

ϵ -robust Nash Equilibria

Patricia Bouyer, Nicolas Markey and *Daniel STAN*

CNRS, LSV, ENS Cachan

Fifth Cassting Meeting 2014



- 1 Concurrent framework
- 2 Existence of equilibria
- 3 Linear and robust equilibria

- 1 Concurrent framework
- 2 Existence of equilibria
- 3 Linear and robust equilibria

Games with mixed strategies

Concurrent non-zero sum games allow

- To modelize heterogeneous systems
- Several events to occur simultaneously
- Agents' goals not to be necessarily antagonistic

Games with mixed strategies

Concurrent non-zero sum games allow

- To modelize heterogeneous systems
- Several events to occur simultaneously
- Agents' goals not to be necessarily antagonistic

whereas mixed strategies enable

- Synthesizing strategies for controllers
- with memory
- Breaking the symmetry (by randomization)
- Equilibrium more likely to occur

Formal model

Definition (Arena)

$$\mathcal{A} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Tab}, (\text{Allow}_i)_{i \in \text{Agt}} \rangle$$

with

- $|\text{States}|, |\text{Agt}|, |\text{Act}| < +\infty$
- $\text{Tab} : \text{States} \times \text{Act}^{\text{Agt}} \rightarrow \text{States}$
- $\forall i \in \text{Agt} \quad \text{Allow}_i : \text{States} \rightarrow 2^{\text{Act} \setminus \{\emptyset\}}$


Formal model

Definition (Arena)

$$\mathcal{A} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Tab}, (\text{Allow}_i)_{i \in \text{Agt}} \rangle$$

with

- $|\text{States}|, |\text{Agt}|, |\text{Act}| < +\infty$
- $\text{Tab} : \text{States} \times \text{Act}^{\text{Agt}} \rightarrow \text{States}$
- $\forall i \in \text{Agt} \quad \text{Allow}_i : \text{States} \rightarrow 2^{\text{Act} \setminus \{\emptyset\}}$



Definition (Game)

$$\mathcal{G} = \langle \mathcal{A}, s, \phi \rangle$$

where


- \mathcal{A} is an arena
- $s \in \text{States}$ is an initial state
- $\phi : \text{States}^\omega \rightarrow \mathbb{R}^{\text{Agt}}$ a utility function

Definition (Game)

$$\mathcal{G} = \langle \mathcal{A}, s, \phi \rangle$$

where

- \mathcal{A} is an arena
- $s \in \text{States}$ is an initial state
- $\phi : \text{States}^\omega \rightarrow \mathbb{R}^{\text{Agt}}$ a utility function



$$\phi(r) = \begin{cases} (1, 0) & \text{if } r \in \text{States}^* w_1^\omega \\ (0, 1) & \text{if } r \in \text{States}^* w_2^\omega \\ (0, 0) & \text{otherwise} \end{cases}$$

Family of utility functions

- Safety condition
- Reachability
- Limit average
- Terminal reachability

Family of utility functions

- Safety condition
- Reachability
- Limit average
- Terminal reachability

Definition (Final states)

Let F denote the set of states that have no successor except themselves. ϕ is a *terminal reachability utility function* if

$$\forall r \phi(r) \neq 0 \Leftrightarrow \exists h \in \text{States}^* \exists f \in F : r = h \cdot f^\omega$$

Family of utility functions

- Safety condition
- Reachability
- Limit average
- Terminal reachability

Definition (Final states)

Let F denote the set of states that have no successor except themselves. ϕ is a *terminal reachability utility function* if

$$\forall r \phi(r) \neq 0 \Leftrightarrow \exists h \in \text{States}^* \exists f \in F : r = h \cdot f^\omega \wedge \phi(r) = \phi(f^\omega)$$


Family of utility functions

- Safety condition
- Reachability
- Limit average
- Terminal reachability

Definition (Final states)

Let F denote the set of states that have no successor except themselves. ϕ is a *terminal reachability utility function* if

$$\forall r \phi(r) \neq 0 \Leftrightarrow \exists h \in \text{States}^* \exists f \in F : r = h \cdot f^\omega \wedge \phi(r) = \phi(f^\omega)$$



Definition (Strategies)

A *strategy* for player i in arena \mathcal{A} is given by σ_i such that for all $h \in \text{States}^+$,

$$\sigma_i(h) \in \text{Dist}(\text{Allow}_i(\text{last}(h)))$$

We call *strategy profile* the data of strategies for all players, and any finite non-empty sequence of states is a *history*.

