Information-Flow Patterns in Games with Imperfect Information

Marie van den Bogaard

LSV, ENS Paris-Saclay

Tuesday, November 29th 2016

Interactive scenarios

Interactive scenarios

Ingredients:

- active agents
- common goal
- partial view
- limited communication
- unpredictable environment

Interactive scenarios

Ingredients:

- active agents
- common goal
- partial view
- limited communication
- unpredictable environment

$\rightsquigarrow$ cooperation is difficult, make it work with computational agents?

Interactive scenarios

Goals:

- *describe*: formal model
- *prescribe*: distributed solutions

$\rightsquigarrow$ Games (with Imperfect Information)!

Games (with perfect information)

Games (with perfect information)

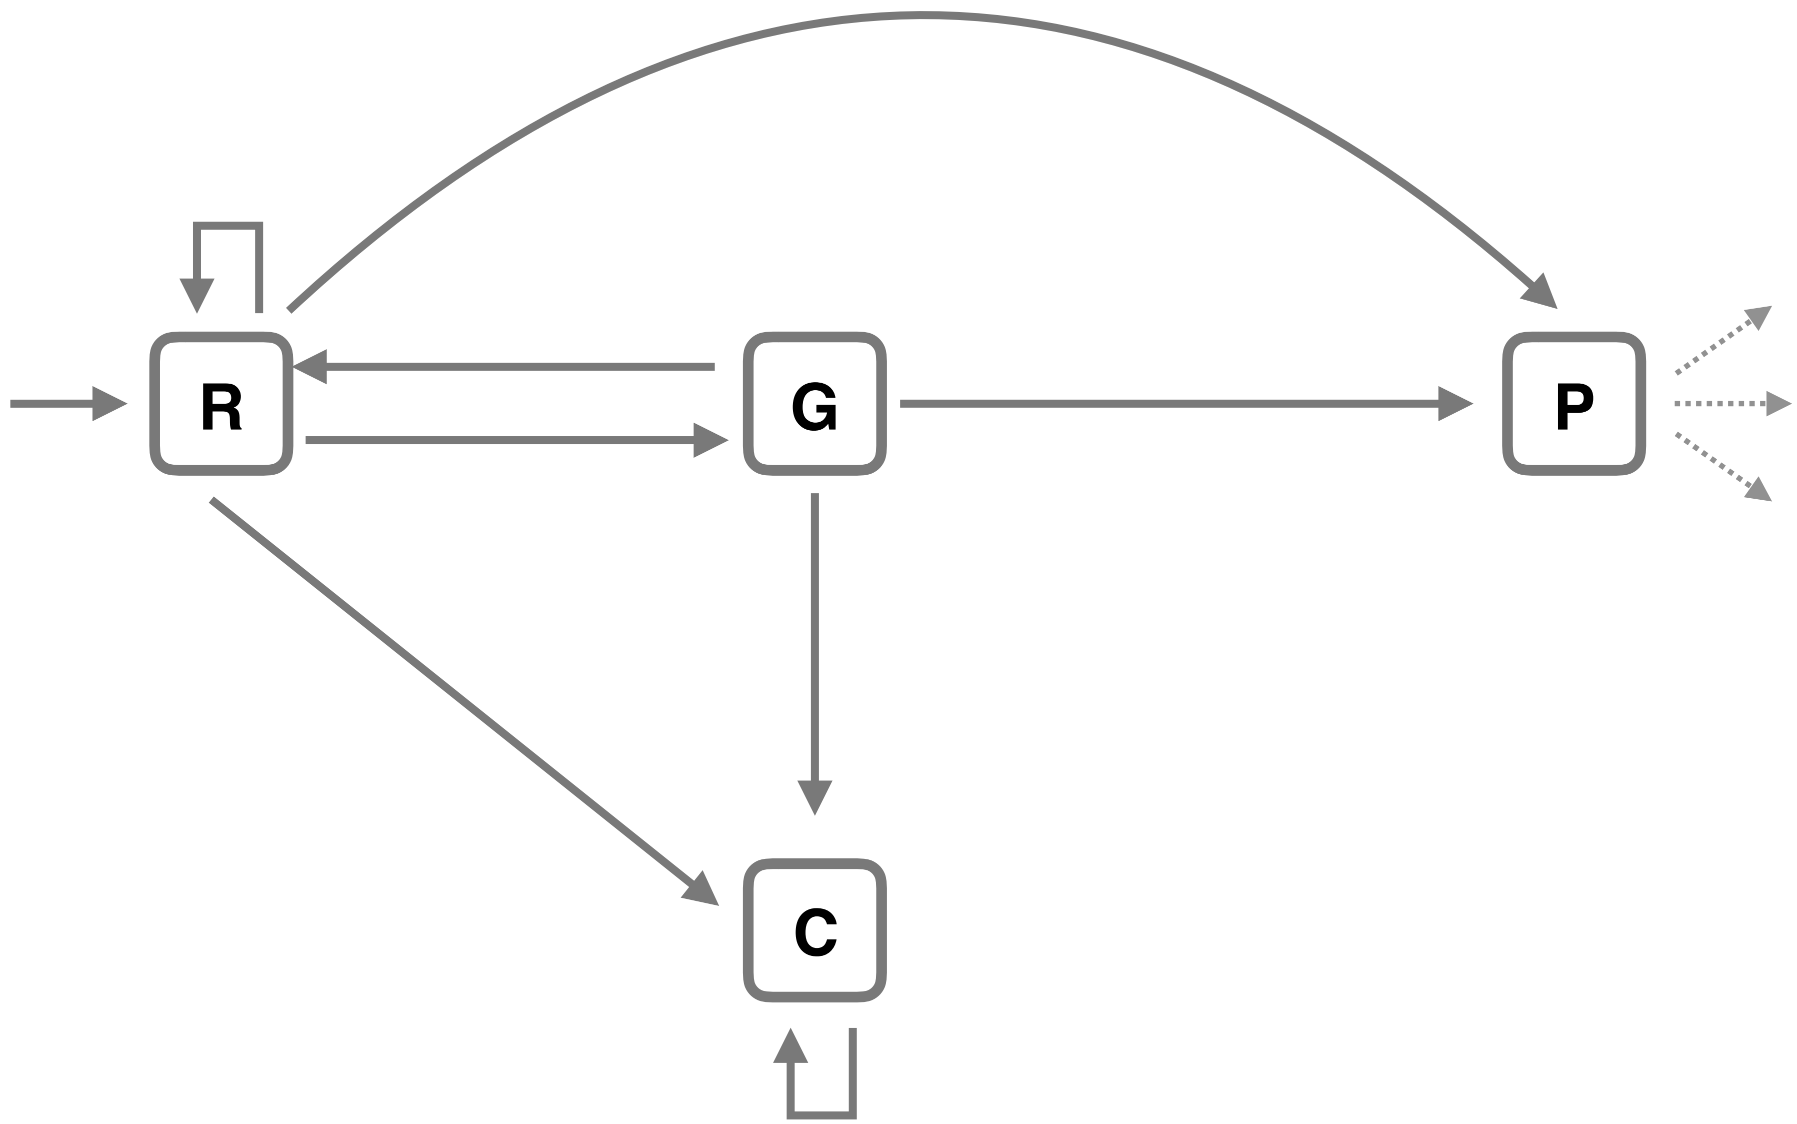

Games (with perfect information)

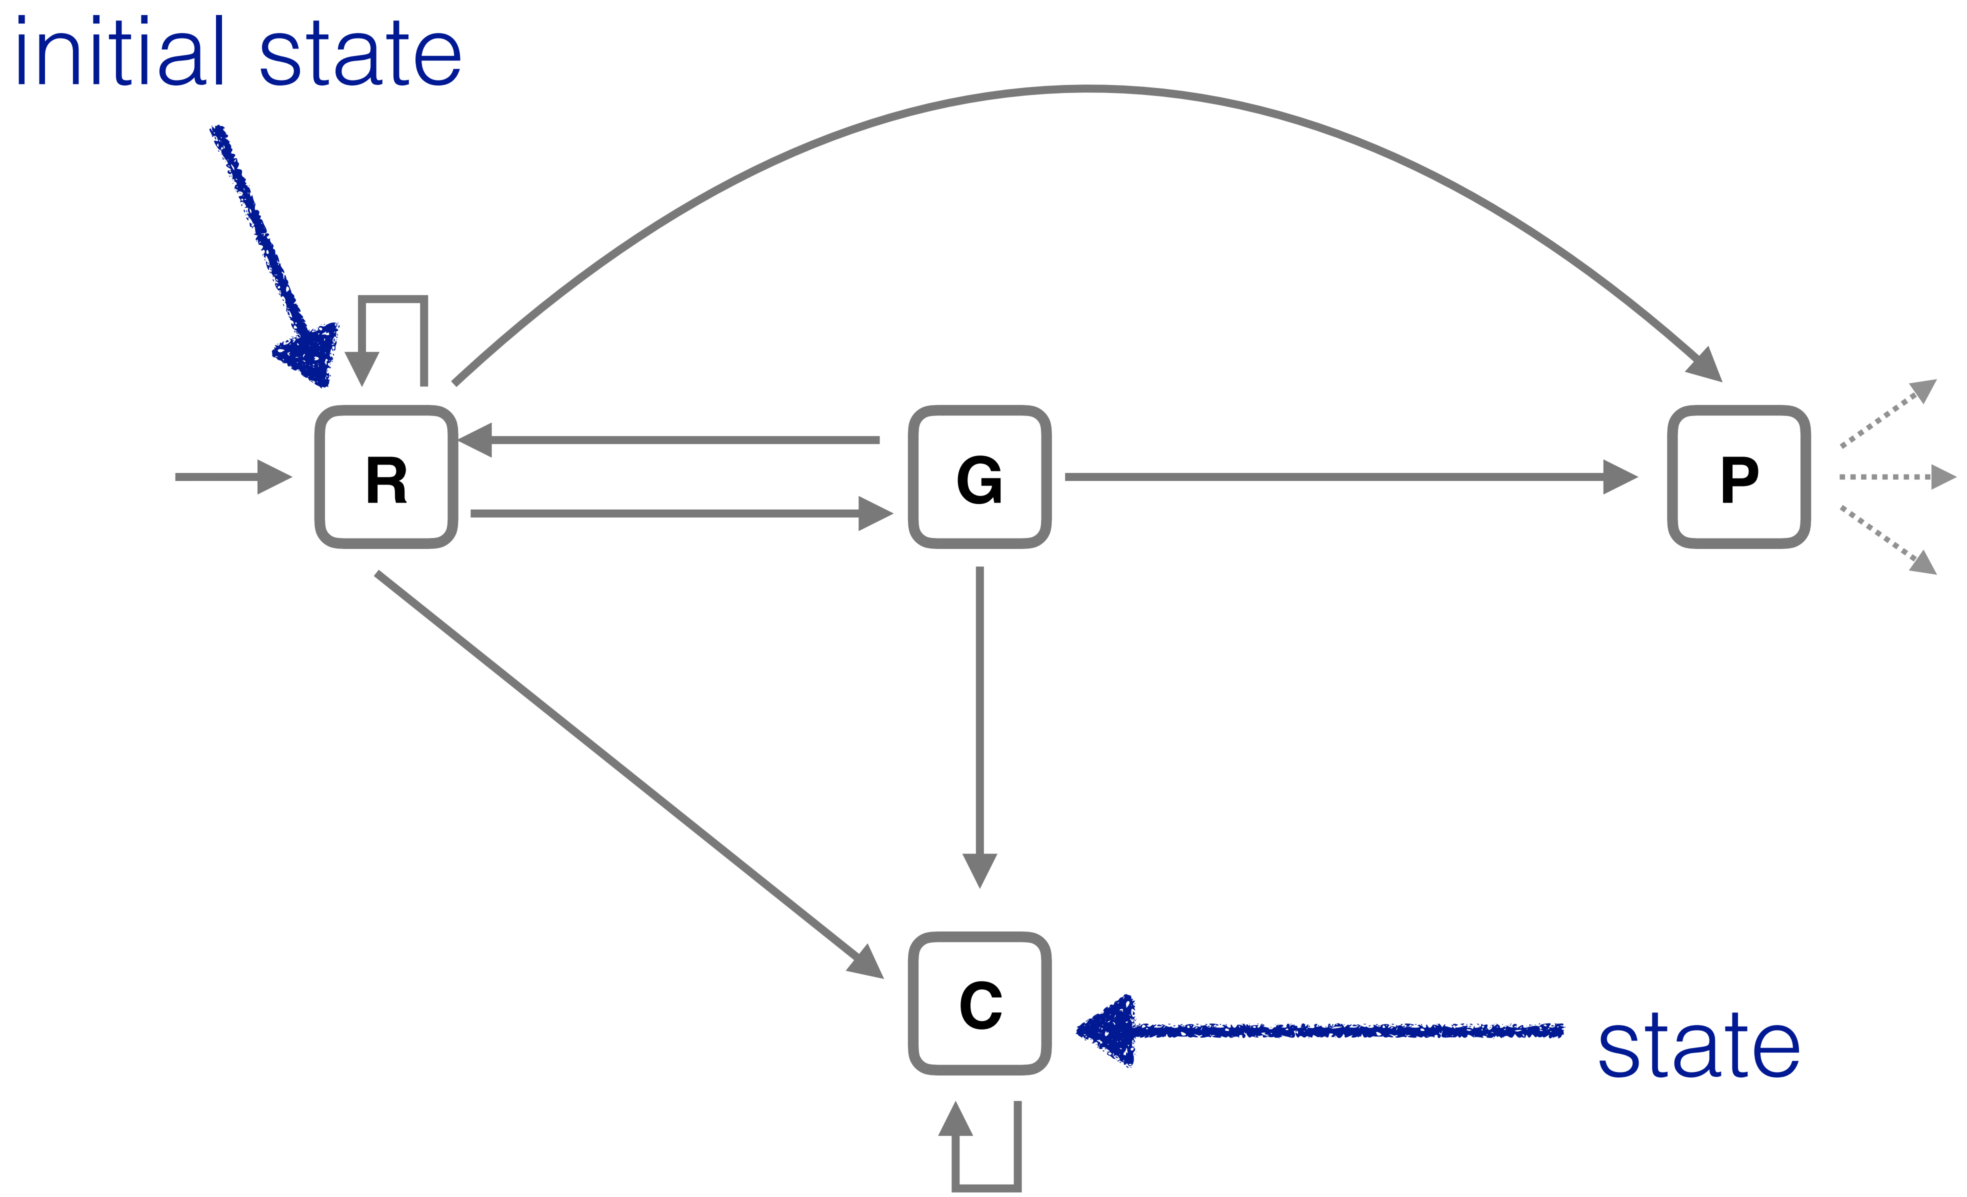

transition

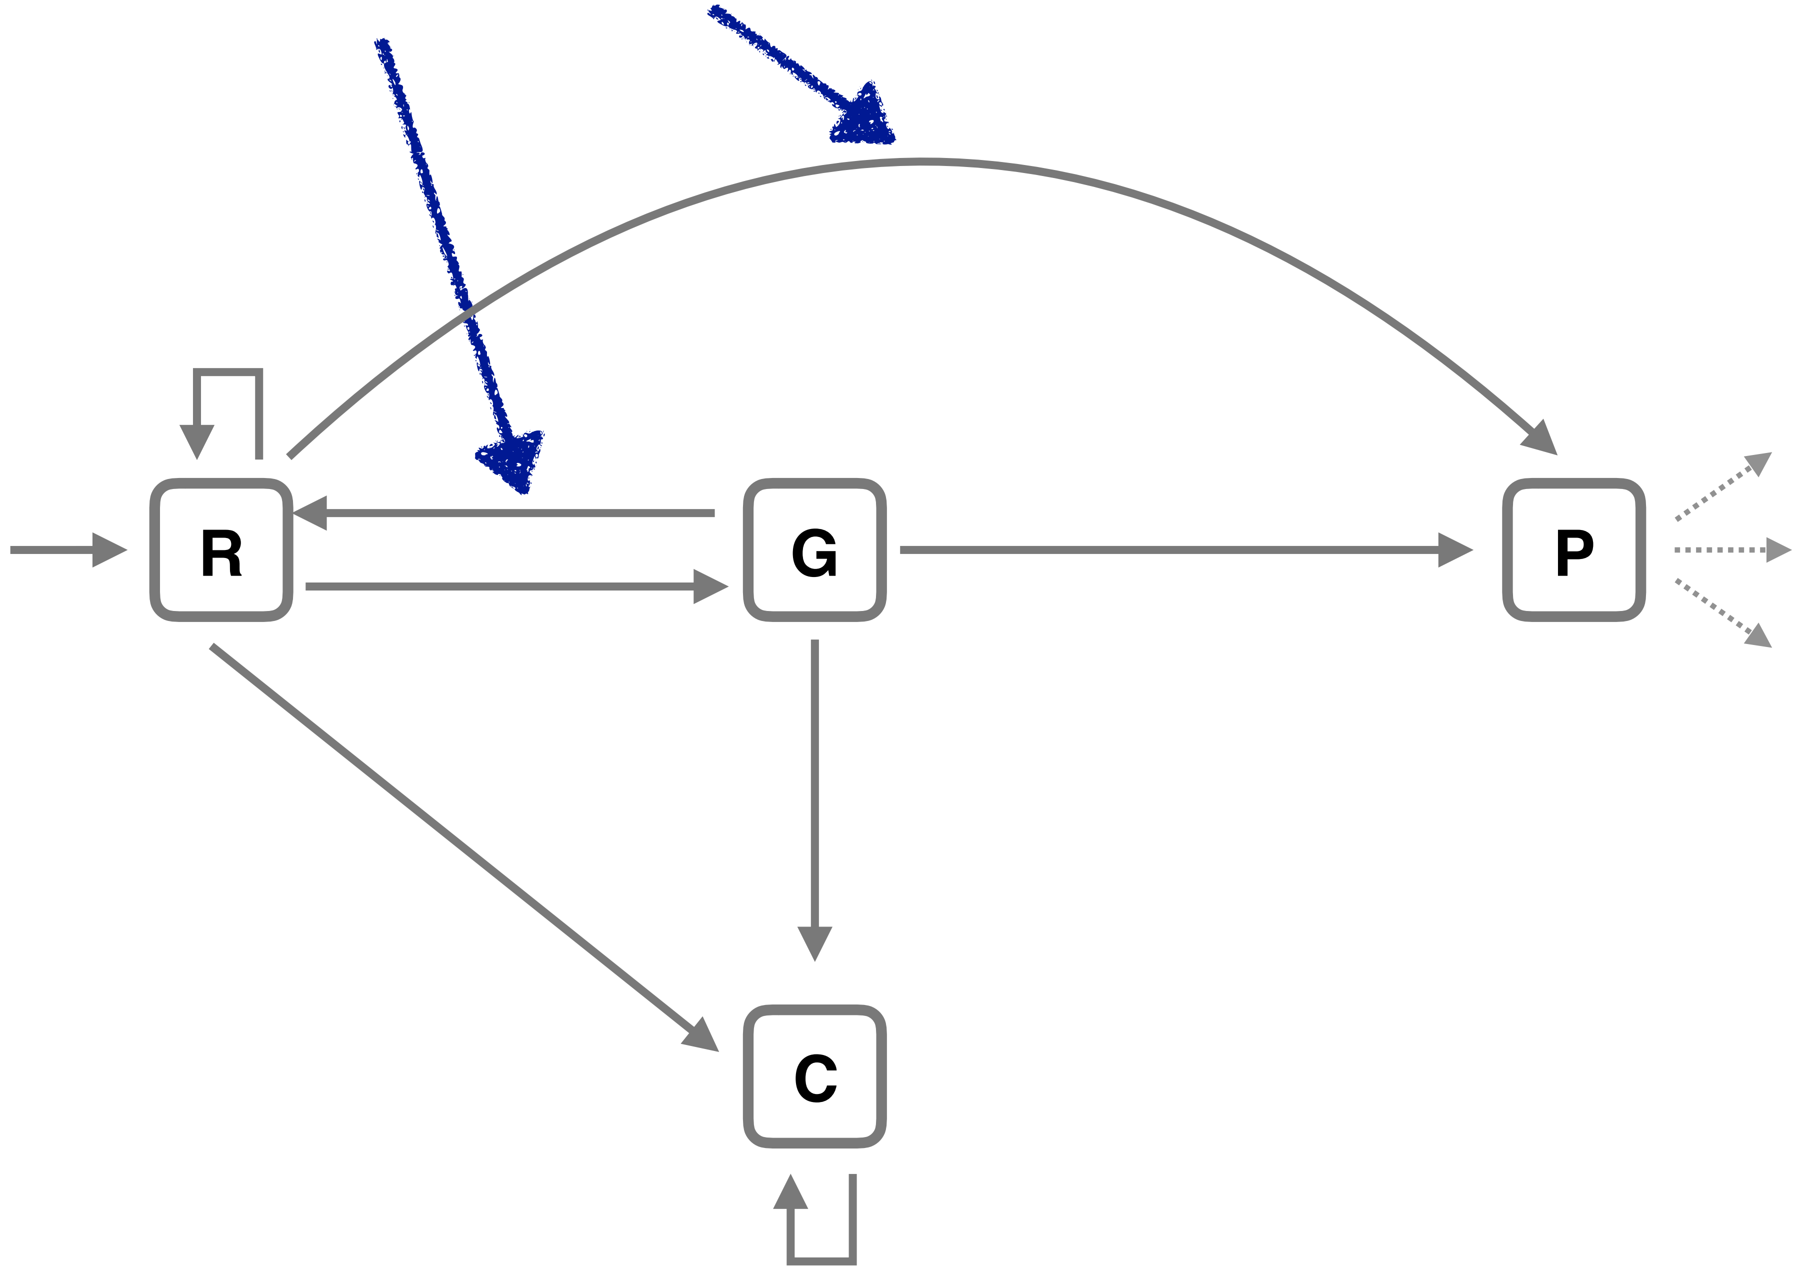

Two players against Nature

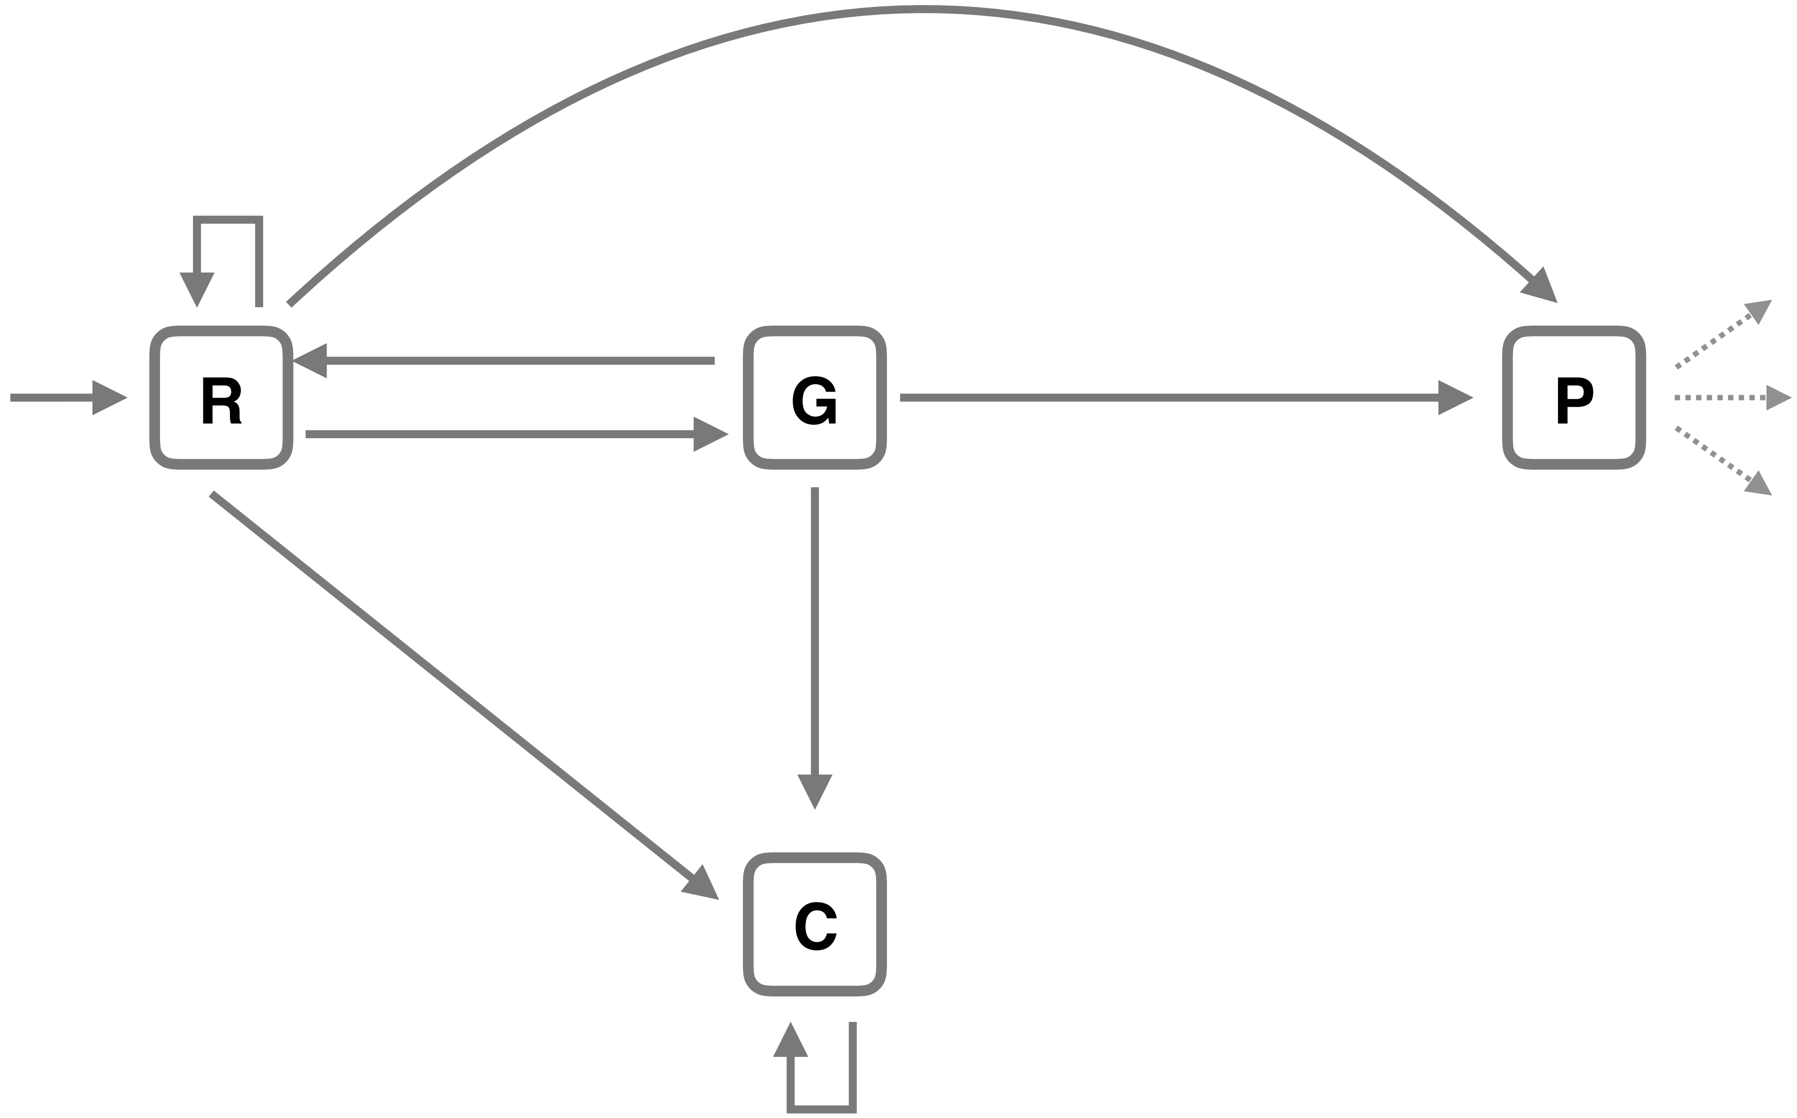

Action sets: $A^1 = \{ W, F \}$ $A^2 = \{ W, F \}$

action profile

Two players against Nature

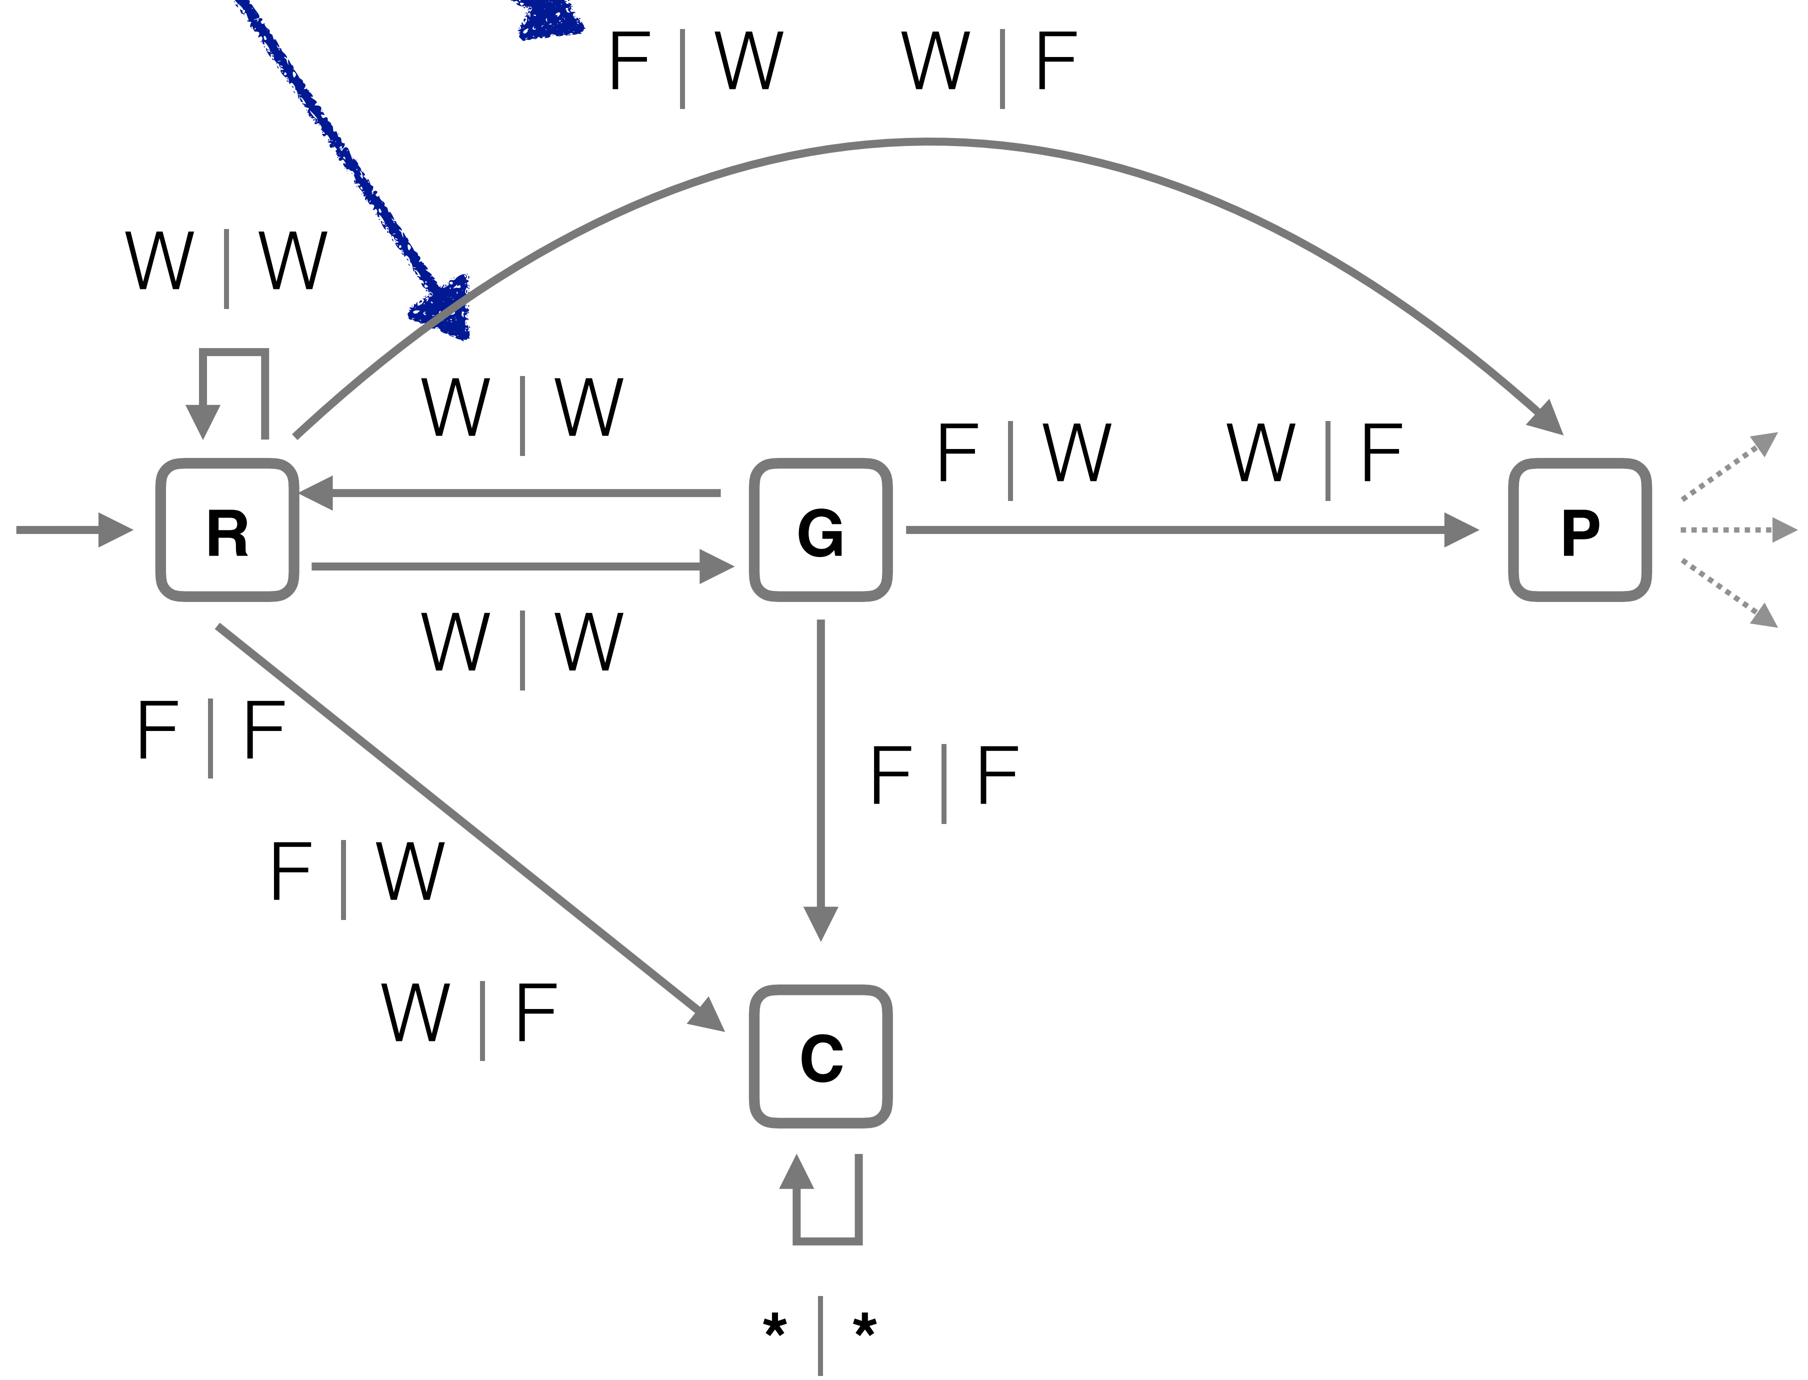

Action sets: $A^1 = \{ W, F \}$ $A^2 = \{ W, F \}$

Synchronous!

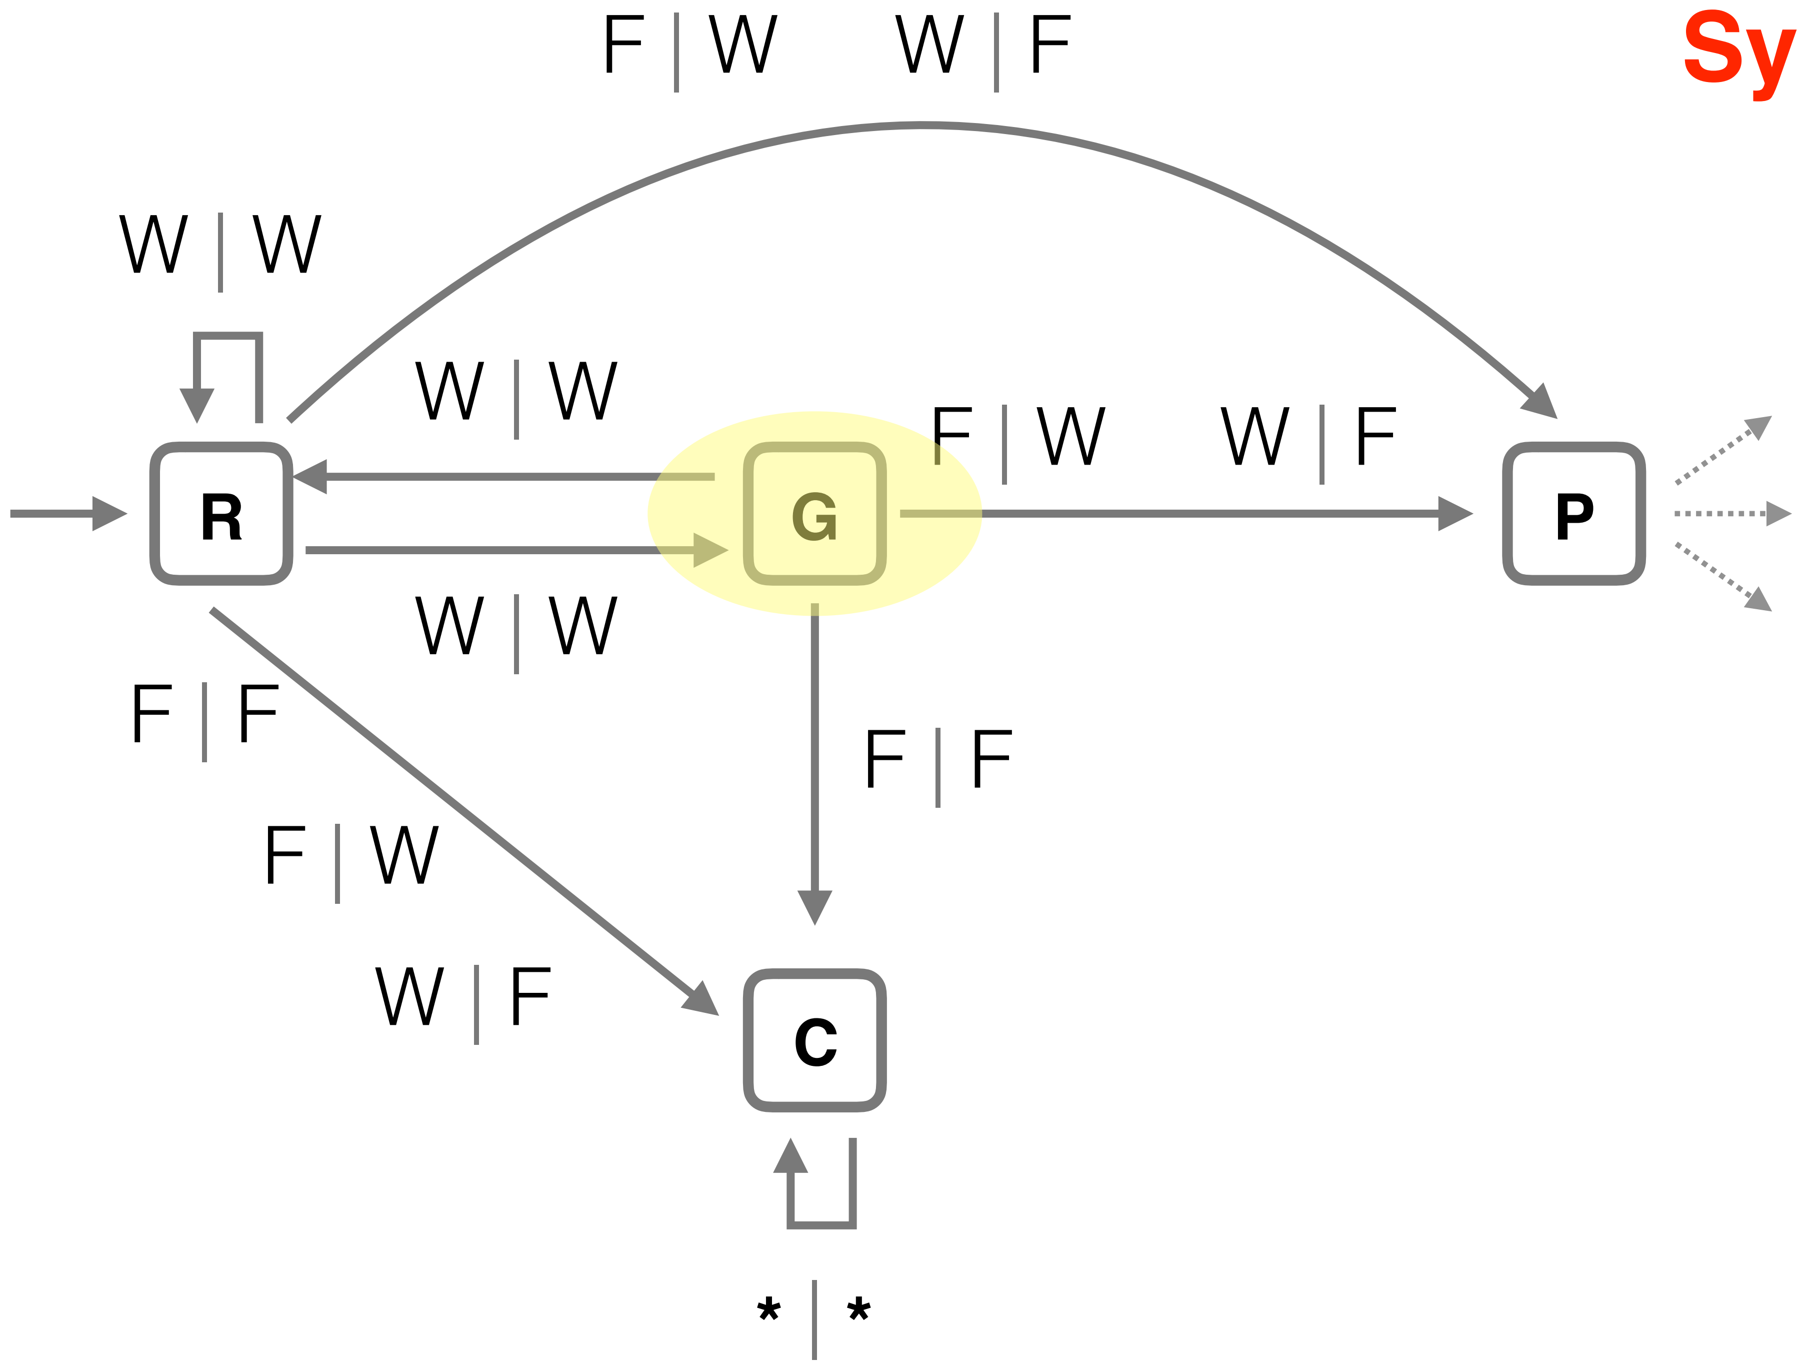

Action sets: $A^1 = \{ W, F \}$ $A^2 = \{ W, F \}$

Synchronous!

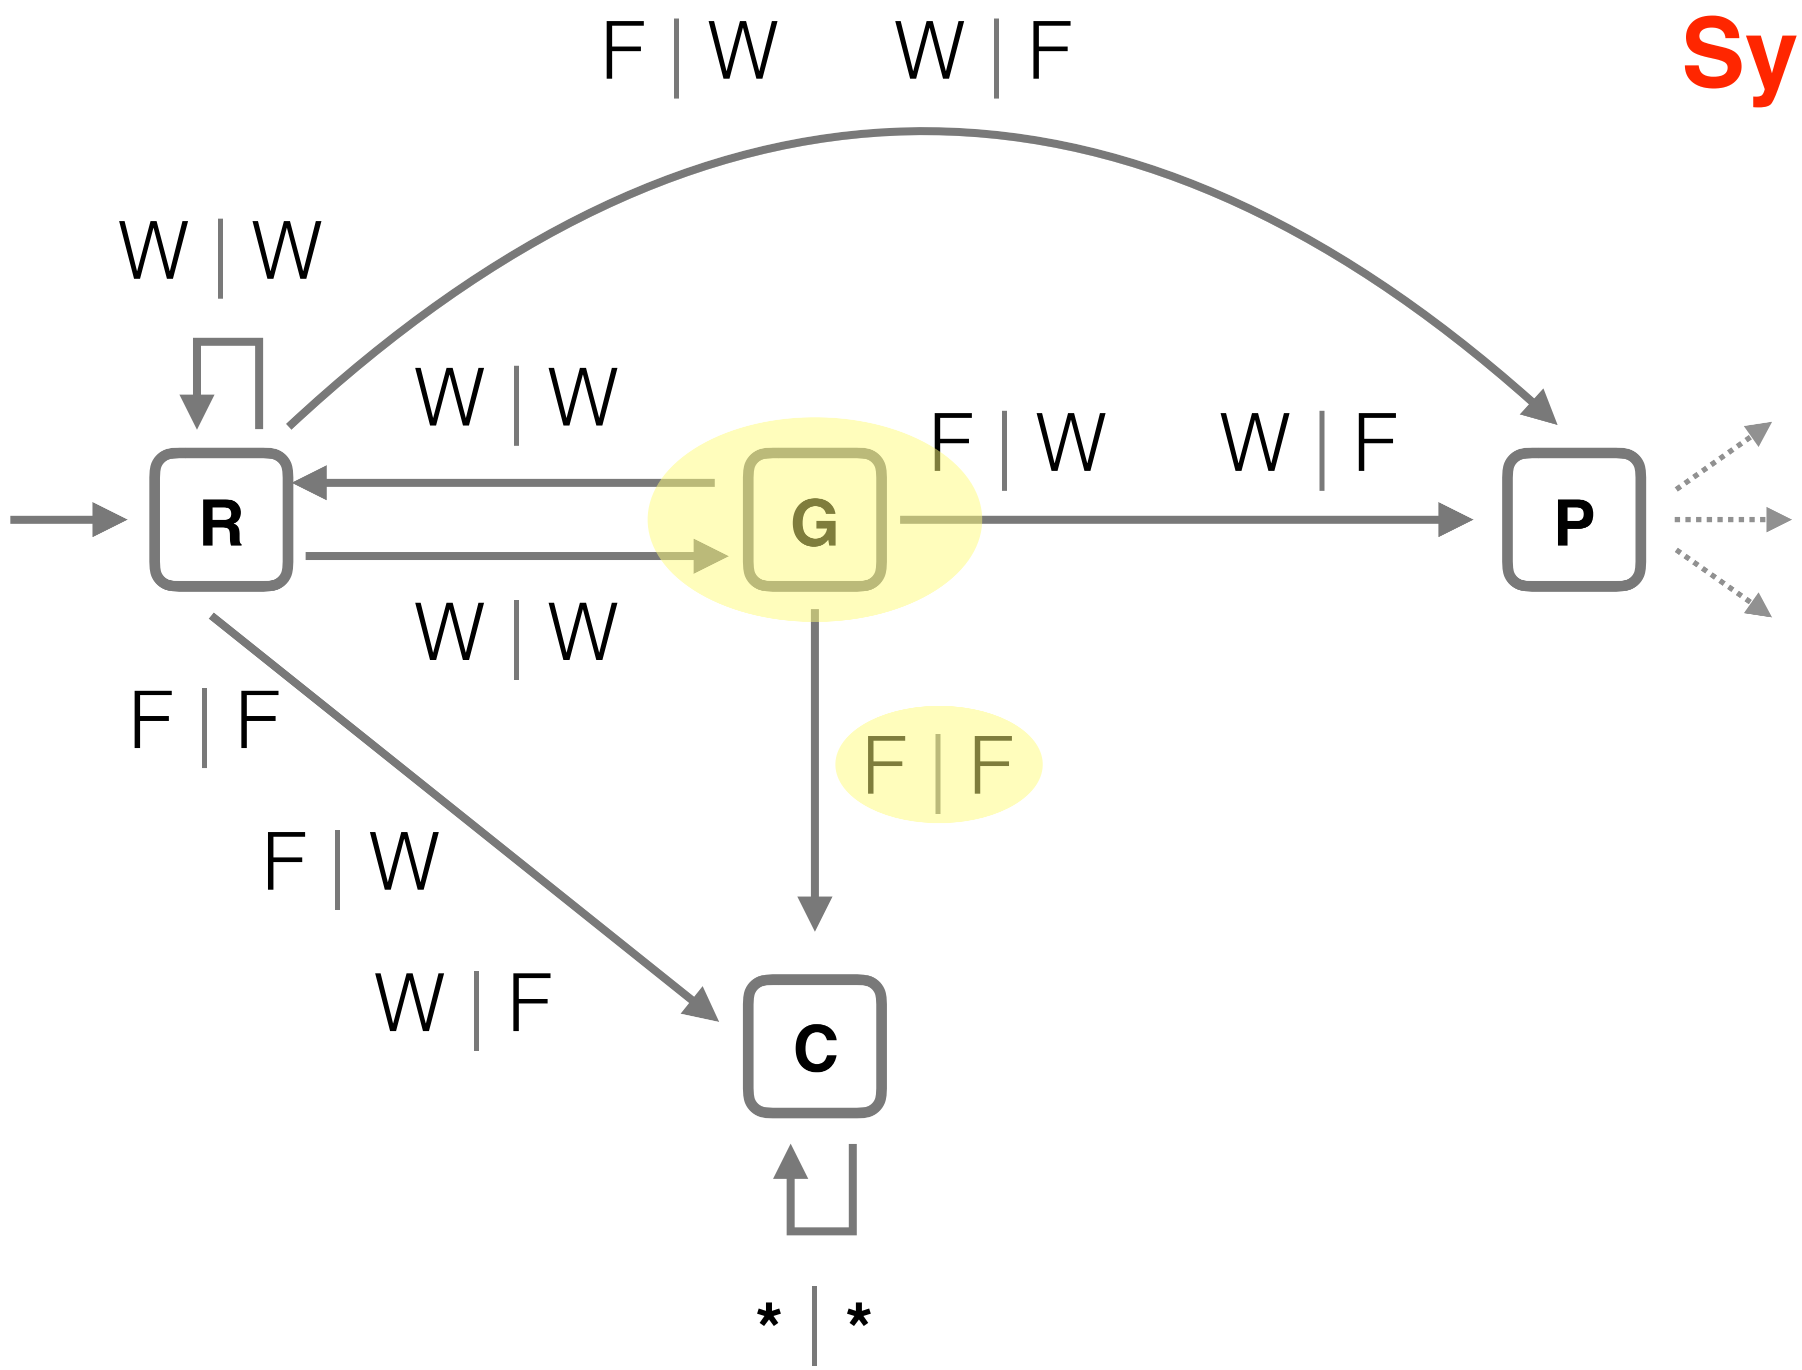

Action sets: $A^1 = \{ W, F \}$ $A^2 = \{ W, F \}$

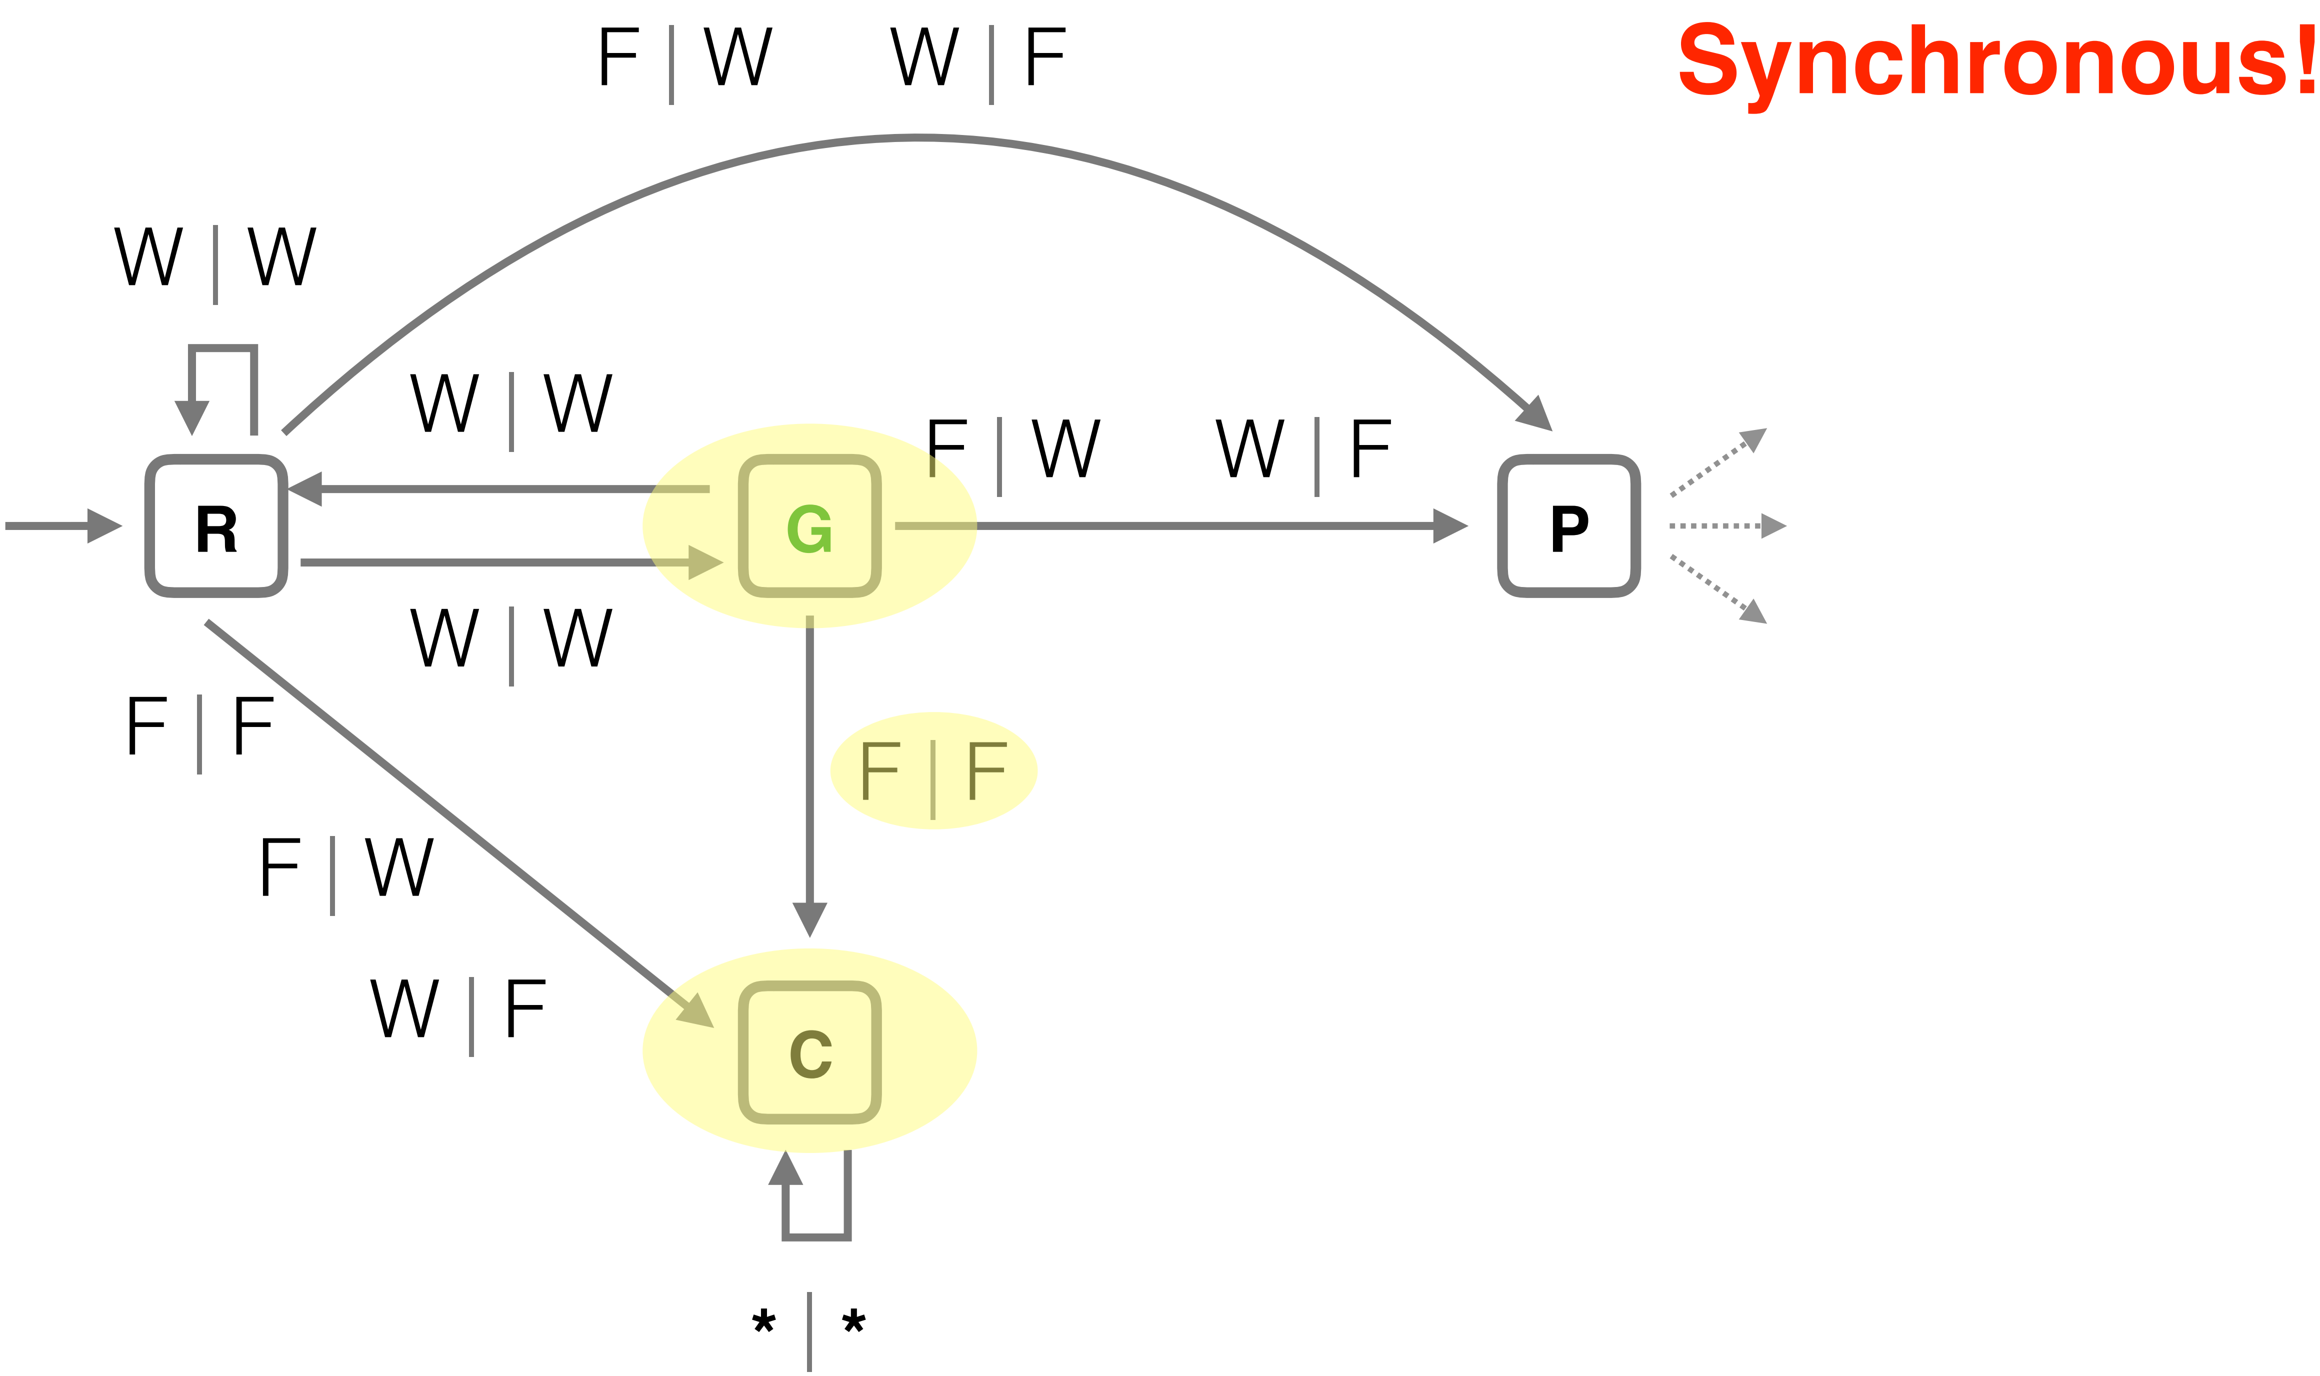

Action sets: $A^1 = \{ W, F \}$ $A^2 = \{ W, F \}$

Synchronous!

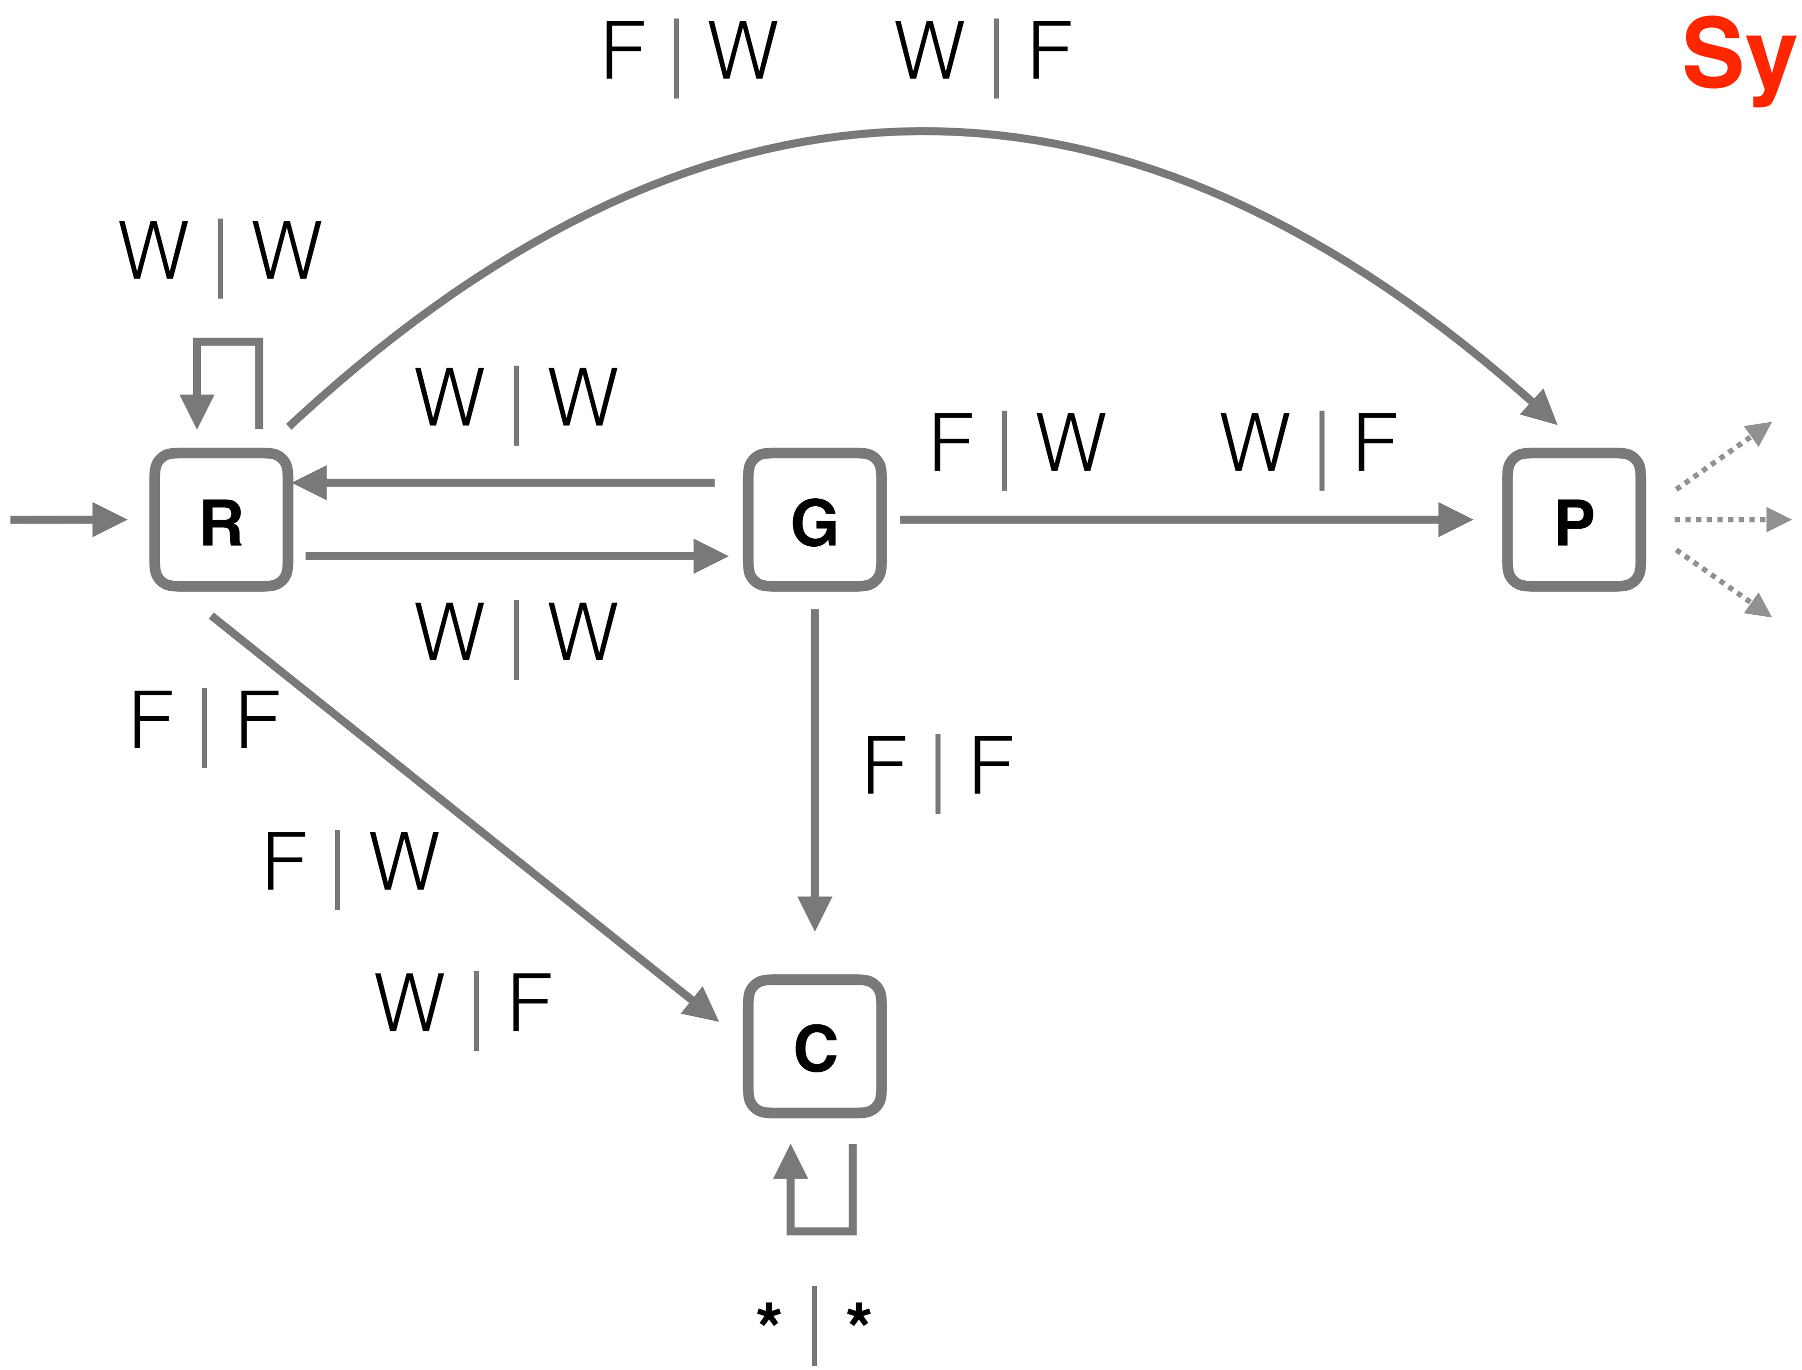

Action sets: A^1={ W, F } A^2={ W, F }

Nature
(nondeterminism)

