

2009

TR-2009012: Rational Decisions in Non-Probabilistic Settings

Sergei Artemov

Follow this and additional works at: http://academicworks.cuny.edu/gc_cs_tr

 Part of the [Computer Sciences Commons](#)

Recommended Citation

Artemov, Sergei, "TR-2009012: Rational Decisions in Non-Probabilistic Settings" (2009). *CUNY Academic Works*.
http://academicworks.cuny.edu/gc_cs_tr/333

This Technical Report is brought to you by CUNY Academic Works. It has been accepted for inclusion in Computer Science Technical Reports by an authorized administrator of CUNY Academic Works. For more information, please contact AcademicWorks@gc.cuny.edu.

Rational decisions in non-probabilistic settings

Sergei Artemov

The CUNY Graduate Center
365 Fifth Avenue, rm. 4329
New York City, NY 10016, USA

October 27, 2009

Abstract

The knowledge-based rational decision model (*KBR*-model), developed in [1], offers an approach to rational decision making in a non-probabilistic setting, e.g., in perfect information games with deterministic payoffs. The *KBR*-model is an epistemically explicit form of standard game-theoretical assumptions, e.g., Harsanyi's Maximin Postulate. This model suggests following maximin strategy over all scenarios which the agent considers possible to the best of his knowledge.

In this paper, we compare *KBR* with other approaches and show that *KBR* is the only non-probabilistic decision-making method which is definitive, rational, and based exclusively on knowledge.

1 Introduction

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




Figure 1: Game One - Apollo 13 Dilemma

What is a rational way for A to resolve this issue? There are well-known models of decision making under uncertainty which assume a priori knowledge/belief of the probability distribution of the consequences of a player’s actions (von Neumann and Morgenstern [21], Savage [17]). These models seem not to apply to this situation for several reasons. First, the payoffs of 0, 1, and 2 are ordinal and do not reflect proportional values of these outcomes to A . What is the monetary equivalent of success in a space-race situation if the mission also has a strong political element? What if it is a high-profile rescue mission? How disastrous is failure? These are the sorts of questions which do not seem to have obvious answers within the given framework. The same holds for probabilities: nothing in the problem description indicates that such probabilities can be assigned responsibly.

With enough good will, Harsanyi’s Maximin Postulate¹ of rational behavior can be applied here:

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

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

1.1 General epistemic approach to games

Epistemic states of players are necessary elements of game analysis. Game theorists have long considered epistemic conditions under which traditional game-theoretical solutions hold (cf. [3, 4, 5, 6, 7, 9, 10, 12, 13, 14] and many others).

In accordance with recent trends in Game Theory, we consider games with arbitrary epistemic conditions of players. In particular, rationality of a player is not necessarily assumed to be known to other players. If it is known, then this should be made an explicit part of the game description.

Games here are presented in extensive finite tree-like form with deterministic payoffs at their terminal nodes (leaves). Game trees, that includes payoffs, are commonly known by the players. A game is called *generic* if there are no indistinguishable payoffs for each player.

¹‘Harsanyi’s postulate’ here means ‘postulate formulated and endorsed by Harsanyi,’ cf. [13] Sections 6.2 and 6.3, Postulate A1. The notion of maximin had appeared, of course, much earlier, e.g. in [20, 21].

²As occurred with *Apollo 13*.

1.2 How to play the Centipede game

The starting point of the knowledge-based rational decision model (*KBR*-model) is the paradigmatic example of Rosenthal's Centipede game ([16])³, Figure 2, where we offer an alternative solution which is much less restrictive and more general than the canonical backward induction solution (cf. [3, 19]).

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.




Figure 2: Centipede game of length 5

The canonical *backward induction* solution of Centipede is covered by Aumann's Theorem on Rationality ([3]) which assumes *common knowledge of rationality* and 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, etc. This analysis has been used in textbooks and expository articles as an illustration of Aumann's Theorem (cf. [19]).

The question is how to play the Centipede game in a general setting when common knowledge of rationality is not necessarily assumed. What if players are rational, but ignorant of each other's rationality? What if *A* knows that *B* is rational, but not vice versa? There are many meaningful versions of the Centipede game corresponding to different epistemic states of players and the backward induction solution addresses only one, rather extreme version of this game with the unrealistically strong assumption of common knowledge of rationality⁴. Theorem 1 answers this question in full.

For the Centipede game, Theorem 1 looks stronger than Aumann's Theorem since Theorem 1 does not require common knowledge of rationality. This may come as a surprise, but the standard solution *down* at each node of the Centipede game does not require any mutual knowledge of rationality at all.