Definition (Strategies)

A *strategy* for player i in arena \mathcal{A} is given by σ_i such that for all $h \in \text{States}^+$,

$$\sigma_i(h) \in \text{Dist}(\text{Allow}_i(\text{last}(h)))$$

We call *strategy profile* the data of strategies for all players, and any finite non-empty sequence of states is a *history*.

Definition (Expectation)

We consider a game \mathcal{G} and a strategy profile σ . $X_0 = s$,

$X_{n+1} = \text{Tab}(X_n, A_n)$ with $A_n \sim \prod_i \sigma_i(X_0 \dots X_n)$.

Let $r = \lim X_0 \dots X_n \in \text{States}^\omega$.

Under some mesurability assumptions, the expectation of $\phi(r)$ exists.

If $\mathbb{P}(r \in h\text{States}^\omega) > 0$, we write $\mathbb{E}^\sigma(\phi \mid h)$ the conditionnal expectation.

Nash Equilibrium

Definition

Let σ a strategy profile and h an history, then (σ, h) is a *Nash Equilibrium* (NE) if for all agent i and any other strategy for i (*deviation*) σ'_i ,

$$\mathbb{E}^{\sigma^{[i/\sigma'_i]}}(\phi_i | h) \leq \mathbb{E}^{\sigma}(\phi_i | h)$$

Nash Equilibrium

Definition

Let σ a strategy profile and h an history, then (σ, h) is a *Nash Equilibrium* (NE) if for all agent i and any other strategy for i (*deviation*) σ'_i ,

$$\mathbb{E}^{\sigma^{[i/\sigma'_i]}}(\phi_i | h) \leq \mathbb{E}^{\sigma}(\phi_i | h)$$

We can show that we can restrict to deterministic deviation only (for terminal reachability objectives).


Nash Equilibrium

Definition

Let σ a strategy profile and h an history, then (σ, h) is a *Nash Equilibrium* (NE) if for all agent i and any other strategy for i (*deviation*) σ'_i ,

$$\mathbb{E}^{\sigma^{[i/\sigma'_i]}}(\phi_i | h) \leq \mathbb{E}^{\sigma}(\phi_i | h)$$


We can show that we can restrict to deterministic deviation only (for terminal reachability objectives).



The uniform strategy for both players is a NE (payoff $(2/3, 1/3)$).

- 1 Concurrent framework
- 2 Existence of equilibria
- 3 Linear and robust equilibria

Does a mixed Nash Equilibrium always exist?



Does a mixed Nash Equilibrium always exist?




Figure: Hide-or-Run game

Does a mixed Nash Equilibrium always exist?




Figure: Hide-or-Run game

Value problem in a zero-sum game is *not* a special case of Nash Equilibrium problem with *positive* terminal rewards

Does a mixed Nash Equilibrium always exist?




Figure: Hide-or-Run game

Value problem in a zero-sum game is *not* a special case of Nash Equilibrium problem with *positive* terminal rewards

Theorem

The existence problem is undecidable for 3-player concurrent games with non-negative terminal rewards and a constrain.

Does a mixed Nash Equilibrium always exist?




Figure: Hide-or-Run game

Value problem in a zero-sum game is *not* a special case of Nash Equilibrium problem with *positive* terminal rewards

Theorem

The existence problem is undecidable for 3-player concurrent games with non-negative terminal rewards and a constrain. Also holds on arbitrary terminal rewards without constrains.

Existence of equilibria

Theorem (Nash 1950)

Every one-stage game has a Nash Equilibrium in mixed strategies.

Existence of equilibria

Theorem (Nash 1950)

Every one-stage game has a Nash Equilibrium in mixed strategies.

Theorem (Secchi and Sudderth 2001)

NE always exists for safety qualitative objectives. Strategies have finite memory.

Existence of equilibria

Theorem (Nash 1950)

Every one-stage game has a Nash Equilibrium in mixed strategies.

Theorem (Secchi and Sudderth 2001)

NE always exists for safety qualitative objectives. Strategies have finite memory.

Theorem (Chatterjee et al. 2004)

For $\epsilon > 0$, ϵ -Nash Equilibrium always exists with terminal reward, and strategies are stationary.

Termination problem

General scheme.

Let \mathbb{M} be the set of stationary strategy profiles.