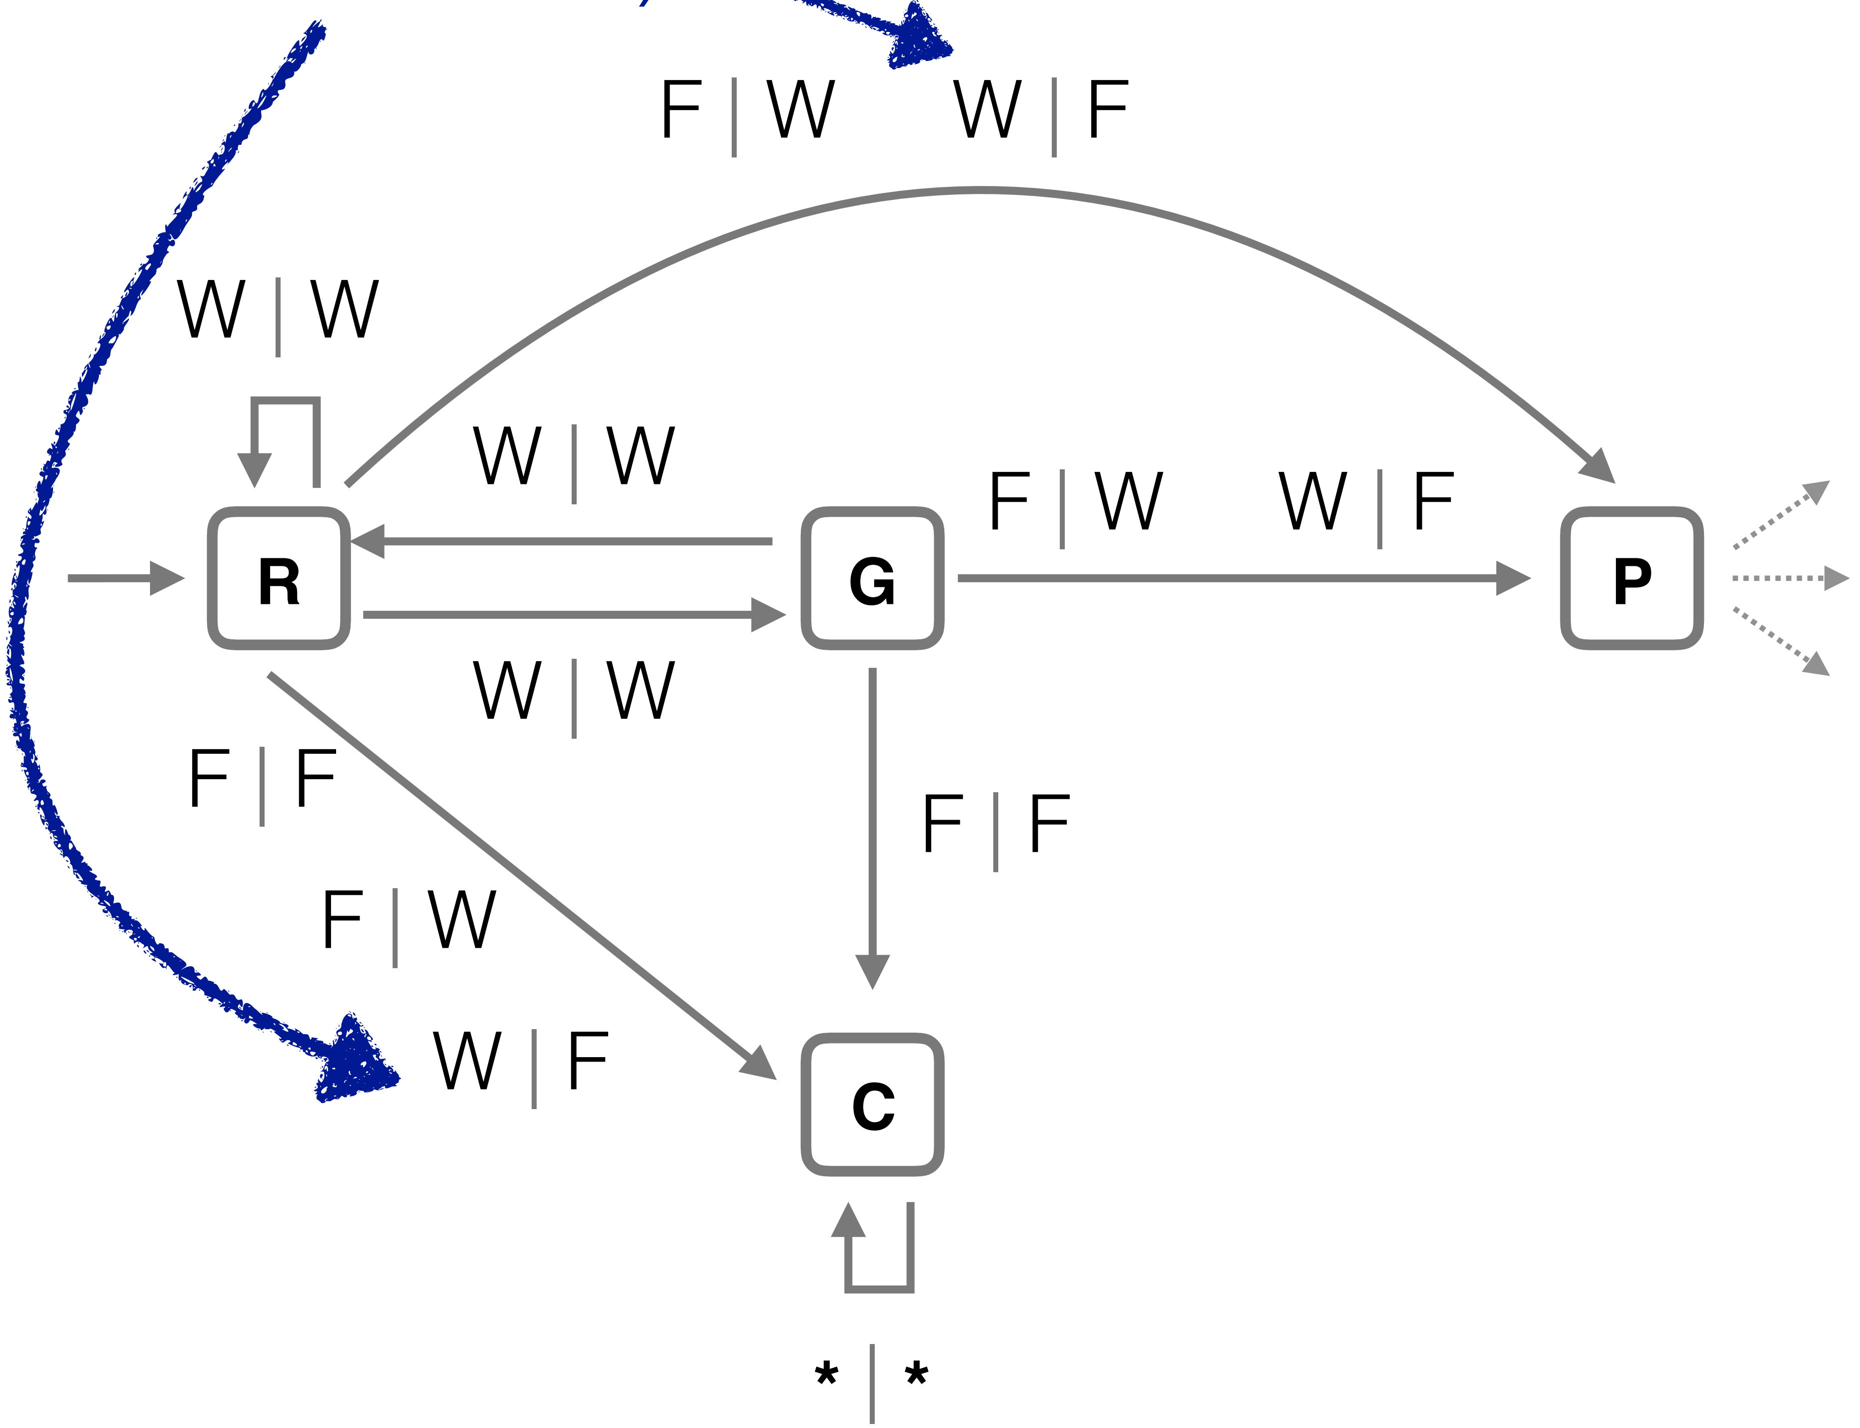

Action sets: $A^1 = \{W, F\}$ $A^2 = \{W, F\}$

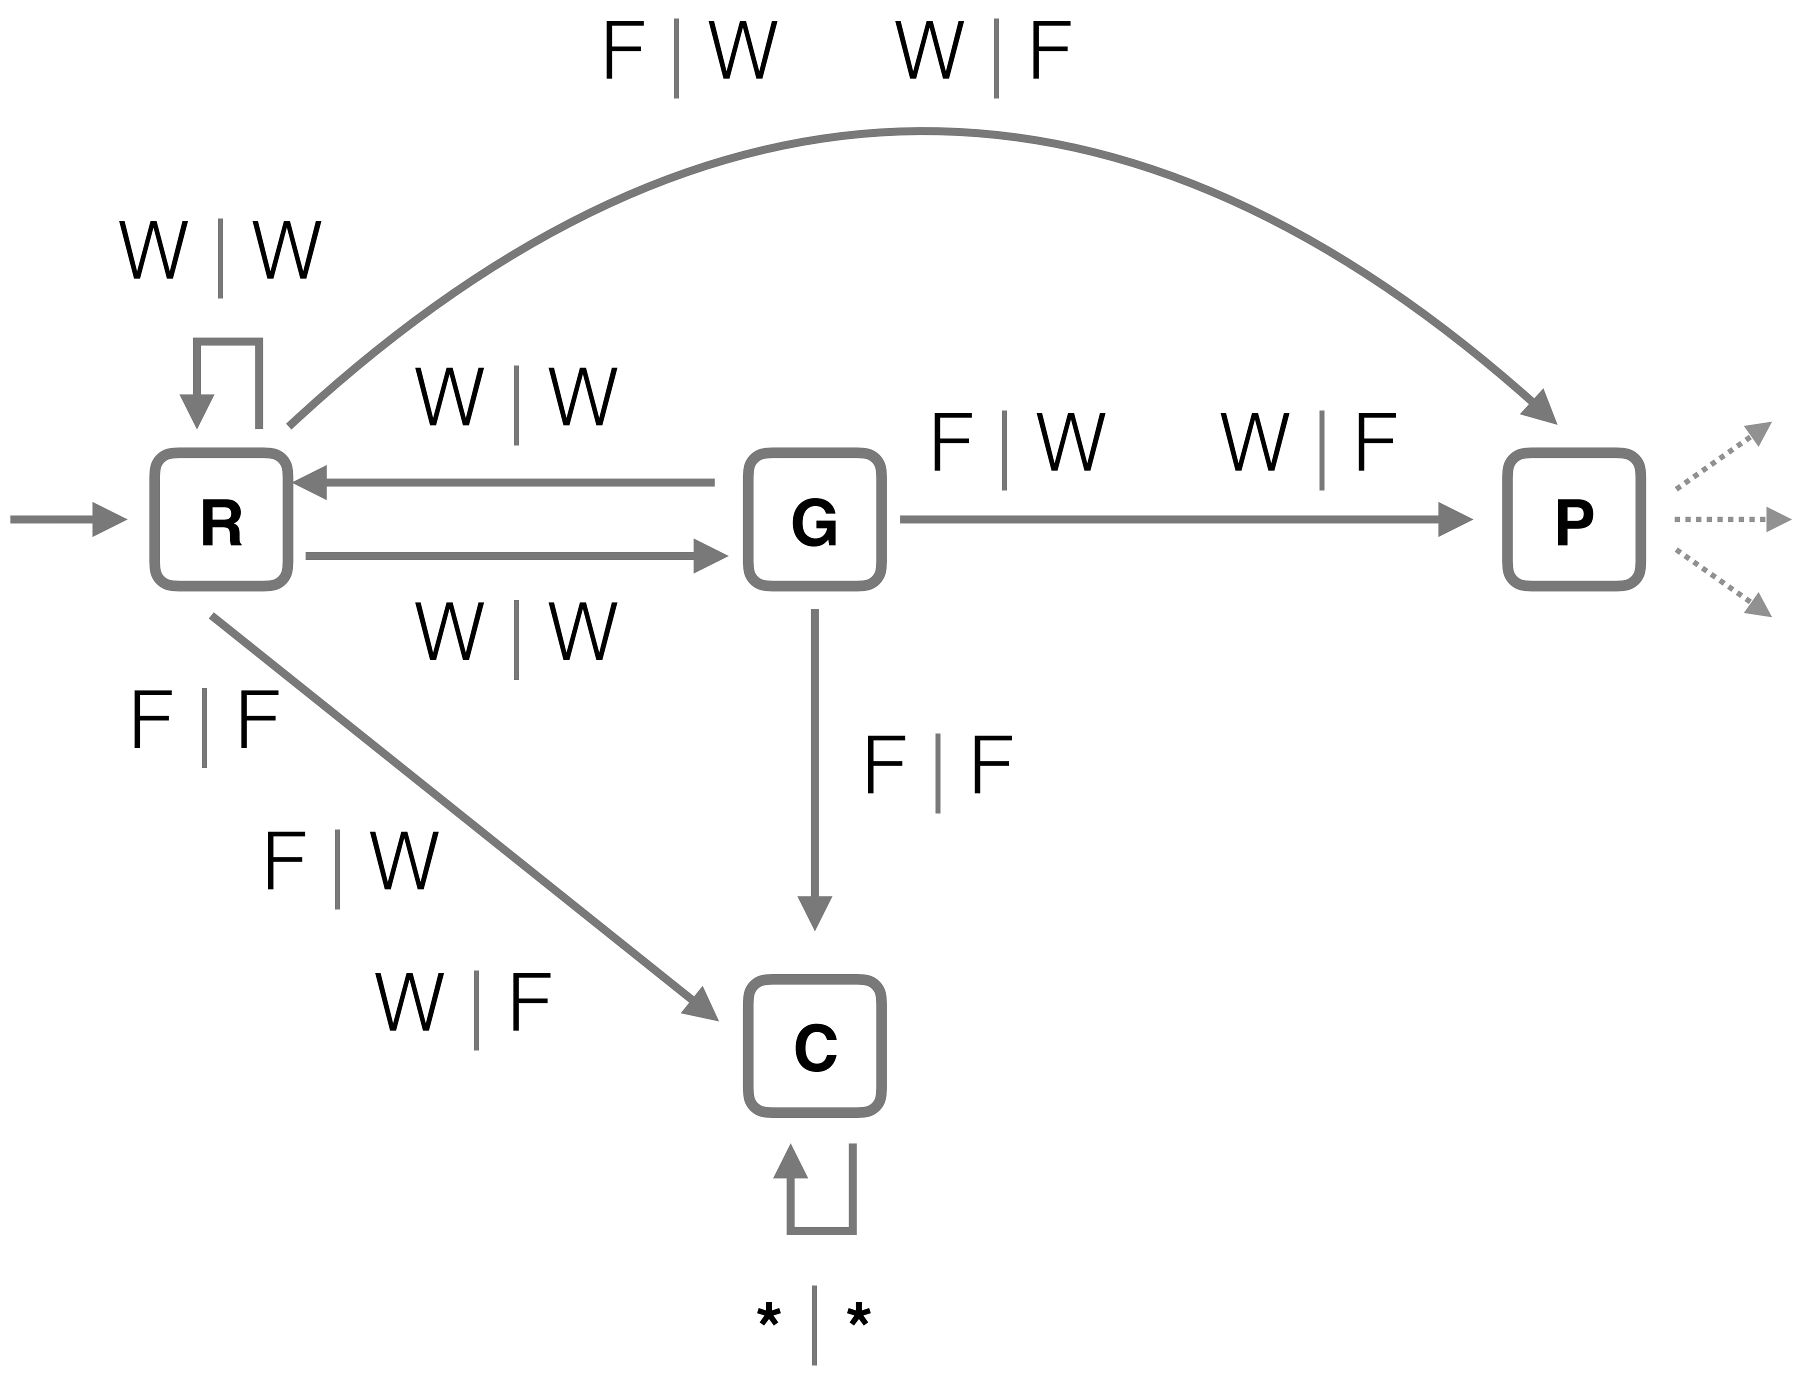

Action sets: A^1={ W, F } A^2={ W, F }

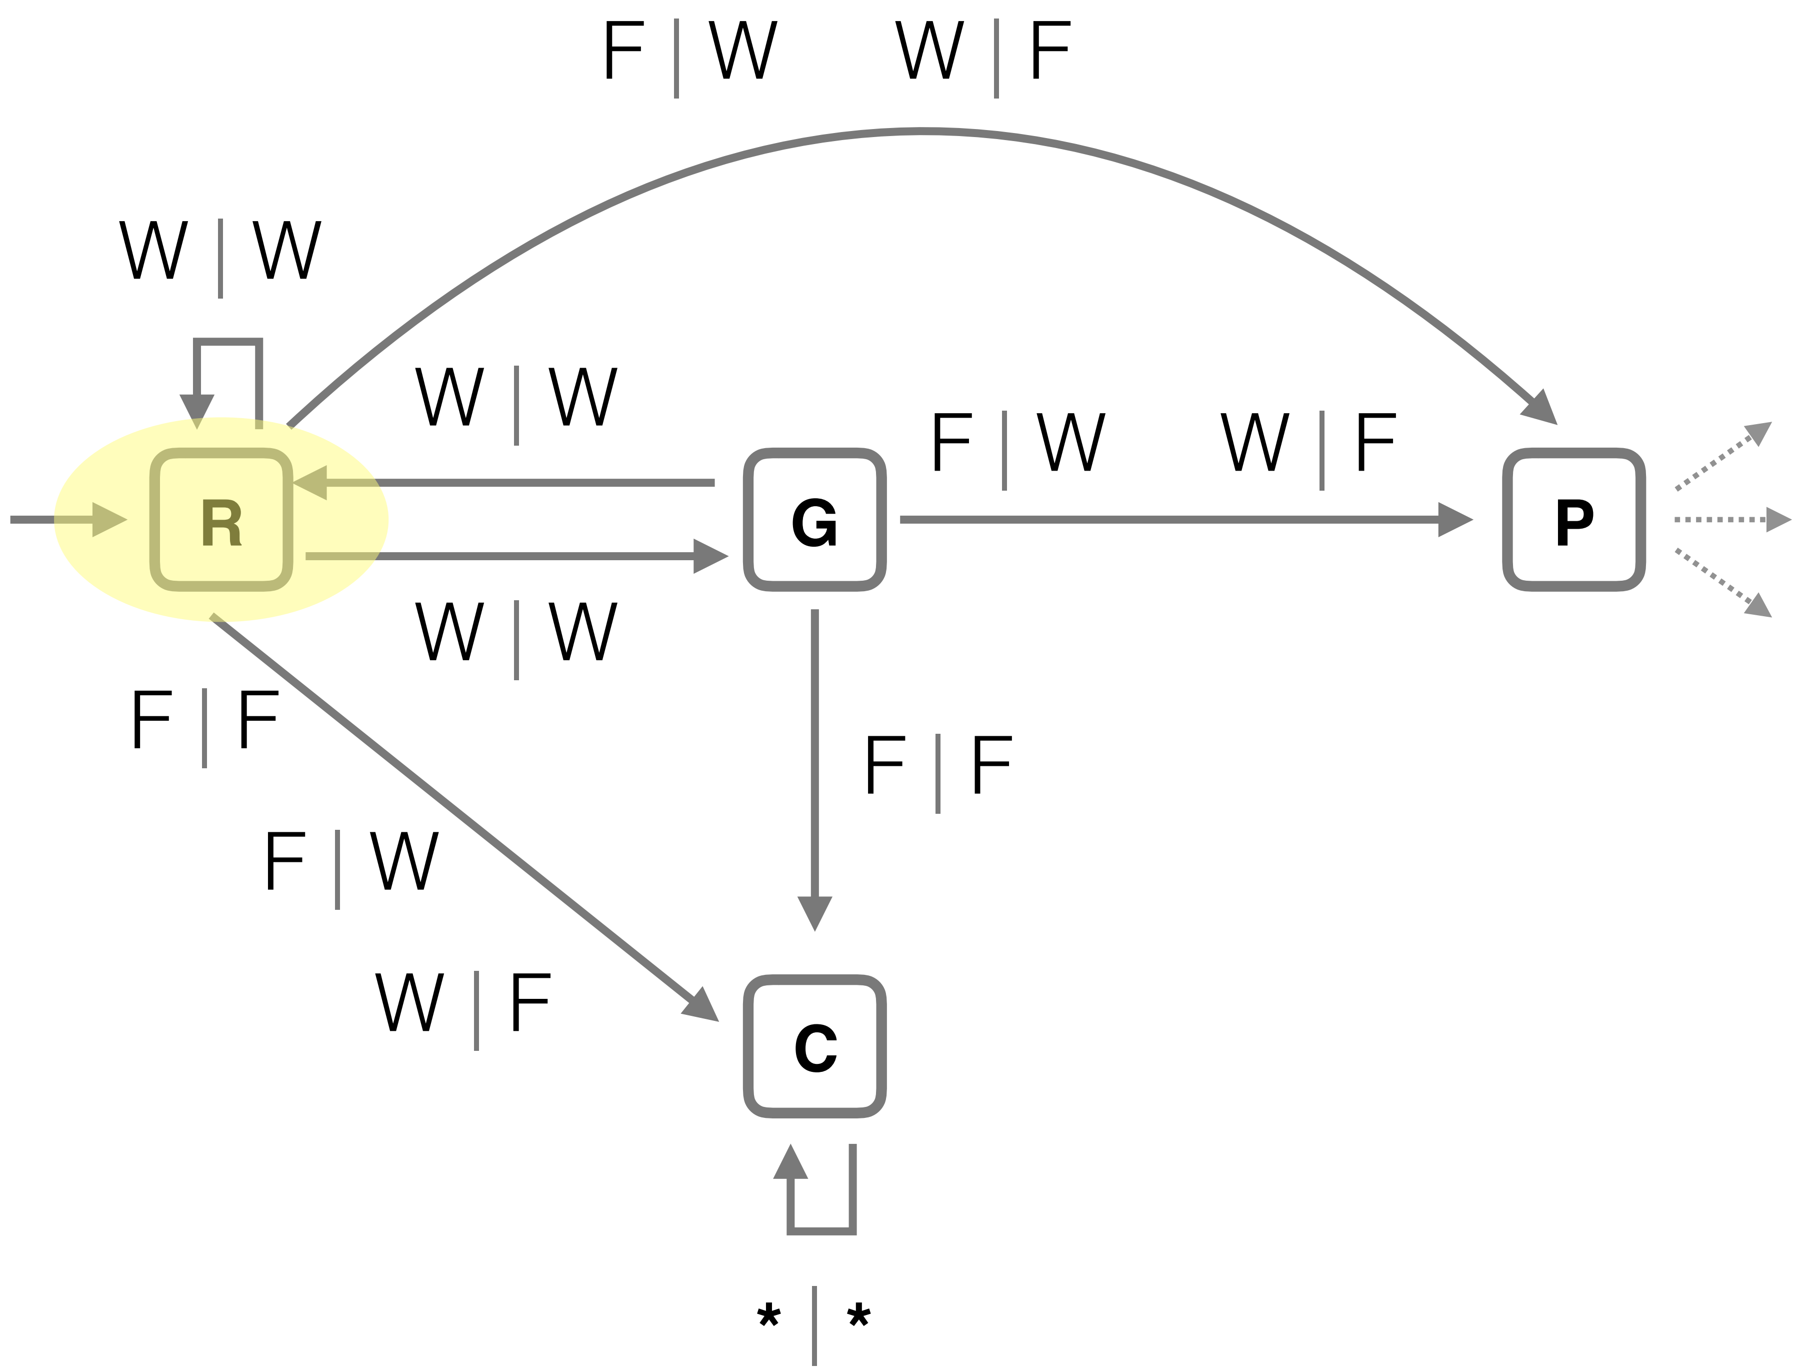

Play: **R**

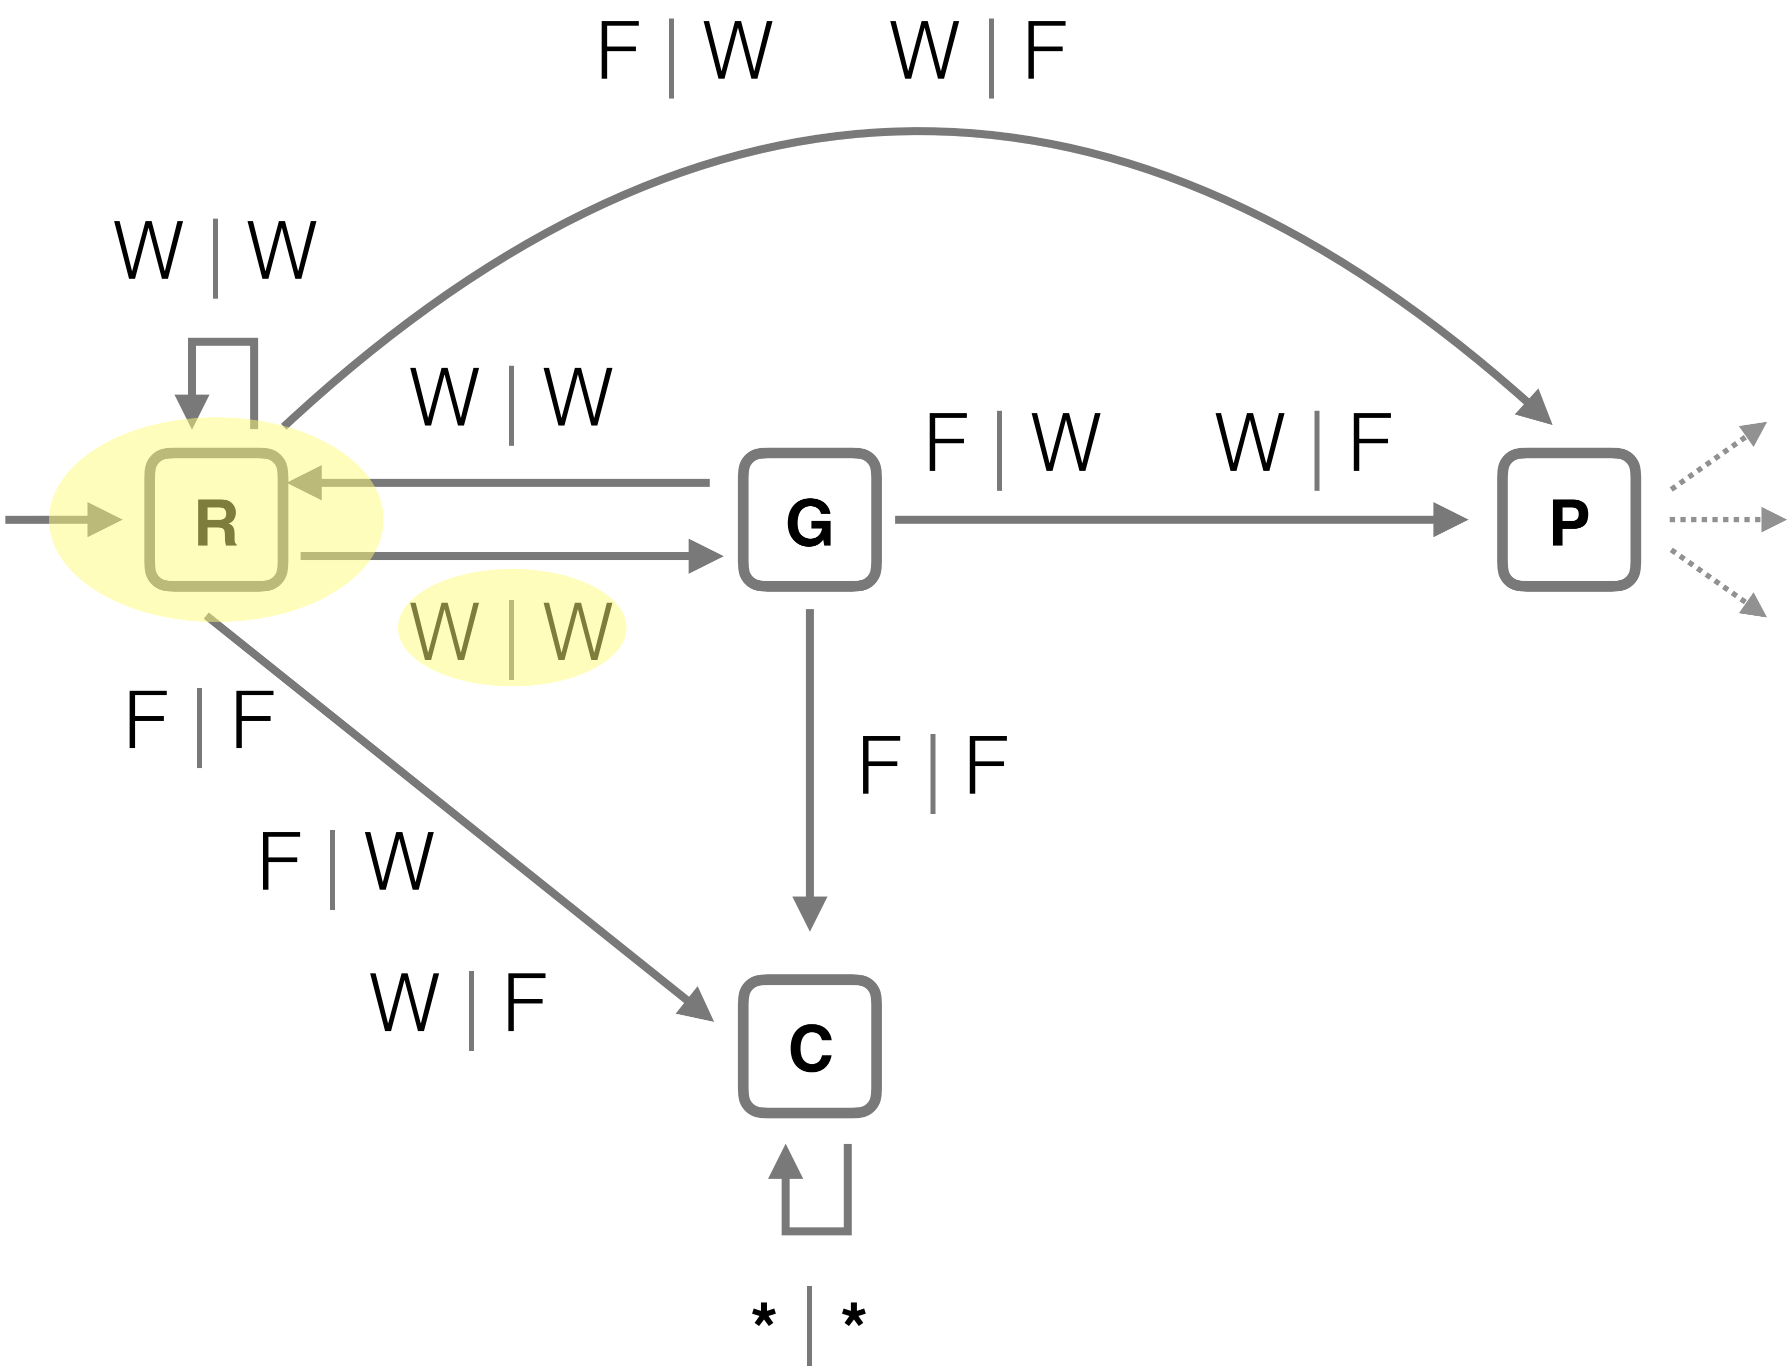

Play: **R**

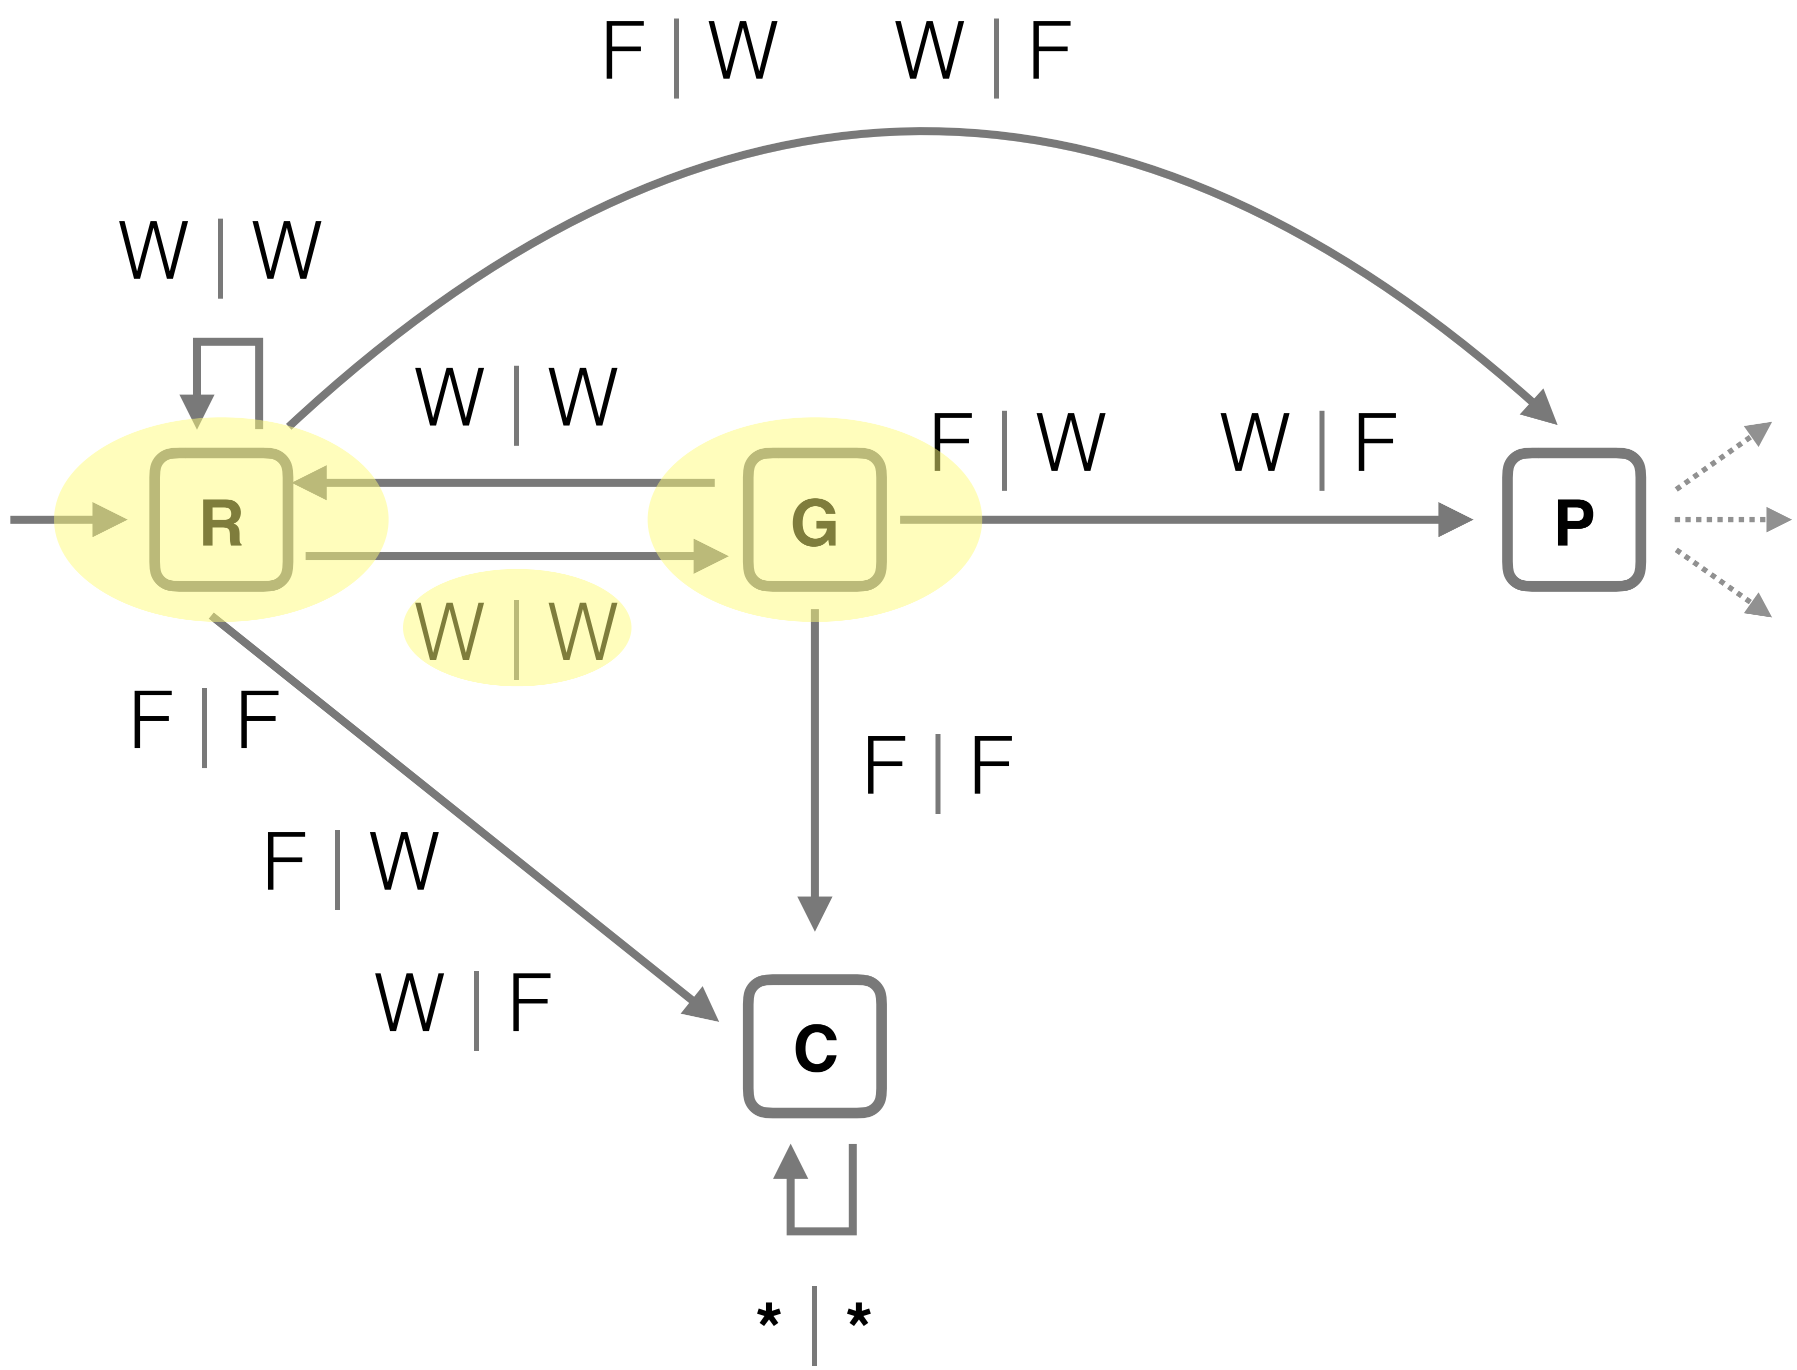

Play: R G

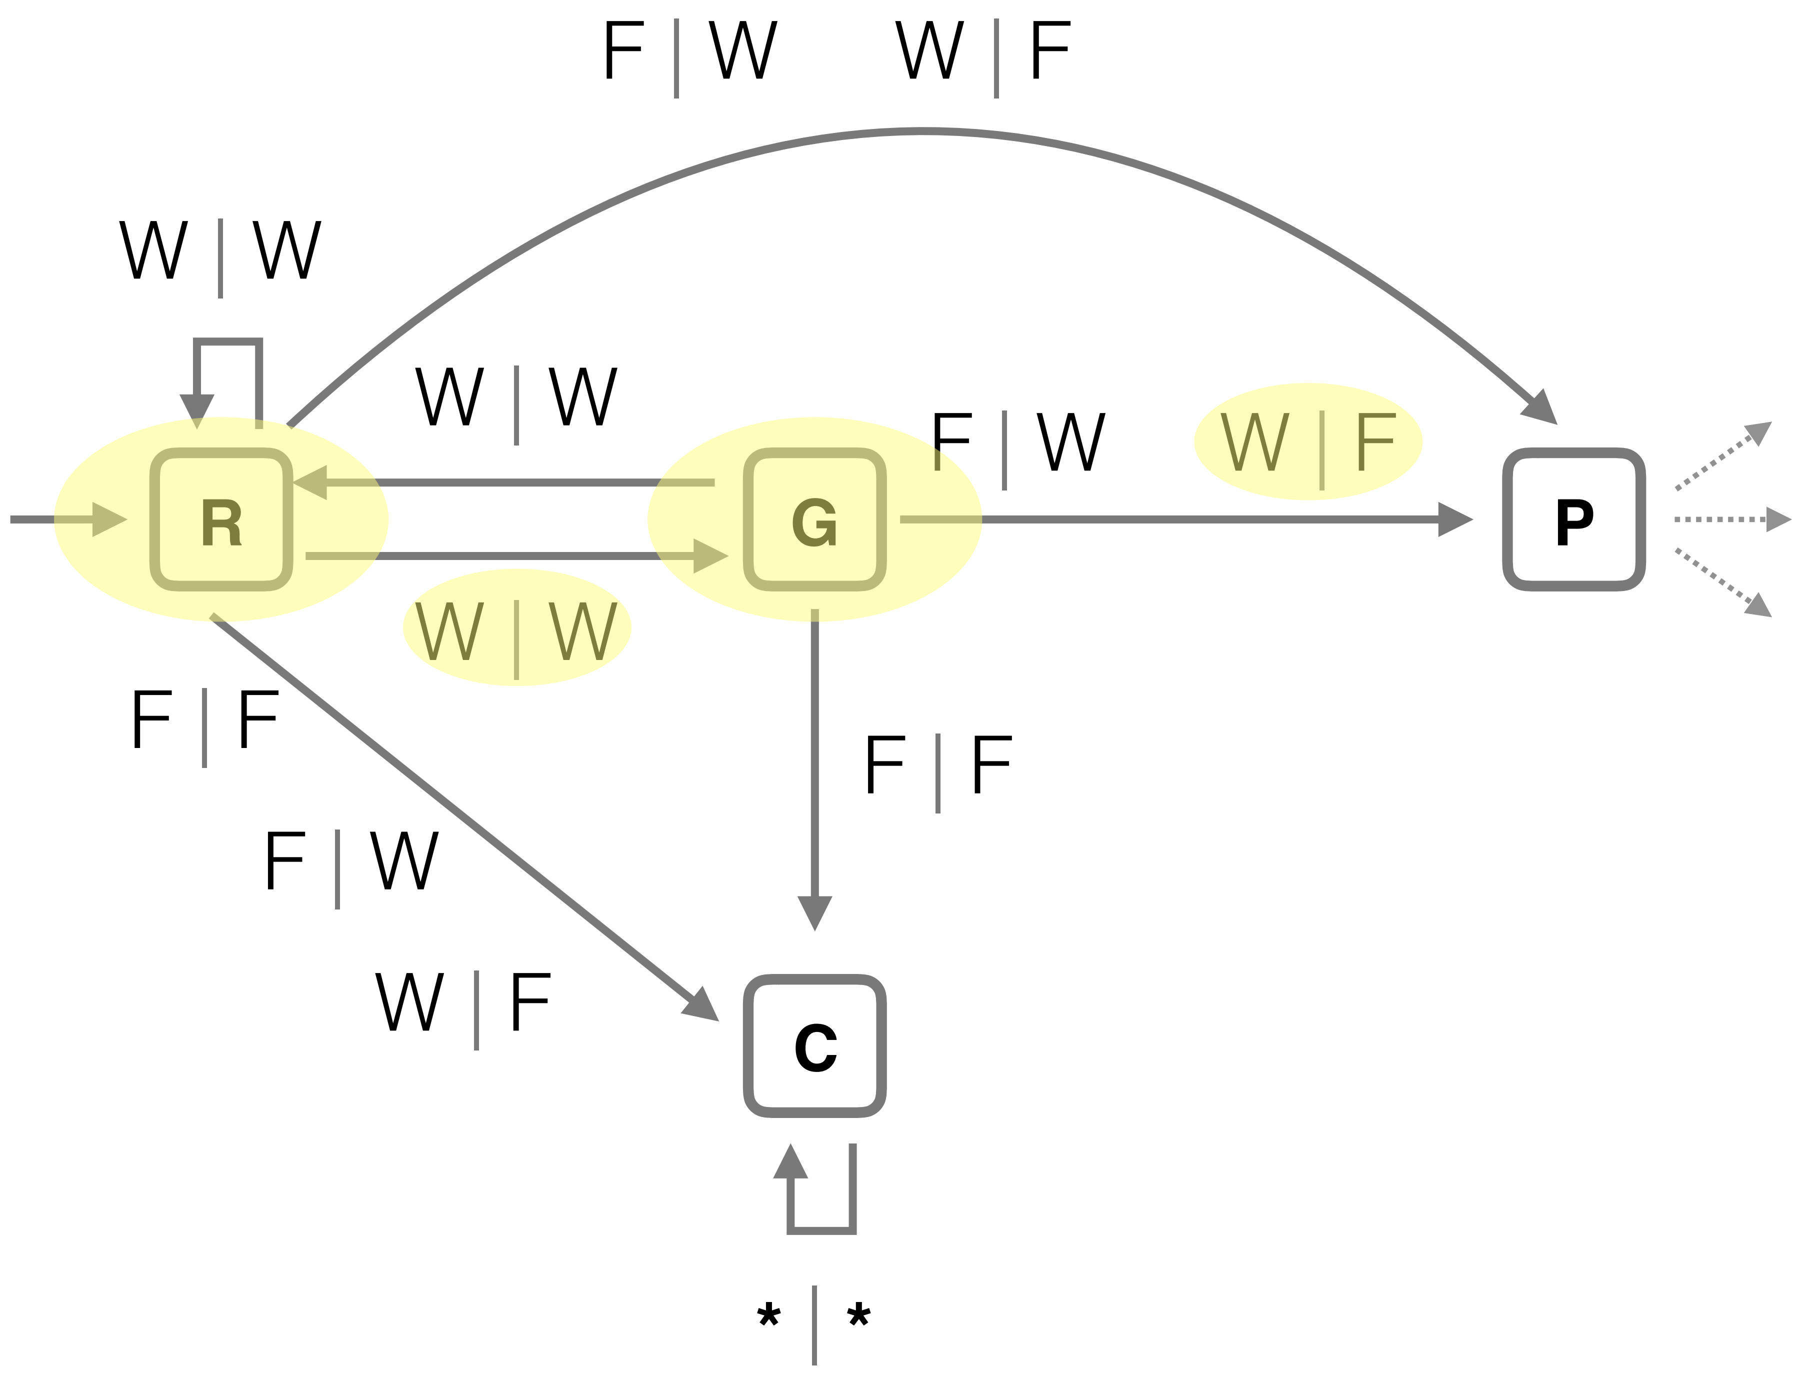

Play: R G

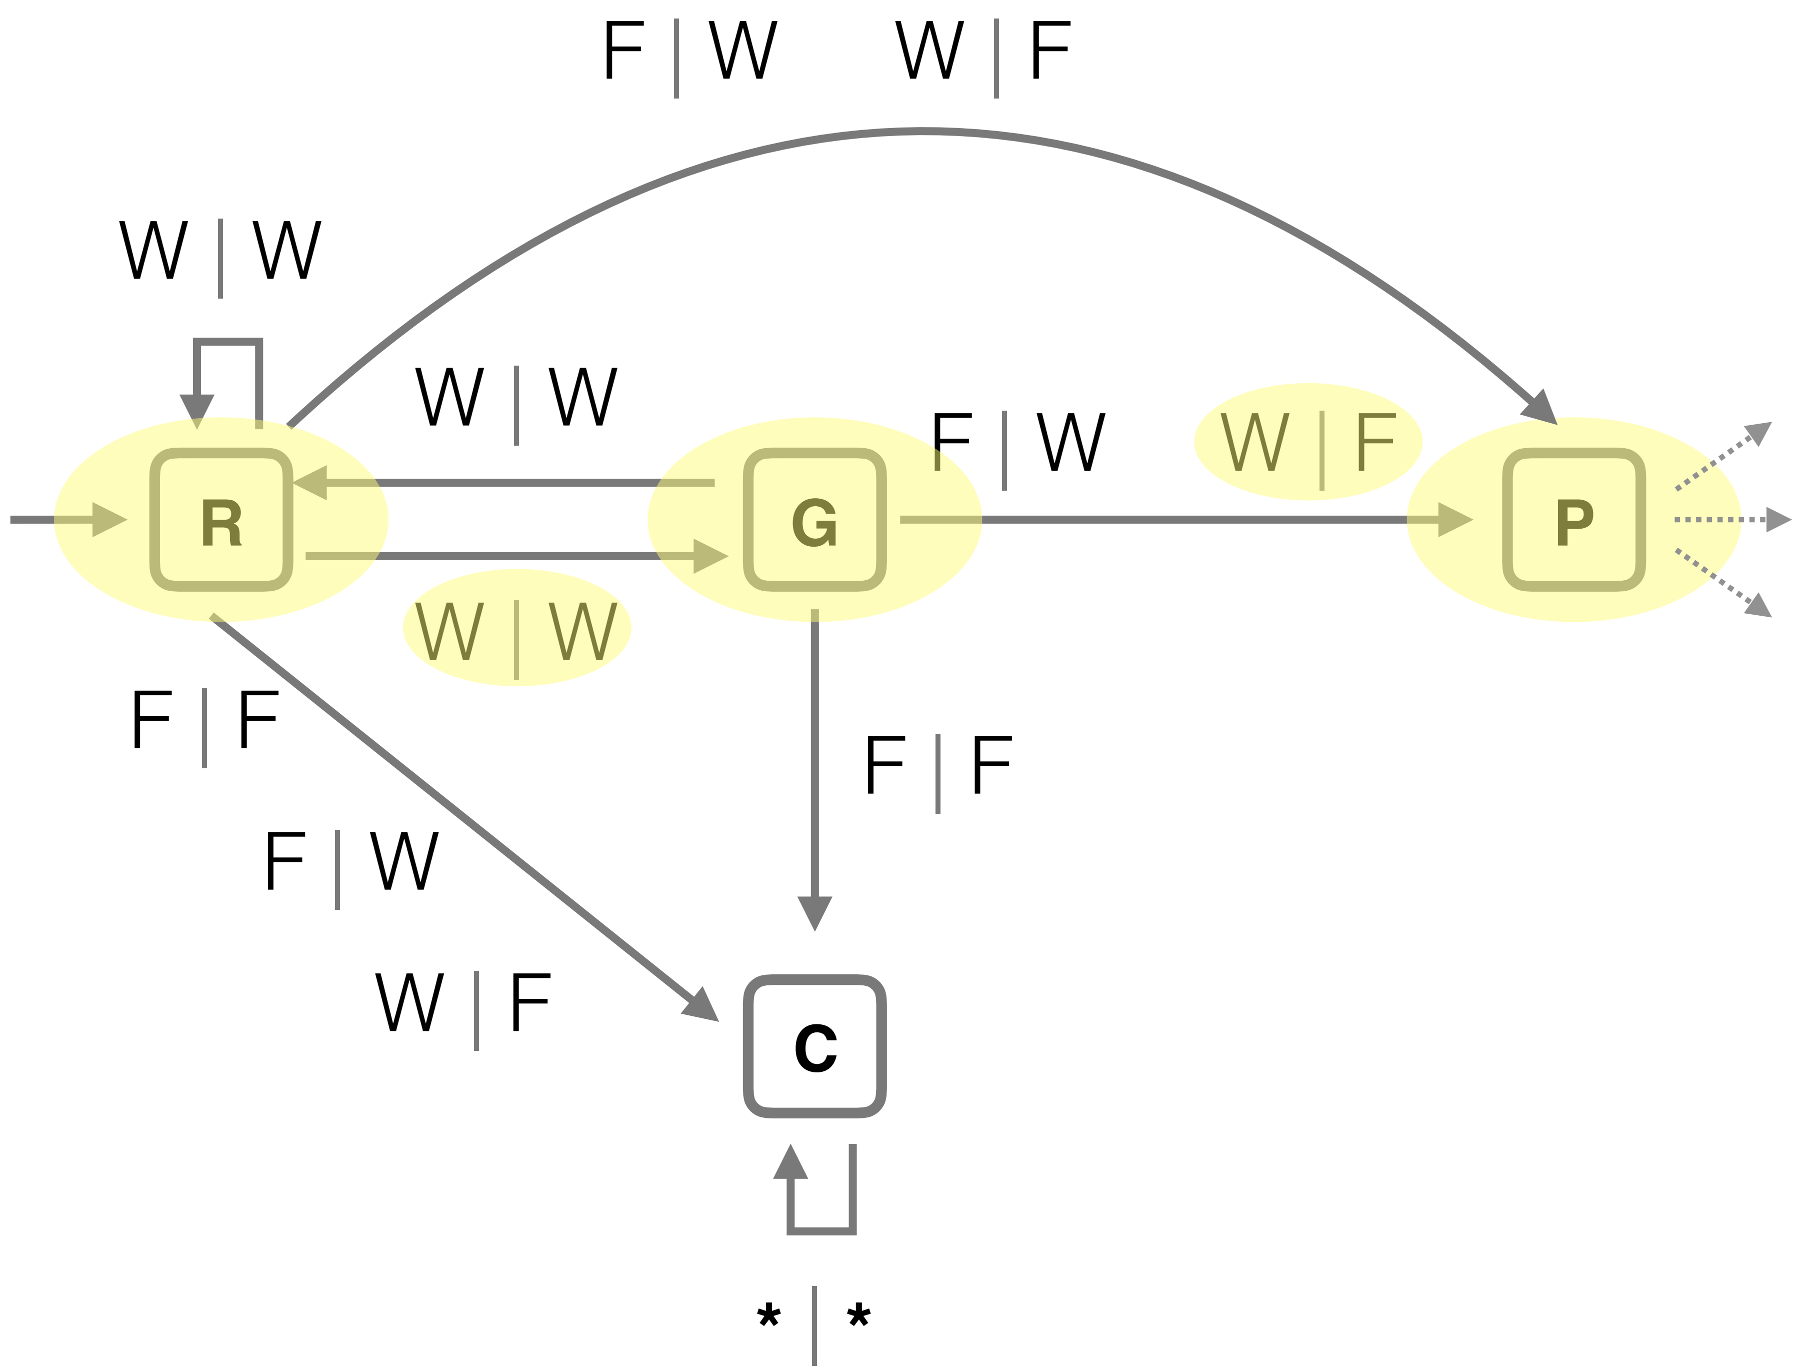

Play: R G P

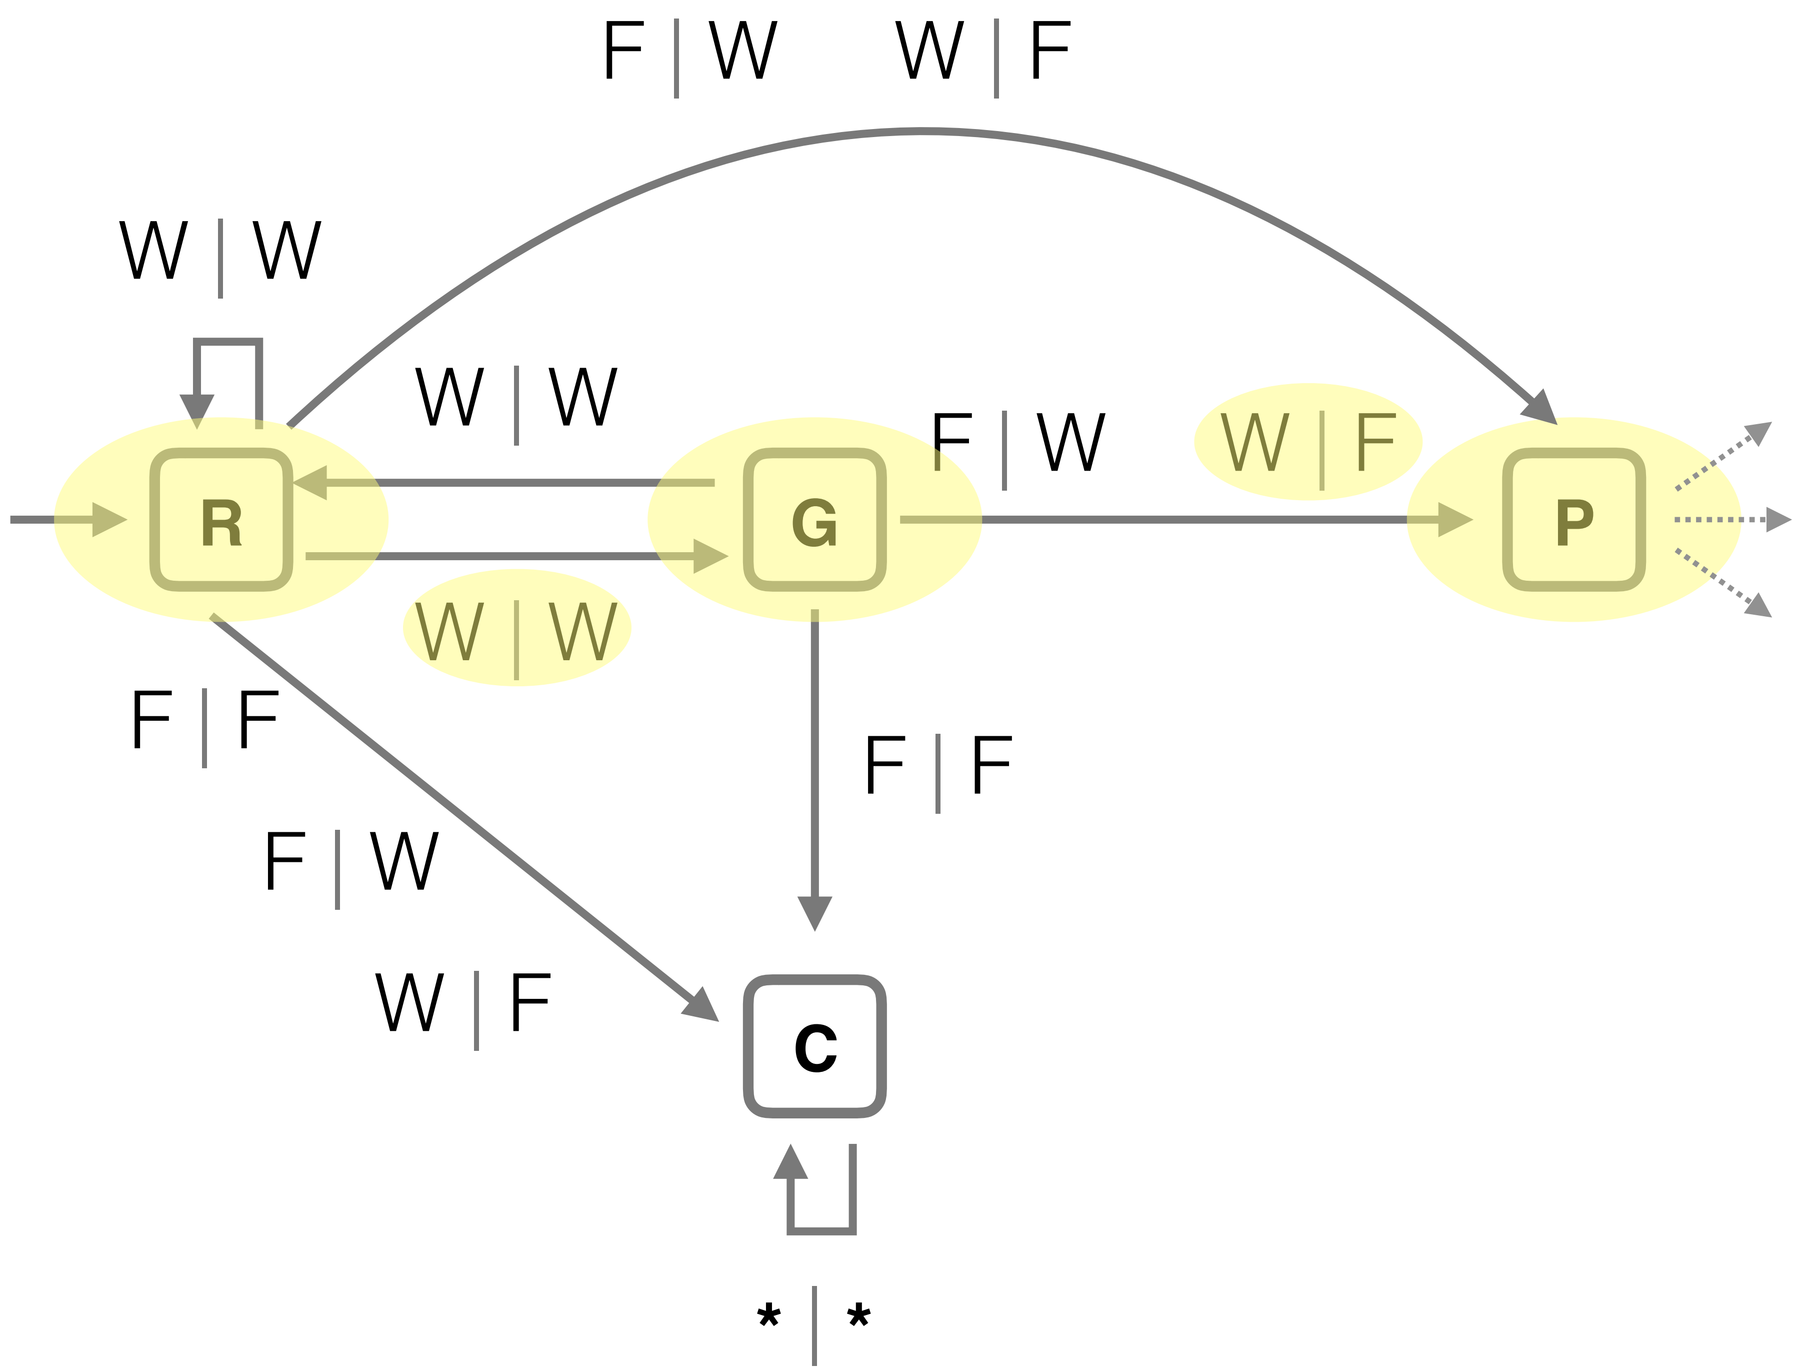

Play: R G P ...

Winning condition:
« avoid C »

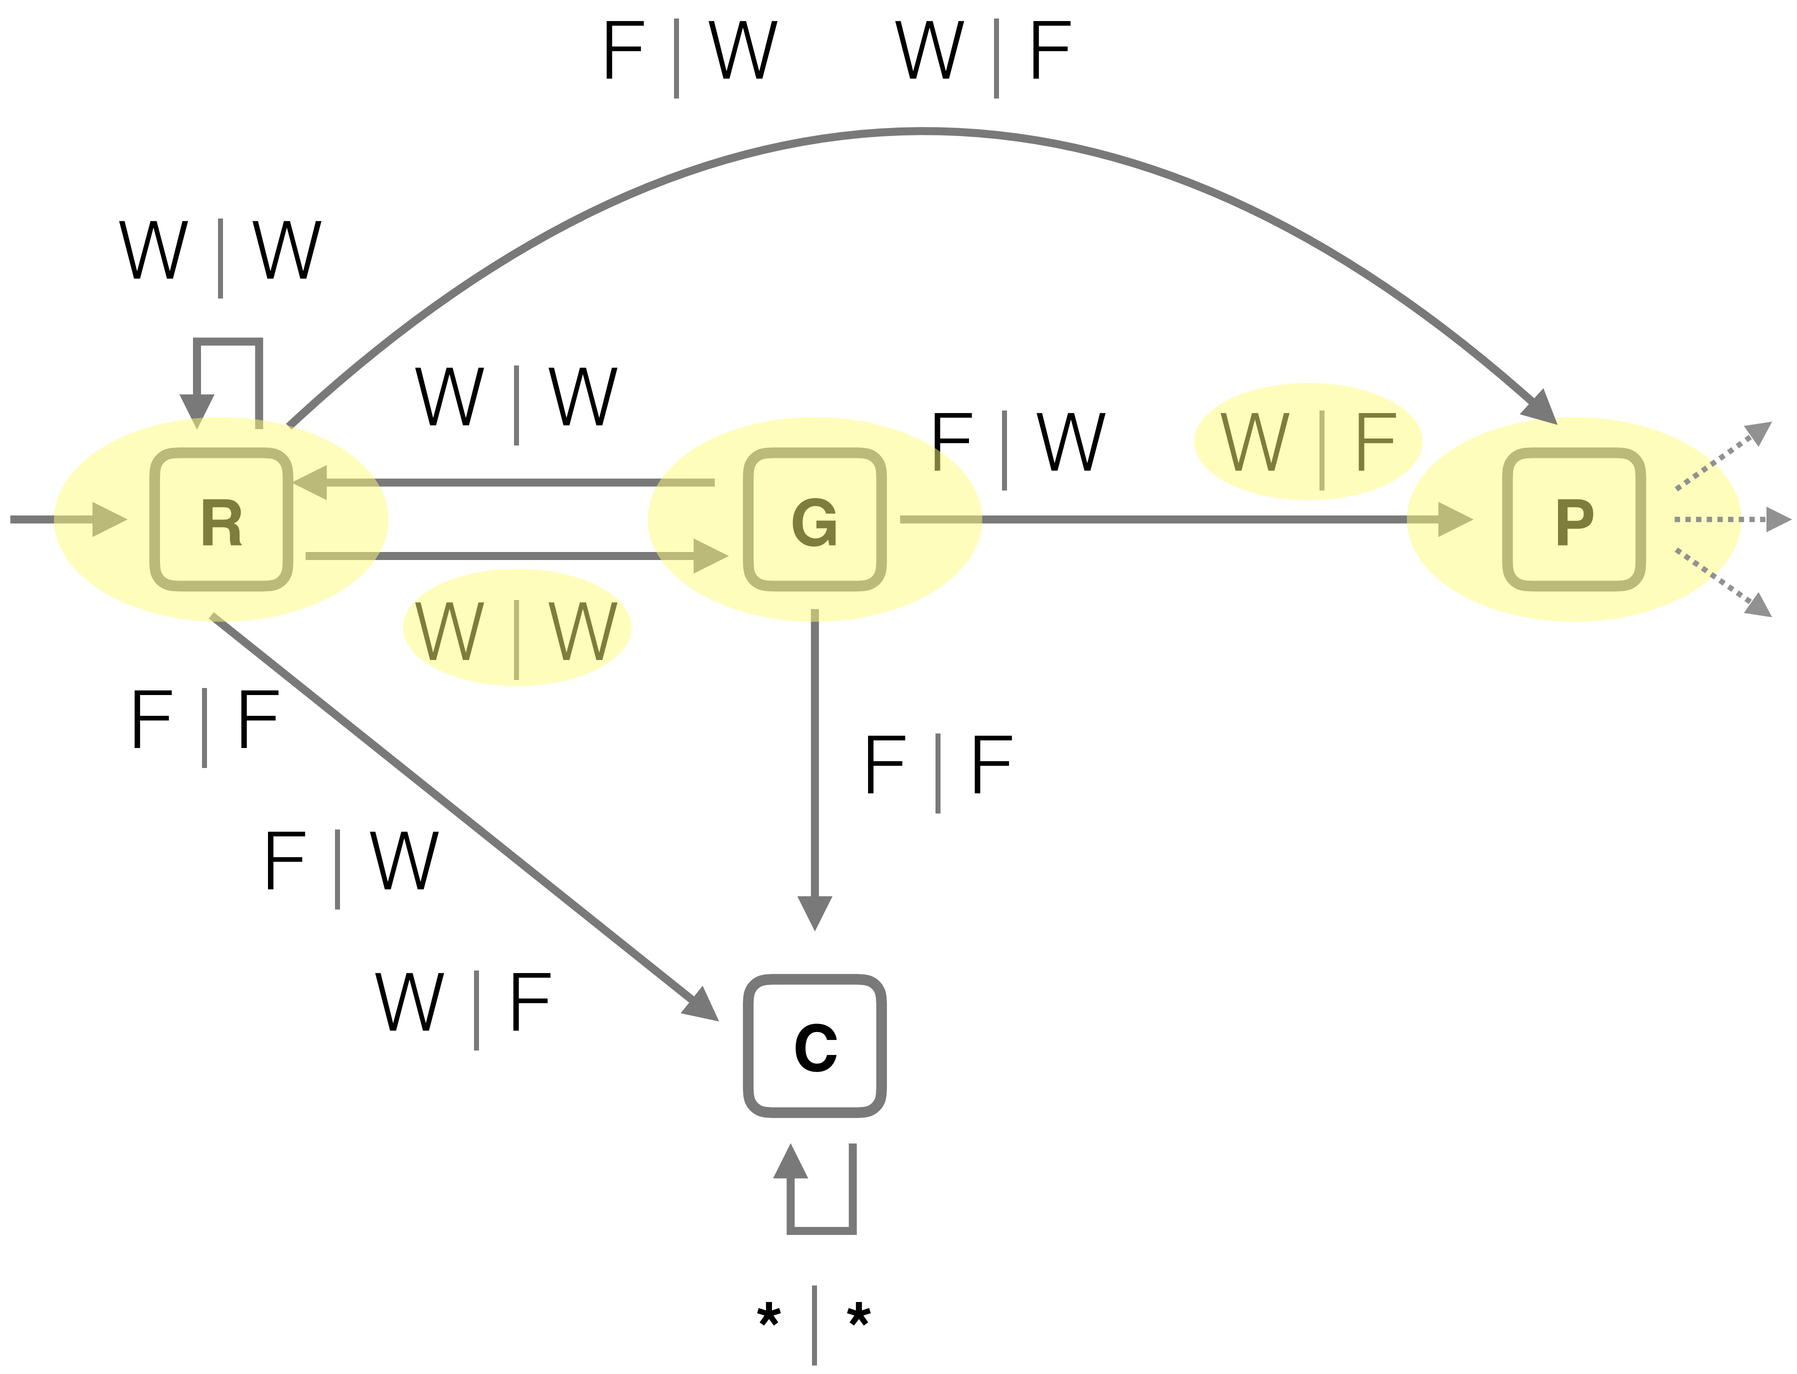

Play: R G P ...

Winning condition: « avoid C »

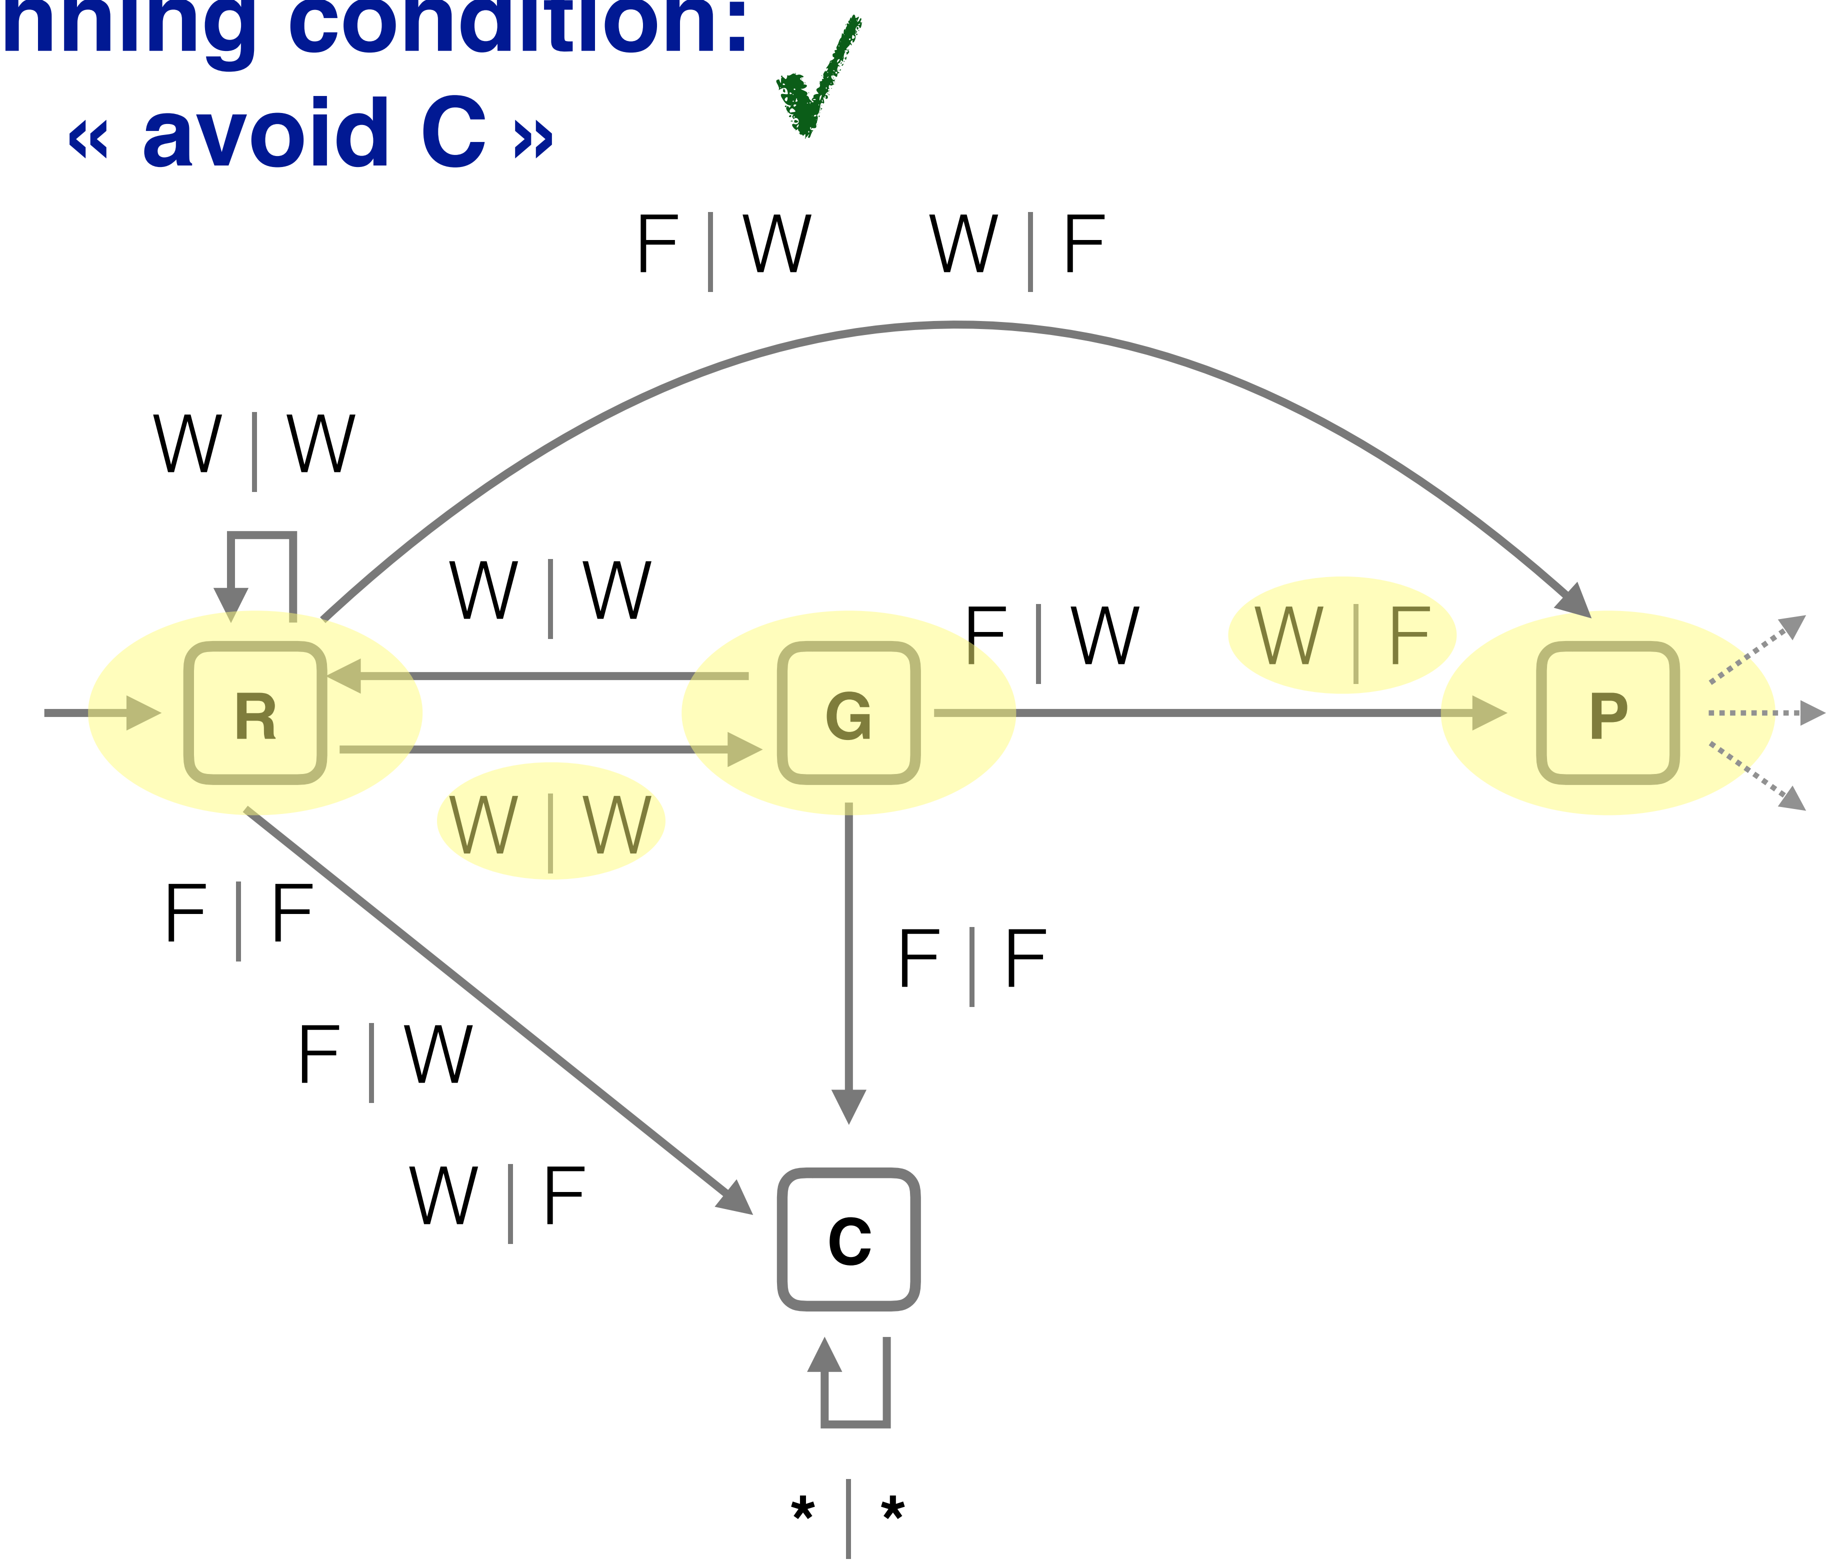

Play: R G P ...

