

City University of New York (CUNY)

CUNY Academic Works

Computer Science Technical Reports

CUNY Academic Works

2009

TR-2009011: Knowledge-Based Rational Decisions

Sergei Artemov

[How does access to this work benefit you? Let us know!](#)

More information about this work at: https://academicworks.cuny.edu/gc_cs_tr/332

Discover additional works at: <https://academicworks.cuny.edu>

This work is made publicly available by the City University of New York (CUNY).
Contact: AcademicWorks@cuny.edu

Knowledge-Based Rational Decisions

Sergei Artemov

The CUNY Graduate Center
365 Fifth Avenue, 4319
New York City, NY 10016, USA
sartemov@gc.cuny.edu

September 4, 2009

Abstract

We outline a mathematical model of rational decision-making based on standard game-theoretical assumptions:

- 1) *rationality yields a payoff maximization given the player's knowledge;*
- 2) *the standard logic of knowledge for Game Theory is the modal logic S5.*

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

We demonstrate that uncertainty in games of perfect information results exclusively from players' different perceptions of the game. 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.

1 Introduction

In this paper, 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 the modal logic S5.*

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

Rendering epistemic conditions explicit is a necessary element of game analysis and recommendations. Without such disclosure, solutions offered by Game Theory would be incomplete or even misleading. Game theorists have long been aware of this and have

studied epistemic conditions under which traditional game-theoretical solutions, e.g., Nash equilibria, backward induction solutions, etc., hold (cf. [5, 6, 7, 8, 9, 11, 13, 16, 17] and many others).

The very notion of rationality carries a strong epistemic element, since a player’s rational choice depends on the player’s knowledge/beliefs. In his lecture [10], Adam Brandenburger says:

“What, then, is the implication of rationality in a game? This is a central question of the epistemic program in game theory.”

In this paper, we offer a mathematical model of rational decision-making based on aforementioned principles 1 and 2. In this model, epistemic states of players are essential elements of the game description.

Strictly speaking, this model could be formalized within a certain theory over the modal logic of knowledge, though we will try to keep the exposition informal for intuitive appeal and comprehension. We will be using knowledge operators as well as other logical connectives as part of the usual mathematical lingo, and will reason informally from the principles of the logic of knowledge **S5**.

The application of epistemic modal logic in Game Theory is an established tradition (cf., for example [8, 9, 12, 15, 20, 22, 23, 24, 25]). In this paper, however, we use the logic of knowledge to offer a new paradigm of rational decision-making, which we suggest calling

knowledge-based rationality,

or *KBR*, for short. The *KBR* paradigm is basically the classical notion of a player’s rationality as payoff maximization, given the player’s state of knowledge within the framework of the corresponding epistemic (modal) logic.

There are well-known models of decision-making under uncertainty which assume *a priori* knowledge/belief of the probability distribution of consequences of a player’s actions (von Neumann and Morgenstern [26] and Savage [21]). Knowledge-based rationality is an alternative mathematical model of decision-making under uncertainty which relies on the traditional understanding of rationality, utilizes a player’s knowledge at each node of the game, and does not require any probabilistic assumptions.

To this end, we use the *KBR*-model to analyze games of perfect information (PI games), though we show that it can be applied to other classes of games as well.

Technical report [3] contains a preliminary account of the *KBR*-approach.

2 Rationality: logical format

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

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

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

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

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

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

$$rP \rightarrow \mathbf{K}_P(rP). \tag{1}$$

3 Best Known Move

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 $1, 2, \dots, m$, and s_i denote

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

In particular, the following holds:

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

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,*

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

In other words, $HKP_A(j)$ is the largest payoff which A **knows** that he gets when playing j . If $p \leq HKP_A(j)$, then A knows that he gets a payoff of at least p when choosing j . If $p > HKP_A(j)$, then A considers it possible that he is not getting payoff p or higher when choosing j .

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

$$\mathbf{K}_A('A \text{ gets at least } p \text{ when choosing } j')$$

and

$$\bigwedge_{q \in G(p)} \neg \mathbf{K}_A(\text{'A gets at least } q \text{ when choosing } j\text{'}) .$$

The following is an easy though fundamental observation.

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 $HKP_A(j)$.*

Proof. Indeed, assuming A knows the game, A knows that his payoff will be one of a finite set of possible outcomes between the worst-case and best-case payoffs. The set of p 's such that P knows that his payoff given j will be at least p is finite, hence it has a maximum. \square

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 (knows) that 1 is no longer possible, then $HKP_A(j) = 20$. If instead A learns that 3 is no longer possible, then $HKP_A(j)$ remains equal to 10.

Note that if we base our analysis on knowledge (e.g., system S5) rather than on belief (for which factivity is not assumed), then A cannot **know** that 2 is no longer possible because this is just not true! Therefore, 2 will always be possible for A , hence $HKP_A(j)$ at all epistemic states associated with this node is less than or equal to the actual payoff.

Definition 2 Best Known Move *for player A at a given node of the game is a move j from $1, 2, \dots, m$ 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 $1, 2, \dots, m$*

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

By

$$kbest_A(j)$$

we denote the proposition

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

¹If, for simplicity's sake, we assume that all payoffs are different and all $HKP_A(j)$'s are different as well, then there is one and only one best known move at a given node.

If the epistemic element of Definition 1, “**implied by A ’s knowledge,**” is ignored, then this definition reflects the usual maximin reasoning: the player chooses a move which maximizes this player’s guaranteed payoff. However, the epistemic component makes all the difference: in *KBR*-reasoning, the player maximizes his guaranteed **known** payoff.

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

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

Let us consider the extensive-form game tree in Figure 1. As usual, we assume that all three players A , B , and C are rational and the game tree is commonly known. Player A moves at node u , players B and C at nodes v and w , respectively. Each player has the option of moving left or right, with indicated payoffs

$$l, m, n$$

where l , m , and n are payoffs for A , B , and C , respectively. The game starts at node u .

