Int J Game Theory
DOI 10.1007/s00182-010-0266-2

ORIGINAL PAPER

Adaptive learning and p-best response sets

J. Durieu · P. Solal · O. Tercieux

Accepted: 22 November 2010

Abstract A product set of strategies is a p-best response set if for each agent it contains all best responses to any distribution placing at least probability p on his opponents' profiles belonging to the product set. A p-best response set is minimal if it does not properly contain another p-best response set. We study a perturbed joint fictitious play process with bounded memory and sample and a perturbed independent fictitious play process as in Young (Econometrica 61:57–84, 1993). We show that in n-person games only strategies contained in the unique minimal p-best response set can be selected in the long run by both types of processes provided that the rate of perturbations and p are sufficiently low. For each process, an explicit bound of p is given and we analyze how this critical value evolves when n increases. Our results are robust to the degree of incompleteness of sampling relative to memory.

Keywords Evolutionary game theory · Fictitious play process · p-Dominance · Stochastic stability

JEL Classification C72 · C73

J. Durieu
University of Grenoble 2, Grenoble, France
e-mail: Jacques.Durieu@upmf-grenoble.fr

P. Solal
University of Saint-Etienne, CNRS UMR 5824 GATE, Lyon-St-Etienne, France
e-mail: solal@univ-st-etienne.fr

O. Tercieux (✉)
Paris School of Economics and CNRS, Paris, France
e-mail: tercieux@pse.ens.fr

Published online: 09 January 2011

1 Introduction

In this paper, we use the concept of minimal p-best response, introduced by Tercieux (2006a, b), to formulate predictions on the strategies selected in the long run by adaptive processes based on a fictitious play process. These predictions hold for the class of finite n-person games and for a large class of fictitious play processes with bounded memory and sample. We consider fictitious play processes in which agents believe that the play of their opponents is either independent or correlated. This is in contrast with most of the literature which deals with independent fictitious play processes. Moreover, we impose no restriction on the relative size of the sample.

Young (1993) considers a fictitious play process with a bounded memory of size m and a sample of size k. He focuses his attention on the case of an incomplete sample: the ratio k/m has to be sufficiently small. When k and m are sufficiently large, Young (1993) shows that in a 2×2 game with two symmetric equilibria, the risk dominant equilibrium by Harsanyi and Selten (1988) corresponds to the unique stochastically stable state of such a process. Maruta (1997) exploits the concept of strict $1/2$-dominant equilibrium which is a generalization of the notion of risk dominance. He focuses on the class of finite two-person games that are weakly acyclic. Maruta (1997) establishes that when a strict $1/2$-dominant equilibrium exists, it corresponds to the unique stochastically stable state of Young's process provided k/m is sufficiently small and k and m are sufficently large. Young (1998) considers a class of generic n-person games which contains games that have a cyclic best-response structure. He extends the definition of the adaptive process to n-person games by considering an independent fictitious play process: each agent believes that each of his opponents plays according to a fixed mixed strategy and that these strategies are independent among players. In the sequel, we refer to such beliefs as independent. In order to formulate results, Young (1998) makes use of the concept of curb set introduced by Basu and Weibull (1991).[1] He shows that if k/m is sufficiently small and k and m are sufficiently large, then the stochastically stable states of the adaptive process correspond to the minimal curb sets minimizing stochastic potential. However, in order to identify such a minimal curb set, it is necessary to compute the minimal stochastic potential of all minimal curb sets. Moreover, a restriction to generic n-person games is needed in order to have no recurrent set of the adaptive process lying outside a minimal curb set.

These results suggest the following questions. Given a fictitious play process with bounded memory and sample, can we formulate selection results for the class of finite n-person games? What are the differences in terms of selection results between an independent fictitious play process and a joint fictitious play process—where each agent believes that the play of his opponents is correlated? In this paper we investigate two possible theories: agents believe either that their opponents act independently or that a correlated behavior is allowed among their opponents. However, we keep the standard assumption that mistakes are independent among agents.

Firstly, we consider the class of adaptive processes based on a joint fictitious play process: each agent estimates the stationary correlated strategy supposedly used by his

[1] Hurkens (1995) also considers an adaptive process with bounded memory and shows that the process converges to a minimal curb set. However, he does not address the question of stochastic stability.

opponents with a sample drawn in the recent history and plays a best response against this belief. We focus on finite n-person games for $n > 2$. Then, the joint fictitious play process is distinct from the independent fictitious play process.[2] We show that the stochastically stable states of a joint fictitious play process with bounded memory m and sample k form a subset of the recurrent sets of the process that are associated with the minimal p-best response set of the game for $p \leq (n+1)^{-1}$ provided k and m are sufficiently large. As a corollary, if the game admits a strict p-dominant equilibrium for $p \leq (n+1)^{-1}$, then the unique stochastically stable state corresponds to this equilibrium. Thus, the critical value of p is of order $1/n$.[3] Another interesting feature of this result is that it holds for any relative size of the sample.

