Multi-Agent Systems

Material used

- Halpern: Reasoning about Uncertainty. Chapter 6-7

1 Epistemic frames
2 Modal epistemic logic
3 Epistemic probability frames
4 The dynamics of multi-agent systems
5 Protocols
6 Using protocols to specify situations

0 Introduction

- Up to now two simplifying assumptions were made
  - Single agents
  - Only static situations (the agents knowledge are independent of the world where the agent lives)
- For modelling interactive situations a more natural framework is required, especially in situations where
  - agents are bargaining
  - playing a game
  - performing a distributed computation
  - performing a conversation

1 Epistemic frames

- In the first chapter we introduced the concept an *epistemic space*: If $W^0 \subseteq W$, then (W, W^0) is called an epistemic space.
- Now we assume that the proposition W^0 can depend on the worlds $w \in W^0$. This leads to the conception of an epistemic frame
- Hence, the knowledge/belief of an agent becomes dependent of the world the agent lives in.
- There are various constraints that restrict the dependencies between the worlds and the agent's beliefs in that world

- Let $\underline{\boldsymbol{K}} \subseteq W \times W$ be a binary relation on W. Then the pair M = $(W, \underline{\boldsymbol{K}})$ is called an epistemic frame. The relation $\underline{\boldsymbol{K}}$ is called an accessibility relation. $(u,v) \in \underline{\boldsymbol{K}}$ says that the agent considers v possible in the world u.

- Define $\underline{\boldsymbol{K}}(u) = \{v \in W: (u,v) \in \boldsymbol{K}\}$

  $\underline{\boldsymbol{K}}$(u) is the set of worlds that the agent considers possible in world u.

- A *situated* epistemic frame is a pair **(***M***,** *w***)** where *M* is an epistemic frame and $w \in W$.

Old Definition 1: Let (W, W^0) be an epistemic space:

- $(W, W^0) \Vdash \text{Possible}_x(U)$ iff $U \cap W^0 \neq \emptyset$ (x considers U possible)
- $(W, W^0) \Vdash \text{Know}_x(U)$ iff $W^0 \subseteq U$ (x knows U)

New Definition 1: Let $(W, \underline{\mathbf{K}}, w)$ be a situated epistemic frame:

- $(W, \underline{\mathbf{K}}, w) \Vdash \text{Possible}_x(U)$ iff $U \cap \underline{\mathbf{K}}(w) \neq \emptyset$
- $(W, \underline{\mathbf{K}}, w) \Vdash \text{Know}_x(U)$ iff $\underline{\mathbf{K}}(w) \subseteq U$

Definition 2: Extending the elementary propositional language by adding sentences of form $P\phi$ and $K\phi$:

- $(W, \underline{\mathbf{K}}, w) \models P\phi$ iff $(W, \underline{\mathbf{K}}, w') \models \phi$ for some $w' \in \underline{\mathbf{K}}(w)$
- $(W, \underline{\mathbf{K}}, w) \models K\phi$ iff $(W, \underline{\mathbf{K}}, w') \models \phi$ for all $w' \in \underline{\mathbf{K}}(w)$

Reflexive	$(u,u)\in \underline{\boldsymbol{K}}$	$u\in \underline{\boldsymbol{K}}(u)$
Transitive	$(u,v)\in \underline{\boldsymbol{K}}$ & $(v,w)\in \underline{\boldsymbol{K}}$ $\Rightarrow (u,w)\in \underline{\boldsymbol{K}}$	$v\in \underline{\boldsymbol{K}}(u) \Rightarrow$ $\underline{\boldsymbol{K}}(v) \subseteq \underline{\boldsymbol{K}}(u)$
Euclidian	$(u,v)\in \underline{\boldsymbol{K}}$ & $(u,w)\in \underline{\boldsymbol{K}}$ $\Rightarrow (v,w)\in \boldsymbol{K}$	$v\in \underline{\boldsymbol{K}}(u) \Rightarrow$ $\underline{\boldsymbol{K}}(v) \supseteq \underline{\boldsymbol{K}}(u)$
Symmetric	$(u,v)\in \underline{\boldsymbol{K}} \Rightarrow (v,u)\in \underline{\boldsymbol{K}}$	$v\in \underline{\boldsymbol{K}}(u) \Rightarrow u\in \underline{\boldsymbol{K}}(v)$
Serial	For each *w* there is some *w*' such that $(w,w')\in \underline{\boldsymbol{K}}$	$\underline{\boldsymbol{K}}(w) \neq \varnothing$
Equivalence relation	Reflexive, symmetric, transitive	$u\in \underline{\boldsymbol{K}}(u)$; $v\in \underline{\boldsymbol{K}}(u) \Rightarrow \underline{\boldsymbol{K}}(v)=\underline{\boldsymbol{K}}(u)$

- Take $\underline{\boldsymbol{K}}$ as a binary relation and show the following:
  a. $\underline{\boldsymbol{K}}$ is reflexive iff $u \in \underline{\boldsymbol{K}}(u)$ for all worlds u
  b. $\underline{\boldsymbol{K}}$ is transitive iff $v \in \underline{\boldsymbol{K}}(u) \Rightarrow \underline{\boldsymbol{K}}(v) \subseteq \underline{\boldsymbol{K}}(u)$ for all worlds u,v
  c. $\underline{\boldsymbol{K}}$ is Euclidian iff $v \in \underline{\boldsymbol{K}}(u) \Rightarrow \underline{\boldsymbol{K}}(v) \supseteq \underline{\boldsymbol{K}}(u)$ for all worlds u,v

- Assume $\underline{\boldsymbol{K}}$ is reflexive. Show the following:
  $(W, \underline{\boldsymbol{K}}, w) \models K\phi$ implicates $(W, \underline{\boldsymbol{K}}, w) \models \phi$ for all w

- Assume the following clause: $(W, \underline{\boldsymbol{K}}, w) \models \phi$ implicates $(W, \underline{\boldsymbol{K}}, w) \models K\phi$ for all w. Formulate a condition for the relation $\underline{\boldsymbol{K}}$ such that this clause becomes true. Do the same for the clause $(W, \underline{\boldsymbol{K}}, w) \models \phi$ implicates $(W, \underline{\boldsymbol{K}}, w) \models P\phi$.

Definition 3: An epistemic frame for n agents is a tuple $(W, \underline{\boldsymbol{K}}_1, \ldots \underline{\boldsymbol{K}}_n)$ where each $\underline{\boldsymbol{K}}_i$ is a binary relation on W.

A situated epistemic frame for n agents is a tuple $(W, \underline{\boldsymbol{K}}_1, \ldots \underline{\boldsymbol{K}}_n, w)$ with $w \in W$

- In general, different agents will consider different worlds possible, that means $\underline{\boldsymbol{K}}_i(w) \neq \underline{\boldsymbol{K}}_j(w)$ for $i \neq j$.
- One of the advantages of an epistemic frame is that it can be viewed as a labelled graph. The nodes are the worlds in W and there is an edge from u to v labeled i iff $(u,v) \in \underline{\boldsymbol{K}}_{\mathbf{i}}$, i.e. $v \in \underline{\boldsymbol{K}}_{\mathbf{i}}(u)$: $u \xrightarrow{i} v$

Example

Suppose that a deck consists of three cards labelled A, B, and C. Agent 1 and 2 each get one of these cards; the third card is left face down. A possible world is describing the cards held by each agent. Give a description of the epistemic situation!

