

City University of New York (CUNY)

CUNY Academic Works

All Dissertations, Theses, and Capstone
Projects

Dissertations, Theses, and Capstone Projects

2-2016

Epistemic Considerations on Extensive-Form Games

Cagil Tasdemir

Graduate Center, City University of New York

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

More information about this work at: https://academicworks.cuny.edu/gc_etds/781

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

EPISTEMIC CONSIDERATIONS ON EXTENSIVE-FORM GAMES

by

ÇAĞIL TAŞDEMİR

A dissertation submitted to the Graduate Faculty in Computer Science in
partial fulfillment of the requirements for the degree of Doctor of
Philosophy, The City University of New York

2016

© 2016

ÇAĞIL TAŞDEMİR

All Rights Reserved

EPISTEMIC CONSIDERATIONS ON EXTENSIVE-FORM GAMES

by

ÇAĞIL TAŞDEMİR

This manuscript has been read and accepted by the Graduate Faculty in
Computer Science in satisfaction of the dissertation requirement for the
degree of Doctor of Philosophy.

Professor Sergei Artemov

Date

Chair of Examining Committee

Professor Robert Haralick

Date

Executive Officer

Professor Rohit Parikh

Professor Melvin Fitting

Professor Eric Pacuit

Supervisory Committee

THE CITY UNIVERSITY OF NEW YORK

EPISTEMIC CONSIDERATIONS ON EXTENSIVE-FORM GAMES

by

ÇAĞIL TAŞDEMİR

Adviser: Professor Sergei Artemov

Abstract

In this thesis, we study several topics in extensive-form games. First, we consider perfect information games with belief revision with players who are tolerant of each other's hypothetical errors. We bound the number of hypothetical non-rational moves of a player that will be tolerated by other players without revising the belief on that player's rationality on future moves, and investigate which games yield the backward induction solution.

Second, we consider players who have no way of assigning probabilities to various possible outcomes, and define players as conservative, moderate and aggressive depending on the way they choose, and show that all such players could be considered rational.

We then concentrate on games with imperfect and incomplete information and study how conservative, moderate and aggressive players might play such games. We provide models for the behavior of a (truthful) knowledge manipulator whose motives are not known to the active players, and look

into how she can bring about a certain knowledge situation about a game, and change the way the game will be played.

Acknowledgments

This thesis would not have been possible without my mentor Sergei Artemov. I am deeply grateful to him for his guidance, patience, vast knowledge and stimulating suggestions. I consider myself very lucky to have had the opportunity to be his student. I would also like to thank Rohit Parikh, who not only has been very encouraging in my studies, but also introduced me to many other related fields. Our weekly meetings were a pleasure and he has been a true inspiration for me. Melvin Fitting's logic class was the first formal logic class I took, and I learned a lot from him. I would also like to thank Eric Pacuit for his time and insightful comments.

I thank Joseph Driscoll and Lina Garcia for all their help at the Graduate Center and Hunter College.

I would like to thank my childhood friend Alper Ayvaci, with whom I made the journey to the US. He has always been available for my long distance questions no matter how busy he was. I would also like to thank Oliver Fluck for his support and encouragement. Over the years I have had many interesting conversations, academic and otherwise, with Yiğit Atılğan, Can Başkent, Loes Olde Loohuis, George Rabanca, Ivo Vigan and Yunqi Xue at the Graduate Center. These memories are invaluable to me. Konstantinos

Pouliasis has been a great friend and a fellow graduate student, and I have learned a lot from him. The pho and beer sessions with Onur Altındağ, the neighborhood outings with Ben Crist and Jen Frank, the unfinished cocktails with Sarah Raskin and the movies I dragged her to, and the inner city travels with Ekin Şentürk made life in New York City so much better, especially when I needed a break from studying. Thanks guys.

Special thanks to Gregory Paquette for being such a good *sport*, playing hangman with me whenever I wanted to, preparing delicious breakfasts, and listening to my goings on with glazed eyes without complaining.

Finally, I would like to express my endless gratitude to my family. My parents are my first educators, and thanks to them I find joy in learning. My sister amazes me every day with her wisdom and her unique approaches to problems. I am thankful for their inexhaustible love and support. I try not to take it for granted. It was great growing up in a family where everybody constantly questioned everything. Thank you.

Contents

1	Introduction	1
2	Background	3
2.1	Models of Knowledge	4
2.1.1	The Set Theoretical Definition of Knowledge	5
2.1.2	Kripke Model	7
2.2	Mutual Knowledge and Common Knowledge	8
2.3	Game Models, Formally	11
3	Tolerance Analysis of Perfect Information Games with Belief	
	Revision	15
3.1	Backward Induction	15
3.2	Perfect Information Games with Common Knowledge of Rationality	16
3.3	Game Models and Rationality	19
3.4	Tolerating Hypothetical Errors	21
3.5	Belief Revision with Tolerance	25
3.6	Discussion and Future Work	33

4	Choice under Uncertainty	35
4.1	Suitable Choice Functions	38
4.2	The IIA Condition	42
4.3	Related Work	44
4.4	Examples	46
5	Knowledge Manipulation	50
5.1	Two Examples	50
5.1.1	Three Envelopes	50
5.1.2	Manipulation by Leaking True Information	52
5.2	Knowledge Creation in Game Theoretic Situations	54
5.2.1	Our Model	54
5.2.2	Abstract Considerations	55
5.2.3	Representing Knowledge, A Third Model: History-Based Semantics	56
5.3	Creating Knowledge States	59
5.4	States of Knowledge	63
5.5	Using Knowledge Manipulation	68
5.5.1	An Example of Manipulation	72
5.5.2	Can KM Always Create Uncertainty?	75

5.5.3	Predicting the Play	77
5.5.4	A Note on Normal Form Games: Correlated Equilibrium	79
5.6	Related Work	81
5.7	Discussion	82
	Bibliography	84

List of Figures

1	A strategic-form game: Bach-Stravinsky	3
2	An extensive-form game	4
3	Stalnaker's game	13
4	Stalnaker's game, repeated	23
5	Stalnaker's game: 5-move version	27
6	Stalnaker's game: Generalized version of length $2n + 3$	33
7	Bach-Stravinsky game in extensive-form with generic payoffs	49
8	Artemov's game	53
9	A game tree and a history-based tree corresponding to it	63
10	Kripke structure M_α associated with node α	66
11	Bach-Stravinsky game in extensive-form with generic payoffs	68
12	Bach-Stravinsky game with KM's signals as moves	71
13	Another game tree	72
14	KM turns the game into a perfect information game	73
15	KM gives away some information but leaves the game as an imperfect information game	73
16	The reduced game	75

17	Prisoner's Dilemma in extensive-form with signals	77
18	Bach-Stravinsky game in normal form	79
19	Bach-Stravinsky game with knowledge manipulator	80

1 Introduction

Game theory studies situations where a person or an agent makes a decision, but the outcome depends also on the choices of the others. Game theory has been used to analyze situations such as auctions, mechanism design, voting systems, bargaining, behavioral economics, linguistics and so on. In all of these cases, before making a decision, in addition to reasoning about your own choices, it is important to reason about your opponents' choices. But of course your opponents are also capable of doing this reasoning; therefore you could reason about their beliefs about your reasoning process as well. Epistemic game theory treats this reasoning process of individuals as an essential component of the game, and draws its roots from epistemic logic – the logic of knowledge and belief – and economics.¹

Epistemic logic is the branch of logic that deals with formalizing knowledge and belief. This could be done for a single player, e.g. “Ann knows that Paris is the capital of France,” or for a group of players, e.g. “Ann believes that Bob believes that if he goes to the picnic, she will come too.” Epistemic logic allows us to assert that Ann is aware of the fact that Paris is the capital of France as in the first example, and that Ann has a belief about Bob's belief about her decision as in the second example. This is important because a player's decision in a game will be affected by what she knows about the

¹Of course one does need to do this back and forth reasoning in order to do game theory. However, epistemic game theory does take this perspective.

game and also by her beliefs about other players' beliefs.

(Non-cooperative) game theory is about strategies for games where players are considered to be rational, self-interested and calculating. It is therefore natural to include the behavior of interactive epistemic individuals in the study of games.

In this thesis, we study several epistemic situations in extensive-form games.

In particular, we

- do a tolerance analysis of perfect information games with belief revision,
- look into the choice functions of players in imperfect information games where probabilities are not available,
- investigate manipulation of players' knowledge in imperfect information games

2 Background

A non-cooperative game (just *game* from now on) is a game where every player makes her decisions independently from other players. No communication is allowed among them. Two models used to represent games are: (1) strategic-form, where games are represented using matrices, and (2) extensive-form where tree structures are used. In general, to define a finite game – without the epistemic component for now – we assume:

- a finite number of players,
- a finite number of choices available to each player,
- and a payoff for each player for every possible combination of choices that could be made by players. (Each such combination is called a *strategy profile*.)