Secondly, we consider the class of independent fictitious play processes with bounded memory m and sample k which contains the model used by Young (1998). We consider the class of finite n-person games where $n \geq 2$. We show that the strategies played in the long run are contained in the minimal p-best response set for $p \leq n^{1-n}$ provided k and m are sufficiently large. Hence it appears that the assumption on agents' beliefs is crucial for the strength of our selection results. Notice that when $n = 2$, the critical value of p is $1/2$. This means that we extend Maruta's result in two ways. Firstly, if the game has a strict $1/2$-dominant equilibrium, then the state corresponding to this equilibrium is the unique stochastically stable state of the process. Thus, the restriction to weakly acyclic games is not necessary. Secondly, if the game does not admit a strict $1/2$-dominant equilibrium, then stochastically stable states form a subset of the recurrent sets of the process that are associated with the minimal $1/2$-best response set of the game. Here again, our results hold for any relative size of the sample.

The concept of minimal p-best response set permits us to establish results concerning stochastic stability for both types of processes without exhibiting the complete list of recurrent sets of each adaptive process. To understand this point, note that the minimal p-best response set is unique provided $p \leq 1/2$ and that at least one recurrent set of each adaptive process is associated with this set. Moreover, using Ellison (2000)'s terminology, our proof only requires us to show that a lower bound of the radius of the recurrent sets associated with the minimal p-best response set is superior to an upper bound of the coradius of these recurrent sets.[4] As a consequence, contrary to Young (1998), it is unnecessary to focus on a class of generic n-person games.

The paper is organized as follows. Section 2 contains notations and preliminaries. In Sect. 3, minimal p-best response sets are formally defined and some of their properties are established. In Sect. 4, we define adaptive processes based on a fictitious play process. Section 5 states selection results for these processes. Section 6 concludes.

[2] Since in two-person games the joint fictitious play process coincides with the independent fictitious play process, the case $n = 2$ is covered by our results for the independent fictitious play process.

[3] Note that this evolution of p with the number of players is usual in the results of equilibrium selection. See for instance Kajii and Morris (1997) for the robustness to incomplete information approach or Oyama et al. (2008) for selections based on perfect foresight dynamics. However, the critical value found in the aforementioned papers—while of the same order—is n^{-1} and not $(n+1)^{-1}$ as in the present paper. We checked, by means of an example, that under the tighter bound of n^{-1}, our result would not hold.

[4] Since the collection of recurrent sets of each process is not necessarily known, it is not possible to apply the concept of modified coradius, that is, to take into account a step-by-step evolution.

2 Notations and definitions

Let $\subseteq$ denote weak set inclusion and $\subset$ denote proper set inclusion.

Let Γ be a finite n-person strategic-form game. Let X_i be the finite set of pure strategies x_i available to player $i \in I = \{1, 2, \ldots, n\}$. For any finite set A, $\Delta(A)$ denotes the set of all probability distributions on A. We write $\Delta(X_i)$ for the set of probability distributions q_i over X_i for each $i \in I$. Let $q_i(x_i)$ denote the probability mass on strategy x_i. Define the product set $X = \prod_{i\in I} X_i$. Let $\Delta(X)$ be the set of probability distributions on X. Let $X_{-i} = \prod_{j\neq i} X_j$ denote the set of all possible combinations of strategies for the players other than i, with generic elements $x_{-i} = (x_j)_{j\neq i}$. Let $\Delta(X_{-i})$ be the set of probability distributions on X_{-i} with generic elements q_{-i}. We sometimes identify the element of $\Delta(X_i)$ that assigns probability one to a strategy in X_i with this strategy in X_i.

In this paper, a player's belief about others' strategies takes the form of a probability measure on the product of all opponents' strategy sets. We assume that each player i has expected payoffs represented by the function $u_i : X_i \times \Delta(X_{-i}) \rightarrow \mathbb{R}$.

For each player i and probability distribution $q_{-i} \in \Delta(X_{-i})$, let

$$Br_i(q_{-i}) = \{x_i \in X_i : u_i(x_i, q_{-i}) \geq u_i(x_i', q_{-i}), \forall x_i' \in X_i\}$$

be the set of pure best responses of i against q_{-i}.

Let $Y = \prod_{i\in I} Y_i$ be a product set where each Y_i is a nonempty subset of X_i. Let $\Delta(Y_{-i})$ denote the set of probability distributions on X_{-i} with support in Y_{-i}. Finally, $Br_i(\Delta(Y_{-i}))$ denotes the set of strategies in X_i that are pure best responses of i against some distribution q_{-i} with support in Y_{-i}, that is

$$Br_i(\Delta(Y_{-i})) = \bigcup_{q_{-i}\in\Delta(Y_{-i})} Br_i(q_{-i}).$$

3 *p*-Best response sets

We will now introduce the concept of (minimal) p-best response set. Let $p \in [0, 1]$, and let Y_{-i} be a nonempty subset of X_{-i}. We write $\Delta(Y_{-i}, p) \subseteq \Delta(X_{-i})$ for the subset of distributions $q_{-i} \in \Delta(X_{-i})$ such that $\sum_{x_{-i}\in Y_{-i}} q_{-i}(x_{-i}) \geq p$. Let $Br_i(\Delta(Y_{-i}, p))$ denote the set of strategies in X_i that are pure best responses by i to some distribution $q_{-i} \in \Delta(Y_{-i}, p)$ (regardless of probability assigned to other possible combinations of strategies), that is

