Microsoft Research

Rational decisions in non-probabilistic setting

Sergei Artemov
The Graduate Center CUNY

January 21, 2010

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).

Probabilistic assumption

Game Tree:

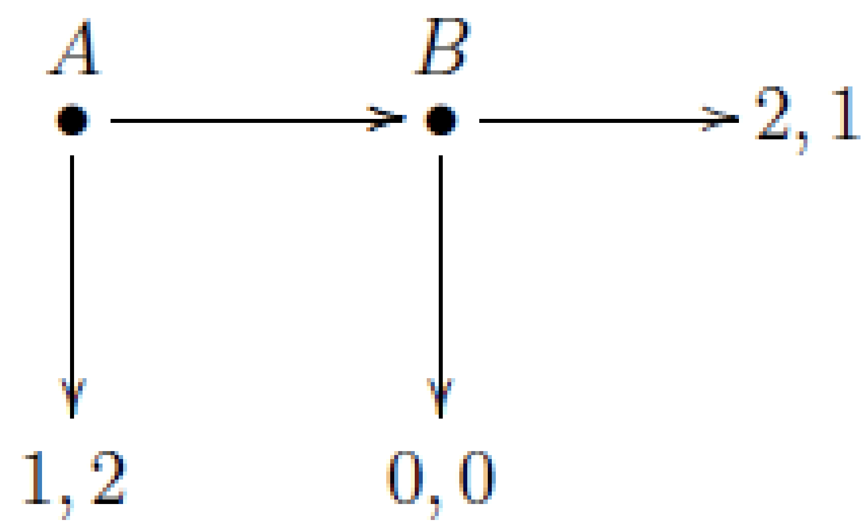

Suppose that *A* knows that *B* is **twice as likely** to play *across* than *down*. Then *A*'s average payoff when *A* plays *across* is 4/3 which is greater than the payoff of 1 when playing *down*.

Hence the best choice for *A* is *across*.

Epistemic assumptions

Game Tree:

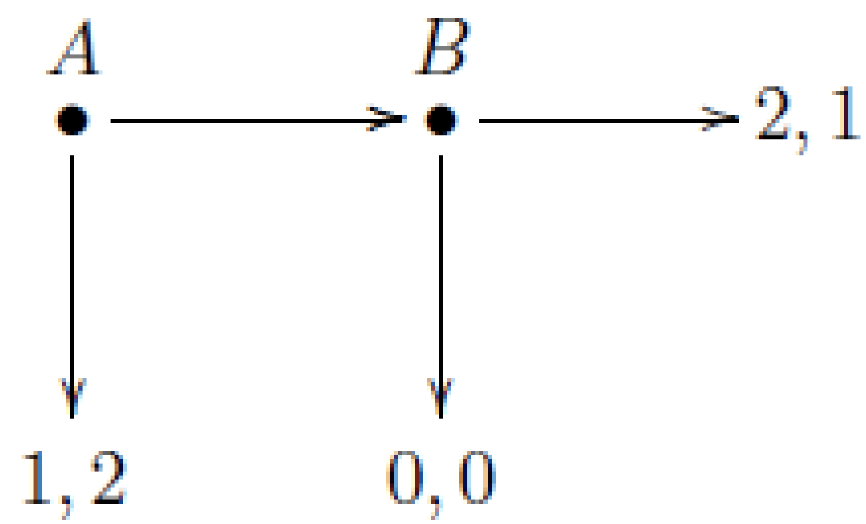

Suppose *A* **knows** that *B* is rational. Then *A* knows that *B* will play *across*, thus delivering payoff 2 to *A*.

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

Uncertainty without probabilities?

