Adaptive learning and p-best response set

J. Durieu[1], P. Solal[1], O. Tercieux [2*]

June 18, 2009

[1] CREUSET, University of Saint-Etienne, Saint-Etienne, France,
e-mail: durieu@univ-st-etienne.fr, solal@univ-st-etienne.fr
[2] Paris School of Economics and CNRS, Paris, France,
e-mail: tercieux@pse.ens.fr

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 *à la* Young (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.

Journal of Economic Literature Classification Numbers: C72, C73.
Key Words: Evolutionary game theory, fictitious play process, p-dominance, stochastic stability.

---

*Corresponding author. We would like to thank Daijiro Okada 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".

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. Young (1993) shows that in a 2×2 game with two symmetric equilibria, the risk dominant equilibrium *à la* Harsanyi and Selten (1988) is associated with the unique stochastically stable state of such a process. Following Young (1993), 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 is associated with the unique stochastically stable state of a fictitious play process with bounded memory m and sample k, provided k/m is sufficiently small.

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 (curb is mnemonic for "closed under rational behavior") introduced by Basu and Weibull (1991). He shows that if k/m is sufficiently small, then the stochastically stable states of an independent fictitious play process with bounded memory m and sample k are associated with the minimal curb sets minimizing stochastic potential. However, in order to identify such a minimal curb set, we need to compute the minimal stochastic potential of all minimal curb sets. One difficulty with this approach is that there exists no fast way to compute minimal stochastic potentials.

These results suggest the following questions:

1. Is it possible to derive selection results in two-person games which do not admit a strict 1/2-dominant equilibrium? More generally, can we formulate results for the class of finite n-person games using the concept of minimal p-best response set?
2. When considering the first question, what are the differences in terms of selection results between a joint fictitious play process – where each agent believes that the play of his opponents is correlated – and the independent fictitious play process?
3. Finally, in general, can we dispense with some usual restrictions such as the restriction to weakly acyclic games and the one on the value of k/m?

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 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[1]. 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 $p \leq (n+1)^{-1}$. As a corollary, if the game admits a strict p-dominant equilibrium for $p \leq (n+1)^{-1}$, then the unique stochastically stable state is the one associated with this equilibrium. Thus, the critical value of p is of order $1/n$.[2] Another interesting feature of this result is that it holds for any relative size of the sample.

[1]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.

[2]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, Takahashi and Hofbauer (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.

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 when p is less than n^{1-n}. Comparing this result with the one obtained with a joint fictitious play process, 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 associated with 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 to establish results concerning stochastic stability for both types of processes for at least two reasons. Firstly, it is unnecessary to exhibit 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. Secondly, our proof only requires to compute a lower bound on the minimum number of perturbations necessary to escape the recurrent sets associated with the minimal p-best response set and an upper bound on the number of perturbations sufficient to reach this collection of recurrent sets from other recurrent sets. Using Ellison's (2000) terminology, we 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. [3]

The paper is organized as follows. Section 2 contains notations and preliminaries. In section 3 minimal p-best response sets are formally defined and some properties are established. In section 4 we define adaptive processes

[3]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.

based on a joint fictitious play process. Section 5 states selection results for these processes. Section 6 contains results obtained for independent fictitious play processes with bounded memory and sample. Section 7 concludes.

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 $u_i : X_i \times \Delta(X_{-i}) \to \mathbb{R}$.

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 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 introduce the concept of (minimal) p-best response set. This concept is a generalization of the criterion of strict p-dominance from a singleton set to a product set of strategies. 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, Rob and Shin (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 $i \in I$. We introduce now the concept of p-best response set which 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*

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

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.

Trivially when a p-best response set is a singleton set, it corresponds to a strict p-dominant equilibrium in the sense of Kajii and Morris (1997). Moreover, by definition, a p-best response set is a 1-best response set.

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.).

•

Example 1

In order to illustrate the concept of minimal p-best response set, consider the two-person game Γ^1 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	$(0,0)$	$(5,5)$	$(8,8)$

$\{(1, A)\}$, $\{(2, B)\}$ and $\{(3, C)\}$ are minimal 1-best response sets. $\{(1, A)\}$ is strictly p-dominant for $p > 7/8$; $\{(2, B)\}$ is strictly p-dominant for $p > 3/5$; $\{(3, C)\}$ is strictly p-dominant for $p > 5/8$. But $\{2, 3\} \times \{B, C\}$ is the minimal 1/2-best response set of Γ^1.

