Information Tracking in Distributed Games

Dietmar Berwanger & Łukasz Kaiser

LSV, CNRS, ENS Cachan & RWTH Aachen

ACTS II, Chennai 2010

No-go Theorem

Peterson and Reif (1979), Multiple-Person Alternation:

Infinite games
with imperfect information
of two or more players
are unsolvable.

Games with imperfect information - structure

- n players $i \in \{0, \dots, n-1\}$, with finite sets $\begin{cases} A^i \text{ actions} \\ B^i \text{ observations} \end{cases}$
- vs Environment

Simultaneous actions $A = \times_{i<n} A^i$

Game graph $G = \big(V, \Delta, (\beta^i)_{i<n} \big)$

- V set of positions
- $\Delta \subseteq V \times A \times V$ transition relation
- $\beta^i : V \longrightarrow B^i$ observation function for player i

Games with imperfect information - playing

Play: path $\pi = v_0, a_0, v_1, a_1, \dots$ from initial position v_0

- at v, players choose an action profile $a = (a^i)_{i<n}$ simultaneously
- Environment picks successor $w \stackrel{a}{\longleftarrow} v$; each player i observes $\beta^i(w)$

▶ observation history $\beta^i(\pi) = \beta^i(v_0), \beta^i(v_1), \dots$

Strategy: $s^i : (VA)^*V \longrightarrow A^i$ such that

$$\beta^i(\pi) = \beta^i(\pi') \implies s^i(\pi) = s^i(\pi')$$

- play π follows s^i $\quad a^i_{\ell+1} = s^i(v_0, a_0, \dots, v_\ell)$

- outcome of a profile $s = (s^i)_{i<n}$ $\quad$ set of all plays that follow each s^i

Games with imperfect information - winning

Winning condition $W \subseteq V^\omega$

... mostly described via finite coloring $\gamma : V \to C$ by regular $W \subseteq C^\omega$

Question: Given game graph (G, γ), winning condition $W \in C^\omega$,
construct/decide existence of
winning strategy profile, with outcome$(s) \subseteq W$.

Distributed: each player aware of strategy of other players

- only not of what they observe

Contrast to individual rationality - coordination

Distributed strategies with finite memory

Strategies s^i implemented by a finite automaton (M, q_0, μ, ν)

- reads observation $b^i \in B^i$
- updates internal state $q' = \mu(q, b^i)$
- outputs action $a^i = \nu(q, b^i)$

Issue #1: Representation

- factual state information explicit: state attributes, available actions
  - future and past only up to bisimulation
  - sufficient under perfect information, zero-sum, regular conditions
    - require little knowledge of history
    - knowledge about other player is irrelevant
- epistemic state information implicit:
  - first-order knowledge space - observation history - already infinite
  - higher-order knowledge may matter as well

How to represent all this information explicitely.

Extensive form of a graph game

Game graph $G = (V, \Delta, \beta^i, \gamma)$ $\rightsquigarrow$ extensive form $\left(T, \hat{\Delta}, \sim^i, \hat{\gamma}\right)$

- $(T, \hat{\Delta}, \hat{\gamma})$ tree unravelling of (V, Δ, γ);
- $\sim^i \in T \times T$ indistinguishability relation:

$$\pi \sim^i \pi' \quad \text{if} \quad \text{same actions and observations for player } i$$

Strategy $s^i : (VA)^* V \to A^i$ with $\pi \sim^i \pi' \implies s^i(\pi) = s^i(\pi')$

- Alternativley, $\mathcal{J}^i$:= partition of information sets induced by $\sim^i$;

$$s^i : \mathcal{J}^i \to A^i - \text{memoryless in the information set}$$

Proposition. Every game is equivalent to its extensive form.

Information vs relevant knowledge

- Information set - maximal support for action
  - generated by observation histories
  - infinite

How to represent all relevant information explicitely.
possibly in a finite way?

- Relevant knowledge - necessary support
  - refinement of observation equivalence depending on history
  - coarsening of information equivalence depending on future

.

Infinite knowledge hierarchies

Two players, 0 and 1,

- each observes a number $n^i > 0$ such that $n^1 = n^0$ or $n^1 = n^0 + 1$
- actions: declare a number m^0, m^1
- outcome winning if $m^1 - m^0 = n^1 - n^0$ and $m^1 = 1$ for $n^1 = 1$