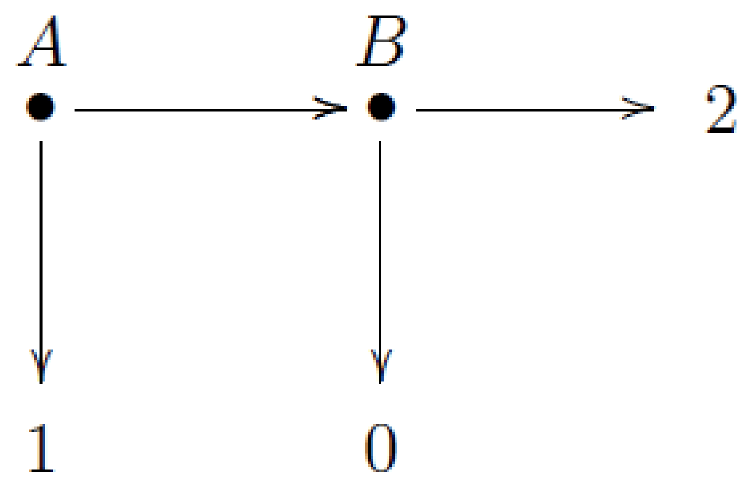

Suppose A is mission control, and has the option of sending into space a specially trained astronaut B who, unfortunately, has been exposed to German measles or a reserve astronaut (payoff 1). If B does not get sick, his mission will be a success (payoff 2), otherwise it will be aborted with failure (payoff 0), cf. Figure 1.

Uncertainty without probabilities?

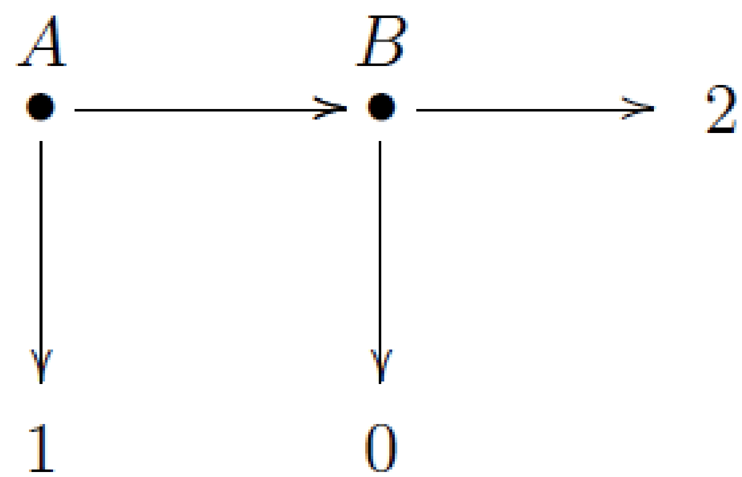

With enough good will, we can apply **Harsanyi's Maximin Postulate** ([12], sections 6.2 and 6.3, postulate A1) here:

> *If you cannot rationally expect more than your maximin payoff, always use a maximin strategy.*

According to our scenario, A can hope, but cannot know for sure, that B does not get sick and delivers payoff 2. Therefore, A has no reason to "rationally expect" more than maximin value 1 when moving *across*, so the rational choice for A is the maximin solution *down*[1].

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?

Non-probabilistic general case

Game Tree:

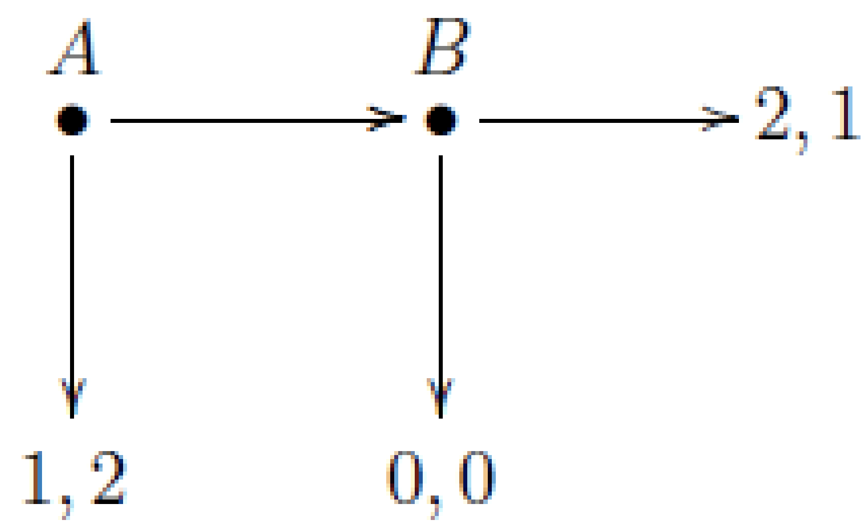

What if (rational) players are not aware of each other rationality and probability distributions?

There were no a canonical answer to this question...