³Rosenthal in [16] introduced Centipede games of arbitrary finite length; we took our formulation of Centipede from [15]. The Centipede game is defined by some inequalities of preferences and admits different payoff realizations, all of which reflect the same idea. The game in Figure 2 provides a representative example of a generic length-5 Centipede game, cf. also [3, 7, 19] for other examples.

⁴Common knowledge of rationality is rarely met in practice ([3]).

Theorem 1 [1] *In the Centipede game, under any states of players' knowledge, rational players play down at each node.*

Proof. Theorem holds for Centipede games of any finite length. We will give a proof for a specific but representative example in Figure 5. This solution can be extended to Centipede games of any finite length in a straightforward way.

At node 5, player *A* chooses *down*.

At node 4, player *B*'s maximin strategy is to play *down*. In addition, *B* **cannot know** that *A* would play *across* because *A* in fact plays *down*⁵. Therefore, *B* cannot rationally expect to get more than his maximin payoff of 4 at node 4 **regardless of his knowledge about *A***. By Harsanyi's Maximin Postulate, *B* chooses maximin strategy *down*.

At node 3, player *A* cannot know that *B* will play *across* at 4, since *B* in fact plays *down*. Therefore, *A* cannot rationally expect to exceed his maximin payoff of 3 and hence, by Harsanyi's Maximin Postulate, plays *down*.

Likewise, *B* plays *down* at 2 and *A* plays *down* at 1. □

This solution of the Centipede game illuminates a synergy between rationality and knowledge and hints at a comprehensive theory of knowledge-based rational decisions under uncertainty which does not rely on probabilistic assumptions.

The Centipede game shows that rationality at each node can lead to strategies which are bad for all players. Alternative (and better) ways to play the Centipede game are discussed in Sections 5.3 and 5.4.

2 Rationality

We abstain from giving a comprehensive definition of what is rational, but instead use natural properties of rationality which follow from common game-theoretical assumptions. These properties turn out to be sufficient for definitive knowledge-based rational decisions.

2.1 Decisions based on knowledge in non-probabilistic setting

How should Harsanyi's Maximin Postulate be read within the context of other principles of rationality, e.g., commonly accepted understanding that *rational decisions should be based on knowledge* rather than on luck, guesswork, sudden opponent cooperation or error, divine intervention, etc.⁶

If rational decisions should be based on players' knowledge, the informal notion

to rationally expect a given payoff

⁵This argument uses the universally accepted *factivity of knowledge* principle: only true propositions can be known.

⁶Unless specifically mentioned in the formulation of a game.

in Harsanyi's Maximin Postulate should be read as

to know that you will get a given payoff.

Indeed, if you know that you will get a certain payoff, you rationally expect it to happen.

Conversely, suppose you don't know that you will get at least this payoff. Then for each of your strategies, there are responses by your opponents which you consider possible to the best of your knowledge, which deny you a given payoff. It does not then seem rational for you to expect this payoff in a non-probabilistic setting.

This knowledge-based reading of Harsanyi's Maximin Postulate is appropriate for games in which probabilistic assumptions are not part of the picture, e.g., perfect information games with deterministic payoffs.

When building the theory of knowledge-based rational decisions, we should certainly take into account another of Harsanyi's postulates, formulated in [13], the **Mutually Expected Rationality Postulate**:

If you are a rational player, you must expect, and act on this expectation, that other rational players will likewise follow these rationality postulates [which include Harsanyi's Maximin Postulate].

Here again we suggest reading

must expect ... that other rational players will likewise follow ...

as

know ... that other rational players will likewise follow

Moreover, traditional game-theoretical reasoning actually uses an iterated form of the Mutually Expected Rationality Postulate which assumes **common knowledge** of the fact that rational players follow rationality postulates.

We do not argue that *KBR* is the only way to play games rationally: traditional game-theoretical approaches, including probabilistic methods, can be more appropriate for specific real games (cf. Section 5.4). However, there is a need for a comprehensive mathematical model of rational **non-probabilistic** decision making under uncertainty. Theoretically, *KBR* covers decision scenarios and games which do not contain probabilistic elements in their descriptions, e.g., perfect information games with deterministic payoffs. To the extent that the notion of perfect information game corresponds to reality, *KBR* theory reflects real games as well.