Figure 1 shows an example of a strategic-form game with two players:

	Bach	Stravinsky
Bach	2,1	0,0
Stravinsky	0,0	1,2

Figure 1: A strategic-form game: Bach-Stravinsky

In the Bach-Stravinsky game, the couple wants to go to a concert but the wife (row) prefers Bach and the husband (column) prefers Stravinsky. But each

would rather go together than listen to their favorite composer by themselves.

In each cell, the wife's payoffs are listed first, and husband's second.

Figure 2 shows an example of an extensive-form game with two players:

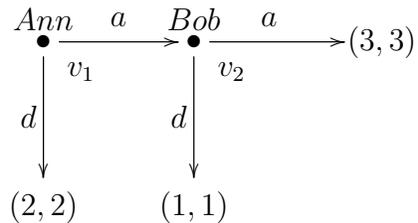


Figure 2: An extensive-form game

In this game, Ann and Bob are the players, each has a choice of playing *across* (*a*) or *down* (*d*) when it is their turn to play. Payoffs are defined on the terminal nodes (leaves), and the first payoff in each pair belongs to Ann and the second one to Bob. It is Ann's turn to play at node v_1 ; she chooses between playing *across* and *down*. If she chooses *down*, each player gets a payoff of 2, and the game ends. If she chooses *across*, the game proceeds to v_2 , which belongs to Bob, so it is now Bob's turn to make a move. If he chooses *down*, the game ends, and each player gets a payoff of 1. If he chooses *across*, the game ends, and each player gets a payoff of 3.

2.1 Models of Knowledge

In order to be able to talk about the epistemic component of games, we will first look at various models used to represent knowledge.

2.1.1 The Set Theoretical Definition of Knowledge

This section relies heavily on [65], [30]. The following model of knowledge is associated with Hintikka [40]. An *information structure* is a pair (Ω, P) where Ω is a set of states. A state is a full description of the world and the states are considered to be mutually exclusive. P is a function that assigns to each state ω a non-empty subset of states, $P(\omega)$. At state $\omega \in \Omega$, the player excludes all the states outside $P(\omega)$, and does not exclude any state in $P(\omega)$.

The three properties of information structures that are commonly associated with the term “rationality” are as follows:

P1. $\omega \in P(\omega)$: The player always considers the true state possible.

P2. If $\omega' \in P(\omega)$, then $P(\omega') \subseteq P(\omega)$: A player cannot hold the view that there exists a state u such that $u \in P(\omega')$ and $u \notin P(\omega)$ if he believes that $\omega' \in P(\omega)$.

P3. If $\omega' \in P(\omega)$, then $P(\omega') \supseteq P(\omega)$: If $\omega' \in P(\omega)$, and there is a state $u \in P(\omega)$ that is not in $P(\omega')$, then at ω , a rational player can conclude, from the fact that he cannot exclude u , that the state is not ω' , a state at which he would be able to exclude u , contradicting the assumption that $\omega' \in P(\omega)$.

The combination of these three properties is equivalent to the assumption that the information structure is *partitional*, namely that there exists a par-

tition of Ω (a collection of mutually exclusive subsets of Ω that completely cover Ω) such that $P(\omega)$ is the set, within the partition, that includes ω .

Let (Ω, P) be an arbitrary information structure (not necessarily partitional).

The event $E \subseteq \Omega$ is *known* at state ω if $P(\omega) \subseteq E$. This means, at ω the player knows E if he can exclude all states that are not in E . The statement *the player knows E* is identified with all states in which E is known, that is the set $K(E) = \{\omega : P(\omega) \subseteq E\}$. This definition implies the following properties:

- If $E \subseteq F$, then $K(E) \subseteq K(F)$.
- $K(E \cap F) = K(E) \cap K(F)$.
- $K(\Omega) = \Omega$.

Additional properties of the operator K are derived from assumptions about the information structure.

- I1. If we assume P1, we obtain $K(E) \subseteq E$.
- I2. If we assume P2, we obtain $K(E) \subseteq K(K(E))$.
- I3. If we assume P3, we obtain $\neg K(E) \subseteq K(\neg K(E))$.

2.1.2 Kripke Model

This model is due to Kripke [44]. A (single)-agent Kripke structure over a set of atomic propositions AP is a triple $\langle \Omega, R, \pi \rangle$ where Ω is a non-empty set of states, $R \subseteq \Omega \times \Omega$ is a binary relation (known as the *accessibility relation*) and $\pi : \Omega \rightarrow 2^{AP}$ is the *interpretation function* which assigns a truth value to every atomic proposition at every state $\omega \in \Omega$. $(M, \omega) \models \varphi$ denotes the notion that the formula φ is satisfied by the Kripke structure $M = \langle \Omega, R, \pi \rangle$ at state ω .

If φ is atomic, say p , $(M, \omega) \models p$ iff π assigns *true* to p , i.e. iff $p \in \pi(\omega)$. For the rest of the formulas, the satisfaction relation is defined inductively as follows:

- $(M, \omega) \models \neg\varphi$ iff $(M, \omega) \not\models \varphi$
- $(M, \omega) \models \varphi \vee \psi$ iff $(M, \omega) \models \varphi$ or $(M, \omega) \models \psi$
- $(M, \omega) \models K(\varphi)$ iff for all $v \in \Omega$ such that $\omega R v$, we have $(M, v) \models \varphi$.

Here $K(\varphi)$ denotes the fact that “ φ is known.”

We could also derive additional properties about knowledge in Kripke structures by imposing some constraints on the accessibility relation R . If R is reflexive, transitive and symmetric (i.e. it is an equivalence relation) we obtain the following for all $\omega \in \Omega$ and for every formula φ :

$$\text{K1. } (M, \omega) \models K(\varphi) \rightarrow \varphi.$$

$$\text{K2. } (M, \omega) \models K(\varphi) \rightarrow K(K(\varphi))$$

$$\text{K3. } (M, \omega) \models \neg K(\varphi) \rightarrow K(\neg K(\varphi))$$

These three properties K1-K3 in Kripke structures correspond to I1-I3 in information structures respectively. Informally, K1 means that if φ is known, φ is true. K2 says that if φ is known, then it is known that φ is known. K3 says that if φ is not known, then it is known that φ is not known.

We can go from a Kripke structure to an information structure by taking $P(\omega) = \{\omega' : (\omega, \omega') \in R\}$. Conversely, given an information structure, we can go to a Kripke structure by letting $R = \{(\omega, \omega') : \omega' \in P(\omega)\}$.

Comparison between the two models. Let φ be a formula. The event $E_\varphi = \{\omega : (M, \omega) \models \varphi\}$ is the set of all states where φ is true. $E_{K\varphi} = K(E_\varphi)$, which means that every state where $K(\varphi)$ is satisfied in Kripke model is a state where E_φ is known in the set-theoretical model of knowledge and vice versa. The proof stating the equality of two definitions can be found in [65]. Also see [30] and [35].

2.2 Mutual Knowledge and Common Knowledge

The concept of common knowledge is due to Lewis [45] and Aumann [7]. See [19] for a comparison between Lewis' and Aumann's definitions of common

knowledge.

To say that a fact is *mutual knowledge* means that everyone knows it. To say that a fact is *common knowledge* means that everyone knows it, everyone knows that everyone knows it, and so on ad infinitum.

In the set-theoretical model of knowledge, to say that “an event is mutual knowledge in some state” means in that state every player knows the event. To say that “an event is common knowledge in some state” means in that state the event is mutual knowledge, every player knows that the event is mutual knowledge, every player knows that every player knows that the event is mutual knowledge and so on. In the case of two players, we say that an event E is common knowledge in state ω if in state ω player 1 knows E , player 2 knows E , player 1 knows 2 knows E , player 2 knows player 1 knows E , and so on indefinitely.

Let Ω be the state space, and K_1 and K_2 two knowledge functions representing the knowledge of two players, 1 and 2. $K_i(E)$ is the set of states in which player i knows that event $E \subseteq \Omega$ occurs.

Define a third partition of Ω given the two partitions of the players. One partition *refines* another if every member of the first partition is a subset of a member of the second partition. The *meet* P_m of two partitions P_1 and P_2 is the partition of Ω such that

- (1) both partitions P_1 and P_2 are refinements of P_m ; and

(2) there is no refinement of P_m that satisfies (1).

For some event $E \subseteq \Omega$, let $\omega \in \Omega$ be the true state of the world. Then E is common knowledge in state ω if and only if E contains the member of P_m that contains ω .

