Playing (Almost-)Optimally in Concurrent Büchi and co-Büchi Games

Benjamin Bordais, Patricia Bouyer, Stéphane Le Roux

Laboratoire Méthodes Formelles
Université Paris-Saclay, CNRS, ENS Paris-Saclay
France

Context

- Two-player games on graphs as a tool for formal verification (e.g. controller synthesis)
- Win/lose games: the objectives of the two players are opposite
- **Concurrent games**, as opposed to turn-based games

Concurrent games

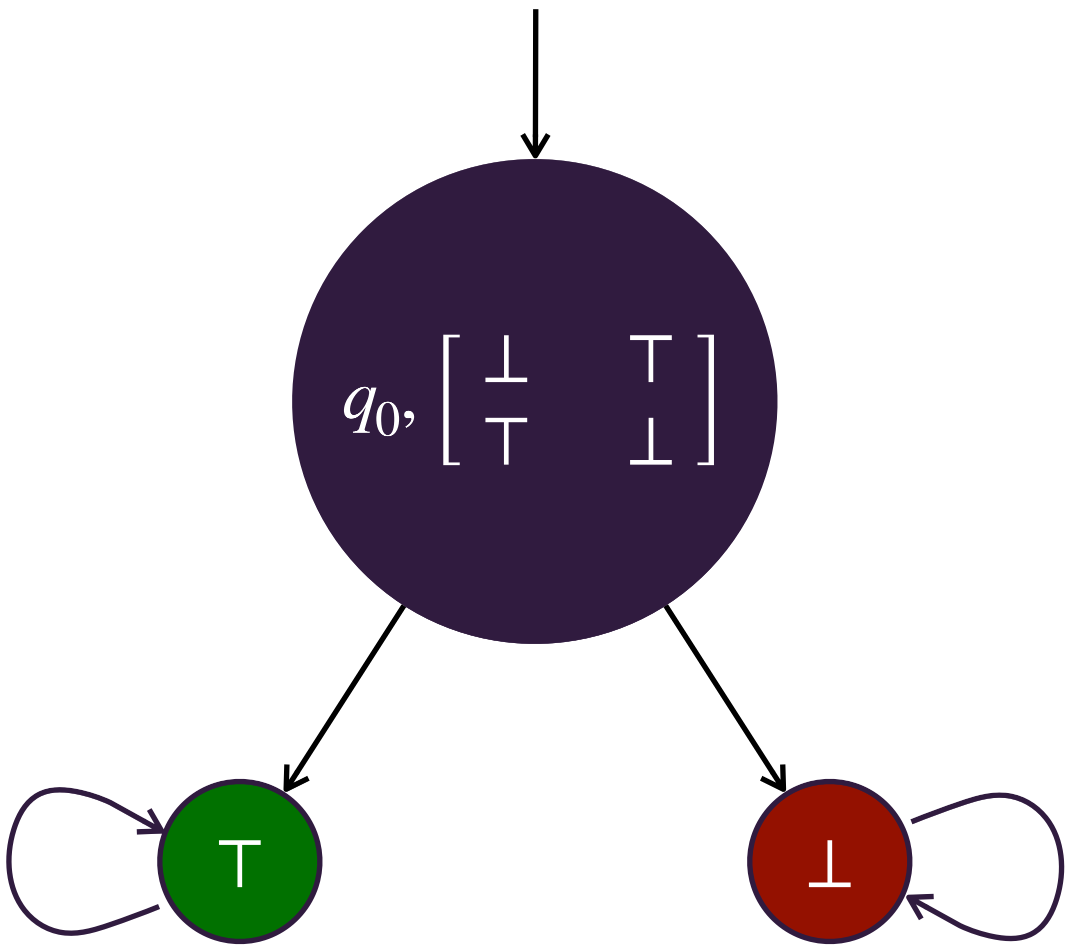

« Matching-penny game »

Concurrent games

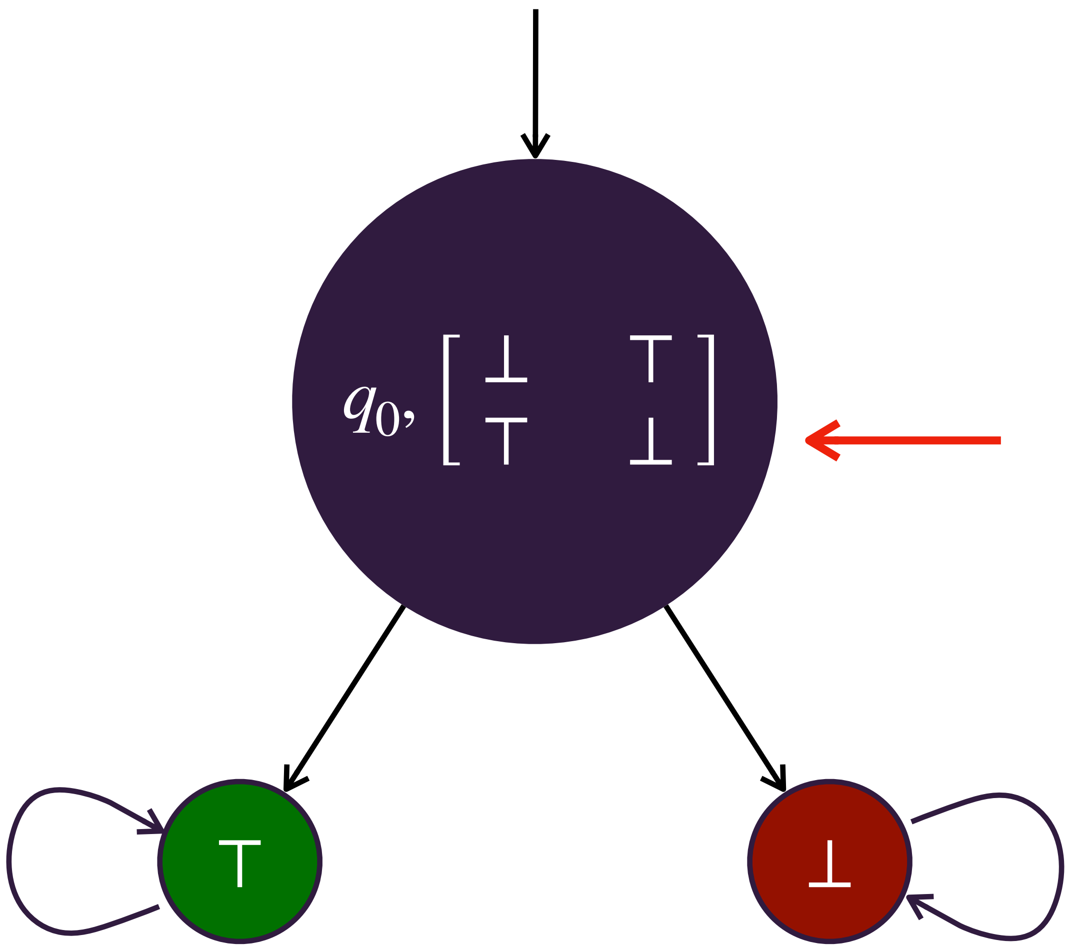

« Matching-penny game »

- Player A chooses a row

Concurrent games

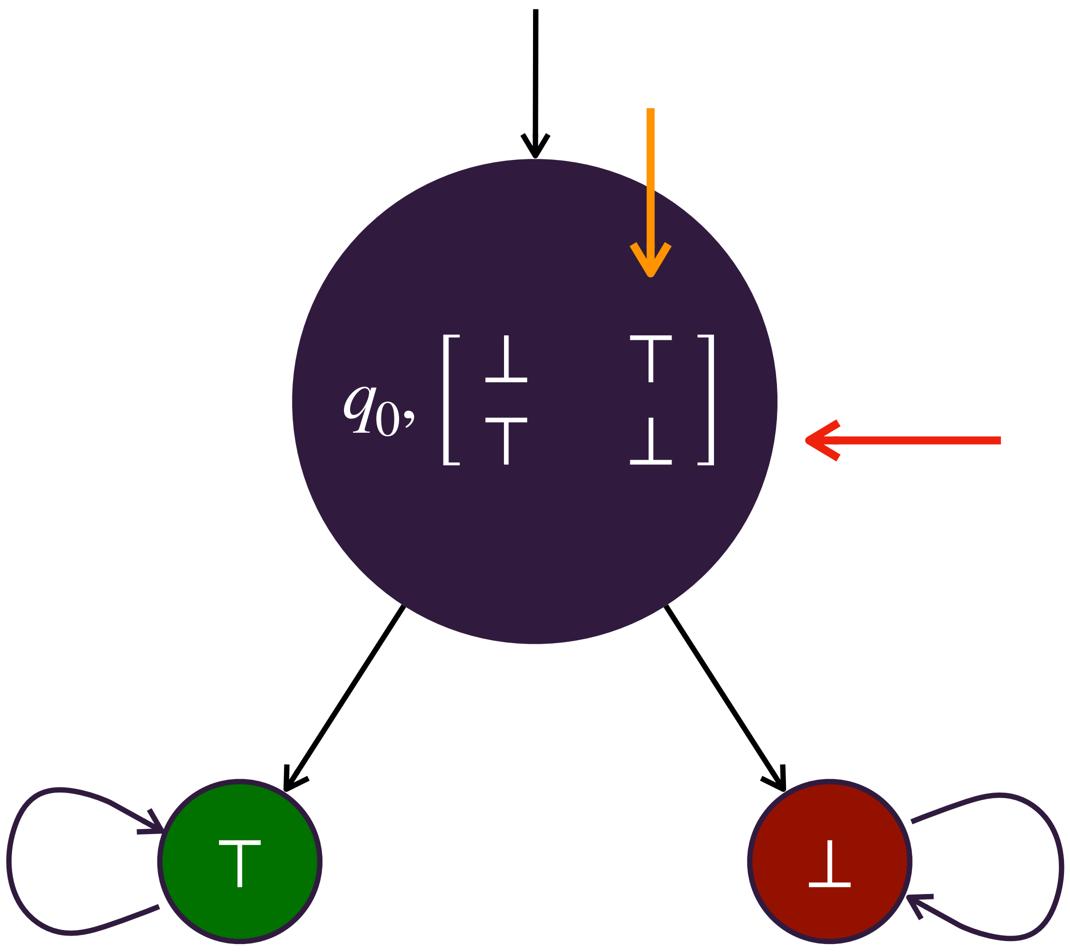

- Player A chooses a row
- Player B chooses a column

« Matching-penny game »

Concurrent games

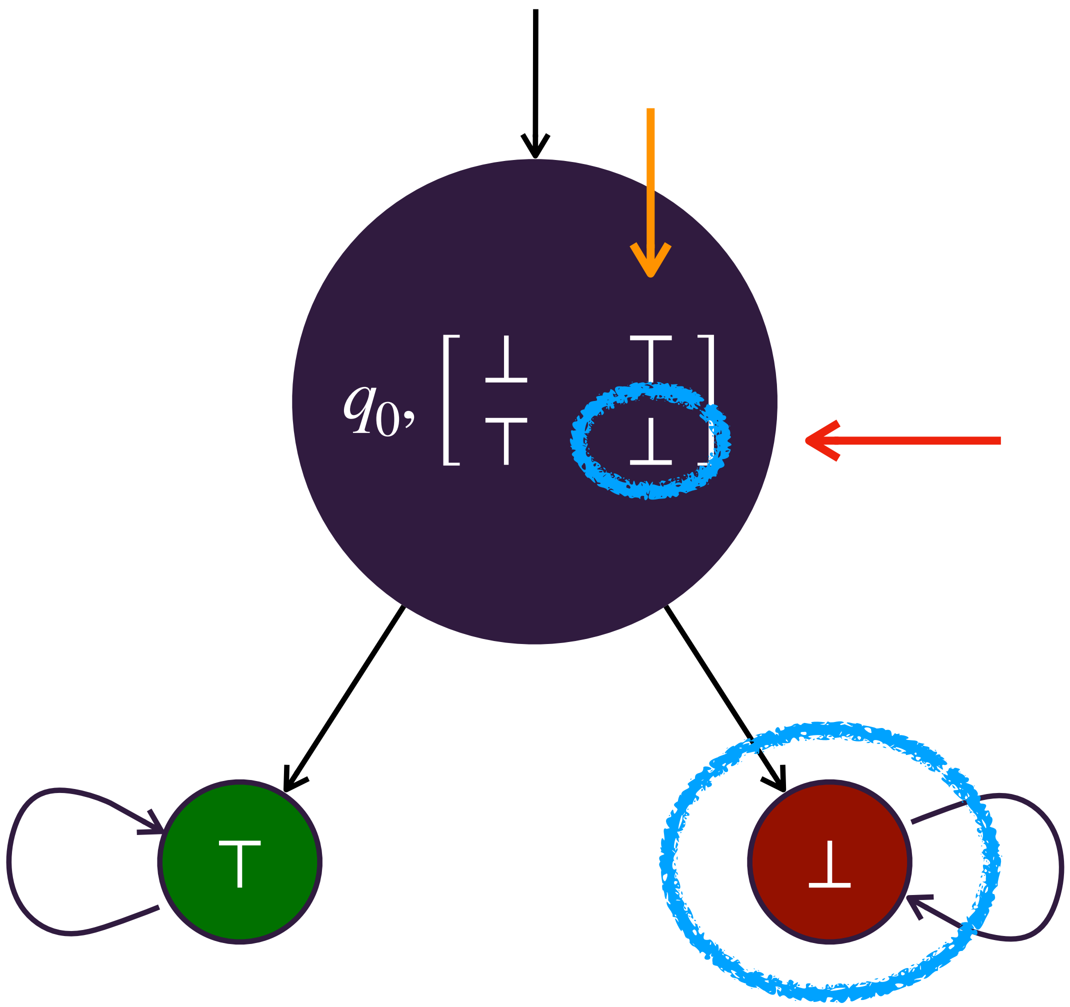

« Matching-penny game »

- Player A chooses a row
- Player B chooses a column
- The game proceeds to the corresponding state

Concurrent games

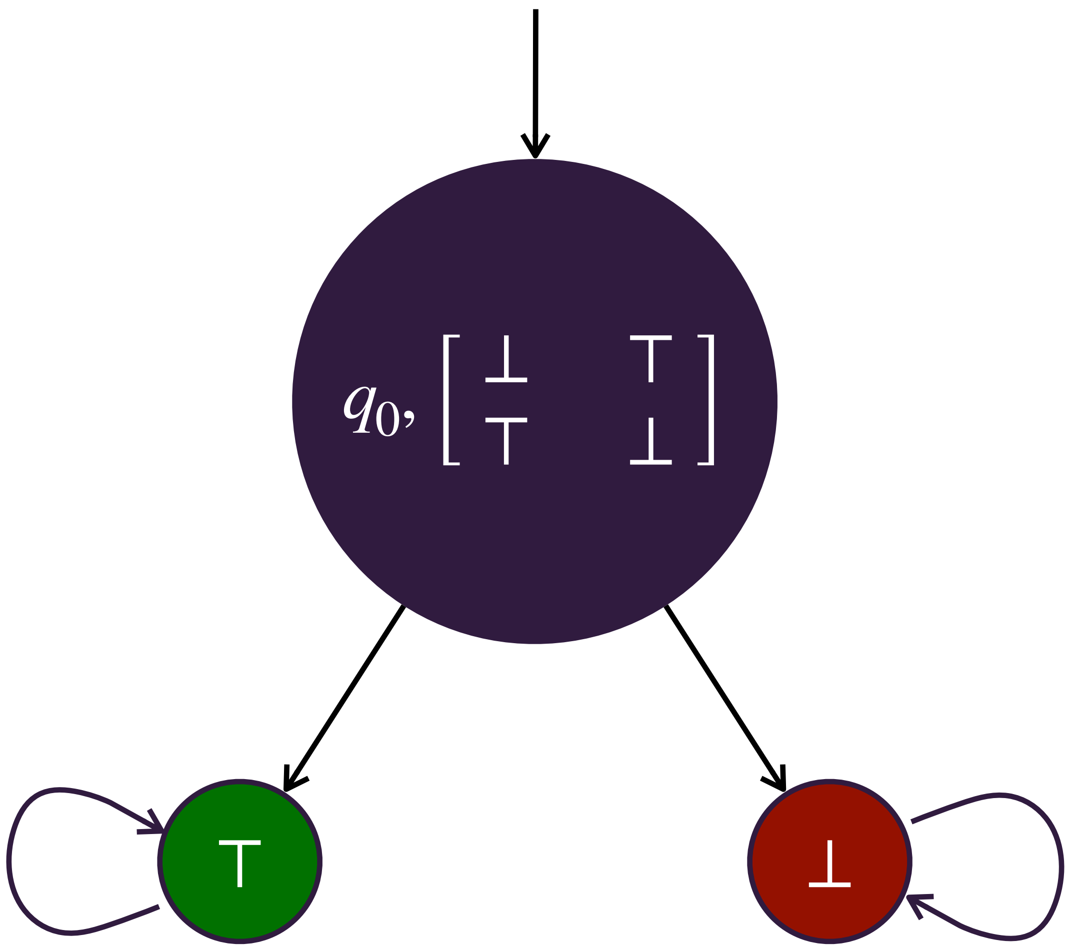

« Matching-penny game »

- Player A chooses a row
- Player B chooses a column
- The game proceeds to the corresponding state
- Strategy for player A: $\sigma_A : Q^+ \to$ row

Concurrent games

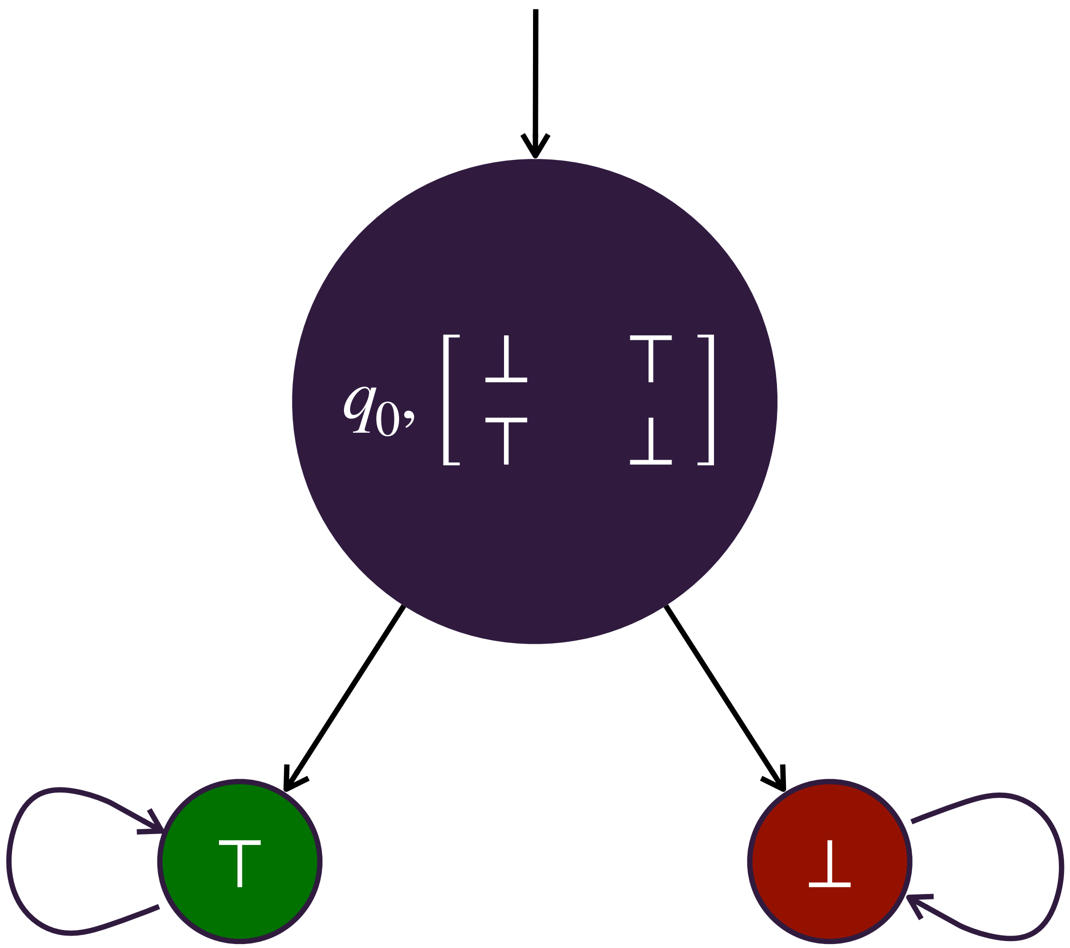

« Matching-penny game »

- Player A chooses a row
- Player B chooses a column
- The game proceeds to the corresponding state
- Strategy for player A: $\sigma_A : Q^+ \to$ row
- Outcome of σ_A: infinite path compatible with σ_A

Concurrent games

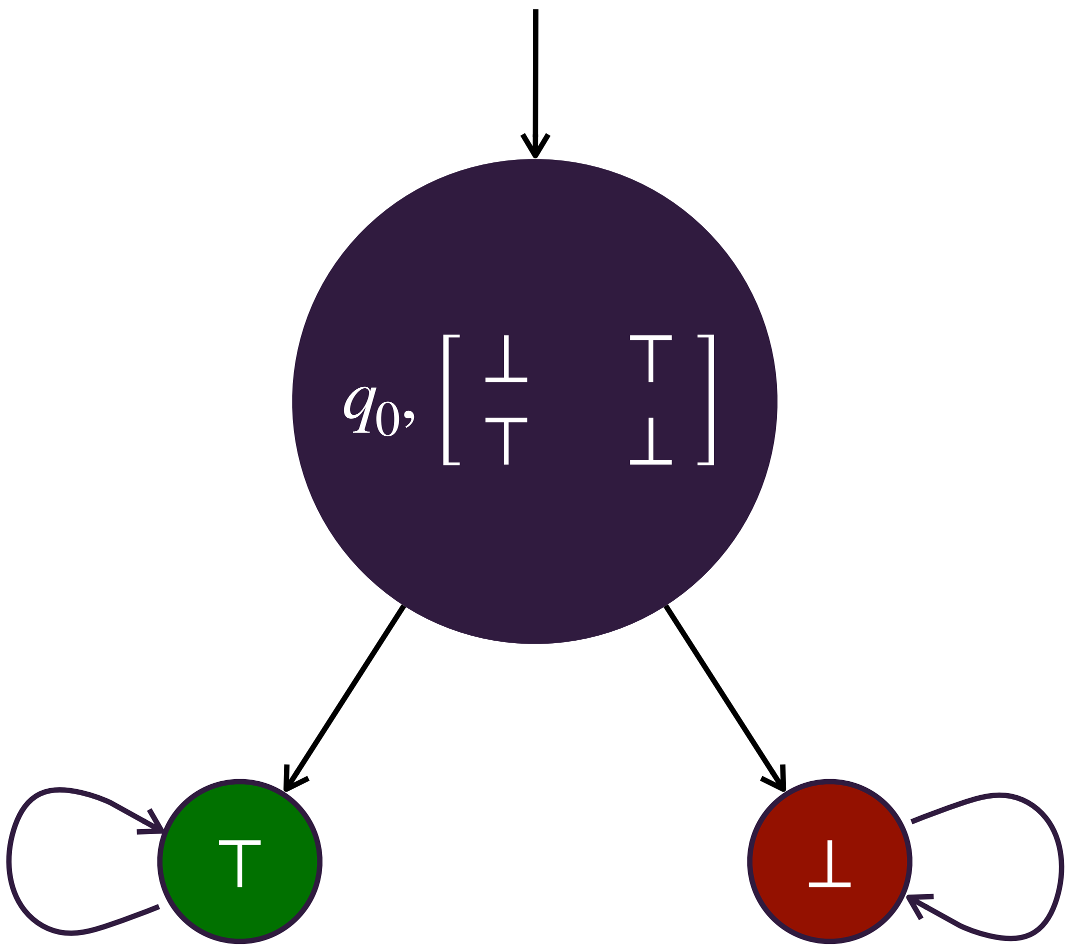

« Matching-penny game »

- Player A chooses a row
- Player B chooses a column
- The game proceeds to the corresponding state
- Strategy for player A: $\sigma_A : Q^+ \to$ row
- Outcome of σ_A: infinite path compatible with σ_A
- Objective for Player A: $W \subseteq Q^\omega$