□

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. In the sequel, we refer to such beliefs as correlated [4]. Since adaptive processes based on an independent fictitious play process are defined in Young (1993, 1998), we only present in this section adaptive processes based on a joint fictitious play process. Adaptive processes based on an independent fictitious play process will be shortly described in section 6.

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 plays 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 $q_{-i}^J \in \Delta(X_{-i})$. 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. In case of multiple best responses, all are played with positive probability. 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. For any minimal p-best response set Y, define $\mathcal{R}_Y^J \subseteq \mathcal{R}^J$ as the collection of recurrent sets of P_J^0 such that $H' \in \mathcal{R}_Y^J$ if and only if $S(H') \subseteq Y$. Observe that $\mathcal{R}_Y^J \neq \emptyset$ since Y is a minimal p-best response set and thus a closed set under the best-response rule.

---

[4] In a different context, Brandenburger and Dekel (1987) already introduce a similar distinction between independent and correlated beliefs.

In addition, suppose that in each time period there is some small probability $\epsilon > 0$ that agent i experiments by choosing a strategy randomly from X_i. The event that i experiments is assumed to be independent of the event that $j \neq i$ experiments. With these experiments 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 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 and that states which have a positive probability in μ_J^* form a subset of $\mathcal{R}^J$. And, 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 experiments. 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 experiments 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, the stochastically stable sets are the recurrent sets with the minimum stochastic potential. This means that each state belonging to a recurrent set with the minimum stochastic potential is stochastically stable.

5 Joint Fictitious play process

We consider the class of finite n-person games with $n > 2$. 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. The following theorem allows to identify a product set of strategies which contains the stochastically stable states relative to a joint fictitious play process with bounded memory m and sample k.

Theorem 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 $S(H^*) \subseteq Y$.*

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

(1) We give a lower bound on the resistances of the transitions that begin at $H' \in \mathcal{R}_Y^J$ and end in any recurrent set $H'' \notin \mathcal{R}_Y^J$. Observe that if agent $j \in I$ experiments one time and plays outside Y_j, then each agent $i \in I \backslash \{j\}$ may hold beliefs q_{-i}^J such that $\sum_{x_{-i} \notin Y_{-i}} q_{-i}^J(x_{-i}) = 1/k$. And, if agent j experiments 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_{-i}^J such that $\sum_{x_{-i} \notin Y_{-i}} q_{-i}^J(x_{-i}) = \alpha/k$. 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_{-i}^J such that $\sum_{x_{-i} \in Y_{-i}} q_{-i}^J(x_{-i}) = \bar{p}$. By definition of a minimal p-best response set, we have $p > \bar{p}$. Thus, a transition from $H' \in \mathcal{R}_Y^J$ to any $H'' \notin \mathcal{R}_Y^J$ requires at least z experiments (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 $\sum_{x_{-i} \in Y_{-i}} q_{-i}^J(x_{-i}) > \bar{p}$. Therefore, for all recurrent sets $H' \in \mathcal{R}_Y^J$ and $H'' \notin \mathcal{R}_Y^J$, 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 $H'' \notin \mathcal{R}_Y^J$ and end at some element in $\mathcal{R}_Y^J$. Note that, in any state $h'' \in H''$, if each agent $i \in I$ experiments and chooses a strategy $x_i \in Y_i$, then each agent $i \in I$ may hold beliefs q_{-i}^J such that $\sum_{x_{-i} \in Y_{-i}} q_{-i}^J(x_{-i}) > 0$. In particular, if each agent $i \in I$ experiments inside Y_i at one period, we may have $\sum_{x_{-i} \in Y_{-i}} q_{-i}^J(x_{-i}) \geq 1/k$ for all $i \in I$. And, if each agent $i \in I$ experiments and plays inside Y_i α times in succession with $0 < \alpha \leq k$, then we may have $\sum_{x_{-i} \in Y_{-i}} q_{-i}^J(x_{-i}) \geq \alpha/k$ 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_{-i}^J such that $\sum_{x_{-i} \in Y_{-i}} q_{-i}^J(x_{-i}) \geq p$. Thus, from any state $h'' \in H''$, the process may move into an element in $\mathcal{R}_Y^J$ with $\alpha^* n$ experiments where α^* is such that $\alpha^*/k \geq p$. This means that, for any recurrent set $H'' \notin \mathcal{R}_Y^J$, there exists $H' \in \mathcal{R}_Y^J$ such that the minimum number of experiments 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 defavorable case where $k = m$. Note that z experiments 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 observe opponents' profiles in Y_{-i}. Thus, we have shown that for all $H'' \notin \mathcal{R}_Y^J$, there exists $H' \in \mathcal{R}_Y^J$ such that $r_{H''H'} \leq \lceil pk \rceil n$. In the following, such a H' is denoted by $\psi(H'')$. Thus, for all $H'' \notin \mathcal{R}_Y^J$, $r_{H''\psi(H'')} \leq \lceil pk \rceil n$.