Winning condition:
« avoid C »

Strategies:
s^1: V* -> A^1
s^2: V* -> A^2

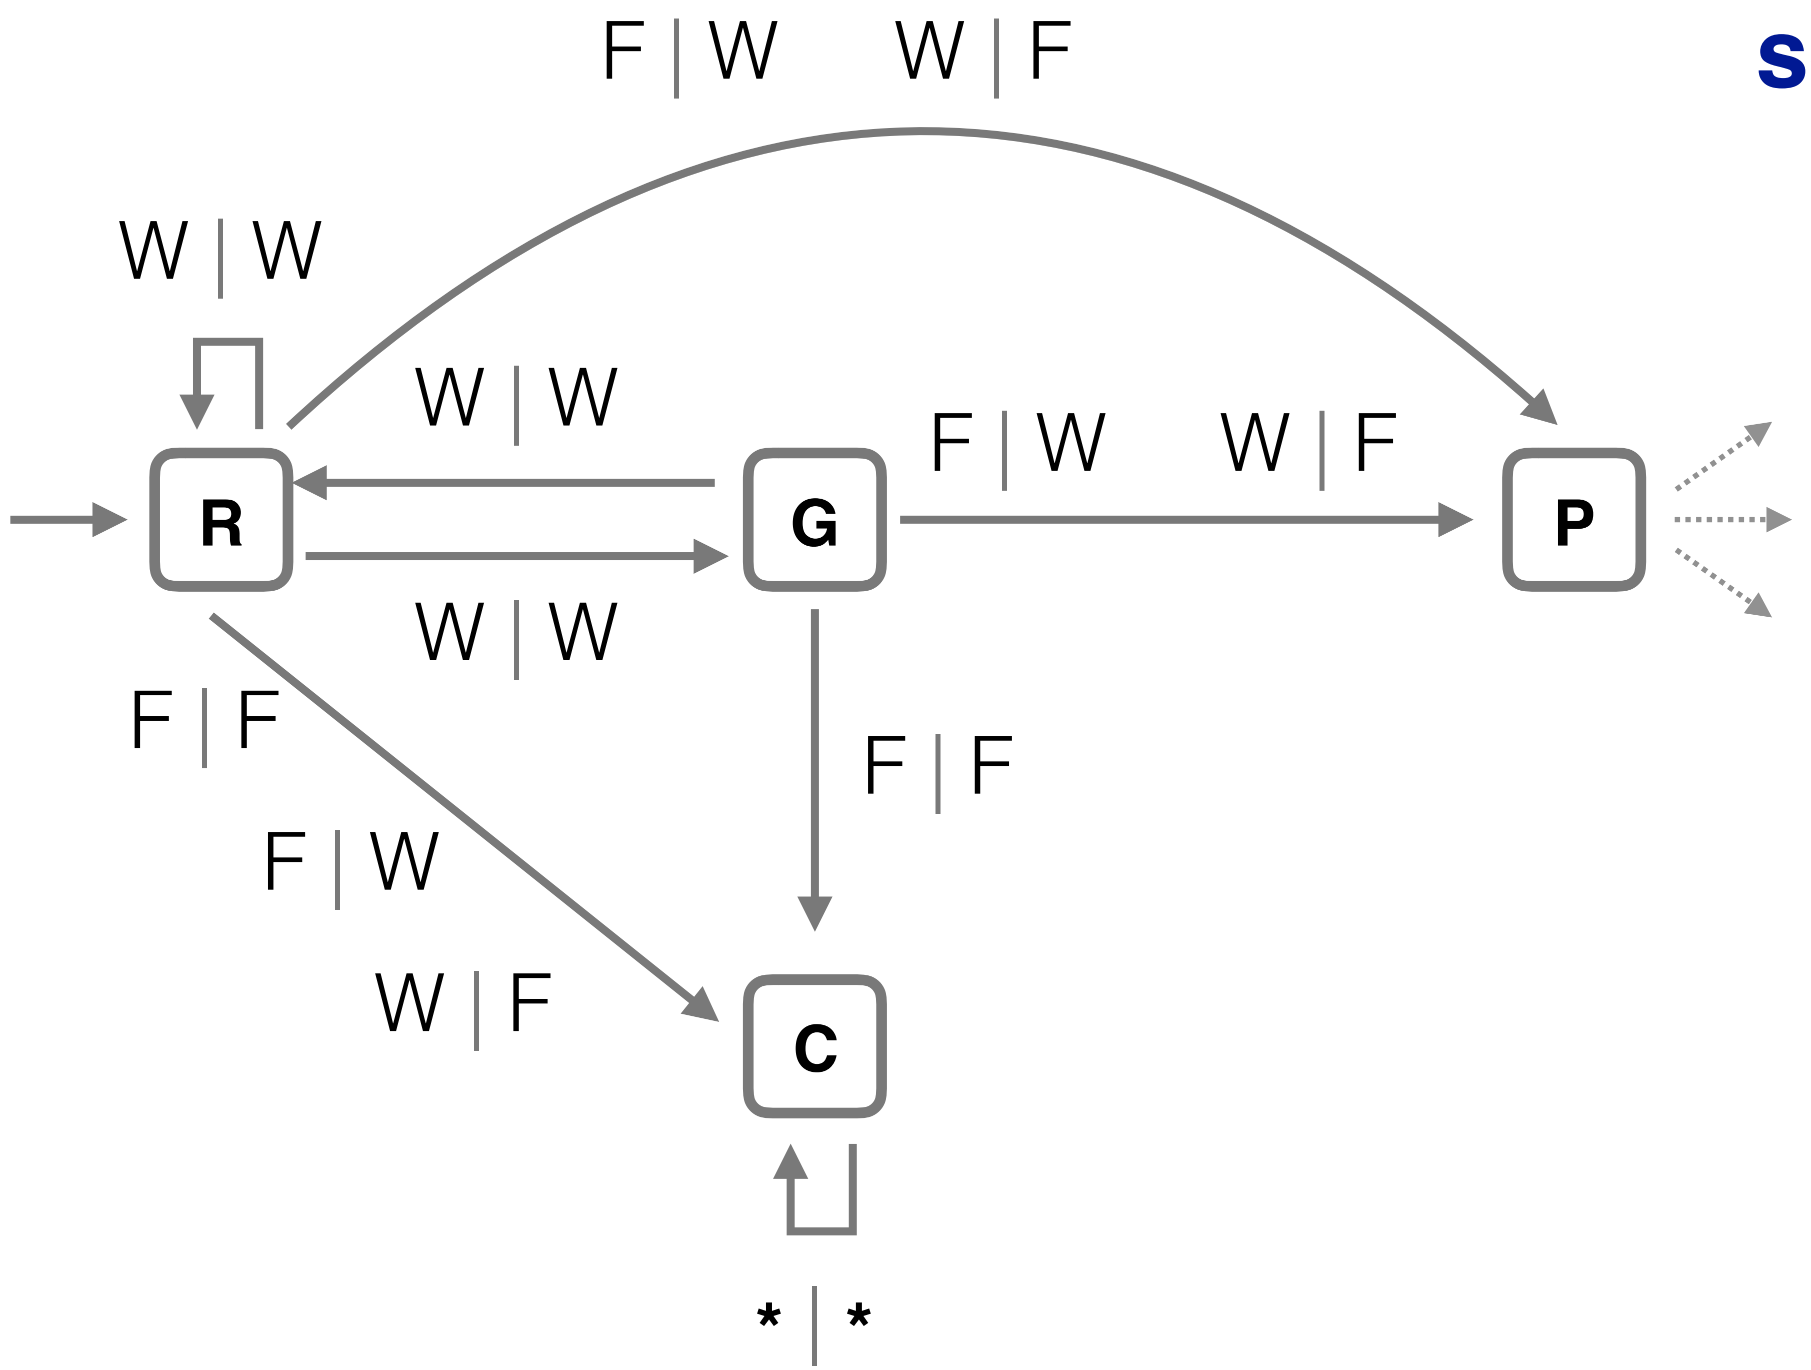

Winning condition:
« avoid C »

Strategies:
s^1: at R play F
else play W
s^2: at R play W
else play F

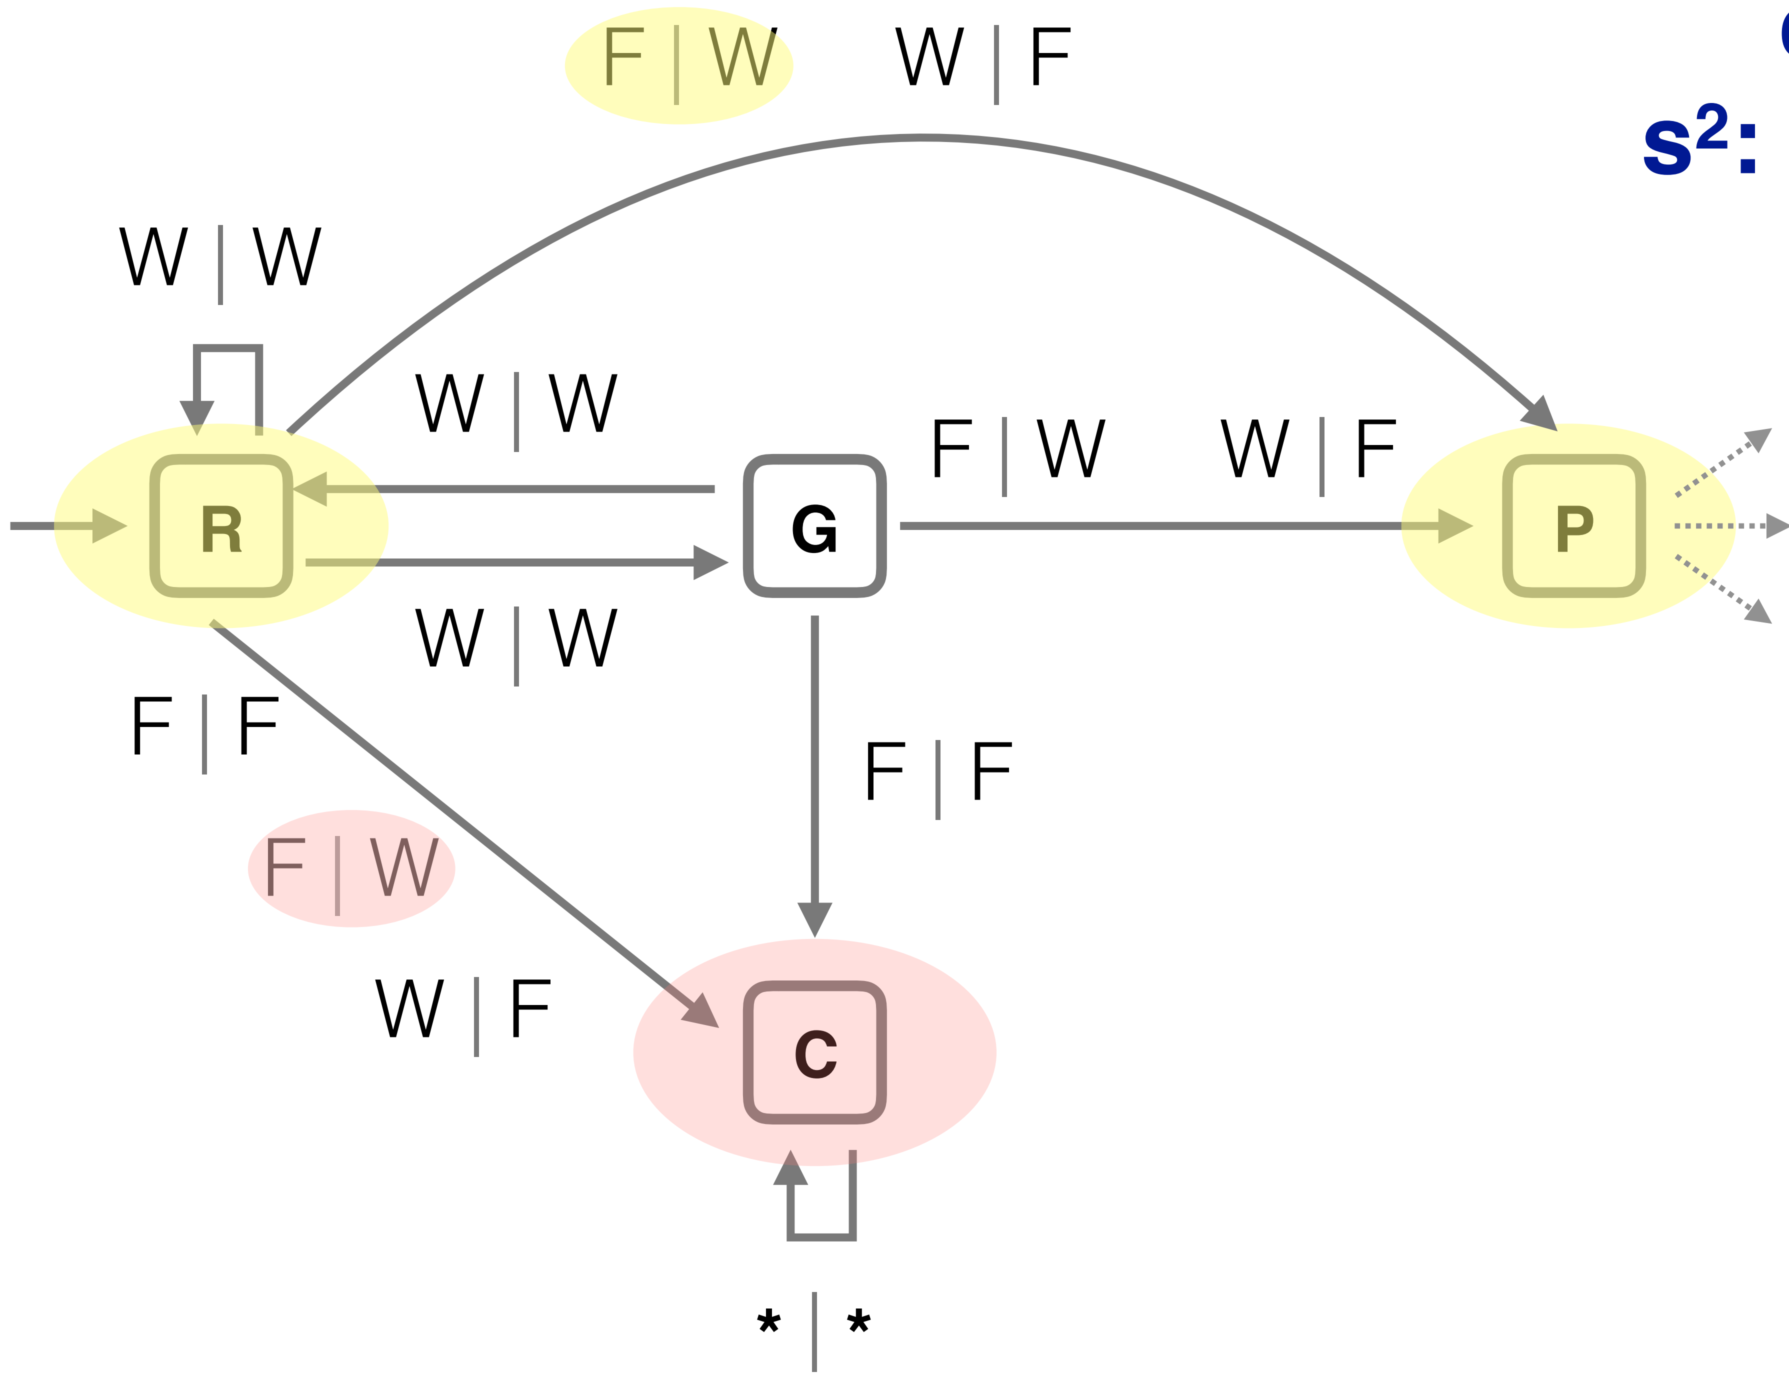

Plays: **R P …**
R C …

Winning condition:
« avoid C »

Strategies:
s^1: at R play F
else play W
s^2: at R play W
else play F

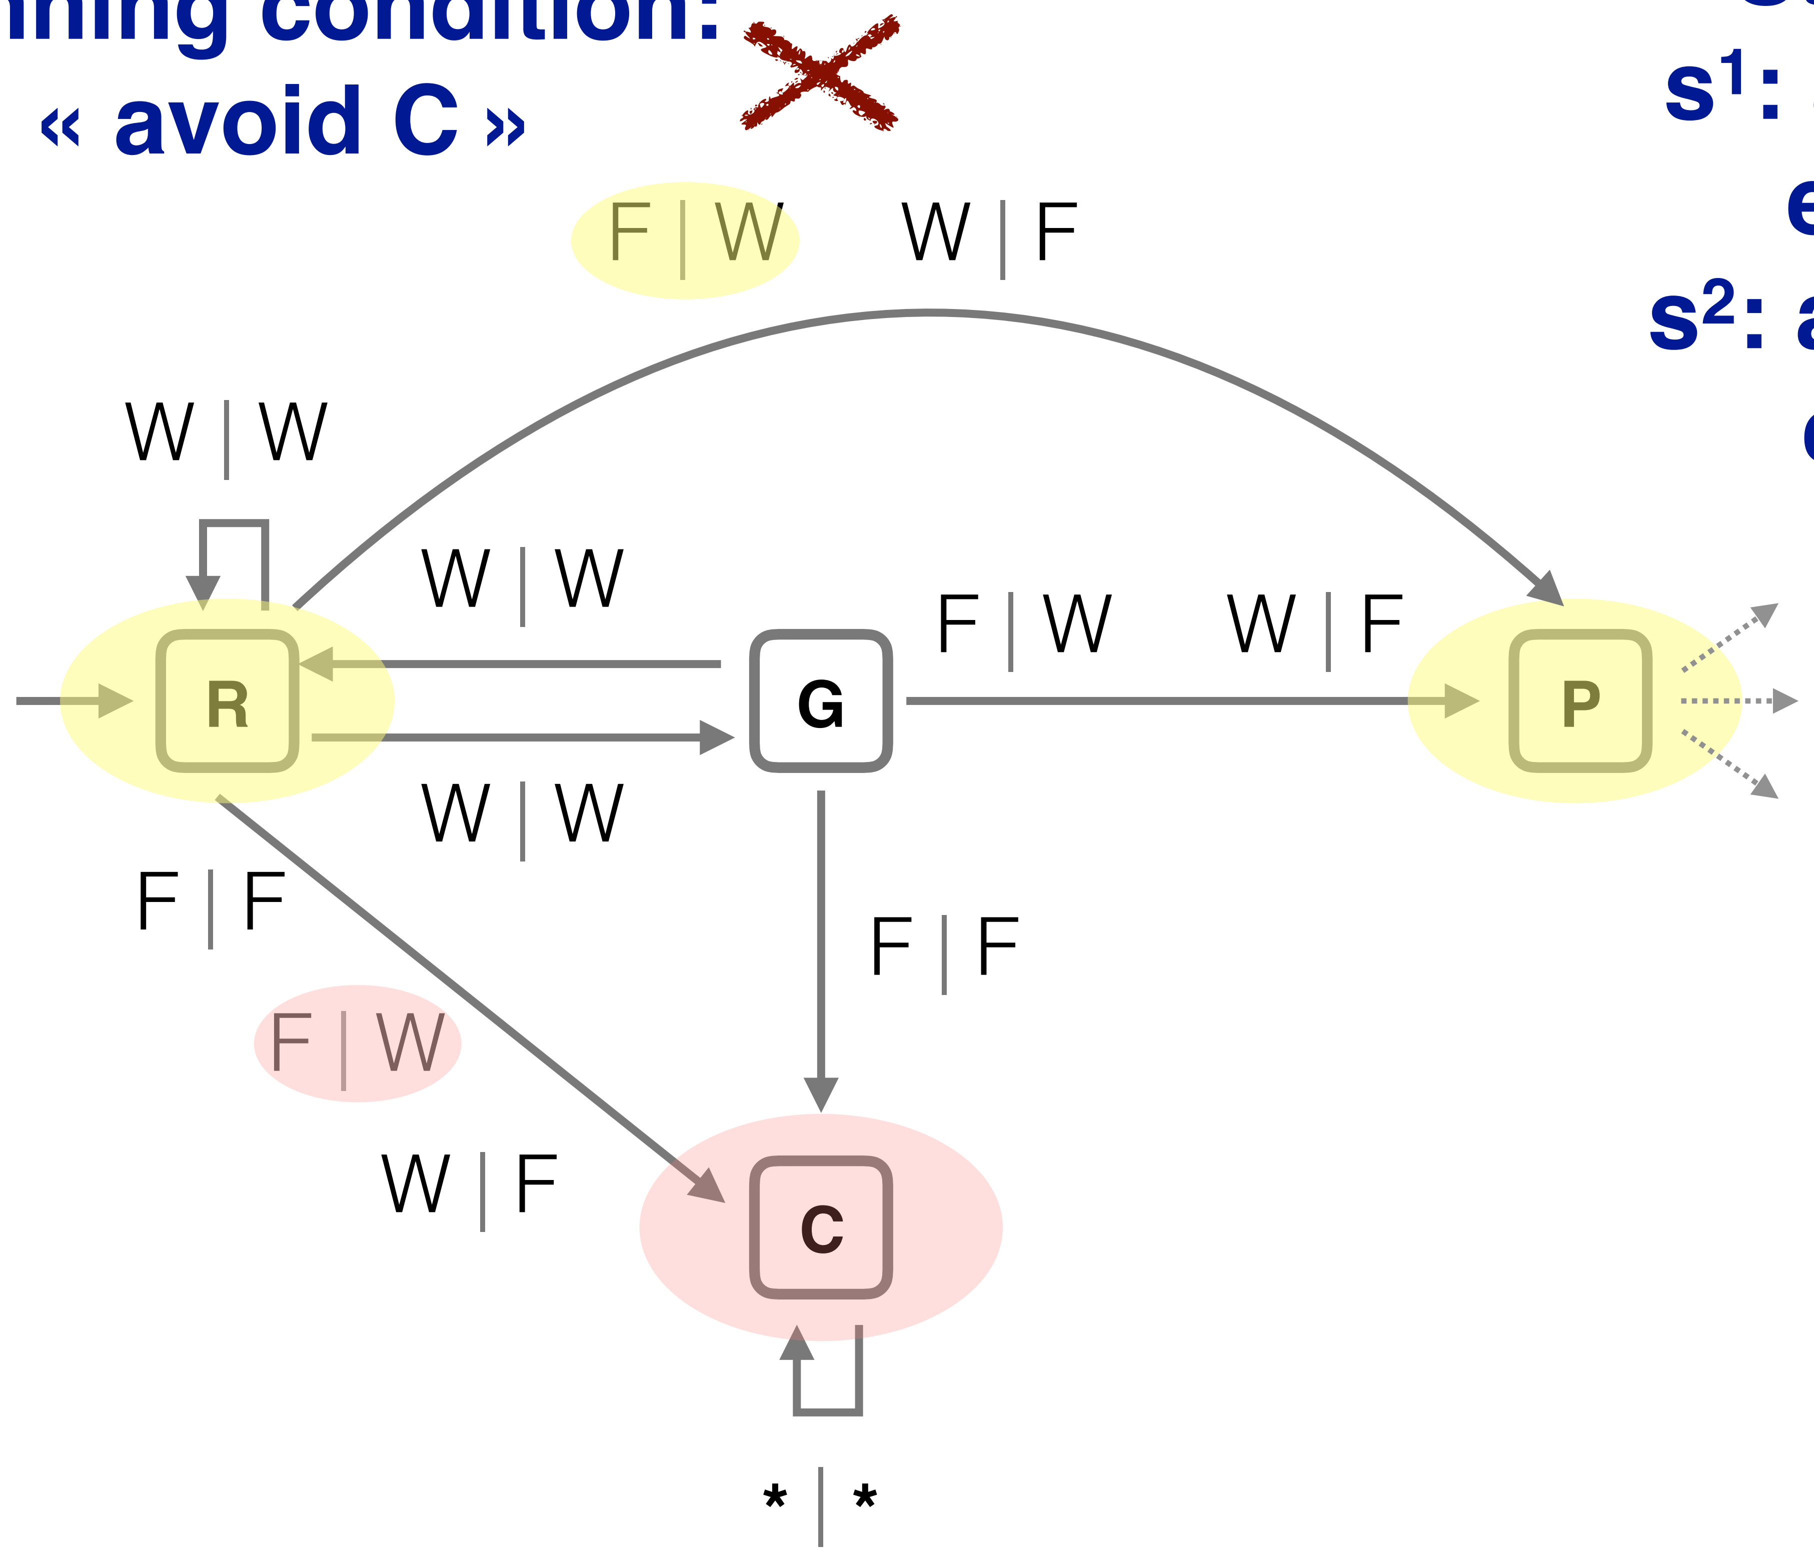

Plays: **R P …**
R C …

Winning condition:
« avoid C »

Strategies:
s^1: always play W
s^2: at R play W
at G play F

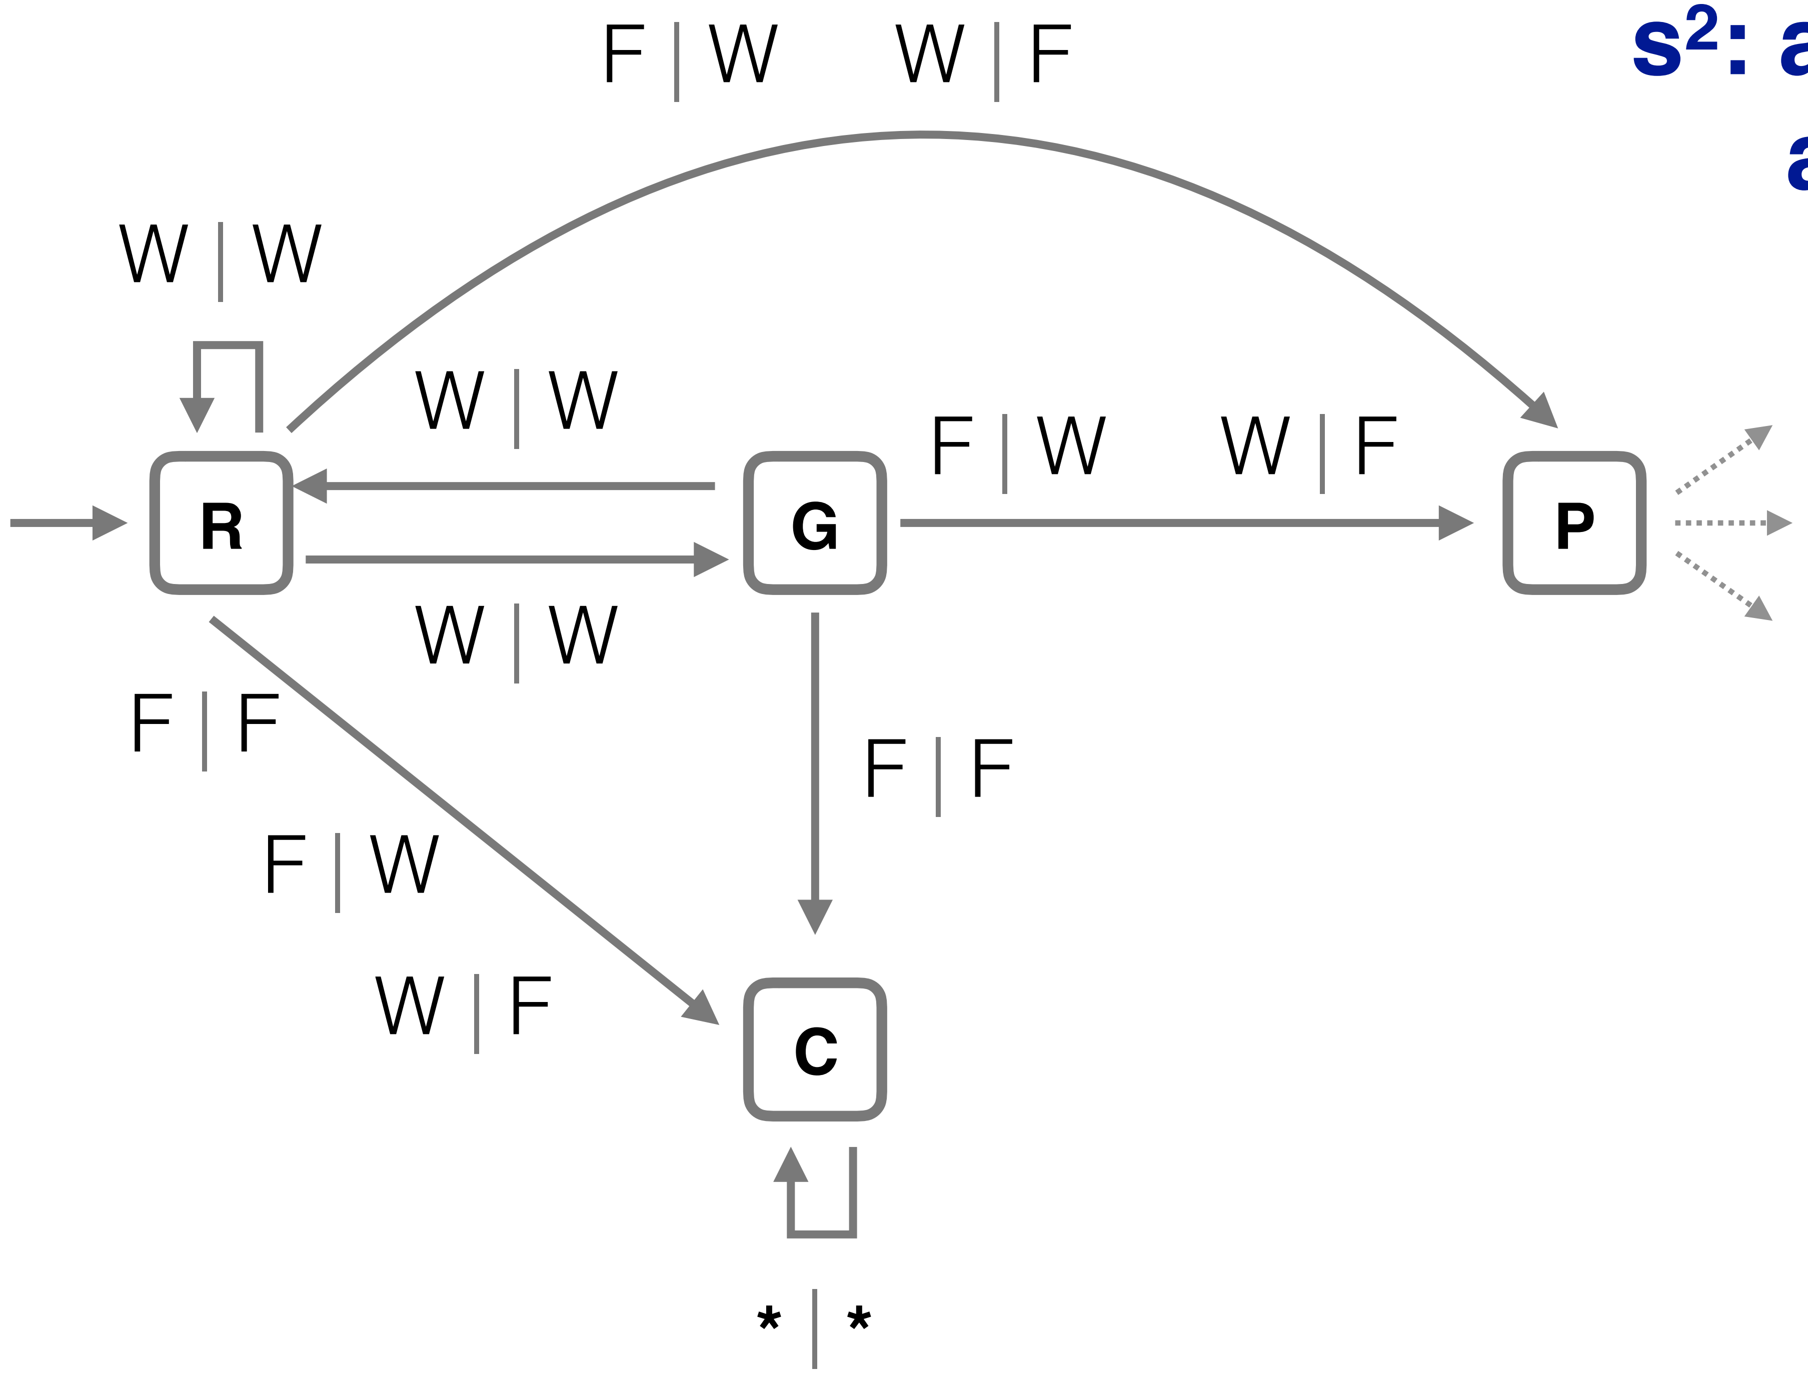

Winning condition:
« avoid C »

Strategies:
s^1: always play W
s^2: at R play W
at G play F

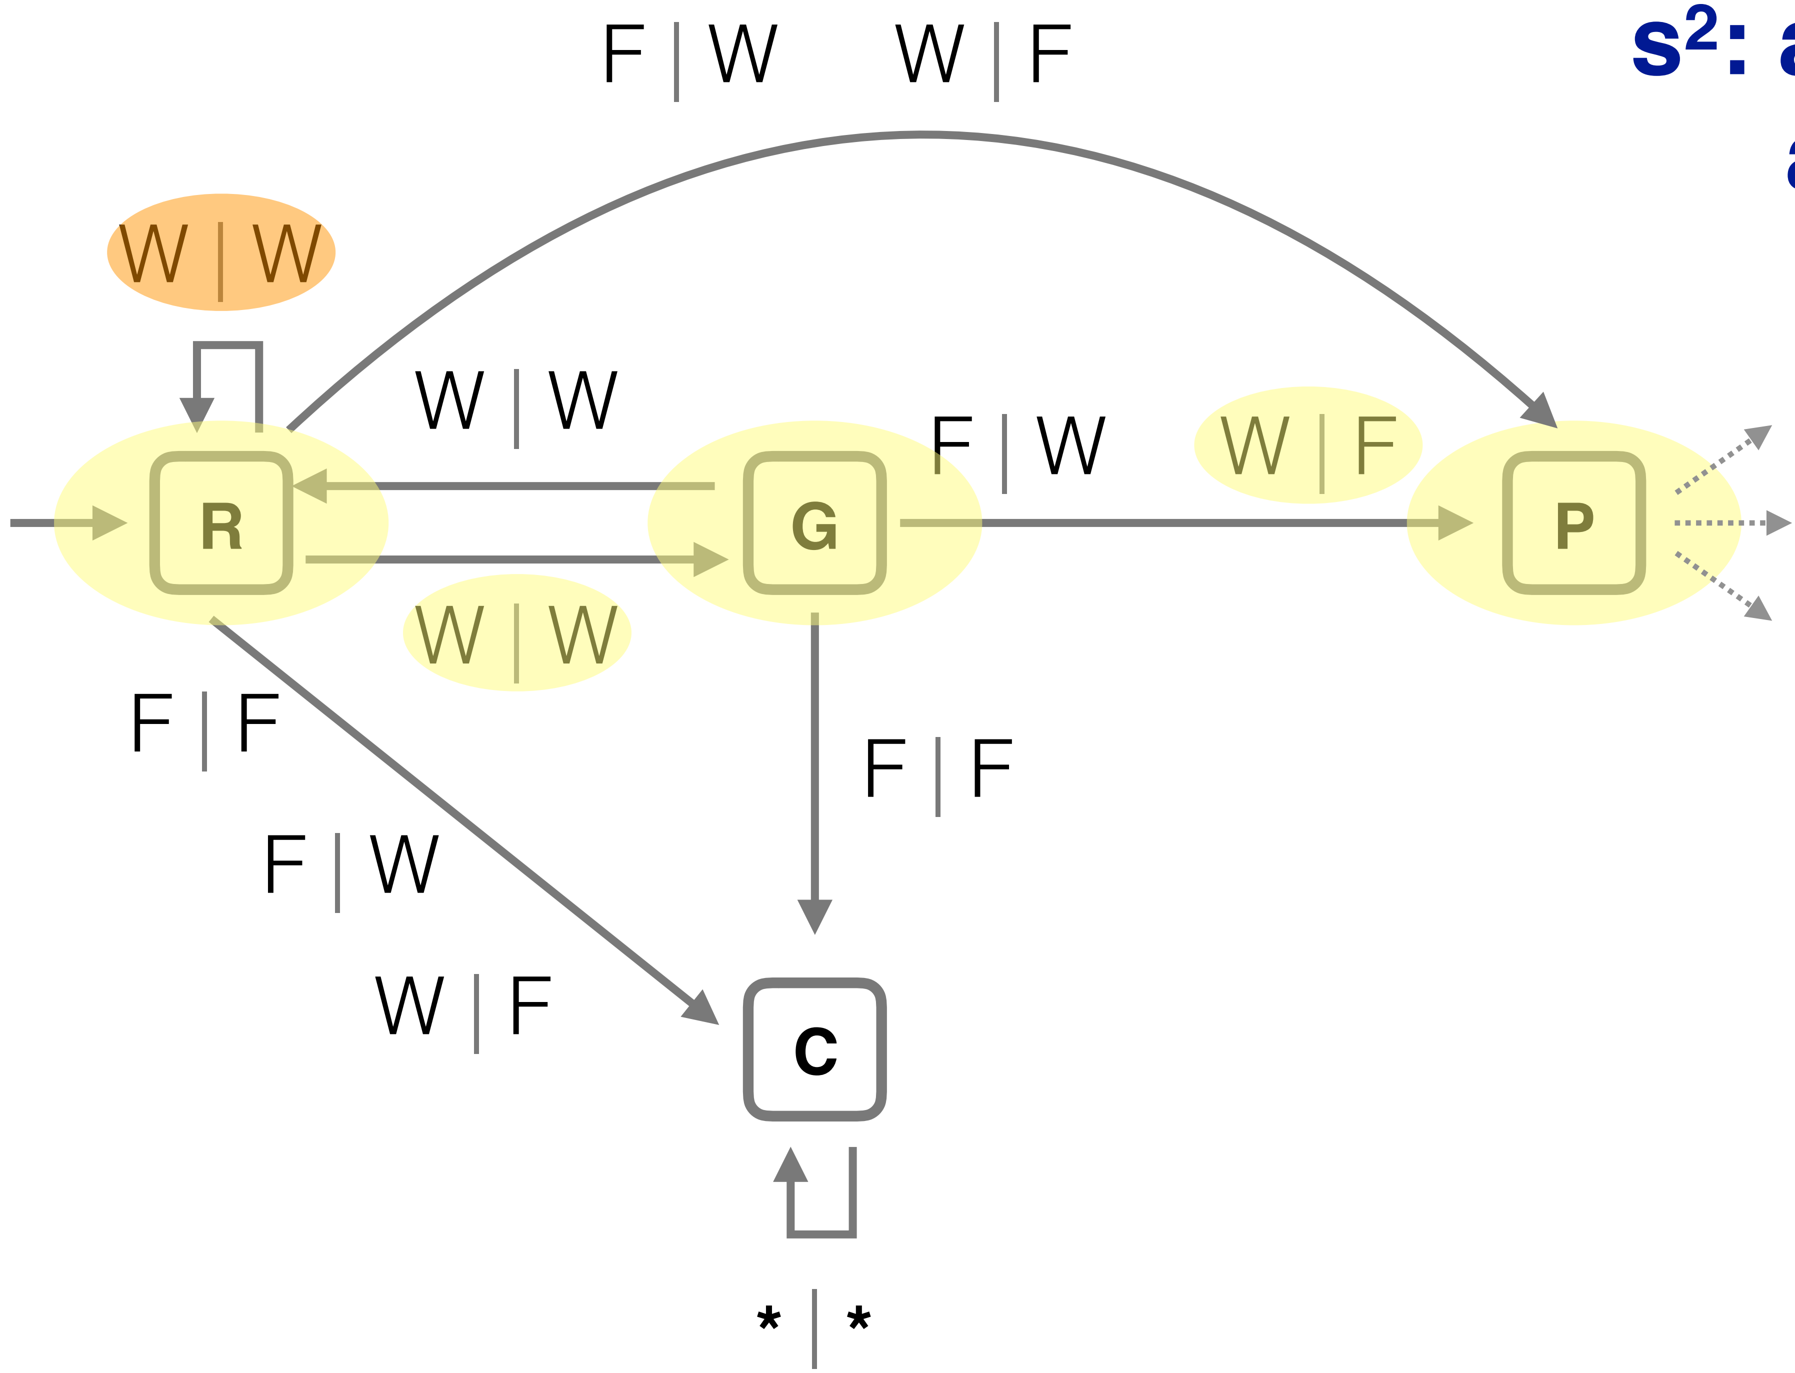

Winning condition:
« avoid C » ✓

Strategies:
s^1: always play W
s^2: at R play W
at G play F

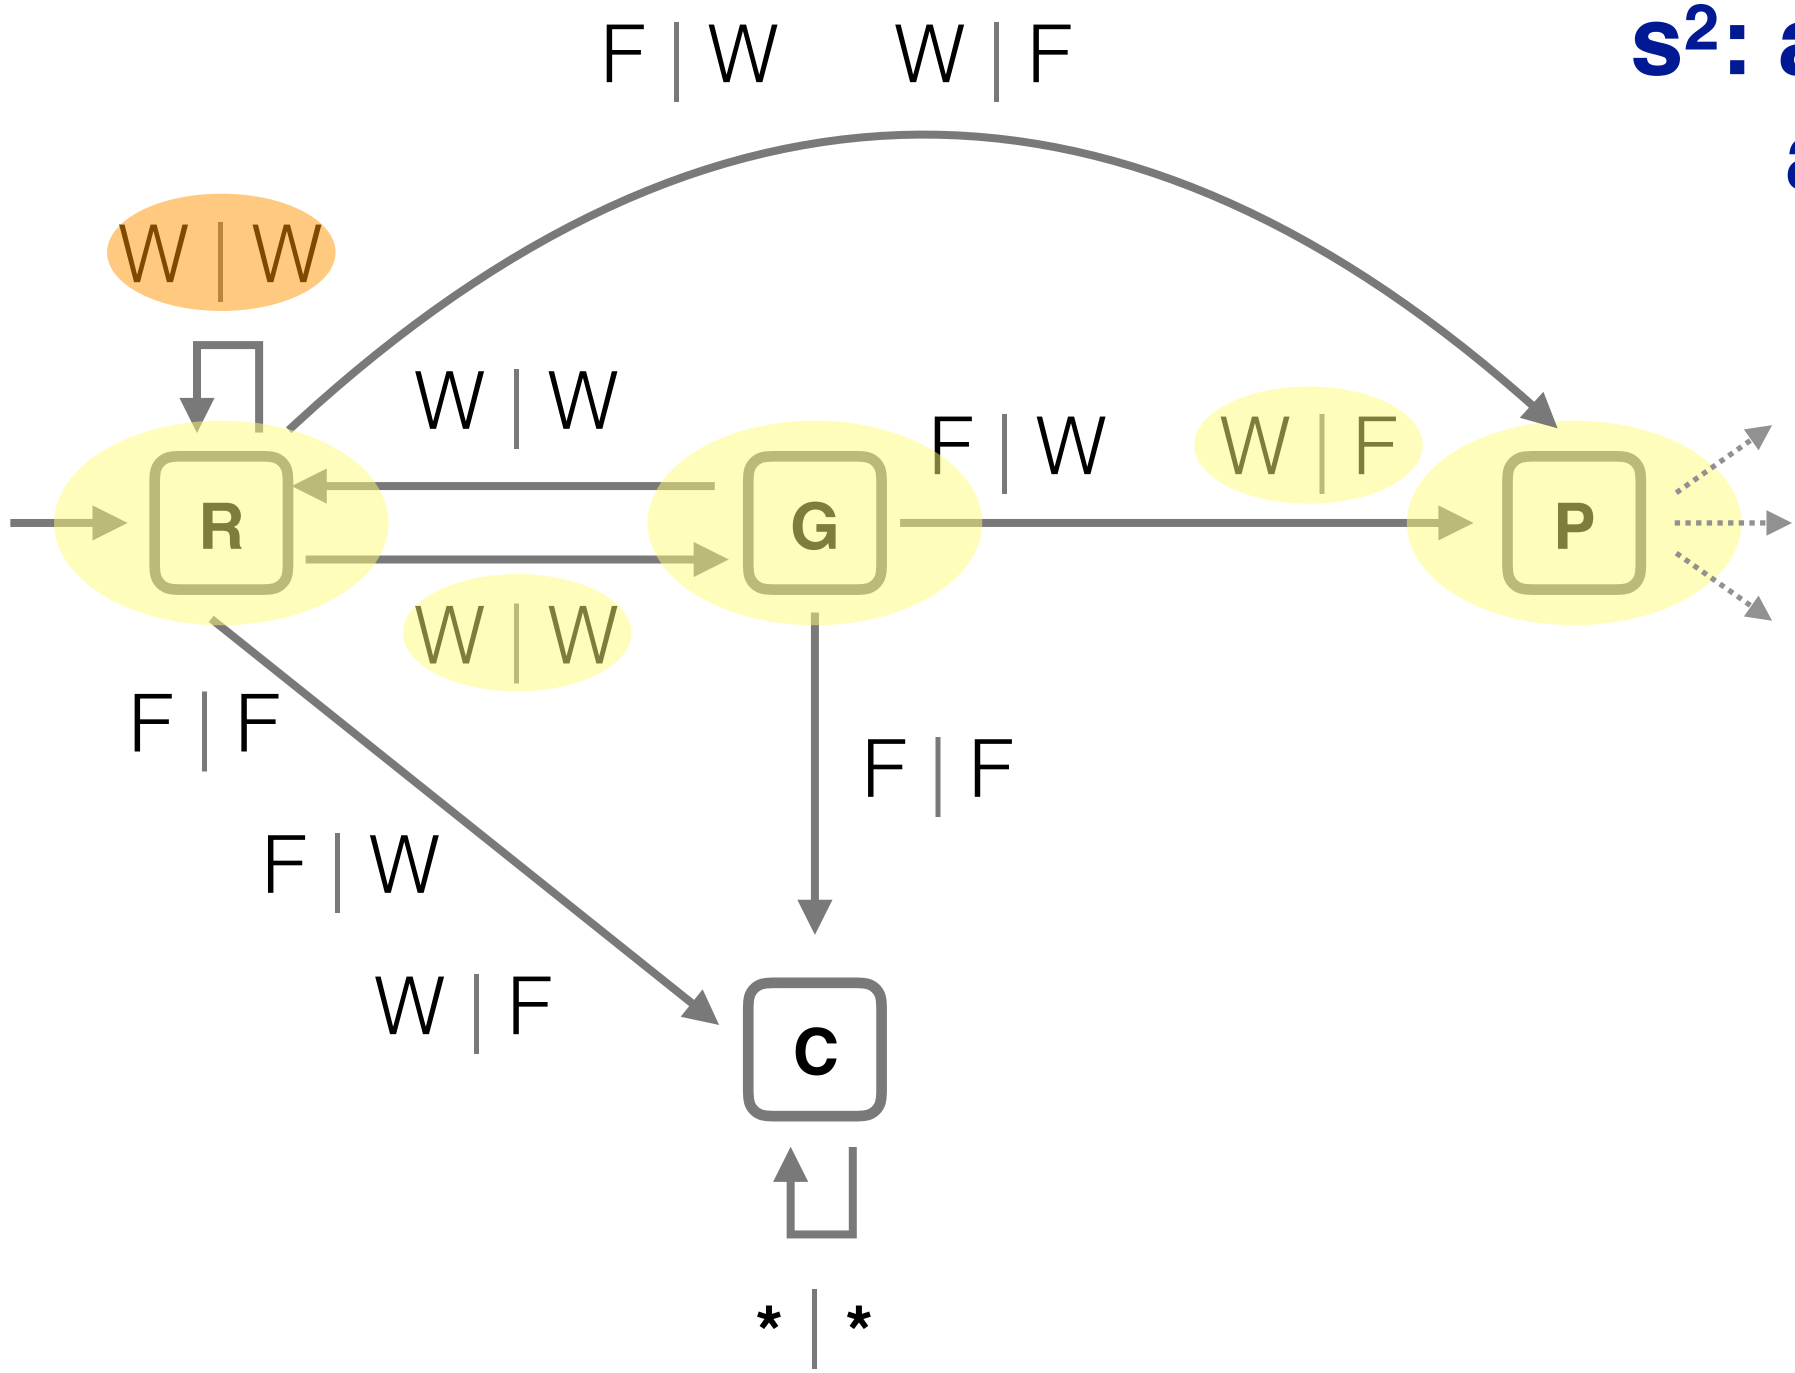

Games with Imperfect Information

Games with Imperfect Information

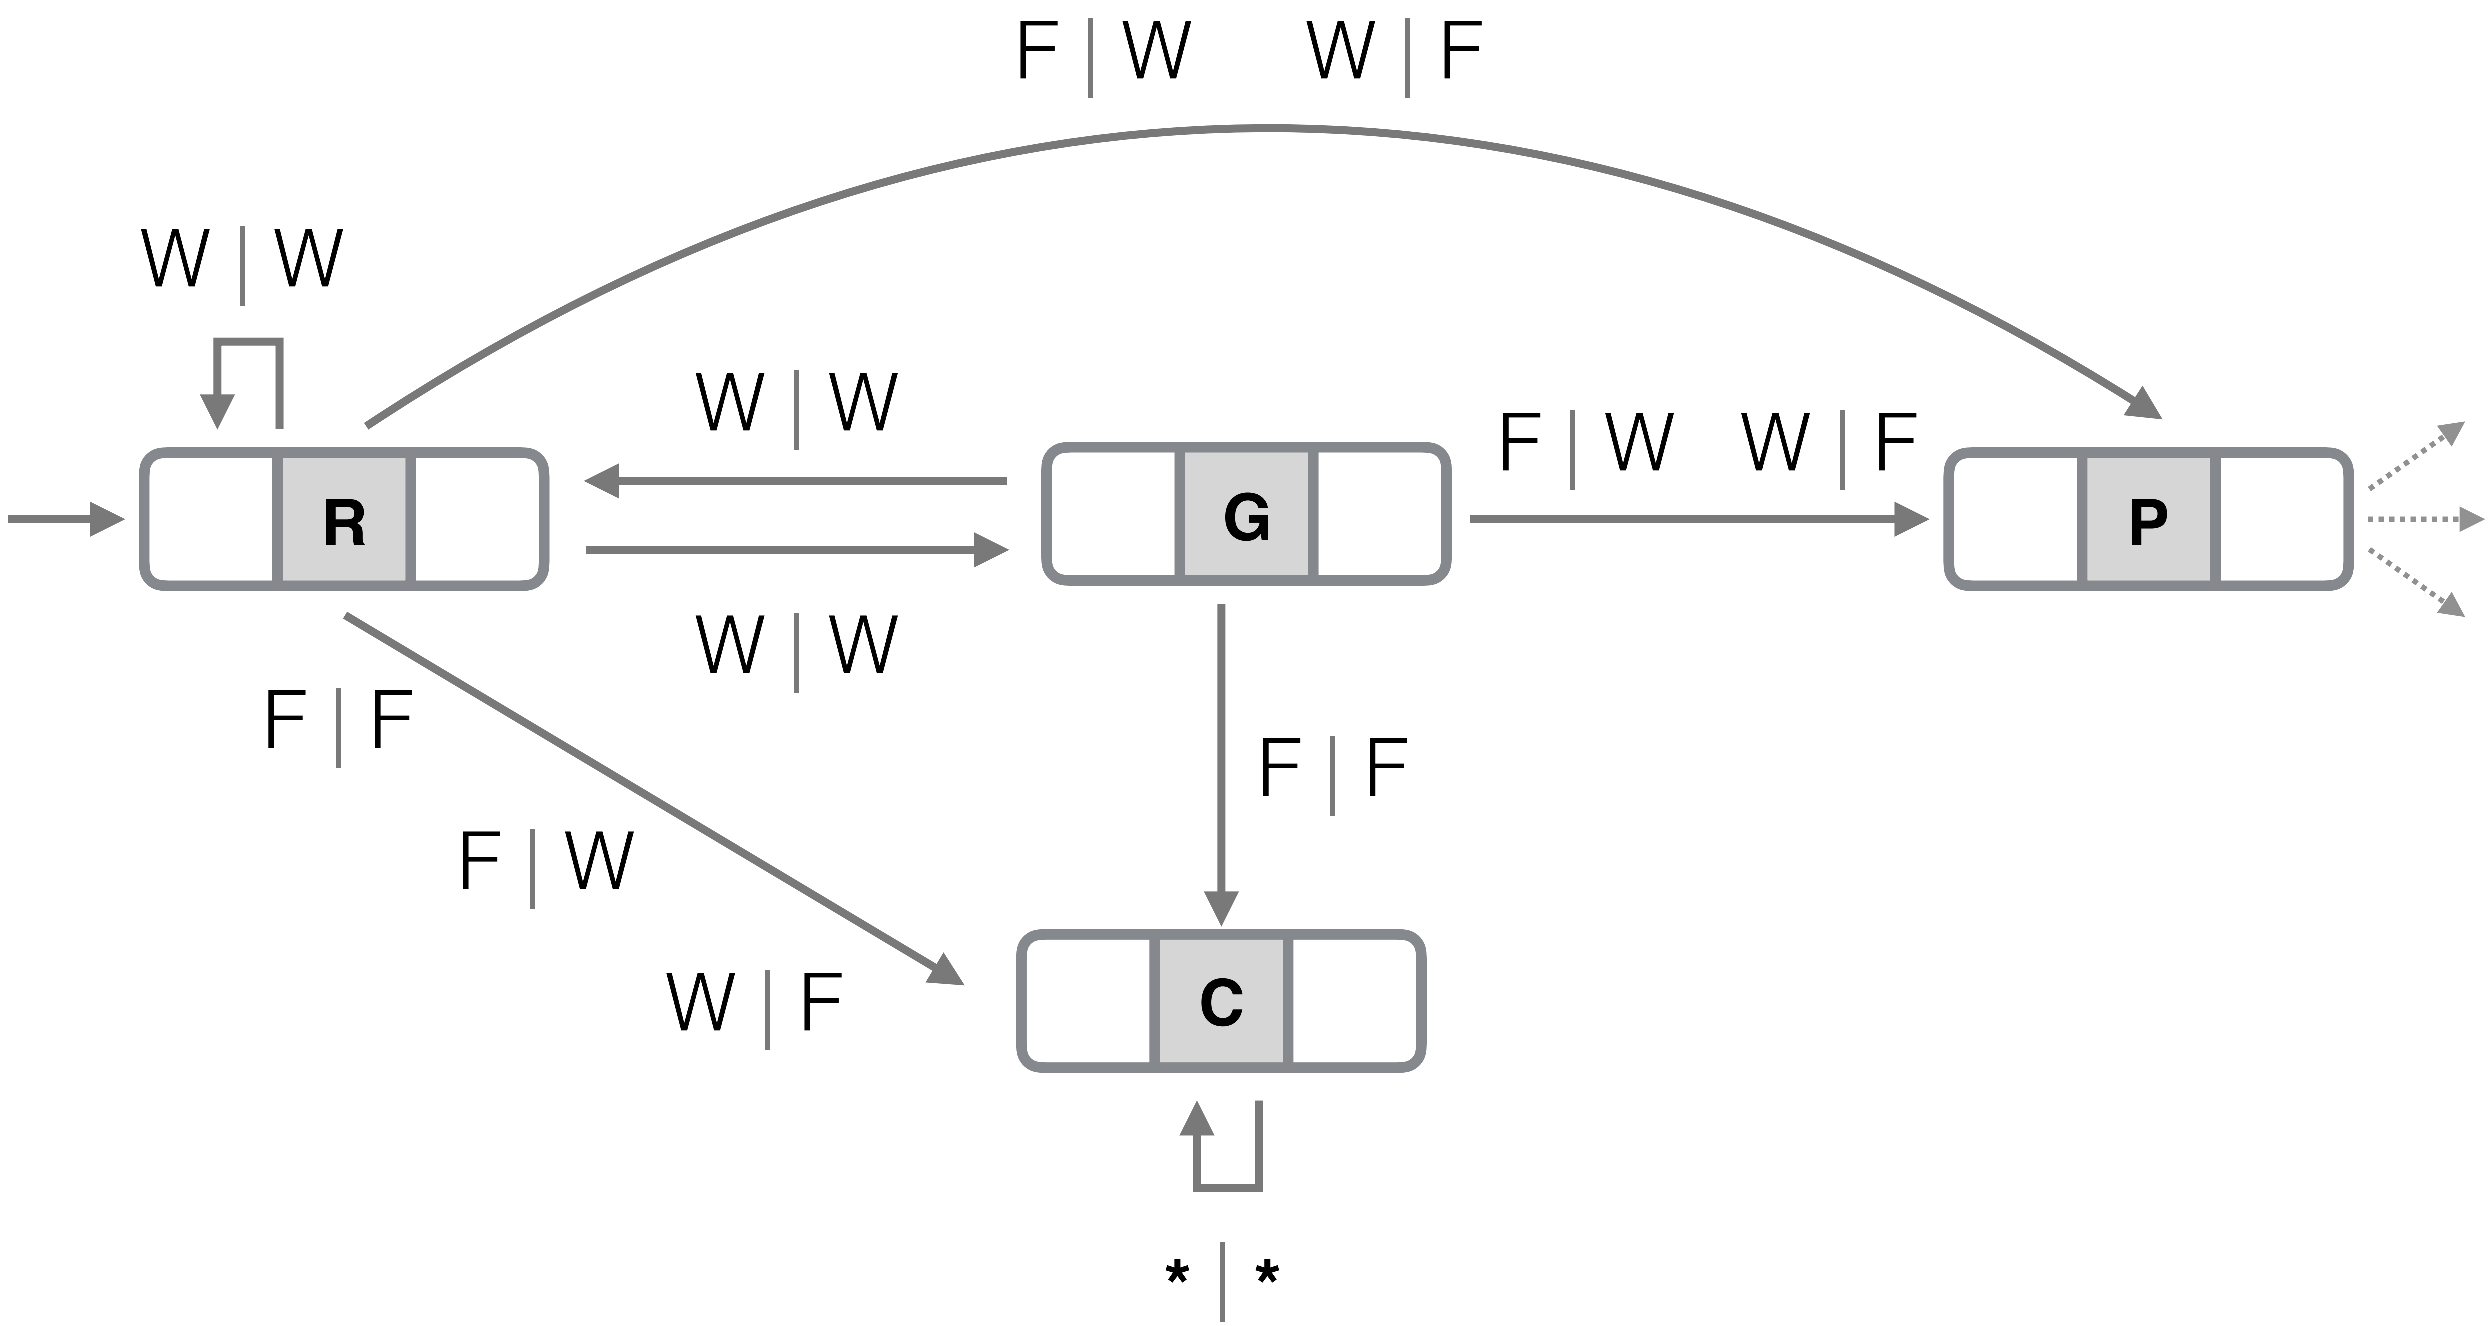

Games with Imperfect Information

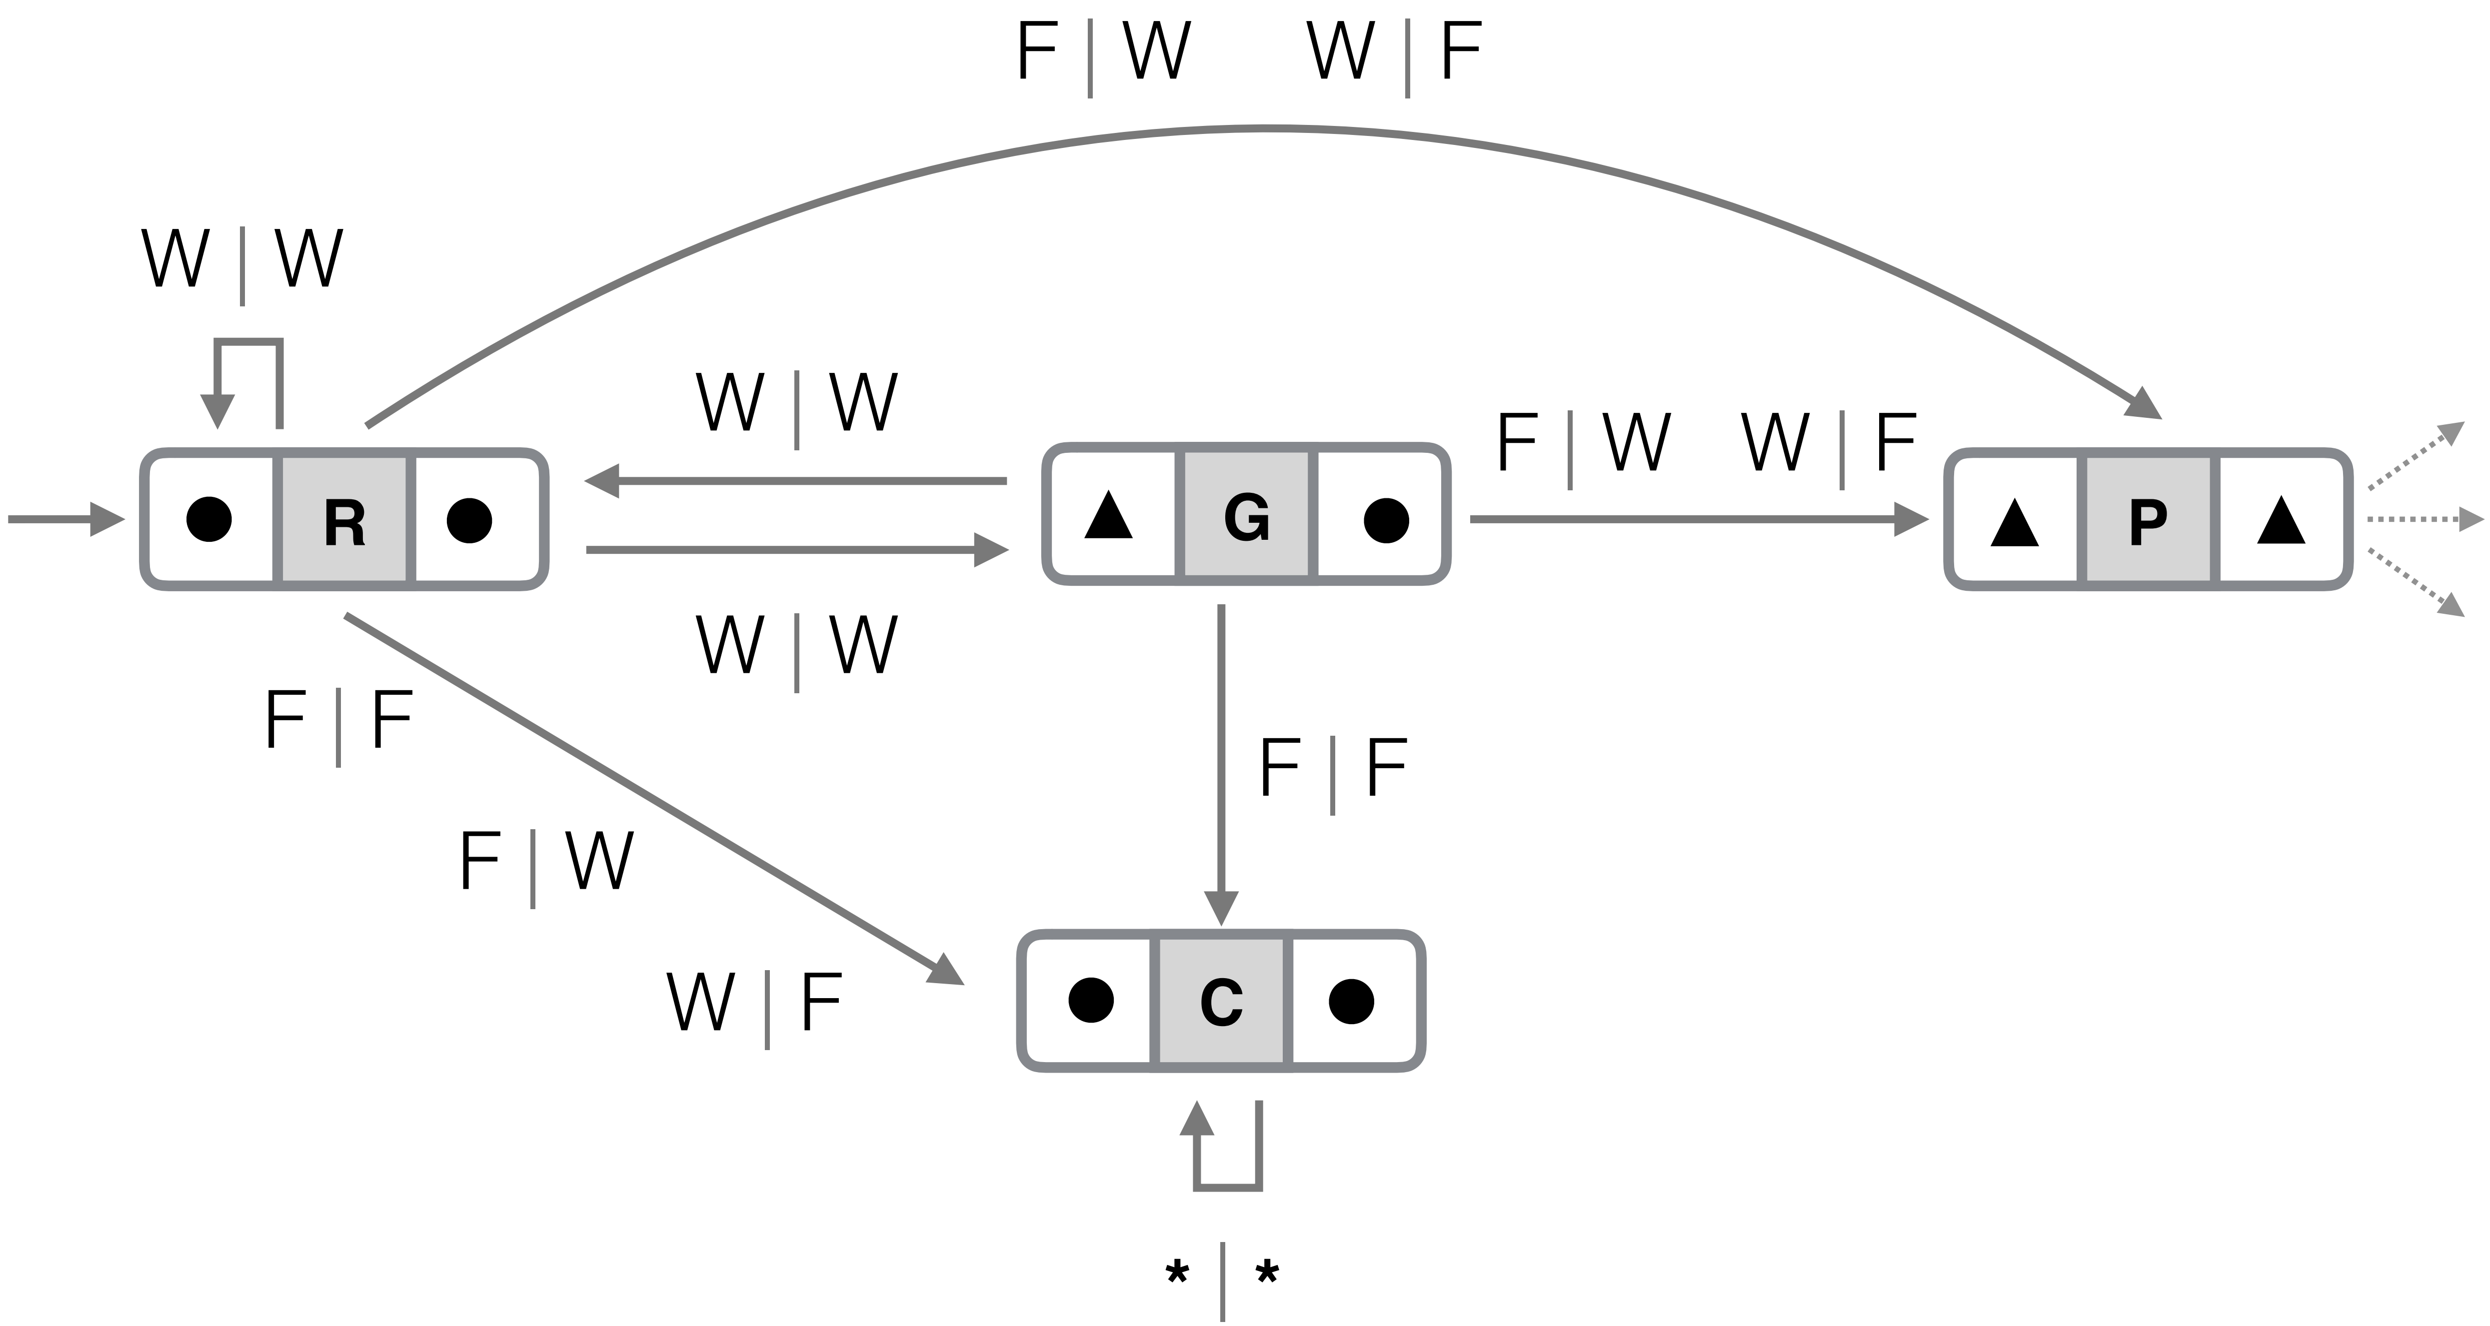

Observation sets: $B^1 = \{\bullet, \blacktriangle\}$ $B^2 = \{\bullet, \blacktriangle\}$

Games with Imperfect Information

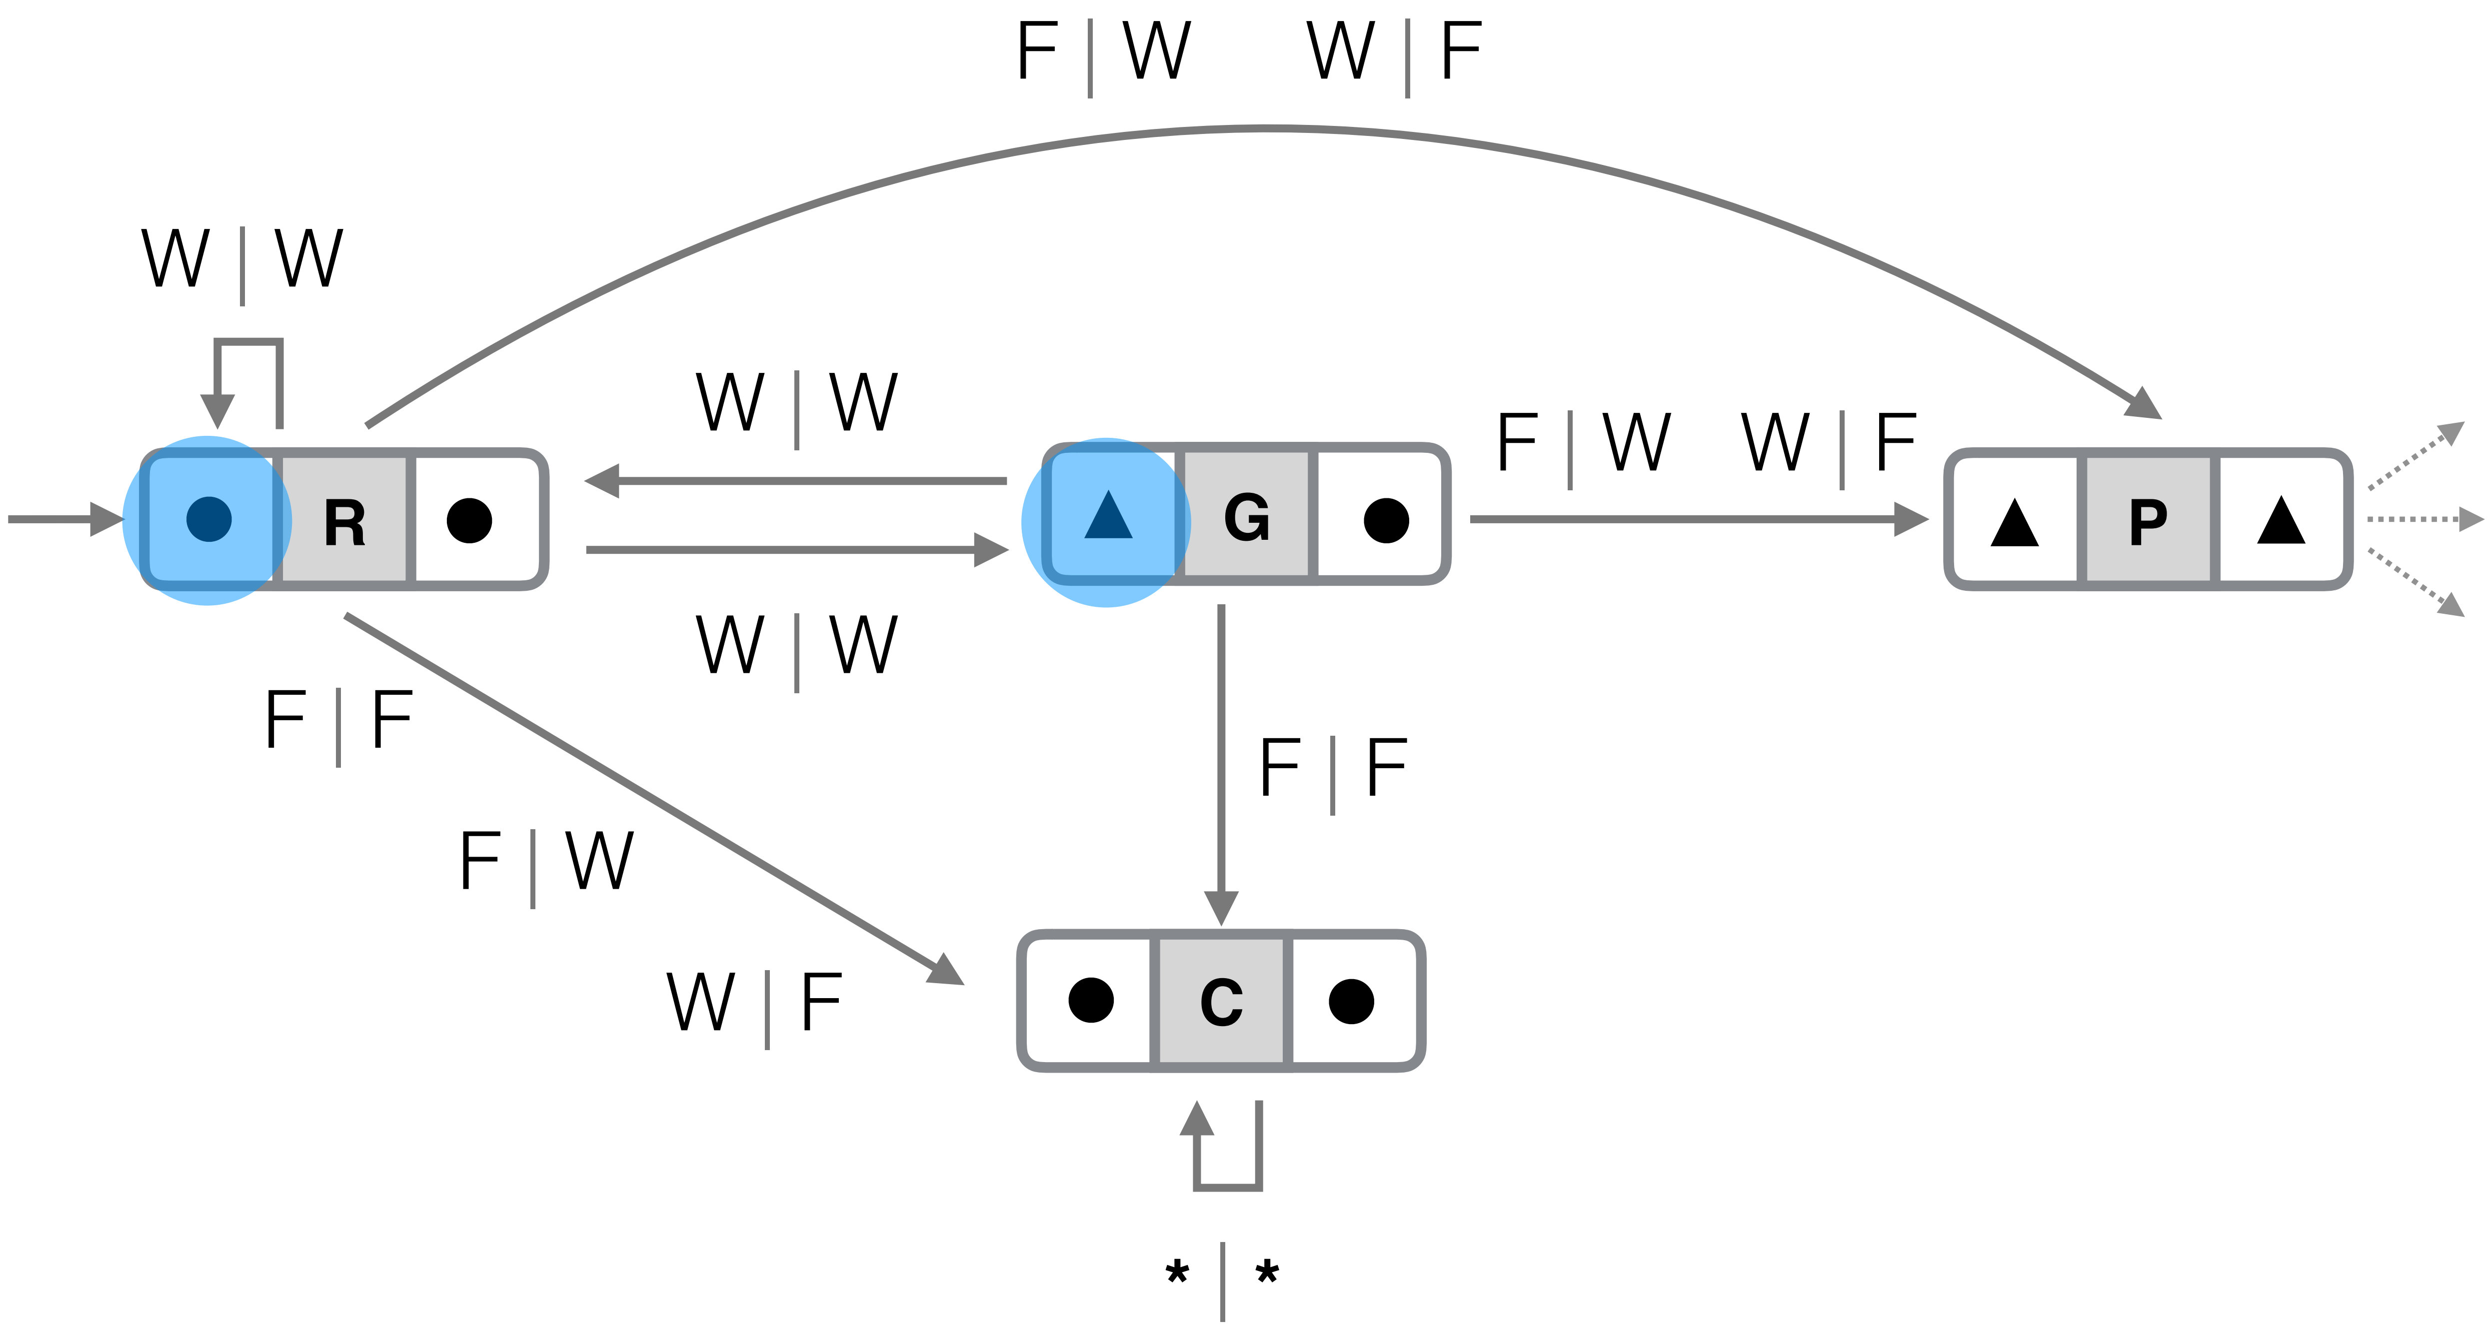

Indistinguishability: $\mathbf{v_L \not\sim^1 v_R}$

Games with Imperfect Information

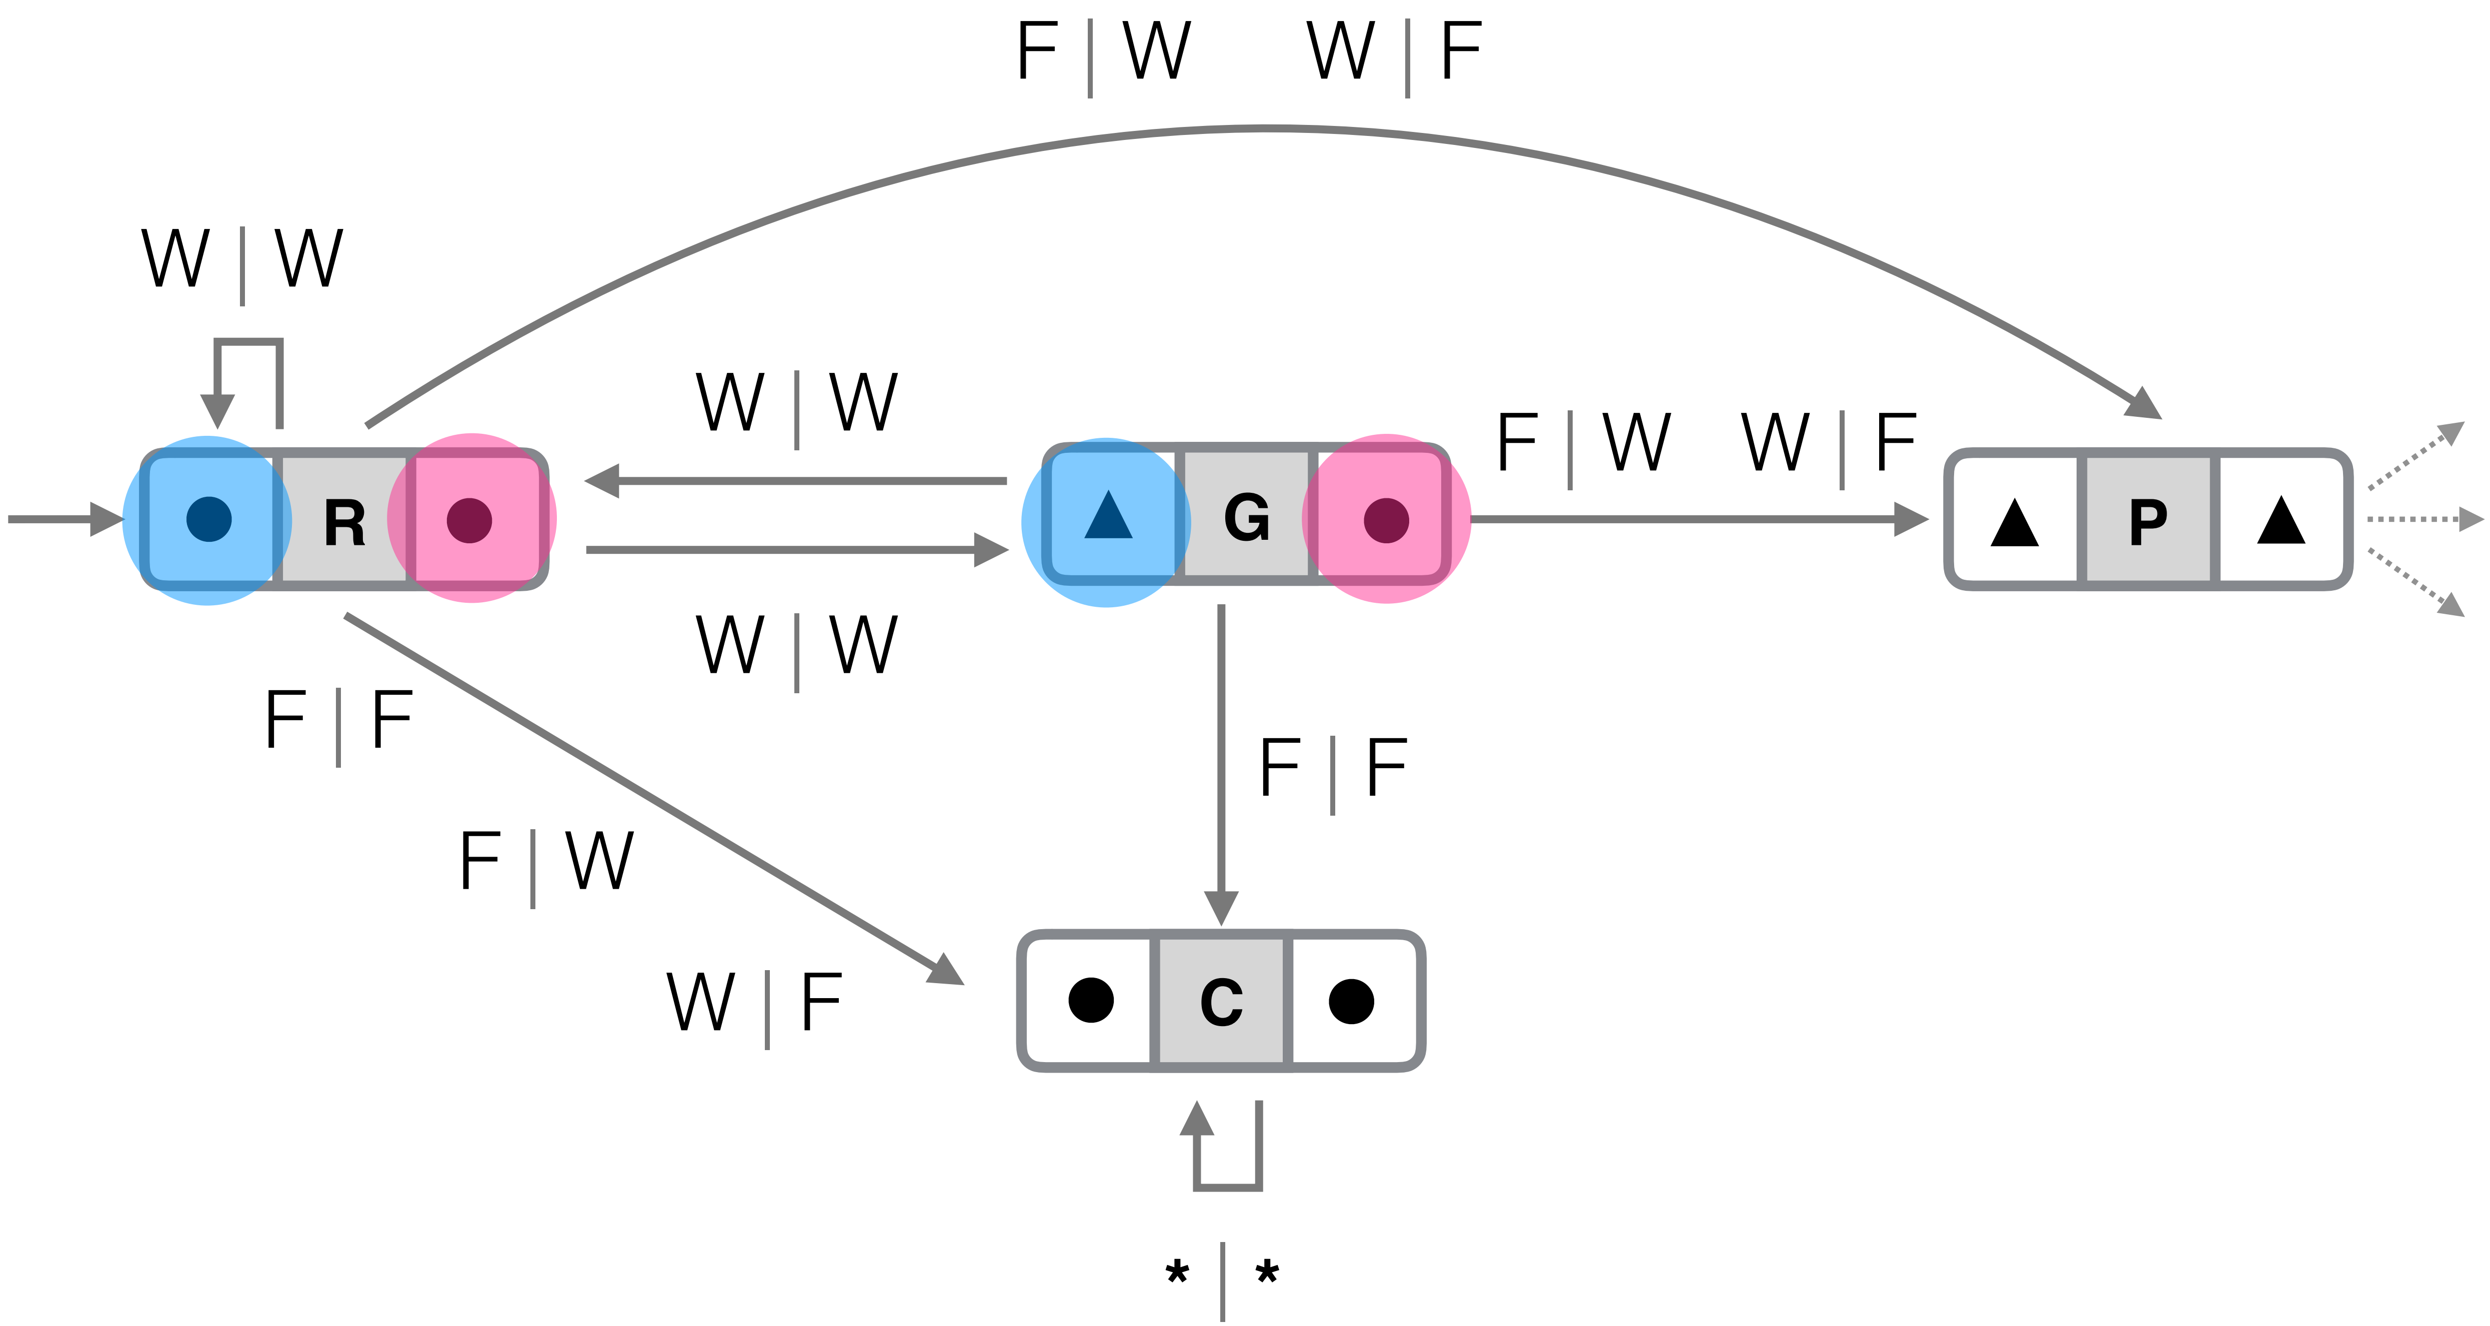

Indistinguishability: $v_L \not\sim^1 v_R$ $v_L \sim^2 v_R$

Winning condition:
« avoid C »

