Computer Science Colloquium, Graduate Center CUNY

Knowledge-based rational decisions

Sergei Artemov

September 24, 2009

Rational decisions, informally

The standard game-theoretical assumption:

the player's rationality yields a payoff maximization given the player's knowledge.

Traditional Game Theory assumes ***enough knowledge*** to either

- avoid uncertainty completely (Aumann);
- deal with uncertainty probabilistically, i.e., when a player knows probability distribution of all consequences of his actions and is willing to take chances
(von Neumann & Morgenstern).

Game Theory

John von Neumann was an Hungarian American mathematician who made major contributions to mathematics, quantum mechanics, economics, and computer science. **Oskar Morgenstern** was an Austrian American economist. In 1944, he and von Neumann co-wrote ***Theory of Games and Economic Behavior***, recognized as the first book on game theory.

Robert Aumann

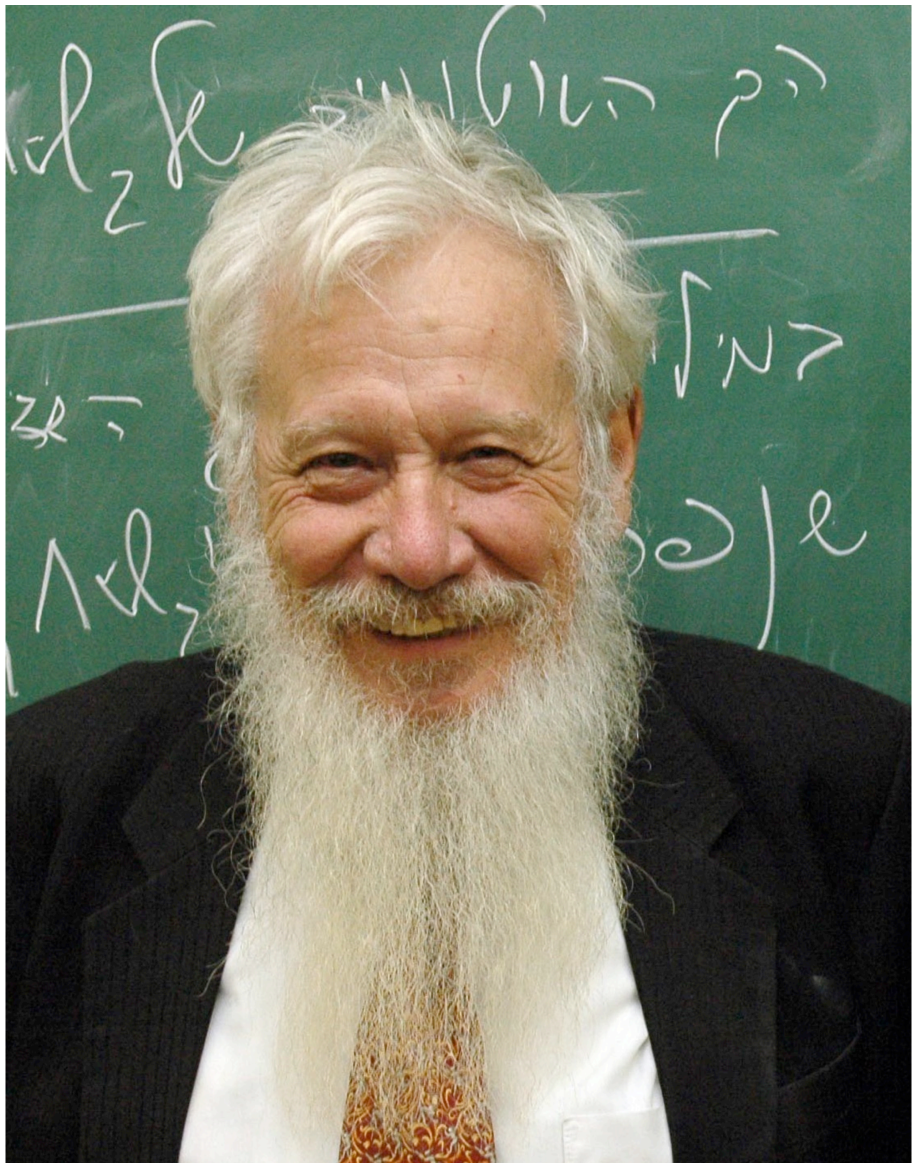

Alma Mater: City College of New York, Nobel Prize of 2005. Pioneered studies of Mathematical theory of ***Rationality and Common Knowledge***.

Rational decisions, informally

The standard game-theoretical assumption:

the player's rationality yields a payoff maximization given the player's knowledge.