Infinite knowledge hierarchies

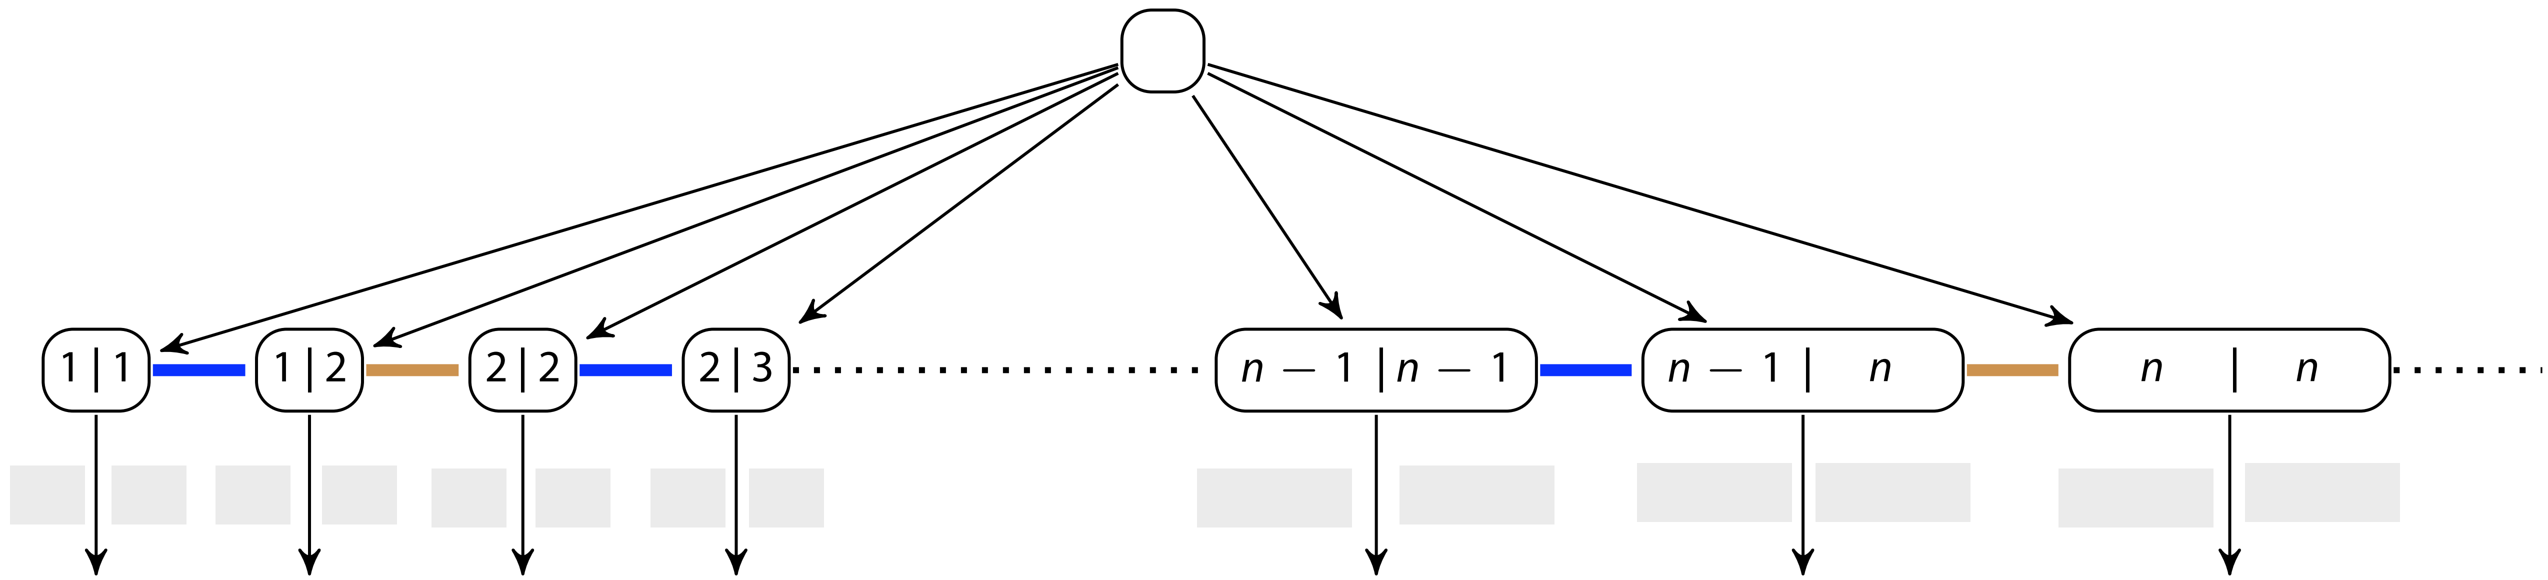

Infinite knowledge hierarchies

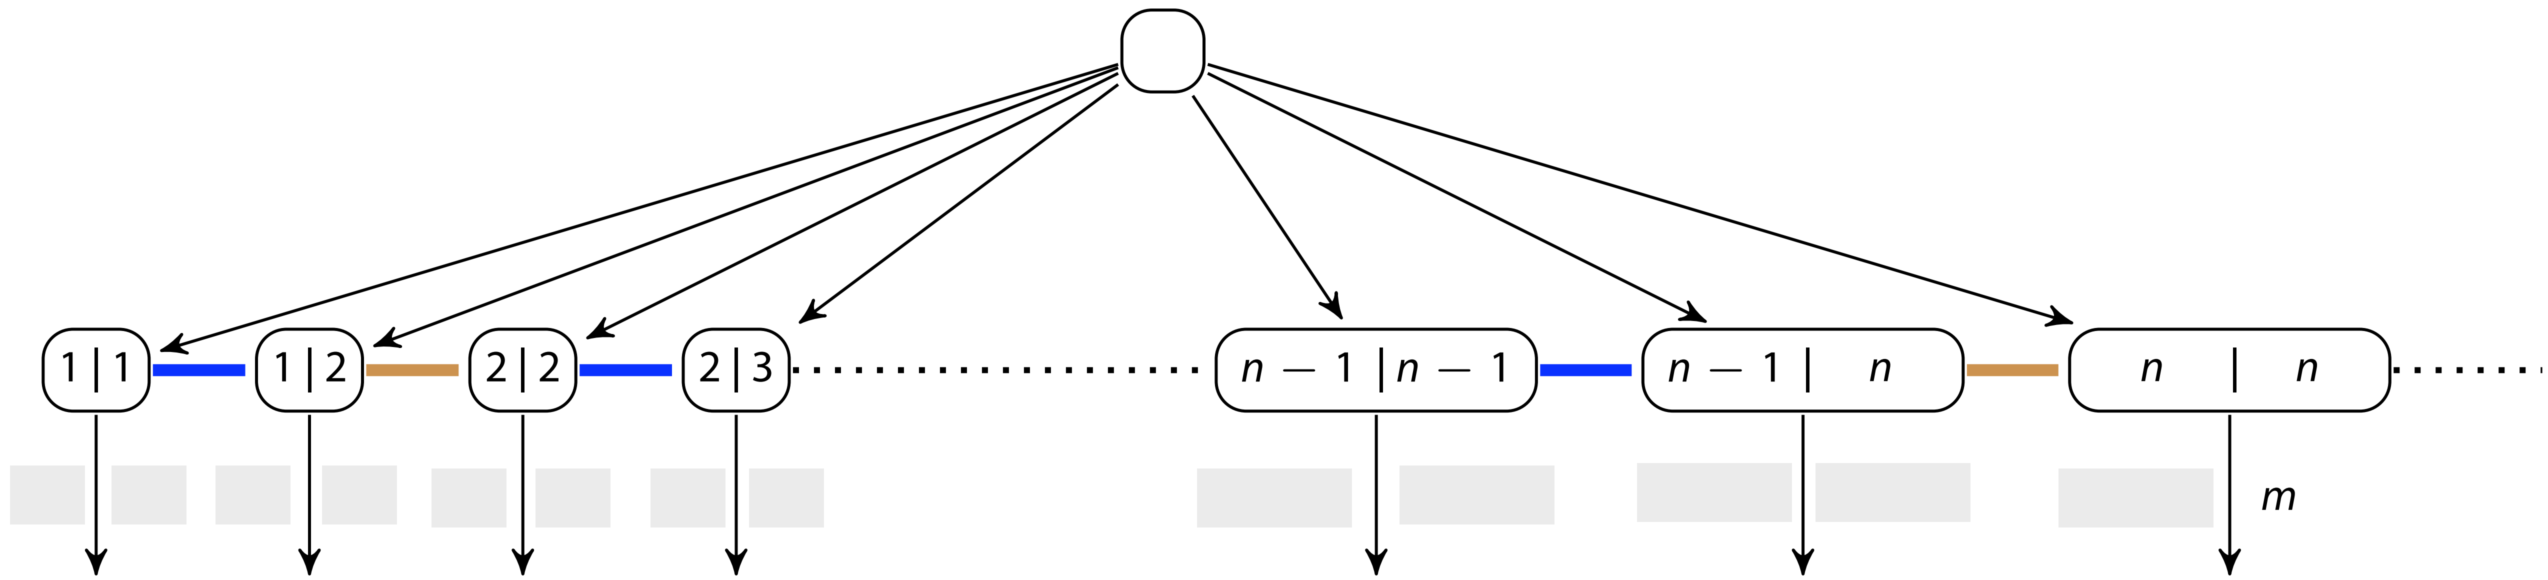

Infinite knowledge hierarchies

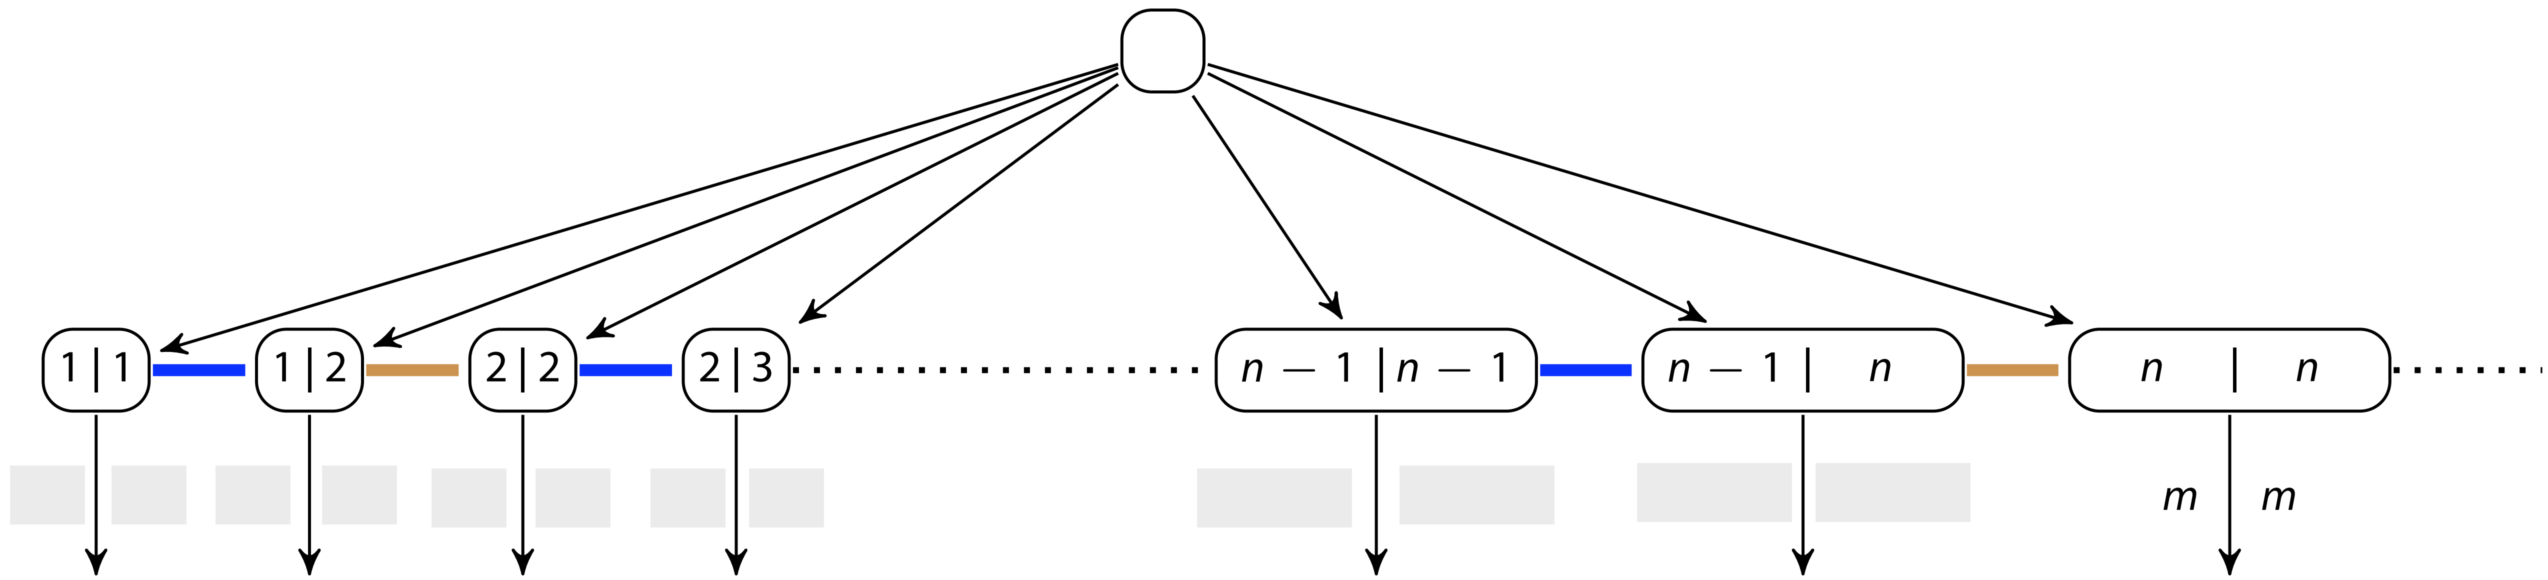

Infinite knowledge hierarchies

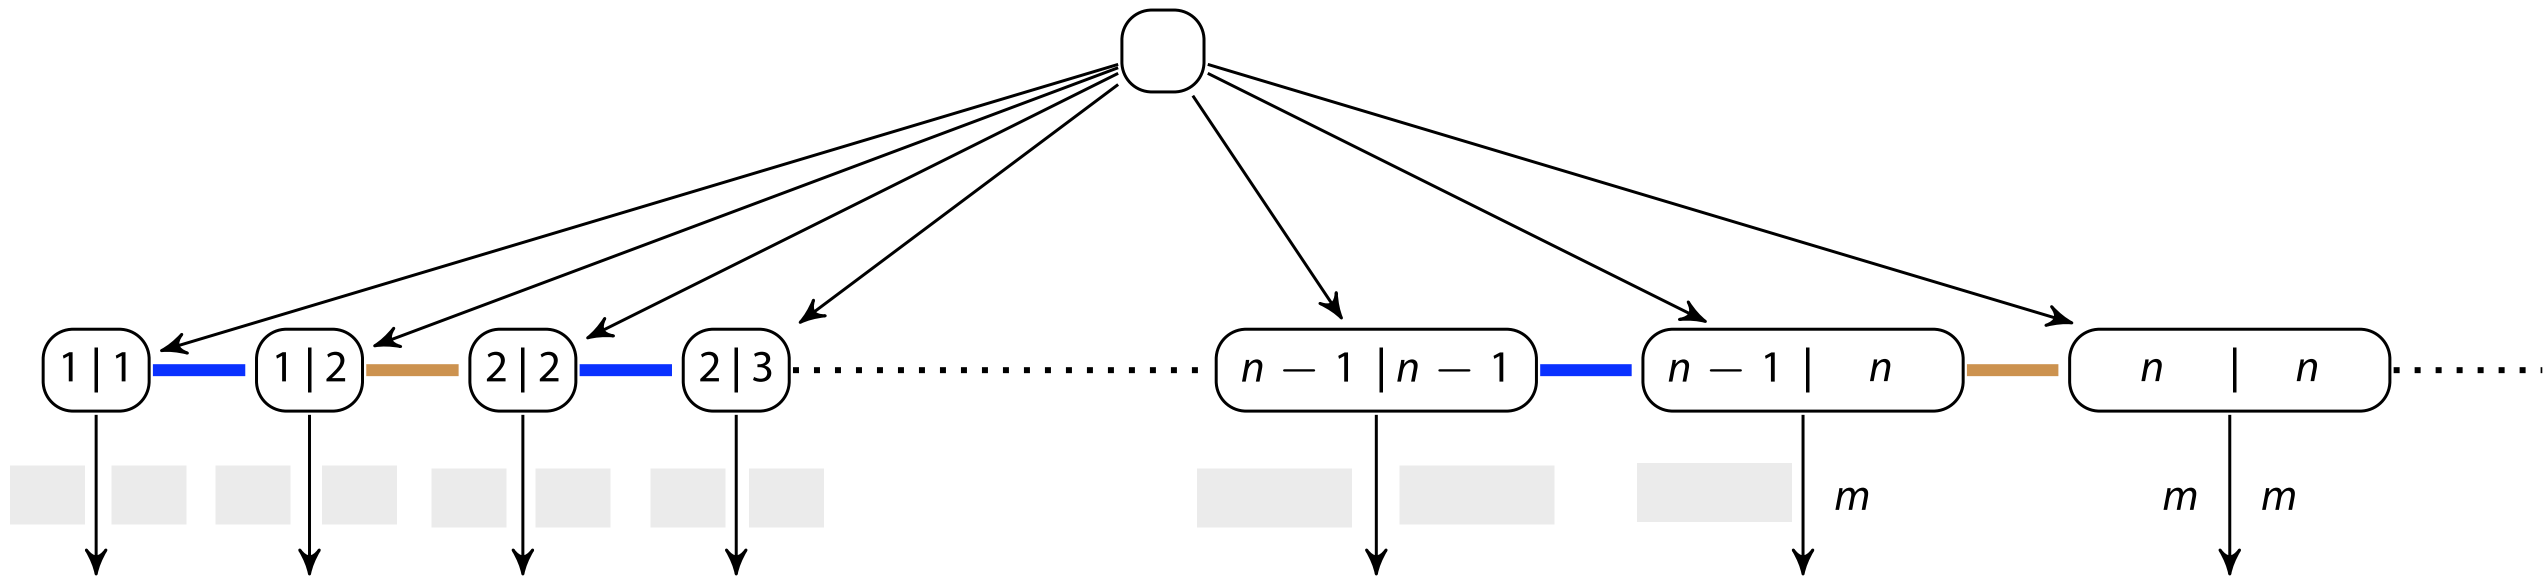

Infinite knowledge hierarchies

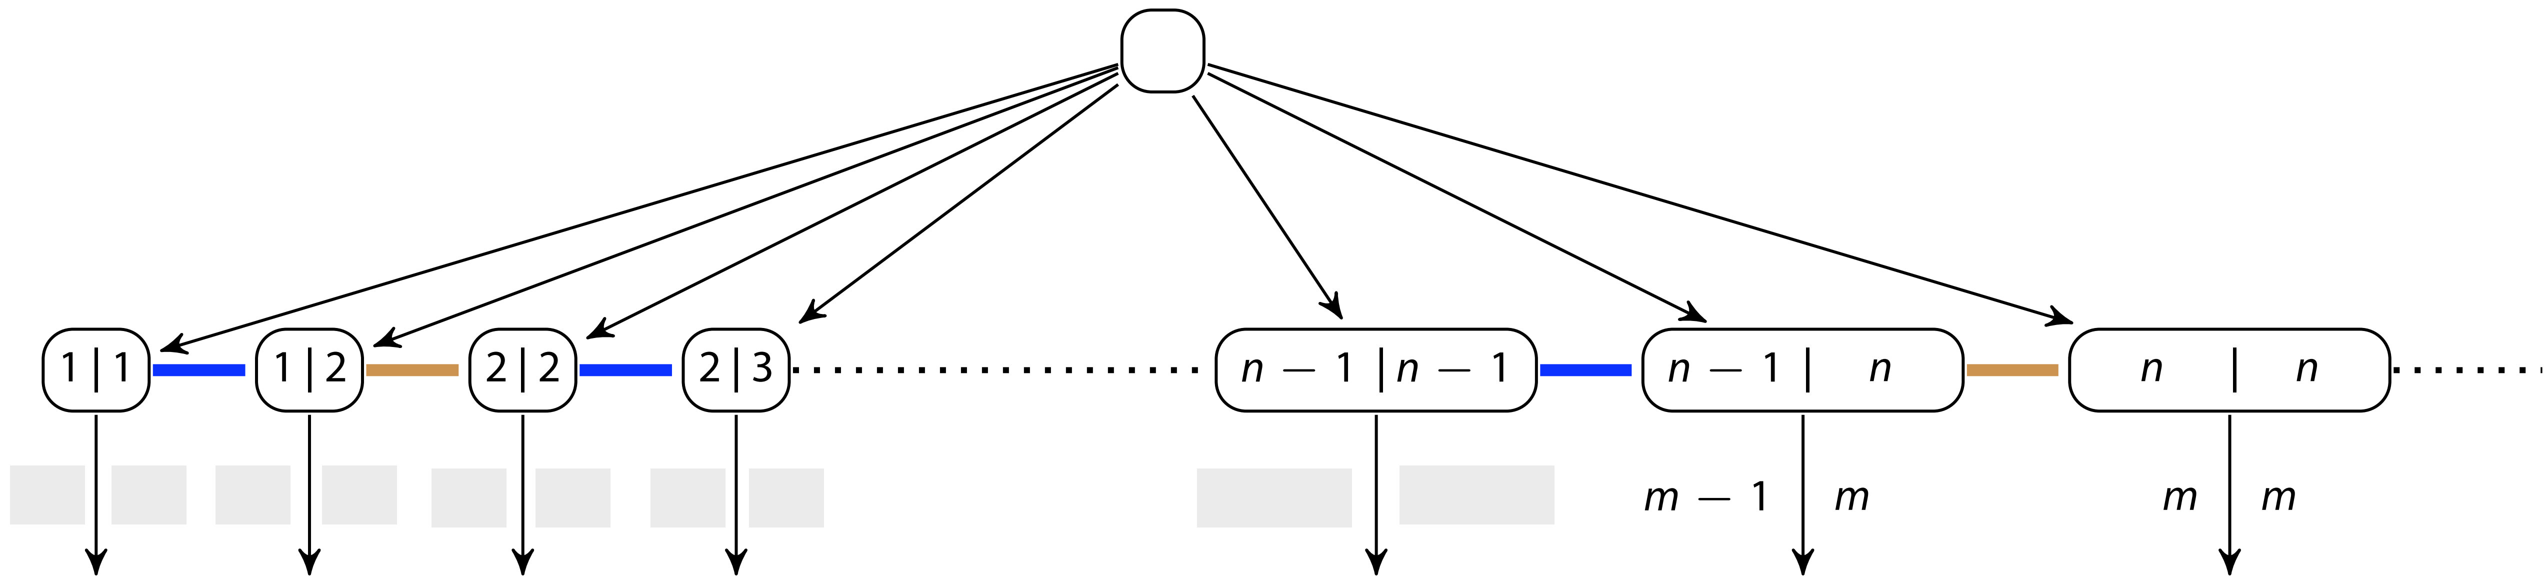

Infinite knowledge hierarchies

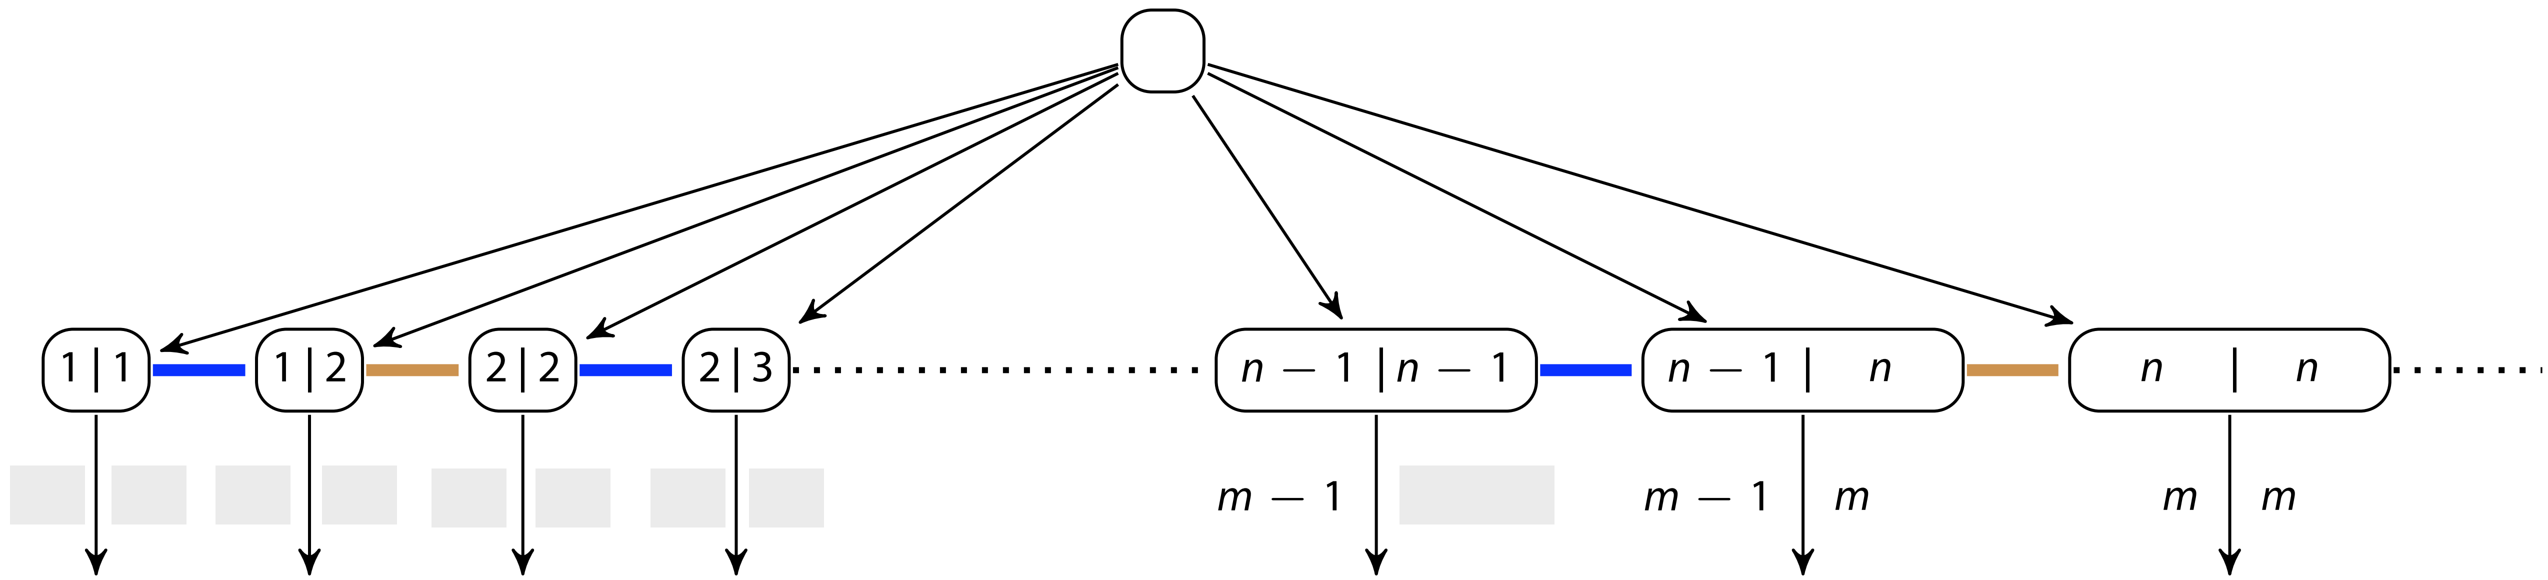

Infinite knowledge hierarchies

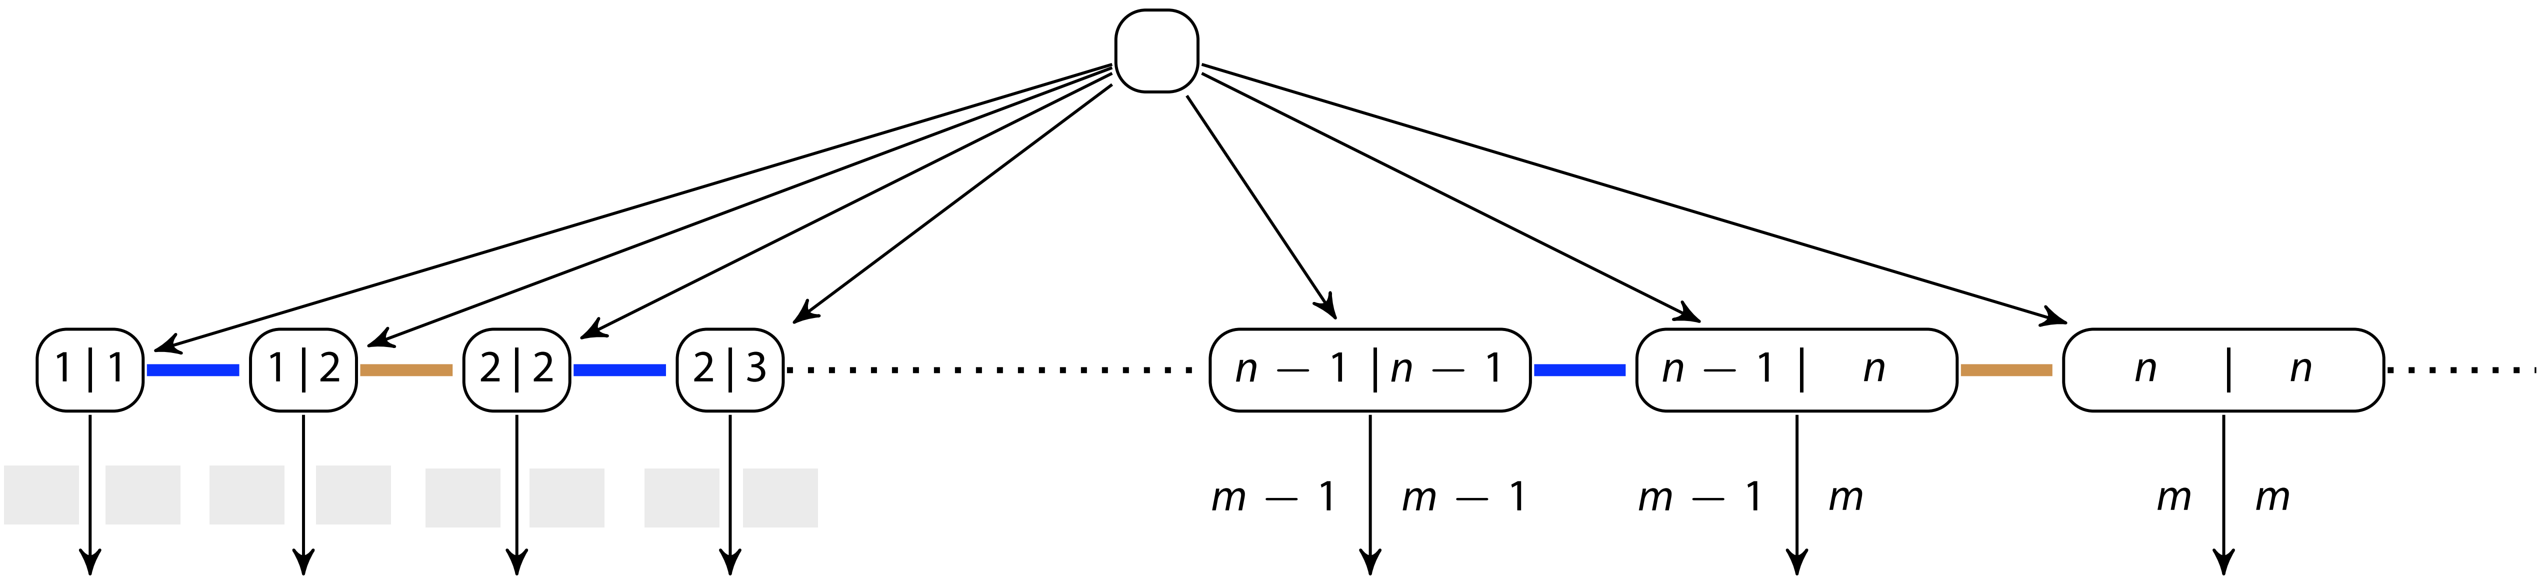

Infinite knowledge hierarchies

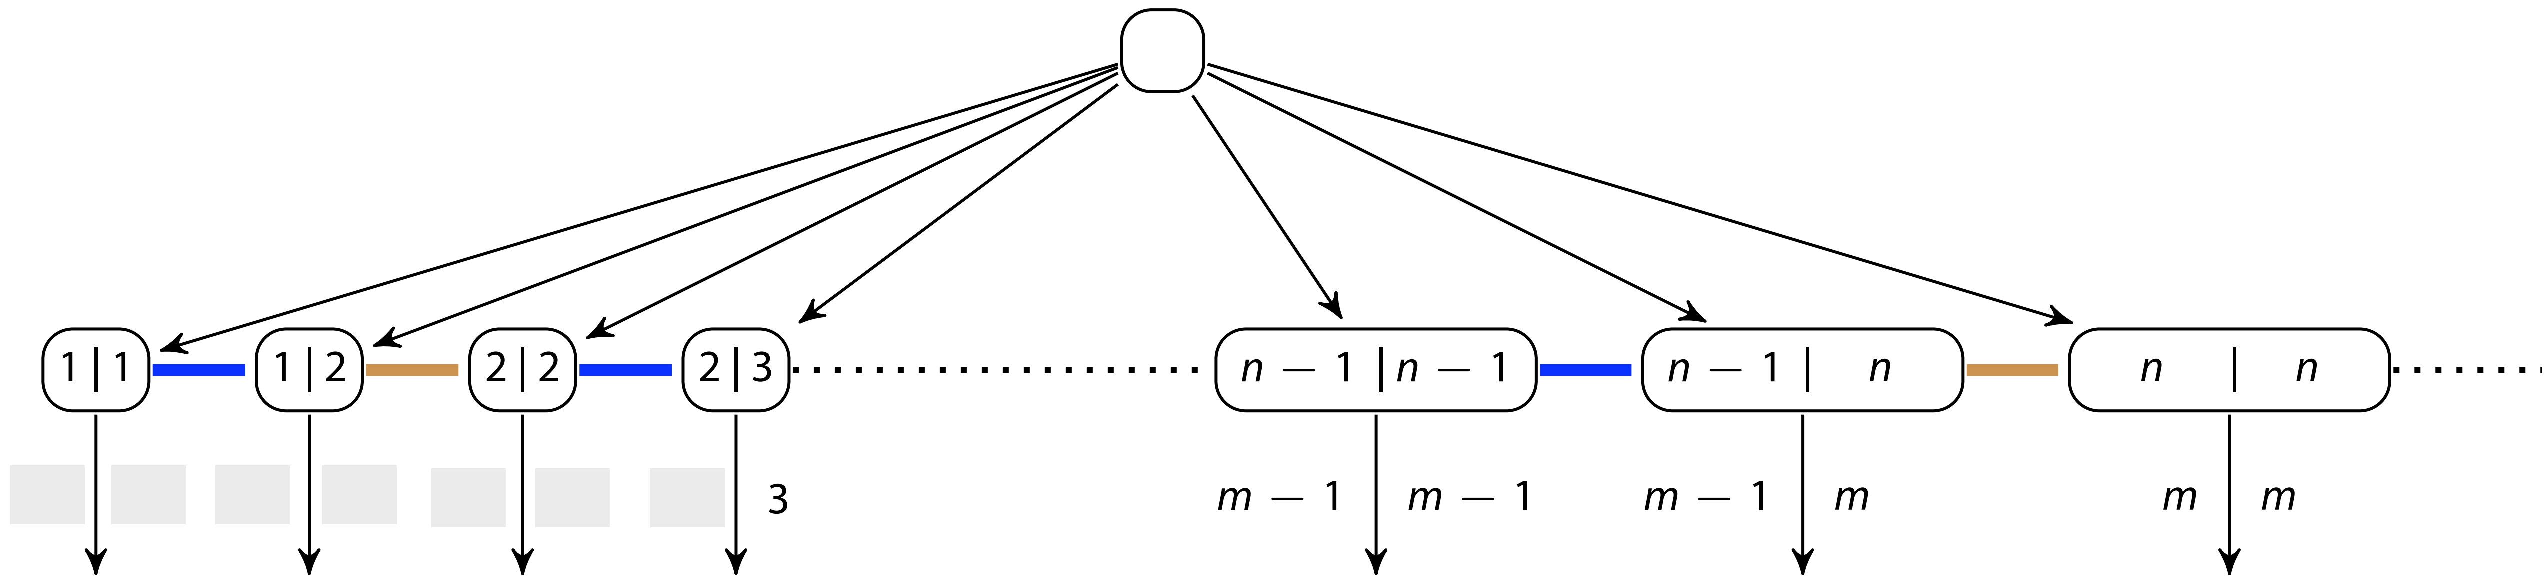

Infinite knowledge hierarchies

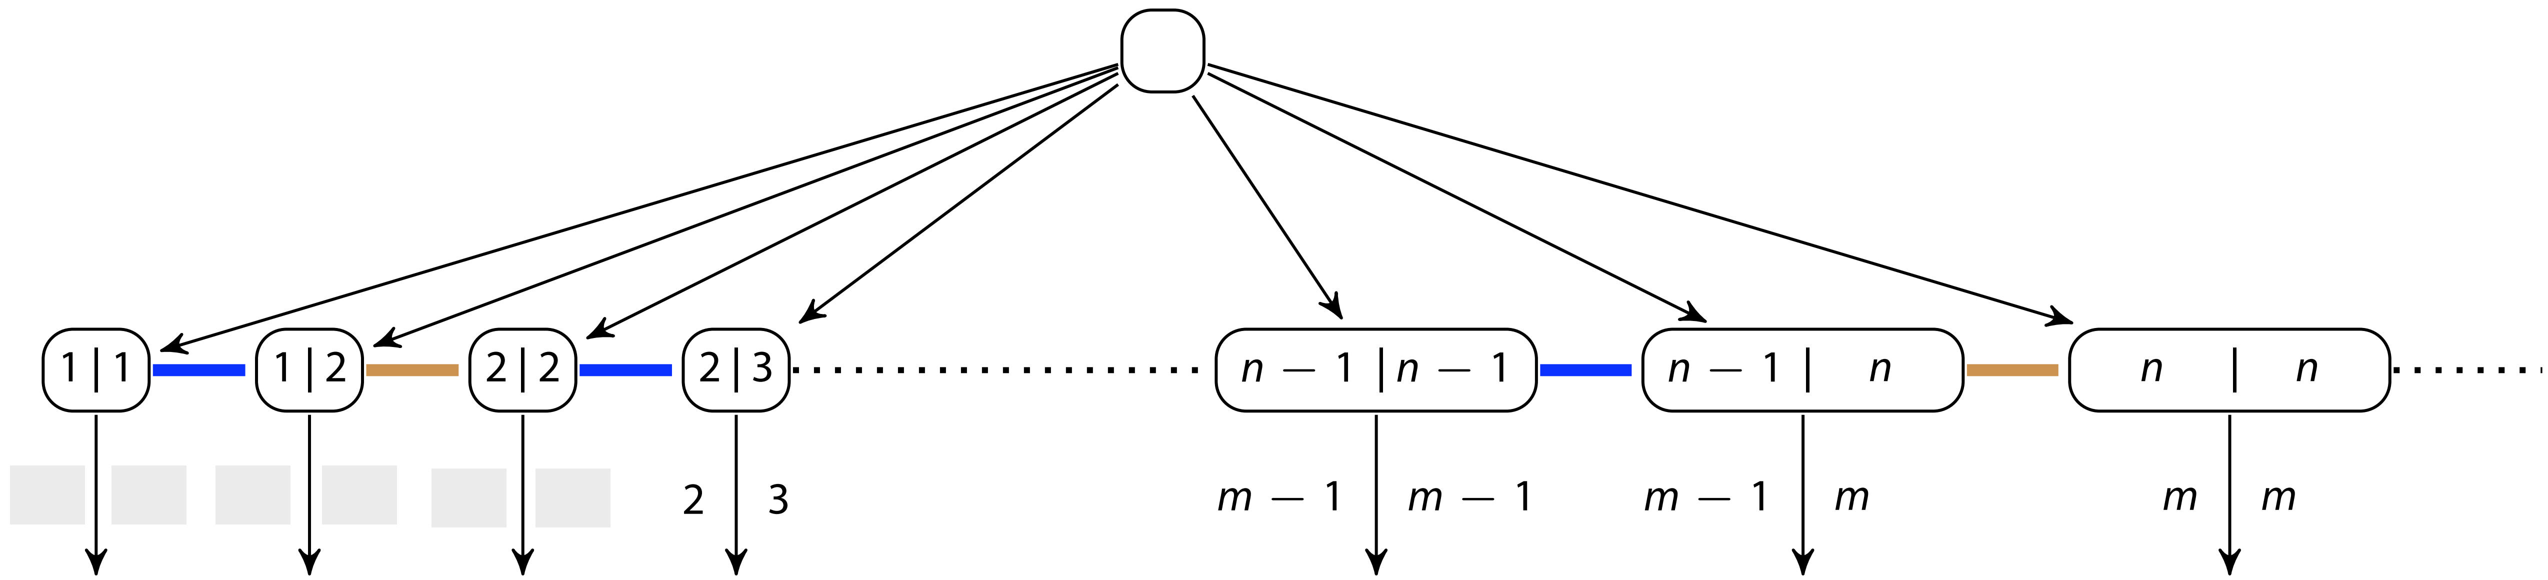

Infinite knowledge hierarchies

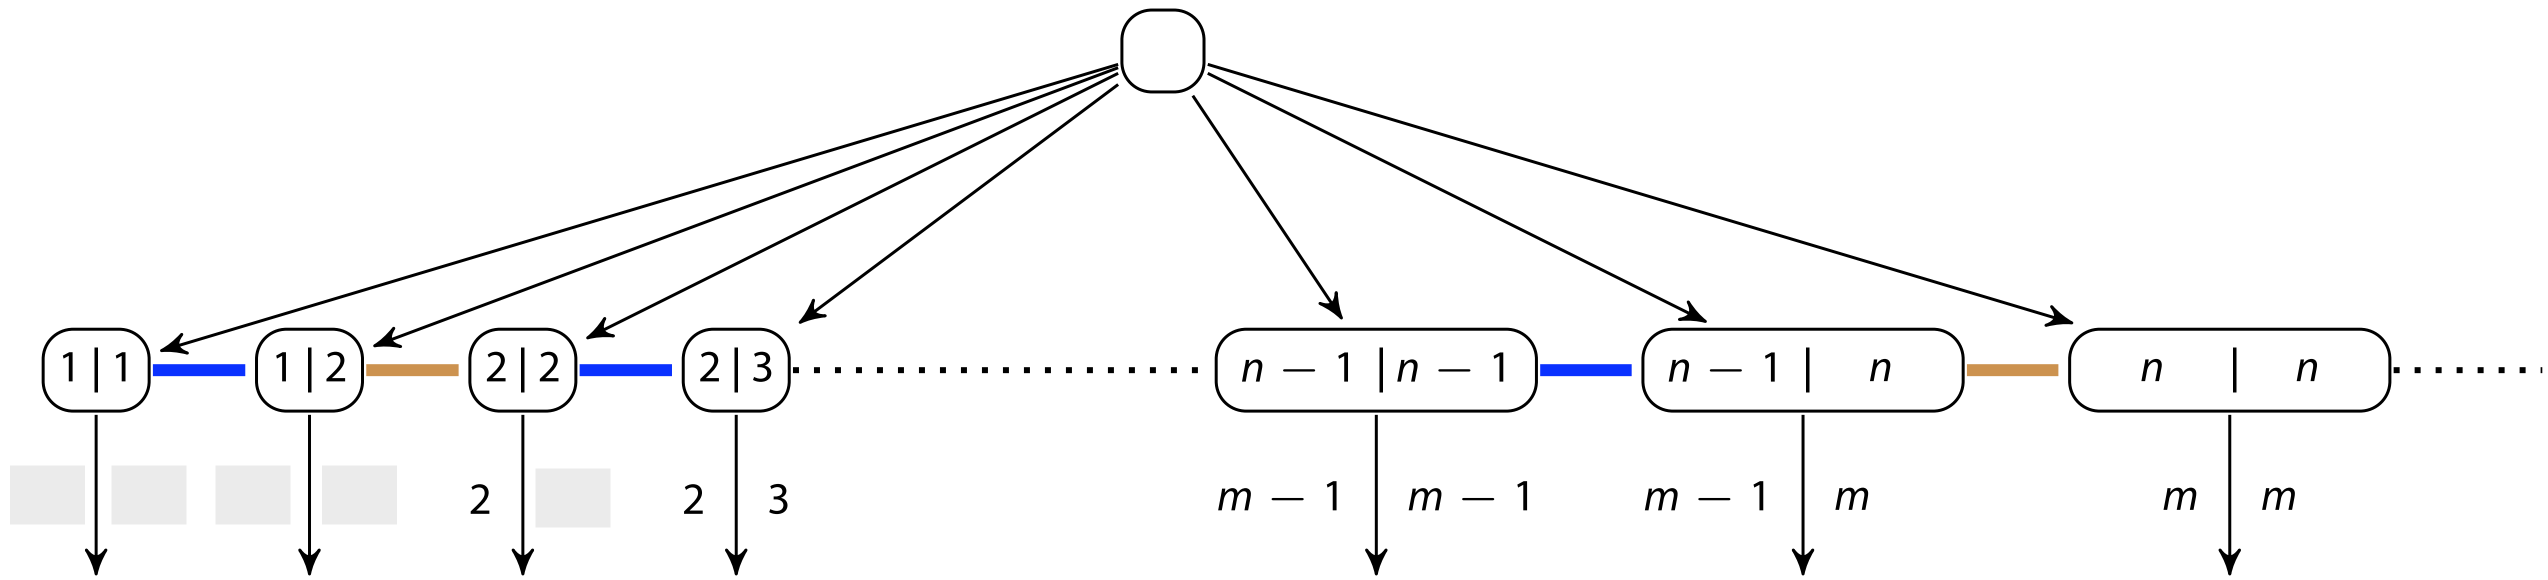

Infinite knowledge hierarchies

