On Tolerance Analysis of Games with Belief Revision

Çağıl Taşdemir

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

Abstract. Aumann's Rationality Theorem claims that in perfect information games, common knowledge of rationality yields backward induction (BI). Stalnaker argued that in the belief revision setting, BI does not follow from Aumann's assumptions. However, as shown by Artemov, if common knowledge of rationality is understood in the robust sense, i.e., if players do not forfeit their knowledge of rationality even hypothetically, then BI follows. A more realistic model would bound the number of hypothetical non-rational moves by player i that can be tolerated without revising the belief in i's rationality on future moves. We show that in the presence of common knowledge of rationality, if n hypothetical non-rational moves by any player are tolerated, then each game of length less than $2n+3$ yields BI, and that this bound on the length of model is tight for each n. In particular, if one error per player is tolerated, i.e., $n = 1$, then games of length up to 4 are BI games, whereas there is a game of length 5 with a non-BI solution.

1 Introduction

Aumann proved that in games of perfect information, common knowledge of rationality yields backward induction [3]. Stalnaker showed that if players are allowed to revise their beliefs in each other's rationality in response to surprising information, this is not the case [10]. In [7], Halpern showed that the difference between the two lies in how they interpret the following counterfactual statement: "If the player were to reach vertex v, then she would be rational at vertex v."

Let us consider the game in Figure 1 which is due to Stalnaker and which Halpern uses to point out the difference in Aumann's and Stalnaker's arguments. Assume that it is common knowledge that the actual state is (dda). This means that Ann plays *down* (d) in vertex v_1, Bob plays *down* (d) in vertex v_2, and Ann plays *across* (a) in vertex v_3, and that all of this is common knowledge between Ann and Bob. To say that some fact F is common knowledge between Ann and Bob means that Ann knows F, Bob knows F, Ann knows that Bob knows F, Bob knows that Ann knows F and so on. So if we assume that the state (dda) is common knowledge, this means that all moves are known right at the beginning of the game. The question is whether (dda), which is different

S. Artemov and A. Nerode (Eds.): LFCS 2013, LNCS 7734, pp. 350–358, 2013.

than the backward induction solution (aaa), can be the solution of the game in the presence of common knowledge of rationality.

Here it should also be noted that Stalnaker has no problem with the formal correctness of Aumann's proof. Aumann's framework does not allow belief revision. Stalnaker, on the other hand, allows players to revise their beliefs after a non-rational move by another player, even if that mentioned non-rational move is hypothetical.

Let us look at the game in Figure 1 from Stalnaker's perspective: At (dda), Ann is rational at v_1, because Bob is playing d at v_2. At v_2, Bob revises his belief on Ann's rationality, due to her hypothetical non-rational move (or the surprising information) a at v_1, and considers Ann's playing d at v_3 also possible as a result of this belief revision. He plays d and he is rational. Ann is rational at v_3 by playing a.

While Halpern layed out the differences in Aumann's and Stalnaker's arguments, Artemov in [2] showed that in perfect information games with Stalnaker-style belief revision setting, if players maintain their beliefs in each other's rationality in all, even hypothetical situations, i.e., if there is so-called *robust knowledge of rationality* in the game, then the only solution of the game is the backward induction. That is, if Bob does not revise his beliefs on Ann's rationality at v_2, (dda) cannot be the solution of the game in the presence of common knowledge of rationality.

Other works on epistemic foundations for backward induction include [1], [4], [5], [8] and [9].

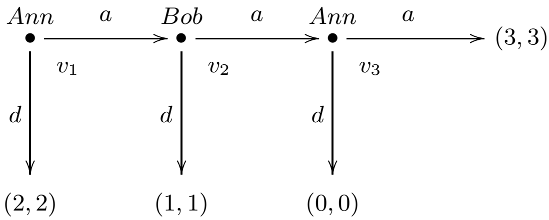

Fig. 1. 3-move game

In this paper, we will investigate the case where it is common knowledge that players are rational (in Stalnaker's sense) at all vertices of the game tree, and they tolerate n hypothetical non-rational moves of other players. If $n = 0$, we end up with Stalnaker's framework where after 1 error, players revise their beliefs.