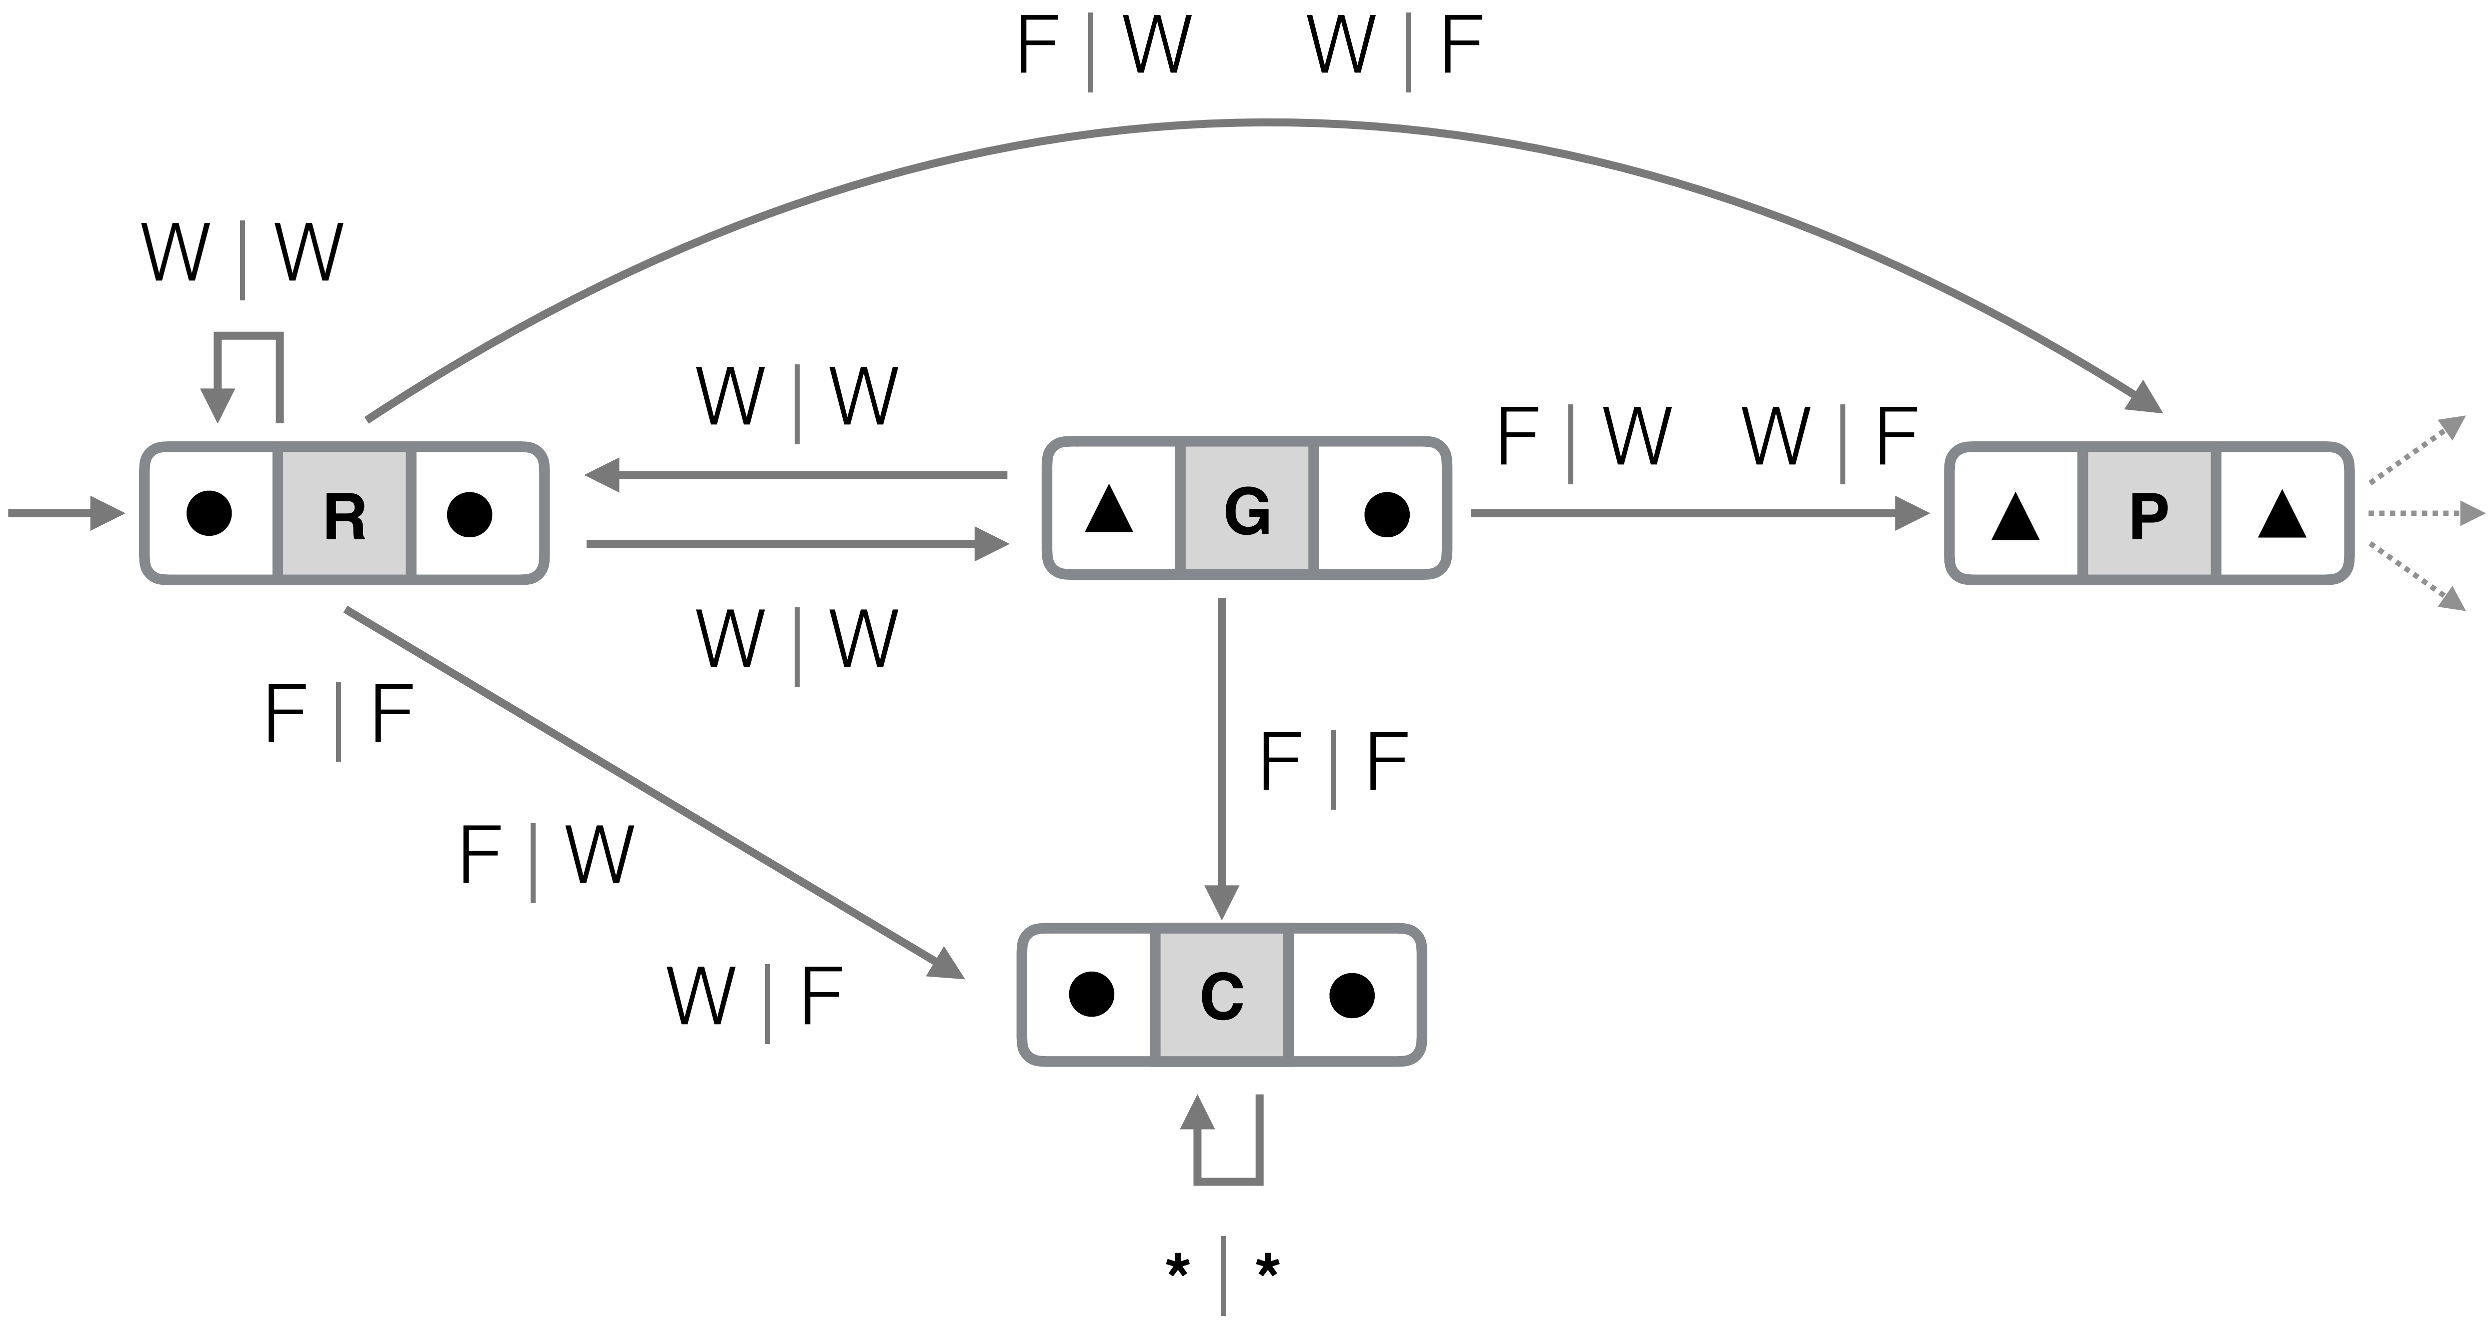

Indistinguishability: $v_L \not\sim^1 v_R$ $v_L \sim^2 v_R$

Winning condition:
« avoid C »

Strategies:
s^1: $V^* \to A^1$?
s^2: $V^* \to A^2$

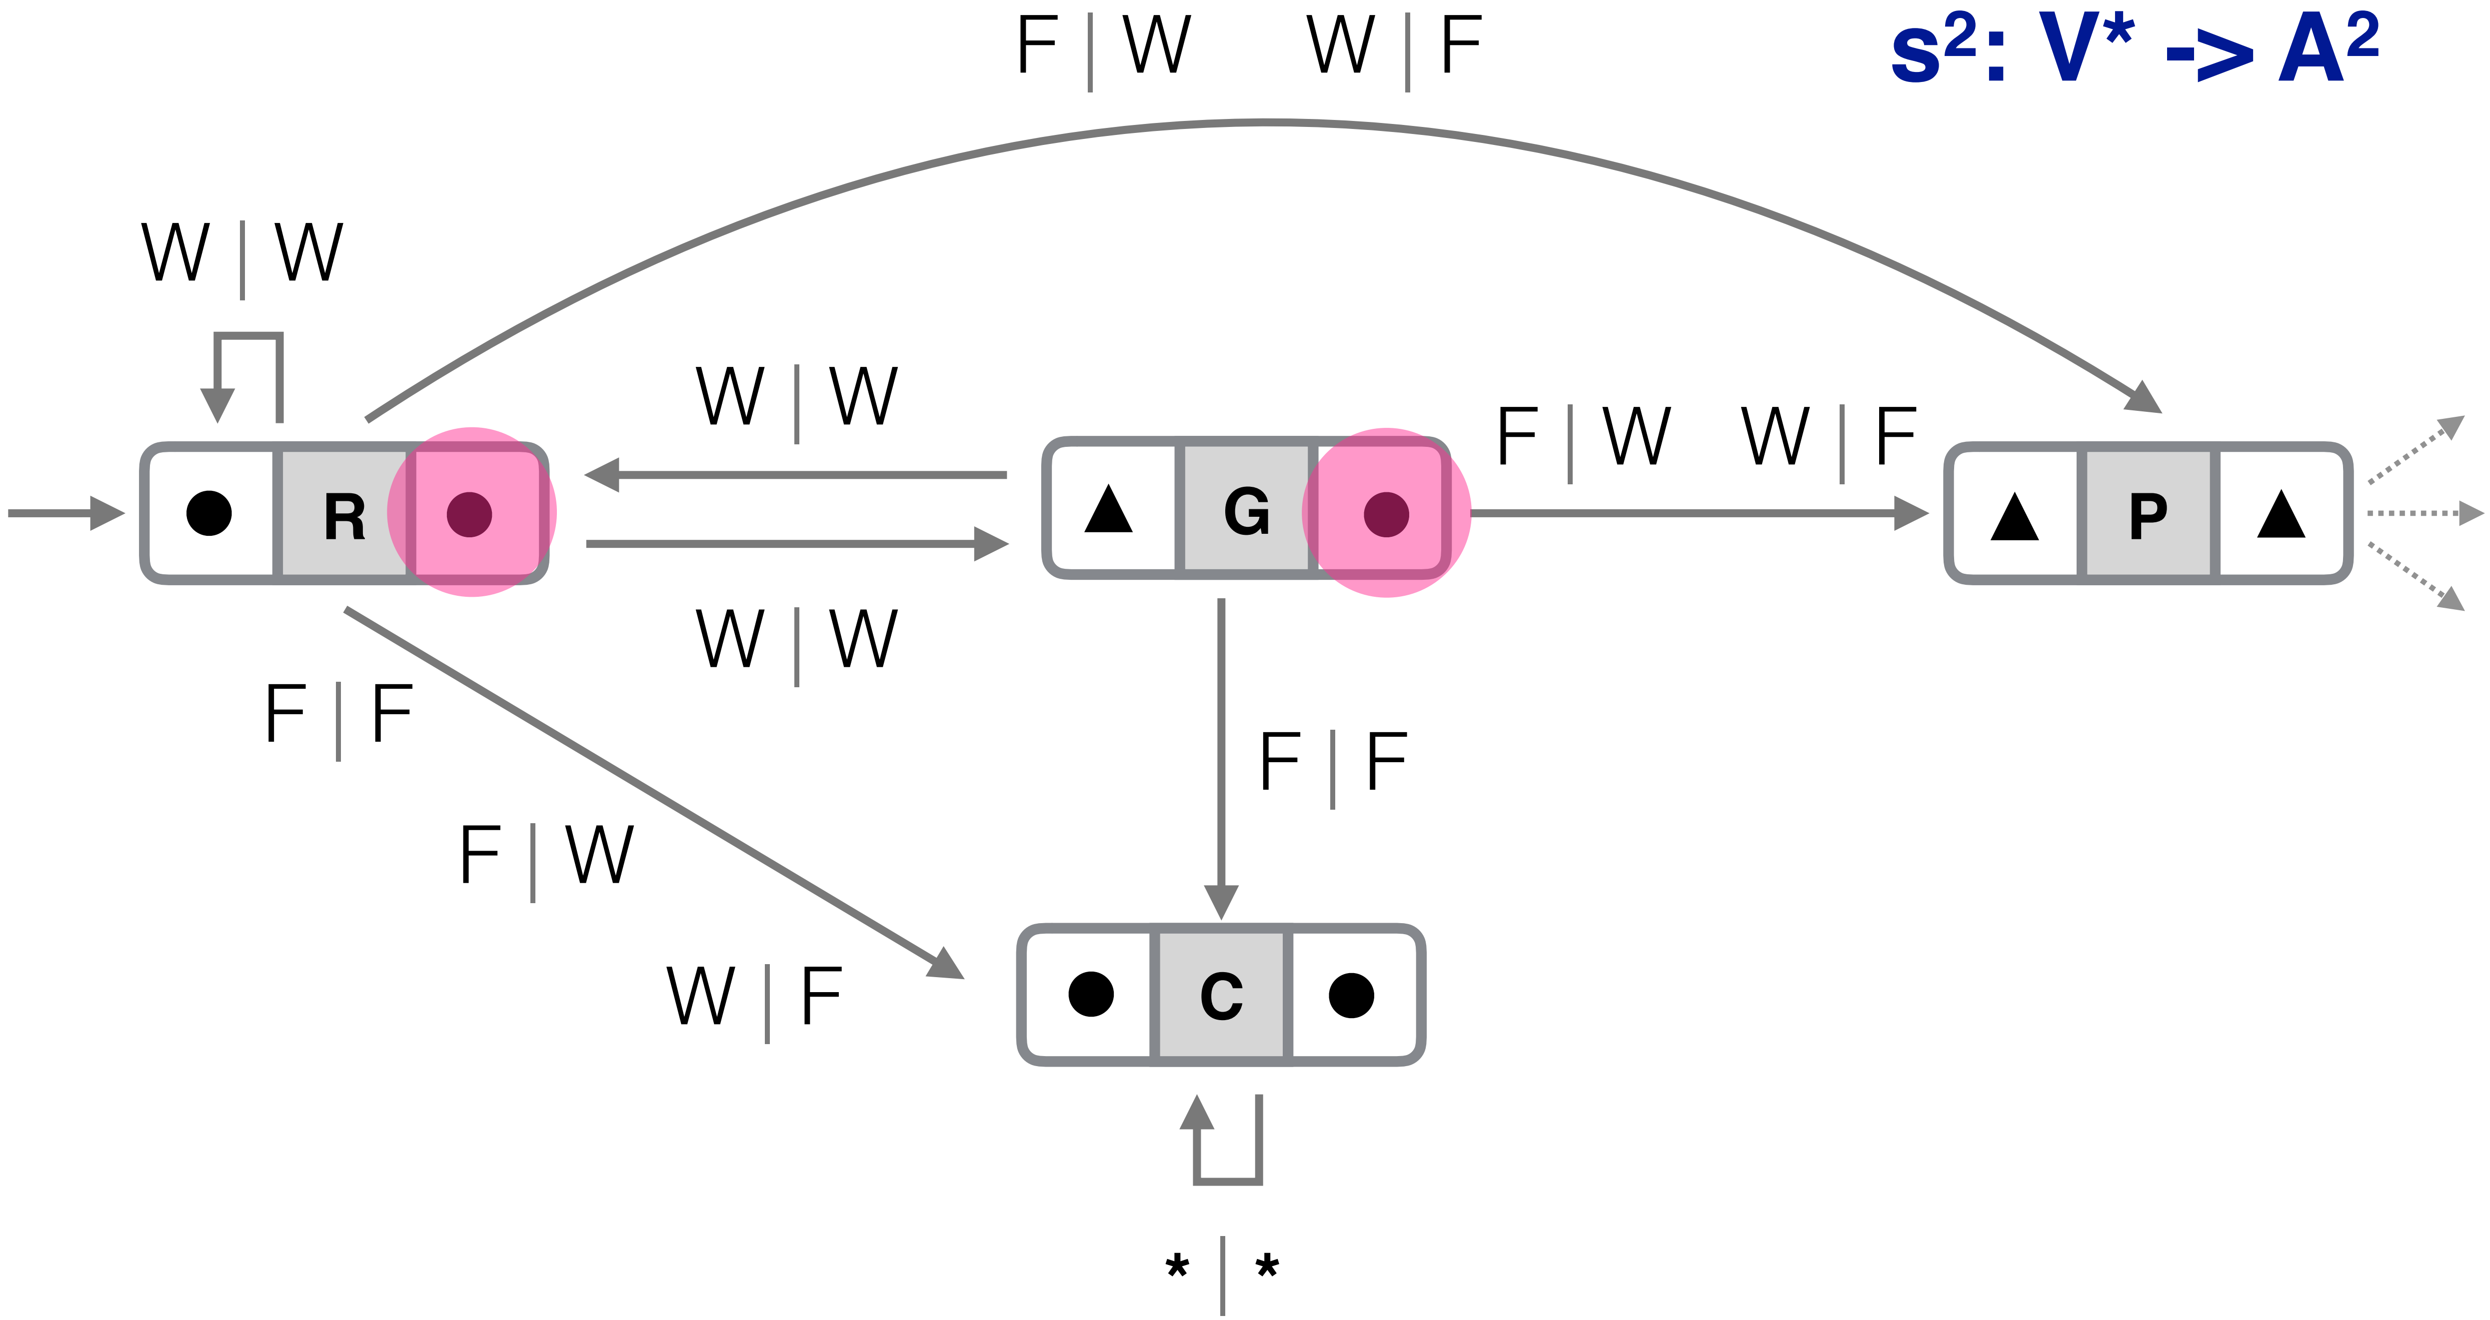

Indistinguishability: $v_L \not\sim^1 v_R$ $v_L \sim^2 v_R$

Winning condition:
« avoid C »

Strategies:
~~$s^1: V^* \to A^1$~~
~~$s^2: V^* \to A^2$~~ ?

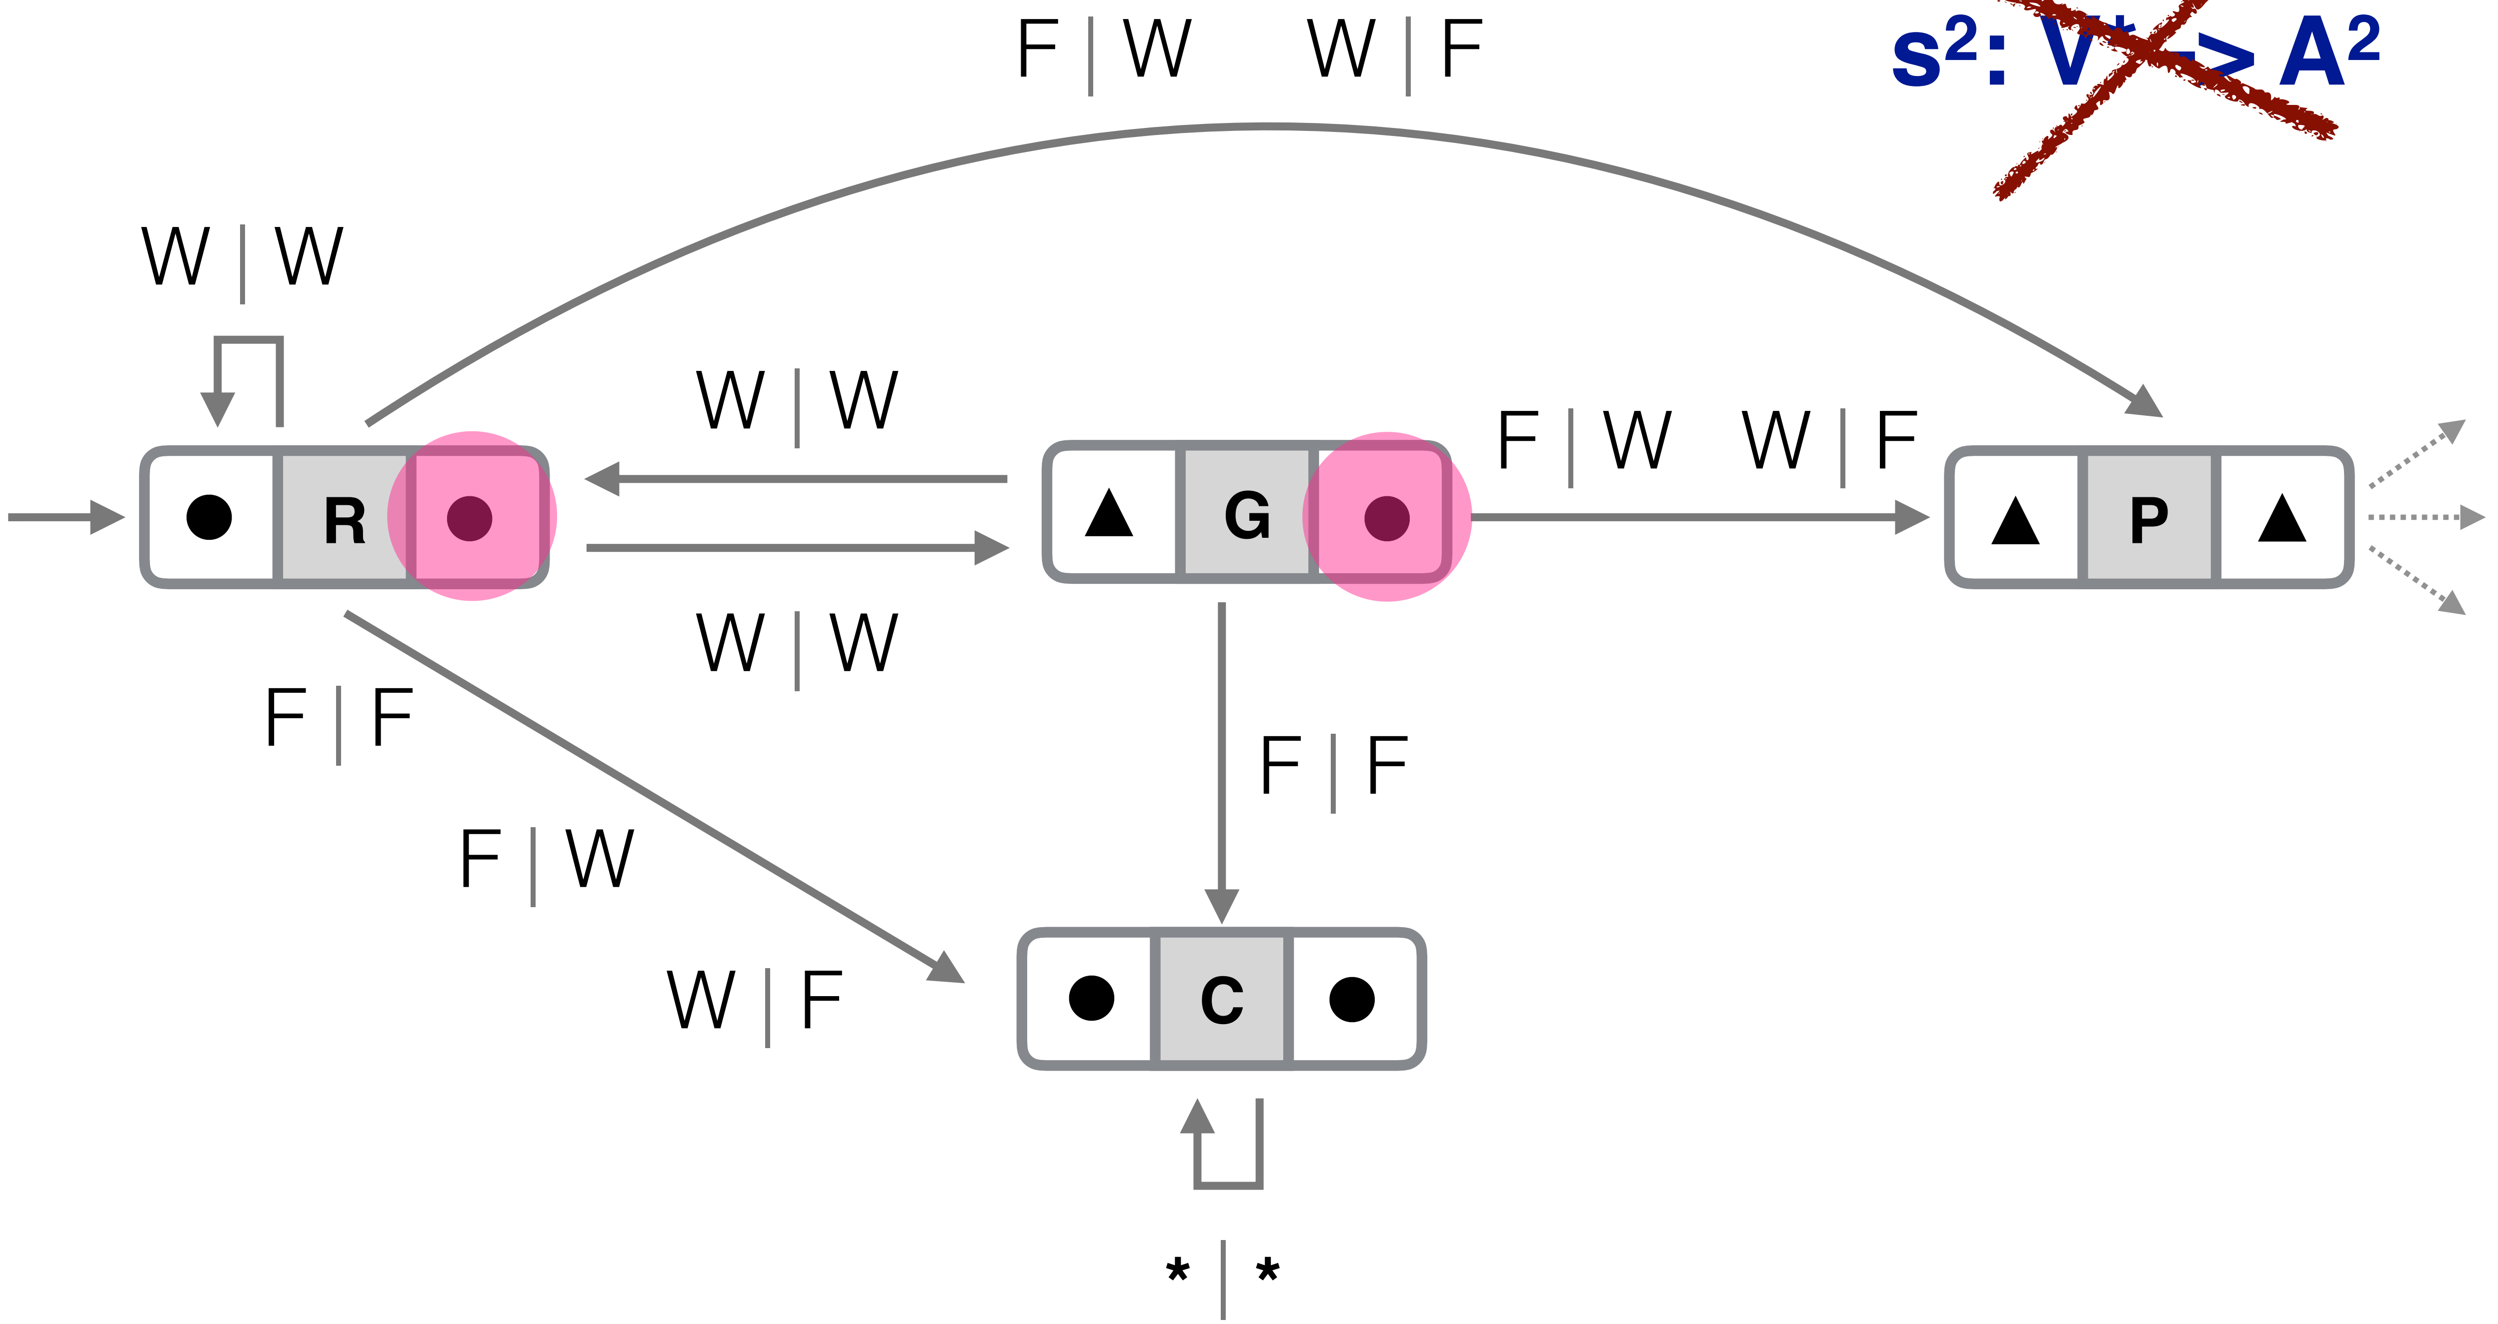

Indistinguishability: $v_L \not\sim^1 v_R$ $v_L \sim^2 v_R$

Winning condition: « avoid C »

Strategies:
s^1: B^{1*} -> A^1
s^2: B^{2*} -> A^2

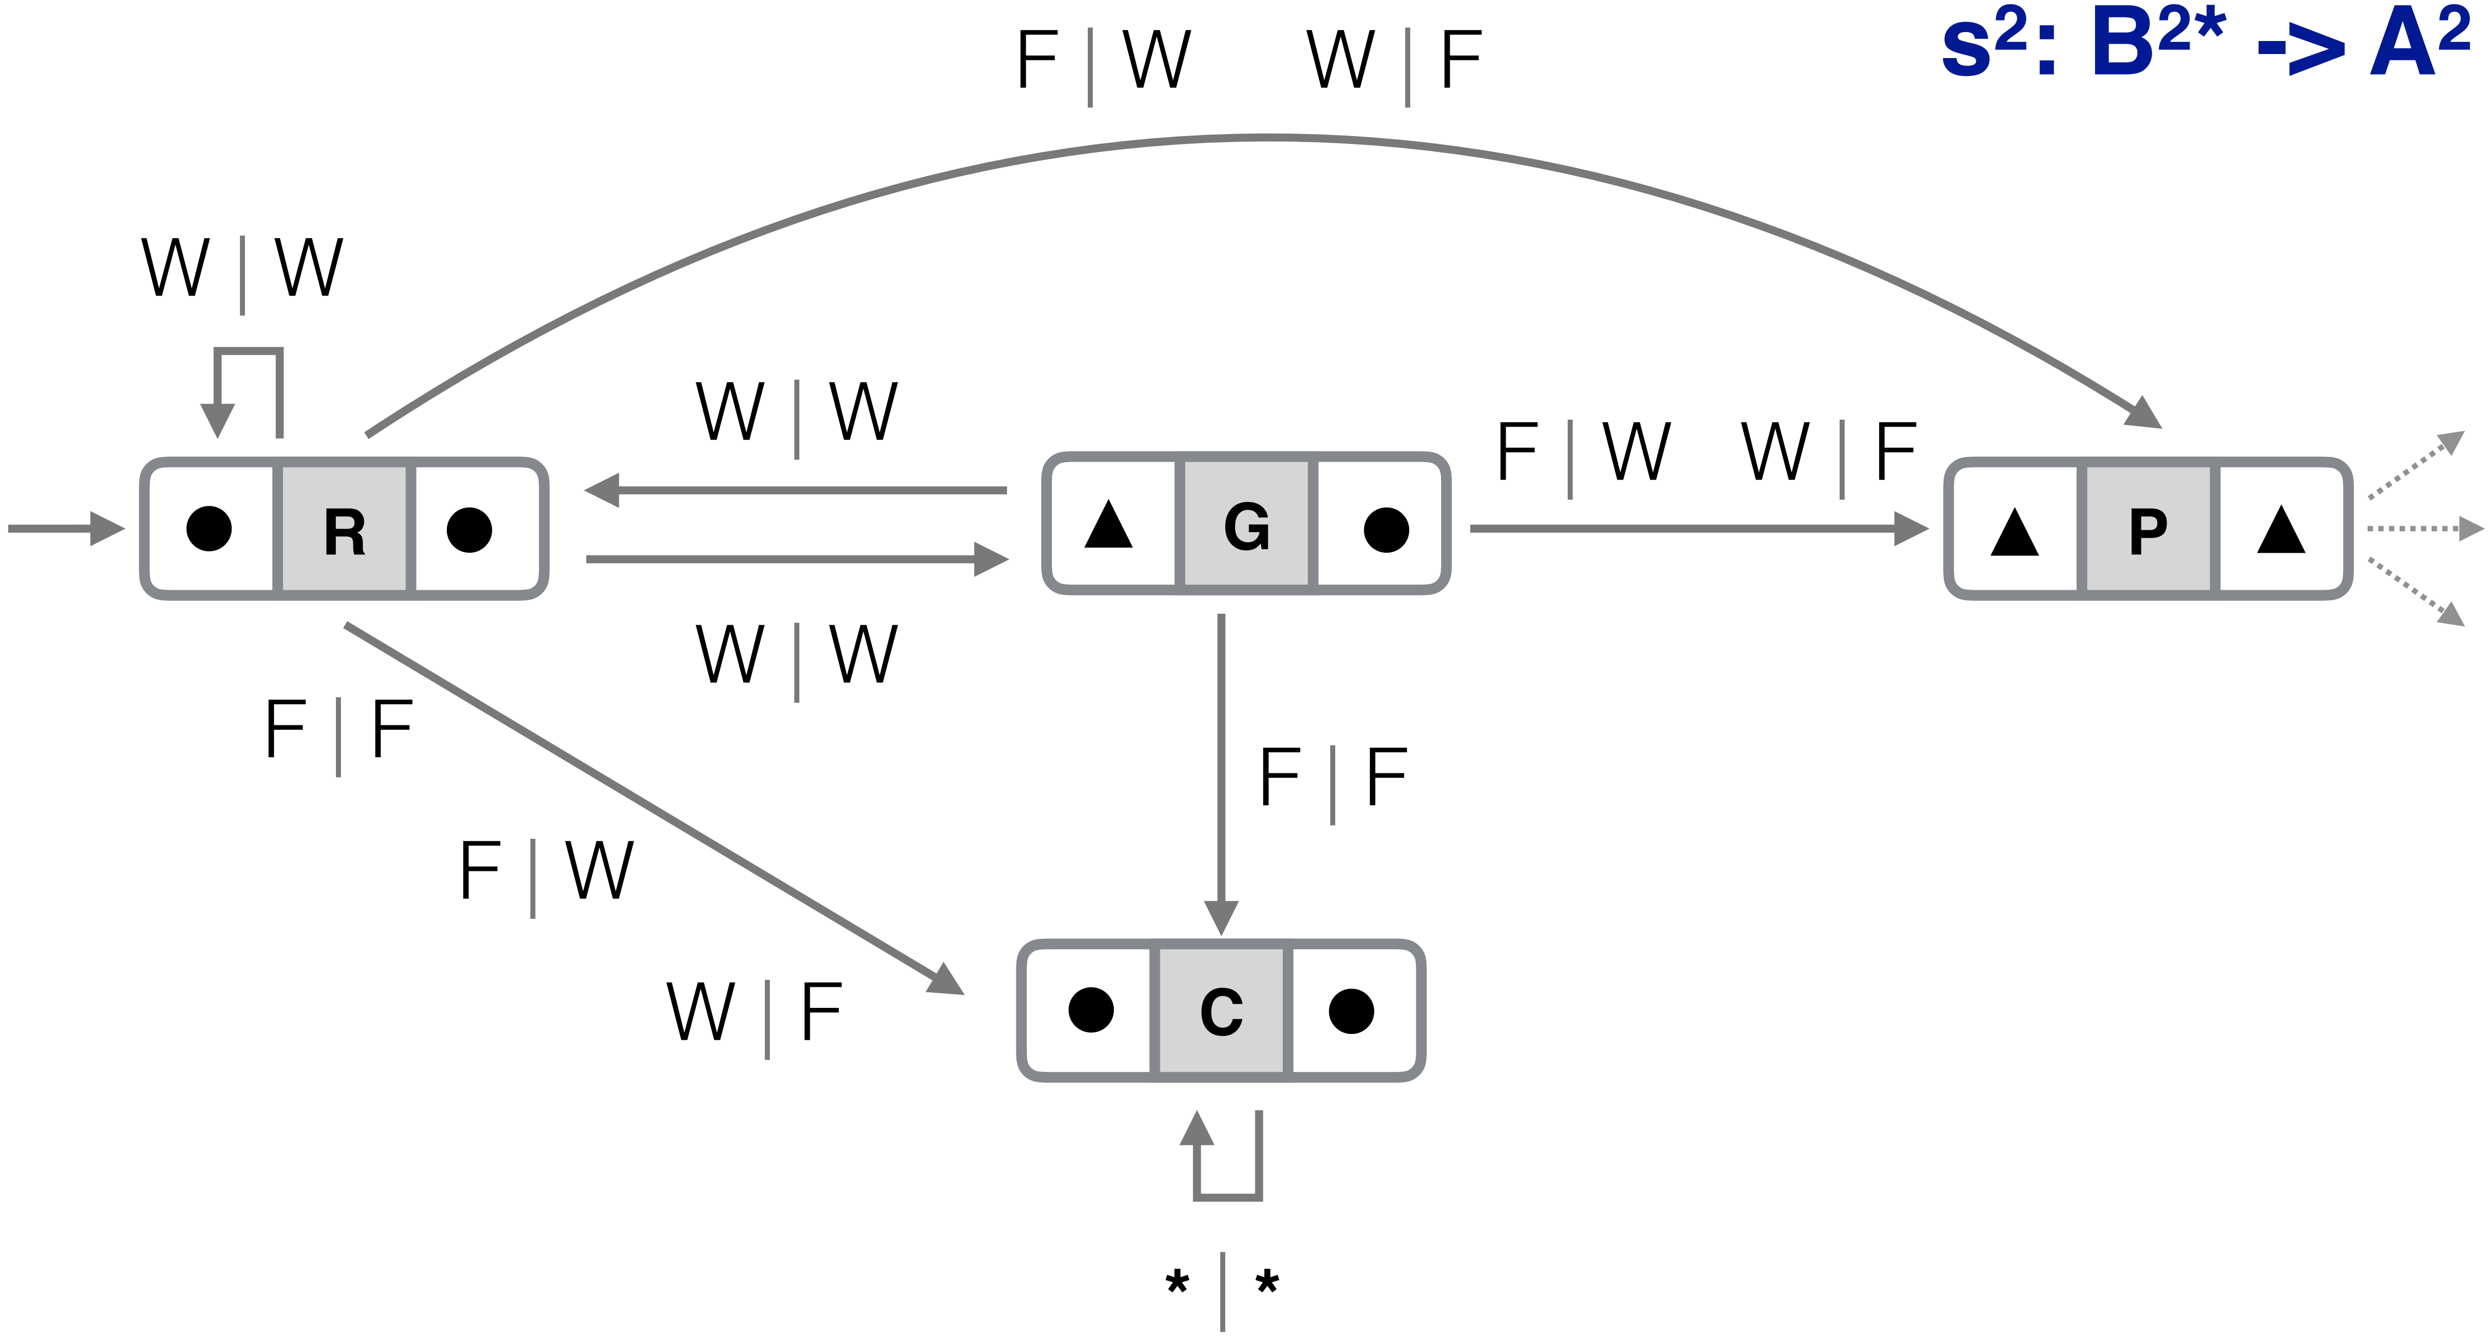

Indistinguishability: $v_L \not\sim^1 v_R$ $v_L \sim^2 v_R$

Winning condition:
« avoid C »

Strategies:
s^1: see ▲ play F
else play W
s^2: always play W

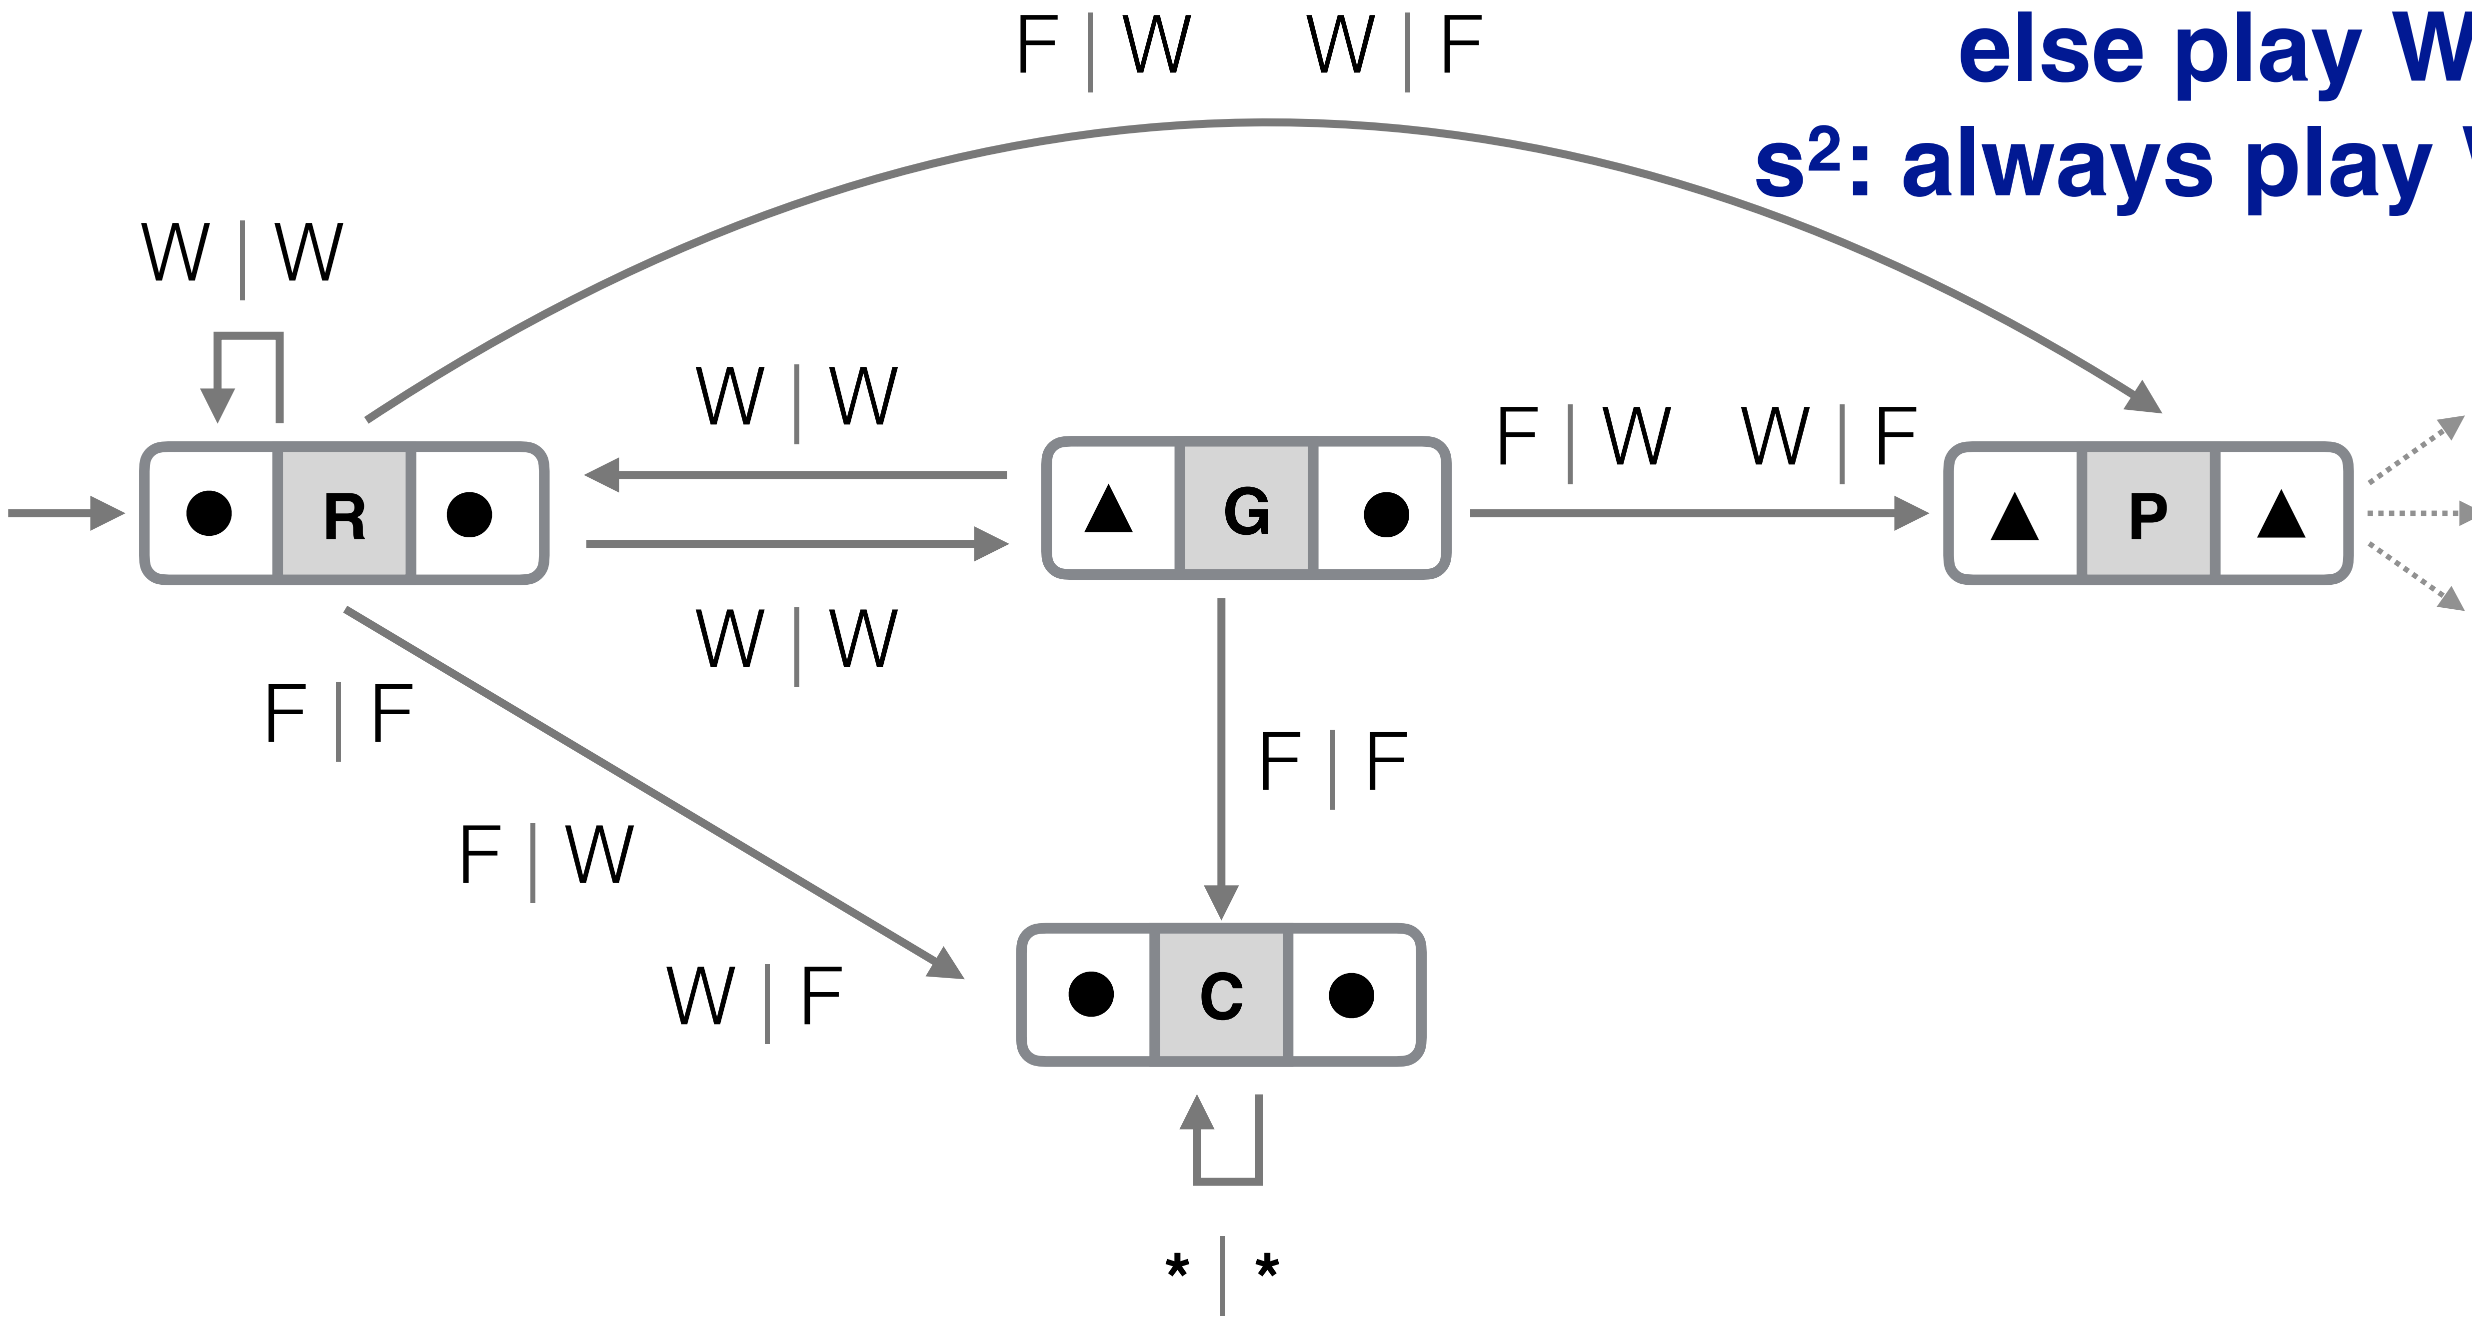

Indistinguishability: $v_L \not\sim^1 v_R$ $v_L \sim^2 v_R$

Winning condition:
« avoid C »

Strategies:
s^1: see ▲ play F
else play W
s^2: always play W

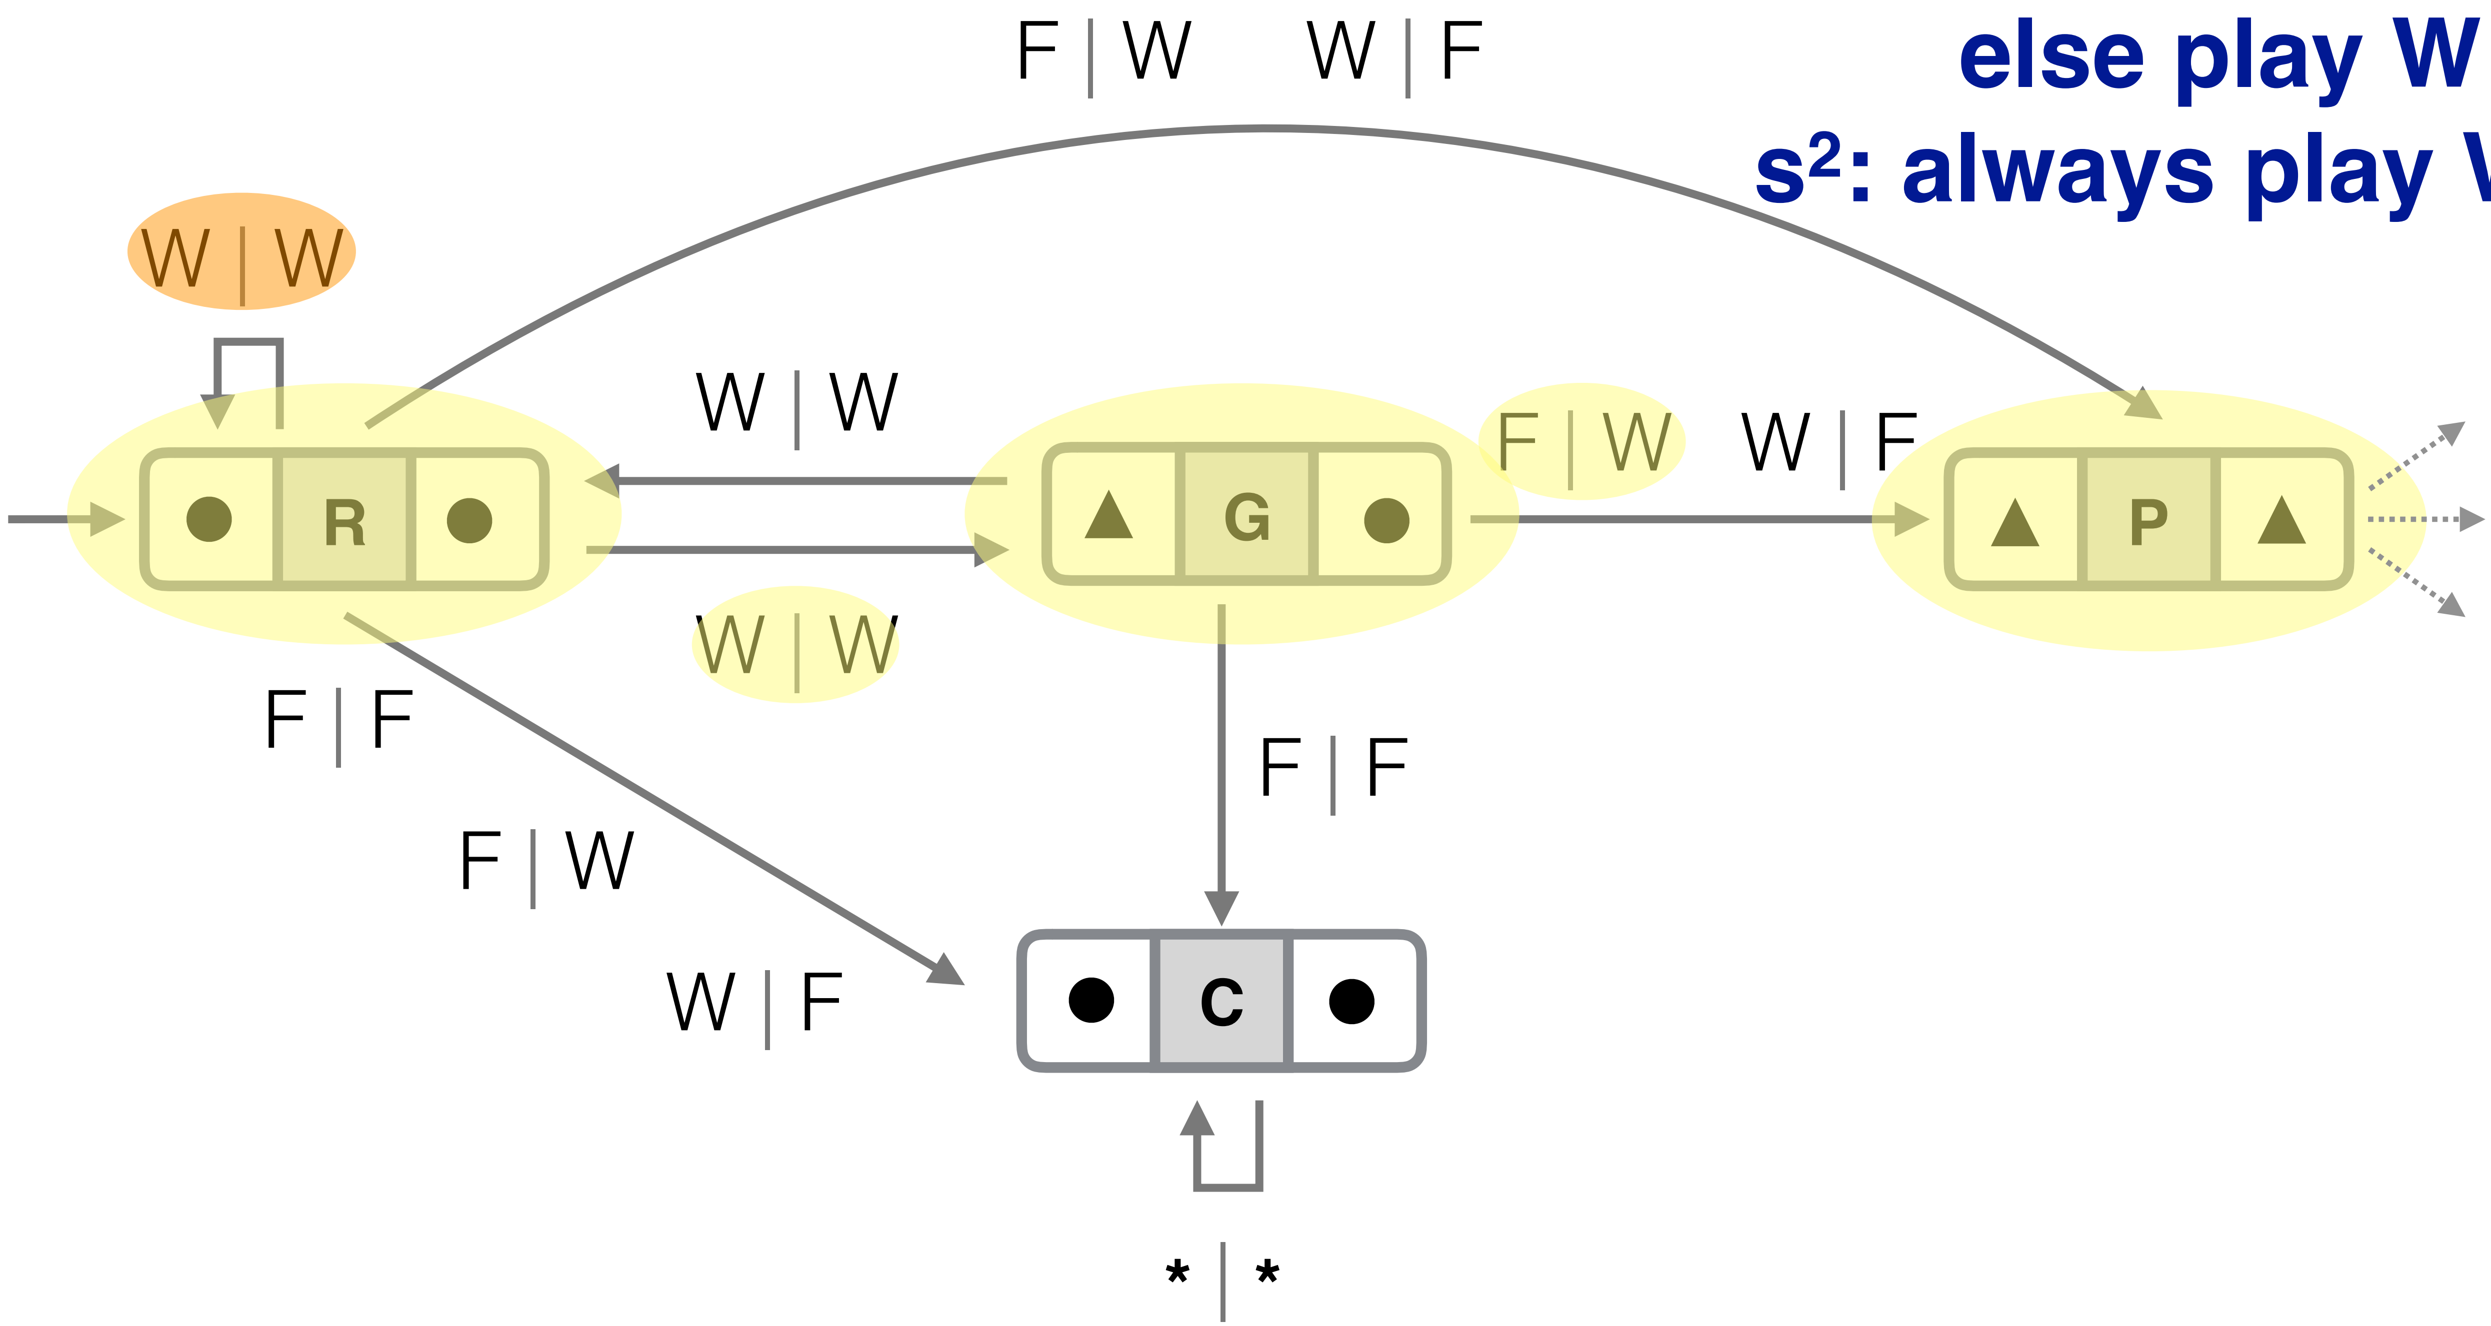

Winning condition:
« avoid C » ✓

Strategies:
s^1: see ▲ play F
else play W
s^2: always play W

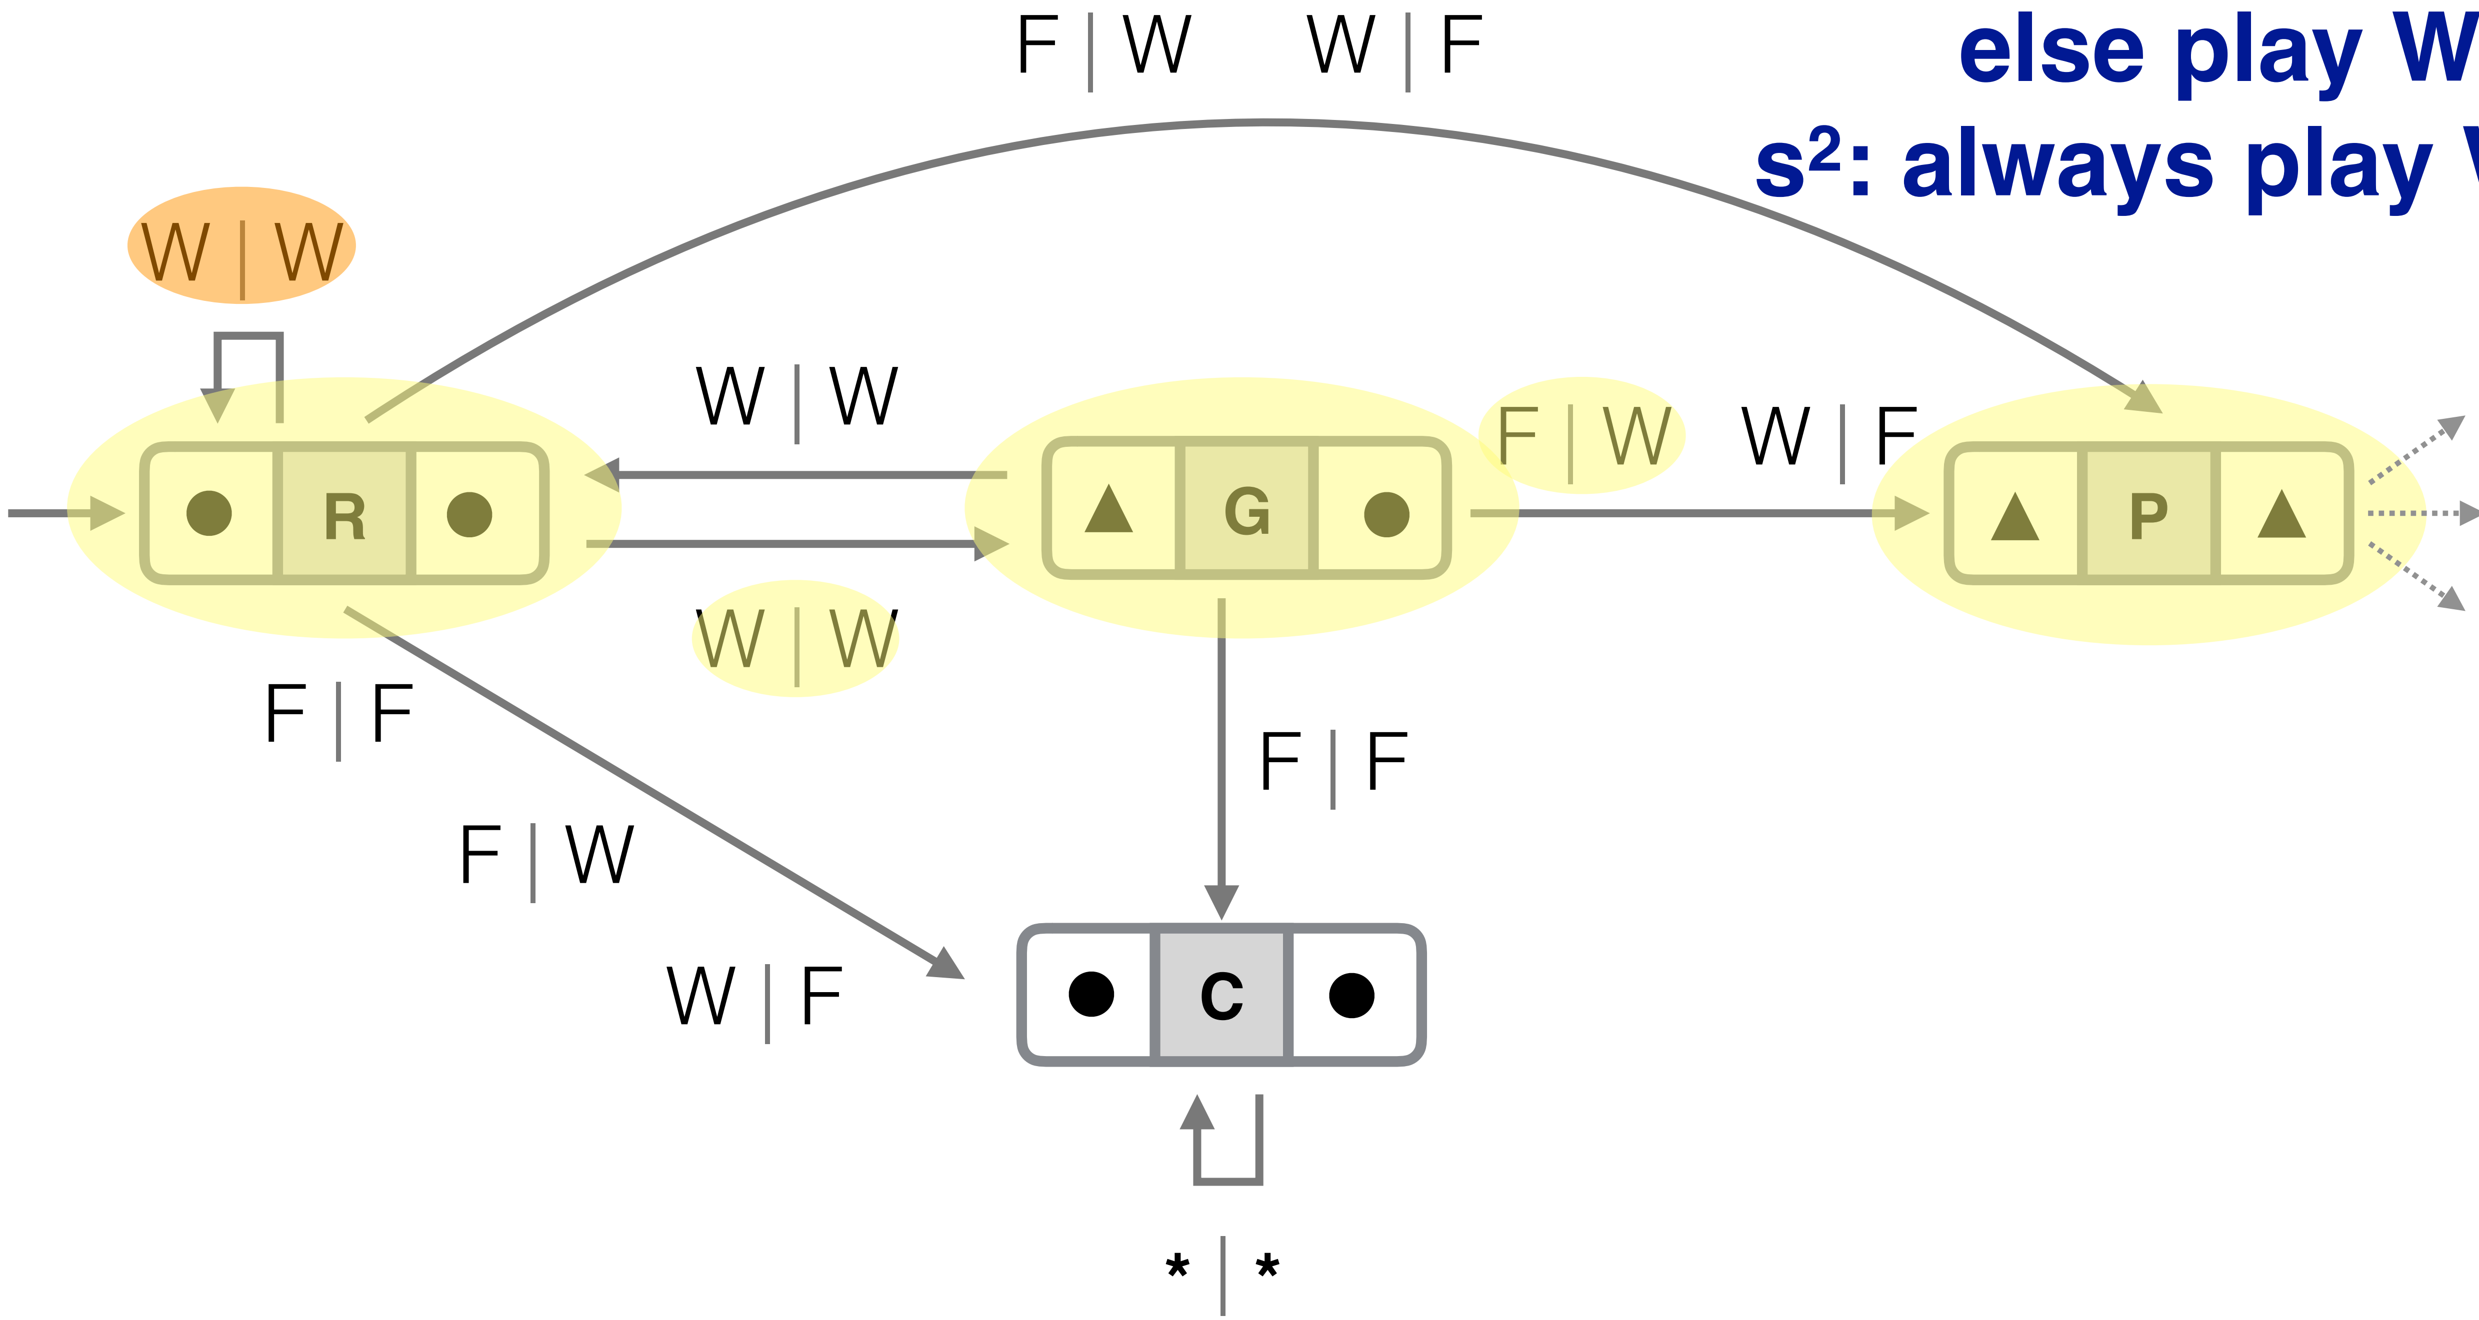

Games with Imperfect Information

$$G = (V, E, (\beta^i)_{i \leq n}, v_0, W)$$

- (V, E) finite graph (states, transitions)
- initial state v_0,
- plays/histories: infinite paths/finite prefixes
- observation functions $\beta^i : V \to B^i$

  induces *indistinguishability* relation $\sim^i$

Games with Imperfect Information

$$G = (V, E, (\beta^i)_{i \le n}, v_0, W)$$

- (V, E) finite graph (states, transitions)
- initial state v_0,
- plays/histories: infinite paths/finite prefixes
- observation functions $\beta^i : V \to B^i$
  induces *indistinguishability* relation $\sim^i$
- winning condition: ω-regular set of plays $W \subseteq V^\omega$

Games with Imperfect Information

$$G = (V, E, (\beta^i)_{i \leq n}, v_0, W)$$

- (V, E) finite graph (states, transitions)
- initial state v_0,
- plays/histories: infinite paths/finite prefixes
- observation functions $\beta^i : V \to B^i$
  induces *indistinguishability* relation $\sim^i$
- winning condition: ω-regular set of plays $W \subseteq V^\omega$
- *strategy*^{i}: function from observation histories to actions
- distributed strategy: $s = (\mathit{strategy}^1, \ldots, \mathit{strategy}^n)$

Distributed strategy synthesis

Challenge: coordinating players towards a *common objective*

Distributed strategy synthesis problem:

1) *Solvability:*

Given a game and a winning condition,
does there exist a distributed winning strategy?

2) *Implementability:*

If the game is solvable, construct
a distributed winning strategy with finite memory.

Background Results

Theorem[Peterson,Reif '79, Pnueli,Rosner '90, Janin '07]

Distributed strategy synthesis is **undecidable**
for games with imperfect information.

Theorem[Peterson,Reif 79, Pnueli,Rosner 90]

Decidable case: **hierarchical** games.

Proposition[Berwanger, Kaiser, Puchala '11]

Distributed winning strategies may require **infinite** memory.

Overview of the Thesis

Approach: better understand interaction mechanisms via analysis of *information flow*

1. Propagation of uncertainty (*why is it so hard?*)
   joint work with D. Berwanger, DLT'15
2. One-Way Information Flow (*decidable cases*)
   joint work with D. Berwanger and A. B. Mathew, ATVA'15
3. Delayed Information Flow (*realistic scenarios*)
   joint work with D. Berwanger, FSTTCS'15

Propagation of uncertainty

General case: Distributed strategy synthesis is **undecidable**
[Peterson,Reif '79, Pnueli,Rosner '90, Janin '07]

But: many causes

- imperfect information and Nature
- arbitrarily many decision points
- complicated winning conditions

Propagation of uncertainty

General case: Distributed strategy synthesis is **undecidable**
[Peterson,Reif '79, Pnueli,Rosner '90, Janin '07]

But: many causes

- imperfect information and Nature
- arbitrarily many decision points
- complicated winning conditions

Propagation of uncertainty

Consensus Game Acceptors:

Two players against Nature

Play:

- Nature picks a **finite** path π
  from initial to final state $\oplus$, $\ominus$ or $\circledast$
- each player sees his observations along π
- at final state, players accept or reject

Winning condition:

- consensus
- path end at $\oplus$, play *accept*, at $\ominus$ play *reject*!