There was no epistemically justified complete framework of making decisions under uncertainty with unknown probability distribution.

What is in the talk?

We outline a mathematical model of rational decision-making under uncertainty which is based on the standard game-theoretical postulates:

- rationality yields a payoff maximization;
- decisions are made by players based on their knowledge;
- the logic of knowledge for Game Theory is the modal logic **S5**.

Probabilistic method

Game Tree:

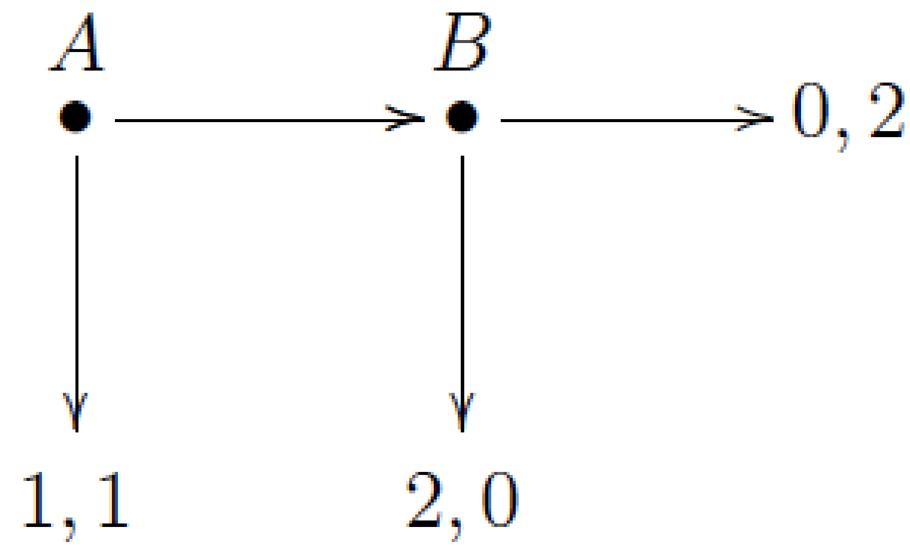

Suppose B is twice as likely to play *across* than *down*. Then A's average payoff when A plays *across* is 2/3 which is less than the payoff of 1 when playing *down*.

Hence the rational choice for A is *down*.

Backward induction method

Game Tree:

$$
\begin{tikzcd}
\overset{A}{\bullet} \arrow[r] \arrow[d] & \overset{B}{\bullet} \arrow[r] \arrow[d] & 0,2 \\
1,1 & 2,0 &
\end{tikzcd}
$$

Suppose *A* knows that *B* is rational. Then *A* knows that *B* will play *across*, thus delivering payoff 0 to *A*.

Hence the rational choice for *A* is *down*.

Uncertainty without probabilities?

Game Tree:

$$
\begin{tikzcd}
\overset{A}{\bullet} \arrow[r] \arrow[d] & \overset{B}{\bullet} \arrow[r] \arrow[d] & 0,2 \\
1,1 & 2,0 &
\end{tikzcd}
$$

What if *A* does not know that *B* is rational? *A* neither knows *a priori* probabilities, nor is *A* himself willing to assign *ad hoc* probabilities.

There is no canonical solution to this game.
There is no mature theory of rational decision making depending on different epistemic states of players.

Epistemic Game Theory

First admitted that epistemic states matter and studied conditions under which standard game theoretical solutions hold (backward induction, Nash, etc.).

Is still on the way towards developing a coherent theory of games in which epistemic states of players are a legitimate part of the game specification?

Foundational Problem

Rational decision making under uncertainty in perfect information games.

$$
\begin{tikzcd}
\overset{A}{\bullet} \arrow[r] \arrow[d] & \overset{B}{\bullet} \arrow[r] \arrow[d] & 2,2 \\
1,1 & 0,0 &
\end{tikzcd}
$$

A is mission control which has the option of sending to space a team which has not yet recovered from its previous mission (1,1), or sending a fresh crew *B* whose captain has been exposed to measles. If *B*'s captain does not gets sick, the mission will be a success (2,2), otherwise the mission should be aborted with failure (0,0).

Knowledge-based decisions

Decisions are made exclusively on the basis of player's knowledge.

In particular, we do not consider decision methods which rely on luck, guesswork, wishful thinking, opponent error, miracles, divine intervention, *etc*.

Definitive decisions

Generic game: no indistinguishable payoffs for each player.

A decision-making method provides a definitive choice at each node of a generic game.

This condition rules out speculative ‘solutions’ such as ‘all moves are rational,’ which are easy to offer.

Rational Decisions

Aumann:

"[A] rational player will not knowingly continue with a strategy that yields him less than he could have gotten with a different strategy."