$$Br_i(\Delta(Y_{-i}, p)) = \bigcup_{q_{-i}\in\Delta(Y_{-i},p)} Br_i(q_{-i}).$$

Let us recall the definition of a strict p-dominant equilibrium first introduced by Morris et al. (1995) in two-person games and extended to n-person games by Kajii and Morris (1997). A profile $x^* \in X$ is a strict $\mathbf{p} = (p_1, \ldots, p_n)$-dominant equilibrium if for each player $i \in I$

$$\{x_i^*\} = Br_i(\Delta(x_{-i}^*, p_i)).$$

In the sequel, we focus on the case where $p_i = p$ for all $i \in I$. The concept of p-best response set extends the concept of strict p-dominant equilibrium to product sets of strategies. Formally, a (minimal) p-best response set is defined as follows.

Definition 1 Let $p \in [0, 1]$. A *p-best response set* is a product set $Y \subseteq X$ where for each player i

$$Br_i\,(\Delta(Y_{-i}, p)) \subseteq Y_i.$$

A p-best response set Y is a *minimal p-best response set* if no p-best response set is a proper subset of Y.

Let $\mathcal{Q}_p$ be the collection of p-best response sets for some $p \in [0, 1]$. The following lemma states some properties of minimal p-best response sets.

Lemma 1 *Let Γ be a finite n-person game.*

1. *Γ has a minimal p-best response set for any $p \in [0, 1]$.*
2. *Fix $p \in [0, 1]$. Then two distinct minimal p-best response sets of Γ are disjoint.*
3. *For $p \leq 1/2$, there exists a unique minimal p-best response set in Γ.*
4. *Let $p \leq p' \leq 1/2$. Let Y and Y' be the minimal p-best response set and p'-best response set, respectively. Then, $Y \supseteq Y'$.*

Proof Proofs of points 1 and 2 are similar to Theorem 1 in Tercieux (2006a).

3. By contradiction, assume that Y and Y' are two minimal p-best response sets for $p \leq 1/2$. By point 2, we know that $Y \cap Y' = \emptyset$. Then there exists at least one player i such that $Y_i \cap Y_i' = \emptyset$. Consider this player i and pick $q_{-i} \in \Delta(Y_{-i}, 1/2) \cap \Delta(Y_{-i}', 1/2) \subseteq \Delta(Y_{-i}, p) \cap \Delta(Y_{-i}', p)$. Because both Y and Y' are two p-best response sets we must have that $Br_i(q_{-i}) \subseteq Y_i \cap Y_i'$. But this contradicts the fact that $Y_i \cap Y_i' = \emptyset$. Observe that, for $p > 1/2$, $\Delta(Y_{-i}, p) \cap \Delta(Y_{-i}', p) = \emptyset$ may occur.

4. By contradiction, assume $Y \not\supseteq Y'$. Observe that since Y is a mimimal p-best response set it is also a p'-best response set. Thus, there exists $\underline{Y} \subseteq Y$ such that $\underline{Y}$ is a minimal p'-best response set. Clearly, $\underline{Y} \neq Y'$, hence this contradicts our uniqueness result (point 3.). □

4 Adaptive processes

As noted by Monderer and Sela (1997), there exist two natural ways to define a fictitious play process in finite n-person games. The first way considers situations where agents believe that each of his opponents is using a stationary mixed strategy. A variant of this process has been studied for example by Young (1993, 1998) and Maruta (1997). The second approach is the joint fictitious play process in which each agent believes that his opponents are using a correlated strategy. We refer to such beliefs as correlated. Since adaptive processes based on an independent fictitious play process are defined in Young (1993, 1998), we only describe in detail adaptive processes based on a joint fictitious play process.

Let $t = 1, 2, \ldots$ denote successive time periods. Define $x^t = (x_1^t, \ldots, x_n^t)$ as the strategy-tuple played at time t. Fix integers k and m such that $k \leq m$. Assume for the sake of generality that the first m profiles are randomly selected. At the beginning of any period $t > m$, each agent inspects k strategy-tuples drawn without replacement from the most recent m periods. Each agent i estimates the correlated strategy supposed to be used by his opponents with the empirical distribution of his opponents' profiles in the sample. In this context, $q_{-i}^J \in \Delta(X_{-i})$ represents the empirical distribution on i's opponents profiles built from the sample. For each $x_{-i} \in X_{-i}$, $q_{-i}^J(x_{-i})$ is a component of q_{-i}^J. In determining his strategy, i chooses a pure best response to q_{-i}^J. We obtain a joint fictitious play process with bounded memory and sample. This process, denoted P_J^0, defines a finite Markov chain on the set of truncated histories $H = X^m$. For any subset $H' \subseteq H$, define $S_i(H')$ as the set of strategies available to i that appear in H'. The span of H', denoted by $S(H')$, is the product set of all strategies that appear in H', that is $S(H') = \prod_{i \in I} S_i(H')$. Let $\mathcal{R}^J$ be the collection of recurrent sets of P_J^0.

