Subgame Perfect Equilibrium Quantitative Reachability Games

Jean-François Raskin
Université libre de Bruxelles

iFiP WG 2.2 meeting 2019
Vienna

Partially based on recent works that appeared in CSL'15, FoSSACS'17, GANDALF'18, CONCUR'19 together with Véronique Bruyère, Thomas Brihaye, Noemie Meunier, Aline Goeminne, Arno Pauly, Stéphane Le Roux, Marie Van den Bogaard.

Non zero-sum reachability games

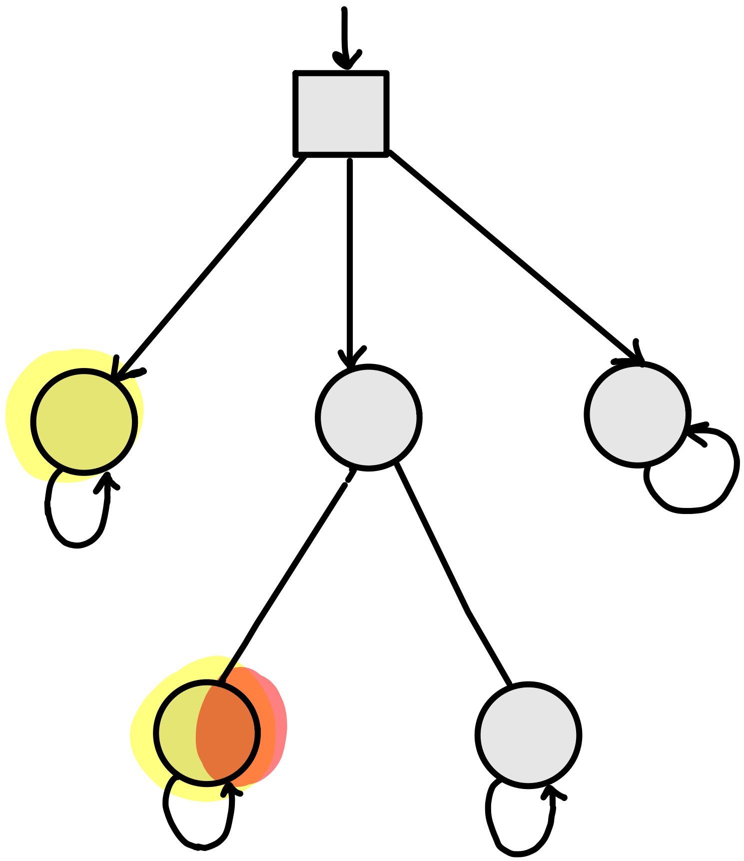

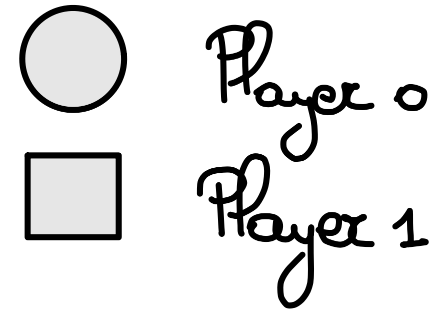

$\Diamond T_0 = \Diamond$

$\Diamond T_1 = \Diamond$

Non zero-sum reachability games

○ Player 0

□ Player 1

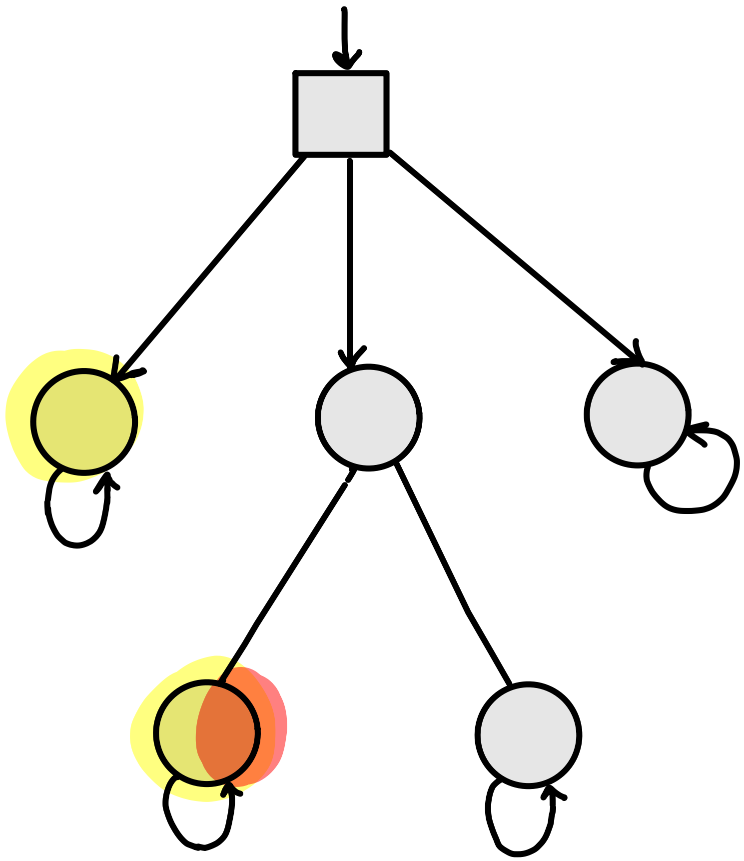

$\Diamond T_0 = \Diamond$

$\Rightarrow$ no winning strategy for Pl. 0

$\Diamond T_1 = \Diamond$

$\Rightarrow$ no winning strategy for Pl. 1

Need for other solution concepts : NE and SPE

Nash equilibrium

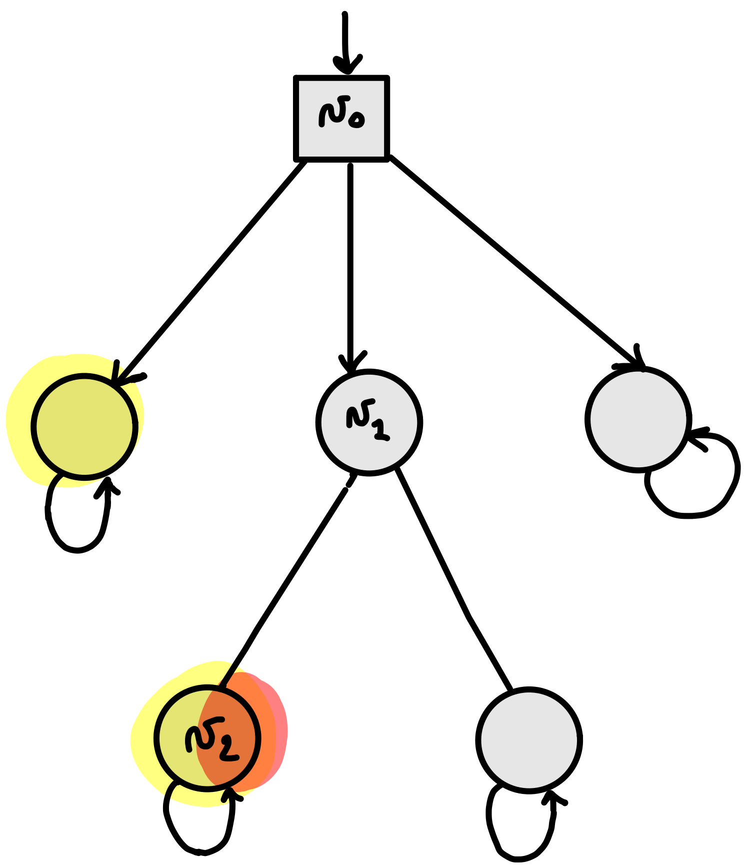

$\underline{\underline{\text{Def}}}$: $\langle \sigma_0, \sigma_1 \rangle$ is a NE if there is no unilateral profitable deviation.

i.e. $Val_0 \langle \sigma_0, \sigma_1 \rangle \geq Val_0 \langle \sigma_0', \sigma_1 \rangle$, $\forall \sigma_0'$

$Val_1 \langle \sigma_0, \sigma_1 \rangle \geq Val_1 \langle \sigma_0, \sigma_1' \rangle$, $\forall \sigma_1'$

Nash equilibrium

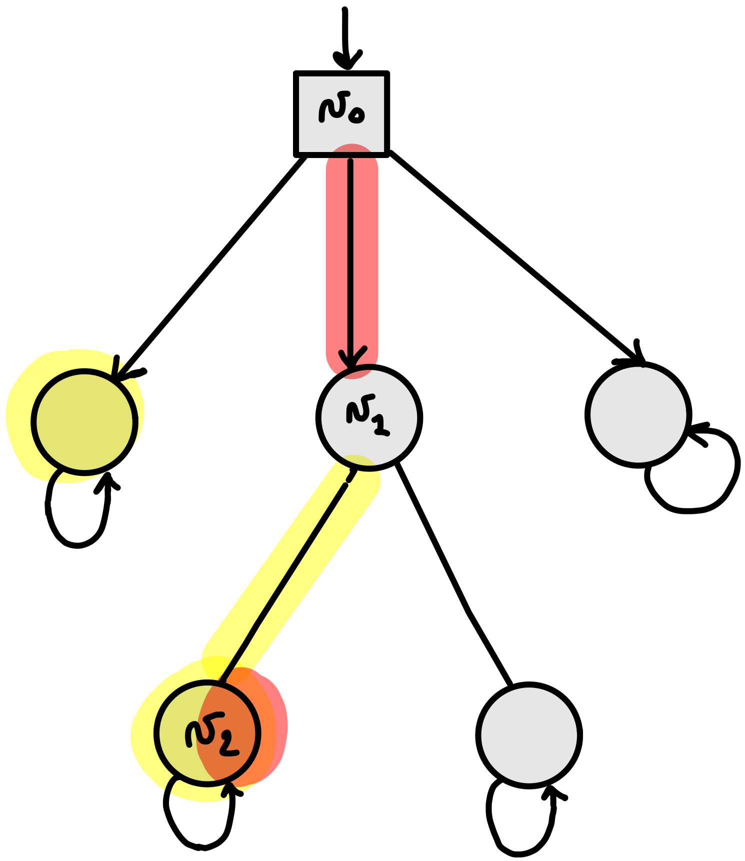

Def: $\langle \sigma_0, \sigma_1 \rangle$ is a NE if there is no unilateral profitable deviation.

i.e. $Val_0 \langle \sigma_0, \sigma_1 \rangle \geq Val_0 \langle \sigma_0', \sigma_1 \rangle$, $\forall \sigma_0'$

$Val_1 \langle \sigma_0, \sigma_1 \rangle \geq Val_1 \langle \sigma_0, \sigma_1' \rangle$, $\forall \sigma_1'$

ex: $v_0\, v_1\, v_2^{\omega}$ is a NE as both players win

Nash equilibrium

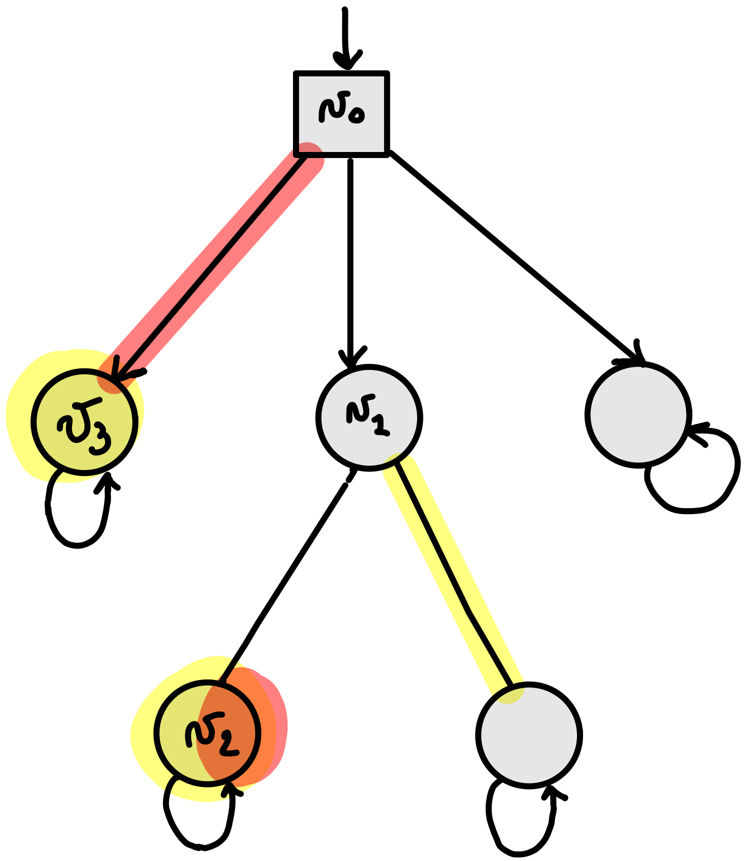

Def: $\langle \sigma_0, \sigma_1 \rangle$ is a NE if there is no unilateral profitable deviation.

i.e. $Val_0 \langle \sigma_0, \sigma_1 \rangle \geq Val_0 \langle \sigma_0', \sigma_1 \rangle, \ \forall \sigma_0'$

$Val_1 \langle \sigma_0, \sigma_1 \rangle \geq Val_1 \langle \sigma_0, \sigma_1' \rangle, \ \forall \sigma_1'$

ex: $v_0 v_3^{\omega}$ is a NE even if both players fail to win

→ no unilateral change is profitable.

