

2009

TR-2009015: Knowledge-Based Rational Decisions and Nash Paths

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-2009015: Knowledge-Based Rational Decisions and Nash Paths" (2009). *CUNY Academic Works*.
http://academicworks.cuny.edu/gc_cs_tr/336

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.

Knowledge-based rational decisions and Nash paths

Sergei Artemov

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

November 24, 2009

Abstract

The knowledge-based rational decision model (*KBR*-model) offers an approach to rational decision making in a non-probabilistic setting, e.g., in perfect information games with deterministic payoffs. The *KBR*-model uses standard game-theoretical assumptions and suggests following a strategy yielding the highest payoff which the agent can secure to the best of his knowledge.

In this report, we prove a conjecture by A. Brandenburger that in perfect information games, each *KBR*-path is a Nash path.

1 Introduction

Our model of rational decision making uses standard game-theoretical assumptions, e.g., Harsanyi's Maximin Postulate ([6]),

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

and the traditional postulate of rational decision-making:

A rational player chooses a strategy that yields the highest payoff to the best of his knowledge.

As noted in [1, 2], if a rational player operates in a non-probabilistic setting and bases his decision on **knowledge** rather than luck, guesswork, sudden opponent cooperation or error, etc., the aforementioned postulates lead to the same mathematical model of decision making that we call the Knowledge-Based Rational decision model (*KBR*-model).

Though Game Theory often considers decisions based on beliefs rather than knowledge (cf. [4]), a special theory of knowledge-based decision making looks to be appropriate

as well. The principal difference between knowledge and belief is the factivity property of knowledge that beliefs do not necessarily possess. In some situations, players seem to make decisions on the basis of their knowledge and not merely on their beliefs: military, high-stakes commercial, juridical decisions, etc. Furthermore, according to commonly accepted properties of knowledge such as positive and negative introspection¹ ([5]), the decision-maker is aware of what he knows and what he does not know and hence is capable of distinguishing what he actually knows from what he merely believes without actual knowledge.

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

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

In a somewhat more formal setting, *KBR* assumes two Rationality Postulates (cf. [1]):

1. A rational player chooses a strategy yielding the highest payoff the agent can secure to the best of his knowledge. Equivalently, a rational player chooses a maximin solution over all strategy profiles the player deems possible.
2. Postulate (1) is commonly known and accepted by rational players.

Postulate (1) is the epistemically explicit form of Harsanyi's Maximin Postulate. Similarly, (2) is merely Harsanyi's Mutually Expected Rationality Postulate ([6]) expressed in epistemic language.

2 Strategies, profiles, paths

In this paper, we consider generic extensive-form perfect information games which include specification of the relevant states of knowledge for each player. In particular, for each player i , it is specified *which strategy profiles σ are known to be impossible* by player i . All other profiles are called *epistemically possible for player i* . By the factivity property of knowledge, no player is playing a strategy known to be impossible by any of the players.

¹A.k.a. the Axiom of Transparency and the Axiom of Wisdom [7].

2.1 Simple examples

Consider Game One in Figure 1 in which both players are rational and aware of each other's rationality². There are two strategies for each player, *down* or *across*, which makes the total number of strategies equal to four:

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




Figure 1: Game One, rationality is commonly known.

Since *B* is rational, he knows that he is not playing *down_B*. *A* knows this, concludes that *B* is playing *across_B*, and rationally decides to play *across_A*. Moreover, *B* knows this as well. So, in Game One, for each player, the only strategy which is epistemically possible is

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

which happens to also be the backward induction solution.

Now consider Game Two in Figure 2 on the same game tree as Game One, in which both players are rational but their rationality is not mutually known. In particular, each player considers each strategy of the other player possible. *B* knows that he is playing




Figure 2: Game Two, rationality is not mutually known.

across_B. The epistemically possible strategy profiles for *B* are

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

²Taking into account the length of the game, this is equivalent to assuming the common knowledge of rationality.

Player A considers either strategy by B possible and cannot rationally expect to get a payoff greater than 1 if he plays $across_A$. By Rationality Postulate 1, A cannot choose $across_A$. Therefore, for A , the epistemically possible strategy profiles are

$$\{down_A, down_B\}, \quad \{down_A, across_B\}.$$

Note that though there is more than one strategy profile epistemically possible for each player, A and B each have a unique, epistemically possible strategy, namely $down_A$ for A and $across_B$ for B ; we call them *KBR-strategies*. All uncertainty concerning possible profiles stems from insufficient information about other players, but each player has a unique rational strategy of his own under these uncertainties.

The *KBR-solution* of this game is the strategy profile consisting of *KBR-strategies* of individual players. In other words, the *KBR-solution* is the profile that will actually be played.

