod_{j \neq i} (1 - p_j) > 0.$$
This revision law is also fully distributed since it does not require any coordination between players.

We are now ready to state the main result of this section.

Theorem 4 (**Convergence to NE for separable revisions).** *Let $\mathcal{N}$ be a set of players and ρ be a revision law over $\mathcal{N}$. Algorithm BRA(ρ) converges a.s. to a NE for all BR-potential games $\mathfrak{G}$ over $\mathcal{N}$ if and only if ρ is separable.*

Proof: We prove the sufficient and necessary conditions separately.

a) Sufficient condition: The sufficient condition is proved by contradiction.

Algorithm BRA(ρ) naturally induces a Markov process over $\mathcal{A}$ (up to adding a self-loop transition at each NE with probability one instead of stopping the algorithm). Let R be the set of recurrent action profiles, i.e. profiles that are visited infinitely often with probability one. All NE are recurrent since the algorithm stays there when a NE is reached. We are going to show that R is only made of NEs.

By contradiction, assume that there exists a BR-potential game together with a separable revision law ρ for which algorithm BRA(ρ) does not always converge to a NE. Let x be one recurrent action profile that is not a NE, with the highest potential. Starting from $X_0 = x$, let us consider a run of the algorithm using the sequence $K_1, K_2, K_3, \dots, K_N$ of sets of players as defined in (8) (having positive probability by separability of ρ) which generates the sequence of profiles $X_1, X_2, \dots, X_N$.

First notice that, if $X_i = x$ for all $i \in \{1, \dots, j\}$, then players $k_1, \dots, k_j$ already play a best response to x. In particular, j cannot be equal to N since this would imply that x is a NE.

So let i be the smallest index in $\{1, \dots, N\}$ such that $X_i \neq x$. We have $X_i = \text{BR}_{k_i}(x)$ since all other players in K_i already played a best response to x. This implies that the potential strictly increases along this deviation. Since X_i is in the same recurrent class as x and x is the recurrent non-NE with the highest potential, we deduce that X_i is a NE. But this contradicts the fact that x is recurrent since the algorithm will never visit x again.

b) Necessary condition: The necessary part is proved by induction on the number of players.

First notice that if the game has a single player, the synchronous and asynchronous algorithms are identical, which implies the convergence in one step to the unique NE.

Let us assume that any revision law that guarantees a.s. convergence on any game with $N-1$ players is separable. This induction property is denoted $\mathscr{P}(N-1)$.

We now consider a revision law ρ over a set $\mathcal{N}$ of players of size N such that for any game $\mathfrak{G}$ over $\mathcal{N}$, every execution of Algorithm BRA(ρ) converges a.s. to a NE. We will show that ρ is separable.

We first show that the support of ρ must contain a singleton. For that, let us construct a specific game over $\mathcal{N}$ with action set $\mathcal{A}_k = \{0, \ldots, p-1\}$ for every player k, where p is a prime number larger than N.

The payoffs are taken identical for every players and equal to the potential $F(x) = -(\sum_k x_k \mod p)$. The maximum of the potential is 0, which is reached when e.g. all players choose action 0.

The best response of any player k in a state with potential $-h$ is to choose action $(x_k - h) \mod p$. Therefore, when m players play simultaneously in a state with potential $-h$, the potential becomes $-((h + m(-h)) \mod p) = -((1-m)h \mod p)$. Starting from a state with potential -1 and using a sequence of revision sets with respective sizes $m_1, \ldots, m_\ell, \ldots$ the successive values of the potential are:

$$\begin{aligned}
&-((1-m_1) \mod p),\\
&-((1-m_1)(1-m_2) \mod p),\\
&\vdots\\
&-((1-m_1)(1-m_2)\cdots(1-m_\ell) \mod p),\\
&\vdots
\end{aligned}$$

Since BRA(ρ) converges a.s. to a NE, this sequence reaches value 0 in finite time. Since p is prime, the only possibility is that there exists a revision set with size 1 in the support of ρ, say $\{k\}$. Hence, any revision law that satisfies property $\mathscr{P}(N)$ must contain a singleton.