In addition, independently across time and players, with a (small) probability $\epsilon > 0$, a player will instead make a "mistake" i.e., he will choose a strategy at random (all strategies having positive probability of being chosen). With these mistakes as part of the process, each state of the system is reachable with positive probability from every other state. Hence, the full process, denoted P_J^ϵ and called the perturbed joint fictitious play process with bounded memory and sample, is an irreducible and aperiodic finite state Markov chain on H. Consequently, for each $\epsilon > 0$, P_J^ϵ has a unique stationary distribution μ_J^ϵ satisfying $\mu_J^\epsilon P_J^\epsilon = \mu_J^\epsilon$. It is well known that the limit stationary distribution $\mu_J^* = \lim_{\epsilon \to 0} \mu_J^\epsilon$ exists (see Young 1993) and that states which have a positive probability in μ_J^* are called stochastically stable states and form a subset of $\mathcal{R}^J$. The recurrent sets appearing in the support of μ_J^* are those which are the easiest to reach from all other recurrent sets, with "easiest" interpreted as requiring the fewest mistakes. Formally, for each pair of states $h^1, h^2 \in H$, a h^1h^2-path is a sequence of states that begins at h^1 and ends at h^2. The resistance of a path is the sum of the numbers of mistakes on the edges that compose it. Define the resistance $r_{h^1h^2}$ of the transition from h^1 to h^2 as the least resistance over all h^1h^2-paths. Since only recurrent sets of P_J^0 may appear in the support of μ_J^*, we can restrict our attention to transitions between recurrent sets. Construct a complete directed graph with one vertex for each recurrent set. For two recurrent sets $H^1 \subseteq H$ and $H^2 \subseteq H$ of P_J^0, the weight on the edge from H^1 to H^2 is the resistance $r_{H^1H^2} = \min_{h^1 \in H^1} \min_{h^2 \in H^2} r_{h^1h^2}$. A tree rooted at H^1 is a collection of directed edges such that from every vertex except H^1 there is a unique path in the tree to H^1. The resistance of a rooted tree is the sum of the resistances on the edges that compose it. The *stochastic potential* $\rho(H^1)$ of H^1 is the minimum resistance over all trees rooted at H^1. Finally, following Theorem 4 in Young (1993), the stochastically stable sets are characterized as the recurrent sets with the minimum stochastic potential. Each state belonging to such a recurrent set is stochastically stable.

Similarly, it is possible to define an adaptive process based on an independent fictitious play process, where each agent assumes that every other agent is choosing a strategy according to a probability distribution and that these distributions are independent among agents. For each $i \in I$, let $\Delta^I(X_{-i})$ denote the set of probability distributions on X_{-i} constructed as a product of probability distributions on X_j, $j \neq i$. At each period, each agent i forms belief $q^I_{-i} \in \Delta^I(X_{-i})$. Following Young (1993, 1998), for m and k fixed and such that $k \leq m$, we consider the independent fictitious play process with bounded memory m and sample k. We denote this process by P^0_I and by $\mathcal{R}^I$ the collection of recurrent sets of P^0_I. When mistakes are allowed, we obtain a perturbed process denoted P^ϵ_I. Let μ^*_I be the limit stationary distribution of the independent fictitious play process with bounded memory m and sample k.

5 Selection results

Fix a minimal p-best response set Y. Consider an adaptive process based on a joint fictitious play process. Define $\mathcal{R}^J_Y \subseteq \mathcal{R}^J$ as the collection of recurrent sets of P^0_J such that $H' \in \mathcal{R}^J_Y$ if and only if $S(H') \subseteq Y$. Similarly, for an adaptive process based on an independent fictitious play process, let $\mathcal{R}^I_Y \subseteq \mathcal{R}^I$ be the collection of recurrent sets of P^0_I such that $H' \in \mathcal{R}^I_Y$ if and only if $S(H') \subseteq Y$. Observe that $\mathcal{R}^J_Y \neq \emptyset$ and $\mathcal{R}^I_Y \neq \emptyset$ since Y is a minimal p-best response set and thus a closed set under the best-response rule.

The following Theorem allows us to identify a product set of strategies which contains the stochastically stable states relative to adaptive processes based on a fictitious play process.

Theorem 1 1. *Let Γ be a finite n-person game with $n > 2$. Let $p \leq (n+1)^{-1}$. Let Y be the minimal p-best response set of Γ. There exists $\bar{k}$ such that for $m \geq k \geq \bar{k}$, it holds that $\mu^*_J(H^*) = 1$ for some $H^* \subseteq H$ such that $S(H^*) \subseteq Y$.*
2. *Let Γ be a finite n-person game with $n \geq 2$. Let $p \leq n^{1-n}$. Let Y be the minimal p-best response set of Γ. There exists $\bar{k}$ such that for $m \geq k \geq \bar{k}$, it holds that $\mu^*_I(H^*) = 1$ for some $H^* \subseteq H$ such that $S(H^*) \subseteq Y$.*

Proof See Appendix A. □

As stated earlier, if the minimal p-best response set Y is a singleton, it corresponds to a strict p-dominant equilibrium in the sense of Kajii and Morris (1997). From this observation, we obtain the following corollary.

