Belief-based rational decisions

Sergei Artemov

September 22, 2009

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.

John Nash

Ph.D. from Princeton,
Nobel Prize of 1994,
Mathematical Game Theory:
a system for predicting the outcome of competitive games, which can also be applied to political and economic conflicts such as labor negotiations, business competition, international political tensions, etc.

Robert Aumann

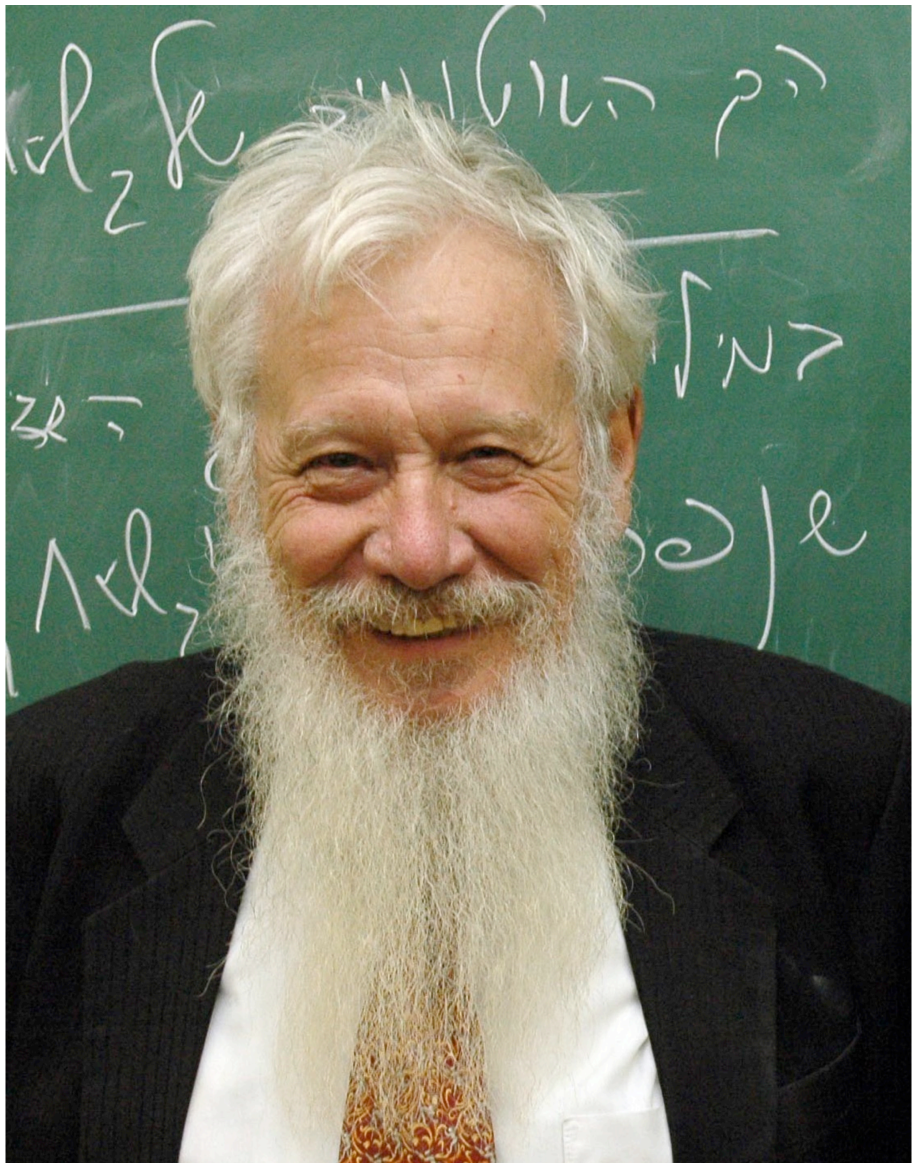

Alma Mater: City College of New York, Nobel Prize of 2005. Pioneered studies of Mathematical theory of ***Rationality and Common Knowledge***, found connections to mathematical logic

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 avoid uncertainty completely;

or

- to deal with uncertainty probabilistically, i.e., when a player ***knows probability distribution*** of all consequences of his actions and is willing to take chances.