Brandenburger:

"Rational player always plays the highest payoff strategy given his knowledge."

Such descriptions contain not fully defined notions like 'strategy yield,' 'highest payoff strategy,' *etc*. Since we consider only decision making based on **knowledge**, it is fair to make these notions well-defined by relativizing them to agent's knowledge: '*known strategy yield,*'*'highest known payoff strategy,' etc*.

Plan

1. We first develop a theory of knowledge-based rational decisions.

2. We then prove a general theorem that for perfect information games, this decision-making method is the only one which is rational, definitive, and based on knowledge.

3. Finally, we will show how the world of perfect information games looks through the prism of knowledge-based rationality.

Player P's rationality will be represented by a special atomic proposition

$$rP \quad - \quad \text{`}P \textit{ is rational.'}$$

Player P's knowledge (or belief) will be denoted by modality $\mathbf{K}_P$, hence

$$\mathbf{K}_P(F) \quad - \quad \text{`}P \textit{ knows (believes) that } F\text{.'}$$

In particular, $\mathbf{K}_P(rQ)$ states that 'player P knows (believes) that player Q is rational.'

In addition, we assume that rationality is self-known:

$$rA \ \rightarrow \ \mathbf{K}_A(rA).$$

We consider games presented in a tree-like extensive form. Let, at a given node of the game, player P have to choose one and only one of moves $1, 2, \ldots, m$, and s_i denote

$$s_i \;\equiv\; P \textit{ chooses i-th move.} \tag{2}$$

In particular, the following holds:

$$s_1 \vee s_2 \vee \ldots \vee s_m, \qquad\qquad s_j \rightarrow \bigwedge_{i \neq j} \neg s_i. \tag{3}$$

Definition 1 *For a given node v of the game, the corresponding player A, and a possible move j by A, the* **Highest Known Payoff**, *$HKP_A(j)$ is the highest payoff* **implied by A's knowledge** *at node v, given j is the move chosen by A. In more precise terms,*

$$HKP_A(j) = \max\{a \mid A \textit{ knows at } v \textit{ that his payoff given } s_j \textit{ is at least } a\}.$$

Let $G(a)$ be the (finite) set of all possible payoffs for A which are greater than a. Then, the highest known payoff can be defined as follows: $HKP_A(j) = a$ if and only if

$$\mathbf{K}_A(\text{‘}A \text{ gets at least } a \text{ when choosing } j\text{’})$$

and

$$\bigwedge_{b \in G(a)} \neg\mathbf{K}_A(\text{‘}A \text{ gets at least } b \text{ when choosing } j\text{’}) \ .$$

Proposition 1 [Correctness of HKP] *For each node of a finite game, corresponding player A, and possible move j by A, there exists a unique* $HKP_A(j)$.

Definition 2 Best Known Move *for player A at a given node of the game is a move j from* $1, 2, \ldots, m$ *which has the largest highest known payoff,* $HKP_A(j)$[1]*. In a more formal setting,* j *is a best known move for* A *at a given node if for all* i *from* $1, 2, \ldots, m$

$$HKP_A(j) \geq HKP_A(i) \ .$$

By

$$kbest_A(j)$$

we denote the proposition

'j is the best known move for A at a given node.'

In a yet even more formal setting, $kbest_A(j)$ can be formally defined as

$$kbest_A(j) \ \equiv \ \bigwedge_i [HKP_A(j) \geq HKP_A(i)] \ . \tag{4}$$

Theorem 1 *A best known move exists at each node and is always known to the player:*

1) If $kbest_A(j)$ holds, then $\mathbf{K}_A[kbest_A(j)]$.

2) If $\neg kbest_A(j)$ holds, then $\mathbf{K}_A[\neg kbest_A(j)]$.

Corollary 1 *At each node, there is always at least one best known move*

$$kbest_A(1) \vee kbest_A(2) \vee \ldots \vee kbest_A(m) \ .$$

If, in addition, all payoffs are different, the best known move is unique

$$kbest_A(j) \ \rightarrow \ \bigwedge_{i \neq j} \neg kbest_A(i) \ .$$

1. Rational player A always plays the highest payoff strategy given A's knowledge (Brandenburger, lectures).

2. " [A] rational player will not knowingly continue with a strategy that yields him less than he could have gotten with a different strategy." (Aumann, [5]).

3. "...a player is irrational if she chooses a particular strategy while believing that another strategy of hers is better." (Bonanno, [9])

4. For a rational player i, "there is no strategy that i knows would have yielded him a conditional payoff ... larger than that which in fact he gets." (Aumann, [5])

5. Rational player A chooses a strategy if and only if A knows that this strategy yields the highest payoff of which A is aware.

