Consistency of Vanishing Smooth Fictitious Play

(In collaboration with Michel Benaïm)

Mathieu Faure

UNINE

June 21th 2011

Two-player normal form game

Let us consider a 2-person game in normal form :

- Player 1 (Decision Maker), Player 2 (Nature),
- finite sets of actions I and L; sets of mixed actions : $\Delta(I)$, $\Delta(L)$,
- payoff function of DM : π ($I \times L$ matrix)

Two-player normal form game

Let us consider a 2-person game in normal form :

- Player 1 (Decision Maker), Player 2 (Nature),
- finite sets of actions I and L; sets of mixed actions : $\Delta(I)$, $\Delta(L)$,
- payoff function of DM : π ($I \times L$ matrix)
- we assume that agents play **repeatedly**. Let $h_n = (i_1, l_1, ..., i_n, l_n)$ be history up to time n.

Two-player normal form game

Let us consider a 2-person game in normal form :

- Player 1 (Decision Maker), Player 2 (Nature),
- finite sets of actions I and L; sets of mixed actions : $\Delta(I)$, $\Delta(L)$,
- payoff function of DM : π ($I \times L$ matrix)
- we assume that agents play **repeatedly**. Let $h_n = (i_1, l_1, ..., i_n, l_n)$ be history up to time n.
- a strategy for DM is a map σ :

$$\cup_n \mathcal{H}_n \to \Delta(I), \;\; h_n \mapsto \sigma(h_n) \in \Delta(I).$$

Two-player normal form game

Let us consider a 2-person game in normal form :

- Player 1 (Decision Maker), Player 2 (Nature),
- finite sets of actions I and L; sets of mixed actions : $\Delta(I)$, $\Delta(L)$,
- payoff function of DM : π ($I \times L$ matrix)
- we assume that agents play **repeatedly**. Let $h_n = (i_1, l_1, ..., i_n, l_n)$ be history up to time n.
- a strategy for DM is a map σ :

$$\cup_n \mathcal{H}_n \to \Delta(I), \;\; h_n \mapsto \sigma(h_n) \in \Delta(I).$$

- A pair of strategies (σ, τ) induces a probability measure $\mathbb{P}$ on $(I \times L)^{\mathbb{N}}$; we assume that agents play **independently** :

$$\mathbb{P}\left(i_{n+1} = i, l_{n+1} = l \mid \mathcal{H}_n\right) = \mathbb{P}\left(i_{n+1} = i \mid \mathcal{H}_n\right) \mathbb{P}\left(l_{n+1} = l \mid \mathcal{H}_n\right).$$

Regret

Empirical distribution of moves and the average realized payoff up to time n :

$$x_n := \frac{1}{n}\sum_{k=1}^{n}\delta_{i_k}, \;\; y_n = \frac{1}{n}\sum_{k=1}^{n}\delta_{l_k}, \;\; \pi_n := \frac{1}{n}\sum_{k=1}^{n}\pi(i_k, l_k).$$

Define

$$\Pi : Y \to \mathbb{R} : \;\; y \mapsto \max_i \pi(i, y).$$

Regret

Empirical distribution of moves and the average realized payoff up to time n :

$$x_n := \frac{1}{n}\sum_{k=1}^{n}\delta_{i_k}, \;\; y_n = \frac{1}{n}\sum_{k=1}^{n}\delta_{l_k}, \;\; \pi_n := \frac{1}{n}\sum_{k=1}^{n}\pi(i_k, l_k).$$

Define

$$\Pi : Y \to \mathbb{R} : \;\; y \mapsto \max_i \pi(i, y).$$

Definition (regret at stage n)

$$e_n = \qquad\qquad .$$

Regret

Empirical distribution of moves and the average realized payoff up to time n :

$$x_n := \frac{1}{n}\sum_{k=1}^{n}\delta_{i_k}, \;\; y_n = \frac{1}{n}\sum_{k=1}^{n}\delta_{l_k}, \;\; \pi_n := \frac{1}{n}\sum_{k=1}^{n}\pi(i_k, l_k).$$

Define

$$\Pi : Y \to \mathbb{R} : \;\; y \mapsto \max_i \pi(i, y).$$

Definition (regret at stage n)

$$e_n = \underbrace{\max_{i\in I}\pi(i, y_n)}_{\Pi(y_n)} \qquad\qquad .$$

Regret

Empirical distribution of moves and the average realized payoff up to time n :

$$x_n := \frac{1}{n}\sum_{k=1}^{n}\delta_{i_k},\ \ y_n = \frac{1}{n}\sum_{k=1}^{n}\delta_{l_k},\ \ \pi_n := \frac{1}{n}\sum_{k=1}^{n}\pi(i_k, l_k).$$