Concurrent games

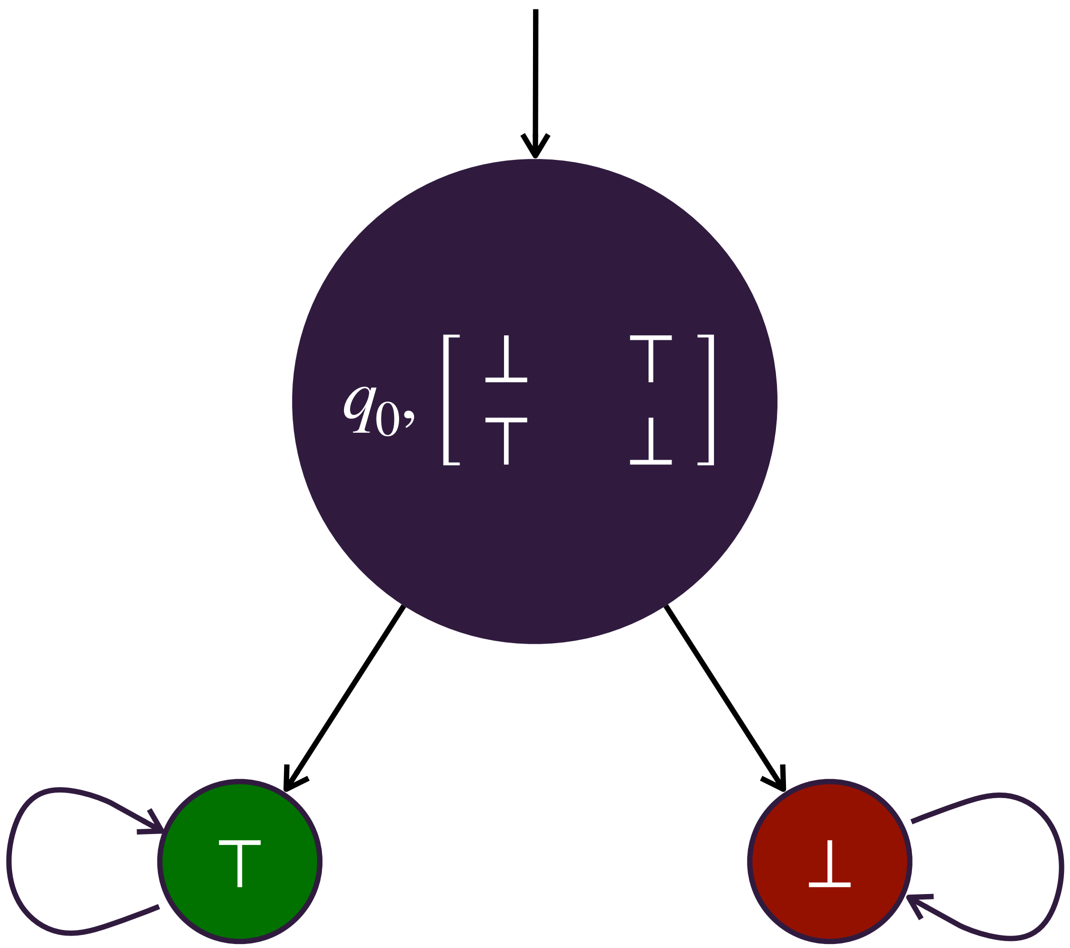

« Matching-penny game »

- Player A chooses a row
- Player B chooses a column
- The game proceeds to the corresponding state
- Strategy for player A: $\sigma_A : Q^+ \to$ row
- Outcome of σ_A: infinite path compatible with σ_A
- Objective for Player A: $W \subseteq Q^\omega$
- Objective for Player B: W^c

Concurrent games

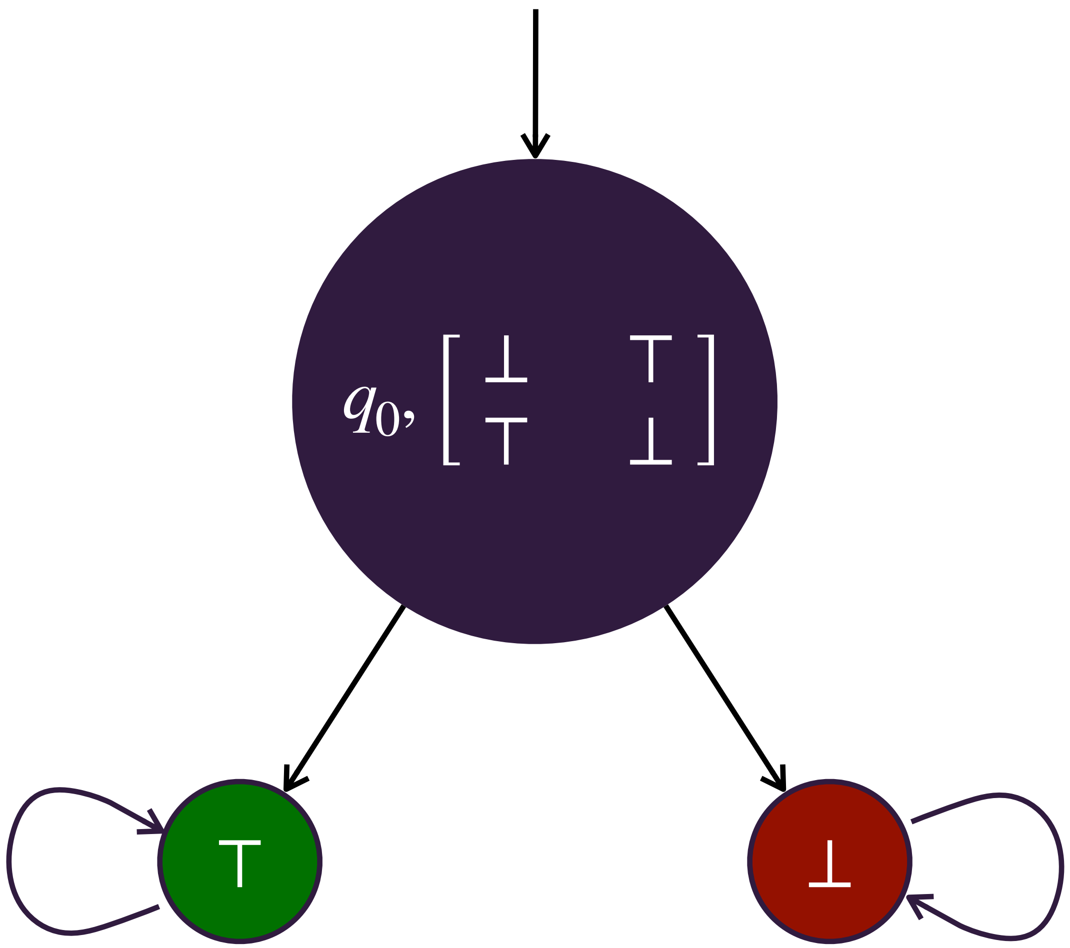

« Matching-penny game »

- Player A chooses a row
- Player B chooses a column
- The game proceeds to the corresponding state
- Strategy for player A: $\sigma_A : Q^+ \to$ row
- Outcome of σ_A: infinite path compatible with σ_A
- Objective for Player A: $W \subseteq Q^\omega$
- Objective for Player B: W^c
- Winning strategy σ_A: all outcomes of σ_A belong to W

Concurrent games

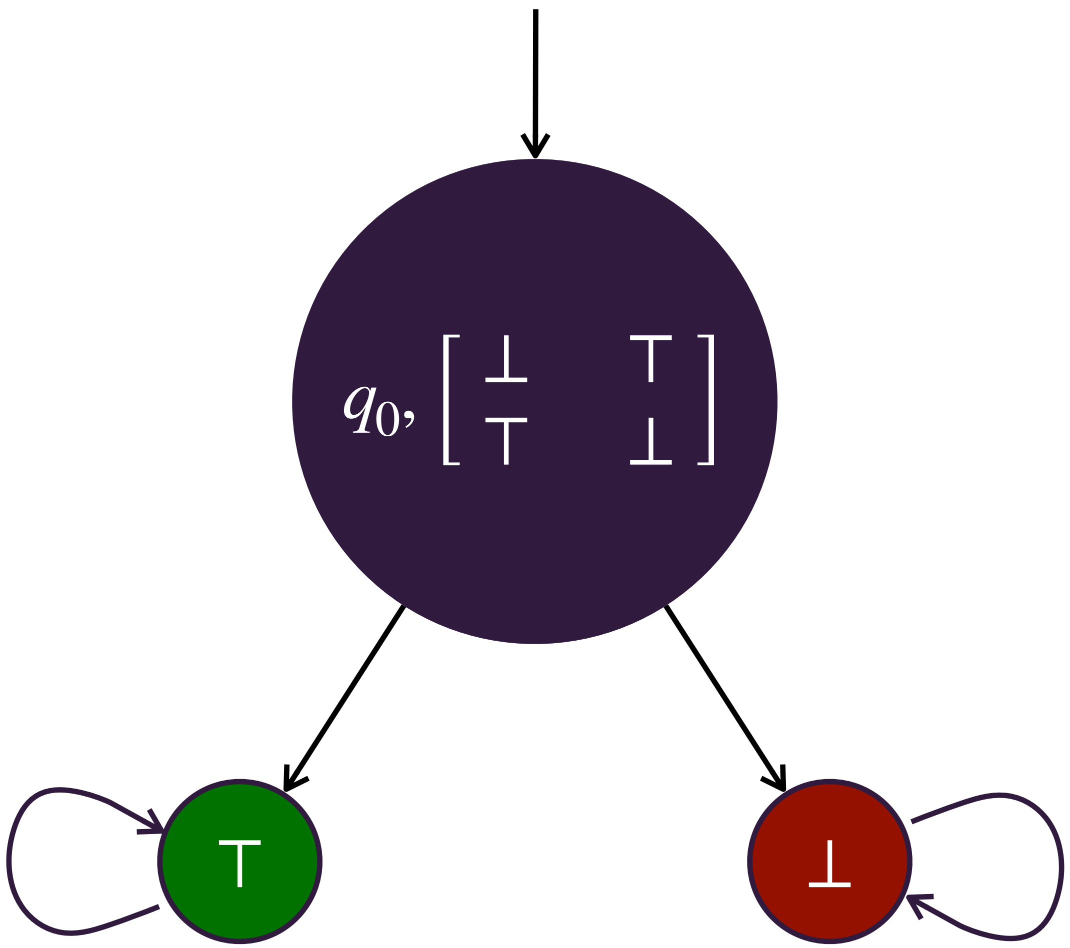

« Matching-penny game »

- Player A chooses a row
- Player B chooses a column
- The game proceeds to the corresponding state
- Strategy for player A: $\sigma_A : Q^+ \to$ row
- Outcome of σ_A: infinite path compatible with σ_A
- Objective for Player A: $W \subseteq Q^\omega$
- Objective for Player B: W^c
- Winning strategy σ_A: all outcomes of σ_A belong to W
- There is no winning strategy, for either of the players

Concurrent games

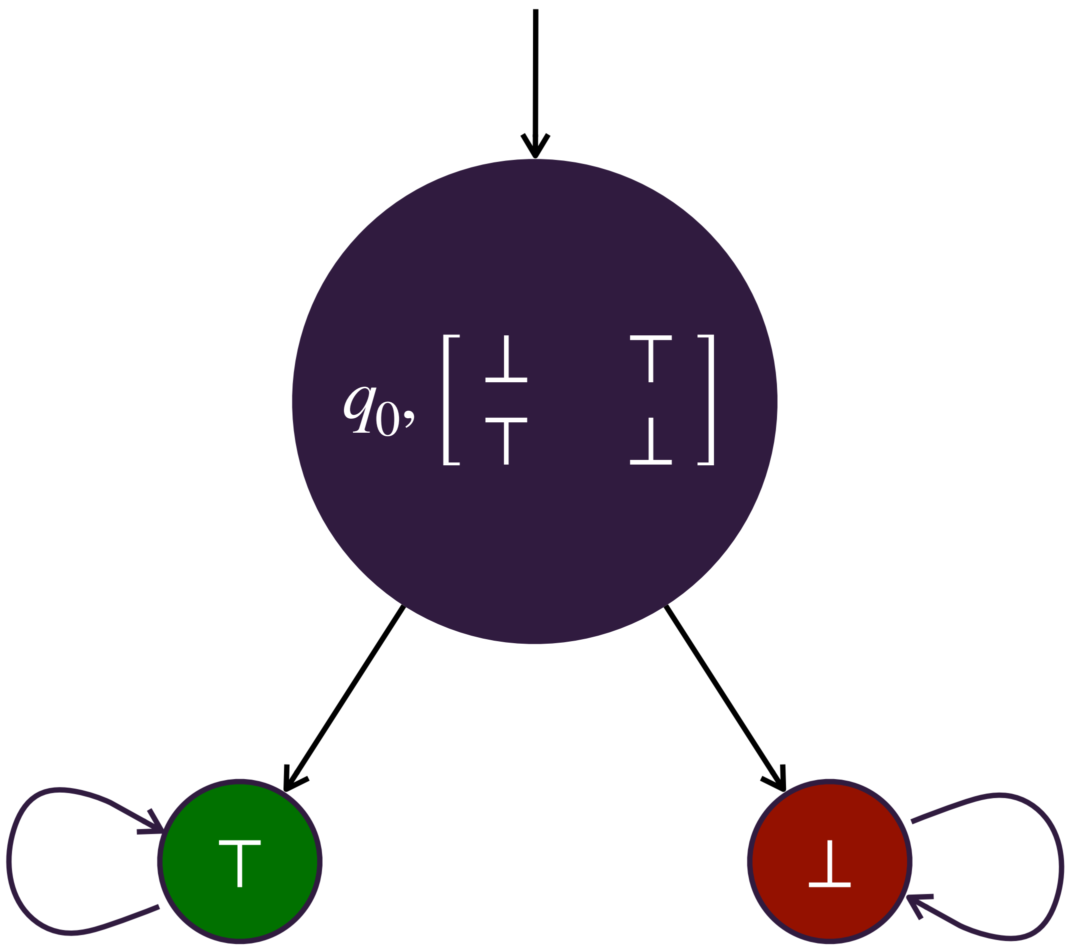

« Matching-penny game »

- Player A chooses a row
- Player B chooses a column
- The game proceeds to the corresponding state
- Strategy for player A: $\sigma_A : Q^+ \to$ row
- Outcome of σ_A: infinite path compatible with σ_A
- Objective for Player A: $W \subseteq Q^\omega$
- Objective for Player B: W^c
- Winning strategy σ_A: all outcomes of σ_A belong to W
- There is no winning strategy, for either of the players

Major difference with turn-based games

Concurrent games

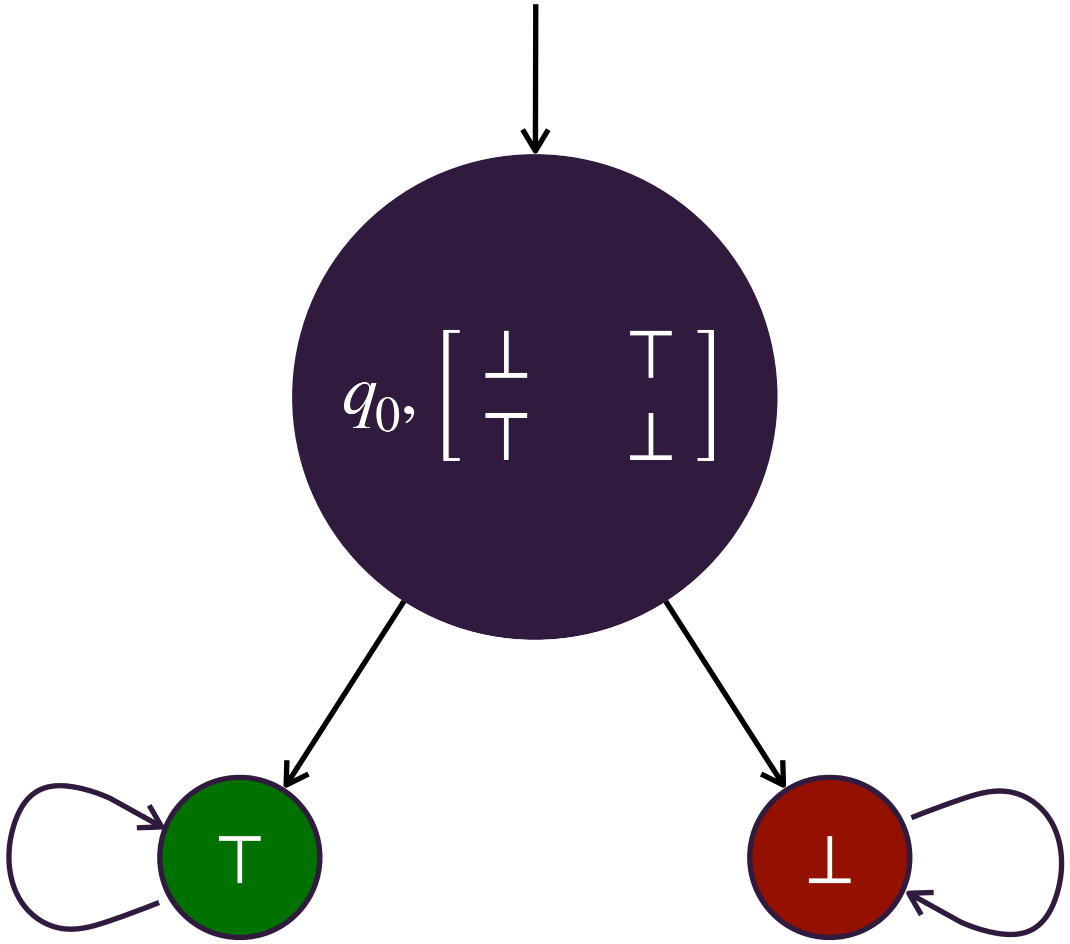

« Matching-penny game »

Concurrent games

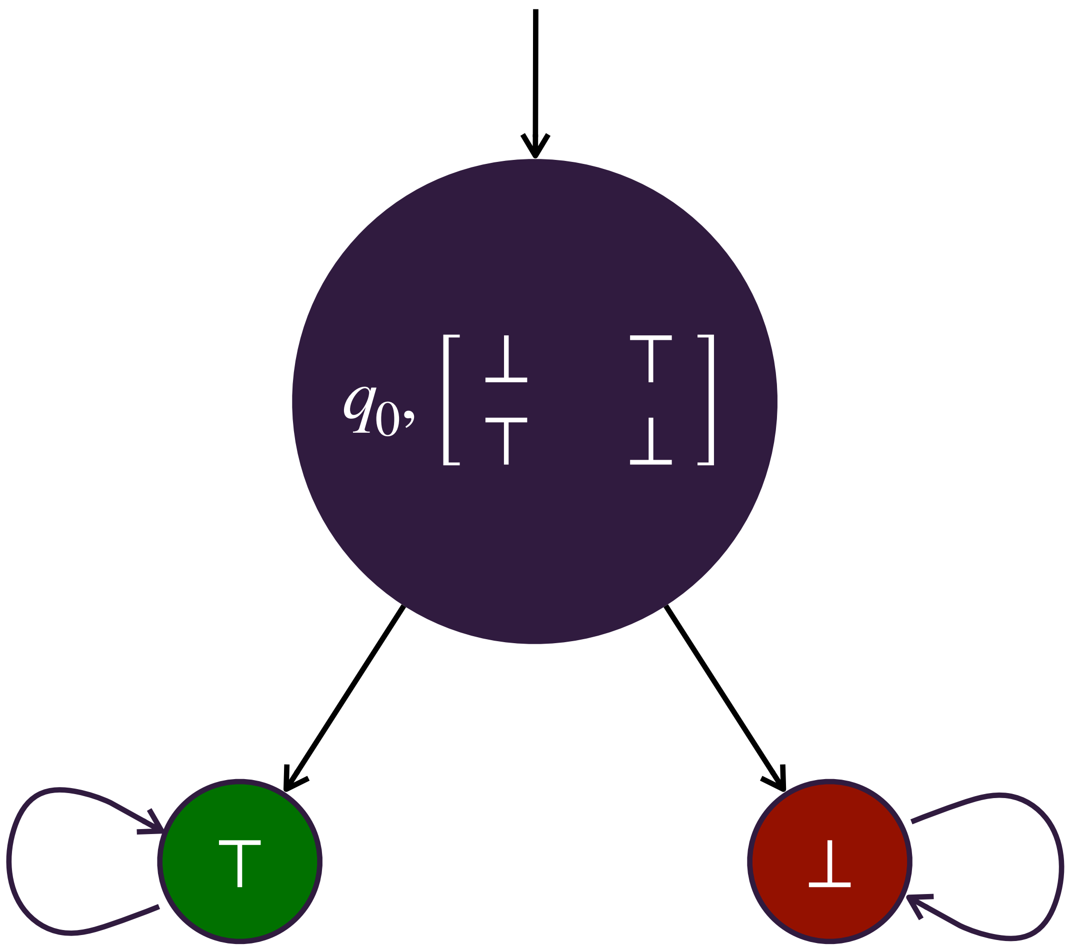

« Matching-penny game »

- Need for **randomization**!
- Randomized strategy:
  choose rows/columns according to a distribution

Concurrent games

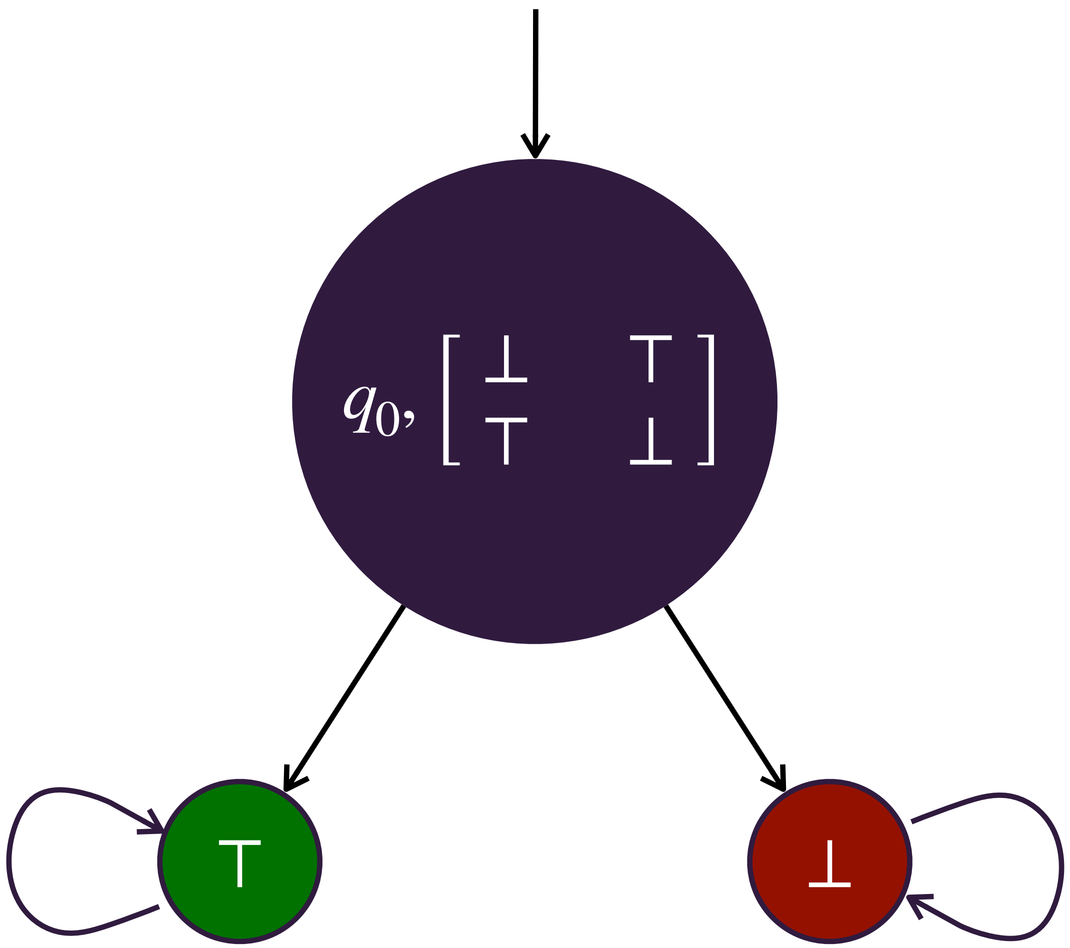

« Matching-penny game »

- Need for **randomization**!
- Randomized strategy:
  choose rows/columns according to a distribution
- Given randomized strategies σ_A and σ_B, the **payoff** (for A) is the probability $\mathbb{P}_{\sigma_A,\sigma_B}(W)$
- Optimal strategy for A:
  σ_A that maximizes $\inf_{\sigma_B} \mathbb{P}_{\sigma_A,\sigma_B}(W)$