Rational decisions, informally

The standard game-theoretical assumption:

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

There was no theory of making decisions under ***uncertainty with unknown probability distribution.***

There is a solution, however, which logically follows from the standard postulates of Game Theory and commonly accepted set of knowledge principles, a.k.a. ***the logic of knowledge*** **S5**.

Rational decisions, informally

Knowledge-Based Rationality models decision-making strictly on the basis of players' knowledge:

at each node, rational players choose the best moves known to them.

New features:

- Clear separation of ***best move*** and ***best known move***.

- ***Players' knowledge*** becomes the key element of game description.

Game Theory

Game Tree:

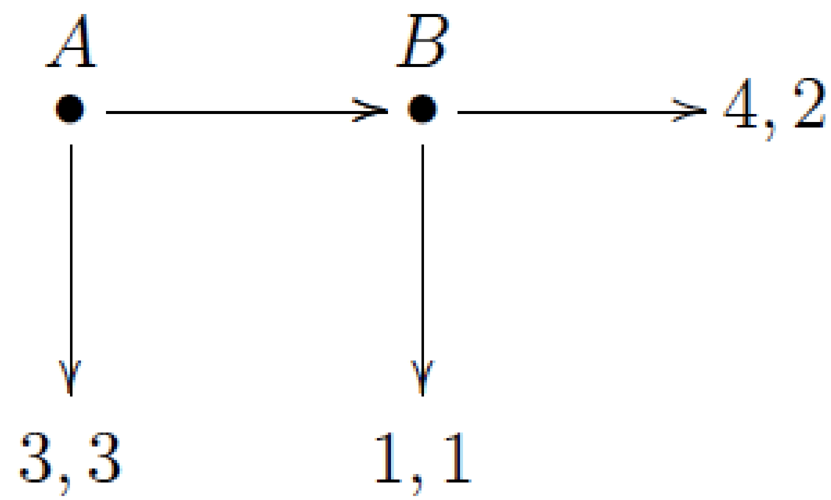

A small company *B* is founded by a scientist who owns a patent. *B* is unable to develop this technology efficiently and hopes to be acquired by a bigger company, *A* (payoffs 3,3). *A* is interested in the patent but not eager to assume responsibility for the entirety of *B*. If *A* refuses to buy *B*, *B* then has a choice to either sell the patent to *A* (payoffs 4,2) or terminate negotiations (payoffs 1,1) and wait for a better offer.

Under normal rational behavior, *A* does not buy the company but purchases the patent. However, if *A* refuses to buy the company, *B* may be dissatisfied and opt to withdraw.

Game Theory (before)

Game Tree:

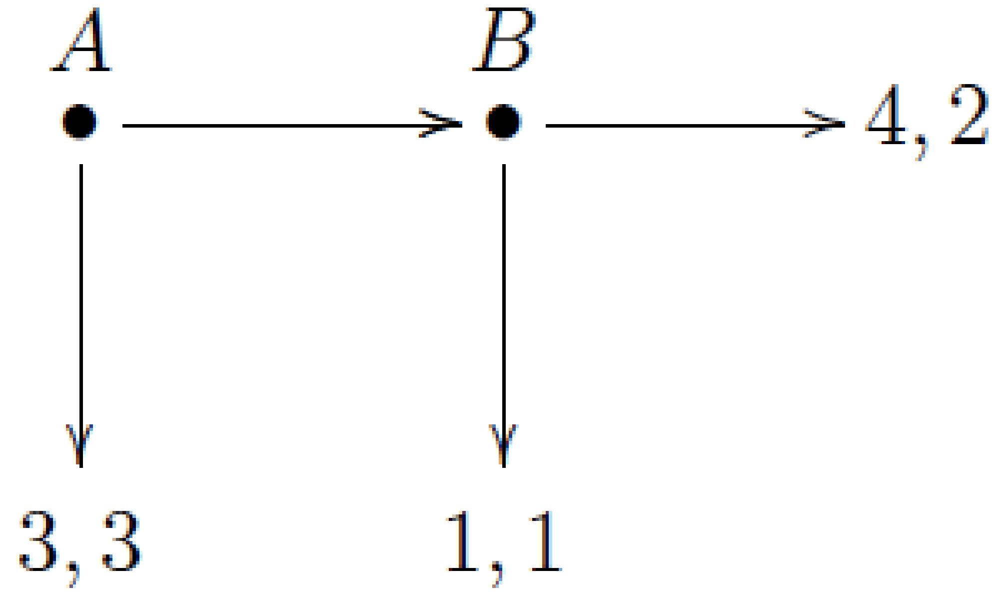

How should this game be played by rational players?

