Subgame perfect equilibrium with an algorithmic perspective

Jean-François Raskin (Université libre de Bruxelles)
joint work with Marie van den Bogaard (Université Gustave Eiffel)
Léonard Brice (ENS Paris-Saclay)

March 26, 2021.

Simons Institute for the Theory of Computing (Berkeley)

Objectives of the talk

① Recall SPE for sequential games

② Expose some recent progresses in algorithms to handle SPEs (for mean-payoff objectives) → left open in the literature

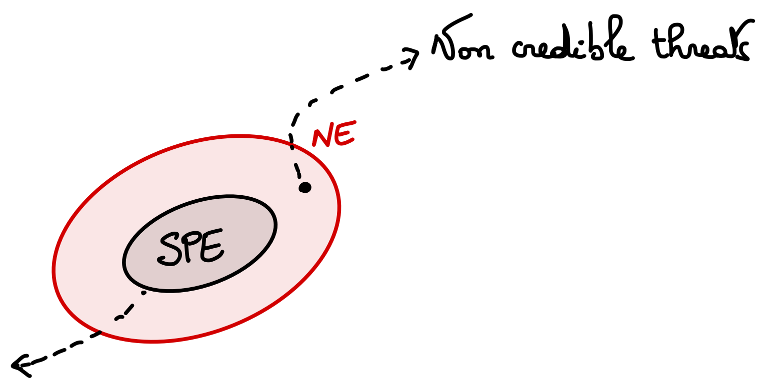

Q: how to obtain an effective representation of all possible outcomes produced by rational players?

N player turn-based graph games

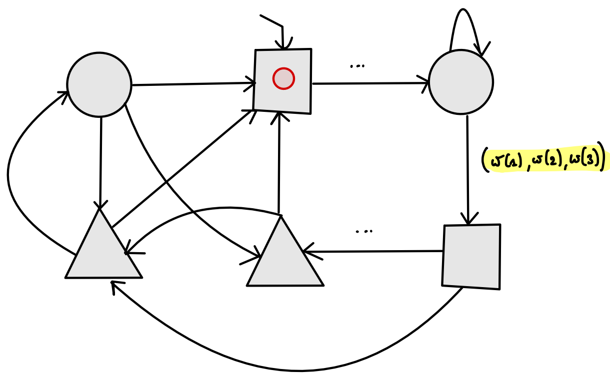

Vertices are partitioned

$$V = V_1 \uplus V_2 \uplus \dots \uplus V_N$$

V_i = vertices of Player $i \in [1, N]$

$$E \subseteq V \times V \;,\; w : E \to \mathbb{Z}^N$$

Play $\rho \in V^\omega$, payoffs given by $\mu_i : V^\omega \to \mathbb{R}$, $i \in [1, N]$

↳ mean-payoff $\mu_i(\rho)$

$$\underline{MP}_i(\rho) = \liminf_{n \to \infty} \frac{1}{n} \cdot \sum_{j=0}^{n-1} w(v_j, v_{j+1})(i)$$

N player turn-based graph games

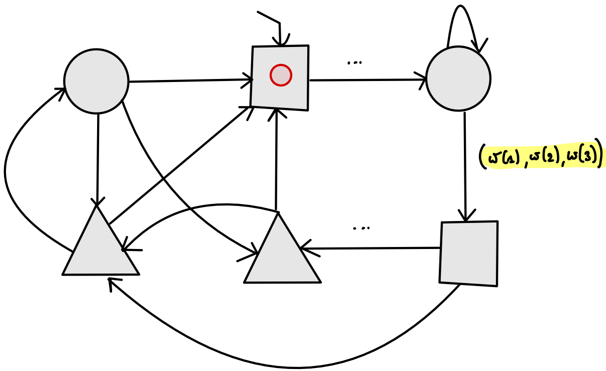

Vertices are partitioned

$$V = V_1 \uplus V_2 \uplus \dots \uplus V_N$$

V_i = vertices of Player $i \in [1, N]$

$E \subseteq V \times V$, $w : E \to \mathbb{Z}^N$

Play $\rho \in V^\omega$, payoffs given by $\mu_i : V^\omega \to \mathbb{R}$, $i \in [1, N]$

↳ mean-payoff $\mu_i(\rho)$

$$MP_i(\rho) = \liminf_{n \to \infty} \frac{1}{n} \cdot \sum_{j=0}^{n-1} w(v_j, v_{j+1})(i)$$

Rationality : each player wants to maximize his own payoff.

Strategies, profiles, outcomes

$$\sigma_i : V^* \cdot V_i \to E$$

s.t. $\forall \pi \cdot v \in V^* \cdot V_i : (v, \sigma_i(\pi \cdot v)) \in E$

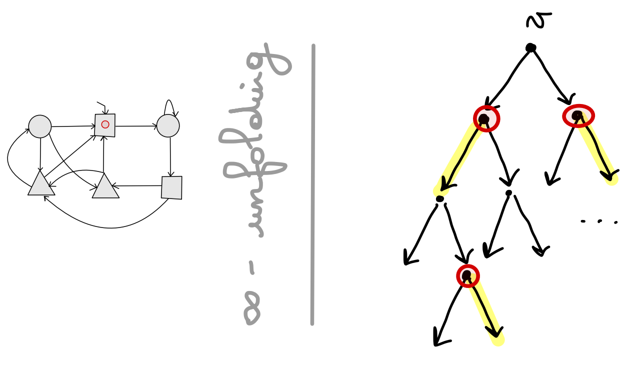

Σ_i = set of strategy of Pl. $i \in [1, N]$

Strategies, profiles, outcomes

$\sigma_i : V^* \cdot V_i \to E$

s.t. $\forall \pi \cdot v \in V^* \cdot V_i : (v, \sigma_i(\pi \cdot v)) \in E$

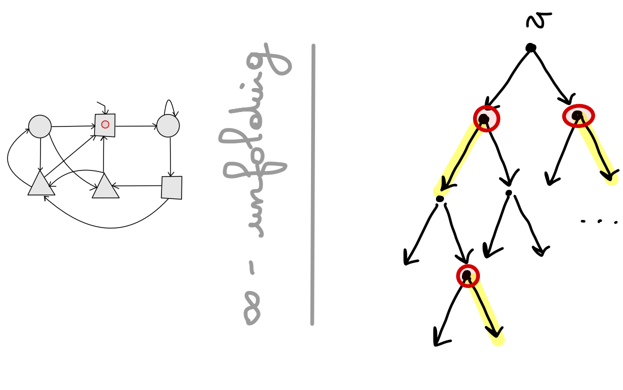

Σ_i = set of strategy of Pl. $i \in [1, N]$

Profiles:

$(\sigma_1, \sigma_2, \dots, \sigma_N) \in \Sigma_1 \times \Sigma_2 \times \dots \times \Sigma_n$

$= (\sigma_i, \bar{\sigma}_{-i})$

all strategies but σ_i

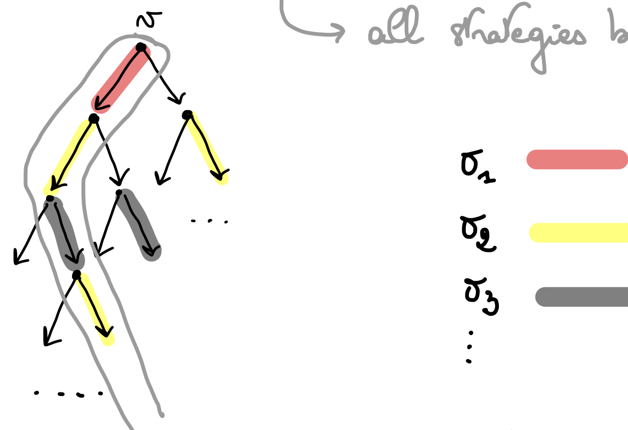

$= Out_v(\sigma_1, \sigma_2, \dots, \sigma_N) = \rho \in V^\omega$

$Out_{v_0}(\bar{\sigma}) = v_0 v_1 v_2 \dots v_n \dots = \rho$

s.t. $v_0 = v$, $\forall j \geq 0$: if $\rho(j) \in V_i$: $v_{j+1} = \sigma_i(\rho(0..j))$

Nash equilibrium and subgame perfect equilibrium

A profile of strategies $(\sigma_1, \sigma_2, \dots, \sigma_N)$ is a Nash equilibrium (NE) in v_0, if for all $i \in [1, N]$, for all $\sigma_i' \in \Sigma_i$: $\mu_i(\mathrm{Out}_{v_0}(\bar{\sigma}_{-i}, \sigma_i')) \leqslant \mu_i(\mathrm{Out}_{v_0}(\bar{\sigma}_{-i}, \sigma_i))$

= No player has an incentive to deviate unilaterally.

Nash equilibrium and subgame perfect equilibrium

A profile of strategies $(\sigma_1, \sigma_2, \dots, \sigma_N)$ is a Nash equilibrium (NE) in v_0, if for all $i \in [1, N]$, for all $\sigma_i' \in \Sigma_i$: $\mu_i(\mathrm{Out}_{v_0}(\bar{\sigma}_{-i}, \sigma_i')) \leq \mu_i(\mathrm{Out}_{v_0}(\bar{\sigma}_{-i}, \sigma_i))$

= No player has an incentive to deviate unilaterally.

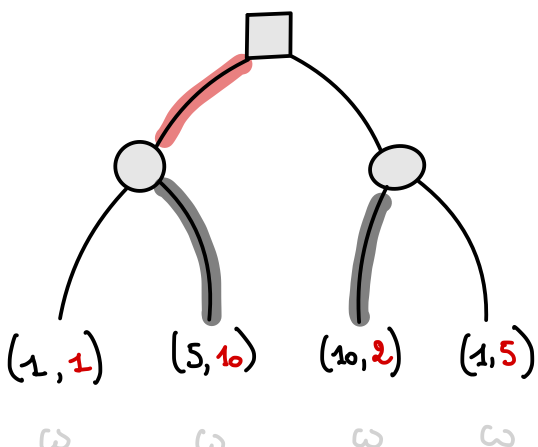

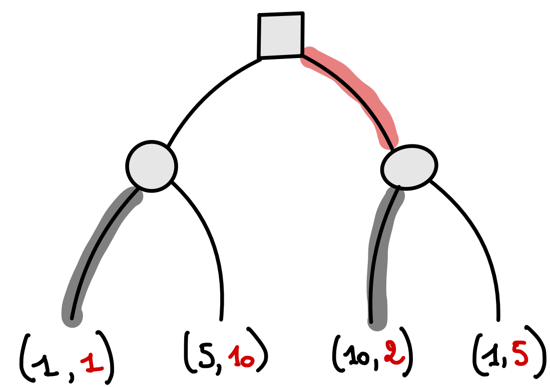

Nash equilibrium and subgame perfect equilibrium

A profile of strategies $(\sigma_1, \sigma_2, \dots, \sigma_N)$ is a Nash equilibrium (NE) in v_0, if for all $i \in [1, N]$, for all $\sigma_i' \in \Sigma_i$: $\mu_i(\mathrm{Out}_{v_0}(\sigma_{-i}, \sigma_i')) \leqslant \mu_i(\mathrm{Out}_{v_0}(\sigma_{-i}, \sigma_i))$

= No player has an incentive to deviate unilaterally.

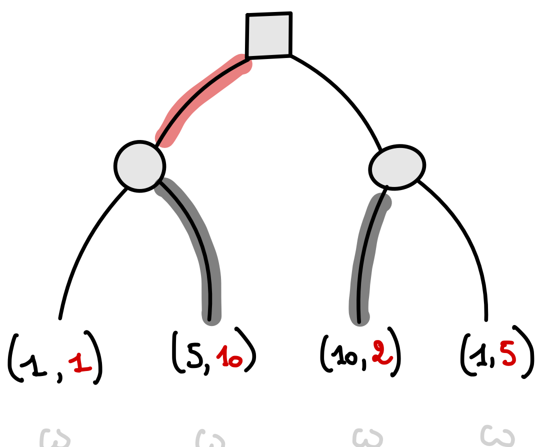

non credible threat

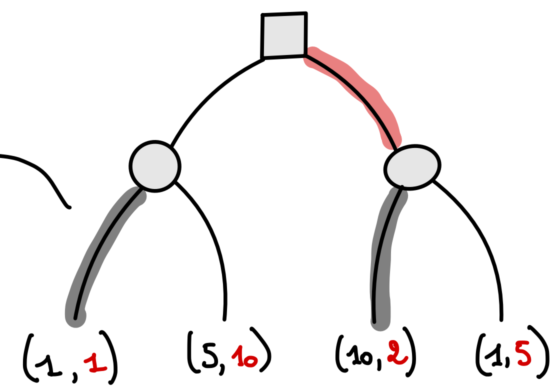

Pl. is not rational in this subgame!