Define

$$\Pi : Y \to \mathbb{R} :\ \ y \mapsto \max_i \pi(i, y).$$

Definition (regret at stage n)

$$e_n = \underbrace{\max_{i\in I}\pi(i, y_n)}_{\Pi(y_n)} - \underbrace{\frac{1}{n}\sum_{k=1}^{n}\pi(i_k, l_k)}_{\pi_n}.$$

Consistency

Definition

Player 1's strategy is *consistent* if, regardless of the strategy τ of nature,

$$\limsup_{n \to +\infty} e_n \leq 0, \text{ almost surely.}$$

It is η-consistent provided

$$\limsup_{n \to +\infty} e_n \leq \eta, \text{ almost surely.}$$

Fictitious play

Let br be the *best response map* :

$$br : Y \rightrightarrows X, \;\; y \mapsto Argmax_{x \in X} \pi(x,y)$$

Fictitious play

Let br be the *best response map* :

$$br : Y \rightrightarrows X, \;\; y \mapsto Argmax_{x \in X} \pi(x, y)$$

Definition

The strategy of DM σ is a *fictitious play* strategy if, $\forall n \in \mathbb{N}$,

$$\sigma(h_n) \in br(\tilde{y}_n),$$

Fictitious play

Let br be the *best response map* :

$$br : Y \rightrightarrows X, \;\; y \mapsto Argmax_{x \in X} \pi(x, y)$$

Definition

The strategy of DM σ is a *fictitious play* strategy if, $\forall n \in \mathbb{N}$,

$$\sigma(h_n) \in br(\tilde{y}_n),$$

with $\tilde{y}_n = \frac{1}{n+1} \underbrace{y_0}_{\text{prior}} + \frac{n}{n+1} y_n.$

Remark

FP is not consistent

Perturbed payoff function and Smooth best response

Consider the perturbed payoff function

$$\tilde{\pi}^{\varepsilon}(x,y) := \pi(x,y) + \varepsilon\rho(x),$$

where

- ε is a positive parameter,
- $\rho(x) = -\sum_k x_k \log x_k$.

Then we have

Perturbed payoff function and Smooth best response

Consider the perturbed payoff function

$$\tilde{\pi}^{\varepsilon}(x,y) := \pi(x,y) + \varepsilon\rho(x),$$

where

- ε is a positive parameter,
- $\rho(x) = -\sum_k x_k \log x_k$.

Then we have

- For all $y \in Y$, $Argmax_{x\in X}\tilde{\pi}^{\varepsilon}(\cdot,y)$ reduces to one point,

Perturbed payoff function and Smooth best response

Consider the perturbed payoff function

$$\tilde{\pi}^{\varepsilon}(x,y) := \pi(x,y) + \varepsilon\rho(x),$$

where

- ε is a positive parameter,
- $\rho(x) = -\sum_k x_k \log x_k$.

Then we have

- For all $y \in Y$, $Argmax_{x \in X} \tilde{\pi}^{\varepsilon}(\cdot, y)$ reduces to one point,
- we can define the smooth best response map $\tilde{\mathbf{br}} : Y \to X$:

$$\tilde{\mathbf{br}}(y,\varepsilon) := Argmax_{x \in X} \tilde{\pi}^{\varepsilon}(x,y).$$

Smooth fictitious play

Definition (Smooth fictitious play)

σ is a smooth fictitious play (SFP(ε)) strategy for player 1 if, for all $n \in \mathbb{N}$

$$\sigma(h_n) = \tilde{\mathbf{br}}(y_n, \varepsilon),$$

Smooth fictitious play

Definition (Smooth fictitious play)

σ is a smooth fictitious play (SFP(ε)) strategy for player 1 if, for all $n \in \mathbb{N}$

$$\sigma(h_n) = \tilde{\mathbf{br}}(y_n, \varepsilon),$$

Theorem (Fudenberg-Levine)

Given $\eta > 0$, SFP(ε) is η-consistent, provided ε is small enough.

Smooth fictitious play

Definition (Smooth fictitious play)

σ is a smooth fictitious play (SFP(ε)) strategy for player 1 if, for all $n \in \mathbb{N}$

$$\sigma(h_n) = \tilde{\mathbf{br}}(y_n, \varepsilon),$$

Theorem (Fudenberg-Levine)

Given $\eta > 0$, SFP(ε) is η-consistent, provided ε is small enough.

Benaïm-Hofbauer-Sorin (2006) gave an alternative proof using stochastic approximations technics.