Concurrent games

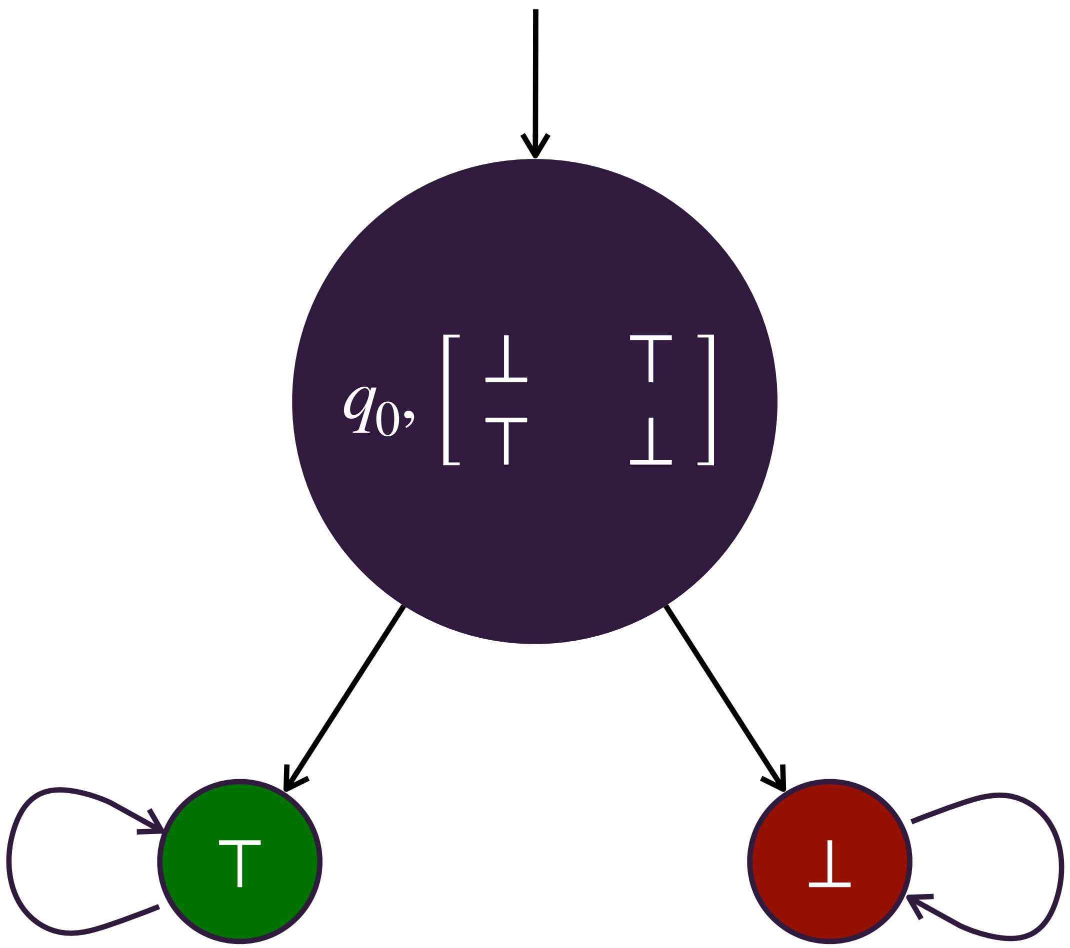

« Matching-penny game »

- Need for **randomization**!
- Randomized strategy:
  choose rows/columns according to a distribution
- Given randomized strategies σ_A and σ_B, the **payoff** (for A) is the probability $\mathbb{P}_{\sigma_A,\sigma_B}(W)$
- Optimal strategy for B
  σ_B that minimizes $\sup_{\sigma_A} \mathbb{P}_{\sigma_A,\sigma_B}(W)$

Concurrent games

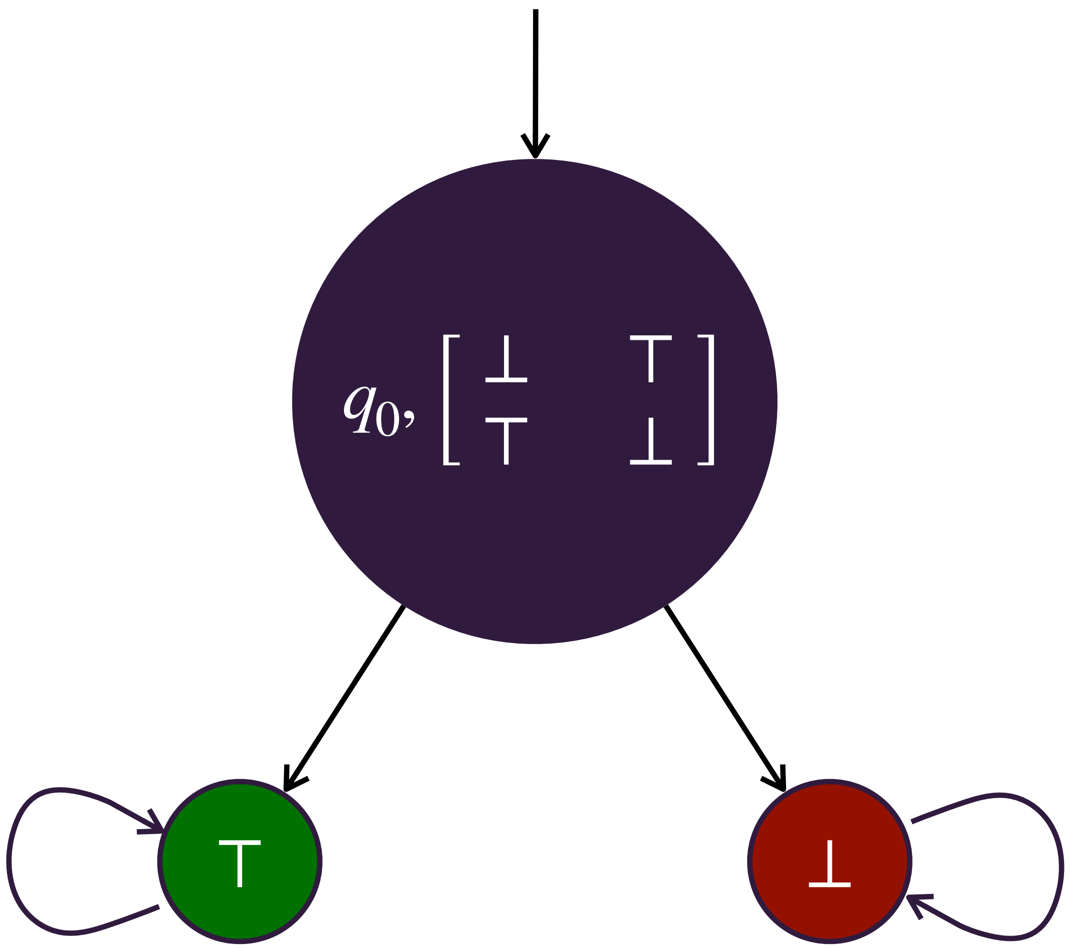

« Matching-penny game »

- Need for **randomization**!
- Randomized strategy:
  choose rows/columns according to a distribution
- Given randomized strategies σ_A and σ_B, the **payoff** (for A) is the probability $\mathbb{P}_{\sigma_A,\sigma_B}(W)$
- Optimal strategy for B
  σ_B that minimizes $\sup_{\sigma_A} \mathbb{P}_{\sigma_A,\sigma_B}(W)$
- ε-optimal strategy for A:
  σ_A that achieves $\sup_{\sigma'_A} \inf_{\sigma_B} \mathbb{P}_{\sigma'_A,\sigma_B}(W)$ up to ε

Concurrent games

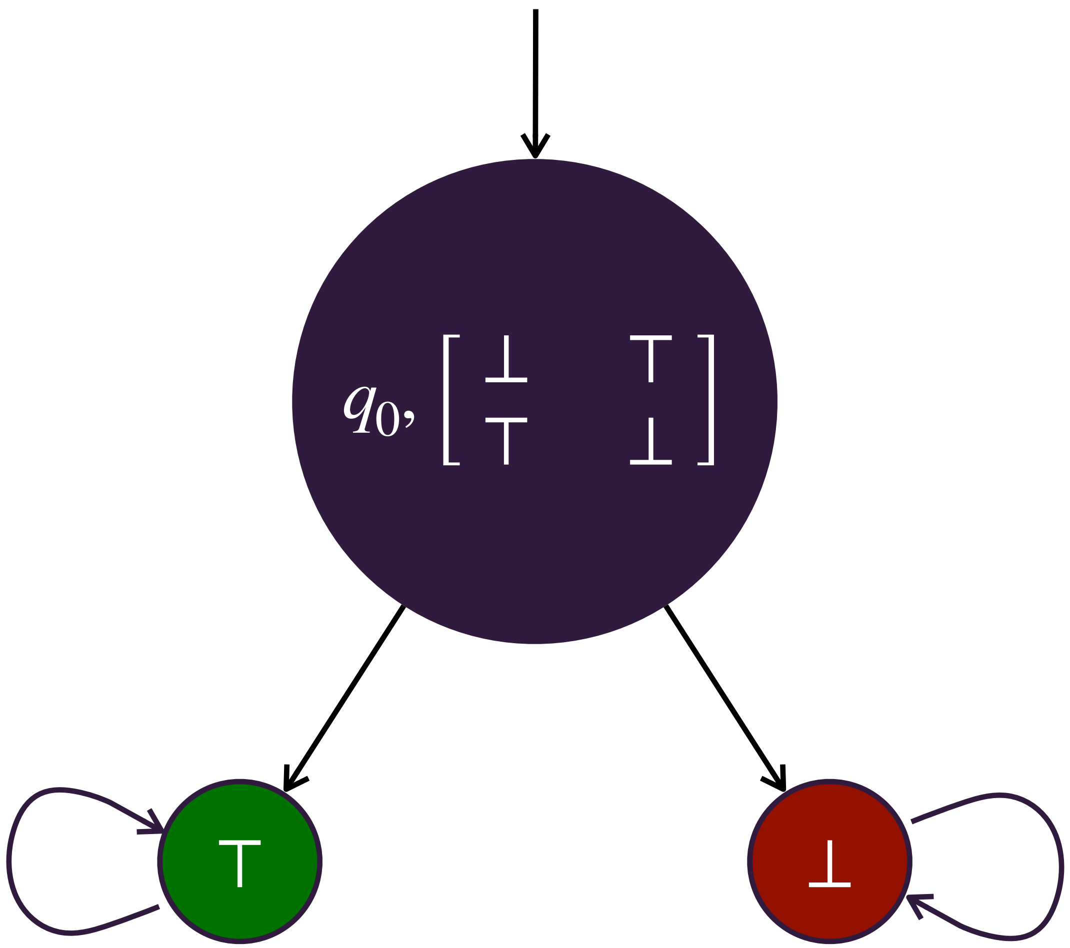

« Matching-penny game »

- Need for **randomization**!
- Randomized strategy:
  choose rows/columns according to a distribution
- Given randomized strategies σ_A and σ_B, the **payoff** (for A) is the probability $\mathbb{P}_{\sigma_A, \sigma_B}(W)$
- Optimal strategy for B
  σ_B that minimizes $\sup_{\sigma_A} \mathbb{P}_{\sigma_A, \sigma_B}(W)$
- ε-optimal strategy for A:
  σ_A that achieves $\sup_{\sigma'_A} \inf_{\sigma_B} \mathbb{P}_{\sigma'_A, \sigma_B}(W)$ up to ε
- There are optimal strategies for both players:

Concurrent games

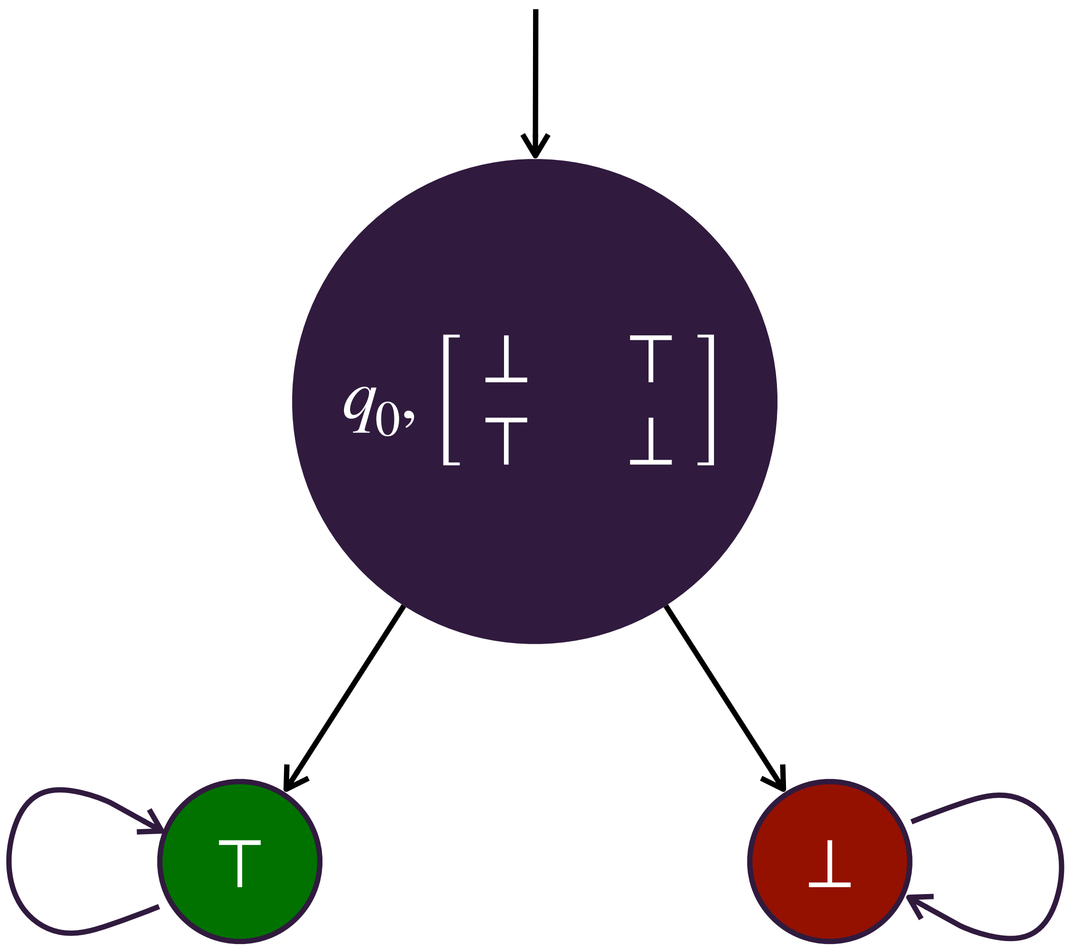

« Matching-penny game »

- Need for **randomization**!
- Randomized strategy:
  choose rows/columns according to a distribution
- Given randomized strategies σ_A and σ_B, the **payoff** (for A) is the probability $\mathbb{P}_{\sigma_A,\sigma_B}(W)$
- Optimal strategy for B
  σ_B that minimizes $\sup_{\sigma_A} \mathbb{P}_{\sigma_A,\sigma_B}(W)$
- ε-optimal strategy for A:
  σ_A that achieves $\sup_{\sigma'_A} \inf_{\sigma_B} \mathbb{P}_{\sigma'_A,\sigma_B}(W)$ up to ε
- There are optimal strategies for both players:
  - Player A: chooses uniformly at random a row

Concurrent games

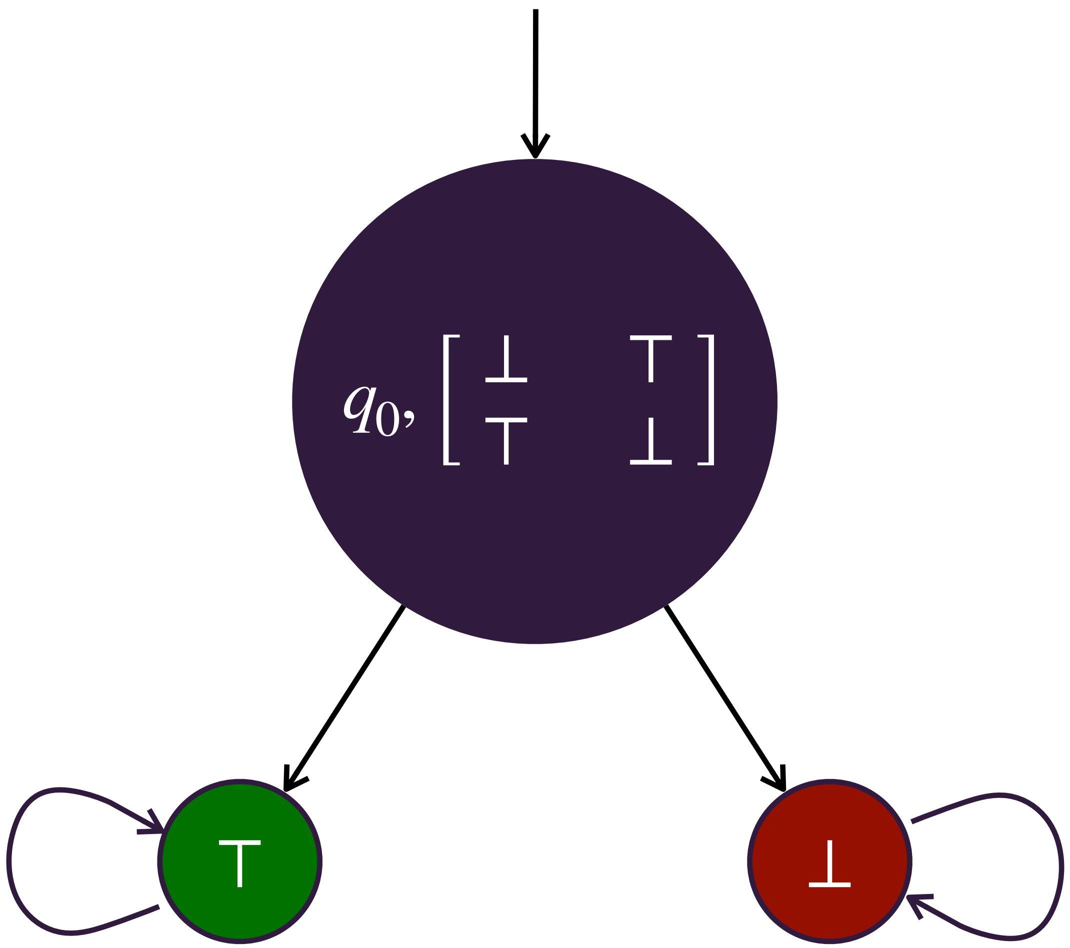

« Matching-penny game »

- Need for **randomization**!
- Randomized strategy:
  choose rows/columns according to a distribution
- Given randomized strategies σ_A and σ_B, the **payoff** (for A) is the probability $\mathbb{P}_{\sigma_A,\sigma_B}(W)$
- Optimal strategy for B
  σ_B that minimizes $\sup_{\sigma_A} \mathbb{P}_{\sigma_A,\sigma_B}(W)$
- ε-optimal strategy for A:
  σ_A that achieves $\sup_{\sigma'_A} \inf_{\sigma_B} \mathbb{P}_{\sigma'_A,\sigma_B}(W)$ up to ε
- There are optimal strategies for both players:
  - Player A: chooses uniformly at random a row
  - Player B: chooses uniformly at random a column

Concurrent games

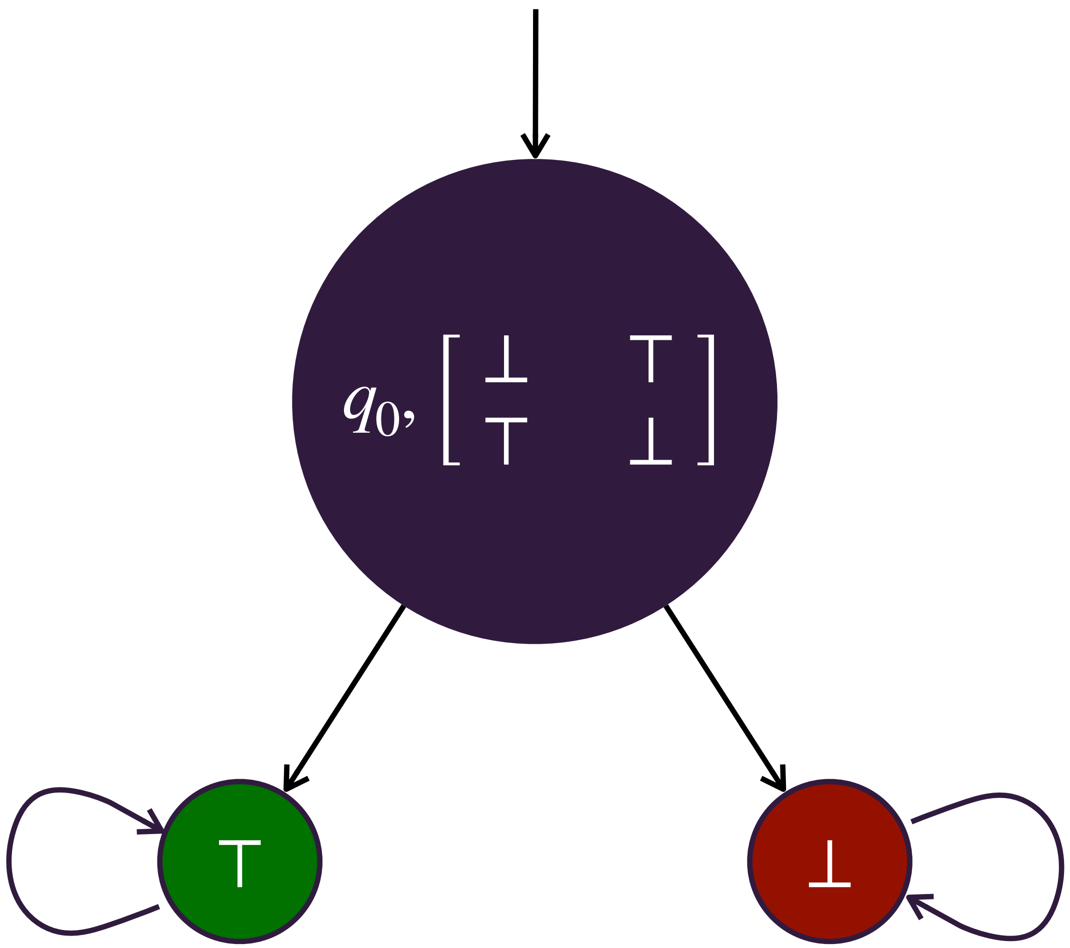

« Matching-penny game »

- Need for **randomization**!
- Randomized strategy:
  choose rows/columns according to a distribution
- Given randomized strategies σ_A and σ_B, the **payoff** (for A) is the probability $\mathbb{P}_{\sigma_A, \sigma_B}(W)$
- Optimal strategy for B
  σ_B that minimizes $\sup_{\sigma_A} \mathbb{P}_{\sigma_A, \sigma_B}(W)$
- ε-optimal strategy for A:
  σ_A that achieves $\sup_{\sigma'_A} \inf_{\sigma_B} \mathbb{P}_{\sigma'_A, \sigma_B}(W)$ up to ε
- There are optimal strategies for both players:
  - Player A: chooses uniformly at random a row
  - Player B: chooses uniformly at random a column
- Value of the game: $\frac{1}{2}$

Properties of concurrent games

Martin's determinacy theorem for Blackwell games

Concurrent games with Borel objectives have values:

$$v(q) = \sup_{\sigma_A} \inf_{\sigma_B} \mathbb{P}_{\sigma_A,\sigma_B}(W) = \inf_{\sigma_B} \sup_{\sigma_A} \mathbb{P}_{\sigma_A,\sigma_B}(W)$$

Properties of concurrent games

Martin's determinacy theorem for Blackwell games

Concurrent games with Borel objectives have values:

$$v(q) = \sup_{\sigma_{\mathsf{A}}} \inf_{\sigma_{\mathsf{B}}} \mathbb{P}_{\sigma_{\mathsf{A}},\sigma_{\mathsf{B}}}(W) = \inf_{\sigma_{\mathsf{B}}} \sup_{\sigma_{\mathsf{A}}} \mathbb{P}_{\sigma_{\mathsf{A}},\sigma_{\mathsf{B}}}(W)$$

- Optimal strategies might not exist in general (except for safety objectives)

Properties of concurrent games

Martin's determinacy theorem for Blackwell games

Concurrent games with Borel objectives have values:

$$v(q) = \sup_{\sigma_{\mathsf{A}}} \inf_{\sigma_{\mathsf{B}}} \mathbb{P}_{\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}}(W) = \inf_{\sigma_{\mathsf{B}}} \sup_{\sigma_{\mathsf{A}}} \mathbb{P}_{\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}}(W)$$

