Knowledge-Based Rational Decisions

S. Artemov

Computational Logic Seminar, September 1, 2009

General Picture

We do not invent new rationality principles, but try rather to reveal what was hidden in standard game-theoretical assumptions concerning rational decision-making:

1) the player's rationality yields a payoff maximization given the player's knowledge;

2) the standard logic of knowledge for Game Theory is S5.

It happens that these principles lead to a meaningful mathematical model which we outline in this paper.

What has been done?

1. Within this model, each game has a solution and rational players know which moves to make at each node.

2. Uncertainty in games of perfect information results exclusively from players' different perceptions of the game.

3. In strictly competitive perfect information games, any level of players' knowledge leads to the backward induction solution which coincides with the maximin solution. The same result holds for the well-known centipede game: its standard 'backward induction solution' does not require any mutual knowledge of rationality.

No probability in the picture

Knowledge-Based Rationality is different from other well-known approaches for handling uncertainty in games:

von Neumann & Morgenstern (1944), which assumes known probability distribution;

Savage (1972), which assumes known subjective probability distribution.

The *KBR*-model which we offer does not make any probabilistic assumptions and models decision-making strictly on the basis of players' knowledge.

Rationality: logical format

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

$$rP \quad - \quad \text{‘}P\ \textit{is rational.}\text{’}$$

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.’

Epistemic Logic for Game Theory

In Game Theory, it is usually assumed that knowledge modalities $\mathbf{K}_P$ satisfy postulates of the modal logic of knowledge $\mathsf{S5}$:

Axioms and rules of classical logic;

$\mathbf{K}_P(F \rightarrow G) \wedge \mathbf{K}_P(F) \rightarrow \mathbf{K}_P(G)$, *epistemic closure principle*;

$\mathbf{K}_P(F) \rightarrow F$, *factivity*;

$\mathbf{K}_P(F) \rightarrow \mathbf{K}_P\mathbf{K}_P(F)$, *positive introspection*;

$\neg\mathbf{K}_P(F) \rightarrow \mathbf{K}_P(\neg\mathbf{K}_P(F))$, *negative introspection*;

Necessitation Rule: if F is derived without hypothesis, then $\mathbf{K}_P(F)$ is also derived.

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

$$rA \;\rightarrow\; \mathbf{K}_A(rA). \tag{1}$$

Highest Known Payoff

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,*

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

Correctness of HKP

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 $HKP_A(j)$.*

Baby example

Example 1 Suppose at a given node of the game, move j by A can be met by three responses by his opponent:

Response 1, with A's payoff 10;

Response 2, with A's payoff 20;

Response 3, with A's payoff 30.

Suppose that the actual response is 2, which is not necessarily known to A. So, the *actual payoff* for A at node j is equal to 20.

If A considers all three responses 1, 2, and 3 possible, then $HKP_A(j) = 10$.

If A learns for sure that 1 is no longer possible, then $HKP_A(j) = 20$.

If instead A learns that 3 is no longer possible, then $HKP_A(j) = 10$.

Best Known Move

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

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

Simple Game Tree

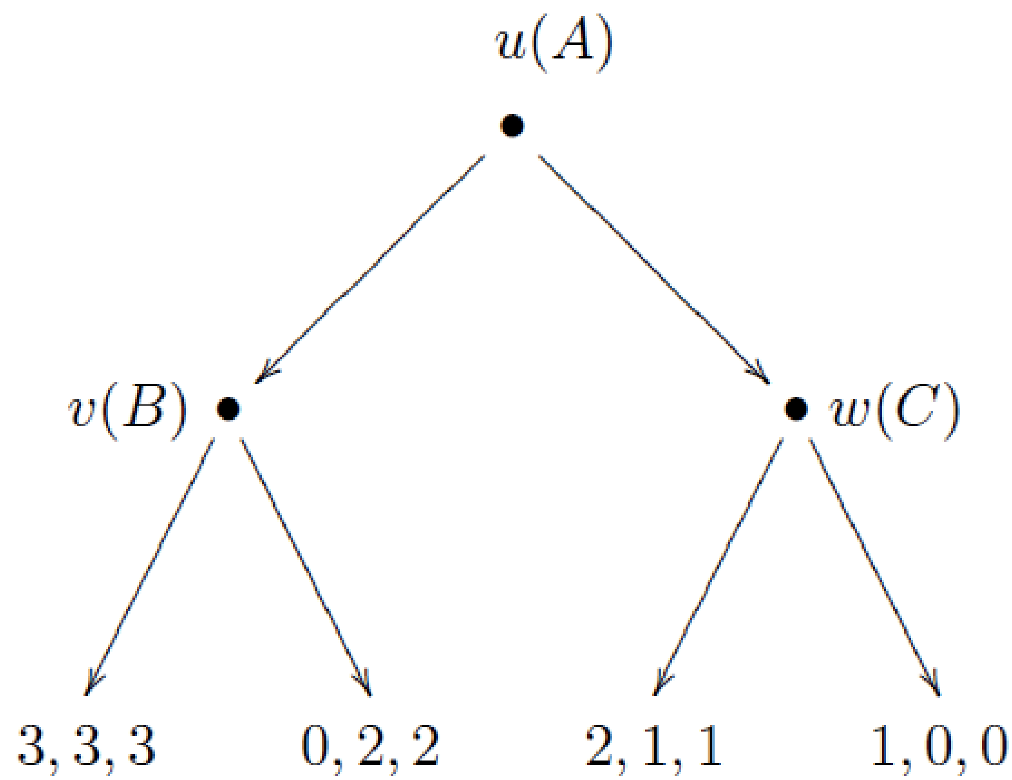

Game I. A is not aware of B's and C's rationality and considers any move for B and C possible.

Game II. A knows that C is rational, but does not know that B is rational.

Game III. A knows that both B and C are rational.