Non-probabilistic general case

Game Tree:

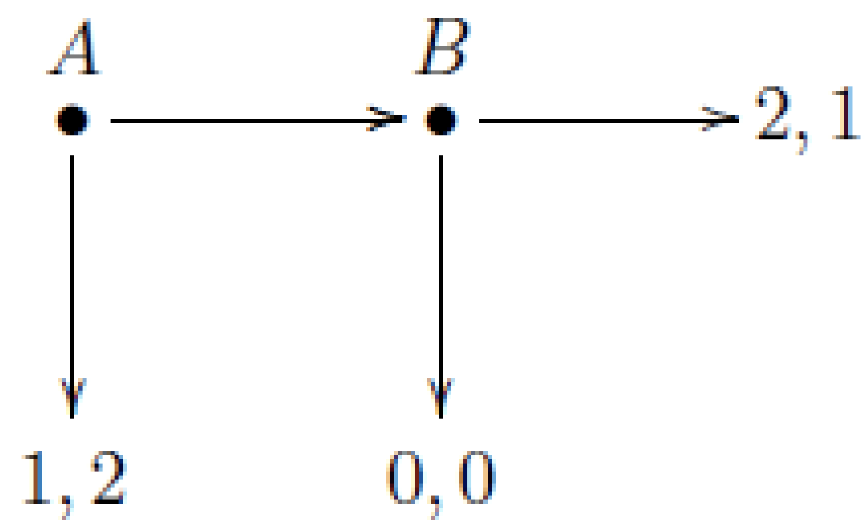

Eliminating strictly dominated strategies.

All four strategies are undominated.

$$\{down_A, down_B\},\ \{across_A, down_B\},\ \{down_A, across_B\},\ \{across_A, across_B\}$$

No definitive answer!

Non-probabilistic general case

Game Tree:

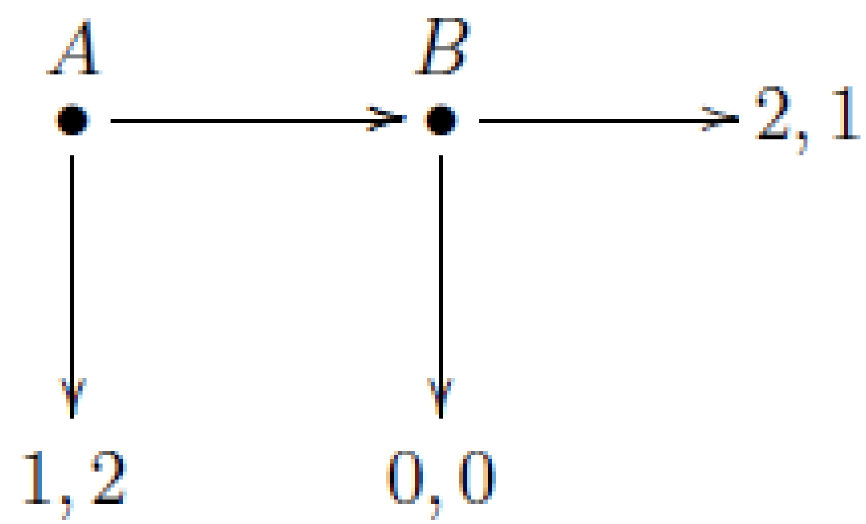

Nash Equilibria:

$$\{down_A, down_B\} \text{ and } \{across_A, across_B\}$$

No definitive answer either!

Harasnyi's Rationality Postulates

Rationality Postulates:

I. *A rational player in perfect information games chooses a maximin solution among all strategies the player deems possible.*

II. *Postulate* (I) *is commonly known and accepted by rational players.*

Postulate I is the epistemically explicit form of Harsanyi's Maximin Postulate. Likewise, (II) is nothing but Harsanyi's Mutually Expected Rationality Postulate expressed in epistemic language.

Harasnyi's Rationality Postulates

Rationality Postulates:

I. *A rational player in perfect information games chooses a maximin solution among all strategies the player deems possible.*

II. *Postulate* (I) *is commonly known and accepted by rational players.*

Postulate I is the epistemically explicit form of Harsanyi's Maximin Postulate. Likewise, (II) is nothing but Harsanyi's Mutually Expected Rationality Postulate expressed in epistemic language.