Consider the best response function: $BR : \mathbb{M} \rightarrow 2^{\mathbb{M}}$ mapping a to a set of strategy profiles improving the payoff of each player and show it is continuous, then apply Kakutani fix-point theorem to show

$\exists \sigma \sigma \in BR(\sigma)$. □

Termination problem

General scheme.

Let \mathbb{M} be the set of stationary strategy profiles.

Consider the best response function: $BR : \mathbb{M} \rightarrow 2^{\mathbb{M}}$ mapping a to a set of strategy profiles improving the payoff of each player and show it is continuous, then apply Kakutani fix-point theorem to show

$\exists \sigma \sigma \in BR(\sigma)$. □

Continuity of BR is based on termination assumptions.

Termination problem

General scheme.

Let \mathbb{M} be the set of stationary strategy profiles.

Consider the best response function: $BR : \mathbb{M} \rightarrow 2^{\mathbb{M}}$ mapping a to a set of strategy profiles improving the payoff of each player and show it is continuous, then apply Kakutani fix-point theorem to show

$\exists \sigma \sigma \in BR(\sigma)$. □

Continuity of BR is based on termination assumptions.

- One-stage termination

Termination problem

General scheme.

Let \mathbb{M} be the set of stationary strategy profiles.

Consider the best response function: $BR : \mathbb{M} \rightarrow 2^{\mathbb{M}}$ mapping a to a set of strategy profiles improving the payoff of each player and show it is continuous, then apply Kakutani fix-point theorem to show

$\exists \sigma \sigma \in BR(\sigma)$. □

Continuity of BR is based on termination assumptions.

- One-stage termination
- Assume no final safety collaboration, bound probability to make someone loose

Termination problem

General scheme.

Let \mathbb{M} be the set of stationary strategy profiles.


Consider the best response function: $BR : \mathbb{M} \rightarrow 2^{\mathbb{M}}$ mapping a to a set of strategy profiles improving the payoff of each player and show it is continuous, then apply Kakutani fix-point theorem to show

$\exists \sigma \sigma \in BR(\sigma)$. □


Continuity of BR is based on termination assumptions.

- One-stage termination
- Assume no final safety collaboration, bound probability to make someone loose
- Consider a discounted version

Limit behaviour




Limit behaviour




- NE strategies (probability of playing b):
 $\{(x, 0), (0, x) \mid 1 \geq x > 0\}$

Limit behaviour



- NE strategies (probability of playing b):
 $\{(x, 0), (0, x) \mid 1 \geq x > 0\}$
- NE payoffs: $\{(1, 0), (0, 1)\}$

Limit behaviour



- NE strategies (probability of playing b):
 $\{(x, 0), (0, x) \mid 1 \geq x > 0\}$
- NE payoffs: $\{(1, 0), (0, 1)\}$
- BR function graph not continuous in $(0, 0)$

- 1 Concurrent framework
- 2 Existence of equilibria
- 3 Linear and robust equilibria**

Assumptions

From now, we consider stationary memoryless strategies (set \mathbb{M})

Assumptions

From now, we consider stationary memoryless strategies (set \mathbb{M})

Definition (Cycling Arena)

Let \mathcal{A} be an arena. Assume there exists an state $s \in \text{States}$ and a mixed strategy profile σ such that no player can enforce reaching a final state:

$$\forall i \in \text{Agt} \forall \sigma'_i \in \mathbb{M}_i; \mathbb{P}^{\sigma[i/\sigma'_i]}(\text{States}^*F^\omega \mid s) = 0$$

Such a state is called cycling.

Note that such a definition implies a Nash Equilibrium with payoff 0 for all players.

Assumptions

From now, we consider stationary memoryless strategies (set \mathbb{M})

Definition (Cycling Arena)

Let \mathcal{A} be an arena. Assume there exists an state $s \in \text{States}$ and a mixed strategy profile σ such that no player can enforce reaching a final state:

$$\forall i \in \text{Agt} \forall \sigma'_i \in \mathbb{M}_i; \mathbb{P}^{\sigma[i/\sigma_i]^s}(\text{States}^*F^\omega \mid s) = 0$$

Such a state is called cycling.

Note that such a definition implies a Nash Equilibrium with payoff 0 for all players.

Lemma (Remark)

One can effectively transform every game \mathcal{G} into a non-cycling game \mathcal{G}' , such that every Nash Equilibrium in \mathcal{G}' can be converted into a Nash Equilibrium in \mathcal{G} .

Strong components

Definition (Strong Component)

Let \mathcal{A} be an arena and $C \subseteq \text{States}$.

