Nash equilibria in games on graphs with a public signal monitoring

Patricia Bouyer

LSV, CNRS & ENS Paris-Saclay
Université Paris-Saclay, Cachan, France

What this talk is about

- pure Nash equilibria in game graphs
- imperfect information monitoring
- public signals
- epistemic abstraction
- computability issues

Two-player turn-based zero-sum games

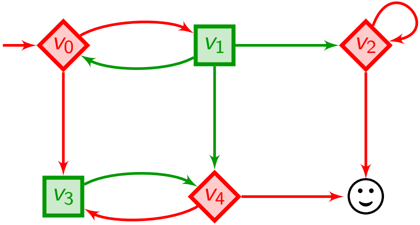

- Game graph $G = (V, E)$
- V partitioned into $V_\Diamond$ and $V_\Box$

Two-player turn-based zero-sum games

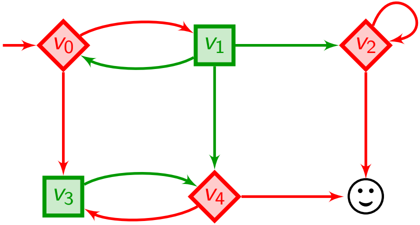

- Game graph $G = (V, E)$
- V partitioned into $V_\diamond$ and $V_\Box$

- Strategy for player i:

$$\sigma_i \colon V^* V_i \to V$$

- $\mathrm{out}(\sigma_\diamond)$: all paths compatible with $\sigma_\diamond$

Two-player turn-based zero-sum games

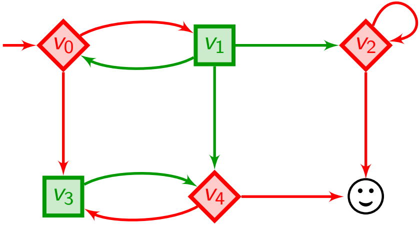

- Game graph $G = (V, E)$
- V partitioned into $V_\Diamond$ and $V_\Box$

- Strategy for player i:

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

- $\mathrm{out}(\sigma_\Diamond)$: all paths compatible with $\sigma_\Diamond$

- Objective for $\Diamond$: $\Omega \subseteq V^\omega$
- $\sigma_\Diamond$ winning strat. if $\mathrm{out}(\sigma_\Diamond) \subseteq \Omega$

Two-player turn-based zero-sum games

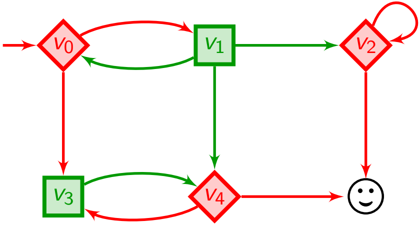

- Objective of $\diamond$: Reach ☺
- $\sigma_\diamond(v_0) = v_3$, $\sigma_\diamond(v_2) = \sigma_\diamond(v_4) =$ ☺ is a winning strategy

- Game graph $G = (V, E)$
- V partitioned into $V_\diamond$ and $V_\square$

- Strategy for player i:

$$\sigma_i \colon V^* V_i \to V$$

- $\mathrm{out}(\sigma_\diamond)$: all paths compatible with $\sigma_\diamond$

- Objective for $\diamond$: $\Omega \subseteq V^\omega$
- $\sigma_\diamond$ winning strat. if $\mathrm{out}(\sigma_\diamond) \subseteq \Omega$

Two-player turn-based zero-sum games

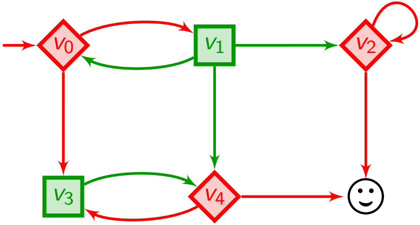

- Objective of $\diamond$: Reach ☺
- $\sigma_\diamond(v_0) = v_3$, $\sigma_\diamond(v_2) = \sigma_\diamond(v_4) = $ ☺ is a winning strategy

- Game graph $G = (V, E)$
- V partitioned into $V_\diamond$ and $V_\square$

- Strategy for player i:

$$\sigma_i \colon V^* V_i \to V$$

- $\mathrm{out}(\sigma_\diamond)$: all paths compatible with $\sigma_\diamond$

- Objective for $\diamond$: $\Omega \subseteq V^\omega$
- $\sigma_\diamond$ winning strat. if $\mathrm{out}(\sigma_\diamond) \subseteq \Omega$

- Determinacy: Either $\diamond$ has a winning strategy for Ω, or $\square$ has a winning strategy for $V^\omega \setminus Omega$

Non-zero-sum multiplayer games

- Several players $\mathrm{Agt} = \{A_1, \ldots, A_N\}$

Non-zero-sum multiplayer games

- Several players $\mathrm{Agt} = \{A_1, \ldots, A_N\}$
- Each player A plays according to a strategy σ_A

Non-zero-sum multiplayer games

- Several players $\mathrm{Agt} = \{A_1, \ldots, A_N\}$
- Each player A plays according to a strategy σ_A
- Each player A has a payoff function

$$\mathsf{payoff}_A : V^\omega \to \mathbb{R}$$

Non-zero-sum multiplayer games

- Several players $\mathsf{Agt} = \{A_1, \ldots, A_N\}$
- Each player A plays according to a strategy σ_A
- Each player A has a payoff function

$$\mathsf{payoff}_A : V^\omega \to \mathbb{R}$$

- Non-zero-sum

Non-zero-sum multiplayer games

- Several players $\mathsf{Agt} = \{A_1, \ldots, A_N\}$
- Each player A plays according to a strategy σ_A
- Each player A has a payoff function

$$\mathsf{payoff}_A : V^\omega \to \mathbb{R}$$

- Non-zero-sum
- Selfishness hypothesis: each player wants to maximize her own payoff!

Non-zero-sum multiplayer games

- Several players $\mathrm{Agt} = \{A_1, \ldots, A_N\}$
- Each player A plays according to a strategy σ_A
- Each player A has a payoff function

$$\mathsf{payoff}_A : V^\omega \to \mathbb{R}$$

- Non-zero-sum
- Selfishness hypothesis: each player wants to maximize her own payoff!
- Need of solution concepts to describe the kind of interactions between the players

Non-zero-sum multiplayer games

- Several players $\mathsf{Agt} = \{A_1, \ldots, A_N\}$
- Each player A plays according to a strategy σ_A
- Each player A has a payoff function

$$\mathsf{payoff}_A : V^\omega \to \mathbb{R}$$

- Non-zero-sum
- Selfishness hypothesis: each player wants to maximize her own payoff!
- Need of solution concepts to describe the kind of interactions between the players
- The simplest: Nash equilibria

Nash equilibria in turn-based games

Nash equilibrium

A strategy profile $(\sigma_A)_{A \in \mathsf{Agt}}$ is a Nash equilibrium if no player can improve her payoff by unilaterally changing her strategy.

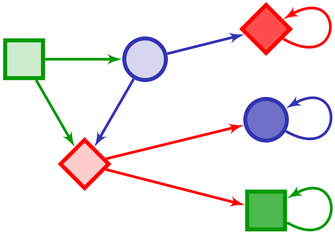

Nash equilibria in turn-based games

Nash equilibrium

A strategy profile $(\sigma_A)_{A \in \mathrm{Agt}}$ is a Nash equilibrium if no player can improve her payoff by unilaterally changing her strategy.

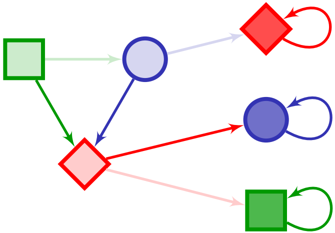

is a Nash equilibrium with payoff $(0, 1, 0)$

Nash equilibria in turn-based games

Nash equilibrium

A strategy profile $(\sigma_A)_{A \in \mathsf{Agt}}$ is a Nash equilibrium if no player can improve her payoff by unilaterally changing her strategy.

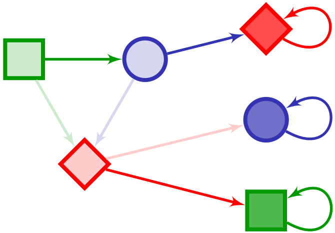

is not a Nash equilibrium

Nash equilibria in turn-based games

Nash equilibrium

A strategy profile $(\sigma_A)_{A \in \mathsf{Agt}}$ is a Nash equilibrium if no player can improve her payoff by unilaterally changing her strategy.

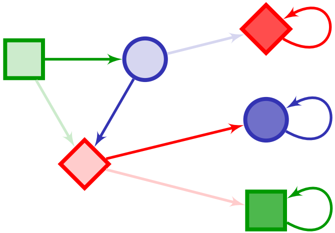

is a Nash equilibrium with payoff $(0, 1, 0)$

Boolean Nash equilibria in turn-based games

ψ_A: objective of player A

Boolean Nash equilibria in turn-based games

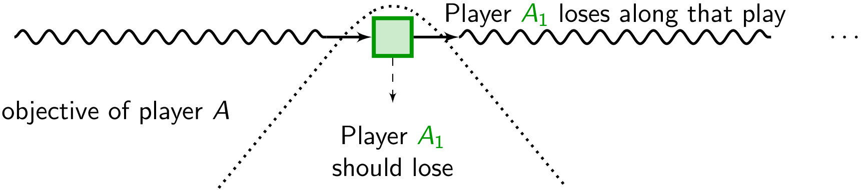

ψ_A: objective of player A

Boolean Nash equilibria in turn-based games

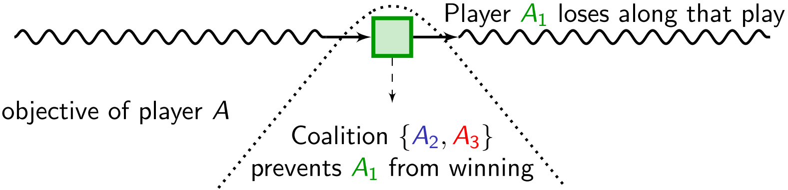

. . .