Game I

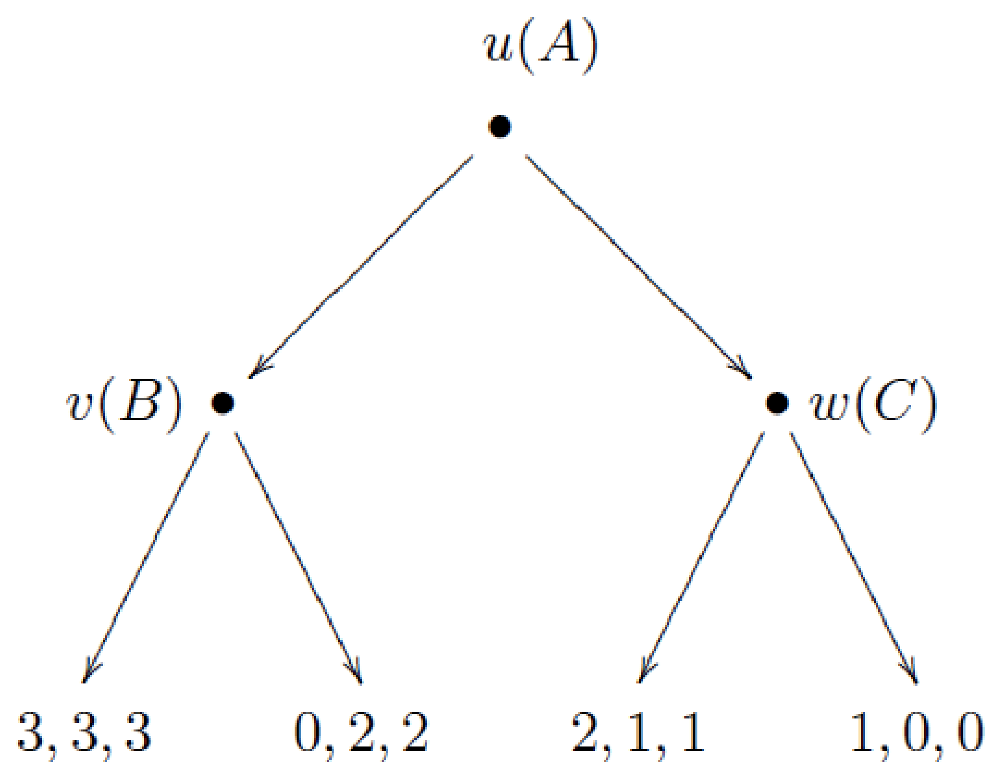

In **Game I**,

$$HKP_A(v) = 0, \qquad HKP_A(w) = 1,$$

$$kbest_A(w)\ .$$

A's actual payoff at u is 2.

Game II

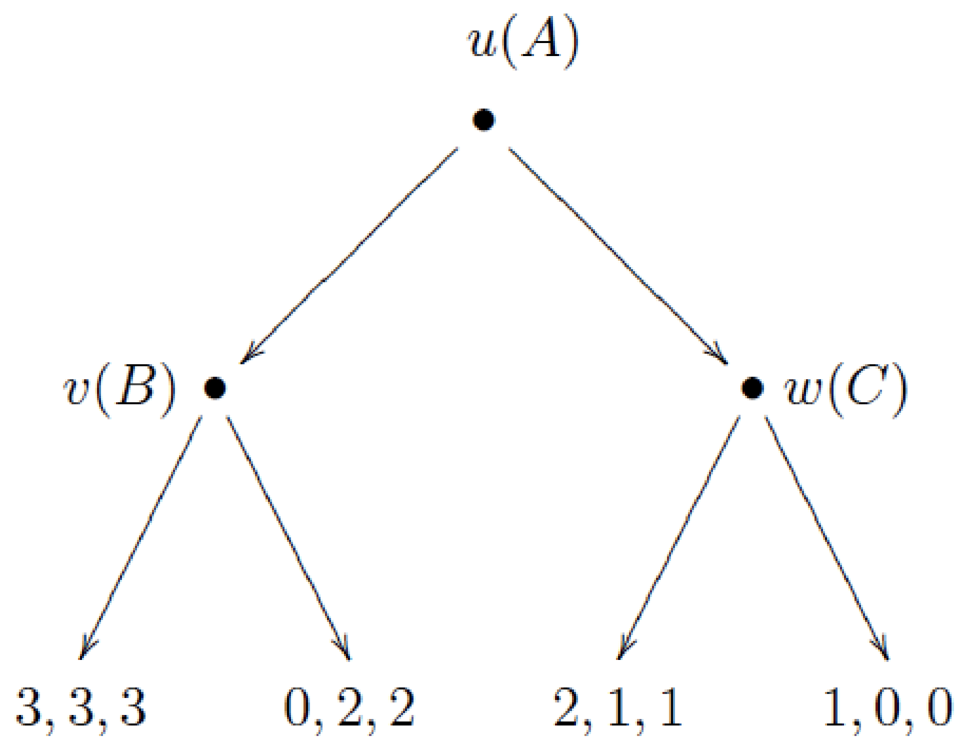

In **Game II**,

$$HKP_A(v) = 0, \qquad HKP_A(w) = 2,$$

$$kbest_A(w) \ .$$

A's actual payoff at u is 2.

Game III

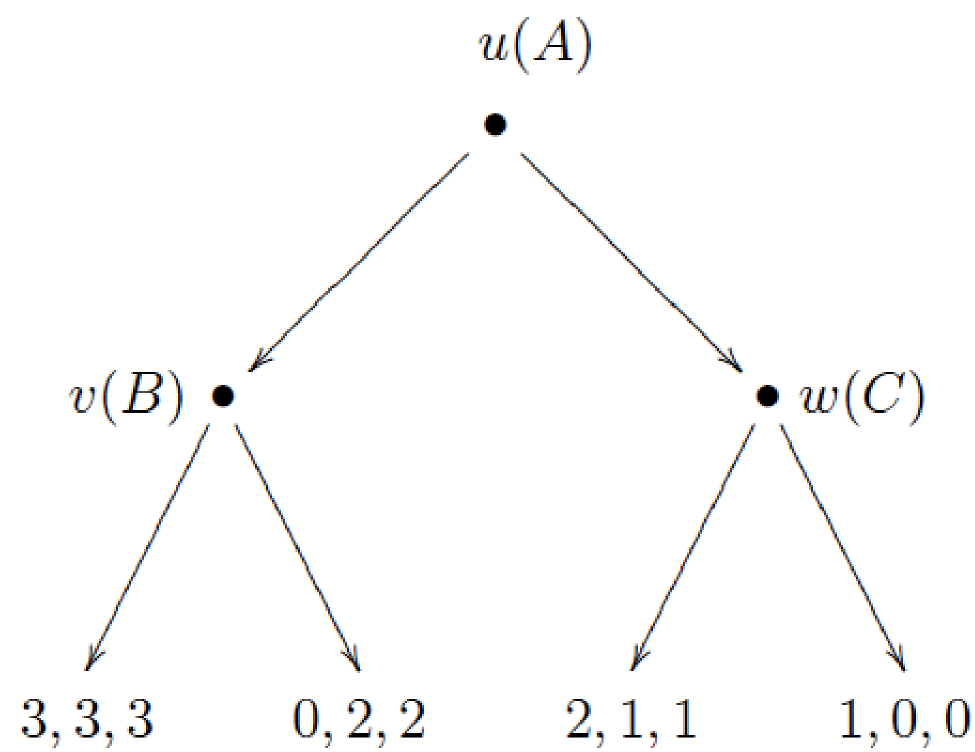

In **Game III**,

$$HKP_A(v) = 3, \qquad HKP_A(w) = 2,$$

$$kbest_A(v) \ .$$

A's actual payoff at u is 3.