(2) Suppose by contradiction that a recurrent set $H'' \notin \mathcal{R}_Y^J$ 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 $H' \in \mathcal{R}_Y^J$ to $\bar{H} \notin \mathcal{R}_Y^J$ (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 $r_{T_{H'}} = r_{T_{H''}} - r_{H'\bar{H}} + r_{H''\psi(H'')}$. And, 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.

•

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 *Suppose 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 $S(H^*) = \{x^*\}$.*

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 in the bounded memory. This result can be explained in the following way. Consider a transition that begins at any recurrent set not contained in $\mathcal{R}_Y^J$ and ends at some recurrent set associated with Y. Due to correlated beliefs, such a transition may require that all agents experiment simultaneously. To understand this, observe that without additional information on the game and on the collection of recurrent sets of the process, it is only possible to state that each agent i has a best-response in Y_i if he has a belief q_{-i}^J such that $\sum_{x_{-i} \in Y_{-i}} q_{-i}^J(x_{-i}) \geq p$. But, if the process is in a state associated with the complement of Y relative to X, it is necessary that all agents experiment simultaneously in Y to have $\sum_{x_{-i} \in Y_{-i}} q_{-i}^J(x_{-i}) > 0$ for each agent i. Thus, to obtain a result that holds for the complete class of finite n-person games, we have to consider both direct transitions and simultaneous experiments among the set of agents to join a recurrent set in $\mathcal{R}_Y^J$. 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. Thus, the sampling procedure, which allows such heterogeneous observations, is not required to establish Theorem 1.

Example 2

Consider the three-person game Γ^2 represented by the payoff matrices

	A	B
1	$(2,2,2)$	$(2,0,2)$
2	$(0,2,2)$	$(0,0,0)$

I

	A	B
1	$(2,2,0)$	$(0,0,0)$
2	$(0,0,0)$	$(9,9,9)$

II

The minimal 1/5-best response set Y of Γ^2 is $\{(2, B, II)\}$ and Theorem 1 applies.

□

6 Independent fictitious play process

In this section, we consider adaptive processes based on an independent fictitious play process. 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 by P^0_I this process. When experiments are allowed, we obtain a perturbed process denoted P^ϵ_I. Denote by μ^*_I the limit stationary distribution of the independent fictitious play process with bounded memory m and sample k.

The following result establishes that, if $p \leq n^{1-n}$, only strategies belonging to the minimal p-best response set are candidates to be selected in the long run. Moreover, in contrast to Young (1993, 1998), this result holds without assumption on the size of the sample k relatively to the size of the memory m. It is sufficient to assume that $k \leq m$.

Let $\mathcal{R}^I$ be the collection of recurrent sets of P^0_I. For any minimal p-best response set Y, define $\mathcal{R}^I_Y \subseteq \mathcal{R}^I$ as the collection of recurrent sets of P^0_I such that $H' \in \mathcal{R}^I_Y$ if and only if $S(H') \subseteq Y$.

Theorem 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.