Non credible threats

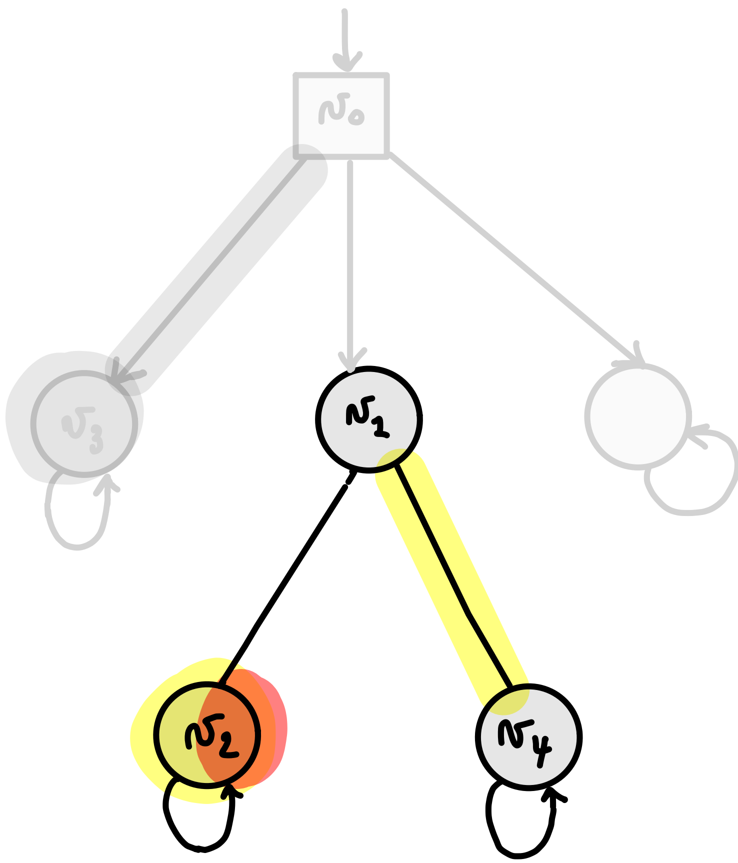

Def: $\langle \sigma_0, \sigma_1 \rangle$ is a NE if there is no unilateral profitable deviation.

i.e. $Val_0 \langle \sigma_0, \sigma_1 \rangle \geq Val_0 \langle \sigma_0', \sigma_1 \rangle$, $\forall \sigma_0'$

$Val_1 \langle \sigma_0, \sigma_1 \rangle \geq Val_1 \langle \sigma_0, \sigma_1' \rangle$, $\forall \sigma_1'$

is NOT SUBGAME PERFECT

$v_1 \longrightarrow v_4$ is not rational

$\equiv$ non credible threat

Subgame perfect equilibrium

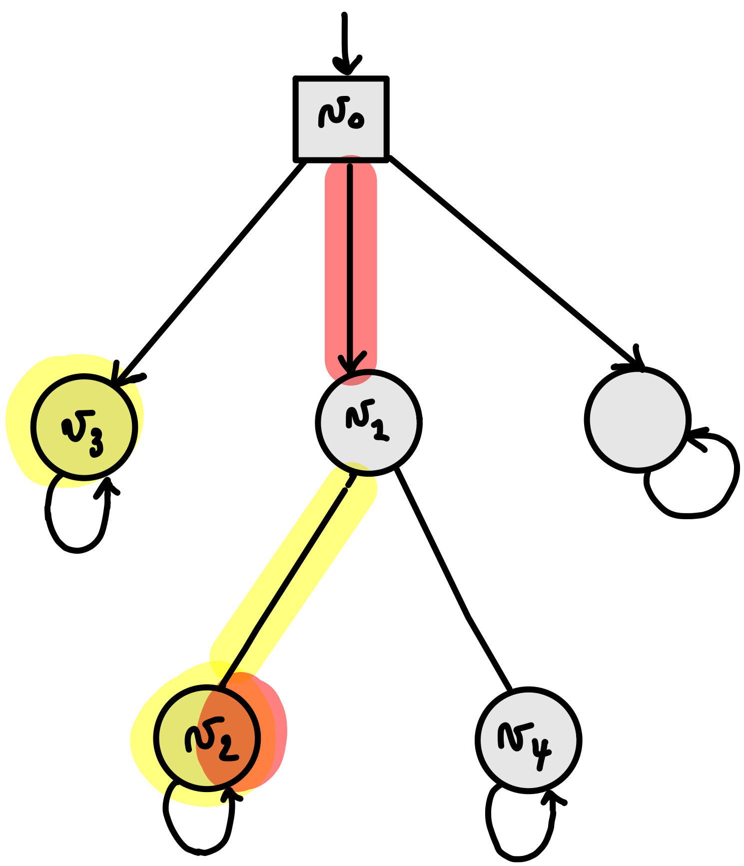

<u>Def</u>: $\langle \sigma_0, \sigma_1 \rangle$ is subgame perfect if there is no unilateral profitable deviation in any subgame

i.e. for all histories h:

$$Val_0^h \langle \sigma_0, \sigma_1 \rangle \geq Val_0^h \langle \sigma_0', \sigma_1 \rangle, \ \forall \sigma_0'$$

$$Val_1^h \langle \sigma_0, \sigma_1 \rangle \geq Val_1^h \langle \sigma_0, \sigma_1' \rangle, \ \forall \sigma_1'$$

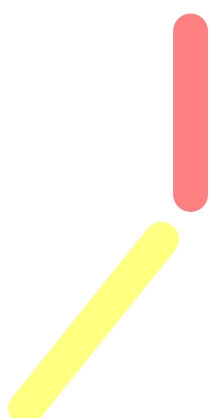 is the only subgame perfect equilibrium

Theorems

NE - Reachability

True for all ω-regular objectives.

THEOREM. NE always exist in reachability games

THEOREM. Constrained existence for NE is NP complete.

P ⊆ Π must win their objectives

SPE - Reachability

True for all ω-regular objectives.

THEOREM. SPE always exist in reachability games

THEOREM. Constrained existence for SPE is PSPACE-C.

for ω-regular objectives some complexity gaps exist ...

Quantitative reachability games

→ minimize the number of steps to reach target

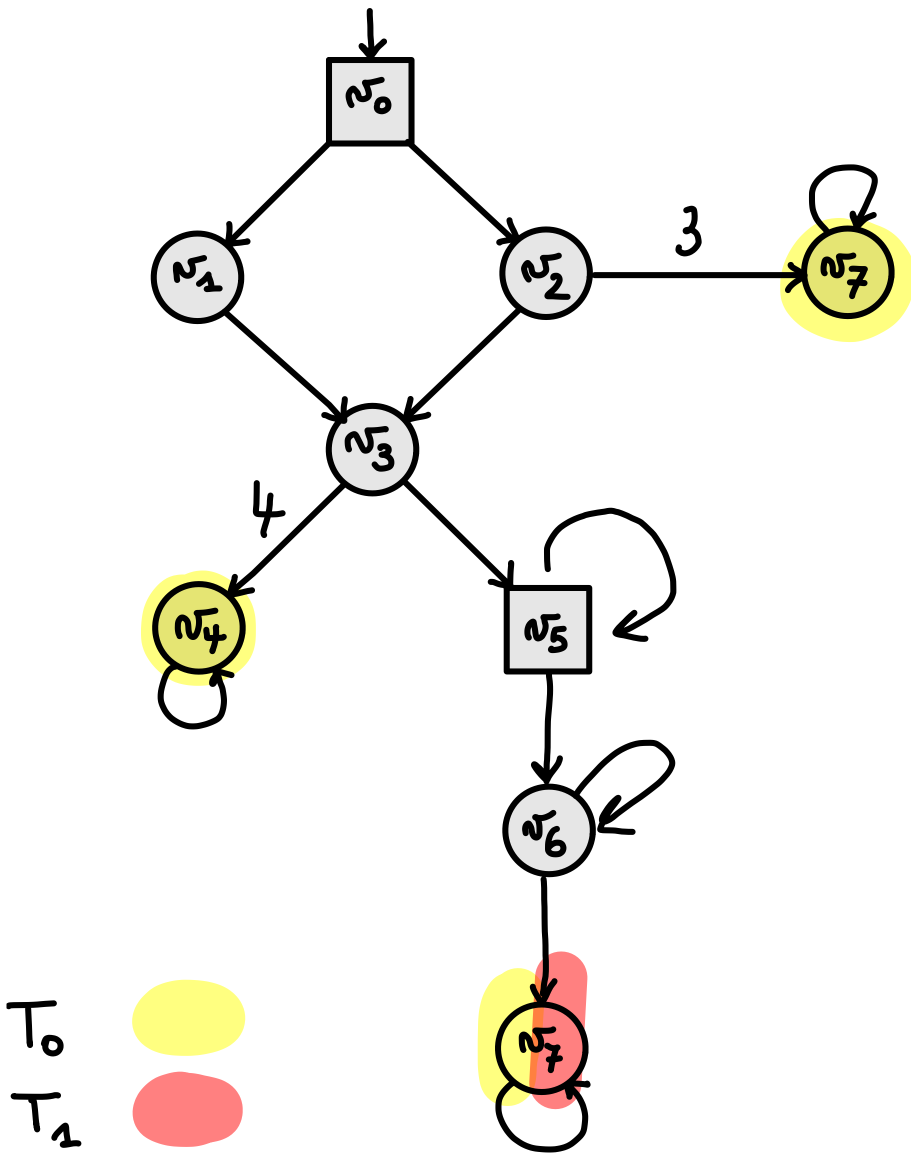

CEP:

"Constrained existence problem"

$\exists \bar{\sigma}$: . $\bar{\sigma}$ is a NE/SPE

. $Val(\bar{\sigma}) \leq \vec{v}$

↳ $\in (\mathbb{N} \cup \{\bot\})^k$

an upper bound for each player

NE - Tool: zero-sum value

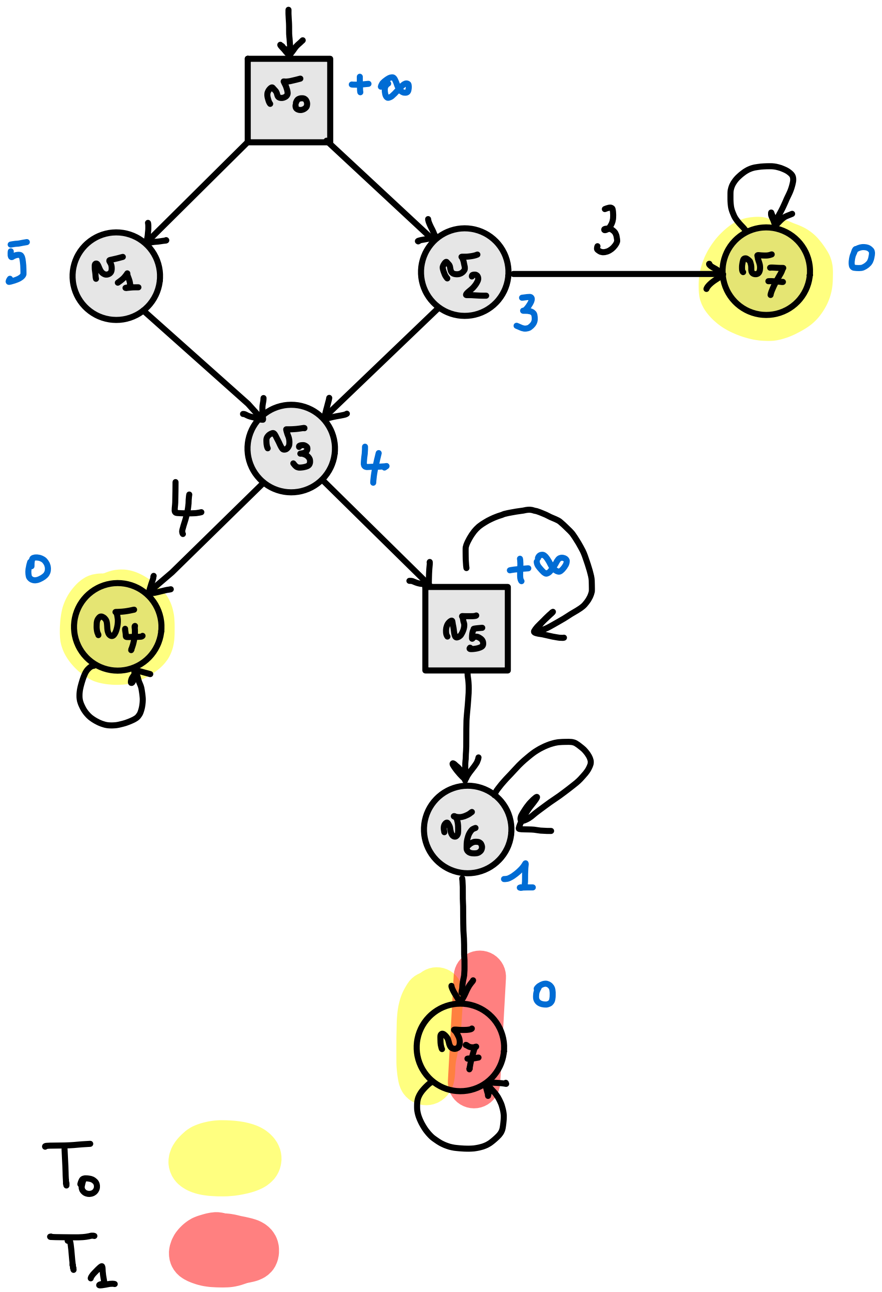

WORST - CASE VALUE

$\lambda : V \to \mathbb{N} \cup \{+\infty\}$

if v belongs to player i

then $\lambda(v)$ = worst-case value that Player i can force from v.

NE - Tool: zero-sum value

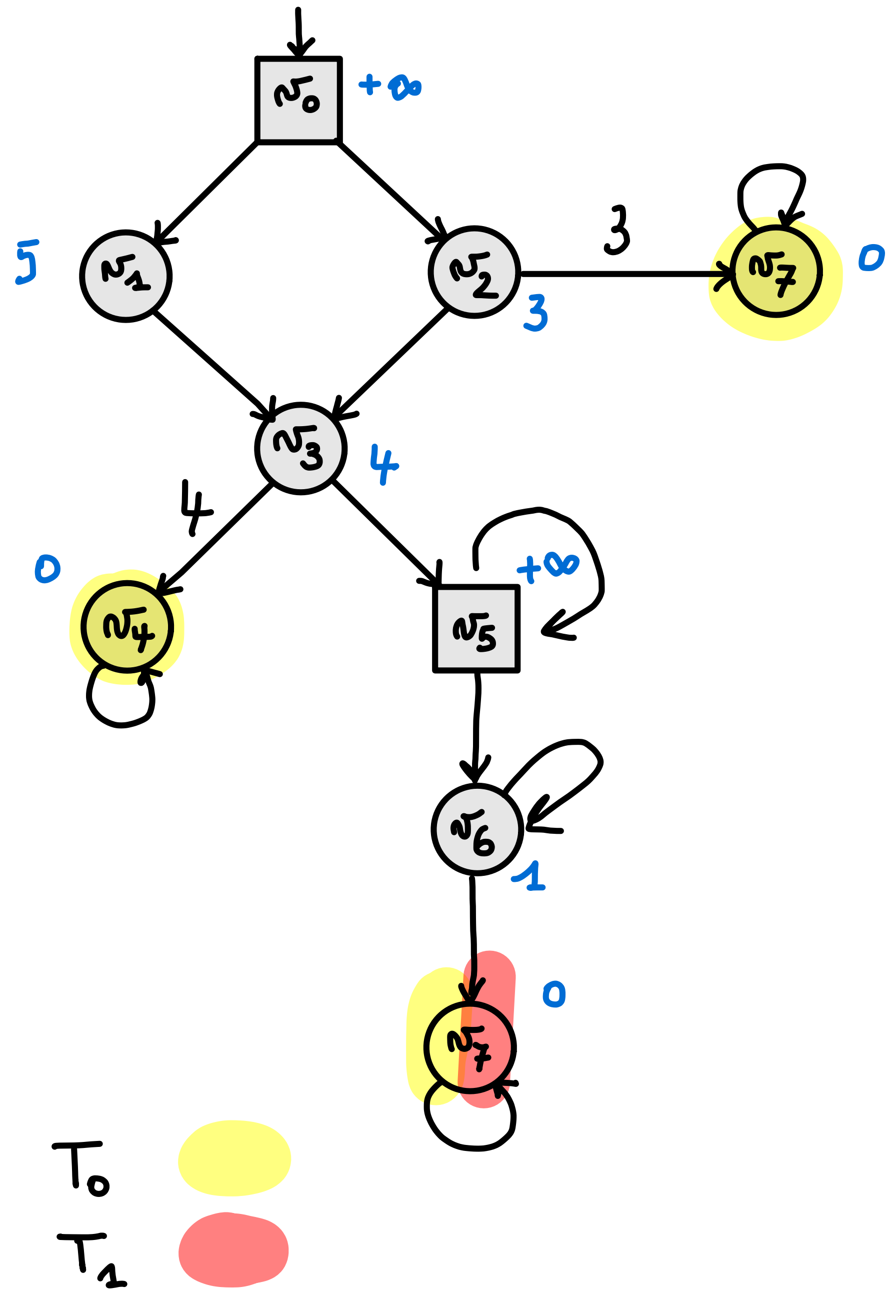

WORST - CASE VALUE

$\lambda : V \to \mathbb{N} \cup \{+\infty\}$

if v belongs to player i
then $\lambda(v)$ = worst-case value
that Player i can force from v.