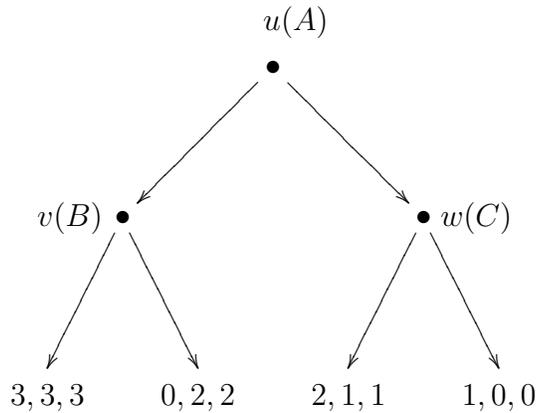


Figure 1: Game Tree 1

Being rational, B and C choose ‘left’ at v and w , but this can be unknown to A . Actually, there are several **different games** behind the game tree on Figure 1 which differ based on A ’s epistemic states, e.g.,

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.

In **Game I**, the highest known payoffs for A when choosing between v and w are

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

therefore the best known move for A at u is w :

$$kbest_A(w) .$$

A 's actual payoff at u is 2.

In **Game II**, the highest known payoffs for A when choosing v or w are

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

therefore, the best known move for A is w :

$$kbest_A(w) .$$

A 's actual payoff at u is 2.

In **Game III**, the highest known payoffs for A when choosing v or w are

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

therefore the best known move for A is v :

$$kbest_A(v) .$$

A 's actual payoff at u is 3.

The following theorem states that the best known move at each node of the game always exists, is unique (given different payoffs), and is always known to the player who is making a decision at this node²: for each possible move, the player knows whether or not it is the best known move.

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 $\mathbf{K}_A[\neg kbest_A(j)]$.*

Proof. We first establish a technical lemma.

Lemma 1 *For each j ,*

- 1) *If $HKP_A(j) = p$, then $\mathbf{K}_A[HKP_A(j) = p]$.*
- 2) *If $HKP_A(j) \neq p$, then $\mathbf{K}_A[HKP_A(j) \neq p]$.*

²Though ‘the best known move is known’ sounds like a tautology, it needs to be stated and proved, since within an epistemic environment, many different shades of knowledge and truth are possible. It is not true in general that F yields $\mathbf{K}_A(F)$. In some epistemic settings, it is possible for some class of propositions F to have F yields $\mathbf{K}_A(F)$, but not $\neg F$ yields $\mathbf{K}_A(\neg F)$, etc.

1) Suppose $HKP_A(j) = p$. Then, from the definition of highest known payoff,

$$\mathbf{K}_A(\text{'A gets at least } p \text{ when choosing } j')$$

and

$$\bigwedge_{q \in G(p)} \neg \mathbf{K}_A(\text{'A gets at least } q \text{ when choosing } j')$$

where $G(p)$ is the finite set of all possible payoffs for A which are greater than p .

In the logic of knowledge **S5**, both positive introspection and negative introspection hold, hence

$$\mathbf{K}_A \mathbf{K}_A(\text{'A gets at least } p \text{ when choosing } j')$$

and

$$\bigwedge_{q \in G(p)} \mathbf{K}_A[\neg \mathbf{K}_A(\text{'A gets at least } q \text{ when choosing } j')].$$

Since the knowledge/belief modality \mathbf{K}_A commutes with the conjunction,

$$\mathbf{K}_A[\bigwedge_{q \in G(p)} \neg \mathbf{K}_A(\text{'A gets at least } q \text{ when choosing } j')].$$

This yields that A knows $HKP_A(j) = p$, i.e., $\mathbf{K}_A[HKP_A(j) = p]$.

2) Suppose $HKP_A(j) = t$ and $t \neq p$. Then there are two possibilities: $t < p$ or $p < t$.
If $t < p$, then

$$\neg \mathbf{K}_A(\text{'A gets at least } p \text{ when choosing } j').$$

By negative introspection,

$$\mathbf{K}_A[\neg \mathbf{K}_A(\text{'A gets at least } p \text{ when choosing } j')],$$

hence A knows that $HKP_A(j) \neq p$, i.e., $\mathbf{K}_A[HKP_A(j) \neq p]$.

If $p < t$, then, since

$$\mathbf{K}_A(\text{'A gets at least } t \text{ when choosing } j'),$$

A knows that p is not the highest known payoff³, i.e., $\mathbf{K}_A[HKP_A(j) \neq p]$.

Corollary 1 For each i, j ,

- 1) If $HKP_A(i) \leq HKP_A(j)$, then $\mathbf{K}_A[HKP_A(i) \leq HKP_A(j)]$.
- 2) If $HKP_A(i) < HKP_A(j)$, then $\mathbf{K}_A[HKP_A(i) < HKP_A(j)]$.

³We naturally assume a certain level of intelligence from A , e.g., A should be able to compare numbers p and t and conclude that p is less than t , etc., in an epistemic **S5**-style environment.

We now proceed to prove Theorem 1.

1) Suppose $kbest_A(j)$. According to (4), $kbest_A(j)$ is the conjunction

$$\bigwedge_i [HKP_A(j) \geq HKP_A(i)] .$$

For each of the conjuncts,

$$[HKP_A(j) \geq HKP_A(i)] \rightarrow \mathbf{K}_A[HKP_A(j) \geq HKP_A(i)] ,$$

hence

$$\bigwedge_i [HKP_A(j) \geq HKP_A(i)] \rightarrow \bigwedge_i [\mathbf{K}_A[HKP_A(j) \geq HKP_A(i)] .$$

Since modality \mathbf{K}_A commutes with conjunctions,

$$\bigwedge_i [HKP_A(j) \geq HKP_A(i)] \rightarrow \mathbf{K}_A \bigwedge_i [HKP_A(j) \geq HKP_A(i)] .$$