We are now ready for the second step of the proof that uses the induction assumption. We first construct a revision law ρ_{-k} over the set of players $\mathcal{N} \setminus \{k\}$ whose support is obtained by removing k from all the revisions sets in $\mathcal{S}(\rho)$. Now, we can claim that $BRA(\rho_{-k})$ converges in all games with $N-1$ players: Indeed, from any game $\mathfrak{G}$ over $N-1$ players, one can construct a game $\mathfrak{G}^+$ over N players by adding a dummy player (say k) with a single action that does not affect the utilities of the other players. Since $BRA(\rho)$ converges on all games with N players, it converges on game $\mathfrak{G}^+$. Since the added player does not play any role in $\mathfrak{G}^+$, this implies that the revision law ρ_{-k} converges on $\mathfrak{G}$. By the induction assumption, $\mathscr{P}(N-1)$, ρ_{-k} is separable. By definition of separability, ρ is also separable. ■

The notion of separability is game-independent. Given a specific game $\mathfrak{G}$, it is possible to define a new version of the separability property (called $\mathfrak{G}$-separability in the following) that depends on $\mathfrak{G}$ or more precisely on the interaction of players in $\mathfrak{G}$.

More formally, players k and ℓ do not interact in game $\mathfrak{G} = (\mathcal{N}, \mathcal{A}, u)$ if for every profile x, functions $u_k(x_\ell, x_{-\ell})$ and $u_\ell(x_k, x_{-k})$ are constant w.r.t. to, respectively, variables x_ℓ and x_k. In the opposite case, we say that players k and ℓ are *neighbors*. We consider the graph whose vertices are the players, and the undirected edges are defined by the neighbor relation. The vertices of this graph can be colored such that no neighbors have the same color. Notice that the number of colors is not unique, for example one can use one color per player whatever the graph.

Definition 5 (**$\mathfrak{G}$-separability).** Let us consider that the interaction graph of game $\mathfrak{G}$ is colored with a set $\mathcal{C}$ of colors. Let $\mathcal{H}$ be a separable family over $\mathcal{C}$. For each set C in $\mathcal{H}$, consider the set $\mathcal{N}(C)$ of all the players whose color is in C. The collection of all such sets $\mathcal{N}(C)$ for all C in $\mathcal{H}$ is called a $\mathfrak{G}$-separable family over $\mathcal{N}$.

It should be clear that, if all players are neighbors in $\mathfrak{G}$, $\mathfrak{G}$-separability coincides with separability, because all players must have a different color. It is also clear that separability implies $\mathfrak{G}$-separability for all games by coloring all players with distinct colors.

Example 3 (**$\mathfrak{G}$-separable, but not separable family**). Here is an example of a $\mathfrak{G}$-separable revision family that is not separable. Consider a game $\mathfrak{G}$ with 4 players such that 1 is a neighbor of 2, 2 is a neighbor of 3, and 3 is a neighbor of 4, with no other neighboring relations. The players can be colored with two colors, *Blue* for players 1 and 3 and *Red* for players 2 and 4. Consider the separable revision family $\mathcal{H}$ over the colors, made of two sets, $\{Blue\}, \{Blue, Red\}$. Now the family $\mathcal{F}$ is made of two sets $\{1,3\}, \{1,2,3,4\}$. By definition, $\mathcal{F}$ is $\mathfrak{G}$-separable but it is not separable because it does not contain any singleton.

The main (straightforward) property of games where two players (say k and ℓ) do not interact with one another is

$$\mathrm{BR}_{\{k,\ell\}}(\alpha) = \mathrm{BR}_k(\mathrm{BR}_\ell(\alpha)) = \mathrm{BR}_\ell(\mathrm{BR}_k(\alpha)). \tag{10}$$

In other words, letting k and ℓ play simultaneously or one after the other leads to the same state.