Existence Theorem

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

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

Rational decision: verbal accounts

For simplicity's sake, we assume here that all payoffs are different and we work under the assumptions of Corollary 1.

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.

Rational decision: formal accounts

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

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

Rationality Thesis

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).*

There is always a solution...

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.*

Proof. It suffices to check that A **knows** s_j, which describes A's best move. By Rationality Thesis:

$$\mathbf{K}_A\{rA \ \rightarrow \ [kbest_A(j) \rightarrow s_j]\} \ .$$

$$\mathbf{K}_A[rA] \ \rightarrow \ \{\mathbf{K}_A[kbest_A(j)] \rightarrow \mathbf{K}_A[s_j]\} \ .$$

By self-knowledge of rationality, $\mathbf{K}_A[rA]$. Let j be the *KBR*-move. Then,

$$kbest_A(j) \ .$$

By Theorem 1, A's best known move is known to A, hence

$$\mathbf{K}_A[kbest_A(j)] \quad \text{and} \quad \mathbf{K}_A[s_j] \ .$$

Actual Payoffs may be higher...

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.

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.'

Solving Game I

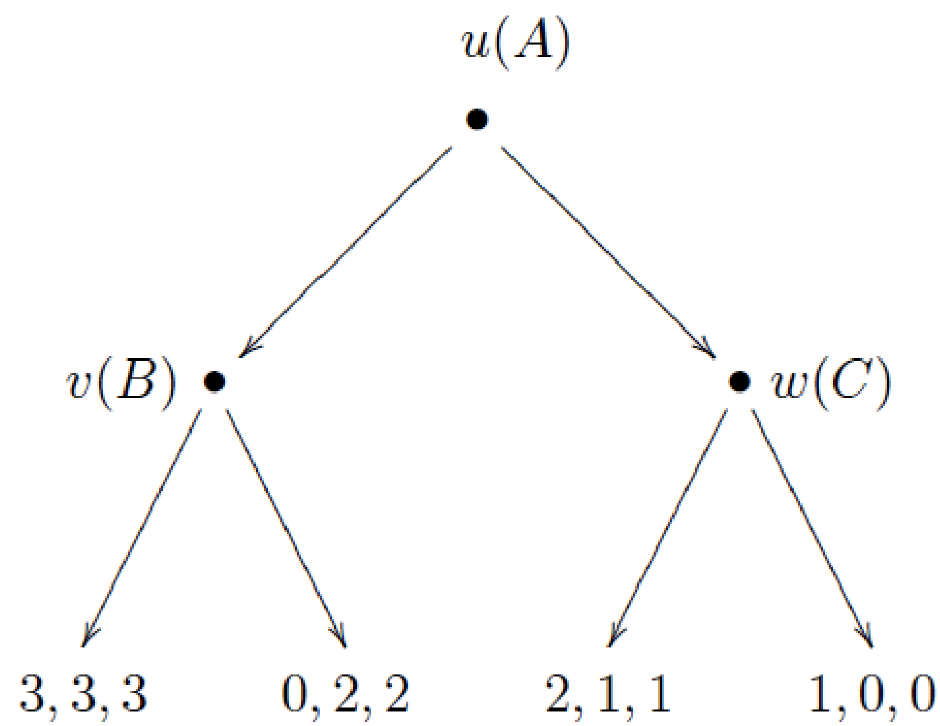

Game I: A is ignorant of B and C's rationality.

$$HKP_A(v) = 0, \quad HKP_A(w) = 1, \quad HKP_A(u) = 1;$$

the KBR-solution: *A plays 'right,' B and C play 'left.'*
Actual payoffs for A, B, and C:

$$AP_{A,B,C}(u) = 2, 1, 1, \quad AP_{A,B,C}(v) = 3, 3, 3, \quad AP_{A,B,C}(w) = 2, 1, 1.$$

Solving Game II

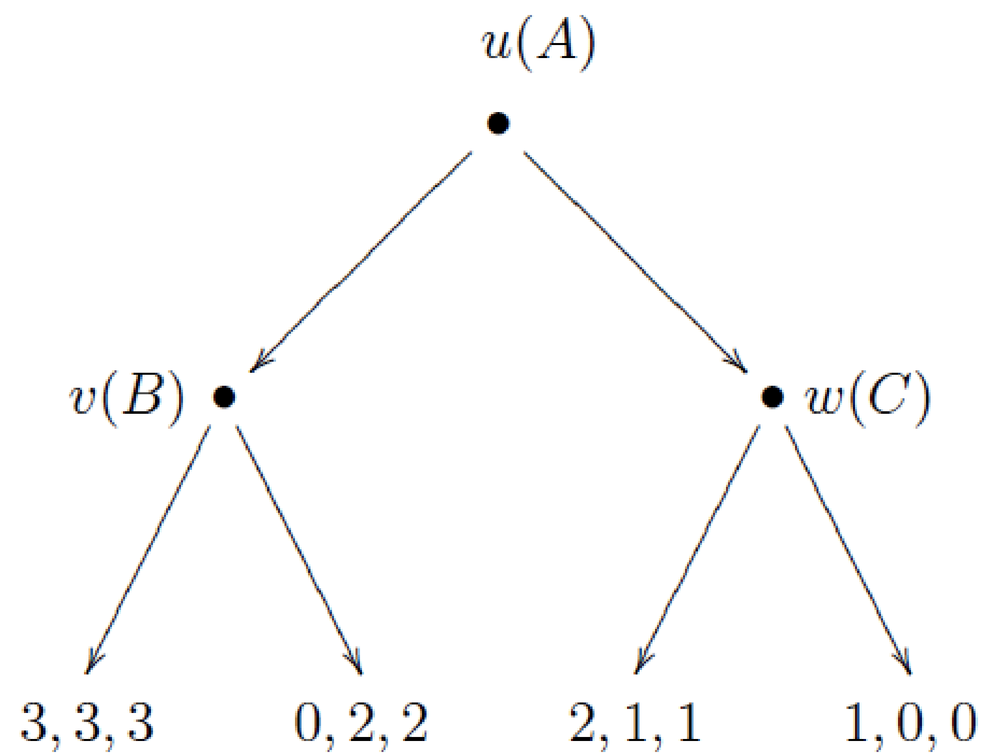

Game II: A knows that C is rational.

$$HKP_A(w) = 2, \quad HKP_A(v) = 0, \quad HKP_A(u) = 2;$$

the same KBR-solution and the same actual payoffs as in Game I.