Therefore,

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

2) Suppose $\neg kbest_A(j)$, which, by Boolean logic and elementary properties of inequalities, is equivalent to

$$\begin{aligned} & \neg \bigwedge_i [HKP_A(j) \geq HKP_A(i)] , \\ & \bigvee_i \neg [HKP_A(j) \geq HKP_A(i)] , \end{aligned}$$

or

$$\bigvee_i [HKP_A(i) > HKP_A(j)] .$$

By Corollary 1,

$$[HKP_A(i) > HKP_A(j)] \rightarrow \mathbf{K}_A[HKP_A(i) > HKP_A(j)] ,$$

hence

$$\bigvee_i [HKP_A(i) > HKP_A(j)] \rightarrow \bigvee_i \mathbf{K}_A[HKP_A(i) > HKP_A(j)] .$$

In modal logic **S5**, for any set of formulas Γ ,

$$\bigvee \mathbf{K}_A \Gamma \rightarrow \mathbf{K}_A \bigvee \Gamma ,$$

hence

$$\bigvee_i \mathbf{K}_A[HKP_A(i) > HKP_A(j)] \rightarrow \mathbf{K}_A \bigvee_i [HKP_A(i) > HKP_A(j)]$$

and

$$\bigvee_i [HKP_A(i) > HKP_A(j)] \rightarrow \mathbf{K}_A \bigvee_i [HKP_A(i) > HKP_A(j)] .$$

This concludes the proof of 2). □

Corollary 2 *At each node, there is always at least one best known move*

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

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

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

Let us extend Definition 1 by defining the **Highest Known Payoff for A at v**, $HKP_A(v)$, to be the highest A's payoff at v which is implied by A's knowledge at v . It is easy to see that if j is the best known move for A at node v , $kbest_A(j)$, then

$$HKP_A(v) = HKP_A(j) .$$

4 Rationality based on knowledge

We consider several verbal accounts of rationality and show that they lead to the same formal model⁴.

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

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

⁴For simplicity's sake, we assume here that all payoffs are different and we work under the assumptions of Corollary 2.

The natural formalization of 3 is the principle

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

The natural formalization of 4 is the principle

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

The natural formalization of 5 is the principle

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

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

Proof. From the rules of logic, (9) implies (5). Furthermore, (6–8) are equivalent in propositional logic.

We now prove that (5) and (6) are equivalent. Indeed, assume (5) and suppose rA and $kbest_A(j)$ hold. Then, by (5), s_j holds as well. However, since the player picks only one move (by (3)), s_i does not hold for any $i \neq j$, hence (6). Assume (6) and let rA and $kbest_A(j)$ both hold. Then, by (6), $\neg s_i$ occurs for all $i \neq j$. Since the player must choose (by (3)), he chooses the only remaining option, namely, j . Hence s_j , and thus (5).

It now suffices to establish that (8) implies (9). Assume (8) with i and j swapped:

$$rA \rightarrow [s_j \rightarrow \neg kbest_A(i)], \text{ when } i \neq j .$$

Given rA and s_j , we now have

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

which, together with Corollary 2, yields

$$kbest_A(j) .$$

Therefore

$$rA \rightarrow [s_j \rightarrow kbest_A(j)]$$

which, together with (8), yields (9). □

Theorem 2 shows that each of (5–9) captures the same robust principle of rational decision-making⁵. As a fundamental principle of rationality, it can be assumed as known by any intelligent agent, in particular, by any player.

⁵Note that Theorem 2 can be established within the basic modal logic of beliefs \mathbf{K} and requires neither factivity nor positive/negative introspection. Therefore, the equivalence of (5–9) can be extended to a variety of logics of belief as well.

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).*

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. The best known move is well-defined at each node, hence the existence of a *KBR*-solution for each well-defined game. The uniqueness is obvious once players have only one best known move at each node. To show that players play according to the *KBR*-solution, it suffices to demonstrate that at each node v , the corresponding player A **knows** proposition s_j (cf. (2)), which describes A 's best move at v .

By the rationality principle (5),

$$rA \rightarrow [kbest_A(j) \rightarrow s_j] .$$

This principle is commonly known, in particular, it is known to A :

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

Distributing the knowledge operator, by the logic of knowledge, we conclude

$$\mathbf{K}_A[rA] \rightarrow \{\mathbf{K}_A[kbest_A(j)] \rightarrow \mathbf{K}_A[s_j]\} . \quad (10)$$

Since players are rational, rA holds. By self-knowledge of rationality (1), $rA \rightarrow \mathbf{K}_A(rA)$, hence

$$\mathbf{K}_A[rA] . \quad (11)$$

Let j be the *KBR*-move by A at a given node. Then,

$$kbest_A(j) .$$

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

$$\mathbf{K}_A[kbest_A(j)] . \quad (12)$$

From (10), (11), and (12), by logical reasoning, we derive

$$\mathbf{K}_A[s_j] .$$

□

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

$$AP_Q(v),$$

is the payoff which Q gets if the game is played from v according to the *KBR-solution* of the game.

Note that according to the traditional game-theoretical approach (cf. [5]), we consider payoffs at all the nodes of the game, including those which will never be reached when the game is played.

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.’

Consider, for example, **Game I** in Figure 1. It has the following highest known payoffs for A :

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

the *KBR-solution*:

A plays ‘right’ at u , B plays ‘left’ at v , and C plays ‘left’ at w ;

and actual payoffs for A , B , and C (denoted $AP_{A,B,C}$):

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

Game II in Figure 1 has the highest known payoffs for A :

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

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

Finally, **Game III** in Figure 1 has the highest known payoffs for A :

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

the *KBR-solution*:

A plays ‘left’ at u , B plays ‘left’ at v , and C plays ‘left’ at w ;

and actual payoffs

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

5 Game Awareness

We will focus on *knowledge of the game*, which includes knowledge of the game tree, e.g., possible moves, payoffs, etc. Common knowledge of the game tree is a reasonable assumption here and looks attainable by some sort of public information, communication about the rules of the game, etc. However, the game is not defined unless epistemic states of players are specified as well (cf. example of three different games on the same game tree in Figure 1). It does not make sense to speak of a solution to the game when the principal ingredients of the game's definition, the epistemic states of players, are not specified properly. There are some traditional defaults, however, such as common knowledge of rationality, which usually make the game well-defined, but only in a specific, usually extreme sense. A good example is given by Aumann's Theorem on Rationality [5] which states that in perfect information games, common knowledge of rationality implies backward induction⁶.