The natural formalization of 1 is the principle

$$rA \quad \rightarrow \quad [kbest_A(j) \rightarrow s_j] \ . \tag{5}$$

The natural formalization of 2 is the principle

$$rA \quad \rightarrow \quad [kbest_A(j) \rightarrow \neg s_i], \quad when\ i \neq j \ . \tag{6}$$

The natural formalization of 3 is the principle

$$[kbest_A(j) \wedge s_i], \ \rightarrow \quad \neg rA, \quad when\ i \neq j \ . \tag{7}$$

The natural formalization of 4 is the principle

$$rA \quad \rightarrow \quad [s_i \quad \rightarrow \quad \neg kbest_A(j)], \quad when\ i \neq j \ . \tag{8}$$

The natural formalization of 5 is the principle

$$rA \quad \rightarrow \quad [kbest_A(j) \leftrightarrow s_j] \ . \tag{9}$$

Theorem 2 *Principles (5–9) are equivalent.*

Definition 3 [Rationality Thesis] *Principles (5–9) are assumed to be commonly known.*

The aforementioned Rationality Thesis provides a method of **decision-making under uncertainty**: a rational player at a given node calculates his highest known payoff and his best known move and chooses accordingly. We propose calling such a decision-making method *knowledge-based rationality, KBR.*

Definition 4 *By a KBR-solution of the game, we mean the assignment of a move to each node according to the Rationality Thesis (Definition 3).*

Back to the Foundational Problem

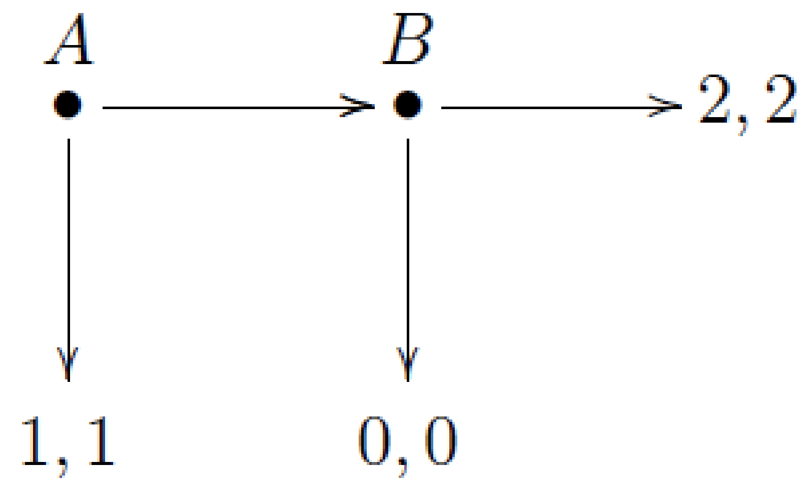

For player *A, HKP(down)=1*.

HKP(across)=0, since *A* ***does not know*** that he will get any higher payoff than 0.

Solution: A plays *down*. (Exactly like in ***Apollo 13***...)

More Games

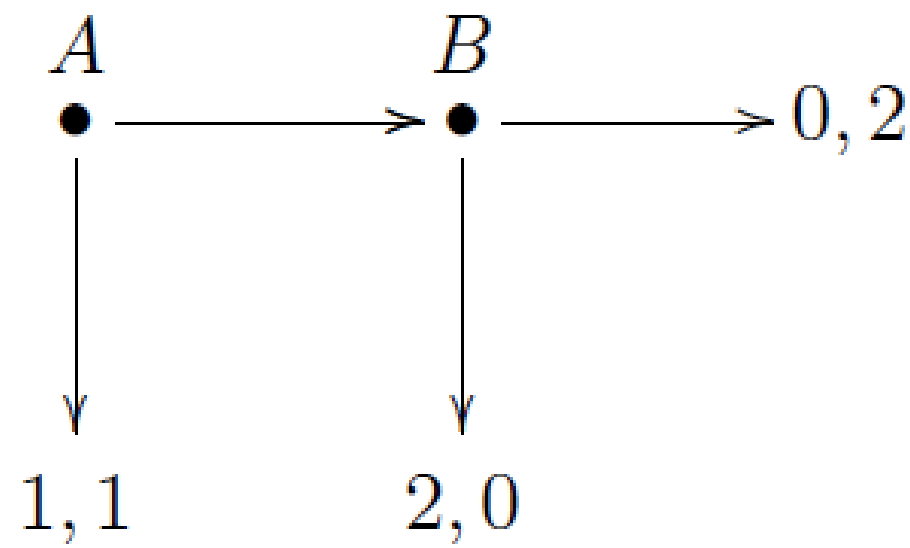

A is not aware of *B*'s rationality.
HKP(down)=1.