ψ_A: objective of player A

Boolean Nash equilibria in turn-based games

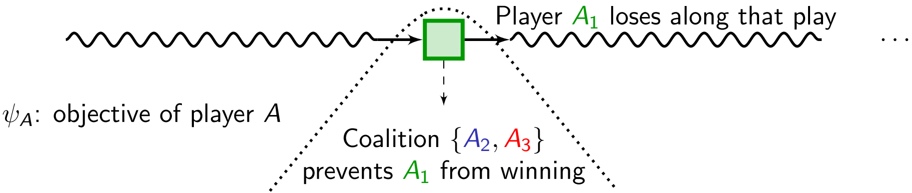

Recipe

- for every $A \in \mathsf{Agt}$, compute the set of winning states W_A
- find a path witness for the formula:

$$\Phi_{\mathsf{NE}} = \bigwedge_{A \in \mathsf{Agt}} \left(\neg\psi_A \Rightarrow \mathbf{G} \neg W_A \right)$$

(valid for tail or reachability objectives)

Existing results in the framework of turn-based games

[UW11,Umm11]

- There always exists a Nash equilibrium for Boolean ω-regular objectives

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Stochastic Multiplayer Games *(LMCS)*
[Umm11] Ummels. Stochastic multiplayer games: theory and algorithms (RWTH Aachen University)

Existing results in the framework of turn-based games

[UW11,Umm11]

- There always exists a Nash equilibrium for Boolean ω-regular objectives
- One can decide the constrained existence of a Nash equilibrium (and compute one!)

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Stochastic Multiplayer Games *(LMCS)*
[Umm11] Ummels. Stochastic multiplayer games: theory and algorithms (RWTH Aachen University)

Existing results in the framework of turn-based games

[UW11,Umm11]

- There always exists a Nash equilibrium for Boolean ω-regular objectives
- One can decide the constrained existence of a Nash equilibrium (and compute one!)
- One cannot decide the existence of a mixed (i.e. stochastic) Nash equilibrium

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Stochastic Multiplayer Games *(LMCS)*
[Umm11] Ummels. Stochastic multiplayer games: theory and algorithms (RWTH Aachen University)

Existing results in the framework of turn-based games

[UW11,Umm11]

- There always exists a Nash equilibrium for Boolean ω-regular objectives
- One can decide the constrained existence of a Nash equilibrium (and compute one!)
- One cannot decide the existence of a mixed (i.e. stochastic) Nash equilibrium

$\rightsquigarrow$ this is why we restrict to pure equilibria

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Stochastic Multiplayer Games *(LMCS)*
[Umm11] Ummels. Stochastic multiplayer games: theory and algorithms (RWTH Aachen University)

What about concurrent games?

[BBMU15] Bouyer, Brenguier, Markey, Ummels. Pure Nash equilibria in concurrent games *(LMCS)*

What about concurrent games?

The matching-penny game:

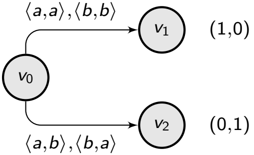

What about concurrent games?

The matching-penny game:

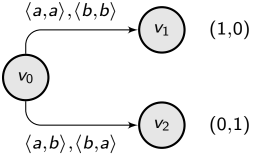

There is no pure Nash eq.

[BBMU15] Bouyer, Brenguier, Markey, Ummels. Pure Nash equilibria in concurrent games *(LMCS)*

What about concurrent games?

The matching-penny game:

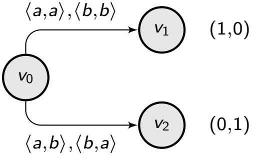

There is no pure Nash eq.

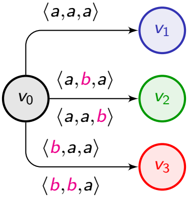

What about concurrent games?

The matching-penny game:

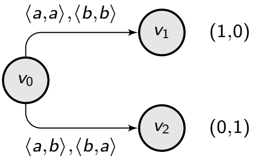

There is no pure Nash eq.

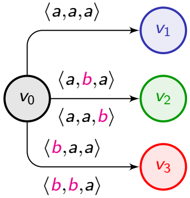

- $\mathrm{susp}\Big((v_0, v_3), \langle a, a, a\rangle\Big) = \{A_1\}$

[BBMU15] Bouyer, Brenguier, Markey, Ummels. Pure Nash equilibria in concurrent games *(LMCS)*

What about concurrent games?

The matching-penny game:

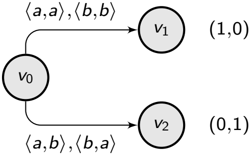

There is no pure Nash eq.

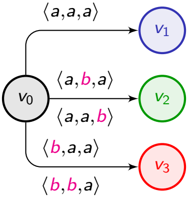

- $\mathrm{susp}\Big((v_0, v_3), \langle a, a, a\rangle\Big) = \{A_1\}$
- $\mathrm{susp}\Big((v_0, v_2), \langle a, a, a\rangle\Big) = \{A_2, A_3\}$

What about concurrent games?

The matching-penny game:

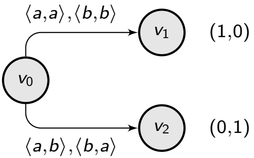

There is no pure Nash eq.

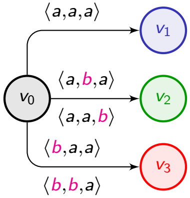

- $\mathrm{susp}\Big((v_0, v_3), \langle a, a, a\rangle\Big) = \{A_1\}$
- $\mathrm{susp}\Big((v_0, v_2), \langle a, a, a\rangle\Big) = \{A_2, A_3\}$

Solution via the suspect game abstraction,
a structure to track suspect players

What about concurrent games?

The matching-penny game:

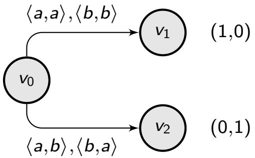

There is no pure Nash eq.

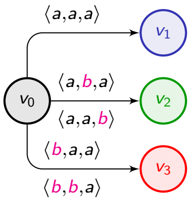

- $\mathrm{susp}\Big((v_0, v_3), \langle a, a, a\rangle\Big) = \{A_1\}$
- $\mathrm{susp}\Big((v_0, v_2), \langle a, a, a\rangle\Big) = \{A_2, A_3\}$

Solution via the suspect game abstraction,
a structure to track suspect players

Can we add more partial information to that framework?

[BBMU15] Bouyer, Brenguier, Markey, Ummels. Pure Nash equilibria in concurrent games *(LMCS)*

Concurrent games with signals

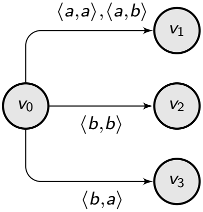

[Tom98] Tomala. Pure equilibria of repeated games with public observation *(International Journal of Game Theory)*

Concurrent games with signals

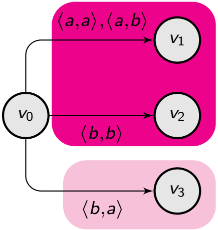

- Signal for player A_1: ● and ●

- On playing a, player A_1 will receive ●
- On playing b, player A_1 will receive either ● or ●

[Tom98] Tomala. Pure equilibria of repeated games with public observation *(International Journal of Game Theory)*

Concurrent games with signals

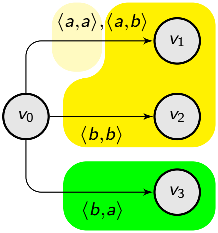

- Signal for player A_1: ● and ●
- Signal for player A_2: ●, ● and ●

- On playing a, player A_1 will receive ●
- On playing b, player A_1 will receive either ● or ●

- On playing a, player A_2 will receive either ● or ●
- On playing b, player A_2 will receive ●

[Tom98] Tomala. Pure equilibria of repeated games with public observation *(International Journal of Game Theory)*

Concurrent games with signals

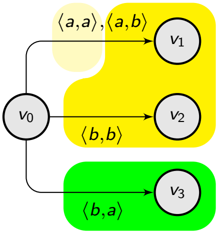

- Signal for player A_1: and
- Signal for player A_2: , and

- On playing a, player A_1 will receive
- On playing b, player A_1 will receive either or

- On playing a, player A_2 will receive either or
- On playing b, player A_2 will receive

Public signal

Same signal to every player!