In [7], Aumann showed that given an event $E \subseteq \Omega$, and a state $\omega \in \Omega$, if two players have a common probability measure over the set of states Ω and if their posterior probabilities of event E are common knowledge at state ω , then they must be equal. The posterior here is defined by conditioning the player's prior on her information cell. In [32], Geanakoplos and Sebenius showed that given a random variable $X : \Omega \rightarrow \mathfrak{R}$ - to be thought of as a bet between two players - and a state $\omega \in \Omega$ it cannot be common knowledge that both players have strictly positive conditional expectations from the bet. In [31], Geanakoplos and Polemarchakis showed that two players will always reach an agreement on their posteriors by announcing them back and forth. In [53], Parikh and Krasucki showed that when there are more than two players, communication done in pairs does not lead to common knowledge for the group. They showed, however, that conditional probabilities will converge in any fair protocol for communication, i.e., when none of the participants is blocked from communication.

2.3 Game Models, Formally

An N -player extensive form game with the epistemic component consists of the following:

- A finite set of N players
- A rooted tree, called the game tree
- One payoff for each player that is associated with every terminal (leaf) node. (i.e. a payoff tuple for every possible play.)
- A partition of the non-terminal nodes of the game tree in N subsets, one for each player. Each player's subset of nodes contains the nodes where it is that player's turn to make a move.
- Each set of nodes of a player is further partitioned in *information sets* such that:
 - There is a one-to-one correspondence between outgoing edges of any two nodes in the same information set.
 - Every path from the root to a terminal node can cross each information set only once.
- The above description of the game is common knowledge among players. [39]

A game is said to be a *perfect information game* if every member of every information set is a singleton. Any game without perfect information has *imperfect information*. A game is called *generic* if payoffs are different at every leaf of the game tree for each player.

To formalize knowledge, Aumann uses the partition model [9]:

A *strategy* of player i is a function s_i that assigns to each node v of i an action at that node. If we denote the set of i 's strategies by S_i , a knowledge system M (for the given game) is a tuple $M = (\Omega, K_1, \dots, K_N, s)$ that consists of

- a set Ω (the states of the world),
- for each player i , a partition K_i of Ω (i 's information partition).
- and a function s from Ω to $\prod_i S_i$

$s(\omega)$ represents the N -tuple of the players' strategies at the state ω . s_i is assumed to be measurable with respect to K_i , which means that players know their own strategy.

We will let P denote the function that maps non-terminal nodes to players to indicate the player moving at a given node.

Example 2.1 This example is due to Stalnaker [71], and it has been formalized by Halpern [36]. We will discuss it in detail in Chapter 3.

In this game,

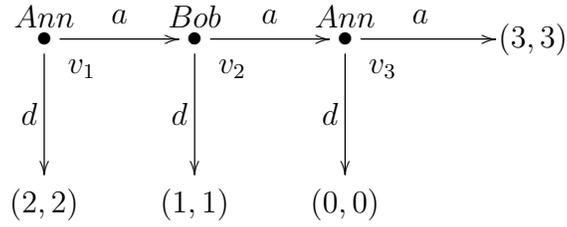


Figure 3: Stalnaker's game

- We have two players: Ann and Bob.
- The game tree is as depicted in Figure 3. The game starts at node v_1 , and at each node, the player whose turn it is to play has two strategies available to her/him: *across* (a) and *down* (d).
- At each terminal node, the first number in the pair denotes Ann's payoff, and the second one denotes Bob's if the game reaches this node.
- Nodes v_1 and v_3 belong to Ann, and node v_2 belongs to Bob.
- If this is a perfect information game, then at each node v the player who moves at v knows that the play is at v .

The epistemic component of this game is given as follows:

- There are five states of the world: $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5\}$
- The function s , that maps each state to a strategy profile is given as follows: $s(\omega_1) = dda$, $s(\omega_2) = ada$, $s(\omega_3) = add$, $s(\omega_4) = aaa$, $s(\omega_5) = aad$.

- Information partitions of Ann and Bob are:

$$K_{Ann} = \{\{\omega_1\}, \{\omega_2\}, \{\omega_3\}, \{\omega_4\}, \{\omega_5\}\}, \text{ and}$$

$$K_{Bob} = \{\{\omega_1\}, \{\omega_2, \omega_3\}, \{\omega_4\}, \{\omega_5\}\}$$

The given information partitions tell us that at all states Ann knows what state she is at, whereas Bob cannot distinguish between states ω_2 and ω_3 , which means that he does not know the action that Ann will select at node v_3 .

3 Tolerance Analysis of Perfect Information Games with Belief Revision

The work described in this chapter has appeared in [72].

3.1 Backward Induction

Backward induction is a common solution concept in extensive-form games. In generic perfect information games, the process is as follows:² Start with the non-terminal nodes whose children are only the terminal nodes (leaves). At each of these non-terminal nodes, choose the move, say m , that leads to a better payoff, say p , for the player who owns that node (m exists since we assume the game is generic); eliminate all the other moves originating from the non-terminal node; and eliminate the payoffs they yield as well, and replace the said node with the payoff p , therefore making it a terminal-node. Note that this is done for all non-terminal nodes whose children are only the terminal nodes. Now, iterate the procedure until there is one node (which is going to be a leaf) left. This gives us the BI *outcome*. [1]

If we look at our example from Figure 3, backward induction works as follows: The only non-terminal node whose children are only terminal nodes is v_3 . So

²Note that the backward induction process is more difficult to define on games with imperfect information.

we first decide what Ann would play at v_3 , then taking her move into account, we look at what Bob would play at the next such node, that is v_2 , and finally taking all of this into account we check what Ann would do at v_1 .

3.2 Perfect Information Games with Common Knowledge of Rationality

Aumann proved that in games of perfect information, common knowledge of rationality yields backward induction [9]. He calls a player rational if he chooses rationally at *all* his nodes, even at the ones that he *knows* will not be reached because he himself has precluded them by a previous choice. In a subsequent paper [1], where a new notion of rationality is introduced, this older version is referred to as *strategic rationality*.

This (first) approach has been criticized, in particular by Stalnaker in [71], who showed that if players are allowed to revise their beliefs in each other's rationality in response to surprising information, backward induction does not follow from common knowledge of rationality.

In [36], Halpern showed that the difference between the two approaches lies in how the following counterfactual statement is interpreted: "If the player were to reach vertex v , then she would be rational at vertex v ." It should also be noted that Stalnaker has no problem with the formal correctness of Aumann's proof. Aumann's framework in [9] talks about knowledge and has

no room for 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.³

Let us consider the game in Figure 3 which is due to Stalnaker and which Halpern uses to point out the difference in Aumann's and Stalnaker's arguments.

We will first examine the backward induction solution: At node v_3 , Ann is better off playing *across* since this move gives her a payoff of 3, which is more than the payoff of 0 she would get if she played *down*. Taking this into account, Bob is better off playing *across* at node v_2 if the game reaches this node. Playing *across* would result in a payoff of 3 for him, which is better than the payoff of 1 he would get if he played *down*. Taking all of this into account, Ann is better off playing *across* at v_1 because she prefers the game to reach v_2 (and then eventually to v_3) in which case her payoff would be a 3, compared to ending the game immediately by playing *down* at v_1 and getting a payoff of 2. So the backward solution would be (aaa) where the players play *across* at the node(s) where it is their turn to make a move.

Let us look at the same game from Stalnaker's perspective: Assume that it is common knowledge that the actual state is (dda) . This means that Ann plays *down* (d) at node v_1 , Bob plays *down* (d) at v_2 , and Ann plays *across* (a) at v_3 , and that all of this is common knowledge between Ann and Bob. This

³It would be worth pointing out that other approaches besides Stalnaker's are possible. See [67] for an example.

means that all moves are known at the beginning of the game. The question is whether (dda) , which is different than the backward induction solution (aaa) , could follow from assuming common knowledge of rationality. Also see [51] for an extended discussion on this.

At (dda) , Ann is rational⁴ at v_1 , because Bob is playing d at v_2 . (Remember that the state (dda) is common knowledge.) 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 laid out the differences in Aumann's and Stalnaker's arguments, Artemov in [4] 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 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. Baltag, Smets and Zvesper also make a similar point in [10].

In a recent work [1], Arieli and Aumann replaced the common "knowledge"

⁴We have not yet given a formal definition of rationality. At this point, it could be thought of as follows: A rational player will not choose a strategy if she *knows* that there is another strategy that will bring her a better payoff.

of rationality requirement of [9] by common “strong belief” of rationality, and also replaced the “strategies” by “plans,” which specify only what a player does at unprecluded nodes. Recall that a strategy of a player is a function that assigns to each of that player’s nodes an action at the nodes that he is active. A *plan* of a player is the restriction of a strategy s_i of that player to the set of the nodes that s_i does not preclude by an action at a previous node. The nodes that are not allowed by a plan are said to be precluded. A player is called *plan-rational* if at each of his unprecluded nodes h , he has no plan that he believes, conditional on h being reached, yields him a higher payoff than his current plan. A player is said to *strongly believe* a proposition X if for each node h that belongs to the said player, either he believes X at node h or X is inconsistent with h being reached. Given these definitions, they show that common strong belief of plan rationality entails the BI outcome. [1]

Epistemic foundations of backward induction has been studied extensively in [5], [10], [27], [59], [67].