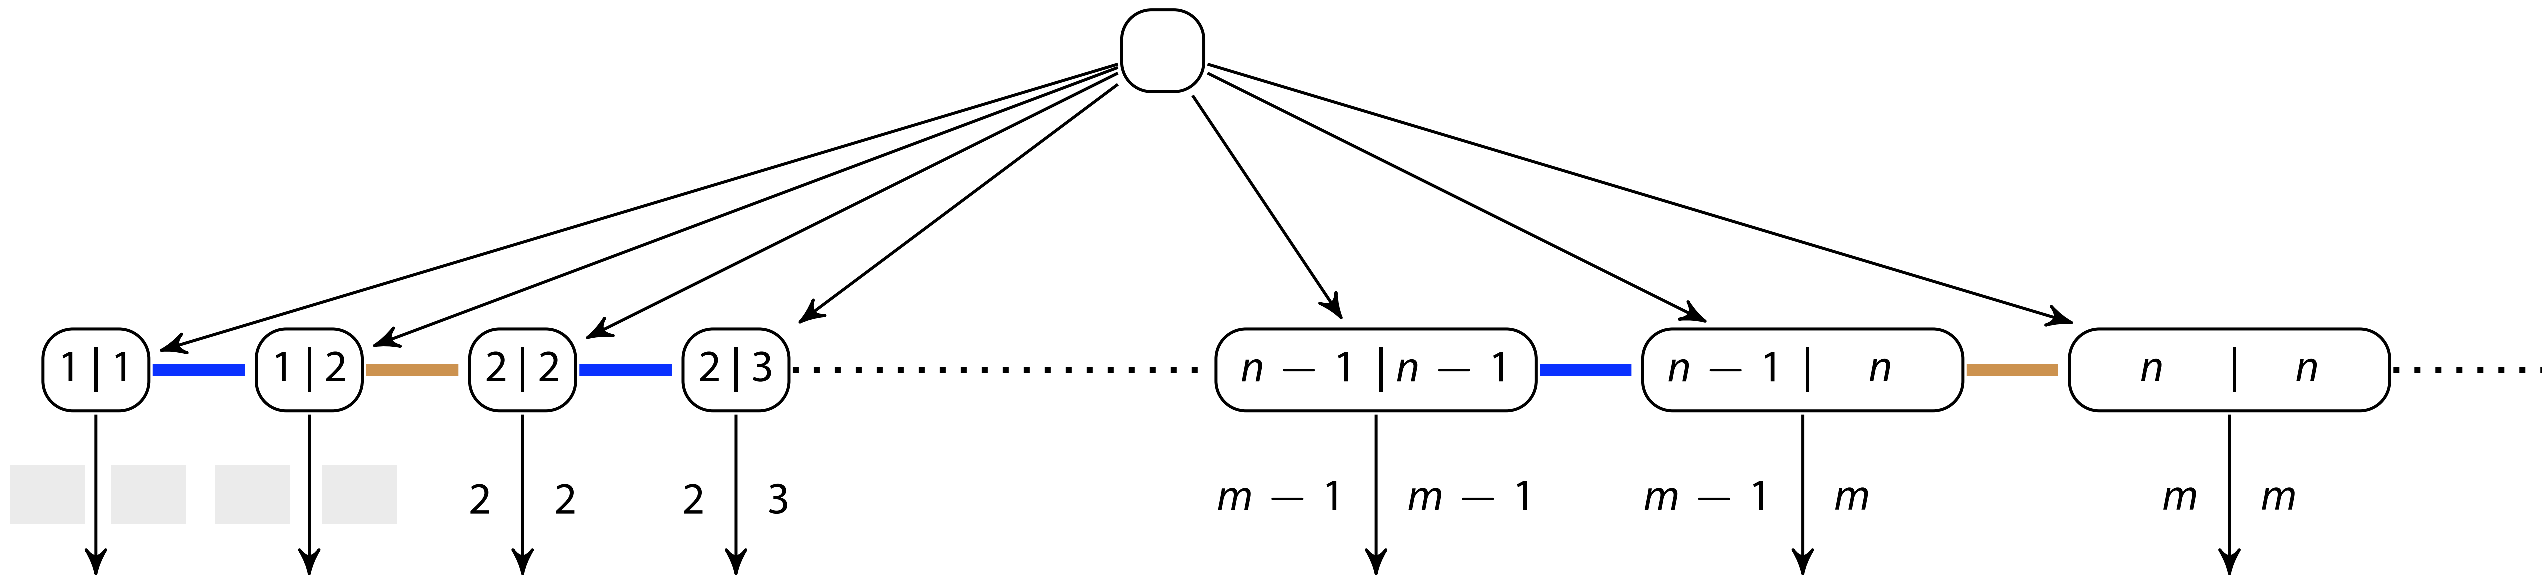

Infinite knowledge hierarchies

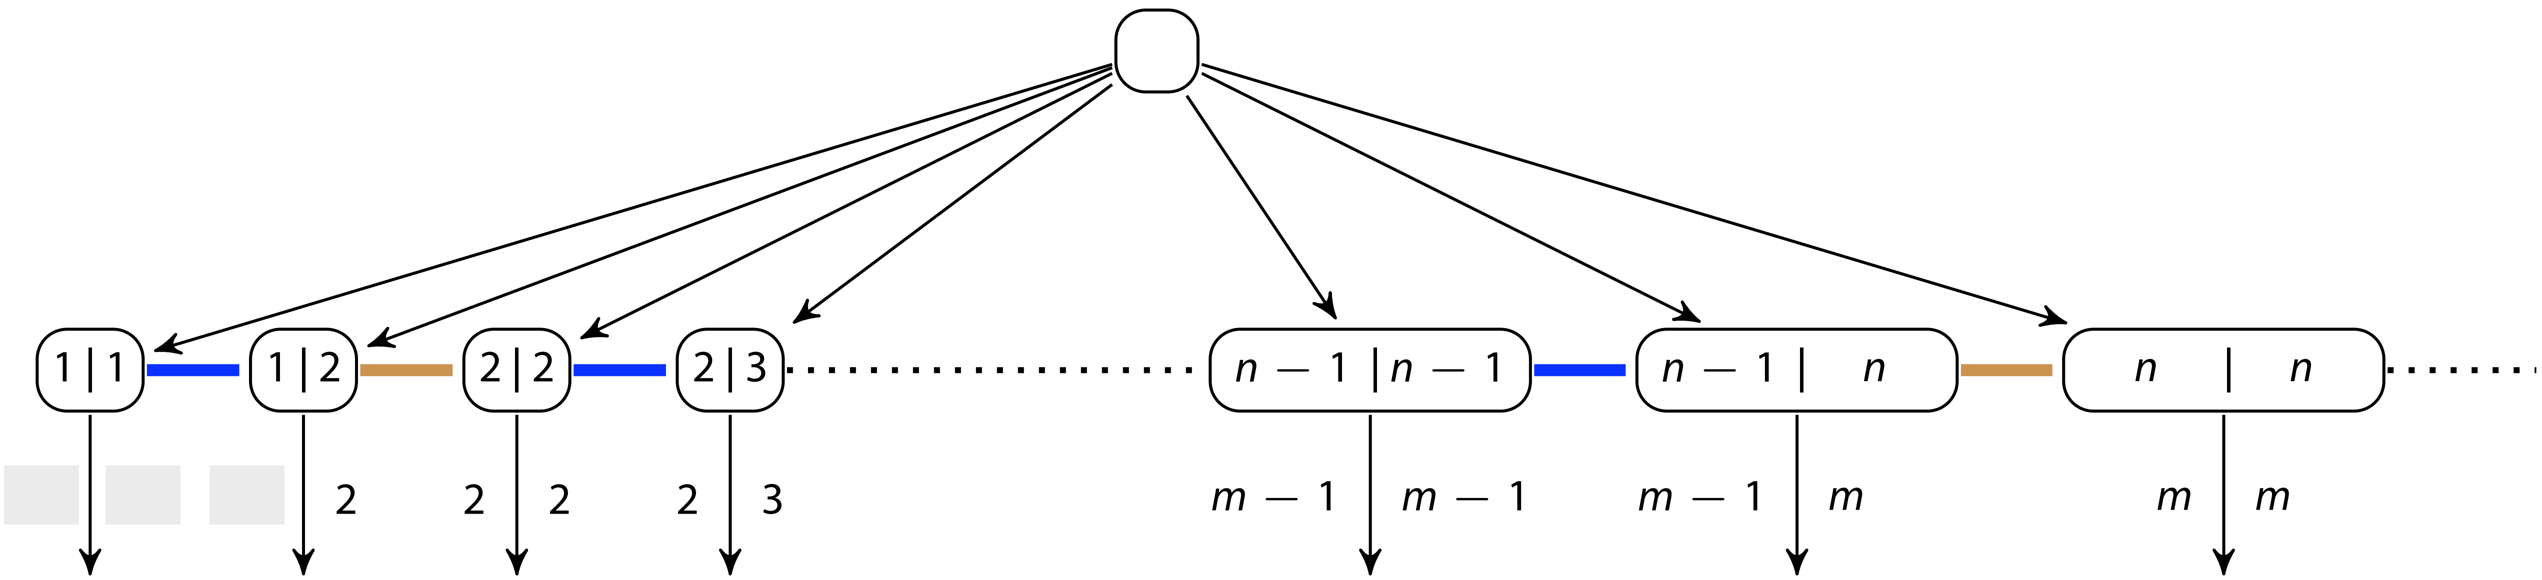

Infinite knowledge hierarchies

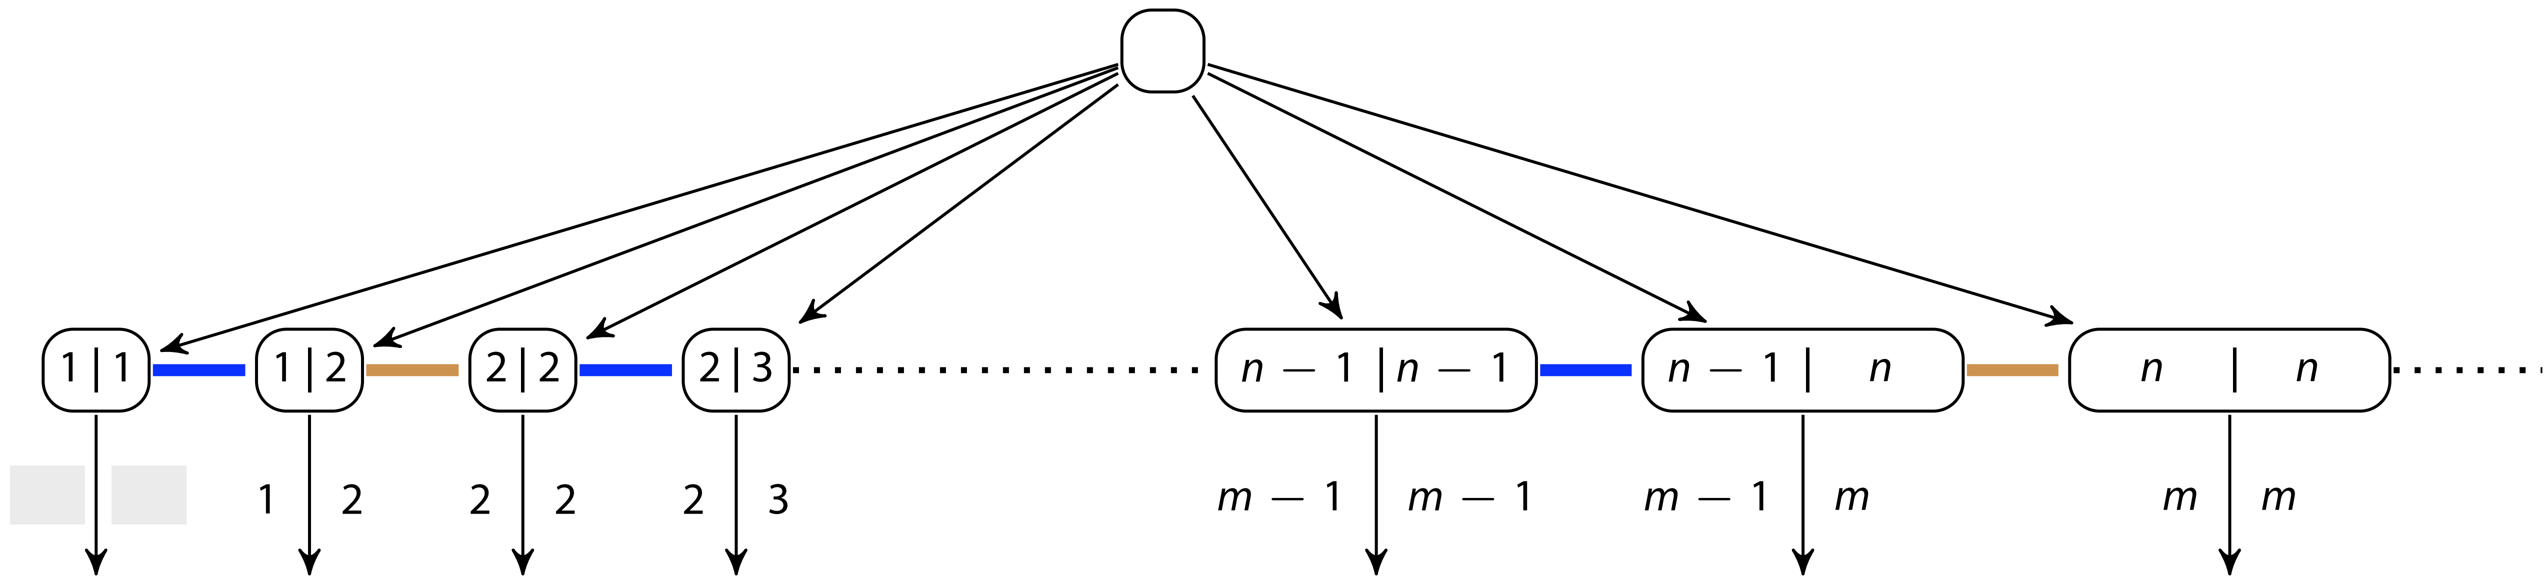

Infinite knowledge hierarchies

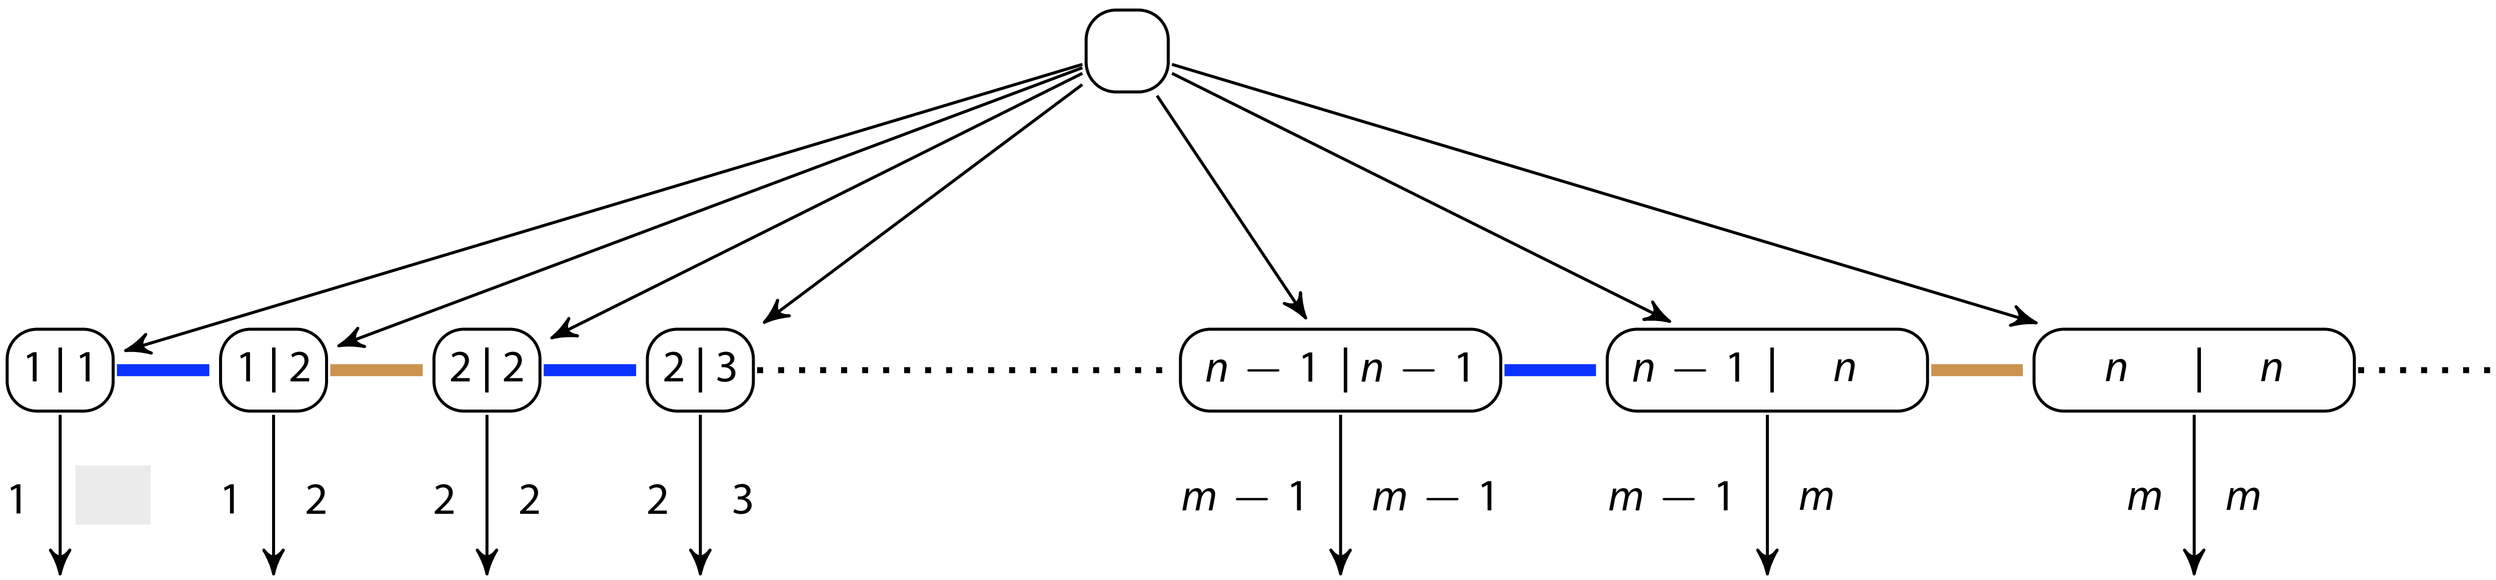

Infinite knowledge hierarchies

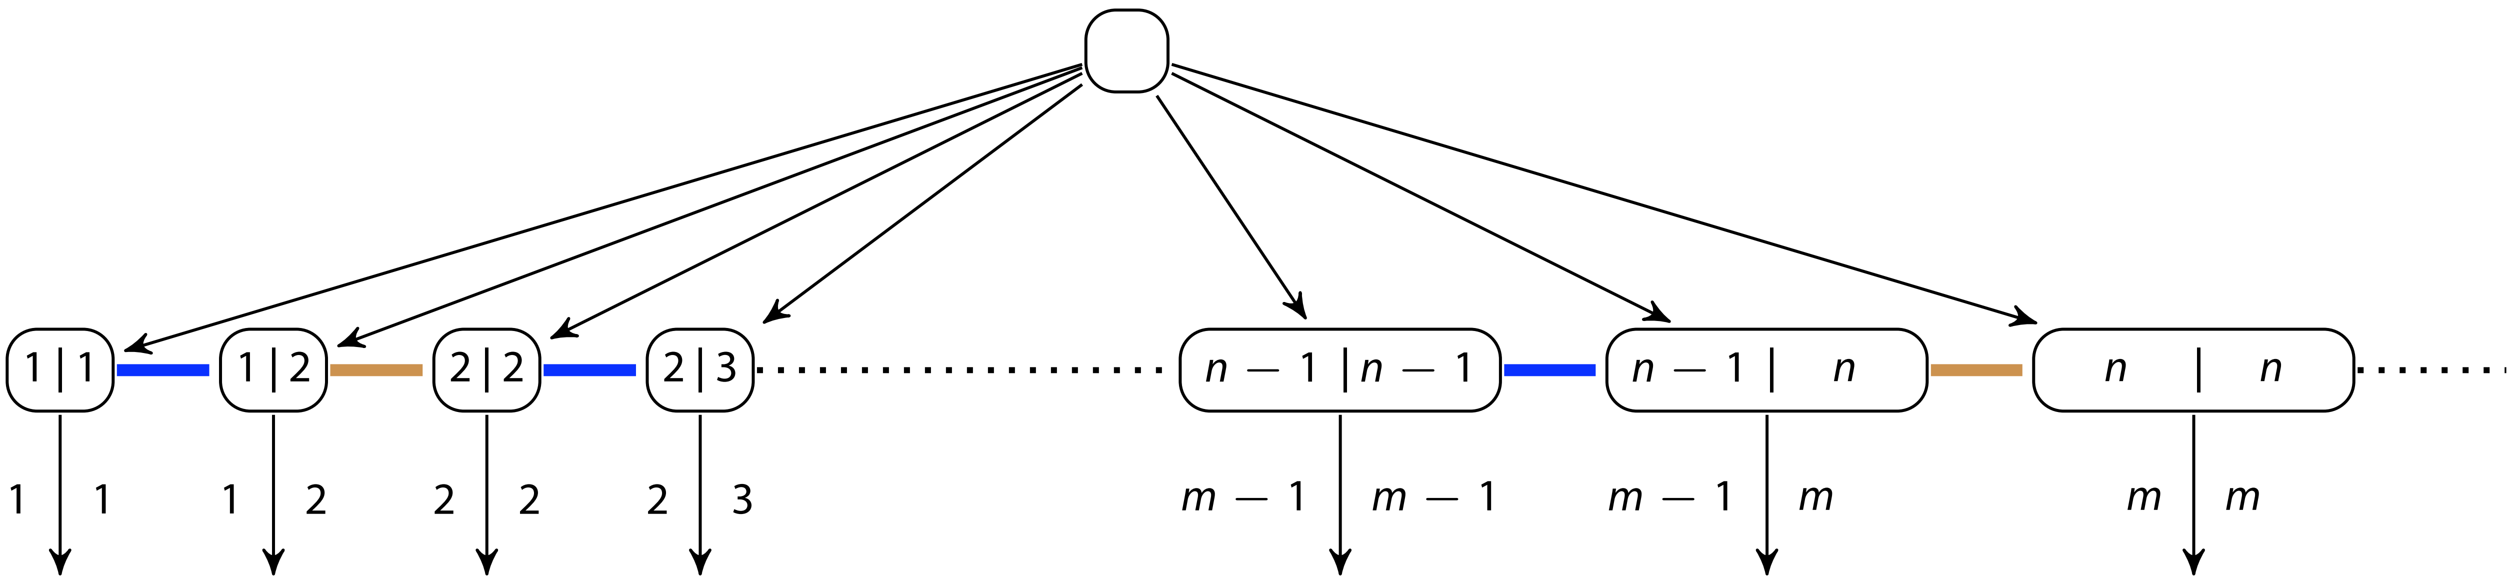

Infinite knowledge hierarchies

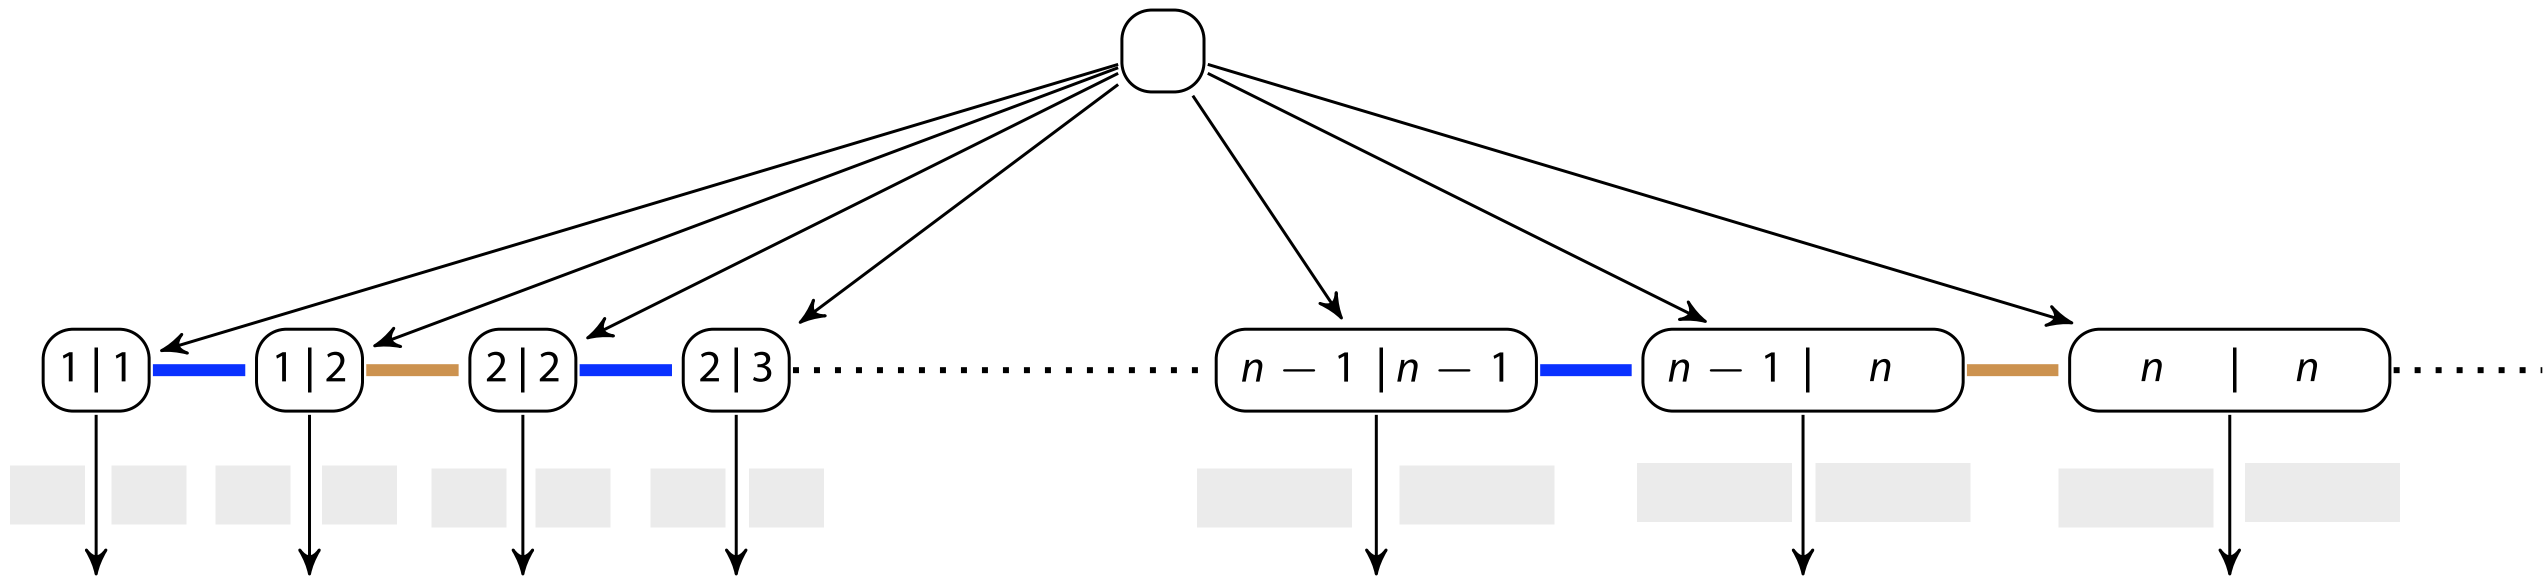

Solution of

$$\begin{cases} s^1(1) & = 1 \\ s^0(n) & = s^1(n) \\ s^1(n+1) & = s^0(n) + 1 \end{cases}$$

unique: $s^i(n) = n$.

Issue #2: Computation

Two-dimensional dynamics when viewed in extensive form:

- sequences of actions (along transitions)
- chains of inference (along indistinguishability)

Both dimensions may be unbounded.