λ- CONSISTENCY

ρ is λ-consistent if:

$\forall i \geq 0$

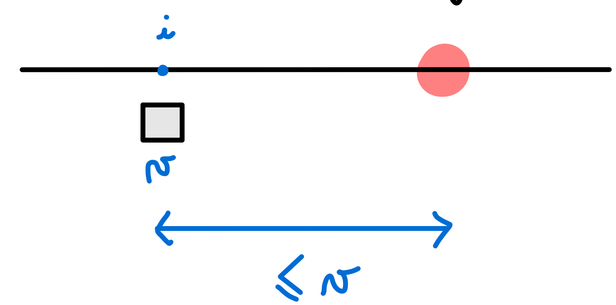

↳ for all players

NE - Tool: zero-sum value

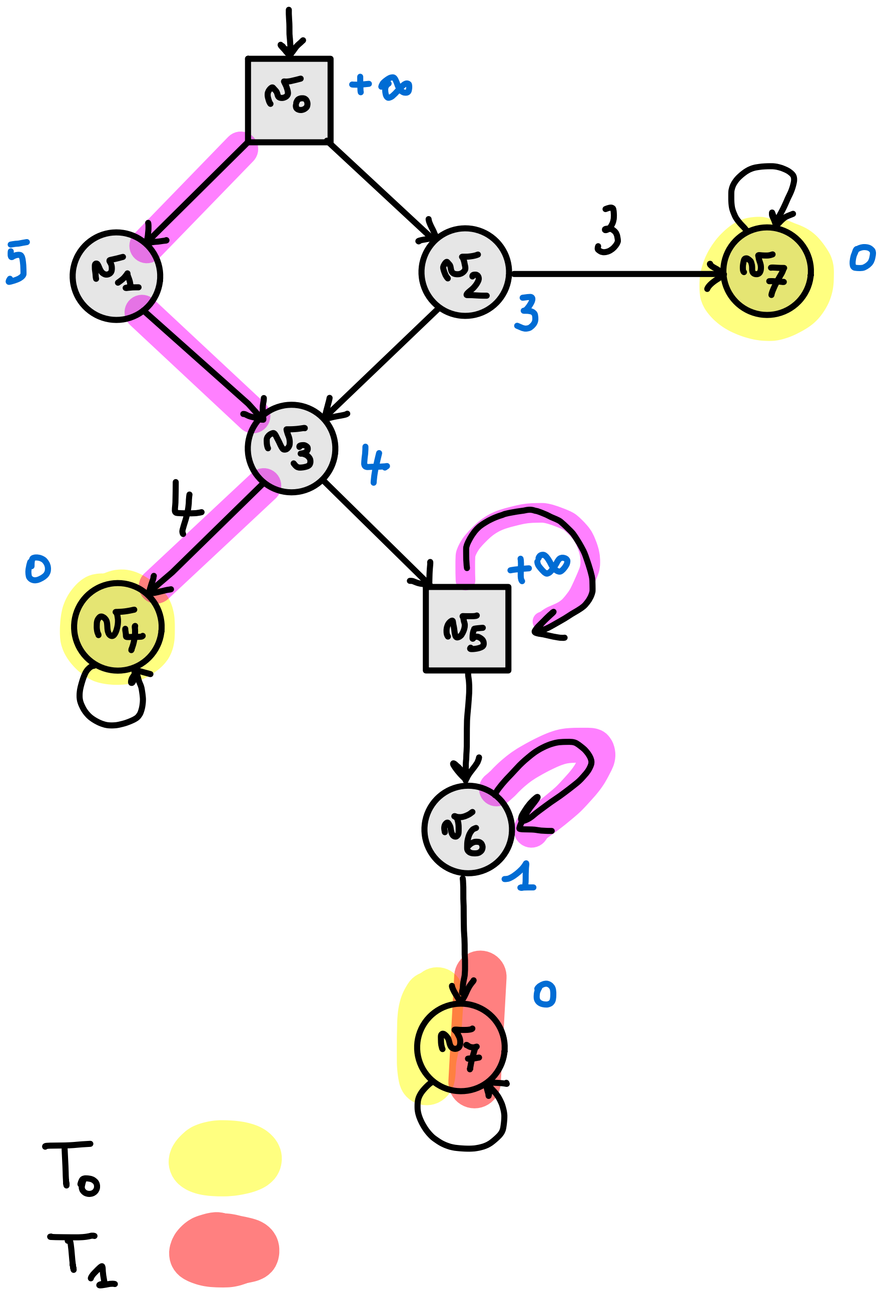

WORST - CASE VALUE

$\lambda : V \to \mathbb{N} \cup \{+\infty\}$

if v belongs to player i

then $\lambda(v)$ = worst-case value that Player i can force from v.

λ- CONSISTENCY

ρ is λ-consistent if:

$\forall i \geq 0$

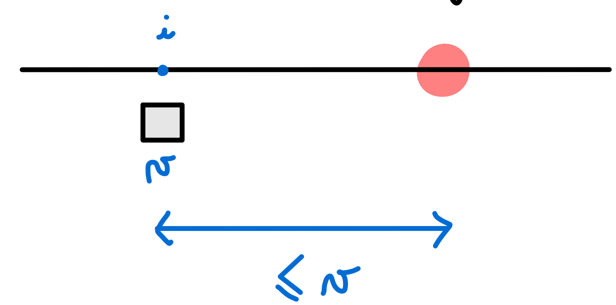

↳ for all players

ex: 7

is λ-consistent and it is a NE.

NE - Tool: zero-sum value

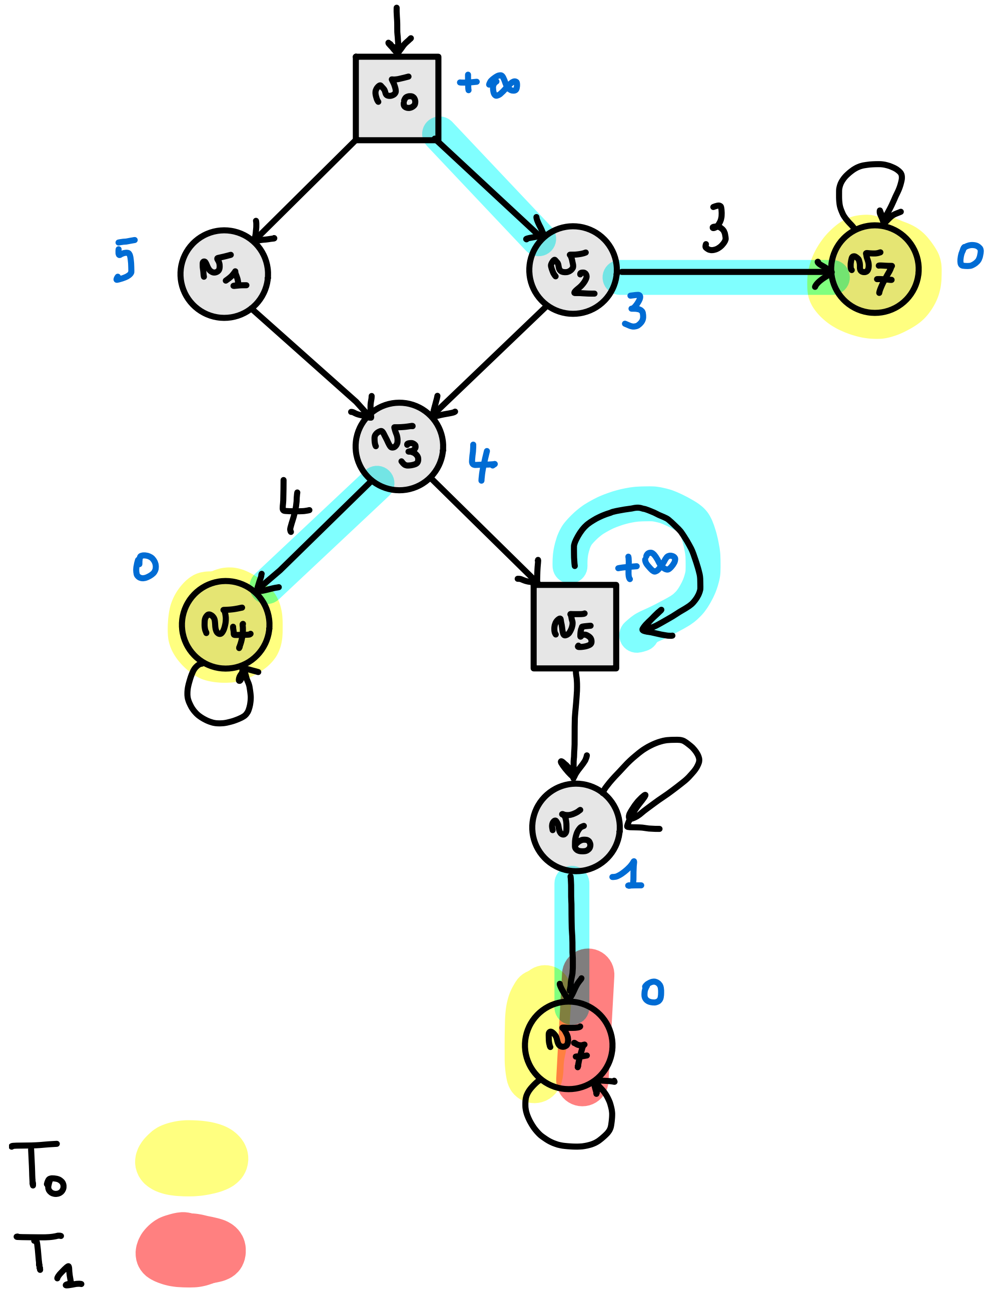

ex:

is λ-consistent and it is a NE.

Why does λ- consistency matter ?

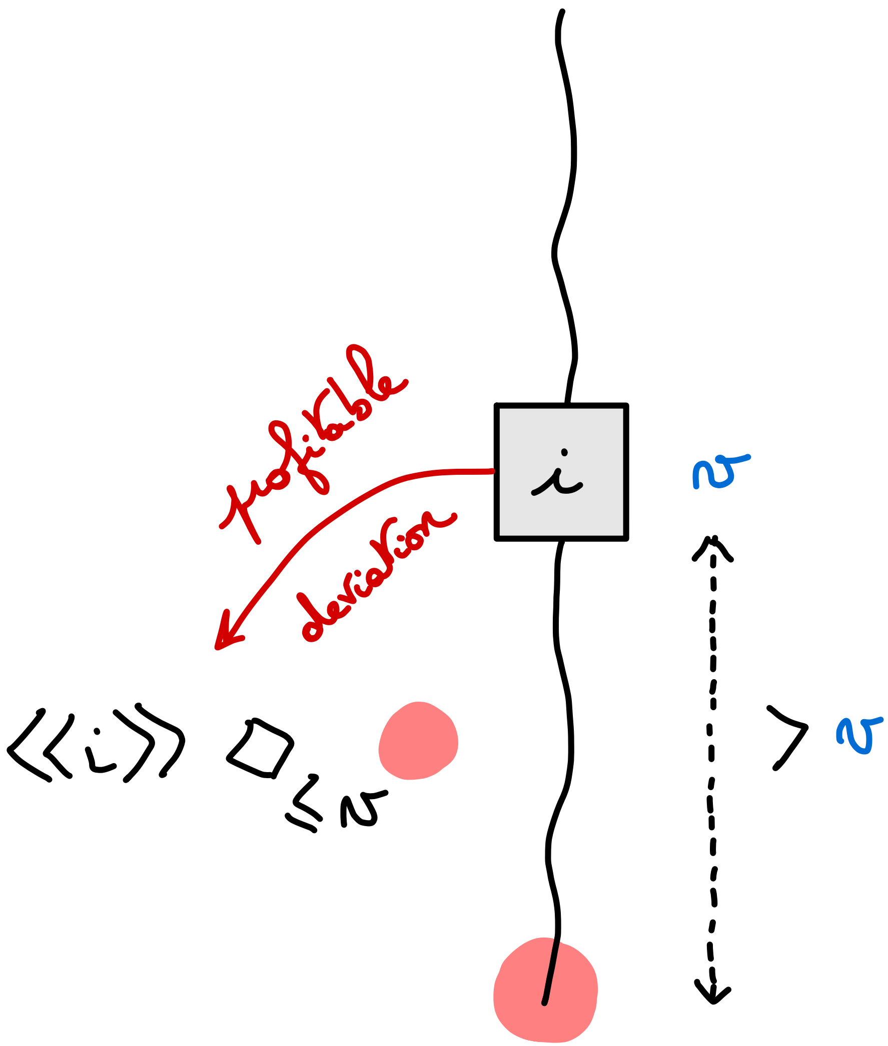

Outcomes of NE are λ-consistent

Why does λ-consistency matter ?