HKP(across)=0, since *A* ***does not know*** that he will get any higher payoff than 0.

Solution: A plays *down.*

Passive manipulation

Game Tree:

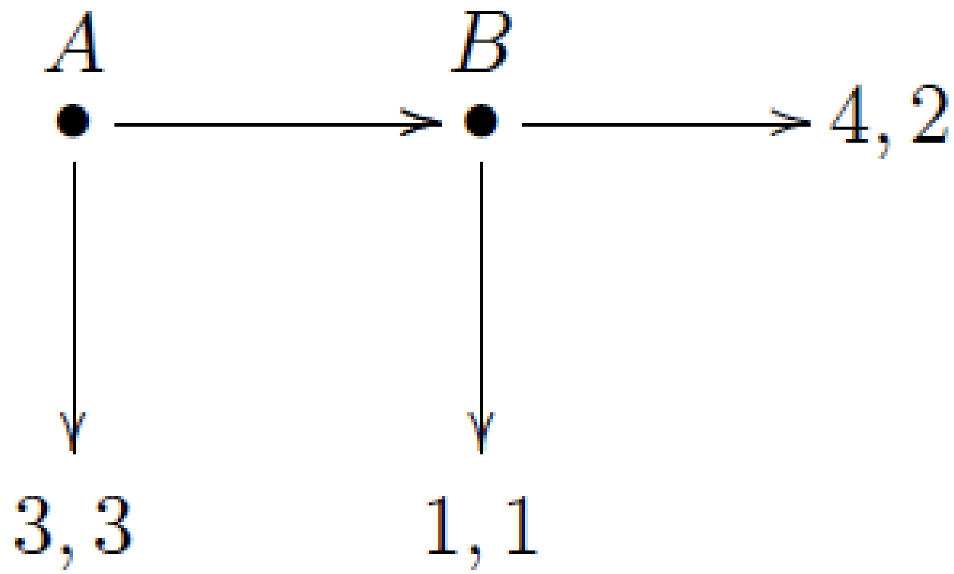

A is not sure of B's rational behavior, *A* plays *down*, payoffs 3,3.

B does not have the incentive to disclose his rationality since B wants A to move down.