•

Theorem 2 states that the restriction on the value of k/m can be ruled out. A similar result has been obtained in Theorem 1, although for partially different reasons. Indeed, in Theorem 1, the restriction on the value k/m is not necessary since, in the transitions ending at some recurrent set associated with the minimal p-best response set Y, all agents experiment simultaneously. By contrast, in Theorem 2, such transitions require, in the most defavorable case, that $n-1$ agents experiment simultaneously. Nevertheless, a sampling procedure is not needed. This is due to the focus on direct transitions and to the relative smallness of the critical value of p. Indeed, assume that the process is in a recurrent set not associated with Y and that each agent $j \in I\backslash\{i\}$ experiments in Y_j. Consider the minimal number of periods in which experiments occur necessary to have a belief q^I_{-i} such that $\sum_{x_{-i}\in Y_{-i}} q^I_{-i}(x_{-i}) \geq p$, and thus to have a strategy in Y_i as a best response for agent i. For every admissible value of p, this number constitutes a small proportion (i.e., less than half) of the length of the memory, even if $k = m$. Then, each agent $j \in I\backslash\{i\}$ may have a strategy in Y_j as a best response since he can observe both the periods in which experiments occur and periods in which agent i chooses a strategy in Y_i.

Theorem 2 suggests that the assumption on agents' beliefs 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 Theorem 1, it is of order $1/n$ in Theorem 2. To have an intuition of this result, note that the impact of n experiments at a given time is "stronger" with correlated beliefs than with independent beliefs. Indeed, consider a state belonging to a recurrent set not associated with the minimal p-best response set Y. Assume that, in period t, n agents experiment and play strategies inside Y. For each $i \in I$ observing period t, we have $\sum_{x_{-i}\in Y_{-i}} q^J_{-i}(x_{-i}) \geq 1/k$ and $\sum_{x_{-i}\in Y_{-i}} q^I_{-i}(x_{-i}) \geq (1/k)^{n-1}$. As a consequence, for p, k and n fixed, the upper bound on the resistances of a transition from a recurrent set not associated with Y to a recurrent set associated with Y is greater in the adaptive process based on independent beliefs than in the adaptive process based on correlated beliefs. Since the lower bounds on the resistances of the converse transitions coincide in both processes, the critical value of p is smaller with independent beliefs than with correlated beliefs.

The following example shows that, in comparison with Theorem 1, the introduction of a more restrictive condition on p is necessary to have all strategies selected in the long run associated with the minimal p-best response set when the adaptive process is based on an independent fictitious play process.

Example 3

Consider the three-person game Γ^2 introduced in Example 2. Assume that agent 1 chooses rows, agent 2 columns, and agent 3 boxes. Observe that

Γ^2 is weakly acyclic and possesses two strict Nash equilibria: $(1, A, I)$ and $(2, B, II)$. Moreover, recall that the minimal 1/5-best response set Y of Γ^2 is $\{(2, B, II)\}$. As shown in Example 2, the state associated with $(2, B, II)$ is the unique stochastically stable relative to a joint fictitious play process with bounded memory and sample. Consider 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 show that the unique stochastically stable state of this process is not associated with $(2, B, II)$. In other words, the result in Theorem 2 does not necessarily hold if we consider minimal p-best response sets where $p \leq (n+1)^{-1}$.

Applying Theorem 1 in Young (1993, p.64), we know that the unperturbed adaptive process admits two recurrent sets that are associated with both strict Nash equilibria of Γ^2. As a consequence, in order to identify the stochastically stable state(s), it is sufficient to consider the transitions from one recurrent set to the other one and to determine the minimal number of experiments allowing each transition.

Firstly, we compute the minimal number of experiments allowing a transition from the recurrent set associated with $(2, B, II)$ to the basin of attraction of the recurrent set associated with $(1, A, I)$. Assume that, in period t, the process is in the state associated with $(2, B, II)$. Assume also that, from periods $t+1$ to $t+\lceil \frac{9}{11}k \rceil$ inclusive, one agent, say agent 1, experiments and plays 1. Then, by using the sampling procedure and the inequality $k/m \leq 1/3$, one can build a sequence of plays that ends at the state associated with $(1, A, I)$ without additional experiments. Thus, a transition from the state associated with $(2, B, II)$ to the state associated with $(1, A, I)$ can be done with $\lceil \frac{9}{11}k \rceil$ experiments. Moreover, it can be easily verified that this number of experiments is minimal.