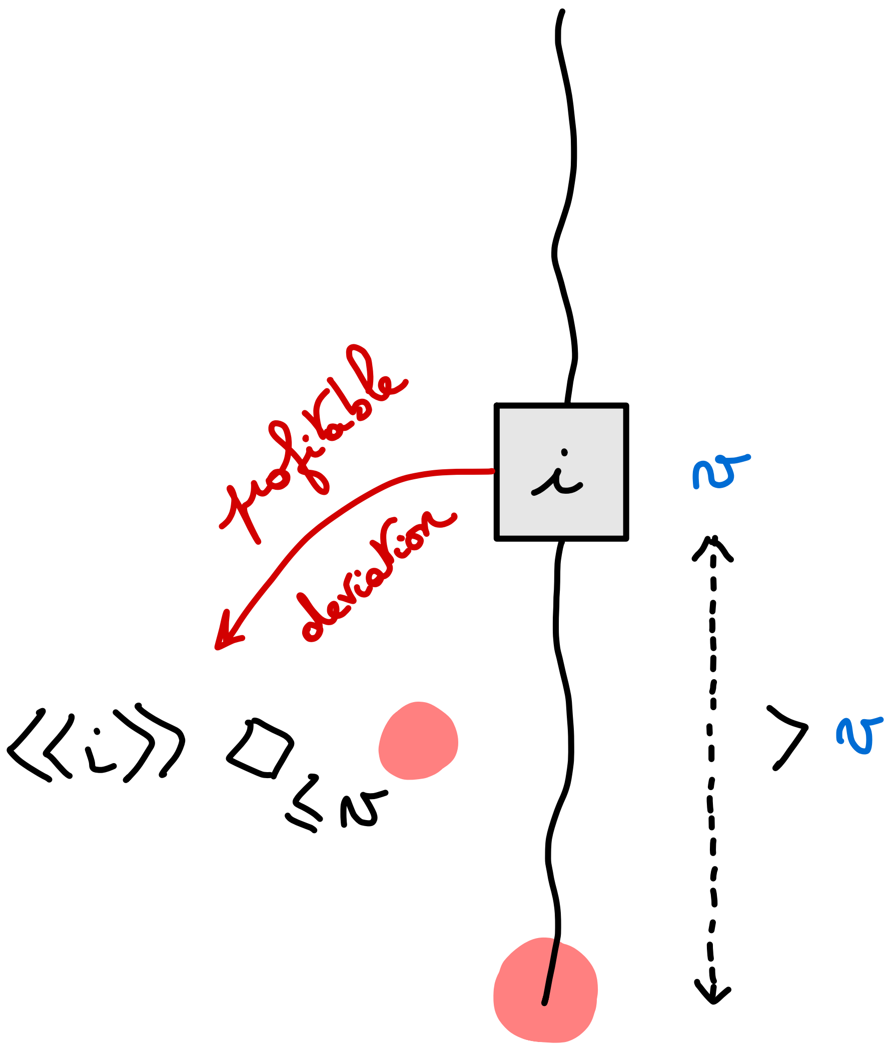

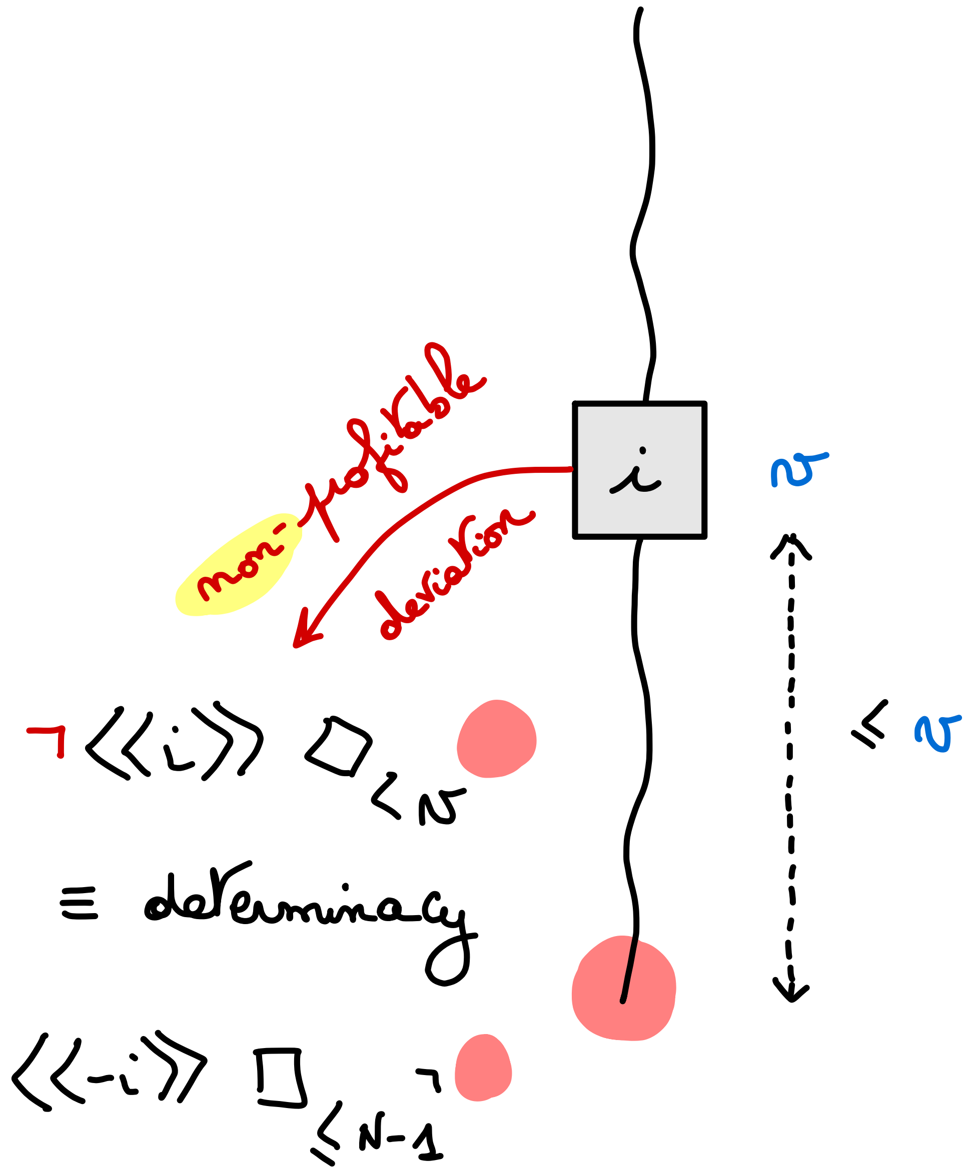

Outcomes of NE are λ-consistent

λ-consistent outcomes are outcomes of NE

Why does λ-consistency matter ?

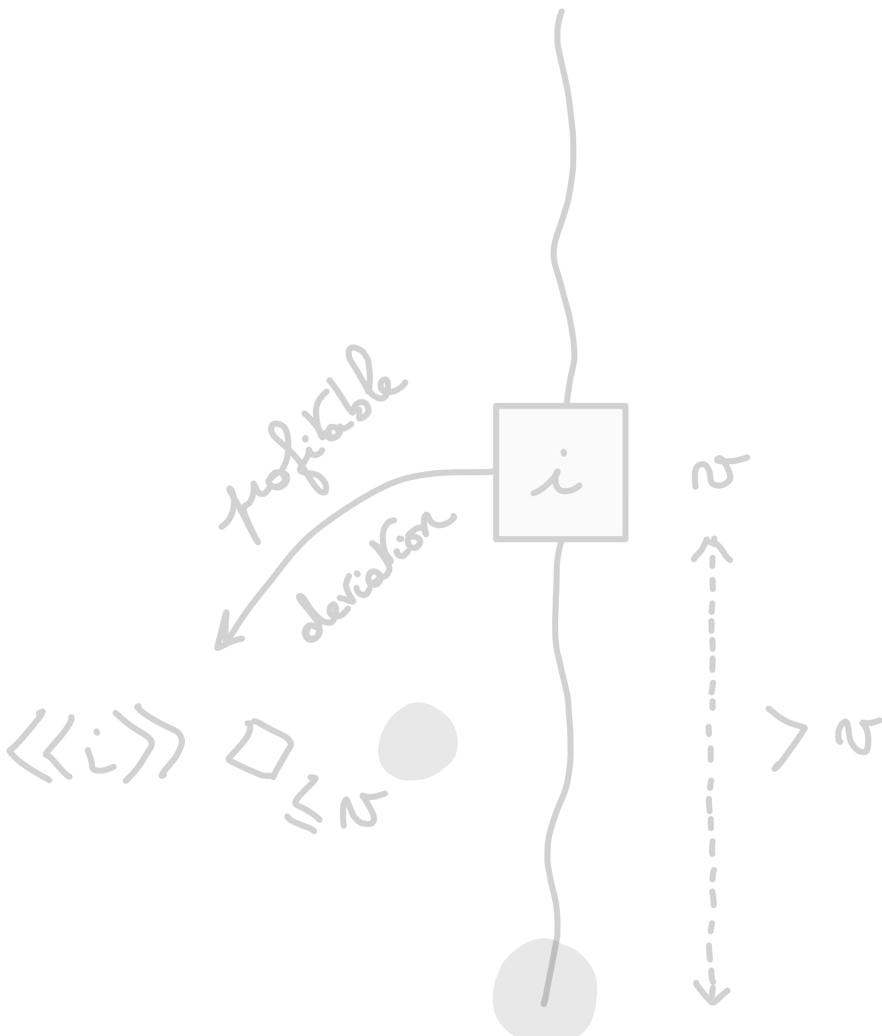

Outcomes of NE are λ-consistent

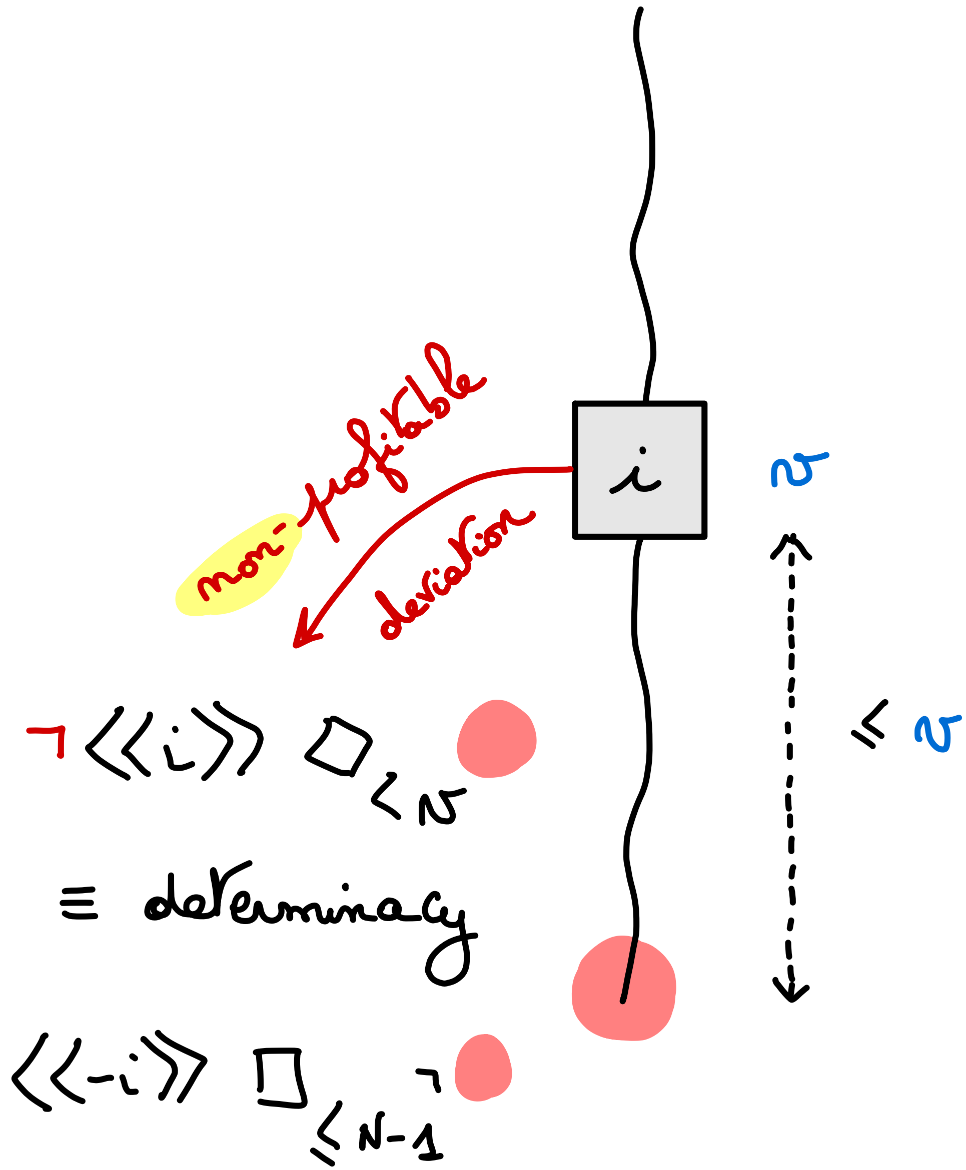

λ-consistent outcomes are outcomes of NE

NE - Tool: zero-sum value

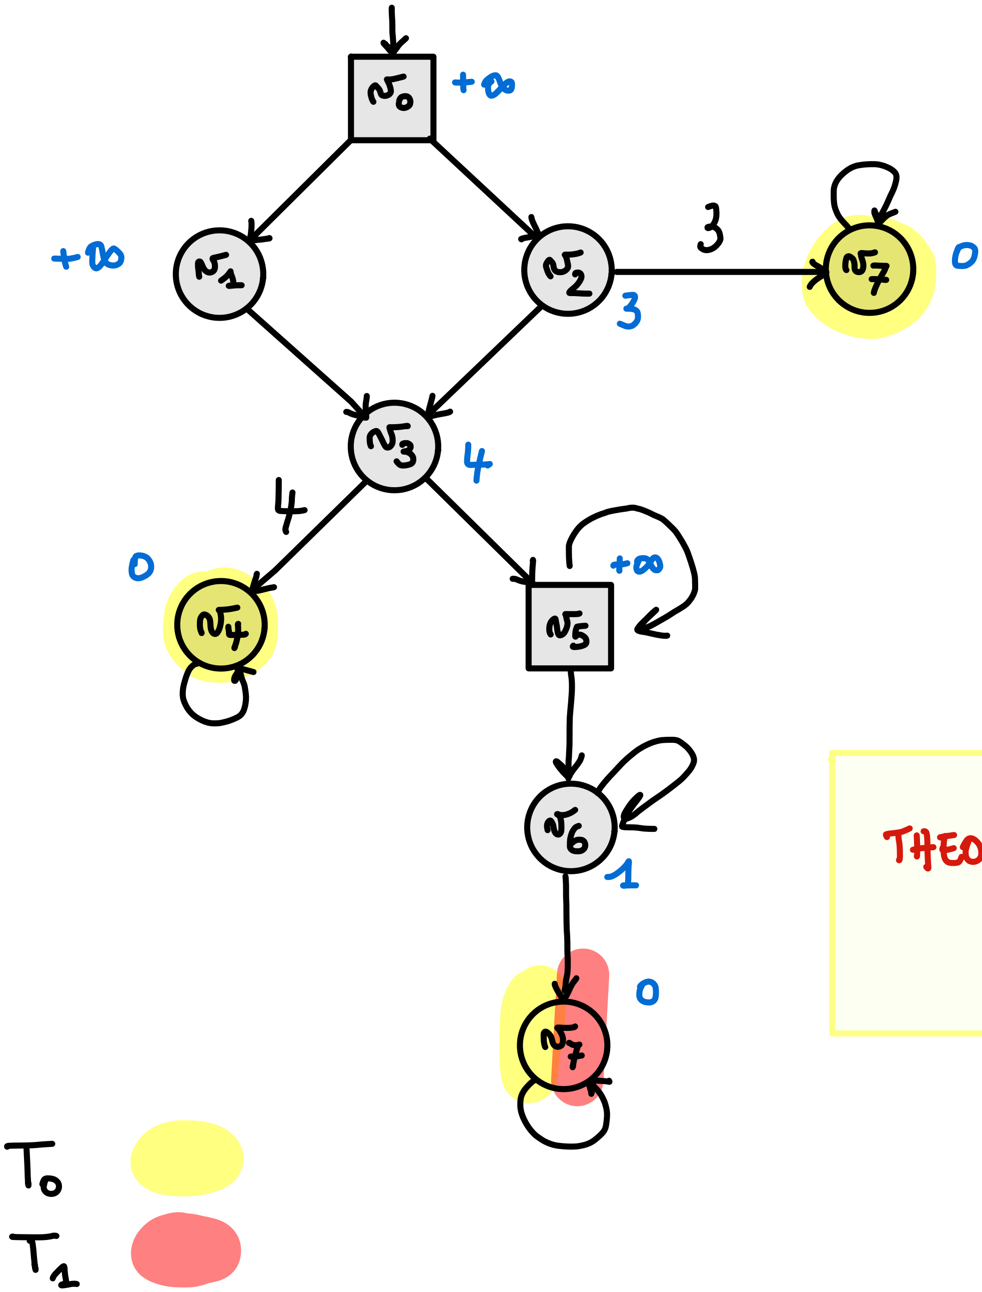

ρ is λ-consistent if:

$\forall i \geq 0$

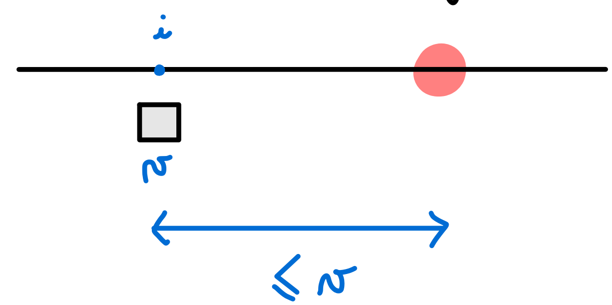

↳ for all players

THEOREM: ρ is the outcome of a NE iff ρ is λ-consistent.