Standard answer: ***assume that A knows that B is rational.***
Then *B* plays *right*. *A* knows this, hence plays *right*, too.
The actual payoffs are 4,2.

But what ***if A is not sure of B's rationality?***
No answer...

Game Theory (after)

Game Tree:

$$
\begin{tikzcd}
\overset{A}{\bullet} \arrow[r] \arrow[d] & \overset{B}{\bullet} \arrow[r] \arrow[d] & 4,2 \\
3,3 & 1,1 &
\end{tikzcd}
$$

Our answer: ***solutions for all states of A's knowledge.***

- ***A is not sure of B's rational behavior*** (e.g., *A* suspects that *B* is angry at *A* for chasing payoff 4 at *B*'s expense): *A* plays *down* since *A's* known payoff there is higher than at *across*. Payoffs 3,3.

- ***A knows that B is rational***: *A* plays *right,* payoffs 4,2.

Game Theory: passive manipulation

Game Tree:

$$\begin{tikzcd} \overset{A}{\bullet} \arrow[r] \arrow[d] & \overset{B}{\bullet} \arrow[r] \arrow[d] & 4,2 \\ 3,3 & 1,1 & \end{tikzcd}$$

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.

Game Theory: active manipulation

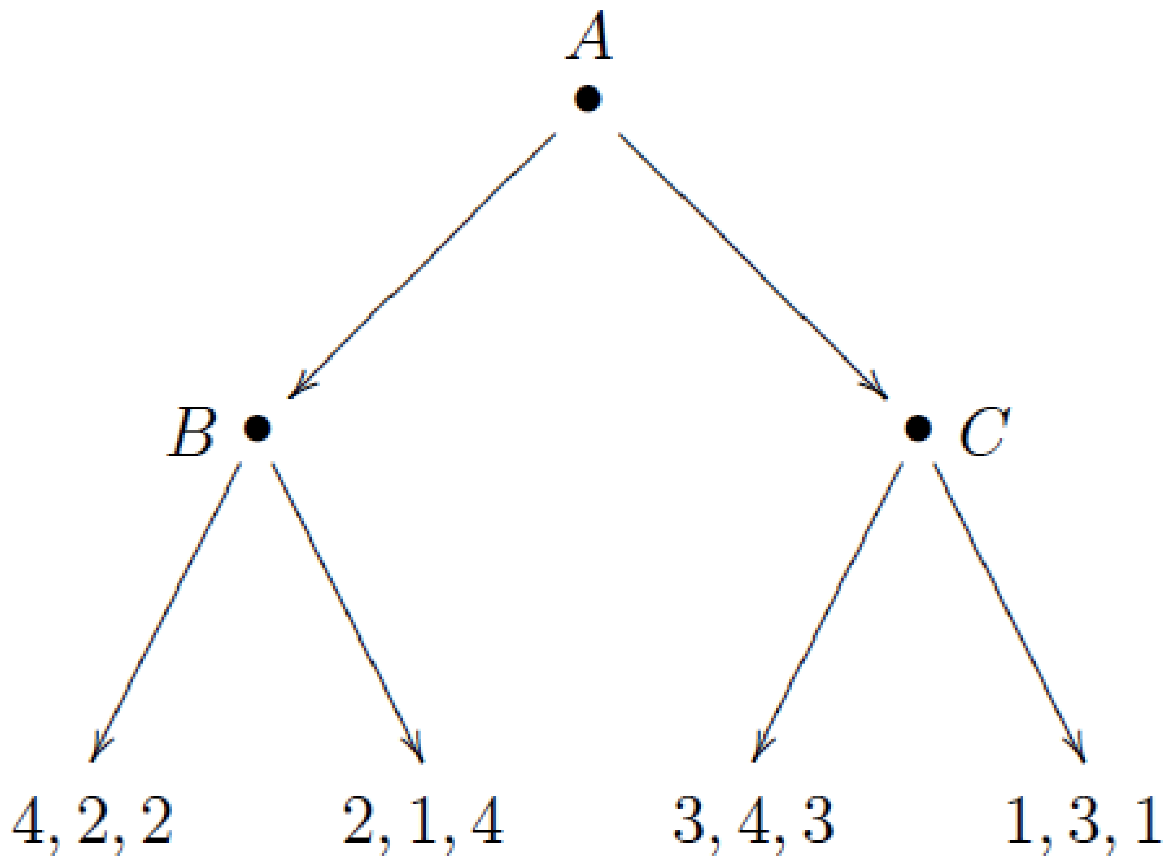

Suppose *A* is not aware of *B* and *C*'s rationality. Then *A* moves *left* to secure payoff 2. Actually, *A* gets 4 which is more than expected. Suppose also that *B* and *C* are smart enough to understand this. Then *B* can manipulate *A* by leaking the true information that ***C* is rational**. *A*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