Subgame

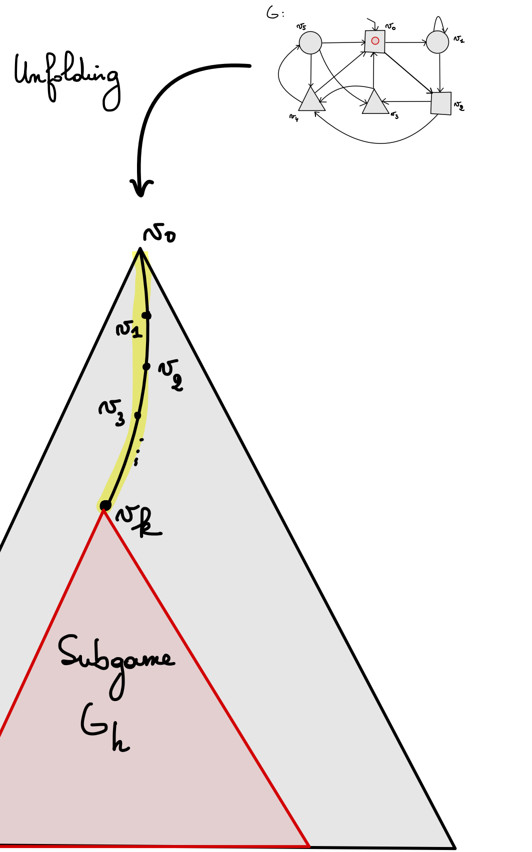

Every history $h = v_0 v_1 \dots v_k$ defines a subgame G_h

Subgame

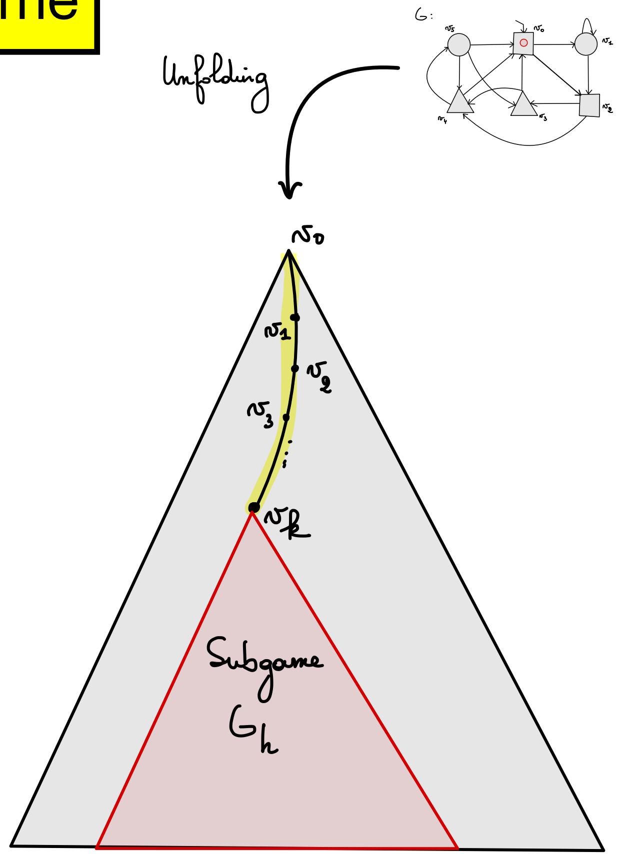

Every history $h = v_0 v_1 \dots v_k$ defines a <u>subgame</u> G_h

ex:

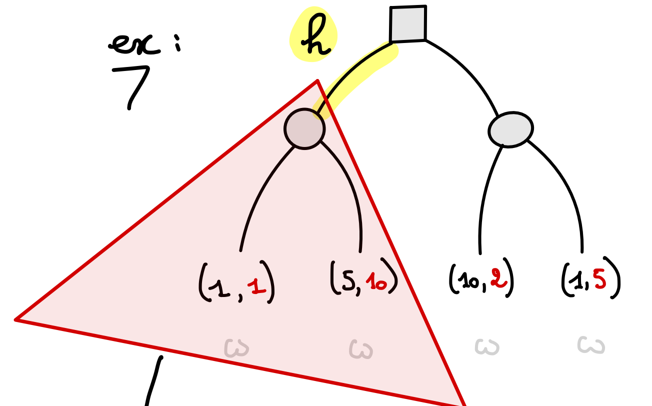

Subgame induced by h

for $h = h' \cdot v$,

$$\mathrm{Out}_h(\bar{\sigma}) = h' \cdot \mathrm{Out}_v(\bar{\sigma}^{h})$$

Nash equilibrium and subgame perfect equilibrium

A profile of strategies $(\sigma_1, \sigma_2, \dots, \sigma_N)$ is a subgame perfect equilibrium (SPE) if for all subgames G_h of G, for all Pl. $i \in [1,N]$, for all $\sigma_i^{h\prime} \in \Sigma_i$:

$$\mu_i(\mathrm{Out}_h(\sigma_{-i}^h, \sigma_i^{h\prime})) \leqslant \mu_i(\mathrm{Out}_h(\sigma_{-i}^h, \sigma_i^h)).$$

→ Players must be rational in all subgames (→ no non-credible threats)

Nash equilibrium and subgame perfect equilibrium

A profile of strategies $(\sigma_1, \sigma_2, \dots, \sigma_N)$ is a **subgame perfect equilibrium (SPE)** if for all subgames G_h of G, for all Pl. $i \in [1,N]$, for all $\sigma_i^{h'} \in \Sigma_i$:

$$\mu_i\left(\mathrm{Out}_h\left(\bar{\sigma}^h_{-i}, \sigma_i^{h'}\right)\right) \leq \mu_i\left(\mathrm{Out}_h\left(\bar{\sigma}^h_{-i}, \sigma_i^{h}\right)\right).$$

→ Players must be rational in all subgames (→ no non-credible threats)

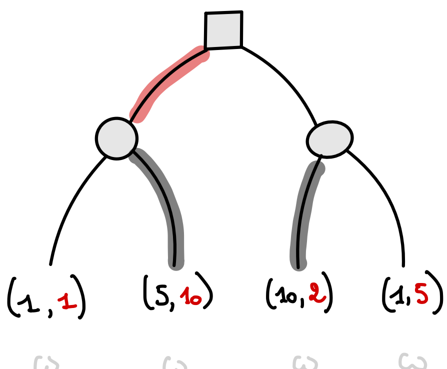

SPE

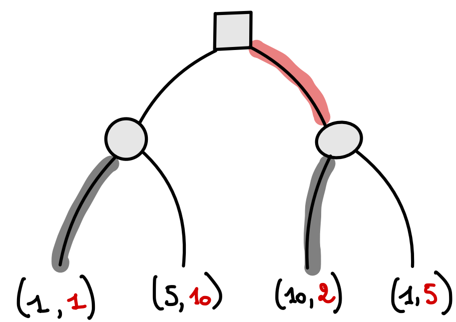

~~SPE~~

Nash equilibrium and subgame perfect equilibrium

A profile of strategies $(\sigma_1, \sigma_2, \dots, \sigma_N)$ is a **subgame perfect equilibrium (SPE)** if for all subgames G_h of G, for all PP. $i \in [1,N]$, for all $\sigma_i^{h\prime} \in \Sigma_i$:

$$\mu_i\left(\mathrm{Out}_h\left(\bar{\sigma}^h_{-i}, \sigma_i^{h\prime}\right)\right) \leqslant \mu_i\left(\mathrm{Out}_h\left(\bar{\sigma}^h_{-i}, \sigma_i^{h}\right)\right).$$