SPE: subgame perfect value

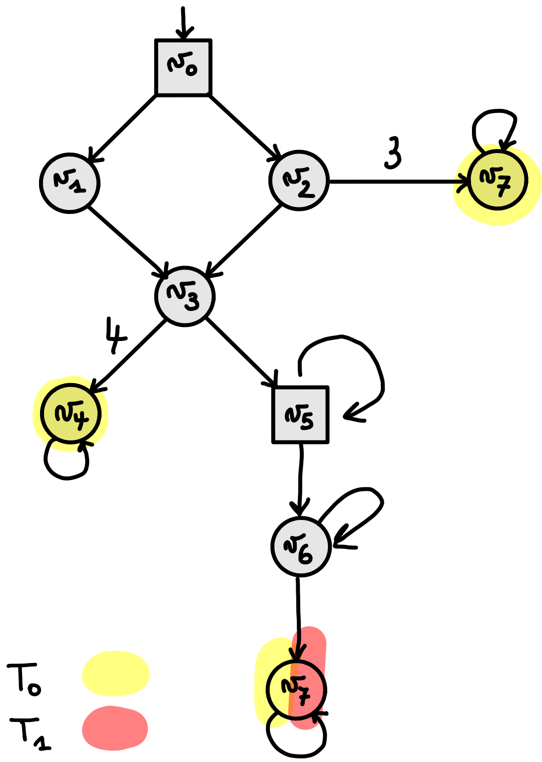

CEP:

"Constrained existence problem"

$\exists \bar{\sigma}$: . $\bar{\sigma}$ is a SPE

. $Val(\bar{\sigma}) \leq \vec{v}$

$\hookrightarrow \in (\mathbb{N} \cup \{\perp\})^k$

SPE: subgame perfect value

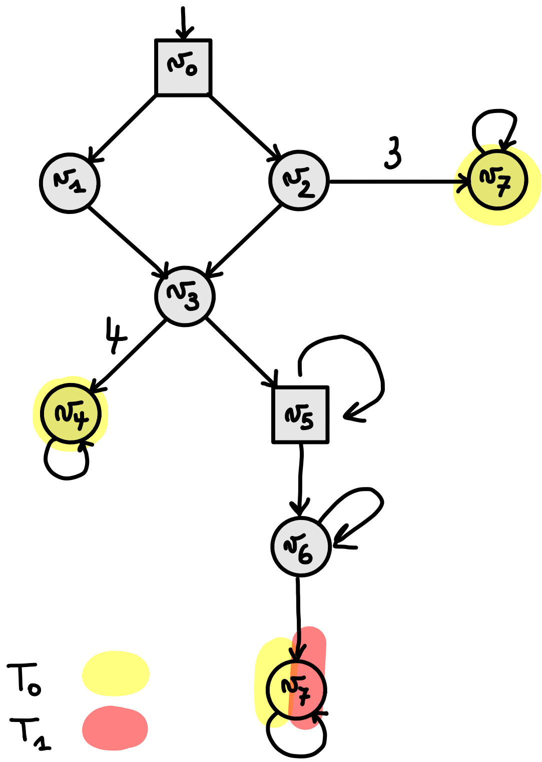

$\lambda_0, \lambda_1, \ldots, \lambda^*$ $\longrightarrow$ seq. of values

SPE: subgame perfect value

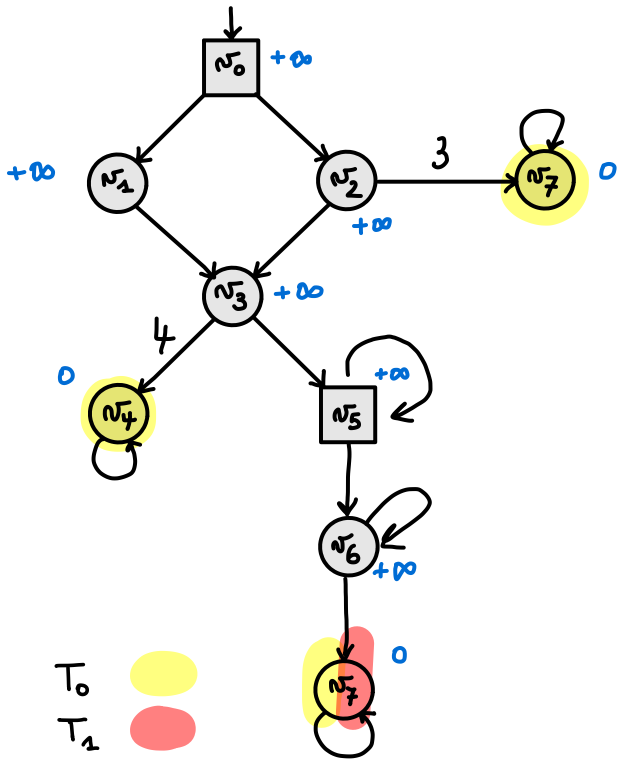

$\lambda_0, \lambda_1, \ldots, \lambda^*$ $\rightarrow$ seq. of values

$$\lambda_0(v) = \begin{cases} 0 & \text{if } v \text{ is in } T_{Owner(v)} \\ +\infty & \text{otherwise.} \end{cases}$$

SPE: subgame perfect value

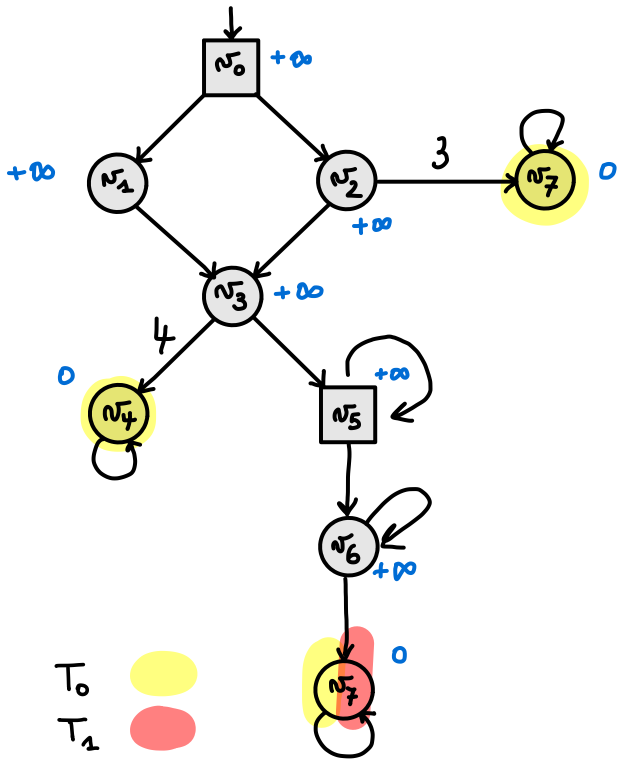

$\lambda_0, \lambda_1, \dots, \lambda^*$ $\longrightarrow$ seq. of values

$$\lambda_0(v) = \begin{cases} 0 & \text{if } v \text{ is in } T_{\mathrm{Owner}(v)} \\ +\infty & \text{otherwise.} \end{cases}$$

Update :

$$\lambda_{k+1}(v) = \begin{cases} 0 & \text{if } v \in T_i \\ 1 + \min\limits_{v' \in \mathrm{Succ}(v)} \max \{ \mathrm{VAL}_i(v'.\rho) \mid v'.\rho \models \lambda_k \} \end{cases}$$

λ_k - consistency

Owner(v) chooses for successor

then he potentially faces worst-case among λ_k consistent paths.

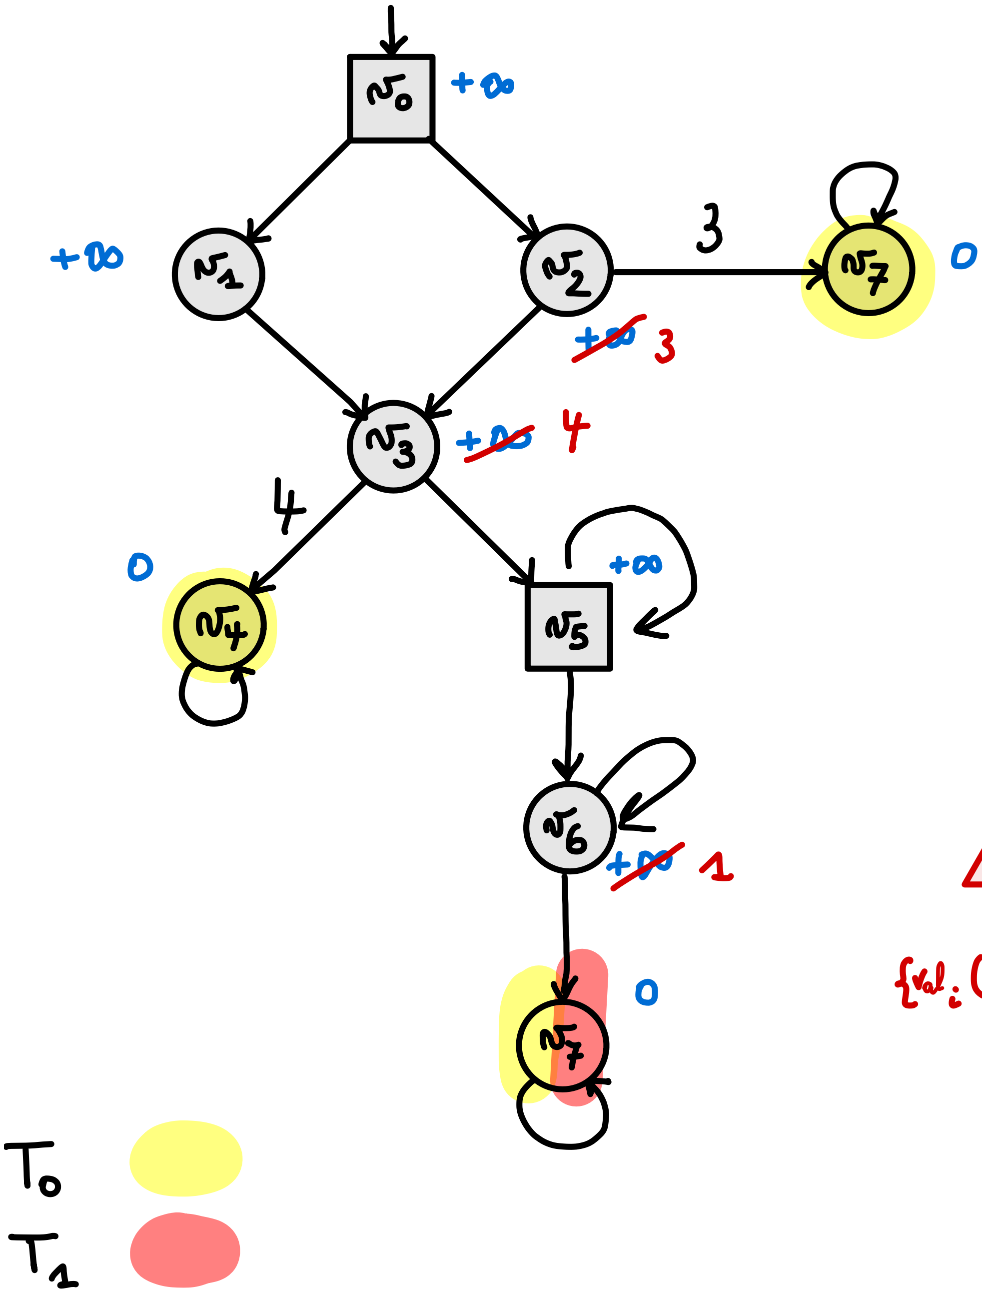

v_0
$+\infty$
v_1
$+\infty$
v_2
$\cancel{+\infty}$ 3
3
v_7
0
v_3
$\cancel{+\infty}$ 4
4
v_4
0
v_5
$+\infty$
v_6
$\cancel{+\infty}$ 1
v_7
0
T_0
T_1