2 Game Models and Rationality

Halpern extends Aumann models to represent N-player extensive form games with perfect information where players can revise their beliefs [7]. An extended model is a tuple

$$M = (\Omega, K_1, ..., K_N, s, f)$$

where Ω is a set of states of the world, K_i is the information partition of player i, and s maps each state $\omega \in \Omega$ to a strategy profile $s(\omega) = (s_1, ..., s_N)$ where s_i is player i's strategy at state ω. Function f, called *selection function*, maps state-vertex pairs to states. Informally, $f(\omega, v) = \omega'$ means that ω' is the closest state to ω where vertex v is reached. Let $h_i^v(s)$ denote player i's payoff if strategy profile s is played starting at vertex v. Let P be the function that maps non-terminal nodes to players to indicate the player moving at a given node.

Definition 1. Player i is *Aumann-rational*, or A-rational, at vertex v in state ω if for all strategies s^i such that $s^i \neq s_i(\omega)$, $h_i^v(s(\omega')) \geq h_i^v(s_{-i}(\omega'), s^i)$ for some $\omega' \in K_i(\omega)$ where $s_{-i}(\omega')$ denotes the strategy profile of the players other than i at state ω'.

Note that according to this definition, a player is rational as long as her strategy in the current state ω yields her a payoff at least as good as any of her other strategies in *some* state that she considers possible at ω.

Definition 2. Player i is *Stalnaker-rational*, or *S-rational*, at vertex v in state ω if i is A-rational at v in state $f(\omega, v)$.

Substantive rationality is rationality (A-rationality or S-rationality, depending on which framework we are working with) at all vertices of the game tree.

The formalization of selection functions is due to Halpern [7], and the main idea of a selection function f is for each state ω and vertex v to indicate the epistemically closest state $f(\omega, v)$ to ω in which v is reached. Halpern assumes that the selection function f satisfies the following requirements:

- F1. Vertex v is reached in $f(\omega, v)$.
- F2. If v is reached in ω, then $f(\omega, v) = \omega$.
- F3. $s(f(\omega, v))$ and $s(\omega)$ agree on the subtree below v.

3 Tolerating Hypothetical Errors

Our model extends Halpern's so that the selection function now satisfies an additional requirement $F4_n$ (given below) in order to represent the n-tolerance of the players. We will give the definitions of an error-vertex and the condition $F4_n$ simultaneously.

The following definition extends Aumann's rationality to hypothetical moves.

Definition 3. We say that a *move m at v in ω is rational* if player $i = P(v)$ is Aumann-rational at v in some state ω' which has the same profile as some $\widetilde{\omega} \in K_i(\omega)$ except, possibly, for the move m which is plugged into v.

Definition 4. Given a state ω, a vertex v is an *n-error vertex*, if n is the least natural number ≥ 0 such that each player makes not more than n non-rational moves (possibly hypothetical) at vertices v' from the root to v in states $f(\omega, v')$. Obviously, the root vertex is always 0-error. If there are $2k$ moves from the root to v, then each player makes $\leq k$ moves there and v is at most a k-error vertex. Other examples will be discussed later in this section.

The following condition reflects the idea of n-tolerance which is built-in into the selection function: sets of future scenarios are not revised after $\leq n$ (possibly hypothetical) non-rational moves of each player.

> **Condition F4$_n$.** For each state ω and k-error vertex v with $k \leq n$ and $i = P(v)$, if $\omega' \in K_i(f(\omega, v))$, then there exists a state $\omega'' \in K_i(\omega)$ such that $s(\omega')$ and $s(\omega'')$ agree on the subtree below v.

Halpern uses a similar condition to model Aumann's framework, which says that players consider at least as many strategies possible at ω as at $f(\omega, v)$; and this applies to all vertices in the game tree. Our condition F4$_n$ says the same thing for $\leq n$-error vertices, hence limiting the tolerance level in the game to n (possibly hypothetical) non-rational moves per player. In other words, in an n-tolerance game, players will not revise their beliefs about rationality for the first n hypothetical non-rational moves of those players.