Propagation of uncertainty

$A^1 = \{ A, R \}$

$B^1 = \{ \bullet, \blacktriangle, \blacksquare \}$

$A^2 = \{ A, R \}$

$B^2 = \{ \bullet, \blacktriangle \}$

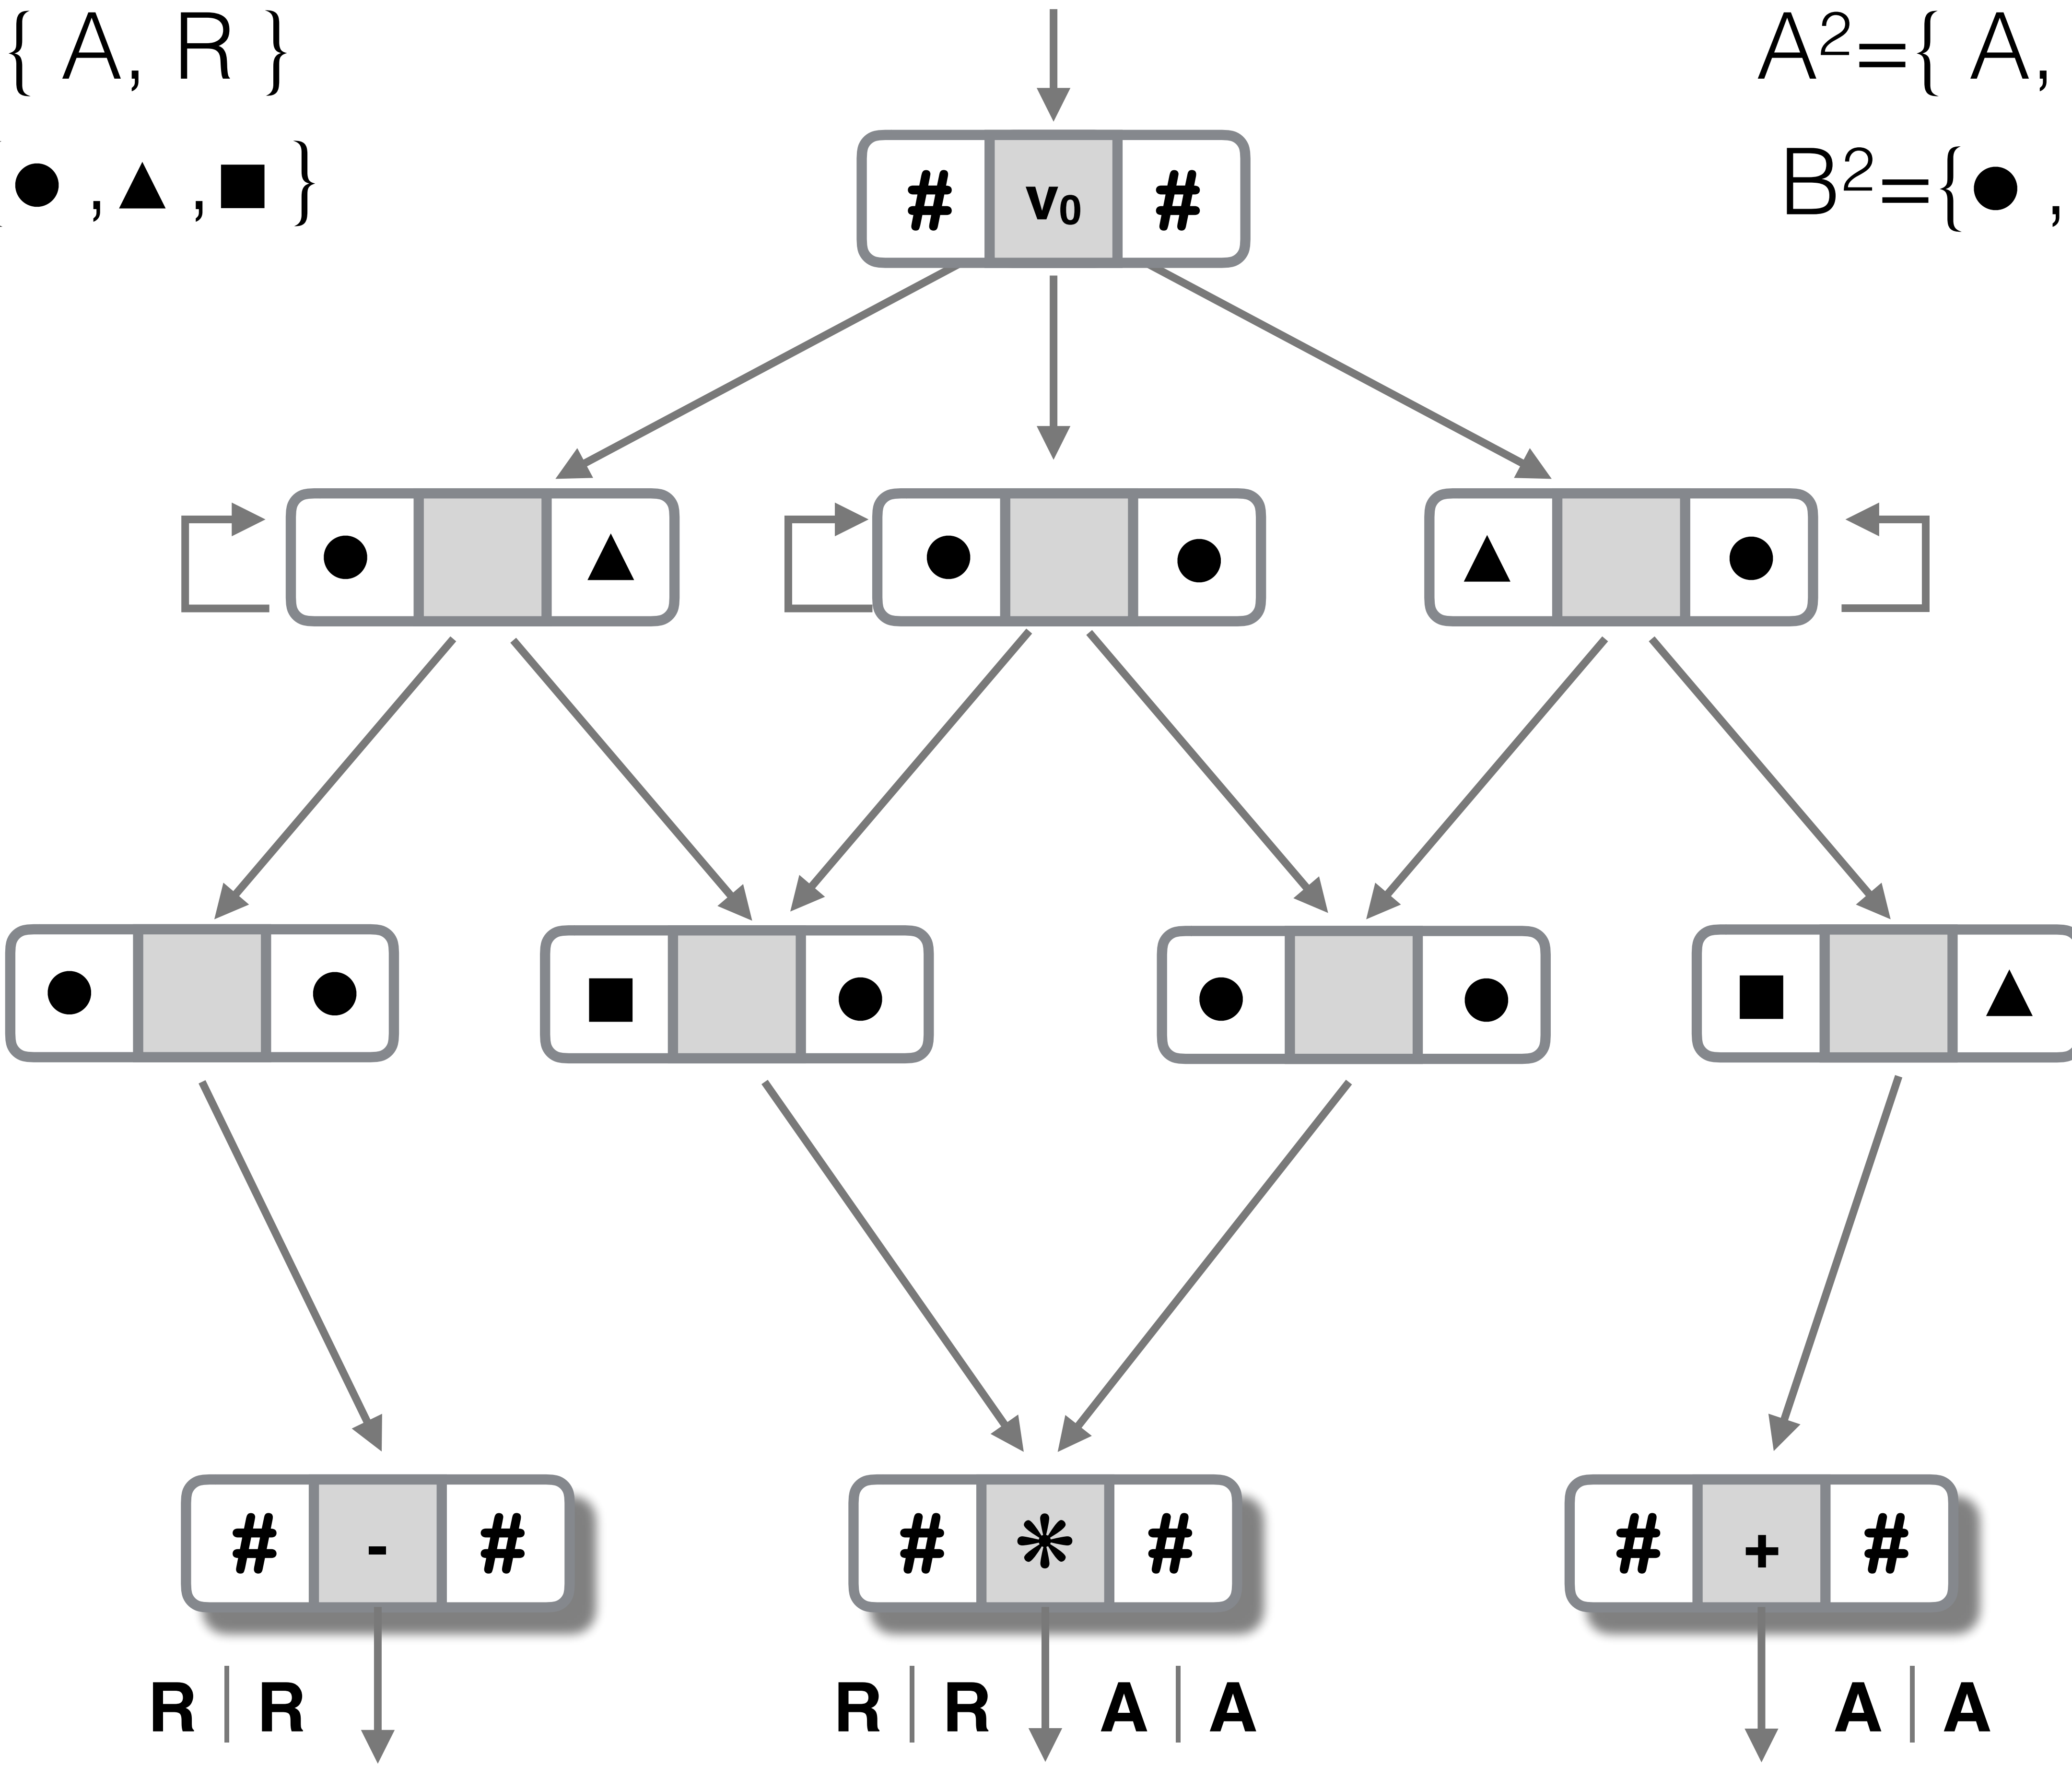

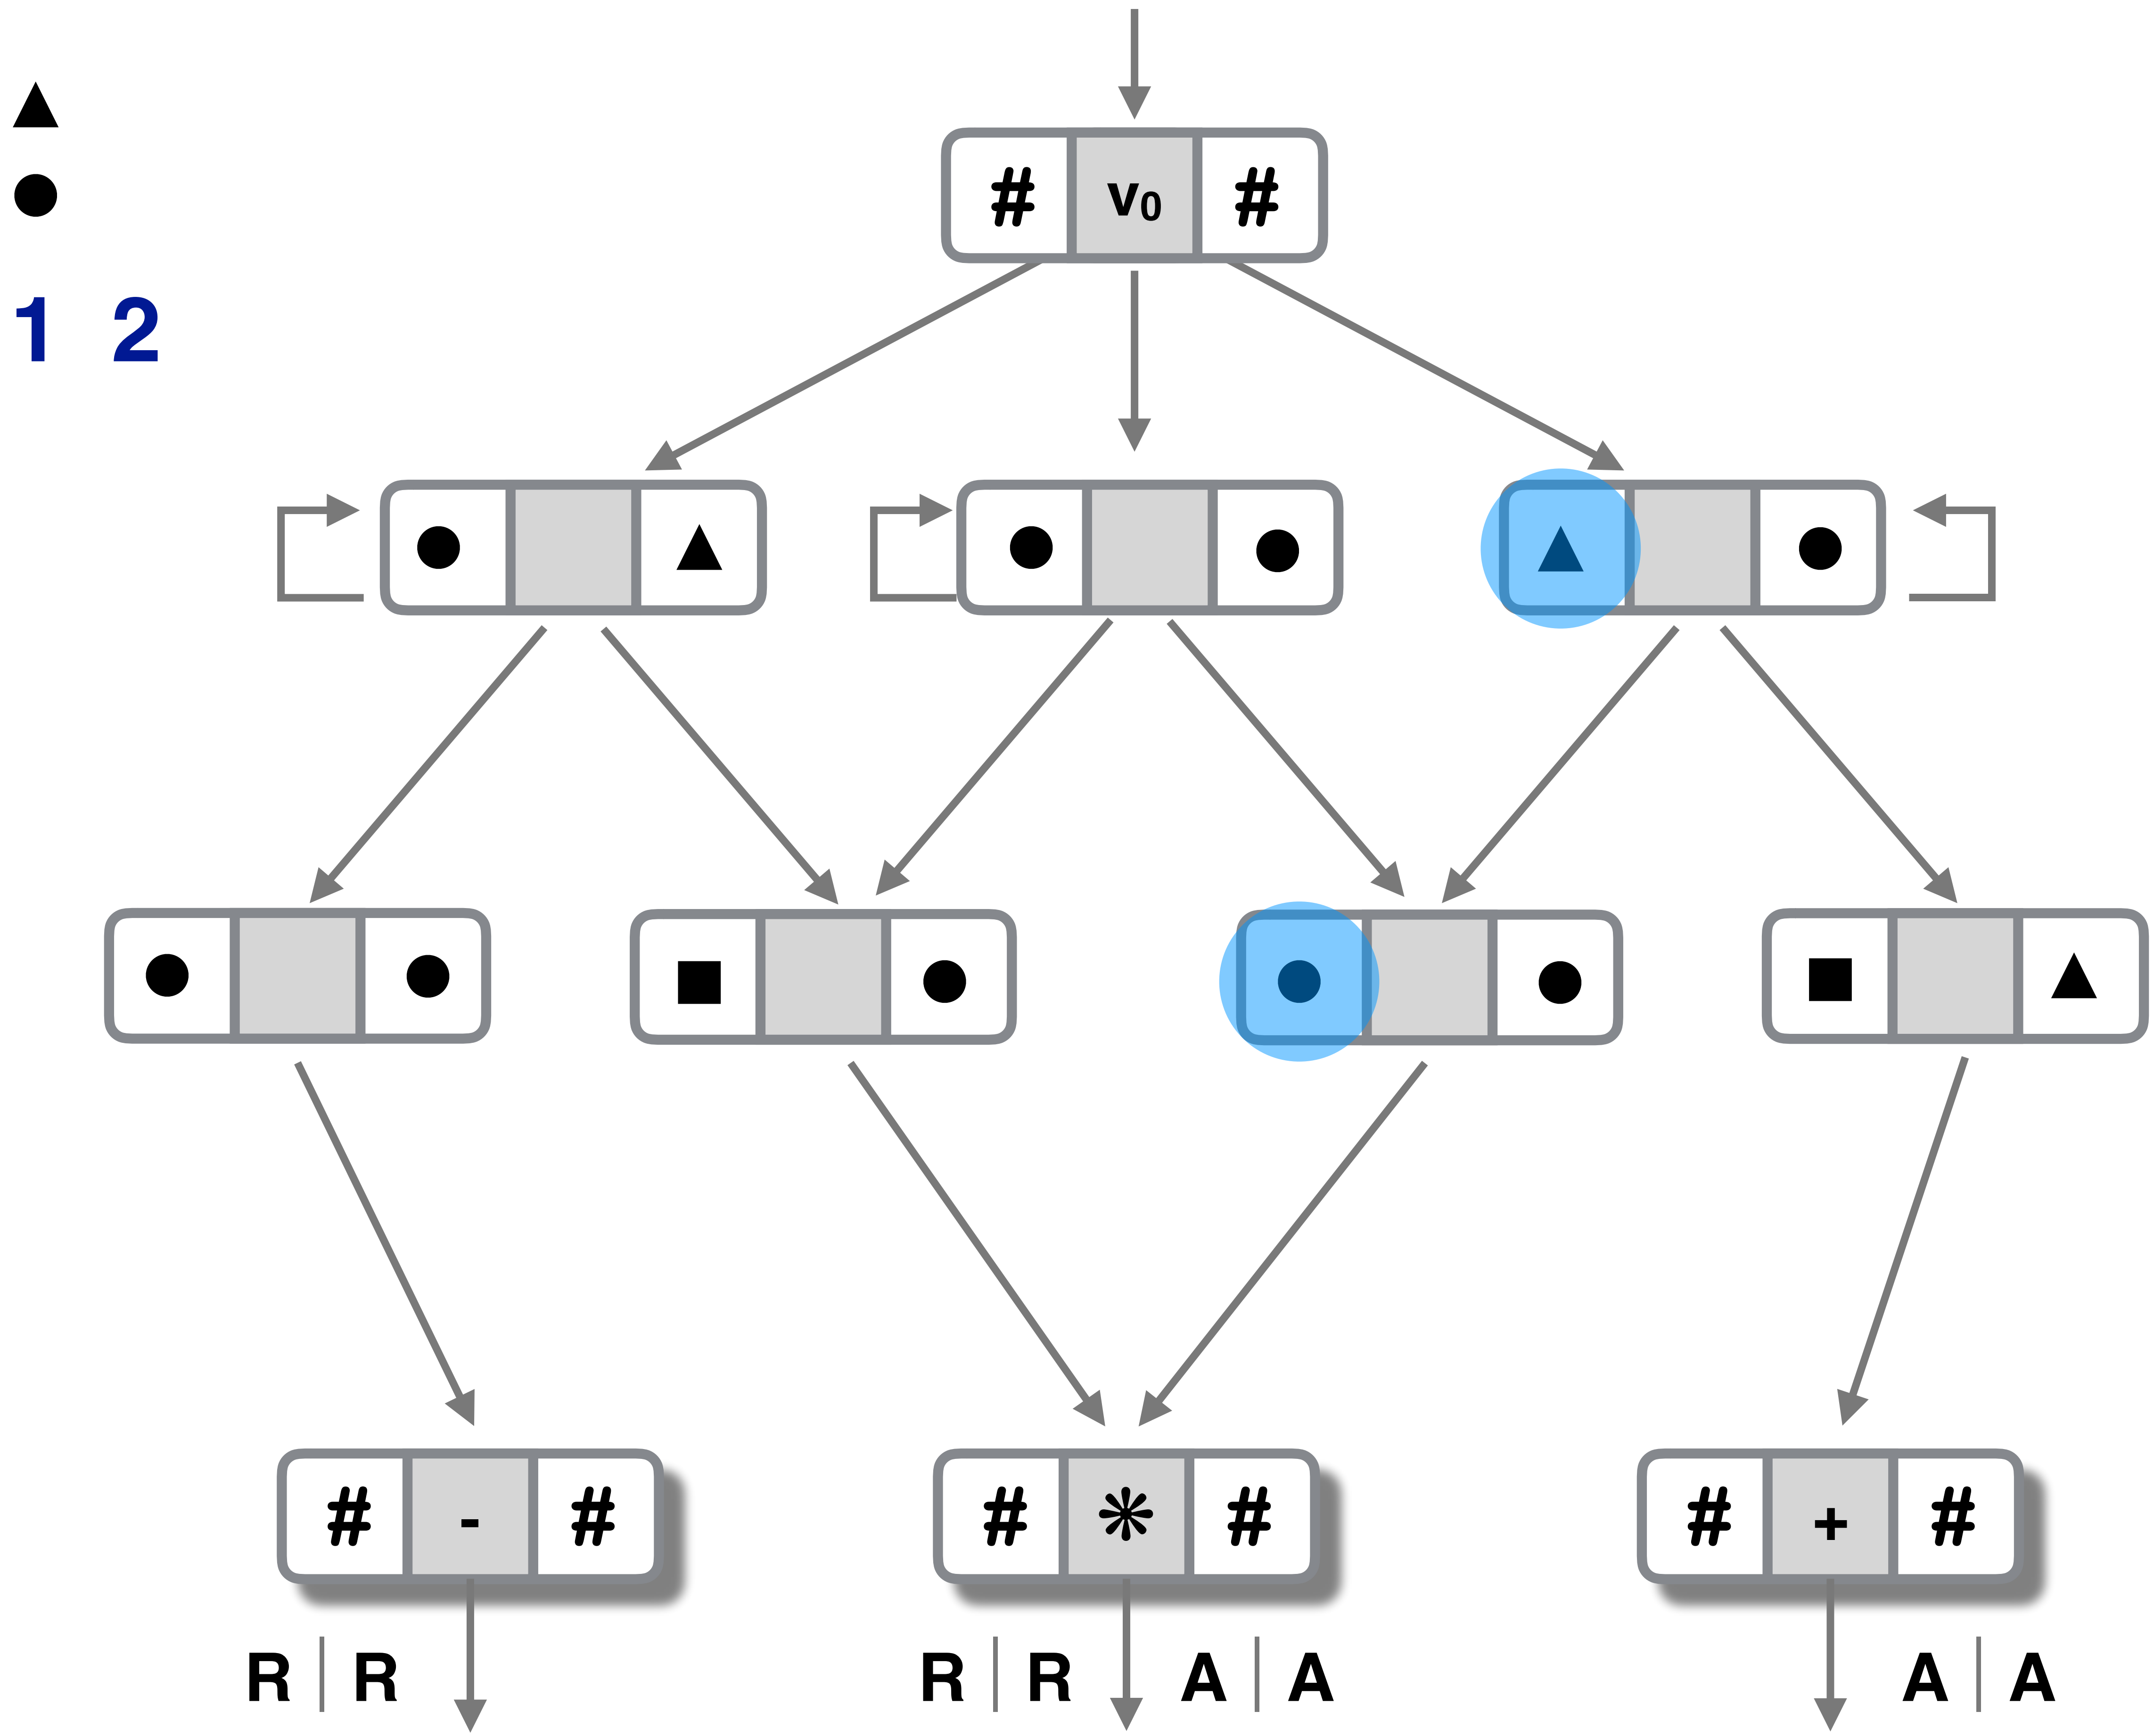
1 2
v0
-
*
+
R | R
R | R A | A
A | A

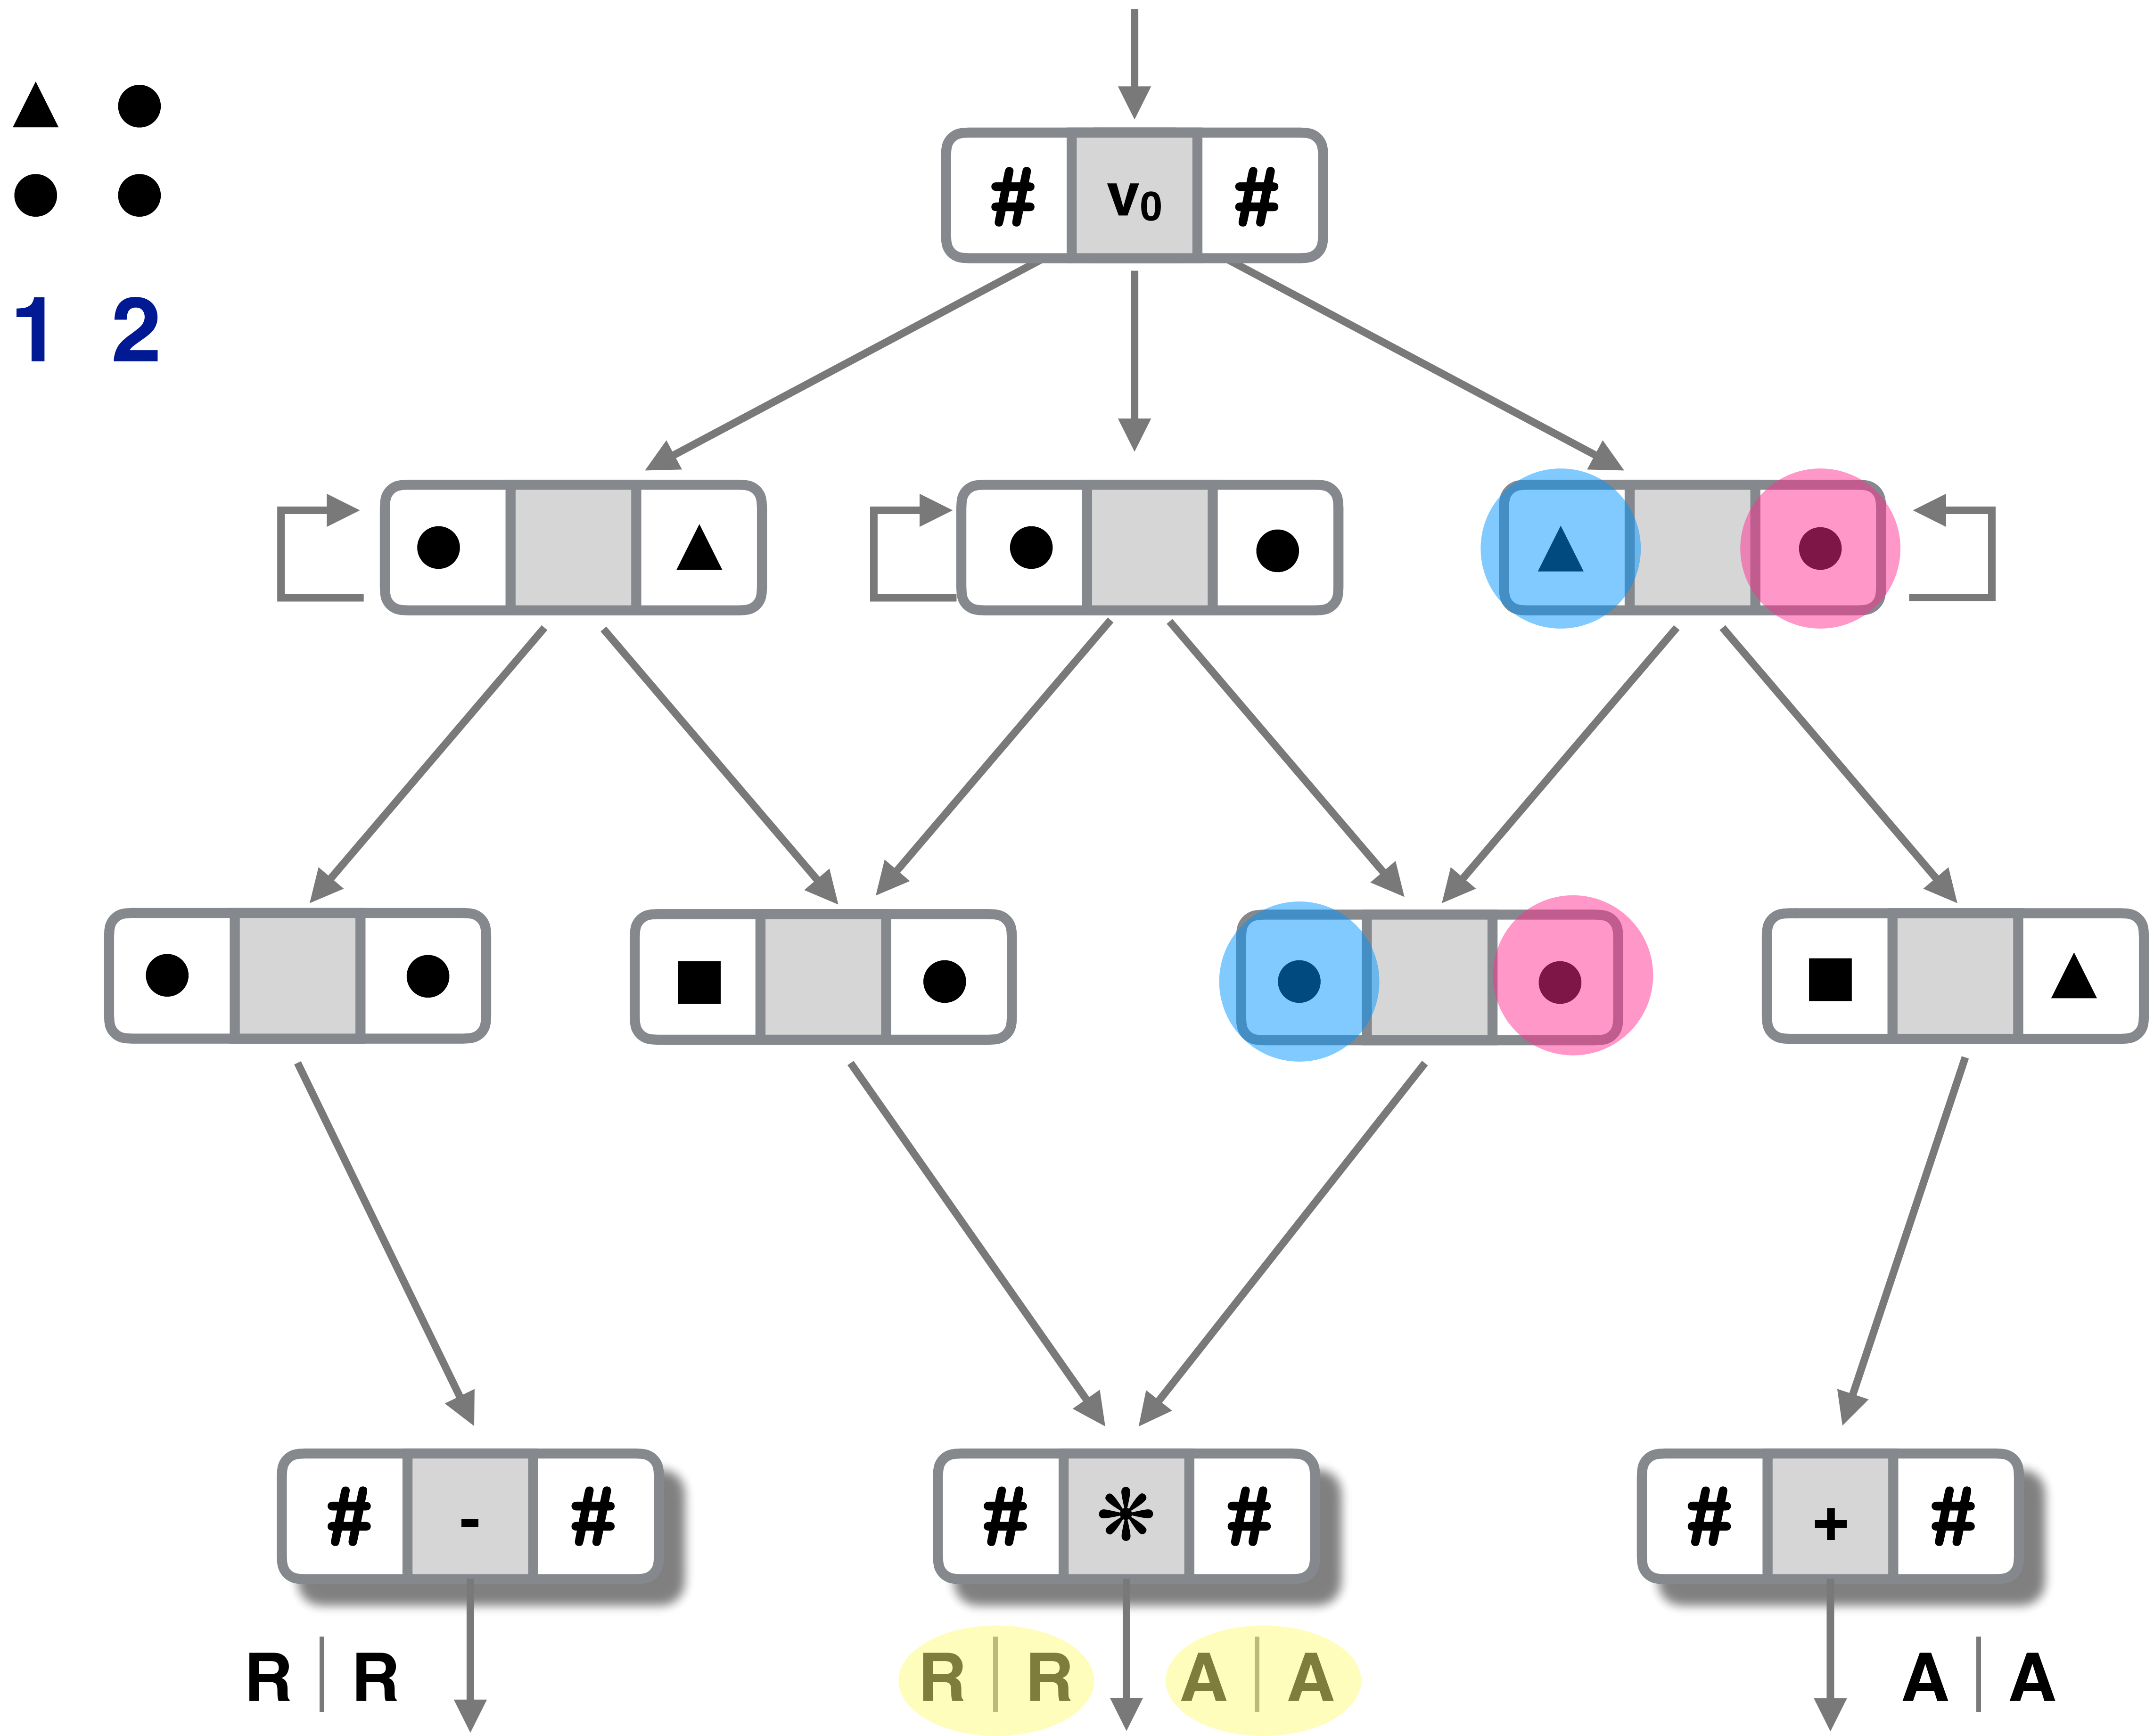
▲ ●
● ●
1 2
#
v_0
#
-
*
+
R | R
R | R
A | A
A | A

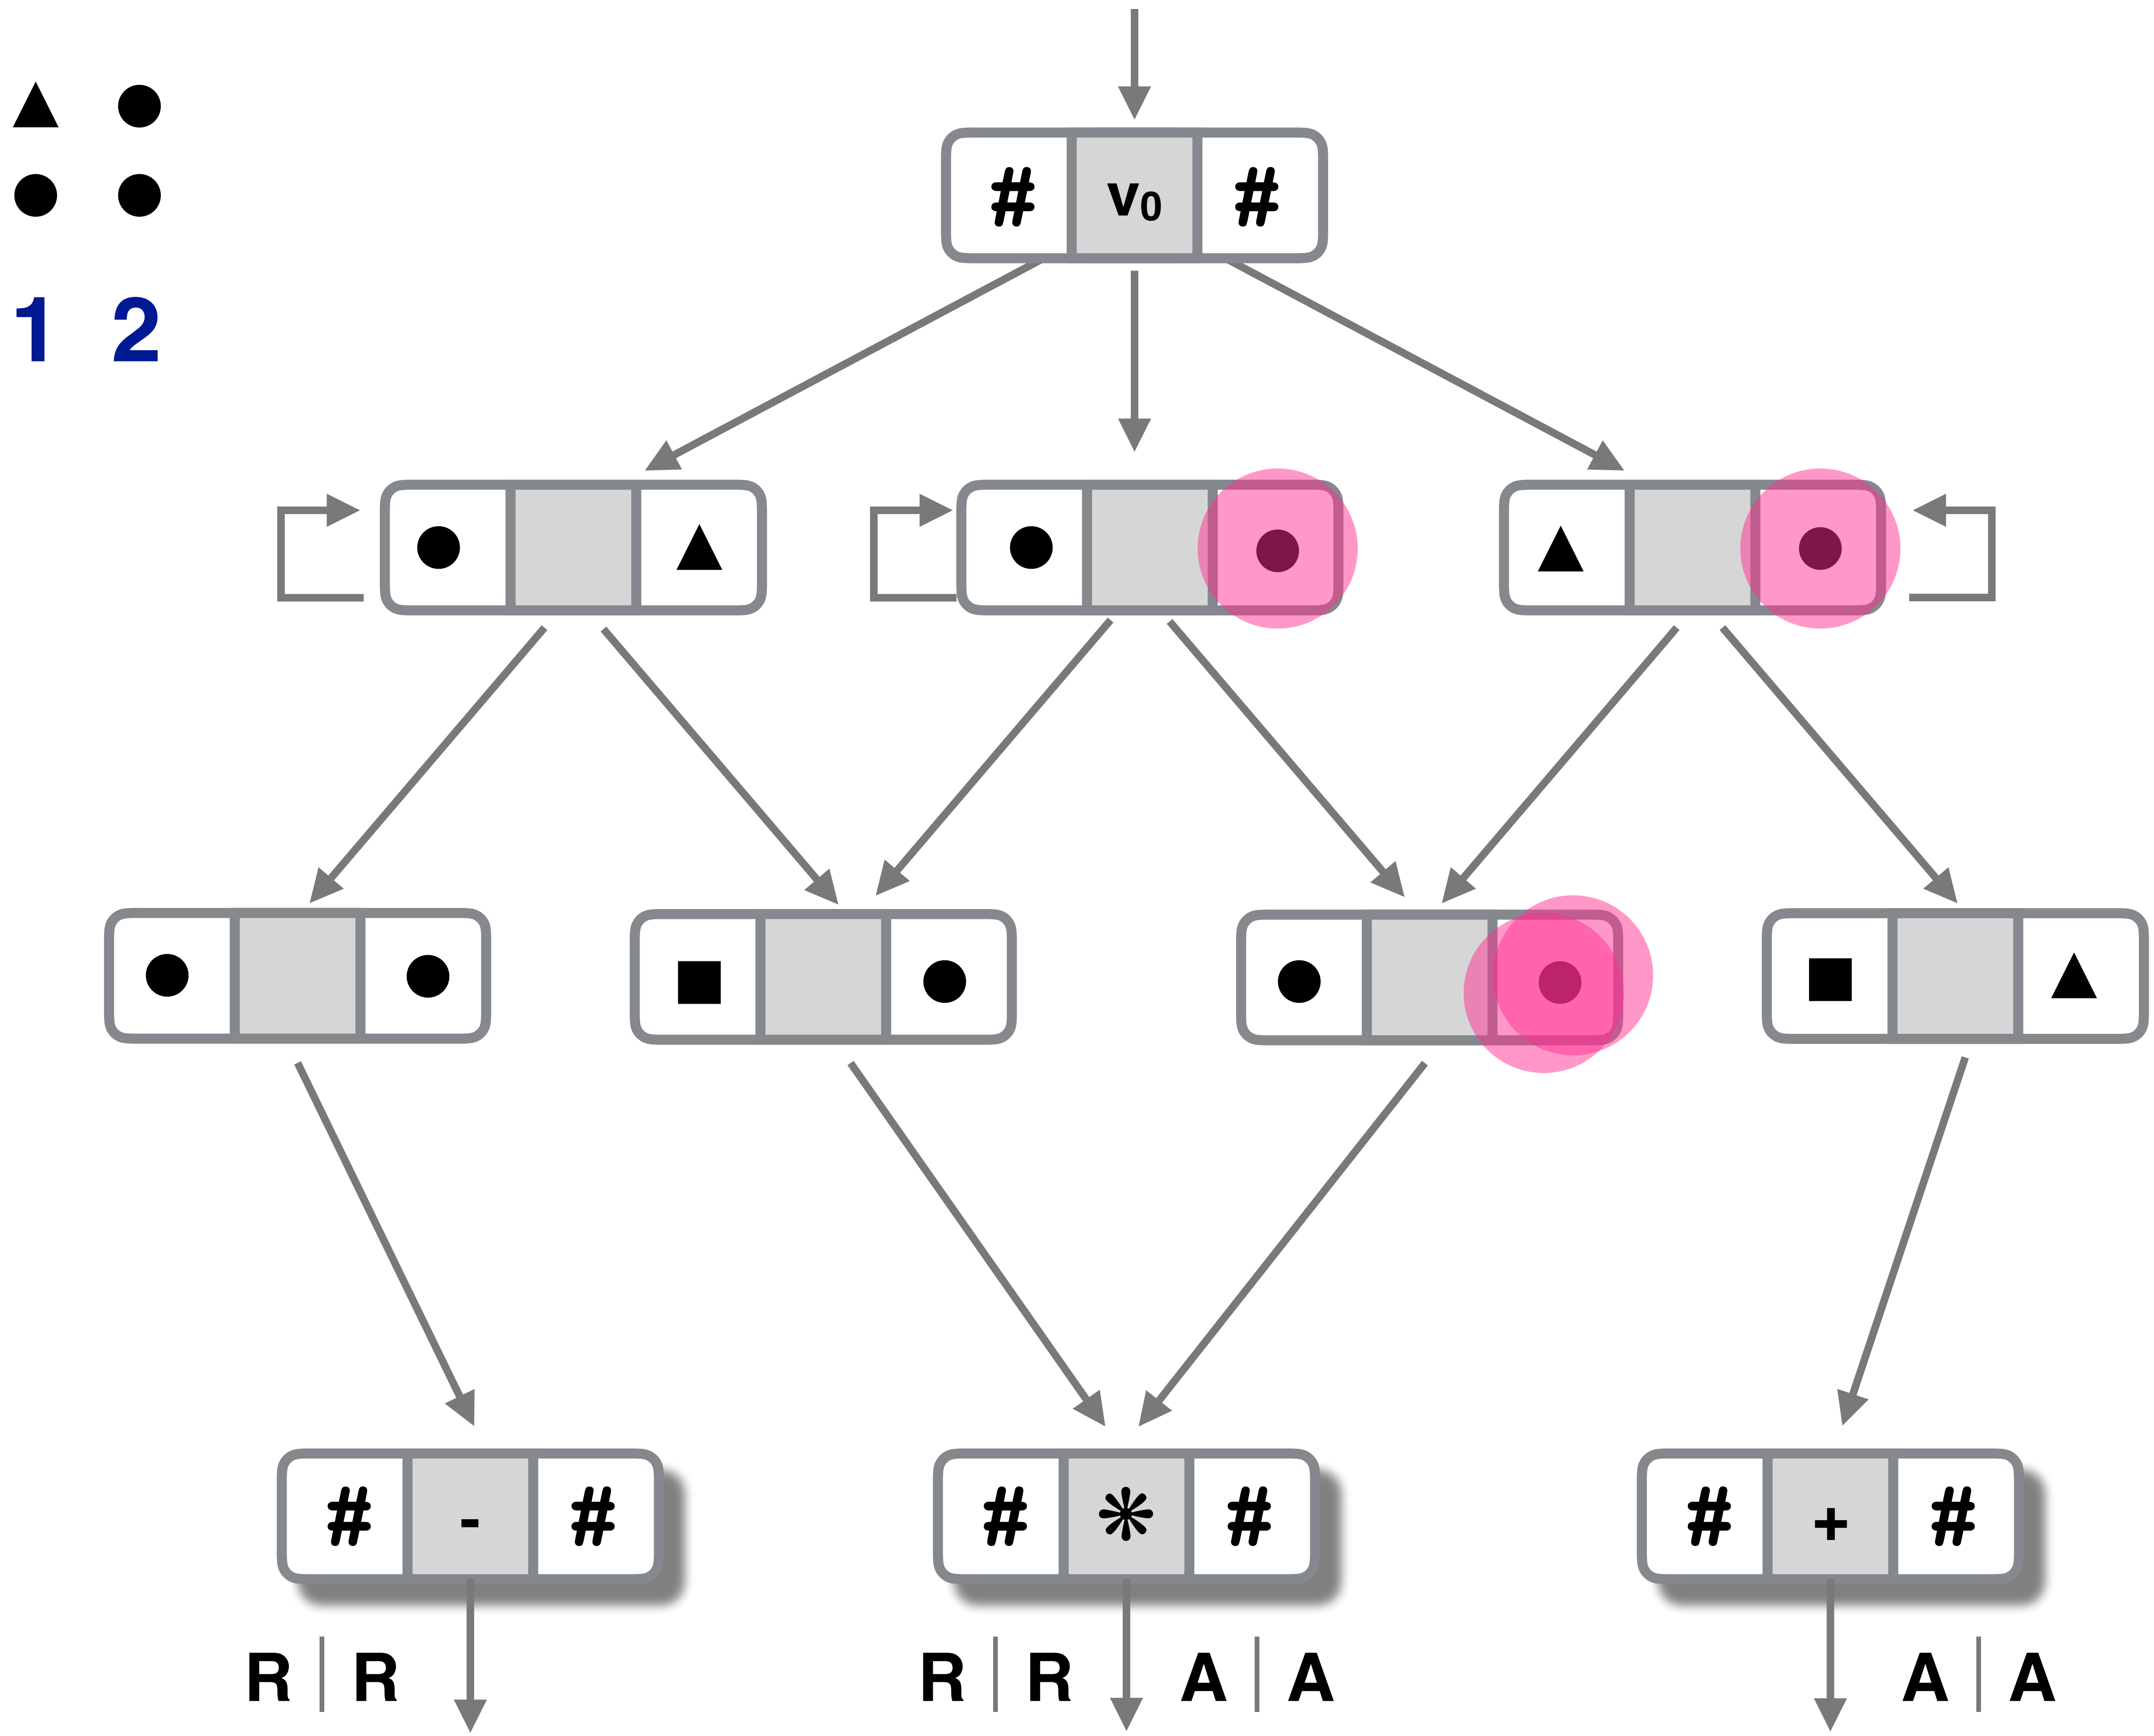
1 2
#
v0
#
-
*
+
R | R
R | R
A | A
A | A

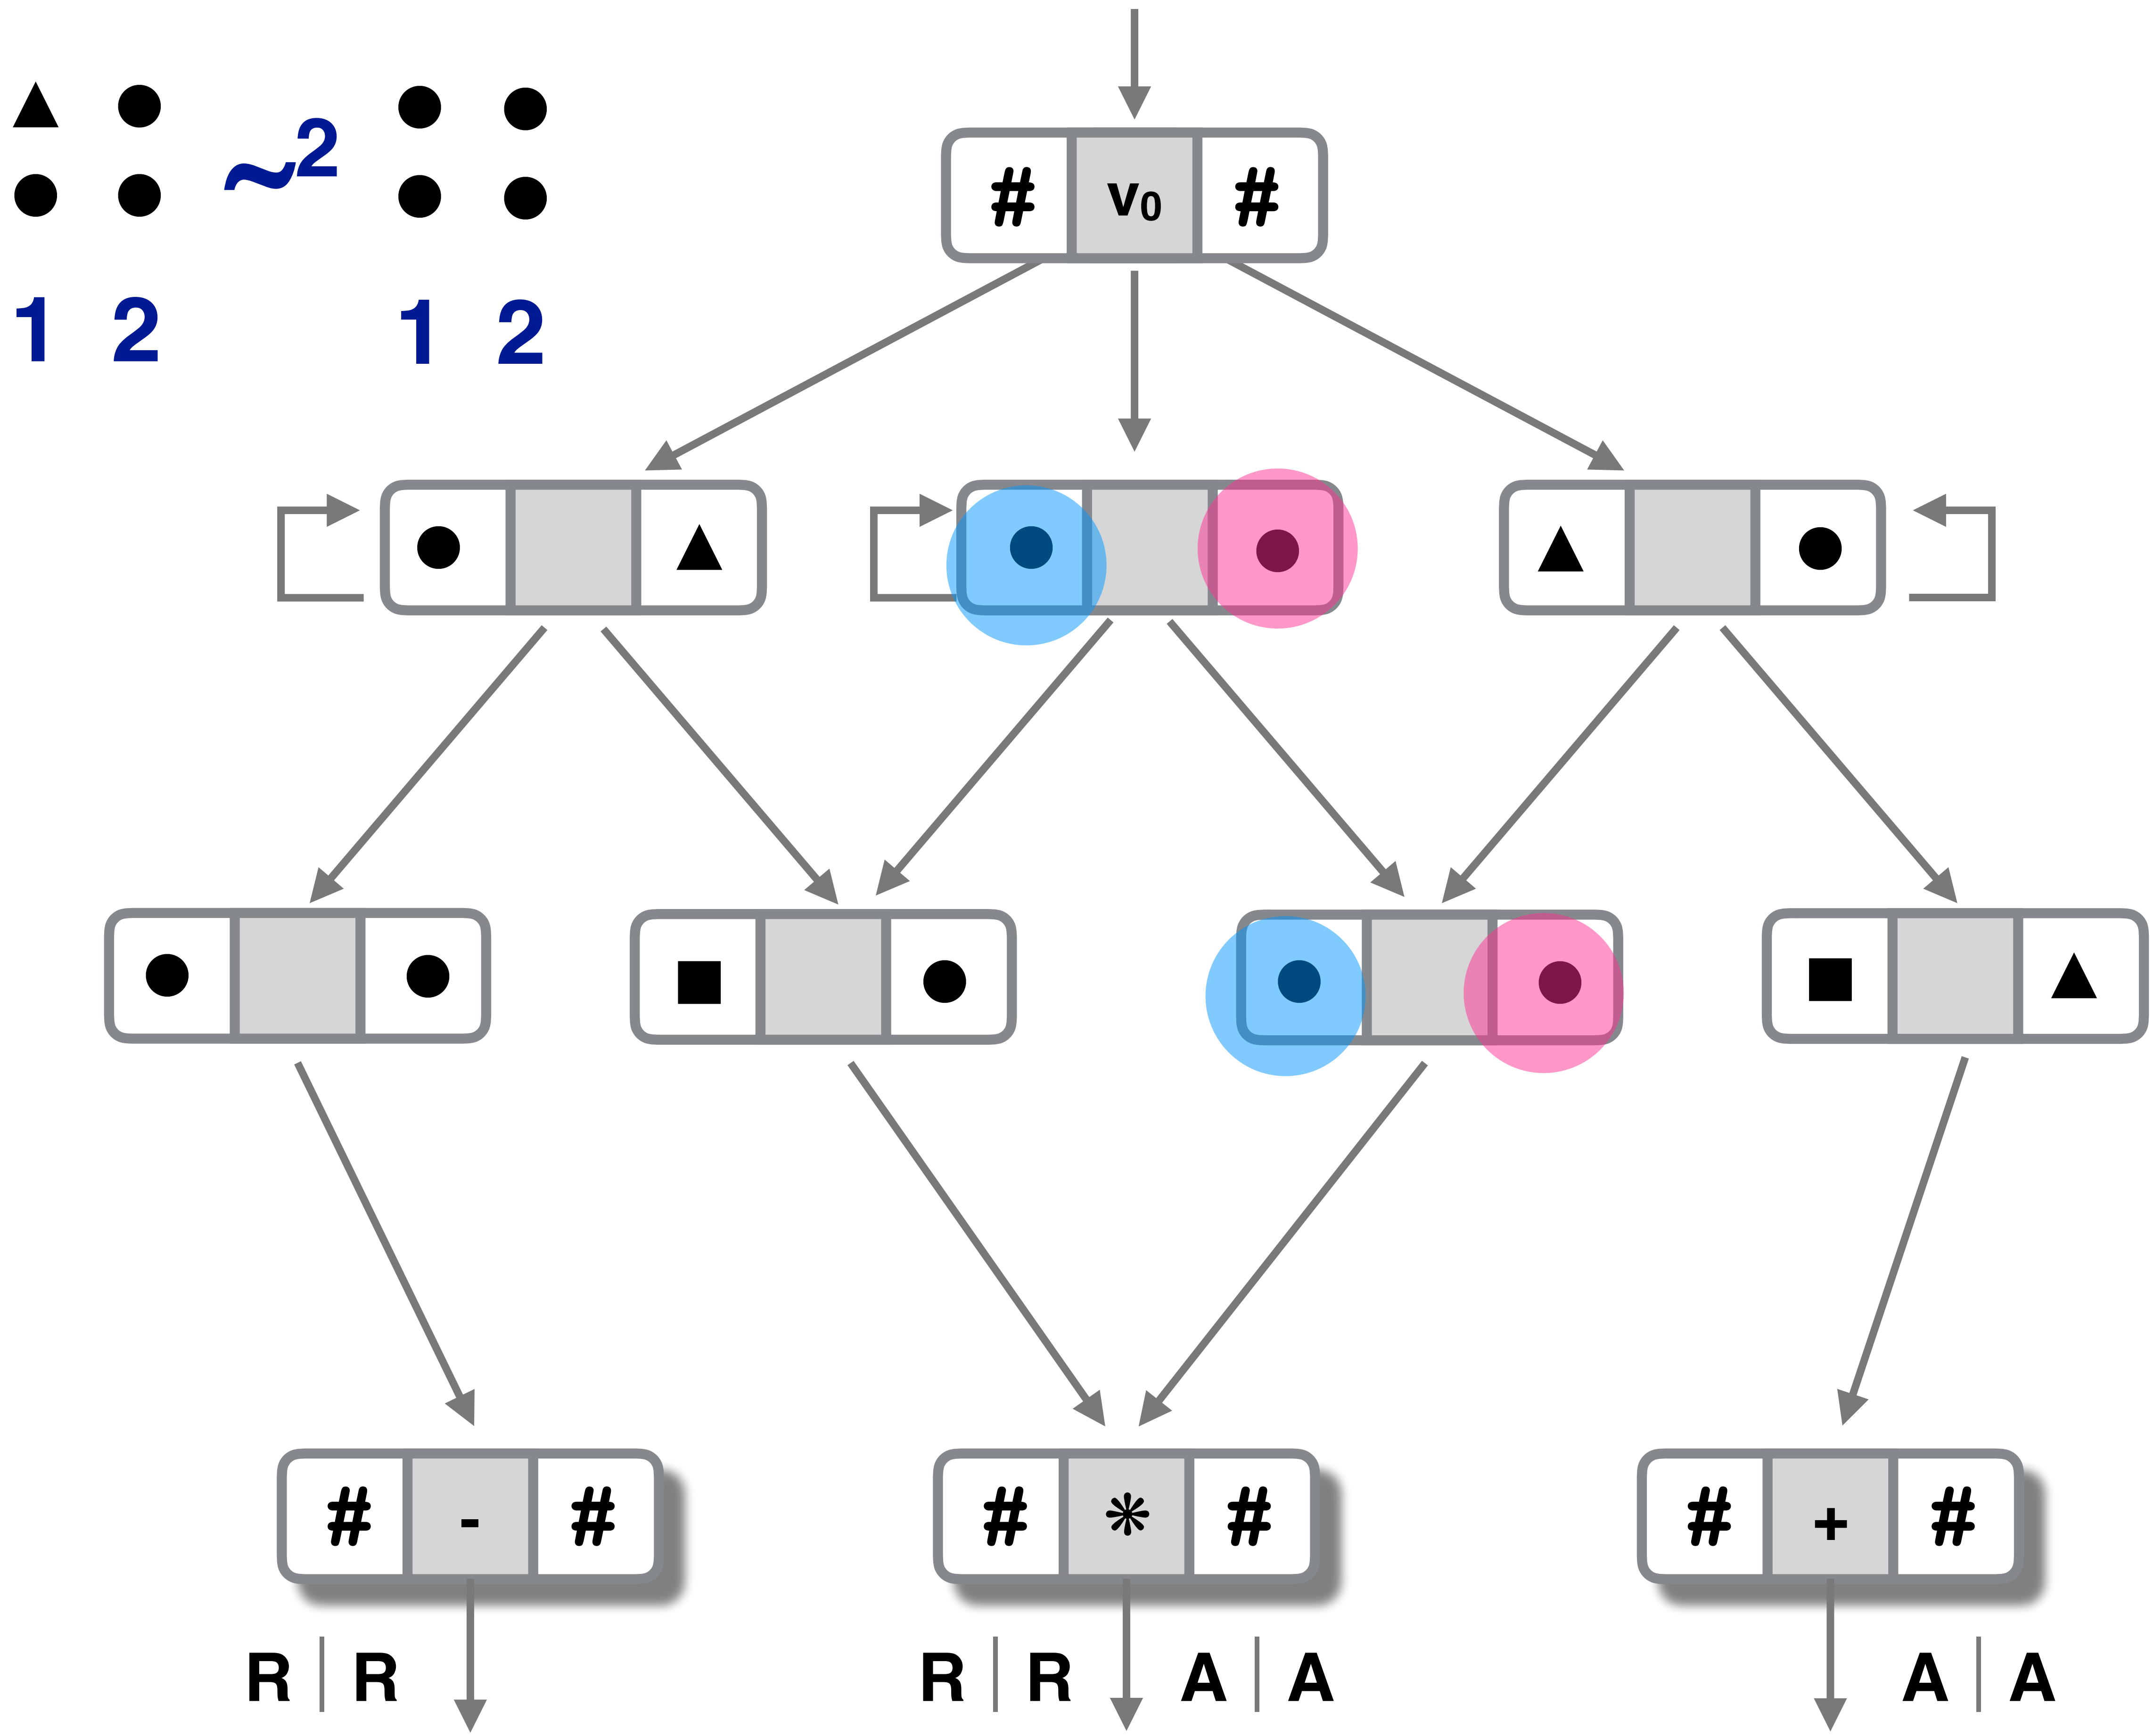
▲ ●
● ●
1 2
∼2
● ●
● ●
1 2
v0
-
*
+
R | R
R | R
A | A
A | A

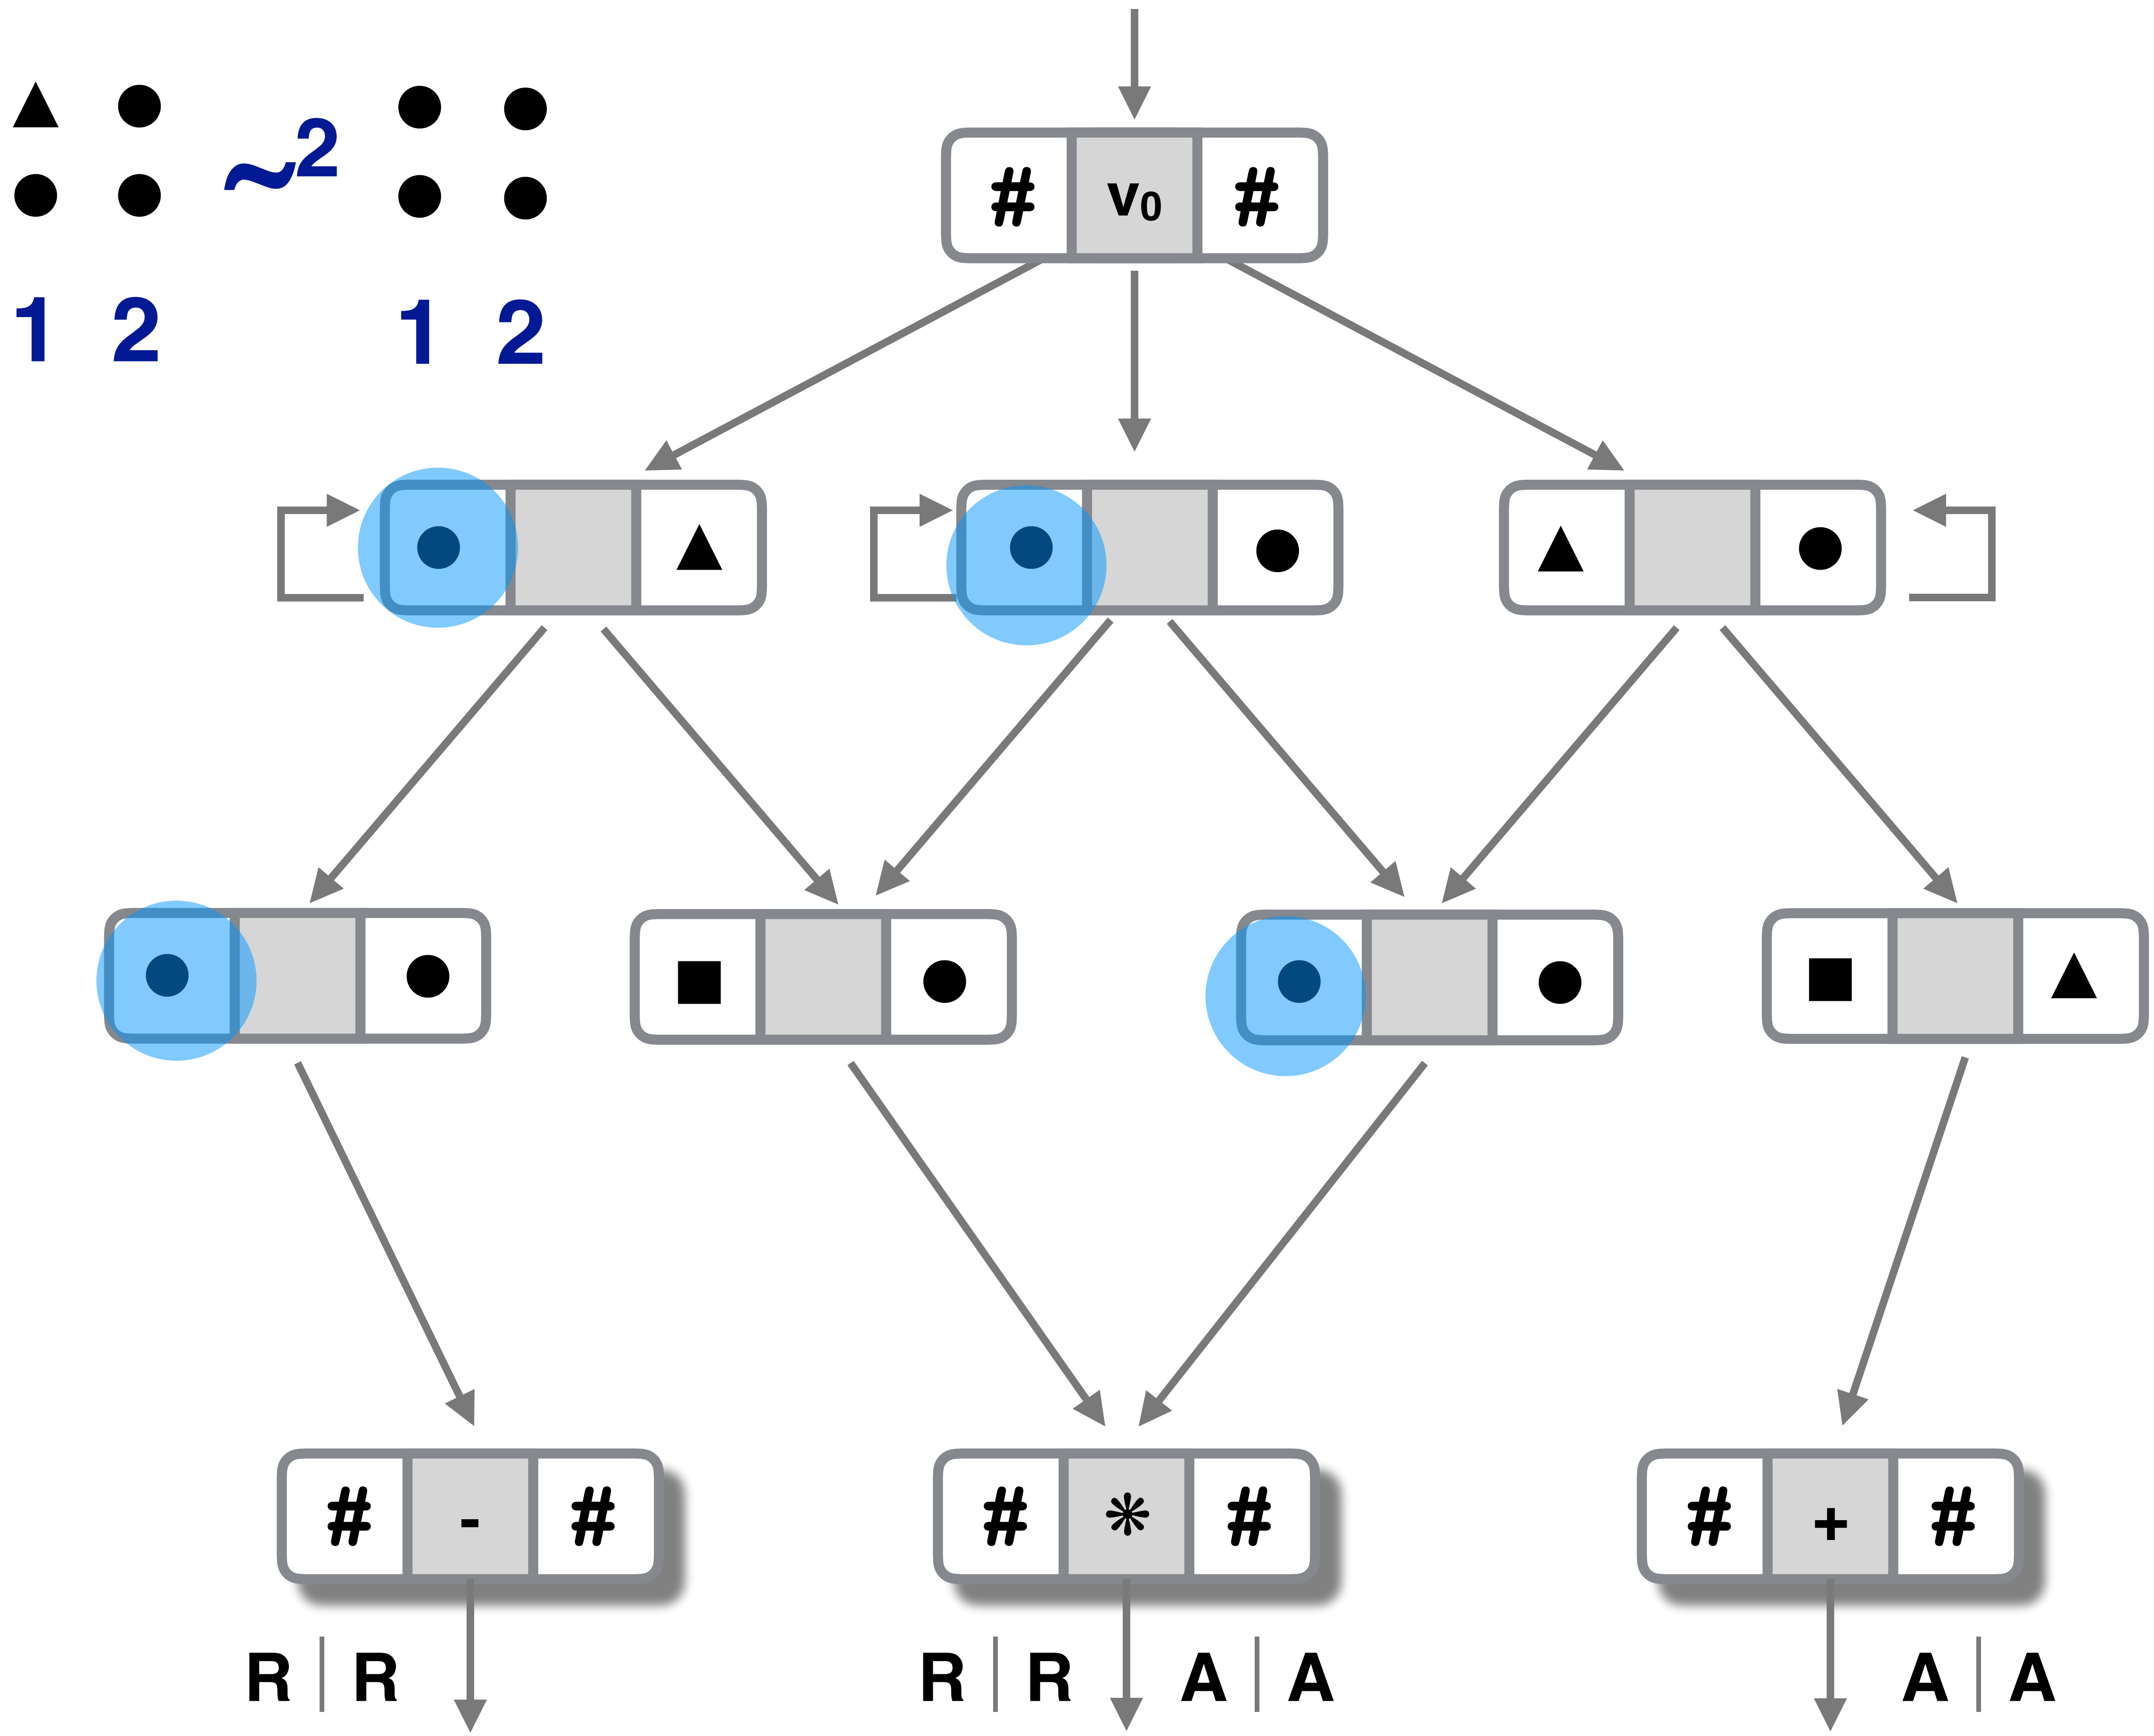

$\sim_2$
1 2
1 2
v_0
-
*
+
R | R
R | R A | A
A | A

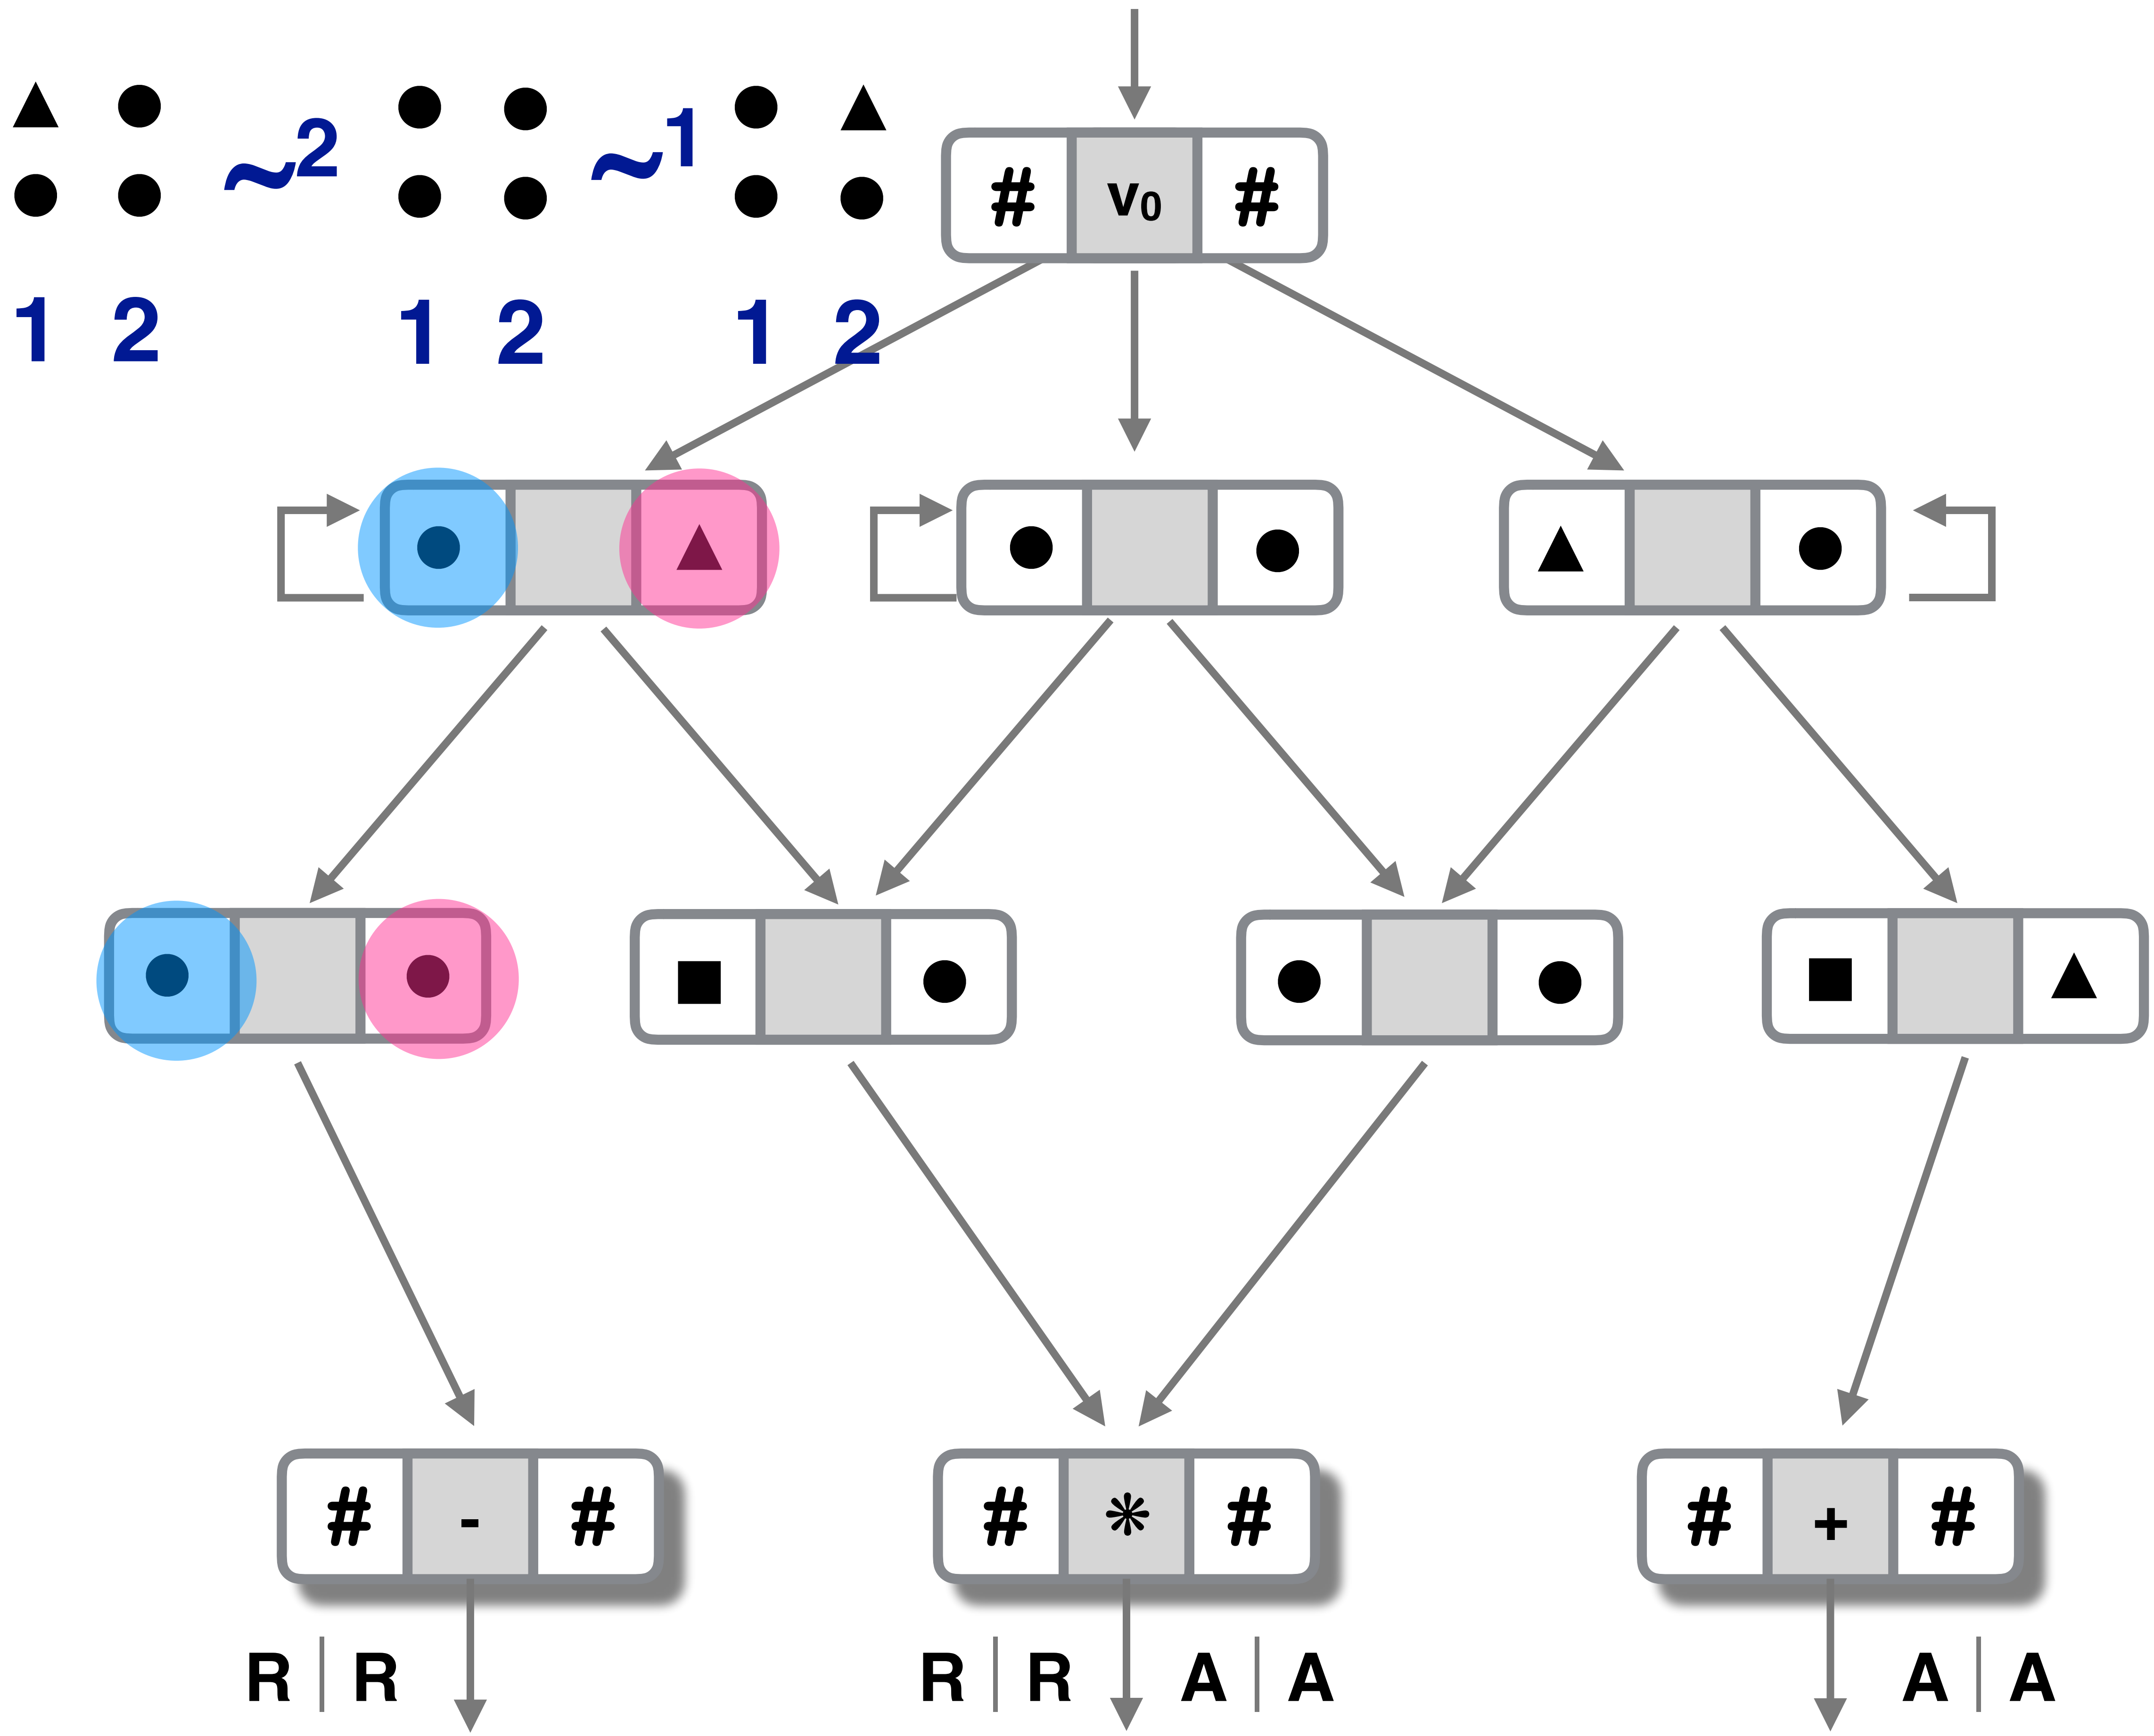
~2
~1
1 2
1 2
1 2
v0
-
*
+
R | R
R | R A | A
A | A