The Grid -- *Undecidability*

Unbounded knowledge hierarchy on a finite graph

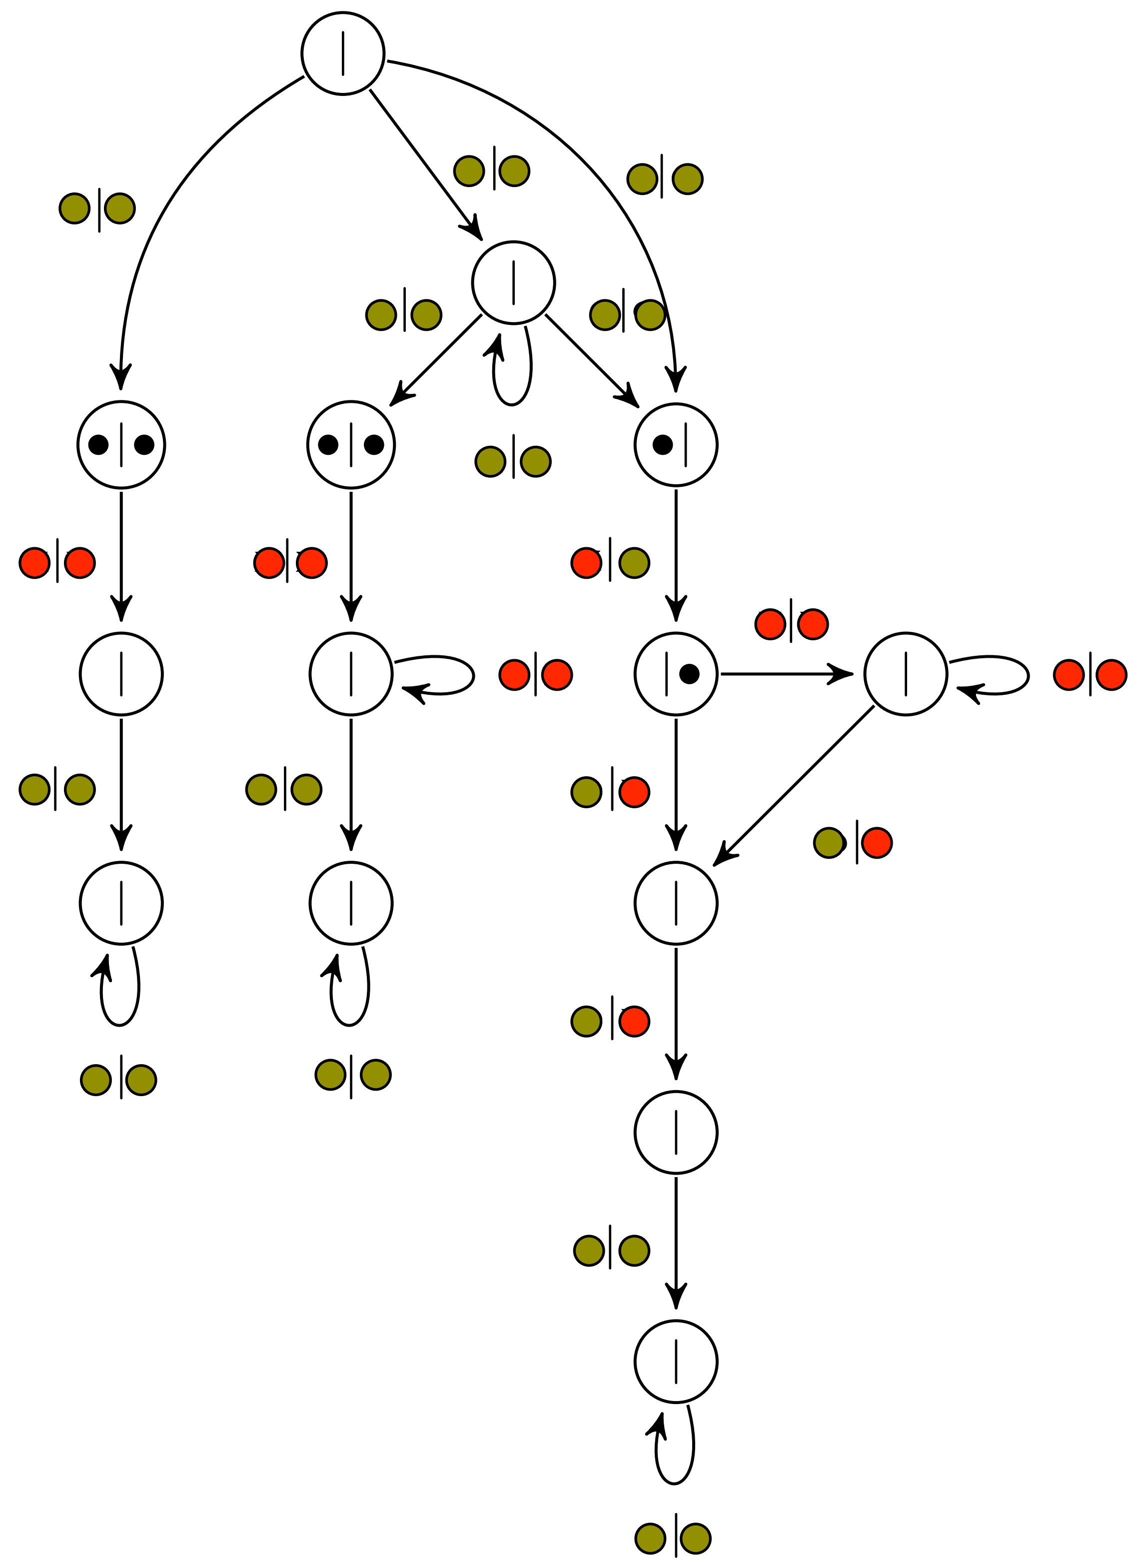

Extensive form

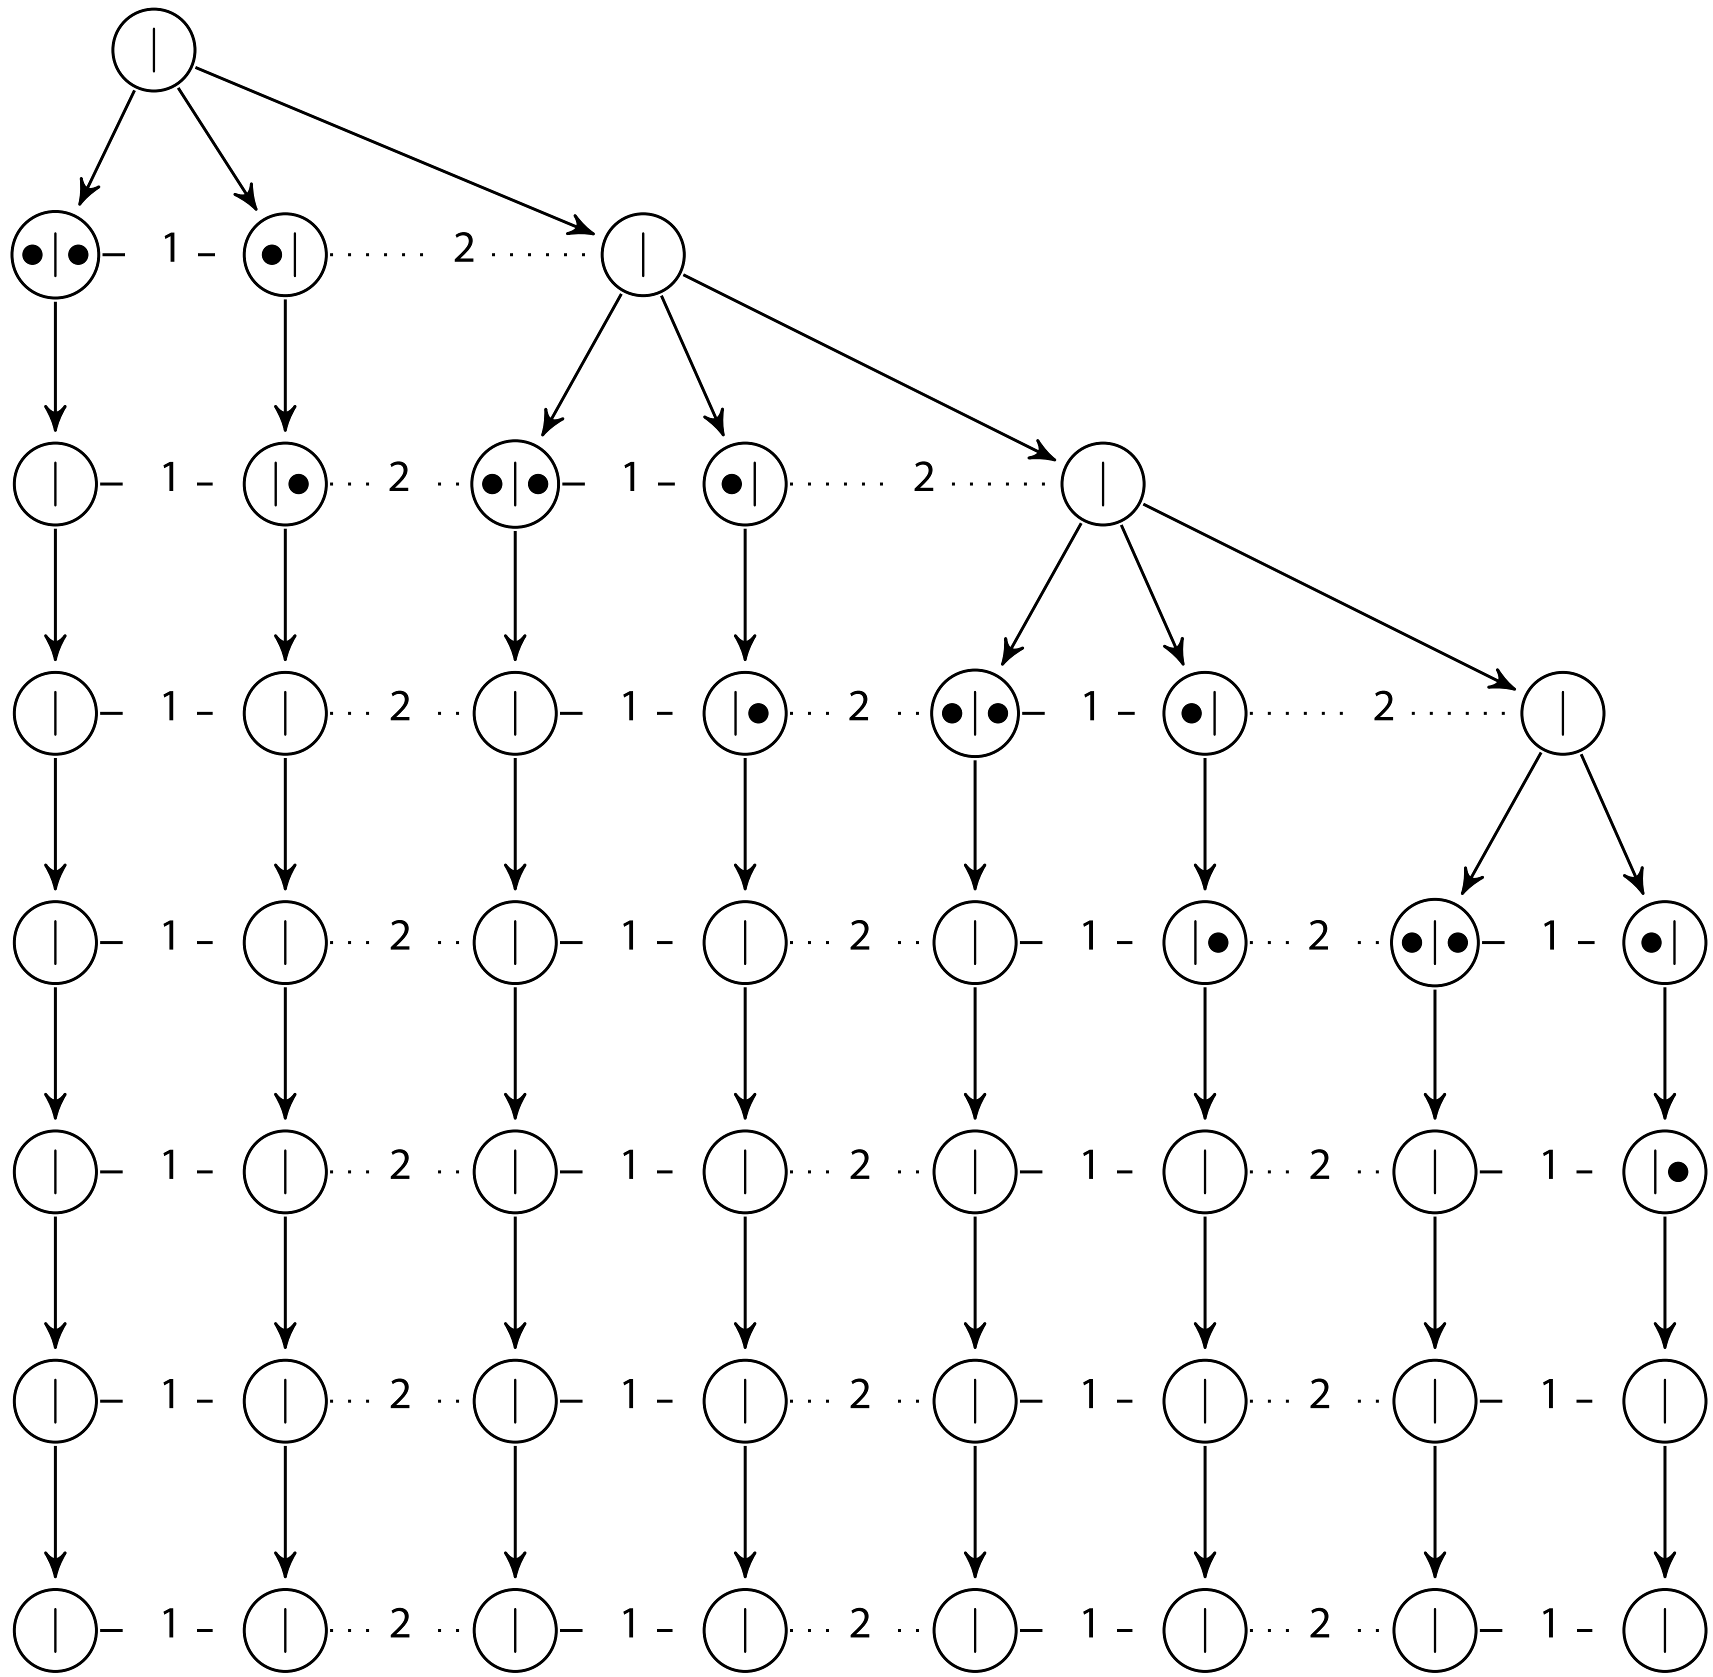

Extensive form

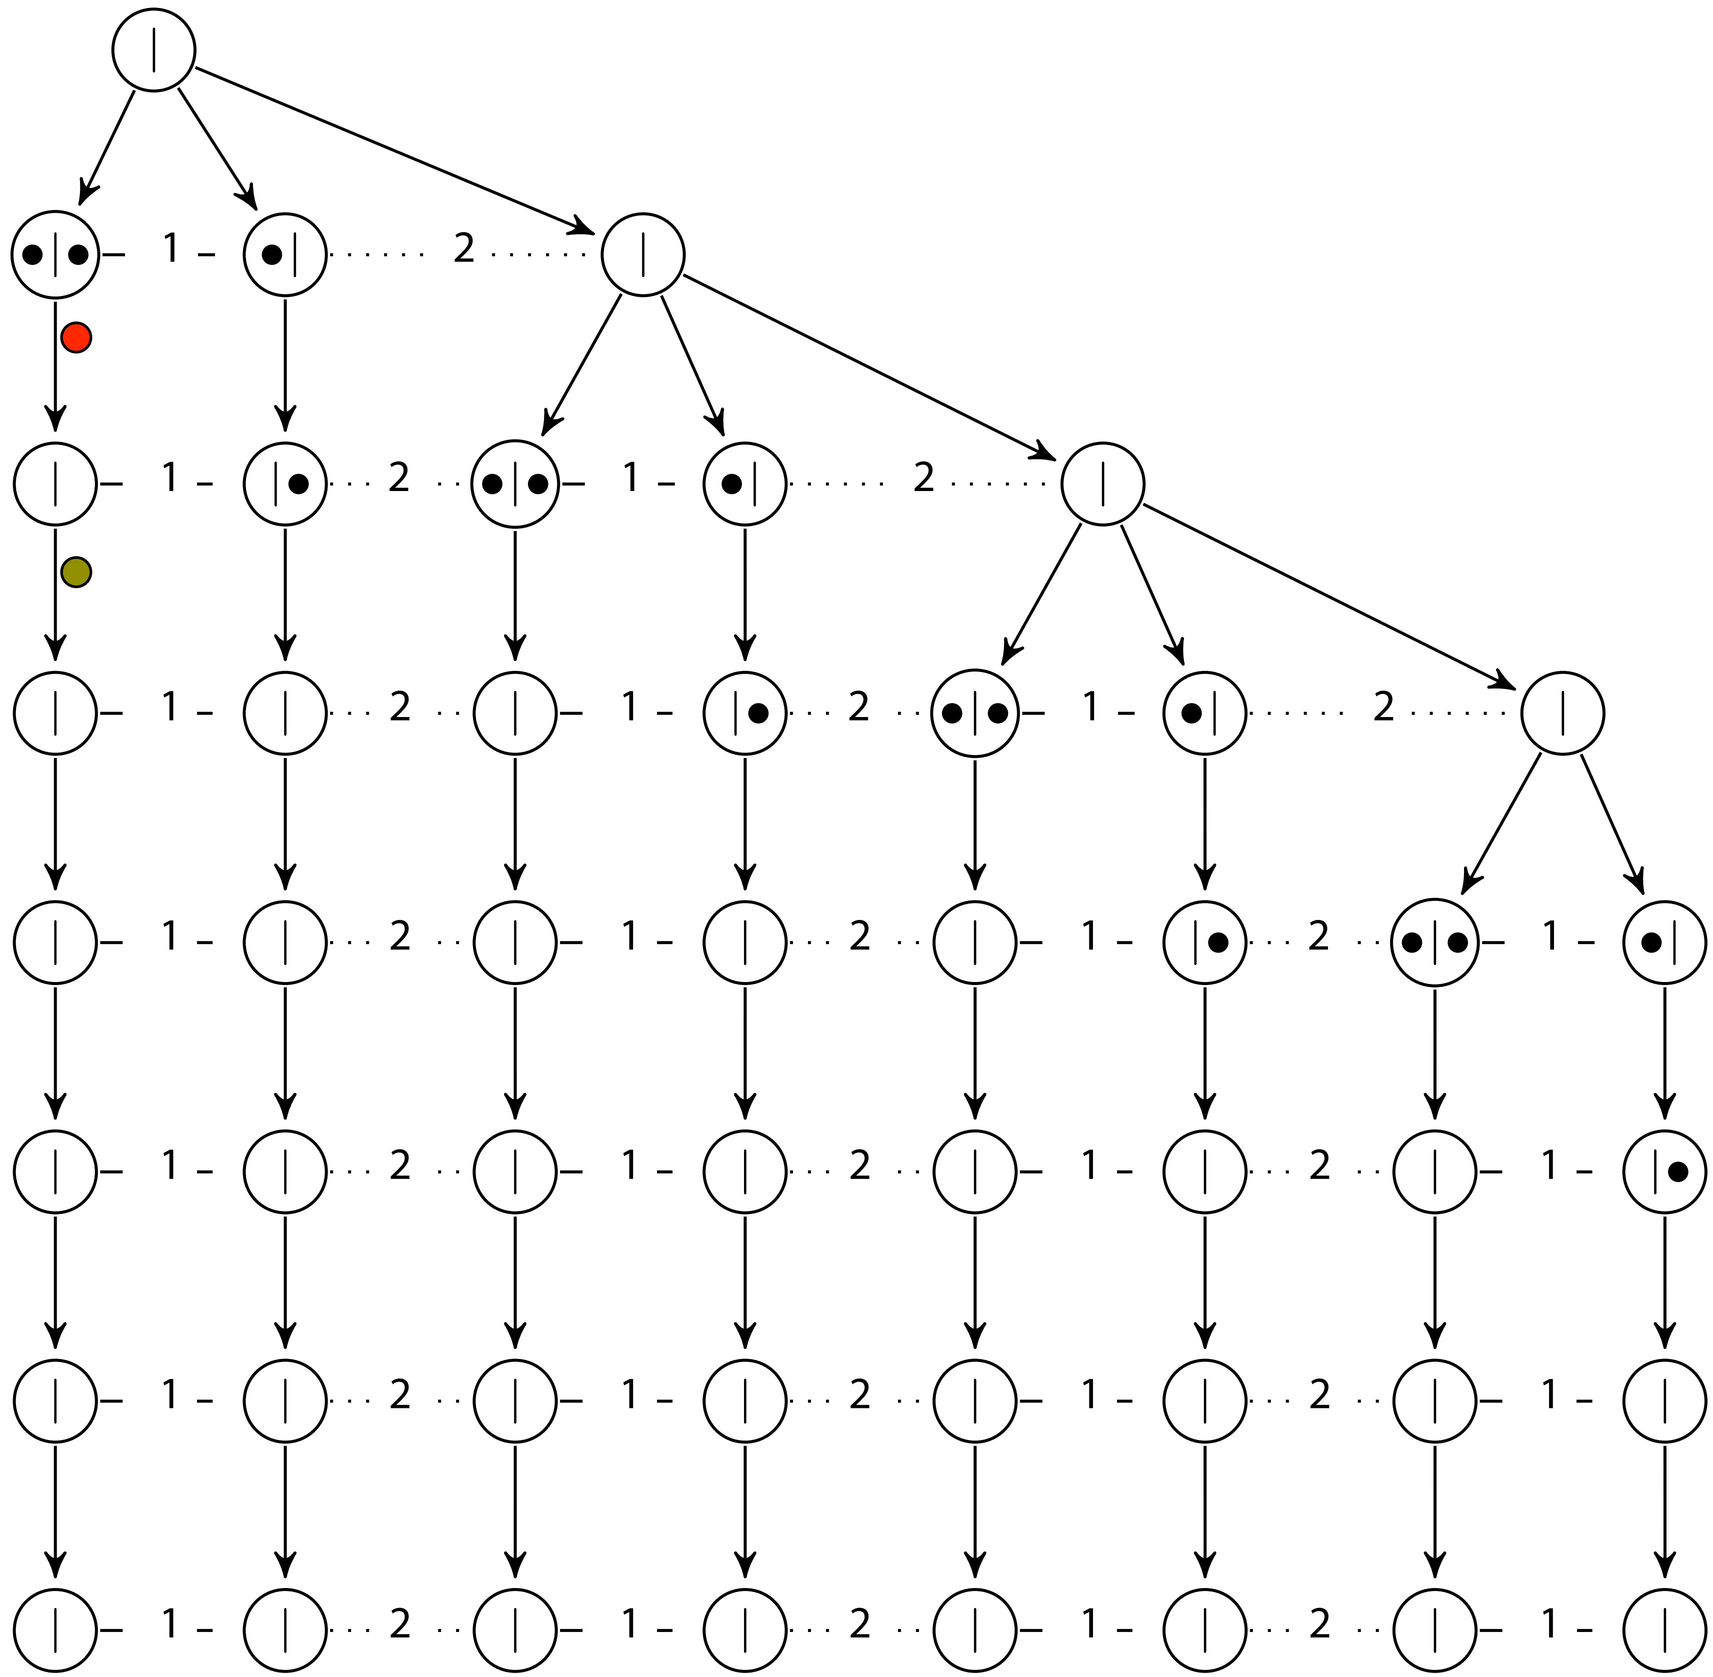

Extensive form

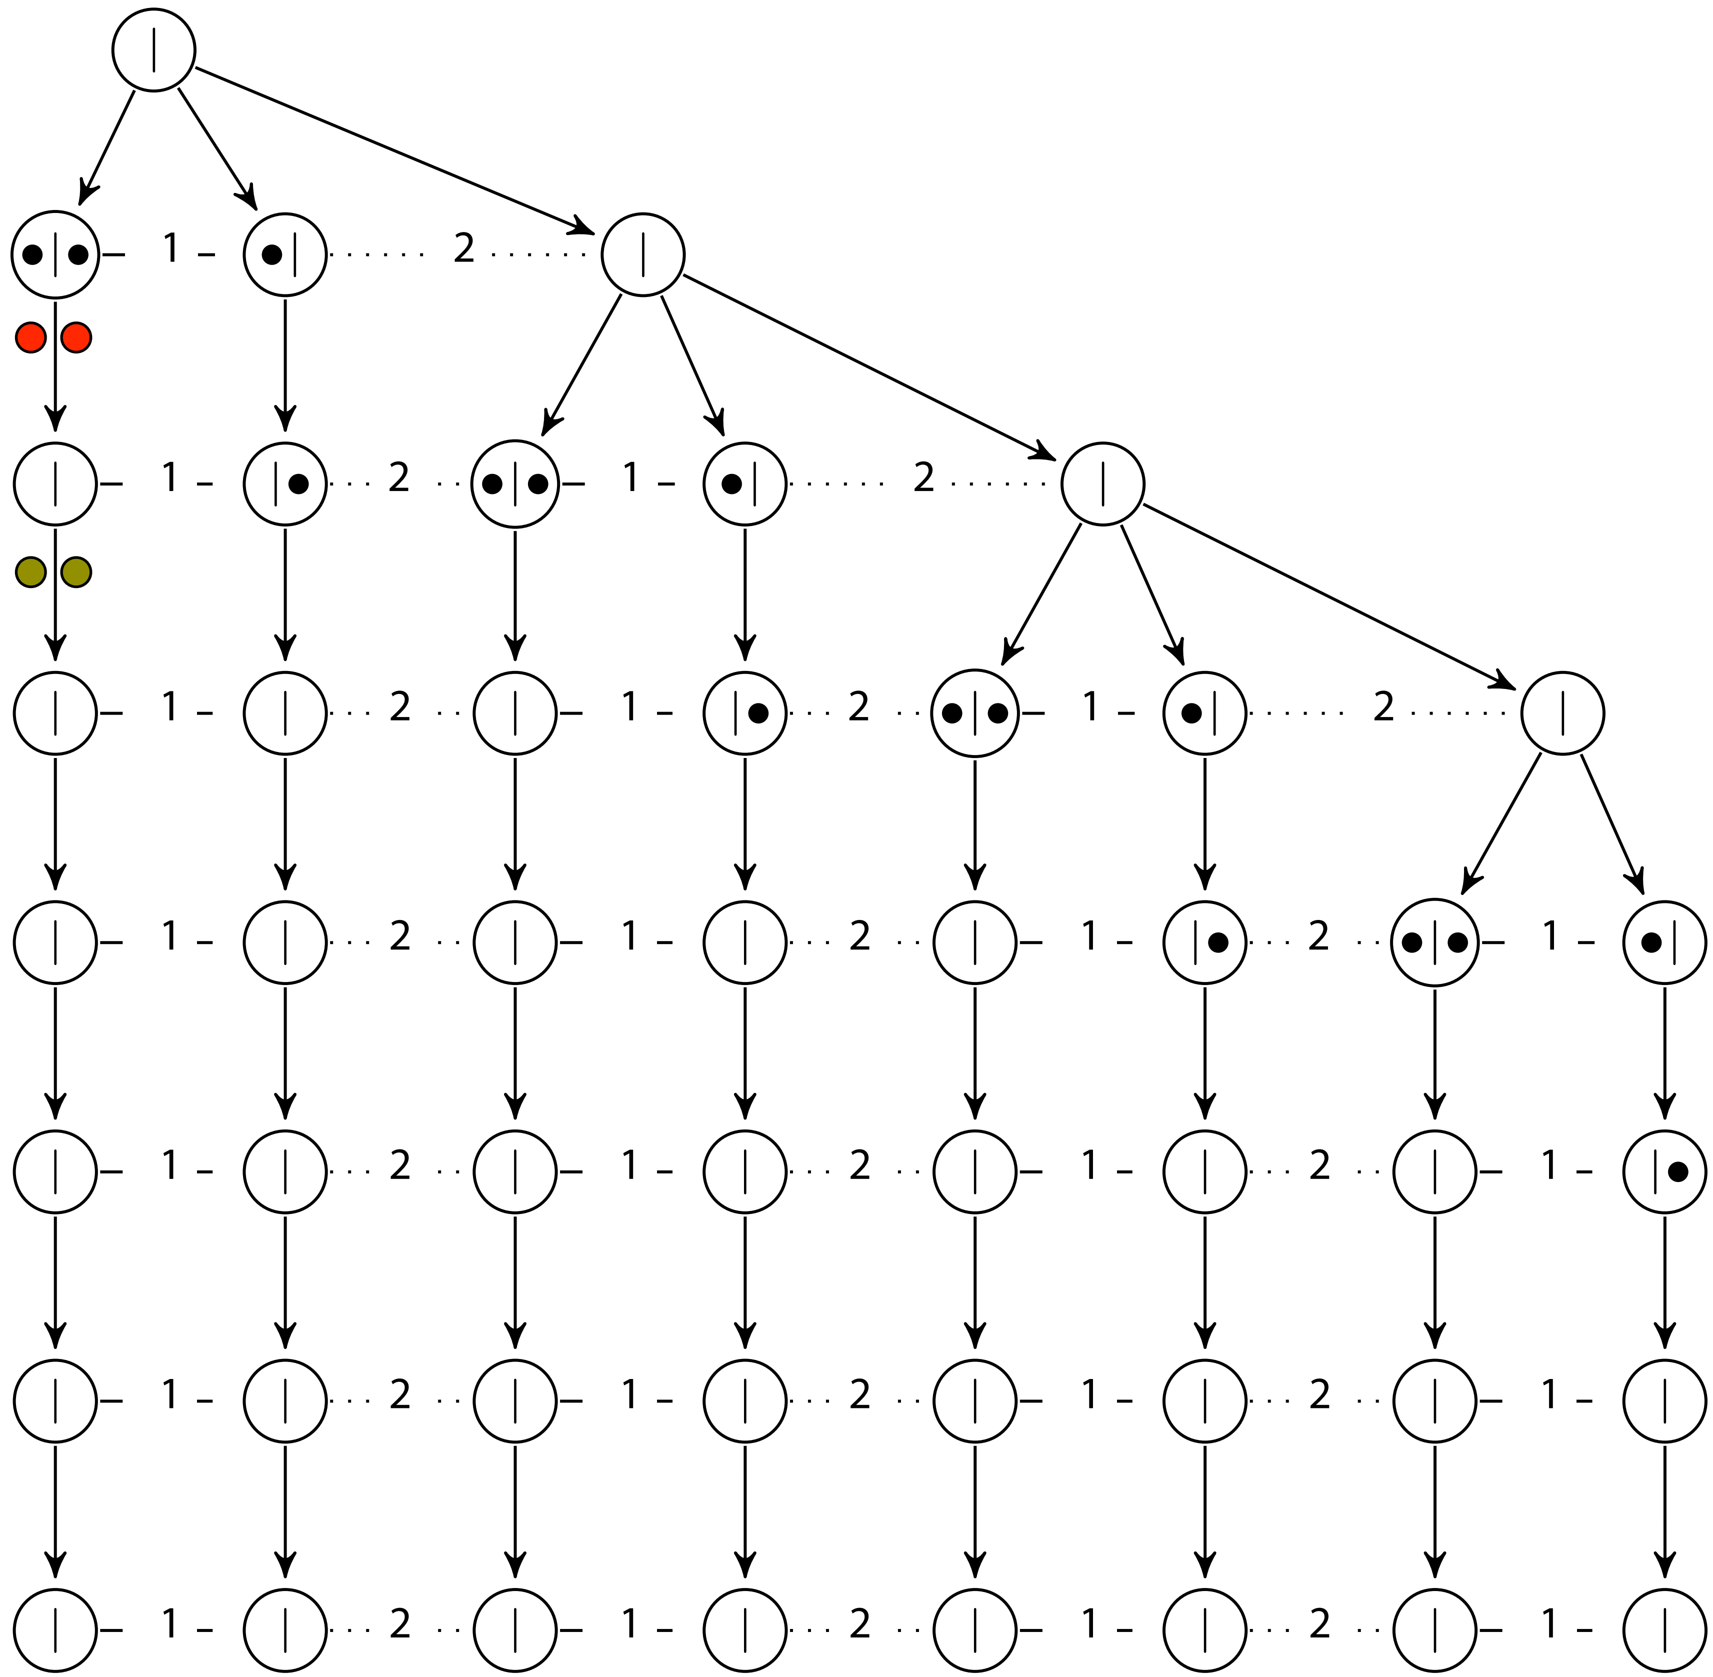

Extensive form

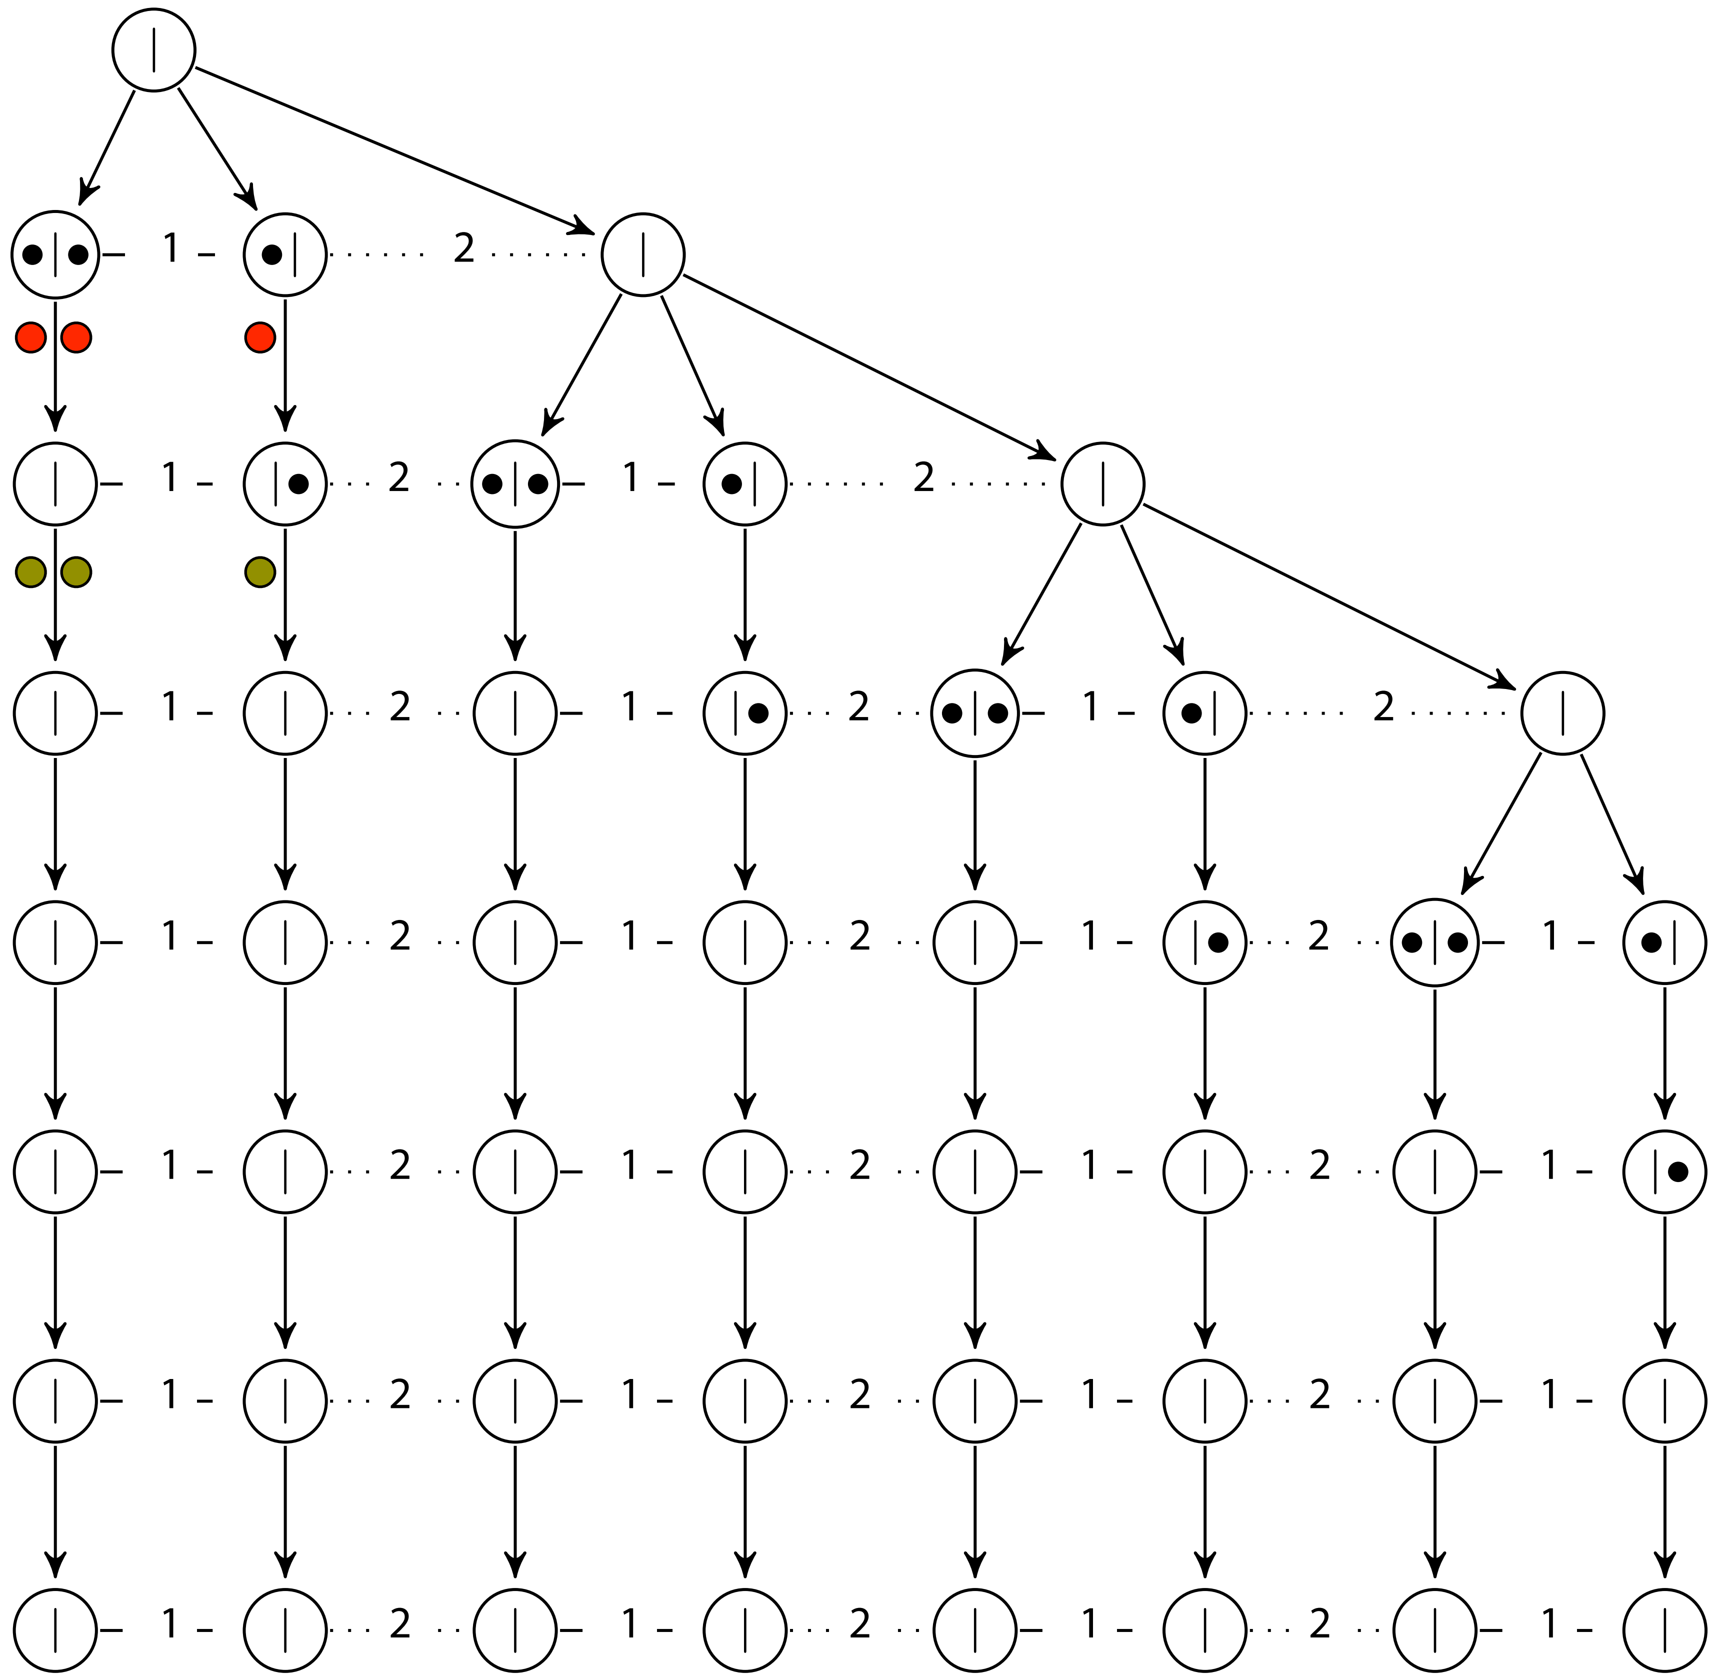

Extensive form

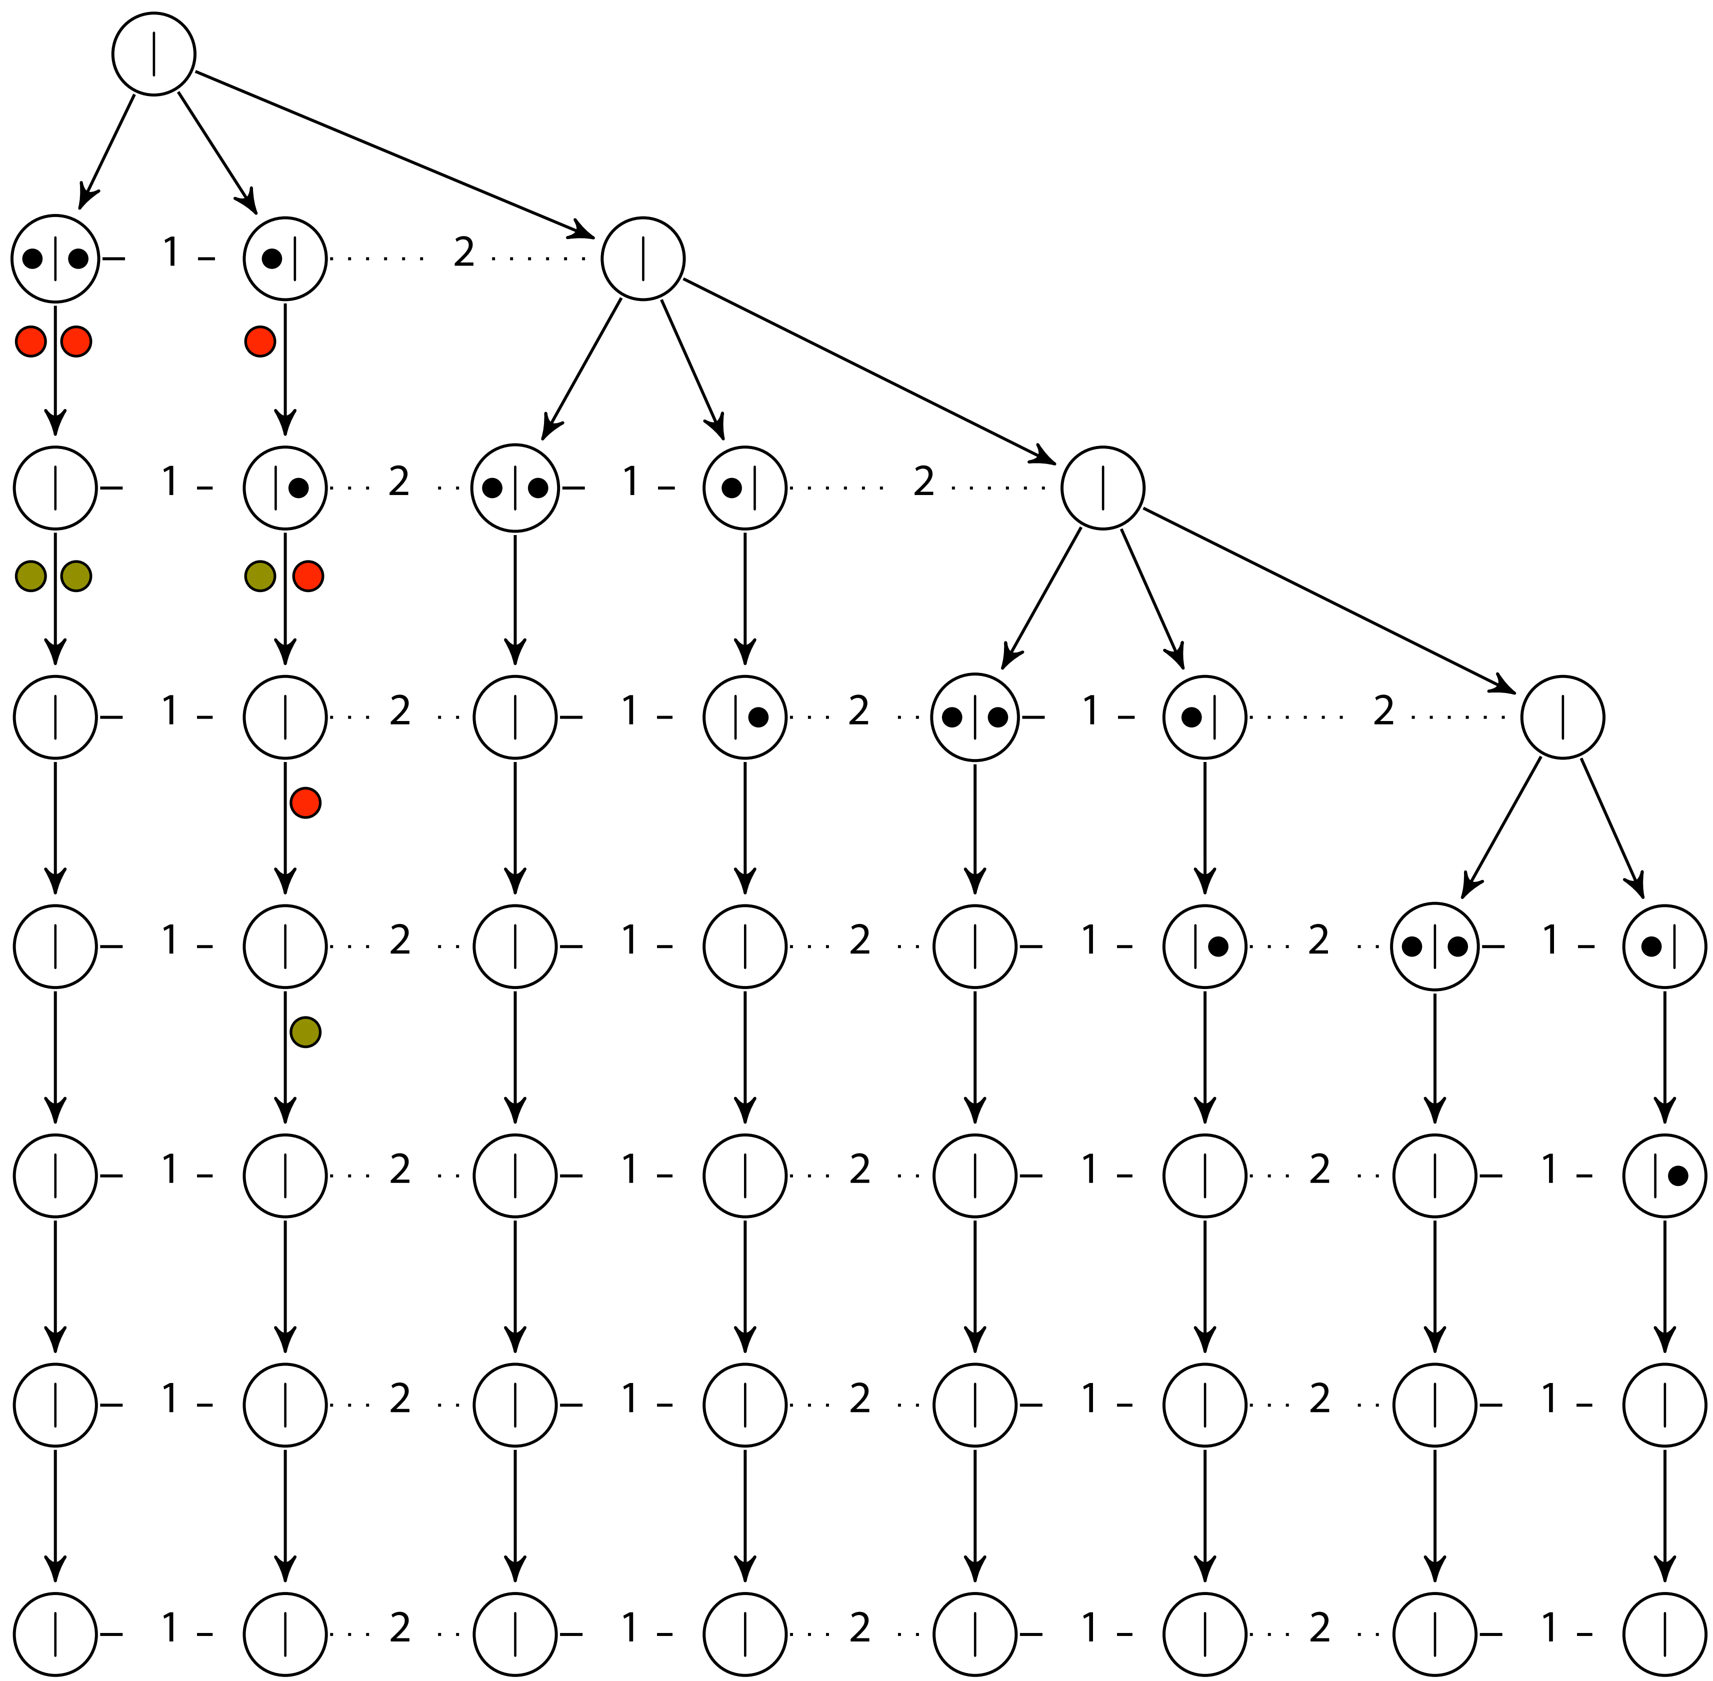

Extensive form

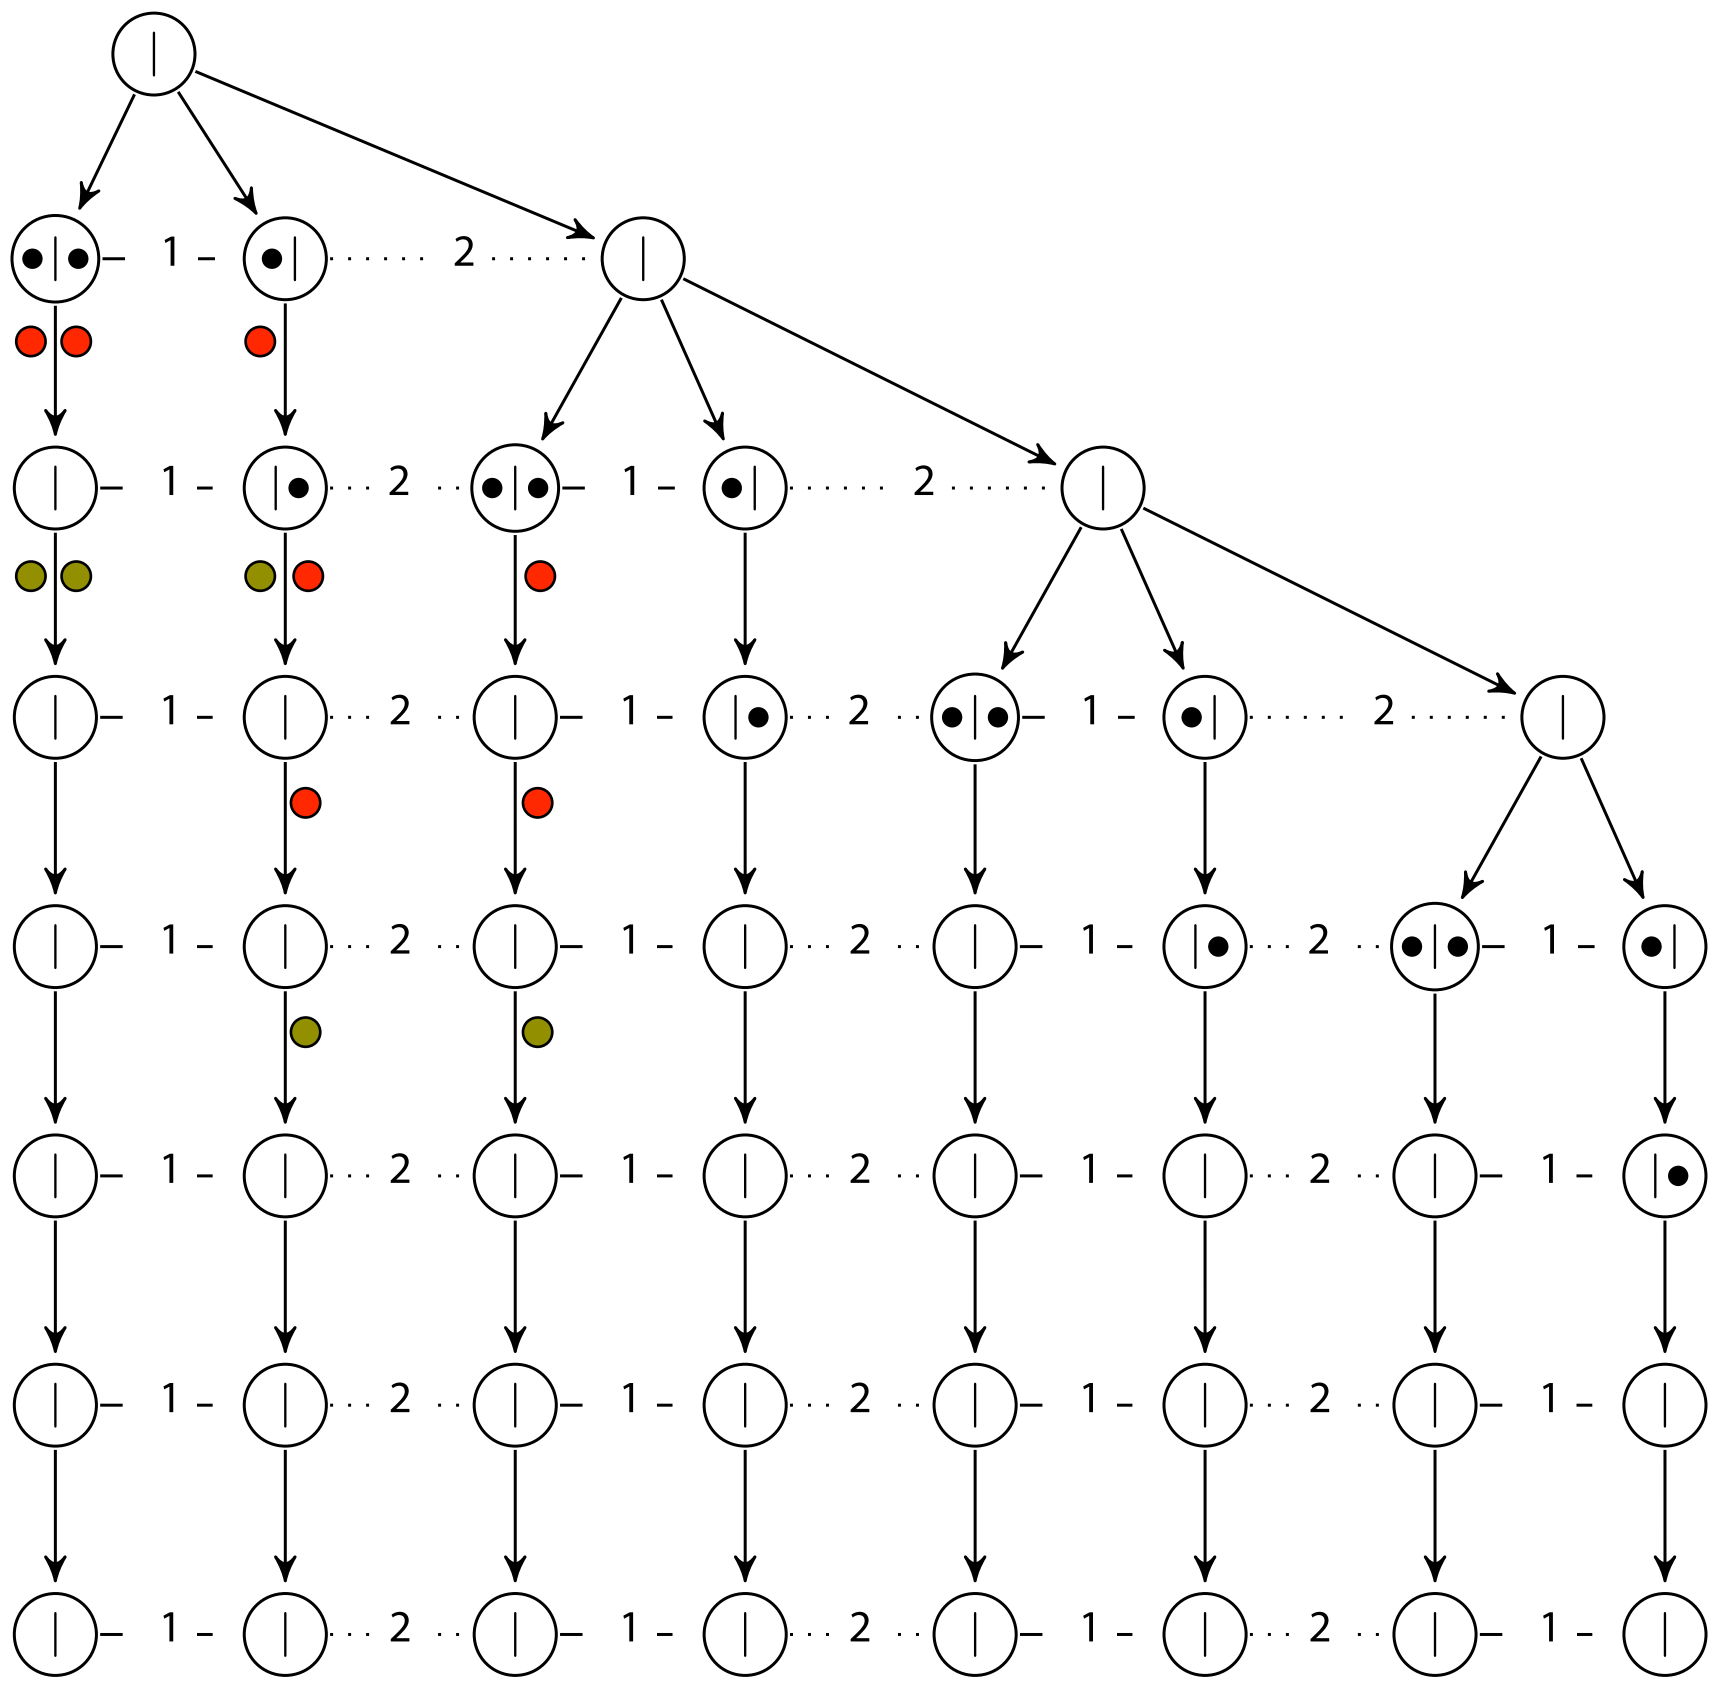

Extensive form

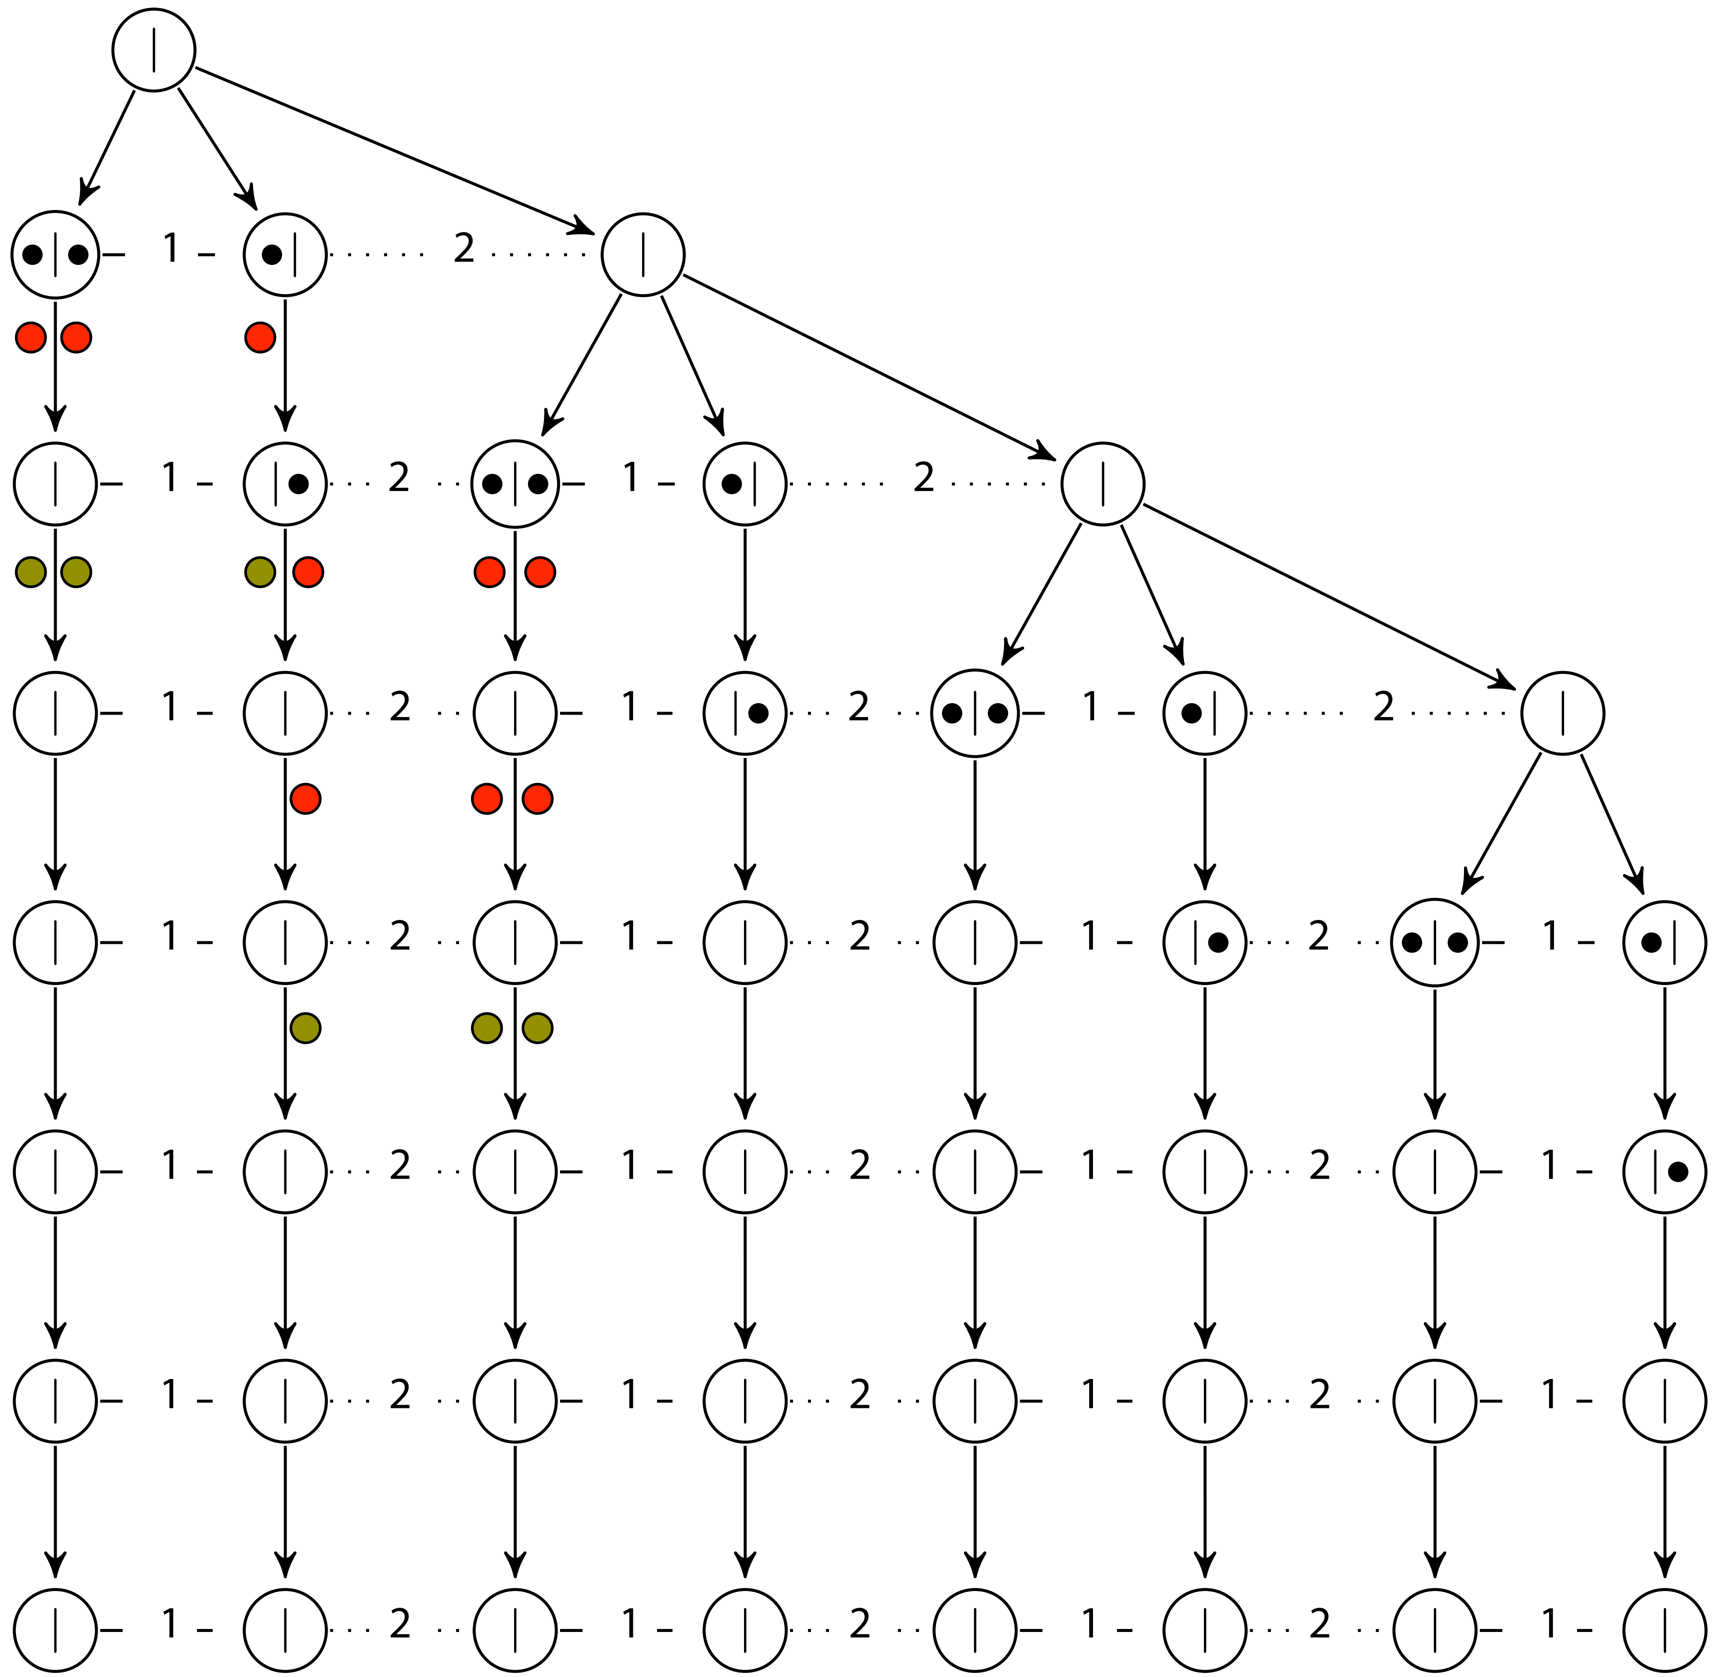

Extensive form

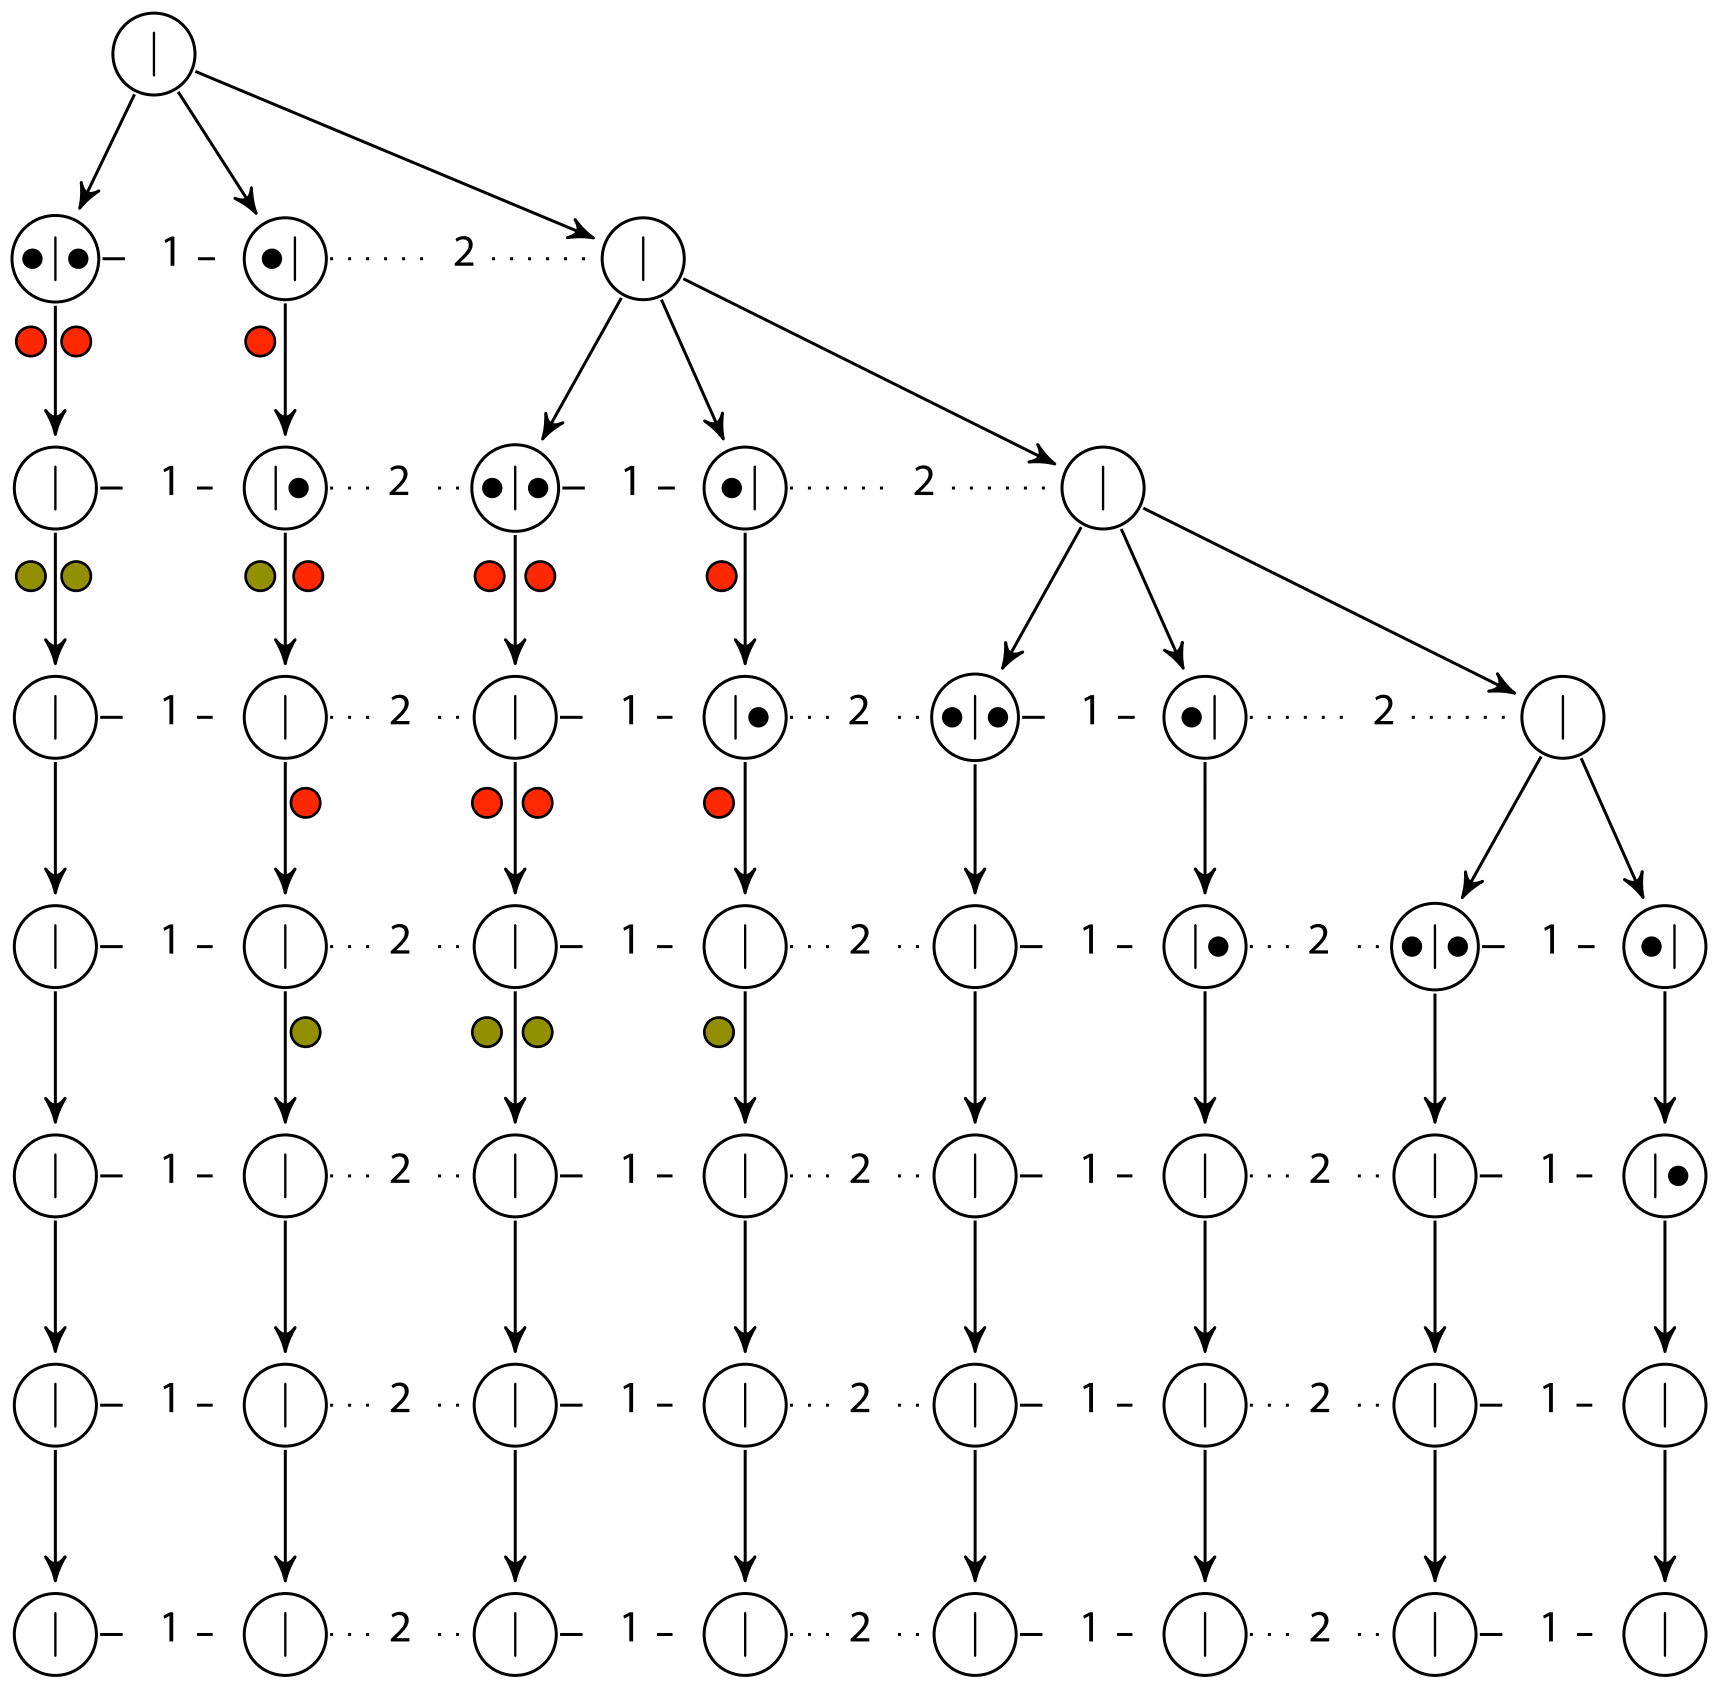

Extensive form

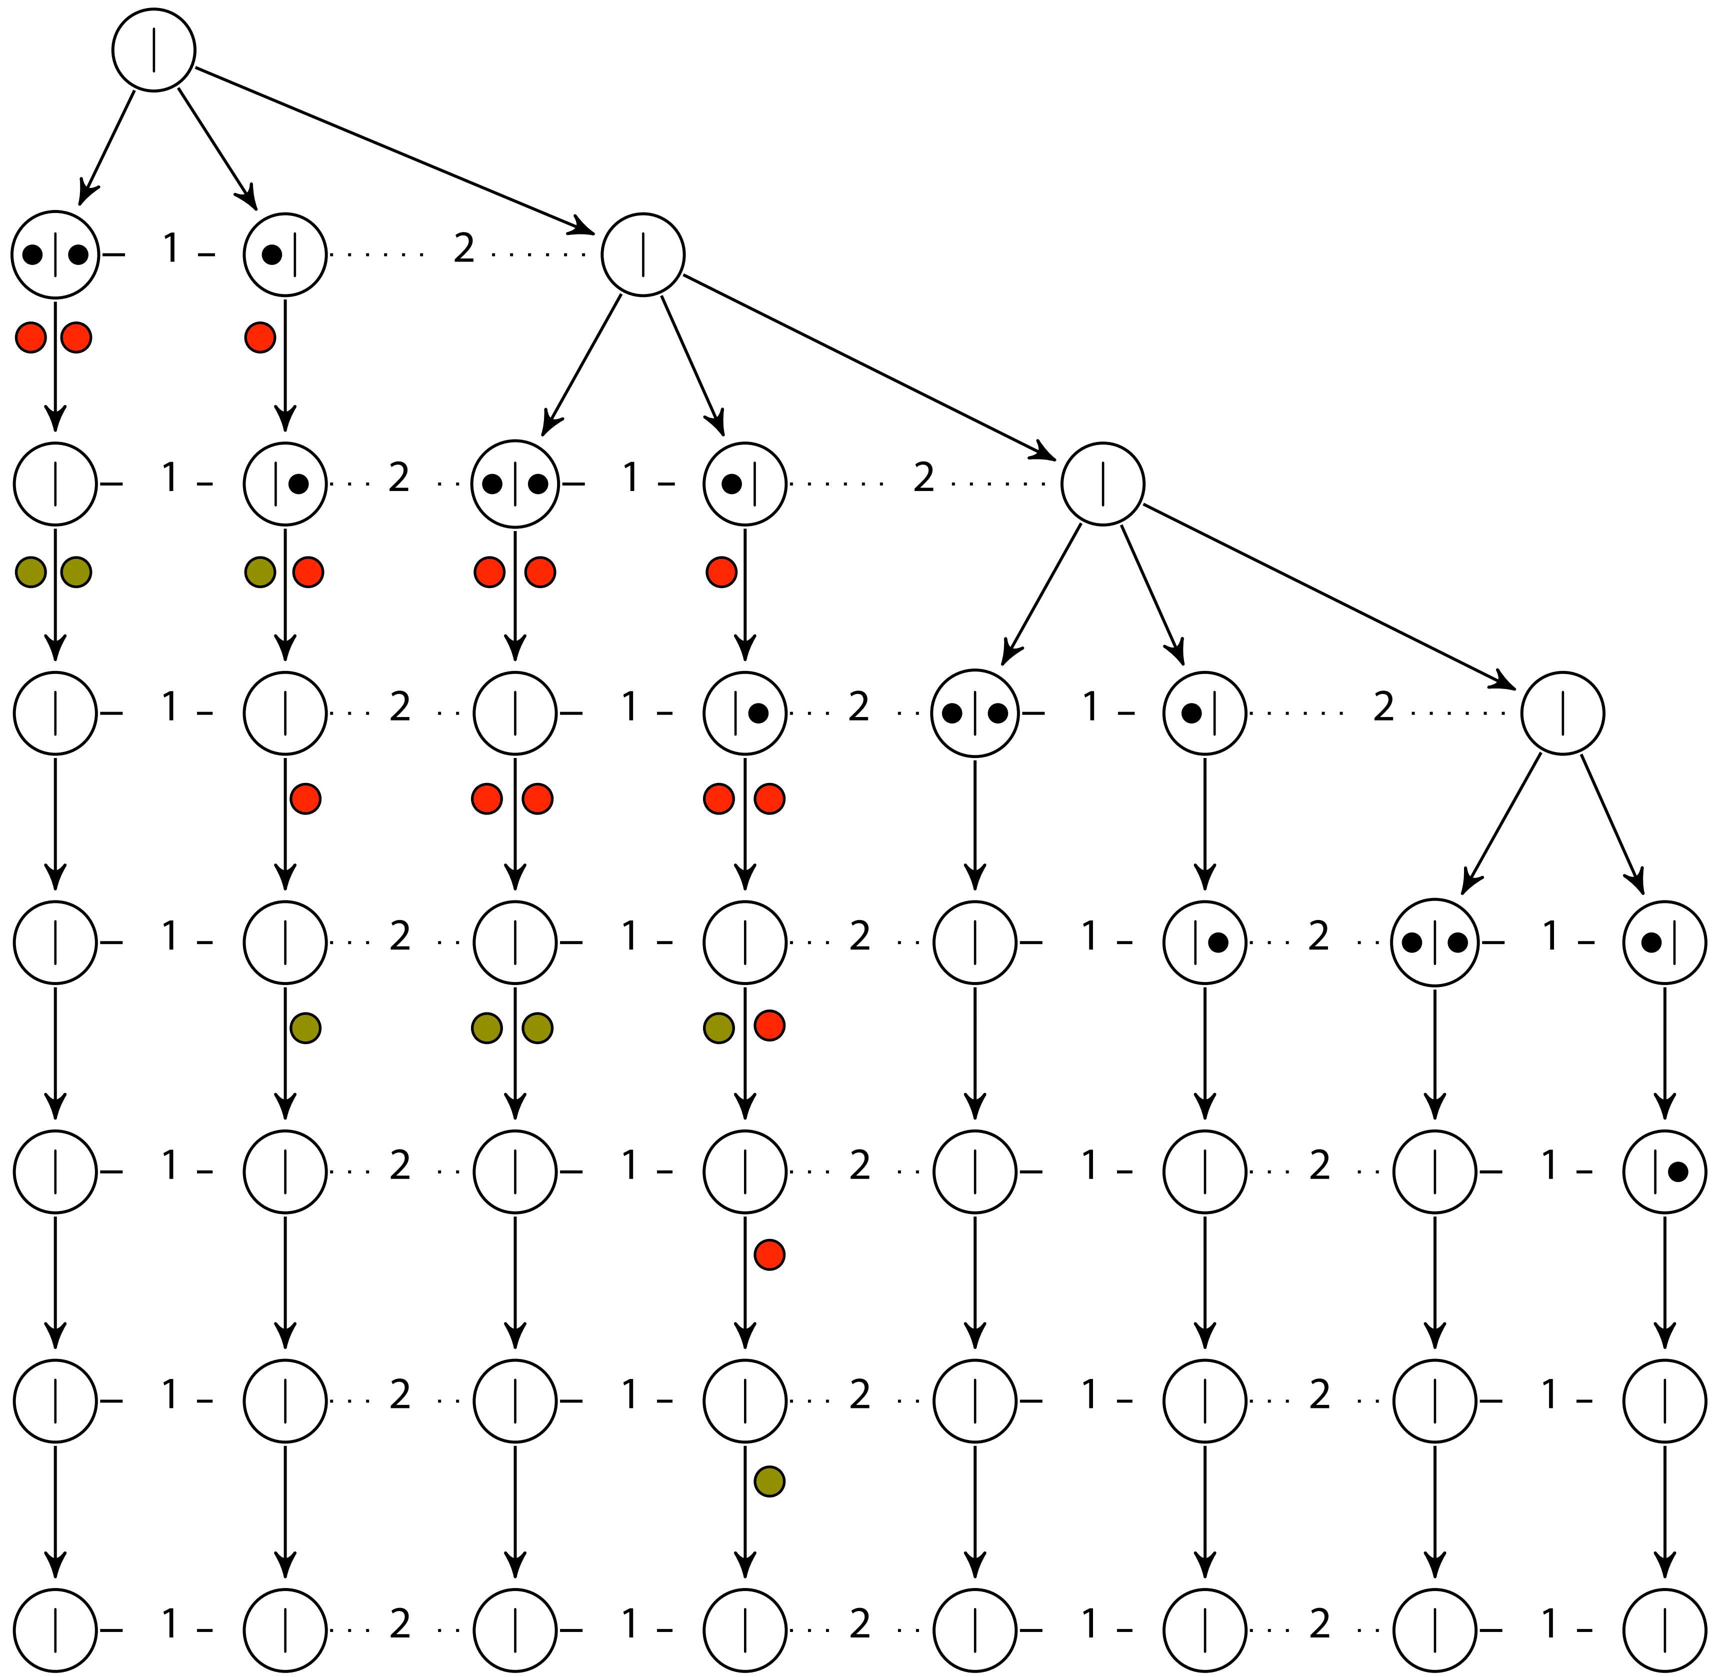

Extensive form

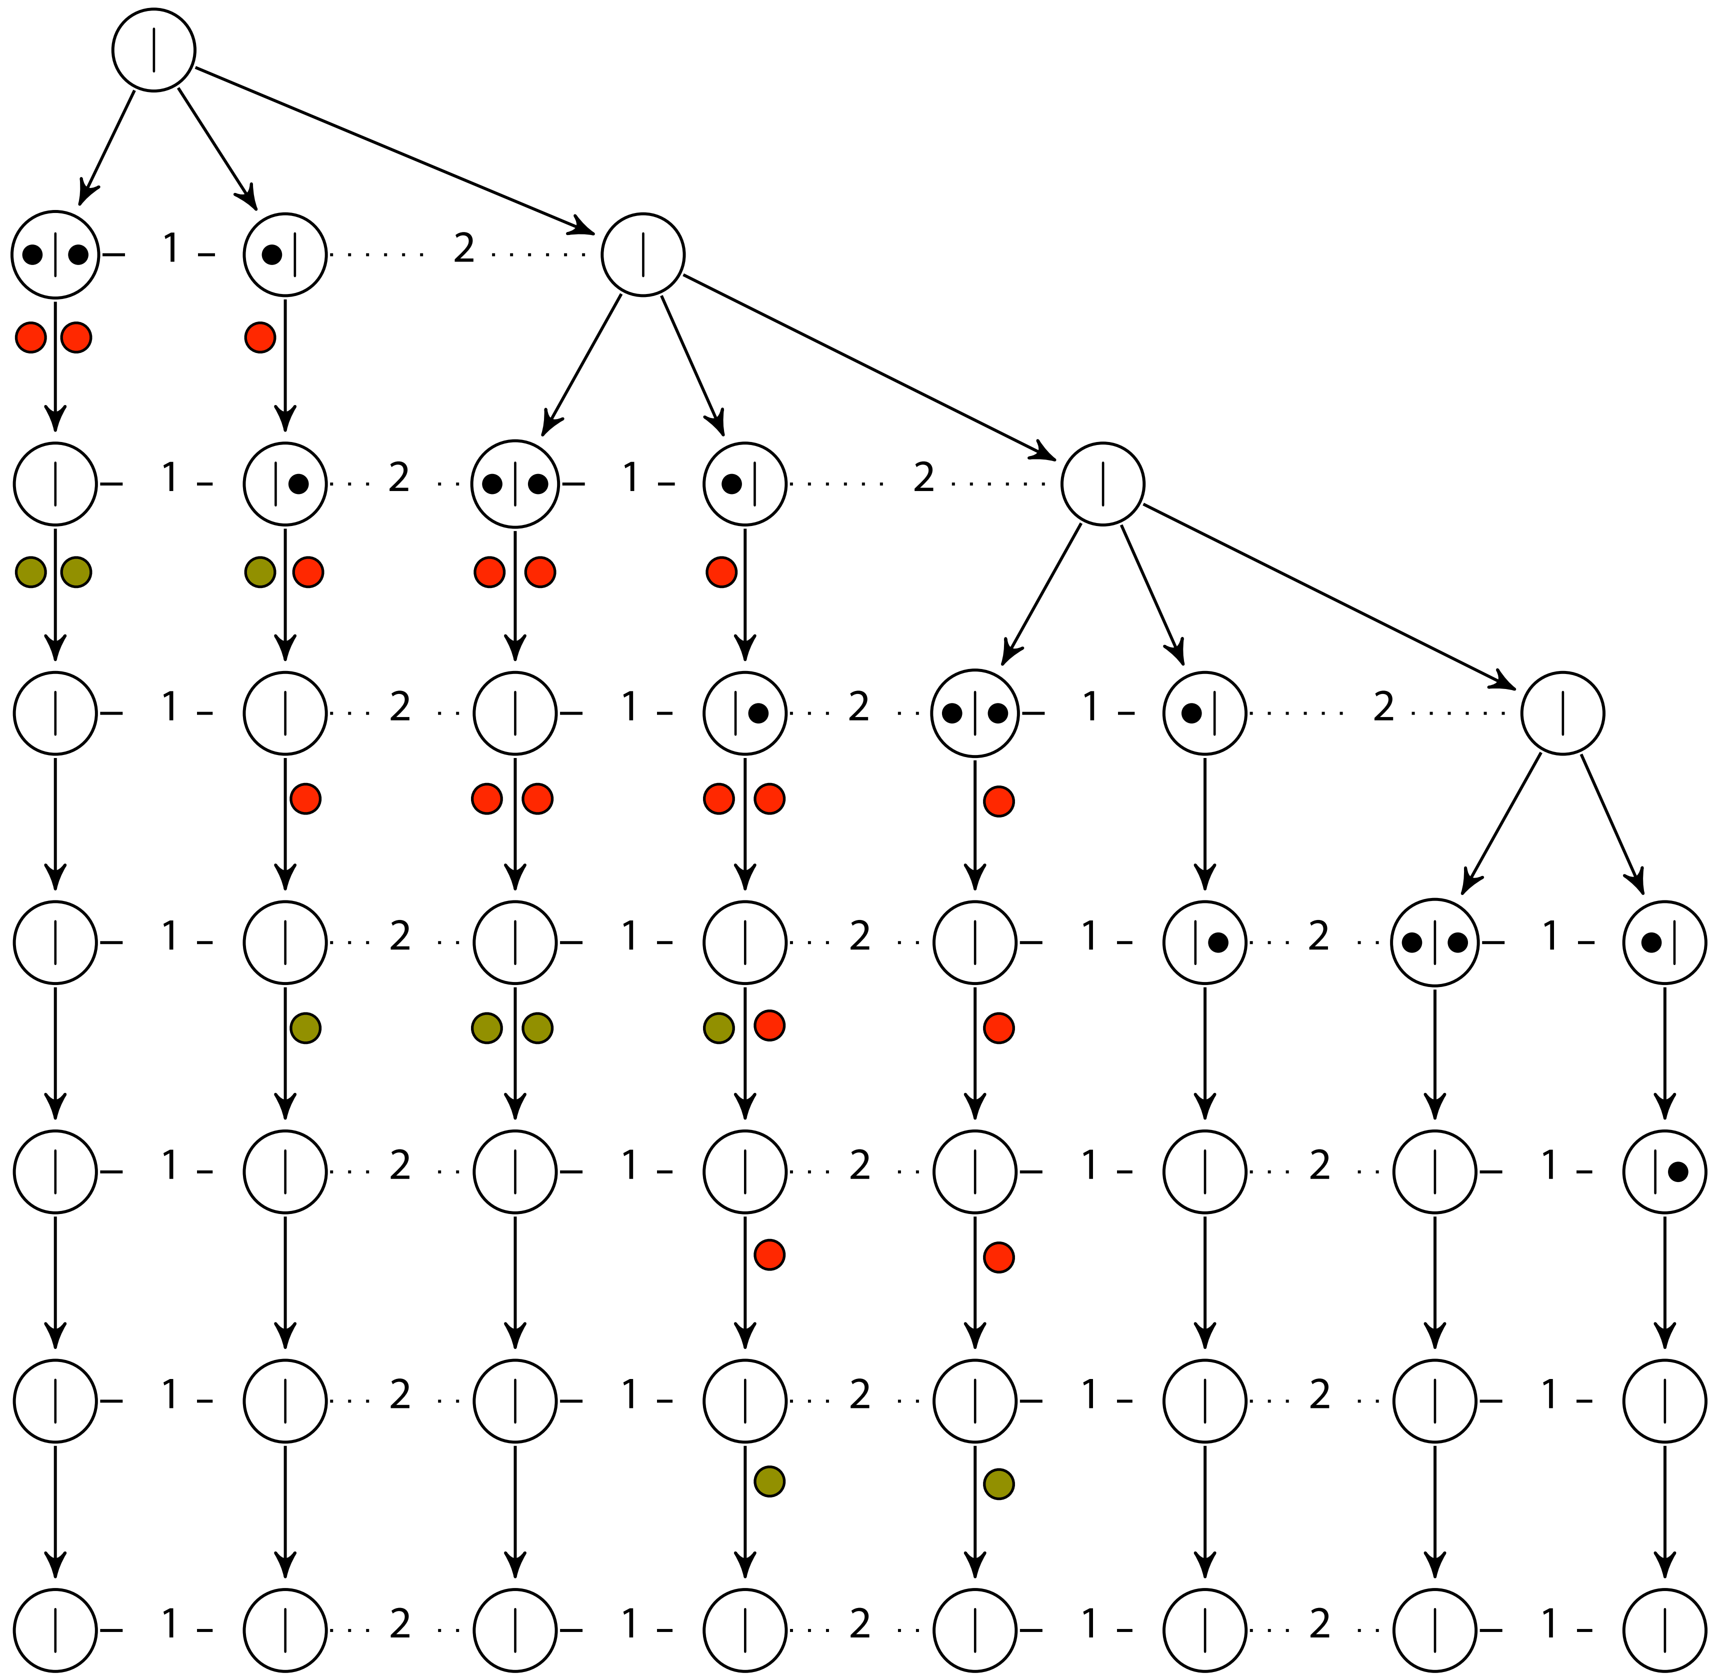

Extensive form

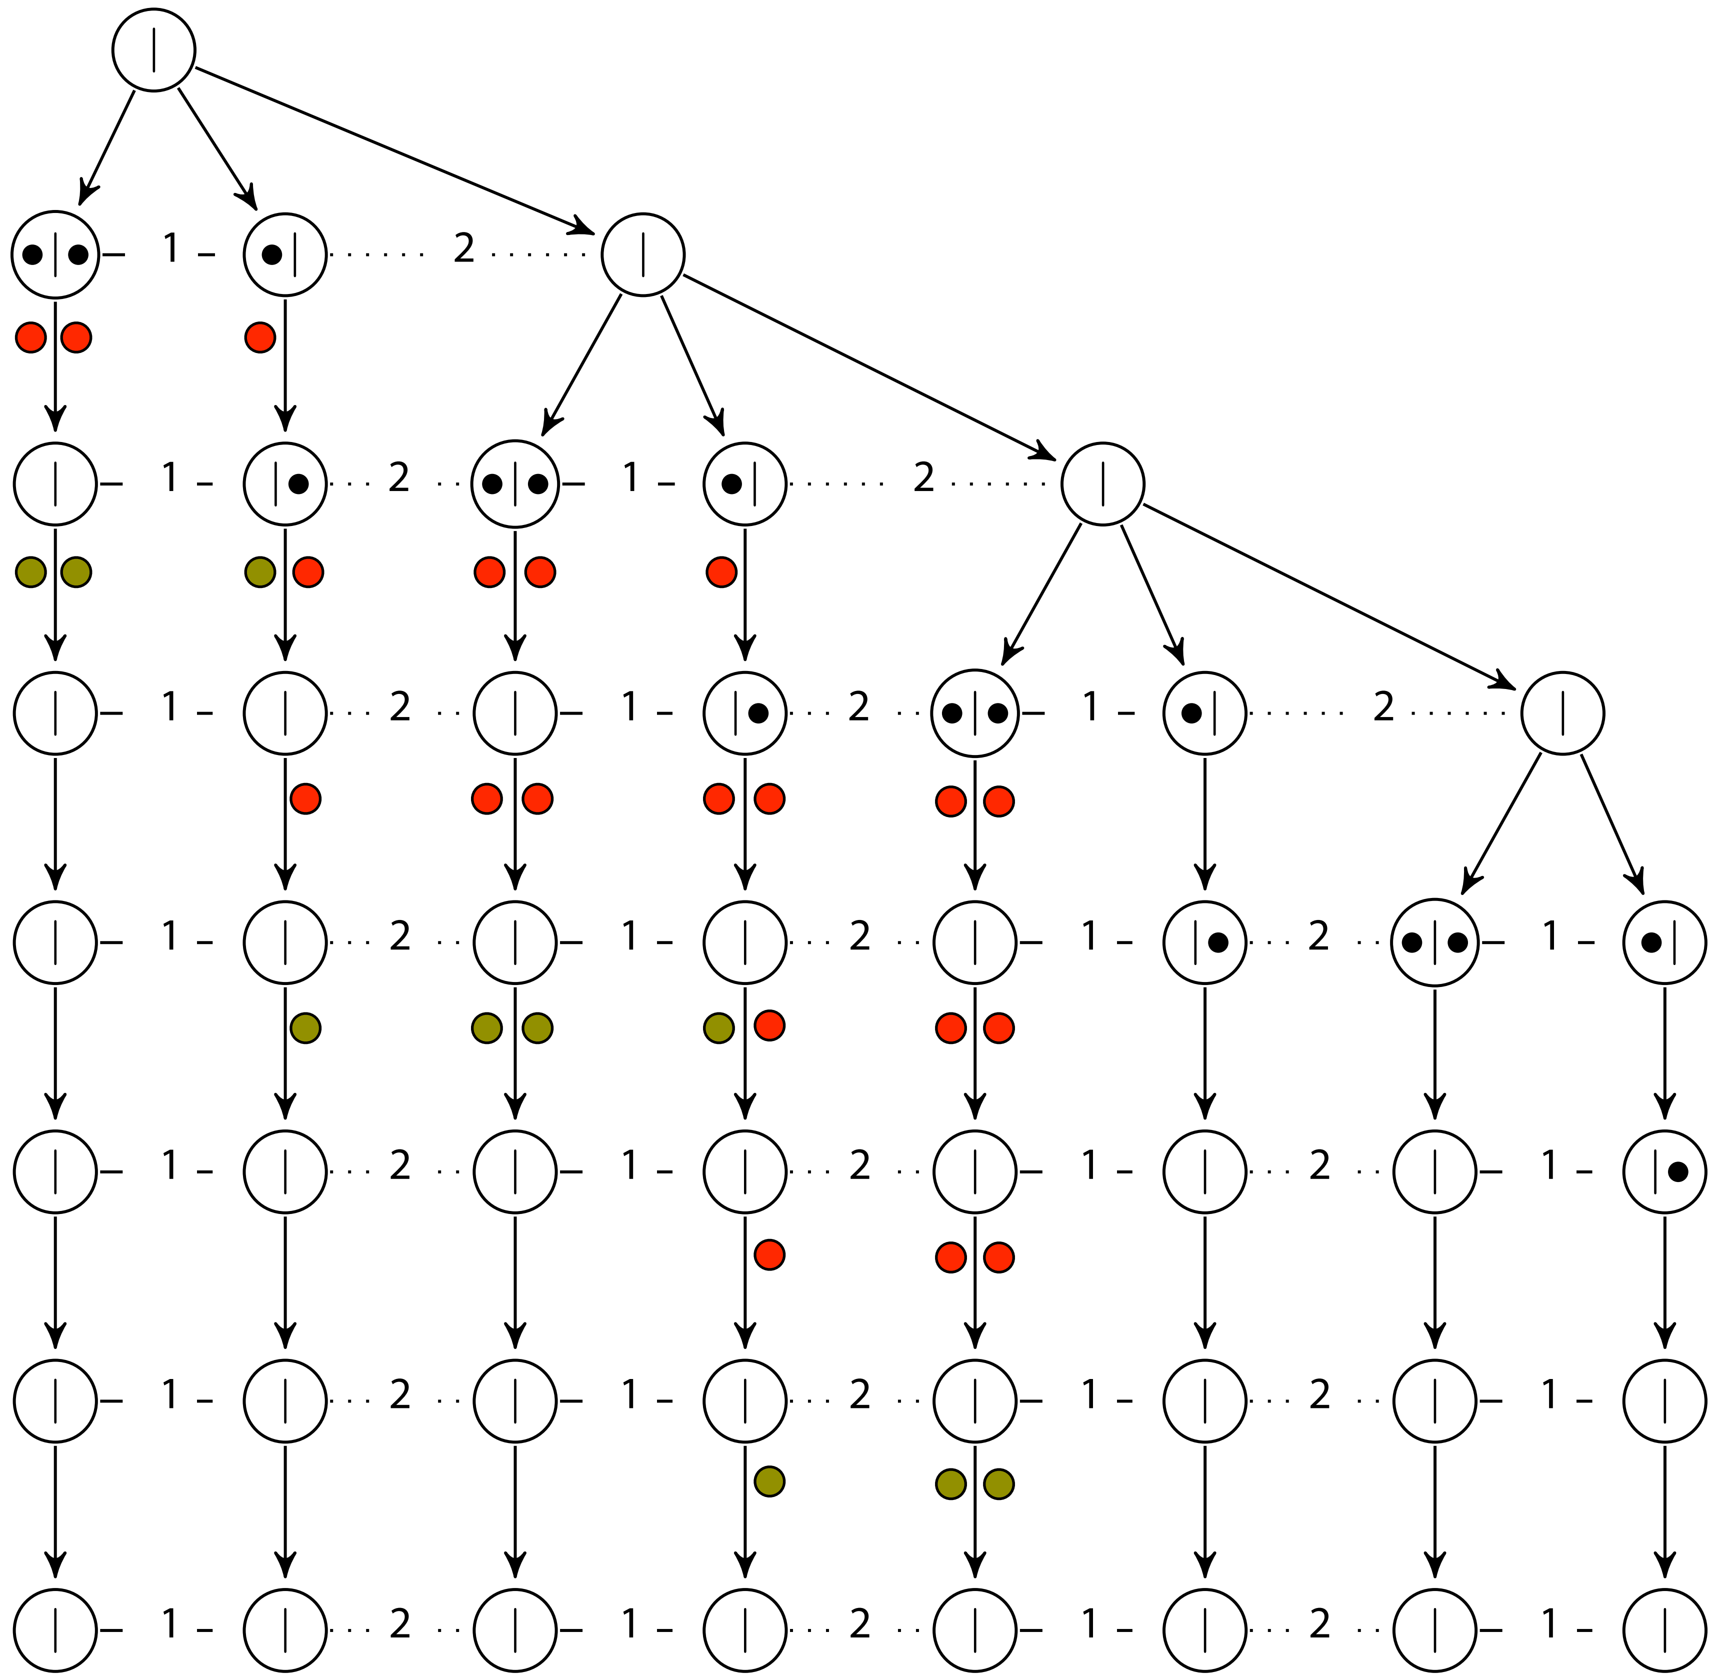

Extensive form

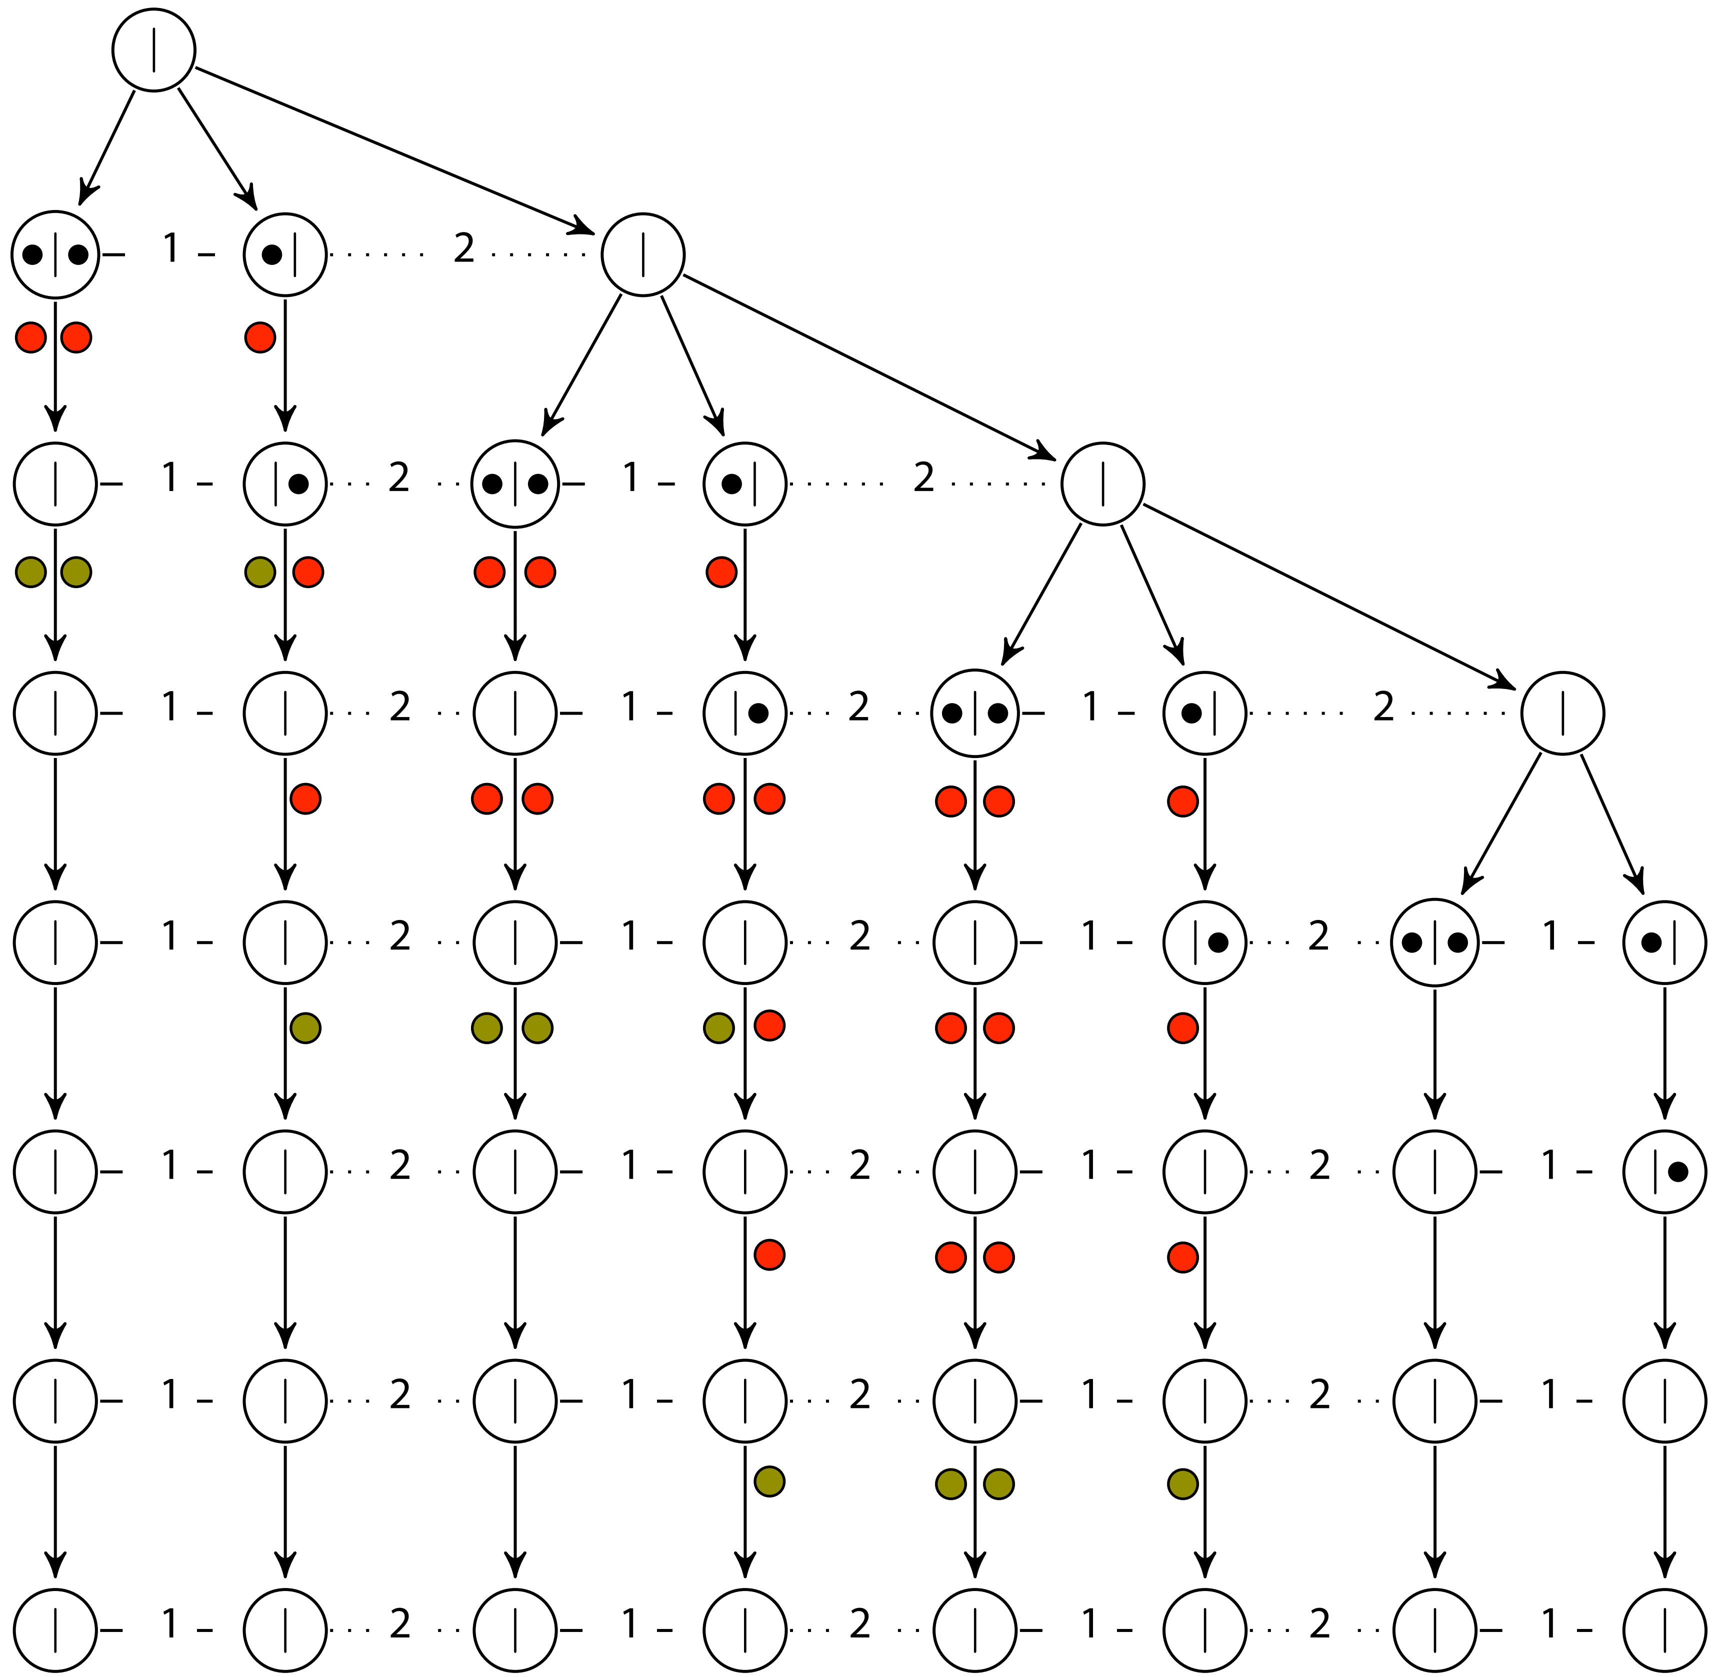

Extensive form

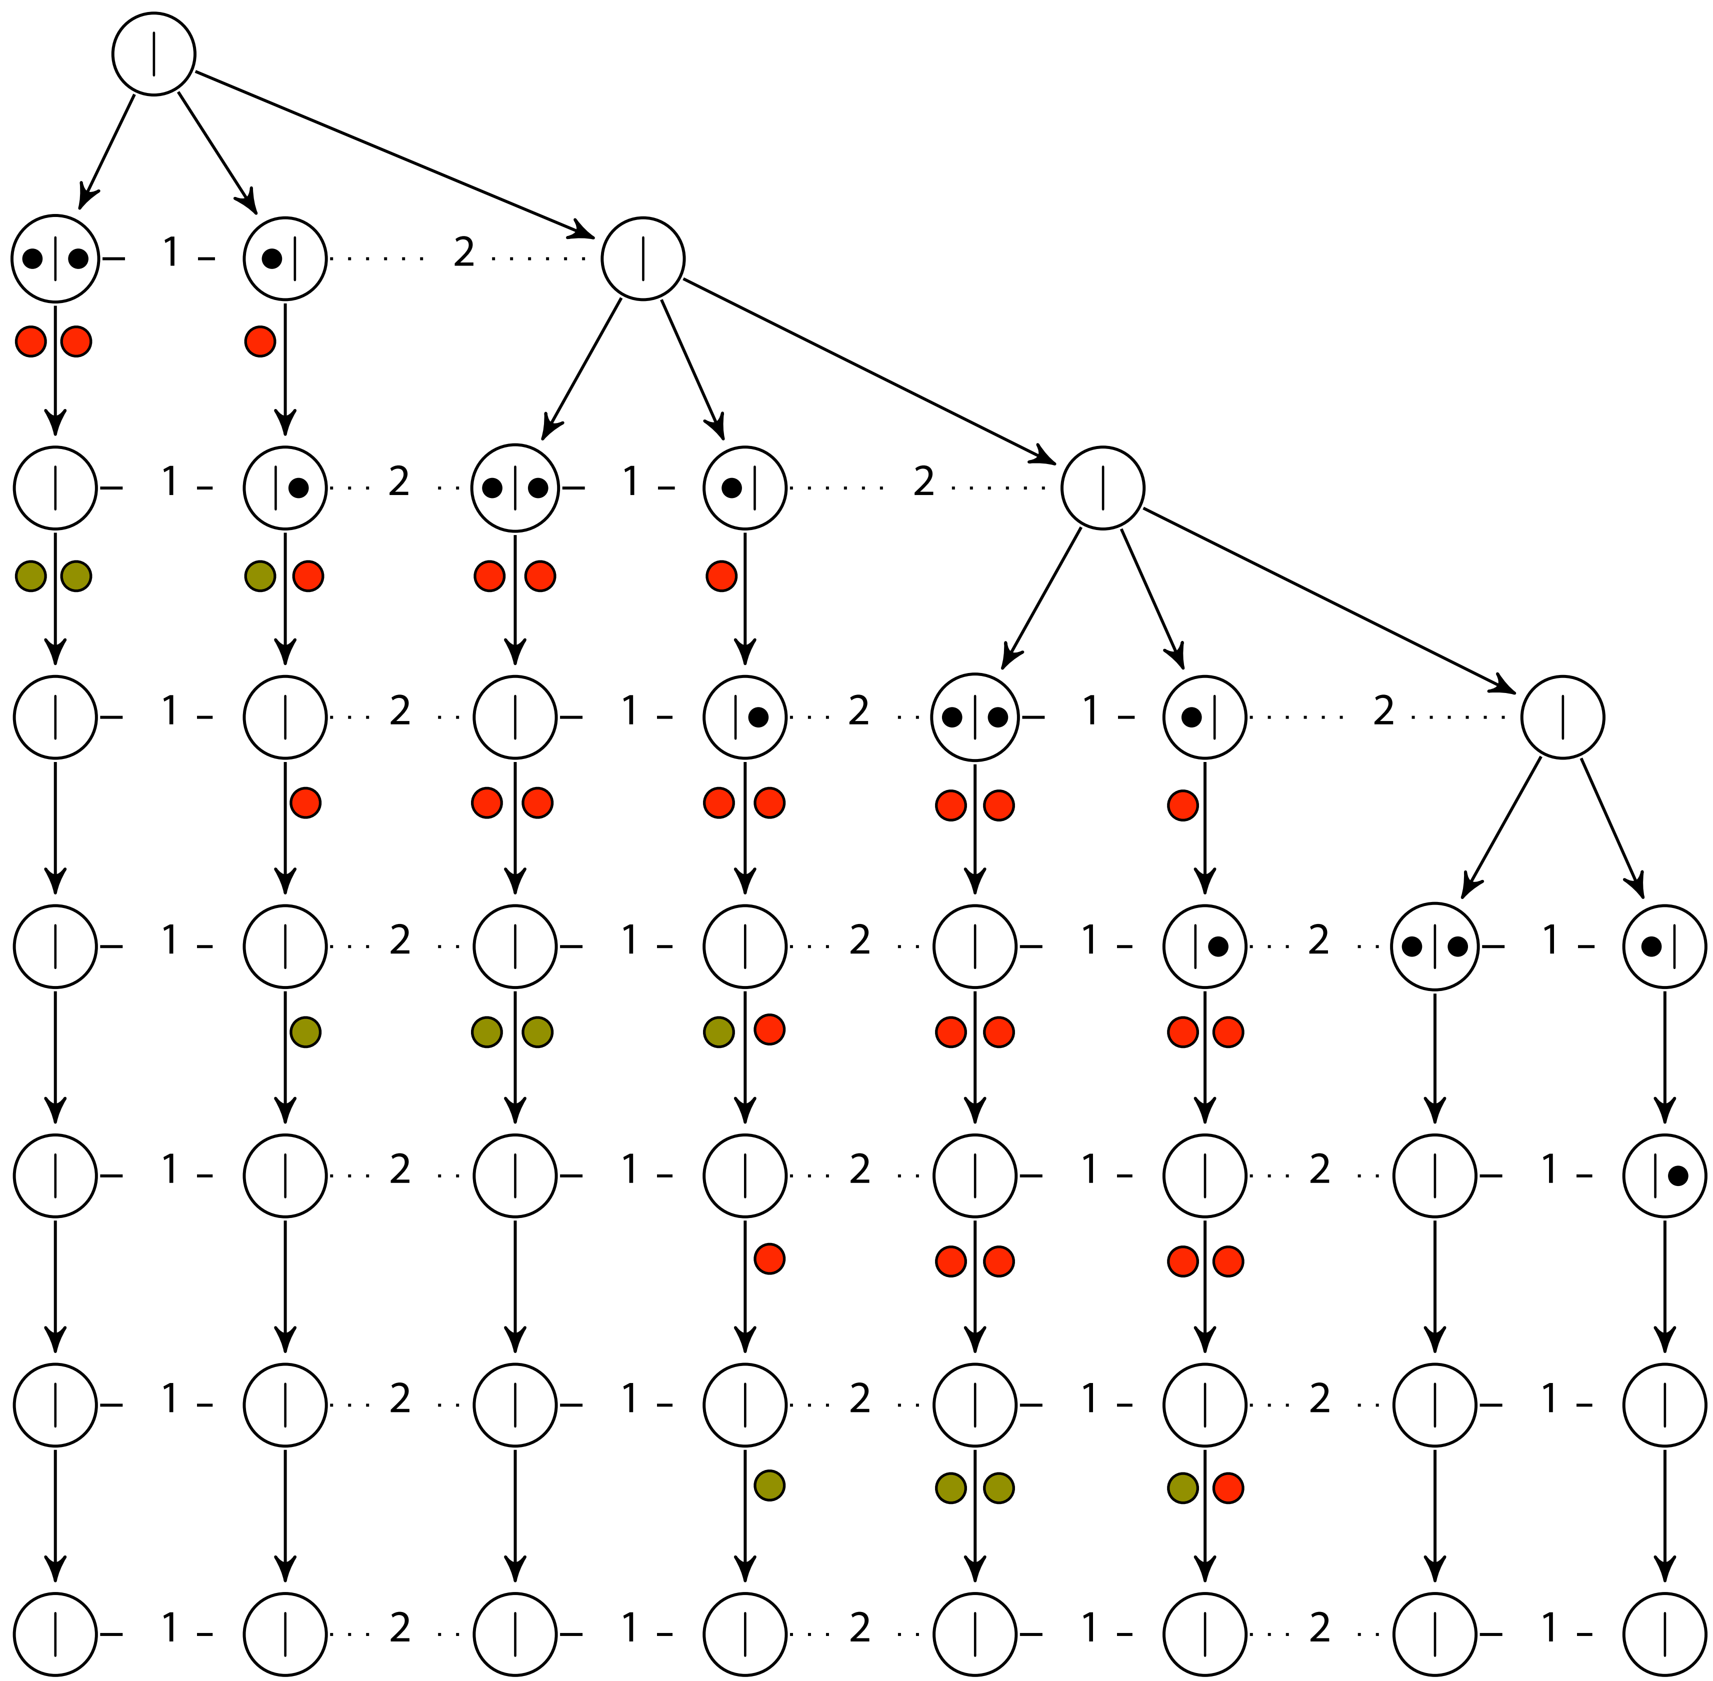

Extensive form

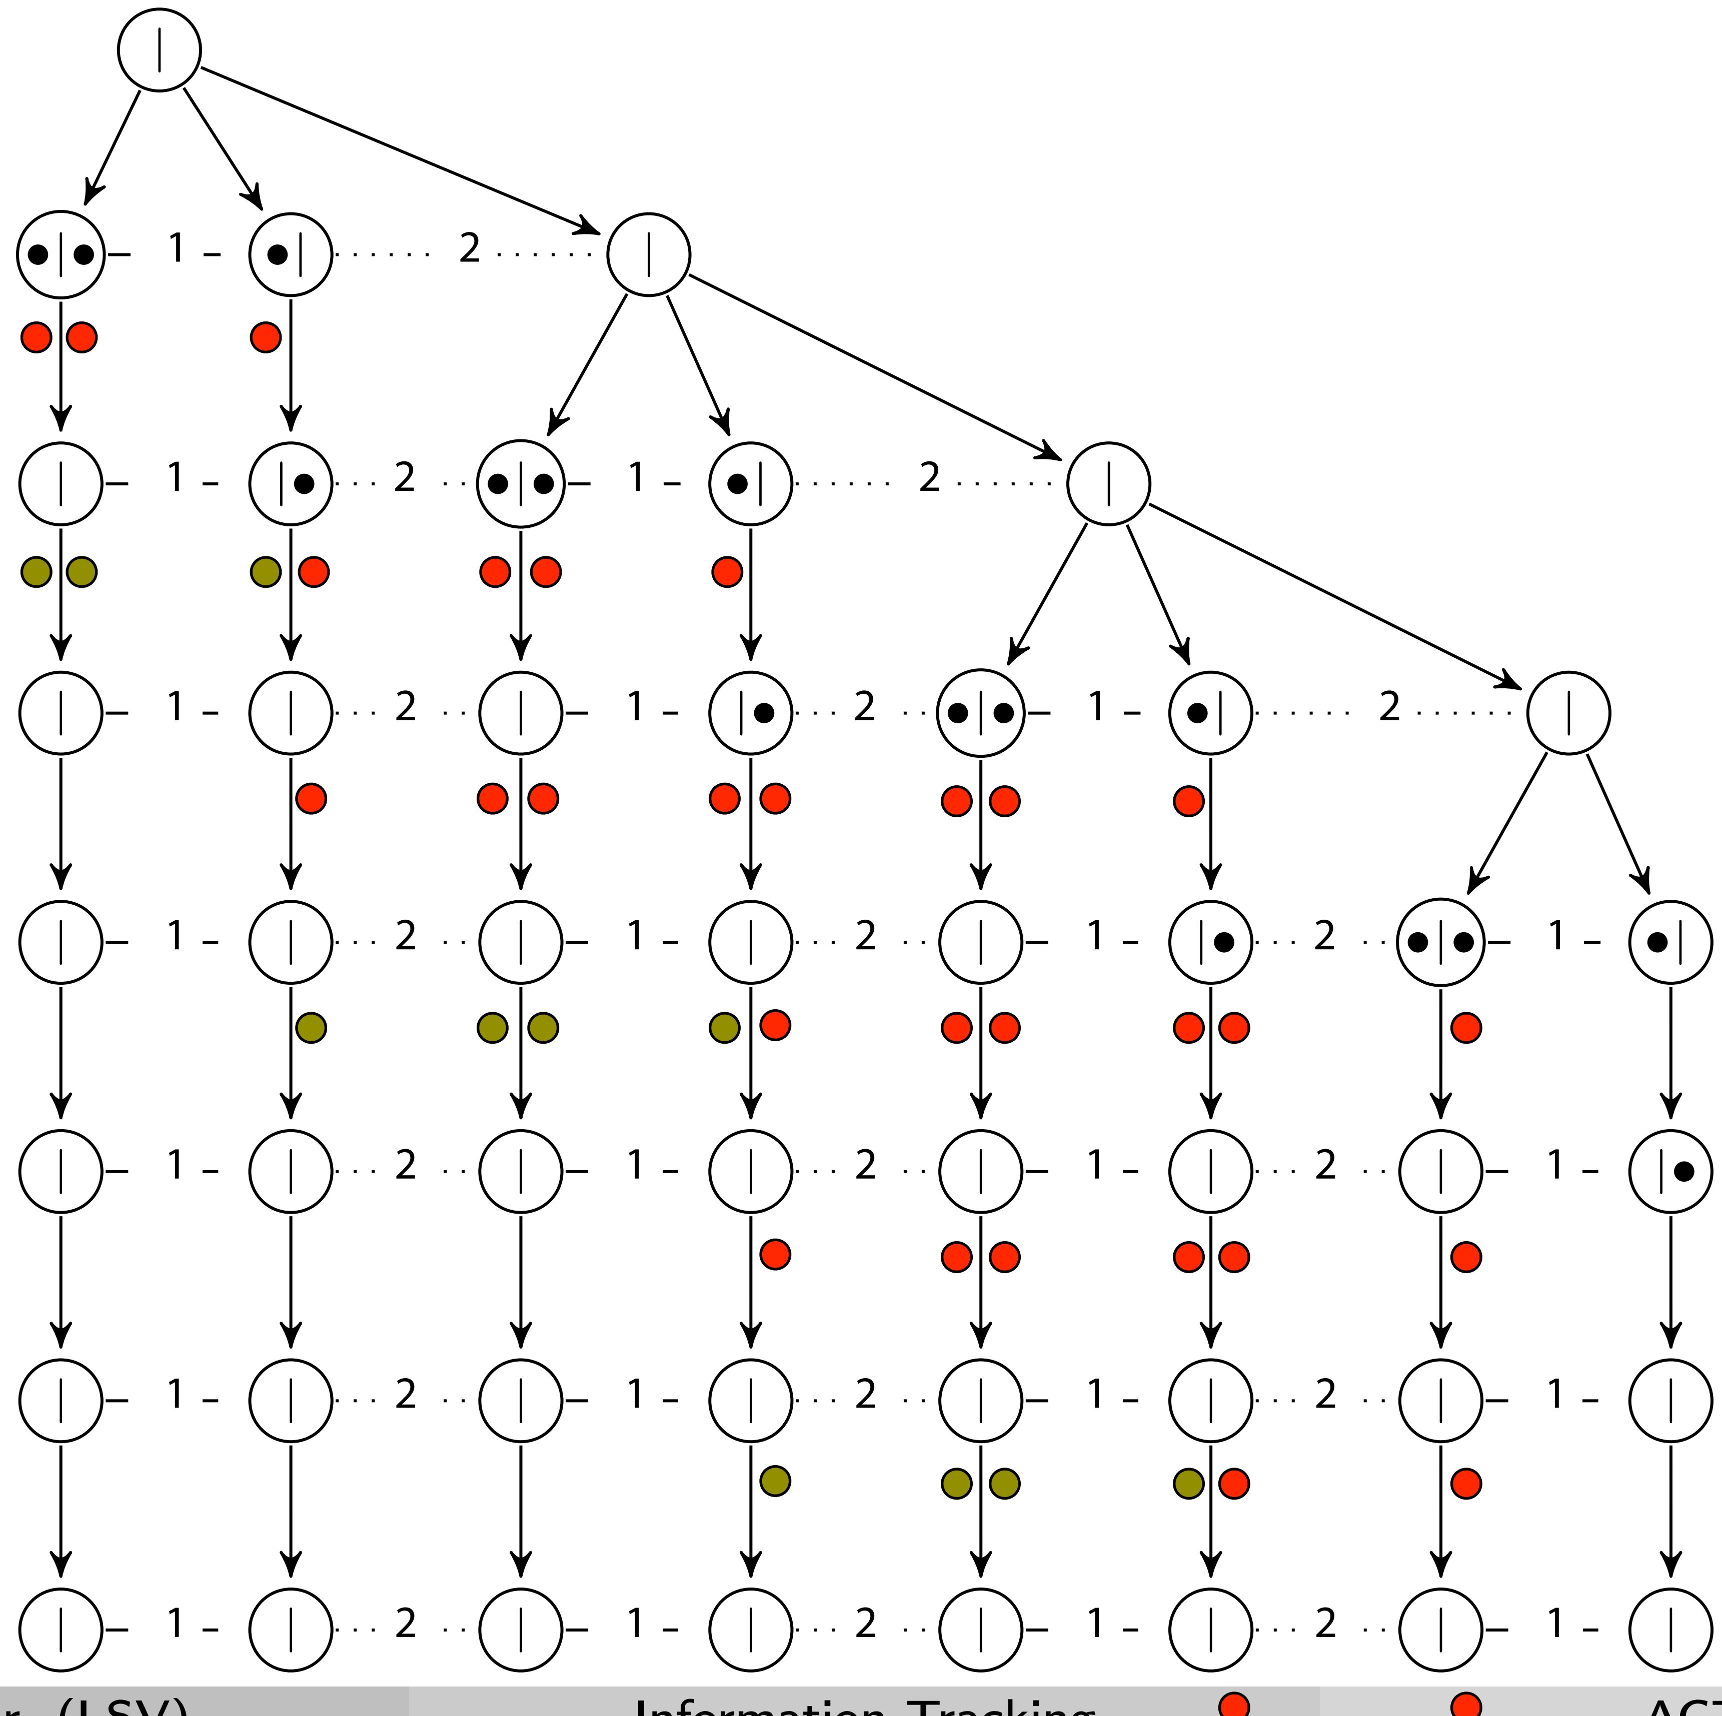

Extensive form

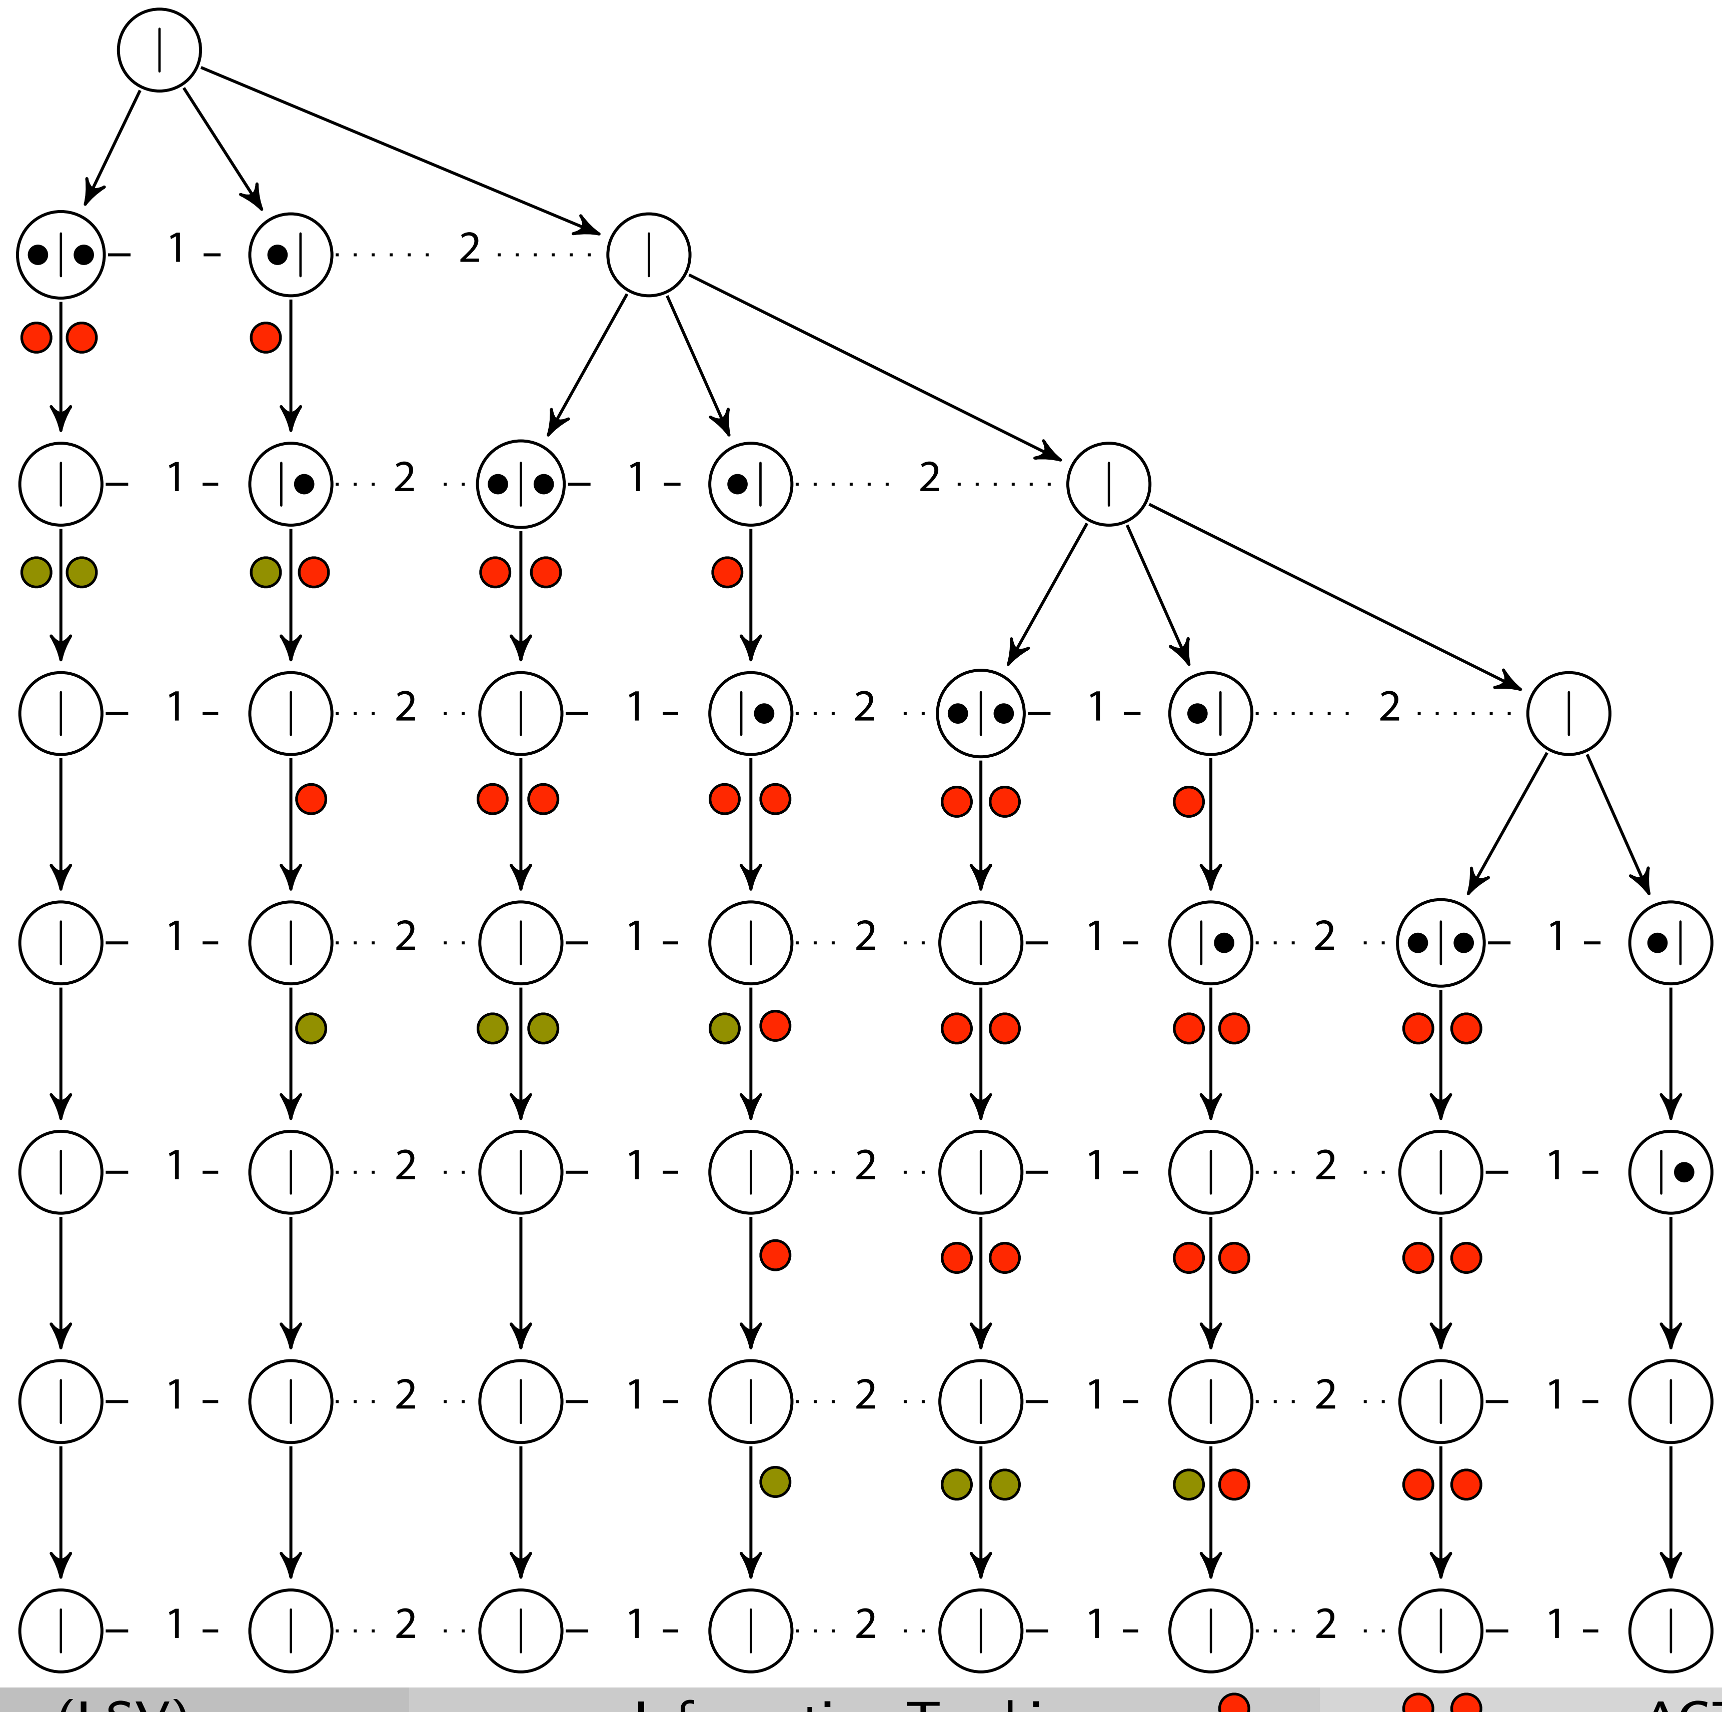