Active manipulation

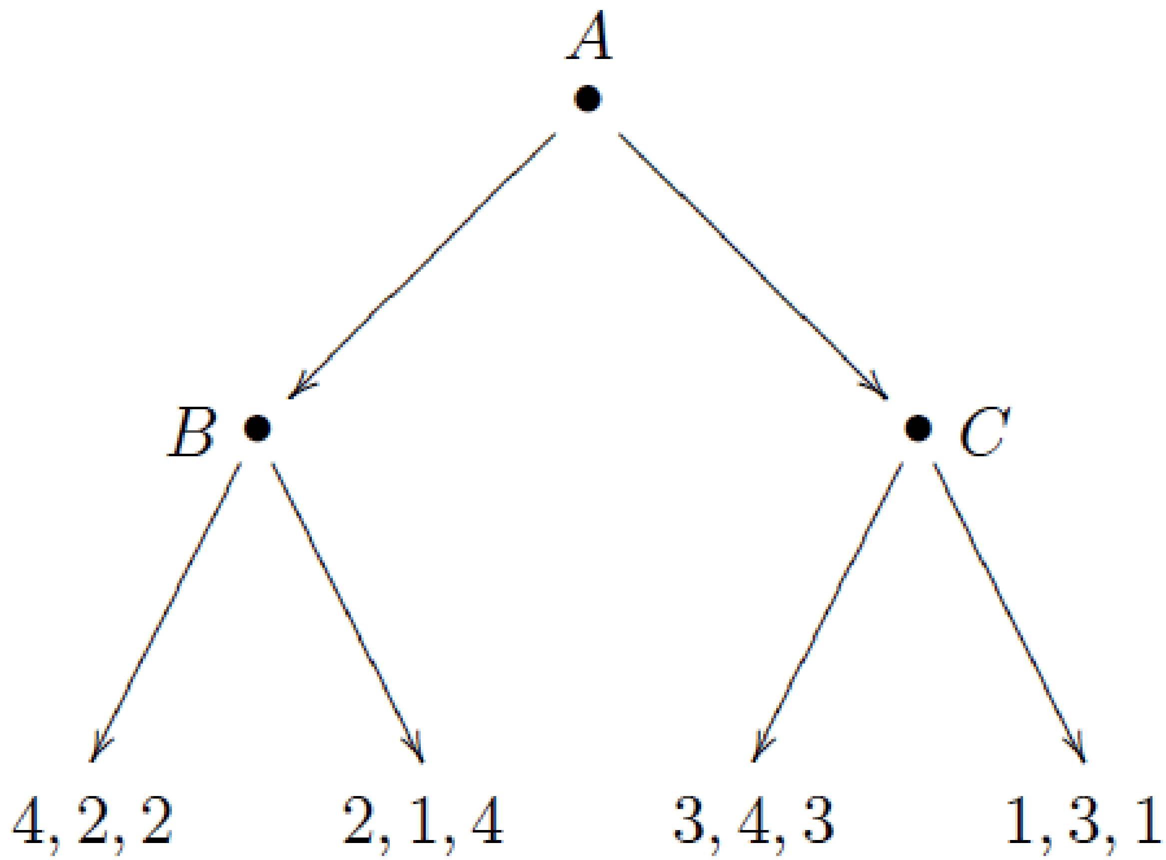

Suppose *A* is not aware of *B* and *C*'s rationality. Then *A* moves *left* to secure payoff 2. Actually, *A* gets 4 which is more than expected. Suppose also that *B* and *C* are smart enough to understand this. Then *B* can manipulate *A* by leaking the true information that ***C is rational***. *A* then knows that *right* secures his payoff 3, which is higher than *A*'s known payoff of *left*: *A* plays *right* and gets 3 (less), *B* gets 4 (much more) and *C* gets 3 (more). *C* does not have an incentive to disclose that *B* is rational, hence

B wins without ever making a move!

The Centipede game (before)