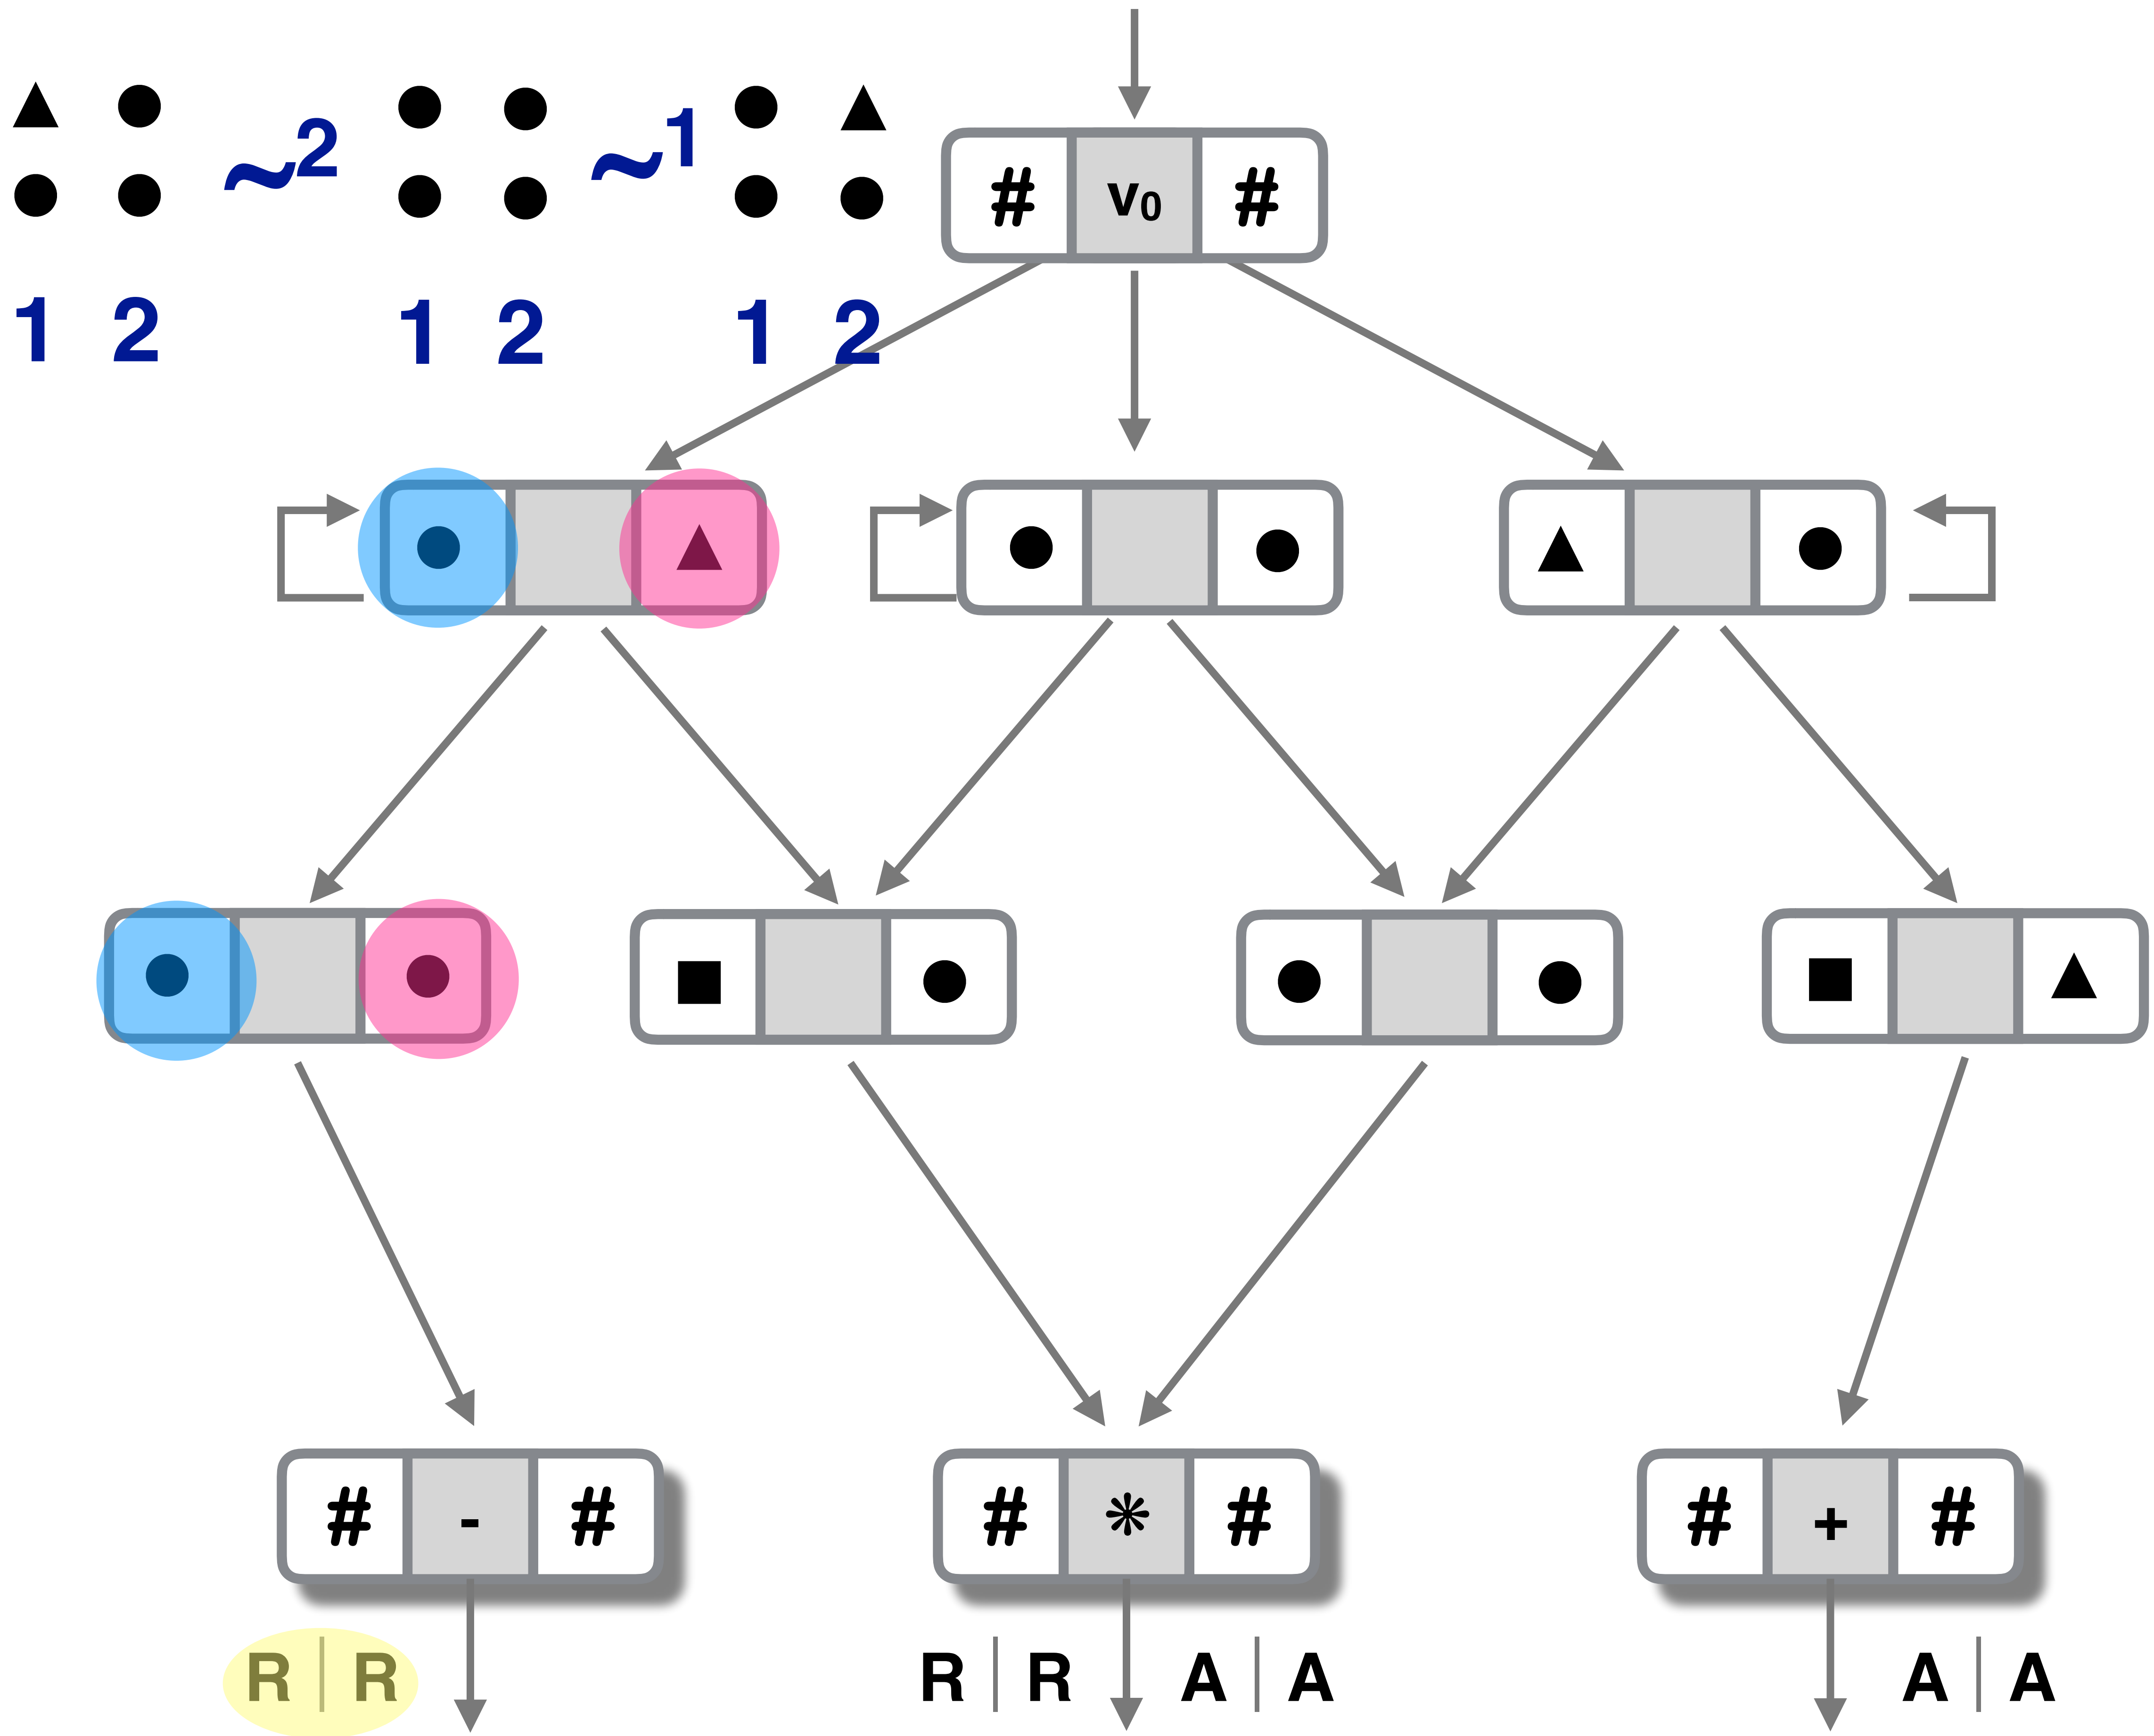
~2
~1
1 2
1 2
1 2
#
v0
#
-
*
+
R | R
R | R
A | A
A | A

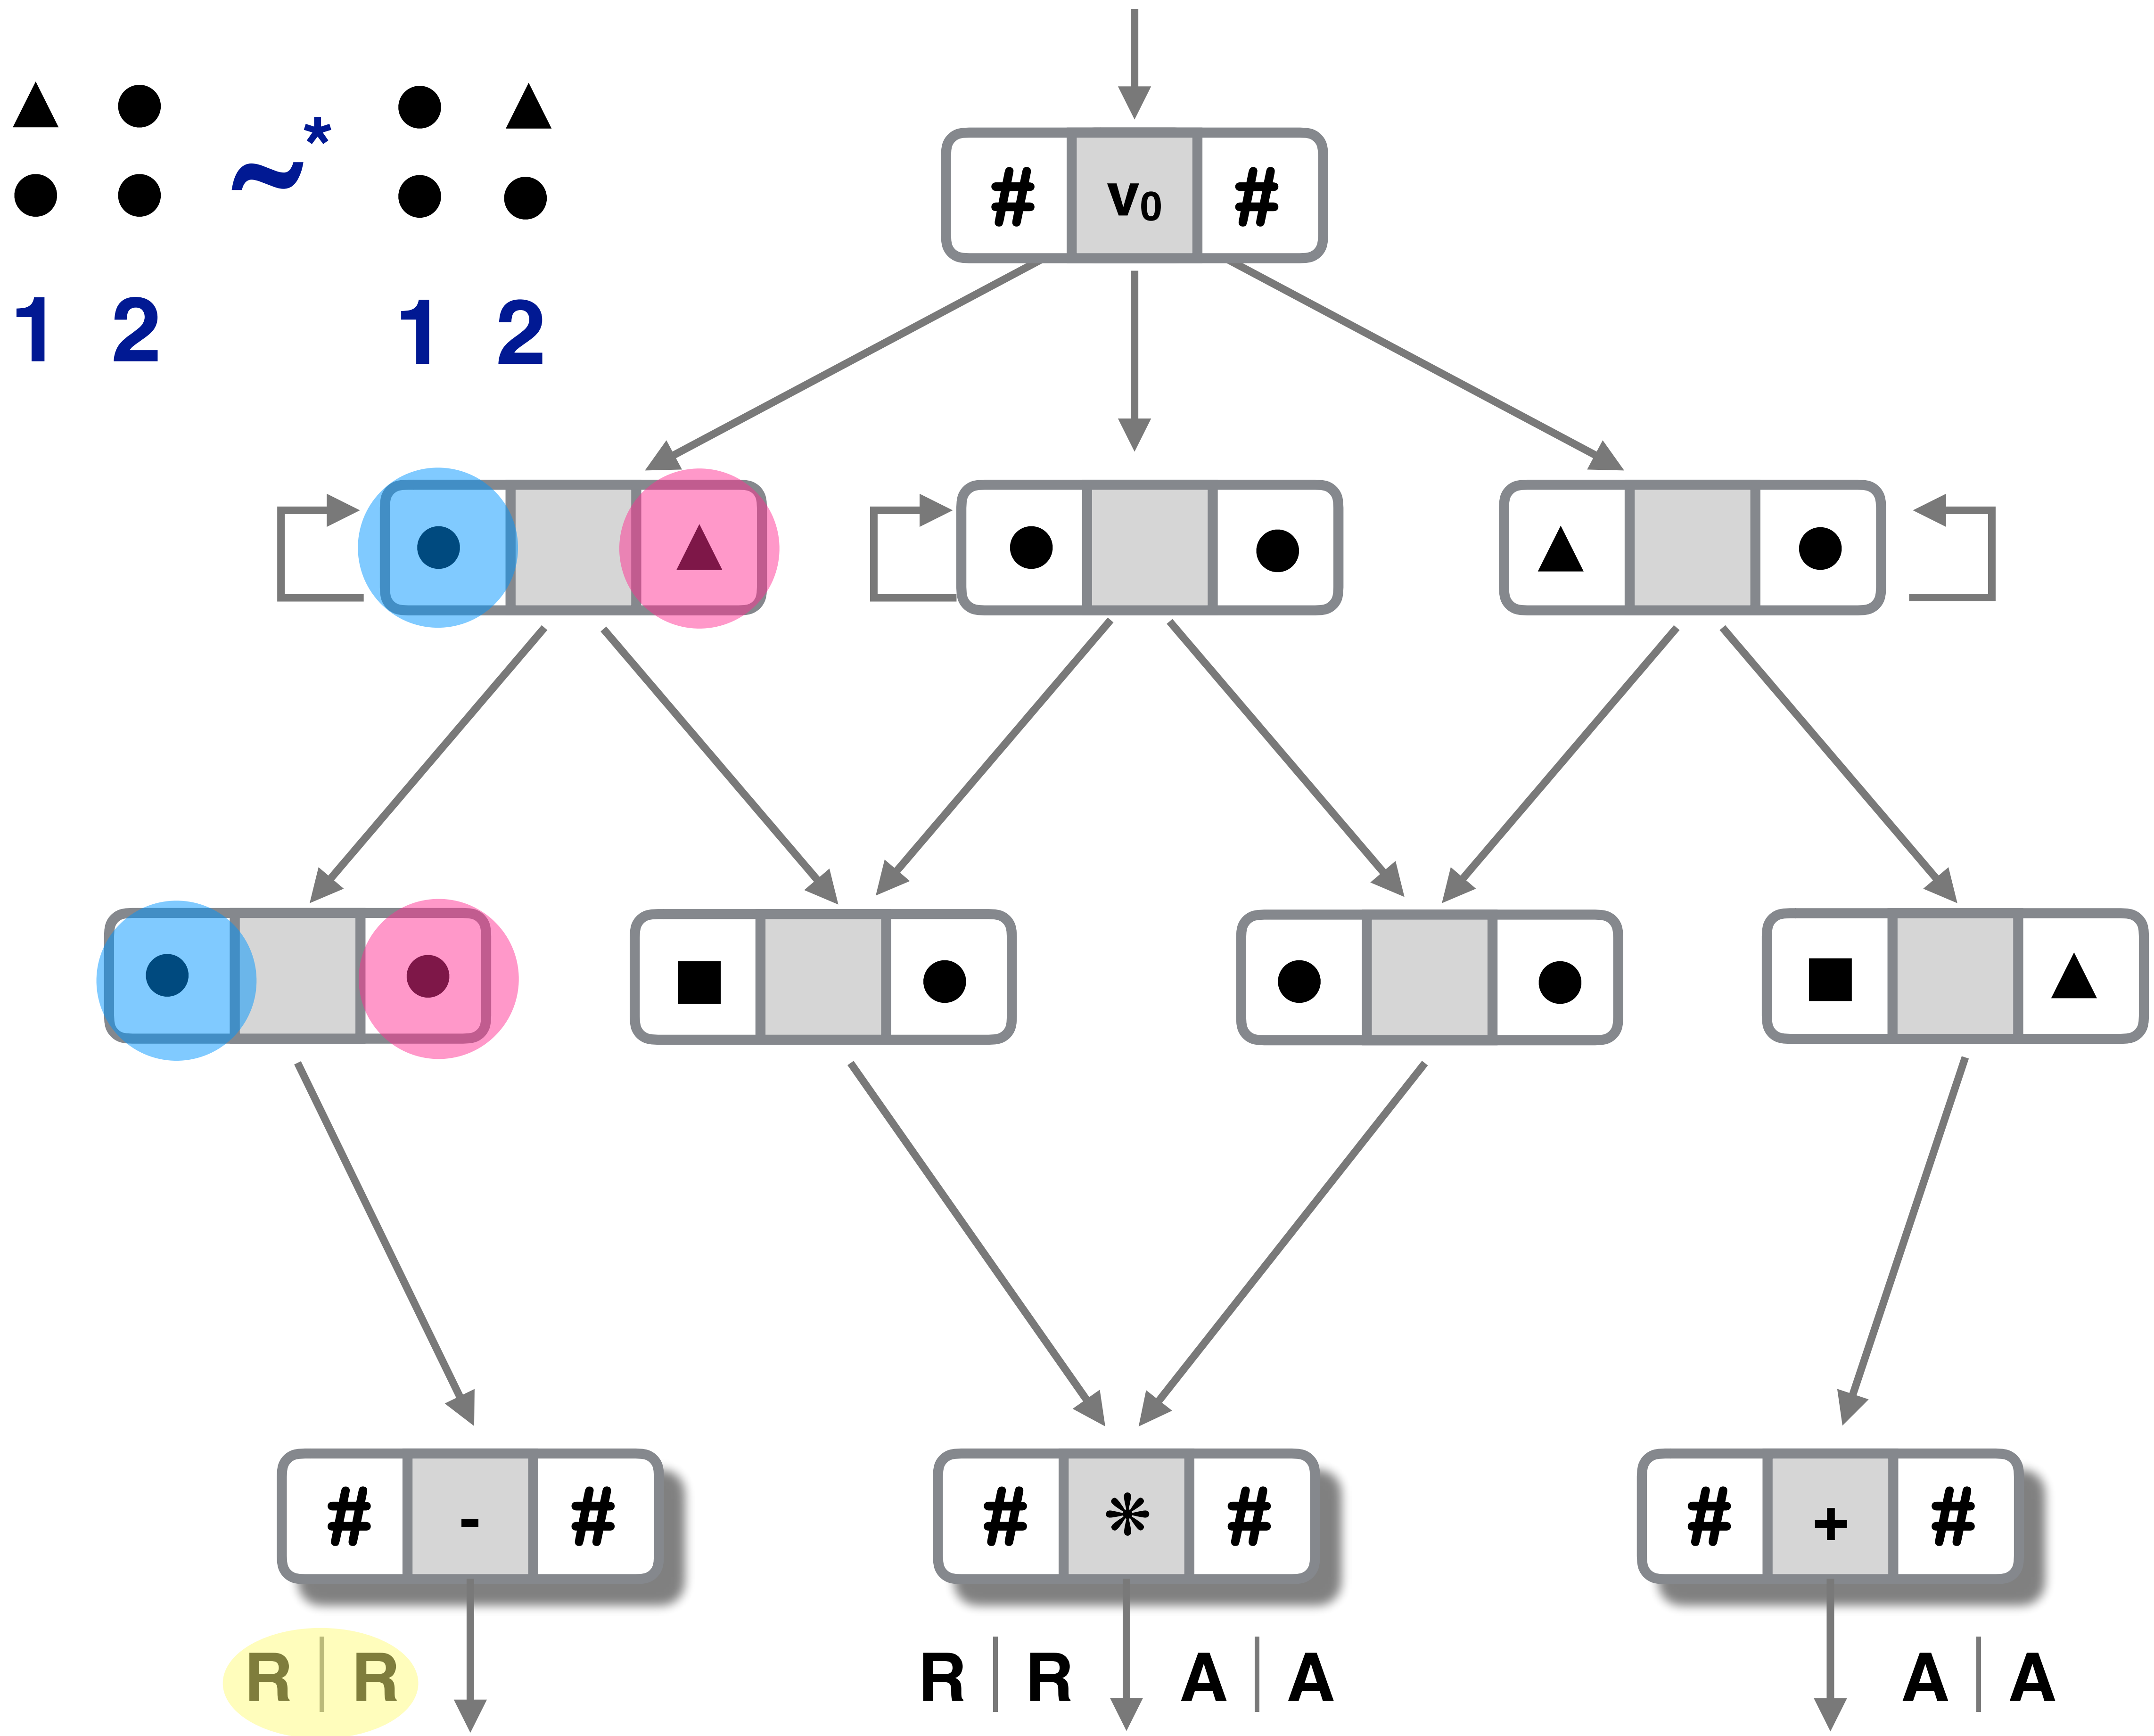
∼*
1 2
1 2
#
v_0
#
-
*
+
R | R
R | R
A | A
A | A

Propagation of uncertainty

Consensus strategies:

Strategy s^i : observed play $\rightarrow \{acc, rej\}$, s.t.
$s^1(\pi) = s^2(\pi)$ for all π.

Key: Not enough to consider $\sim^1$ or $\sim^2$

Necessary for consensus: $\sim^{*} := (\sim^1 \cup \sim^2)^*$

$$\rightsquigarrow \text{ if } \pi \sim^{*} \pi', \text{ then } s(\pi) = s(\pi')$$

"play accept" must be common knowledge

Propagation of uncertainty

Theorem Distributed strategy synthesis is **undecidable**
for consensus game acceptors.

Tool: correspondence with context-sensitive languages
via domino tilings results [Latteux, Simplot '97]

Insight: cooperation is intrinsically hard

One-way Information Flow

How to avoid uncertainty propagation?

Idea: Restrict the shape of information flow

Known decidable case: hierarchical observation pattern
[PR'79; PR '90; Kupferman,Vardi '01]

$\rightsquigarrow$ Absence of *information fork* criterion [Finkbeiner, Schewe '05]

One-way Information Flow

How to avoid uncertainty propagation?

Idea: Restrict the shape of information flow

Known decidable case: hierarchical observation pattern
[PR'79; PR '90; Kupferman,Vardi '01]

$\rightsquigarrow$ Absence of *information fork* criterion [Finkbeiner, Schewe '05]

Relaxations of strict hierarchy principle

Hierarchical Observation

Hierarchical Observation

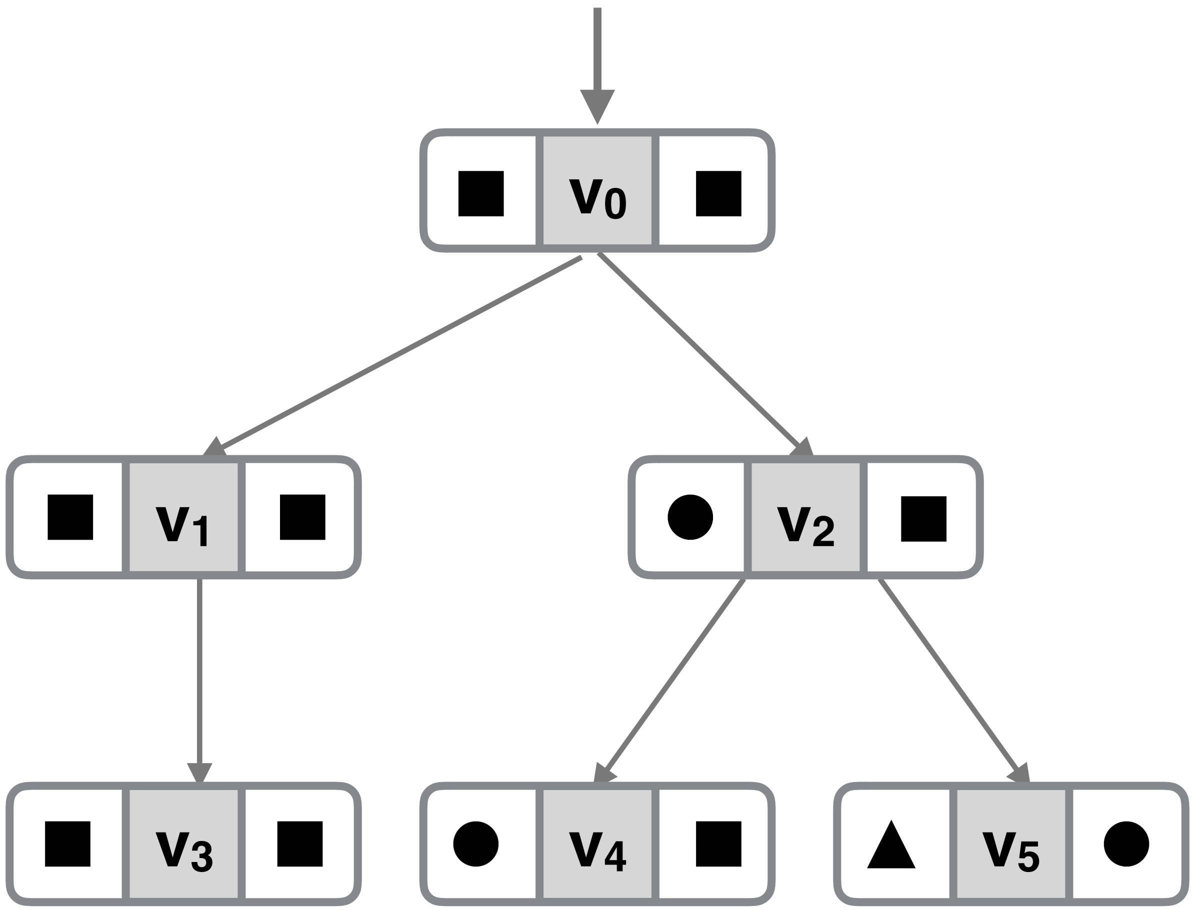

Hierarchical Observation

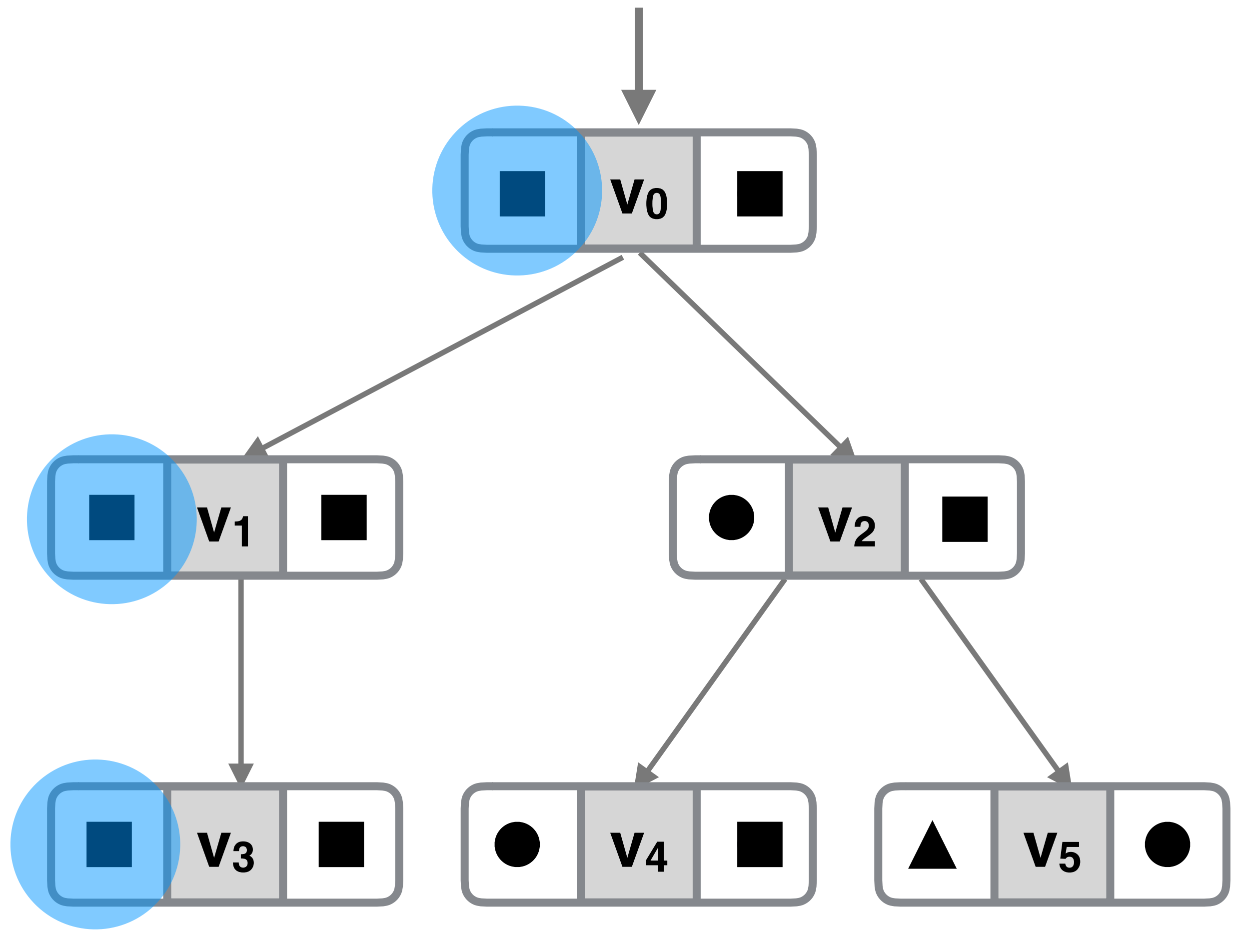

Hierarchical Observation

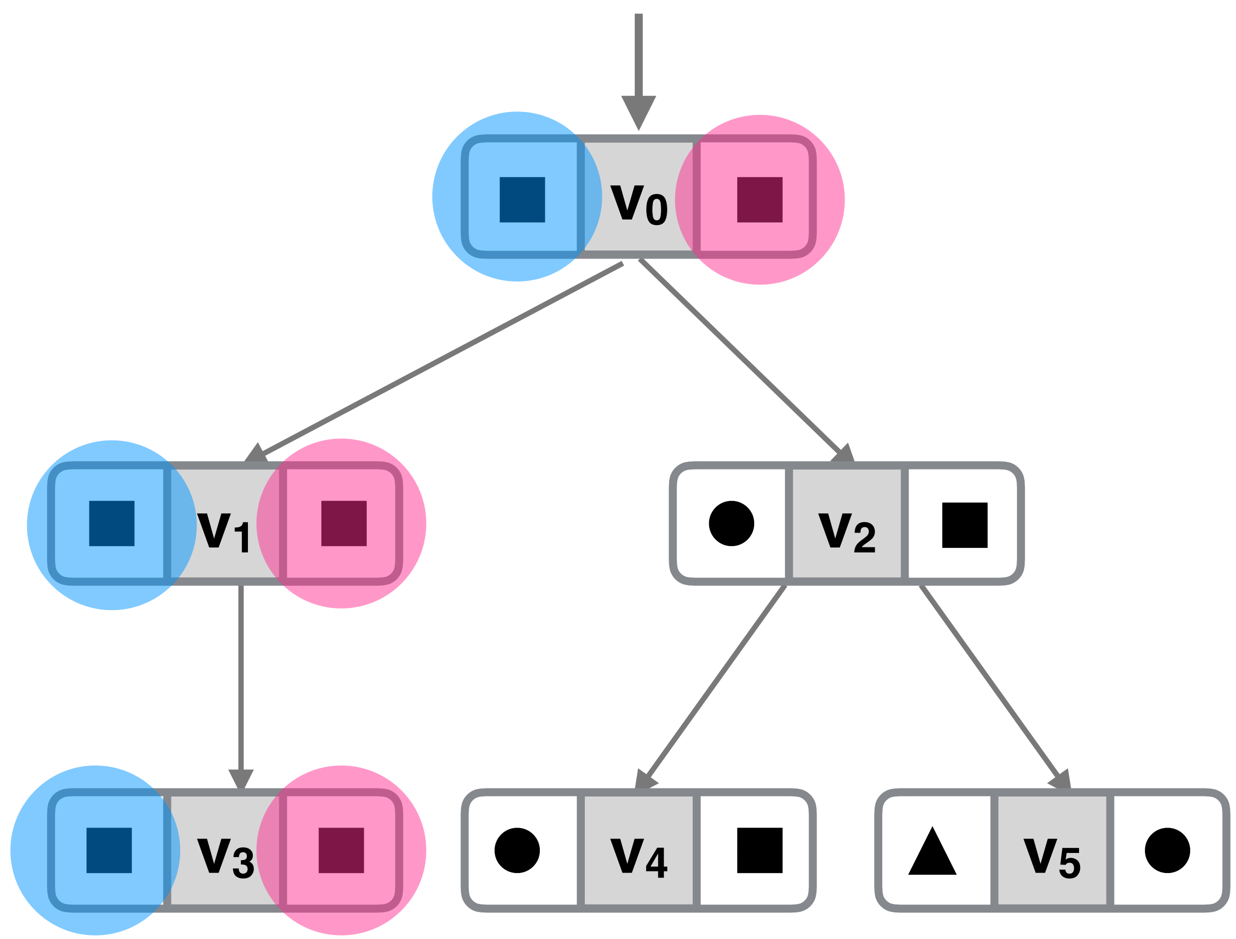

Hierarchical Observation

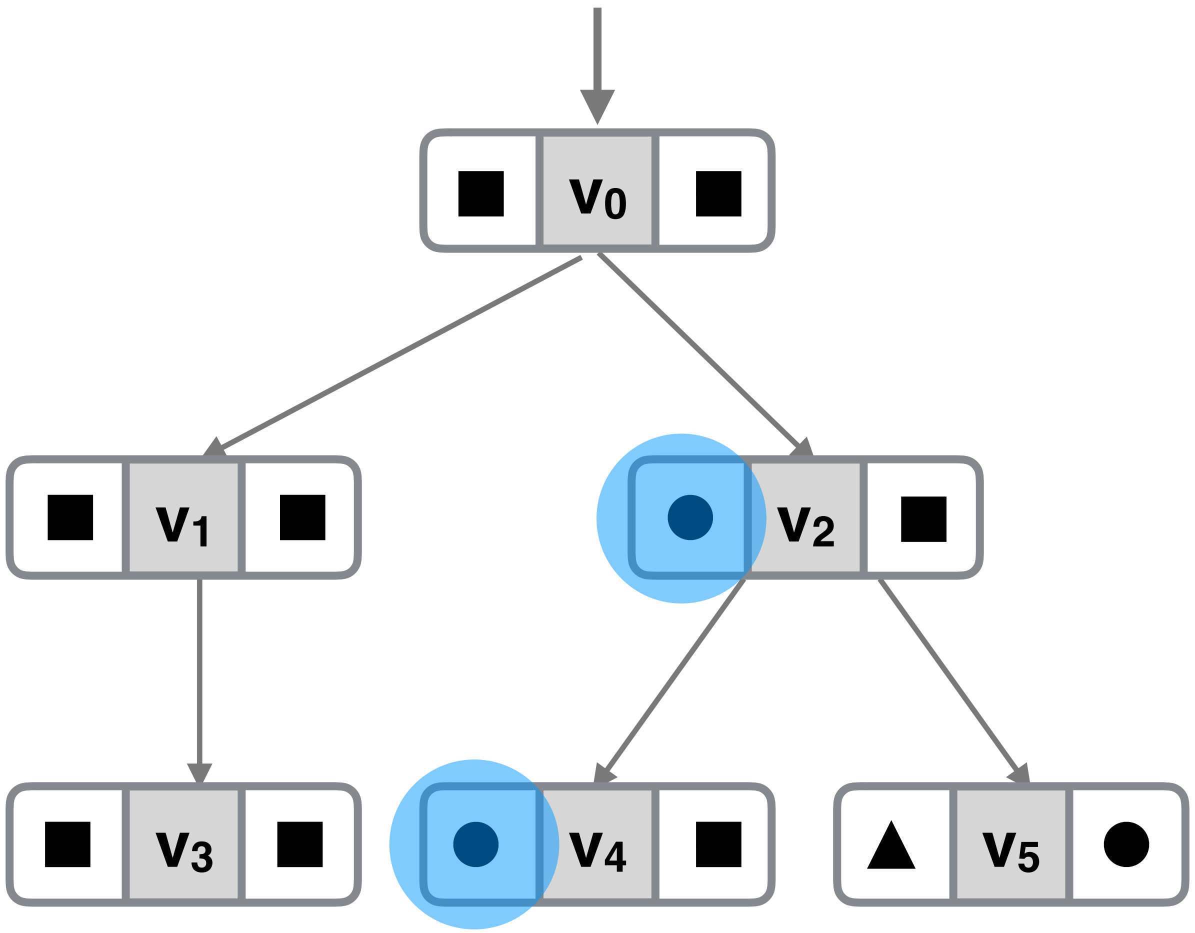

Hierarchical Observation

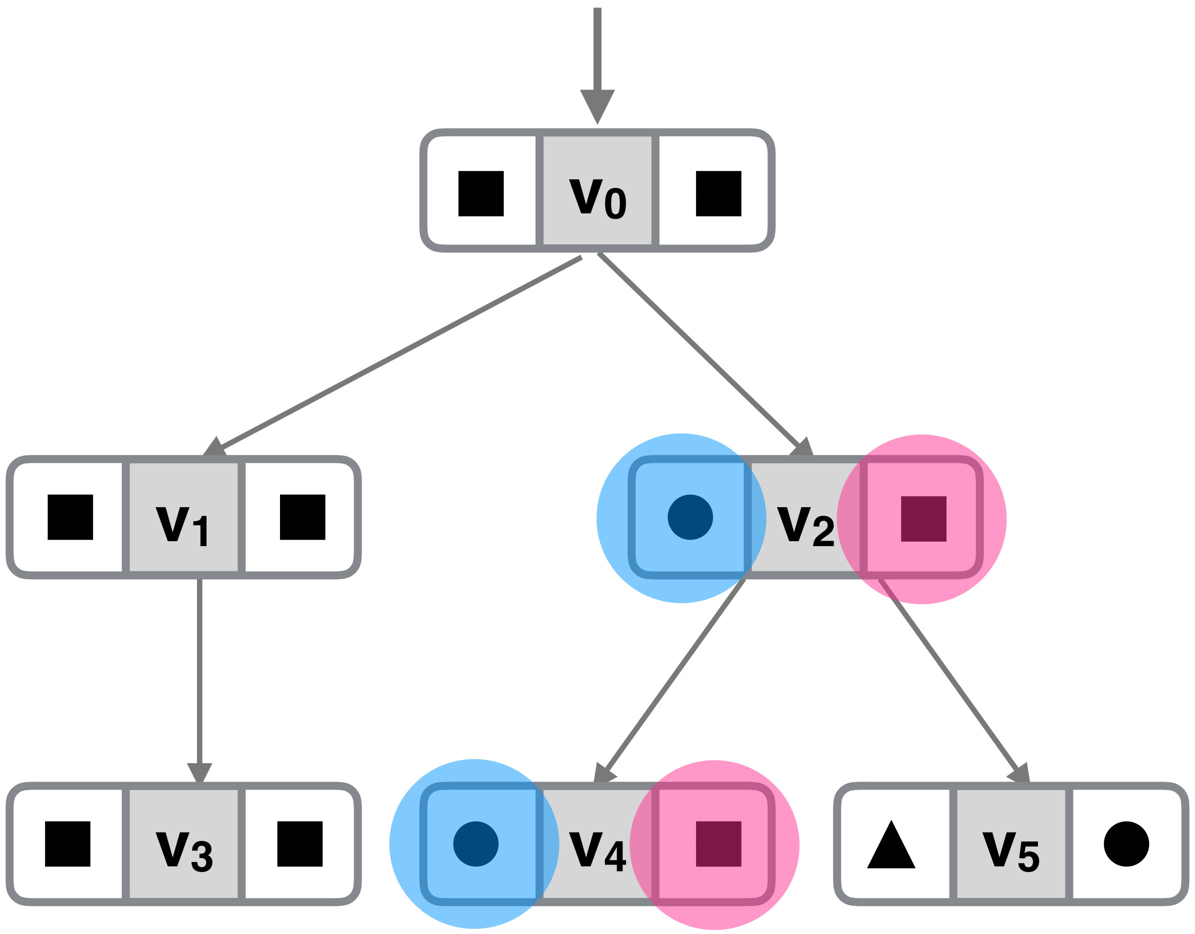

Hierarchical Observation

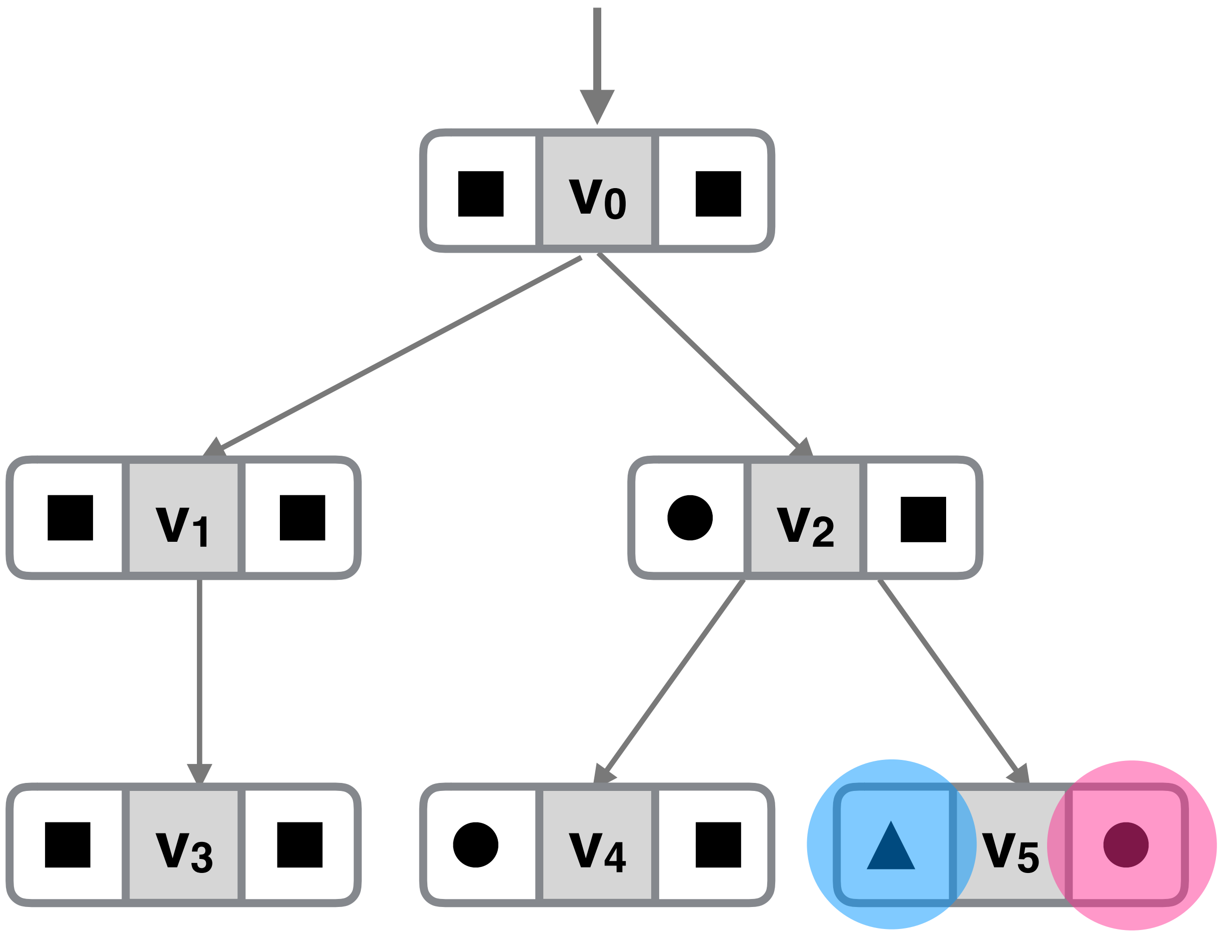

Hierarchical Observation

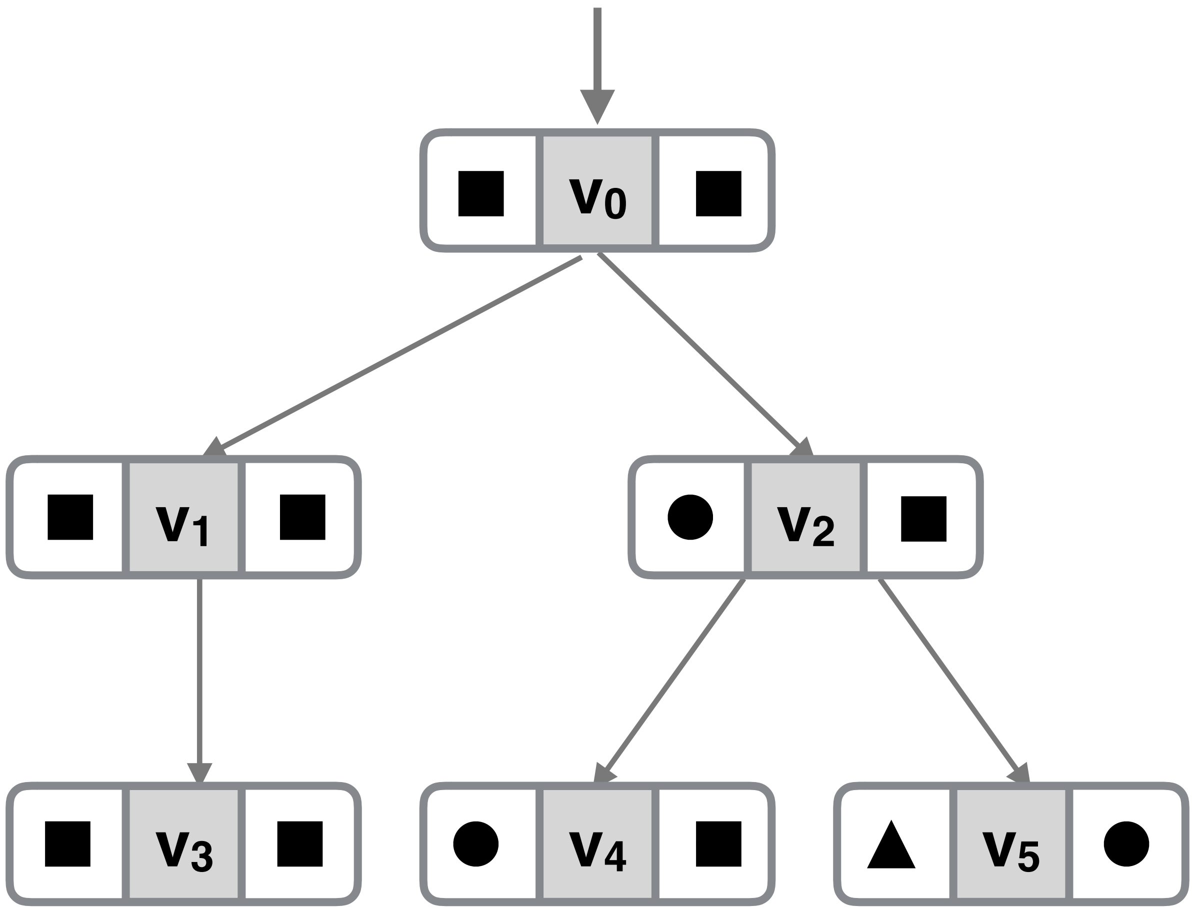

$\mathbf{1} <_{\text{obs}} \mathbf{2}$

Incorporating Perfect Recall

Compare knowledge of players accumulated from the start

"observe" $\rightsquigarrow$ "deduce"

Incorporating Perfect Recall

Compare knowledge of players accumulated from the start

"observe" $\rightsquigarrow$ "deduce"

Information set of player i at history π:

$$P^i(\pi) = \{\pi' \in Hist(G) | \pi' \sim^i \pi\}$$

Incorporating Perfect Recall

Compare knowledge of players accumulated from the start

"observe" $\rightsquigarrow$ "deduce"

Information set of player i at history π:

$$P^i(\pi) = \{\pi' \in Hist(G) | \pi' \sim^i \pi\}$$

Player $i \leq_{inf}$ Player k: for all histories π, $P^i(\pi) \subseteq P^k(\pi)$

Hierarchical *information*: $\leq_{inf}$ total order

Hierarchical Information

Hierarchical Information

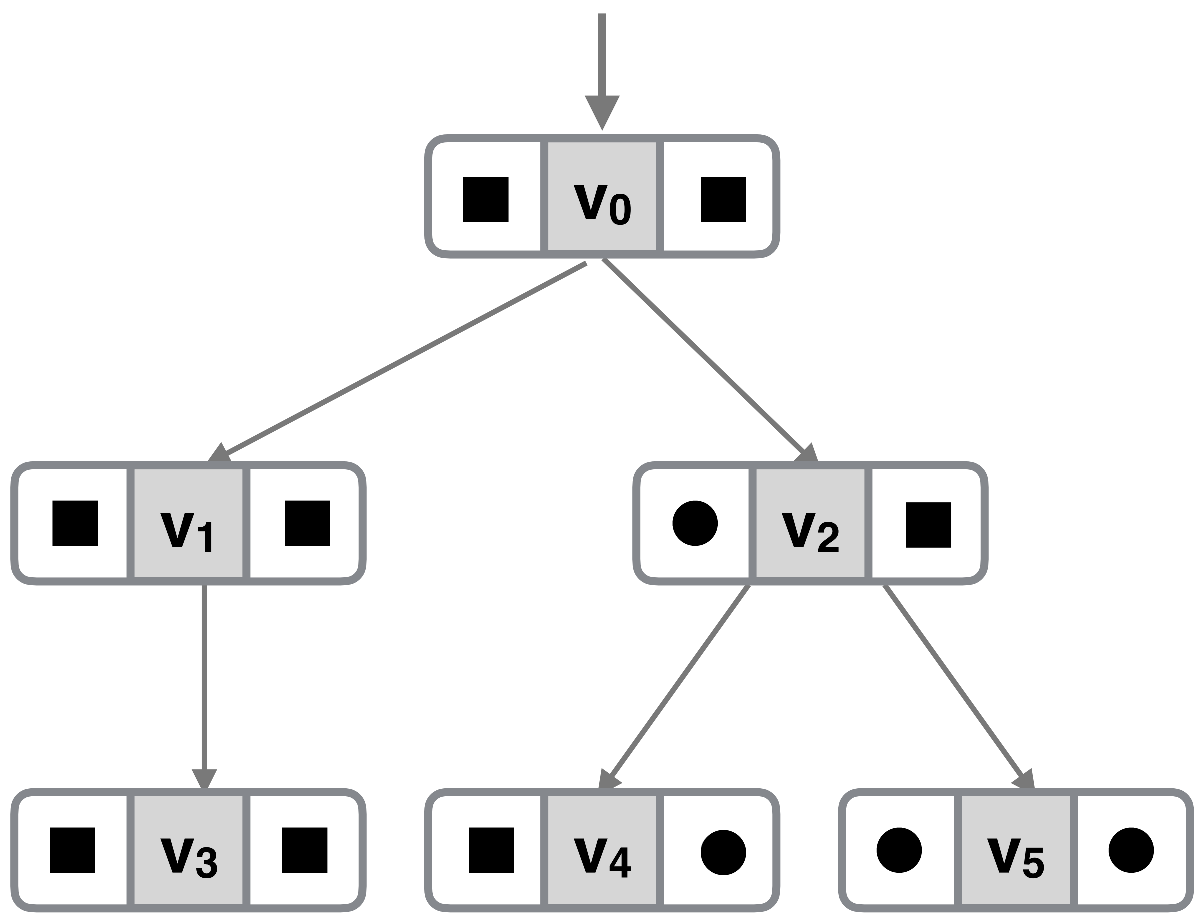

Hierarchical Information

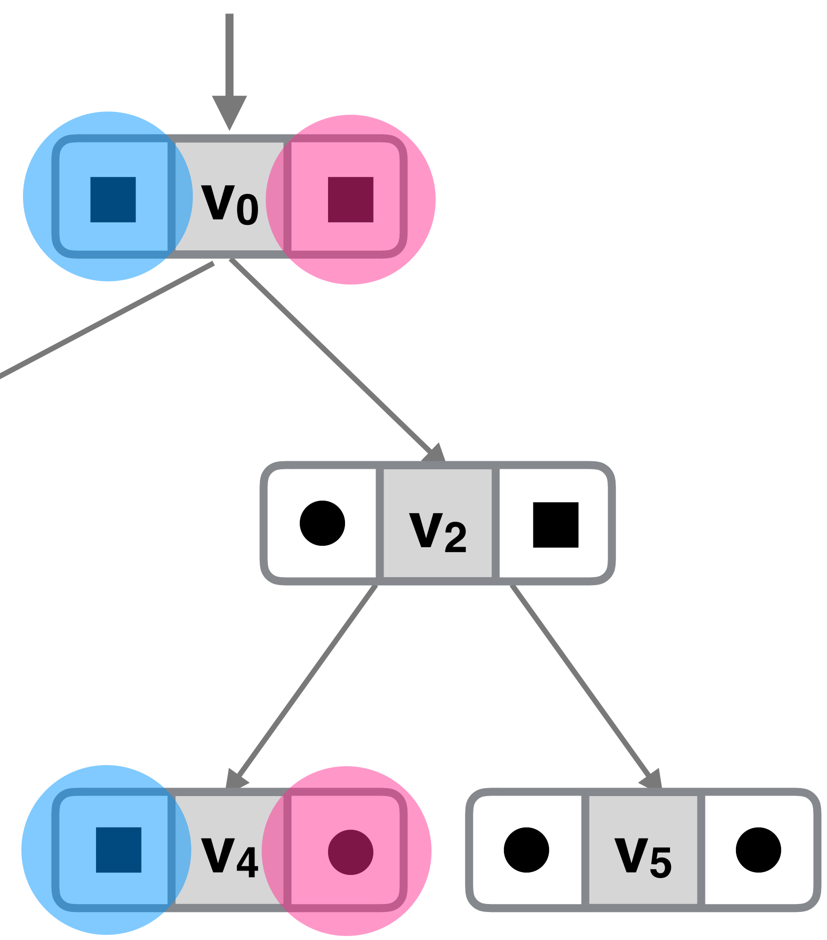

Hierarchical Information

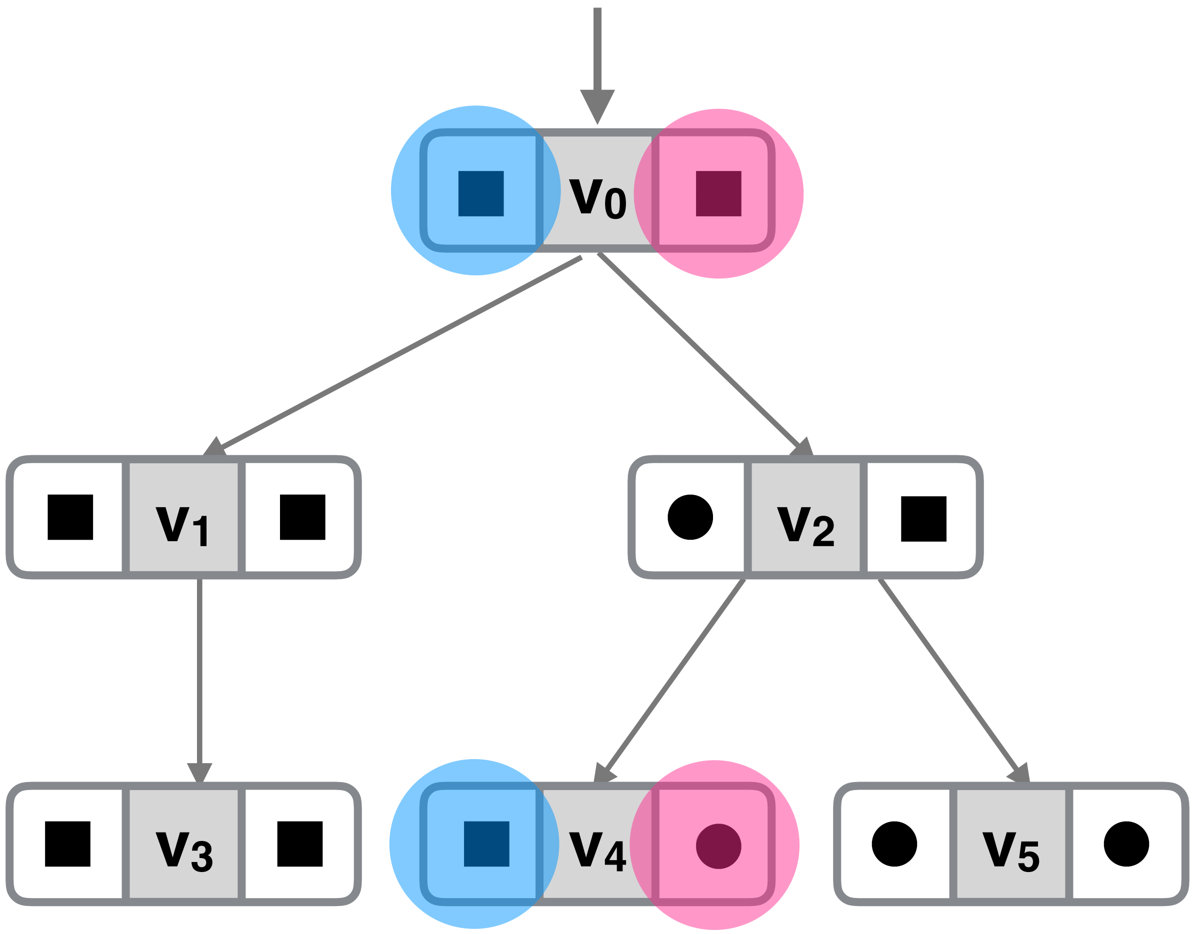

$\mathbf{1} \not\prec_{obs} \mathbf{2}$

Hierarchical Information

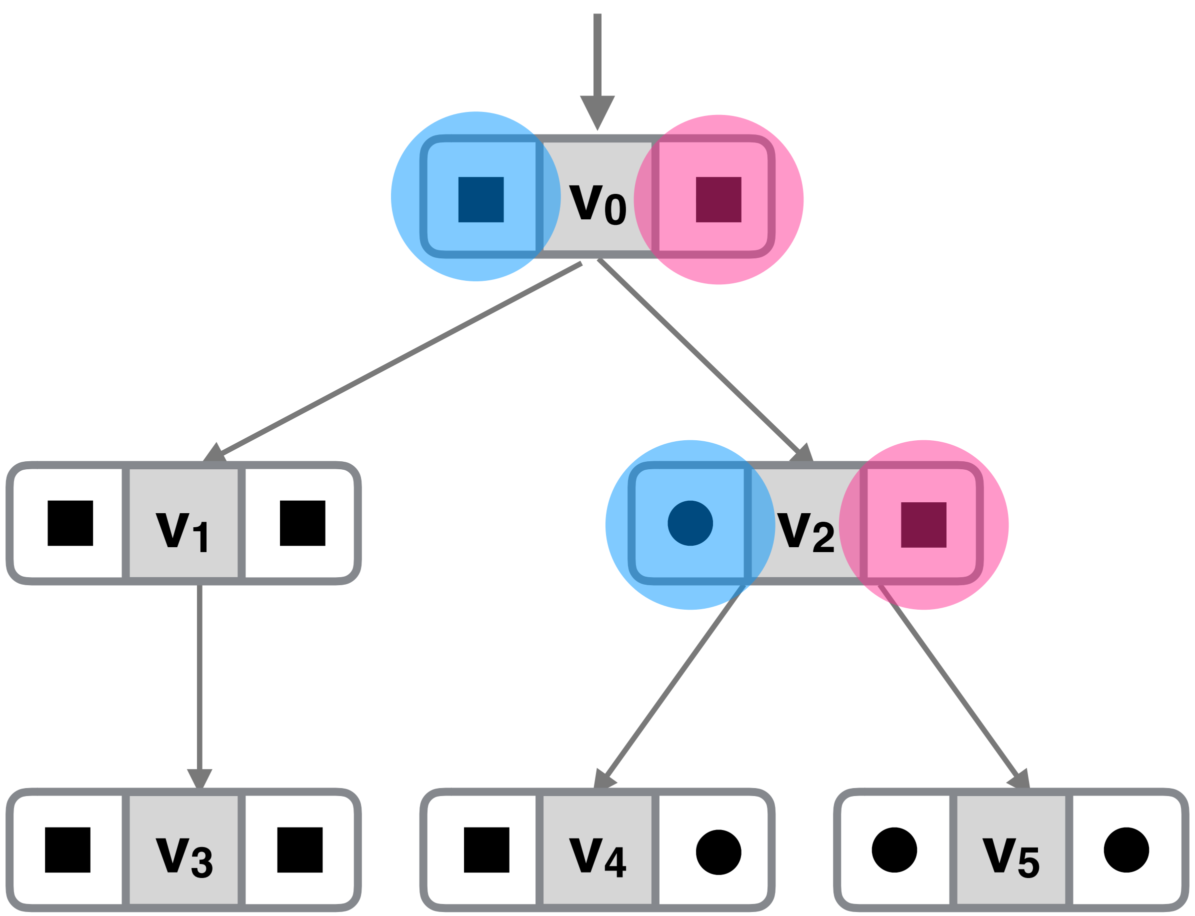

Hierarchical Information

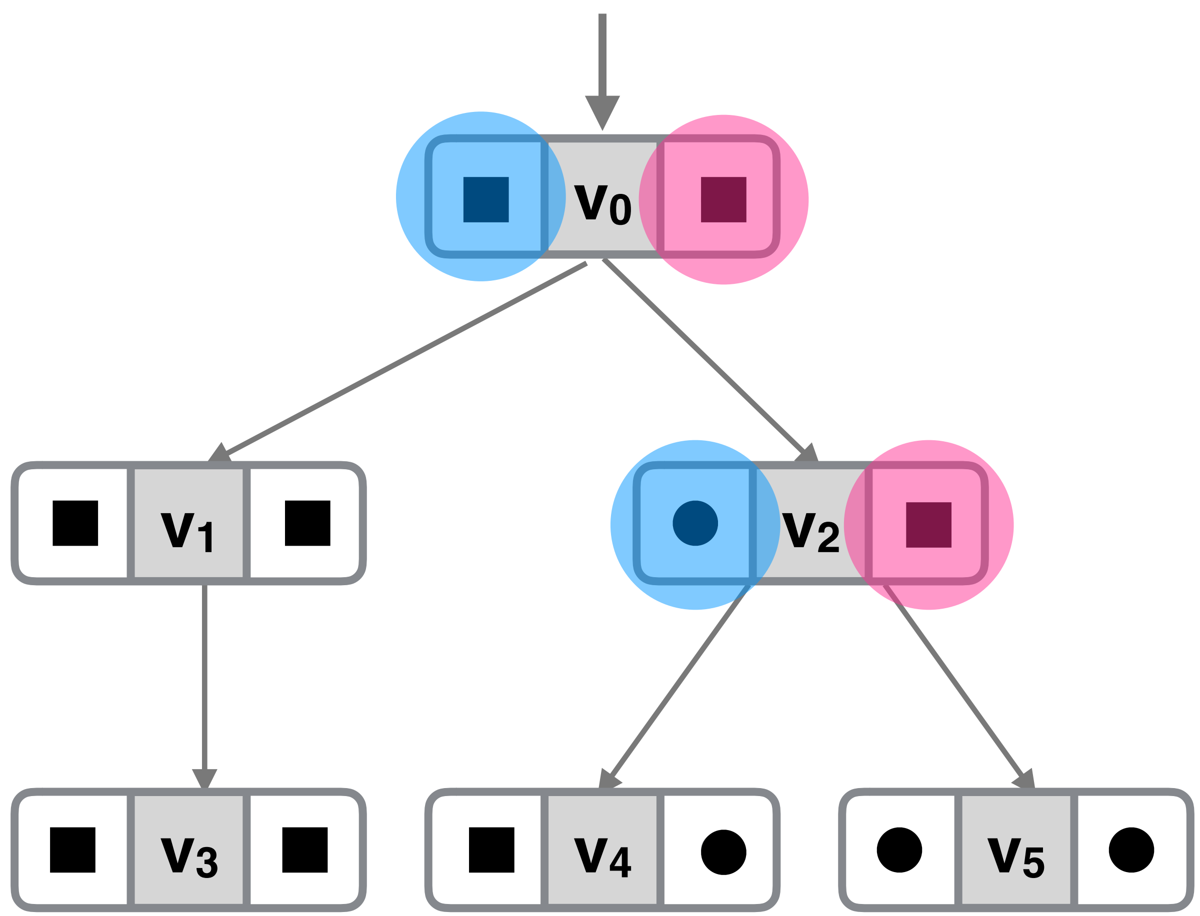

$2 \not\preccurlyeq_{obs} 1$

Hierarchical Information

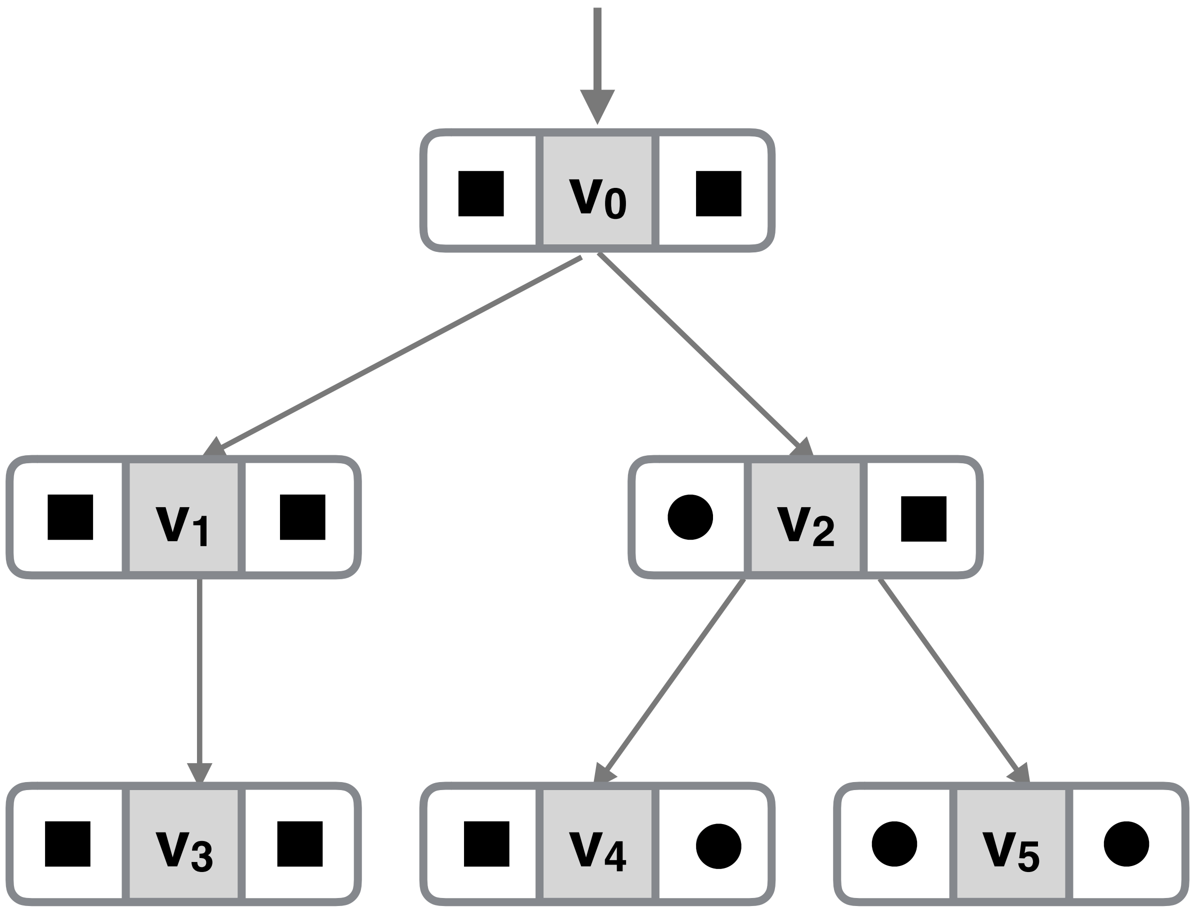

Hierarchical Information

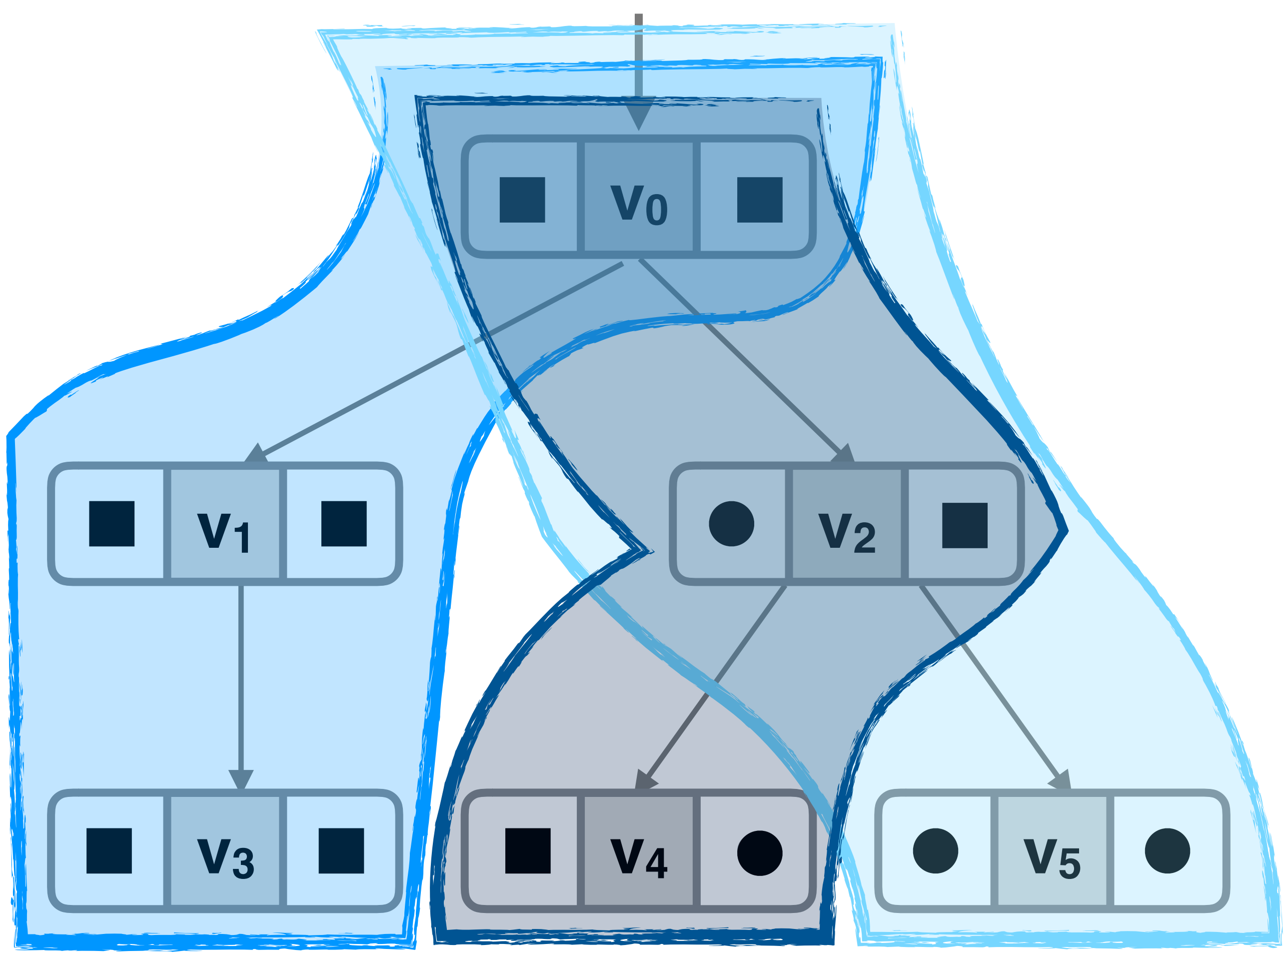

Hierarchical Information

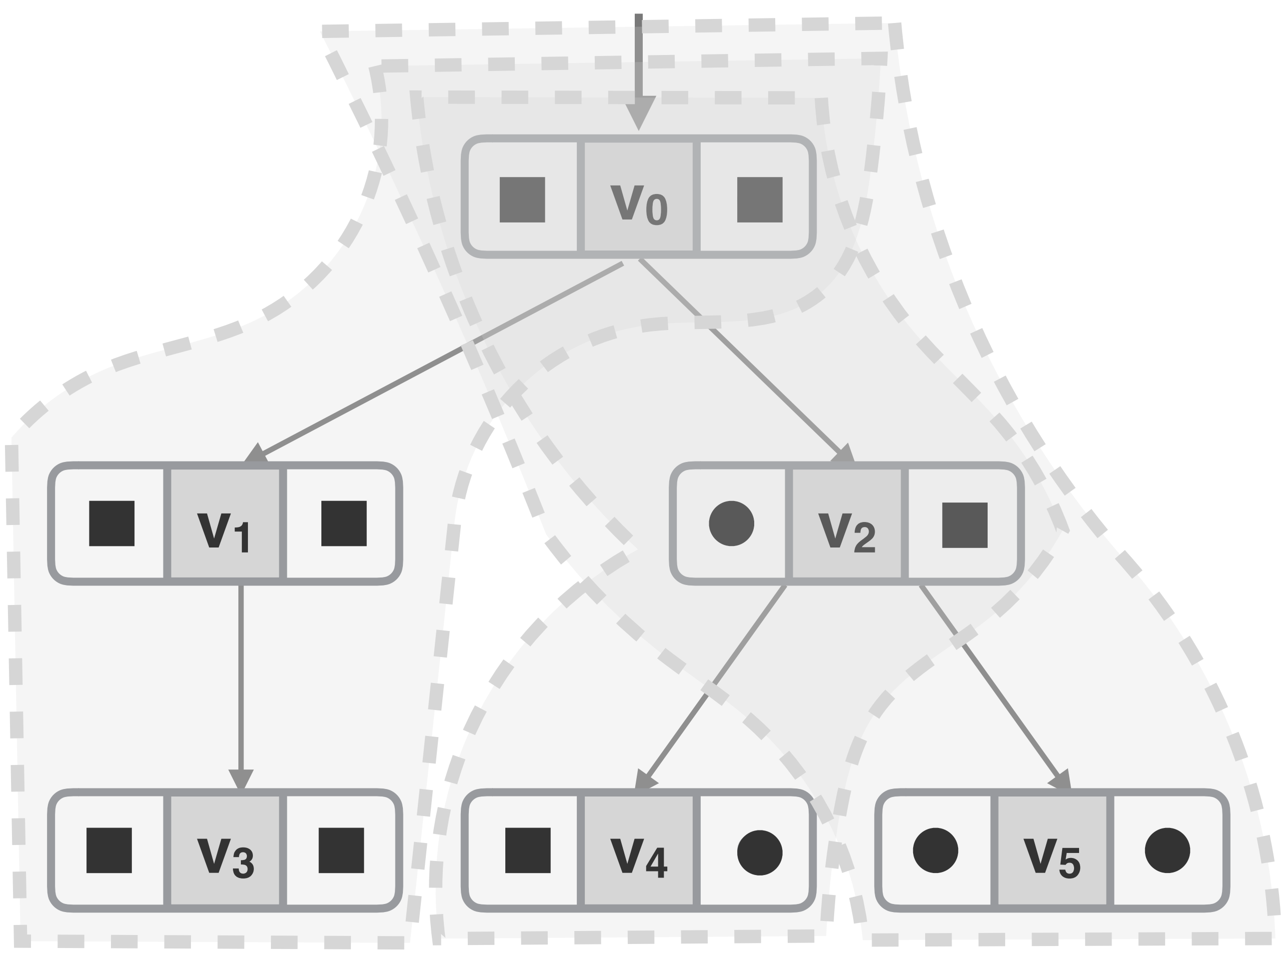

Hierarchical Information

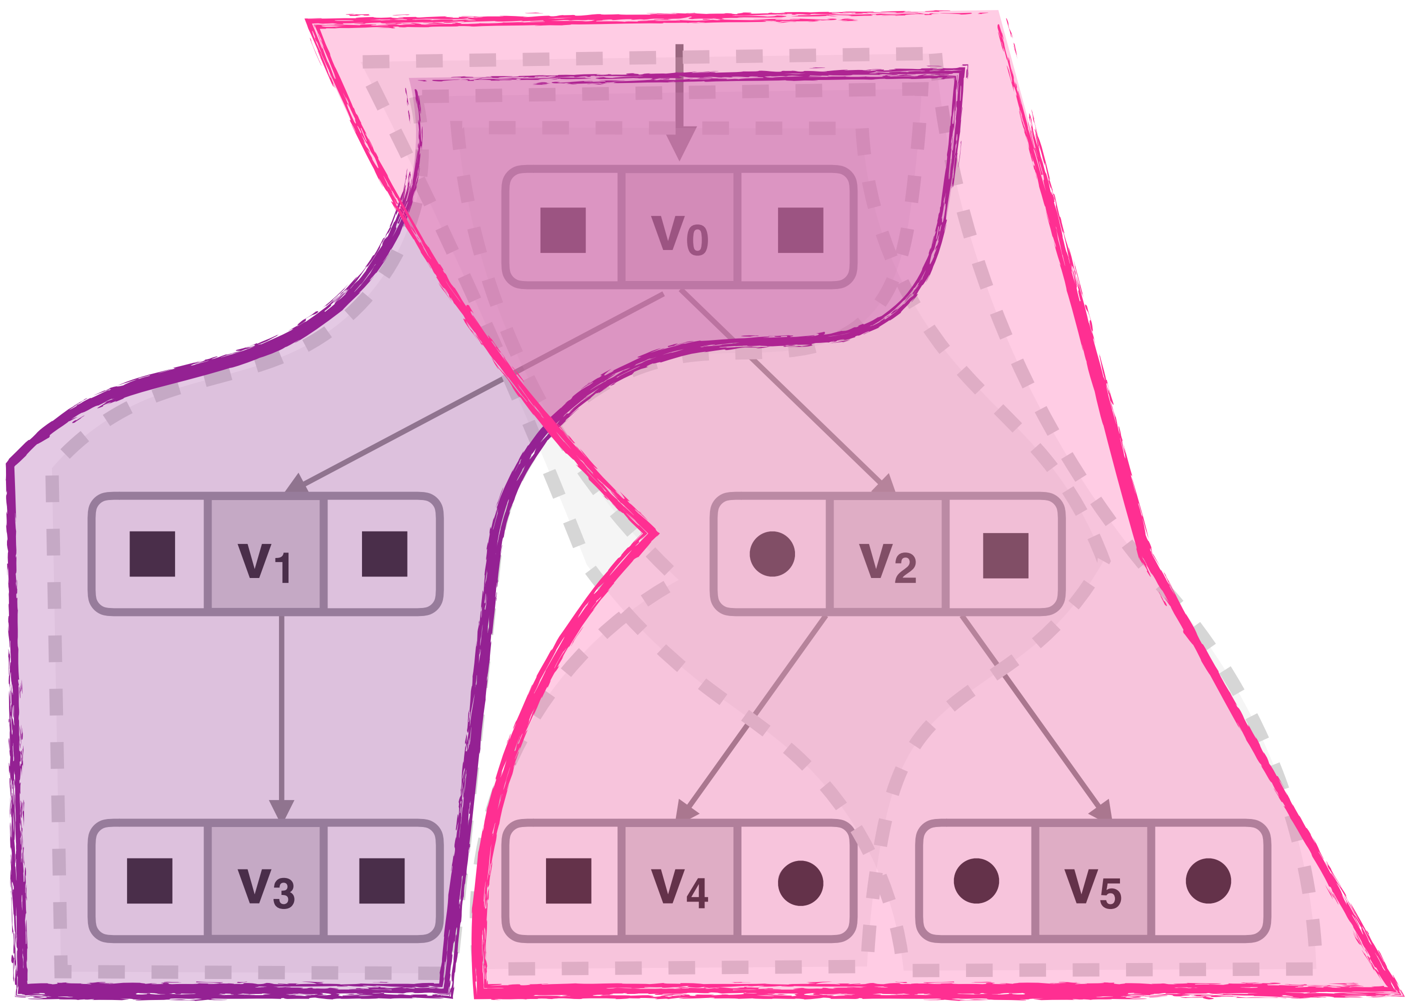

Hierarchical Information

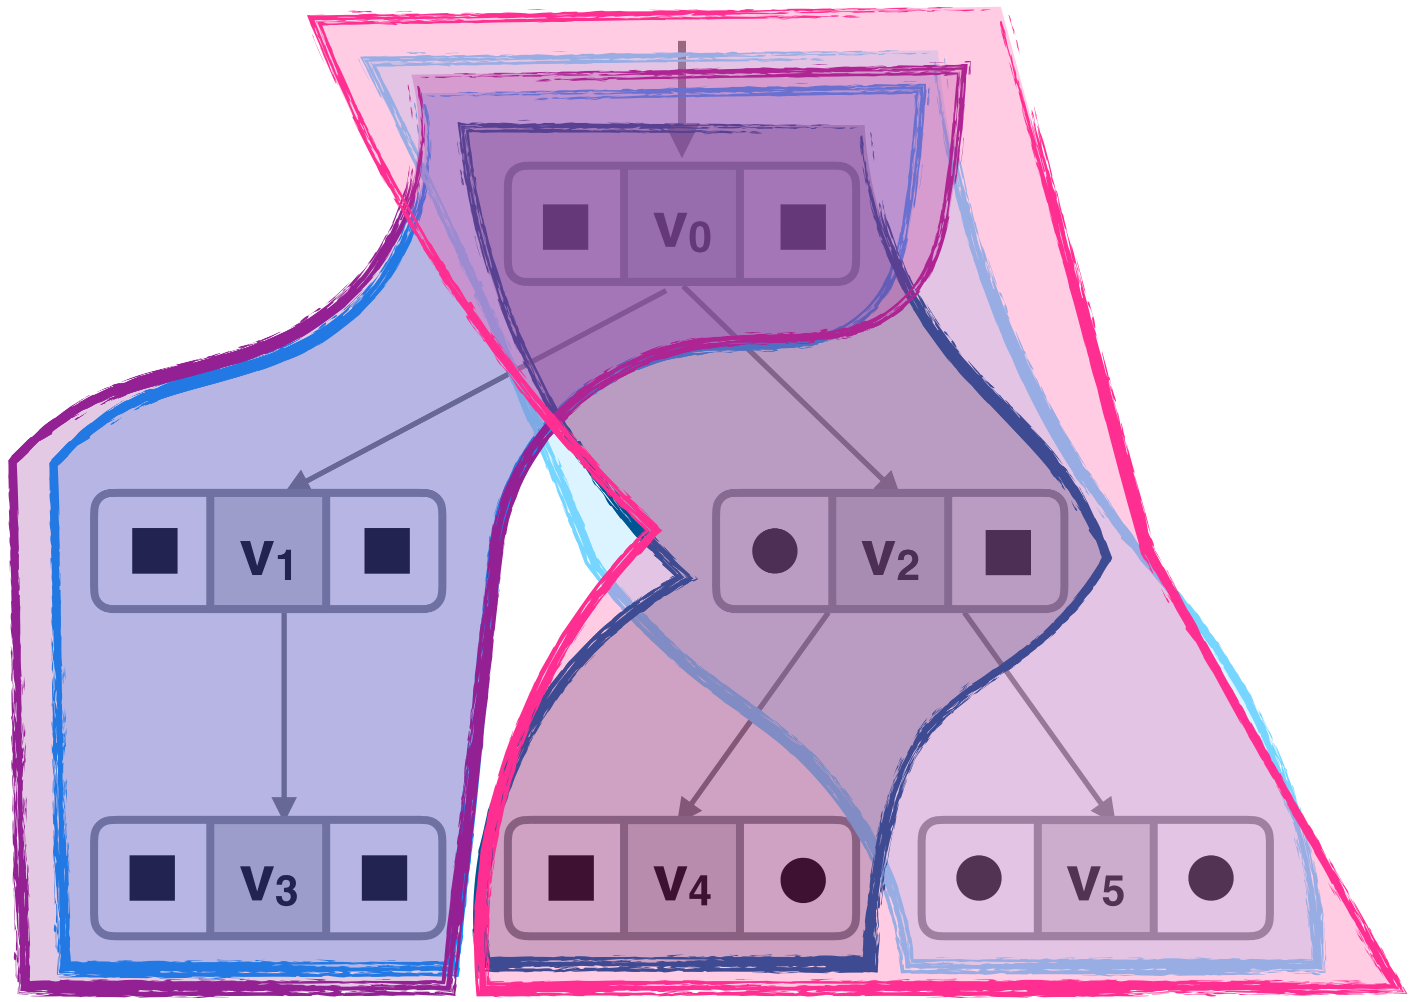

$\mathbf{1} <_{\inf} \mathbf{2}$

Hierarchical Information

Proposition: With hierarchical *information*,
distributed strategy synthesis is decidable.