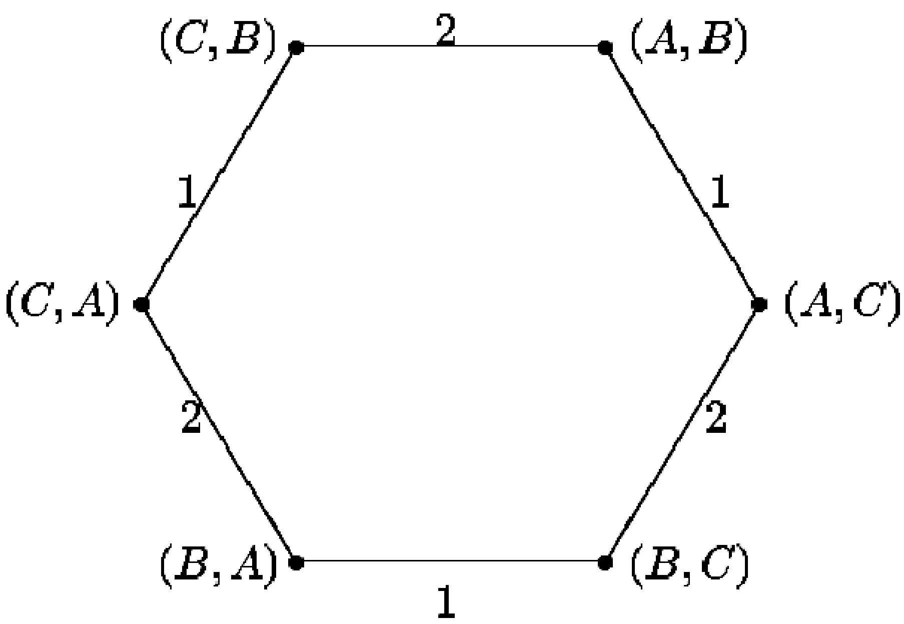

- 6 possible worlds
- in the world (A,B) agent 1 thinks two worlds are possible: (A,B) and (A,C)
- That means agent 1 knows that he has card A but considers it possible that agent 2 could hold either card B or card C.
- The relation $\boldsymbol{K}_1$ and $\boldsymbol{K}_2$ are equivalence relations (loops and arrows on edges are omitted)

2 Modal epistemic logic

Perhaps the simplest kind of reasoning about uncertainty involves reasoning about whether certain situations are possible or impossible. The following provides a logic of knowledge that allows just this kind of reasoning.

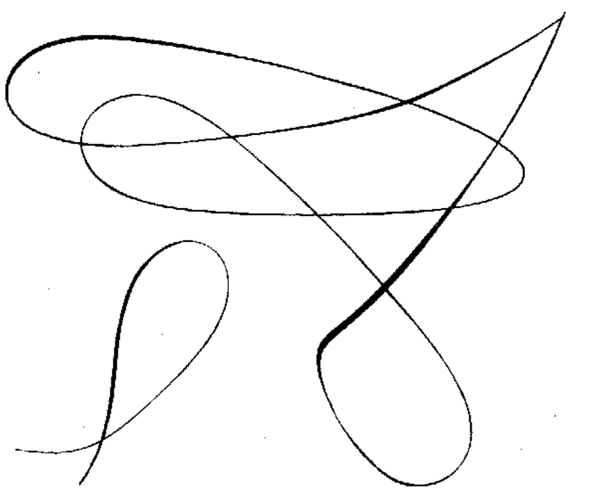

Syntax of propositional epistemic logic

Giving a set *At* of primitive (atomic) formula, the language $\mathscr{L}_n^K(At)$ is formed with the help of modal operators $K_{\mathbf{1}}, \ldots K_{\mathbf{n}}$, one for each agent. Formulas are formed by starting with primitive (propositional) formulas and closing off under negation and conjunction and the application of modal operators $K_{\mathbf{i}}$, so that if ϕ is a formula, so is $K_{\mathbf{i}}\phi$.

- $K_1K_2\,p\wedge\neg K_2K_1K_2\,p$ is a well-formed formula of $\mathscr{L}_n^K(At)$
- $K_1K_2\,pq$, $K_1\vee K_2\,p\rightarrow q$ are not
- Important definition: $P_{\mathbf{i}}\phi =_{def} \neg K_{\mathbf{i}}\neg\phi$

Definition 4: Let $(W, \underline{K}_1, \ldots \underline{K}_n)$ be an epistemic frame for n agents, and π be an interpretation -- assigning propositions to the atomic formulas of the language $\mathscr{L}_n^K(At)$. Then $M = (W, \underline{K}_1, \ldots \underline{K}_n, \pi)$ is called an epistemic structure.

a. $(M, w) \models p$ (for $p \in At$) iff $w \in \pi(p)$
b. $(M, w) \models \phi \wedge \psi$ iff $(M, w) \models \phi$ and $(M, w) \models \psi$
c. $(M, w) \models \neg\phi$ iff $(M, w) \not\models \phi$
d. $(M, w) \models K_i\phi$ iff $(M, w') \models \phi$ for all $w' \in \underline{\boldsymbol{K}}_i(w)$
e. $(M, w) \models P_i\phi$ iff $(M, w') \models \phi$ for some $w' \in \underline{\boldsymbol{K}}_i(w)$

A formula ϕ is considered *valid in an epistemic structure M*, denoted $M \models \phi$ iff $(M, w) \models \phi$ for all w in M.

Fact 1: Suppose $M = (W, \underline{K}_1, \ldots \underline{K}_n, \pi)$ is an epistemic structure

a. $M \models (K_i\phi \wedge K_i(\phi \rightarrow \psi)) \rightarrow K_i\psi$ (distribution axiom)
b. if $M \models \phi$ then $M \models K_i\phi$ (rule of knowledge generalization)
c. if $\underline{K}_i$ is transitive then $M \models (K_i\phi \rightarrow K_i K_i \phi)$ (positive IA*)
d. if $\underline{K}_i$ is Euclidian then $M \models (P_i\phi \rightarrow K_i P_i \phi)$ (negative IA*)
e. if $\underline{K}_i$ is serial then $M \models \neg K_i$ *false* (Consistency axiom)
f. if $\underline{K}_i$ is reflexive then $M \models (K_i\phi \rightarrow \phi)$ (Knowledge axiom)

* IA = introspection axiom

Sample proof

$M \models (K_i\phi \wedge K_i(\phi \rightarrow \psi)) \rightarrow K_i\psi$ (distribution axiom)

Proof:

1. assume $(M, w) \models (K_i\phi \wedge K_i(\phi \rightarrow \psi))$
2. then $(M, w') \models \phi$ and $(M, w') \models (\phi \rightarrow \psi)$ for all $w' \in \underline{\boldsymbol{K}}_i(w)$
3. hence, $(M, w') \models \psi$ for all $w' \in \underline{\boldsymbol{K}}_i(w)$
4. Therefore $(M, w) \models K_i\psi$
5. Thus, $(M, w) \models (K_i\phi \wedge K_i(\phi \rightarrow \psi)) \rightarrow K_i\psi$
6. Since this is true for all $w \in W$, it follows that
7. $M \models (K_i\phi \wedge K_i(\phi \rightarrow \psi)) \rightarrow K_i\psi$

Consider the following collection of axioms and inference rules:

Prop. All substitution instances of tautologies of prop. logic

K1. $(K_i\phi \wedge K_i(\phi\rightarrow\psi)) \rightarrow K_i\psi$ (distribution axiom)

K2. $K_i\phi \rightarrow \phi$ (Knowledge axiom)

K3. $\neg K_i$ *false* (Consistency axiom)

K4. $K_i\phi \rightarrow K_i K_i \phi)$ (positive inspection axiom)

K5. $P_i\phi \rightarrow K_i P_i \phi)$ (negative inspection axiom)

MP. From ϕ and $\phi\rightarrow\psi$ infer ψ(modus ponens)

Gen. from ϕ infer $K_i\phi$ (rule of knowledge generalization)

K = {Prop, MP, Gen, K1}, T = {Prop, MP, Gen, K1, K2},
S4 = {Prop, MP, Gen, K1, K2, K4}, S5 = {Prop, MP, Gen, K1, K2, K4, K5}
KD45 = {Prop, MP, Gen, K1, K3, K4, K5}

Fact 2: Consider sets of epistemic structures where the accessibility relations satisfy particular conditions.

a. K is a sound and complete axiomatization with respect to all epistemic structures
b. T is a sound and complete axiomatization with respect to all *reflexive* epistemic structures
c. S4 is a sound and complete axiomatization with respect to all *reflexive* & *transitive* epistemic structures
d. S5 is a sound and complete axiomatization with respect to all *reflexive* & *transitive* & *symmetric* epistemic structures
e. KD45 is a sound and complete axiomatization with respect to all *Euclidian* & *transitive* & *serial* epistemic structures

3 Epistemic probability frames

Epistemic frames in multi agent systems are used for modelling the simplest kind of reasoning about uncertainty where certain situations are considered possible or impossible. Epistemic probability frames introduce degrees of certainty into this picture.

- 6 possible worlds
- in the world (A,B) agent 1 thinks two worlds are possible: (A,B) and (A,C)
- an epistemic probability frame assigns probabilities to these possibilities.