3.3 Game Models and Rationality

Halpern extends Aumann models to represent N -player extensive form games with perfect information where players can revise their beliefs [36]. An extended model is a tuple $M' = (\Omega, K_1, \dots, K_N, s, f)$ where

- Ω is a set of states of the world,
- K_i is the information partition of player i ,
- s maps each state $\omega \in \Omega$ to a strategy profile $s(\omega) = (s_1, \dots, s_N)$ 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 . Again, let P be the function that maps non-terminal nodes to players to indicate the player moving at a given node.

Definition 3.1 *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 .

Also note that we assume that every player knows her own strategy, so $s_i(\omega) = s_i(\omega')$ if $\omega' \in K_i(\omega)$.

Definition 3.2 *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 [36], 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 .

3.4 Tolerating Hypothetical Errors

We would like to 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. So if $n = 0$, we would end up with Stalnaker's framework where after 1 error, players revise their beliefs.

Our model will extend Halpern's so that the selection function will now satisfy 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 3.3 *We say that a move m at vertex v in state ω is rational if player $i = P(v)$ is Aumann-rational at v in some state ω' which has the same profile as some $\tilde{\omega} \in K_i(\omega)$ except, possibly, for the move m which is plugged into v .⁵*

Definition 3.4 *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.

⁵By "a move m being plugged into v ", we mean that the move at vertex v is m rather than the one that is dictated by $\tilde{\omega}$.

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 3.1 We repeat Stalnaker's game tree from Chapter 2.3 below. The following extended model is the same as in Example 2.1.

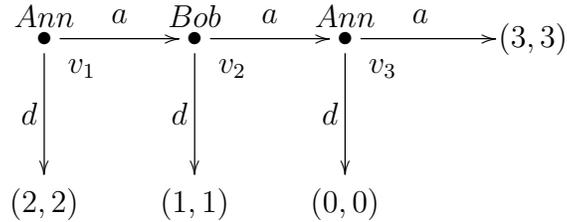


Figure 4: Stalnaker's game, repeated

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.

In state ω_1 :

- 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 . This is because (dda) would bring her a payoff of 2 whereas with (ada) she would only get 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 .

3.5 Belief Revision with Tolerance

Example 3.2 Let us consider the game in Figure 3 one more time. 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 of 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 $j = 1 - 3$
- $f(\omega_1, v_2) = \omega_2$, $f(\omega_1, v_3) = \omega_3$, $f(\omega_2, v_3) = \omega_3$, and $f(\omega, v) = \omega$ for all other ω and v .

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 3.1.
- v_3 is again 1-error by trivial combinatorial reasons, as before. Moreover, Bob's hypothetical 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 3.3 Figure 5 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 $(dddda)$.

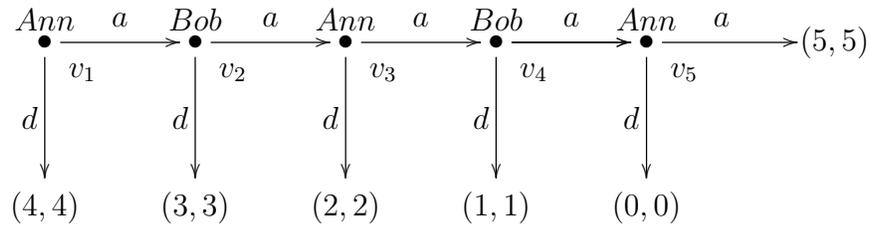


Figure 5: Stalnaker's game: 5-move version

The strategy profiles are as follows:

- 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 $j = 1 - 7$
- $f(\omega_1, v_2) = \omega_2, f(\omega_1, v_3) = \omega_3, f(\omega_1, v_4) = \omega_4, f(\omega_1, v_5) = \omega_5,$
 $f(\omega_2, v_3) = \omega_3, f(\omega_2, v_4) = \omega_4, f(\omega_2, v_5) = \omega_5,$
 $f(\omega_3, v_4) = \omega_4, f(\omega_3, v_5) = \omega_5,$
 $f(\omega_4, v_5) = \omega_5,$
 $f(\omega_7, v_5) = \omega_6,$ and
 $f(\omega, v) = \omega$ for all other ω and v .

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 (therefore subscribing to a payoff of 3 instead of 4). 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 (yet). So he looks at what will happen 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 what will happen 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 3.1 *In perfect information games with common knowledge of rationality and of n -tolerance, each game of length less than $2n + 3$ yields BI.*

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 $\leq 2n + 1$ steps), 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 [4], this yields BI solution for the rest of the game. \square

Something puzzles me about this: Consider a Centipede game where the BI outcome is that player 1 goes out at the first node. Suppose that the player

makes a “mistake” and stays in, but no further mistakes. I agree that the BI solution is played from then on, but how can you guarantee the BI outcome is reached?

Theorem 3.2 *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 3.3 (which has length $5 = 2 \times 1 + 3$, i.e., corresponds to $n = 1$) to an arbitrary n as in Figure 6. 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 3.3 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 . \square

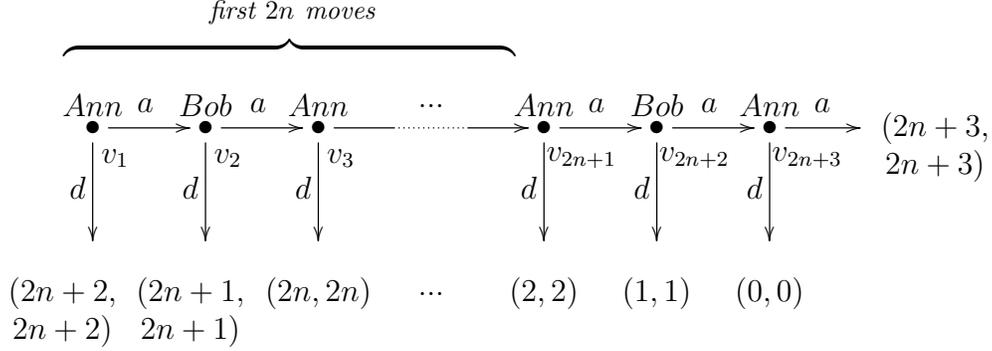


Figure 6: Stalnaker's game: Generalized version of length $2n + 3$

3.6 Discussion and Future Work

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 [34] to small numbers like one – two. In order to calculate the backward induction solution, 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. According to Theorem 3.1, assuming 1-tolerance of players (which we regard as a meaningful assumption for humans) the only solution is backward induction.

However, in the games we have discussed in this chapter, 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 who makes the hypothetical error is trying to reach the pareto-optimal payoff pair, which is the BI solution. A logical next step could be to look into this direction. This is similar to *forward induction*. See [13] and [43] for related discussions.

Another possible direction would be to incorporate trembling-hand equilibrium into our framework without making use of probabilities. Trembling-hand equilibrium is a model in which there is some small probability that a player will “tremble” and make an unintended, possibly non-rational move [69].

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. A natural problem of what happens in between, when the level of tolerance to hypothetical errors is a parameter of the game is discussed in this chapter.

4 Choice under Uncertainty

This chapter discusses how players might choose their strategies in imperfect information games. It is joint work with Parikh and Witzel, and has appeared in [57].

Recall that a game that does not have perfect information is an imperfect information game. Consider a player choosing between two acts A and B , whose outcomes are uncertain and depend on factors that the player does not fully know. However, for each pair of possible outcomes the player does know how she would choose. We would like to investigate whether the player then has a way of choosing between the acts which will work at least some of the time.

Suppose a player is in a situation of uncertainty where she has to choose between two moves L and R but does not know for sure what the outcome will be with either choice. Assume moreover that the player has no way of assigning probabilities to the various outcomes. In the absence of such information, how might the player choose?

One option is the maxmin route. A player can choose L if the *worst* possible outcome caused by L is better than the *worst* outcome with R . We will describe such a player as *conservative*. However, an ambitious player may choose R if the *best* outcome under R is better than the *best* outcome with L . We will describe such a player as *aggressive*. If we had cardinal utilities

and probabilities we would have used the expressions *risk averse* and *risk loving*.