Our specific framework

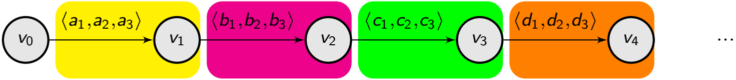

Our specific framework

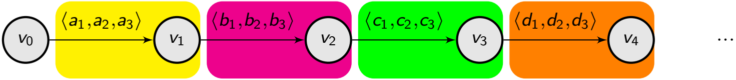

- What player A_i sees:

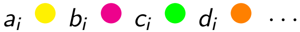

Our specific framework

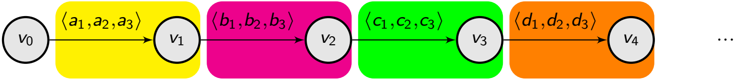

- What player A_i sees:

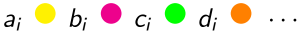

$\rightsquigarrow$ induces undistinguishability relation $\sim_{A_i}$

Our specific framework

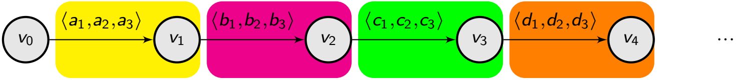

- What player A_i sees:

$$a_i \quad b_i \quad c_i \quad d_i \quad \cdots$$

$\rightsquigarrow$ induces undistinguishability relation $\sim_{A_i}$

- Strategy of player A_i has to respect $\sim_{A_i}$

Our specific framework

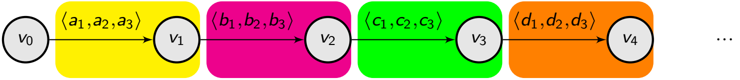

- What player A_i sees:

$$a_i \quad b_i \quad c_i \quad d_i \quad \cdots$$

$\rightsquigarrow$ induces undistinguishability relation $\sim_{A_i}$

- Strategy of player A_i has to respect $\sim_{A_i}$
- Privately visible payoff: based on

$$a_i \quad b_i \quad c_i \quad d_i \quad \cdots$$

Our specific framework

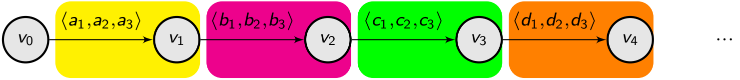

- What player A_i sees:

$$a_i \quad b_i \quad c_i \quad d_i \quad \cdots$$

$\rightsquigarrow$ induces undistinguishability relation $\sim_{A_i}$

- Strategy of player A_i has to respect $\sim_{A_i}$
- Privately visible payoff: based on

$$a_i \quad b_i \quad c_i \quad d_i \quad \cdots$$

- Publicly visible payoff: based on sequences of colors

$$\ldots$$

Payoff functions of interest

- Boolean ω-regular payoff function (for Ω):

$$\mathsf{payoff}(\rho) = \begin{cases} 1 & \text{if } \rho \in \Omega \\ 0 & \text{otherwise} \end{cases}$$

Payoff functions of interest

- Boolean ω-regular payoff function (for Ω):

$$\mathsf{payoff}(\rho) = \begin{cases} 1 & \text{if } \rho \in \Omega \\ 0 & \text{otherwise} \end{cases}$$

- Mean-payoff (limsup or liminf) w.r.t. weight function w:

$$\begin{cases} \underline{\mathsf{MP}}_w(\rho) = \liminf_{n\to\infty} \sum_{i=0}^{n} w\left(v_i \xrightarrow{m_i} v_{i+1}\right) \\ \overline{\mathsf{MP}}_w(\rho) = \limsup_{n\to\infty} \sum_{i=0}^{n} w\left(v_i \xrightarrow{m_i} v_{i+1}\right) \end{cases}$$

Payoff functions of interest

- Boolean ω-regular payoff function (for Ω):

$$\mathsf{payoff}(\rho) = \begin{cases} 1 & \text{if } \rho \in \Omega \\ 0 & \text{otherwise} \end{cases}$$

- Mean-payoff (limsup or liminf) w.r.t. weight function w:

$$\begin{cases} \underline{\mathsf{MP}}_w(\rho) = \liminf_{n\to\infty} \sum_{i=0}^{n} w\left(v_i \xrightarrow{m_i} v_{i+1}\right) \\ \overline{\mathsf{MP}}_w(\rho) = \limsup_{n\to\infty} \sum_{i=0}^{n} w\left(v_i \xrightarrow{m_i} v_{i+1}\right) \end{cases}$$

For public visibility, we will assume that atomic propositions/atomic weights are defined w.r.t. the signal alphabet Σ.

An example

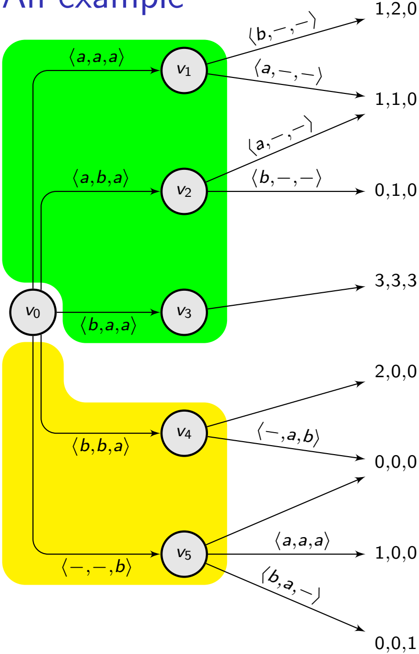

- Three players concurrent game with public signal

An example

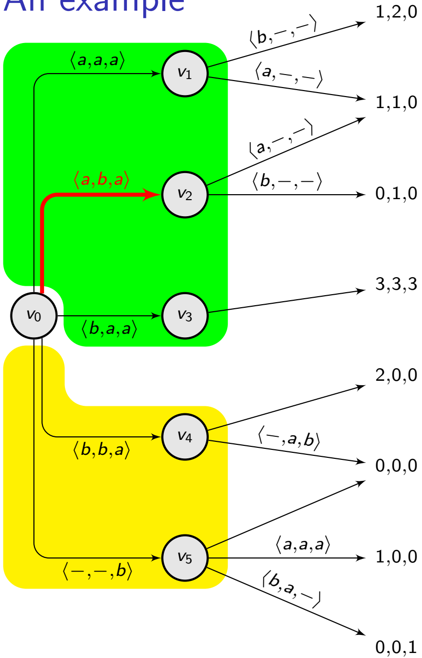

- Three players concurrent game with public signal
- Consider the (partial) strategy profile σ_{Agt}. Can we complete it into a Nash equilibrium?

An example

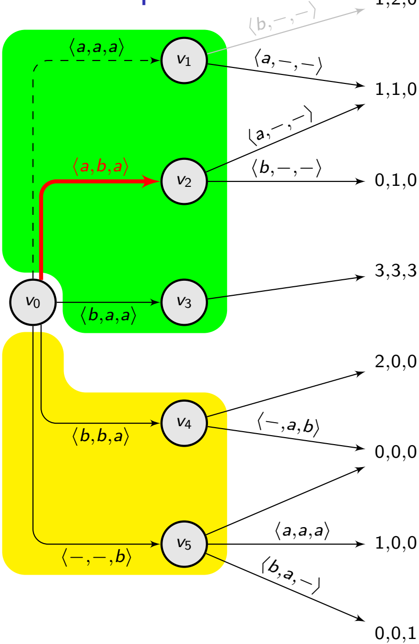

- Three players concurrent game with public signal
- Consider the (partial) strategy profile σ_{Agt}. Can we complete it into a Nash equilibrium?
- This is an A_2-deviation, which is invisible to both A_1 and A_3. A_1 has to play a.

An example

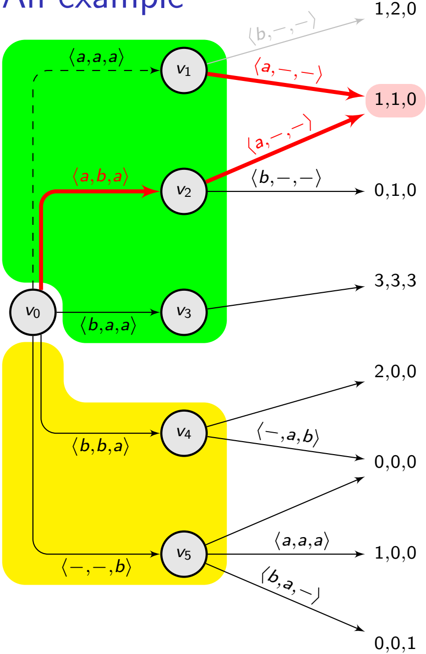

- Three players concurrent game with public signal
- Consider the (partial) strategy profile σ_{Agt}. Can we complete it into a Nash equilibrium?
- This is an A_2-deviation, which is invisible to both A_1 and A_3. A_1 has to play a. Which we propagate out of v_2.

An example

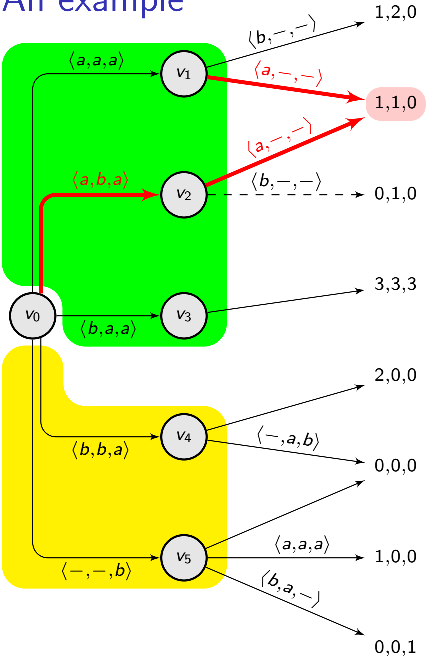

- Three players concurrent game with public signal
- Consider the (partial) strategy profile σ_{Agt}. Can we complete it into a Nash equilibrium?
- This is an A_2-deviation, which is invisible to both A_1 and A_3. A_1 has to play a. Which we propagate out of v_2.
- This is a non-profitable A_1-deviation.

An example

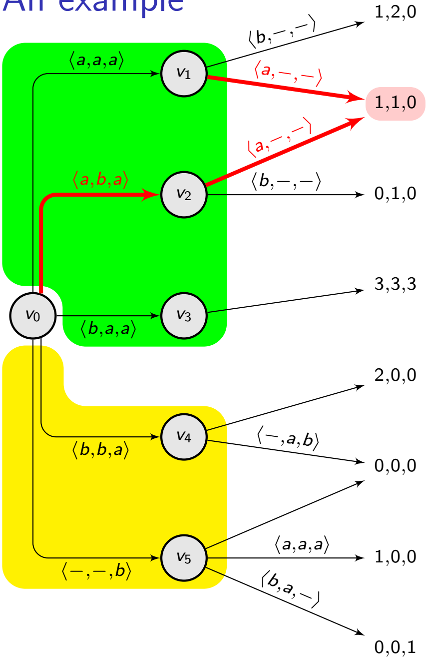

- Three players concurrent game with public signal
- Consider the (partial) strategy profile σ_{Agt}. Can we complete it into a Nash equilibrium?
- This is an A_2-deviation, which is invisible to both A_1 and A_3. A_1 has to play a. Which we propagate out of v_2.
- This is a non-profitable A_1-deviation.
- No one (alone) can deviate to v_3.