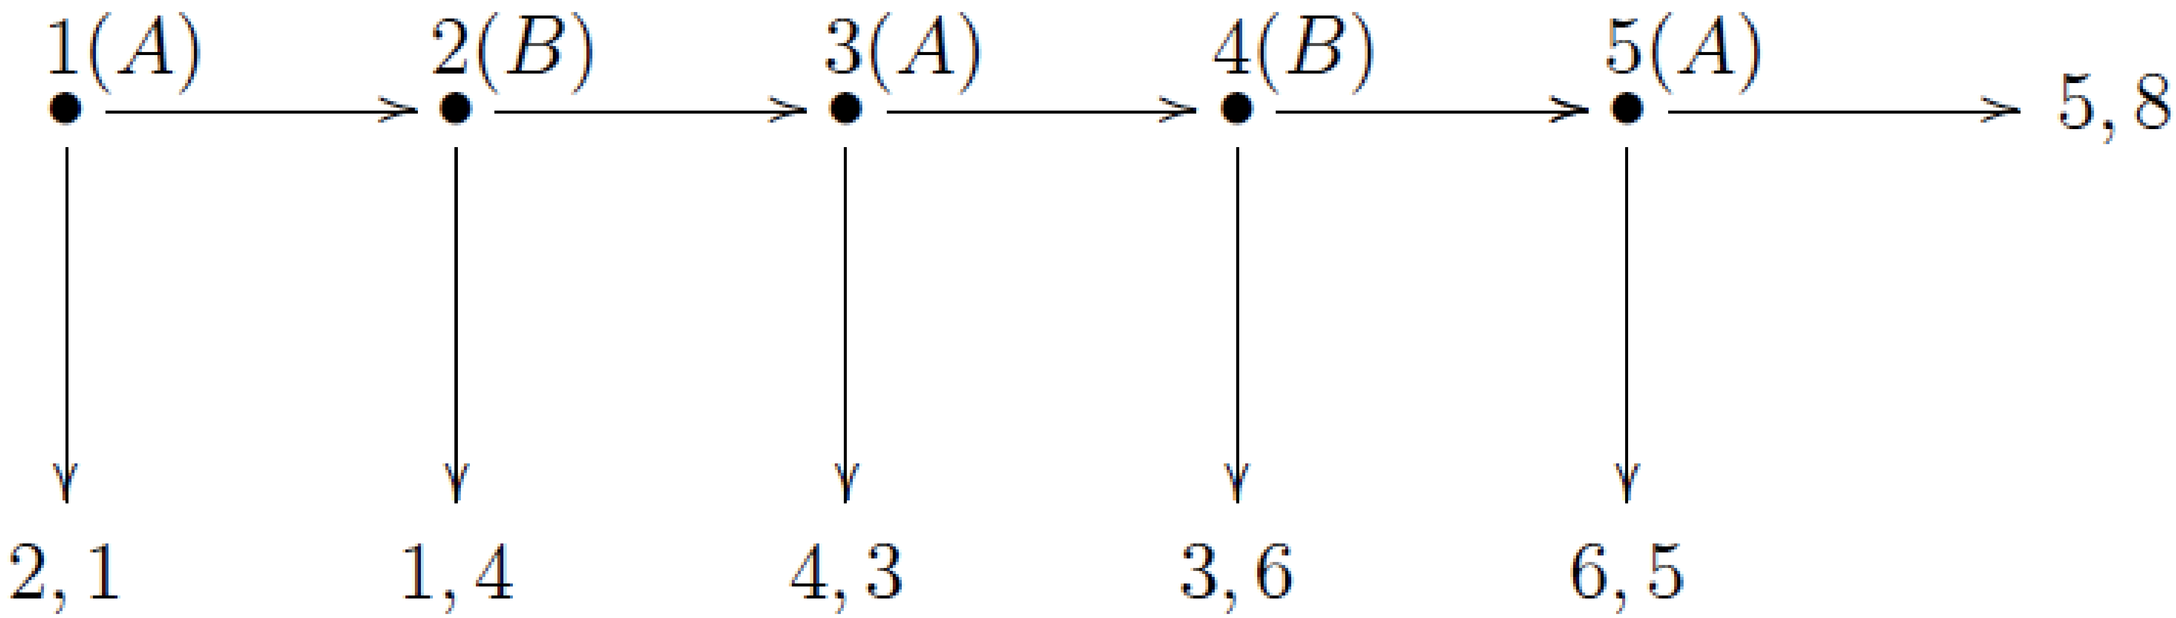

A is rational, hence at node 5, *A*'s choice is *down*.
B knows that A is rational, hence *B* plays *down* at 4.
A knows that B knows that A is rational, hence *A* plays *down* at 3.
B knows that A knows that B knows that A is rational ...

Unbounded nested knowledge of rationality is assumed!

The Centipede game (after)