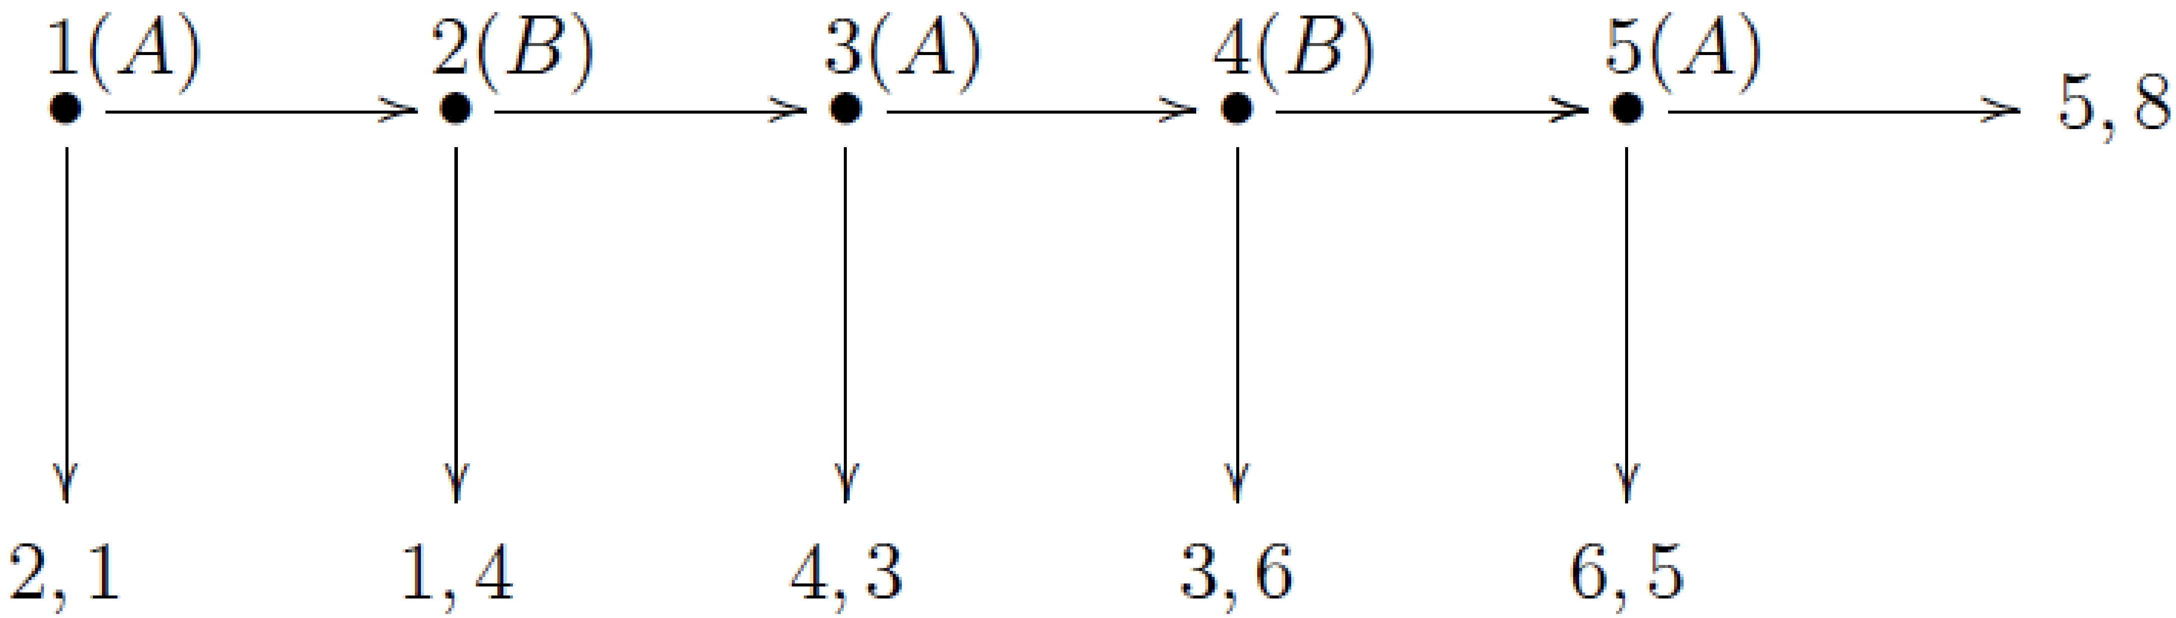

A is rational, hence at node 5, *A*'s choice is *down*.
B knows that A is rational, hence *B* plays *down* at 4.
A knows that B knows that A is rational, hence *A* plays *down* at 3.
B knows that A knows that B knows that A is rational ...

Unbounded nested knowledge of rationality is assumed!

The Centipede game (after)