An example

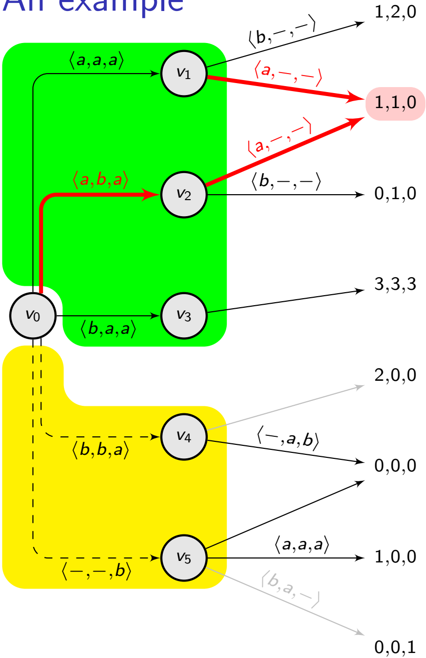

- Three players concurrent game with public signal
- Consider the (partial) strategy profile σ_{Agt}. Can we complete it into a Nash equilibrium?
- This is an A_2-deviation, which is invisible to both A_1 and A_3. A_1 has to play a. Which we propagate out of v_2.
- This is a non-profitable A_1-deviation.
- No one (alone) can deviate to v_3.
- A_1 can deviate to v_4 and A_3 can deviate to v_5: A_2 knows there has been a deviation, but (s)he doesn't know whether A_1 or A_3 did so, and whether the game proceeds to v_4 or v_5. On the other hand, both A_1 and A_3 know!

An example

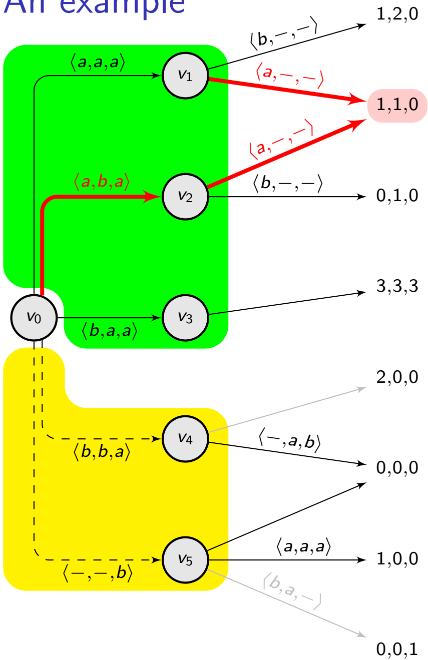

- Three players concurrent game with public signal
- Consider the (partial) strategy profile σ_{Agt}. Can we complete it into a Nash equilibrium?
- This is an A_2-deviation, which is invisible to both A_1 and A_3. A_1 has to play a. Which we propagate out of v_2.
- This is a non-profitable A_1-deviation.
- No one (alone) can deviate to v_3.
- A_1 can deviate to v_4 and A_3 can deviate to v_5: A_2 knows there has been a deviation, but (s)he doesn't know whether A_1 or A_3 did so, and whether the game proceeds to v_4 or v_5. On the other hand, both A_1 and A_3 know! But if the game proceeds to v_4, A_3 can help A_2 punishing A_1, and if the game proceeds to v_5, A_1 can help A_2 punishing A_3.

An example

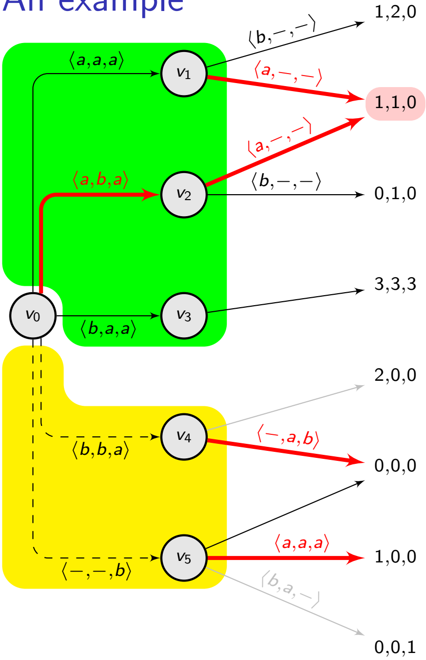

- Three players concurrent game with public signal
- Consider the (partial) strategy profile σ_{Agt}. Can we complete it into a Nash equilibrium?
- This is an A_2-deviation, which is invisible to both A_1 and A_3. A_1 has to play a. Which we propagate out of v_2.
- This is a non-profitable A_1-deviation.
- No one (alone) can deviate to v_3.
- A_1 can deviate to v_4 and A_3 can deviate to v_5: A_2 knows there has been a deviation, but (s)he doesn't know whether A_1 or A_3 did so, and whether the game proceeds to v_4 or v_5. On the other hand, both A_1 and A_3 know! But if the game proceeds to v_4, A_3 can help A_2 punishing A_1, and if the game proceeds to v_5, A_1 can help A_2 punishing A_3.

An example

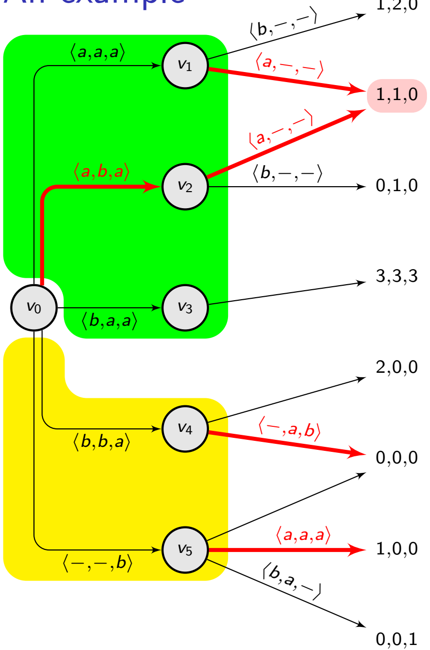
1,2,0
⟨b,−,−⟩
⟨a,a,a⟩
v1
⟨a,−,−⟩
1,1,0
⟨a,−,−⟩
⟨a,b,a⟩
v2
⟨b,−,−⟩
0,1,0
3,3,3
v0
v3
⟨b,a,a⟩
2,0,0
v4
⟨−,a,b⟩
⟨b,b,a⟩
0,0,0
v5
⟨a,a,a⟩
1,0,0
⟨−,−,b⟩
⟨b,a,−⟩
0,0,1

What we learn from that example

- The main outcome of a Nash equilibrium has to be robust to invisible deviations

What we learn from that example

- The main outcome of a Nash equilibrium has to be robust to invisible deviations
- Visible deviations may induce some uncertainty on possible deviators (no common knowledge)

What we learn from that example

- The main outcome of a Nash equilibrium has to be robust to invisible deviations
- Visible deviations may induce some uncertainty on possible deviators (no common knowledge)

- How to systematically track all individual deviations and uncertainty induced by imperfect information monitoring?

What we learn from that example

- The main outcome of a Nash equilibrium has to be robust to invisible deviations
- Visible deviations may induce some uncertainty on possible deviators (no common knowledge)

- How to systematically track all individual deviations and uncertainty induced by imperfect information monitoring?
- Is that always possible?

What we learn from that example

- The main outcome of a Nash equilibrium has to be robust to invisible deviations
- Visible deviations may induce some uncertainty on possible deviators (no common knowledge)

- How to systematically track all individual deviations and uncertainty induced by imperfect information monitoring?
- Is that always possible?
- Can we build a finite epistemic structure?

The epistemic game abstraction

Inspired by:

- the standard powerset construction [Rei84]
- the epistemic unfolding for coordination/distributed synthesis [BKP11]
- the suspect game [BBMU15]
- the deviator game [Bre16]

[Rei84] Reif. The complexity of two-player games of incomplete information *(J. Comp. and Syst. Sc.)*
[BKP11] Berwanger, Kaiser, Puchala. Perfect-information construction for coordination in games *(FSTTCS'11)*
[BBMU15] Pure Nash equilibria in concurrent games *(Log. Meth. in Comp. Sc.)*
[Bre16] Brenguier. Robust equilibria in mean-payoff games *(FoSSaCS'16)*

The epistemic game abstraction

Inspired by:

- the standard powerset construction [Rei84]
- the epistemic unfolding for coordination/distributed synthesis [BKP11]
- the suspect game [BBMU15]
- the deviator game [Bre16]

The idea is to track all possible undistinguishable behaviours, including the single-player deviations

[Rei84] Reif. The complexity of two-player games of incomplete information *(J. Comp. and Syst. Sc.)*
[BKP11] Berwanger, Kaiser, Puchala. Perfect-information construction for coordination in games *(FSTTCS'11)*
[BBMU15] Pure Nash equilibria in concurrent games *(Log. Meth. in Comp. Sc.)*
[Bre16] Brenguier. Robust equilibria in mean-payoff games *(FoSSaCS'16)*

Epistemic states

s

$\perp \mapsto \{v\}$ — vertex the game is in if no deviation has occurred

$A_1 \mapsto V_{A_1}$

$A_2 \mapsto V_{A_2}$ — vertices the game might be in if A_2 has deviated (invisible deviation)