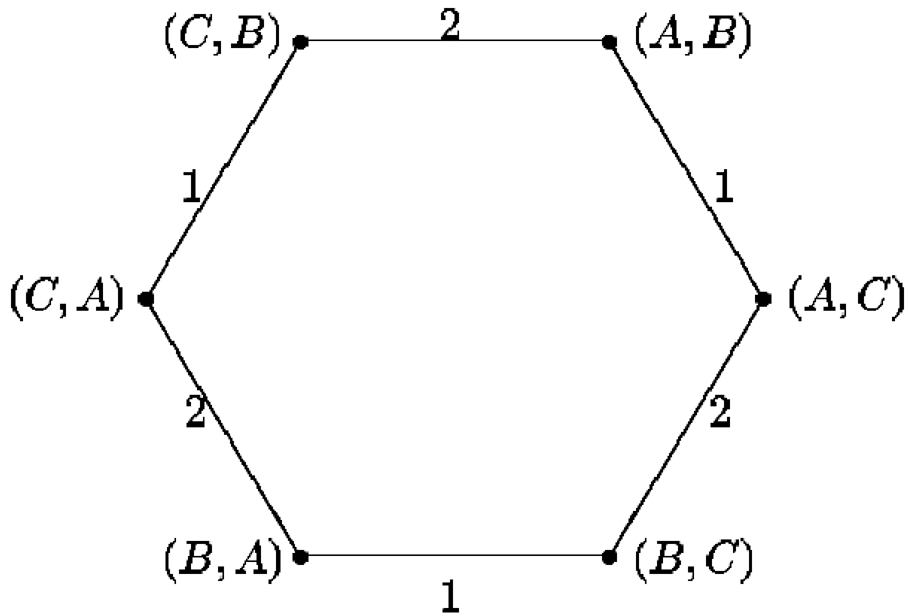

Definition 5: An epistemic probability frame for n agents is a tuple

$(W, \underline{\boldsymbol{K}}_1, \ldots \underline{\boldsymbol{K}}_n, \{(\mu_{w,1}, \ldots \mu_{w,n}): w \in W\})$ where

1. $(W, \underline{\boldsymbol{K}}_1, \ldots \underline{\boldsymbol{K}}_n)$ is an epistemic frame and
2. $(\underline{\boldsymbol{K}}_i(w), \mu_{w,i})$ is a classical probability space for each $1 \leq i \leq n$ (i.e. $\mu_{w,i}(U)$ is a defined probability for all propositions $U \subseteq \underline{\boldsymbol{K}}_i(w)$.)

Example

Suppose that a deck consists of three cards labelled *A*, *B*, and *C*. Agent 1 and 2 each get one of these cards; the third card is left face down. A possible world is describing the cards held by each agent. Give a description of the epistemic probability frame if we assume that agent 1 saw card 3 for 5 ms *before* she saw her own card and assigned probability ½ that card 3 was an *A*.

- 6 possible worlds
- in the world *CB* agent 1 thinks two worlds are possible: *CB* and *CA*.
- She assigns the probabilities $\mu_{CB,1}(CB)$ =? and $\mu_{CB,1}(CA)$ =?
- What about the other probabilities?

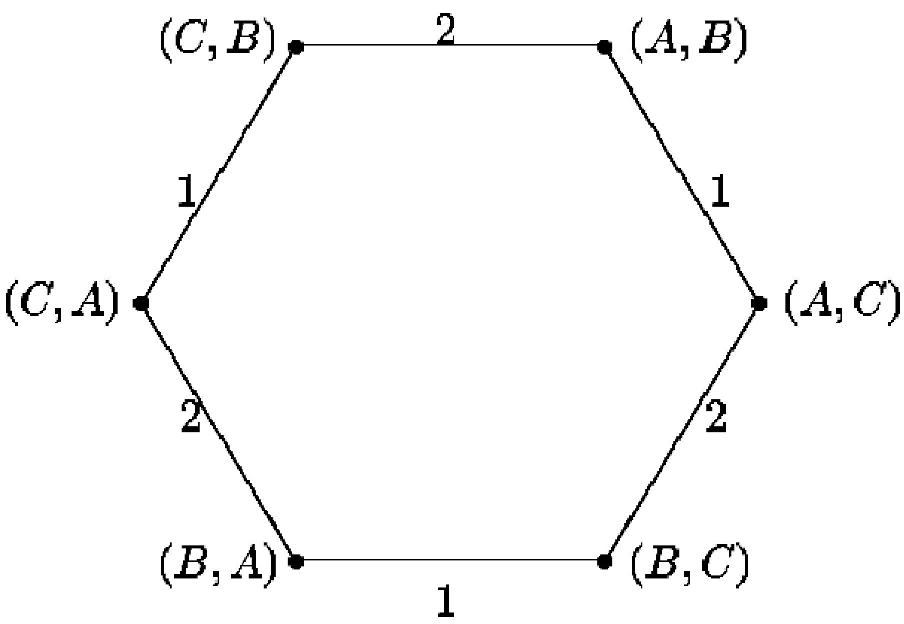

Suppose that a deck consists of three cards labelled *A*, *B*, and *C*. Agent 1 and 2 each get one of these cards; the third card is left face down. A possible world is describing the cards held by each agent. Give a description of the epistemic probability frame if we assume that agent 1 saw card 3 for 5 ms *before* she saw her own card and assigned probability ½ that card 3 was an *A*.

- 6 possible worlds
- in the world *CB* agent 1 thinks two worlds are possible: *CB* and *CA*.
- $\mu_{CB,1}(CB)=2/3$, $\mu_{CB,1}(CA)=1/3$
- $\mu_{AB,1}(AB)=1/2$, $\mu_{AB,1}(AC)=1/2$ $\mu_{CB,2}(CB)=1/2$, $\mu_{CB,2}(CA)=1/2$, …

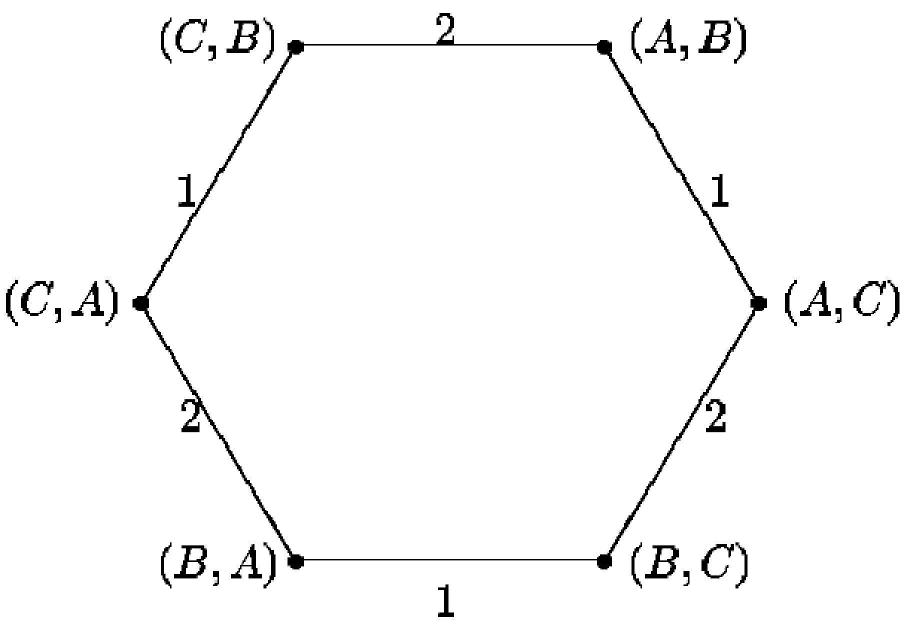

The following two constraints on epistemic probability frames are very natural for frames where the accessibility relations are equivalence relations.

SDP (*state-determined probability*)

For all i, v and w: if $v \in \underline{\boldsymbol{K}}_i(w)$ then $\mu_{v,i} = \mu_{w,i}$

CP (*common prior assumption*)

There exist a probability space (W, μ) such that

$\mu_{w,i} = \mu | \underline{\boldsymbol{K}}_i(w)$, i.e. $\mu_{w,i}(U) = \mu(U | \underline{\boldsymbol{K}}_i(w))$ for $U \subseteq \underline{\boldsymbol{K}}_i(w)$

CP says that differences in beliefs among agents can be completely explained by differences in information. Agents start out with identical prior beliefs and the condition on the information they later receive.

Which constraint is satisfied in our example?

4 The Dynamics of multi-agent systems

This section introduces multi-agent systems which are more realistic for modelling interacting agents (players in a poker game, robots interacting to clean a house, …).

– discriminating global and local (internal) states
– global state changes as a result of individual actions.