Solving Game III

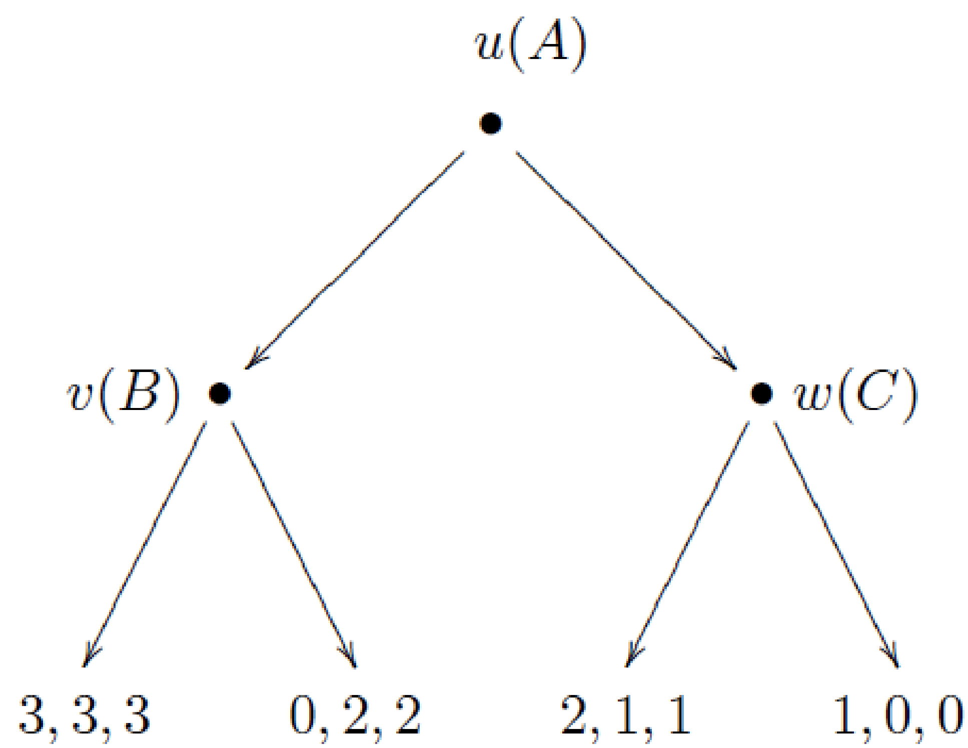

Game III: A knows that both B and C are rational.

$$HKP_A(w) = 2, \quad HKP_A(v) = 3, \quad HKP_A(u) = 3;$$

the KBR-solution: *A, B, and C play 'left.'*

Actual payoffs

$$AP_{A,B,C}(u) = 3,3,3, \quad AP_{A,B,C}(v) = 3,3,3, \quad AP_{A,B,C}(w) = 2,1,1.$$

The Centipede Game

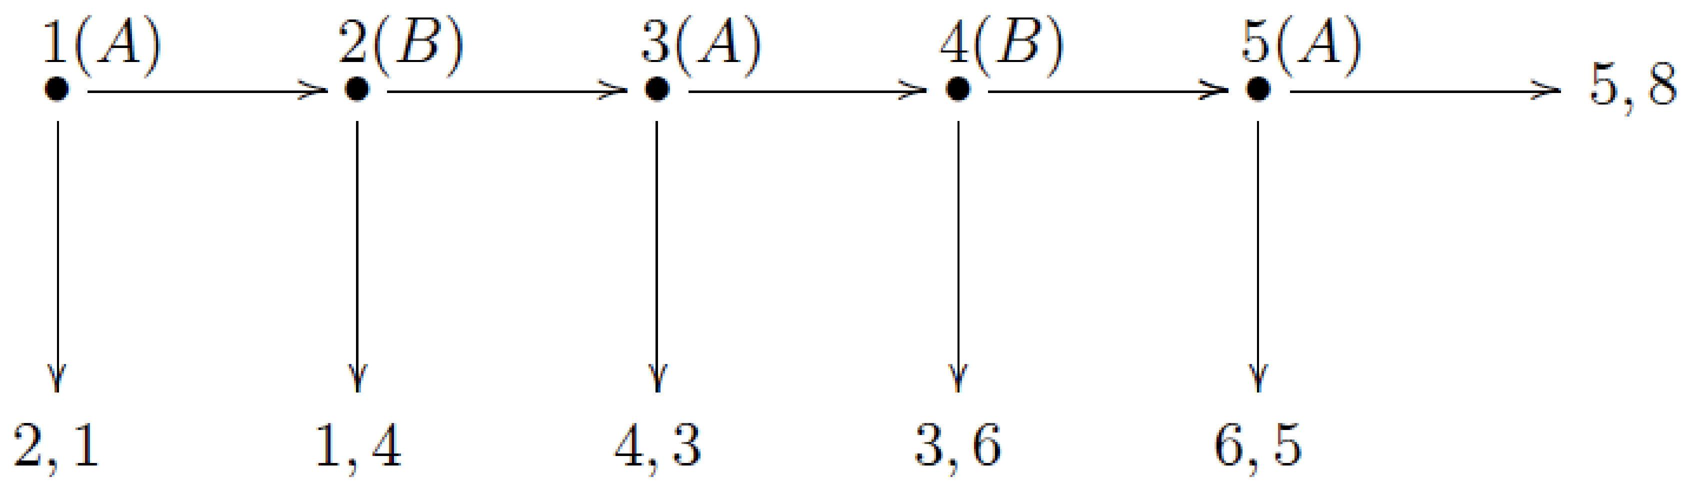

The classic backward induction solution (BI) predicts playing down at each node. Indeed, at node 5, player A's rational choice is down. Player B is certainly aware of this and, anticipating A's rationally playing down at 5, would himself play down at 4. Player A understands this too, and would opt down at 3 seeking a better payoff, etc. The backward induction solution is the unique Nash equilibrium of this game.

Epistemic analysis of Centipede

The question we try to address now is that of finding solutions for the centipede game under a reasonable variety of epistemic assumptions about players A and B. We assume common knowledge of the game tree and concentrate on tracking knowledge of rationality. The classical analysis states that it takes common knowledge of players' rationality (or, at least, as many levels of knowledge as there are moves in the game) to justify backward induction in perfect information games, with the centipede game serving as an example. We will try to revise the perception that stockpiling of mutual knowledge assumptions are needed for solving the centipede game.

KBR *vs.* BI

There is a unique *KBR*-solution to the centipede game for each set of epistemic states of players. We show that each of them leads to the backward induction solution: players choose 'down' at each node.