- Optimal strategies might not exist in general (except for safety objectives)
- (Infinite) Memory is sometimes needed by optimal and almost-optimal strategies
    - Parity games require infinite memory for both optimal and almost-optimal strategies

Properties of concurrent games

Martin's determinacy theorem for Blackwell games

Concurrent games with Borel objectives have values:

$$v(q) = \sup_{\sigma_{\mathsf{A}}} \inf_{\sigma_{\mathsf{B}}} \mathbb{P}_{\sigma_{\mathsf{A}},\sigma_{\mathsf{B}}}(W) = \inf_{\sigma_{\mathsf{B}}} \sup_{\sigma_{\mathsf{A}}} \mathbb{P}_{\sigma_{\mathsf{A}},\sigma_{\mathsf{B}}}(W)$$

- Optimal strategies might not exist in general (except for safety objectives)
- (Infinite) Memory is sometimes needed by optimal and almost-optimal strategies
  - Parity games require infinite memory for both optimal and almost-optimal strategies
- Note: this is **specific to concurrent games**!
  (as compared to turn-based)

An example of a Büchi game

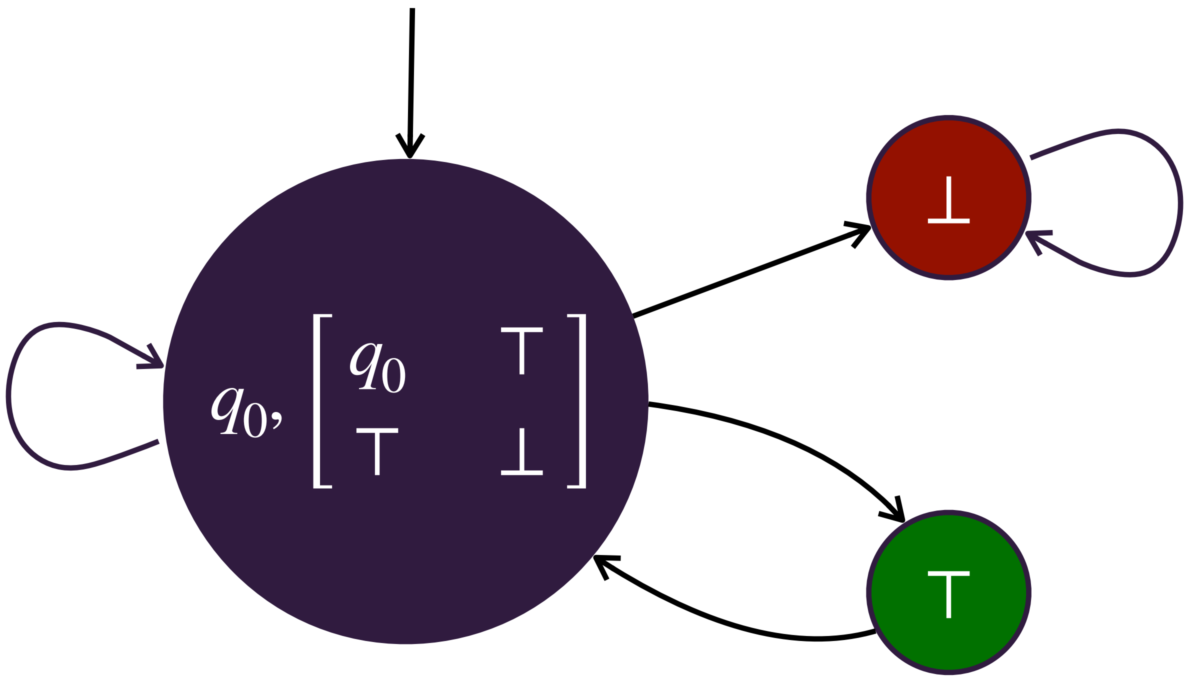

« The snowball game »

[AH00] L. De Alfaro, T. Henzinger. Concurrent omega-regular games (LICS'00)

An example of a Büchi game

- Objective is to visit ⊤ infinitely often

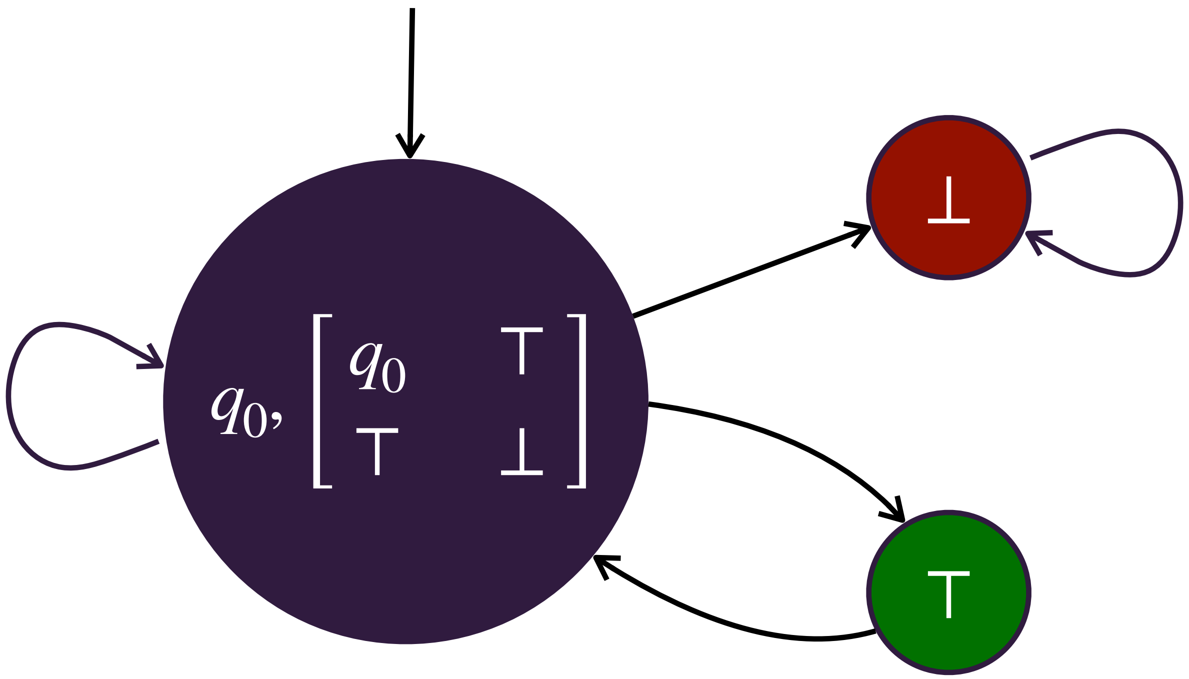

« The snowball game »

[AH00] L. De Alfaro, T. Henzinger. Concurrent omega-regular games (LICS'00)

An example of a Büchi game

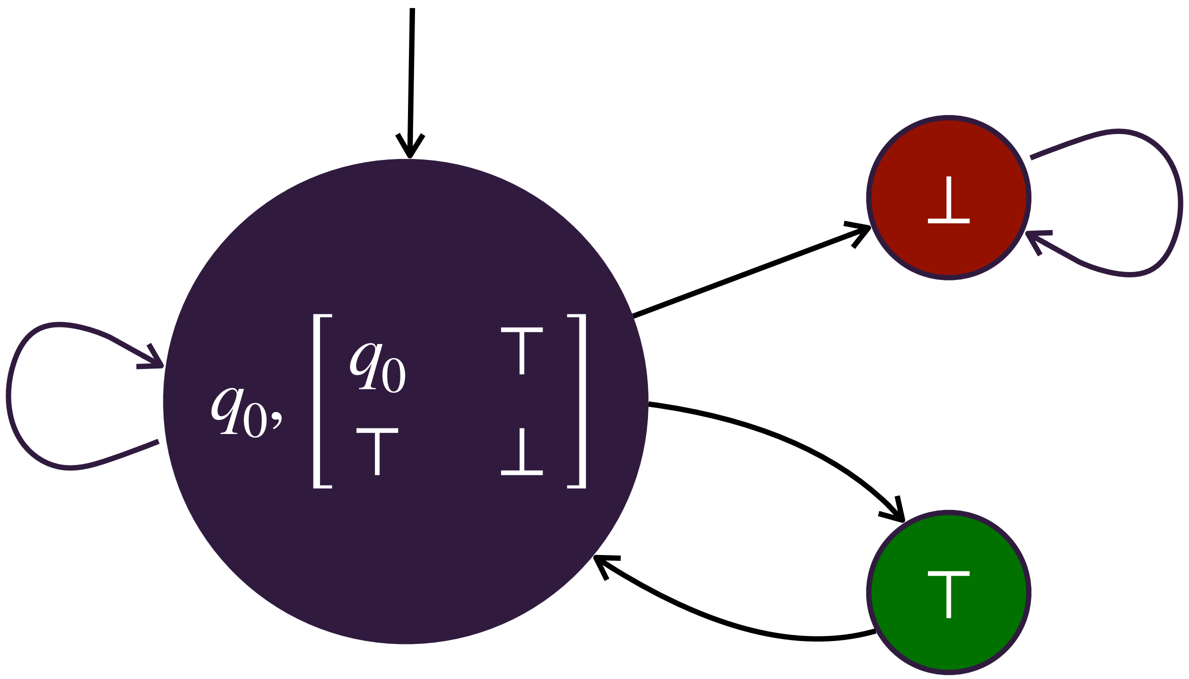

« The snowball game »

- Objective is to visit $\top$ infinitely often
- Value of the game is 1

[AH00] L. De Alfaro, T. Henzinger. Concurrent omega-regular games (LICS'00)

An example of a Büchi game

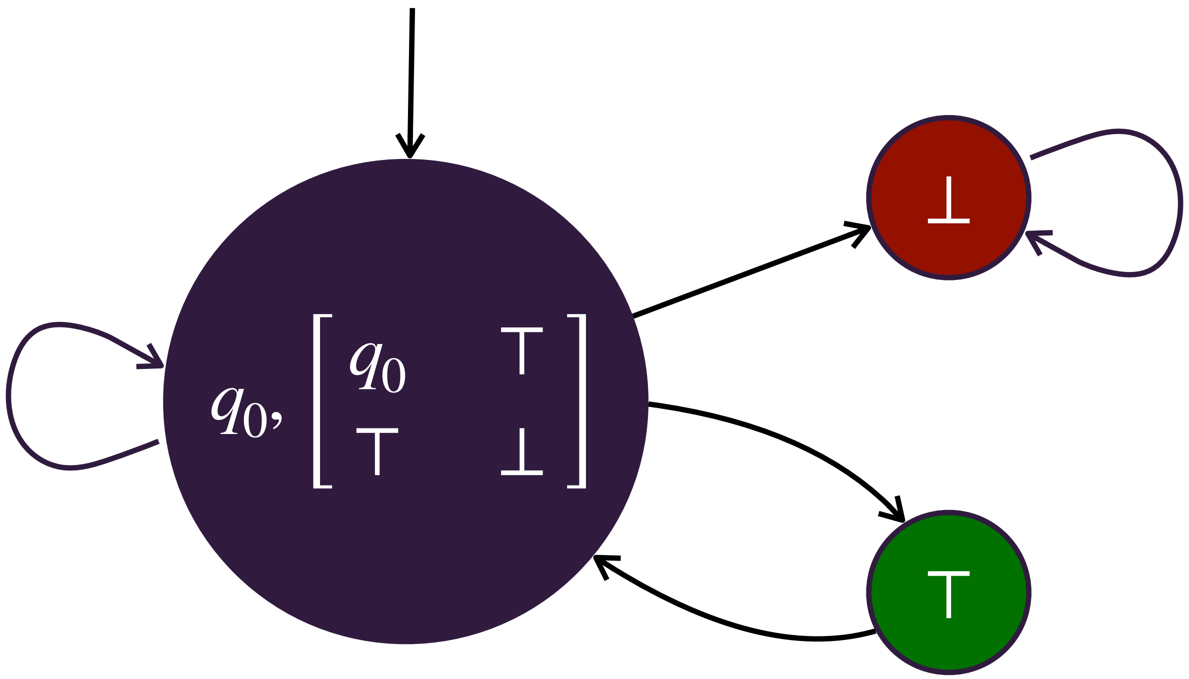

« The snowball game »

- Objective is to visit ⊤ infinitely often
- Value of the game is 1
- Player A (rows) has no optimal strat.

[AH00] L. De Alfaro, T. Henzinger. Concurrent omega-regular games (LICS'00)

An example of a Büchi game

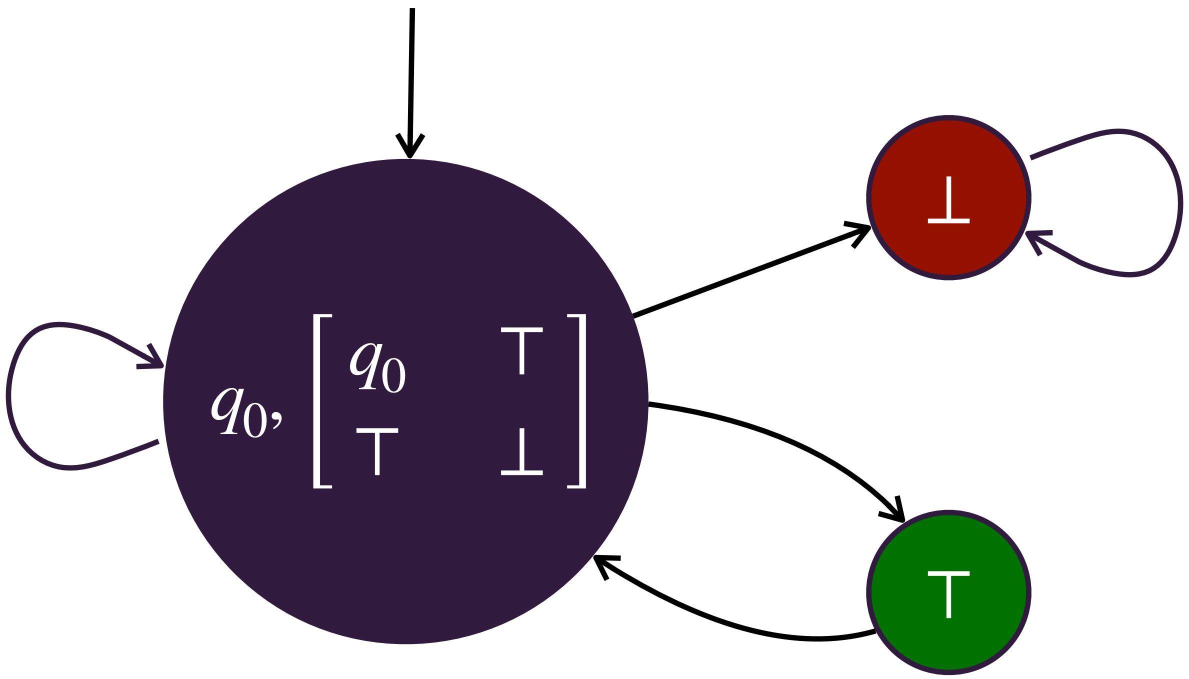

« The snowball game »

- Objective is to visit ⊤ infinitely often
- Value of the game is 1
- Player A (rows) has no optimal strat.
- Every finite-memory strat. has value 0

[AH00] L. De Alfaro, T. Henzinger. Concurrent omega-regular games (LICS'00)

An example of a Büchi game

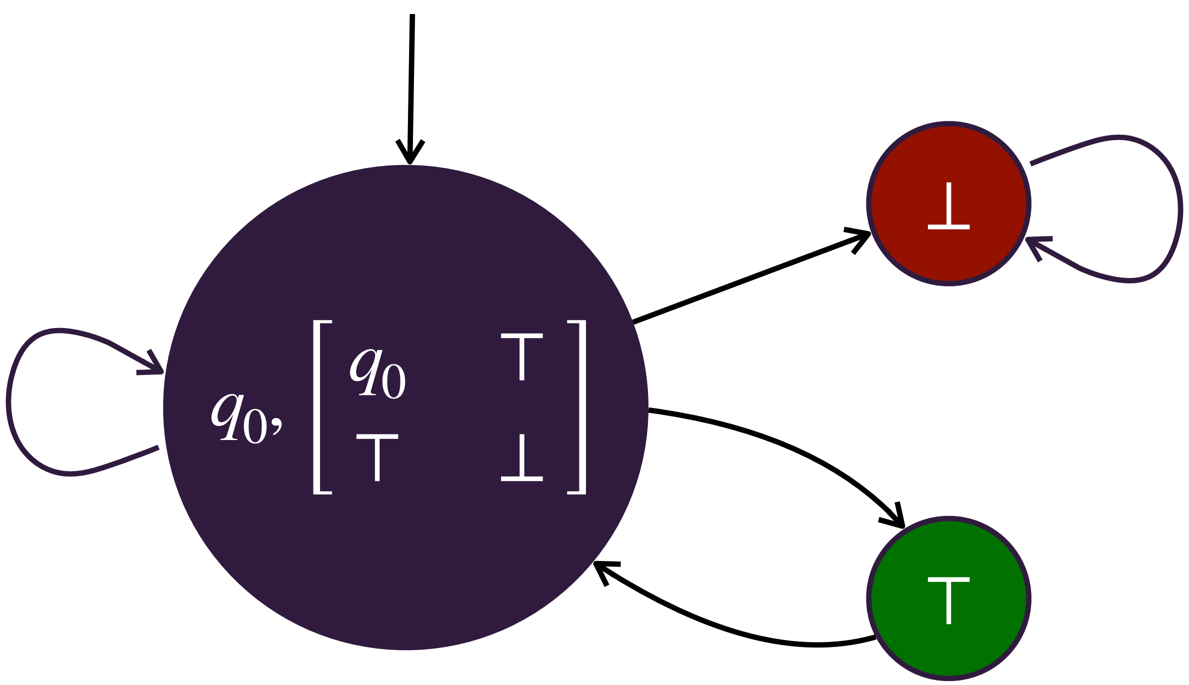

« The snowball game »

- Objective is to visit ⊤ infinitely often
- Value of the game is 1
- Player A (rows) has no optimal strat.
- Every finite-memory strat. has value 0
- Player A needs infinite memory to play ε-optimal for every $\varepsilon > 0$:
  - Play first row with probability $1 - \varepsilon_k$ and second row with probability ε_k
  - k is the number of visits to ⊤
  - $(\varepsilon_k)_k$ quickly decreases to 0

[AH00] L. De Alfaro, T. Henzinger. Concurrent omega-regular games (LICS'00)

The approach of this work

The approach of this work

- We are interested in **low memory requirements** for optimal and almost-optimal strategies in concurrent games with parity objectives in general, and more specifically Büchi and co-Büchi objectives

The approach of this work

- We are interested in **low memory requirements** for optimal and almost-optimal strategies in concurrent games with parity objectives in general, and more specifically Büchi and co-Büchi objectives

- Low memory requirement = **positional** strategies

The approach of this work

- We are interested in **low memory requirements** for optimal and almost-optimal strategies in concurrent games with parity objectives in general, and more specifically Büchi and co-Büchi objectives
- Low memory requirement = **positional** strategies
- σ_{A} is positional if it depends only on the last visited state

The approach of this work

- We are interested in **low memory requirements** for optimal and almost-optimal strategies in concurrent games with parity objectives in general, and more specifically Büchi and co-Büchi objectives