$A_3 \mapsto V_{A_3}$

Epistemic states

s

$\perp \mapsto \{v\}$ ← vertex the game is in if no deviation has occurred

$A_1 \mapsto V_{A_1}$

$A_2 \mapsto V_{A_2}$ ← vertices the game might be in if A_2 has deviated (invisible deviation)

$A_3 \mapsto V_{A_3}$

s

$\perp \mapsto \emptyset$ ← a visible deviation has for sure occurred

$A_1 \mapsto V_{A_1}$

$A_2 \mapsto V_{A_2}$ ← vertices the game might be in if A_2 has deviated (visible deviation)

$A_3 \mapsto V_{A_3}$

The example again

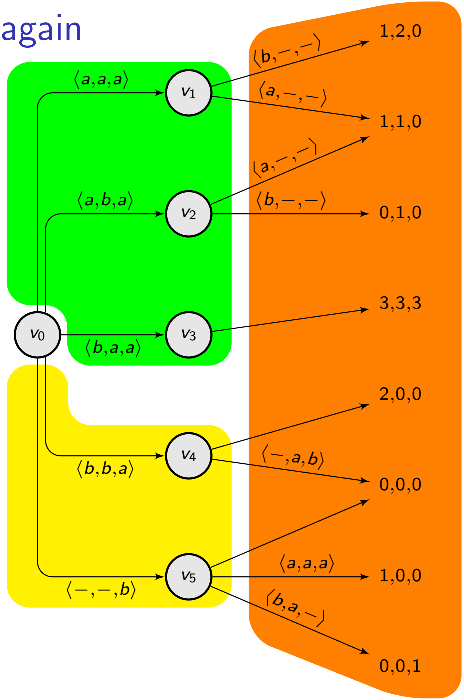

The example again

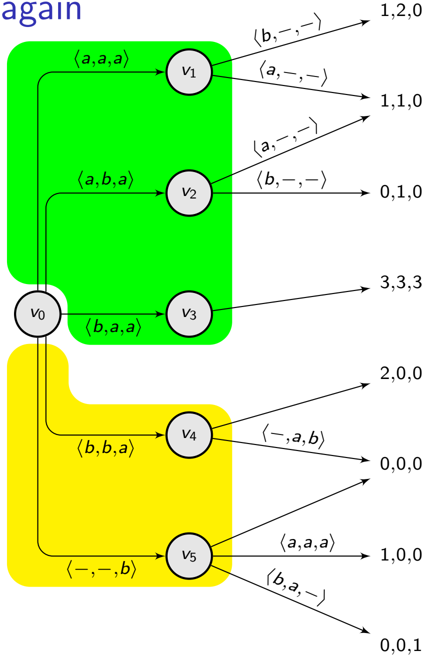

Example of construction

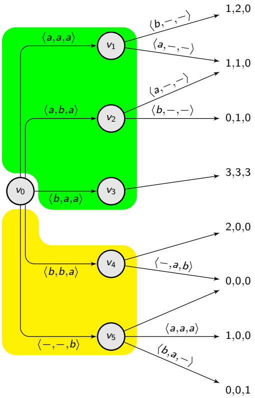

Example of construction

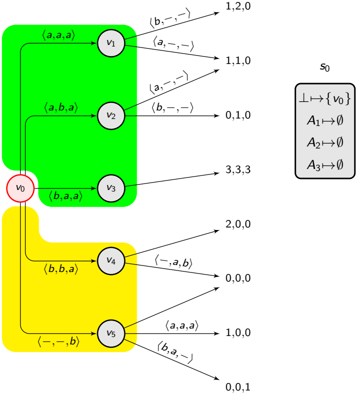

Example of construction

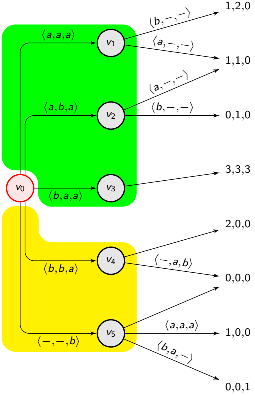

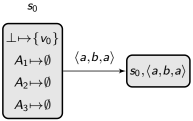

Example of construction

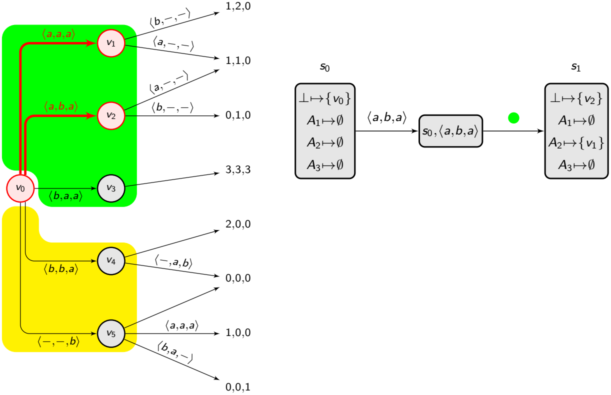

Example of construction

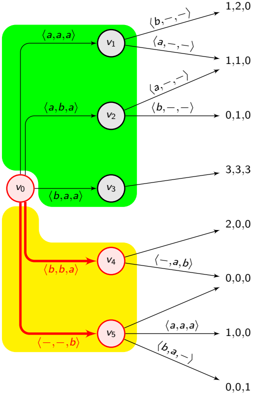

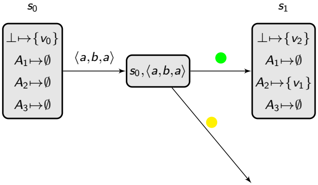

Example of construction

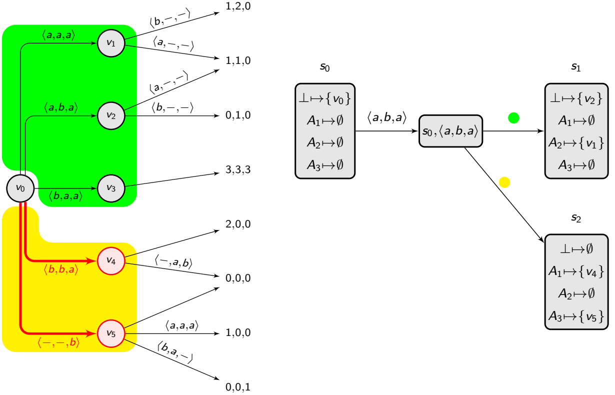

Example of construction

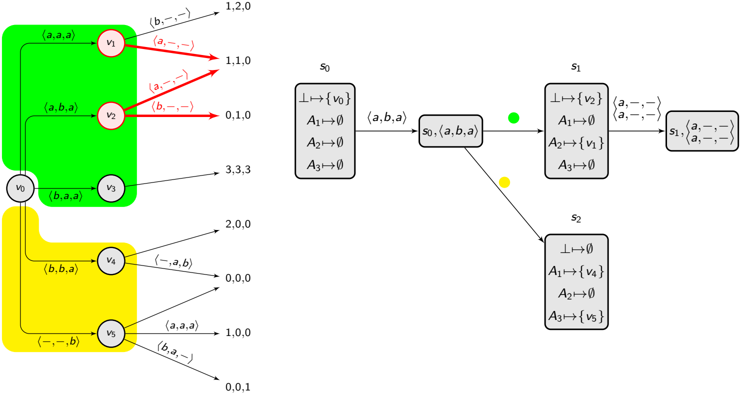

Example of construction

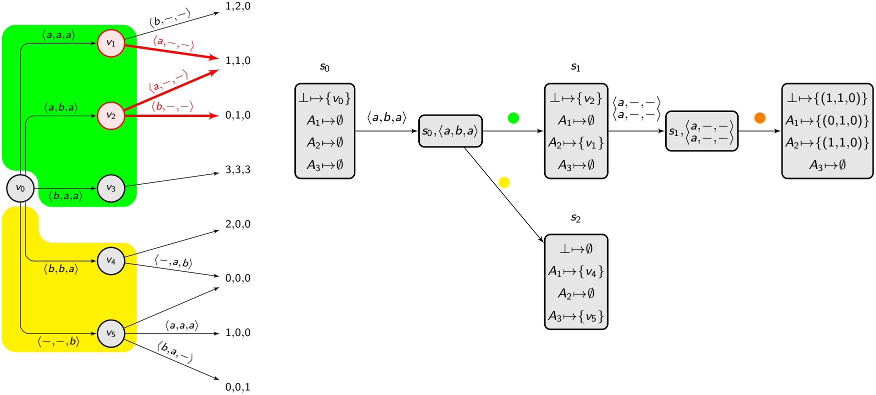

Example of construction

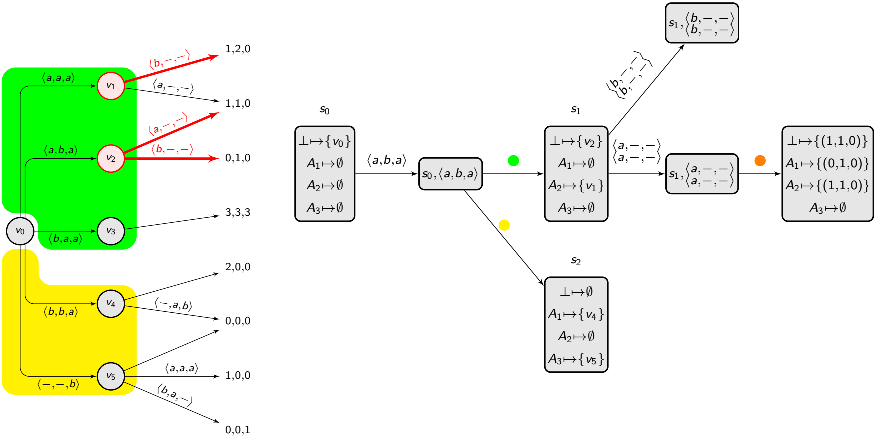

Example of construction

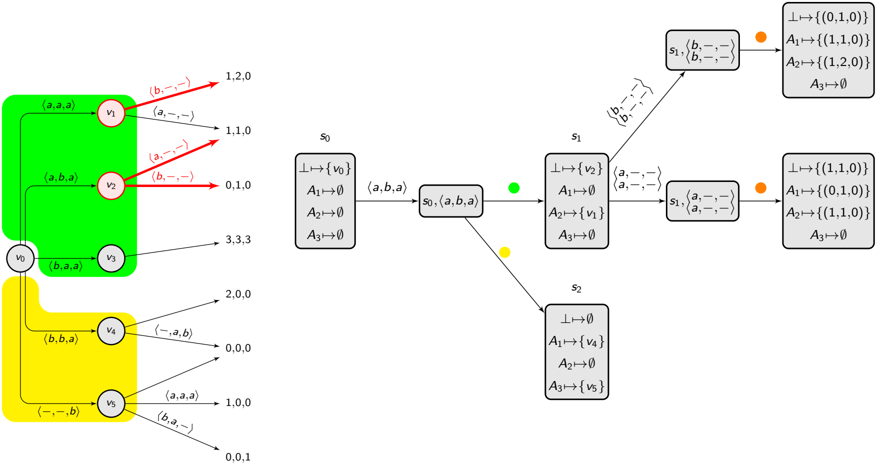

Example of construction

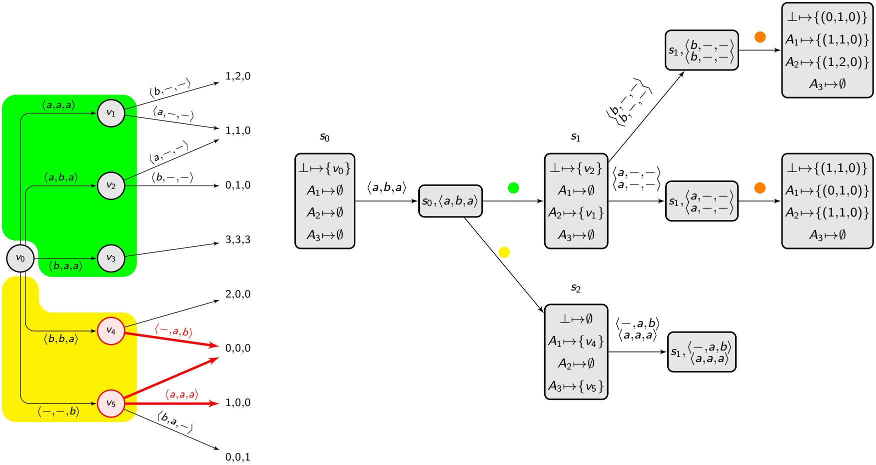

Example of construction

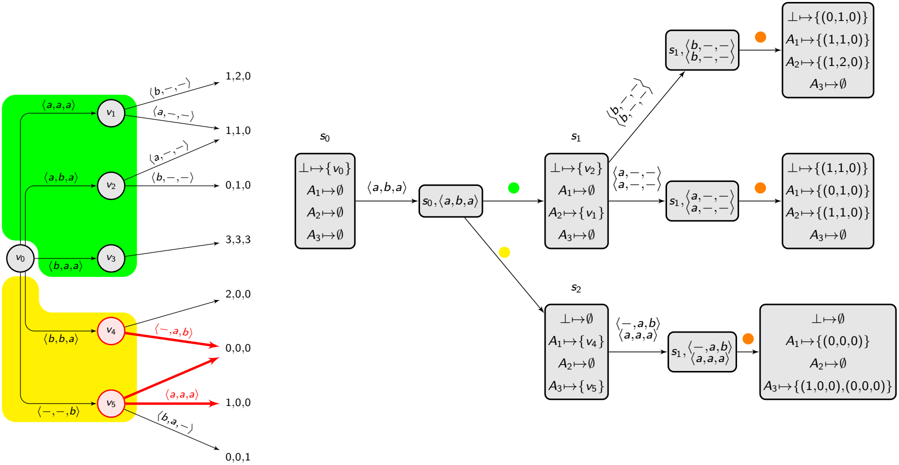

Example of construction

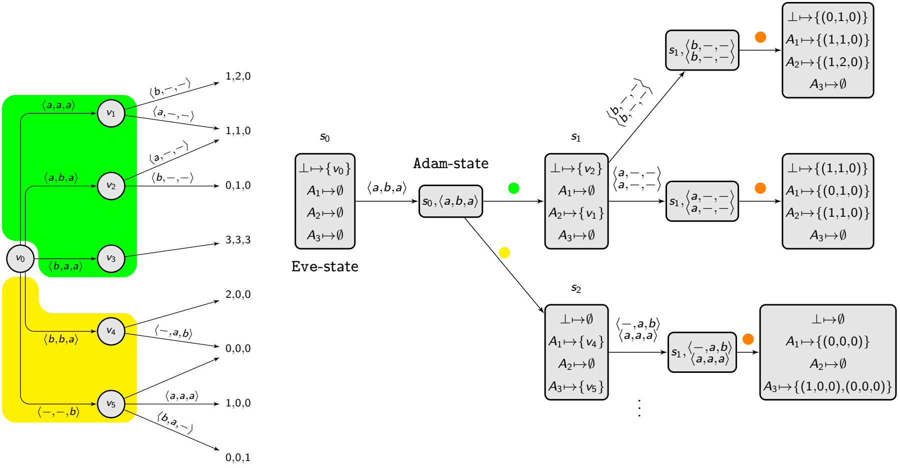

Properties of the epistemic game

- To every history H in the epistemic game, one can associate sets
  - $\mathit{concrete}_{\perp}(H)$: at most one concrete real history (if no deviation)
  - $\mathit{concrete}_{A}(H)$: all possible A-deviations
  - $\mathit{concrete}(H) = \bigcup_{A \in \mathrm{Agt} \cup \{\perp\}} \mathit{concrete}_{A}(H)$