Corollary 6 (**$\mathfrak{G}$-separability implies convergence to NE).** *Let $\mathfrak{G}$ be a BR-potential game with N players and let ρ be $\mathfrak{G}$-separable. Then BRA(ρ) converges to a NE of the game.*

Proof: The proof is similar to the proof of Theorem 4.

Let us assume that there is a set of recurrent profiles R that does not contain any NE. Let x be a strategy with the largest potential in R. Starting from x, let us consider a trajectory of the algorithm using the sequence $K_1, K_2, \ldots, K_m$ of revision sets, ordered according to (8) for the associated separable family over the colors. The first time that the action profile becomes different from x (to a new profile y), the potential increases because of the order chosen on the sets. Since x is the non NE state with the largest potential in R, then $y \notin R$ unless it is a NE. Since x is recurrent, y cannot be a NE. And $y \in R$ because there is a path coming from the recurrent state x.

Therefore, applying the sequence of revision sets $K_1, \ldots, K_m$ starting in x, the algorithm remains in x.

This implies that $x_k = BR_k(x)$ for all players $k \in \mathcal{N}$, so x is a NE of the game.

By contradiction, this implies convergence to NE of Algorithm 2.

■

IV. Extension to Smoothed Best Response

Theorem 4 says that Algorithm BRA(ρ) converges to a NE in potential games under the separability condition. This implies convergence to a local minimum of the potential. However its potential can be arbitrarily far from the global maximum.

To ensure convergence to an optimal NE (maximizing the potential), one can replace the greedy best response (used so far) by a smoothed best response allowing the algorithm to escape from local maxima.

We denote by $\mathbf{u}_k(x_{-k})$ the payoff vector of player k under action profile x: $\mathbf{u}_k(x_{-k}) = (u_k(\alpha, x_{-k}))_{\alpha \in \mathcal{A}_k}$.

A *random choice* Q is a random variable over $\mathcal{A}_k$, the actions of player k whose law only depends on the payoff vector $\mathbf{u}_k(x_{-k})$. This random choice is used to modify algorithm BRA(ρ) in the following way.

Algorithm 3: Smoothed BR algorithm with revision law ρ and random choice Q, SmoothBRA(ρ, Q)

```
repeat
    Pick a set of players K ⊂ N according to ρ_x;
    foreach player k ∈ K simultaneously do
        Select action α_k := Q(u_k(x_{-k}));
    x := α;
until infinity;
```

Due to its random nature, and unlike BRA(ρ), this algorithm never ends. Its convergence properties will only be given in terms of its asymptotic distribution[1].

In the following, we will focus on the classical *logit choice*, parametrized by the temperature $1/\theta$. For each player k and profile x, the law is given by:

$$\mathbb{P}\left[Q(\mathbf{u}_k(x_{-k})) = \alpha\right] = \frac{\exp(\theta u_k(\alpha, x_{-k}))}{\sum_{\beta \in \mathcal{A}(k)} \exp(\theta u_k(\beta, x_{-k}))}.$$

Note that the logit choice is a close approximation of the BR mapping when θ goes to infinity.

Under this random choice, the sequence $(X_n)_{n\in\mathbb{N}}$ of action profiles computed by SmoothBRA(ρ, Q) forms a Markov chain whose transition matrix can be constructed as follows.

Let Diff$(x, y) \stackrel{\text{def}}{=} \{k : x_k \neq y_k\}$ be the set of players that must have played if the sequence $(X_n)_{n\in\mathbb{N}}$ jumps from $X_n = x$ to $X_{n+1} = y$ in one step. Of course, the set of players that actually played can be larger because some of them may have chosen to not change their action. This is the reason why we introduce the intermediate matrix P^V, defined for all sets V of players by:

$$P^V_{x,y} = \begin{cases} \displaystyle\prod_{k\in V} \frac{\exp(\theta u_k(y_k, x_{-k}))}{\sum_{\alpha\in A(k)} \exp(\theta u_k(\alpha, x_{-k}))} & \text{if Diff}(x,y) \subseteq V \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

Then the transition matrix P is

$$P_{x,y} = \sum_{V \supseteq \text{Diff}(x,y)} \rho_x(V) P^V_{x,y}.$$

The asymptotic behavior of algorithm SmoothBRA(ρ, Q) is given by the stationary distribution π of this ergodic Markov chain on $\mathcal{A}$.