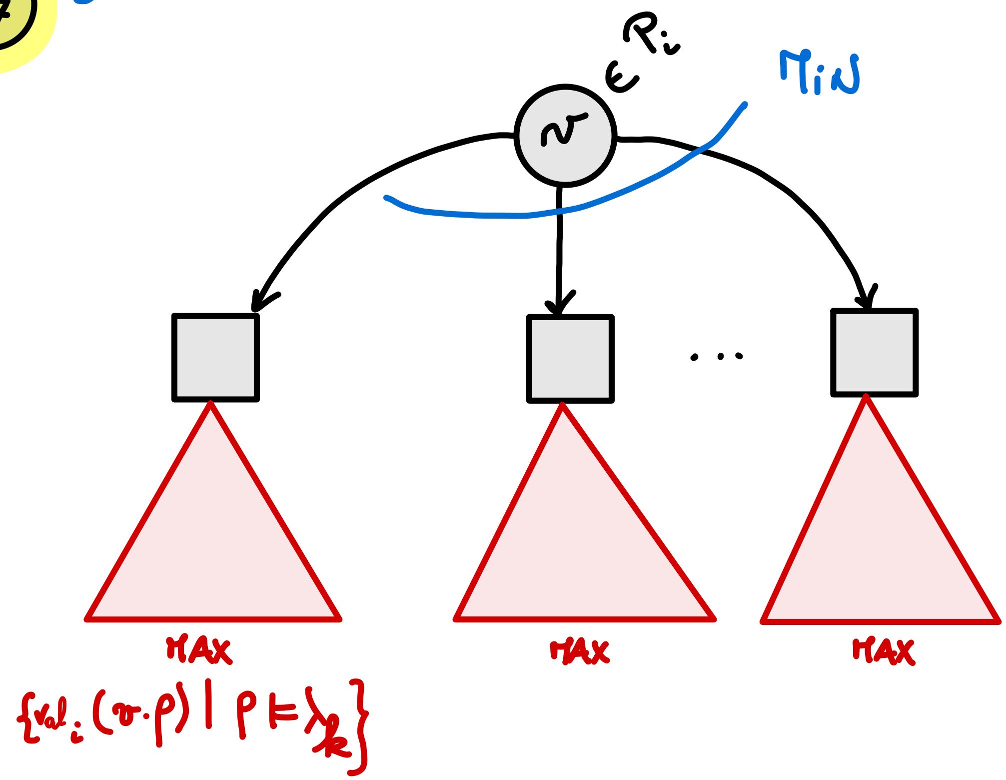

v
$\in P_i$
MIN
MAX
MAX
MAX
$\{val_i(v.\rho) \mid \rho \models \lambda_k\}$

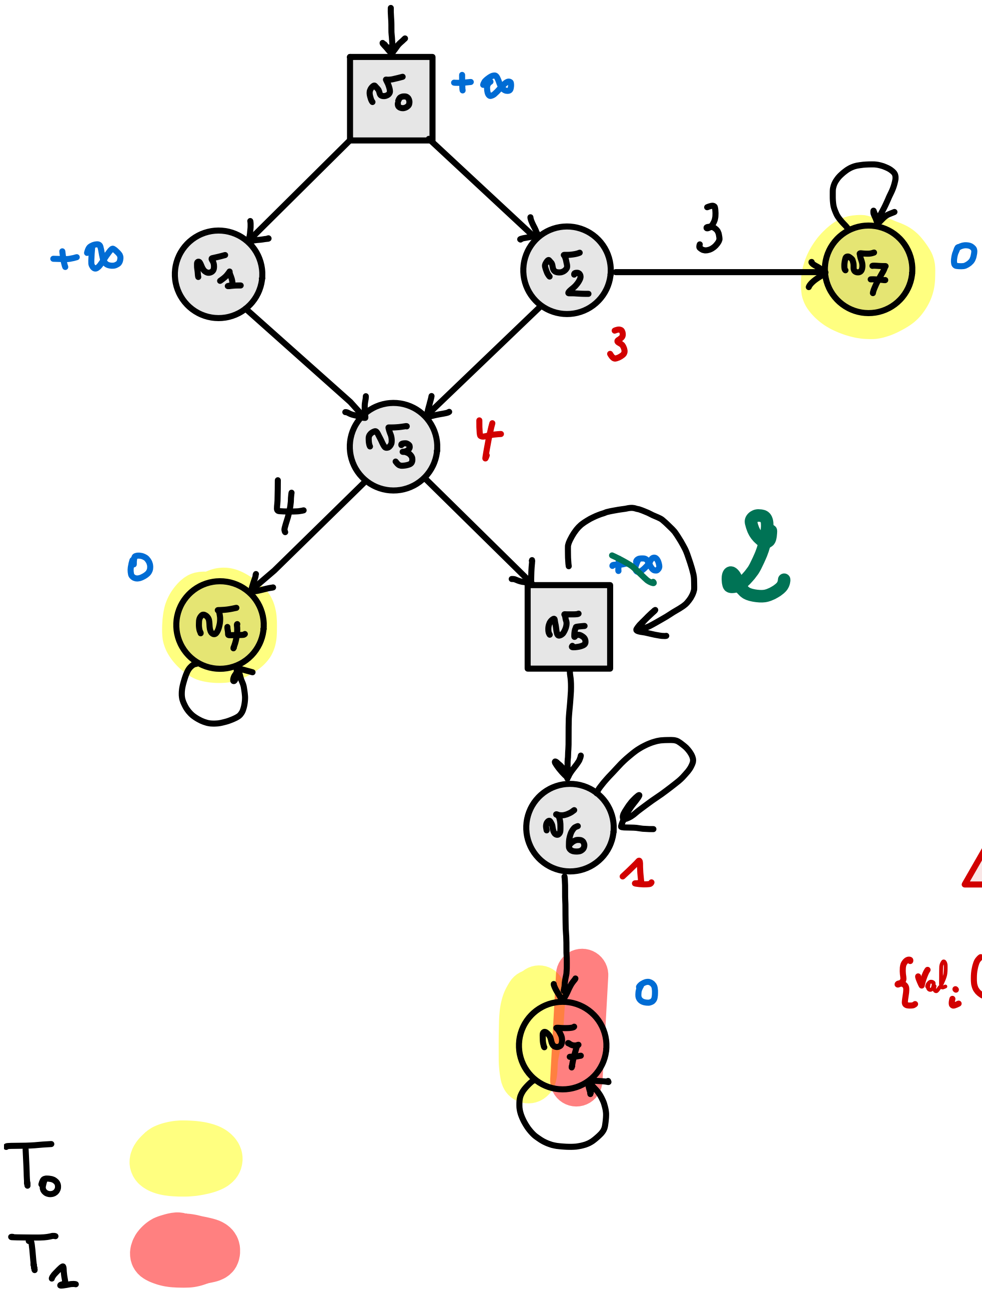
v0
+∞
v1
+∞
v2
3
3
v7
0
v3
4
4
v4
0
v5
+∞
2
v6
1
v7
0
T0
T1

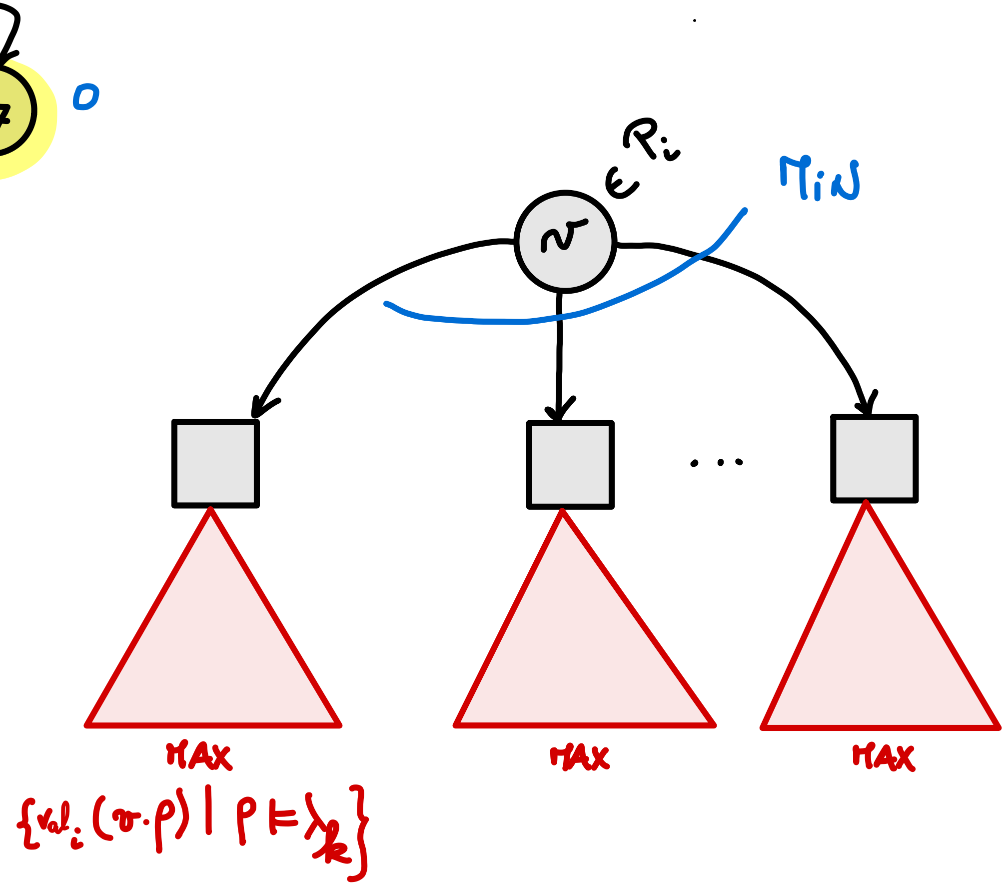
∈ Ri
v
MIN
MAX
MAX
MAX

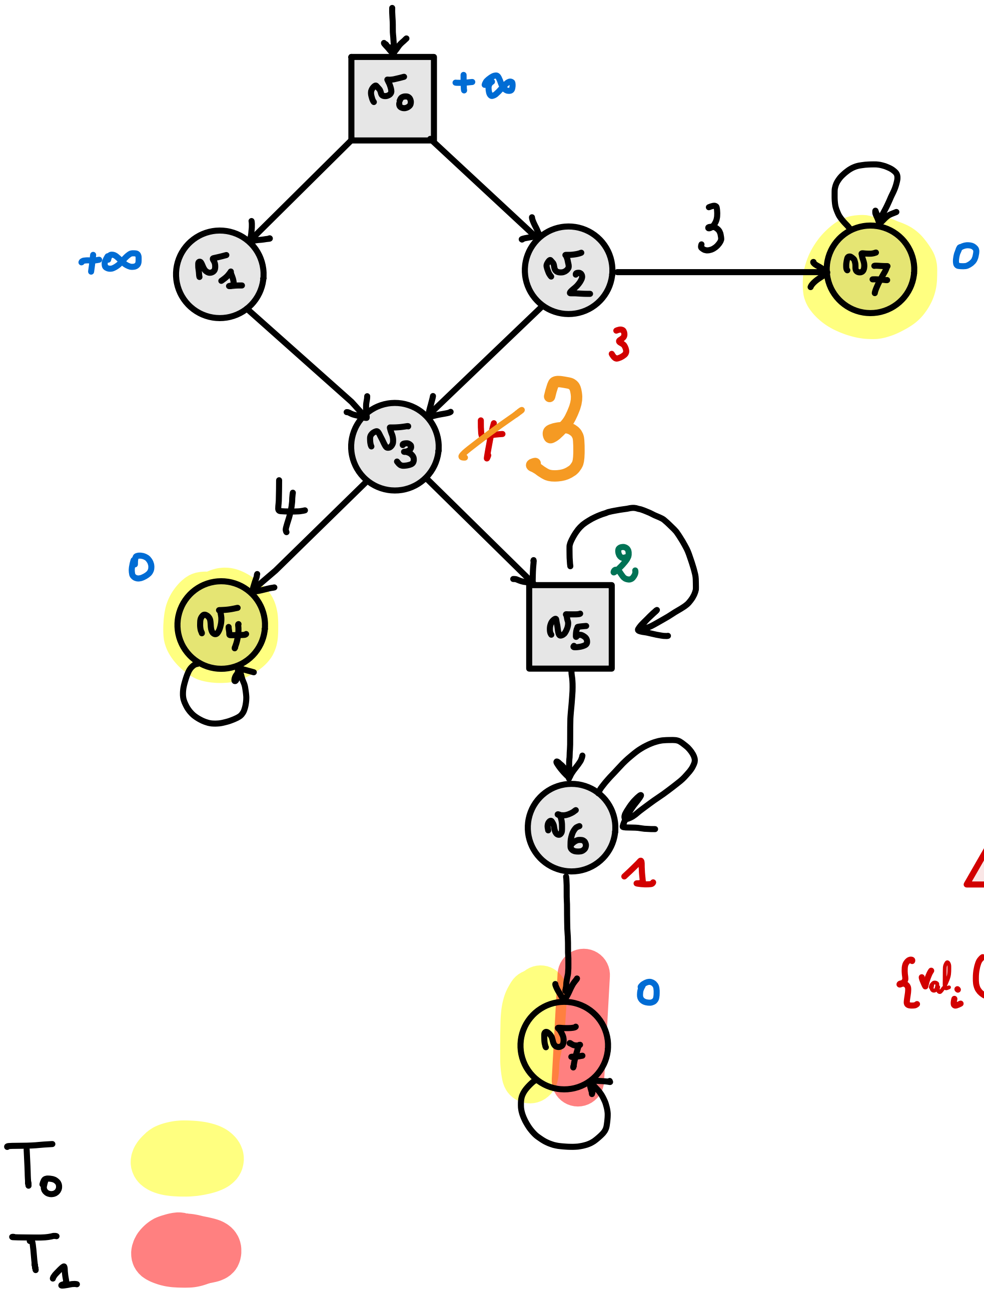

v_0
$+\infty$
v_1
$+\infty$
v_2
3
3
v_7
0
v_3
~~4~~ 3
4
v_4
0
v_5
2
v_6
1
v_7
0
T_0
T_1