→ Players must be rational in all subgames (→ no non-credible threats)

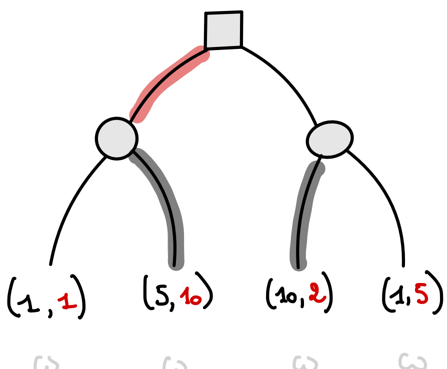

SPE

in this subgame PP. ◯ has an incentive to deviate
↓
non credible threat

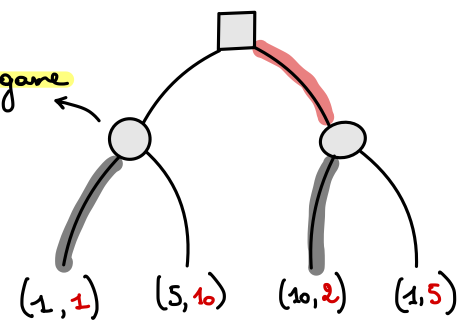

~~SPE~~

Outcomes supported by equilibria

$$OutNE(G) = \bigcup_{\bar{\sigma} \in NE} \{Outcome_{v_0}(\bar{\sigma})\}$$

$$OutSPE(G) = \bigcup_{\bar{\sigma} \in SPE} \{Outcome_{v_0}(\bar{\sigma})\}$$

only supported by non-credible threats

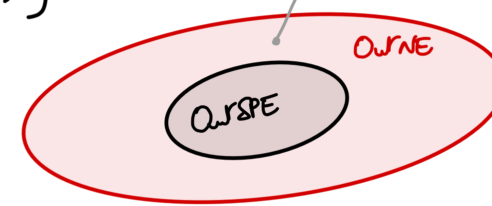

↳ Sets of behaviors induced by NE / SPE - rational players in G

SPE - Known results

Existence

- guaranteed when $(\mu_i)_{i \in [1,N]}$ are continuous (ex: discounted sum, quantitative reach.)

  can be extended to lower semi-continuous [Flesch et al.]

  or $(\mu_i)_{i \in [1,N]}$ are B-omega-regular objectives (ex: parity)

  Established by Ummels (2006) using transfinite induction

- not guaranteed for mean-payoff objectives (see example)

Effective Representation

OutSPE(G)

- for quantitative reachability [CONCUR'19]
- ... also for B-omega regular obj. (ex: parity) [Ummels 06]

  alternating tree automata

- TODAY: mean-payoff objectives [arXiv : 2101.10685]

Effective representation of the set of outcomes of NE/SPE

Why?

Effective representation of the set of outcomes of NE/SPE

why ?

1 Existence problem for SPE : $\mathrm{Out\text{-}SPE}(G) \overset{?}{\neq} \emptyset$

Effective representation of the set of outcomes of NE/SPE

why?

(1) Existence problem for SPE : $\mathrm{Out.SPE}(G) \overset{?}{\neq} \phi$

(2) Quantitative rational verification [Wooldridge et al. 17] for NE.

B, NE

→ ∀ ∃ NE/SPE : do all Some behaviors by rational agents satisfy some spec. ψ?

all: $\mathrm{Out.NE}(G) \overset{?}{\subseteq} [\![\psi]\!]$ $\mathrm{Out.SPE}(G) \overset{?}{\subseteq} [\![\psi]\!]$

Some: $\mathrm{Out.NE}(G) \cap [\![\psi]\!] \overset{?}{\neq} \phi$ $\mathrm{Out.SPE}(G) \cap [\![\psi]\!] \overset{?}{\neq} \phi$

→ idem for $M_{\sigma_0} \otimes G_{0 \cup [1,N]}$

System , [1,N] = Env.

Effective representation of the set of outcomes of NE/SPE

why ?

① Existence problem for SPE : $\mathrm{Out}\text{-}\mathrm{SPE}(G) \overset{?}{=} \emptyset$

② Quantitative rational verification [Wooldridge et al. 17] for NE.

→ $\forall$ / $\exists$ NE/SPE : do all / some behaviors by rational agents satisfy some spec. ψ ?

all: $\mathrm{Out}\text{-}\mathrm{NE}(G) \overset{?}{\subseteq} [\![\psi]\!]$ $\qquad$ $\mathrm{Out}\text{-}\mathrm{SPE}(G) \overset{?}{\subseteq} [\![\psi]\!]$

some: $\mathrm{Out}\text{-}\mathrm{NE}(G) \cap [\![\psi]\!] \overset{?}{\neq} \emptyset$ $\qquad$ $\mathrm{Out}\text{-}\mathrm{SPE}(G) \cap [\![\psi]\!] \overset{?}{\neq} \emptyset$

→ idem for $M_{\sigma_0} \otimes G_{0 \cup [1,N]}$

↳ System , $[1,N]$ = Env.

③ Cooperative rational synthesis : $\exists \rho \in \mathrm{Out}\text{-}\left\{\begin{smallmatrix}\mathrm{SPE}\\ \mathrm{NE}\end{smallmatrix}\right\}(G) : \rho \models \varphi_0$

[Kupferman et al. 16]

(first seminar)

↳ constrained-$\exists$ $\qquad$ ↳ Spec. synt.

How to reason on SPE ?