Idea: Monitor observations of the less informed players
attach to states these "new" observations

$\rightsquigarrow$ "equivalent" game with hierarchical *observation*

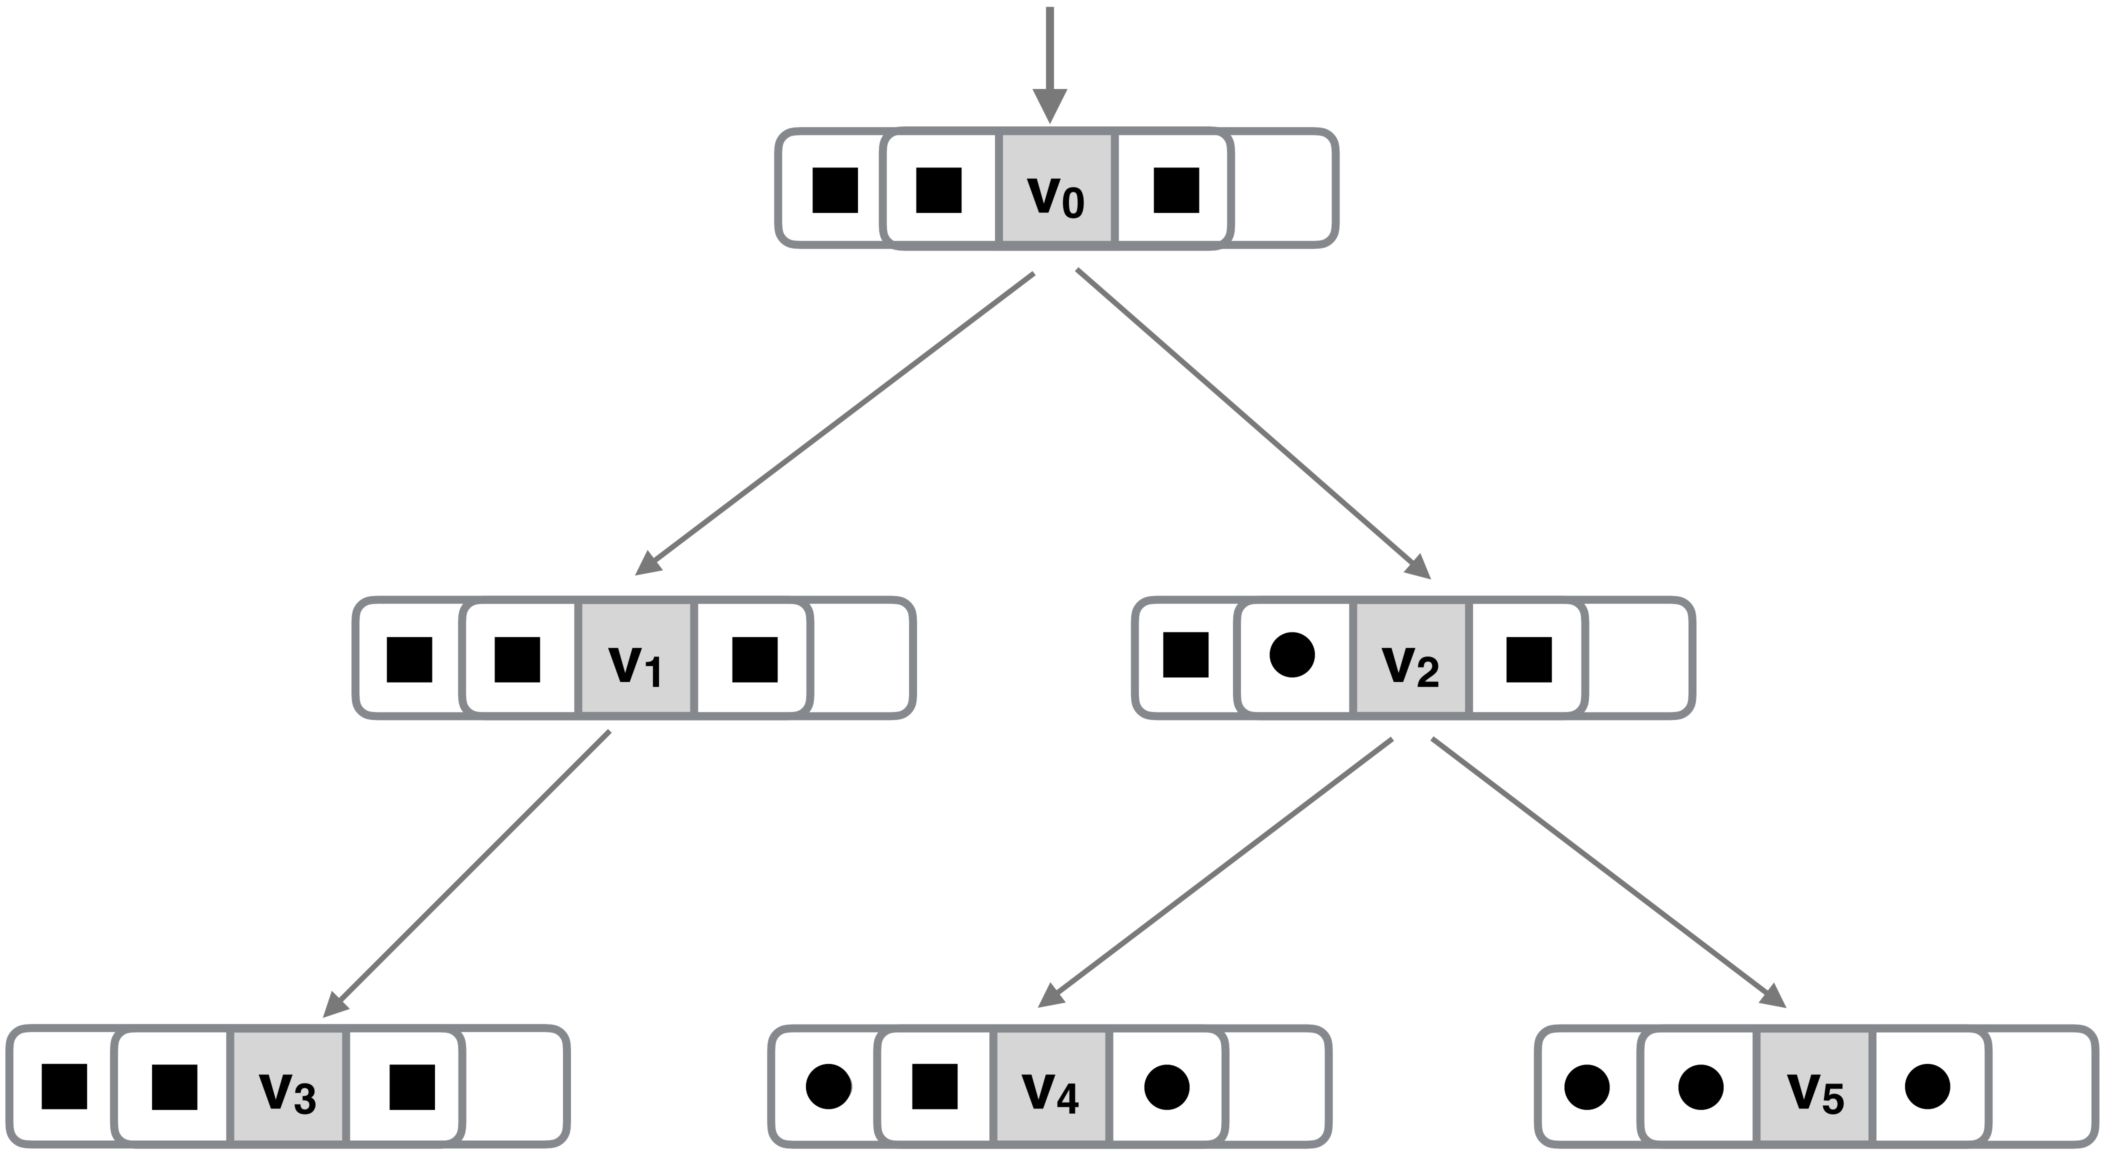
v0
v1
v2
v3
v4
v5

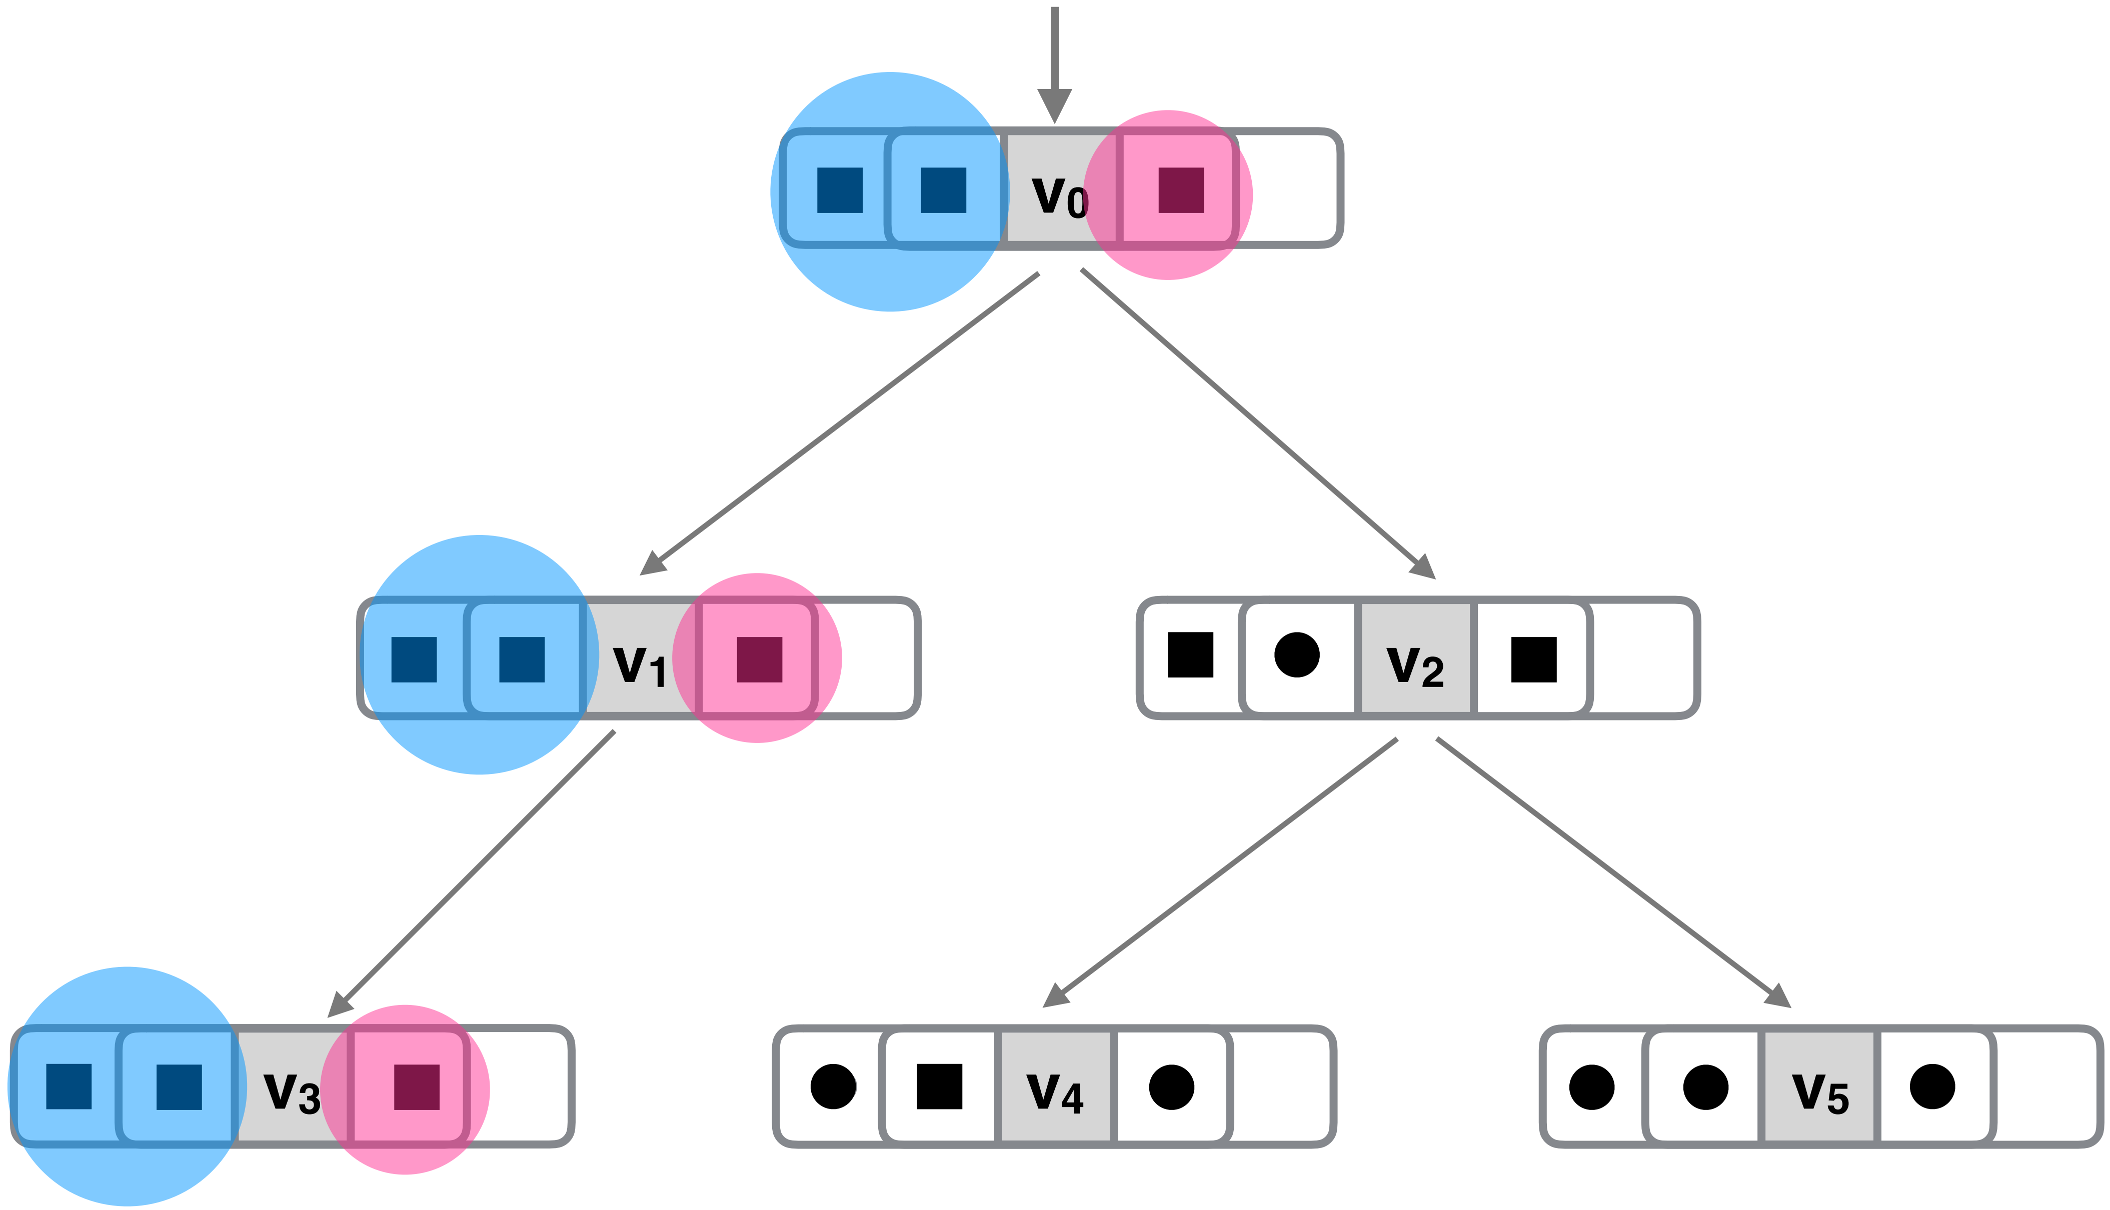
v0
v1
v2
v3
v4
v5

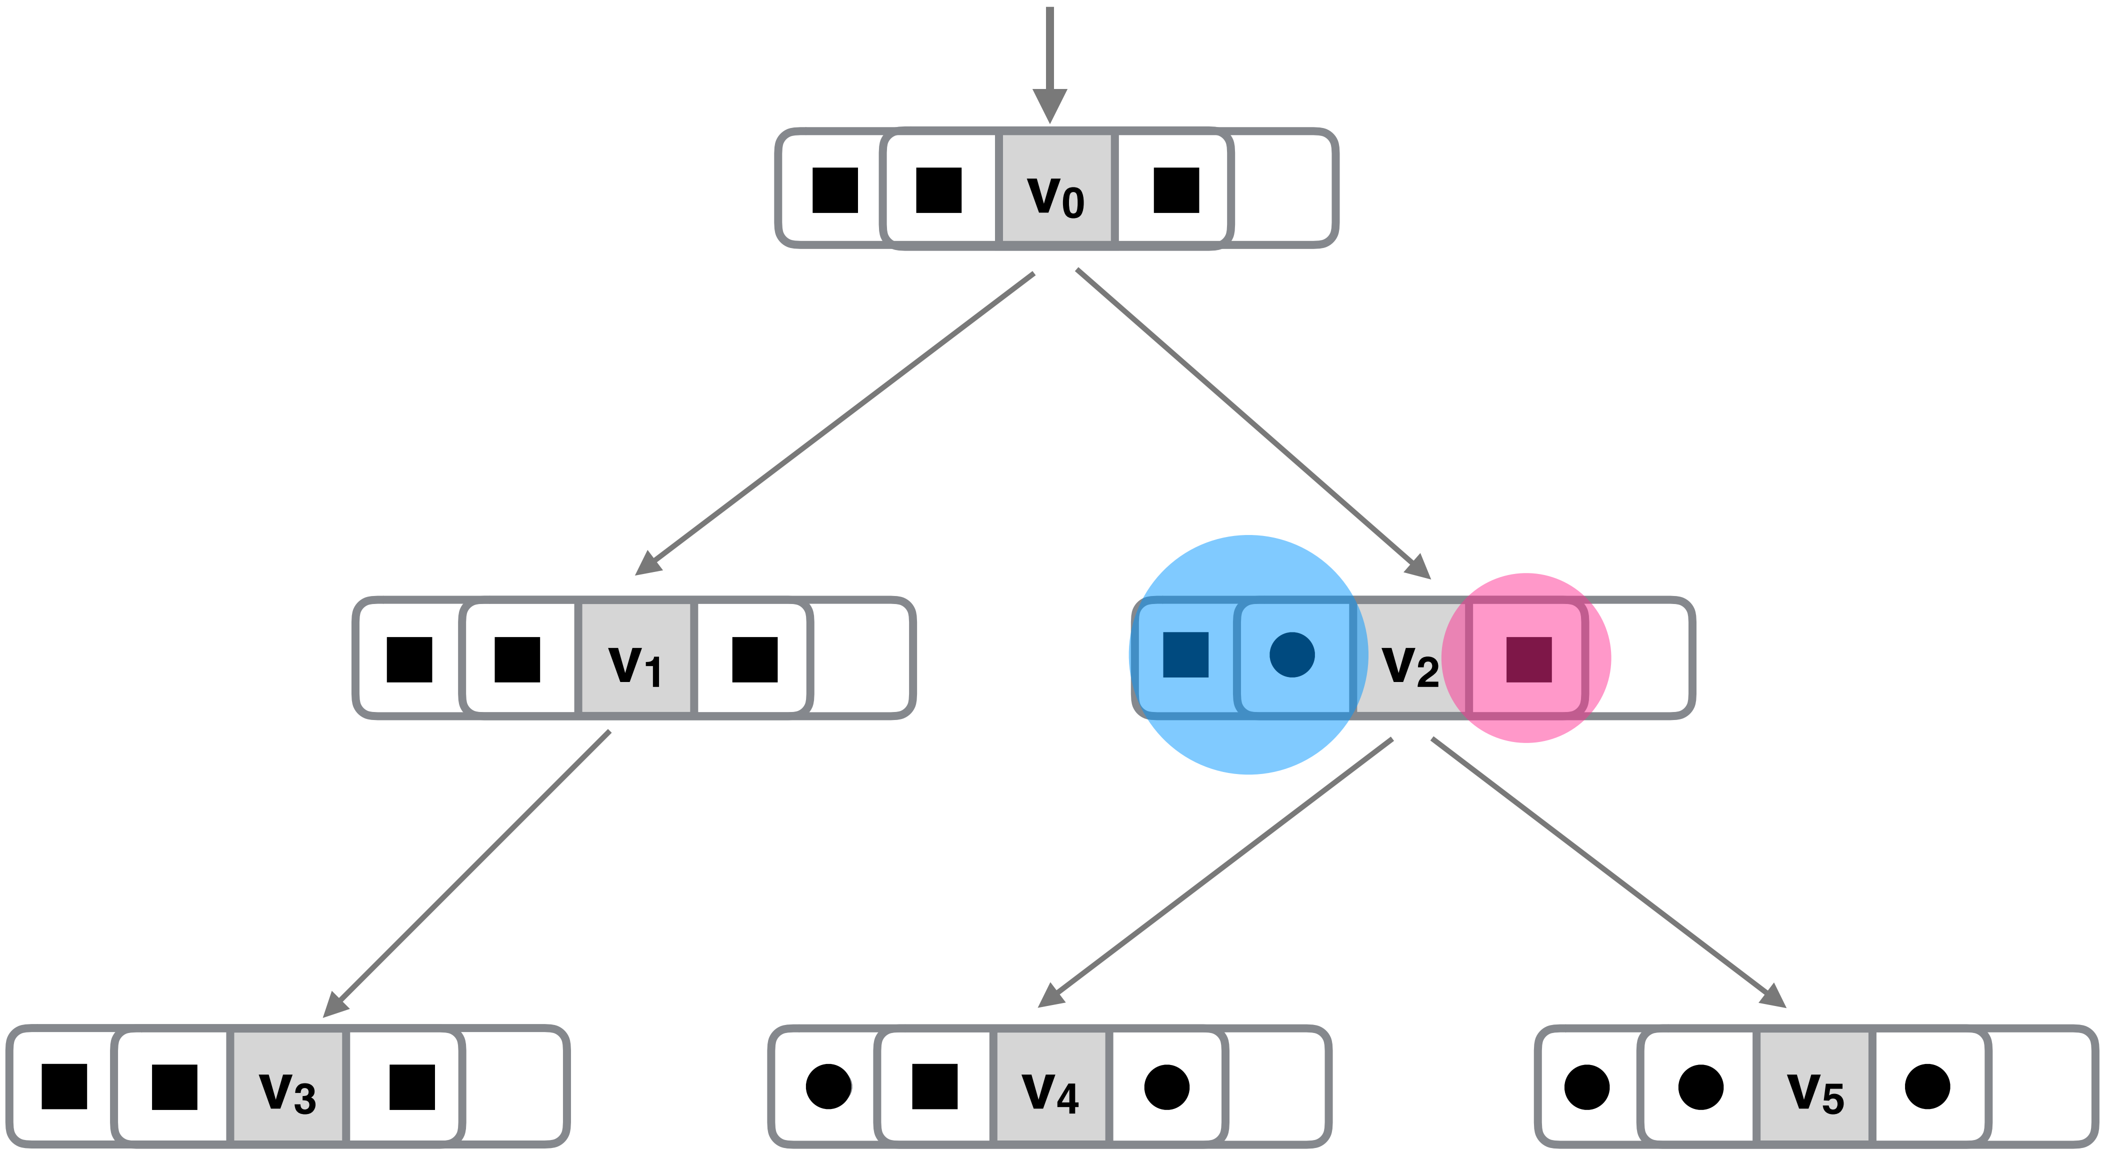
v0
v1
v2
v3
v4
v5

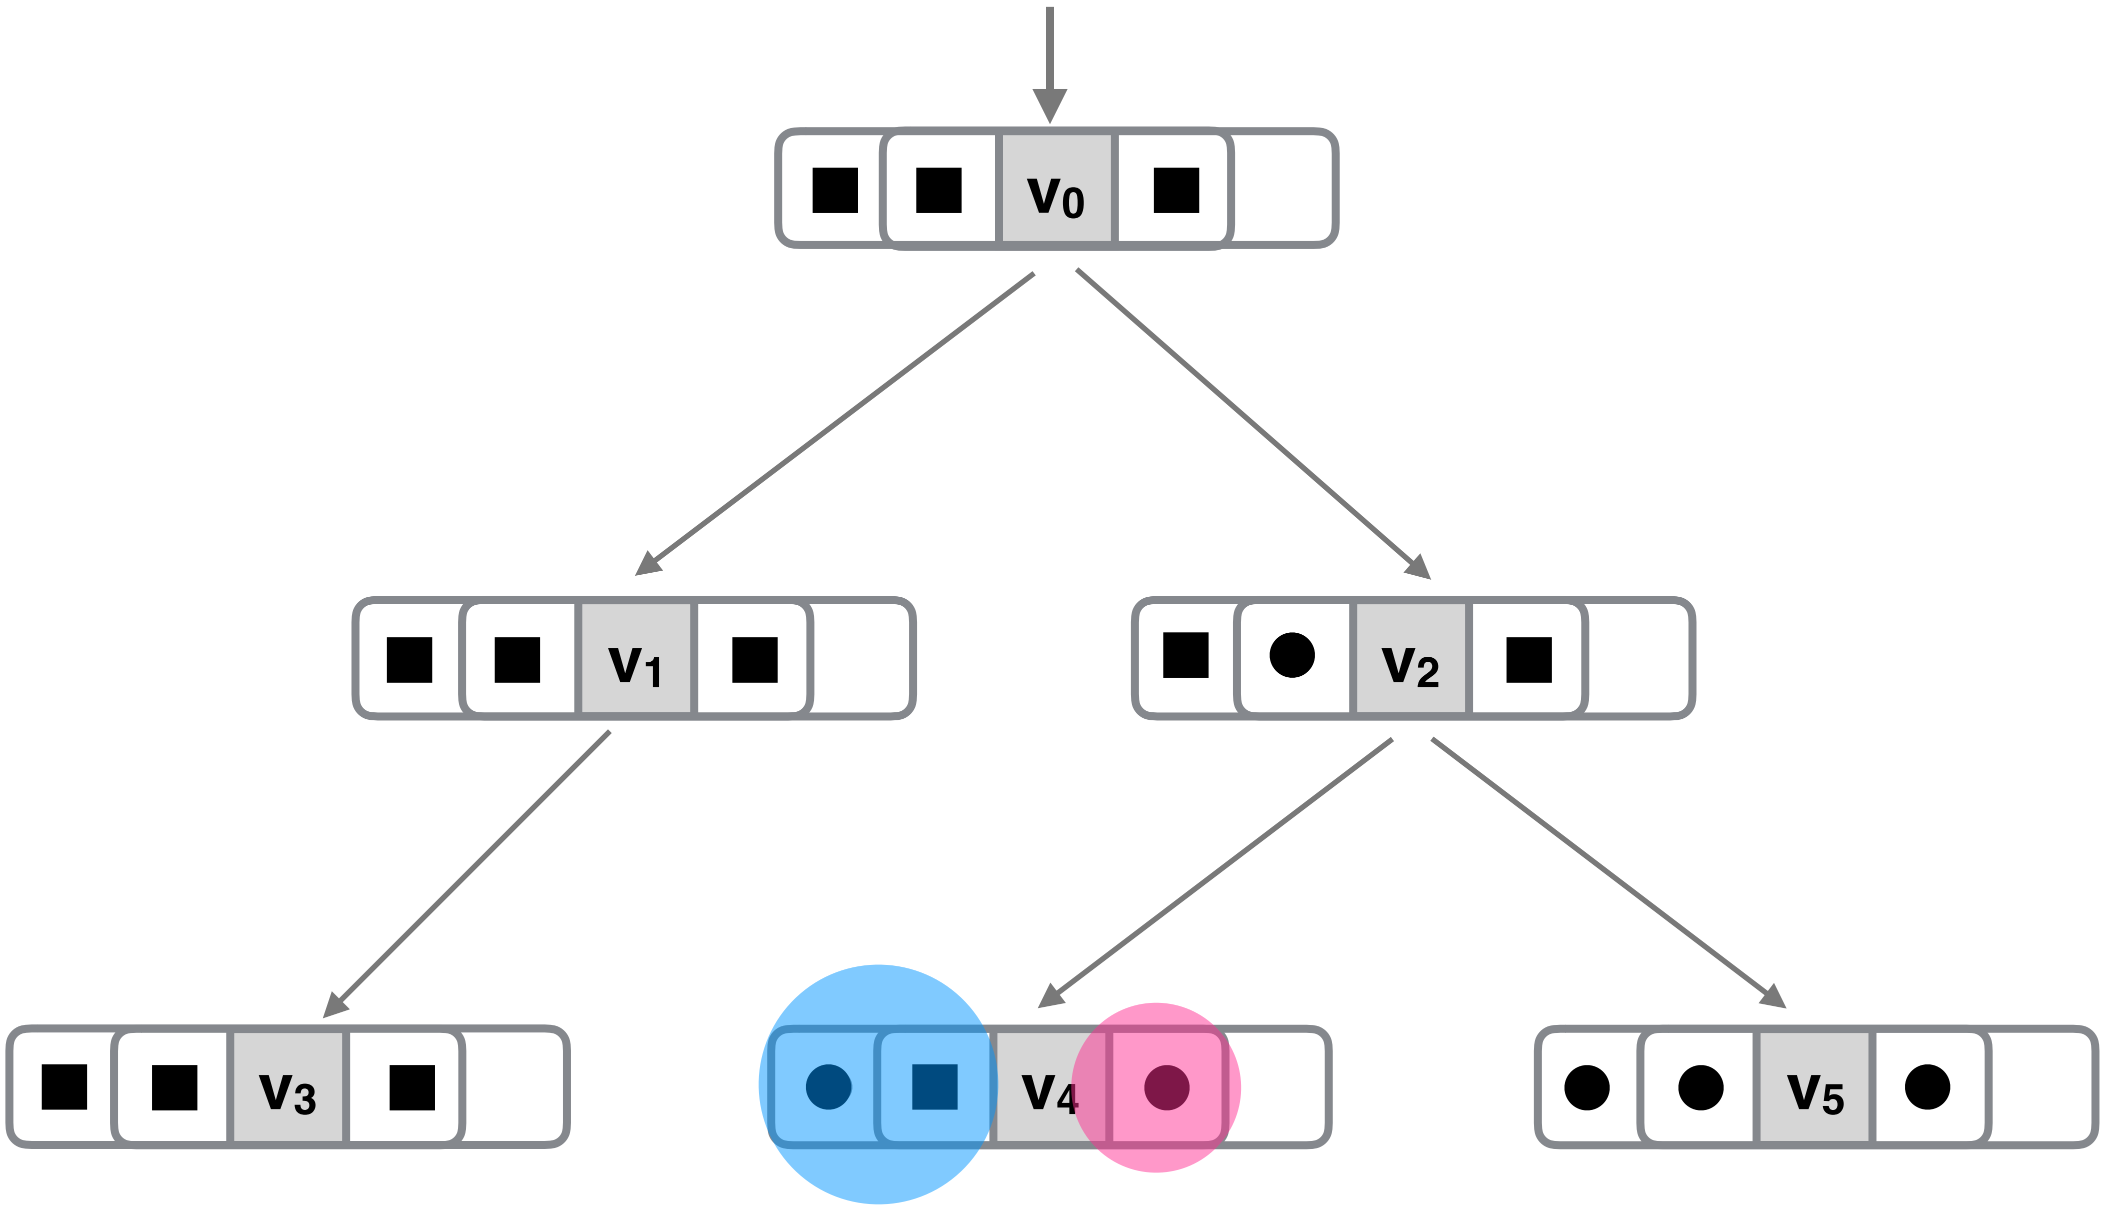

v0
v1
v2
v3
v4
v5

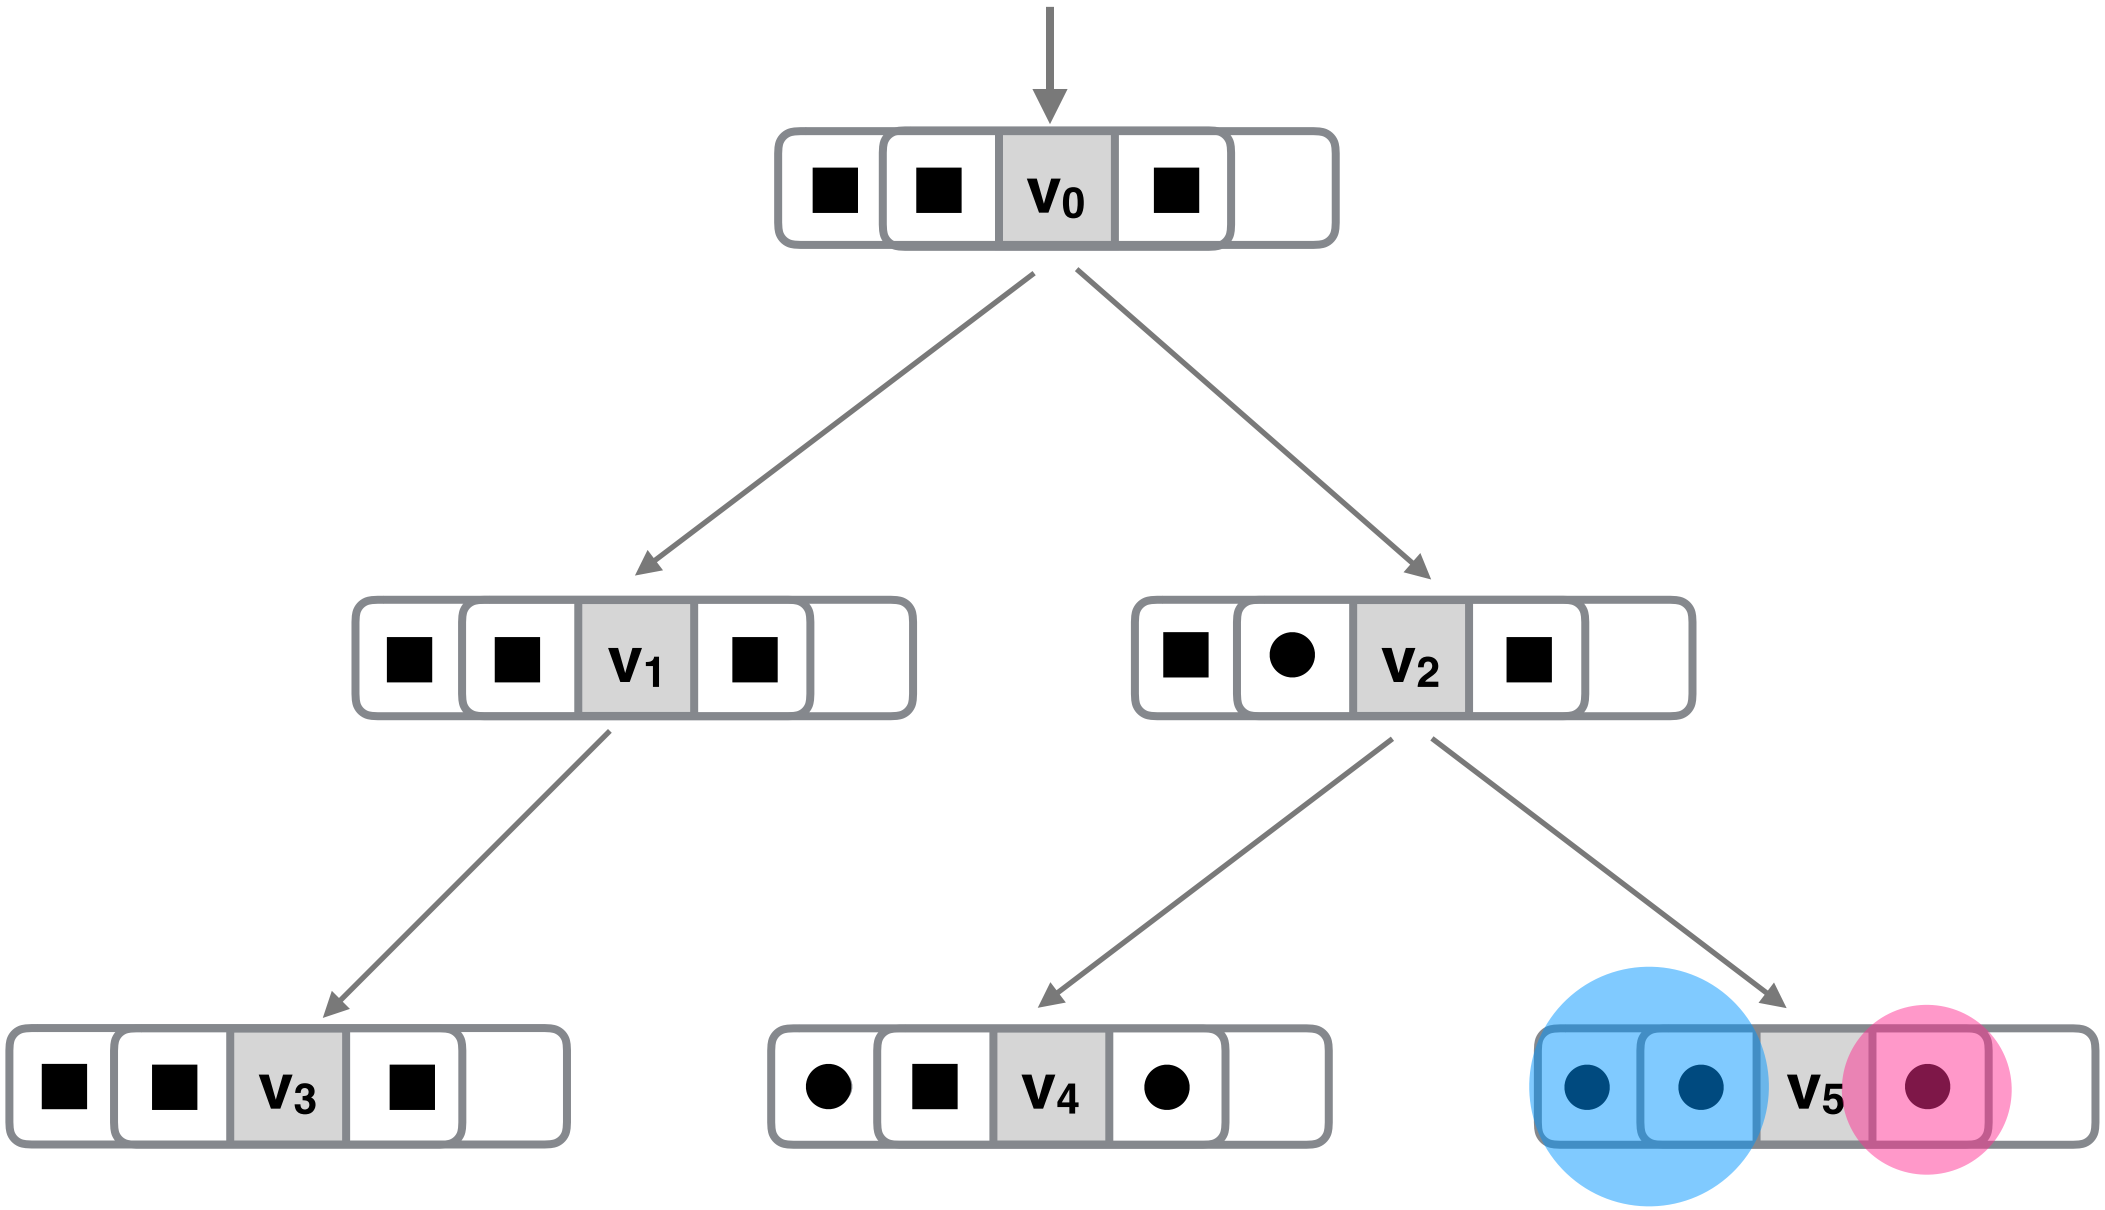

v_0
v_1
v_2
v_3
v_4
v_5

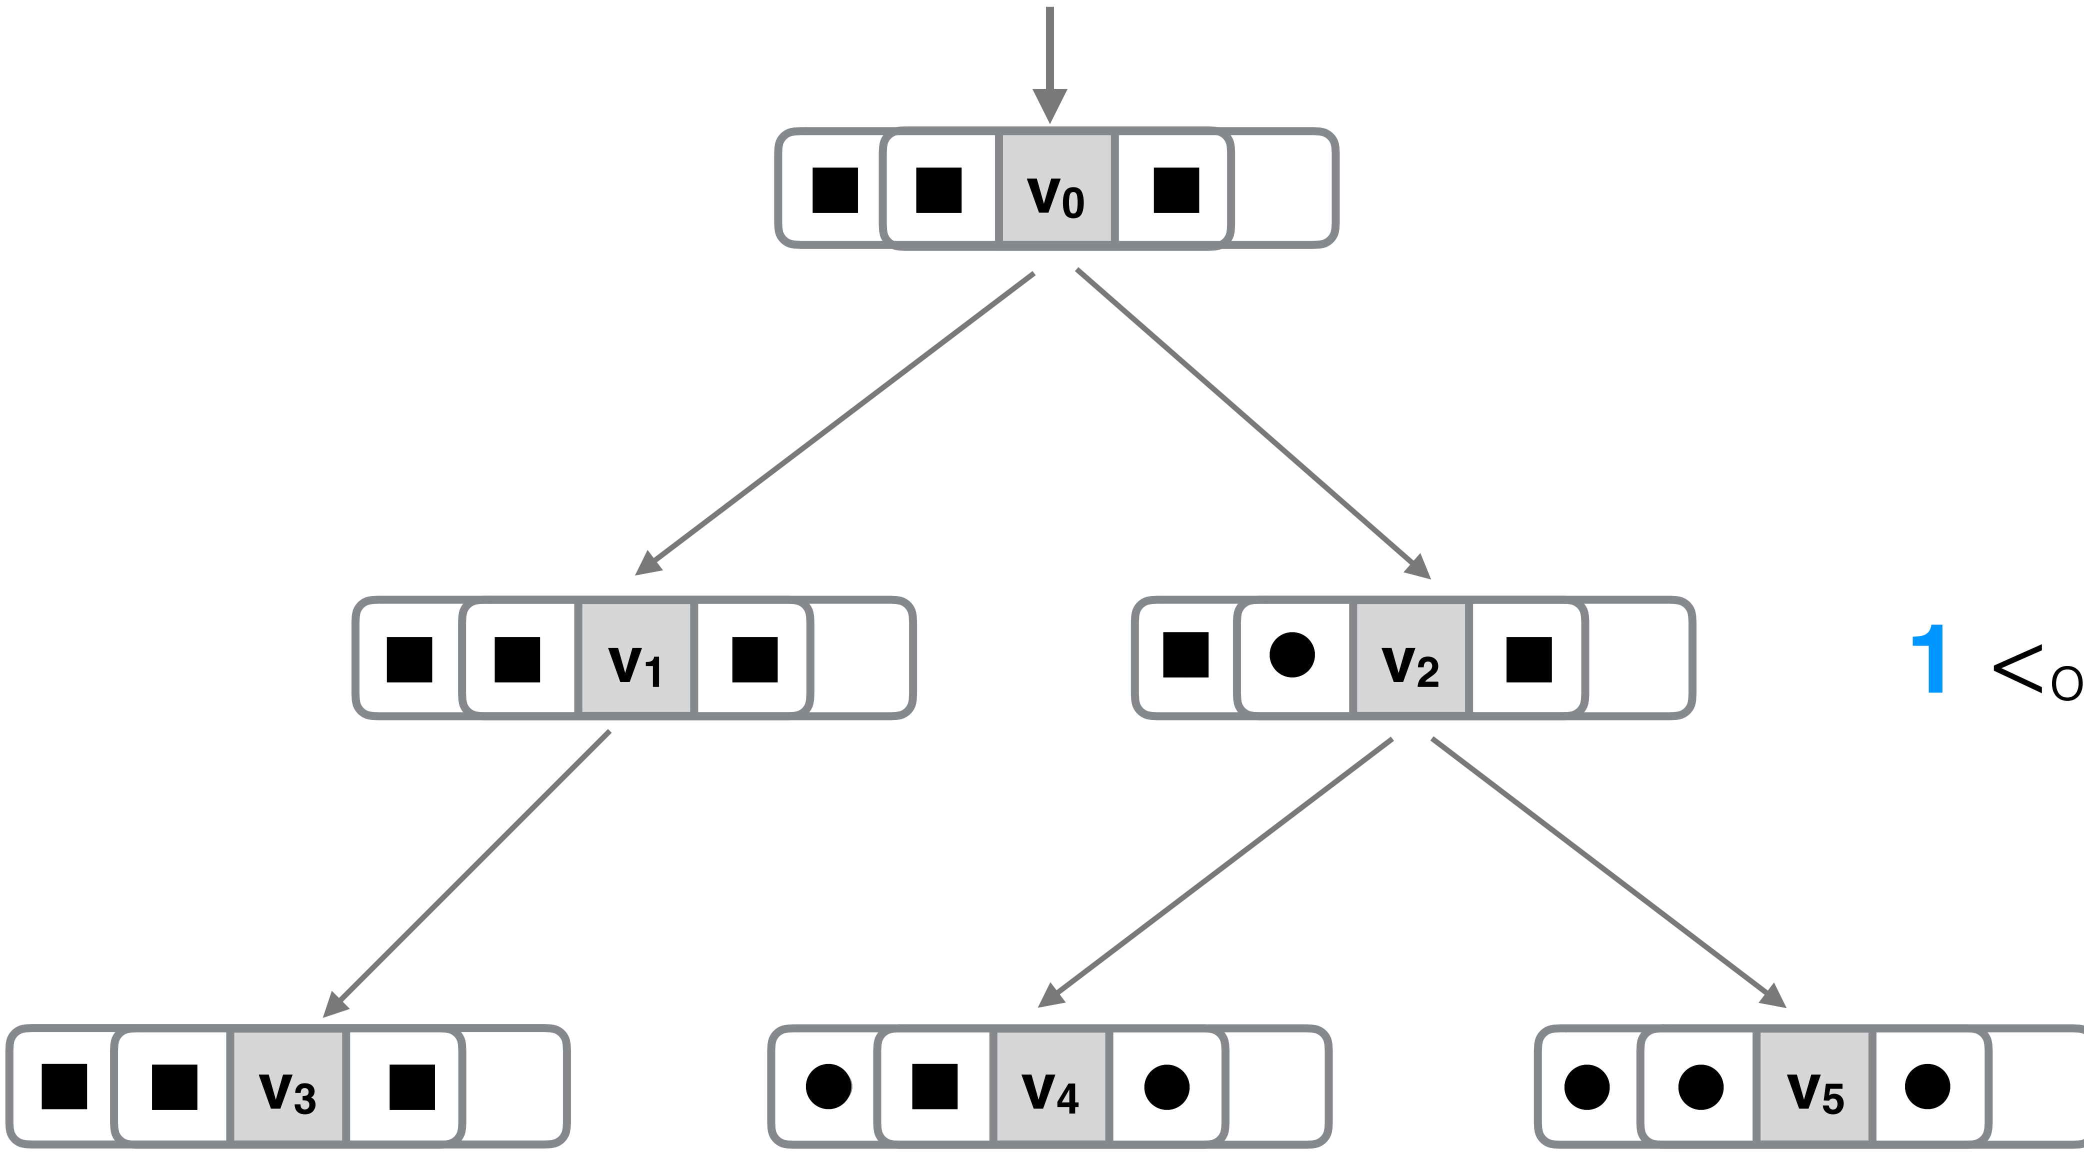

$1 <_{obs} 2$

Finite-State Signals

Signal: for a game G, colours C,

$$f : \text{history } \pi \to C$$

- *information-consistent* for player i: $\pi \sim^i \pi' \to f(\pi) = f(\pi')$
- *finite-state*: Mealy machine – over *global* states

Finite-State Signals

Signal: for a game G, colours C,

$$f : \text{history } \pi \to C$$

- *information-consistent* for player i: $\pi \sim^i \pi' \to f(\pi) = f(\pi')$
- *finite-state*: Mealy machine – over *global* states

Lemma: Any finite-state signal that is information-consistent for i is implementable by a Mealy machine M^i – over observations of i.

Finite-State Signals

Signal: for a game G, colours C,

$$f : \text{history } \pi \to C$$

- *information-consistent* for player i: $\pi \sim^i \pi' \to f(\pi) = f(\pi')$
- *finite-state*: Mealy machine – over *global* states

Lemma: Any finite-state signal that is information-consistent for i is implementable by a Mealy machine M^i – over observations of i.

Synchronisation $G \times M^i$: same game as G,
but signals f^i now state observations

Landscape of Hierarchical Patterns

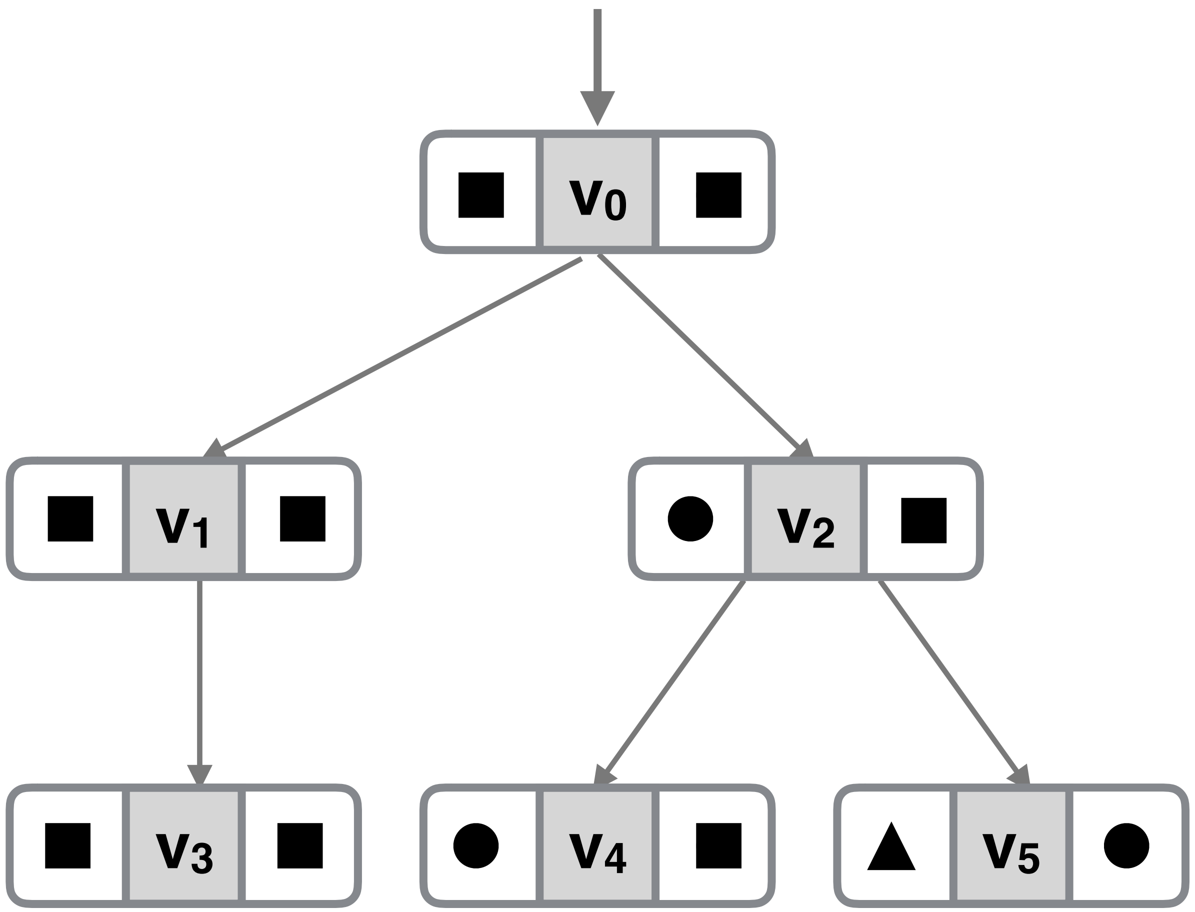

Hierarchical Observation

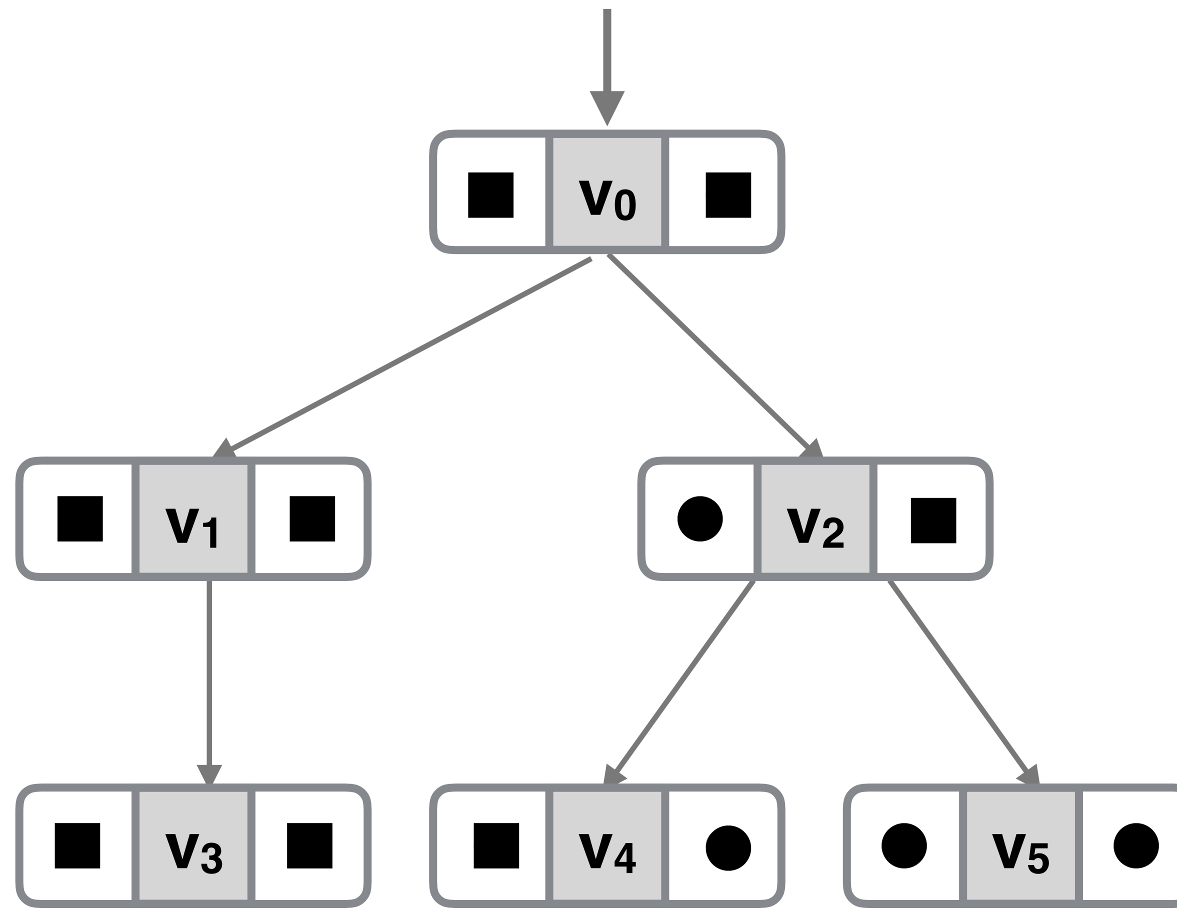

Hierarchical Information

Landscape of Hierarchical Patterns

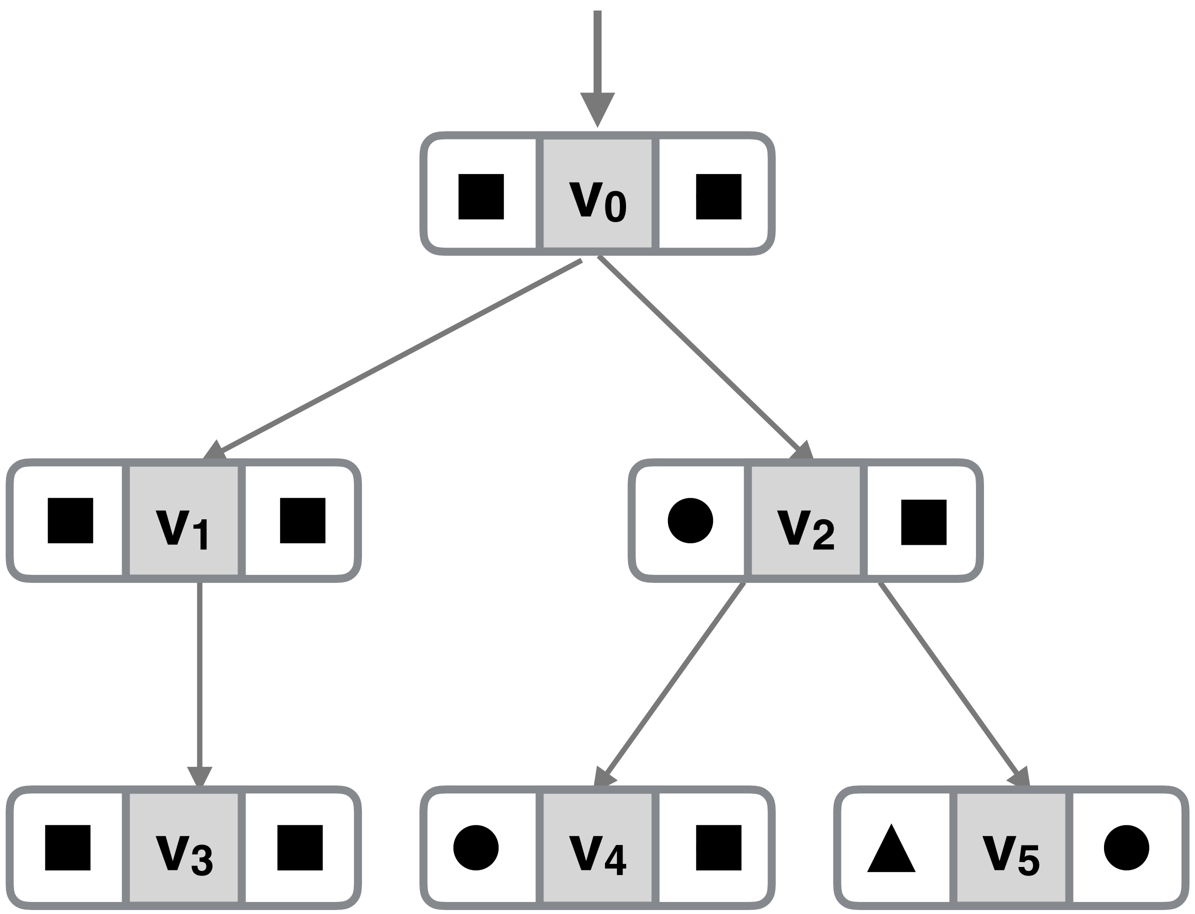

Hierarchical Observation

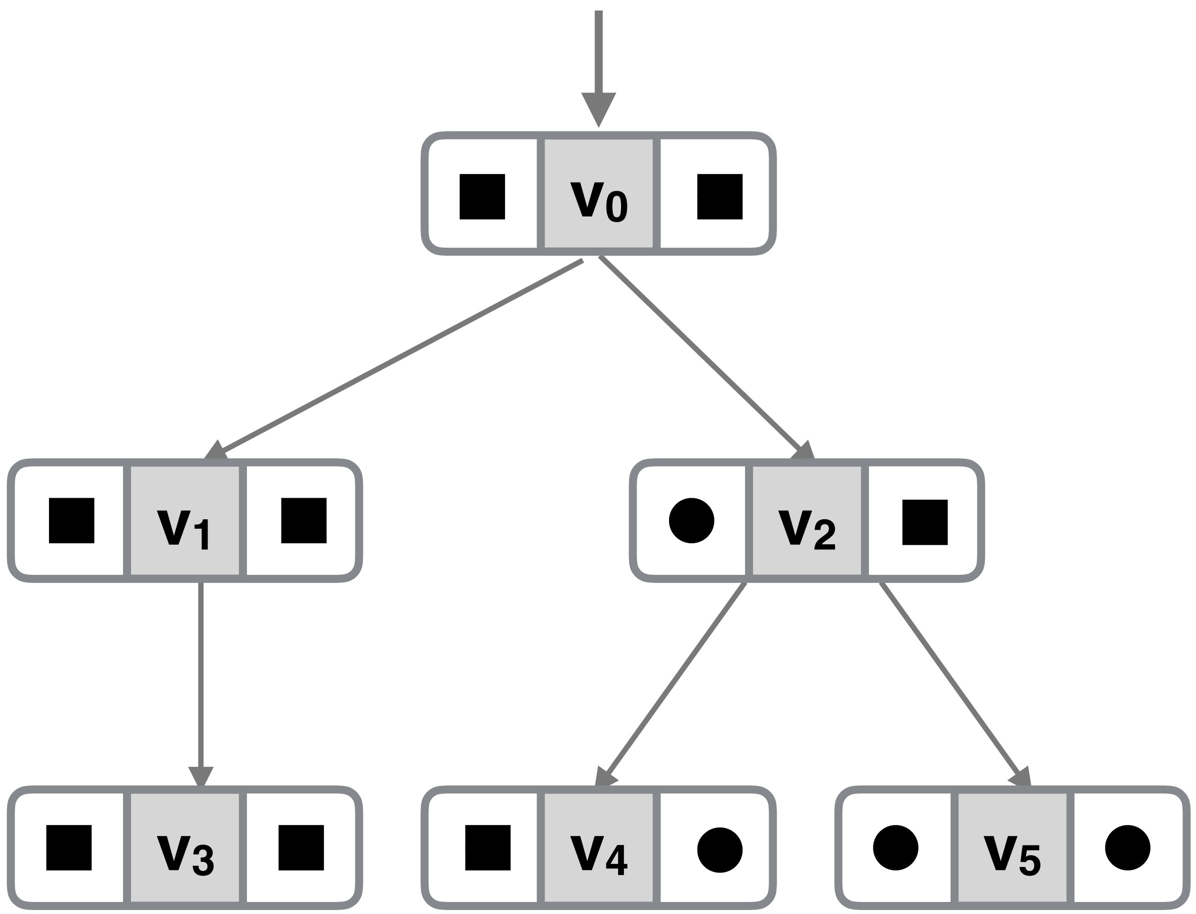

(static) Hierarchical Information

Landscape of Hierarchical Patterns

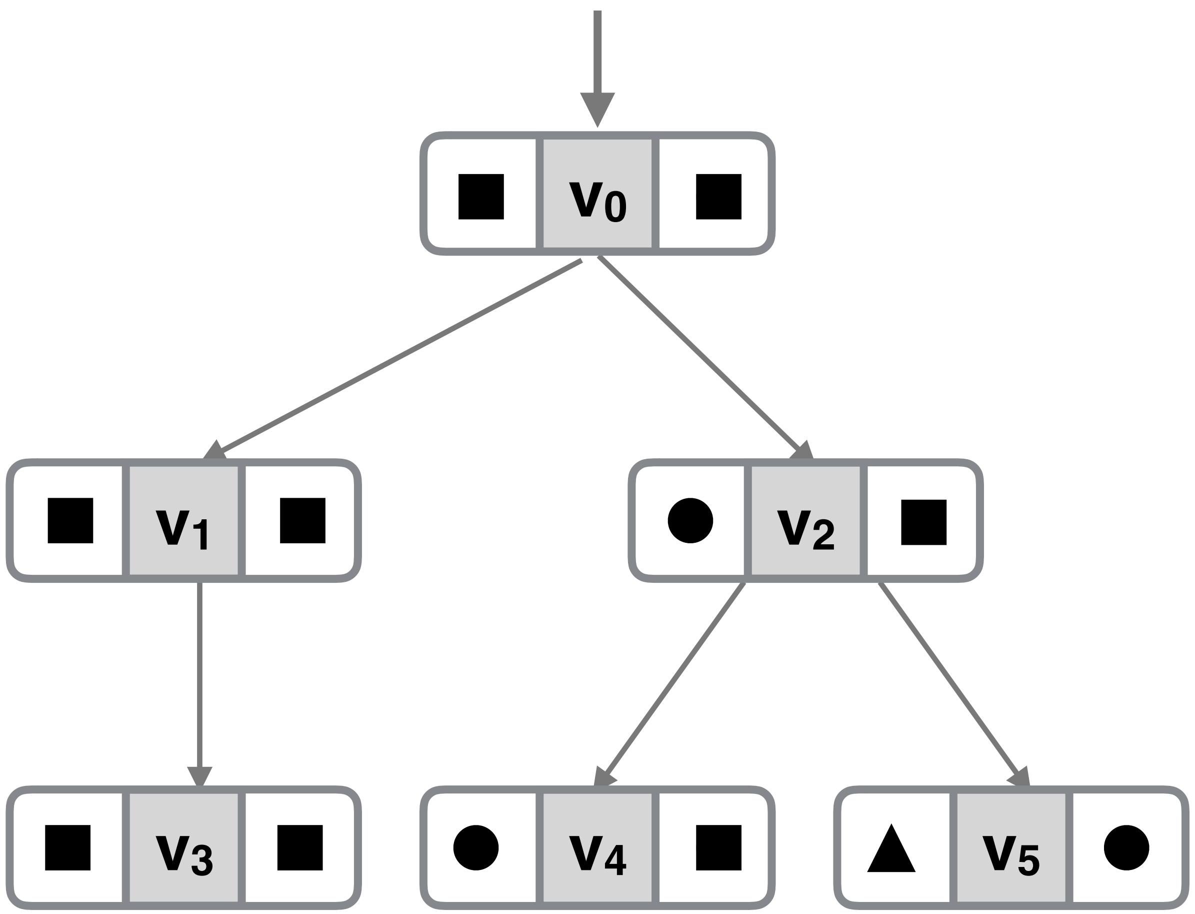

Hierarchical Observation

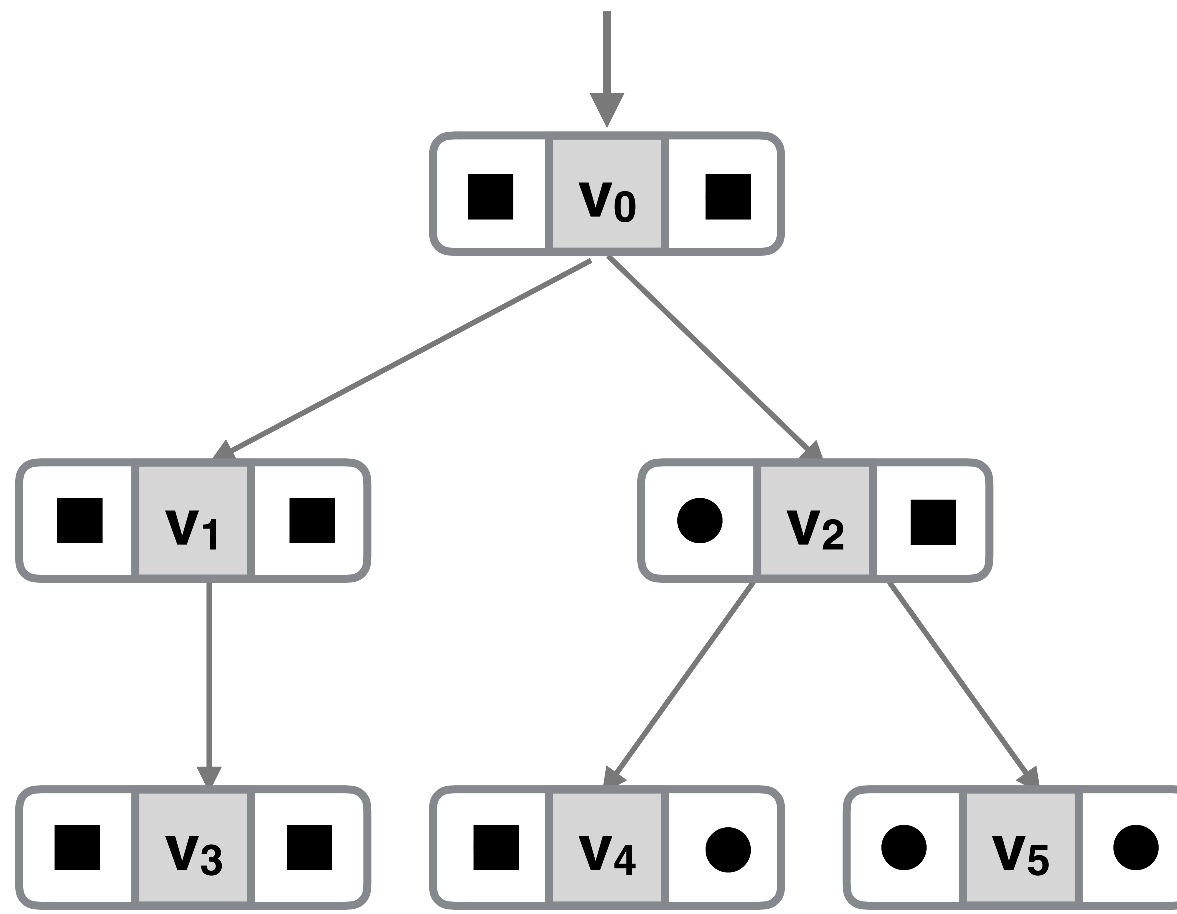

(static) Hierarchical Information

$\rightsquigarrow$ Possible to go *dynamic*?