- Low memory requirement = **positional** strategies

- σ_A is positional if it depends only on the last visited state

Our approach: focus on interactions, and characterize well-behaved interactions

A tool to apprehend concurrent interactions: game forms

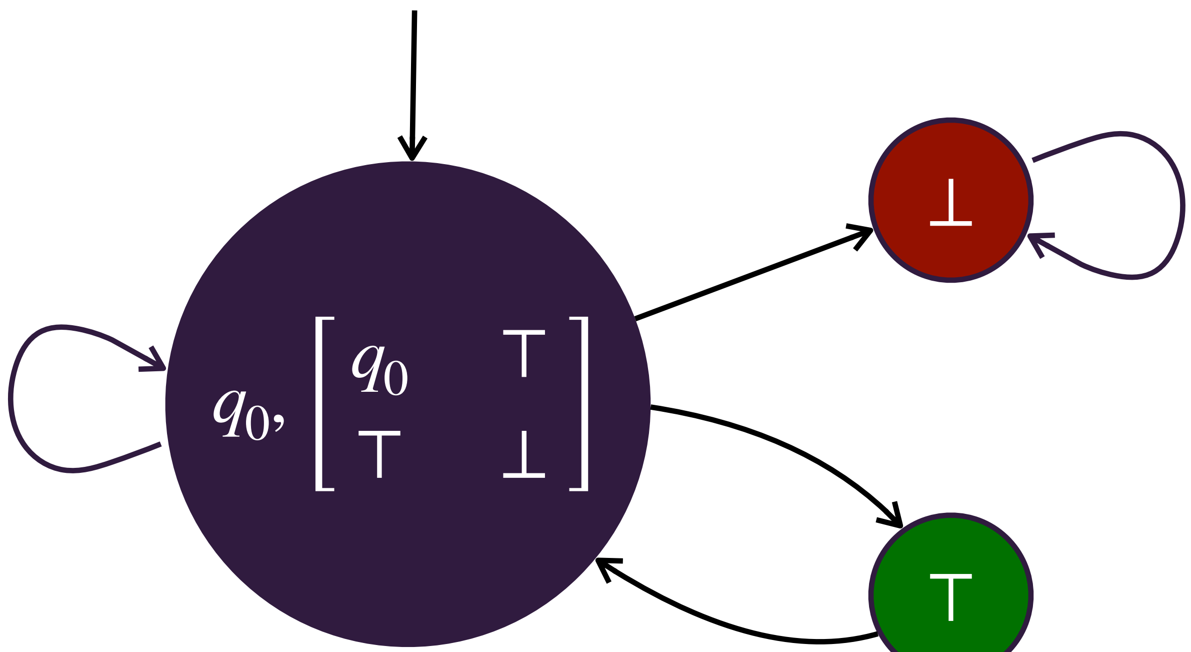

A tool to apprehend concurrent interactions: game forms

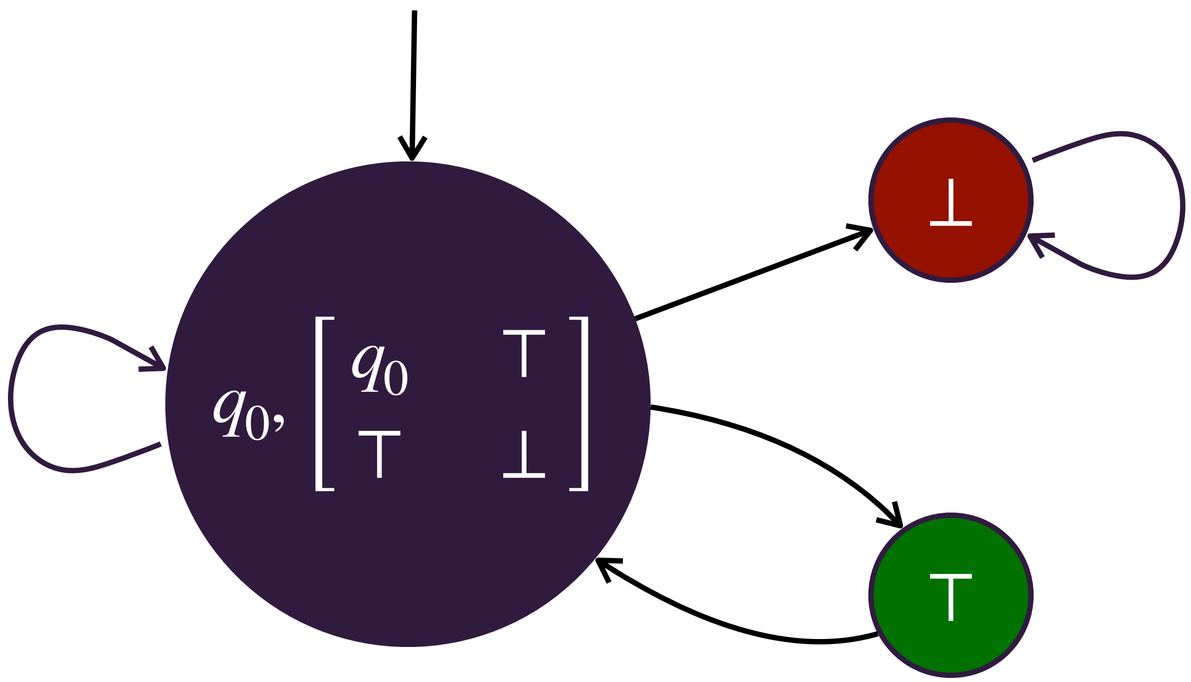

Game form

$$\begin{bmatrix} x & y \\ y & z \end{bmatrix}$$

A tool to apprehend concurrent interactions: game forms

Elementary brick

A tool to apprehend concurrent interactions: game forms

Elementary brick

Nice constructions ⟵ Nice bricks

A tool to apprehend concurrent interactions: game forms

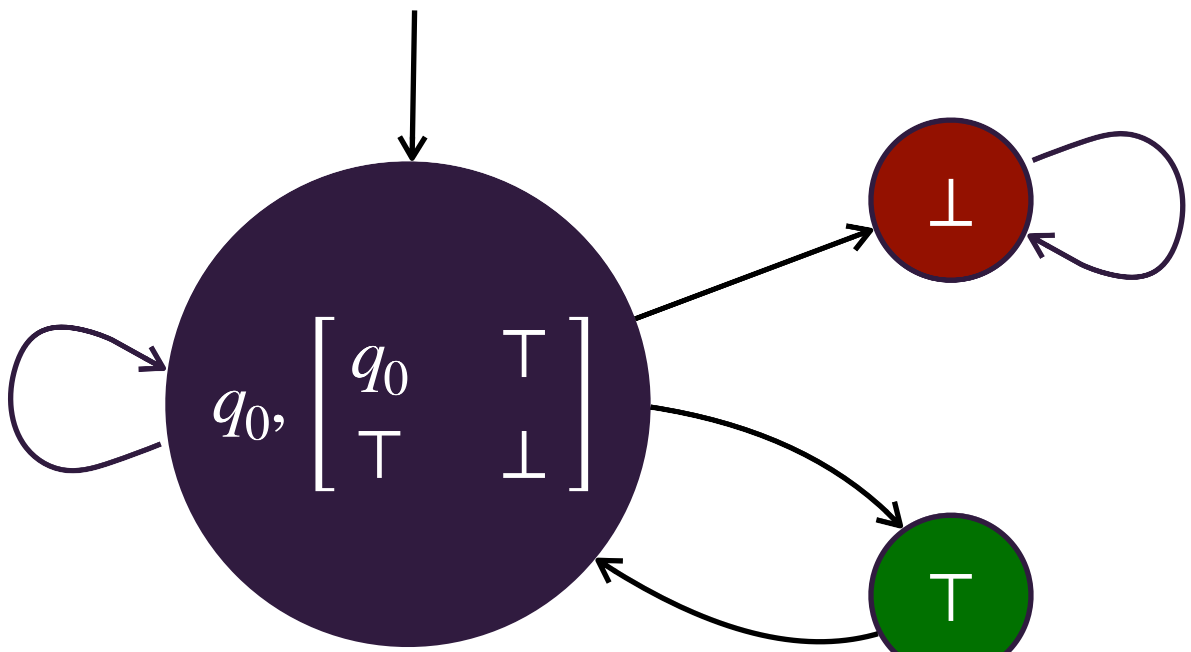

Game form

$$\begin{bmatrix} x & y \\ y & z \end{bmatrix}$$

Games on graphs with good properties ← Game forms with good properties

Approach in this work

Approach in this work

- $\mathscr{I}$ set of game forms

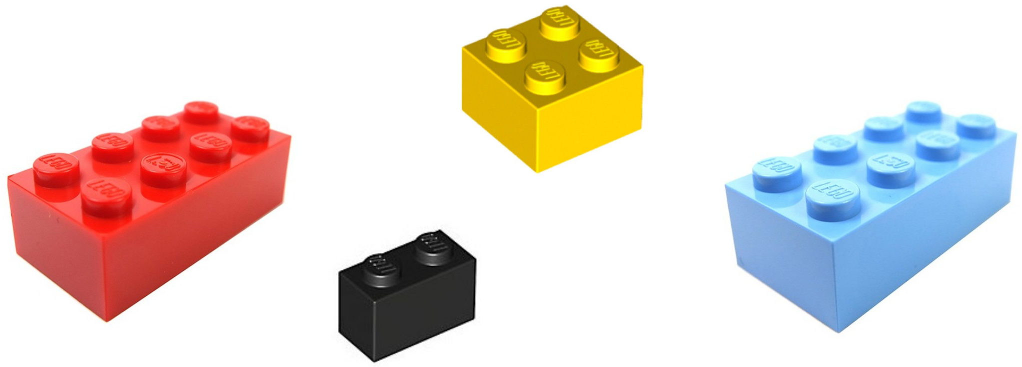

Approach in this work

- $\mathcal{I}$ set of game forms
- Identify properties of $\mathcal{I}$ so that all concurrent games built using game forms $\mathcal{I}$ behave well

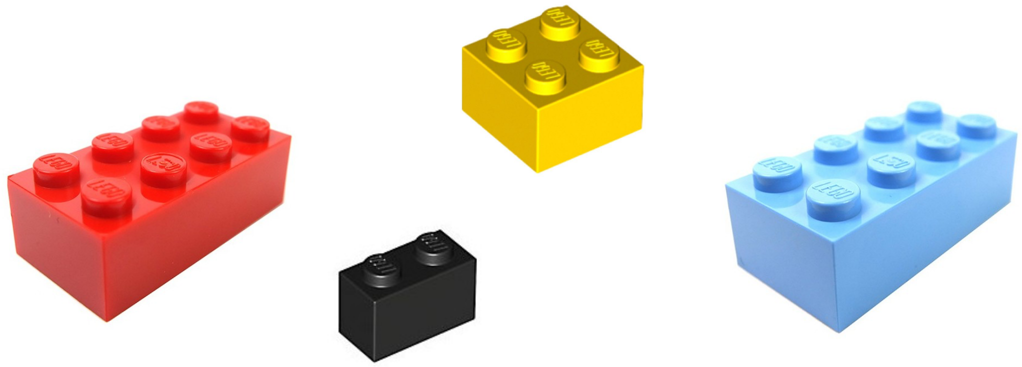

Approach in this work

- $\mathscr{I}$ set of game forms
- Identify properties of $\mathscr{I}$ so that all concurrent games built using game forms $\mathscr{I}$ behave well

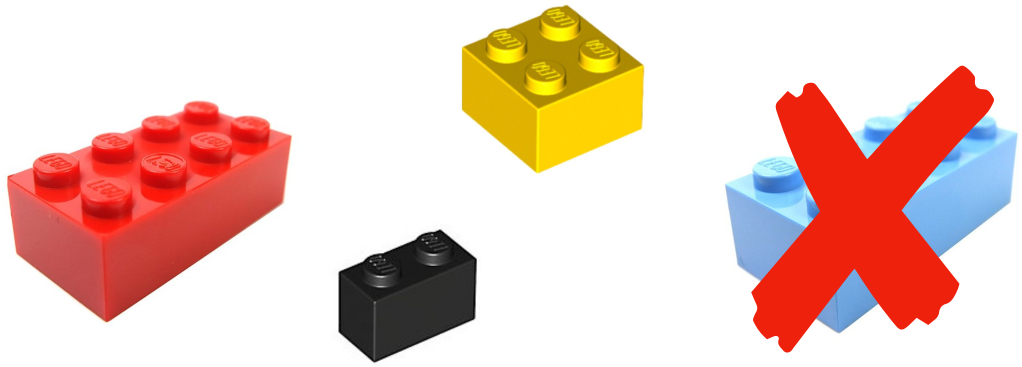

Approach in this work

- $\mathscr{I}$ set of game forms
- Identify properties of $\mathscr{I}$ so that all concurrent games built using game forms $\mathscr{I}$ behave well

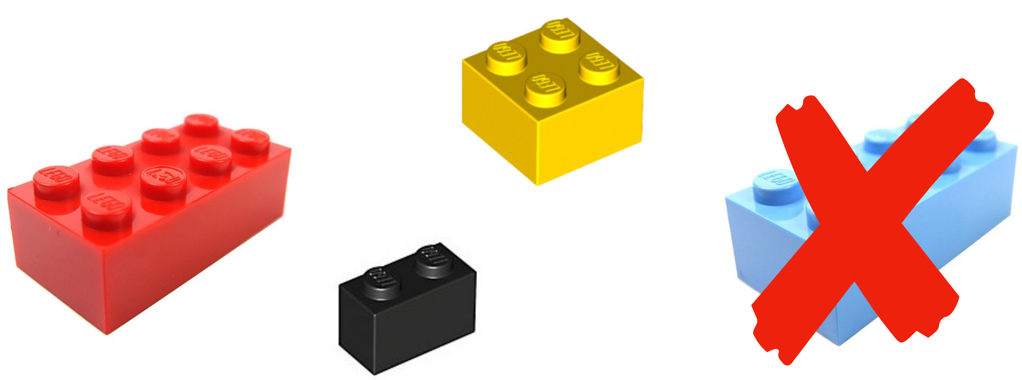

Approach in this work

- $\mathscr{I}$ set of game forms
- Identify properties of $\mathscr{I}$ so that all concurrent games built using game forms $\mathscr{I}$ behave well

Behave well = positional (almost-)optimal strategies are sufficient

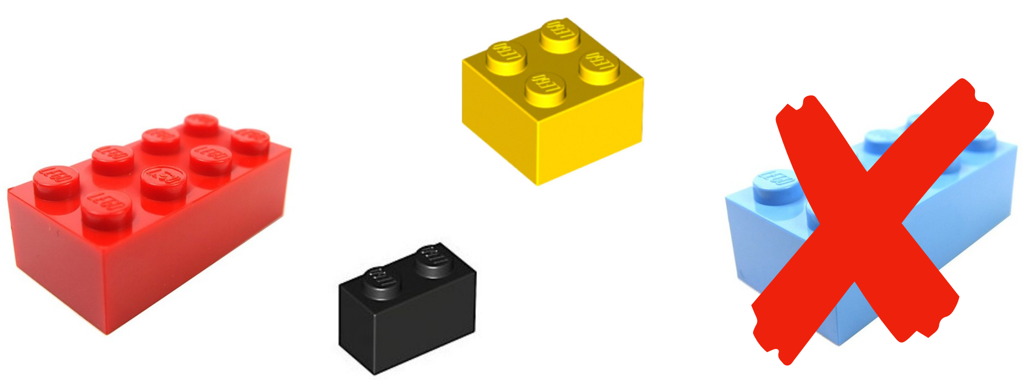

Approach in this work

- $\mathcal{I}$ set of game forms
- Identify properties of $\mathcal{I}$ so that all concurrent games built using game forms $\mathcal{I}$ behave well

Behave well = positional (almost-)optimal strategies are sufficient

(Co-)Büchi conditions

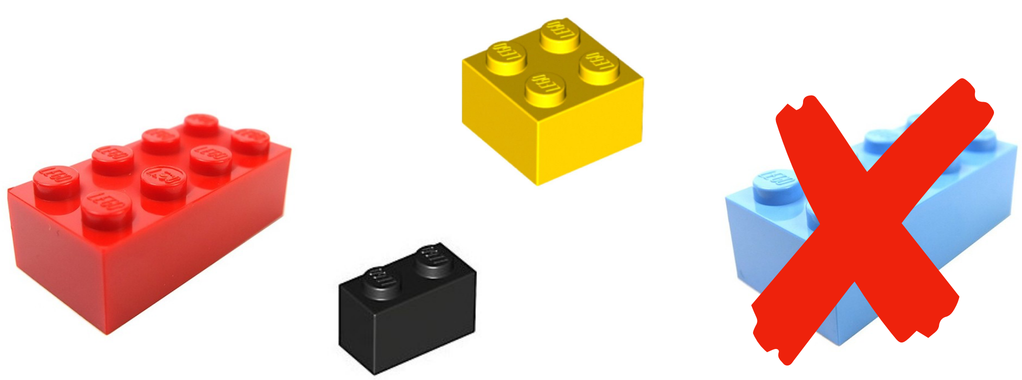

Previous works with a similar methodology

[BBL21] Bordais, Bouyer, Le Roux. From local to global determinacy in concurrent graph games (FSTTCS'21)
[BBL22] Bordais, Bouyer, Le Roux. Optimal Strategies in Concurrent Reachability Games (CSL'22)

Previous works with a similar methodology

- Determinacy of deterministic games [BBL21]
  - The matching-penny is not a good game form
  - Local determinacy condition on game forms

[BBL21] Bordais, Bouyer, Le Roux. From local to global determinacy in concurrent graph games (FSTTCS'21)
[BBL22] Bordais, Bouyer, Le Roux. Optimal Strategies in Concurrent Reachability Games (CSL'22)

Previous works with a similar methodology

- Determinacy of deterministic games [BBL21]
  - The matching-penny is not a good game form
  - Local determinacy condition on game forms
- Reachability objectives [BBL22]
  - Optimal and almost-optimal strategies can be chosen positional (when they exist)
  - Local condition (called RM) on game forms to ensure existence (and therefore positionality) of optimal strategies everywhere

[BBL21] Bordais, Bouyer, Le Roux. From local to global determinacy in concurrent graph games (FSTTCS'21)
[BBL22] Bordais, Bouyer, Le Roux. Optimal Strategies in Concurrent Reachability Games (CSL'22)

What game theory tells us

- One can associate to each state q of the game its value $v(q)$, and these values satisfy local optimality equations

What game theory tells us

- One can associate to each state q of the game its value $v(q)$, and these values satisfy local optimality equations

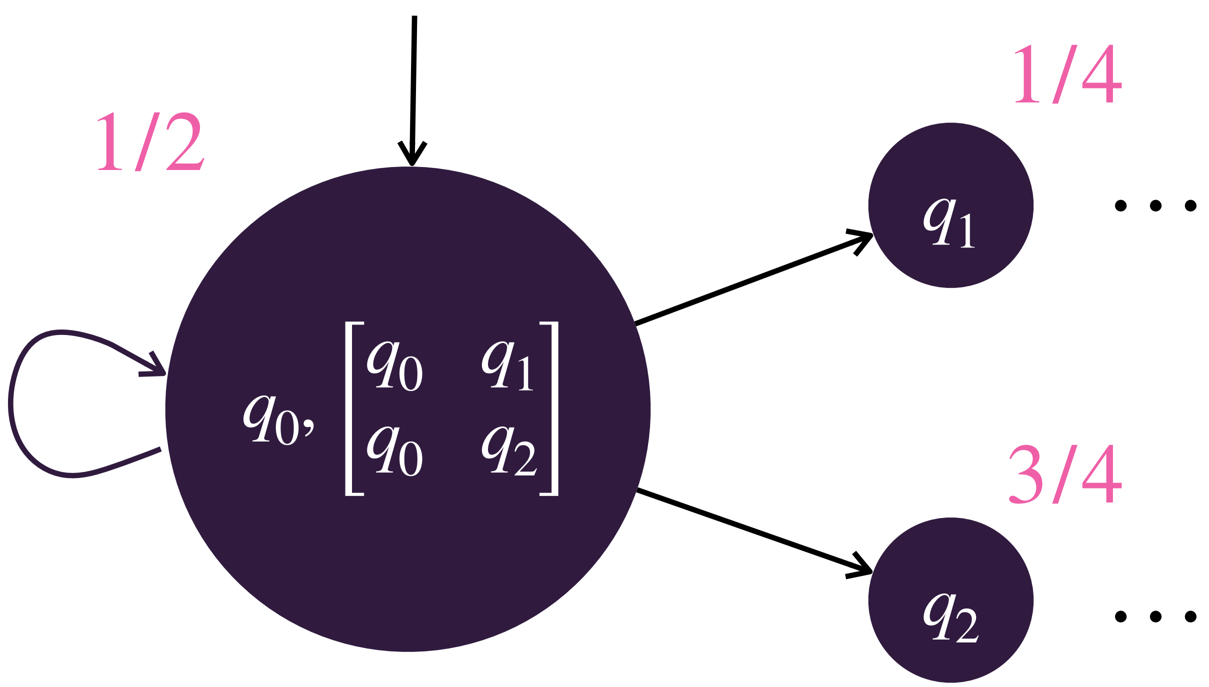

What game theory tells us

- One can associate to each state q of the game its value $v(q)$, and these values satisfy local optimality equations

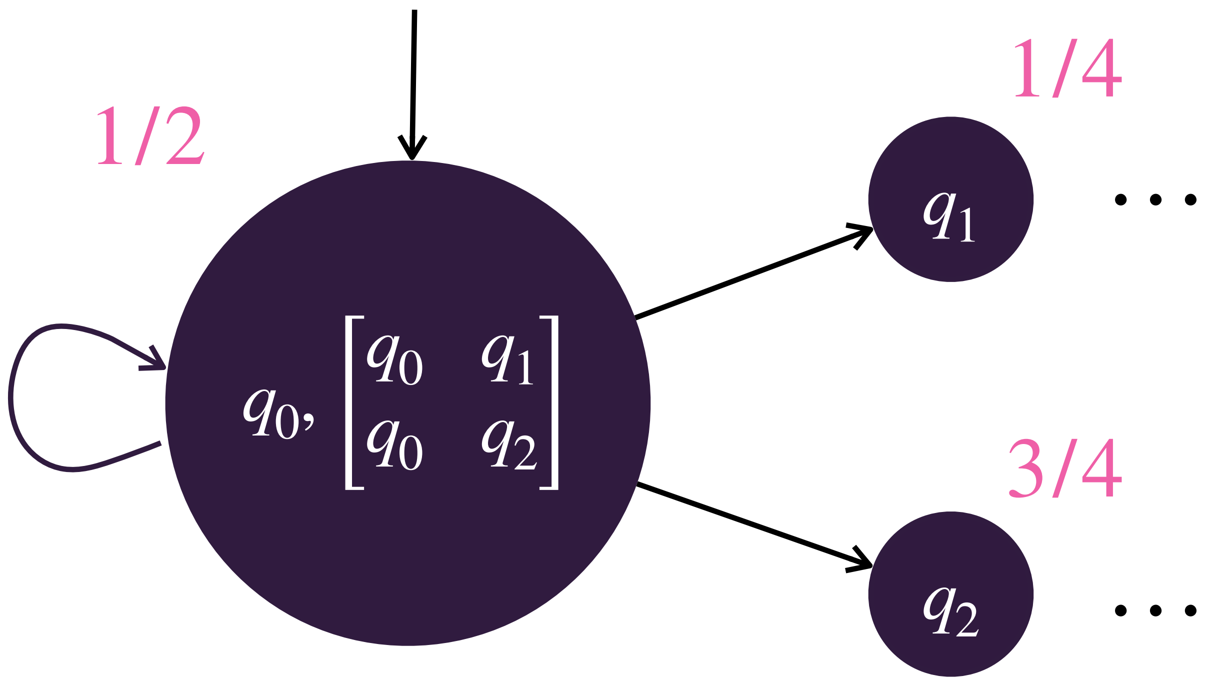

Game in normal form

$$\begin{bmatrix} 1/2 & 1/4 \\ 1/2 & 3/4 \end{bmatrix}$$

What game theory tells us

- One can associate to each state q of the game its value $v(q)$, and these values satisfy local optimality equations

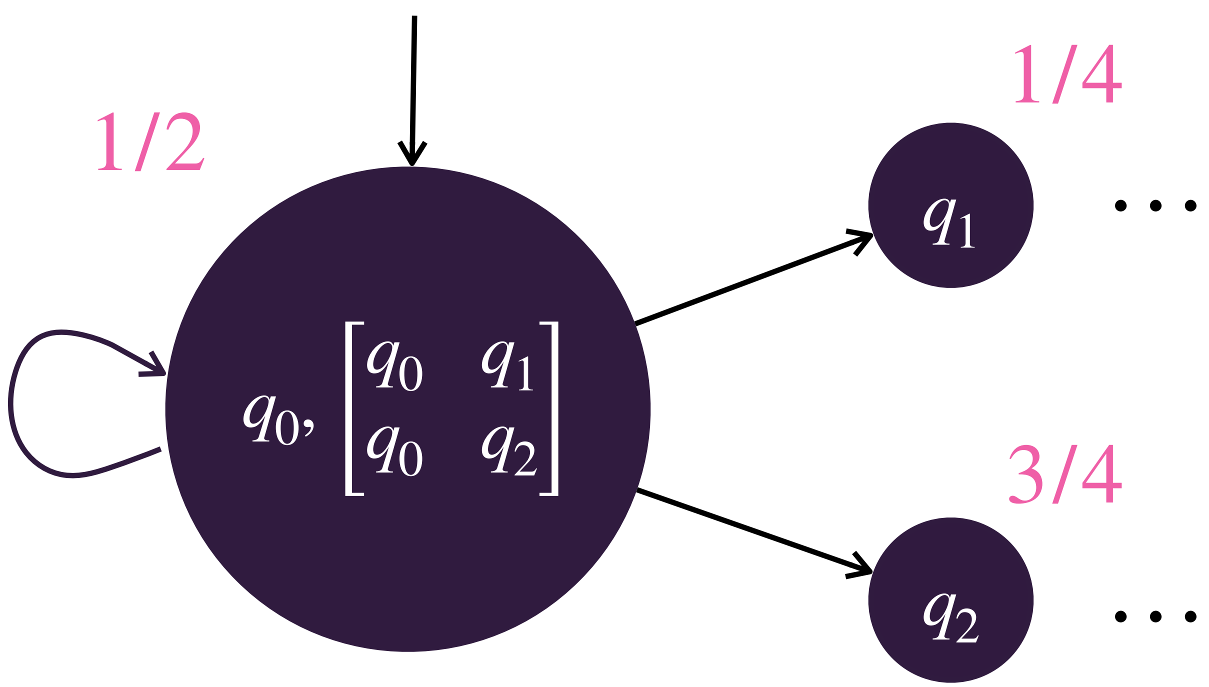

Game in normal form

$$\begin{bmatrix} 1/2 & 1/4 \\ 1/2 & 3/4 \end{bmatrix}$$

MinMax theorem
(van Neumann)

- Both players have (local) optimal strategies in this game in normal form

What game theory tells us

- One can associate to each state q of the game its value $v(q)$, and these values satisfy local optimality equations

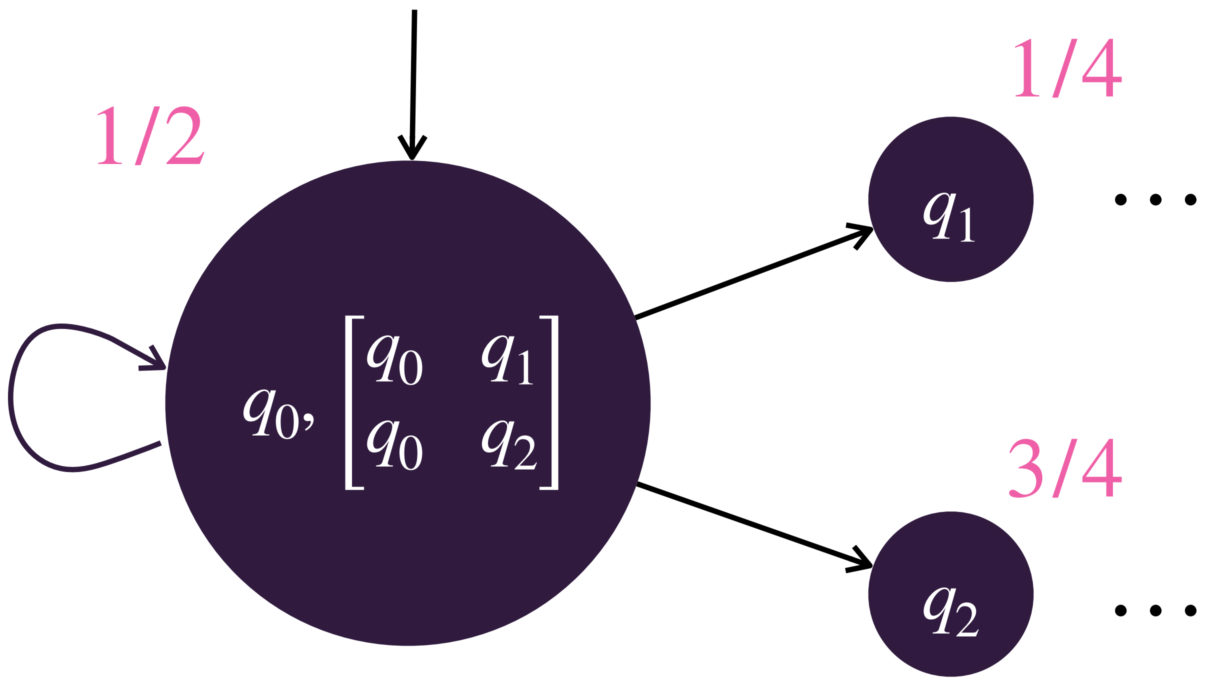

Game in normal form

$$\begin{bmatrix} 1/2 & 1/4 \\ 1/2 & 3/4 \end{bmatrix}$$

MinMax theorem (van Neumann)

- Both players have (local) optimal strategies in this game in normal form

- All globally optimal strategies (in the graph) are locally optimal

What game theory tells us

- One can associate to each state q of the game its value $v(q)$, and these values satisfy local optimality equations