It is clear then that in the same situation, an aggressive player with the same options and the same preferences as a conservative one may still make a different choice. Some people never buy lottery tickets on the ground that the worst outcome under buying, namely losing one's money, is worse than the certain outcome (no gain, no loss) under not buying. But those who do buy such tickets are clearly judging by the best outcome.

In this chapter we will assume that utilities are ordinal. In other words, between any two choices a and b , the agent may be neutral, prefer a or prefer b . Numbers $u(a), u(b)$ can be assigned to a and b so that $u(a) < u(b)$ iff b is preferred to a . However, ordinal utilities are preserved by all order preserving transformations. If c is preferred to b and b to a (which we may write $c > b > a$) then there is no difference between utility assignments to a, b, c of 1, 2, 10 or 1, 9, 10.

If we had access to cardinal utilities and a subjective probability were available, we could have used expected value as a basis for comparison of choices. However, these tools are not always available. Suppose for instance, that a voter has clear preferences among three candidates A, B and C. She prefers A to B to C. She still might not have a clear intuition in response to the question, "Do you prefer B or a 50-50 chance of A versus C?" Such choices between bets have been used by both Ramsey and de Finetti [61, 28], but

they do not always make sense to the person being asked. And without clear and consistent answers to such questions, we cannot have access to subjective probabilities (or to cardinal utilities). Savage in [68] imposes rationality conditions on choices between bets in order to derive subjective probabilities and utilities.

The general issue is that a player in uncertainty is choosing between two sets (or sequences) of payoffs. The payoffs with L are, say, a_1, a_2, \dots, a_k and the payoffs with R are b_1, b_2, \dots, b_m such that $a_1 > a_2 > \dots > a_k$ and $b_1 > b_2 > \dots > b_m$. A conservative player chooses L over R if a_k is preferred to b_m . An aggressive player chooses R over L if b_1 is better than a_1 . In addition to conservative and aggressive players, we can also consider *moderate* players who try to find the middle way, staying away from the maximum or the minimum.

More generally, let a player use a *choice function* f to represent a sequence of outcomes by a single element. f takes a finite ordered set (*list* from now on)⁶ and chooses a representative element from that list. A conservative player uses the minimum, an aggressive player uses the maximum, and a moderate player uses (say) the median. (In case the number of elements is even we can use the higher of the two medians.)

If the choice function f is used then we let $X \preceq Y$ iff $f(X) \leq f(Y)$, where

⁶[66] also discuss choosing from lists, but their notion of list is different from ours. We are assuming that before choosing a list has already been arranged as (a_1, \dots, a_n) so that $a_i < a_{i+1}$ for all $i \leq n$. They allow for the possibility that the same objects are offered in different orders and different elements may be chosen from the same set.

X, Y are finite lists of outcomes. It is easily seen that the relation \preceq defined this way will be transitive, although it is not anti-symmetric.

The function f should satisfy some rationality conditions.

4.1 Suitable Choice Functions

Definition 4.1 *A choice function f is suitable if it has the properties below.*

1. *If lists A and B are isomorphic by an order preserving map g , then $g(f(A)) = f(B)$.*
2. *If list A is augmented to a list B by adding an element x which exceeds all elements of A , then $f(B) \geq f(A)$.*
3. *If list A is augmented to a list B by adding an element x which is less than all elements of A , then $f(B) \leq f(A)$.*
4. *If lists A and B overlap but all elements in $B - A$ exceed those in $A \cap B$ which exceeds all elements in $A - B$, then $f(A) \leq f(B)$.*

Note that conditions (2) and (3) together imply (for finite lists) condition (4) which implies both (2) and (3).

Let us define $s(n) = f(\{1, \dots, n\})$.

Because of the order isomorphism property (1) above, the function s completely characterizes f , for any finite list is order isomorphic to some list

$\{1, \dots, n\}$. By abuse of language we will say that s is suitable iff the corresponding f is. Note that s takes a single natural number as argument whereas f takes an ordered list.

Theorem 4.1 s is suitable iff

$$s(n) \leq s(n+1) \leq s(n) + 1 \text{ for all } n. \quad (*)$$

Proof:

Necessity: Note that $f(\{1, \dots, n\}) \leq f(\{1, \dots, n+1\})$ since the second list is obtained from the first by adding a larger element. This yields $s(n) \leq s(n+1)$.

Also, $f(\{1, \dots, n+1\}) \leq f(\{2, \dots, n+1\})$ since the first list is obtained from the second by adding a smaller element. But the second value is just $s(n) + 1$. This follows from the isomorphism condition as the list $\{2, \dots, n+1\}$ is isomorphic to the list $\{1, \dots, n\}$ via the function $g(n) = n + 1$. So we get $s(n+1) \leq s(n) + 1$.

Before proving sufficiency we remark that $(*)$ yields $s(n+m) \leq s(n) + m$. This is easily shown using induction on m .

Sufficiency: Suppose that s satisfies $s(n) \leq s(n+m) \leq s(n) + m$ for all m, n . It is easy to see that the first three conditions above will hold for the corresponding f .

To see that the fourth condition also holds, suppose that list A is $\{1, \dots, n\}$ and list B is $\{m+1, \dots, n, \dots, m+r\}$ so that the elements $1, \dots, m$ are in A

and below B , and elements $n + 1, \dots, m + r$ are elements of B above A . Since we assume overlap as in property (4), we may assume that $m + r$ is at least n .

Now $f(B) = s(r) + m$ since B consists of r elements *in order* but shifted rightward from $1, \dots, r$ by an amount of m . Since $n \leq r + m$, we have $s(n) \leq s(r + m) \leq s(r) + m$. Now $f(A) = s(n)$, and $f(B) = s(r) + m$. So indeed $f(A) \leq f(B)$. \square

This shows that there are uncountably many suitable choice functions since $g(n) = s(n + 1) - s(n)$ can be zero or one, infinitely many times. Since there are not that many human beings, the conservative humans with $s(n) = 1$, aggressive humans with $s(n) = n$ and moderate humans with $s(n)$ approximately equal to $n/2$ are the typical cases to appear in practice.

Theorem 4.2 *The minimum, the median and the maximum are all suitable choice functions (SCFs) in the sense above (and the corresponding notions of f -rationality are equivalent to being conservative, moderate, and aggressive respectively).*

Proof: It is obvious that the median, the maximum and the minimum are preserved by isomorphism. We check the fourth condition in Definition 4.1 just for the median.

Suppose that X and Y overlap so that X is $a_1 > a_2 > \dots > a_k > b_1 > \dots > b_m$ and Y is $b_1 > b_2 > \dots > b_m > c_1 > \dots > c_p$. $X - Y$ is above $Y - X$. Clearly

if the median of X is an a_i or the median of Y is a c_i then we are done. If both medians are b_i and b_j respectively. Since b_i is the median of X , we have $i + k = m - i + 1$. Similarly, because b_j is the median of Y , we have and $j = p + m - j + 1$. Thus we get $2i = m + 1 - k$ and $2j = p + m + 1$. Thus $i < j$ and $b_i > b_j$. \square

Note that we could also conclude this from Theorem 4.1 since the median $med(n)$ on $\{1, \dots, n\}$ obviously satisfies the condition $med(n) \leq med(n+1) \leq med(n) + 1$ for all n .

Definition 4.2 *Given an SCF f , an f -rational player is a player who, when uncertain between lists X and Y of alternatives, always picks X if $f(X) > f(Y)$.*

Sometimes we can speak of one strategy being dominated by another. Suppose that there is a set of possible worlds $\{\omega_1, \dots, \omega_n\}$ and for each world ω_i actions L and R yield payoffs p_i and q_i respectively. Then we say that L is dominated by R if for each i , $p_i \leq q_i$ and for at least one i , $p_i < q_i$.⁷ In such

⁷The notion of *strictly dominated* strategy that we intend is *subjective*. Thus we actually mean a strategy which is dominated by another pure strategy that the agent *considers possible* at the time of play. Consider the following scenario: Suppose Ann is giving a dinner party and she has to make a decision between serving two dinners d and d' . Assume that three guests have been invited, and guest 1 and guest 2 prefer d to d' whereas guest 3 not only prefers d' to d but also, he is allergic to one of the ingredients in d . If Ann is a conservative agent, she will serve d' to avoid the worst case scenario. Now let's say guest 3 is not going to the dinner party but Ann is unaware of this. In this case she will still serve d' even though it is a strictly dominated strategy if only guests 1 and 2 will be there. But since guest 3 was expected to show up in the original scenario and Ann does not know that he is not coming, we do not consider serving d' as strictly dominated in this example. Note that the notion of a dominated strategy makes sense only when two choices lead to

a situation, the number of possible outcomes with either L or R will be the same, namely n .

It is easily seen that all three kinds of players, conservative, moderate and aggressive will never pick a strictly dominated strategy.