We believe, however, that a serious approach to Game Theory would be to adopt a standard of game specification which, in addition to a complete description of the game tree – moves, payoffs, etc., includes a sufficient specification of epistemic states of players at each node. The goal of Game Theory then is to find and analyze solutions of games depending on all (reasonable) epistemic assumptions.

In the following section, we will try to provide examples of such an approach.

Definition 6 *We distinguish the following notions. Knowledge of the game tree, which includes knowledge of the rules – possible moves, payoffs, etc., but does not necessarily include knowledge of epistemic states of players, which should be specified separately. Knowledge of the game, which is knowledge of the game tree and of epistemic states of all players prior to the game.*

Our default requirement for analyzing the game is rationality of players and common knowledge of the game tree, which does not exclude considering irrational players or players who are not completely aware of the game tree when needed.

6 Epistemic analysis of the Centipede Game

Figure 2 illustrates the centipede game suggested by Rosenthal, 1982, [19] and studied in an epistemic context by Aumann, 1995 [5]. Player *A* makes moves at nodes 1, 3, and 5, player *B* at nodes 2 and 4. Each player has the option of moving across or down, with indicated payoffs m, n where m is *A*'s payoff, and n is the payoff for *B*. The game starts at node 1.

⁶Common knowledge of rationality (or its finite-nesting versions) has been widely adopted as an epistemic condition for backward induction in perfect information games ([5, 7, 24]). In the same paper [5], Aumann states that common knowledge of rationality is an idealized condition that is rarely met in practice.

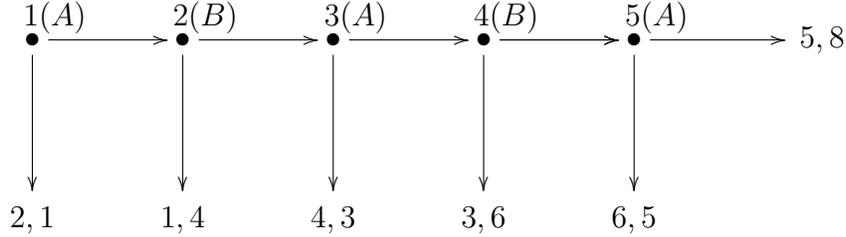


Figure 2: Centipede game of length 5

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.

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. This is a well-known issue (cf. [5, 7, 8, 24]) and 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. In this section, we will try to revise the perception that stockpiling of mutual knowledge assumptions are needed for solving the centipede game.

According to Sections 3 and 4, 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.

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.*

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,

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

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

Node 4, player B . Obviously $HKP('down') = 6$. On the other hand, $HKP('across') = 5$, 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

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

Since B is rational, d_4 .

Node 3, player A . $HKP('down') = 4$, whereas $HKP('across') = 3$, since A considers d_4 possible. Hence

$$kbest_A('down') \text{ holds at node 3.}$$

Since A is rational, d_3 .

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

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

Since B is rational, d_2 .

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

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

Since A is rational, d_1 . □

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.'

7 Strictly competitive games of perfect information

The proof of the main result of Section 6 that under any epistemic conditions, a *KBR*-solution coincides with the *BI*-solution, hence rational players play the backward induction strategy, can be extended to strictly competitive two-person games⁷ of perfect information.

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

$$X \succsim_A Y \Leftrightarrow Y \succsim_B X . \quad (13)$$

Since possible outcomes in extensive-form games are normally associated with payoffs at terminal nodes, we can reformulate (13): for each possible outcomes m_1, n_1 and m_2, n_2 ,

$$m_1 \leq m_2 \Leftrightarrow n_2 \leq n_1 . \quad (14)$$

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.*

Proof. The idea of the proof is similar to that of Theorem 4. Again, for simplicity's sake, we assume that for each player, his payoffs at terminal nodes are different.

Lemma 2 *The KBR-solution coincides with the maximin solution.*

Proof. For a player P , let

$$\text{maximin}_P(v)$$

be P 's maximin payoff at node v .

We show, by backward induction, that the highest known payoff at each node is equal to the player's maximin payoff.

Induction Base: pre-terminal nodes⁸. At such a node t , the active player, A , being rational, picks the move with the highest payoff, hence

$$\text{HKP}_A(t) = \text{maximin}_A(t) .$$

The other player, B , knows the game tree and knows that at least the worst-case payoff at t is guaranteed. On the other hand, he cannot **know** that he gets any higher payoff, since the actual choice of A is possible for B and brings B the minimal possible payoff at t :

$$\text{HKP}_B(t) = \text{maximin}_B(t) .$$

⁷In particular, zero-sum games.

⁸A node is pre-terminal if all of its successors in the game tree are terminal nodes.