Properties of the epistemic game

- To every history H in the epistemic game, one can associate sets
  - $\mathit{concrete}_\perp(H)$: at most one concrete real history (if no deviation)
  - $\mathit{concrete}_A(H)$: all possible A-deviations
  - $\mathit{concrete}(H) = \bigcup_{A \in \mathsf{Agt} \cup \{\perp\}} \mathit{concrete}_A(H)$

H history in the epistemic game. For every $h_1 \neq h_2 \in \mathit{concrete}(H)$,

$$h_1 \sim_A h_2 \qquad \text{iff} \qquad h_1, h_2 \notin \mathit{concrete}_A(H)$$

Properties of the epistemic game (cont'd)

Winning condition for Eve

A strategy $\sigma_{\texttt{Eve}}$ is said winning for payoff $p \in \mathbb{R}^{\mathrm{Agt}}$ from s_0 whenever $\mathrm{payoff}(\mathit{concrete}_\perp(\mathrm{out}_\perp(\sigma_{\texttt{Eve}}, s_0))) = p$, and for every $R \in \mathrm{out}(\sigma_{\texttt{Eve}}, s_0)$, for every $A \in \mathrm{Agt}$, for every $\rho \in \mathit{concrete}_A(R)$, $\mathrm{payoff}_A(\rho) \le p_A$.

Properties of the epistemic game (cont'd)

Winning condition for $\texttt{Eve}$ (publicly visible payoffs)

A strategy $\sigma_{\texttt{Eve}}$ is said winning for p from s_0 whenever $\mathsf{payoff}'(\mathsf{out}_\perp(\sigma_{\texttt{Eve}}, s_0)) = p$, and for every $R \in \mathsf{out}(\sigma_{\texttt{Eve}}, s_0)$, for every $A \in \mathsf{susp}(R)$, $\mathsf{payoff}'_A(R) \leq p_A$.

Properties of the epistemic game (cont'd)

Winning condition for Eve (publicly visible payoffs)

A strategy $\sigma_{\texttt{Eve}}$ is said winning for p from s_0 whenever $\mathsf{payoff}'(\mathsf{out}_\perp(\sigma_{\texttt{Eve}}, s_0)) = p$, and for every $R \in \mathsf{out}(\sigma_{\texttt{Eve}}, s_0)$, for every $A \in \mathsf{susp}(R)$, $\mathsf{payoff}'_A(R) \leq p_A$.

Proposition

There is a Nash equilibrium in $\mathcal{G}$ with payoff p from v_0 if and only if Eve has a winning strategy for p in $\mathcal{E}_{\mathcal{G}}$ from s_0.