Within the *BI*-solution, the players actually **avoid** making decisions under uncertainty by assuming enough knowledge of rationality to **know** exactly all the opponent's moves. In the *KBR*-solution, the players make decisions under uncertainty by calculating their highest known payoffs and determining their best moves. So the *BI*-solution is a special extreme case of the *KBR*-solution. For the centipede game, however, both methods bring the same answer: playing down at each node.

No extra knowledge is needed

Consider a natural formalization of the centipede game in an appropriate epistemic modal logic with two agents A and B and rationality propositions rA and rB.

$rA = A$ *is rational,*

$rB = B$ *is rational,*

$a_i =$ *'across' is chosen at node i,*

$d_i =$ *'down' is chosen at node i.*

Theorem 4 *In the centipede game, under any states of players' knowledge, the KBR-solution coincides with the BI-solution, hence rational players play the backward induction strategy.*

Just don't do anything stupid

Proof. The proof consists of calculating the best known move at each node. Note that since epistemic states of players at each node do not contain false beliefs, the actual moves of players are considered possible, otherwise a corresponding player would have a false belief that some actual move is impossible.

Node 5, player A. Obviously,

$$kbest_A(\text{'}down\text{'}) \text{ holds at node 5.}$$

Indeed, A knows that playing 'down' yields 6, whereas playing 'across' yields 5. Since A is rational, d_5.

Node 4, player B. Obviously $HKP(\text{'}down\text{'}) = 6$. On the other hand, $HKP(\text{'}across\text{'}) = 5$, since B considers d_5 possible. If B would deem d_5 impossible, B would know $\neg d_5$, which is false and hence cannot be known. Therefore

$$kbest_B(\text{'}down\text{'}) \text{ holds at node 4.}$$

Since B is rational, d_4.

Just don't do anything stupid

Since A is rational, d_3.

Node 2, player B. $HKP(\text{'}down\text{'}) = 4$, $HKP(\text{'}across\text{'}) = 3$, since B considers d_3 possible. Hence

$$kbest_B(\text{'}down\text{'}) \text{ holds at node 2.}$$

Since B is rational, d_2.

Node 1, player A. $HKP(\text{'}down\text{'}) = 2$, $HKP(\text{'}across\text{'}) = 1$, since A considers d_2 possible. Hence

$$kbest_A(\text{'}down\text{'}) \text{ holds at node 1.}$$

Since A is rational, d_1. $\square$

Just don't do anything stupid

In this solution, the players calculate their best known moves without using any epistemic assumptions about other players. It so happens that this *KBR*-solution coincides with the *BI*-solution, since the worst-case in the centipede game is exactly the *BI*-choice at each node.

This theorem establishes that in the centipede game, the level of knowledge of players does not matter: any states of knowledge of players lead to the same solution, 'down at each node.'

Strictly competitive games

A two-person game is called *strictly competitive* if for any two possible outcomes (histories) X and Y, player A prefers X to Y if and only if player B prefers Y to X. Using standard notation (cf., for example, [18]) for preference relation of player P, $\precsim_P$, we can present this as

$$X \precsim_A Y \quad \Leftrightarrow \quad Y \precsim_B X \; . \tag{10}$$

Since possible outcomes in extensive-form games are normally associated with payoffs at terminal nodes, we can reformulate (10): for each possible outcomes m_1, n_1 and m_2, n_2,

$$m_1 \leq m_2 \quad \Leftrightarrow \quad n_2 \leq n_1 \; . \tag{11}$$

Theorem 5 *In strictly competitive games of perfect information, under any states of players' knowledge, the KBR-solution coincides with the maximin solution and with the BI-solution.*

Nested knowledge matters!

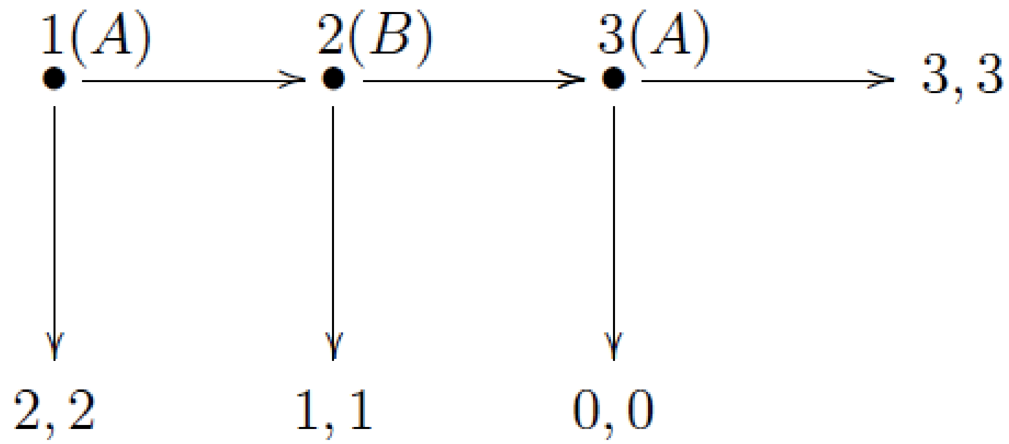

Game I: both players are rational

$$rA \quad \text{and} \quad rB \tag{12}$$

but neither $\mathbf{K}_A(rB)$ nor $\mathbf{K}_B(rA)$ hold: both A and B consider possible any move by their opponent at any node. Let $\Diamond_P(F)$ stand for $\neg\mathbf{K}_P(\neg F)$. Then

$$\Diamond_B(a_3) \wedge \Diamond_B(d_3) \quad \text{and} \quad \Diamond_A(a_2) \wedge \Diamond_A(d_2)\ . \tag{13}$$

Game I is defined by the game tree in Figure 1, rationality of players (12), and epistemic conditions (13).