4.2 The IIA Condition

Definition 4.3 *An SCF f satisfies the IIA condition whenever $a = f(X)$, $Y \subseteq X$, and $a \in Y$ then $a = f(Y)$.*

This definition occurs in [48] but the version we give is not the only one in the literature. (See [62]). Of course there is no particular reason why IIA *should* be obeyed in such a case. The role of $f(X)$ is to play the role of an element which in some sense *represents* X rather than that of a *most preferred* element of X . Thus the median is probably the closest to the expected value which we tend to use when we do have cardinal utilities and a subjective probability.

Informally the IIA condition says that if an element x is chosen from a list S and $x \in T \subseteq S$ then x must be chosen from T . The idea here is that if x is the best element of S and is in a subset T then it is also the best element of T .

different results but over the same set of possible worlds. When one is choosing between L and R in a game tree, the two actions lead to different sets of nodes and the notion of domination does not have an obvious intuitive meaning.

The choice functions we have spoken about do not choose the best elements of a list but rather *typical* elements. Thus the intuitive justification for IIA does not apply. Still it would be worth finding out which choice functions satisfy the IIA condition.

Theorem 4.3 *The minimum and the maximum are the only suitable choice functions which satisfy IIA.*

Proof: That the minimum and the maximum satisfy IIA is clear: Suppose that x is the minimum element of S , and $x \in T \subseteq S$ then x is also the minimum element of T . Similarly, if x is the maximum element of S , and $x \in T \subseteq S$ then x is also the maximum element of T .

Suppose now that f is a suitable choice function other than the minimum and the maximum. Then the corresponding function s is also not the minimum or the maximum and there must be an n such that $1 < s(n) < n$.

We will now show that $f(\{2, \dots, n+1\})$ must be $s(n)$ as well as $s(n)+1$, clearly impossible. Note that $s(n+1) \leq s(n) + 1 \leq n$ and so $s(n+1) \in \{1, \dots, n\}$. By IIA, $s(n+1)$ must be $s(n)$. Also $s(n+1) \geq 2$, so $s(n+1) \in \{2, \dots, n+1\}$, therefore $f(\{2, \dots, n+1\}) = s(n+1) = s(n)$. But $\{2, \dots, n+1\}$ is isomorphic to $\{1, \dots, n\}$ by the successor function, so $f(\{2, \dots, n+1\}) = s(n) + 1$. \square

4.3 Related Work

The issue we have discussed so far is: if a player has preferences among *elements* of a certain list, how does that preference translate to preferences among *sublists*? In the previous discussion we used a choice function f to convert an ordering between elements to an ordering between lists. But the problem has been looked at more abstractly without relying on a choice function.

Suppose for instance that a voter has preferences among candidates for an election, how will that translate to preferences among slates of candidates? Two conditions, dominance and independence, which relate preferences among individual elements to preferences among subsets are discussed in [42] and [11]. See also [15] and [29].

Let R be a quasi-linear order (i.e., reflexive, transitive, and total) on a set X and let P be the strict order part of R , i.e. xPy iff $xRy \wedge \neg yRx$. Let Ξ be the set of all finite subsets of X and let \succeq be an order on Ξ . \succ is the strict part of \succeq . For singleton sets, we assume that $\{x\} \succeq \{y\}$ iff xRy . So singletons follow their (unique) elements.

Dominance: For all $A \in \Xi$ and $x \in X$

$$(i) [xPy \text{ for all } y \in A] \rightarrow A \cup \{x\} \succ A$$

$$(ii) [yPx \text{ for all } y \in A] \rightarrow A \succ A \cup \{x\}$$

Independence: For all $A, B \in \Xi$, for all $x \in X - (A \cup B)$

$$A \succ B \rightarrow A \cup \{x\} \succeq B \cup \{x\}$$

(Note: we are using the symbol “ \succ ” also for set subtraction.)

The conditions Dominance and Independence are called (G) and (M) respectively in [42] by Kannai and Peleg who showed that if X has at least six elements then these very natural conditions are incompatible.

If we use the technique of comparing the minimum, the maximum or the median as ways of comparing sets, then all three techniques satisfy a weaker form of the dominance condition.

Weak dominance: For all $A \in \Xi$ and $x \in X$

$$(i) [xPy \text{ for all } y \in A] \rightarrow A \cup \{x\} \succeq A$$

$$(ii) [yPx \text{ for all } y \in A] \rightarrow A \succeq A \cup \{x\}$$

For instance, suppose we say that $A \succeq B$ iff $\min(A) \geq \min(B)$ then this definition of \succeq satisfies both independence and weak dominance. \max also obeys independence and weak dominance.

However, independence fails for all SCFs except the \min and the \max . Suppose that there is an n such that $1 < s(n) < n$. It follows that there is an n such that $s(n) < s(n+1) = s(n) + 1 < n + 1$. (Take the smallest n such that $s(n) < s(n+1)$). Let $m = s(n)$ and $m + 1 = s(n+1)$.

We will make the proof clearer by an example. Say $n = 5$, $m = 2$ (So the lists in our example will have 5 elements each, and f picks the second element from each list). We tacitly assume that larger numbers are preferred

to smaller numbers. Let $X = \{1, 5, 6, 8, 9\}$ and $Y = \{1, 4, 7, 8, 9\}$. Then clearly f will pick 5 from X and 4 from Y so X will be preferred to Y .

Add 10 to both lists. $X' = \{1, 5, 6, 8, 9, 10\}$ and $Y' = \{1, 4, 7, 8, 9, 10\}$. Now f will pick 6 from X' and 7 from Y' so Y' will be preferred to X' . Adding the single element 10 to both X and Y reversed the preference between the two lists.

4.4 Examples

We now give some applications of the technical work so far.

Example 4.1 The Asian Disease – Tversky and Kahneman 1981 [73]

Imagine that the United States is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the programs are as follows:

- If Program A is adopted, 200 people will be saved.
- If Program B is adopted, there is a one-third probability that 600 people will be saved and a two-thirds probability that no people will be saved.

Which of the two programs would you favor?

In this version of the problem, a substantial majority (72%) of 152 respondents favor program A, indicating risk aversion.

Other 155 respondents, selected at random, receive a question in which the same cover story is followed by a different description of the options:

- If Program A' is adopted, 400 people will die.
- If Program B' is adopted, there is a one-third probability that nobody will die and a two-thirds probability that 600 people will die.

Again, the same question is asked:

Which of the two programs would you favor?

A clear majority (78%) now favor B'.

One way to understand this phenomenon is that the way the problem is stated causes a change from *min*-rationality to *max*-rationality. Programs A and A' result in 400 deaths. Programs B and B' amount to no deaths if we are lucky and to 600 deaths if we are unlucky. A cautious person would prefer program A and an aggressive (optimistic) person would prefer B. The way the question is posed causes a shift from caution to optimism. But both forms of rationality are 'rational'.

Example 4.2 Rationality can enter at two different levels. A purely decision theoretic level, where a single agent is trying to make the 'best' choice; or a game theoretic level where two or more agents are involved, and we not only

have to think about the ‘rational’ choice but also about what another agent thinks is the rational choice.

The previous example was in a purely decision theoretic mode. We now consider how the notion of ‘temperament-based’ rationality can be helpful in understanding a classic game theoretic example. We consider a player (the husband) in a state of uncertainty and show how the play is affected by his temperament.

Figure 7 presents the Bach-Stravinsky game as an extensive form game with the wife choosing first and the husband next, but we leave it open whether the husband knows the wife’s choice. We have made the payoffs for Bach-Stravinsky and Stravinsky-Bach different so that the game is generic.

We consider various scenarios involving the husband’s knowledge and temperament. We assume that the wife knows the husband’s payoffs and temperament and he does not know hers.

Case 1) The husband does not know the wife’s move and she knows this.

a) He is aggressive. Then being aggressive, he will choose S (Stravinsky) for his move since the highest possible payoff (for him) is 3. Anticipating his move, she will also choose S , and they will end up with payoffs of (2,3).

b) The husband is conservative. Then not knowing what his wife chose, he will choose B since the minimum payoff of 1 is better than the minimum payoff of 0. Anticipating this, the wife will also choose B and they will end up with (3,2).

Case 2) Finally if the husband *will* know what node he is at (and the wife knows this), then the wife, regardless of the husband's temperament, will choose B . The husband will also choose B and they will end up at $(3,2)$.