Example 1. Consider the game in Figure 1. The following extended model is from [7]

The strategy profiles are as follows:

- $s^1 = (dda)$
- $s^2 = (ada)$
- $s^3 = (add)$
- $s^4 = (aaa)$: this is the BI solution.
- $s^5 = (aad)$

The extended model is $M_1 = (\Omega, K_{Ann}, K_{Bob}, s, f)$ where

- $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5\}$
- $K_{Ann} = \{\{\omega_1\}, \{\omega_2\}, \{\omega_3\}, \{\omega_4\}, \{\omega_5\}\}$
- $K_{Bob} = \{\{\omega_1\}, \{\omega_2, \omega_3\}, \{\omega_4\}, \{\omega_5\}\}$
- $s(\omega_j) = s^j$ for $j = 1 - 5$
- $f(\omega_1, v_2) = \omega_2$, $f(\omega_1, v_3) = \omega_4$, $f(\omega_2, v_3) = \omega_4$, $f(\omega_3, v_3) = \omega_5$, and $f(\omega, v) = \omega$ for all other ω and v.

It is assumed that the actual state is ω_1 with $s(\omega_1) = (dda)$, and this is commonly known to players. Let us check which vertices are erroneous:

- v_1 is a 0-error vertex.
- v_2 is a 1-error vertex since Ann's move from v_1 to v_2 is not rational in ω_1. She only considers (dda) possible at this node and changing the move at v_1 to a results in (ada), which would make Ann non-rational at v_1.
- v_3 is a 1-error vertex, by an easy combinatorial argument. Ann was not rational at v_1, so v_3 is at least 1-error vertex. However, it is at most 1-error, since the move at v_2 is made by Bob, and it cannot change the maximum of error counts at v_3. However, let us check that Bob is rational in moving from v_2 to v_3 at $f(\omega_1, v_2) = \omega_2 = ada$. In (ada), Bob considers both (ada) and

(add) possible, and plugging the hypothetical move a into vertex v_2 would result in strategy profile (aaa) that corresponds to state ω_4 in which Bob is rational at v_2.

4 Belief Revision with Tolerance

Example 2. Let us consider the game in Figure 1 again. Note that in this game, with the model M_1, in the presence of common knowledge of substantive rationality the realized strategy profile, i.e., (dda), is different than the backward induction solution (aaa). We will also assume common knowledge of substantive rationality, and show that (dda) cannot be the solution of the 1-tolerant version of this game.

Since players are 1-tolerant, the selection function f should satisfy the condition F4_n with $n = 1$. This means that the first hypothetical error for each player is tolerated. In particular, even if Ann and Bob make one hypothetical error each, those will be tolerated and beliefs in rationality will not be revised.

Therefore in a 1-tolerance game, if we assume that the state (dda) is common knowledge, we need to consider only three strategy profiles:

- $s^1 = (dda)$: This is the original strategy profile which is commonly known.
- $s^2 = (ada)$: This is the revised state at v_2.
- $s^3 = (aaa)$: This is the revised state at v_3.

The extended 1-tolerance game model is $M_2 = (\Omega, K_{Ann}, K_{Bob}, s, f)$ where

- $\Omega = \{\omega_1, \omega_2, \omega_3\}$
- $K_{Ann} = K_{Bob} = \{\{\omega_1\}, \{\omega_2\}, \{\omega_3\}\}$
- $s(\omega_j) = s^j$ for $j = 1 - 3$
- $f(\omega_1, v_2) = \omega_2,\ f(\omega_1, v_3) = \omega_3,\ f(\omega_2, v_3) = \omega_3$.

The actual state is ω_1 with $s(\omega_1) = (dda)$. Let us count the number of errors in this model.

- v_1 is 0-error.
- v_2 is 1-error, since in order to (hypothetically) get from v_1 to v_2, Ann has to make a non-rational move, by the same reasoning as in Example 1.
- v_3 is again 1-error by trivial combinatorial reasons, as before. Moreover, Bob's move from v_2 to v_3 is rational at ω_2, by the same reasoning as before.