Solving Game I

$$
\begin{tikzcd}
\overset{1(A)}{\bullet} \arrow[r] \arrow[d] & \overset{2(B)}{\bullet} \arrow[r] \arrow[d] & \overset{3(A)}{\bullet} \arrow[r] \arrow[d] & 3,3 \\
2,2 & 1,1 & 0,0 &
\end{tikzcd}
$$

As a rational player, A plays 'across' at node 3. However, at node 2, B considers it possible that A plays 'down' at 3. Therefore,

$$HKP_B(\text{`}across\text{'}) = 0 \ , \text{ whereas } \ HKP_A(\text{`}down\text{'}) = 1 \ ,$$

and B chooses 'down' at 2. Likewise, by (13), A considers either of a_2 and d_2 possible. Therefore, at root node 1, A chooses 'down.' The solution of the game is

$$d_1, \ d_2, \ a_3 \ .$$

Game II

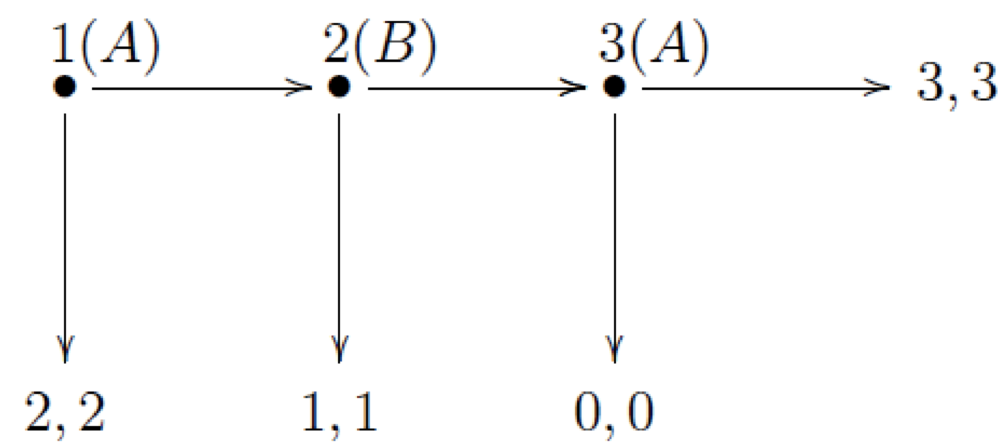

Game II: level 1 mutual knowledge of rationality is assumed:

$$\mathbf{K}_A(rB) \quad \text{and} \quad \mathbf{K}_B(rA) \ , \tag{14}$$

but neither $\mathbf{K}_B\mathbf{K}_A(rB)$ nor $\mathbf{K}_A\mathbf{K}_B(rA)$. In particular, A does not know that B knows that A plays 'across' at 3, or, symbolically,

$$\Diamond_A \Diamond_B (d_3) \ . \tag{15}$$

Game II is defined by epistemic conditions (14) and (15).

Solving Game II

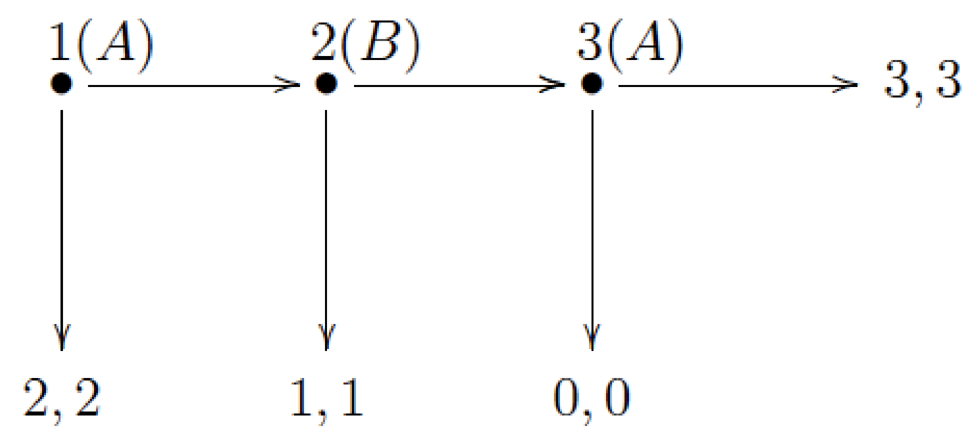

A plays 'across' at node 3, hence a_3. B knows that A, as a rational player, chooses 'across' at 3, i.e., B knows that a_3. As a rational player, B chooses 'across' at 2, hence a_2.

At node 1, A considers it possible that B considers d_3 possible. A knows that B is rational, hence A considers it possible that B plays 'down' at 2:

$$HKP_A(\text{'}across\text{'}) = 1 \text{ , whereas, by the game tree, } HKP_A(\text{'}down\text{'}) = 2 \ .$$

As a rational player, A chooses 'down,' hence d_1. The solution of Game II is represented as

$$d_1, \ a_2, \ a_3 \ .$$

Game III

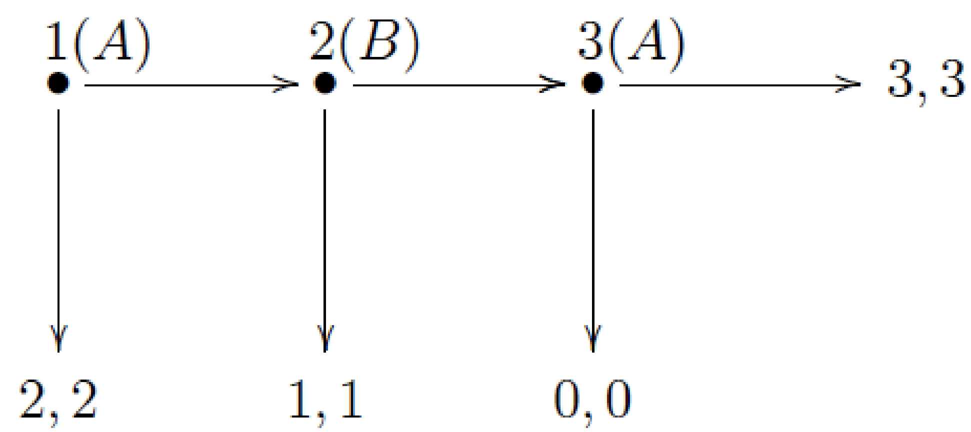

Game III: Common knowledge of rationality is assumed. This level of knowledge is already sufficient for backward induction reasoning. Indeed, A plays 'across' at node 3, B knows that A as a rational player chooses 'across' at 3, hence B chooses 'across' at 2. A knows that B knows that A plays 'across' at 3, hence A knows that the best known move for B is 'across.' Moreover, since A knows that B is rational, A knows that B plays 'across' at node 2. Therefore, the best known move for A at 1 is 'across,' hence A chooses 'across' at 1. The solution of Game III is

$$a_1,\ a_2,\ a_3\ .$$

Anti-Centipede Game

It is clear how to generalize the **anti-centipede game** to any finite length in such a way that a shift from solution 'down' to solution 'across' at node 1 happens only at the nested depth of mutual knowledge of rationality which is equal to the length of the game.