Cf. Halpern,
Chapter 6.3

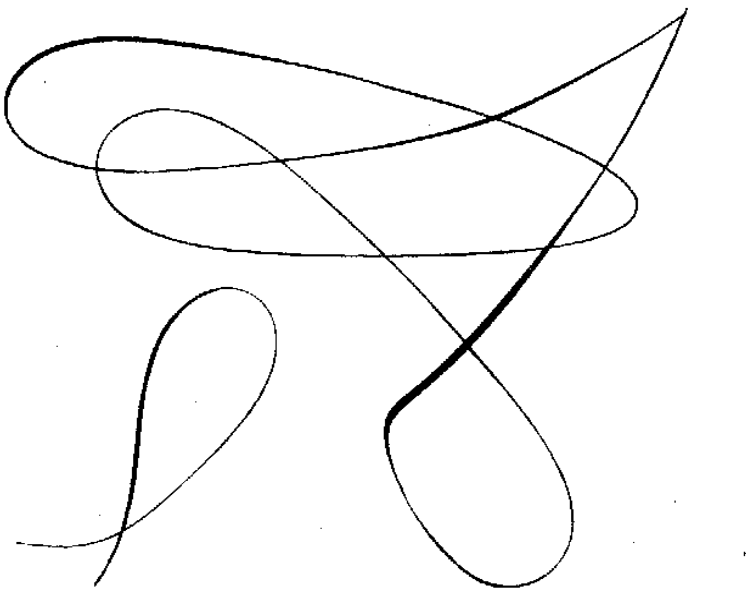

Local and global states, runs and points

- A *global state* describes the system at a given point t of time
- A global state is structured as follows: $(s_e, s_1, \ldots, s_n)$
  - s_e is the environment's state
  - s_i is agent i's local state
- A *run* is a complete description of one possible way in which the system's state can evolve over time. Formally, a run r is a function from time to global states, hence $r(t)$ describes the global state at time t.
- If $r(t) = (s_e, s_1, \ldots, s_n)$, then define
  - $r_e(t) = s_e$, the environment's state at point (r, t)
  - $r_i(t) = s_i$, the agent i's local state at point (r, t)

Systems of runs as epistemic frames

We can identify systems of runs $\mathfrak{R}$ with epistemic frames where the accessibility relations are equivalence relations:

$\boldsymbol{F}_{\mathfrak{R}} = (W, \underline{\boldsymbol{K}}_1, \ldots \underline{\boldsymbol{K}}_n)$ where

- W = $\{(r, t): r \in \mathfrak{R}$ and $t \in$ time$\}$, worlds are the points in $\mathfrak{R}$
- $\underline{\boldsymbol{K}}_i\ (r, t) = \{(r', t'): r_i(t) = r'_i(t')\}$, two points are equivalent for agent i iff they are indistinguishable to i (they identify the same local state)

To model a dynamic system as a multi-agent system requires deciding how to model the local states. In the case of multi-agent systems this is a harder task than in the single-agent case because now the uncertainty includes what agents are thinking about one another.

Suppose Alice tosses two coins and sees how the coins land. Bob learns how the first coin landed after the second coin is tossed, but does not learn the outcome of the second coin toss. There are exactly 7 possible global states:

- One initial state: ([.], [.], [.])
- Two time-1 states of the form ([X], [X], [*tick*]), where $X \in \{H,T\}$
- Four time-2 states of the form ([X_1, X_2], [X_1, X_2], [X_1, *tick*]), where $X_i \in \{H,T\}$

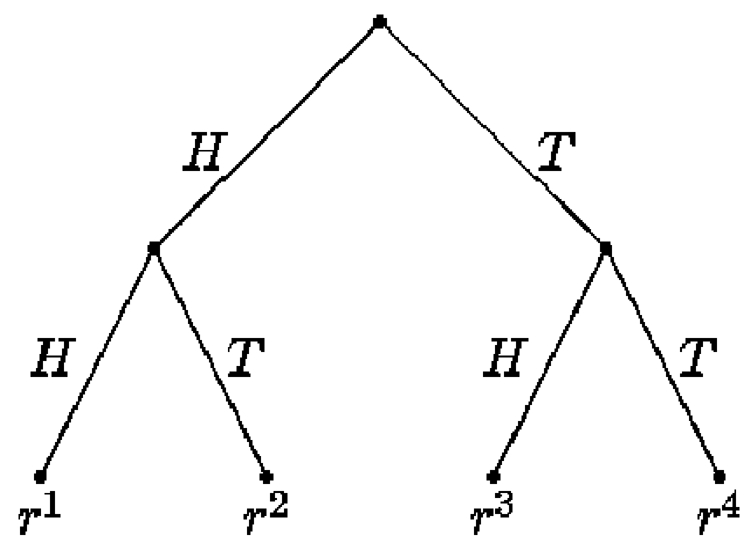

These states can be identified with the nodes of the shown tree.

Determine the equivalence classes

$\underline{\boldsymbol{K}}_A (r_1, 0) =$?
$\underline{\boldsymbol{K}}_B (r_1, 0) =$?

$\underline{\boldsymbol{K}}_A (r_1, 1) =$?
$\underline{\boldsymbol{K}}_B (r_1, 1) =$?

$\underline{\boldsymbol{K}}_A (r_1, 2) =$?
$\underline{\boldsymbol{K}}_B (r_1, 2) =$?

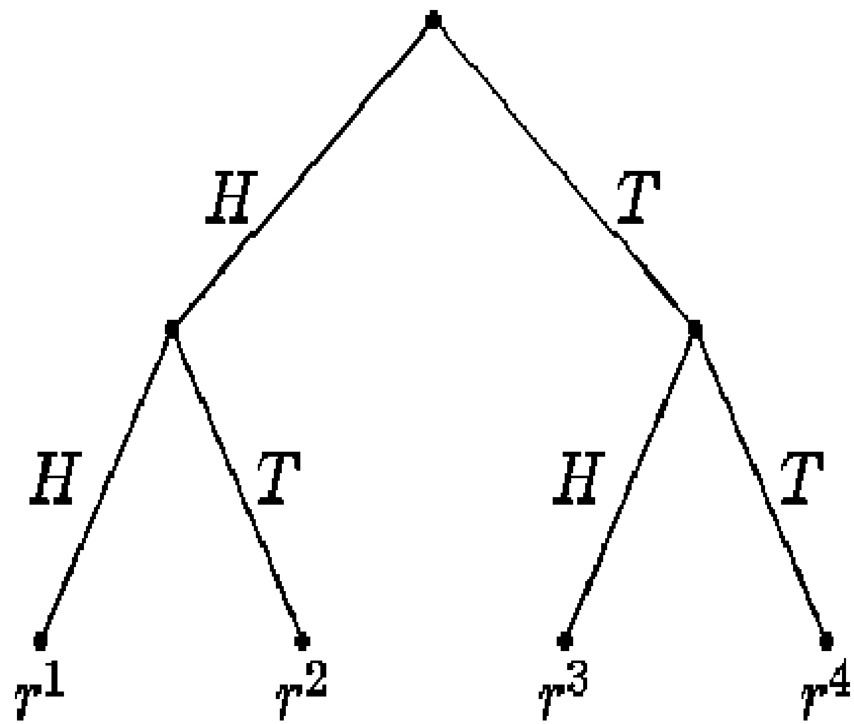

Determine the equivalence classes

$\underline{\boldsymbol{K}}_A\ (r_1, 0) = \{(r_1, 0), (r_2, 0), (r_3, 0), (r_4, 0)\}$
$\underline{\boldsymbol{K}}_B\ (r_1, 0) = \{(r_1, 0), (r_2, 0), (r_3, 0), (r_4, 0)\}$

$\underline{\boldsymbol{K}}_A\ (r_1, 1) = \{(r_1, 1), (r_2, 1)\}$
$\underline{\boldsymbol{K}}_B\ (r_1, 1) = \{(r_1, 1), (r_2, 1), (r_3, 1), (r_4, 1)\}$

$\underline{\boldsymbol{K}}_A\ (r_1, 2) = \{(r_1, 2)\}$
$\underline{\boldsymbol{K}}_B\ (r_1, 2) = \{(r_1, 2), (r_2, 2)\}$

Determining probabilities with uniform runs

- In the example under discussion we can assume that $\mu(r_i, t) = \mu(r_j, t)$ for all runs r_i and r_j. Obviously, the probabilities are independent of *t* and we can ask for the probabilities of *runs* given some local state.
- For example, agent *A* can be in state [H] – conforming to the situation $\underline{\boldsymbol{K}}_A(r_1, 1) = \{(r_1, 1), (r_2, 1)\}$. Or agent *B* can be in state [H, *tick*]. This conforms to the situation $\underline{\boldsymbol{K}}_B(r_1, 2) = \{(r_1, 2), (r_2, 2)\}$.
- Generally, we can assume CP (Section 3) and we get $\mu_{w,A} = \mu | \underline{\boldsymbol{K}}_A(w)$ and $\mu_{w,B} = \mu | \underline{\boldsymbol{K}}_B(w)$
- With $\mu(r_i) = 1/4$ we get $\mu_{(r1,\,1),\,A}(r_1) = \mu_{(r1,\,1),\,A}(r_2) = 1/2$ and $\mu_{(r1,\,1),\,B}(r_1) = \mu_{(r1,\,1),\,B}(r_2) = 1/4$. What about $\mu_{(r1,\,2)}$?