Dynamic hierarchies

Dynamic hierarchies

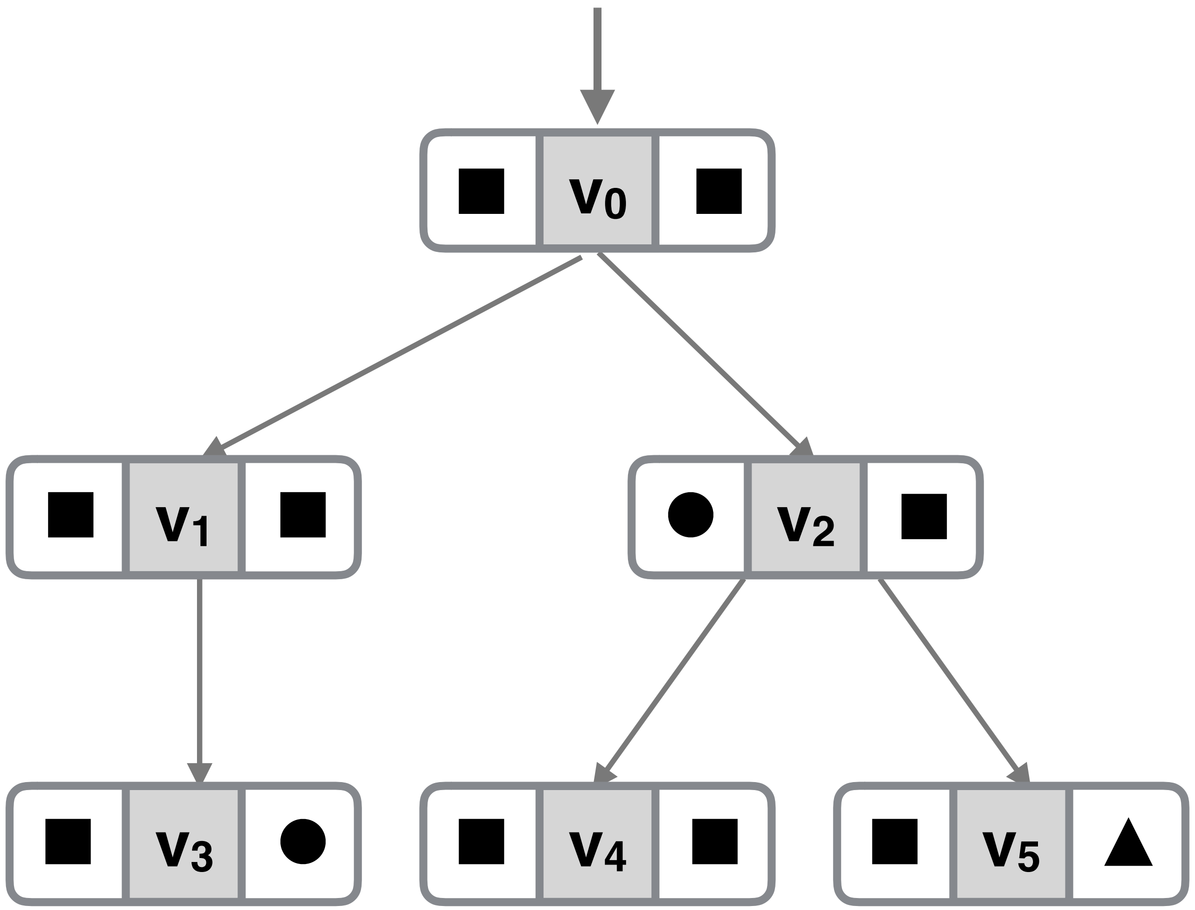

Dynamic hierarchies

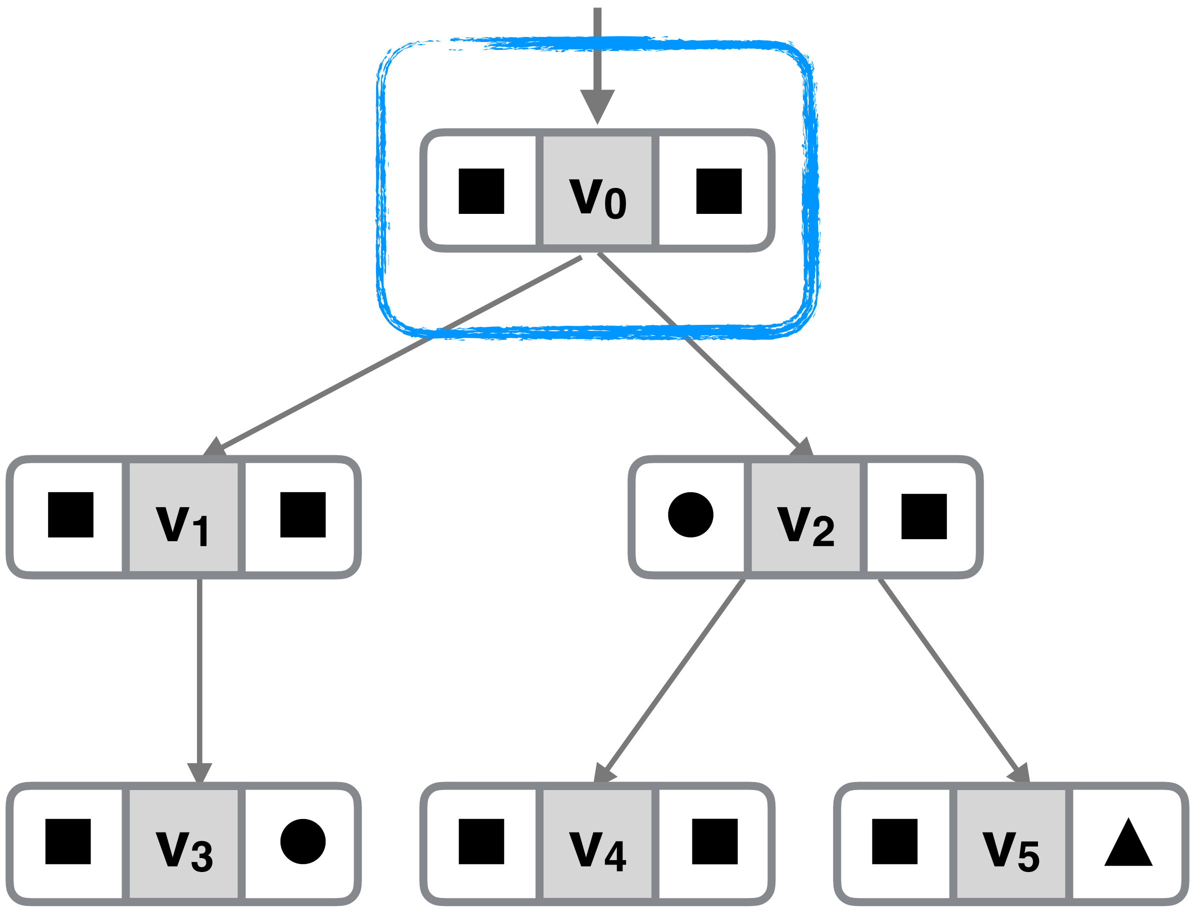

Dynamic hierarchies

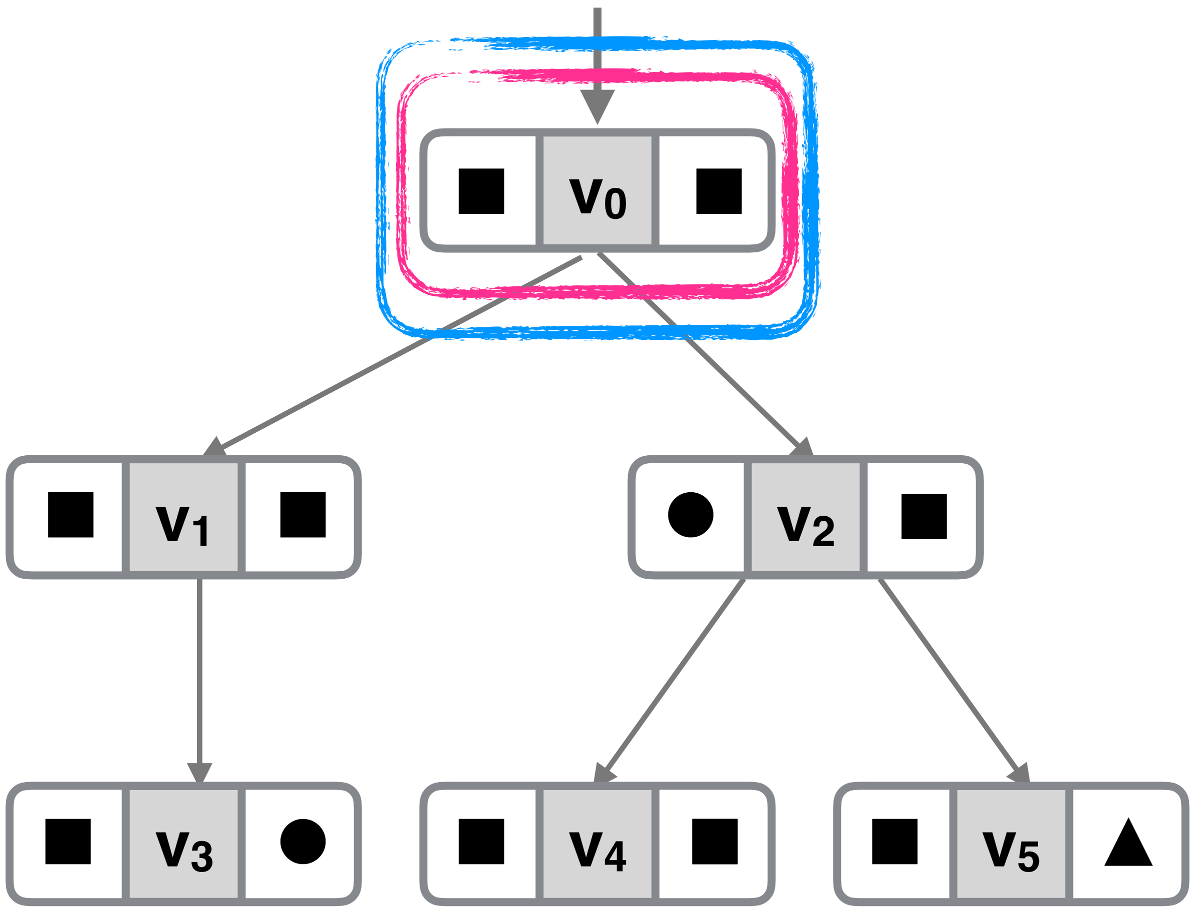

1 $=_{\text{inf}}$ **2**

Dynamic hierarchies

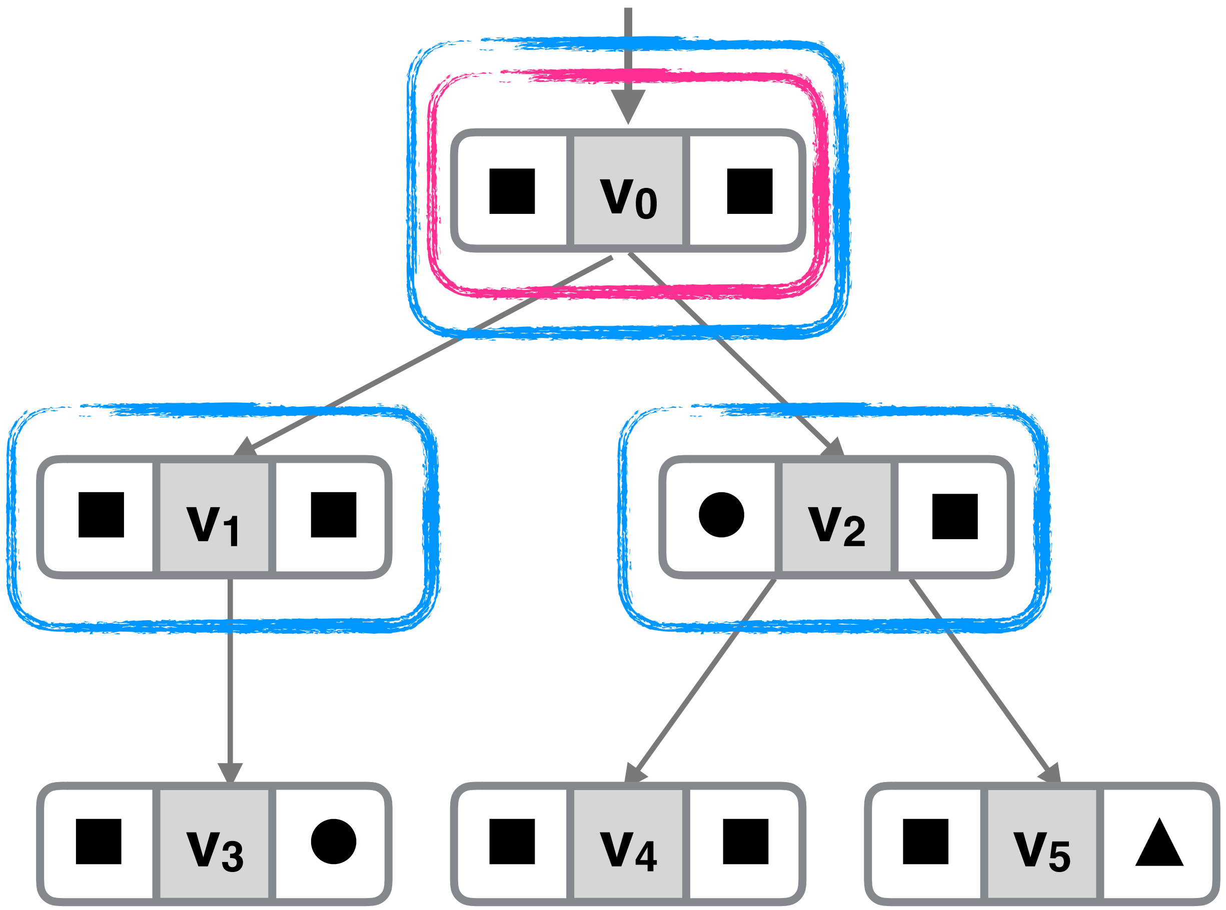

1 $<_{\text{inf}}$ **2**

Dynamic hierarchies

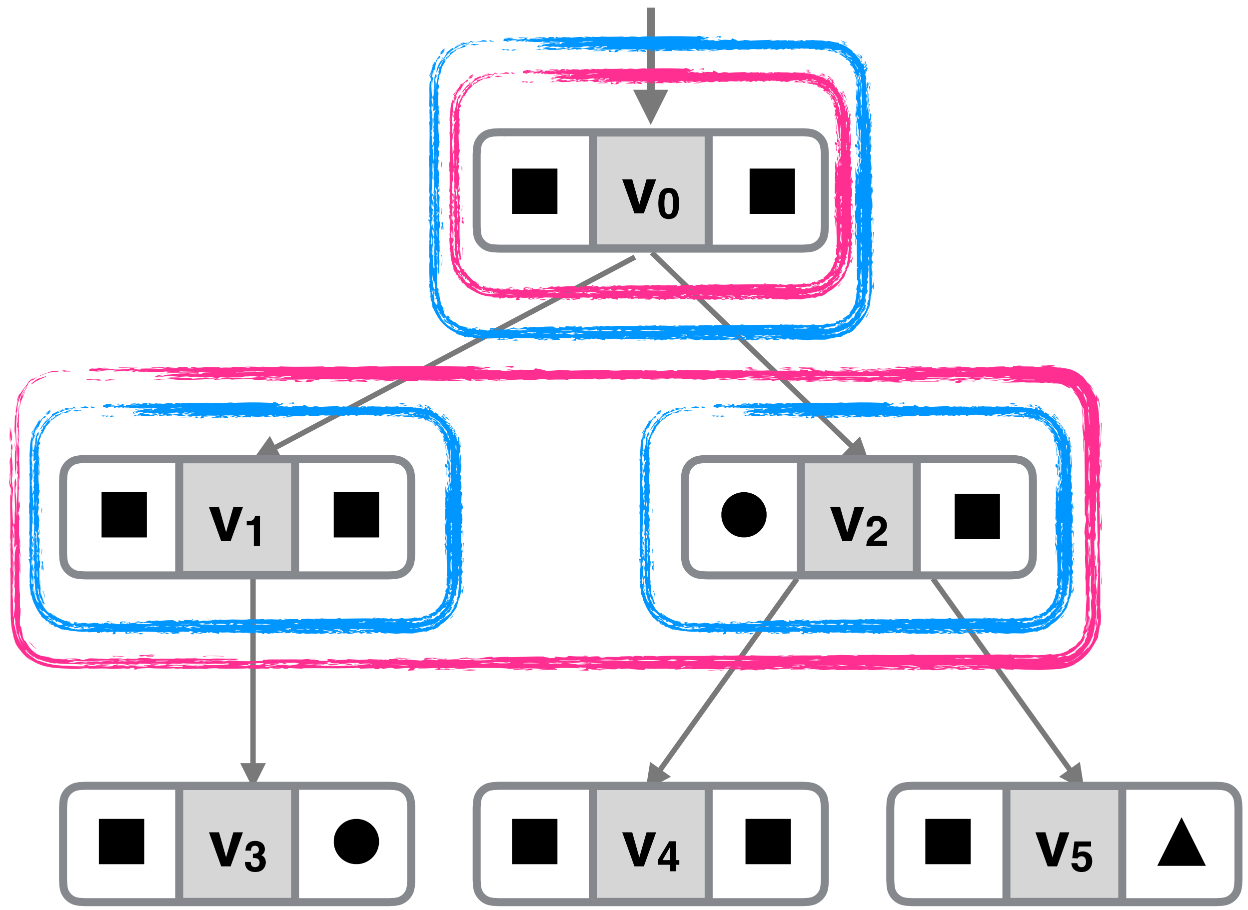

$1 <_{\text{inf}} 2$

$1 <_{\text{inf}} 2$

Dynamic hierarchies

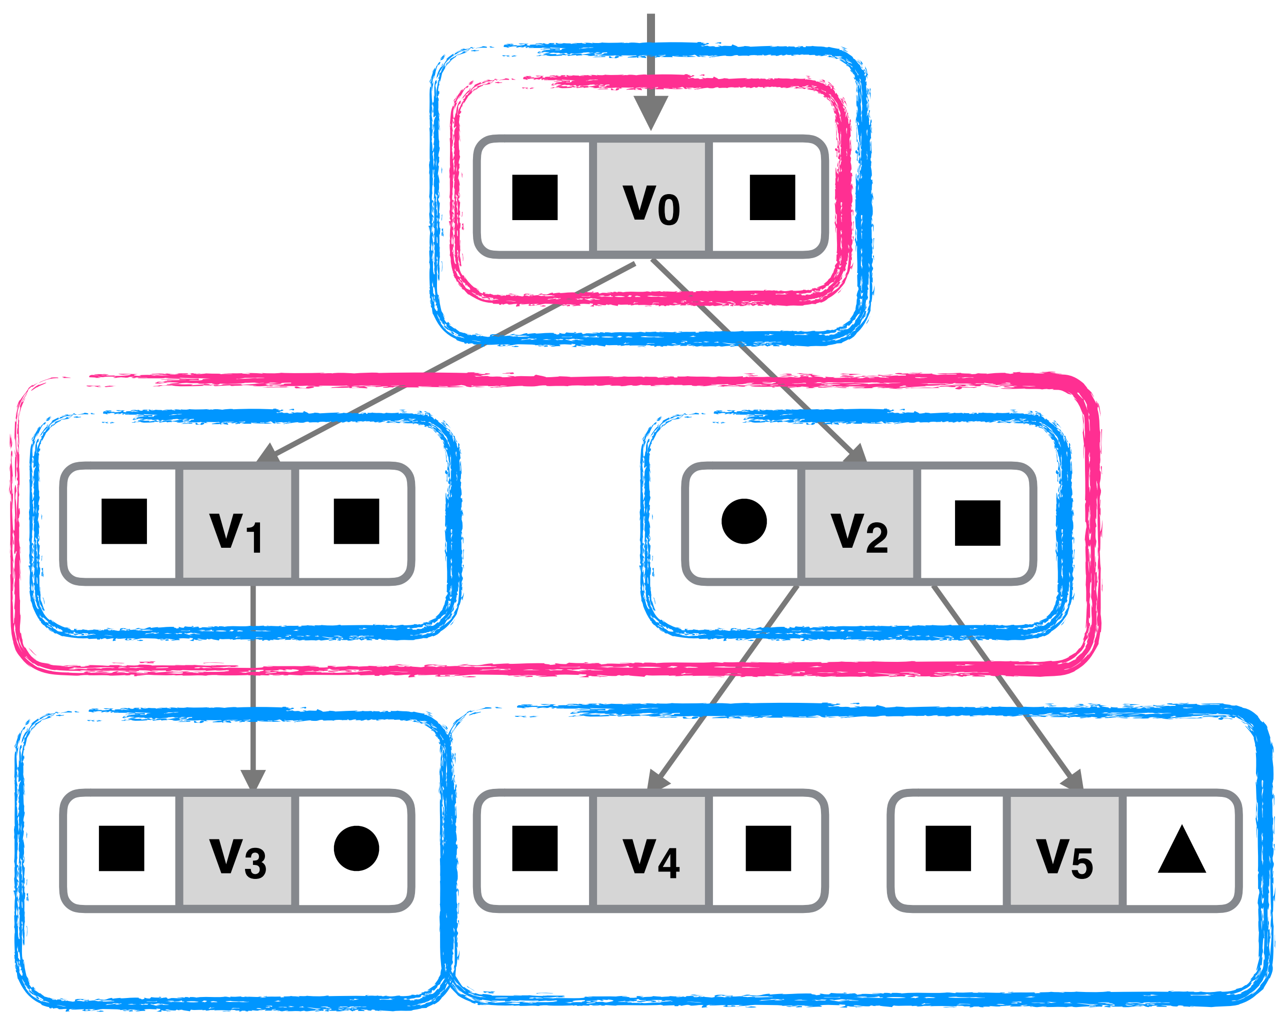

$1 <_{\mathrm{inf}} 2$

$1 <_{\mathrm{inf}} 2$

Dynamic hierarchies

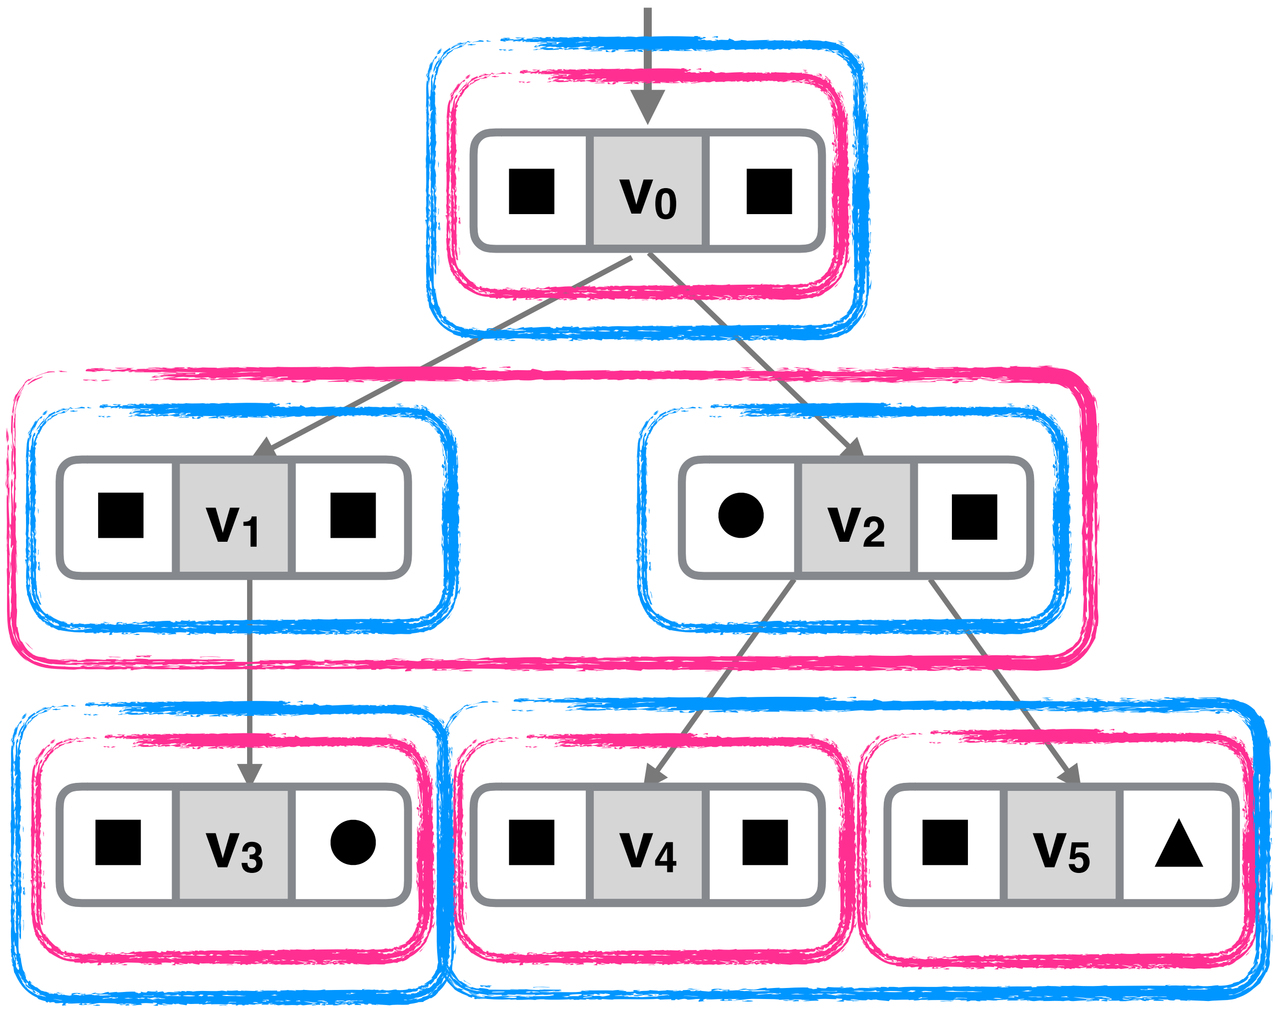

$1 <_{\mathrm{inf}} 2$

$1 <_{\mathrm{inf}} 2$

$2 <_{\mathrm{inf}} 1$

Dynamic hierarchies

Theorem With *dynamic* hierarchical information,
distributed strategy synthesis is decidable.

Dynamic hierarchies

Theorem With *dynamic* hierarchical information,
distributed strategy synthesis is decidable.

Idea: **Shadow game** with static hierarchical information

From G with dynamic HI, game $Sh(G)$ with *n shadow* players $1', \ldots, n'$:

- each i' plays the role of i-most informed
- at every history: shadow i' same information set and same action choice as i

Dynamic hierarchies

Theorem With *dynamic* hierarchical information,
distributed strategy synthesis is decidable.

Idea: **Shadow game** with static hierarchical information

From G with dynamic HI, game $Sh(G)$ with n *shadow* players $1', \ldots, n'$:

- each i' plays the role of i-most informed
- at every history: shadow i' same information set and same action choice as i

Key: Information rank

Idea: Shadow Game

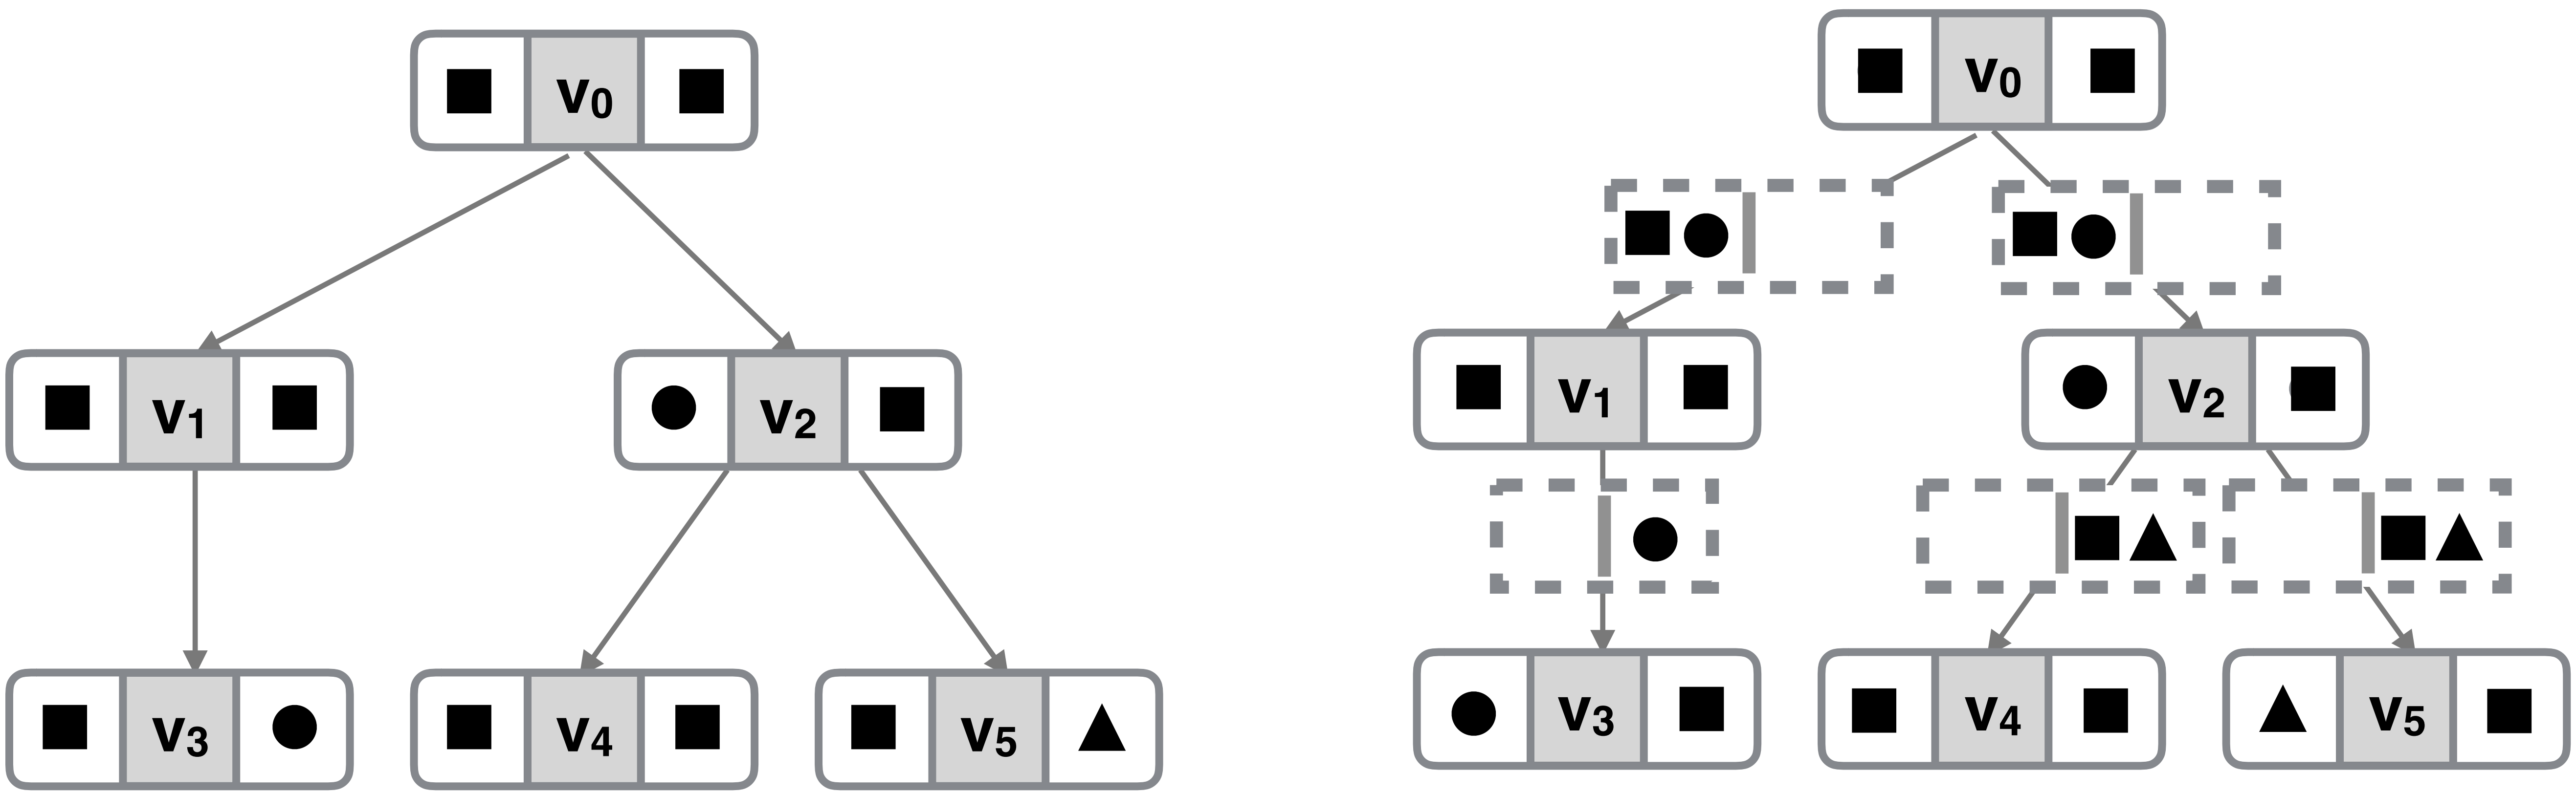

Idea: Shadow Game

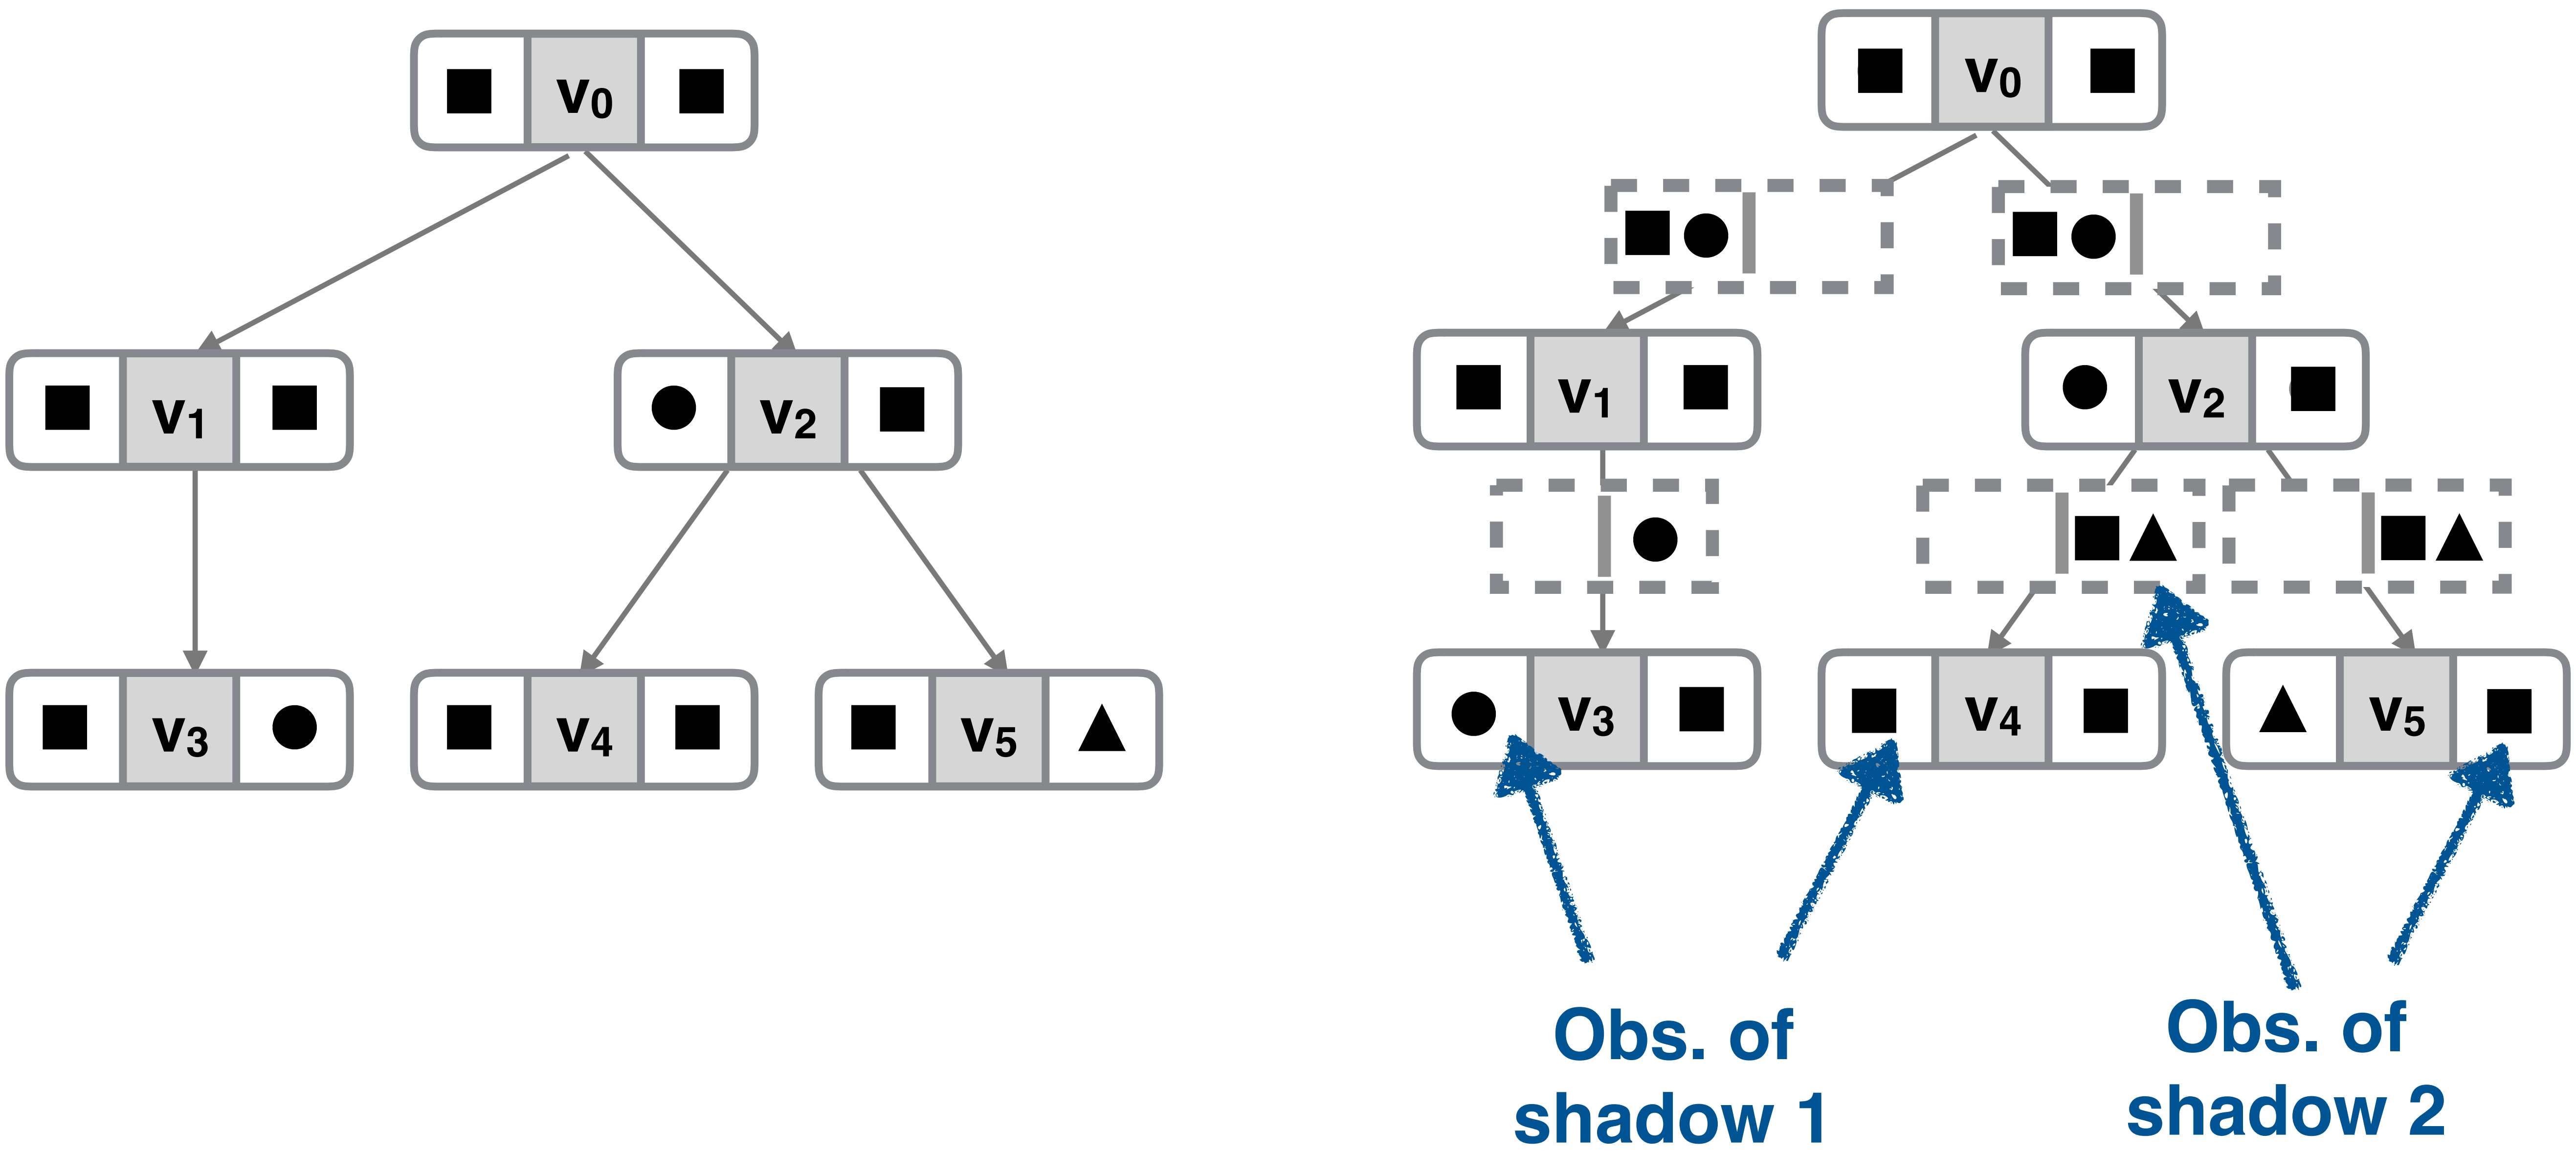

Shadow Game

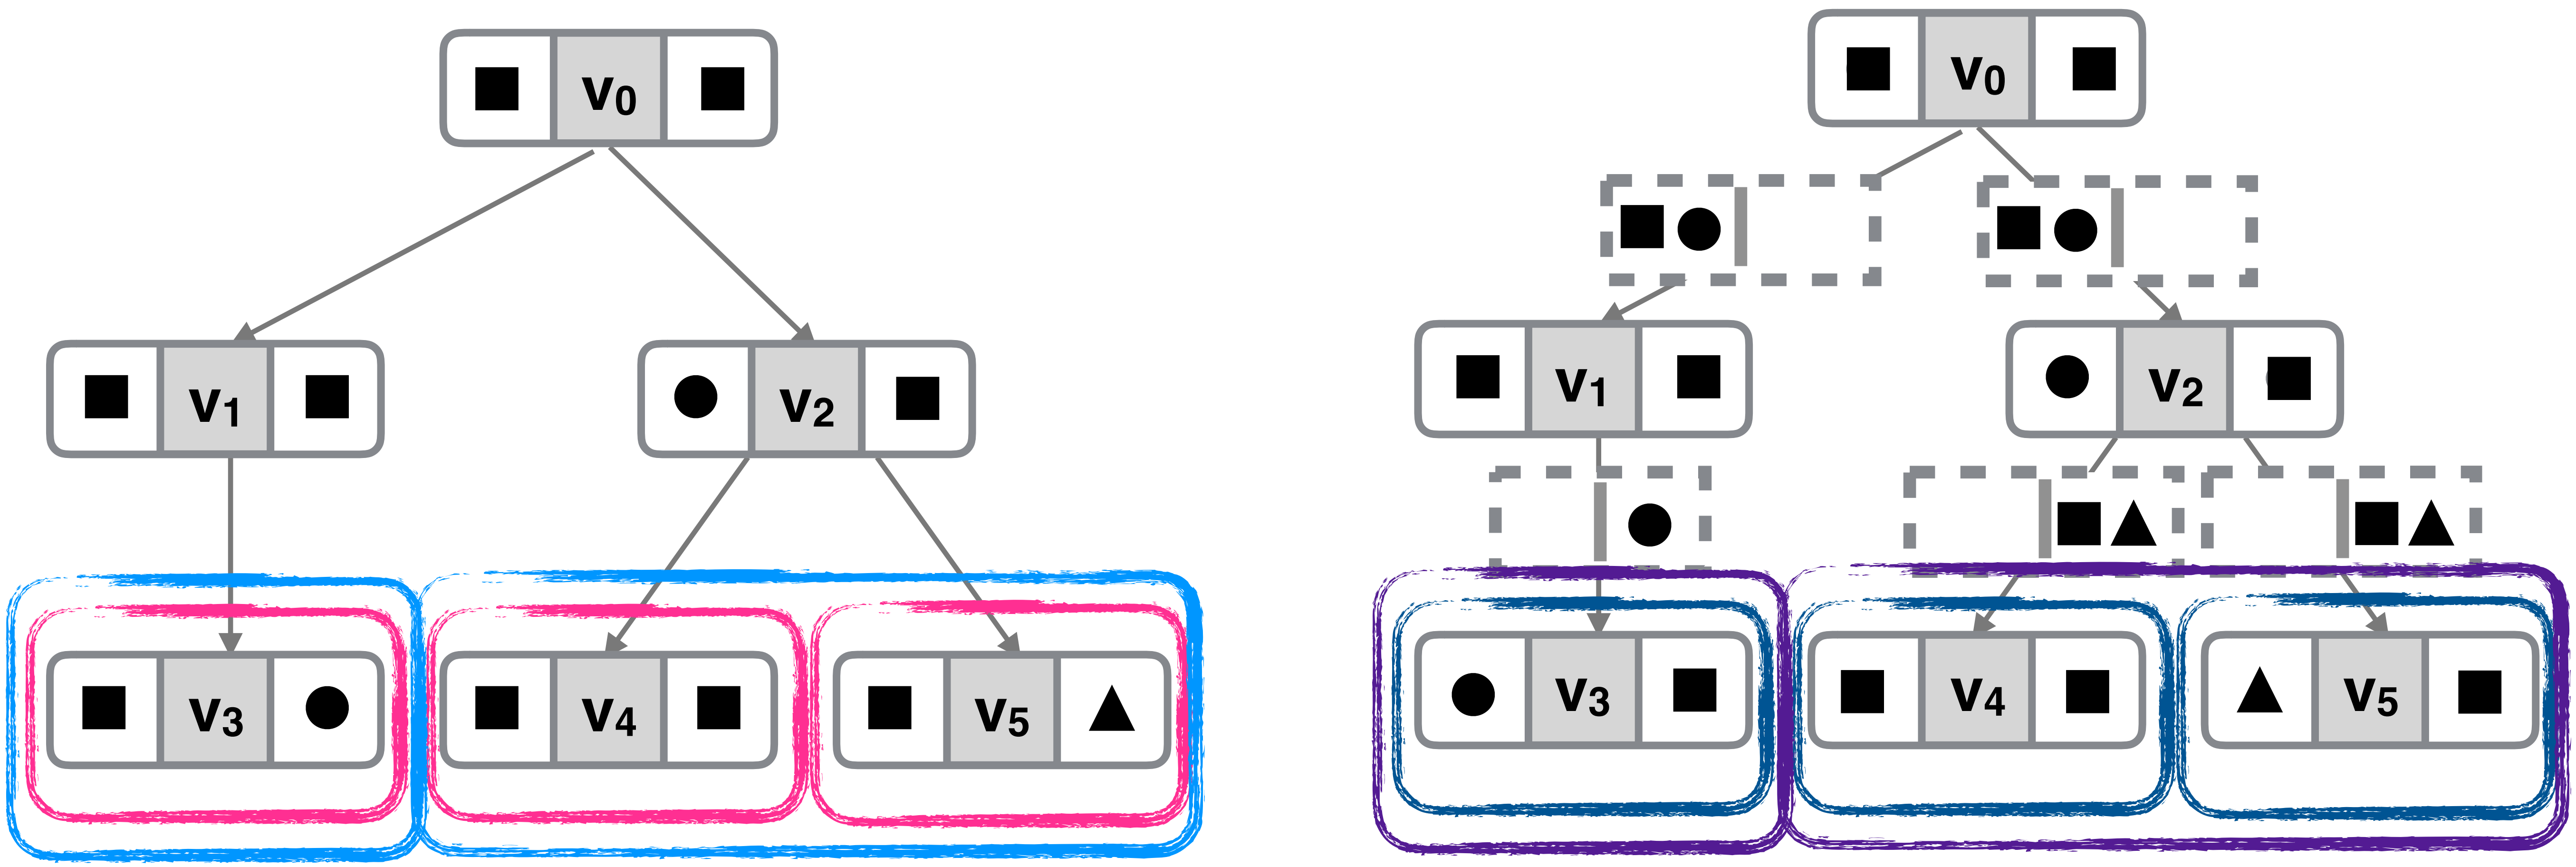

Dynamic hierarchies

Information rank (as a finite-sate signal)

Dynamic hierarchies

Information rank (as a finite-sate signal)

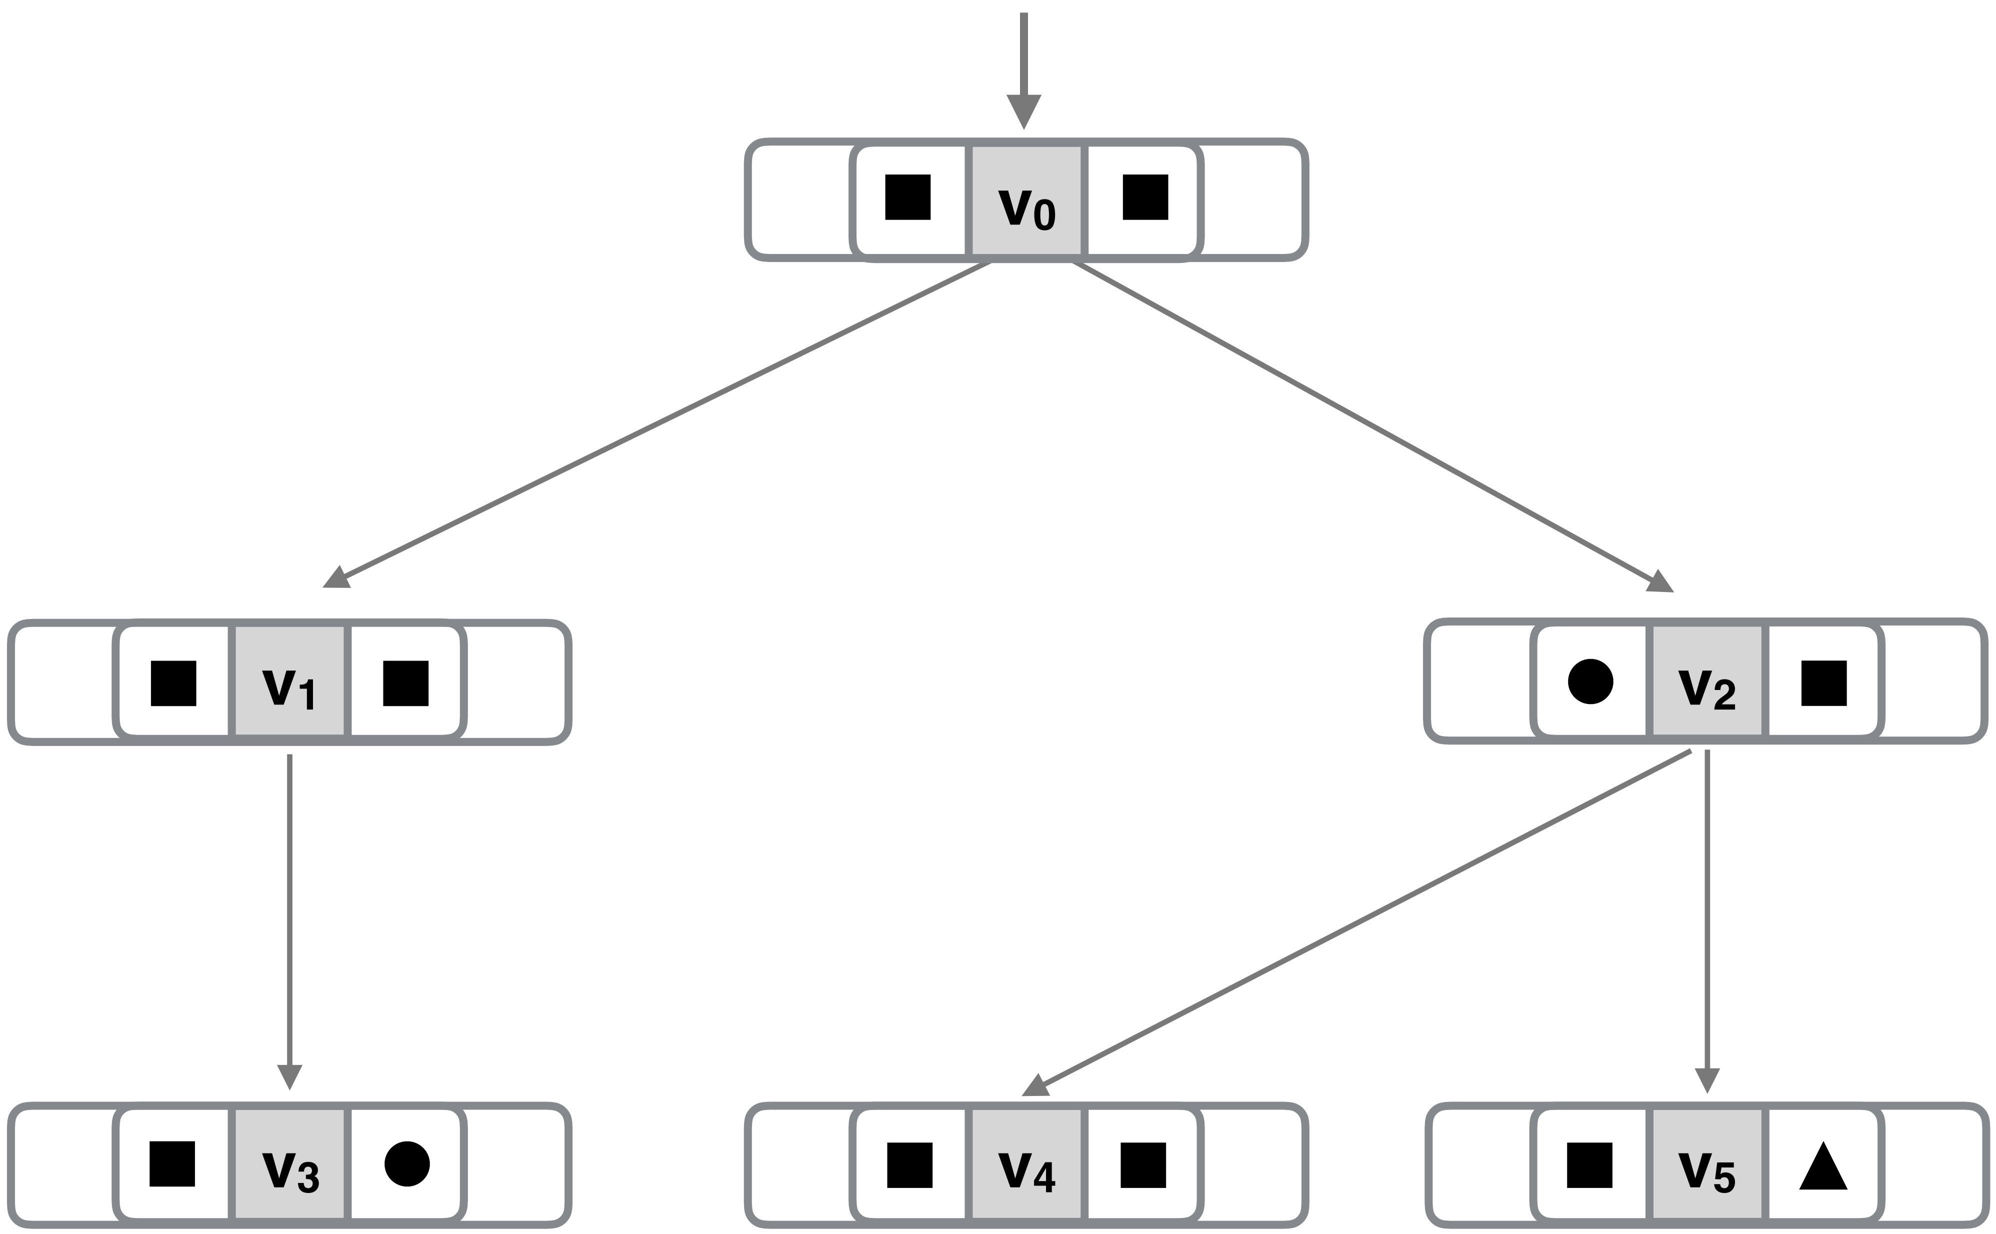

Dynamic hierarchies

Information rank (as a finite-sate signal)

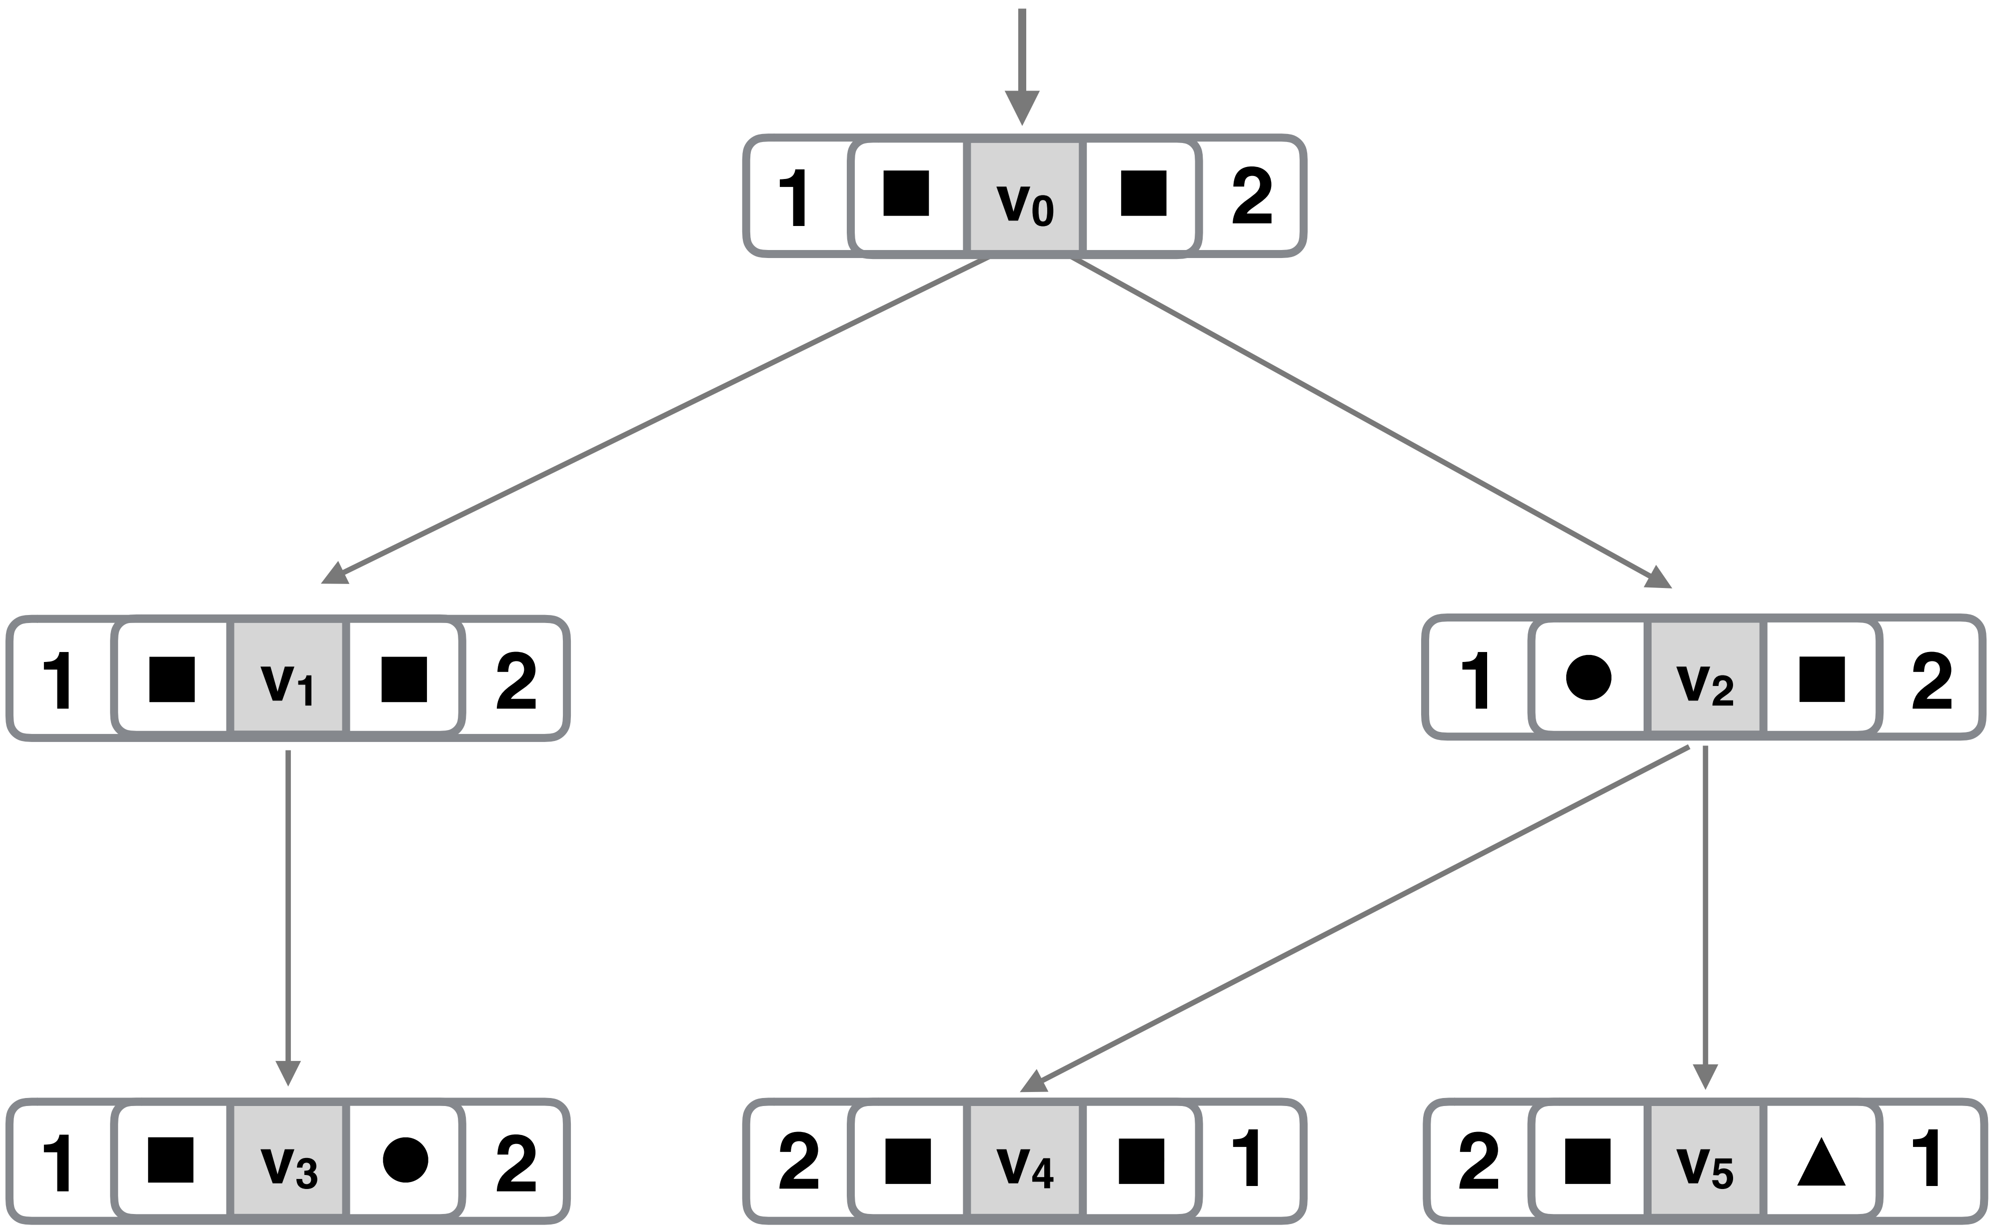

Transient Perturbations

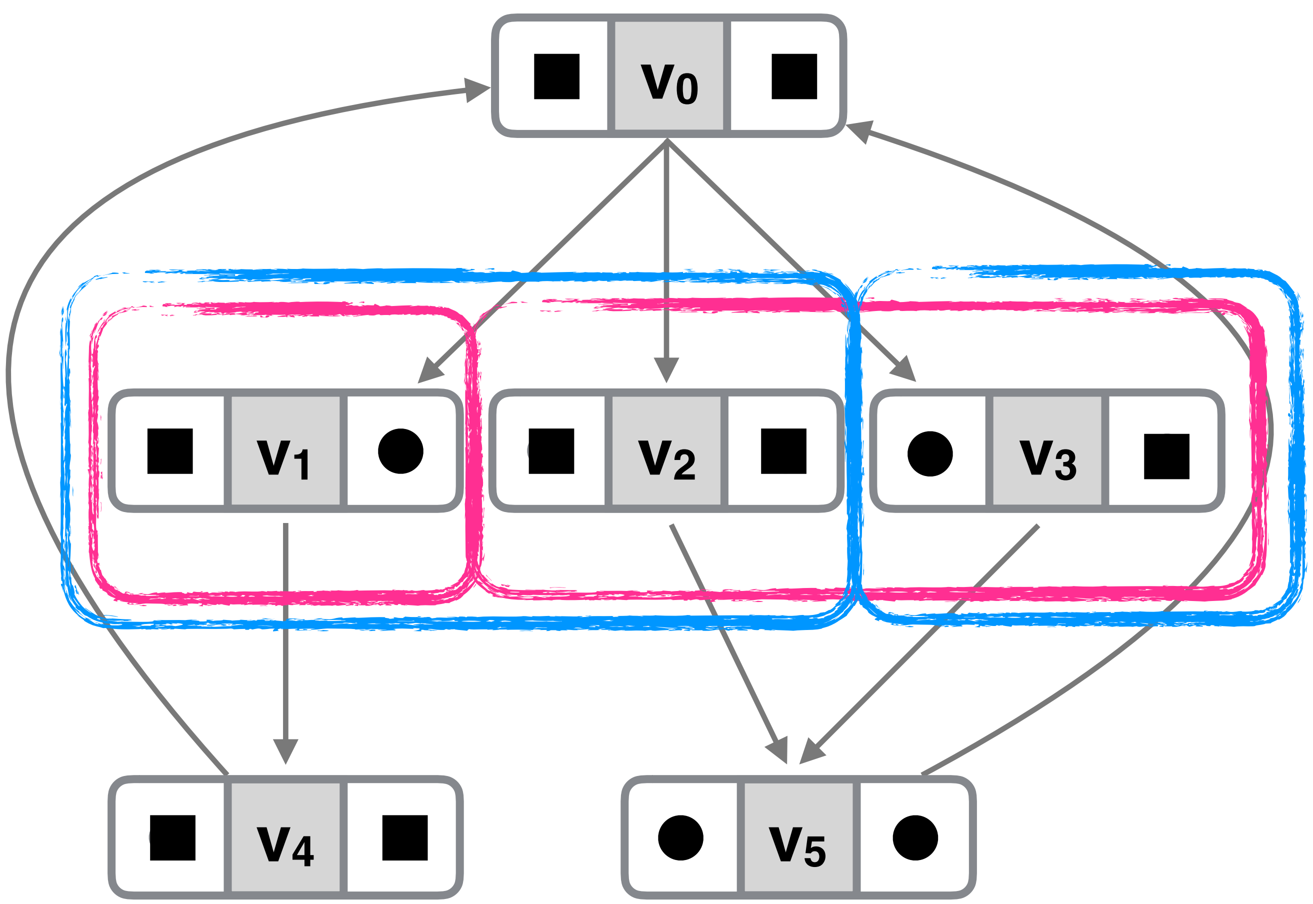

1 $<_{inf}$ **2**

incomparable information

1 $<_{inf}$ **2**

Recurring Hierarchical Information

On every play, infinitely often hierarchical information

Theorem With *recurring* hierarchical information,
distributed strategy synthesis is decidable,
for observable ω-regular winning conditions.

Idea: Information tracking construction
[Berwanger, Kaiser, Puchala '11]

Landscape of Hierarchical Patterns

Hierarchical Observation

$\cap$

(static) Hierarchical Information

$\cap$

(dynamic) Hierarchical Information

$\cap$

Recurring Hierarchical Information

1. Context
2. Model and Background
3. Information-Flow Patterns
   - Propagation of uncertainty
   - One-way Information Flow
   - Delayed Information Flow
4. Conclusion

Delayed Information Flow

Some uncertainty sources are difficult to embed at the design level

Example: can we cope with a flawed information delivery mechanism?

$\rightsquigarrow$ Distributed games with delayed monitoring

Instant monitoring:

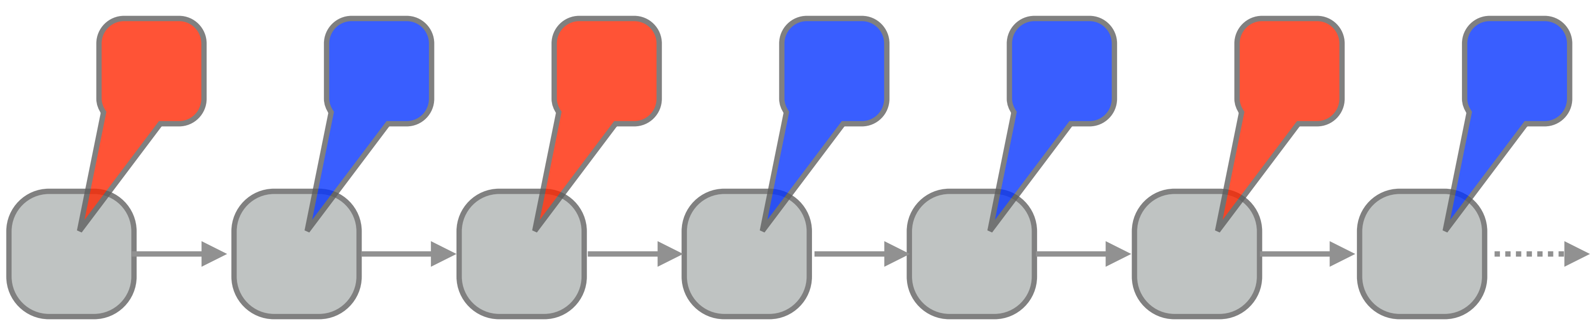

Delayed monitoring:

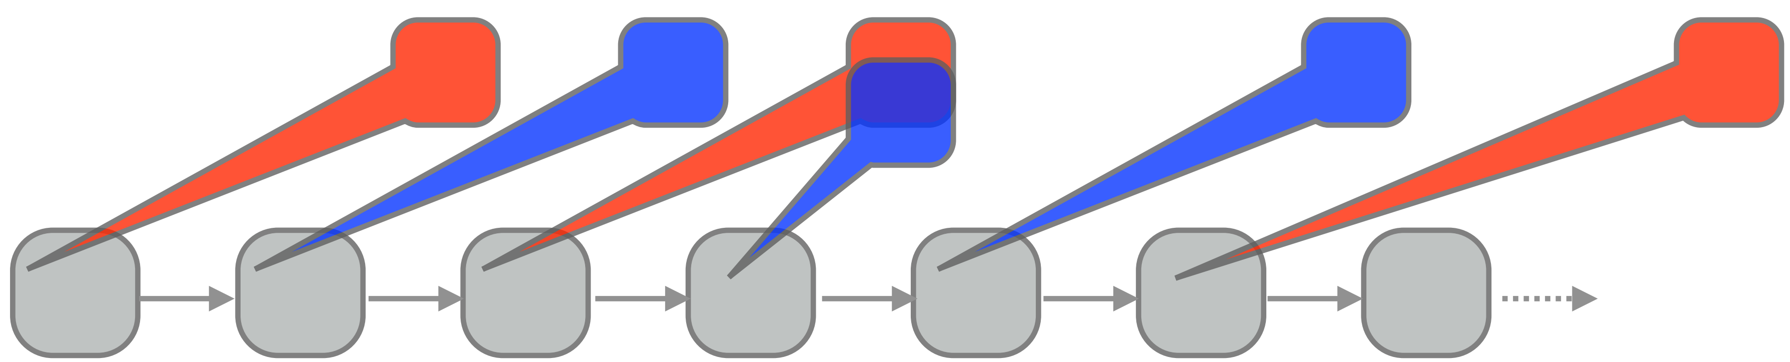

Delayed Information Flow

For games s.t.: perfect information about the states,
particular winning conditions,
solvable

Theorem Winning strategies for the instant-monitoring game can be adapted to the delayed-monitoring variant.

Key: delayed-response strategies for repeated games [Fudenberg, Ishii, Kominers '14]

Looking at Information-Flow Patterns

Insights on interaction mechanisms:

- cooperation is intrinsically hard
  consensus enough to cause undecidability
- shape information flow to ensure decidability
  hierarchical information patterns
- method to cope with realistic communication failures
  delayed response strategies

What's next?

- Consensus game acceptors
  beyond one decision
  refine further execution complexity classification

What's next?

- Consensus game acceptors

  beyond one decision

  refine further execution complexity classification

- Hierarchical Patterns

  remove restriction of observable winning condition
  (Recurring Hierarchical Information case)

  monitored architectures

What's next?

- Consensus game acceptors

  beyond one decision

  refine further execution complexity classification

- Hierarchical Patterns

  remove restriction of observable winning condition
  (Recurring Hierarchical Information case)

  monitored architectures

- Delayed Monitoring

  relax perfect information about states assumption

  extend class of winning conditions

What's next?

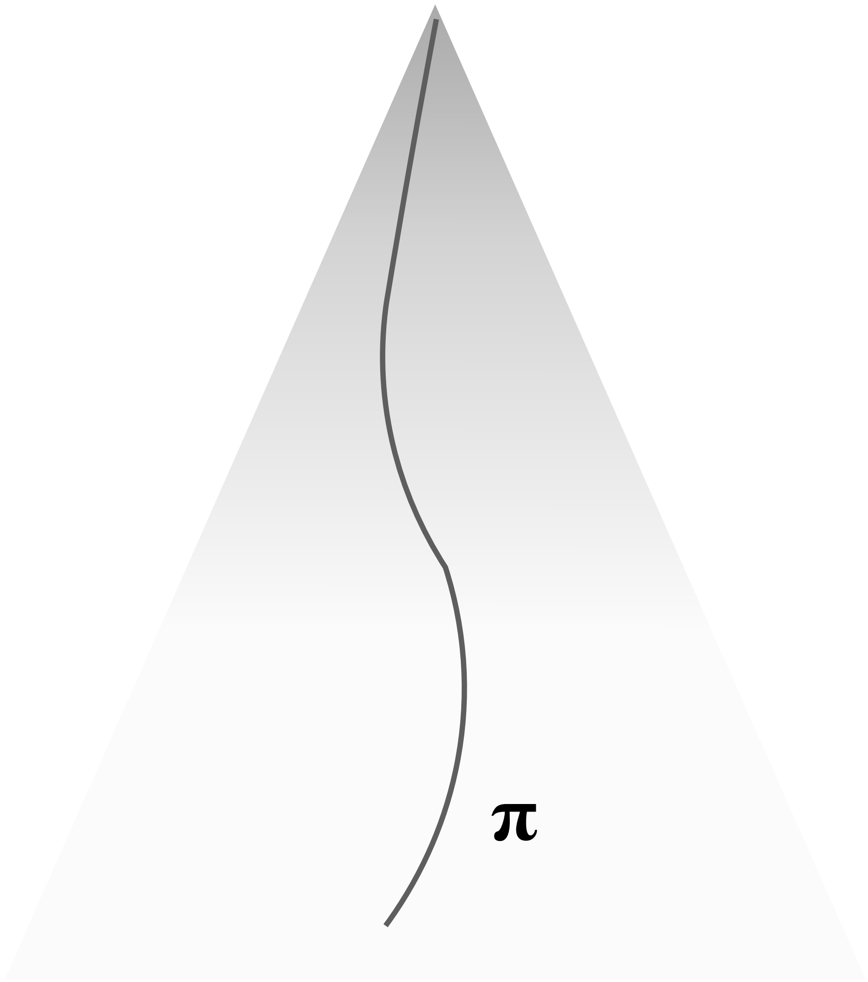

Multi-level synthesis:

- start with arbitrary game

What's next?

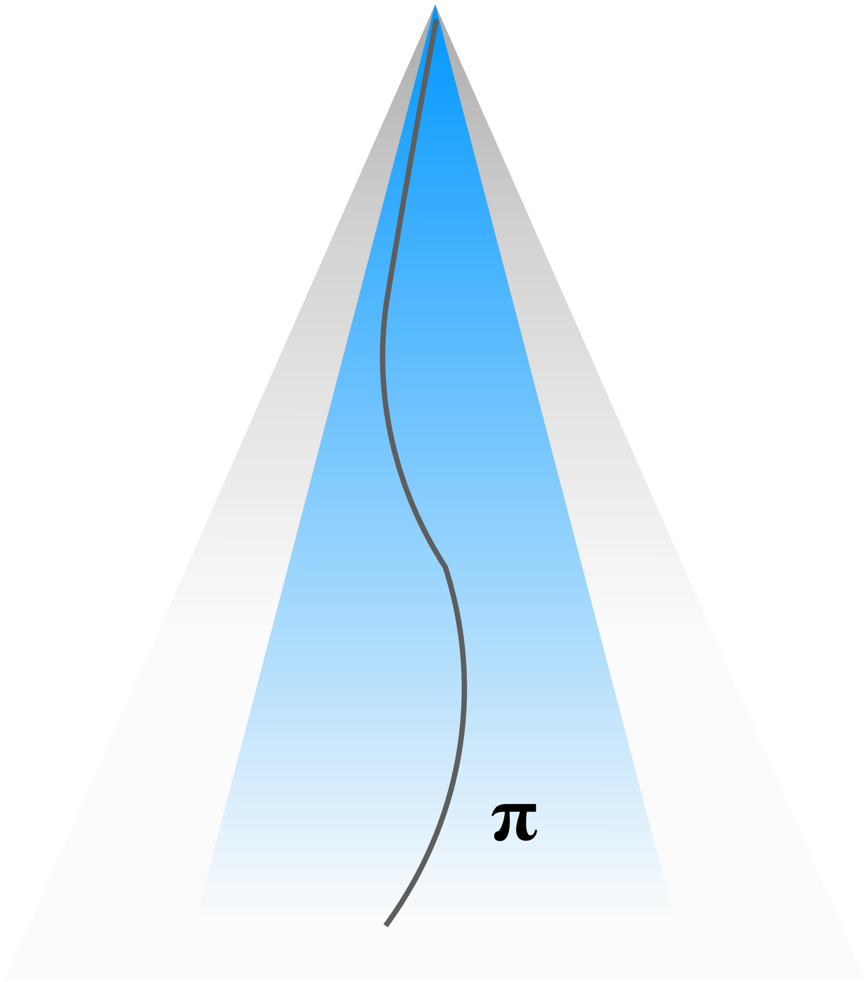

Multi-level synthesis:

- start with arbitrary game
- avoid propagation of uncertainty:

find strategies to enforce pattern ensuring decidability

What's next?

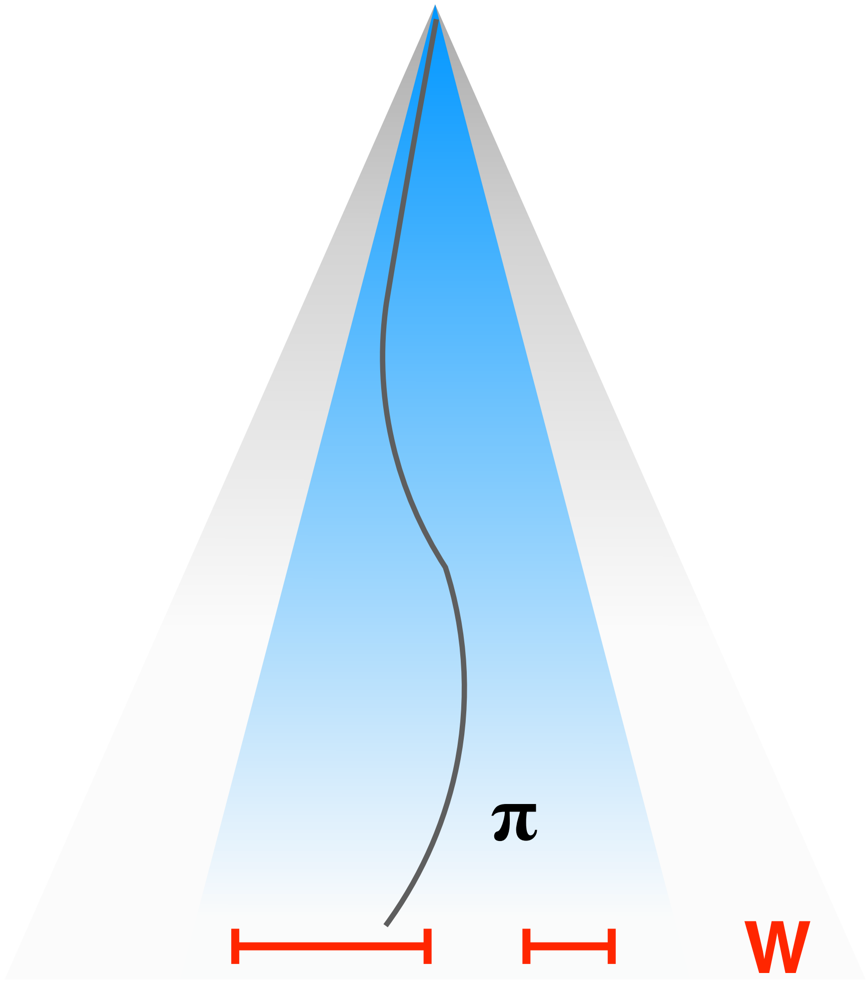

Multi-level synthesis:

- start with arbitrary game
- avoid propagation of uncertainty:

  find strategies to enforce pattern ensuring decidability

- refine from there:

  winning condition
  actions and observations