Secondly, we compute the minimal number of experiments allowing a transition from the recurrent set associated with $(1, A, I)$ to the basin of attraction of the recurrent set associated with $(2, B, II)$. Assume that, in period t, the process is in the state associated with $(1, A, I)$. Assume also that, from periods $t+1$ to $t+\lceil \frac{\sqrt{2}}{\sqrt{11}}k \rceil$ inclusive, two agents, say agents 1 and 2, experiment and plays 2 and B respectively. Then, by using the sampling procedure and the inequality $k/m \leq 1/3$, one can build a sequence of plays that ends at the state associated with $(2, B, II)$ without additional experiments. Thus, a transition from the state associated with $(1, A, I)$ to the state associated with $(2, B, II)$ can be done with $2\lceil \frac{\sqrt{2}}{\sqrt{11}}k \rceil$ experiments.

Moreover, it can be easily verified that this number of experiments is minimal.

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 $k > 29$. This means that a transition from the state associated with $(2, B, II)$ to the state associated with the state $(1, A, I)$ can be done with a smaller number of experiments than the converse transition. Thus, the state associated with $(1, A, I)$ is the unique stochastically stable state relative to each independent fictitious play process with bounded memory m and sample k where $k/m \leq 1/3$.

This example allows to draw two conclusions. First, by point 3 in Lemma 1, we know that, for $p \leq 1/2$, $\{(2, B, II)\}$ is the unique minimal p-best response set of Γ^2. Thus, an independent fictitious play process with bounded memory and sample where $k/m \leq 1/3$ may select, in the long run, the state associated with the equilibrium which has the higher value of p. Second, 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.

□

Observe that several recurrent sets may be associated with the minimal p-best response set, particularly when p is small. In such a case, Theorem 2 does not allow a complete identification of the stochastically stable states. However, it makes easier this identification as it excludes states that do not belong to the minimal p-best response set. Hence, less resistance trees have to be calculated. This fact is illustrated in the following example.

Example 4

Consider the two-person game Γ^1 defined in Example 1. The minimal 1/2-best response set of Γ^1 is $\{2, 3\} \times \{B, C\}$. By Theorem 2, the set of stochastically stable states H^* is such that $S(H^*) \subseteq \{2, 3\} \times \{B, C\}$. Note that this identification of the stochastically stable states is partial. Indeed, Young (1993) shows that $S(H^*) = \{(2, B)\}$. Nevertheless, the concept of minimal p-best response set helps to determine the stochastically stable state of the fictitious play process with bounded memory and sample as it excludes the equilibrium $\{(1, A)\}$. So, it is unnecessary to compute the minimum resistance of trees rooted at $\{(1, A)\}$.

□

Finally, considering the class of games that admit a strict p-dominant equilibrium, we can state the following corollary.

Corollary 2 *Suppose 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 $S(H^*) = \{x^*\}$.*

This result can be contrasted with Maruta's main result (1997, p.228). Maruta focuses on the class of weakly acyclic two-person games. He establishes that the strict 1/2-dominant equilibrium is the unique stochastically stable state relative to the Young's process provided it exists. Corollary 2 extends this result to the class of finite n-person games that admit a strict p-dominant equilibrium. This means that it is not necessary to restrict our attention to the class of weakly acyclic games. Moreover, our result holds for a class of adaptive processes defined without restriction on k/m.

7 Conclusion

The concept of (minimal) p-best response set extends the property of strict p-dominance from strict Nash equilibria to product sets of pure strategies. In the class of finite n-person games, it 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 provide a study of the evolution of the critical value of p, helpful to identify the stochastically stable states, as n increases. These results suggest that the way by which agents form beliefs has a crucial impact on the selection results obtained with the concept of minimal p-best response set. The collection of recurrent sets that are not associated with the minimal p-best response set and thus are not candidates for stochastic stability is smaller in independent fictitious play processes than in joint fictitious play processes. Our results are robust to the ratio between the size of the sample and the size of the memory and, therefore apply to a class of fictitious play processes with bounded memory and sample defined without restriction on the relative size of the sample.

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, Takahashi and Hofbauer (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 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[5].