In Game Two the *KBR-solution* is

$$\{down_A, across_B\}$$

and the *KBR-path* is

$$down_A.$$

2.2 Rational player's view of a perfect information game

In general, we assume the usual understanding of strategies: a player's strategy specifies what he does at each of his nodes, if reached (cf. [3]). To render this precise, we need the notion of a subgame in an epistemic setting.

For each node v of game G , a *subgame* G_v is determined by the rooted subtree with root v : epistemically possible strategy profiles for i in G_v are epistemically possible strategy profiles for i in game G relativized to the nodes from the subtree with root v .




Figure 3: Subgame G_v of game G .

Lemma 1 [1, 2] *At each node of a generic perfect information game, there is a unique move (called a KBR-move) by the corresponding player that yields the highest payoff that player can secure to the best of his knowledge (called highest known payoff).*

Proof. Consider a node v and player i making a move at v . Given epistemically possible-for- i strategy profiles, there is a unique highest payoff h_v which i can secure by one of his strategies at node v . Note that h_v is the unique maximin value over all epistemically possible profiles for player i . Since the game is generic, there is at node v one move for player i through which he can possibly receive payoff h_v ; indeed, a different move will generate a game path in a subtree which does not contain payoff h_v for i at all. \square

Corollary 1 *In a generic game with rational players, there is a unique KBR-move at each node.*

Definition 1 *A KBR-strategy for a given player i is a collection of KBR-moves at nodes where i makes a move.*

Corollary 2 *For each generic game with rational players, there is a unique KBR strategy profile and players actually play this profile.*

These observations lead to the following informal picture of epistemically possible strategy profiles for each rational player A ; here B is any player other than A . At a node at




Figure 4: Strategy profiles that (rational) player A considers possible.

which A makes a move, only the KBR -move is epistemically possible for A . At a node at which some other player makes a move, A may consider multiple moves as epistemically possible. All epistemically possible strategy profiles for A are constituted from A 's unique KBR -strategy σ_A and strategies by others considered epistemically possible by A .

Corollary 3 [1, 2] *The real payoff for each player at a given node is greater than or equal to the highest known payoff at this node.*

2.3 Pure maximin and backward induction solution

Pure maximin strategy for a given player i corresponds to the reading of a game in which i has no information whatsoever about other players' epistemic states. Then i considers all moves by opponents epistemically possible. Under these conditions, pure maximin is a special case of KBR .

Another special case of the KBR -solution is given by the backward induction solution BI under Aumann's conditions of common knowledge of rationality ([3]). In this case, each

player has sufficient information to exactly determine his opponents' move at each node. For each player, there is only one epistemically possible strategy profile: the *KBR*-solution of the game.

2.4 A subpath of a *KBR*-path is a *KBR*-path as well

Each strategy profile σ determines a unique path P associated with σ : P starts at the root node and moves according to σ .

It follows from the definition of a subgame that for each player, epistemically possible profiles in game G and its subgame G_v coincide on nodes from G_v . This observation leads to the following lemma.

Lemma 2 *Let G be an extensive-form perfect information game, P be a *KBR*-path in G , and v be a node in P . Then the part P_v of P starting from v is the *KBR*-path of the subgame G_v in which v is the root node.*

We refer the reader to Figure 5.




Figure 5: Subpath of a *KBR*-path.

3 Each *KBR*-path is Nash

In this section, we prove a conjecture by A. Brandenburger, stated in a private communication, that in perfect information games with rational players, each *KBR*-path is a Nash path. A path P is Nash iff there is a Nash strategy profile σ such that P is the σ -path. This result enables us to compare *KBR* with such classical decision-making methods as iterated elimination of strictly dominated strategies *IESDS*, Nash Equilibria *NE*, and the backward induction solution *BI* (cf. discussion in Section 4).

We first observe that the (unique) *KBR* strategy profile for a generic PI game is not necessarily a Nash profile. A *KBR*-player plays to the best of his knowledge, which may

be limited: there might be better moves unknown to him. Consider Game Two in Figure 2. The *KBR* strategy profile is

$$\sigma = \{down_A, across_B\},$$

and the *KBR*-path is

$$P = down_A.$$

It is easy to see that σ is not a Nash profile, since A can unilaterally improve his payoff by playing *across*. Indeed, the strategy profile will then be

$$\sigma' = \{across_A, across_B\}$$

and the path

$$P' = (across_A, across_B)$$

that yields A 's payoff 2.

On the other hand, there is a Nash strategy profile

$$\sigma'' = \{down_A, down_B\}$$

that has the same path P as σ .

Theorem 1 *In a PI game with rational players, the KBR-path is a Nash path.*

Proof. Induction on maximal game length $n(G)$. The base: $n(G)=1$. Then the *KBR*-path consists of one rational move which constitutes a Nash profile.