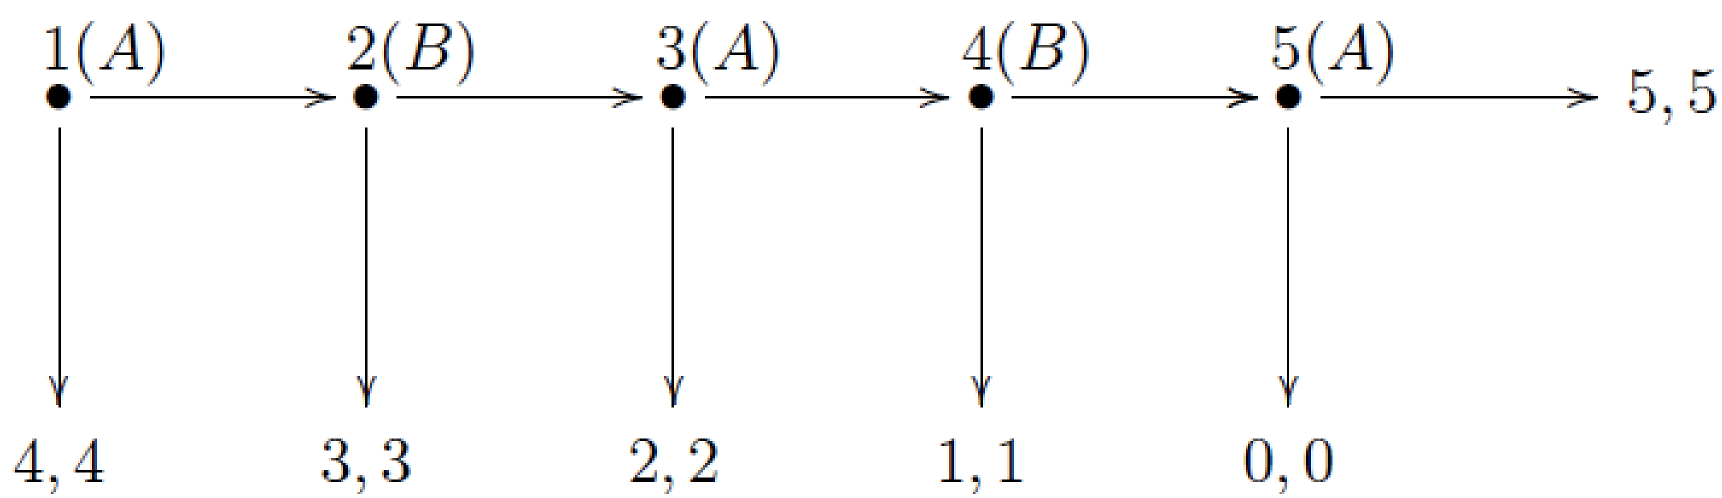

Knowledge of the Game

We claim that in Games I and II, players A and B do not have knowledge of the corresponding game in its entirety. Indeed, the complete description of a game includes

1) a *Game Tree*, which is commonly known;

2) *Rationality*: propositions rA and rB are assumed true (but not necessarily assumed mutually known).

3) *Epistemic Conditions* $\mathcal{E}$ describing what is specifically known by players, in addition to general knowledge from 1 and 2.

Knowledge of the game consists of knowing 1, 2, 3 and basic mathematical facts, together with whatever follows from them in the logic of knowledge. E.g., each player knows that he is rational: $\mathbf{K}_A(rA)$, $\mathbf{K}_A\mathbf{K}_A(rA)$, etc.

Knowledge of the Game I

$$\begin{tikzcd}
\overset{1(A)}{\bullet} \arrow[r] \arrow[d] & \overset{2(B)}{\bullet} \arrow[r] \arrow[d] & \overset{3(A)}{\bullet} \arrow[r] \arrow[d] & 3,3 \\
2,2 & 1,1 & 0,0 &
\end{tikzcd}$$

In Game I, $\mathcal{E}$ contains condition $\neg\mathbf{K}_B(a_3)$. From this we can logically derive $\neg\mathbf{K}_B(rA)$.

As we see, A knows rA, since $\mathbf{K}_A(rA)$ holds, but B does not know rA, since $\mathbf{K}_B(rA)$ does not hold. Therefore, A and B have a **different understanding of Game I**, and B's knowledge is not complete.

Knowledge of the Game II

$$\begin{tikzcd}
\overset{1(A)}{\bullet} \arrow[r] \arrow[d] & \overset{2(B)}{\bullet} \arrow[r] \arrow[d] & \overset{3(A)}{\bullet} \arrow[r] \arrow[d] & 3,3 \\
2,2 & 1,1 & 0,0 &
\end{tikzcd}$$

In Game II, proposition $\mathbf{K}_A\mathbf{K}_B(rA)$ does not hold. On the other hand, we conclude

$$\mathbf{K}_B\mathbf{K}_B(rA) \ ,$$

by positive introspection of $\mathbf{K}_B$. Therefore, B knows $\mathbf{K}_B(rA)$. However, A does not know $\mathbf{K}_B(rA)$, since $\mathbf{K}_A\mathbf{K}_B(rA)$ does not hold. Again, players A and B have different accounts of the rules of Game II.

Knowledge of the Game III

Game III is mutually known to its players A and B in its entirety because the game description is common knowledge. Indeed, in Game III, the complete description includes

1) the *Game Tree* (commonly known);

2) *Rationality*: rA and rB;

3) *Epistemic Conditions*: $\mathcal{E}$ = *Common Knowledge of Rationality*.

Since, for each player P,

$$\textit{Common Knowledge that } F \;\rightarrow\; \mathbf{K}_P(\textit{Common Knowledge that } F)\,, \tag{16}$$

A's and B's knowledge of Game III is complete. Indeed, A and B each know the *Game Tree*, which is common knowledge. A and B also know *Rationality*, which is common knowledge. Finally, A and B both know *Epistemic Conditions* $\mathcal{E}$ because of (16).

Full knowledge is power

Proposition 2 *Any intelligent agent (observer) who knows the game in full, knows the KBR-solution of the game and actual payoffs.*

Indeed, suppose B's best known move is j, then B concludes

$$\text{`}\textit{Epistemic State of B}\text{'} \rightarrow kbest_B(j) \ .$$

The laws of logic are known to each intelligent agent, hence:

$$\mathbf{K}_A[\text{`}\textit{Epistemic State of B}\text{'} \ \rightarrow \ kbest_B(j)] \ ,$$

$$\mathbf{K}_A[\text{`}\textit{Epistemic State of B}\text{'}] \ \rightarrow \ \mathbf{K}_A[kbest_B(j)] \ .$$

By Rationality Thesis,

$$\mathbf{K}_A[rB \ \rightarrow \ (kbest_B(j) \rightarrow s_j)] \ ,$$

$$\mathbf{K}_A[rB] \ \rightarrow \ (\mathbf{K}_A[kbest_B(j)] \rightarrow \mathbf{K}_A[s_j]) \ .$$

Since A knows the epistemic state of B and *Rationality*, $\mathbf{K}_A[s_j]$.