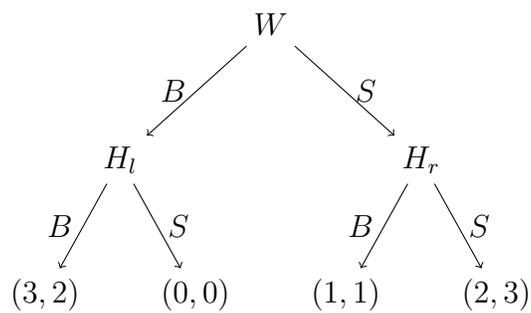


Figure 7: Bach-Stravinsky game in extensive-form with generic payoffs

5 Knowledge Manipulation

Some of the ideas in this chapter were expressed in a preliminary form in a conference presentation [56] but the precise semantics was left out. A more detailed version appeared in [58].

5.1 Two Examples

Providing knowledge is the removal of uncertainty and not providing it means that some uncertainty will remain. Thus the issue of knowledge is tied intimately with that of action under uncertainty. Sometimes we have to choose between two or more actions, but lacking information, we are not sure what outcome will come about as a result of our choice. Chapter 4 has concentrated on choices under uncertainty and the properties of various choice functions. We now turn to the effect of such choice functions and the states of knowledge on agents' behavior.

We give two examples of how influencing someone's knowledge or beliefs influences their actions.

5.1.1 Three Envelopes

The following example is from Gerbrandy [33].

Ann has to choose one of three closed envelopes. One envelope is empty, one envelope contains three yuan for Ann and three yuan for Bill, the third contains 6 yuan for Bill (Ann gets nothing). These facts are commonly known. In addition, Bill knows which envelope contains the money. Ann has no idea.

In this game, Ann can do no better than to choose an envelope at random. She may then expect a payoff of 1 yuan, and Bill may expect 3 (the average of the money in the three envelopes).

Suppose Bill is allowed to communicate with Ann, but only by saying things that are true (say he communicates via an independent referee). What is smart for Bill to say? He could tell Ann that the second envelope contains 3 yuan for each of them; Ann, believing him, would then choose this envelope. This particular communication act secures, but does not improve, Bill's expected payoff of 3 yuan. Bill can do better by providing Ann with less information, and tell her, for example, that it is either the second or the third envelope that contains money for both of them, without specifying which. Ann can now improve her chances by choosing one of these two envelopes, and Bill can expect a payoff of 4.5.

Suppose now that there is no referee: there is no way for Ann to check the truth of what Bill says (apart from opening an envelope, of course). If Ann is gullible, Bill's best option is to lie,

and tell her that there is 3 yuan for her in the third envelope.

So Ann is suspicious of what Bill says. She does not believe Bill when he says that the third envelope contains 3 yuan for her, and, by a symmetric argument, she will not believe him either if he, this time truthfully, claims that it is the second. But she probably should believe him if he says that it is either the second or the third: Bill could not possibly gain anything by lying about this.

5.1.2 Manipulation by Leaking True Information

The following example is from Artemov [3]. Consider the game tree in Figure 8. Players are A, B and C. The root node A belongs to player A. If A plays left (L), it is then B's turn to make a move. If A plays right (R), it is C's turn next. The first payoff in each leaf belongs to A, the second one to B, and the third one to C.

In this game the initial assumption is that all three players are rational but A is not aware of B's and C's rationality. The knowledge based rationality (KBR) solution by Artemov [2] chooses the move with the highest *known* payoff for each player. Therefore it suggests A's choosing left to secure a payoff of 2 which is the highest known payoff if she plays left as opposed to 1 if she plays right. Actually, A gets 4 which is more than expected since B is rational and will play left. Payoff for B, as well as for C, is 2.

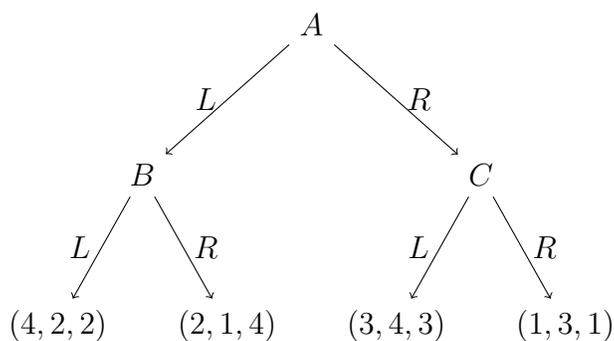


Figure 8: Artemov's game

However, if B is aware of C's rationality, at the start of the game he can manipulate A by (anonymously) leaking the true information that C is rational. A then knows that C will play left, so he plays right and gets a payoff of 3. In that case the highest expected payoff is the same as A's actual payoff. B's and C's payoffs will be 4 and 3 respectively, both higher than what they would get in the initial scenario. Here C does not have an incentive to disclose to A that B is also rational, and B ends up with a higher payoff without 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.

5.2 Knowledge Creation in Game Theoretic Situations

5.2.1 Our Model

In our model we have a number of active players as well as a knowledge manipulator (KM). The knowledge manipulator arranges for the players to have certain restricted amounts of knowledge, both about the situation and about the knowledge of the other players. But she makes no moves herself. When the game ends, all the players including KM receive payoffs.

As we show later, our games *could* be reduced to more familiar forms treating KM as yet another active player. We choose not to do that since the role of the manipulator in real life is different, whether we are speaking about Julian Assange revealing certain secret messages or the government of some country restricting access to the internet. Iago in Shakespeare's play *Othello* is also a knowledge manipulator, although what he supplies to Othello is false beliefs rather than knowledge. It is important that Othello trusts Iago rather than questioning his motives. So in this chapter, we will assume that the active players do not concern themselves with the motives of KM.

5.2.2 Abstract Considerations

Let us consider a game tree for two⁸ players with a set A of nodes, divided into A_1 , the set of nodes where player 1 moves, A_2 where player 2 moves, and T the set of terminal nodes, so that A is the disjoint union of A_1 , A_2 , and T . Moreover payoff functions p_1 and p_2 are defined on T . To simplify matters we will usually assume that both p_1 and p_2 are 1-1. (*i.e., the payoffs at distinct leaves are distinct, i.e., the tree is generic.*)

In that case we know that if we have a *perfect information* game, then backward induction yields a unique way in which the game is played, and according to Aumann, that will indeed be the way the game will be played if there is common knowledge of rationality, see [9, 3].

But of course a perfect information game might be played differently from an imperfect information game with the same structure, same moves, and the same payoffs. As we saw in Example 4.2, this matters, because someone who can manipulate the knowledge of others can also affect the way they play some particular game. If the game has payoffs not only for the active players, but also for KM, then KM will seek to manipulate the active players' knowledge in such a way as to maximize her own payoff.

⁸The number two has no special significance and is only used to simplify notation.

5.2.3 Representing Knowledge, A Third Model: History-Based Semantics

Before we explain how KM can create the knowledge states of the active players, we will briefly discuss another model that is commonly used to represent knowledge: History-based models. We have already covered the set theoretical definition of knowledge and Kripke structures in Section 2.1. While the former uses information partitions and the latter uses accessibility relations to model static knowledge, history-based structures can represent evolving knowledge which we need for our applications.

We will now consider an abstract extensional presentation of a game with communication in which the game is described as a set of *global histories*, each of which represents one possible system evolution given by a sequence of global events. A global history is all that happens. The individual players only see their own local history which is their own perception of the global history. For each system, the set of *players* who participate in the play is assumed to be a fixed finite set $[n] = \{1, 2, \dots, n\}$ (KM is not a member of this set.) Similarly, for each game, the set E of possible events is fixed and consists of signals for the individual players, n -tuples of signals (sent by KM) and moves (by the actual players). This material is adapted from [54].

The set of global histories E^* is the set of all finite sequences over E . H, H' etc. denote elements of E^* . We will let $H \preceq H'$ denote that H is a prefix of

H' . We write $H;H'$ or just HH' to denote the concatenation of the history H with the history H' . When H is of length $\geq k$, we let H_k denote the prefix of H consisting of the first k elements. For a set \mathcal{H} of histories, let $\mathcal{P}(\mathcal{H})$ denote the set $\{H' \mid H' \preceq H \text{ for some } H \in \mathcal{H}\}$ containing all prefixes of sequences in \mathcal{H} .

Definition 5.1 *A system is a tuple $S = (\mathcal{H}, \lceil_1, \dots, \lceil_n)$, where $\mathcal{H} \subseteq E^*$ (our protocol) is the set of all possible histories of S , and for $i \in [n]$, $\lceil_i: \mathcal{P}(E^*) \rightarrow E^*$ is the projection map for i . $\mathcal{H}_i \stackrel{\text{def}}{=} \{\lceil_i(H) \mid H \in \mathcal{P}(\mathcal{H})\}$ is the set of local histories of i .*

Local histories $\lceil_i(H)$ are got by ‘projecting’ global histories H to local components. Each player sees her own moves and when KM sends out an n -tuple of signals, player i sees the i -th component of that n -tuple. Thus if (s_1, \dots, s_n) is the n -tuple signal sent out by KM at some node of the game tree, player i only sees s_i . She may *infer* some s_j , but only if s_i occurs only in conjunction with s_j .