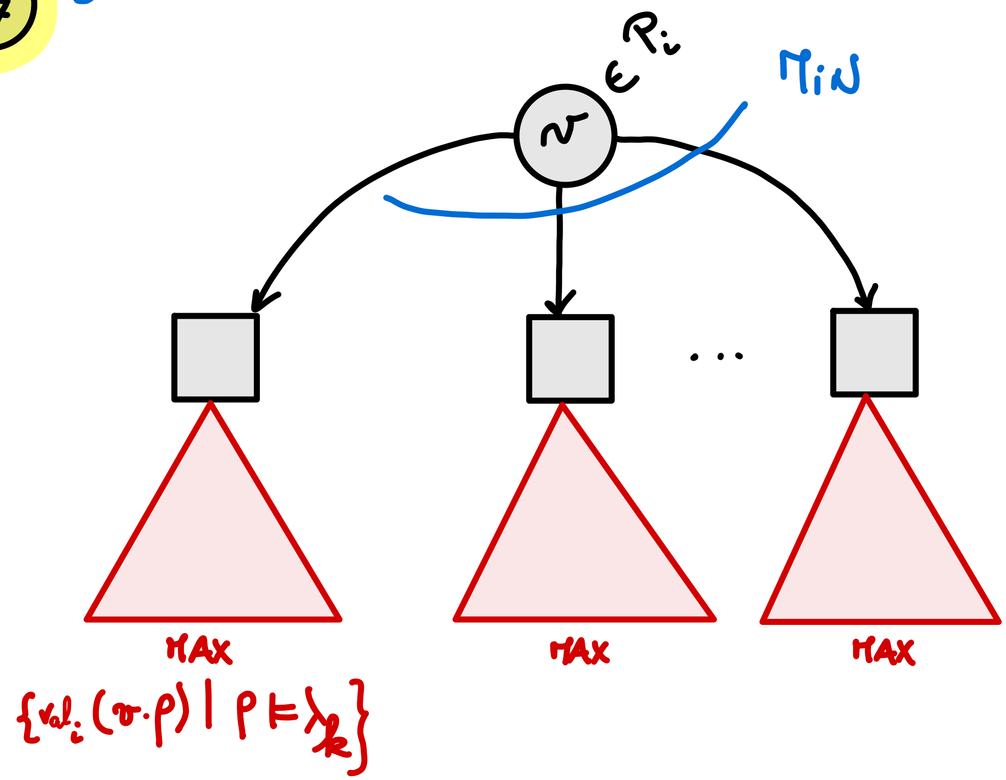

$v \in P_i$
MIN
MAX
MAX
MAX
$\{val_i(v\cdot\rho) \mid \rho \models \lambda_k\}$

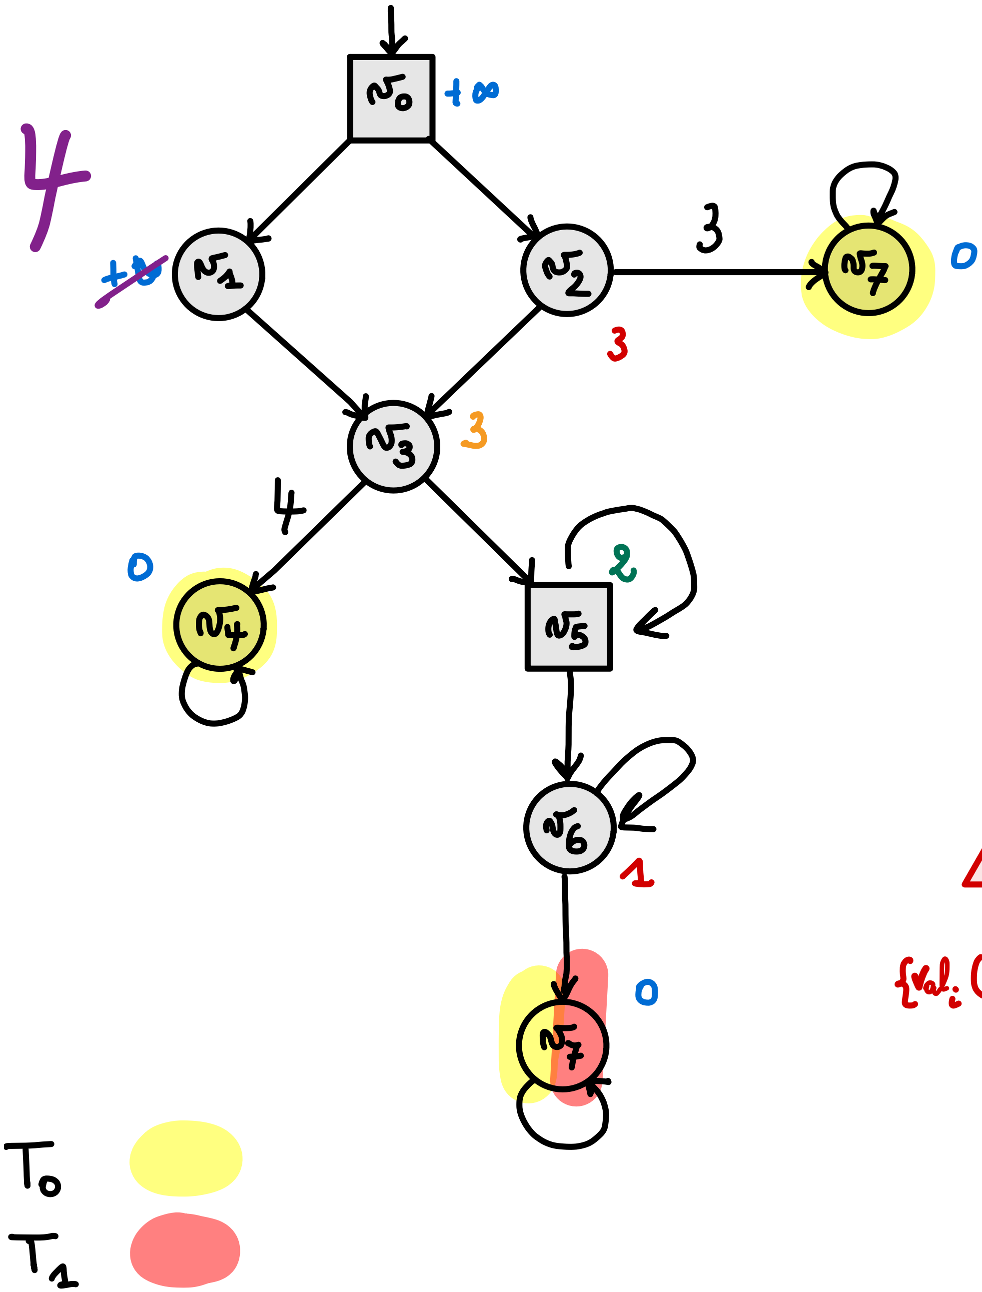

4
v_0 +∞
v_1
v_2 3
3
v_7 0
v_3 3
4
v_4 0
v_5 2
v_6 1
v_7 0
T_0
T_1

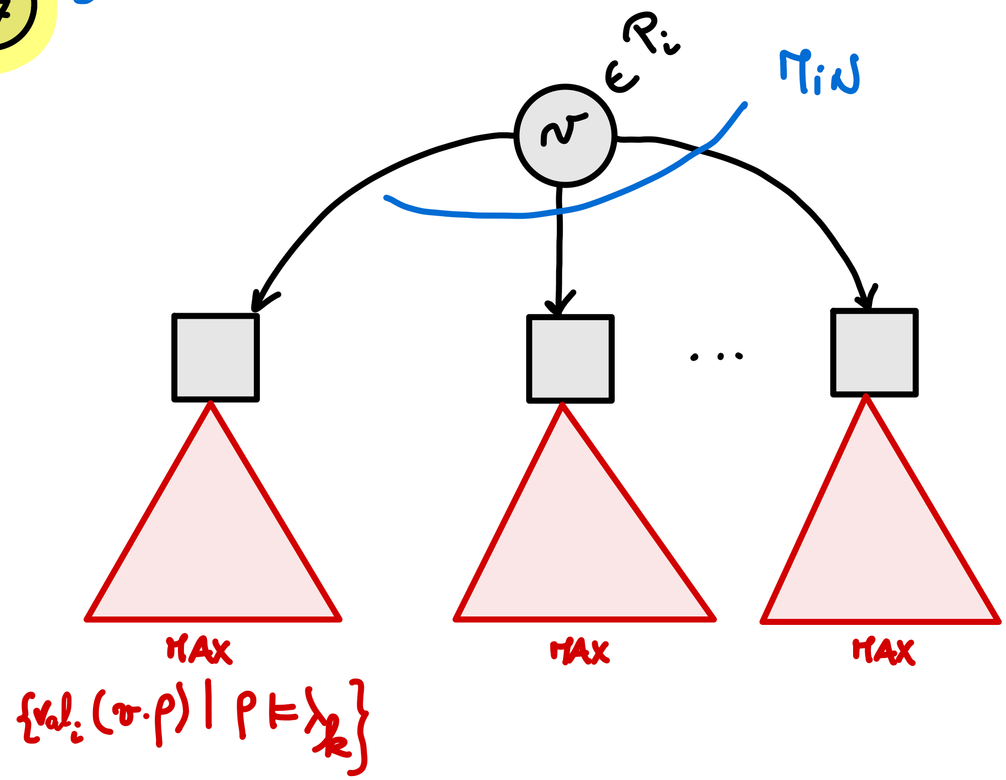

$v \in P_i$
MIN
MAX
MAX
MAX
$\{val_i(v\cdot\rho) \mid \rho \models \lambda_k\}$

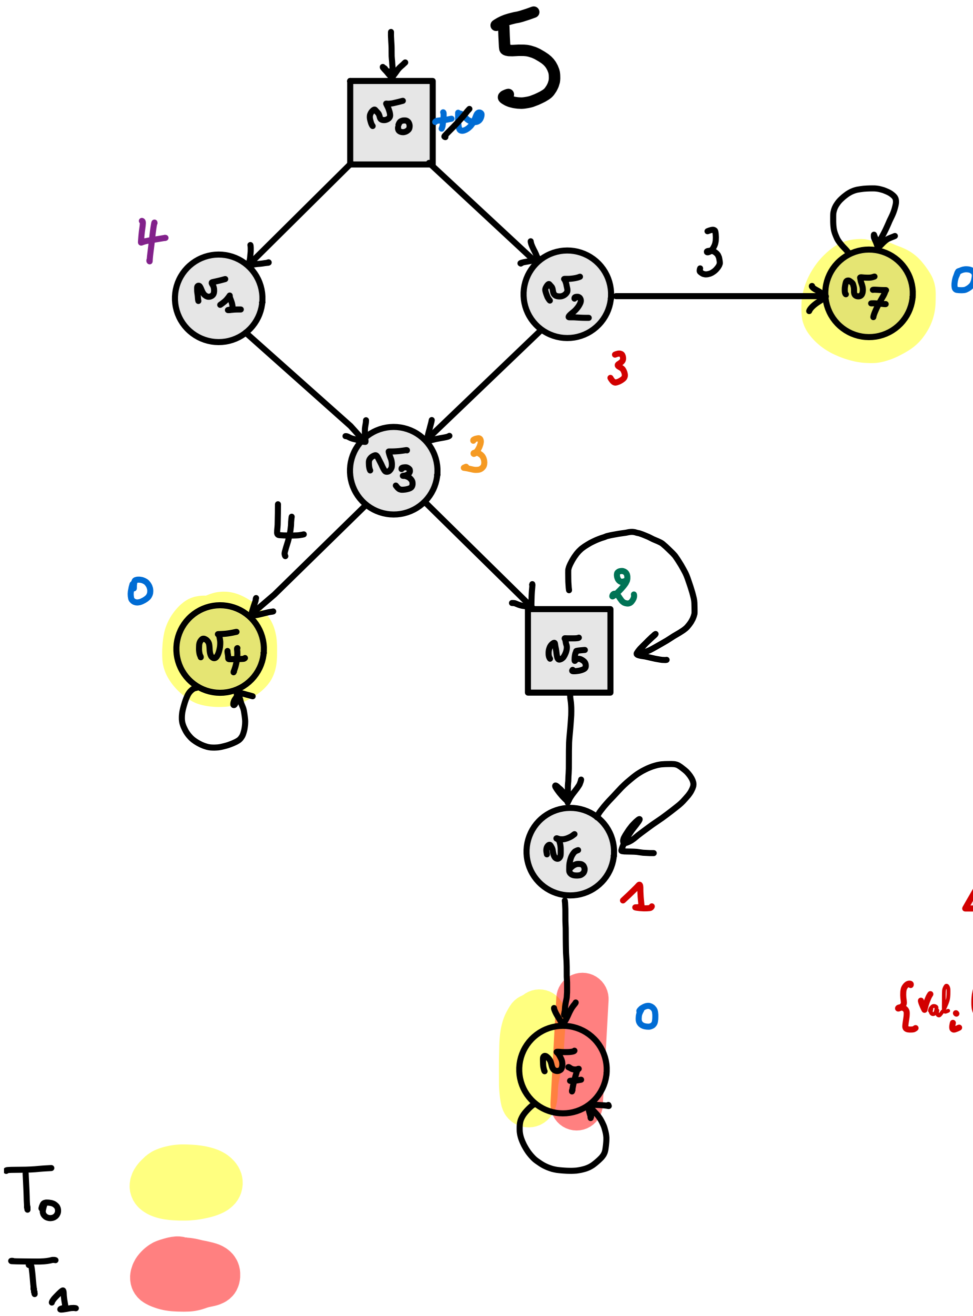

5
v0
4
v1
v2
3
v7
0
3
v3
3
4
0
v4
v5
2
v6
1
v7
0
T0
T1

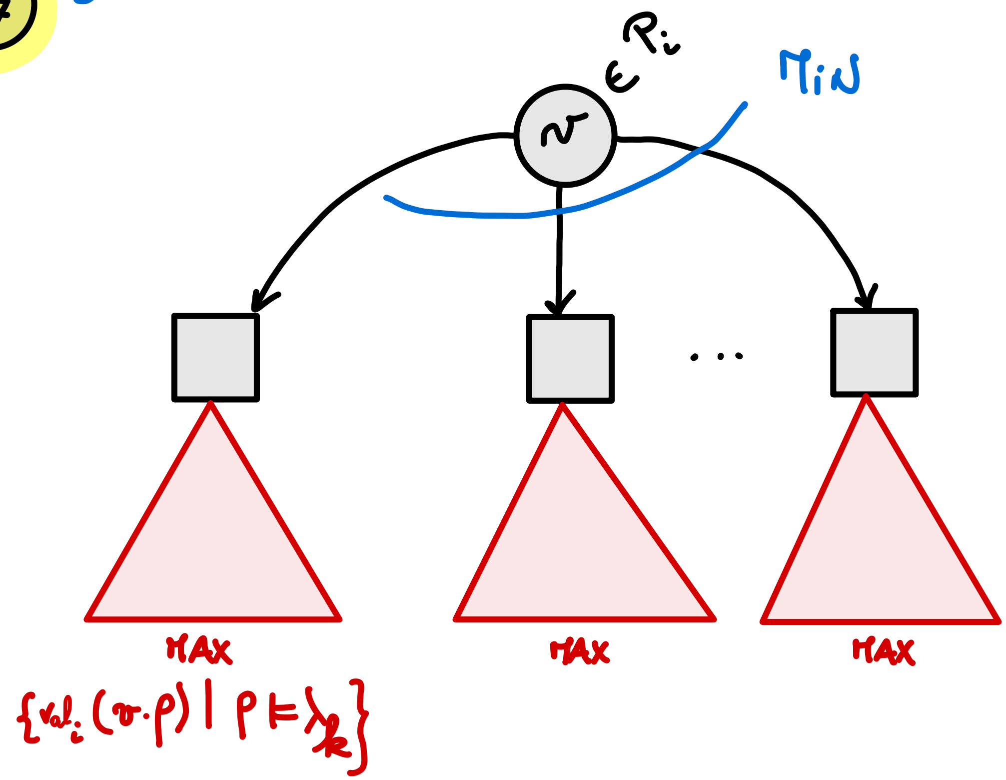

v
∈ Pi
MIN
MAX
MAX
MAX
{val_i(v·ρ) | ρ ⊨ λ_k}

fixed point reached

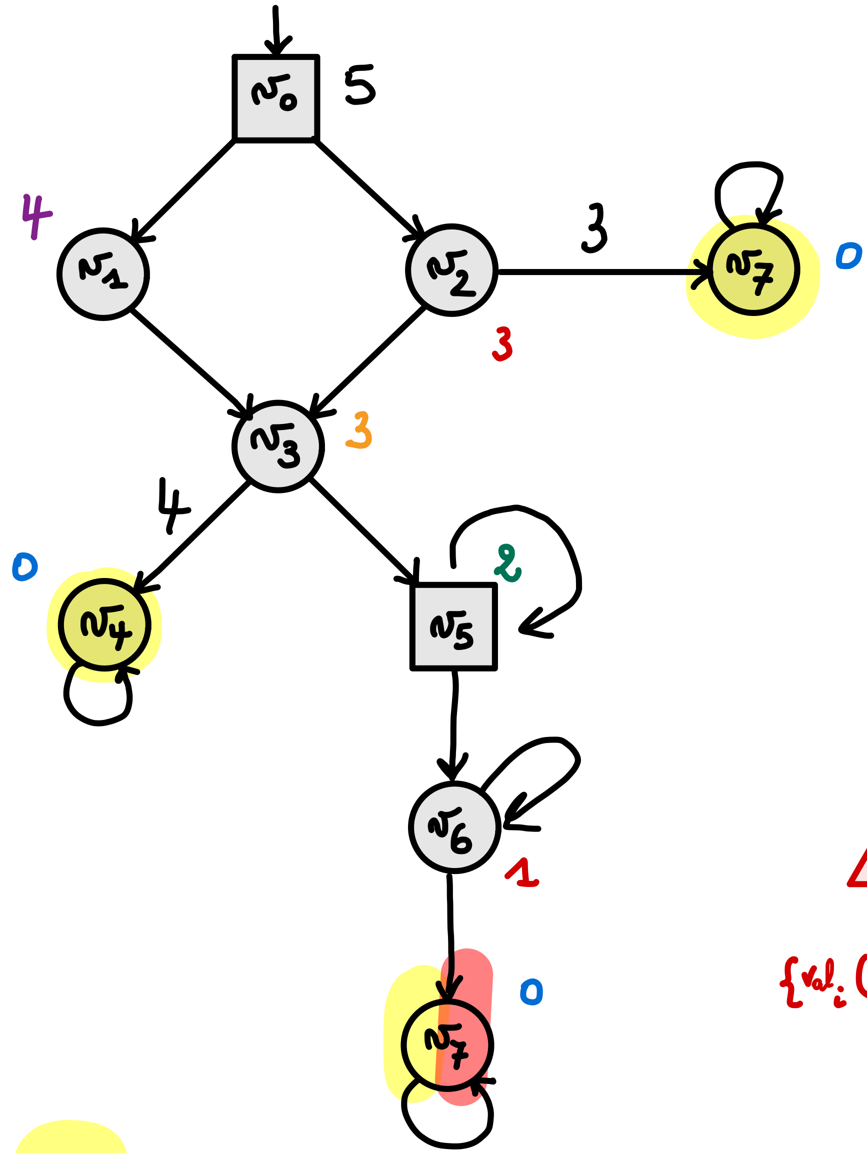

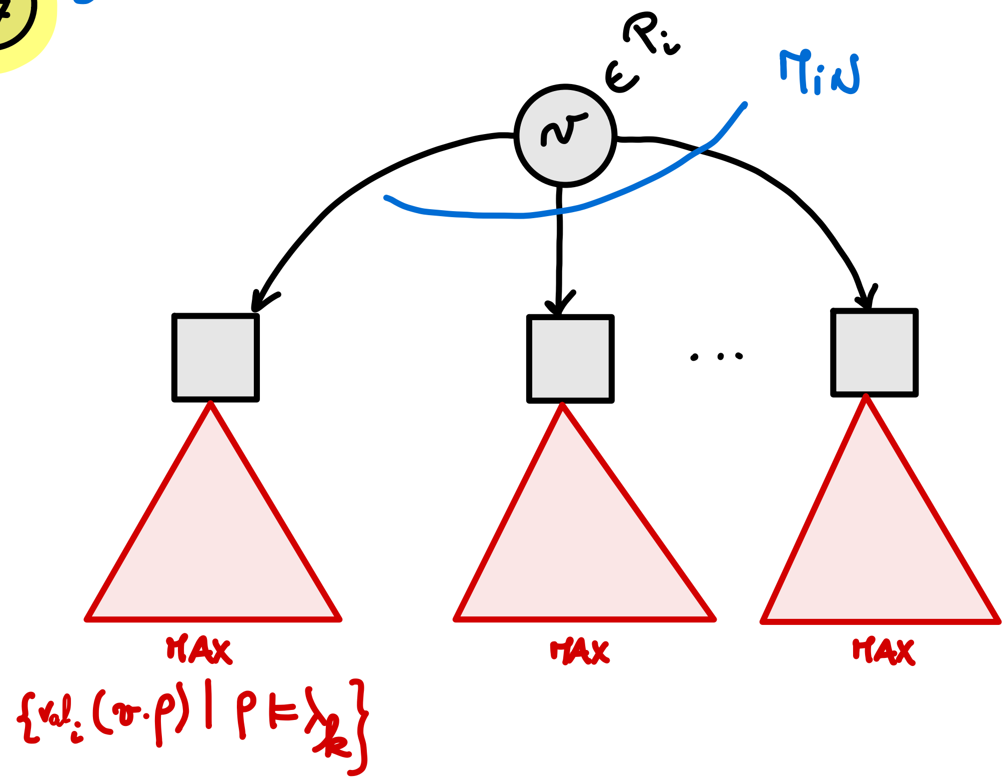

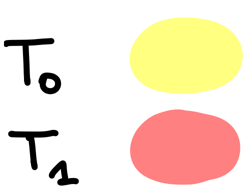

fixed point reached

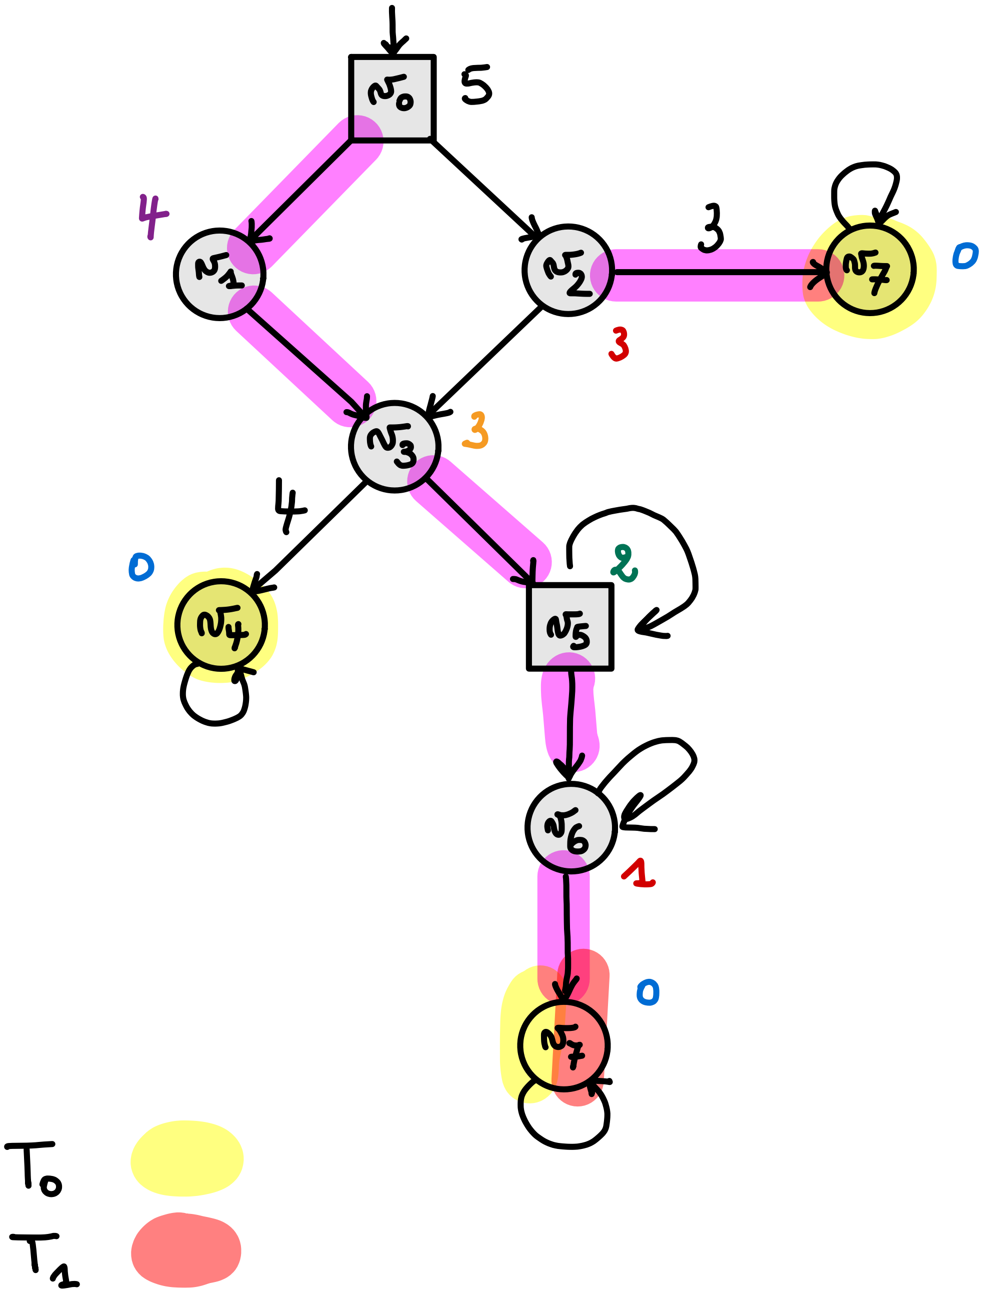

is the λ^*-consistent outcome

fixed point reached

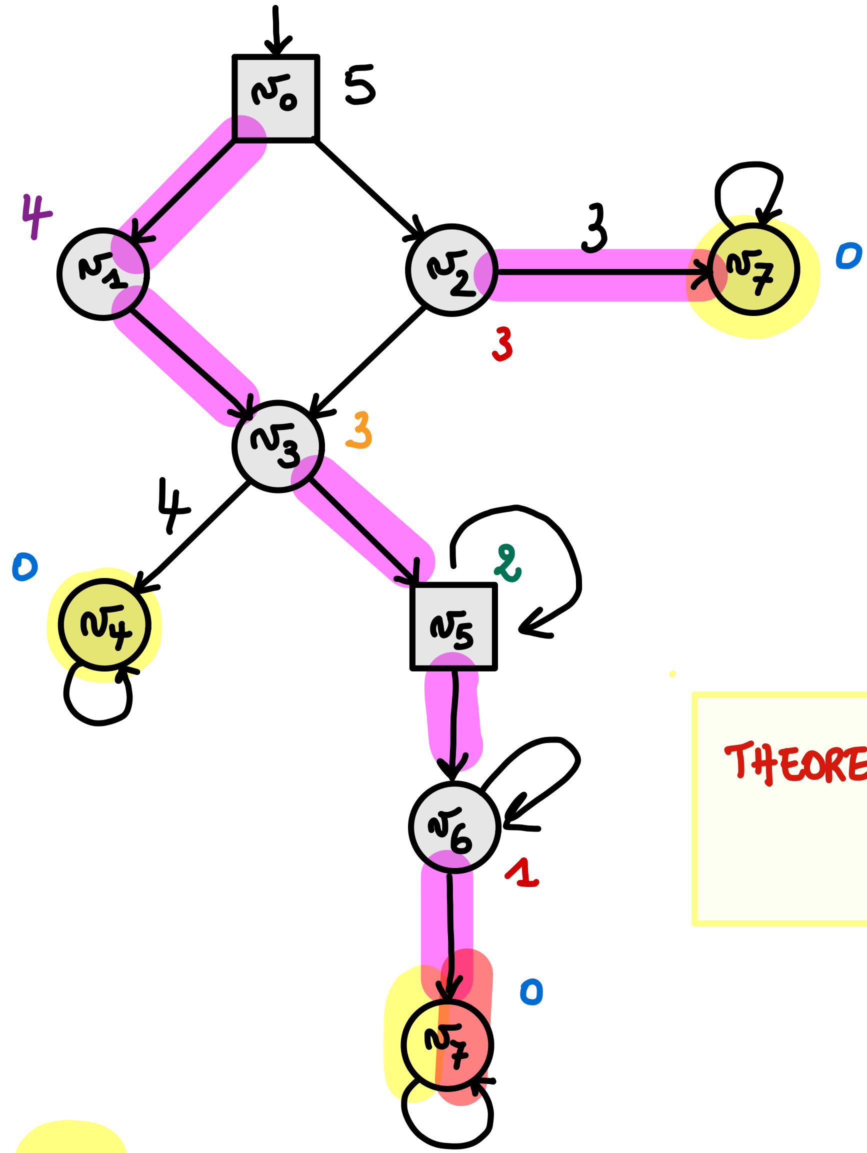

THEOREM: ρ is the outcome of a SPE iff ρ is λ^*-consistent.

T_0

T_1

fixed point reached

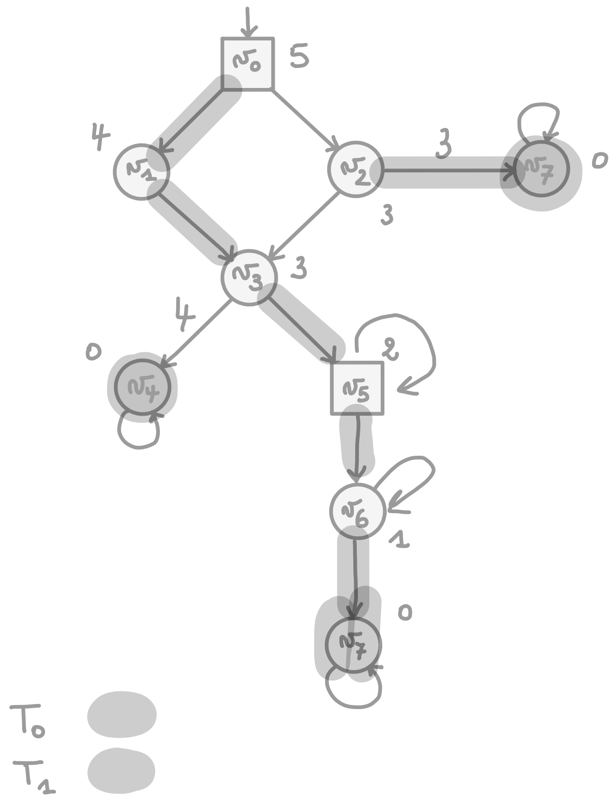

Termination

$$\lambda : V \to \mathbb{N} \cup \{+\infty\}$$

$$\lambda \leqslant \lambda' \text{ if } \forall v \in V : \lambda(v) \leq \lambda'(v)$$

↳ well quasi order

Update : $\leqslant$-monotone

Better complexity (Pspace) through bounds on values (exponential).

$$O\left(|V|^{(|V|+3)(|\Pi|+2)}\right)$$

Theorems

NE - Quantitative Reachability

THEOREM. NE always exist in reachability games

THEOREM. Constrained existence for NE is NP complete.

SPE - Quantitative Reachability