Application to ω-regular objectives

Player A_1 loses along that play

ψ_A: objective of player A

Application to ω-regular objectives

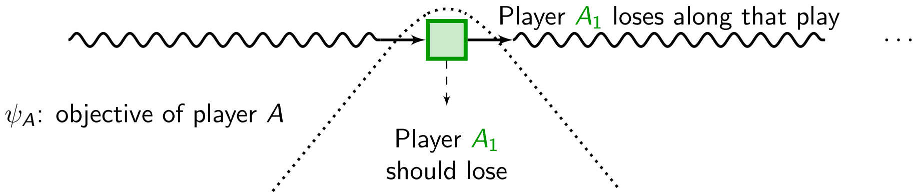

Application to ω-regular objectives

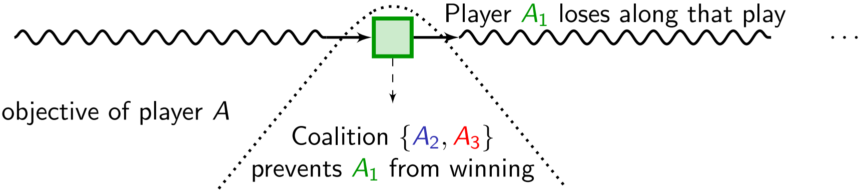

ψ_A: objective of player A

Application to ω-regular objectives

ψ_A: objective of player A

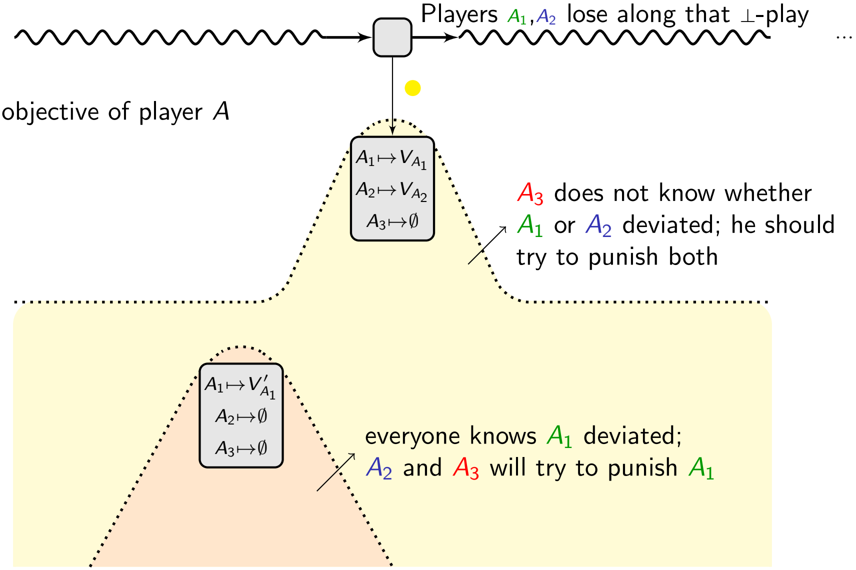

Application to ω-regular objectives (cont'd)

- This amounts to solving two-player turn-based games with generalized (i.e. conjunctions of) ω-regular objectives

Application to ω-regular objectives (cont'd)

- This amounts to solving two-player turn-based games with generalized (i.e. conjunctions of) ω-regular objectives

Theorem

One can decide the (constrained) existence of a Nash equilibrium in a game with public signal and publicly visible payoff functions associated with parity conditions in EXPSPACE. It is EXPTIME-hard.

Application to ω-regular objectives (cont'd)

- This amounts to solving two-player turn-based games with generalized (i.e. conjunctions of) ω-regular objectives

Theorem

One can decide the (constrained) existence of a Nash equilibrium in a game with public signal and publicly visible payoff functions associated with parity conditions in EXPSPACE. It is EXPTIME-hard.

- By reduction from the distributed synthesis problem (proof of [BK10]):

Theorem

One cannot decide the existence of a Nash equilibrium in a game with private signals and publicly visible ω-regular payoff functions. Even for three players.

[BK10] Berwanger, Kaiser. Information Tracking in Games on Graphs *(Journal of Logic, Language and Information)*

Application to mean-payoff functions

- Using results on the polyhedron problem [BR15,Bre16]:

Theorem

One can decide the (constrained) existence of a Nash equilibrium in a game with public signal and publicly visible mean-payoff functions, in NP, with a coNEXPTIME oracle. This in particular can be solved in EXPSPACE. It is EXPTIME-hard.

[BR15] Brenguier, Raskin. Pareto curves of multidimensional mean-payoff games *(CAV'15)*
[Bre16] Brenguier. Robust equilibria in mean-payoff games *(FoSSaCS'16)*

Application to mean-payoff functions

- Using results on the polyhedron problem [BR15,Bre16]:

Theorem

One can decide the (constrained) existence of a Nash equilibrium in a game with public signal and publicly visible mean-payoff functions, in NP, with a coNEXPTIME oracle. This in particular can be solved in EXPSPACE. It is EXPTIME-hard.

- By reduction from blind mean-payoff games (proven undecidable in [DDG+10])

Theorem

One cannot decide the constrained existence of a Nash equilibrium in a game with public signal and privately visible mean-payoff functions. Even for two players.

[BR15] Brenguier, Raskin. Pareto curves of multidimensional mean-payoff games *(CAV'15)*
[Bre16] Brenguier. Robust equilibria in mean-payoff games *(FoSSaCS'16)*
[DDG+10] Degorre, Doyen, Gentilini, Raskin, Toruńczyk. Energy and Mean-Payoff Games with Imperfect Information *(CSL'10)*

Application to mean-payoff functions: decidability

The polyhedron problem

In a multi-dimensional mean-payoff two-player turn-based game, the polyhedron problem aks, given a polyhedron π, whether there is a strategy for $\mathtt{Eve}$ which ensures a payoff vector which belongs to π.

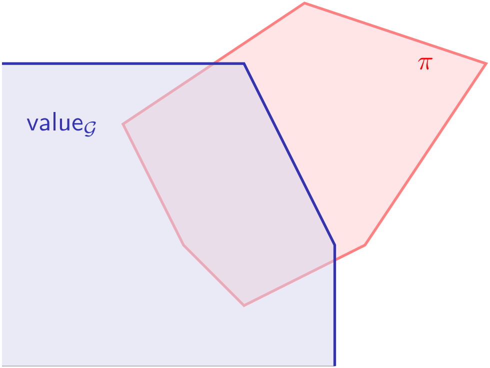

$$\text{value}_\mathcal{G} = \{v \in \mathbb{R}^d \mid \exists \sigma \forall \rho \in \text{out}(\sigma),\ \forall i,\ \text{MP}_i(\rho) \geq v_i\}$$

Application to mean-payoff functions: decidability

The polyhedron problem

In a multi-dimensional mean-payoff two-player turn-based game, the polyhedron problem aks, given a polyhedron π, whether there is a strategy for $\mathtt{Eve}$ which ensures a payoff vector which belongs to π.

Theorem [BR15]

If there is a solution to the polyhedron problem, there is one solution with a payoff of polynomial size.
The polyhedron problem is $\Sigma_2\mathsf{P}$-complete ($\Sigma_2\mathsf{P} = \mathsf{NP}^{\mathsf{NP}}$)