Determining probabilities with non-uniform runs

Consider the situation described before but assume now that the first coin has bias 2/3, the second coin is fair, and the tosses are independently as shown in the figure.

Calculate the probability distributions

(1) $\mu_{(r1, 1), A}(r_i)$

(2) $\mu_{(r1, 2), B}(r_i)$

(3) $\mu_{(r1, 2), A}(r_i)$

(4) $\mu_{(r1, 0), A}(r_i)$

Start with calculating $\mu(r_i)$!

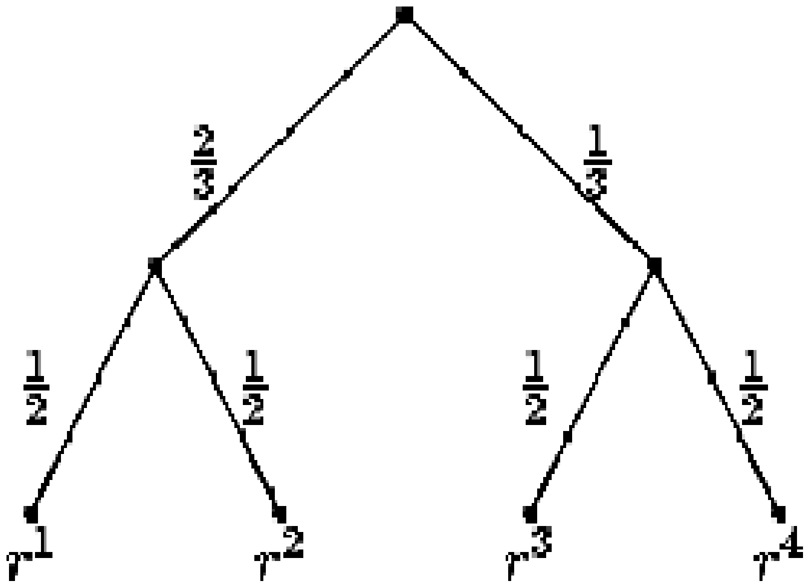

Determining probabilities with non-uniform runs

r_1	r_2	r_3	r_3
1/3	1/3	1/6	1/6

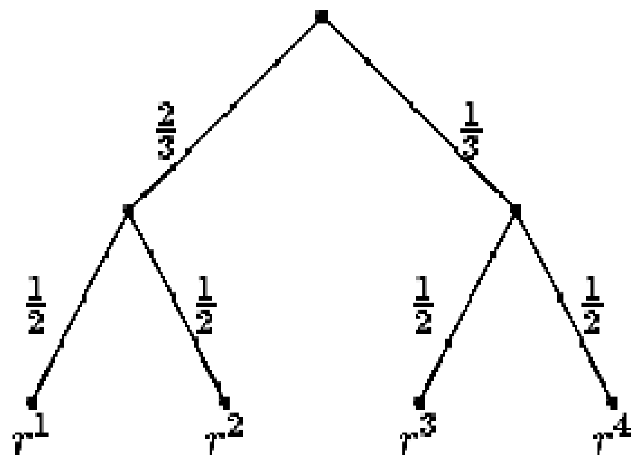

(1) $\mu_{(r1,\ 1),\ A}(r_i) = \mu(r_i | \{(r_1, 1), (r_2, 1)\}) = ½$ for i=1 and i=2

(2) $\mu_{(r1,\ 2),\ B}(r_i) = \mu(r_i | \{(r_1, 2), (r_2, 2)\}) = ½$ for i=1 and i=2

(3) $\mu_{(r1,\ 2),\ A}(r_i) = \mu(r_i | \{(r_1, 2)\}) = 1$ for i=1

(4) $\mu_{(r1,\ 0),\ A}(r_i) = \mu(r_i | \{(r_1, 0), (r_2, 0), (r_3, 0), (r_4, 0)\}) =$

$$\begin{cases} 1/3 \text{ for i=1 and i=2} \\ 1/6 \text{ for i=3 and i=4.} \end{cases}$$

5 Protocols

Multi-agent systems provide a useful way of representing complex situations. But where does the system come from? Changes often occur as a result of *actions*. These actions, in turn, are often performed as a result of agents using a *protocol* or *strategy*. Cf. Halpern, Chapter 6.6.

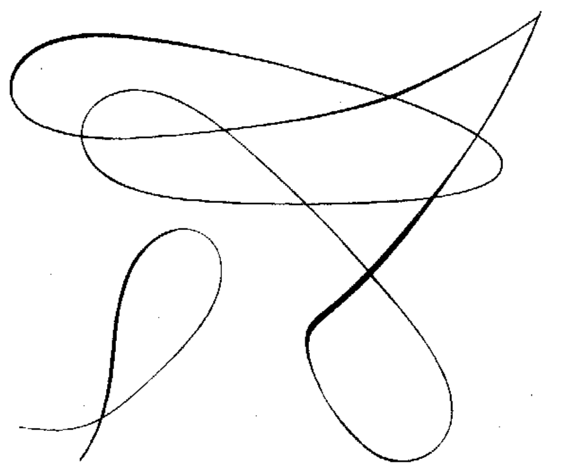

- A protocol P_i for agent i is a function that associates with every local state in L_i a nonempty subsets of actions in ACT_i
- If P_i is deterministic then P_i prescribes a unique action for i at each local state, that is $| P_i (l) | = 1$ for each $l \in L_i$.
- If P_i is a probabilistic protocol, then $| P_i (l) | > 1$ and each local state l is associated with a probability measure over a subset of action in ACT_i. We drop protocols P_e for the environment in this presentation (see Halpern, p 208)
- A joint protocol $(P_1, \ldots P_i)$ consists of a protocol for each of the agents and associates with each global state a subset of possible joint actions $ACT_1 \times \ldots \times ACT_n$.
- If the local protocols are probabilistic then a probability on the joint actions can be calculated by treating the local protocols as independent

Given a joint protocol and a set of initial global states, it is possible to generate a system of runs in a straightforward way.

Example:

On sonny days Alice tosses a coin with bias 2/3, on cloudy days Alice tosses a coin with bias 1/4, and sunny days happen with probability 3/4.

- Alice' local states: [sonny], [cloudy], [H], [T].
- Alice' actions: P_A([sonny]) = {toss-heads, toss-tails}: 2/3,1/3
  P_A([cloudy]) = {toss-heads, toss-tails}: 1/4, 3/4
- Resulting systems of runs (the nodes represent the states of the system)

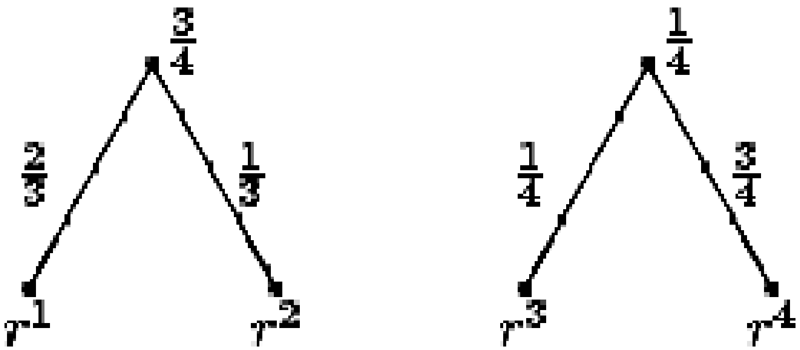

Calculations

Alice' local states (repeated):
[sonny], [cloudy], [H], [T].