Corollary 1 1. *Suppose that Γ is a finite n-person game with $n > 2$ and that Γ admits a strict p-dominant equilibrium $x^* \in X$ for $p \leq (n+1)^{-1}$. There exists $\bar{k}$ such that for $m \geq k \geq \bar{k}$, it holds that $\mu^*_J(H^*) = 1$ for $H^* \subseteq H$ such that $S(H^*) = \{x^*\}$.*
2. *Suppose that Γ is a finite n-person game with $n \geq 2$ and that Γ admits a strict p-dominant equilibrium $x^* \in X$ for $p \leq n^{1-n}$. There exists $\bar{k}$ such that for $m \geq k \geq \bar{k}$, it holds that $\mu^*_I(H^*) = 1$ for $H^* \subseteq H$ such that $S(H^*) = \{x^*\}$.*

The above results suggest four comments. First, we note that for the second part of the Theorem, we could have used a weaker definition of the concept of p-best response set where $\Delta(Y_{-i}, p) \subseteq \Delta(X_{-i})$ would have been defined as the subset of distributions $q_{-i} \in \Delta^I(X_{-i})$ such that $\sum_{x_{-i} \in Y_{-i}} q_{-i}(x_{-i}) \geq p$. It is clear that this definition of an "independent p-best response set" relying only on independent beliefs would have been enough to prove our selection result for the independent fictitious play process: the strategies played in the long run are contained in the minimal independent p-best response set for $p \leq n^{1-n}$. Since, for a given $p \leq 1/2$, the minimal independent p-best response set is a subset of the minimal p-best response set, this result is sharper than point 2 in Theorem 1.

Second, it is worthwhile to notice that the above results hold without assuming that the sample is sufficiently incomplete compared to the memory. In fact, it is unnecessary to introduce a sampling procedure. To understand this point, consider a transition from a recurrent set not associated with the minimal p-best response set to a recurrent set associated with Y. In a joint fictitious play process with bounded memory and sample, such a transition is accomplished when all agents make a mistake simultaneously. As a consequence, at any period, observation of heterogeneous periods in the memory among the agents is not necessary for such a transition: it is sufficient that each agent observes the more recent periods. By contrast, in an independent fictitious play process with bounded memory and sample, that transition requires only that $n-1$ agents make a mistake simultaneously. Nevertheless, a sampling procedure is not needed. This is due to the focus on direct transitions between recurrent sets and to the relative smallness of the critical value of p. Indeed, for every admissible value of p, the number of periods in which mistakes occur constitutes a small proportion (i.e., less than half) of the length of the memory. On a related topic, note that for each process and for a given number of agents, as the minimal value p of the p-best response set becomes closer and closer to the critical value exhibited in Theorem 1, the minimal size of the sample (or of the memory) $\bar{k}$ must be larger and larger.

Third, another comment concerns the assumption on agents' beliefs. Theorem 1 reveals that this assumption is crucial for the use of the concept of minimal p-best response set. Whereas the greatest admissible value for p evolves as $(1/n)^{n-1}$ in point 1, it is of order $1/n$ in point 2. To have an intuition of this gap, note that the impact of simultaneous mistakes is "stronger" with correlated beliefs than with independent beliefs. As a consequence, for p, k and n fixed, it is easier to reach a state associated with the minimal p-best response set in an adaptive process based on a joint fictitious play process than in an adaptive process based on an independent fictitious play process. The following example illustrates this point.

Example 1 Consider the three-person game Γ^1 represented by the payoff matrices

I

	A	*B*
1	2, 2, 2	2, 0, 2
2	0, 2, 2	0, 0, 0

II

	A	*B*
1	2, 2, 0	0, 0, 0
2	0, 0, 0	9, 9, 9

Assume that agent 1 chooses row, agent 2 column, and agent 3 matrix. Observe that Γ^1 is weakly acyclic and possesses two strict Nash equilibria: $(1, A, I)$ and $(2, B, II)$. Moreover, the minimal 1/5-best response set Y of Γ^1 is $\{(2, B, II)\}$. By point 1 in Theorem 1, the state corresponding to $(2, B, II)$ is the unique stochastically stable relative to a joint fictitious play process with bounded memory and sample. Consider now an adaptive process based on an independent fictitious play process with bounded memory m and sample k where $k/m \leq 1/3$, as in Young (1993) and Maruta (1997). We claim that the unique stochastically stable state of this process does not correspond to $(2, B, II)$. Indeed, a straightforward computation shows that a transition from the state corresponding to $(2, B, II)$ to the state corresponding to $(1, A, I)$ can be done with $\lceil \frac{9}{11} k \rceil$ mistakes. A similar computation shows that a transition from the state corresponding to $(1, A, I)$ to the state corresponding to $(2, B, II)$ can be done with $2\lceil \frac{\sqrt{2}}{\sqrt{11}} k \rceil$ mistakes. Finally, observe that $2\lceil \frac{\sqrt{2}}{\sqrt{11}} k \rceil > \lceil \frac{9}{11} k \rceil$ provided k is sufficiently large, that is, when $k > 29$. Thus, the state corresponding to $(1, A, I)$ is the unique stochastically stable state relative to each independent fictitious play process with bounded memory m and sample k where $k/m \leq 1/3$ and $k > 29$.