The well-known Markov chain tree theorem provides an explicit formula for π (up to a multiplicative factor), based on spanning trees over the Markov chain transition graph.

Theorem 7 (Markov Chain Tree Theorem [10]). *Let $\mathcal{T}_x$ be the set of spanning in-trees of the transition graph, with root in x. The stationary probability π_x is proportional to the sum of the probability weights of all the spanning trees T in $\mathcal{T}_x$:*

$$\pi_x \propto \sum_{T\in\mathcal{T}_x} \prod_{(y,z)\in T} P_{y,z}.$$

The stationary distribution π puts a positive probability on all action profiles x, meaning that all action profiles will be visited infinitely often during the execution of SmoothBRA(ρ, Q). However, when θ goes to ∞, the probability mass will concentrate on some profiles. Such profiles are called *stochastically stable* under ρ.

Stochastic stability can be asserted using the following lemma, based on the *orders* (w.r.t. θ) of the transition probabilities. When θ goes to ∞, $\varepsilon \stackrel{\text{def}}{=} \exp(-\theta)$ goes to 0.

Let us express the transition probabilities as a function of ε instead of θ:

$$P_{x,y} = \sum_{V\supseteq \text{Diff}(x,y)} \rho_x(V) \prod_{k\in V} \frac{\varepsilon^{-u_k(y_k,x_{-k})}}{\sum_{\alpha\in A(k)} \varepsilon^{-u_k(\alpha,x_{-k})}}. \quad (12)$$

which can be written under the following first order development w.r.t. ε:

$$P_{x,y} = c_{x,y}\varepsilon^{q_{x,y}} + o(\varepsilon^{q_{x,y}}),$$

where $q_{x,y}$ is called the *order* of $P_{x,y}$ w.r.t. ε (or, equivalently, w.r.t. θ).

Lemma 8 (Stochastic stability characterization [11]). *State x is stochastically stable if and only if the order of its minimal in-tree is the smallest, among all in-trees.*

Proof: This characterization has first been given in [11]. The following proof is similar to the original proof. It is helpful to detail it here to further highlight the notion of orders, used in the rest of the paper.

[1]For practical purpose, one needs to stop SmoothBRA(ρ, Q) after a finite number of iterations. Many heuristic stopping rules have been proposed (see [9] for example), but in practice there is a fair amount of "black magic" involved and no single stopping rule provides guarantees in the general case.

Using Equation (12), for any pair of states x and y, one can compute the order $q_{x,y}$ as

$$q_{x,y} = \min_{V \supseteq \text{Diff}(x,y) \cap \mathcal{S}(\rho)} \left(\sum_{k \in V} \left(\max_{\alpha \in \mathcal{A}_k} u_k(\alpha, x_{-k}) - u_k(y_k, x_{-k}) \right) \right). \tag{13}$$

Note that the orders of the transitions do not depend on the values of ρ but only of its support. Therefore, the limit value of π only depends on $\mathcal{S}(\rho)$, as well. By using the Markov chain tree theorem (Theorem 7), the order q_x of π_x w.r.t. ε is

$$q_x \stackrel{\text{def}}{=} \min_{T \in \mathcal{T}_x} \sum_{(y,z) \in T} q_{y,z}$$

Therefore, the only components in π that do not go to 0 when ε goes to 0 are those with the smallest order:

$$(\lim_{\theta \to \infty} \pi_x > 0) \Leftrightarrow (q_x = \min_{y \in \mathcal{A}} q_y). \tag{14}$$

■

Equation (13) provides an explicit formula of the order of the transition probability from x to y. In particular, one may observe that the order is always non-negative, and equal to zero if and only if there is a revision set that only involves best responses of the players in that set for going from x to y.