C is called a strong component if there exists $\sigma \in \mathbb{M}$ such that every state of C is reachable from another with strategy profile σ :

$$\forall s, s' \in C \mathbb{P}^\sigma(\text{States}^* s' \mid s) > 0$$

Such σ will be said to stabilize C . We denote with SC the set of strong components.

Note that it is equivalent to say that the previous probability is equal to 1. We can also remark that every strong component intersecting F is reduced to a singleton.

Strong component escaping

From now on, we consider non-cycling games.

Strong component escaping

From now on, we consider non-cycling games.

Definition (Exiting actions)

Let $C \in \text{SC}$. $a \in \text{Act}$ is an exiting action from C for state $s \in C$ and player i if:

$$\mathbb{P}^{\sigma}[i/(s \rightarrow a)]^s (s \cdot (\text{States} \setminus C) \mid s) > 0$$

for some stationary σ stabilizing C

We define

$$\text{Exit}(C) = \{(a, i, s) \mid a \text{ is an exiting action from } C \text{ state } s \in C \text{ for player } i\}$$

Strong component escaping

From now on, we consider non-cycling games.

Definition (Exiting actions)

Let $C \in \text{SC}$. $a \in \text{Act}$ is an exiting action from C for state $s \in C$ and player i if:

$$\mathbb{P}^{\sigma}[i/(s \rightarrow a)]^s (s \cdot (\text{States} \setminus C) \mid s) > 0$$

for some stationary σ stabilizing C

We define

$$\text{Exit}(C) = \{(a, i, s) \mid a \text{ is an exiting action from } C \text{ state } s \in C \text{ for player } i\}$$

Lemma

For any $C \in \text{SC}$, $\text{Exit}(C) \neq \emptyset$.

Reduced State space

Definition

For any $C \in \text{SC}$ strong component and $\epsilon > 0$, we define

$$\Delta_\epsilon(C) = \left\{ \sigma \in \mathbb{M} \mid \sum_{(a,i,s) \in \text{Exit}(C)} \sigma_i(a \mid s) \geq \epsilon \right\}$$

We also denote $\Delta_\epsilon = \bigcap_{S \in \max \text{SC}} \Delta_\epsilon(C)$.

Reduced State space

Definition

For any $C \in \text{SC}$ strong component and $\epsilon > 0$, we define


$$\Delta_\epsilon(C) = \left\{ \sigma \in \mathbb{M} \mid \sum_{(a,i,s) \in \text{Exit}(C)} \sigma_i(a \mid s) \geq \epsilon \right\}$$

We also denote $\Delta_\epsilon = \bigcap_{S \in \max \text{SC}} \Delta_\epsilon(C)$.


Lemma

For $\epsilon \leq \frac{1}{|\text{Act}|}$, $\Delta_\epsilon \neq \emptyset$. It is also convex.

Limit behaviour



Limit behaviour



$$\sigma_1(b | s_1) + \sigma_2(b | s_2) \geq \epsilon$$

Bounding probability of termination

Lemma

*If $\sigma \in \Delta_\epsilon$, then $\Pr^\sigma(\text{States}^*F \mid s) = 1$*

Bounding probability of termination

Lemma

If $\sigma \in \Delta_\epsilon$, then $\Pr^\sigma(\text{States}^*F \mid s) = 1$

Theorem

For $\epsilon > 0$, there exists $p > 0$ and $k \in \mathbb{N}$ such that for any $\sigma \in \Delta_\epsilon$,

$$\forall s \in \text{States} \mathbb{P}^\sigma(\text{States}^k \cdot F \mid s) \geq p$$

That is to say, after k iterations, there is a bounded probability that a final state is reached.

Existence theorem

Definition (Best response function)

Let $BR_\epsilon : \Delta_\epsilon \rightarrow 2^{\Delta_\epsilon}$ with

$$BR_\epsilon(\sigma) = \left\{ \sigma' \in \Delta_\epsilon \mid \forall i \in \text{Agt} \forall \mathbf{s} \in \text{States}, \sigma'_i \in \operatorname{argmax}_{\sigma'_i} \mathbb{E}^{\sigma^{[i/\sigma'_i]}}(\phi_i \mid \mathbf{s}) \right\}$$

Existence theorem

Definition (Best response function)

Let $BR_\epsilon : \Delta_\epsilon \rightarrow 2^{\Delta_\epsilon}$ with

$$BR_\epsilon(\sigma) = \left\{ \sigma' \in \Delta_\epsilon \mid \forall i \in \text{Agt} \forall s \in \text{States}, \sigma'_i \in \operatorname{argmax}_{\sigma'_i} \mathbb{E}^{\sigma^{[i/\sigma'_i]}}(\phi_i \mid s) \right\}$$

Theorem

For $0 < \epsilon \leq \frac{1}{|\text{Act}|}$, BR_ϵ has a fixed point.