Induction Step. Let v be a non-pre-terminal node, and let A be the player who has to choose one of the moves $1, 2, \dots, m$ at v . By the Induction Hypothesis, for all such j 's,

$$HKP_A(j) = \text{maximin}_A(j) ,$$

hence

$$\max\{HKP_A(j) \mid j = 1, 2, \dots, m\} = \max\{\text{maximin}_A(j) \mid j = 1, 2, \dots, m\} = \text{maximin}_A(v) .$$

So, since A chooses rationally,

$$HKP_A(v) = \text{maximin}_A(v) .$$

For player, B , by the Induction Hypothesis, for all $j = 1, 2, \dots, m$,

$$HKP_B(j) = \text{maximin}_B(j) .$$

Therefore,

$$\min\{HKP_B(j) \mid j = 1, 2, \dots, m\} = \min\{\text{maximin}_B(j) \mid j = 1, 2, \dots, m\} = \text{maximin}_B(v) .$$

Since B knows the game tree, B knows that A has to choose one of $1, 2, \dots, m$, hence B 's payoff at v cannot be less than the minimal of payoffs at $1, 2, \dots, m$:

$$HKP_B(v) \geq \min\{HKP_B(j) \mid j = 1, 2, \dots, m\} .$$

On the other hand, B considers it possible that A makes the best move for A (the actual choice of A) which is the worst move for B . Therefore,

$$HKP_B(v) = \min\{HKP_B(j) \mid j = 1, 2, \dots, m\}$$

and

$$HKP_B(v) = \text{maximin}_B(v) .$$

□

It now remains to check that the BI -solution actually coincides with the maximin solution. Indeed, both solutions depend on the game tree and do not depend on epistemic states of players. In the BI -solution, A chooses the highest payoff given B will choose his highest payoffs, etc. In the maximin solution, A chooses the highest payoff given B will choose A 's minimal payoff. Since the game is strictly competitive, the minimal payoff for A occurs together with B 's maximal payoff. Therefore, the algorithm for calculating the BI -solution coincides with the algorithm of calculating the maximin solution. □

8 When nested knowledge matters

In this section, we offer an alternative to the centipede game as an illustration of Aumann's Theorem on Rationality. In the new game, to justify the backward induction solution, one really needs mutual knowledge of rationality, in which nested depth is the length of the game.

Figure 3 shows a game which we provisionally call *anti-centipede game of length 3*.

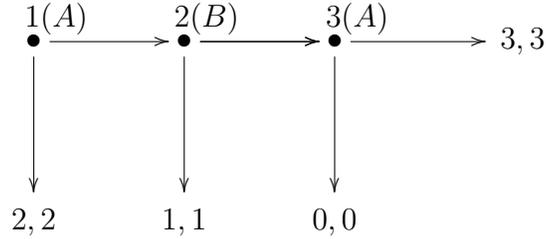


Figure 3: Anti-centipede game of length 3

Of course, the tree in Figure 3 does not define the game completely: the epistemic states of players remain to be specified. As usual, we assume rationality of players and sufficient knowledge of the game tree of both players.

Game I: both players are rational

$$rA \quad \text{and} \quad rB \tag{15}$$

but ignorant of each other's rationality: neither $\mathbf{K}_A(rB)$ nor $\mathbf{K}_B(rA)$ hold. As a result, both A and B consider possible any move by their opponent at any node. Let $\diamond_P(F)$ stand for $\neg\mathbf{K}_P(\neg F)$, hence \diamond_P is a *possibility operator* associated with the knowledge operator \mathbf{K}_P . Then

$$\diamond_B(a_3) \wedge \diamond_B(d_3) \quad \text{and} \quad \diamond_A(a_2) \wedge \diamond_A(d_2) . \tag{16}$$

Game I is defined by the game tree in Figure 1, rationality of players (15), and epistemic conditions (16).

Let us solve Game I. 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,

$$HKP_B(\text{'across'}) = 0 ,$$

whereas

$$HKP_A(\text{'down'}) = 1 ,$$

and B chooses 'down' at 2. Likewise, by (16), A considers either of a_2 and d_2 possible. Therefore, at root node 1, A chooses 'down.' The solution of the game is

$$d_1, d_2, a_3 ,$$

and both players get payoff 2.

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

$$\mathbf{K}_A(rB) \quad \text{and} \quad \mathbf{K}_B(rA) , \quad (17)$$

but not level 2 mutual knowledge of rationality: neither $\mathbf{K}_B\mathbf{K}_A(rB)$ nor $\mathbf{K}_A\mathbf{K}_B(rA)$ hold. In particular, it would follow from $\mathbf{K}_A\mathbf{K}_B(rA)$ that A knows that B knows that A plays ‘across’ at 3. But A in Game II is not so well-informed and does not know that B knows that A plays ‘across’ at 3, or, symbolically,

$$\neg\mathbf{K}_A\mathbf{K}_B(a_3) . \quad (18)$$

Taking into account that it is common knowledge that a_3 is logically equivalent to $\neg d_3$, condition (18) can be equivalently presented as

$$\neg\mathbf{K}_A\mathbf{K}_B(\neg d_3) ,$$

or,

$$\diamond_A\diamond_B(d_3) . \quad (19)$$

Game II is defined by epistemic conditions (17) and (19). Let us find its solution.

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 . Therefore, B ’s best known move at 2 is ‘across,’ since it yields payoff 3 *vs.* payoff 1 when playing ‘down.’ As a rational player, B chooses ‘across’ at 2, hence a_2 .

At node 1, by (19), A considers it possible that B considers d_3 possible. As a rational player who considers d_3 possible, B must choose ‘down’ at 2, hence d_2 . A knows that B is rational, hence A considers it possible that B plays ‘down’ at 2 and delivers payoff 1 for A . Therefore, the highest known payoff for A when playing ‘across’ is 1:

$$HKP_A(\text{‘across’}) = 1 ,$$

whereas, by the game tree,

$$HKP_A(\text{‘down’}) = 2 .$$

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

$$d_1, a_2, a_3 ,$$

and both players get payoff 2.

Game III: Common knowledge of rationality is assumed⁹. 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

⁹Actually, it suffices to assume $\mathbf{K}_A(rB)$ and $\mathbf{K}_A\mathbf{K}_B(rA)$.

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

$$a_1, a_2, a_3 ,$$

and both players get payoff 3.

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 minus one. Figure 4 shows a game tree for the anti-centipede of length 5.