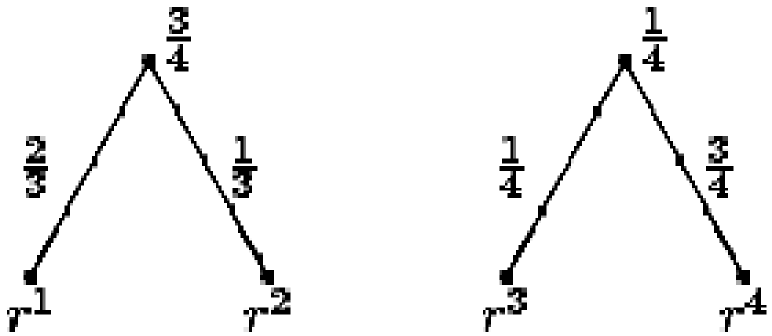

Exercise: equivalence classes and probabilities:

1. $\underline{\boldsymbol{K}}_A$ $(r_1, 0)$ = ?; $\underline{\boldsymbol{K}}_A$ $(r_3, 0)$ = ?
   $\underline{\boldsymbol{K}}_A$ $(r_1, 1)$ = ?; $\underline{\boldsymbol{K}}_A$ $(r_2, 1)$ = ?
2. $\mu(r_i)$ = ?
3. $\mu_{(r1,\ 1),\ A}\ (r_i)$ = ?

Calculations

Alice' local states (repeated):
[sonny], [cloudy], [H], [T].

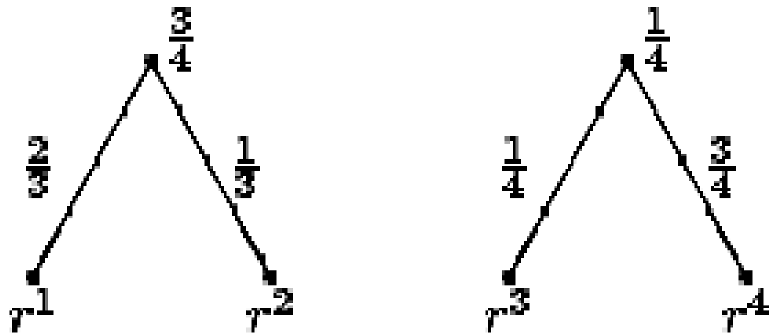

Exercise: equivalence classes and probabilities:

1. $\underline{\boldsymbol{K}}_A\ (r_1, 0) = \{(r_1, 0), (r_2, 0)\}$; $\underline{\boldsymbol{K}}_A\ (r_3, 0) = \{(r_3, 0), (r_4, 0)\}$
   $\underline{\boldsymbol{K}}_A\ (r_1, 1) = \{(r_1, 1), (r_3, 1)\}$; $\underline{\boldsymbol{K}}_A\ (r_2, 1) = \{(r_2, 1), (r_4, 1)\}$
2. $\mu(r_1) = ½$; $\mu(r_2) = ¼$; $\mu(r_3) = 1/16$; $\mu(r_4) = 3/16$;
3. $\mu_{(r1,\ 1),\ A}\ (r_1) = ½\ /\ (1/2+1/16) = 8/9$; $\mu_{(r1,\ 1),\ A}\ (r_3) = 1/9$
   $\mu_{(r1,\ 1),\ A}\ (r_2) = 4/7$, $\mu_{(r1,\ 1),\ A}\ (r_4) = 3/7$.

6 Example protocols to specify situations

The use of protocols helps clarify what is going on in many examples. Because protocols specify the possible actions (and their probabilities) for each local state, a concise description of the local states is required. Cf. Halpern, Chapter 6.7.

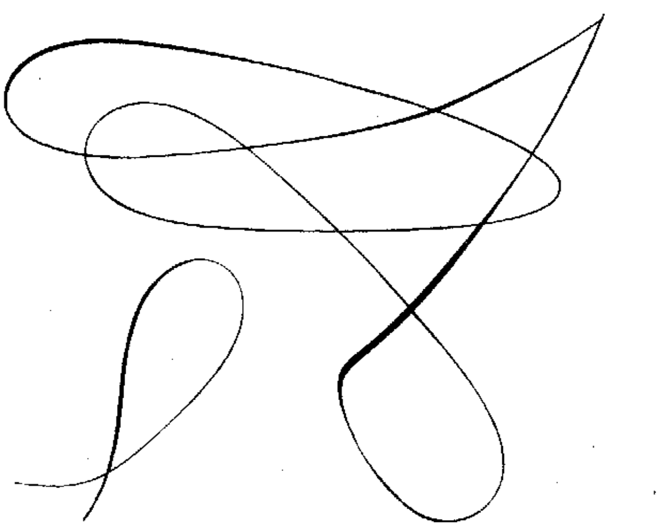

- Suppose the world is characterized by n binary random variables $X_1, \ldots X_n$. There are two agents, a *Teller*, who knows what the true values of the variables are, and a *Listener*, who initially has no idea what they are.
- In each round, the Teller gives the listener very limited information: She describes (truthfully) one world that is not the true world. For instance, if n=2, The Teller can say "not 10" to indicate that the true world is not (10)
- How can this situation be modelled as a system of runs? This depends on what he local states are and what protocol the Teller is following.
- The local states have to represent what the Teller/Listener remembers. We assume that the Teller remembers everything she has said.

A deterministic protocol

W = $\{(x_1, \ldots, x_n): x_i \in \{0,1\}\}$

T: $(w_0, [w_1, \ldots, w_k])$; $0 \leq k \leq n$, $w_0, w_i \in W$, $w_i \neq w_0$,
K=0 conforms to the initial state (w_0)

L: (w_k); (the listener remembers only the last message that T said)

Specify the system of runs for n=2 in case T is using a deterministic protocol (she follows the ordering 00,01,10,11 in generating her utterances)! Numbering of runs: r_{00}, r_{01}, r_{10}, r_{11}.

a. Determine the equivalence classes for L! E.g. $\underline{\boldsymbol{K}}_L(r_{11},1)$=?

b. Assume all worlds have equal probability. What is L's probability that the world is 11 – given his evidence at time 1, i.e. calculate $\mu_{(r11,\ 1),\ L}(r_{11})$ = ?

A deterministic protocol: solution

W = $\{(x_1, \ldots, x_n): x_i \in \{0,1\}\}$

T: $(w_0, [w_1, \ldots, w_k])$; $0 \le k \le n$, $w_0, w_i \in W$, $w_i \neq w_0$,
K=0 conforms to the initial state (w_0)

L: (w_k); (the listener remembers only the last message that T said). For the example assume n=2.

round 0	**00**	**01**	**10**	**11**
round 1	**--**	**not 00**	**not 00**	**not 00**
round 2	**--**	**--**	**not 01**	**not 01**
round 3	**--**	**--**	**--**	**not 10**
	r_{00}	r_{01}	r_{10}	r_{11}

a. $K_L(r_{11},1) = \{r_{01}, r_{10}, r_{11}\}$; $K_L(r_{00},1) = \{r_{00}\}$;

b. $\mu_{(r11,\,1),\,L}(r_{11}) = 1/3$; $\mu_{(r00,\,1),\,L}(r_{00}) = 1$

Another deterministic protocol

W = $\{(x_1, \ldots, x_n): x_i \in \{0,1\}\}$

T: $(w_0, [w_1, \ldots, w_k])$; $0 \le k \le n$, $w_0, w_i \in W$, $w_i \ne w_0$,
K=0 conforms to the initial state (w_0)

L: $[w_1, \ldots, w_k]$ (the listener remembers everything that T said)

- Taking the same system of runs as before, what are the equivalence classes for L? In particular, consider $\underline{\boldsymbol{K}}_L(r_{00}, 3)$?

T can also follow a non-deterministic strategy, for instance by considering all possible messages as equally distributed. What do you expect in this case? (start with considering a Listener with complete memory)

round 0	**00**	**01**	**10**	**11**
round 1				
round 2				
round 3				
Runs??				

Monty Hall puzzle

Suppose you're in a game show and given a choice of three doors. Behind one is a Car. Behind the others are Goats. You pick door 1. Before opening door 1, host Monty Hall (who knows what is behind each door) opens door 3, which has a goat. He then asks you if you still want to take what's behind door 1, or to take instead what's behind door 2. Should you switch?