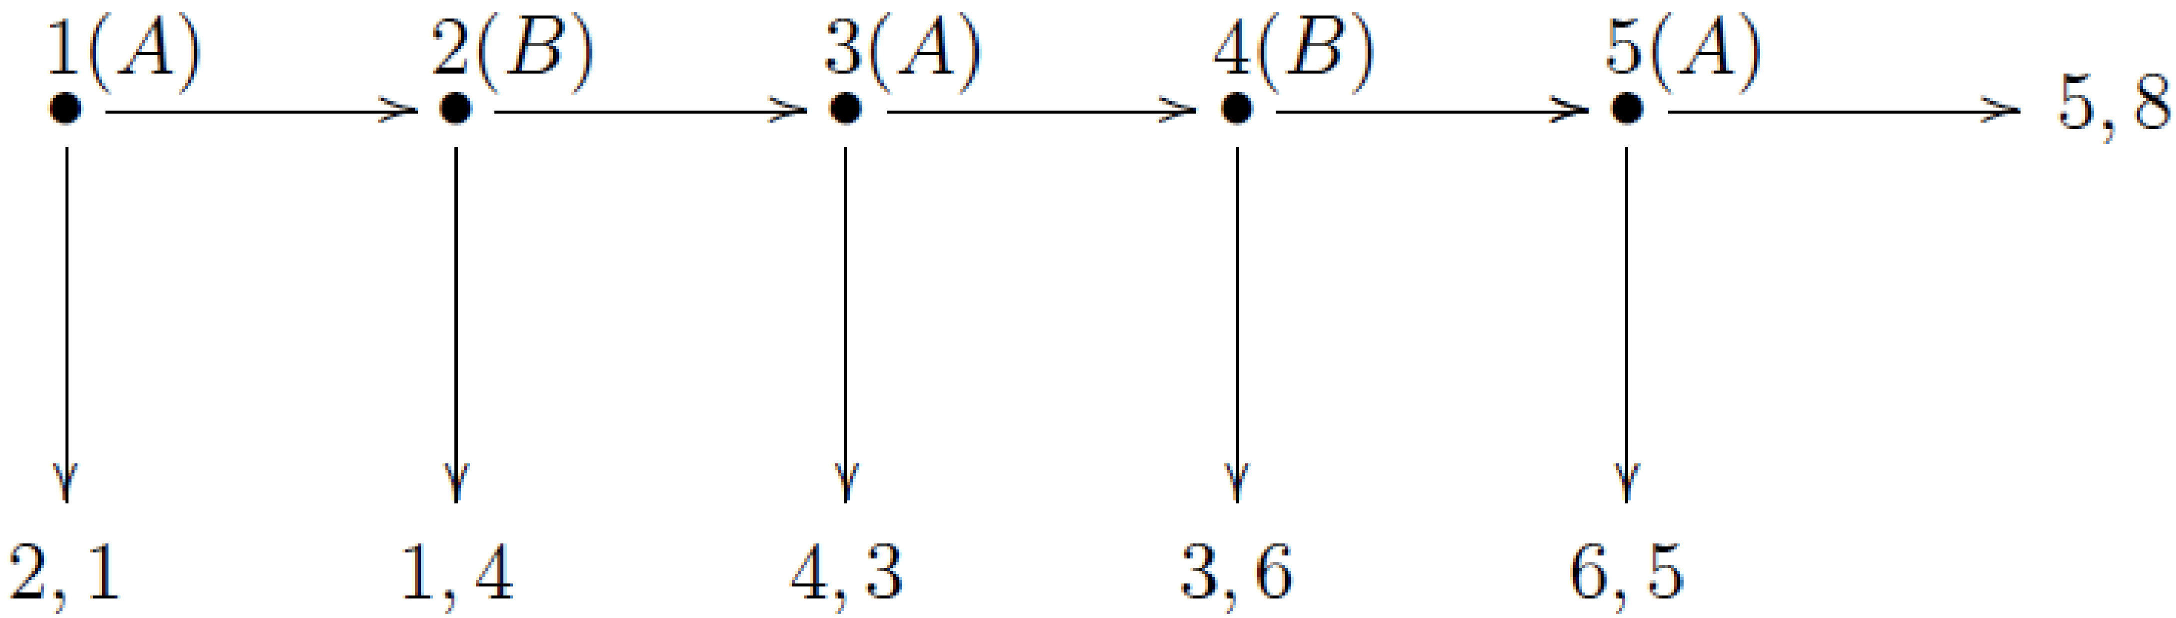

A plays *down* at 5. Consider the latest node where *across* is played (if there is none, we are done). Suppose this is node 1. Since *A* plays *across* at 1, *A* knows that *across* is the better choice, hence *A* knows that *B* plays *across* at 2. But this is impossible, since *B* actually plays *down* at 2, hence *A* plays *down* at 1 as well.

No knowledge about other players is needed!

KBR Theorem

For games of perfect information, KBR is the only decision-making method which is

1. ***Based on knowledge,***
2. ***Definitive,***
3. ***Rational.***

KBR Theorem

For games of perfect information, KBR is the only decision-making method which is

1. Based on knowledge,

2. Definitive,

3. Rational.

Examples of decision methods.

- *Assigning subjective probabilities* - (1) fails.
- *Maximin* - (3) fails, since it underutilizes knowledge.
- *Maximax* - (1) & (3) fails, the player cannot know that he gets maximum possible value.
- *Eliminating dominated strategies* - (1) and (3) hold, but not (2)

KBR Theorem

For games of perfect information, KBR is the only decision-making method which is

1. ***Based on knowledge,***
2. ***Definitive,***
3. ***Rational.***

KBR method can be viewed as 'epistemic maximin,' which has a natural explanation. ***Knowledge*** is defined as ***MIN*** of truth values over all possible states. ***Rationality*** is defined as ***MAX*** over all possible choices. Therefore, since we agree to base rationality on knowledge, we inevitably end up with ***MAXIMIN***, where ***MAXI comes from rationality*** and ***MIN comes from knowledge***.

Full knowledge is power

Model predictions:

- ***Every game with rational players has a solution.*** Rational players know which moves to make at each node.
- ***Those who know the game in full know its solution***, i.e., know everybody's moves.

Partial knowledge can hurt

Model predictions:

- More knowledge yields a ***higher known payoff*** but not necessarily a ***higher actual payoff***. So

  nothing but the truth

  can be misleading.

- Knowing

  the whole truth

  however, yields a higher actual payoff.

When knowledge does not matter

Model predictions:

In ***strictly competitive (e.g. zero-sum) games***, all players' epistemic states lead to the same (maximin) solution.

Maybe this is why military actions (typical zero-sum games) do not require sophisticated reasoning about other players:

just do it

normally suffices.

Belief *vs*. Knowledge check-up

Logic of Knowledge in Game Theory: **S5**.

Logic of Beliefs in Game Theory? **KD45** - the logic of consistent beliefs with positive and negative introspection.

Main theorems of KBR hold for beliefs as well...

Conclusions

Do we recommend playing perfect information games using KBR strategy?

1. Not if you can responsibly assign probabilities to your opponents' responses, otherwise

2. To the best of your knowledge, rule out all impossible paths of the game. If some uncertainly remains, it's this: You cannot know more. Deal with this uncertainty using KBR; this is the only rational method of playing PI games.