Stochastic Approximation Algorithms, the ODE method

M compact subset of $\mathbb{R}^d$. Let $(v_n)_n$ be a M-valued stochastic process governed by the recursive formula

$$v_{n+1} - v_n - \frac{1}{n+1}U_{n+1} = \frac{1}{n+1}f(v_n),$$

where

- f is a *Lipschitz* vector field, inducing a flow Φ,
- $(U_n)_n$ is a *bounded* sequence of random variables.

Question : can we say anything about the (qualitative) asymptotic behavior of $(v_n)_n$?

Mean ODE

$$v_{n+1} = v_n + \frac{1}{n+1}(f(v_n) + U_{n+1}), \tag{1}$$

Consider the mean ODE :

$$\dot{v} = f(v). \tag{2}$$

Mean ODE

$$v_{n+1} = v_n + \frac{1}{n+1}(f(v_n) + U_{n+1}), \tag{1}$$

Consider the mean ODE :

$$\dot{v} = f(v). \tag{2}$$

Link between (1) and (2) : if $(U_n)_n$ is a martingale difference : $\mathbb{E}(U_{n+1} \mid \mathcal{F}_n) = 0$, the asymptotic behavior of the paths $(v_n(\omega))_n$ should be related to the solution curves of (2).

Convergence of Stochastic Approximation Algorithms

Definition

A compact set A is an *attractor* for Φ if there exists an open set $U \supset A$ such that, $\forall \varepsilon > 0$,

$$\Phi_t(U) \subset N^{\varepsilon}(A),$$

provided t is large enough. If $U = M$, A is a *global attractor*.

Convergence of Stochastic Approximation Algorithms

Definition

A compact set A is an *attractor* for Φ if there exists an open set $U \supset A$ such that, $\forall \varepsilon > 0$,

$$\Phi_t(U) \subset N^{\varepsilon}(A),$$

provided t is large enough. If $U = M$, A is a *global attractor*.

Theorem (Limit set theorem, Benaïm, 1996)

$a)$ *The limit set of $(v_n)_n$ is almost surely internally chain transitive for Φ, i.e. a compact, invariant and attractor-free,*

$b)$ *if A is a global attractor, $\mathcal{L}((v_n)_n) \subset A$ almost surely.*

Stochastic approximations related to a differential inclusion

Recursive form :

$$v_{n+1} - v_n - \frac{1}{n+1} U_{n+1} = \frac{1}{n+1} f(v_n), \tag{3}$$

Stochastic approximations related to a differential inclusion

Recursive form :

$$v_{n+1} - v_n - \frac{1}{n+1} U_{n+1} \in \frac{1}{n+1} F(v_n), \tag{3}$$

where F is a sufficiently regular set-valued map.

Stochastic approximations related to a differential inclusion

Recursive form :

$$v_{n+1} - v_n - \frac{1}{n+1} U_{n+1} \in \frac{1}{n+1} F(v_n), \tag{3}$$

where F is a sufficiently regular set-valued map.

Mean Differential Inclusion :

$$\frac{d}{dt} v(t) \in F(v(t)) \tag{4}$$

Limit set theorem

Definition

A solution is an absolutely continuous map such that, for almost every t, (4) holds. The flow induced by (4) is given by

$$\Phi_t(v_0) := \{v(t) : \; v(\cdot) \text{ solution }, \; v(0) = v_0\}.$$

Limit set theorem

Definition

A solution is an absolutely continuous map such that, for almost every t, (4) holds. The flow induced by (4) is given by

$$\Phi_t(v_0) := \{v(t) : \; v(\cdot) \text{ solution }, \; v(0) = v_0\}.$$

Theorem (Benaïm-Hofbauer-Sorin, 2005)

a) The limit set of $(v_n)_n$ is almost surely internally chain transitive for Φ, i.e. a compact, invariant and attractor-free,

b) if A is a global attractor, $\mathcal{L}((v_n)_n) \subset A$ almost surely.