This example allows us to draw two conclusions. First, the restriction $p \leq (n+1)^{-1}$ is not sufficient to have stochastically stable states associated with the strategies contained in the minimal p-best response set for each independent fictitous play process with bounded memory and sample. Second, an independent fictitious play process with bounded memory and sample may select, in the long run, a state such that its span does not coincide with a minimal p-best response set for $p < 1/2$, even if it is strictly contained in a minimal p-best response set for $p \leq n^{1-n}$. □

Finally, one can wonder about the predictive value of our results. Firstly, observe that for a given process and a given p, a minimal p-best response set is associated with at least one recurrent set of the unperturbed process. Consider the span of this recurrent set, say $Z \subseteq X$. For each agent i, a strategy $x_i \in Z_i$ is such that there exists a probability distribution $q_{-i} \in \Delta(Z_{-i})$ ($q^I_{-i} \in \Delta^I(Z_{-i})$ resp.) such that x_i is a best response to q_{-i} (q^I_{-i} resp.). Hence, predictions given by our Theorem will be a subset of the set of correlated (independent resp.) rationalizable strategy profiles. Secondly, Theorem 1 establishes that when the number of players increases, the critical value p becomes smaller and smaller in both processes. As already stated, the decrease in the critical value for the joint fictitious play is slower than for the independent fictitious play. This is due to the fact that the impact of mistakes is "stronger" with correlated beliefs than with independent beliefs. To illustrate that correlation is a useful device to get positive results, we could go even further and consider a process identical in all respect to the joint fictitious play process except that agents would be able to correlate their mistakes.[5] At an intuitive level, allowing agents to make mistakes in a correlated way makes the "mistake counting argument" close to the 2-player case. Indeed, one can show a version of our Theorem for this modified process where the critical value on p would be 1/2 (i.e., the threshold found in the 2-player case). Here the critical

[5] More precisely, in each time period, there would be some small probability $\epsilon > 0$ that all agents make a mistake by choosing a strategy profile in X and this would be counted as a single mistake. We thank an anonymous referee for this suggestion.

value would not depend on the number of players. In any case, for this process or the processes under study in the paper, several recurrent sets may be associated with the minimal p-best response set. Then, Theorem 1 does not allow a complete identification of the stochastically stable states. However, in many examples, as it excludes states that do not belong to the minimal p-best response set, less resistance trees have to be calculated. In other words, the knowledge of a minimal p-best response is often useful.[6] This fact is illustrated in the following example.

Example 2 Consider the two-person game Γ^2 borrowed from Young (1993) and represented by the following payoff matrix:

	A	B	C
1	6, 6	0, 5	0, 0
2	5, 0	7, 7	5, 5
3	0, 0	5, 5	8, 8

The unique minimal 1/2-best response set of Γ^2 is $\{2, 3\} \times \{B, C\}$. Consider Young's process. By point 2 in Theorem 1, the set of stochastically stable states H^* is such that $S(H^*) \subseteq \{2, 3\} \times \{B, C\}$. This identification of the stochastically stable states is partial since Young (1993) shows that $S(H^*) = \{(2, B)\}$. Nevertheless, Theorem 1 indicates that it is unnecessary to compute the minimum resistance of trees rooted at $\{(1, A)\}$.[7] □

6 Conclusion

The concept of minimal p-best response set is helpful for identifying the stochastically stable states of two classes of adaptive processes based either on a joint or an independent fictitious play process with bounded memory and sample. For each class of processes, we show that the critical value of p, helpful to identify the stochastically stable states, becomes smaller as n increases. This means that the predictive value of our results decreases when n increases. However, the paper also suggests that the way by which agents form beliefs has a crucial impact on the selection results. More severe restrictions have to be introduced on the concept of minimal p-best response set when an adaptive process based on an independent fictitious play is studied.

Minimal p-best response sets are relevant for many equilibrium selection methods. For instance, Tercieux (2006b) shows that any equilibrium that is robust to incomplete information in the sense of Kajii and Morris (1997) must be included in the minimal p-best response set where $p \leq 1/n$. Another instance is Matsui and Matsuyama (1995) perfect foresight dynamics setting. Using Tercieux (2006a) or Oyama et al.

[6] It is worthwhile to notice that, using an argument similar to Klimm et al. (2010), one can show that for a given p, deciding whether a product set of strategies is a p-best response set can be done in polynomial time.

[7] This example shows that our Theorem does not provide a tight characterization of stochastically stable states, i.e., there are strategy profiles non-associated with a stochastically stable state included in the minimal 1/2-best response set.

(2008) or Kojima and Takahashi (2008), it can be shown that any state that is globally accessible or absorbing (for small frictions) must be included in the minimal p-best response set where $p \leq 1/n$. Hence, minimal p-best response sets allow us to generalize previous results based on p-dominance. Interestingly, Morris and Ui (2005) have proposed the notion of generalized potential maximizer sets that is also a generalization of p-dominance. While this notion is very powerful for the aforementioned selection methods, one can check that their maximizer set need not be stochastically stable in Young's process.[8]

Acknowledgments We would like to thank Daijiro Okada, the editor and two anonymous referees for helpful comments. Support by the French Ministry for Youth, Education and Research, through project SCSHS 2004-04 is gratefully acknowledged. Earlier versions of this paper were circulated under the title "Adaptive learning and curb set selection".