The Induction Hypothesis: suppose the theorem claim holds for all games with length less than k .

The Induction Step. Consider a PI game G in an extensive tree-like form such that $n(G) = k$. Let P be its *KBR*-path, A be the player who is making a move at root node r , and $1, \dots, m$ be immediate successors to r . By G_1, \dots, G_m , we denote subgames of G with roots at $1, \dots, m$ respectively.

Let b be the highest known payoff for player A at root node r (cf. [1]), i.e., the highest payoff that A knows he can secure at r :

$$b = \mathbf{HKP}_A(r).$$

Then for any strategy σ_A by A , there is a strategy profile σ containing σ_A and epistemically possible for A such that A 's payoff of σ , $U_A(\sigma)$ is less than or equal to b . By Corollary 3, A 's payoff on path P , $U_A(P)$ is greater than or equal to b .

Without loss of generality, assume that A 's root move is $(r, 1)$, and that the rest of P , P_1 occurs within G_1 . By Lemma 2, P_1 is the *KBR*-path in G_1 . By the Induction Hypothesis, since $n(G_1) < k$, P_1 is a Nash path in G_1 , i.e., there is a Nash strategy profile σ^1 such that P_1 is its path in G_1 .




Figure 6: Game G

Our goal now is to extend σ^1 to a Nash strategy profile σ for all of G without changing its path P . For this, we have to define the moves of each player at nodes other than those from G_1 .

At root node r , A 's move is $(r, 1)$ as suggested by P : make it part of σ . It now remains to define moves at all nodes of games G_2, \dots, G_m .

Pick subgame G_i , $i = 2, \dots, m$ and consider the following auxiliary ‘maximin game’ on the same tree. In this maximin game, player A tries to win more than his highest known payoff b , and all other players are playing against this goal. Label a leaf **S** (for **S**uccess) if A 's payoff at this leaf is greater than b , and **F** (for **F**ailure) otherwise. Backward induct to label all other nodes of G_i and define moves for each node v of G_i .

Case 1. A makes a move at node v , and all immediate successors to v are labeled **F**. Then label v as **F** and pick an arbitrary move for A at v .

Case 2. A makes a move at node v , and there is an immediate successor to v that is labeled **S**. Then label v as **S** and pick a move for A from v to one of its immediate **S**-successors.

Case 3. A player other than A makes a move at v , and there is an immediate successor to v labeled **F**. Then label v as **F** and pick a move from v to one of its immediate **F**-successors.

Case 4. A player other than A is making a move at v , and all immediate successors to v are labeled **S**. Then label v as **S** and pick an arbitrary move at v .

Let us denote σ_A^i the strategy by A , and σ_{-A}^i the collection of strategies by all other players in the maximin game on G_i .

The following lemma shows that A cannot win the maximin game.

Lemma 3 *The root node i of G_i is labeled **F**.*

Proof. Since b is the highest known payoff for A at the root node, given σ_A^i , there should be a collection δ_{-A}^i of strategies for other players in G_i (deemed possible by A) such that

A 's payoff of the profile $\{\sigma_A^i, \delta_{-A}^i\}$ is less than or equal to b . Let P' be the path of $\{\sigma_A^i, \delta_{-A}^i\}$ in G_i . We claim that each node of P' is labeled **F**. Backward induction on the length of P' . The leaf of P' is labeled **F** since it indicates A 's payoff on P' which is not greater than b . Consider a node v of P' whose immediate successor in P' is labeled **F**. If v is an A -node³, all immediate successors to v in G_i are labeled **F**, hence v is labeled **F**. If v is a non- A -node⁴, v is labeled **F** as well. So all nodes of P' are labeled **F**, including the root node i of G_i . \square

Now we define the desired strategy profile σ on G_i -nodes:

For each $i = 2, \dots, m$, σ restricted to G_i -nodes coincides with $\{\sigma_A^i, \sigma_{-A}^i\}$.

This concludes the construction of σ and it remains to be shown that

1. σ 's path is P ;
2. σ is a Nash profile in game G .

Item (1) is obvious, since the first move of σ is $(r, 1)$ and the rest of the path is P_1 .

Lemma 4 *σ is a Nash strategy profile.*

Proof. Present σ as a collection of A 's strategy σ_A and non- A -strategies σ_{-A} .

Players other than A cannot improve their payoff by unilaterally deviating from σ_{-A} given σ_A . Indeed, changes outside G_1 do not alter the outcome. Changes inside G_1 cannot improve the payoff since within G_1 , σ is a Nash strategy profile.

Fix σ_{-A} and consider an arbitrary strategy σ'_A for A .