→ finite trees : conceptually easy using backward induction

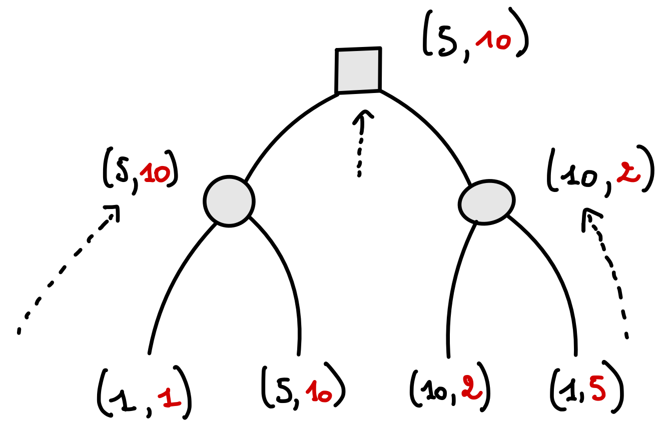

→ infinite trees : backward induction does not generalize well...

Starting point:
NE in infinite duration games

NE - Deviation - Punishment

$\rho \in NE$

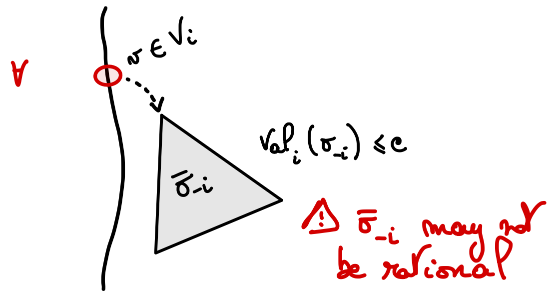

$\rho \notin NE$

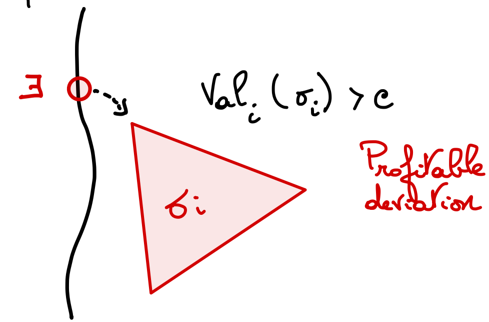

$Val_i(\rho) = c$ and

$c \geq \inf_{\bar{\sigma}_{-i}} \sup_{\sigma_i} Val(Out_v(\sigma_i, \bar{\sigma}_{-i}))$

$Val_i(\rho) = c$ and

$\sup_{\sigma_i} \inf_{\bar{\sigma}_{-i}} Val(Out_v(\sigma_i, \bar{\sigma}_{-i})) > c$

$=$ worst-case value for P_i

→ This works for any prefix independent payoff function for which worst-case value exists and can be realized

[Brihaye et al. 13]

→ OK for MP

Set of outcomes supported by NE - MP

→ Requirement : $\lambda : V \to \mathbb{R} \cup \{-\infty, +\infty\}$

→ A path $\rho = v_0 v_1 \dots v_n \dots$ is **λ-consistent** if

$$\forall i \in [1,N] : \underline{MP}_i(\rho) \geq \max_{v \in Visit(\rho) \cap V_i} \lambda(v_i)$$

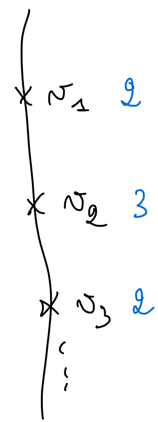

Set of outcomes supported by NE - MP

→ Requirement : $\lambda : V \to \mathbb{R} \cup \{-\infty, +\infty\}$

→ A path $\rho = v_0 v_1 \dots v_n \dots$ is λ-consistent if

$$\forall i \in [1, N] : \underline{MP}_i(\rho) \geq \max_{v \in Visit(\rho) \cap V_i} \lambda(v_i)$$

→ Key : the worst-case value that Pl. i can force against all the other players.

→ $\forall v \in V_i$: let $\lambda_1(v) = \sup_{\sigma_i \in \Sigma_i} \inf_{\sigma_{-i} \in \Sigma_{-i}} \underline{MP}_i(Out_v(\sigma_i, \sigma_{-i}))$ 1st requirement induced by WCV

Theorem: $\rho = v_0 v_1 \dots v_n \dots$, $\rho \in OutNE(G)$ iff ρ is λ_1-consistent.

↳ ρ gives Pl. i at least the wcv that he can force along ρ → if not → profitable deviation.

Set of outcomes supported by NE - MP

- an example

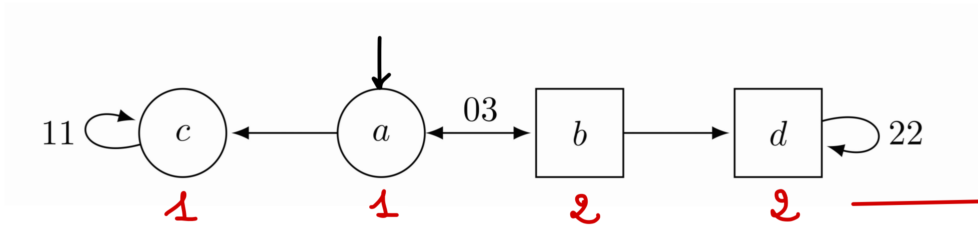

λ_1

worst-case values

↳ The set of λ_1-consistent paths in G are :

$$\{a \to c^\omega\} \cup \bigcup_{k\in\mathbb{N}} \{a \to (b\to a)^k \to b \to d^\omega\}$$

Corollary : The set of λ-consistent paths is recognized by a multi-MP automaton. This language is not necessarily ω-regular.

[Alur et al. 09, Chatterjee et al. 10]

A MP game without SPE

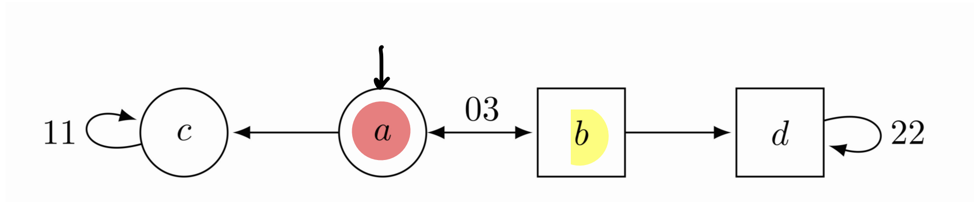

- → Pl. ○ can secure 1 from a (a → c)
- → Pl. □ can secure 2 from b (b → d)
- → So there is no NE in which a ↔ b is taken for ever as Pl. ○ would have an incentive to leave (a → c) but then Pl. □ would prefer to leave before Pl. ○

A MP game without SPE

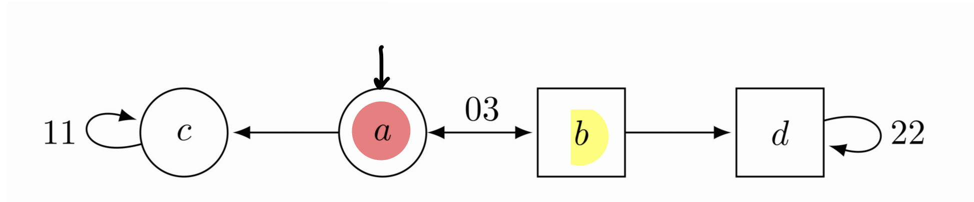

→ Pl. ○ can secure 1 from a (a → c)

→ Pl. □ can secure 2 from b (b → d)

→ So there is no NE in which a ↔ b is taken for ever as Pl. ○ would have an incentive to leave (a → c) but then Pl. □ would prefer to leave before Pl. ○

So → From a Pl. ○ knows that Pl. □ will leave, Pl. ○ has then no incentive to do it before (as he will then get 2 instead of 1)

But → then Pl. □ has no interest to leave as he receives 3 on the cycle.