What is the game we are playing?

Definition 7 *We say that A is certain at a given node if A knows KBR-solutions for subgames at each subsequent node.*

Corollary 2 *Any player who knows the game in full is certain at each node of the game.*

Proposition 3 *In a PI game, certainty at each node yields the BI-solution which coincides with the KBR-solution.*

Proof. By backward induction. At pre-terminal nodes, all players move rationally, which is both a *KBR*- and *BI*-solution. If a player at a given node v knows *KBR*-solutions at all later nodes, he knows actual payoffs and his *KBR*-move at v coincides with the *BI*-move. □

Uncertainty in PI games occurs only because **players do not know the game in full**.

Are we playing the same game?

Definition 8 *For players A and B, by* **iterated rationality** *assertions $\mathcal{IR}$, we understand the set of propositions of the sort 'A knows that B knows that A knows … that A is rational':*

$$
\begin{aligned}
\mathcal{IR} = \{ & rA,\ rB, \\
& \mathbf{K}_B(rA),\ \mathbf{K}_A(rB), \\
& \mathbf{K}_A\mathbf{K}_B(rA),\ \mathbf{K}_B\mathbf{K}_A(rB), \\
& \mathbf{K}_B\mathbf{K}_A\mathbf{K}_B(rA),\ \mathbf{K}_A\mathbf{K}_B\mathbf{K}_A(rB), \\
& \mathbf{K}_A\mathbf{K}_B\mathbf{K}_A\mathbf{K}_B(rA),\ \mathbf{K}_B\mathbf{K}_A\mathbf{K}_B\mathbf{K}_A(rB), \\
& \ldots\}.
\end{aligned}
$$

This definition naturally extends to more than two players.

Theorem 6 *If all players in a PI game are rational and have the same knowledge of iterated rationality, then there is no uncertainty in the game.*

Are we playing the same game?

Lemma 3 *If players have the same knowledge of iterated rationality, then each of them knows the whole set $\mathcal{IR}$ of iterated rationality assertions.*

Indeed, since A is rational, rA holds. Since the rationality of A is self-known, $rA \to \mathbf{K}_A(rA)$,

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

By assumptions, B knows the same $\mathcal{IR}$ assertions as A, in particular,

$$\mathbf{K}_B(rA) \ .$$

By positive introspection of B's knowledge,

$$\mathbf{K}_B\mathbf{K}_B(rA) \ ,$$

which means that $\mathbf{K}_B(rA)$ is known to B; by assumptions, it is known to A as well:

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

etc.

Are we playing the same game?

Now proceed with the usual backward induction reasoning to show that at each node, the player knows all KBR-moves and actual payoffs at all later nodes. At pre-terminal nodes, the players move rationally according to the Game Tree. At the next nodes moving towards the root, players determine the moves of players at the previous nodes using level 1 mutual knowledge of rationality. At the next layer of nodes towards the root, players use level 2 mutual knowledge of rationality to determine all the moves at the successor nodes, etc. The only epistemic condition which is needed for the backward induction reasoning at a node of depth n is level n mutual knowledge of rationality, which is guaranteed by Lemma 3.

Rationality assertions are special

It follows from the proof that to achieve complete certainty in a given game of length n, it is sufficient for players to agree on a finite set of iterated rationality assertions with nested knowledge depth not exceeding n. Such an agreement is only possible when all iterated rationality assertions of nested knowledge depth not exceeding n are actually known to all players. We can formulate the same observation in a dual manner: if a player faces uncertainty in a perfect information game, then there should be an iterated rationality assertion of nested depth not exceeding the length of the game, which is unknown to the player.

Aumann's Theorem on Rationality

> *In PI games, common knowledge of rationality yields backward induction*

easily follows from Theorem 6 and Proposition 3. Indeed, common knowledge of rationality immediately yields that each player knows the whole set of iterated rationality assertions $\mathcal{IR}$.

Altogether, Corollary 2 and Theorem 6 reveal that different and incomplete knowledge of the game form the basis for uncertainty in perfect information games. If uncertainty occurs in a perfect information game, players have different knowledge of the game. The player who faces uncertainty does not have complete knowledge of the game.

Other Epistemic Models

The logic of knowledge approach adopted in this paper provides a flexible and competitive apparatus for specifying and solving games. It has certain advantages over other well-known approaches for tracking epistemic conditions in games, such as protocols and possible paths ([14]), and set-theoretical Aumann structures ([4]). In particular, logical language can deal with incomplete specifications of (possibly infinite) state spaces, which are yet sufficient for solving the game. The aforementioned model-theoretical and set-theoretical approaches, on the other hand, require *a priori* complete specification of state spaces, which may happen to be too hard if at all possible.

What do we actually assume?

We offer a specific, logic-based approach. In our model, we try to accommodate the intellectual powers of players who are considered not to be mere finite-automata payoff maximizers but rather intellectual agents capable of analyzing the game and calculating payoffs conditioned to the rational behavior of all players. In particular, we assume that players have common knowledge of the laws of logic, foundations of knowledge-based rational decision making, and that they follow these principles. We believe that such assumptions about the intellectual powers of players are within the realm of both epistemic and game-theoretical reasoning.

Do we need S5?

The full power of the logic S5 was used in Theorem 3, which states that a *KBR*-solution always exists and that rational and intelligent players follow this solution. However, in specific games, *KBR*-solutions can be logically derived by much more modest epistemic means. For example, in Game I of Section 8, it suffices to apply negative introspection to epistemic conditions (13) to derive the *KBR*-solution and to conclude that players will follow this solution. Roughly speaking, it suffices to add to the game specification that epistemic conditions (13) are known to corresponding players and to reason in the logic S4, which is S5 without the negative introspection principle. These considerations could appeal to epistemologists and modal logicians who might have reservations concerning the use of powerful epistemic principles such as negative introspection. Using S4 has some additional advantages, e.g., it renders the reasoning monotonic in a logical sense, admits natural evidence analysis in the style of [1, 2] where one could hope to produce verified best known strategies for players, etc.

Learning

It is obvious that more player knowledge at a given node yields a greater highest known payoff. However, a greater actual payoff is not guaranteed at this node, which might depend on other players' choices.

What's next?

- Studying specific games in their entirety, with epistemic conditions.
- Incorporating probabilities into the *KBR*-model.
- Justification-tracking in game-theoretical reasoning, introducing tools to control logical omniscience hidden in the *KBR*-approach.
- Capturing the process of acquiring knowledge during games.
- Incorporating other epistemic notions into the model.
- Developing a theory of decision-making based on beliefs.
- Studying logical properties and new principles of Rationality.