A Appendix

Proof of Theorem 2. Firstly, observe that, by definition of a minimal p-best response set Y, for each agent $i \in I$, all strategies that are best responses to a probability distribution in $\Delta(Y_{-i}, p)$ are contained in Y_{-i}. Hence, a minimal p-best response set is also a minimal p'-best response set where $p' < p$ and p' is sufficiently close to p. In the sequel, we consider a minimal p-best response set where $p < n^{1-n}$. Secondly, note that if $\mathcal{R}^I_Y = \mathcal{R}^I$, then Theorem 2 is trivial. In the sequel, we assume $\mathcal{R}^I_Y \subset \mathcal{R}^I$. The proof is broken into two parts. (1) We bound the resistances of transitions beginning at $H' \in \mathcal{R}^I_Y$ and the resistances of transitions ending at $H' \in \mathcal{R}^I_Y$. (2) We show that the recurrent set(s) minimizing the stochastic potential belongs to $\mathcal{R}^I_Y$.

(1) We give a lower bound on the resistances of the transitions that begin at $H' \in \mathcal{R}^I_Y$ and end in any recurrent set $H'' \notin \mathcal{R}^I_Y$. By using a reasoning similar to the one in Theorem 2, we obtain that $r_{H'H''} \geq \lceil (1-\bar{p})k \rceil$ for all recurrent sets $H' \in \mathcal{R}^I_Y$ and $H'' \notin \mathcal{R}^I_Y$, where $\bar{p}$ is 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^I_{-i} such that $\sum_{x_{-i} \in Y_{-i}} q^I_{-i}(x_{-i}) = \bar{p}$.

We now give an upper bound on the resistances of the transitions that begin at any recurrent set $H'' \notin \mathcal{R}^I_Y$ and end to 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 $j \neq i$ experiments and chooses a

[5]For instance, by Oyama, Takahashi and Hofbauer (2008), we know that in Example 1, $\{(3, C)\}$ is such a maximizer while it is not stochastically stable.

strategy $x_j \in Y_j$. In particular, with $n-1$ experiments uniformly distributed over the agents in $I \backslash \{i\}$, we may have $\sum_{x_{-i} \in Y_{-i}} q^I_{-i}(x_{-i}) \geq (1/k)^{n-1}$. And, if each agent j in $I \backslash \{i\}$ experiments and plays inside Y_j α times in succession with $0 < \alpha \leq k$, then we may have $\sum_{x_{-i} \in Y_{-i}} q^I_{-i}(x_{-i}) \geq (\alpha/k)^{n-1}$. Observe that any other distribution of experiments among agents in $I \backslash \{i\}$ is less powerful in the sense that (for $\alpha(n-1)$ experiments) we always obtain $\sum_{x_{-i} \in Y_{-i}} q^I_{-i}(x_{-i}) < (\alpha/k)^{n-1}$. 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^I_{-i} such that $\sum_{x_{-i} \in Y_{-i}} q^I_{-i}(x_{-i}) \geq p$. Thus, from any state $h'' \in H''$, the process may move into an element in $\mathcal{R}^I_Y$ with $\alpha^*(n-1)$ experiments 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 experiments sufficient to transit from H'' to H' is bounded above by $z = \lceil p^{\frac{1}{n-1}} k \rceil (n-1)$.

In order to verify that such a transition is possible for any $k \leq m$, consider the most defavorable case where $k = m$. Consider a state $h'' \in H''$ and fix $i \in I$. Assume that each agent j in $I \backslash \{i\}$ experiments and chooses a strategy $x_j \in Y_j$ from periods $t+1$ to $t + \lceil p^{\frac{1}{n-1}} m \rceil$ inclusive. From periods $t + \lceil p^{\frac{1}{n-1}} m \rceil + 1$ through $t + 2\lceil p^{\frac{1}{n-1}} m \rceil$, agent i has a strategy inside Y_i as a best response. Observe that $t + 2\lceil p^{\frac{1}{n-1}} m \rceil \leq t + m$ provided m is sufficiently large. In particular, if $n = 2$, then $p < 1/2$ implies that $t + 2\lceil pm \rceil \leq t + m$ provided m is sufficiently large. Then, from $t + 2\lceil p^{\frac{1}{n-1}} m \rceil + 1$, each agent $i \in I$ chooses a strategy $x_i \in Y_i$ since $\sum_{x_{-i} \in Y_{-i}} q^I_{-i}(x_{-i}) \geq p$. Thus, for all $H'' \notin \mathcal{R}^I_Y$, there exists $H' \in \mathcal{R}^I_Y$ such that $r_{H''H'} \leq \lceil p^{\frac{1}{n-1}} k \rceil (n-1)$. In the following, such a H' is denoted by $\psi(H'')$. Thus, for all $H'' \notin \mathcal{R}^I_Y$, $r_{H''\psi(H'')} \leq \lceil p^{\frac{1}{n-1}} k \rceil (n-1)$.