→ Need to iterate the reasoning on worst-case value.

Generalization :
The negotiation function

? : Given λ_1 and v, can the player that controls v improve the value that she can obtain against the other players if the other players are not willing to give away their worst-case value (λ_1)?

The negotiation function

? : Given λ_1 and v, can the player that controls v improve the value that she can obtain against the other players if the other players are not willing to give away their worst-case value (λ_1)?

$$\text{Nego} : [\lambda \to \mathbb{R} \cup \{-\infty, +\infty\}] \to [\lambda \to \mathbb{R} \cup \{-\infty, +\infty\}]$$

$\inf \emptyset = +\infty$

let $v \in V_i$, $\text{Nego}(\lambda)(v) = \inf\limits_{\bar{\sigma}_{-i} \in \lambda\text{Rat}(v)} \ \sup\limits_{\sigma_i \in \Sigma_i} \ \underline{\text{MP}}_i\left(\text{Out}_v(\sigma_i, \bar{\sigma}_{-i})\right)$

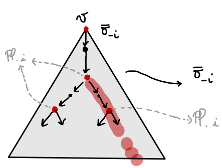

$\bar{\sigma}_{-i}$ is $\lambda\text{Rat}(v)$: in all $\bar{\sigma}_{-i}$ consistent history h, there is a $\bar{\sigma}_{-i}$ continuation that is λ-consistent.

% worst case value against λ-rational players.

How to compute Nego(.) ?

$v \in V_i$

$$\mathrm{Nego}(\lambda)(v) = \inf_{\bar{\sigma}_{-i} \in \lambda\mathrm{Rat}(v)} \sup_{\sigma_i \in \Sigma_i} \underline{MP}_i\left(\mathrm{Out}_v(\sigma_i, \bar{\sigma}_{-i})\right)$$

if the inf is always realizable → STEADY NEGOCIATION

↳ ok for $\underline{MP}$

What is the relative worst-case value that Pl. i can force against λ- rational adversaries ?

Prover want to prove $\mathrm{Nego}(\lambda)(v) \leq \alpha$ to Challenger

P

C

P C game to determine if $\mathcal{N}_{ego}(\lambda)(v) \leq \alpha$?

P proposes outcomes $\rho = v_0 v_1 \dots v_n \dots$ $\in \lambda$-consistent if this is possible

$v_0 = v$

(if not possible FAIL)

C either accepts and the game ends (ACCEPT)
or proposes a deviation after a history h that belongs to Pl. i (DEV)

$\rho = v_0 v_1 \dots v_k v_{k+1} \dots$ $\longrightarrow$ current prefix $\pi = v_0 v_1 \dots v_k v'$

$v_k \in V_i$, $v' \neq v_{k+1}$

P proposes an outcome $\rho = \pi\, v_0 v_1 \dots v_k \dots$ that extends π and such that $v_0 v_1 \dots v_k \dots$ is λ-consistent

(if not possible FAIL)

⋮

$\boxed{P}\ \boxed{C}$ game — winning condition

$= \underline{MP}_i(\rho')$

either $\boxed{C}$ puts an end to deviations $\rightarrow \rho = \pi_1 \cdot \pi_2 \cdot \ldots \cdot \pi_n \cdot \rho'$ and Pl. i receives $\underline{MP}_i(\rho)$

(ACCEPT)

or $\boxed{C}$ deviates ∞-many times and the resulting play is $\rho^\infty = \pi_1 \cdot \pi_2 \cdot \ldots \cdot \pi_n \ldots$ and Pl. i receives $\underline{MP}_i(\rho^\infty)$. (DEV $^\infty$)

or $\boxed{P}$ fails at some round to propose a λ-consistent continuation (FAIL) and Pl. i receives $+\infty$.

$\boxed{P}$ wins iff Pl. i receives a payoff $\leq \alpha$.

(This zero-sum game is determined)

How to compute Nego(.) ?

– an example – $Nego(\lambda_1)$

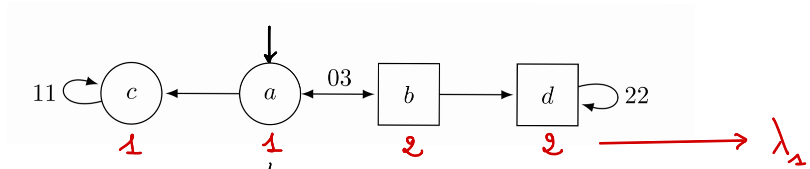

$Nego(\lambda_1)(a) \leq 1$?

$\mathbb{P}$: $a \cdot c^\omega$, $a \cdot c^\omega$ is λ_1-consistent and $\underline{MP}_{\bigcirc}(a \cdot c^\omega) = 1$

$\mathbb{C}$: deviation : $a \to b$

$\mathbb{P}$: from b, the only λ_1-consistent paths are $(b\,a)^* \cdot d^\omega$

even if $(a\,b)^\omega$ is tempting

it does not give 1 to $\bigcirc$

and $\underline{MP}_1((b\,a)^* d^\omega) = 2$ $\Rightarrow$ $\mathbb{C}$ wins

$\cancel{Nego(\lambda_1)(a) \leq 1}$

g

→ generalisation : $Nego(\lambda_1)(a) = 2$

How to compute Nego(.) ?

– an example – $\mathrm{Nego}(\lambda_2)$

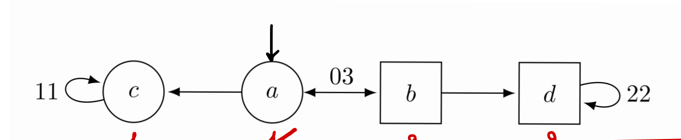

$1 \quad \cancel{1}\ 2 \quad 2 \quad 2 \longrightarrow \lambda_2$

$\hookrightarrow \mathrm{Nego}(\lambda_2)(b) \leq 2$?