The next theorem states that the total probability mass of the NEs goes to one when θ goes to ∞. This implies that Algorithm SmoothBRA(ρ, Q) will only visit NE with a high probability when θ is large.

Theorem 9 (Convergence to NE). *Let $\mathfrak{G}$ be a BR-potential game. If ρ is $\mathfrak{G}$-separable, then for all action profiles x that are not NE, the stationary probability π_x goes to 0 as θ goes to ∞.*

Proof: The proof is based on Lemma 8 that allows to characterize the states with the smallest order. From the remark following the lemma, $q_{x,y}$ is non-negative. It is equal to 0 if and only if there is a set V in Diff$(x,y) \cap \mathcal{S}(\rho)$ such that $y = \text{BR}_V(x)$.

From Corollary 6, if ρ is $\mathfrak{G}$-separable, and starting from any action profile x that is not a NE, there is a finite sequence $(V_n)_{0 \leqslant n < H}$ of sets of players in $\mathcal{S}(\rho)$ and actions profiles $(X_n)_{0 \leqslant n \leqslant H}$ such that

$$\begin{aligned} X_0 &= x \\ X_{n+1} &= \text{BR}_{V_n}(X_n), \quad \forall\, 0 \leqslant n < H, \end{aligned}$$

and X_H is a NE. Using these sequences for all x constructs a path whose order is 0, leading to the NE X_H.

Let T_x^* be the tree with minimal order, routed in x. From T_x^*, it is possible to construct a tree routed in X_H by adding the path from x to X_H with order 0 and removing the arc in T_x^* starting in X_H. This arc has a strictly positive order because X_H is a NE. The new tree has an order strictly smaller that T_x^*, so x cannot achieve the minimum in (14). Therefore, only NE may have positive stationary probabilities. ■

Using the same construction, one can show more generally that any stochastically stable state is included in the recurrent set of algorithm SmoothBRA(ρ, Q).

This result is to be compared with [11] where a similar result is proved. It seems however that the absence of separability and of Assumption 1 in [11] jeopardizes their result.

A stronger result can be proved for revisions that do not allow for simultaneous revisions in the case of exact potential games.

Theorem 10 (Convergence to optimal NE for asynchronous revisions). *Let $\mathfrak{G}$ be a game with an exact potential F. If the revision law does not contain any simultaneous play (in other words, $\mathcal{S}(\rho) = \{\{k\}, k \in \mathcal{N}\}$) then the only stochastically stable profiles are the optimal NE.*

Proof: The proof of this result is a direct consequence of the results in Chapter 12 of [12]. We provide a short proof in the sake of completeness. First note that the stochastically stable profiles do not depend on the actual values of $\rho(\{k\})$, as mentioned before. Let us consider the case where all of them are equal: $\rho(\{k\}) = 1/N$ for all $k \in \mathcal{N}$. In that uniform case, the Markov chain (X_n) is reversible and the stationary probability is explicitly known: for all profiles x, $\pi_x \propto \exp(\theta F(x))$. Therefore, when θ goes to ∞, the total stationary probability of the profiles with optimal potential will go to one. ■

V. Examples

The following examples show that there is little hope to prove more precise results, at least in the general case.

Example 4 shows that Theorem 9 is not true when the revision protocol is not $\mathfrak{G}$-separable. This example shows a game where SmoothBRA has stable non-NE points under a non-separable protocol.

Example 5 shows that Theorem 10 is not true if all players can play alone but simultanenous plays are also allowed. This example has 3 players and it is easy to find examples with any number of players larger than 3. However with two players, Theorem 11 says that if both players can play alone then simultaneous play will not jeopardize convergence to an optimal NE. Example 6 shows that Theorem 11 is not true for all separable protocols: if only one player among the two can play alone, then convergence to the optimal NE is not garanteed.

Finally, Example 7 shows that Assumption 1 cannot be relaxed in Theorem 9 by exhibiting a game with two players (where both can play alone) that admits stable states that are not NE.