Undecidability

Game for a Turing machine $\mathcal{M} = (Q, \Sigma, q_0, \delta, F)$

Two players, action sequences should encode configurations $\in \Sigma^*(q, c)\Sigma^*$

- Actions $\Sigma \cup (Q \times \Sigma)$

- Observations $\circ, \bullet$

Game graph almost as before.

On $(\circ^n)\bullet$, produce configuration $s^i(n)$, such that $\begin{cases} s^1(1) = \text{Init} \\ s^0(n) = s^1(n) \\ s^1(n+1) \vdash s^0(n) \end{cases}$

- Final states unsafe: winning strategy, if machine never halts.
- Final states loop: finite-memory winnning strategy if machine halts.

Interim Summary

- There are finite games where the grand coalition can win, but not with a finite-memory strategy.
- The question of whether the winning coalition has a winning strategy is undecidable, even when restricted to finite-memory strategies.
- There are classes of finite games where grand coalition has finite-memory winning strategies, but their memory requirement cannot be bounded by any computable function.

Bisimulation

Bisimulation on a game graph $G = (V, \Delta, \beta^i, \gamma)$.

- relation $Z \subseteq V \times V$ such that, whenever $(v, v') \in Z$,

- $\gamma(v) = \gamma(v')$; $\beta^i(v) = \beta^i(v')$ for all i
- for all w with $v \xrightarrow{a} w$ there exists w' with $v' \xrightarrow{a} w'$ and $(w, w') \in Z$
- for all w' with $v' \xrightarrow{a} w'$ there exists w with $v \xrightarrow{a} w$ and $(w, w') \in Z$

Maximal bisimulation $\simeq$ -- equivalence, quotient $G/\simeq$.

The tracking of a game

Idea: Regard indistinguishability $\sim^i$ in extensive form as a (symmetric) edge relation $\overset{i}{\rightsquigarrow}$.

- Expand unravelling of G with $\overset{i}{\rightsquigarrow}$;
- Take maxmial bisimulation $\simeq$ on this expansion;
- Tracking: quotient of unravelling of G under $\simeq$:

$$\mathrm{Tr}(G) := \mathrm{Unr}(G)/\simeq$$

Remark.

- $\mathrm{Tr}(G)$ is bisimilar to G;