2.2 Highest payoff principle

A naïve, pre-epistemic understanding of rational decisions is that

A rational player chooses a strategy which yields the highest payoff. (1)

This formulation captures the ‘greedy’ element of rationality, i.e., going for the highest payoff, but totally ignores its epistemic component: as we have already agreed, a rational player decides not on the basis of what is true in the world, but rather on the basis of what he knows/believes. In particular, the highest actual payoff associated with a given strategy can be unknown to the player who, therefore, will not be able to take this payoff into account. The knowledge requirement naturally leads to the following epistemically explicit reading of (1):

*A rational player chooses a strategy which yields the highest **known** payoff.* (2)

2.3 Knowledge vs. maximin

The question is how (2) corresponds to the aforementioned knowledge-based form of Harsanyi’s Maximin Postulate that states

If you do not know that you will get more than your maximin payoff, always use a maximin strategy.

It turns out that the two fundamental elements of rational decision making, knowledge and maximin, essentially coincide within the standard game-theoretical setting of finite extensive-form games with deterministic payoffs. In the rest of this section, we will show that

the maximin solution among all strategies the player deems possible

corresponds to

the highest known payoff to the best of player’s knowledge.

2.4 Highest Known Payoff of a given strategy/move

A *strategy* for player i is a function that assigns an *action* (a.k.a. *move*) to each node of a game in which i is making a decision. A strategy profile

$$\sigma = \{\sigma_1, \sigma_2, \dots, \sigma_n\}$$

is a collection of strategies σ_i for players $i = 1, 2, \dots, n$. Each strategy profile σ uniquely determines the *outcome* $O(v, \sigma)$ associated with a node v and σ : $O(v, \sigma)$ is the *terminal history* (i.e., the sequence of moves by players from v to a terminal node) in which each

move is made according to σ . Likewise, everyone who knows the game tree can calculate i 's payoff determined by σ and a node v : $U_i(v, \sigma)$.

The following observation, first made in [1]⁷, is the foundation of *KBR*-theory:

Proposition 1 *Let v be a node and i the player who makes a move at v ⁸. Then for every strategy σ_i of i , there is a unique Highest Known Payoff, $HKP_i(v, \sigma_i)$, such that*

1. *i knows that he gets at least $HKP_i(v, \sigma_i)$ when playing strategy σ_i ;*
2. *i does not know whether he gets any payoff greater than $HKP_i(v, \sigma_i)$.*

This is a very generic observation. Indeed, player i knows the finite game tree and hence a finite set of his payoffs, some of which are known to i as secured when playing strategy σ_i . Naturally, there is a highest payoff, $HKP_i(v, \sigma_i)$, which player i **knows** he will get when playing σ_i from node v .

For example, in the decision making schema in Figure 1, A has two strategies: $down_A$ with payoff 1, or $across_A$ which passes the choice to B whose intentions are unknown to A . Under these conditions,

$$HKP_A(A, down_A) = 1,$$

whereas

$$HKP_A(A, across_A) = 0,$$

since A considers both strategies by B , $across_B$ and $down_B$, possible, hence the strategy profile

$$\{across_A, down_B\}$$

is possible for A and brings A payoff 0. Given strategy $across_A$, out of two payoffs, 0 and 2, A knows that he gets at least 0, but does not know whether he gets payoff 2.

The situation is different in Game Two in Figure 3. B is now assumed to be a rational player who has his own payoffs (coinciding with those of A). Suppose also that A knows that B is rational.




Figure 3: Game Two

⁷Strictly speaking, [1] defines *HKP* for a move, but its extension from moves to strategies is straightforward.

⁸The condition that i makes a move at v is not really necessary.

Under these assumptions, B would play across seeking payoff 2 for himself (and for A as well). Moreover, this is now known to A ! Therefore, A knows that B will not play down, hence for A the only possible strategy profile containing $across_A$ is

$$\{across_A, across_B\},$$

with A 's payoff 2. Under these circumstances,

$$HKP_A(A, across_A) = 2.$$

In a version of Game Two in which A is not aware of B 's rationality, A considers both moves by B possible and again $HKP_A(A, across_A) = 0$.

Theorem 2 $HKP_i(v, \sigma_i)$ is equal to the minimum of i 's payoffs for all strategy profiles σ containing σ_i which are deemed possible by i at v :

$$HKP_i(v, \sigma_i) = \min\{U_i(v, \sigma) \mid \sigma \text{ is a possible strategy profile containing } \sigma_i\}.$$

Proof. To the best of his knowledge, player i considers some strategy profiles possible and some not. For each strategy σ_i , and each possible strategy profile σ containing σ_i , i can calculate i 's payoff $U_i(v, \sigma)$ of σ starting from v , hence i can calculate their minimum

$$m = \min\{U_i(v, \sigma) \mid \sigma \text{ is a possible strategy profile containing } \sigma_i\}. \quad (3)$$

In particular, there is a possible strategy profile σ' containing σ_i such that $U_i(v, \sigma') = m$. In order to establish that $m = HKP_i(v, \sigma_i)$, it suffices to establish two things:

1. i knows that he gets at least m when playing strategy σ_i ;
2. i does not know whether he gets any payoff greater than m .