Example 4 (**No convergence to NE for a non-separable revision law**). Let us consider a 2-player game with 2 actions each, $\mathcal{A} := \{a, b\}$. The support of the revision law is made of a single set $\{1, 2\}$ (both players always play together). This revision law is not separable.

The payoffs of both players coincide with the potential, given by the following matrix:

$F =$

$1 \backslash 2$	a	b
a	1	0.5
b	0	1

This game has two NE (a,a) and (b,b). The order of all the transition probabilities are given in Figure 1, computed using (13).

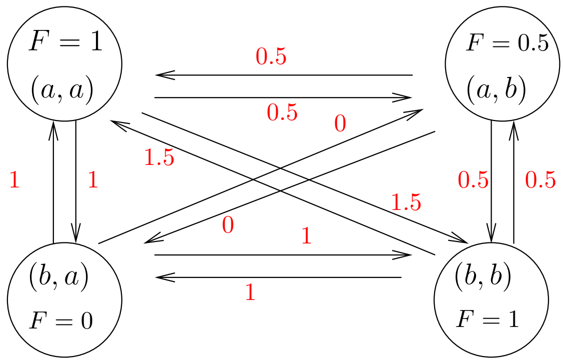

Figure 1. The orders of all the arcs in the transition graph of a 2-player potential game, when the revision process always makes both players play simultaneously

The minimal tree with root in (a,a), $T^*_{(a,a)}$ can be computed from Figure 1. Its order is $q_{(a,a)} = 1$. The minimal tree with root in (b,b), $T^*_{(b,b)}$, can also be found from Figure 1, with order $q_{(b,b)} = 1$. The minimal trees for the profiles (a,b) and (b,a) are both of order 1 as well: $q_{(a,b)} = 1$ and $q_{(b,b)} = 1$. Therefore, the stochastically stable state of SmoothBRA(ρ, Q) are all the states, (NE as well as non-NE). The algorithm will visit non-NE states with probabilities that do not vanish when the parameter θ becomes large. Actually, one can also compute the exact stationary distribution for all θ: $\pi((a,a),(a,b),(b,a),(b,b))$ is proportional to:

$$(s^4+s^3+es^2+es,\, s^4+2s^3+s^2,\, s^4+2es^2+e^2,\, s^4+s^3+es^2+es),$$

where $s \stackrel{\text{def}}{=} e^{\theta/2}$. Its limit when the parameter θ goes to infinity is $(1/4, 1/4, 1/4, 1/4)$ so all states are uniformly selected regardless of their potential. Even more surprisingly, notice that when θ is larger than 2, the state with the largest probability is (a,b), which is not a NE.

Example 5 (**Convergence to non-optimal NE when players can play alone**). Let us consider a 3-player game $\mathcal{N} = \{x,y,z\}$ with 2 actions each, $\mathcal{A} := \{0,1\}$. The support of the revision law is the separable family $\{x\},\{y\},\{z\},\{x,y,z\}$.

The payoffs of the players coincide with the potentials, given by: $F(0,0,0) = 3$, $F(1,0,0) = 0$, $F(0,1,0) = 1$, $F(0,0,1) = 2$, $F(1,1,0) = -10$, $F(0,1,1) = -11$, $F(1,0,1) = -12$, $F(1,1,1) = -1$.

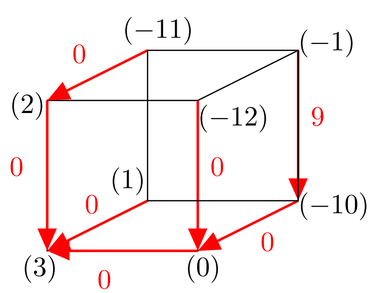

Figure 2. The minimal in-tree rooted in $(0,0,0)$ has order $q_{(0,0,0)} = 9$ (the potentials are given in parenthesis and the orders of the transitions are in red).

This game has two NE $(0,0,0)$ and $(1,1,1)$. The optimal one, with maximal potential is $(0,0,0)$. The minimal tree with root in $(0,0,0)$, $T^*_{(0,0,0)}$, is of order $q_{(0,0,0)} = 9$ (see Figure 2) and does not use the revision set $\{x,y,z\}$.