Condition F4_1 is obviously met, so this is a 1-tolerant model in which strategy profile (dda) is common knowledge. We'll see, however, that the substantive S-rationality condition is violated in this model, namely, Bob is not rational at v_2. Indeed, S-rationality in ω_1 at v_2 reduces to (Aumann-)rationality in $f(\omega_1, v_2)$ at v_2, i.e., in ω_2 at v_2. Since $s(\omega_2) = (ada)$, the real move at v_2 is *down* which is not rational because of the better alternative *across*.

Example 3. Figure 2 shows an extensive form 1-tolerance game of length 5. Assuming common knowledge of substantive rationality, we will show that there exists a non-BI solution, namely ($dddda$).

The strategy profiles are as follows:

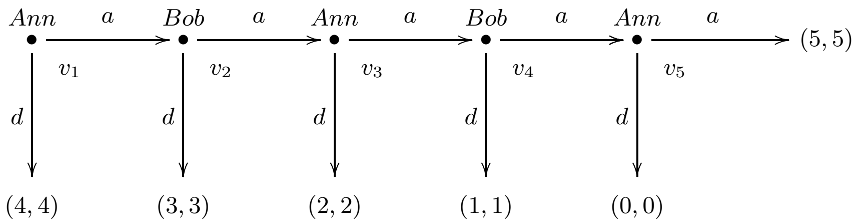

Fig. 2. 5-move game

- s^1 is the strategy profile $(dddda)$
- s^2 is the strategy profile $(addda)$
- s^3 is the strategy profile $(aadda)$
- s^4 is the strategy profile $(aaada)$
- s^5 is the strategy profile $(aaaaa)$
- s^6 is the strategy profile $(aaaad)$
- s^7 is the strategy profile $(aaadd)$.

Consider the extended model $A = (\Omega, K_{Ann}, K_{Bob}, s, f)$ where

- $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5, \omega_6, \omega_7\}$
- $K_{Ann} = \{\{\omega_1\}, \{\omega_2\}, \{\omega_3\}, \{\omega_4\}, \{\omega_5\}, \{\omega_6\}, \{\omega_7\}\}$
- $K_{Bob} = \{\{\omega_1\}, \{\omega_2\}, \{\omega_3\}, \{\omega_4, \omega_7\}, \{\omega_5\}, \{\omega_6\}\}$
- $s(\omega_j) = s^j$ for $j = 1 - 7$
- $f(\omega_1, v_2) = \omega_2,\ f(\omega_1, v_3) = \omega_3,\ f(\omega_1, v_4) = \omega_4,\ f(\omega_1, v_5) = \omega_5,$
  $f(\omega_2, v_3) = \omega_3,\ f(\omega_2, v_4) = \omega_4,\ f(\omega_2, v_5) = \omega_5,$
  $f(\omega_3, v_4) = \omega_4,\ f(\omega_3, v_5) = \omega_5,$
  $f(\omega_4, v_5) = \omega_5,$
  $f(\omega_7, v_5) = \omega_6.$

The actual state is $(dddda)$. Let us count the number of errors.

- v_1 is 0-error.
- Ann is not rational moving from v_1 to v_2 in state $f(\omega_1, v_2) = \omega_2$. Indeed, $s(\omega_2) = (addda)$, hence Ann moves *across* at v_1 while knowing that Bob will play *down* at v_2. Hence v_2 is 1-error.
- Bob, is not rational when moving from v_2 to v_3 in state $f(\omega_1, v_3) = \omega_3$. Indeed, $s(\omega_3) = (aadda)$, hence Bob moves *across* at v_2 while knowing that Ann will play *down* at v_3. Hence v_3 is 1-error by both Ann's and Bob's accounts.
- Ann is not rational moving from v_3 to v_4 in state $f(\omega_1, v_4) = \omega_4$. Indeed, $s(\omega_4) = (aaada)$, hence Ann moves *across* at v_3 while knowing that Bob will play *down* at v_4. Hence v_4 is 2-error on Ann's account.
- v_5 is 2-error by trivial combinatorial reasons. However, it is worth mentioning that Bob is rational when moving *across* from v_4 to v_5.

To secure 1-tolerance, we have to check the conclusion of $F4_1$ at all 0-error and 1-error vertices, in this case at vertices v_1, v_2, and v_3 which is quite straightforward. Indeed, selection function f does not add new indistinguishable states at these vertices, but just makes the corresponding vertex accessible in the revised state.