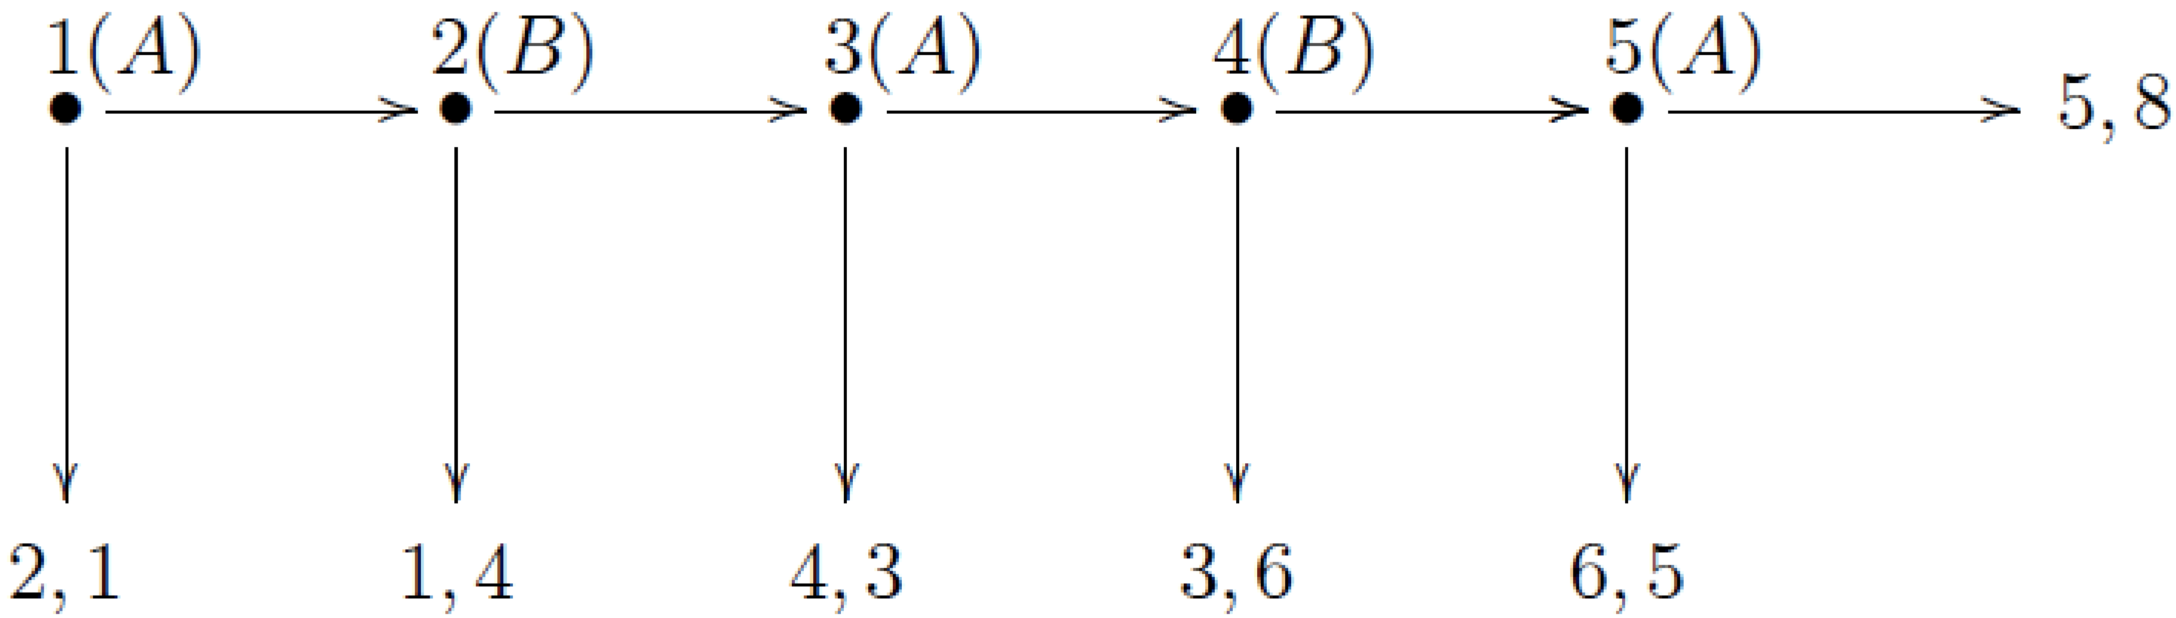

A plays *down* at 5. Consider the latest node where *across* is played (if there is none, we are done). Suppose this is node 1. Since *A* plays *across* at 1, *A* knows that *across* is the better choice, hence *A* knows that *B* plays *across* at 2. But this is impossible, since *B* actually plays *down*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

No knowledge about other players is needed!

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. So, for strictly competitive games,

learning is irrelevant.

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

just do it

normally works.

Belief *vs*. Knowledge check up

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

Logic of Beliefs in Game Theory? **K45D** - the logic of consistent beliefs with positive and negative introspection.

Belief is factive, if

F is believed yields ***F is true***.

For our purposes this is as good as knowledge...

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

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

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,*

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

Highest Believed Payoff, *HBP*:
a similar definition, with belief instead of knowledge.

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 $HKP_A(j)$.*

This also holds, given **knowledge** of basic Math, logic, the (finite) game tree, and **belief** about certain payoffs, not necessarily factive.

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

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

Becomes ***MBB*** = the move which brings ***HBP***.

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

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

The proof does not use factivity, hence fits for beliefs as well.

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

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

The natural formalization of 2 is the principle

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

The natural formalization of 3 is the principle