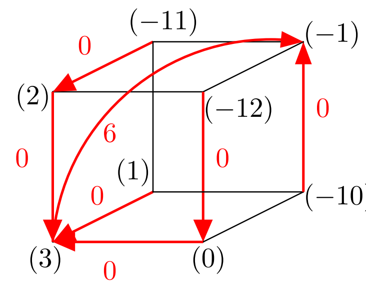

Figure 3. The minimal in-tree rooted in $(1,1,1)$ has order $q_{(1,1,1)} = 6$

The minimal tree with root in $(1,1,1)$, $T^*_{(1,1,1)}$, is of order $q_{(1,1,1)} = 6$ (see Figure 3) and uses the revision set $\{x,y,z\}$ to jump from $(0,0,0)$ to $(1,1,1)$ with order $1+2+3=6$, so it is minimal.

Therefore, the only stochastically stable profile of SmoothBRA(ρ, Q) will be $(1,1,1)$, whose potential is not optimal.

As mentioned before, convergence to optimal NE can be proved in the case of two players who can play alone:

Theorem 11 (Convergence to optimal NE with two players). *Let $\mathfrak{G}$ be a two-players game with exact potential F. If the support of the revision process is $\{\{1\},\{2\},\{1,2\}\}$ then the only stochastically stable states are optimal NE.*

Proof: We will first prove that there exists a tree with minimal order ending in each NE whose arcs only use singletons as revision sets.

Let us notice that, from any state that is not a NE, there is an outgoing arc of minimal order (0), that uses a unilateral deviation. In the case where there are only two NE x^* and y^* (the general case being treated similarly by induction), the non-NE states are all covered by two sets S_1 and S_2: those connected to x^* (resp. y^*) with paths of order 0 that involves single players (some states may belong to both sets).

Now, to get the minimal tree rooted in x^*, only one path is missing, from y^* to one state in S_1. Let us assume that this path contains one diagonal arc (that involves two players), from state u to state v. With no loss of generality, the next vertex after v in this path, denoted w, shares the same action for player 1: $v_1 = w_1$.

The order of the arc from u to v is $F(u_1, \mathrm{BR}_2(u_1)) - F(u_1, v_2) + F(\mathrm{BR}_1(u_2), u_2) - F(v_1, u_2)$, and the order of the arc from v to w is $F(v_1, \mathrm{BR}_2(v_1)) - F(w)$.

Let us replace the path $u \to v \to w$ by the path from $u \to (v_1, u_2) \to w$ that does not contain any diagonal arc. The order of the arc from u to (v_1,u_2) is $F(\mathrm{BR}_1(u_2), u_2) - F(v_1, u_2)$, and that of the arc from (v_1, u_2) to w is $F(v_1, \mathrm{BR}_2(v_1)) - F(w)$. Summing both values gives a path with smallest order.

As for node v, which is no longer on the path, it is in set S_2, so that it belongs to a path with order 0 to y^*. As for the order for all other states, it is unchanged.

In total, the new tree has a smaller order as the previous one, so that all diagonal arcs can be removed on the minimal tree.

Therefore, there exists a minimal tree rooted in x^* that only uses single revision sets. Now, Theorem 10 says that a minimal tree with single revision sets rooted in an optimal NE has the smallest order. ■

The condition on the revision process that all players can play alone, used in Theorem 11, is stronger than separability. Example 6 shows that separability is not enough to guarantee convergence to optimal NE, even with two players.

Example 6 (**Convergence to non-optimal NE for two separable players**). Let us consider the following 2-player game with respective actions $\mathcal{A}_1 = \{a, b\}$ and $\mathcal{A}_2 = \{a, b, c\}$. The support of the revision law is $\{\{2\}, \{1, 2\}\}$, hence it is separable.