Existence theorem

Definition (Best response function)

Let $BR_\epsilon : \Delta_\epsilon \rightarrow 2^{\Delta_\epsilon}$ with

$$BR_\epsilon(\sigma) = \left\{ \sigma' \in \Delta_\epsilon \mid \forall i \in \text{Agt} \forall s \in \text{States}, \sigma'_i \in \operatorname{argmax}_{\sigma'_i} \mathbb{E}^{\sigma^{[i/\sigma'_i]}}(\phi_i \mid s) \right\}$$

Theorem

For $0 < \epsilon \leq \frac{1}{|\text{Act}|}$, BR_ϵ has a fixed point.

Proof sketch.

Δ_ϵ is a non-empty compact convex subset of \mathbb{R}^N where $N = \text{Agt} \times \text{States}$. Moreover $BR_\epsilon(\sigma)$ is a non-empty convex set and the graph of BR_ϵ is continuous. □

ϵ -robust equilibria

Definition (robust equilibria)

Let $\sigma \in \mathbb{M}$, σ is a ϵ -robust Nash Equilibrium if for any player i ,

$$\forall \sigma'_i \exists \sigma''_i \ d(\sigma'_i, \sigma''_i) \leq \epsilon \ \mathbb{E}^{\sigma''_i}(\phi_i | h) \leq \mathbb{E}^{\sigma}(\phi_i | h)$$

with $d(\sigma, \sigma')$ the maximal distance between distributions.

ϵ -robust equilibria

Definition (robust equilibria)

Let $\sigma \in \mathbb{M}$, σ is a ϵ -robust Nash Equilibrium if for any player i ,

$$\forall \sigma'_i \exists \sigma''_i d(\sigma'_i, \sigma''_i) \leq \epsilon \mathbb{E}^{\sigma''_i}(\phi_i | h) \leq \mathbb{E}^{\sigma}(\phi_i | h)$$

with $d(\sigma, \sigma')$ the maximal distance between distributions.

- 1 $\sigma \in \Delta_{\alpha\epsilon}(\sigma)$ is a ϵ -robust NE
- 2 This is not a NE (but the converse is false)
- 3 This is not a ϵ -NE
- 4 This is stationary (stationary NE may not exist for 3 players.)

ϵ -robust equilibria

Definition (robust equilibria)

Let $\sigma \in \mathbb{M}$, σ is a ϵ -robust Nash Equilibrium if for any player i ,

$$\forall \sigma'_i \exists \sigma''_i d(\sigma'_i, \sigma''_i) \leq \epsilon \mathbb{E}^{\sigma''_i}(\phi_i | h) \leq \mathbb{E}^{\sigma}(\phi_i | h)$$

with $d(\sigma, \sigma')$ the maximal distance between distributions.

- 1 $\sigma \in \Delta_{\alpha\epsilon}(\sigma)$ is a ϵ -robust NE
- 2 This is not a NE (but the converse is false)
- 3 This is not a ϵ -NE
- 4 This is stationary (stationary NE may not exist for 3 players.)
- 5 No computation method yet

Overview

- Getting closer to exact NE existence problem (2 players)

Overview

- Getting closer to exact NE existence problem (2 players)
- Using of linear constrains to enforce a non-linear property

Overview

- Getting closer to exact NE existence problem (2 players)
- Using of linear constrains to enforce a non-linear property
- New notion of equilibria, not equivalent to previous ones

Overview

- Getting closer to exact NE existence problem (2 players)
- Using of linear constrains to enforce a non-linear property
- New notion of equilibria, not equivalent to previous ones
- Non-constructive proof (ongoing work)

Thank you for your attention

Questions ?

Bibliography I

Krishnendu Chatterjee, Marcin Jurdziński, and Rupak Majumdar. On Nash equilibria in stochastic games. In *CSL'04*, volume 3210 of *LNCS*, pages 26–40. Springer, 2004.

John F. Nash. Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences of the United States of America*, 36(1): 48–49, 1950.

Piercesare Secchi and William D. Sudderth. Stay-in-a-set games. *International Journal of Game Theory*, 30:479–490, 2001.