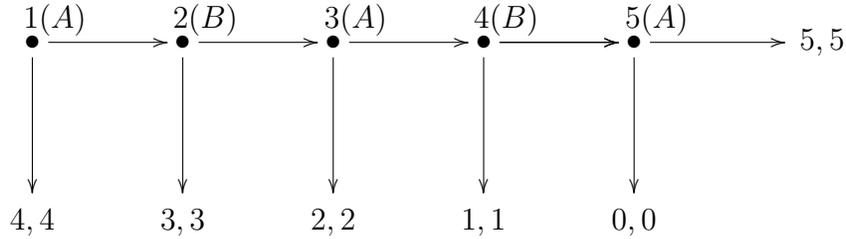


Figure 4: Anti-centipede game of length 5

9 Knowledge of the game

In this section, we will revisit the notion of the knowledge of the game and show that in its entirety, including all necessary epistemic conditions, it is only possible in somewhat special cases. In all other games, players have necessarily **different and incomplete knowledge** of the game which they are playing.

We start with examples. Solving Games I, II, and III from Section 8, we deliberately did not concentrate on knowledge of the game; it is time to analyze it in more detail.

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. In particular, each player knows that he is rational: $\mathbf{K}_P(rP)$, $\mathbf{K}_P\mathbf{K}_P(rP)$, etc.

In Game I, Section 8, \mathcal{E} consists of conditions (16). From the game description, we can logically derive $\neg\mathbf{K}_B(rA)$. Indeed, it is common knowledge (from the Game Tree and general understanding of rationality) that if A is rational, then A plays ‘across’ at node 3, hence

$$rA \rightarrow a_3 . \quad (20)$$

In particular, B knows (20):

$$\mathbf{K}_B(rA \rightarrow a_3) ,$$

hence

$$\mathbf{K}_B(rA) \rightarrow \mathbf{K}_B(a_3) ,$$

and

$$\neg\mathbf{K}_B(a_3) \rightarrow \neg\mathbf{K}_B(rA) .$$

Since, by (16), $\neg\mathbf{K}_B(a_3)$ holds,

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

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. It is easy to see that A ’s knowledge of the game is not complete either, otherwise A would be able to calculate the best move for B at 2 and predict¹⁰ either a_2 or d_2 , which, by (16), is not the case.

In Game II, Section 8, \mathcal{E} consists of conditions (17) and (19). Proposition $\mathbf{K}_A\mathbf{K}_B(rA)$ does not hold. Indeed, it follows from the Game Tree and the Rationality Thesis (Definition 3), that (20) is commonly known. In particular,

$$\mathbf{K}_A\mathbf{K}_B(rA \rightarrow a_3) .$$

From this, by S5-reasoning, we conclude

$$\mathbf{K}_A\mathbf{K}_B(rA) \rightarrow \mathbf{K}_A\mathbf{K}_B(a_3) ,$$

hence

$$\neg\mathbf{K}_A\mathbf{K}_B(a_3) \rightarrow \neg\mathbf{K}_A\mathbf{K}_B(rA) .$$

By (19), $\neg\mathbf{K}_A\mathbf{K}_B(a_3)$, hence

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

On the other hand, from (17), we conclude

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

¹⁰Cf. Proposition 2.

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.

Game III, Section 8 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*: $\mathcal{E} = \text{Common Knowledge of Rationality}$.

Since, for each player P ,

$$\text{Common Knowledge that } F \rightarrow \mathbf{K}_P(\text{Common Knowledge that } F) , \quad (21)$$

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 (21).

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

Proof. From the definitions. Since an agent A knows the game, including epistemic states of all players, A can calculate the highest known payoff and the best known move for each player, say B , at each given node. Indeed, suppose B 's best known move is j . This means that B logically concludes from his epistemic state at a given node that $kbest_B(j)$:

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

The laws of logic are known to each intelligent agent, hence A can reproduce B 's reasoning from the same set of assumptions:

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

By logic of knowledge,

$$\mathbf{K}_A[\text{'Epistemic State of } B'] \rightarrow \mathbf{K}_A[kbest_B(j)] . \quad (22)$$

In addition, A knows that B is rational, since A knows *Rationality*:

$$\mathbf{K}_A[rB] . \quad (23)$$

Furthermore, A , of course, knows how *KBR* works, in particular,

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

from which, by the laws of logic, it follows that

$$\mathbf{K}_A[rB] \rightarrow (\mathbf{K}_A[kbest_B(j)] \rightarrow \mathbf{K}_A[s_j]) . \quad (24)$$

Taking into account that A knows the epistemic state of B at this node, from (22), (23), and (24), we conclude

$$\mathbf{K}_A[s_j] ,$$

meaning A knows B 's move at this node.

Therefore, A knows the moves of all players, i.e., A knows the KBR -solution of the game and can calculate actual payoffs. \square

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

Naturally, if A is certain at v , then A knows all actual payoffs at nodes that immediately follow v and hence can calculate the actual payoffs at node v .

Corollary 3 *Any player who knows the game in full is certain at each node of the game.*

Proof. This follows easily from Proposition 2. The only new feature mentioned here is certainty, which within the current context is a tautology: if a player **knows** the moves of all other players, he is certain at each node. \square

Proposition 3 *In a PI game, rationality of players and 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. \square

Uncertainty in PI games occurs only because players do not know the game in full. Moreover, if uncertainty occurs in a PI game, then some players necessarily have **different understanding of the game** as well.

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’:*

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

This definition naturally extends to more than two players.

Assertions from \mathcal{IR} are important epistemic conditions of a game. For example, our analysis of Games I, II, and III earlier in this section uses \mathcal{IR} specification in an essential way.

\mathcal{IR} assertions play a special role in the theory of perfect information games. In particular, Aumann's Theorem on Rationality refers to PI games with common knowledge of rationality, which in our context yields that \mathcal{IR} holds in its entirety for such games.

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.*

Proof. Will be confined to the case of two players, without loss of generality.

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.*

Proof. Indeed, since A is rational, the rationality of A assertion holds:

$$rA .$$

Since the rationality of A is self-known, $rA \rightarrow \mathbf{K}_A(rA)$,

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

holds, i.e., A knows rA . By assumptions, B knows the same \mathcal{IR} assertions as A , in particular, B knows rA :

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

By positive introspection of B 's knowledge, we conclude

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

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

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

By positive introspection of A 's knowledge,

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

meaning that A knows $\mathbf{K}_A\mathbf{K}_B(rA)$. By assumptions, B knows that as well:

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

etc.