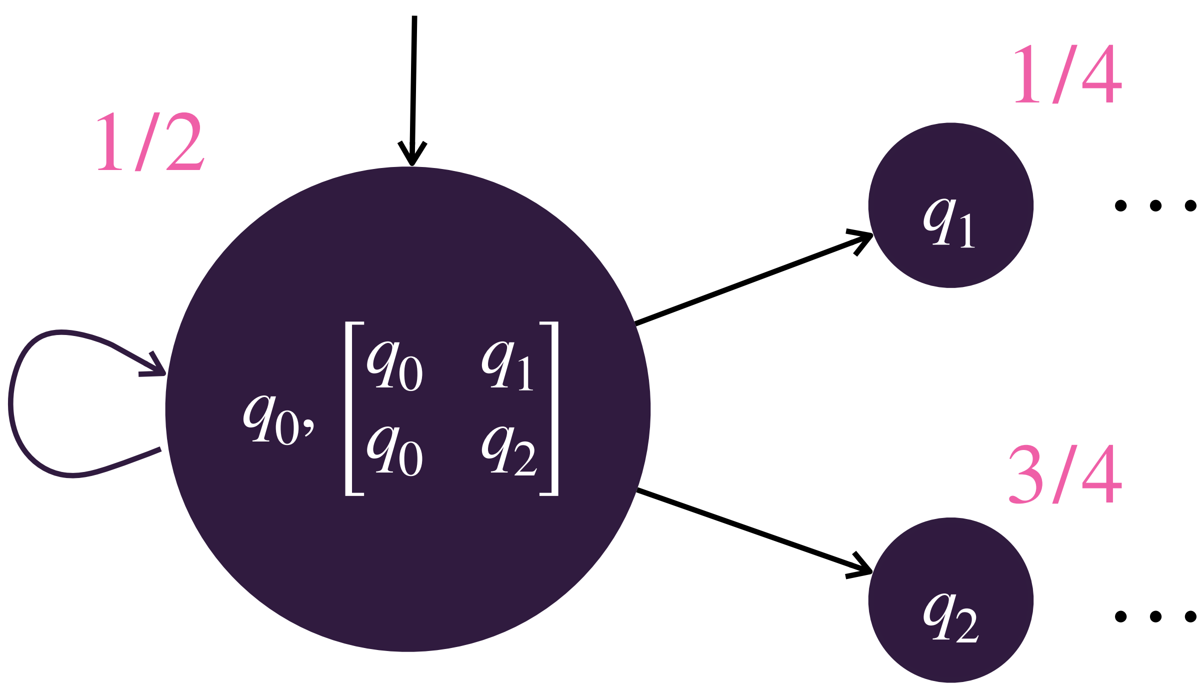

Game in normal form

$$\begin{bmatrix} 1/2 & 1/4 \\ 1/2 & 3/4 \end{bmatrix}$$

MinMax theorem (van Neumann)

- Both players have (local) optimal strategies in this game in normal form
- All globally optimal strategies (in the graph) are locally optimal
- Locally optimal strategies may not be globally optimal (in the graph)

Example

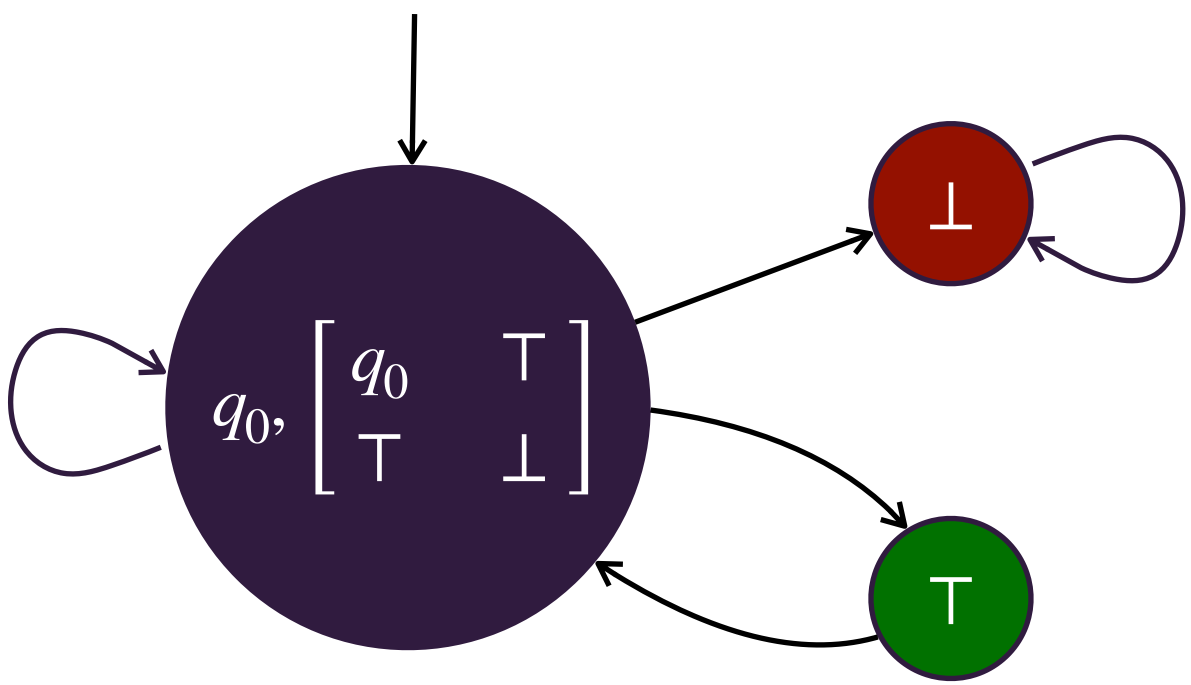

Example

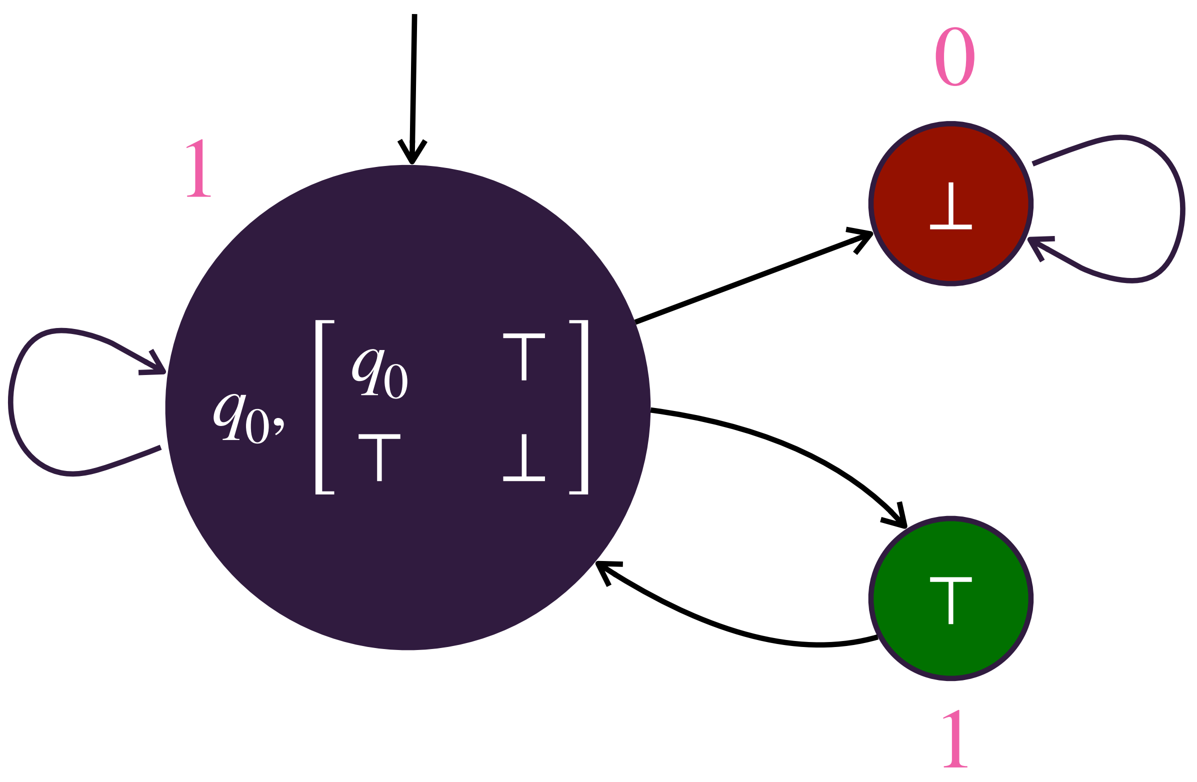

Example

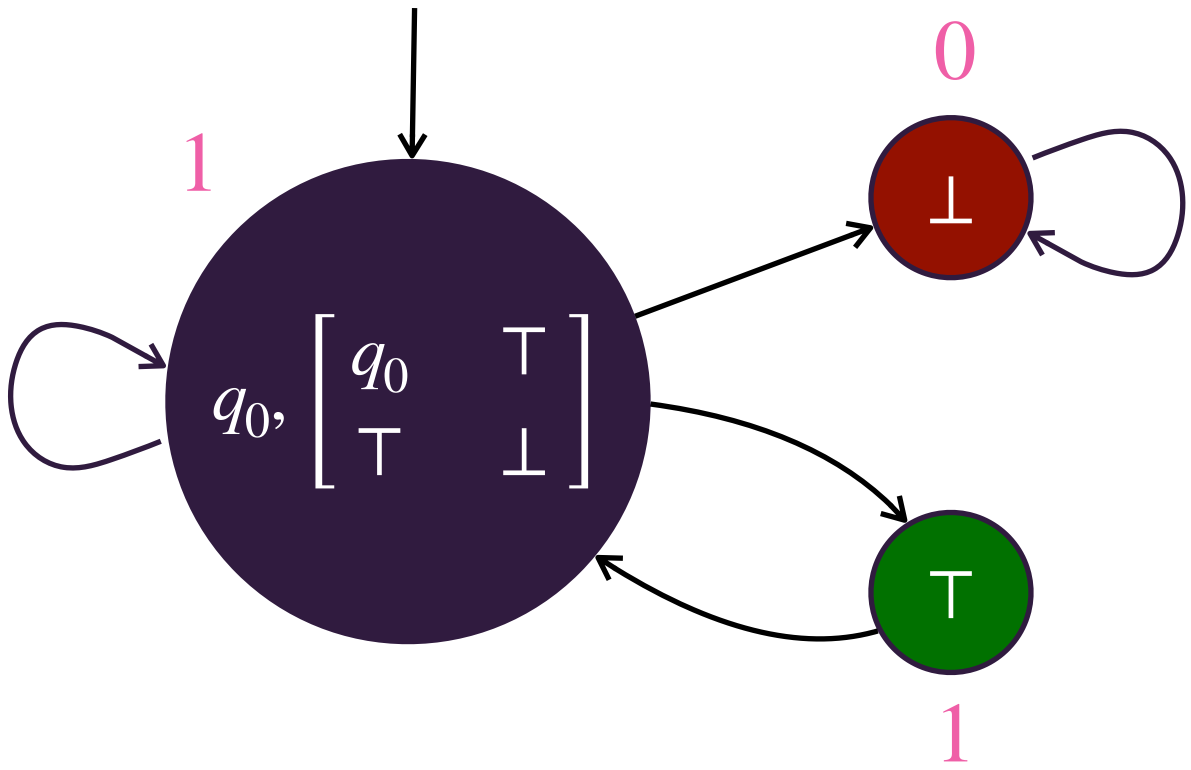

- Locally optimal strategy $\boldsymbol{\sigma}_{\mathsf{A}}$:

Example

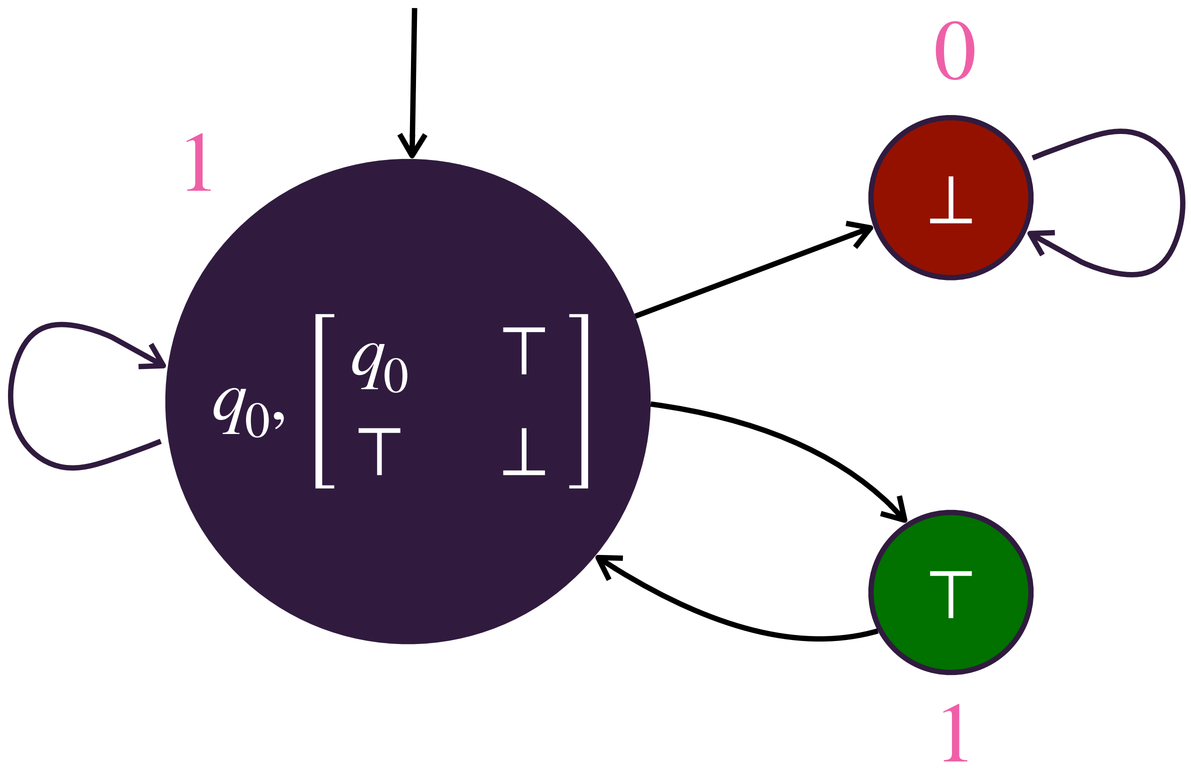

- Locally optimal strategy $\boldsymbol{\sigma}_{\mathsf{A}}$:
  - Player A chooses the first row

Example

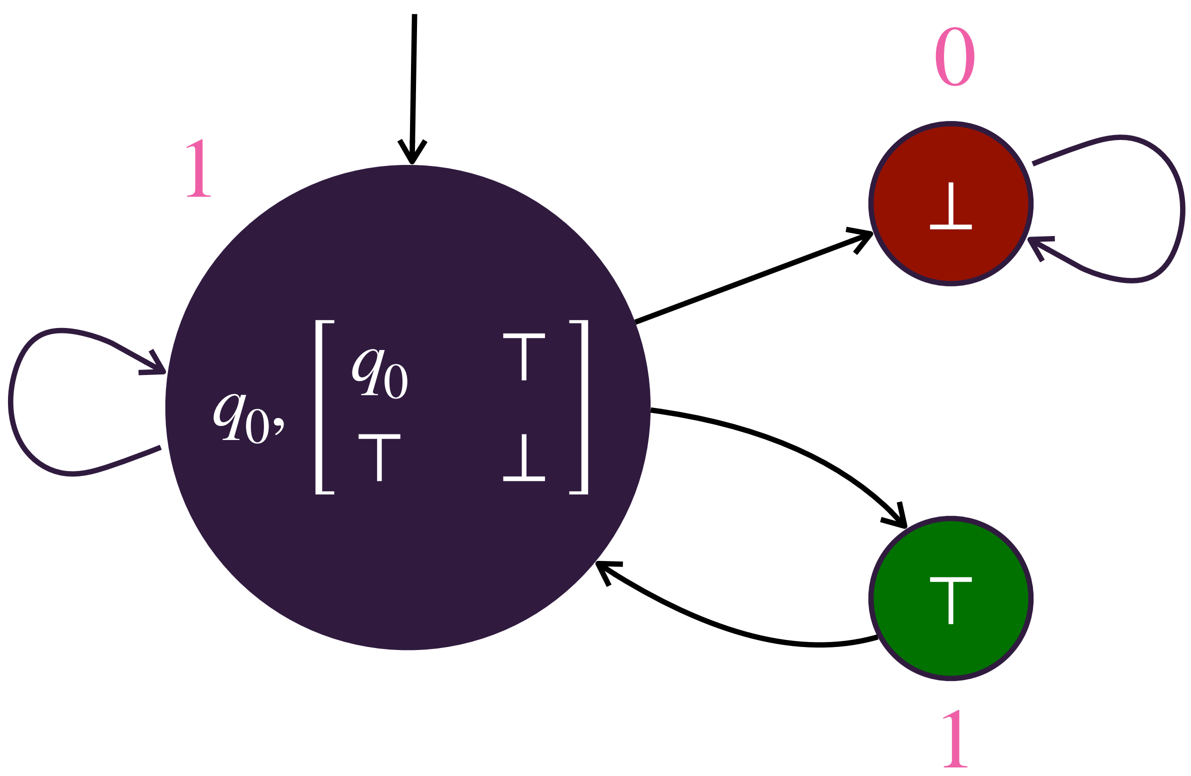

- Locally optimal strategy $\boldsymbol{\sigma}_{\mathsf{A}}$:
  - Player A chooses the first row
- This is obviously not globally optimal

Example

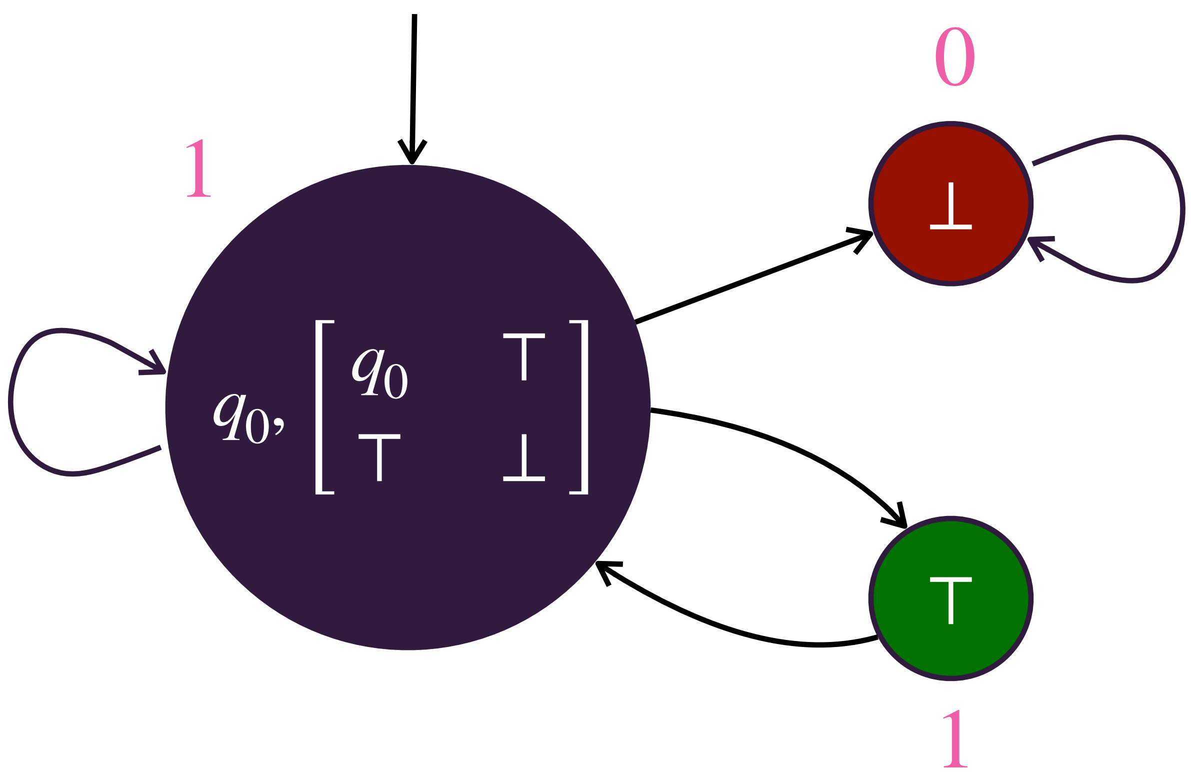

- Locally optimal strategy $\boldsymbol{\sigma}_{\mathsf{A}}$:
  - Player A chooses the first row
- This is obviously not globally optimal
- What is wrong?

Example

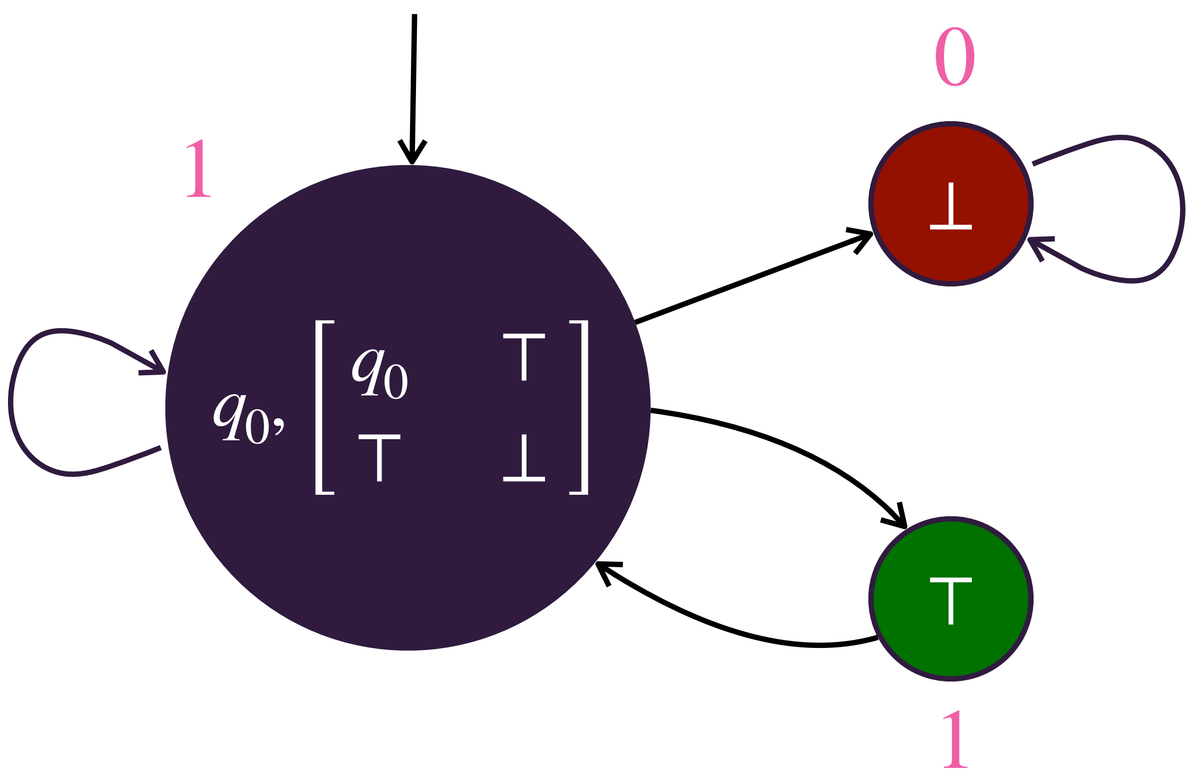

- Locally optimal strategy $\boldsymbol{\sigma}_{\mathsf{A}}$:
  - Player A chooses the first row
- This is obviously not globally optimal
- What is wrong?
  - In the MDP generated by $\boldsymbol{\sigma}_{\mathsf{A}}$, there is an end-component which is losing

Our contributions

Our contributions

Characterize positional (almost-)optimal strategies using locally (almost-)optimal strategies (applies to tail objectives)

Our contributions

Characterize positional (almost-)optimal strategies using locally (almost-)optimal strategies (applies to tail objectives)

Büchi objectives

- Optimal strategies may not exist (known)
- When optimal strategies exist from all states, they can be chosen positional (inherited from reachability games)
- Almost-optimal strategies may require infinite memory (known)
- Characterization of nice game forms (aBM) for ensuring:
  - Positional almost-optimal strategies

Our contributions

Characterize positional (almost-)optimal strategies using locally (almost-)optimal strategies (applies to tail objectives)

Büchi objectives

- Optimal strategies may not exist (known)
- When optimal strategies exist from all states, they can be chosen positional (inherited from reachability games)
- Almost-optimal strategies may require infinite memory (known)
- Characterization of nice game forms (aBM) for ensuring:
  - Positional almost-optimal strategies

co-Büchi objectives

- Optimal strategies may not exist (known)
- When optimal strategies exist, they may require infinite memory
- Almost-optimal strategies can be chosen positional (known [CDAH06])
- Characterization of nice game forms (coBM) for ensuring:
  - Positional optimal strategies

[CDAH06] K. Chatterjee, L. De Alfaro, T. Henzinger. The complexity of quantitative concurrent parity games (SODA'06).