The payoffs of both players coincide with the potential, given by the following matrix:

$$F = \begin{array}{|c||c|c|c|} \hline 1\backslash 2 & a & b & c \\ \hline a & 11 & 0 & 5 \\ b & 5 & 10 & 8 \\ \hline \end{array}$$

This game has two NE (a, a) and (b, b). The minimal tree with root in (a, a) has order 7. Indeed, the path with smallest order to join (b, b) to (a, a) is $(b, b) \to (b, c) \to (a, c) \to (a, a)$ whose order is $2 + 5 + 0$. All other states can be added to this path with the order 0 thanks to unilateral best response of y. The minimal tree with root in (b, b) has order 6. Indeed, the path with the minimal order to join (a, a) to (b, b) is $(a, a) \to (b, a) \to (b, b)$ whose order is $6 + 0$. All other states can be added to this path with order 0 thanks to unilateral best response of y.

Finally, in the same scenario as in Theorem 11 (two players and a revision process that contains both singletons), even worse things can occur when NE are not strict. Non strict NE can only exist when the best response is not unique. This possibility has been discarded up to now in this paper (Assumption 1). In the next example, we consider a case where the NE are not strict so that Assumption 1 is violated.

Example 7 (**No convergence to NE with two separable players if NE are not strict**). Consider a separable revision process, where both players can play alone, with support $\{1\}, \{2\}$ and $\{1, 2\}$ and a game with two actions per player given by the potentials (payoffs are equal to potentials):

$$F = \begin{array}{|c||c|c|} \hline 1\backslash 2 & a & b \\ \hline a & 1 & 1 \\ b & 1 & 0 \\ \hline \end{array}$$

States $(a, a), (a, b), (b, a)$ are non-strict NE and all states (including (b, b)) have a minimal in-tree of order 0, so even when the temperature $1/\theta$ goes to 0, the non-NE state (b, b) has a non-vanishing probability of being chosen. More precisely, by computing the exact stationary distribution π for all temperatures, one can check that when θ goes to infinity, $\pi((a, a), (a, b), (b, a), (b, b)) \to (36/79, 20/79, 20/79, 3/79)$ if the revision sets $\{1\}, \{2\}$ and $\{1, 2\}$ are chosen uniformly.

REFERENCES

[1] R. W. Rosenthal, "A class of games possessing pure-strategy nash equilibria," *Int. J. of Game Theory, Springer*, vol. 2, no. 1, pp. 65–67, 1973.

[2] M. Beckman, C. B. McGuire, and C. B.Winsten, *Studies in the Economics of Transportation*. Yale University Press, 1956.

[3] A. Orda, R. Rom, and N. Shimkin, "Competitive routing in multeuser communication networks," *IEEE/ACM Trans. on Networking*, vol. 1, no. 5, pp. 510–521, 1993.

[4] R. G. Gallager, "A minimum delay routing algorithm using distributed computation," *IEEE Transactions on Communications*, vol. 25, no. 1, pp. 73–85, 1977.

[5] J. Wardrop, "Some theoretical aspects of road traffic research. part ii," *Proc. of the Institute of Civil Engineers*, vol. 1, pp. 325–378, 1954.

[6] T. Roughgarden, *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.

[7] D. Monderer and L. Shapley, "Potential games," *Games and economic behavior, Elsevier*, vol. 14, no. 1, pp. 124–143, 1996.

[8] M. Voorneveld, "Best-response potential games," *Economics letters*, vol. 66, no. 3, pp. 289–295, 2000.

[9] J. S. Liu, "The collapsed gibbs sampler in bayesian computations with applications to a gene regulation problem," *Journal of the American Statistical Association*, vol. 89, no. 427, pp. pp. 958–966, 1994.

[10] F. T. Leighton and R. L. Rivest, *The Markov chain tree theorem*. Massachusetts Institute of Technology, 1983.

[11] C. Alos-Ferrer and N. Netzer, "The logit-response dynamics," *Games and Economic Behavior*, vol. 68, no. 2, pp. 413–427, 2010.

[12] W. H. Sandholm, *Population Games and Evolutionary Dynamics*. MIT Press, 2010.