Property 1 is secured since i knows that he gets at least his maximin payoff m when playing strategy σ_i . In order to establish 2, consider any payoff $q > m$. Then player i cannot know that i gets q when playing strategy σ_i since there is a possible strategy profile σ' containing σ_i which brings payoff m , $m < q$. \square

Corollary 1 *The highest known payoff $HKP_i(v, \sigma_i)$ is always known to player i .*

Proof. Indeed, as an intelligent player, i knows Theorem 2 and is capable of computing his maximin value (3) which is equal to $HKP_i(v, \sigma_i)$. \square

Another way of proving Corollary 1 by referring to general properties of knowledge was used in [1].

The epistemic role of the Highest Known Payoff of a strategy can be illustrated by Figure 4. For each node v , player i , and strategy σ_i of i , the (finite) set of all possible payoffs for i divides naturally into two intervals.

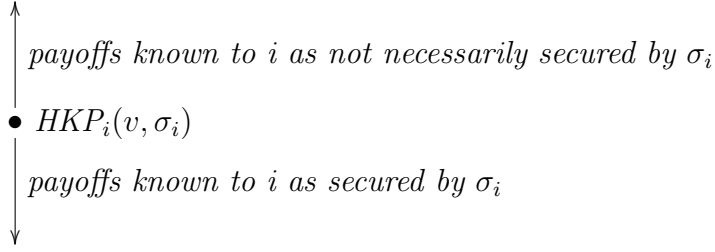


Figure 4: Highest Known Payoff of a strategy

Interval 1: $\{p \mid p \leq HKP_i(v, \sigma_i)\}$. For each payoff p from this interval, player i knows that he will get at least p when playing σ_i from v .

Interval 2: $\{q \mid q > HKP_i(v, \sigma_i)\}$. For each payoff q from this interval, player i knows that if he plays σ_i , then his opponents have response strategies which i cannot rule out to the best of his knowledge and which bring i a payoff strictly less than q . So payoff q is known to i as not necessarily secured by strategy σ_i .

Definition 1 *The Highest Known Payoff of a move \mathcal{M} at a node v , $HKP_i(v, \mathcal{M})$, is the maximum of Highest Known Payoffs of strategies that start with move \mathcal{M} at v :*

$$HKP_i(v, \mathcal{M}) = \max\{HKP_i(v, \sigma_i) \mid \text{for all strategies } \sigma_i \text{ making move } \mathcal{M} \text{ at } v\}.$$

It is clear that $HKP_i(v, \mathcal{M})$ is known to i and attainable when i plays strategy σ_i which realizes the maximum of $HKP_i(v, \sigma_i)$'s.

2.5 Best Known Strategy/Move

Definition 2 *Strategy σ_i is a best known strategy for player i at a given node v if σ_i has the greatest highest known payoff among all i 's strategies, i.e.,*

$$HKP_i(v, \sigma'_i) \leq HKP_i(v, \sigma_i)$$

for all i 's strategies σ'_i .

The best known strategy is not necessarily unique even for generic games, e.g., if strategies differ at some node which is not accessible during the game, then these strategies are different but apparently yield the same payoff under each response strategy and hence have the same HKP . However, if we limit our attention to the **first move** of a strategy, then we come to the notion of *the best known move* for a given player at a given node, which is already unique for generic games. *This notion reflects the idea of a definitive rational choice.*

Definition 3 A move \mathcal{M} for player i at a node v is the **best known move** if it is the first move of a best known strategy for i at v .

Equivalently, the best known move is that which has the greatest *HKP*, cf. Definition 1 and [1].

Theorem 3 [1] *At each node v of a generic game, there exists a unique best known move.*

Proof. Let i be the player who makes a move at v . Existence follows from the fact that for each strategy σ_i , there is a well-defined highest known payoff $HKP_i(v, \sigma_i)$. To prove uniqueness, note that best known strategies are those which have the greatest $HKP_i(v, \sigma_i)$, which is, of course, unique for i at node v , by definition. We claim that all best known strategies at a given node of a generic game start with the same move. Indeed, if two strategies start with different moves in the game tree, they have disjoint sets of terminal nodes and hence disjoint sets of payoffs for a given player. Such strategies could not have the same *HKP*. \square

2.6 Highest Known Payoff at a node

Definition 4 *The Highest Known Payoff for player i at a node v , $\mathbf{HKP}_i(v)$, is the maximum of $HKP_i(v, \sigma_i)$'s for all strategies σ_i by i :*

$$\mathbf{HKP}_i(v) = \max\{HKP_i(v, \sigma_i) \mid \text{for all strategies } \sigma_i \text{ by } i\}.$$

Naturally, $\mathbf{HKP}_i(v)$ equals $HKP_i(v, \sigma_i)$ for the best known strategy σ_i at v .