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

The natural formalization of 4 is the principle

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

The natural formalization of 5 is the principle

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

Generic game: no indistinguishable payoffs for each player.

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

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

Now: commonly believed!

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).*

Theorem 3 *Each perfect information game with rational players who know the game tree has a KBR-solution. Furthermore, if all payoffs are different, then such a solution is unique, each player knows his move at each node, and therefore the game is actually played according to this solution.*

Actually holds for a broader class of games.

Perfect information game - versions:

1. Includes knowledge of the game tree with payoffs, but only beliefs of epistemic states of players.

2. Includes knowledge of the game tree, but allows non-factive beliefs about the payoffs and epistemic states.

Definition 5 Actual Payoff *for a given player X at a given node v,*

$$AP_X(v),$$

is the payoff which X wins if the game is played from v according to the KBR-solution of the game.

This definition stays as is.

It is easy to see that actual payoffs at each node are greater or equal to the best-known payoffs since otherwise, a corresponding player would 'know' the false statement 'he is guaranteed a payoff greater than the one he is actually getting.'

This observation no longer holds since players may be delusional, unjustified wishful thinking, etc.

Belief/Knowledge of the game tree

Main Lemma: Second order belief is knowledge.

Corollary: Self-belief of rationality is factive (knowledge).

Corollary: Common belief of rationality yields common knowledge of rationality.

Aumann's Theorem on Rationality

In PI games, common knowledge of rationality yields backward induction

Remains valid for beliefs as well even with non-factive beliefs about the payoffs as soon as they are commonly believed.