THEOREM. SPE always exist in reachability games

THEOREM. Constrained existence for SPE is PSPACE-C.

→ fixed point is computed on extended graph

exponentially large

$(v, \mathcal{P})$

Players that have already seen their objective (no rationality assumed)

→ bounds on values: $\Theta\left(|V|^{(|V|+3)(|\Pi|+2)}\right)$

Open questions

1. Better bounds on values? $\rightarrow$ we have no examples with $c \geqslant |V| . |\mathfrak{S}|$

Open questions

1. Better bounds on values ? → we have no examples with $c \geq |V| \cdot |\mathcal{S}|$
2. Is the problem FPT in the number of players ?

Open questions

1. Better bounds on values? → we have no examples with $c \geqslant |V| . |\mathcal{S}|$
2. Is the problem FPT in the number of players?
3. What about mean-payoff?

   - value approach extends readily to NE, not to SPE

Open questions

1. Better bounds on values? → we have no examples with $c \geq |V| . |\mathcal{P}|$
2. Is the problem FPT in the number of players?
3. What about mean-payoff?
   - value approach extends readily to NE, not to SPE
   - SPE may not exist:

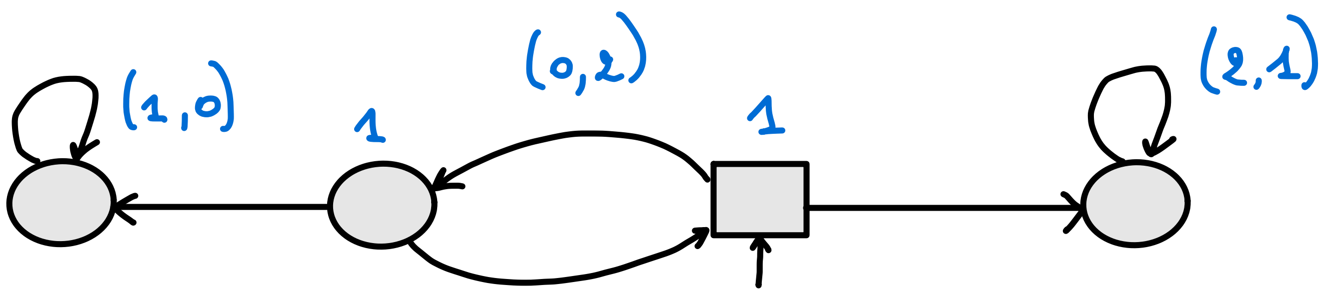

Open questions

1. Better bounds on values? → we have no examples with $c \geq |V|.|\mathcal{P}|$
2. Is the problem FPT in the number of players?
3. What about mean-payoff?
   - value approach extends readily to NE, not to SPE
   - SPE may not exist:

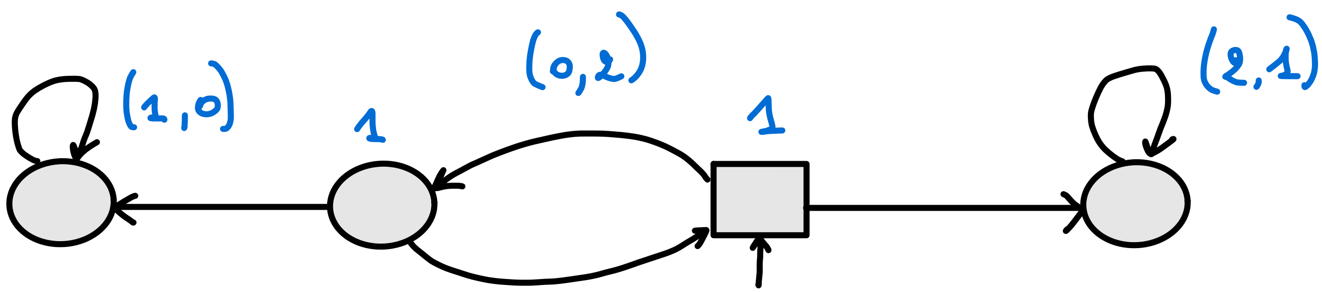

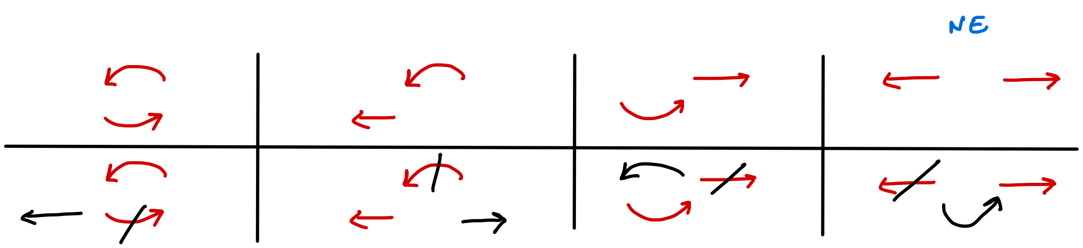

∃-problem is open.