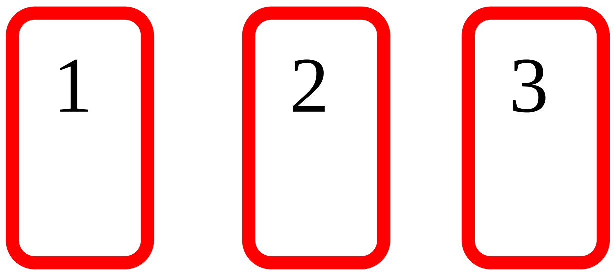

Construct (a) a deterministic (b) a nondeterministic protocol and calculate your subjective probability for a car being behind door 1 (door 2).

The Monty Hall puzzle: deterministic protocol

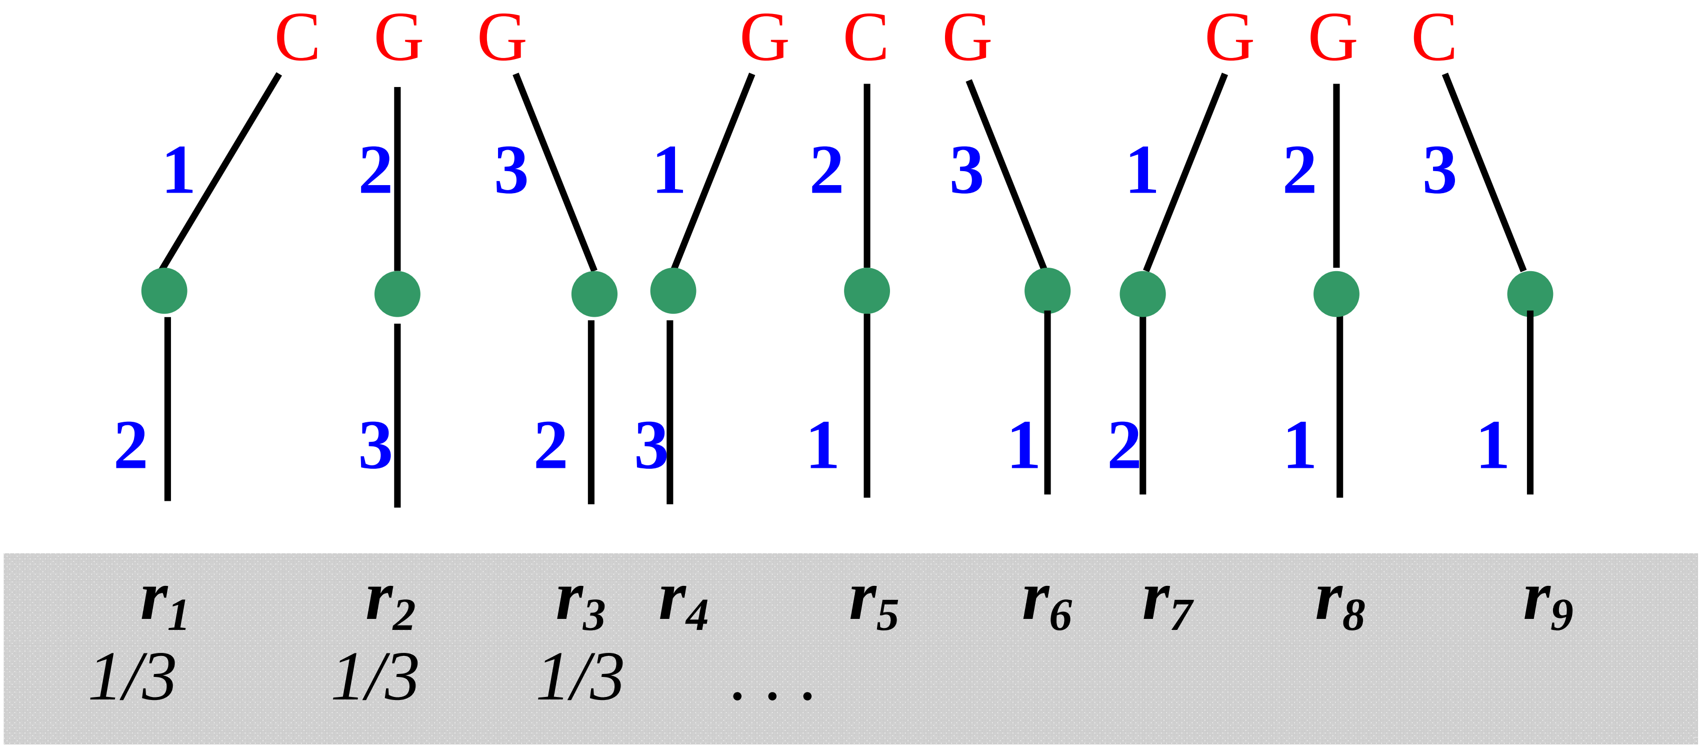

$\underline{\boldsymbol{K}}_{you}(r_1, 1) = \{r_1, r_4, r_7\}$; $\mu_{(r1, 1), you}(\{r_1\}) = 1/3$

$\underline{\boldsymbol{K}}_{you}(r_1, 2) = \{r_1, r_7\}$; $\mu_{(r1, 2), you}(\{r_1\}) = (1/3)/(2/3) = 1/2$

The Monty Hall puzzle: nondeterministic protocol

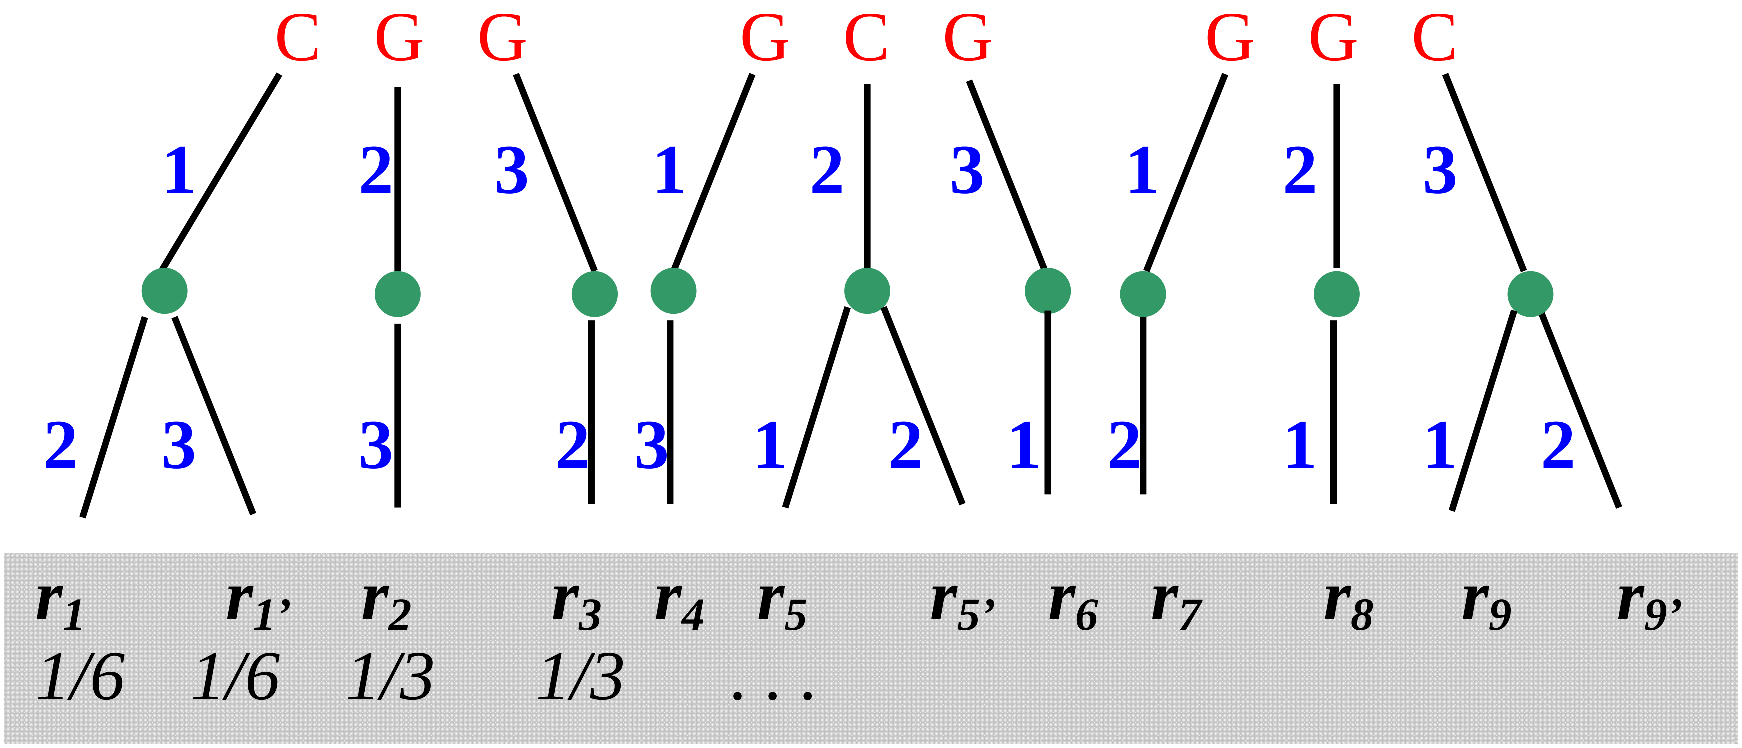

$\underline{\boldsymbol{K}}_{you}(r_1, 1) = \{r_1, r_{1'}, r_4, r_7\}$; $\mu_{(r1, 1), you}(\{r_1, r_{1'}\}) = 1/3$