Our contributions

Characterize positional (almost-)optimal strategies using locally (almost-)optimal strategies (applies to tail objectives)

Büchi objectives

- Optimal strategies may not exist (known)
- When optimal strategies exist from all states, they can be chosen positional (inherited from reachability games)
- Almost-optimal strategies may require infinite memory (known)
- Characterization of nice game forms (aBM) for ensuring:
  - Positional almost-optimal strategies

co-Büchi objectives

- Optimal strategies may not exist (known)
- When optimal strategies exist, they may require infinite memory
- Almost-optimal strategies can be chosen positional (known [CDAH06])
- Characterization of nice game forms (coBM) for ensuring:
  - Positional optimal strategies

[CDAH06] K. Chatterjee, L. De Alfaro, T. Henzinger. The complexity of quantitative concurrent parity games (SODA'06).

From bricks to nice constructions

- Recall: there exist Büchi games where infinite memory is required to play ε-optimally

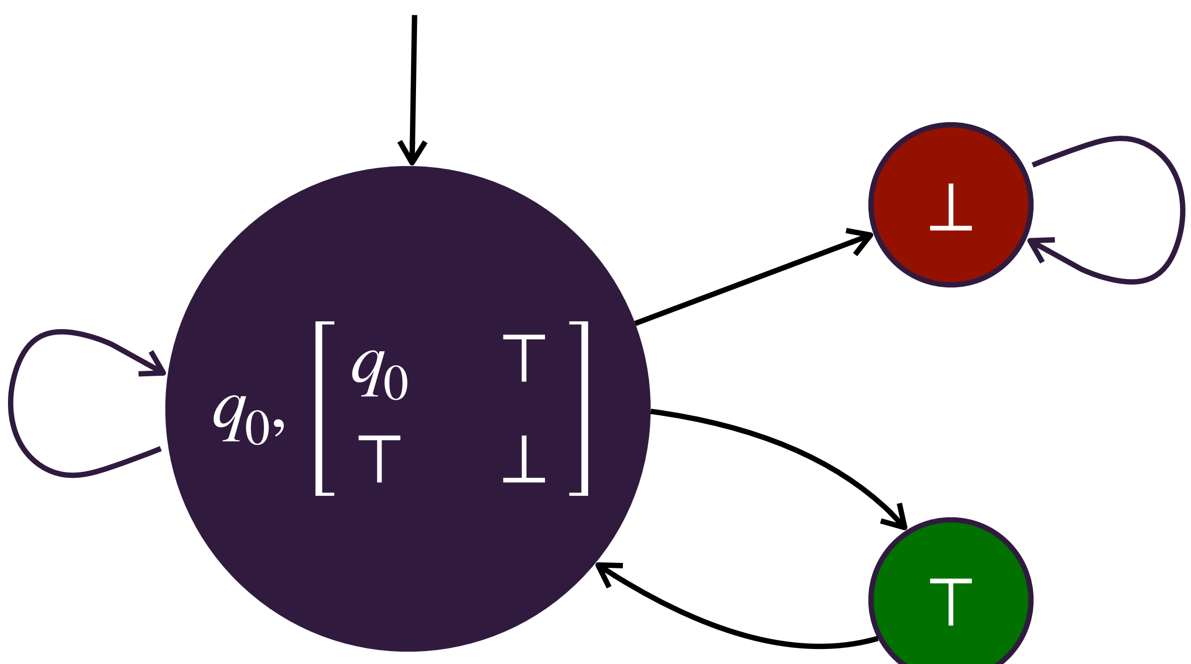

From bricks to nice constructions

- Recall: there exist Büchi games where infinite memory is required to play ε-optimally

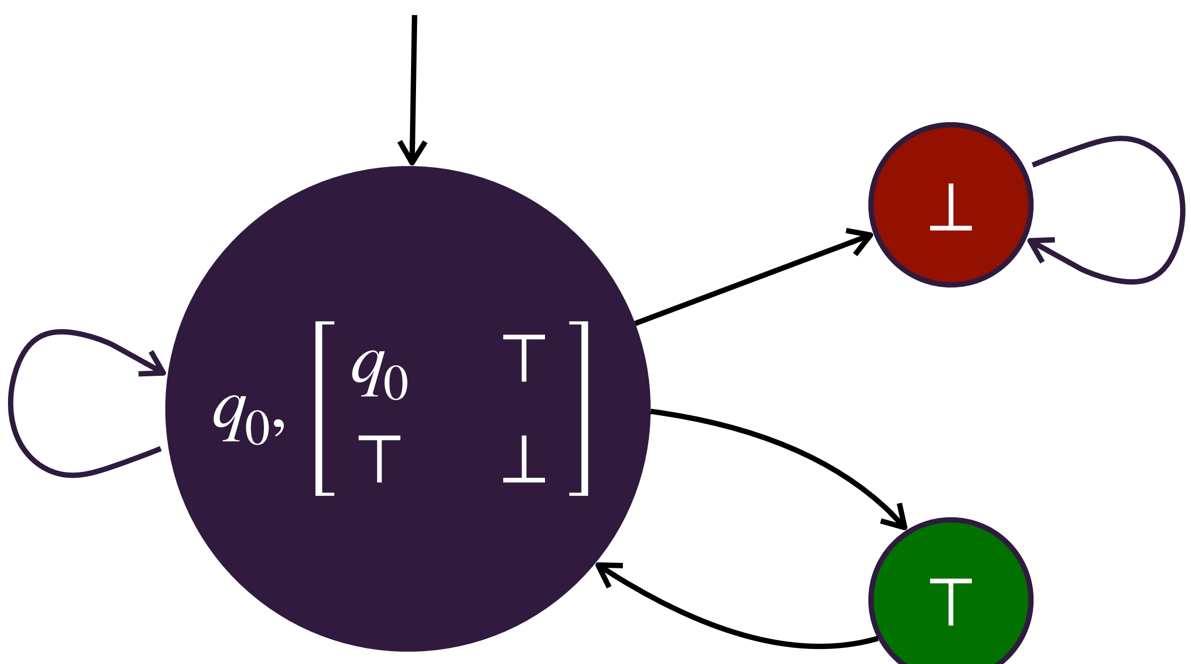

- How should we restrict interactions to avoid this phenomenon?

From bricks to nice constructions

- Recall: there exist Büchi games where infinite memory is required to play ε-optimally

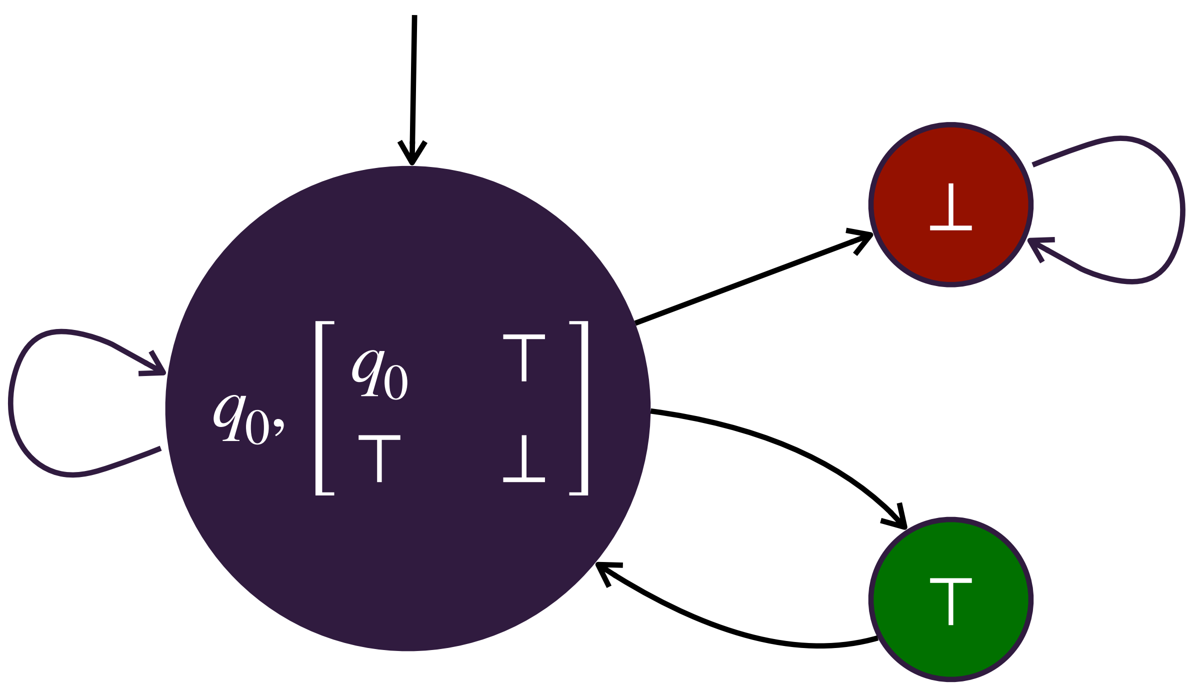

- How should we restrict interactions to avoid this phenomenon?

$$\mathscr{F} = \begin{bmatrix} x & y \\ y & z \end{bmatrix}$$

From bricks to nice constructions

- Recall: there exist Büchi games where infinite memory is required to play ε-optimally

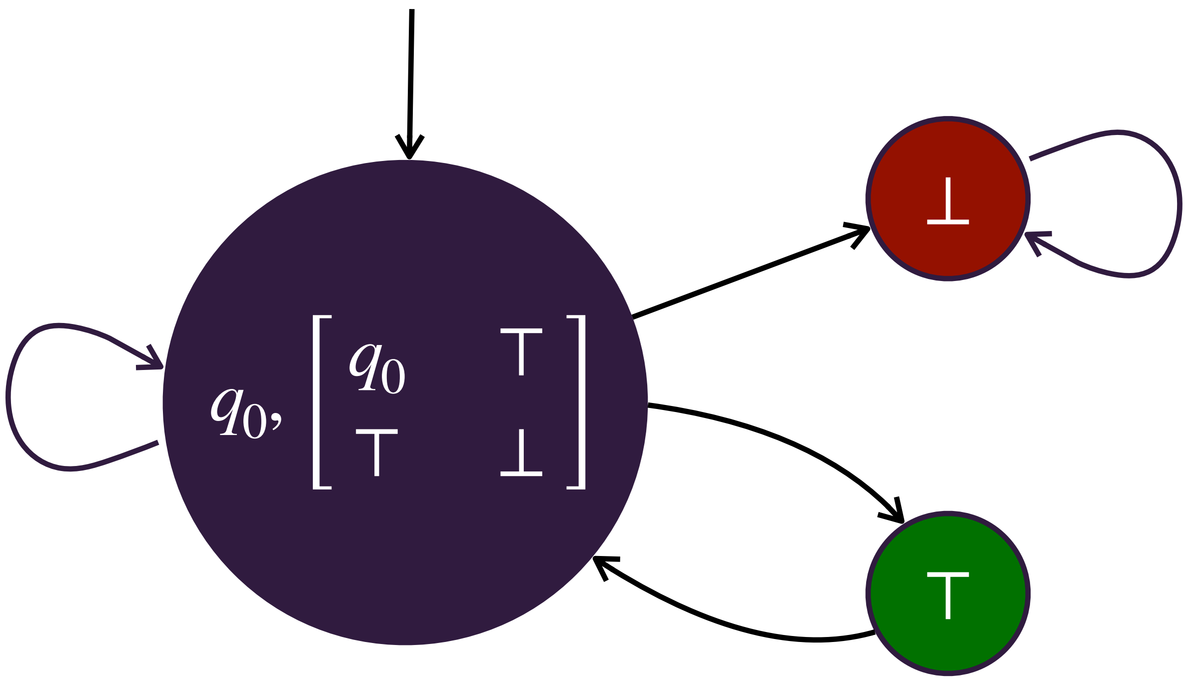

- How should we restrict interactions to avoid this phenomenon?

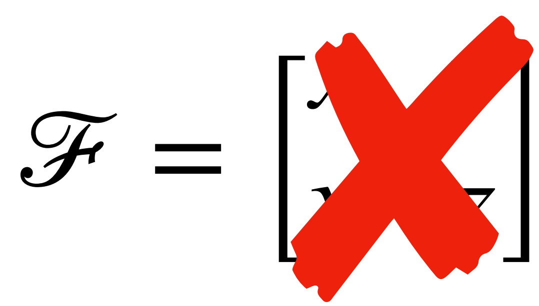

Focus on Büchi conditions

$$\begin{bmatrix} x & y \\ y & z \end{bmatrix}$$

Focus on Büchi conditions

$$\begin{bmatrix} x & y \\ y & z \end{bmatrix} \xrightarrow[\text{« environments »}]{\text{Several local}}$$