Figure 5 shows that $\mathbf{HKP}_i(v)$ separates two intervals.

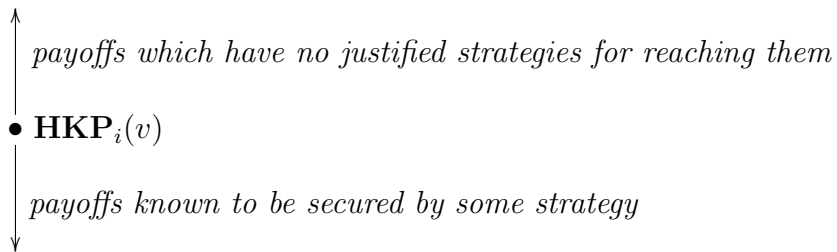


Figure 5: Highest Known Payoff at a node

Interval 1: $\{p \mid p \leq \mathbf{HKP}_i(v)\}$. These payoffs are known to i to be secured by a certain strategy, and i knows exactly which strategy to use to secure these payoffs.

Interval 2: $\{q \mid q > \mathbf{HKP}_i(v)\}$. Neither of the payoffs from this region has a justified strategy for player i to secure this payoff. Unless the game description says otherwise, i cannot rationally expect more than the maximin payoff $\mathbf{HKP}_i(v)$ at node v .

2.7 Knowledge meets maximin

Theorem 2 shows that the maximin solution among all strategies the player deems possible indeed corresponds to the highest known payoff to the best of player's knowledge.

Our analysis of knowledge and rationality in deterministic perfect information games can be summarized by the following two principles, which we accept as basic postulates of the *KBR*-model (cf. [1] for a more detailed analysis).

Rationality Postulates:

I. *A rational player chooses a maximin solution over all strategies of others which the player deems possible.*

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

Postulate I is the epistemically explicit form of Harsanyi's Maximin Postulate. Likewise, (II) is merely Harsanyi's Mutually Expected Rationality Postulate expressed in epistemic language.

As a corollary of Theorem 2, Rationality Postulate I can be equivalently formulated as

I'. *A rational player chooses a move with the best known payoff.*

For the rest of the paper, we will be using Postulates I and I' interchangeably.

2.8 Knowledge-based rational decision method

A *decision method* for a certain class of games is a recipe for choosing moves at game nodes. There are several well-known decision methods used for perfect information games: Nash equilibrium and subgame perfect equilibrium, backward induction solution, pure maximin, (iterated) elimination of dominated strategies, etc. A decision method is *knowledge-based* if decisions are made exclusively on the basis of a player's knowledge at the moment of making this decision. In particular, since none of the factors other than knowledge is explicitly given in the description of perfect information games, we assume that all rational players are making knowledge-based decisions. In such a case, we expect epistemic states of players to be sufficiently specified, otherwise the game is considered under-defined.

Definition 5 *The knowledge-based rational decision (KBR) method chooses a move with the best known payoff at each node.*

Theorem 4 is the principal result of *KBR*-theory ([1]).

Theorem 4 [1] *Any generic perfect information game with rational players and arbitrary epistemic states of players has a unique KBR solution.*

Proof. At any node of any such game, there is a unique best known move which is known to the player who makes a decision. By Rationality Postulate I', the player moves accordingly. □

3 Uniqueness of knowledge-based rationality

The goal of Section 2 has been to argue that in perfect information games, with their finite game trees and deterministic payoffs and no other assumptions for making a decision, the *KBR* decision method of choosing a move with the best known payoff at each node coincides with the epistemic form of Harsanyi's Maximin Postulate.

A decision method is *definitive* for a certain class of games if it provides a unique choice of action (move) at each node of every generic game in this class. This condition rules out speculative 'solutions' such as 'all moves are rational,' etc.

In this Section, we compare *KBR* with other decision methods and show that *KBR* is the only definitive method of playing perfect information games which is consistent with Rationality Postulates I and II. A decision method is *knowledge-based rational* if its choice of action (move) is consistent with Rationality Postulates I and II.

3.1 Uniqueness Theorem

Theorem 5 [KBR Theorem] *For perfect information games, knowledge-based rational decision is the only decision method which is definitive and knowledge-based rational.*

Before we proceed with proving this theorem, consider some examples of decision methods.