A Appendix

Proof of Theorem 1 Let $\lceil e \rceil$ (resp. $\lfloor e \rfloor$) denote the least integer greater (resp. greatest integer smaller) than or equal to e for any real number e.

Firstly, we establish point 1. Note that if $\mathcal{R}^J_Y = \mathcal{R}^J$, then point 1 in Theorem 1 is trivial. In this sequel, we assume $\mathcal{R}^J_Y \subset \mathcal{R}^J$. The proof is divided into two parts. (1) We bound the resistances of transitions beginning at $H' \in \mathcal{R}^J_Y$ and the resistances of transitions ending at $H' \in \mathcal{R}^J_Y$. (2) We show that the recurrent set(s) minimizing the stochastic potential belongs to $\mathcal{R}^J_Y$.

(1) We give a lower bound on the resistances of the transitions that begin at $H' \in \mathcal{R}^J_Y$ and end in any recurrent set $H'' \notin \mathcal{R}^J_Y$. Observe that if agent $j \in I$ makes a mistake and plays outside Y_j, then each agent $i \in I \backslash \{j\}$ may hold beliefs q^J_{-i} such that $\sum_{x_{-i} \notin Y_{-i}} q^J_{-i}(x_{-i}) = 1/k$. And, if agent j makes a mistake and plays outside Y_j α times in succession with $0 < \alpha \leq k$, then each agent $i \in I \backslash \{j\}$ may hold beliefs q^J_{-i} such that $\sum_{x_{-i} \notin Y_{-i}} q^J_{-i}(x_{-i}) = \alpha/k$. Let $\bar{p}$ be the greatest probability such that at least one agent i has a strategy $x_i \notin Y_i$ as a pure best response to q^J_{-i} such that $\sum_{x_{-i} \in Y_{-i}} q^J_{-i}(x_{-i}) = \bar{p}$. By definition of a minimal p-best response set, we have $p > \bar{p}$. Thus, a transition from $H' \in \mathcal{R}^J_Y$ to any $H'' \notin \mathcal{R}^J_Y$ requires at least z mistakes (outside Y_j) where z is such that $z/k \geq 1 - \bar{p}$; otherwise, each agent $i \in I$ has only best response(s) in Y_i since $\sum_{x_{-i} \in Y_{-i}} q^J_{-i}(x_{-i}) > \bar{p}$. Therefore, for all recurrent sets $H' \in \mathcal{R}^J_Y$ and $H'' \notin \mathcal{R}^J_Y$, we necessarily have $r_{H'H''} \geq \lceil (1 - \bar{p})k \rceil$.
We now give an upper bound on the resistances of the transitions that begin at any recurrent set $H'' \notin \mathcal{R}^J_Y$ and end at some element in $\mathcal{R}^J_Y$. Note that, in any state $h'' \in H''$, if each agent $i \in I$ makes a mistake and chooses a strategy $x_i \in Y_i$, then each agent $i \in I$ may hold beliefs q^J_{-i} such that $\sum_{x_{-i} \in Y_{-i}} q^J_{-i}(x_{-i}) > 0$. In particular, if each agent $i \in I$ makes a mistake and plays inside Y_i at one period,

[8] For instance, by Oyama et al. (2008), we know that in Example 2, $\{(3, C)\}$ is such a maximizer while it is not stochastically stable.

we may have $\sum_{x_{-i}\in Y_{-i}} q^J_{-i}(x_{-i}) \geq 1/k$ for all $i \in I$. And, if each agent $i \in I$ makes a mistake and plays inside Y_i α times in succession with $0 < \alpha \leq k$, then we may have $\sum_{x_{-i}\in Y_{-i}} q^J_{-i}(x_{-i}) \geq \alpha/k$ for all $i \in I$. By definition of a minimal p-best response set, we know that any agent i has at least one strategy $x_i \in Y_i$ as a pure best response to q^J_{-i} such that $\sum_{x_{-i}\in Y_{-i}} q^J_{-i}(x_{-i}) \geq p$. Thus, from any state $h'' \in H''$, the process may move into an element in $\mathcal{R}^J_Y$ with $\alpha^* n$ mistakes where α^* is such that $\alpha^*/k \geq p$. This means that, for any recurrent set $H'' \notin \mathcal{R}^J_Y$, there exists $H' \in \mathcal{R}^J_Y$ such that the minimum number of mistakes sufficient to transit from H'' to H' is bounded above by $z = \lceil pk \rceil n$.
In order to verify that such a transition is possible for any $k \leq m$, consider the most unfavourable case where $k = m$. Note that z mistakes imply that each agent $i \in I$ observes a number of opponents' profiles in Y_{-i} sufficient to have a best response in Y_i. As a consequence, after some periods, each agent $i \in I$ only observes opponents' profiles in Y_{-i}. Thus, we have shown that for all $H'' \notin \mathcal{R}^J_Y$, there exists $H' \in \mathcal{R}^J_Y$ such that $r_{H''H'} \leq \lceil pk \rceil n$. In the following, such a H' is denoted by $\psi(H'')$. Thus, for all $H'' \notin \mathcal{R}^J_Y$, $r_{H''\psi(H'')} \leq \lceil pk \rceil n$.