If a rational player operates in a non-probabilistic setting and bases his decision on knowledge rather than luck, guesswork, sudden opponent cooperation or error, etc., the aforementioned postulates lead to the mathematical model of decision making that we call the Knowledge-Based Rational decision model (*KBR*-model).

Knowledge vs Beliefs

Though Game Theory often considers decisions based on beliefs rather then knowledge, a special theory of **knowledge-based decision making** looks to be appropriate as well. In some situations, players seem to make decisions on the basis of their knowledge and not merely on their beliefs: military, high-stakes commercial, juridical decisions, etc.

KBR theory is not universal, but seems to do the job in non-probabilistic settings.

The principal difference between knowledge and belief is the factivity property of knowledge that beliefs do not necessarily possess.

Highest Known Payoff

KBR suggests following a strategy that yields the highest payoff the agent can secure to the best of his knowledge. Equivalently, within the *KBR* approach, a rational player chooses a maximin solution over all strategies of others the player deems possible. These two seemingly different approaches produce the same result: a maximin choice over the set of all strategies a player considers possible (i.e., that cannot be ruled out as impossible) is a strategy yielding the highest guaranteed payoff to the best of that player's knowledge.

> Indeed, let $\mathbf{m}$ be the maximin payoff at a given node v over the set of all strategies that a player i considers possible. Then i knows a strategy that guarantees him payoff $\mathbf{m}$. On the other hand, for any other payoff $p > \mathbf{m}$, i knows that there is no strategy by i that could guarantee him payoff p. Therefore, $\mathbf{m}$ is the highest payoff that i knows he has a strategy for getting and he cannot rationally expect a payoff greater than $\mathbf{m}$.

Highest Known Payoff

↑

payoffs known to i as not necessarily secured by σ_i

• $HKP_i(v, \sigma_i)$

payoffs known to i as secured by σ_i

↓

Always exists and is unique and known to the corresponding player at each node of the game!

Best Known Strategy, Move

There may be many strategies corresponding to the Highest Known Payoff, e.g., the ones which differ only at unattainable nodes.

However, for generic games (all terminal payoffs are different), all such strategies start with the same move:

there is always a unique Best Known Move.

KBR decision method

Rational players always choose the move which yields the Highest Known Payoff, i.e., the Best Known Move.

In generic perfect information games, each player has one only one *KBR* strategy.

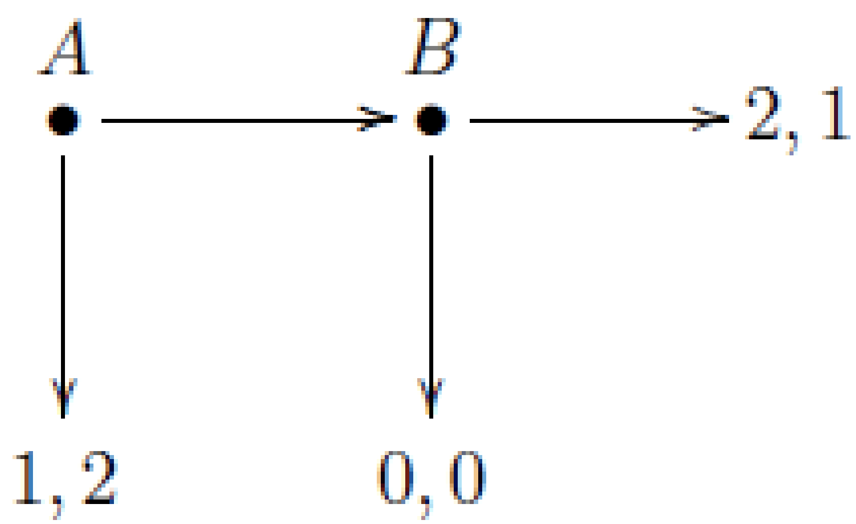

1. *A* knows that *B* is rational:

$$\{across_A, across_B\}$$

2. A does not that B is rational:

$$\{down_A, across_B\}$$

KBR view of the game

A is a rational player. At each *A*-node

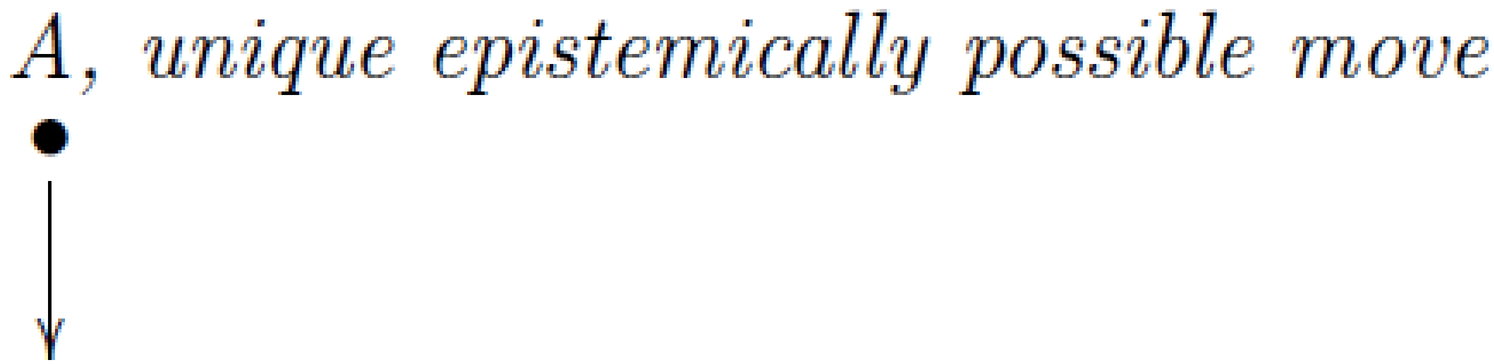

At a *B* node where some other player makes a move

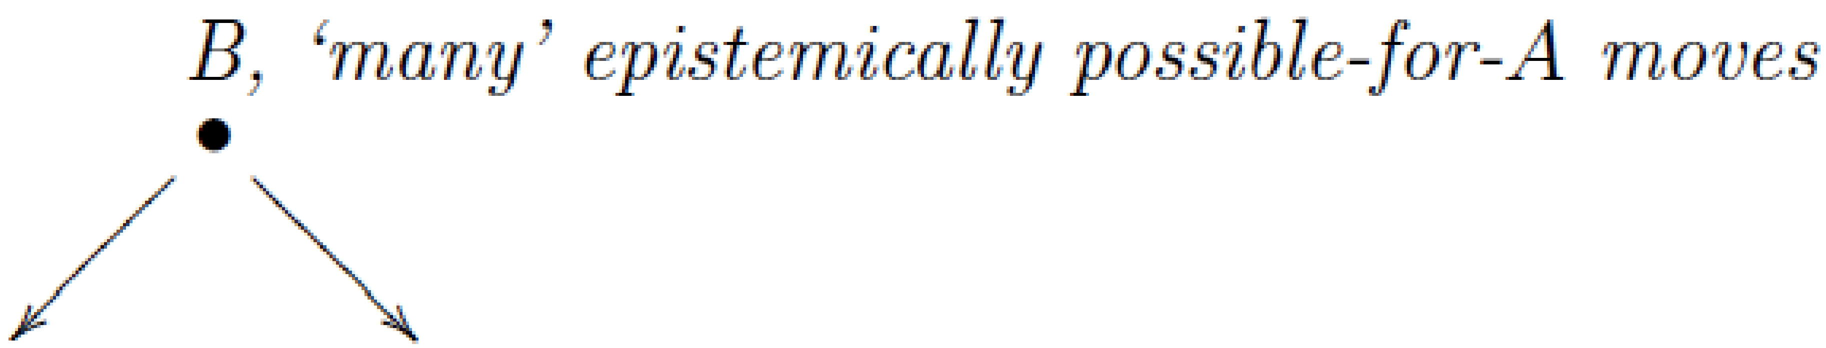

Why *KBR* is so special for PI games?

For perfect information games, KBR is the only decision method which is definitive, rational, and based on knowledge.

Nash Equilibrium - not definitive;
Iterated Undominance - not definitive;
Backward Induction - special case of *KBR*;
Subgame Perfect Equilibrium - special case of *KBR*;
etc.

KBR yields Aumann's rationality

Aumann's approach considers irrational a knowingly dominated strategy. *KBR* also considers such a strategy irrational. Hence,

KBR solution is always Aumann rational as well.

However, Aumann rationality is not definitive!