1. Nash equilibrium and subgame perfect equilibrium;
2. Backward induction solution;
3. Pure maximin;
4. Eliminating dominated strategies.

3.2 Nash equilibrium and subgame perfect equilibrium

Nash equilibrium for extensive form games (cf. [15]) is neither definitive, nor rational, nor knowledge-based. First, we notice that Nash equilibrium and subgame perfect equilibrium are calculated on the basis of a game tree without taking into account epistemic states of players. Therefore, we should not expect either to be based on knowledge.

Consider Game Three in Figure 6.




Figure 6: Game Three

There are two Nash equilibria in this game: $\{across_A, across_B\}$ and $\{down_A, down_B\}$ which allow A to move either way, hence this method is not definitive.

Neither of these equilibria is uniformly rational. Indeed, if A knows that B is rational, A 's rational choice is across since A knows that B will then play across and deliver a payoff of 2 to A . Therefore, the second of the equilibria, $\{down_A, down_B\}$, is not rational for A . If A does not know that B is rational, then A considers both moves by B possible and, by Rationality Postulate I, A should play down. In this case, equilibrium $\{across_A, across_B\}$ is not rational for A .

Subgame perfect equilibrium (cf. [15]) when applied to Game Three eliminates equilibrium $\{down_A, down_B\}$, which is the only rational solution for A if A considers both moves by B possible. Hence this method is not necessarily rational either.

3.3 Backward induction

The backward induction solution is rational and knowledge-based, but not definitive. Moreover, backward induction is a method of **avoiding uncertainty** by calculating opponents' strategies. In particular, backward induction does not provide answers under uncertainty.

In Game Three, if A knows that B is rational, then the backward induction solution is $\{across_A, across_B\}$. However, if B is rational, but A does not know this and considers both moves by B possible, then backward induction does not work, thus leaving A without any recommendation at all.

However, the backward induction solution, when it exists, coincides with the *KBR*-solution. To make this claim precise, let us use the standard common knowledge requirement for backward induction.

Theorem 6 *In perfect information generic games, common knowledge of rationality yields the backward induction solution which coincides with the KBR-solution.*

Proof. Indeed, backward induction reasoning yields that at each node, the corresponding player **knows** the actual strategy used by his opponents from this node onwards. There is no uncertainty at any node. Naturally, both backward induction and *KBR* choose the same uniquely determined rational move. \square

3.4 Pure maximin

Pure maximin is not necessarily rational. In Game Three, if A knows that B is rational, then A 's best strategy is $across_A$ and this strategy is known by A to bring him a payoff of 2. The pure maximin solution $down_A$ brings A payoff 1 and hence is not rational.

3.5 Eliminating dominated strategies

Eliminating dominated strategies is an epistemically correct method which, however, is not definitive. We refer the reader to [15] for exact definitions and restrict our attention here to an example.

In Game Three, where A considers both moves by B possible, neither of A 's strategies, $across_A$ and $down_A$, is dominated and hence cannot be eliminated. Therefore, eliminating dominated strategies alone does not provide a definitive answer here.

However, if eliminating dominated strategies provides a definitive answer, the remaining strategy yields the *KBR* solution.

3.6 Proof of the KBR Theorem

Proof. Fix a generic game and a node v . The *KBR*-method consists of choosing the move which yields the highest known payoff for a given player i at v , $\mathbf{HKP}_i(v)$. Such a move exists and is unique (Theorem 3).

Suppose i has to make a move at v and chooses a unique move \mathcal{M} . Let m be the highest payoff which i knows is secured by \mathcal{M} . By Theorem 3, there is a unique such m .

Case 1. $m = \mathbf{HKP}_i(v)$. Then \mathcal{M} is the *KBR* choice since for generic games, by Theorem 3, different moves have different \mathbf{HKP} 's.

Case 2. $m < \mathbf{HKP}_i(v)$. Choosing \mathcal{M} contradicts Rationality Postulate I' since \mathcal{M} is not the best known move for i at v . Indeed, the best known move at v corresponds to $\mathbf{HKP}_i(v)$, which is different from m .

Case 3. $m > \mathbf{HKP}_i(v)$. This case is impossible, by definition of $\mathbf{HKP}_i(v)$. Indeed, the highest payoff m which i knows is secured by move \mathcal{M} cannot be higher than the highest known payoff at v . \square

4 Relationships to Aumann's model of rationality

In this section, we compare the *KBR* model (which may be called Harsanyi/*KBR* rationality) with Aumann's knowledge-related model of rationality [3].