$\underline{\boldsymbol{K}}_{you}(r_1, 2) = \{r_1, r_7\}$; $\mu_{(r1, 2), you}(\{r_1\}) = (1/6)/(3/6) = 1/3$

A deck has four cards: the ace and deuce of hearts A♡,2♡, and the ace and deuce of spades A♠, 2♠. After a fair shuffle of the deck, two cards are dealt to Alice. At this point there is a probability of 1/6 that Alice has both aces.

Alice now says (truthfully): “I have an ace”. Conditioning on this information, Bob computes the probability that Alice holds both aces to be 1/5. This seems reasonable.

Next, Alice says “I have the ace of spades”. Conditioning on this new information, Bob now computes the probability that Alice holds both aces to be 1/3. Of the three deals in which Alice holds the ace of spades, she holds both aces in one of them. Hence, as a result of learning that the ace Alice holds is the ace of spades, the conditional probability that Alice holds *both* aces goes up from 1/5 to 1/3.

But suppose Alice had instead said, “I have the ace of hearts”. A similar argument shows that the probability that Alice holds both aces is 1/3.

Is this reasonable? When Bob learns that Alice has an ace, he knows that she must have the ace of hearts or the ace of spades (or both). Why should finding out which particular ace it is raise the conditional probability of Alice having two aces?

The first step in analyzing this puzzle in the systems of runs framework is to specify the global states and the exact protocol being used. Surprisingly, there are two protocols that are consistent with the story, one deterministic, the other probabilistic

A deterministic protocol for the second ace puzzle

r_1	r_2	r_3	r_4	r_5	r_6
A♡,A♠	A♡,2♠	A♡,2♡	A♠,2♠	A♠,2♡	2♡,2♠
1/6	1/6	1/6	1/6	1/6	1/6
says Ace	says Ace	says Ace	says Ace	says Ace	says No
says A♠	says A♡	says A♡	says A♠	says A♠	says No

$\mu_{(r1,\,1),\,B}(r_1) = \mu(r_1 | \{r_1, r_2, r_3, r_4, r_5\}) = 1/5$

$\mu_{(r1,\,2),\,B}(r_1) = \mu(r_1 | \{r_1, r_4, r_5\}) = 1/3$

A probabilistic protocol for the second ace puzzle

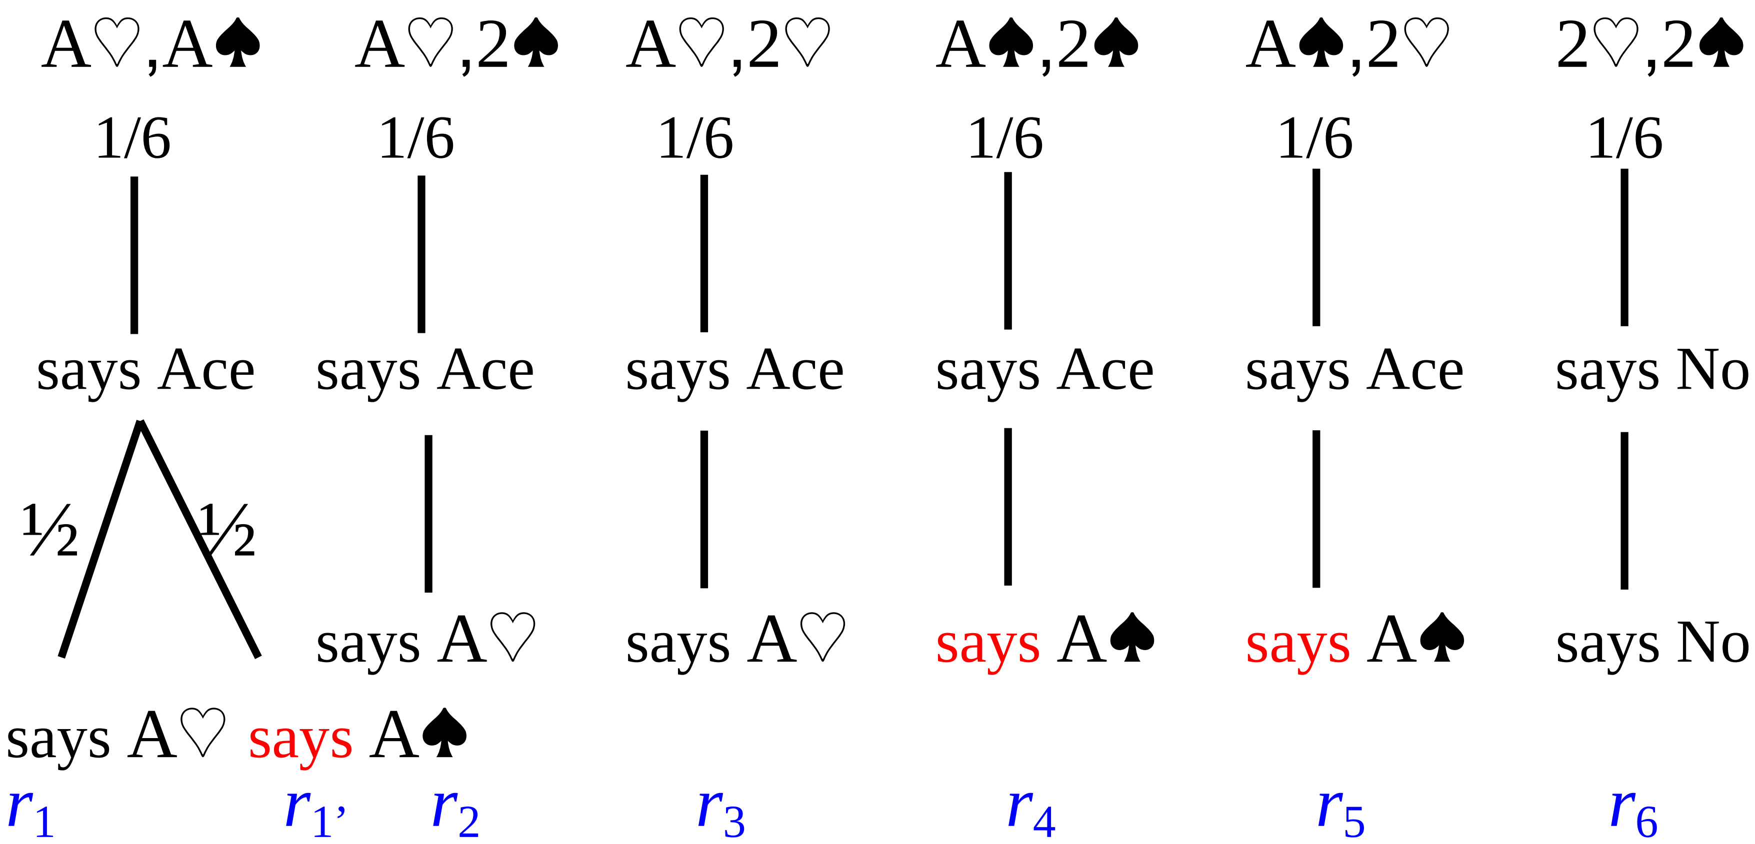

$\mu_{(r1,\,1),\,B}(\{r_1, r_{1'}\}) = \mu(\{r_1, r_{1'}\} \mid \{r_1, r_{1'}, r_2, r_3, r_4, r_5\}) = 1/5$

$\mu_{(r1',\,2),\,B}(\{r_1, r_{1'}\}) = \mu(\{r_{1'}\} \mid \{r_{1'}, r_4, r_5\}) = 1/5$