(2) Suppose by contradiction that a recurrent set $H'' \notin \mathcal{R}^J_Y$ is stochastically stable. Denote by $T_{H''}$ (one of) the tree(s) rooted at H'' that minimizes resistance. Consider the path from $\psi(H'')$ to H'' in $T_{H''}$. This path contains a transition from $H' \in \mathcal{R}^J_Y$ to $\bar{H} \notin \mathcal{R}^J_Y$ (possibly with $H' = \psi(H'')$ and/or $\bar{H} = H''$). Delete this transition and add a transition from H'' to $\psi(H'')$. We obtain a tree $T_{H'}$ rooted at H'. By construction, the resistance of $T_{H'}$ is $r_{T_{H'}} = r_{T_{H''}} - r_{H'\bar{H}} + r_{H''\psi(H'')}$. Moreover, for k sufficiently large, we have $r_{H'\bar{H}} > r_{H''\psi(H'')}$, that is

$$\lceil (1-\bar{p})k \rceil > \lceil pk \rceil n. \tag{1}$$

To see this, note that a sufficient condition for (1) to hold is

$$1 - \bar{p} > pn + \frac{n}{k}. \tag{2}$$

And, (2) is true for k sufficiently large since, by definition of the concept of minimal p-best response set, $1-\bar{p} > 1-p$ and, by hypothesis, $p \leq (n+1)^{-1}$. Thus, $r_{T_{H'}} < r_{T_{H''}}$ and $\rho(H') < \rho(H'')$. This contradicts the fact that H'' is stochastically stable.

The proof of point 2 is similar to that of point 1. The main difference concerns the upper bound on the resistances of the transitions that begin at any recurrent set $H'' \notin \mathcal{R}^I_Y$ and end at some element in $\mathcal{R}^I_Y$. Note that, in any state $h'' \in H''$, there exists one agent $i \in I$ such that $\sum_{x_{-i}\in Y_{-i}} q^I_{-i}(x_{-i}) > 0$ only if each agent $j \neq i$ makes a mistake and chooses a strategy $x_j \in Y_j$. If each agent j in $I\backslash\{i\}$ makes a mistake and plays inside Y_j α times in succession with $0 < \alpha \leq k$, then we may have $\sum_{x_{-i}\in Y_{-i}} q^I_{-i}(x_{-i}) \geq (\alpha/k)^{n-1}$. Thus, from any state $h'' \in H''$, the process may move into an element in $\mathcal{R}^I_Y$ with $\alpha^*(n-1)$ mistakes where α^* is such that $(\alpha^*/k)^{n-1} \geq p$. This means that, for any recurrent set $H'' \notin \mathcal{R}^J_Y$, there exists $H' \in \mathcal{R}^J_Y$ such that the minimum number of mistakes sufficient to transit from H'' to H' is bounded above by $\lceil p^{\frac{1}{n-1}} k \rceil (n-1)$. The same argument as in the proof of point 1 allows us to conclude. □

References

Basu K, Weibull JW (1991) Strategy subsets closed under rational behavior. Econ Lett 36:141–146
Ellison G (2000) Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. Rev Econ Stud 67:17–45
Harsanyi JC, Selten R (1988) A general theory of equilibrium selection in games. MIT Press Books, Cambridge, MA
Hurkens S (1995) Learning by forgetful players. Games Econ Behav 11:304–329
Kajii A, Morris S (1997) The robustness of equilibria to incomplete information. Econometrica 65:1283–1309
Klimm M, Sandholm T, Weibull JW (2010) Finding all minimal sCURB sets in finite games. Working Paper, Stockholm School of Economics
Kojima F, Takahashi S (2008) p-Dominance and perfect foresight dynamics. J Econ Behav Organ 67:689–701
Maruta T (1997) On the relationship between risk-dominance and stochastic stability. Games Econ Behav 19:221–234
Matsui A, Matsuyama K (1995) An approach to equilibrium selection. J Econ Theory 65:415–434
Monderer D, Sela A (1997) Fictitious play and no-cycling conditions. Working Paper, The Technion
Morris S, Rob R, Shin HS (1995) p-Dominance and belief potential. Econometrica 63:145–157
Morris S, Ui T (2005) Generalized potentials and robust sets of equilibria. J Econ Theory 124:45–78
Oyama D, Takahashi S, Hofbauer J (2008) Monotone methods for equilibrium selection under perfect foresight dynamics. Theor Econ 3:155–192
Tercieux O (2006a) p-Best response set. J Econ Theory 131:45–70
Tercieux O (2006b) p-Best response sets and the robustness of equilibria to incomplete information. Games Econ Behav 56:371–384
Young HP (1993) The evolution of conventions. Econometrica 61:57–84
Young HP (1998) Individual strategy and social structure. An evolutionary theory of institutions. Princeton University Press, Princeton