By similar reasoning, we can show that iterated knowledge assertions of rB are also all known. \square

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 $n - 1$ mutual knowledge of rationality, which is guaranteed by Lemma 3. \square

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 less than n . Such an agreement is only possible when all iterated rationality assertions of nested knowledge depth less than 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 less than 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 Proposition 3 and Theorem 6. Indeed, common knowledge of rationality immediately yields that each player knows the whole set of iterated rationality assertions \mathcal{IR} .

Altogether, Corollary 3 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.

10 Strategic games and knowledge-based rationality

In this section, we give an example of KBR -reasoning in strategic-form games.

Imagine two neighboring countries: a big, powerful B , and a small S . Each player has the choice to wage war or keep the peace. The best outcome for both countries is peace. However, if both countries wage war, B wins easily and S loses everything, which is the second-best outcome for B and the worst for S . In situation $(war_B, peace_S)$, B loses internationally, which is the second-best outcome for S . In $(peace_B, war_S)$, B 's government loses national support, which is the worst outcome for B and the second-worst for S .

		war_S		$peace_S$
	war_B	2,0	\Rightarrow	1,2
		\uparrow		\downarrow
	$peace_B$	0,1	\Rightarrow	3,3

There is one Nash equilibrium, $(peace_B, peace_S)$, consisting of the best outcomes for both players. It might look as though they should both play accordingly. However, such a prediction is not well-founded unless certain epistemic conditions are met.

Theorem 7 *In the War and Peace Dilemma, suppose the game matrix is mutually known and the players are rational. Then,*

i) If B is not aware of the other player's rationality, S chooses 'peace' and B chooses 'war.'

ii) If S 's rationality is known to B , both players choose 'peace.'

Proof. Let us analyze this game in epistemic logic with two modalities, \mathbf{K}_B and \mathbf{K}_S , and propositions rB and rS for rationality assertions. We define propositions:

- w_B - B chooses to wage war,
- p_B - B chooses to keep peace,
- w_S - S chooses to wage war,
- p_S - S chooses to keep peace.

For both (i) and (ii), we consider the Rationality assumption

$$rB, rS. \tag{25}$$

(i) **Game I** is defined by an epistemic condition stating that B considers all moves of S epistemically possible:

$$\mathcal{E} = \{\diamond_B(w_S), \diamond_B(p_S)\}. \tag{26}$$

This is a maximin case, and the corresponding best known moves can be directly calculated, or derived, from epistemic conditions and the game description: w_B and p_S .

Informally, S has the dominant strategy, $peace_S$, whereas B lacks one, hence B 's choice actually depends on his expectations of S 's move. Since B considers both moves by S possible, B counts on the worst cases and hence picks war_B .

(ii) **Game II** is defined by adopting the condition that B knows that S is rational:

$$\mathcal{E} = \{\mathbf{K}_B(rS)\}.$$

S 's *KBR*-reasoning does not change, and S 's *KBR*-strategy is 'peace':

$$p_S .$$

From mutual knowledge of the game matrix, B easily concludes that S knows that 'peace' is the dominant strategy for S . Since, in addition, B knows that S is rational (condition \mathcal{E} of Game II), from the Rationality Thesis (Definition 3), B concludes that S will play 'peace,' i.e., that p_S holds. Since S plays 'peace,' B 's moves yield the following payoffs:

$$HKP_B(\text{war}_B) = 1 , \quad HKP_B(\text{peace}_B) = 3 .$$

So the best known move for B is 'peace,' and, since B is rational, he chooses 'peace,' hence p_B . □

In the War and Peace Dilemma, our logical analysis shows that despite that

- a) for both countries, the best choice is 'peace';
- b) it is the only Nash equilibrium in the game;
- c) both countries behave rationally;

to secure the Nash equilibrium outcome, an additional epistemic condition should be met, e.g., the big country should know that its small neighbour will behave rationally.

11 Discussion

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.

11.1 Other models of epistemic rationality.

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.

11.2 Rationality Theorem for PI Games

Corollary 3 and Theorem 6 do not really depend on our theory of knowledge-based decisions under uncertainty developed in Sections 1–4, since they provide sufficient conditions under which uncertainty in the game can be eliminated completely.

11.3 Relations to Aumann’s Rationality Theorem

What do Corollary 3 and Theorem 6 add to Aumann’s Rationality Theorem? Theorem 6 states that uncertainty can be eliminated (the claim of Aumann’s Theorem) under somewhat more general assumptions: instead of Aumann’s common knowledge of rationality requirement, Theorem 6 only requires that players agree, one way or another, on a certain (finite) set of iterated rationality assertions. However, the significance of Theorem 6 and Corollary 3 is more conceptual. They demonstrate the real power of knowledge of the game. Theorem 6 and Corollary 3 show that uncertainty in perfect information games appears only as the result of different epistemic perceptions by players. This makes the case that epistemic conditions of players should be objects of game-theoretical studies.

11.4 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.

11.5 Do we need the full power of 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 more modest epistemic means. For example, in Game I of Section 8, it suffices to apply negative introspection to epistemic conditions (16) 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 (16) 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.

11.6 Learning

It is obvious that more player knowledge at a given node yields a greater highest known payoff. However, a greater actual payoff, which could depend on other players' choices, is not guaranteed at this node. Moreover, it is possible that a player A with superior knowledge of the game can manipulate B 's rational choices by leaking truthful, but incomplete, information, which changes B 's best known move to A 's advantage.

11.7 More rationality principles

In this paper, we do not provide a complete logical description of rationality but rather use some of its commonly accepted features, e.g., the Rationality Thesis (Definition 3), to find solutions for games. The question of what the logical properties of rationality are remains open.