Since $K_{Ann}(\omega_1) = K_{Bob}(\omega_1) = \{\omega_1\}$, everything that is true at ω_1 will be common knowledge to Ann and Bob at that state. To check substantive rationality at ω_1, we need to check players' rationality in the following situations:
$S = \{(\omega_1, v_1), (\omega_2, v_2), (\omega_3, v_3), (\omega_4, v_4), (\omega_5, v_5)\}$

- Ann is rational at (ω_1, v_1). Since Bob plays d at vertex v_2, d is the rational move for Ann at (ω_1, v_1).
- Bob is rational at (ω_2, v_2). At (ω_2, v_2), Bob thinks Ann was not rational at v_1 but since we assume that each player tolerates one error, he does not revise his beliefs on her future rationality. So he looks at node v_3. Seeing that Ann is playing d at that node, he himself chooses to play d at v_2, which is the rational thing to do. So we can conclude that Bob is rational at (ω_2, v_2).
- Ann is rational at (ω_3, v_3). At (ω_3, v_3), Ann thinks Bob was not rational at node v_2. This time Ann tolerates Bob's error and does not revise her beliefs about his future rationality. She looks at node v_4. Seeing that Bob is playing d at that node, she chooses to play d at v_3, which is the rational thing to do.
- Bob is rational at (ω_4, v_4). At (ω_4, v_4), Bob thinks that Ann was not rational at node v_3. Since he will not tolerate one more error, he revises his beliefs and takes into the account the possibility of Ann's playing d at v_5. In this case, it is rational for him to play d.
- Ann is rational at (ω_5, v_5) regardless of her beliefs about Bob.

If we count the length of the game as the number of moves in its longest path in the game tree, this example shows that, assuming common knowledge of rationality and 1-tolerance, there exists a game of length 5, where a non-BI solution is realized.

Theorem 1. *In perfect information games with common knowledge of rationality and of n-tolerance, each game of length less than $2n + 3$ yields BI.*

Sketch of Proof: Let $m \leq 2n + 2$. We will show that all m-tolerant games are BI-games. At a vertex that at which the last move of a given path is made (such vertex is reachable from the root in $\leq 2n + 1$), Aumann-rationality yields the move that is dictated by the backward induction solution. Any other vertex v is reachable from the root in $\leq 2n$ steps. So there are at most $2n$ previous nodes prior to reaching v. Since no player makes two moves in a row, each player makes at most n moves prior to v, and even if all of these moves were erroneous, player $i = P(v)$ will tolerate them and do not revise his assumption of the common belief of rationality till the end of the game. By Artemov's argument in [2], this yields BI solution for the rest of the game.

Theorem 2. *The upper bound* $2n+3$ *from Theorem 1 is tight. Namely, for each* n*, there exists a perfect information game with common knowledge of rationality and of* n*-tolerance of length* $2n+3$ *which does not yield BI.*

Proof: Consider the straightforward generalization of Example 3 (which has length $5 = 2 \times 1 + 3$, i.e., corresponds to $n = 1$) to an arbitrary n in Figure 3. In particular, the profile

$$(dd \dots da)$$

is assumed to be commonly known and players to be n-tolerant.

The same reasoning as in Example 3 shows that this profile $(dd \dots da)$ is both rational and not BI. Since the strategy profile $(dd \dots da)$ is commonly known, and players are n-tolerant, Ann and Bob will not revise their beliefs in each other's rationality during the first $2n$ moves. The move at v_{2n+1} belongs to Ann. She is playing d according to the strategy profile $(dd \dots da)$ and she is rational (because Bob is playing d at v_{2n+2}). However, in order to decide whether Bob is rational at v_{2n+2}, we need to take into account Ann's hypothetical non-rational moves a to reach v_{2n+2}, and there are $n+1$ such moves. Therefore Bob may revise his beliefs on her rationality, consider Ann's playing d possible at v_{2n+3}, and this makes Bob's move d at v_{2n+2} rational. At the very last vertex, Ann is also rational since she plays a.