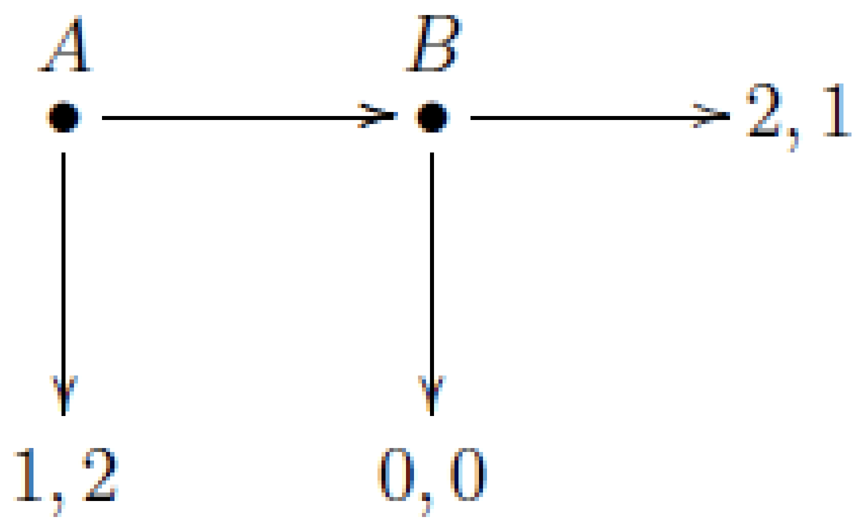

If *A* is not aware of *B*'s rationality, then both strategies for *A* are Aumann's rational whereas only

$$down_A$$

is KBR-rational!

Each *KBR*-path is a Nash path.

This was a (rather surprising) Brandenburger Conjecture: Nash strategy profiles/paths are calculated without any epistemic conditions. However, it turns out that Nash Paths provide a correct approximation to knowledge-based rationality uniformly for all possible epistemic stated of players.

Nash Paths decision method is necessarily non-definitive since it accommodated all KBR solutions for a given Game Tree.