$\mathbb{P}$: $b \cdot d^{\omega}$

$\mathbb{C}$: deviation $b \to a$

$\mathbb{P}$: $(a \cdot b)^{*} \cdot d^{\omega}$

$\mathbb{C}$: deviation $b \to a$

ω

$\rightarrow$ outcome : $(b\,a)^{\omega}$ and $\underline{MP}_{\square}\left((b\,a)^{\omega}\right) = 3$

$\rightarrow$ $\cancel{\mathrm{Nego}(\lambda_2)(b) \leq 2}$

$\rightarrow$ generalisation : $\mathrm{Nego}(\lambda_2)(b) = 3$

How to compute Nego(.) ?

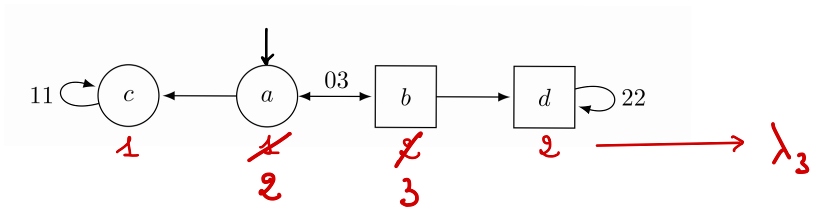

There is no λ_3- consistent path from a nor from b !

(FAIL)

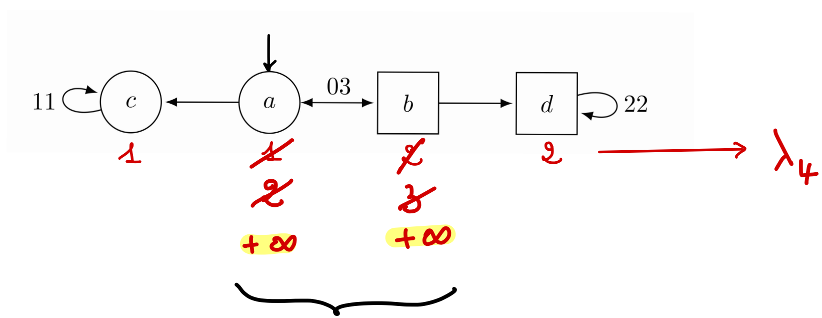

no SPE starting from a or b !

Properties of the negociation function

Let λ_0 be s.t. $\lambda_0(v) = -\infty$, $\forall v \in V$. & no constraint !

Then $\mathrm{Nego}(\lambda_0)(v) = \lambda_1(v) = \inf_{\sigma_{-i} \in \Sigma_{-i}} \sup_{\sigma_i \in \Sigma_i} \underline{MP}_i(\mathrm{Out}_v(\sigma_i, \sigma_{-i}))$

all profiles are λ_0-rational

$$= \sup_{\sigma_i \in \Sigma_i} \inf_{\sigma_{-i} \in \Sigma_{-i}} \underline{MP}_i(\mathrm{Out}_v(\sigma_i, \sigma_{-i}))$$

"worst-case value"

Theorem. $\mathrm{Nego}(\lambda_0)$ characterizes NEs !

Properties of the negociation function

Let λ^* be s.t. Nego$(\lambda^*) = \lambda^*$, i.e. λ^* is a fixed point of Nego.

Lemma 1. $\forall$ λ^*-consistent paths ρ, $\exists\, \bar{\sigma} \in SPE$: $\rho = Out(\bar{\sigma})$.

Lemma 2. $\forall \bar{\sigma} \in SPE$: $\exists\, \lambda^*$ s.t. Nego$(\lambda^*) = \lambda^*$ and $Out(\bar{\sigma})$ is λ^*-consistent.

Theorem. The set of fixed points of the function Nego is a characterization of outcomes of SPEs.

Because Nego is monotone and the set of requirements form a complete lattice and in addition the set of λ-consistent paths is upward-closed then we have the following stronger result:

Corollary. The set of outcomes of SPEs is characterized by the LFP of Nego.

LFP and SPE

ρ is λ^* consistent $\iff$ $\rho \in SPE$

ρ' is choosen such that ρ' minimize the payoff of player ○ using λ^*- profiles

↳ such ρ' exists by hyp. of steady negociations.