Consider rationality principle (9). It is tempting to extend it to a complete-looking definition of rationality:

$$rA \leftrightarrow [kbest_A(j) \leftrightarrow s_j] \quad (27)$$

or something of this sort. We believe this is not a good idea. First, condition

$$[kbest_A(j) \leftrightarrow s_j]$$

describes rational behavior of the player at one particular node of the game, which itself is not sufficient for stating rationality in general. So, condition (27), as stated, is just wrong.

One can also argue that reducing rationality to the making of rational choices at all nodes of the game does not work either. If the player is not rational, it does not mean that he must make irrational choices at any of the nodes. He might as well choose the same moves as a rational player by some other means, e.g., by chance, etc.

However, irrationality seems to possess a definite epistemic meaning: if, by the rules of the game, B does not know that A is rational, B cannot apply the Rationality Principle (Definition 3) to predict A 's behavior. Under these assumptions, unless some specific information tells B otherwise, a risk-averse B has strong reason to assume, as part of the game description, that at a given node, **any legitimate move by A is possible**.

$$\bigwedge_{j \leq m} \diamond_B(s_j) . \quad (28)$$

According to (28), rationality is more than just a conjunction of rational choices in a given game. Rationality is, rather, a reputation: if A possesses it (in B 's eyes), then B counts on A 's rational behavior at any node. Otherwise, if B does not know that A is rational, B

just does not try to predict A 's behavior on grounds of rationality and considers any move by A possible.

Note that as stated, (28) is not a new logical rule, but rather a formalization methodology, which was used to describe games in Sections 3 and 8. Adopting (28) has no bearing on the main results of this paper, which describe how games are played **after** formalization.

11.8 What's next?

The knowledge-based theory of rationality decisions advocates studying games in their entirety, including comprehensive epistemic conditions. Here are some possible avenues of research for logic-based game-theoretical models.

- Studying specific games in their entirety, with various epistemic conditions.
- Incorporating a probabilistic approach and mixed strategies into the *KBR*-model.
- Developing a mechanism of 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 rather than knowledge.
- Studying logical properties and new principles of rationality.

12 Acknowledgments

The author is grateful to Adam Brandenburger, Mel Fitting, Saul Kripke, Vladimir Krupski, Elena Nogina, Olivier Roy, and Cagil Tasdemir for useful discussions. Many thanks to Karen Kletter for editing this paper.

References

- [1] S. Artemov. Explicit provability and constructive semantics. *Bulletin of Symbolic Logic*, 7(1):1–36, 2001.
- [2] S. Artemov. The logic of justification. *The Review of Symbolic Logic*, 1(4):477–513, 2008.
- [3] S. Artemov. *Intelligent Players*. Technical Report TR-2009006, CUNY Ph.D. Program in Computer Science, 2009.

- [4] R. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976.
- [5] R. Aumann. Backward Induction and Common Knowledge of Rationality. *Games and Economic Behavior*, 8:6–19, 1995.
- [6] R. Aumann and A. Brandenburger. Epistemic conditions for Nash equilibrium. *Econometrica: Journal of the Econometric Society*, 64(5):116–1180, 1995.
- [7] C. Bicchieri. *Rationality and Coordination*. Cambridge University Press, Cambridge, 1993.
- [8] C. Bicchieri. The epistemic foundations of Nash equilibrium. In Daniel Little, editor, *On the Reliability of Economic Models: Essays in the Philosophy of Economics*, Springer, pp. 91–137, 1995.
- [9] G. Bonanno. A syntactic approach to rationality in games with ordinal payoffs. In: G. Bonanno, W. van der Hoek and M. Wooldridge (eds.), *Logic and the Foundations of Game and Decision Theory, Texts in Logic and Games Series*, Amsterdam University Press, pp. 59–86, 2008.
- [10] A. Brandenburger. *Epistemic Game Theory*. Lecture of June 22, 2008, <http://pages.stern.nyu.edu/~abranden> .
- [11] A. Brandenburger. Origins of Epistemics. *Theoretical Aspects Of Rationality And Knowledge. Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge*, ACM, p.2, 2009. <http://pages.stern.nyu.edu/~abranden> .
- [12] B. de Bruin. Common Knowledge of Rationality in Extensive Games. *Notre Dame J. Formal Logic*, 49(3):261–280, 2008.
- [13] D. Ellsberg. Recurrent Objections to the Minimax Strategy: Rejoinder. *The Review of Economics and Statistics*, 41(1): 42-43, 1959.
- [14] R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.
- [15] J. Halpern. Substantive Rationality and Backward Induction. *Games and Economic Behavior*, 37:425–435, 2001.
- [16] J.C. Harsanyi. Games with Incomplete Information Played by “Bayesian” Players, I-III. *Management Science*, 14:159-182, 320-334, and 486-502, 1967/68.
- [17] J. Nash. *Non-Cooperative Games*. Ph.D. Thesis, Princeton University, Princeton, 1950.
- [18] M. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, 1994.

- [19] R. Rosenthal. Games of Perfect Information, Predatory Pricing, and the Chain Store Paradox. *Journal of Economic Theory*, 25: 92–100, 1982.
- [20] D. Samet. Hypothetical Knowledge and Games with Perfect Information. *Games and Economic Behavior*, 17: 230–251, 1996.
- [21] L.J. Savage. *The Foundations of Statistics* (Second Revised Edition). New York: Dover, 1972.
- [22] R.C. Stalnaker. Knowledge, Belief and Counterfactual Reasoning in Games. *Economics and Philosophy*, 12:133–163, 1996.
- [23] J. van Benthem. Rational Dynamics and Epistemic Logic in Games. *International Game Theory Review*, 9(1):13–45, 2007.
- [24] P. Vanderschraaf and G. Sillari. *Common Knowledge*. Stanford Encyclopedia of Philosophy, 2007.
- [25] R. Vestergaard, P. Lescanne, and H. Ono. *The Inductive and Modal Proof Theory of Aumann's Theorem on Rationality*. Technical report: JAIST/IS-RR-2006-009.
- [26] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.