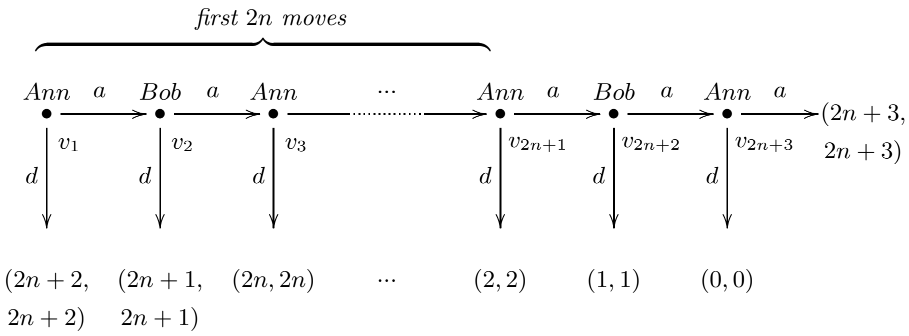

Fig. 3. Game of length $2n+3$

5 Conclusion and Future Work

We see a conceptual contribution of this work in stressing the role of tolerance in the analysis of perfect information games in the belief revision setting that accommodates both Aumann's and Stalnaker's paradigms. Stalnaker's players are zero-tolerant and give up their "knowledge of rationality" in hypothetical reasoning after the first hypothetical non-rational move of other players. Aumann's players are infinitely tolerant, and never give up their knowledge of rationality even when confronted, hypothetically, with strong evidence of the contrary. A natural problem of what happens in between, when the level of tolerance to hypothetical errors is a parameter of the game is addressed is this paper. Our findings indicate that for a given tolerance level n, short games, up to length $2n+2$ are Aumann's games, i.e., yield backward induction solutions only. Longer games of length $2n+3$ and greater can show Stalnaker's behavior based on the revision of player's belief of each other's rationality.

What does it say about games with human players who can be tolerant to some limited degree? One more parameter intervenes here: the nested epistemic depth of reasoning, which is remarkably limited for humans ([6]) to small numbers like one – two. In order to calculate BI, players have to possess the power of nested epistemic reasoning of the order of the length of the game. So, realistically, the BI analysis of human players applies to rather short games, say, of length three – four. According to Theorem 1, assuming 1-tolerance of players (which we regard as a meaningful assumption for humans) the only solution is backward induction.

In this game, with the given definition of rationality of hypothetical moves and common knowledge of a non-BI strategy profile as the actual state, we see no way to interpret the hypothetical errors as a move (signal) where the player is trying to reach the pareto-optimal payoff pair, which is the BI solution. A next logical step could be to look into this direction.

Acknowledgments. The author is deeply indebted to Sergei Artemov for his guidance, stimulating suggestions and encouragement.

References

1. Asheim, G.B.: On the epistemic foundation for backward induction. Mathematical Social Sciences 44, 121–144 (2002)
2. Artemov, S.: Robust knowledge and rationality. Technical Report TR-2010010, CUNY Ph.D. Program in Computer Science (2010)
3. Aumann, R.: Backward induction and common knowledge of rationality. Games and Economic Behavior 8, 6–19 (1995)
4. Baltag, A., Smets, S., Zvesper, J.A.: Keep "hoping" for rationality: A solution to the backward induction paradox. Synthese 169, 301–333 (2009); Knowledge, Rationality and Action, 705–737
5. Feinberg, Y.: Subjective reasoning: Dynamic games. Games and Economic Behavior 52, 54–93 (2005)
6. Ghosh, S., Meijering, B., Verbrugge, R.: Empirical reasoning in games: Logic meets cognition. In: Ågotnes, T., Alechina, N., Logan, B. (eds.) Proceedings Third Logics for Resource Bounded Agents Workshop (LRBA 2010), pp. 15–34 (2010)
7. Halpern, J.: Substantive rationality and backward induction. Games and Economic Behavior 37, 425–435 (2001)
8. Perea, A.: Epistemic game theory: Reasoning and choice. Cambridge University Press (2012)
9. Samet, D.: Hypothetical knowledge and games with perfect information. Games and Economic Behavior 17, 230–251 (1996)
10. Stalnaker, R.: Belief revision in games: forward and backward induction. Mathematical Social Sciences 36, 31–56 (1998)