Application to mean-payoff functions: decidability

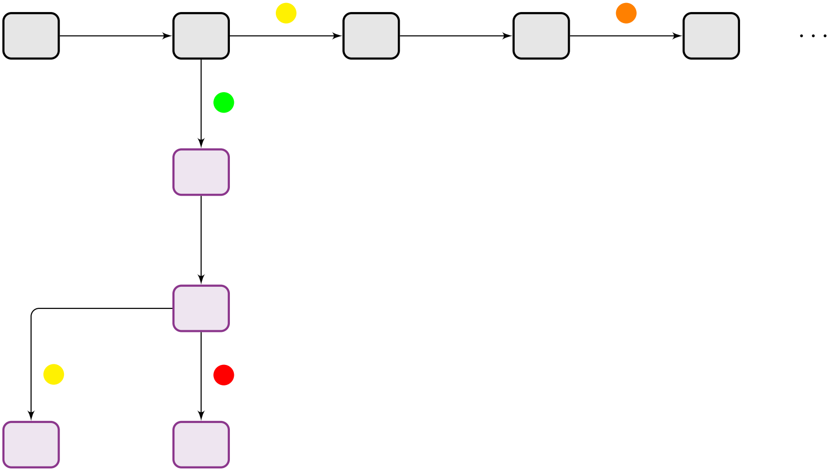

Application to mean-payoff functions: decidability

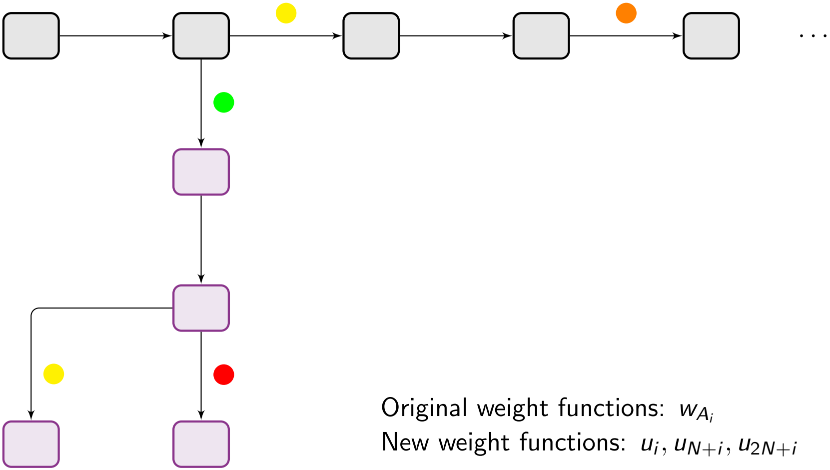

Original weight functions: w_{A_i}
New weight functions: u_i, u_{N+i}, u_{2N+i}

Application to mean-payoff functions: decidability

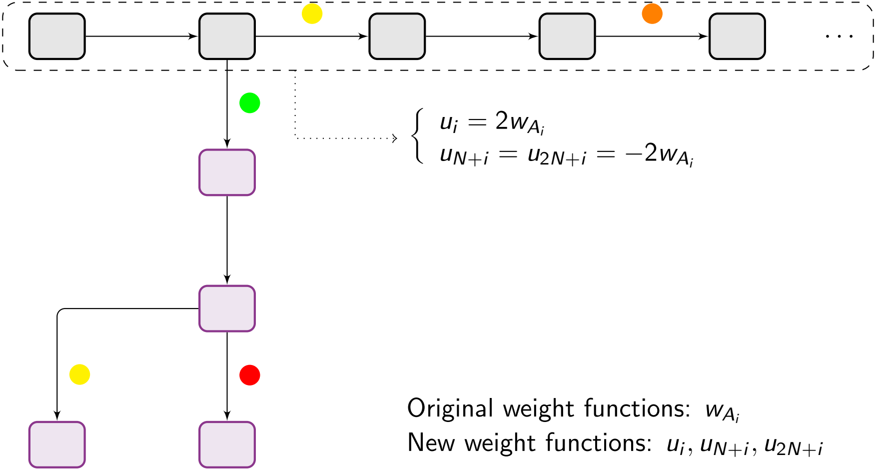

$$\begin{cases} u_i = 2w_{A_i} \\ u_{N+i} = u_{2N+i} = -2w_{A_i} \end{cases}$$

Original weight functions: w_{A_i}
New weight functions: u_i, u_{N+i}, u_{2N+i}

Application to mean-payoff functions: decidability

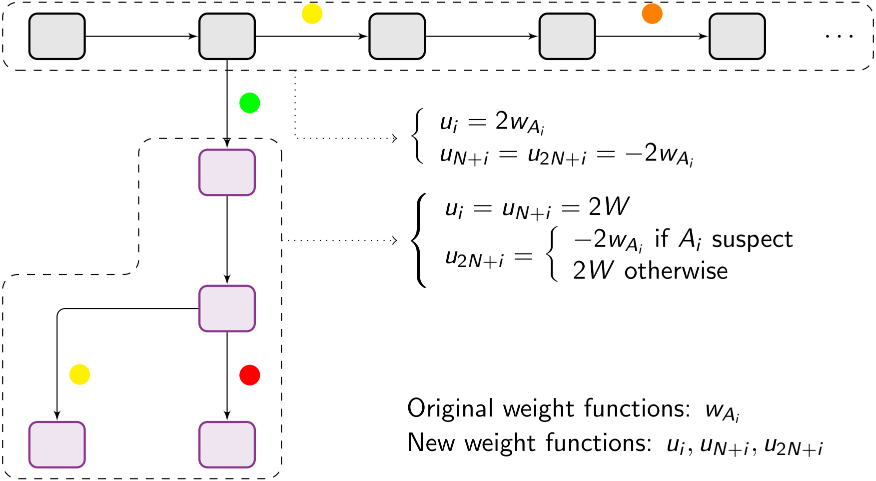

$$\begin{cases} u_i = 2w_{A_i} \\ u_{N+i} = u_{2N+i} = -2w_{A_i} \end{cases}$$

$$\begin{cases} u_i = u_{N+i} = 2W \\ u_{2N+i} = \begin{cases} -2w_{A_i} & \text{if } A_i \text{ suspect} \\ 2W & \text{otherwise} \end{cases} \end{cases}$$

Original weight functions: w_{A_i}

New weight functions: u_i, u_{N+i}, u_{2N+i}

Application to mean-payoff functions: decidability

$$\begin{cases} u_i = 2w_{A_i} \end{cases}$$

There is a Nash equilibrium in the original game with payoff p if and only if there is a strategy for `Eve` in the epistemic game such that for every outcome ρ, for every $1 \leq i \leq N$,

$$\begin{array}{l} \mathsf{MP}_{u_i}(\rho) \geq p_{A_i} \\ \mathsf{MP}_{u_{N+i}}(\rho) \geq -p_{A_i} \\ \mathsf{MP}_{u_{2N+i}}(\rho) \geq -p_{A_i} \end{array}$$

Original weight functions: w_{A_i}

New weight functions: u_i, u_{N+i}, u_{2N+i}

Application to mean-payoff functions: decidability

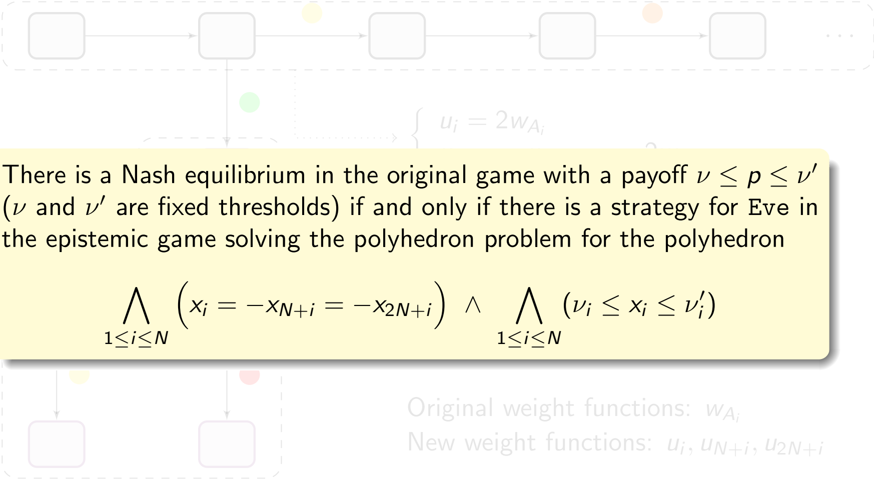

There is a Nash equilibrium in the original game with a payoff $\nu \leq p \leq \nu'$ (ν and ν' are fixed thresholds) if and only if there is a strategy for $\mathtt{Eve}$ in the epistemic game solving the polyhedron problem for the polyhedron

$$\bigwedge_{1 \leq i \leq N} \Big(x_i = -x_{N+i} = -x_{2N+i} \Big) \;\wedge\; \bigwedge_{1 \leq i \leq N} (\nu_i \leq x_i \leq \nu'_i)$$

Application to mean-payoff functions: undecidability

Application to mean-payoff functions: undecidability

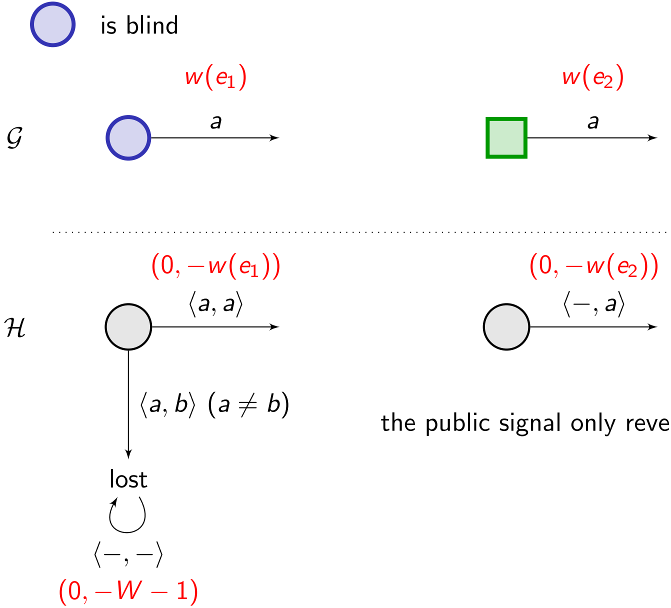

the public signal only reveals lost

Application to mean-payoff functions: undecidability


is blind

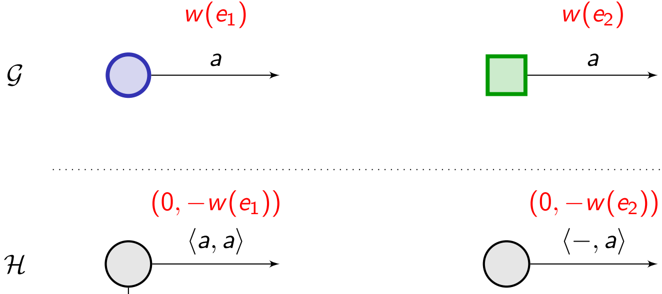

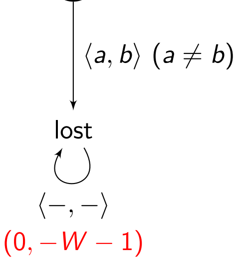

the public signal only reveals lost
but player A_2 has full information

Application to mean-payoff functions: undecidability

has a winning strategy in $\mathcal{G}$ ensuring $\mathsf{MP} > 0$

iff

there is an NE in $\mathcal{H}$ such that player A_2 has a payoff < 0

Application to mean-payoff functions: undecidability


is blind

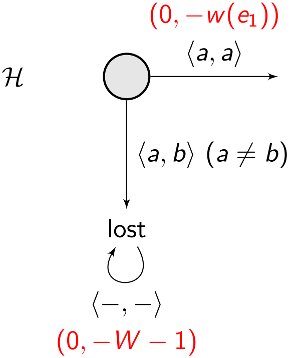

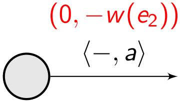

the public signal only reveals lost
but player A_2 has full information

Conclusion

We have:

- proposed a framework for games over graphs with a public signal monitoring Note: framework inspired by [Tom98]
- proposed an abstraction called the epistemic game abstraction, which allows to detect deviators and tocharacterize Nash equilibria in the original game
- used it to show several decidability results.

[Tom98] Tomala. Pure equilibria of repeated games with public observation *(International Journal of Game Theory)*
[RT98] Renault, Tomala. Repeated proximity games *(International Journal of Game Theory)*

Conclusion

We have:

- proposed a framework for games over graphs with a public signal monitoring Note: framework inspired by [Tom98]
- proposed an abstraction called the epistemic game abstraction, which allows to detect deviators and tocharacterize Nash equilibria in the original game
- used it to show several decidability results.

We want to:

- work out the precise complexities
- understand whether one can extend the approach to other communication architectures ([RT98]??)
- understand whether other multiagent frameworks (like fragments of Strategy Logic) can be handled under the assumption of public signal

[Tom98] Tomala. Pure equilibria of repeated games with public observation *(International Journal of Game Theory)*
[RT98] Renault, Tomala. Repeated proximity games *(International Journal of Game Theory)*