Focus on Büchi conditions

$$\begin{bmatrix} x & y \\ y & z \end{bmatrix}$$

Several local « environments » →

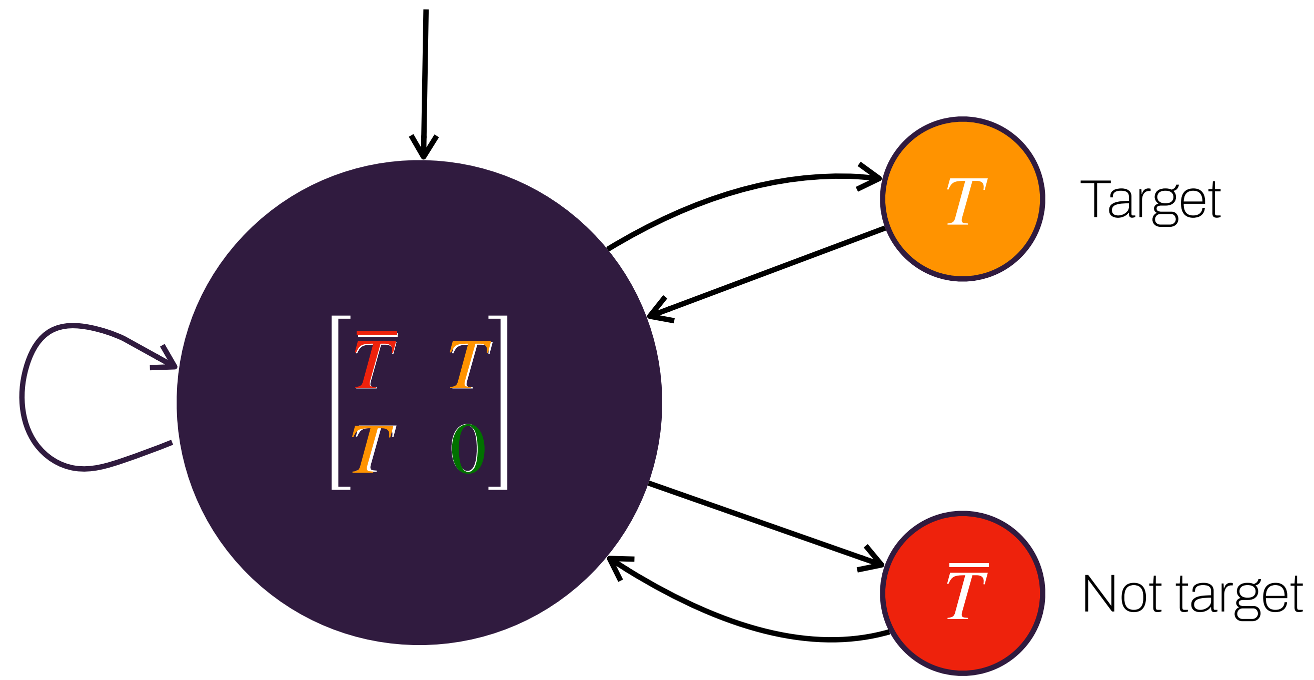

Focus on Büchi conditions

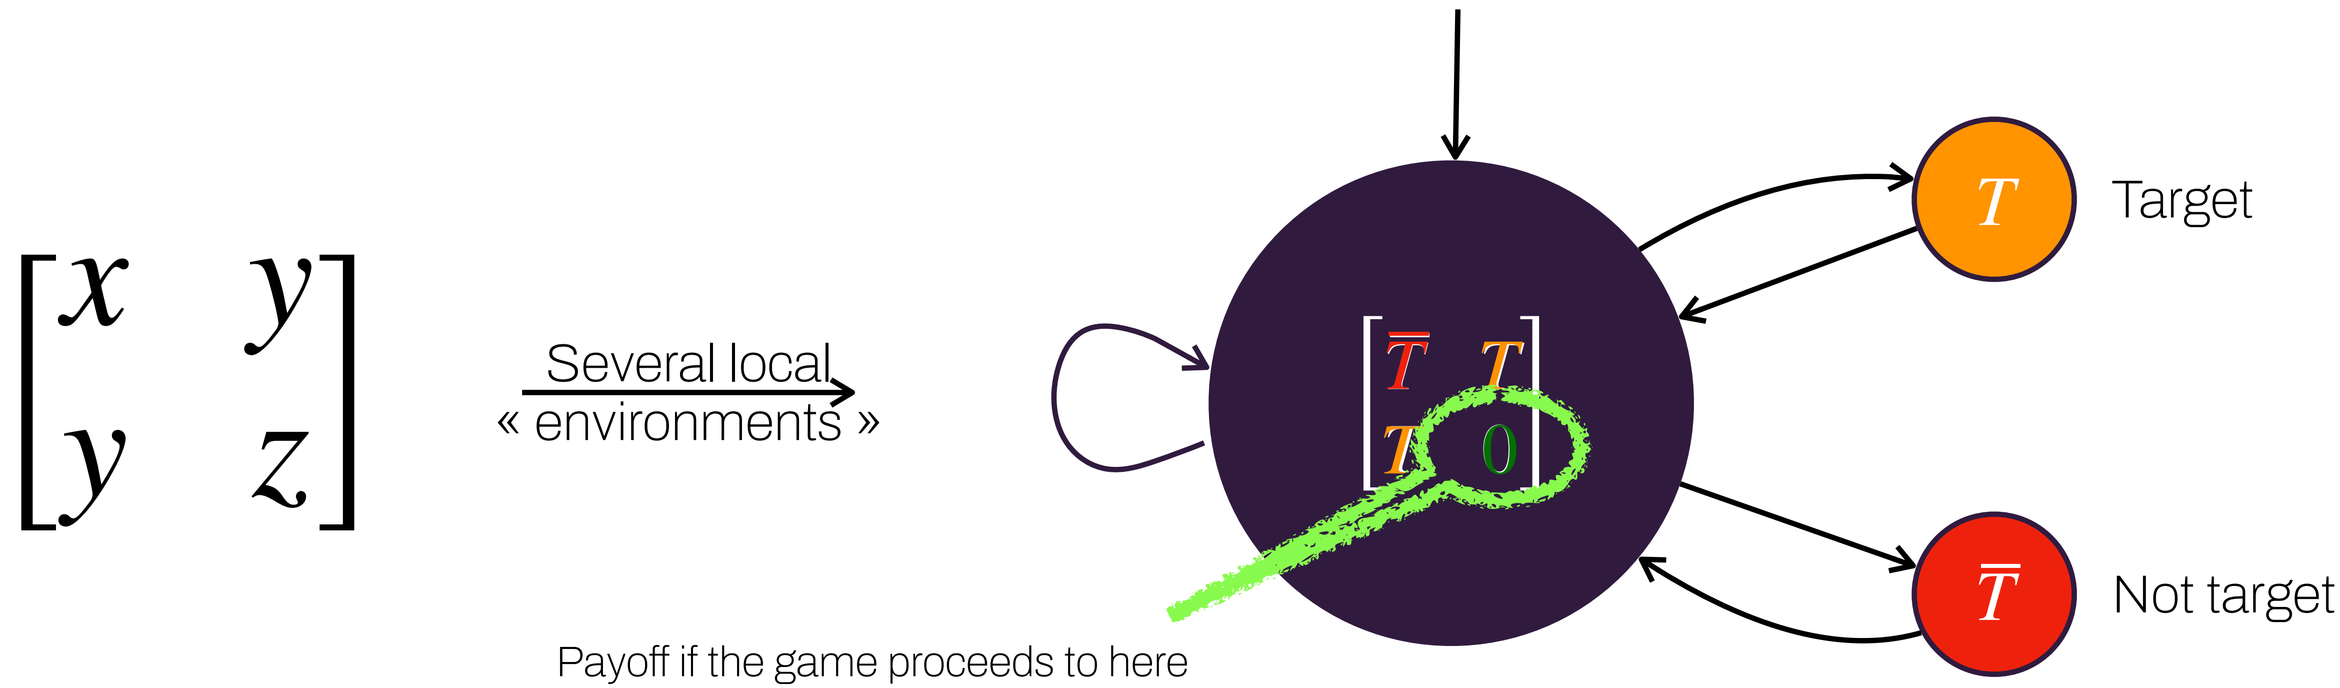

Focus on Büchi conditions

$$\begin{bmatrix} x & y \\ y & z \end{bmatrix}$$

Several local « environments » →

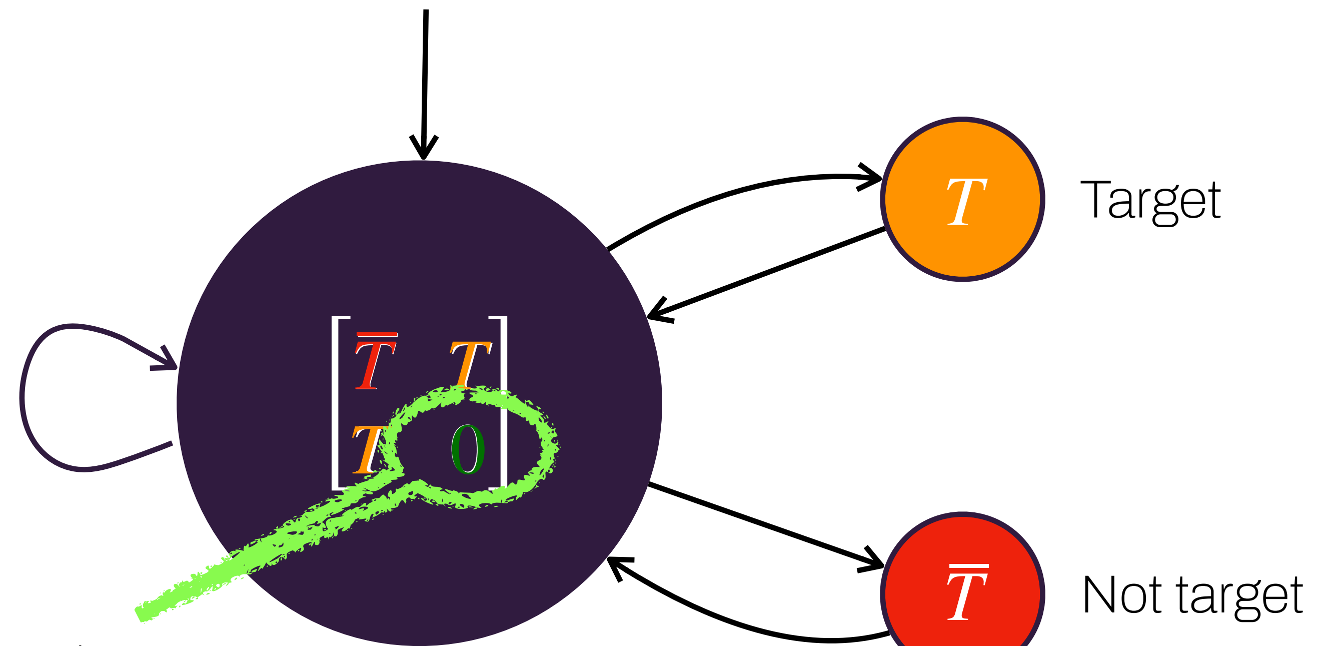

Payoff if the game proceeds to here

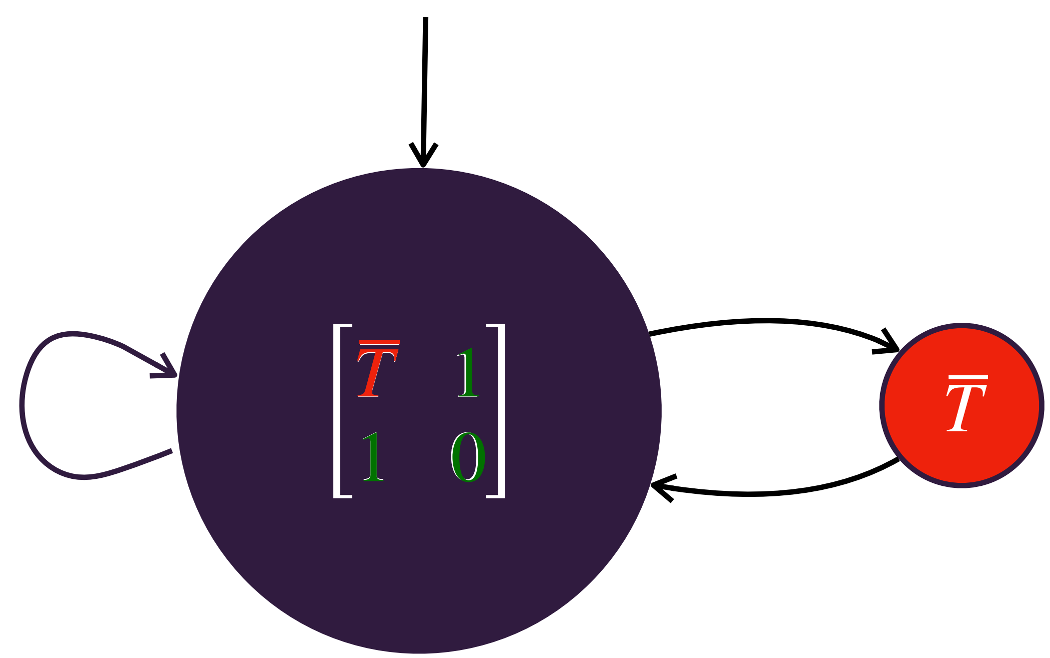

Focus on Büchi conditions

$$\begin{bmatrix} x & y \\ y & z \end{bmatrix}$$

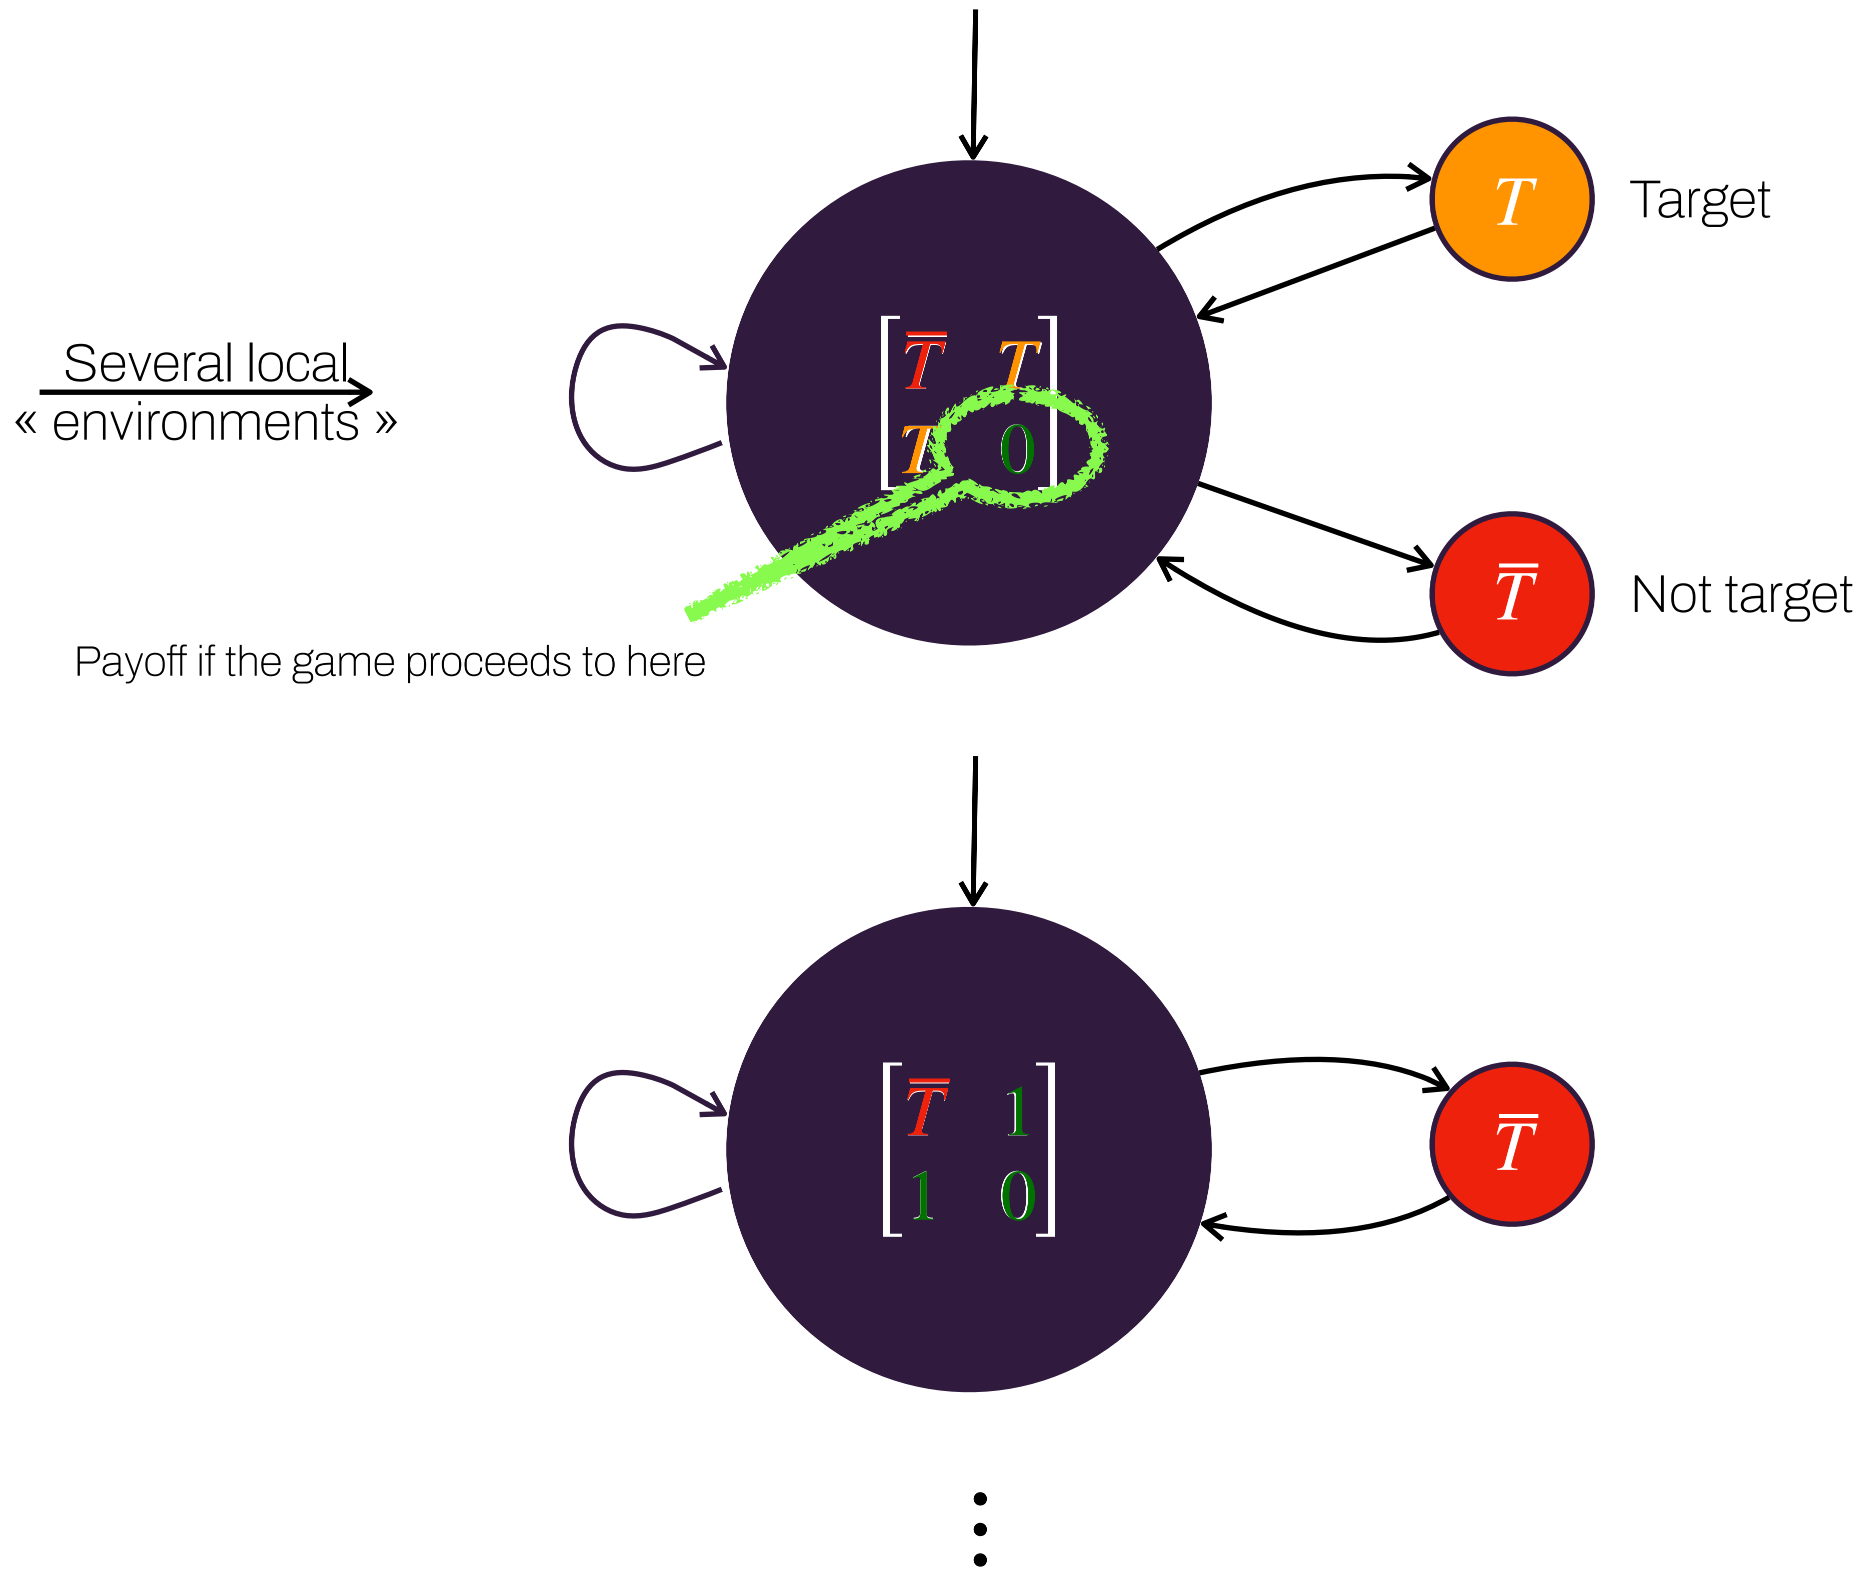

Focus on Büchi conditions

$$\begin{bmatrix} x & y \\ y & z \end{bmatrix}$$

Several local « environments »

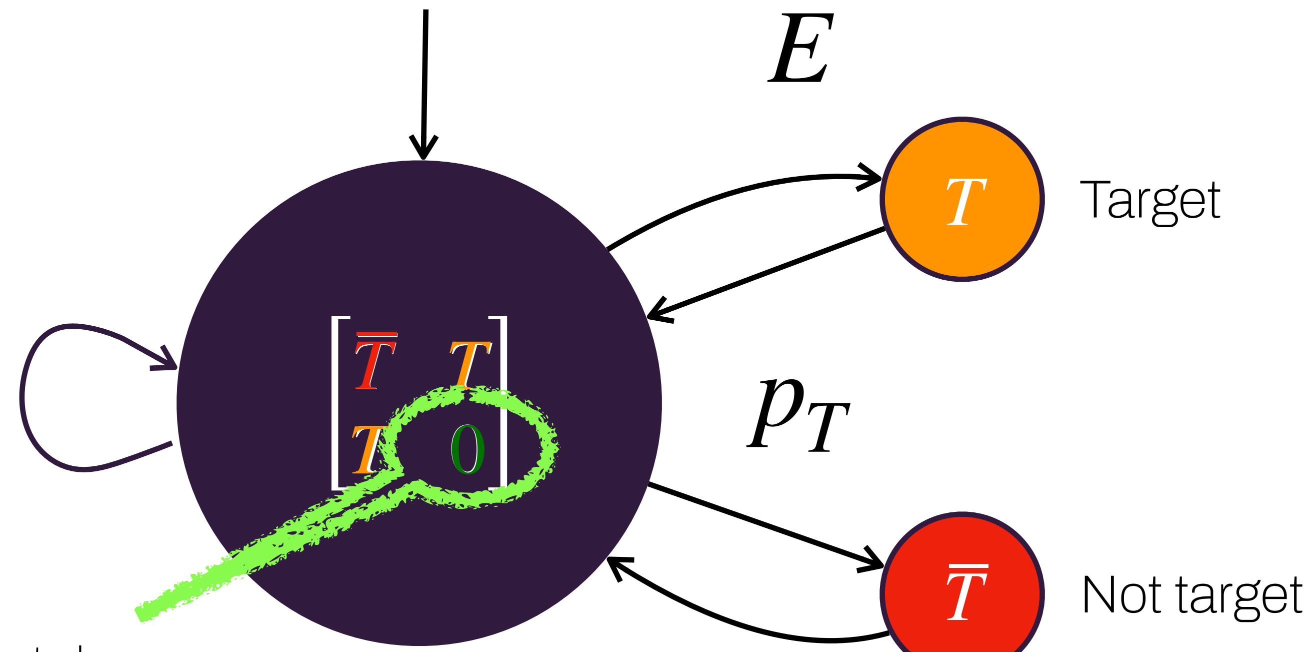

Payoff if the game proceeds to here

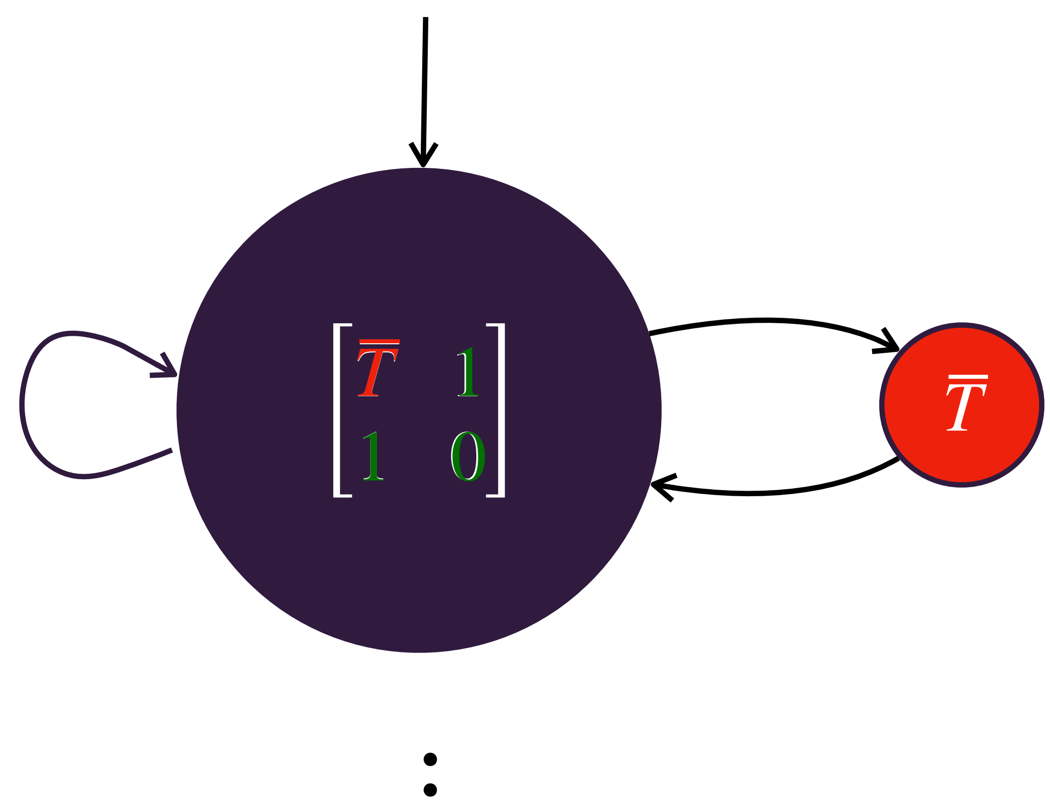

Local environment

- O set of variables ($\{x, y, z\}$ in the example)
- One small game
  - for every $E \subseteq O$,
  - for every $p_T : E \to [0,1]$, and
  - for every $\alpha : O \backslash E \to [0,1]$

Characterization

Definition of aBM (almost-Büchi maximizable)

- A game form $\mathcal{F}$ is aBM whenever every embedding of $\mathcal{F}$ into a local environment admits a positional ε-optimal strategy for every $\varepsilon > 0$.

Characterization

Definition of aBM (almost-Büchi maximizable)

- A game form $\mathscr{F}$ is aBM whenever every embedding of $\mathscr{F}$ into a local environment admits a positional ε-optimal strategy for every $\varepsilon > 0$.

- An aBM game form can be characterized and decided (it can be encoded as a formula of the first-order theory of the reals)

Characterization

Definition of aBM (almost-Büchi maximizable)

- A game form $\mathscr{F}$ is aBM whenever every embedding of $\mathscr{F}$ into a local environment admits a positional ε-optimal strategy for every $\varepsilon > 0$.

- An aBM game form can be characterized and decided (it can be encoded as a formula of the first-order theory of the reals)

Characterization

- If all game forms used in a concurrent game $\mathscr{G}$ are aBM, then $\mathscr{G}$ admits positional ε-optimal strategies for every $\varepsilon > 0$
- If a game form is not aBM, then there is a concurrent game which does not admit a positional ε-optimal strategy for some $\varepsilon > 0$.

How to ensure positional (almost-)optimal strategies?

Existence of positional optimal or ε-optimal strategies under the following restrictions on game forms:

	Positional opt. strat.		Positional almost-opt. strat.	
	Target	Not target	Target	Not target
Safety obj.	No restr.	No restr.	No restr.	No restr.
Reach. obj.	No restr.	RM	No restr.	No restr.
Büchi obj.	No restr.	RM	No restr.	aBM
co-Büchi obj.	RM	coBM	No restr.	No restr.

How to ensure positional (almost-)optimal strategies?

Existence of positional optimal or ε-optimal strategies under the following restrictions on game forms:

	Positional opt. strat.		Positional almost-opt. strat.	
	Target	Not target	Target	Not target
Safety obj.	No restr.	No restr.	No restr.	No restr.
Reach. obj.	No restr.	RM	No restr.	No restr.
Büchi obj.	No restr.	RM	No restr.	aBM
co-Büchi obj.	RM	coBM	No restr.	No restr.

If game forms satisfy the properties below, then positional strategies exist and can be chosen positional

How to ensure positional (almost-)optimal strategies?

Existence of positional optimal or ε-optimal strategies under the following restrictions on game forms:

	Positional opt. strat.		Positional almost-opt. strat.	
	Target	Not target	Target	Not target
Safety obj.	No restr.	No restr.	No restr.	No restr.
Reach. obj.	No restr.	RM	No restr.	No restr.
Büchi obj.	No restr.	RM	No restr.	aBM
co-Büchi obj.	RM	coBM	No restr.	No restr.

If game forms satisfy the properties below, then positional strategies exist and can be chosen positional

If game forms at states not in target are coBM and in targets are RM, then optimal strategies exist and can be chosen positional

How to ensure positional (almost-)optimal strategies?

Existence of positional optimal or ε-optimal strategies under the following restrictions on game forms:

	Positional opt. strat.		Positional almost-opt. strat.	
	Target	Not target	Target	Not target
Safety obj.	No restr.	No restr.	No restr.	No restr.
Reach. obj.	No restr.	RM	No restr.	No restr.
Büchi obj.	No restr.	RM	No restr.	aBM
co-Büchi obj.	RM	coBM	No restr.	No restr.

If game forms satisfy the properties below, then positional strategies exist and can be chosen positional

If game forms at states not in target are coBM and in targets are RM, then optimal strategies exist and can be chosen positional

If game forms at states not in target are aBM then ε-optimal strategies can be chosen positional

How to ensure positional (almost-)optimal strategies?

Existence of positional optimal or ε-optimal strategies under the following restrictions on game forms:

	Positional opt. strat.		Positional almost-opt. strat.	
	Target	Not target	Target	Not target
Safety obj.	No restr.	No restr.	No restr.	No restr.
Reach. obj.	No restr.	RM	No restr.	No restr.
Büchi obj.	No restr.	RM	No restr.	aBM
co-Büchi obj.	RM	coBM	No restr.	No restr.

If game forms satisfy the properties below, then positional strategies exist and can be chosen positional

If game forms at states not in target are coBM and in targets are RM, then optimal strategies exist and can be chosen positional

If game forms at states not in target are aBM then ε-optimal strategies can be chosen positional

ε-optimal strategies can always be chosen positional

How to ensure positional (almost-)optimal strategies?

Existence of positional optimal or ε-optimal strategies under the following restrictions on game forms:

	Positional opt. strat.		Positional almost-opt. strat.	
	Target	Not target	Target	Not target
Safety obj.	No restr.	No restr.	No restr.	No restr.
Reach. obj.	No restr.	RM	No restr.	No restr.
Büchi obj.	No restr.	RM	No restr.	aBM
co-Büchi obj.	RM	coBM	No restr.	No restr.

Properties of game forms

Properties of game forms

- All these notions RM, coBM, aBM, ... can be decided (can be expressed in FO$(\mathbb{R})$)

Properties of game forms

- All these notions RM, coBM, aBM, ... can be decided (can be expressed in FO$(\mathbb{R})$)
- coBM $\subseteq$ RM $\subseteq$ aBM

Properties of game forms

- All these notions RM, coBM, aBM, ... can be decided (can be expressed in FO$(\mathbb{R})$)
- coBM ⊆ RM ⊆ aBM
- These game forms are coBM:
  - « Turn-based » game forms:

$$\begin{bmatrix} x & y & z \\ x & y & z \end{bmatrix}$$

  - Two-variable game forms:

$$\begin{bmatrix} x & y & x \\ y & x & x \end{bmatrix}$$

  - Permutation game forms:

$$\begin{bmatrix} x & y & z \\ z & x & y \\ y & z & x \end{bmatrix}$$

What you can bring home

What you can bring home

- **Concurrent games** behave much less smoothly that turn-based games
  - Optimal strategies might not exist
  - (Almost-)Optimal strategies might require infinite memory

What you can bring home

- **Concurrent games** behave much less smoothly that turn-based games
  - Optimal strategies might not exist
  - (Almost-)Optimal strategies might require infinite memory
- Methodology:
  - Study interactions (**game forms**) as first-class citizens

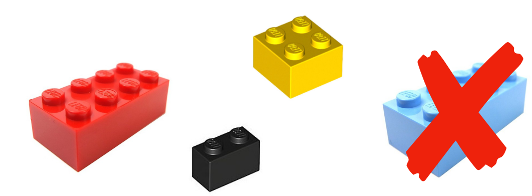

  - Identify interactions (game forms) that are well-behaved (with a property in mind)
  - Show that, all games on graphs with interactions taken in the set of well-behaved game forms behave well; and that this set is maximal

What you can bring home

- **Concurrent games** behave much less smoothly that turn-based games
  - Optimal strategies might not exist
  - (Almost-)Optimal strategies might require infinite memory
- Methodology:
  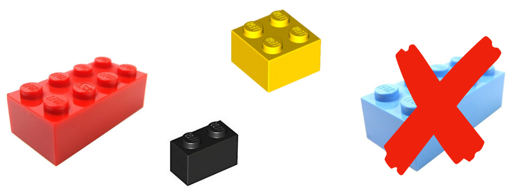
  - Study interactions (**game forms**) as first-class citizens
  - Identify interactions (game forms) that are well-behaved (with a property in mind)
  - Show that, all games on graphs with interactions taken in the set of well-behaved game forms behave well; and that this set is maximal
  
- Going further:
  - Understand beyond (co-)Büchi conditions, e.g. parity conditions
  - (Ongoing work) A different approach, which should be able to deal with parity conditions