(2) Suppose by contradiction that $H'' \notin \mathcal{R}^I_Y$ 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 $H' \in \mathcal{R}^I_Y$ to $\bar{H} \notin \mathcal{R}^I_Y$ (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 $r_{T_{H'}} = r_{T_{H''}} - r_{H'\bar{H}} + r_{H''\psi(H'')}$. And, for k sufficiently large, we have $r_{H'\bar{H}} > r_{H''\psi(H'')}$, that is

$$\lceil (1-\bar{p})k \rceil > \lceil p^{\frac{1}{n-1}} k \rceil (n-1). \tag{3}$$

To see this, note that a sufficient condition for inequality (3) to hold is

$$1 - \bar{p} > p^{\frac{1}{n-1}}(n-1) + \frac{n-1}{k}. \qquad (4)$$

And, (4) 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, $p \leq n^{1-n}$. Thus, $r_{T_{H'}} < r_{T_{H''}}$ and $\rho(H') < \rho(H'')$. This contradicts the fact that H'' is stochastically stable.

•

References

Basu, K., and Weibull, J.W. (1991). "Strategy Subsets Closed Under Rational Behavior", *Economics Letters* **36**, 141-146.

Brandenburger, A., and Dekel, E. (1987). "Rationalizability and Correlated Equilibria", *Econometrica* **55**, 1391-1402.

Ellison, G. (2000). "Basins of Attraction, Long-Run Stochastic Stability, and the Speed of Step-by-Step Evolution", *Review of Economic Studies* **67**, 17-45.

Harsanyi, J.C, and Selten, R., (1988). *A General Theory of Equilibrium Selection in Games*, MIT Press Books.

Kajii, A., and Morris, S. (1997). "The Robustness of Equilibria to Incomplete Information", *Econometrica* **65**, 1283-1309.

Kojima, F., and Takahashi S. (2008). "p-dominance and perfect foresight dynamics", *Journal of Economic Behavior and Organization*, **67**, 689-701.

Maruta, T. (1997). "On the Relationship between Risk-Dominance and Stochastic Stability", *Games and Economic Behavior* **19**, 221-234.

Matsui, A., and Matsuyama, K. (1995). "An Approach to Equilibrium Selection", *Journal of Economic Theory* **65**, 415-434.

Monderer, D., and Sela, A. (1997). "Fictitious Play and No-Cycling Conditions", Working Paper, The Technion.

Morris, S., Rob, R. and Shin, H.S. (1995). "p-dominance and Belief Potential", *Econometrica* **63**, 145-157.

Morris, S., and Ui, T. (2005). "Generalized Potentials and Robust Sets of Equilibria" *Journal of Economic Theory* **124**, 45-78.

Oyama, D., Takahashi, S. and Hofbauer, J. (2008). "Monotone Methods for Equilibrium Selection under Perfect Foresight Dynamics", *Theoretical Economics*, **3**, 155-192.

Tercieux, O. (2006a). "p-Best Response Set", *Journal of Economic Theory*, **131**, 45-70.

Tercieux, O. (2006b). "p-Best Response Sets and the Robustness of Equilibria to Incomplete Information", *Games and Economic Behavior* , **56**, 371-384.

Young, H.P. (1993). "The Evolution of Conventions", *Econometrica* **61**, 57-84.

Young, H.P. (1998). *Individual Strategy and Social Structure. An Evolutionary Theory of Institutions*. Princeton: Princeton University Press.