Decision Methods in PI games

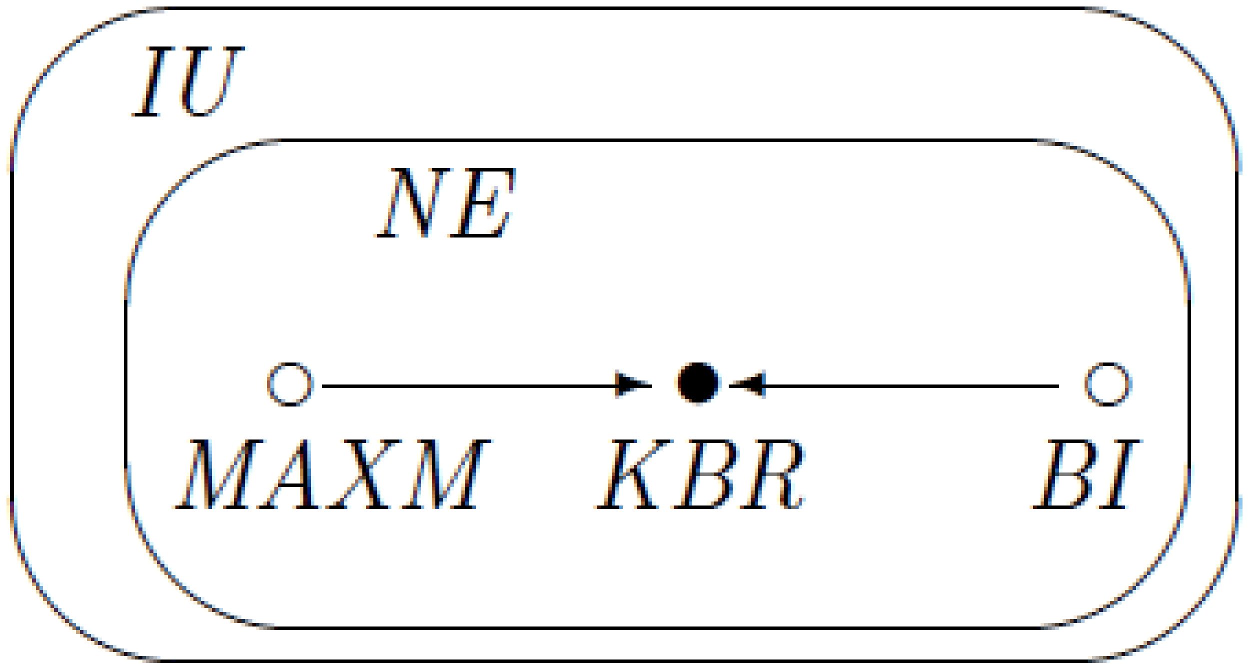

IU - Iterated Undominance paths;
NE - Nash Equilibrium paths;
BI - Backward Induction path;
MAXM - Maximin path;
BKR - Knowledge-based rational path.
Empty circles indicate that the corresponding paths are justified only under additional epistemic conditions.

Active manipulation

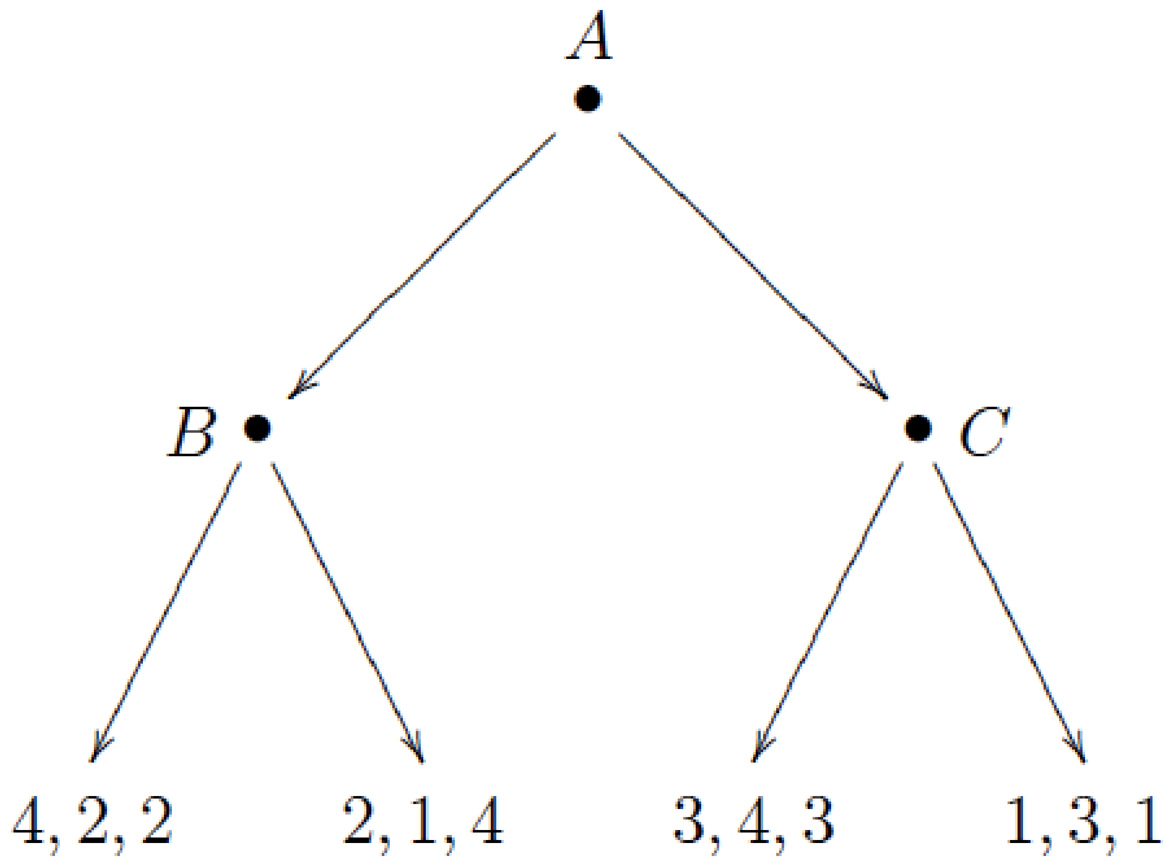

Suppose *A* is not aware of *B* and *C*'s rationality. Then *A* moves *left* to secure payoff 2. Actually, *A* gets 4 which is more than expected. Suppose also that *B* and *C* are smart enough to understand this. Then *B* can manipulate *A* by leaking the true information that ***C is rational***. *A* 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!

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.

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 strategies 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.