Aumann's rationality has been mathematically formalized in set-theoretical Aumann structures, and we refer the reader to [3]. The verbal account of Aumann's rationality has been summarized in [3] as follows:

“... call i *rational at v* if there is no strategy that i knows would have yielded him a conditional payoff at v larger than that which in fact he gets.”

The mathematical formulation of Aumann's rationality considers *irrational* any choice by i of a strategy σ_i in a situation when i **knows** that there is another strategy σ'_i which strictly dominates σ_i . All strategies which are not irrational are considered *rational*.

This definition of rationality works well when uncertainty at a node can be completely resolved by epistemic reasoning, e.g., under common knowledge of rationality in perfect information games, or when the player knows that one strategy strictly dominates all others. However, Aumann’s rationality does not help to make decisions in general situations, e.g., when there is a choice of several strategies, none of which strictly dominates the others.

For example, in Game One (Apollo 13 Dilemma in Figure 1), Aumann’s approach considers rational both strategies by A , *down* and *across*, since neither dominates the other. However, as we have already discussed in the Introduction, the intuitive solution *down* is well-justified and supported by both Harasnyi’s Maximin Postulate and its epistemic form, the *KBR*-method.

In Game Two (Figure 3) and Game Three (Figure 6), when A is not aware of B ’s rationality, Aumann’s approach again considers rational both strategies by A , *down* and *across*, and does not give a definitive solution, whereas there is a unique *KBR*-solution *down*.

Likewise, in the Centipede game, when players are not aware of each other’s rationality, Aumann’s definition considers rational any move at nodes 1–4, which does not lead to a solution.

On the other hand, however, Aumann’s approach is consistent with that of *KBR*.

Proposition 2 *Any KBR-rational strategy is Aumann-rational, but not vice versa.*

Proof. A *KBR*-rational strategy σ_i yields the highest known payoff at a given node v . Therefore, there cannot be another strategy σ'_i about which player i knows that σ'_i strictly dominates σ_i since otherwise,

$$HKP_i(v, \sigma'_i) > HKP_i(v, \sigma_i)$$

and hence σ_i cannot be *KBR*-rational.

Examples of Games One, Two, and Three show that there are Aumann-rational strategies which are not *KBR*-rational. □

Finally, if Aumann’s rationality yields a definitive answer, then this answer is the *KBR*-solution. In this respect, *KBR*-rationality may be regarded as a definitive extension of Aumann’s rationality.

5 Discussion

5.1 Logic of Knowledge

As in [1], we base our informal reasoning on the properties of knowledge using the principles of modal logic **S5**, which is the standard choice for Game Theory (cf. [2, 3, 4, 5, 6, 8, 9, 11, 15, 18, 19] and many others). In particular, knowledge modalities \mathbf{K}_P satisfy the postulates

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.

5.2 Manipulation by leaking true information

Tracking epistemic states of players helps to describe and study learning in games and manipulation by controlling information flow.

Consider the following game tree in Figure 7.




Figure 7: Game Four

In Game Four, A is not aware of B 's and C 's rationality. The *KBR* solution suggests A 's choosing left to secure payoff 2. Actually, A gets 4 which is more than expected since B will play left. Payoff for B , as well as for C , is 2.

Suppose also that B and C are smart enough to understand this. B can then manipulate A by leaking the **true information** that C is rational. A then knows that C will play left, rationally plays right and gets payoff 3 which is what A expected. B 's and C 's payoffs are 4 and 3 respectively, both higher than before. C does not have an incentive to disclose to A that B is rational as well, and B wins without ever making a move!

This example demonstrates an interesting principle: *more knowledge yields a higher known payoff but not necessarily a higher actual payoff. So 'nothing but the truth' can be misleading. Knowing 'the whole truth' however, yields a higher actual payoff.*

5.3 Relaxing rationality conditions

Of course, there are better ways to play the Centipede game, e.g., if some of the conditions of Theorem 1 are relaxed.

Game Five. Consider a version of the Centipede game in Figure 2 in which A decides to make one irrational move across at node 5, and this, along with rationality of players at all other nodes, is commonly known. Then both players choose across at all nodes, A 's payoff is 4, and B 's payoff is 5. Indeed, by assumption, A plays across at 5. At node 4, B knows that A will play across at 5, hence B rationally plays across at 4 as well. B 's reasoning is known to A , hence A plays across at 3, etc.

Game Six. Player B decides to move across at node 4 (and this, along with the rationality of players at all other nodes, is commonly known). A then plays down at node 5 and secures payoff 5 for himself and payoff 3 for B . By assumption, B plays across at node 4. A knows this and plays across at node 3 after rationally choosing payoff 5 over payoff 3. Similarly, B plays across at 2 since B 's payoff is then equal to 3, which is higher than B 's payoff 2. Finally, A chooses across at 1 to secure his payoff 5 and hence B 's payoff 3.