Back to SFP(ε)

State variable : $v_n = (x_n, y_n, \pi_n)$. We have

$$x_{n+1} - x_n = \frac{1}{n+1}\left(\delta_{i_{n+1}} - x_n\right).$$

We have

$$x_{n+1} - x_n - \frac{1}{n+1}\underbrace{\left[\delta_{i_{n+1}} - \tilde{\mathbf{br}}^{\varepsilon}(y_n)\right]}_{\delta_{i_{n+1}} - \mathbb{E}\left(\delta_{i_{n+1}} | \mathcal{H}_n\right)} = \frac{1}{n+1}\left(\tilde{\mathbf{br}}^{\varepsilon}(y_n) - x_n\right)$$

Proof of η-consistency via stochastic approximations

We can prove that

$$v_{n+1} - v_n - \frac{1}{n+1} U_{n+1} \in \frac{1}{n+1} F^{\varepsilon}(v_n),$$

where

- The set-valued map F^{ε} is very regular,
- As a consequence, BHS results apply and, if the differential inclusion $\dot{v}(t) \in F^{\varepsilon}(v(t))$ admits a global attractor A then $\mathcal{L}((v_n)_n) \subset A$.

Proof of η-consistency via stochastic approximations

Theorem (BHS, 2006)

Given $\eta > 0$, for ε small enough we have

- *the set $A := \{v = (x, y, \pi) : \ \Pi(y) - \pi \leq \eta\}$ contains a global attractor for the differential inclusion $\dot{v}(t) \in F^{\varepsilon}(v(t))$,*
- *consequently*

$$\limsup_n \Pi(y_n) - \pi_n \leq \eta \text{ almost surely.}$$

A natural question

What happens when the parameter ε is replaced by a vanishing sequence $\varepsilon_n \downarrow_n 0$?

A natural question

What happens when the parameter ε is replaced by a vanishing sequence $\varepsilon_n \downarrow_n 0$?

Definition (Vanishing Smooth fictitious play)

Given a sequence $\varepsilon_n \downarrow_n 0$, we say that Player 1 plays accordingly to a vanishing smooth fictitious play strategy (VSFP(ε_n)) if, for all $n \in \mathbb{N}$,

$$\sigma(h_n) = \tilde{\mathbf{br}}(y_n, \varepsilon_n).$$

A natural question

What happens when the parameter ε is replaced by a vanishing sequence $\varepsilon_n \downarrow_n 0$?

Definition (Vanishing Smooth fictitious play)

Given a sequence $\varepsilon_n \downarrow_n 0$, we say that Player 1 plays accordingly to a vanishing smooth fictitious play strategy (VSFP(ε_n)) if, for all $n \in \mathbb{N}$,

$$\sigma(h_n) = \tilde{\mathbf{br}}(y_n, \varepsilon_n).$$

Remark

VSFP is not consistent, if $\varepsilon_n = \frac{1}{n}$

Statement of the main result

Theorem (Benaïm, Faure)

If $\varepsilon_n = \frac{1}{n^\alpha}$ $(\alpha < 1)$ *then VSFP*(ε_n) *is consistent.*

- The proof relies on set-valued dynamical systems approach, similarily to BHS,
- unfortunately, we have to enter the realm of *nonautonomous* differential inclusions,

Sketch of the proof, nonautonomous DI

The stochastic process $(v_n)_n$ can be written as

$$v_{n+1} - v_n - \frac{1}{n+1} U_{n+1} \in \frac{1}{n+1} F_n(v_n), \tag{5}$$

where

- $F_n = F^{\varepsilon_n}$,
- $(U_n)_n$ is a bounded martingale difference.

Sketch of the proof, nonautonomous DI

The stochastic process $(v_n)_n$ can be written as

$$v_{n+1} - v_n - \frac{1}{n+1} U_{n+1} \in \frac{1}{n+1} F_n(v_n), \tag{5}$$

where

- $F_n = F^{\varepsilon_n}$,
- $(U_n)_n$ is a bounded martingale difference.

Mean nonautonomous Differential Inclusion :

$$\frac{d}{dt} v(t) \in F(t, v(t)) := F^{\varepsilon_{m(t)}}(v(t)), \tag{6}$$

with $m(t) := \sup \left\{ n \in \mathbb{N} : \sum_{i=1}^{n} \frac{1}{n+1} < t \right\}$

Sketch of the proof, why ε_n matters ?

- The qualitative results proved by BHS on the asymptotic behavior of $(v_n)_n$ in the autonomous case do not apply anymore,
- the good news is that the nonautonomous map F governing the system is very regular (Hausdorff-Lipschitz) :

$$d_H(F(t,v),F(t,v')) \leq L(t)\|v-v'\|, \ \forall t>0, \ v,v' \in M$$

with Lipschitz "constant" $L(t)$ going to infinity at a speed which depends on the sequence $(\varepsilon_n)_n$. More precisely $L(t)=\mathcal{O}(e^{\alpha t})$ if $\epsilon_n = \frac{1}{n^\alpha}$.

- if $L(t)$ does not go to infinity too fast, we can prove that the asymptotic behavior of $(v_n)_n$ is again related to (6).