Case 1. The first move of σ'_A is $(r, 1)$. Then the consequences of σ'_A are limited to changes in A 's strategy within G_1 that cannot yield a better payoff for A , since σ is a Nash profile on G_1 .

Case 2. The first move of σ'_A is (r, i) with some $i = 2, \dots, m$. Suppose, en route to contradiction, that

$$U(\{\sigma'_A, \sigma_{-A}\}) = b' > b$$

and let P' be the path in G_i corresponding to $\{\sigma'_A, \sigma_{-A}\}$. By backward induction on the node depth, we show that all nodes of P' are labeled **S**. Base: the leaf node of P' is labeled **S** since P' delivers A 's payoff $b' > b$. Let v be a node in P' whose immediate successor in P' is labeled **S**. If v is an A -node, then v should be labeled **S** by definition of the labeling process. If v is a non- A -node, then P' 's move at v is made according to σ , which indicates that all immediate successors to v in G_i are labeled **S**, hence v is labeled **S** as well.

We have arrived at a contradiction to Lemma 3 that states i is labeled **F**. This proves Lemma 4. \square

³I.e., A is making a move at v .

⁴I.e., some player other than A is making move at v .

This completes the proof of Theorem 1. □

As an easy corollary to this theorem, we conclude that each backward induction path is a Nash path. Indeed, apply Theorem 1 to the variant of the game in which common knowledge of rationality is assumed. For such games, the resulting *BI*-path will be the *KBR*-path. By Theorem 1, this path is Nash. Likewise, each pure maximin path is Nash as well, since the maximin profile is the *KBR*-profile with players ignorant of each others' rationality.

4 Discussion

It follows easily from definitions that in PI games, *KBR* strategy profiles survive iterated elimination of strictly dominated strategies (*IESDS*). The backward induction solution *BI* operates under the common knowledge of rationality assumption (cf. [3]). The pure maximin solution *MAXM* is justified for players ignorant of each other's rationality. Both *BI* and *MAXM* are special cases of *KBR*, which lacks such limitations. These observations, together with Theorem 1, suggest that major methods such as *IESDS*, Nash Equilibria *NE*, *BI*, *MAXM*, and *KBR* are compatible. However, only *KBR* always produces a justified unique solution (for generic games). *IESDS* and *NE* may be regarded as nondefinitive approximations to *KBR*.

For example, consider Game Two in Figure 2. None of the strategies is strictly dominating, hence the *IESDS* strategy profiles here are

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

There are two *NE* strategy profiles

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

and two *NE*-paths

$$down_A \text{ and } (across_A, across_B).$$

As always, there is a unique *KBR* strategy profile, that coincides here with the *MAXM* profile:

$$\{down_A, across_B\},$$

and one *KBR*-path that happens to be the *MAXM*-path as well:

$$down_A.$$

There is one *BI*-path

$$(across_A, across_B),$$

that is also the *KBR*-path in a version of Game One in which common knowledge of rationality of players is assumed.

Some correlations between methods can be seen in this example: all *NE*-profiles are *IESDS*-profiles and the *BI*-path, *MAXM*-path, and *KBR*-path are all *NE*-paths.

Figure 7 illustrates relationships between the aforementioned solution sets (ovals) and paths (bullets). Ovals represent sets of *IESDS*- and *NE*-paths, respectively. Bullets represent the *BI*-path, *MAXM*-path, and *KBR*-path.




Figure 7: Comparing methods

The empty bullets are a reminder that *MAXM*- and *BI*-paths are justified only under special conditions, e.g., complete ignorance of each others' rationality (*MAXM*), or common knowledge of rationality (*BI*). The arrows indicate that under corresponding conditions, the *MAXM*-path and *BI*-path become the *KBR*-path.

5 Acknowledgments

The author is grateful to Adam Brandenburger for stating the conjecture, and his permanent attention to this project. The author is also indebted to Vladimir Krupski, Elena Nogina, Rohit Parikh, and Cagil Tasdemir for useful and inspiring discussions.

Special thanks to Karen Kletter for editing this report.

References

- [1] S. Artemov. *Rational decisions in non-probabilistic settings*. Technical Report TR-2009012, CUNY Ph.D. Program in Computer Science, 2009.
- [2] S. Artemov. *Knowledge-based rational decisions*. Technical Report TR-2009011, CUNY Ph.D. Program in Computer Science, 2009.
- [3] R. Aumann. Backward Induction and Common Knowledge of Rationality. *Games and Economic Behavior*, 8:6–19, 1995.

- [4] A. Brandenburger. The power of paradox: some recent developments in interactive epistemology. *International Journal of Game Theory*, 35:465-492, 2007.
- [5] R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.
- [6] J.C. Harsanyi. *Rational behaviour and bargaining equilibrium in games and social situations*. Cambridge Books, 1986.
- [7] M. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, 1994.