The aforementioned solutions of Games Five and Six are *KBR*-solutions, which also can be regarded as backward induction solutions given common knowledge of specific irrational moves by players.

5.4 Introducing probabilities

Game Seven. Player A can do even better if he is allowed to play probabilistically. Suppose, in Figure 2, at node 5, A plays across with probability 0.51 and down with probability 0.49, and this, along with the rationality of players at all other nodes, is commonly known. Then B 's expected payoff when playing across at node 4 is 4.02 which is a higher payoff than playing down at node 4. So B rationally chooses across at 4. Likewise, A 's expected payoff at node 3 when playing across is 4.49, which is higher than his payoff of down at node 3. Therefore, A plays across at 3. The same reasoning justifies playing across at nodes 1 and 2 as well. As a result, both players play across at nodes 1–4, and A tosses a 0.51/0.49 coin at node 5. The average payoff for A is 4.49 (and can be made arbitrary close to 4.5 by playing with ‘almost’ 0.5 probability at node 5) and the average payoff for B is 4.02.

This solution may be regarded as an example of probabilistic *KBR*-reasoning that is yet to be worked out in detail. The example shows that if players have reliable information about a priori probabilities of other players' strategies, they can do considerably better than following a pure *KBR* solution which, in this example, would still suggest playing down at each node⁹.

Another possibility is allowing a player to make a probabilistic assessment of uncertainty s/he faces; in particular, of the strategy choices made by the other players. Such

⁹Such a gap between pure and probabilistic *KBR* reasoning was anticipated by A. Brandenburger in a private communication.

an epistemic approach has been studied within a framework of belief-based interactive epistemology, cf. a survey [7].

6 Acknowledgments

The author is grateful to Adam Brandenburger for kind attention to this work. The author is also indebted to Mel Fitting, Vladimir Krupski, Florian Lengyel, Loes Olde Loohuis, Wojtek Moczydlowski, Elena Nogina, Rohit Parikh, Graham Priest, Olivier Roy, and Cagil Tasdemir for useful and inspiring discussions.

Special thanks to Karen Kletter who has been selflessly editing numerous versions of this paper.

References

- [1] S. Artemov. *Knowledge-based rational decisions*. Technical Report TR-2009011, CUNY Ph.D. Program in Computer Science, 2009.
- [2] R. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976.
- [3] R. Aumann. Backward Induction and Common Knowledge of Rationality. *Games and Economic Behavior*, 8:6–19, 1995.
- [4] R. Aumann and A. Brandenburger. Epistemic conditions for Nash equilibrium. *Econometrica: Journal of the Econometric Society*, 64(5):116–1180, 1995.
- [5] C. Bicchieri. *Rationality and Coordination*. Cambridge University Press, Cambridge, 1993.
- [6] 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.
- [7] A. Brandenburger. The power of paradox: some recent developments in interactive epistemology. *International Journal of Game Theory*, 35:465–492, 2007.
- [8] A. Brandenburger. *Epistemic Game Theory*. Lecture of June 22, 2008, <http://pages.stern.nyu.edu/~abranden> .
- [9] 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> .

- [10] D. Ellsberg. Recurrent Objections to the Minimax Strategy: Rejoinder. *The Review of Economics and Statistics*, 41(1):42–43, 1959.
- [11] R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.
- [12] 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.
- [13] J.C. Harsanyi. *Rational behaviour and bargaining equilibrium in games and social situations*. Cambridge Books, 1986.
- [14] J. Nash. *Non-Cooperative Games*. Ph.D. Thesis, Princeton University, Princeton, 1950.
- [15] M. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, 1994.
- [16] R. Rosenthal. Games of Perfect Information, Predatory Pricing, and the Chain Store Paradox. *Journal of Economic Theory*, 25(1):92–100, 1981.
- [17] L.J. Savage. *The Foundations of Statistics* (Second Revised Edition). New York: Dover, 1972.
- [18] J. van Benthem. Rational Dynamics and Epistemic Logic in Games. *International Game Theory Review*, 9(1):13–45, 2007.
- [19] P. Vanderschraaf and G. Sillari. *Common Knowledge*. Stanford Encyclopedia of Philosophy, 2007.
- [20] J. von Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100:295–320, 1928. Translated as “On the Theory of Games of Strategy,” pp. 13–42 in *Contributions to the Theory of Games*, v. IV, Annals of Mathematical Studies, 40, A.W. Tucker and R.D. Luce, eds., Princeton University Press, 1959.
- [21] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.