and $\underline{MP}_i(\rho') \leq \underline{MP}_i(\rho)$

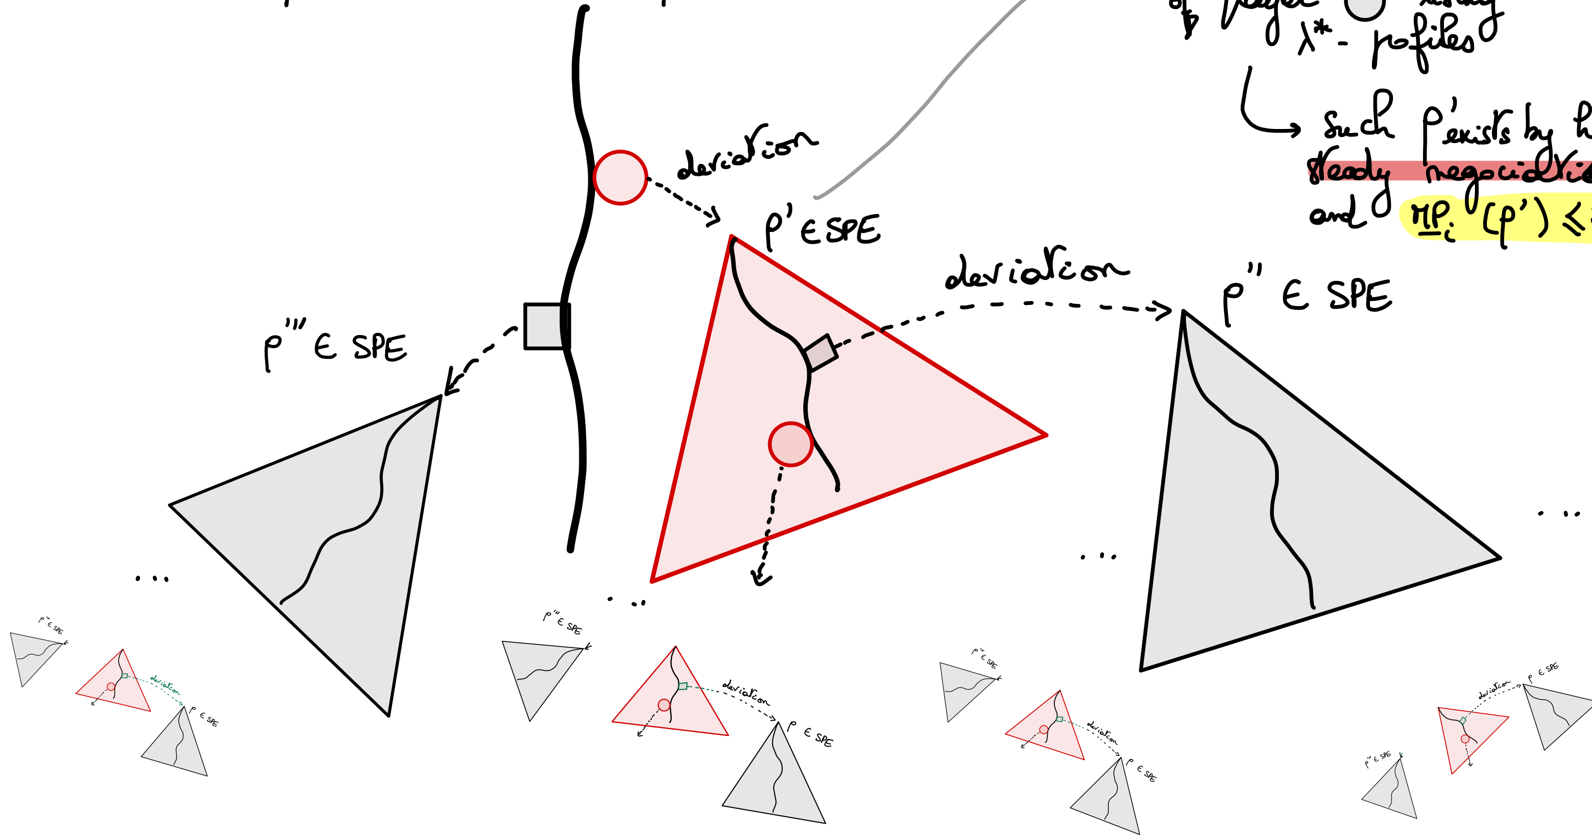

Additional properties

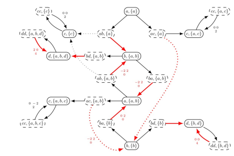

Fig. 10. A concrete negotiation game

1. We can transform the P C game into a finite state multi-mean payoff game (ic'15)

2. This multi-mean payoff game allows us to effectively compute Nego(.)

3. λ^* may not be reached from λ_0 by Kleene-Tarski iteration in finitely many steps

Non finite convergence

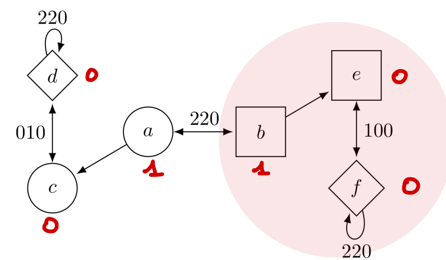

λ_1

$\textcolor{red}{\text{Nego}}(\lambda_1)(a) = 1\,1/2$

P cannot propose to go to the left with a value $\leq 1\,1/2$ as C would deviate to "b". From "b" P can propose on the right part $1\,1/2$ to Pl. ○ (as he needs to give at least 1 to Pl. □)

Symmetrically: $\textcolor{red}{\text{Nego}}(\lambda_1)(b) = 1\,1/2$

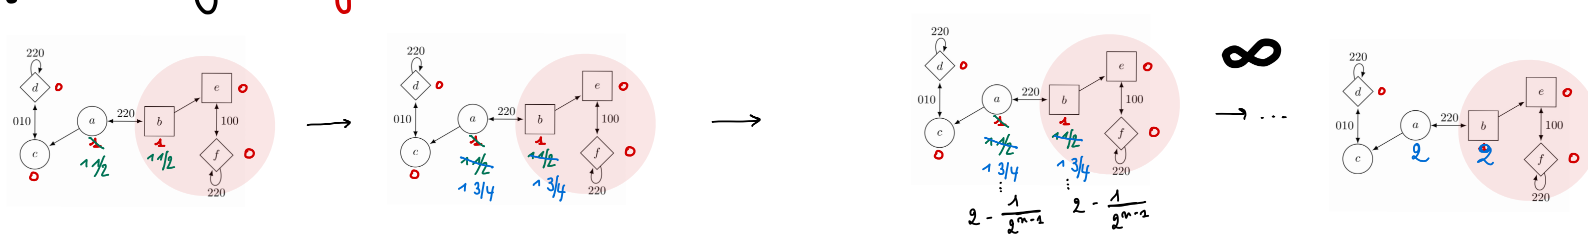

Additional properties

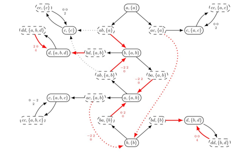

Fig. 10. A concrete negotiation game

① We can transform the $\boxed{P}\boxed{C}$ game into a finite state multi-mean payoff game (ic'15)

② This multi-mean payoff game allows us to effectively compute $\boxed{\text{Nego}(.)}$

③ λ^* may not be reached from λ_0 by Kleene-Tarski iteration in finitely many steps

④ ... But thanks to good properties of multi-mean payoff games, we can show that $\boxed{\text{Nego}(.)}$ is effectively piecewise linear and λ^* can be obtained using linear algebraic techniques.

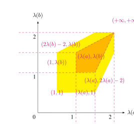

Conclusions and perspectives

→ SPE provides a natural notion of rational behaviors in infinite duration games played on graphs

→ Worst-case value relative to rational adversary formalized by the fixed points of Nego(·) leads to an effective representation of OutSPE(G) for MP games (multi mean-payoff automata)

characterized by worst-case value

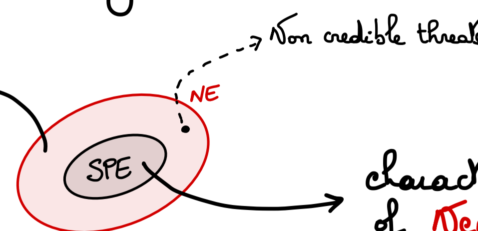

characterized by fixed-points of Nego(·)

Conclusions and perspectives

→ Nego(·) is also applicable to parity games (omega regular obj.)

↳ useful to close complexity gaps

ex : Constrained existence for SPE is in ExpTime (emptiness of alternating automata) and NP-hard. [Ummels '06]

→ Our previous algorithm for quantitative reachability can be rephrased with Nego(·) [CONCUR '19]

→ Nego(·) provides a new algorithmic basis to do rational verification and synthesis based on SPEs.