- the two games have the same extensive form

Main result

In every game with finite tracking
the grand coalition has a winning strategy
iff it has one with finite memory

Proof (1) - Knowledge equivalence in extensive form

In the extensive form:

- winning strategies need not distinguish between $\simeq$- bisimilar position.
- cannot distinguish between $\sim^i$-indistinguishable positions.

Take $\approx^i := (\sim^i \cup \simeq)^*$- transitive closure.

Lemma. If there exists a winning strategy profile, there also exists one s with $s^i(x) = s^i(y)$ whenever $x \approx^i y$.

Proof (2) - Projection to tracking

$\mathcal{K}^i$ = partition induced by $\approx^i$ on unravelling.

$K^i(\pi) \in \mathcal{K}^i$ class of initial play π

Lemma. $K^i(\pi)$ is positional in any tracked game:
if π, π' end in the same position, then $K^i(\pi) = K^i(\pi')$.

- Project K^i to $\mathrm{Tr}(G)$.

Corollary. In $\mathrm{Tr}(G, (K^i)_{i<n})$ the grand coalition has a memoryless (observation-based) strategy.

Proof (3) - automata for knowledge tracking

Lemma. For any game with finite tracking,
there exists an automaton that recognises $K^i(\pi)$
upon input of the action and observation sequence of player i.

States are $\approx^i$-classes; construction on-the-fly.

▶ If the coalition has a winning strategy on G, this automaton yields a finite-memory implementation.

Conclusion

- Semidecision algorithm for n-player games with imperfect information
- Explanation for some known solvable instances:
  - one player against environment;
  - players with hierarchical observations: $\beta^i(v) = \beta^i(v') \implies \beta^j(v) = \beta^j(v')$, for all $j > i$;
  - halting Turing machines
- The question whether the tracking of the game is finite is undecidable.

Outlook

Result is not yet tight.

Revision of classical impossibility notion:

- optimal strategies do exist and can be constructed for any finite stage;
- infinite memory may be simple (one-counter?)