Definition 5.2 *Let H, H' be global histories in \mathcal{H} . For $i \in [n]$, define $H \sim_i H'$ iff $\lceil_i(H) = \lceil_i(H')$.*

\sim_i is an equivalence relation, and it gives the *indistinguishability relation* for i . We can consider this relation as giving the information partition for i in the system S ; that is, given the information available to i , the histories H and H' cannot be distinguished.

The properties of such systems can be studied in a logical language. Let L be a language which has formulae expressing (time dependent) properties of global histories. Then we can write $H_k \models A$, for A belonging to L , to mean that the history H satisfies formula A at stage k . We expand L to a larger language LK by closing under boolean connectives and operators K_i and C . Thus if A is a formula of LK and i is a player, then $K_i(A)$, meaning i knows A , and $C(A)$, meaning A is common knowledge are also in LK .

We can then define $H_k \models K_i(A)$ to hold if for all m and all $H' \in \mathcal{H}$, if $H'_m \sim_i H_k$ then $H'_m \models A$. What player i knows at stage k depends on his local history.⁹ Moreover, the laws of logic $LK5$ (the $S5$ version of the logic of knowledge) are valid.

If $P = \{p_0, p_1, \dots\}$ is a countable set of atomic propositions. then the syntax of the logic is given by:

$$\phi, \psi \in LK ::= p \in P \mid \neg\phi \mid \phi \vee \psi \mid K_i\phi \mid C(\phi)$$

A model is a pair $M = (S, \pi)$, where $\pi : \mathcal{P}(\mathcal{H}) \rightarrow 2^P$ is a valuation map on finite prefixes of global histories which gives the truth values of some atomic predicates at the states. We can now inductively define the notion $H_k \models \phi$, for $H \in \mathcal{H}$, $k \geq 0$ and $\phi \in LK$:

1. $H_k \models p$ iff $p \in \pi(H_k)$, for $p \in P$.

⁹We can define $H_k \models C(A)$ analogously using the reflexive transitive closure of $\bigcup \sim_i$.

2. $H_k \models \neg\phi$ iff $H_k \not\models \phi$.
3. $H_k \models \phi \vee \psi$ iff $H_k \models \phi$ or $H_k \models \psi$.
4. $H_k \models K_i\phi$ iff for all $m \geq 0$, for all $H' \in \mathcal{H}$ such that $H_k \sim_i H'_m$, $H'_m \models \phi$.

Also see [37] for a discussion of topics similar to what is found in this section.

5.3 Creating Knowledge States

Now we return to the topic of our particular application of history based structures and develop a model specifically designed for that application.

How would KM create a knowledge situation? One way that KM can create such a situation is, at each node she sends signals to the active players. The signal function S , whose values are sets of n -tuples of signals is common knowledge. The signal function does not usually determine the actual signals that KM sends but only a set of possible tuples of signals she *can* send. When sending, she picks a particular n -tuple of signals, where n is the number of active players.

Based on the signal he receives, an active player can infer something about the node he is at. Suppose that at some node, KM could send *one* of the three pairs of signals $(x, y), (x, z), (u, v)$ and that this fact is common knowledge. Then, when player 1 receives an x she does not know whether player

2 received a y or a z , whereas if she receives a u she does know that player 2 received a v . On the other hand, player 2 can infer player 1's signal from his own. We shall later see how such facts can be used to create complex states of knowledge.

So we augment a history-based model with signal functions. In this model, only the active players make moves and only KM sends signals. A history is basically a sequence of moves by players and signals sent by KM. Each active player perceives his own local history consisting of his own moves and the particular signals received by him.

Definition 5.3 *A game tree τ with signal function is a standard extensive-form game tree with a set of nodes A , along with a set of signals Σ and a signal function $S : A \rightarrow \mathcal{P}(\Sigma^n)$ where n is the number of players and \mathcal{P} stands for the power set.*

The associated protocol $H(A)$ consists of all sequences

$$(a_1, \sigma_1, a_2, \sigma_2, \dots, a_{k-1}, \sigma_{k-1}, a_k) \in (A \times \Sigma^n)^* \times A$$

such that a_1, \dots, a_k is a path in the game tree starting at the root, for all $i < k$, a_{i+1} is a child of a_i , and $\sigma_i \in S(a_i)$ for each $1 \leq i < k$.

We define a node function $N : H(A) \rightarrow 2^A$ by setting

$$N(a_1, \sigma_1, \dots, a_{k-1}, \sigma_{k-1}, a_k, \sigma_k) := N(a_1, \sigma_1, \dots, a_{k-1}, \sigma_{k-1}, a_k) := a_k$$

For each $\sigma_i = (s_1, \dots, s_n)$, player j observes s_j and moreover the player observes all the moves which were his own.¹⁰

Definition 5.4 We define the set of histories for a game tree with signal function as follows. (For simplicity of notation we assume that the tree is a full binary tree, with two branches at each node and all branches of the same length. Thus all moves will be labeled L and R). A history is an element of

$$(A \times \Sigma^n \times \{L, R\})^* \times A$$

i.e. something which looks like

$$((a_1, \sigma_1, d_1), (a_2, \sigma_2, d_2) \dots, (a_{k-1}, \sigma_{k-1}, d_{k-1}), (a_k, \sigma_k, d_k))$$

which we might write simply as

$$(a_1, \sigma_1, d_1, a_2, \sigma_2, d_2 \dots, a_{k-1}, \sigma_{k-1}, d_{k-1}, a_k, \sigma_k, d_k)$$

We do not insist that a sequence end with some d_k as prefixes of allowed histories also have to be allowed.

Here each d_i is a member of $\{L, R\}$, and σ_i is an element of $S(a_i)$. Two sequences α_1, α_2 are equivalent¹¹ for agent j , if j receives the same signal

¹⁰Note that the node function N does not have to do with payoffs and a_k need not be a terminal node.

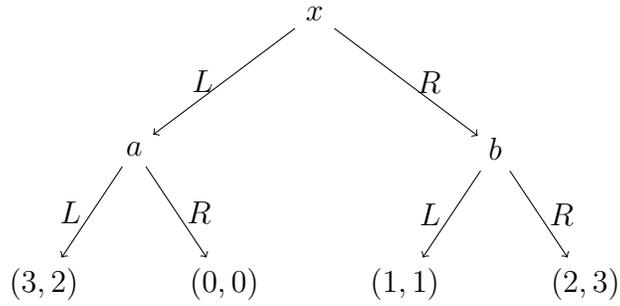
¹¹These sequences α etc are playing the same role as the H_k of the previous subsection.

from KM at each stage p (i.e. $(\sigma_p)_j$ is the same in both cases) and moreover at each stage p in α_1 and α_2 if it was j 's move then j made the same move from $\{L, R\}$.

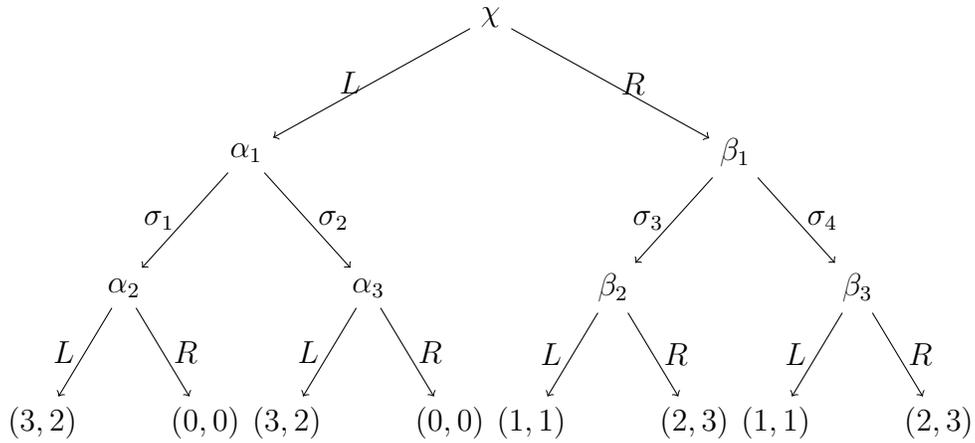
The set of histories themselves as just defined can be seen as nodes of another (much larger) tree than the game tree. Let $H(\tau)$ be this second tree. We will call this $H(\tau)$ *the history-based tree corresponding to τ* . A physical node a of τ can correspond to two or more nodes $\alpha_1, \alpha_2, \dots, \alpha_k$ in $H(\tau)$ where $N(\alpha_1) = N(\alpha_2) = \dots = N(\alpha_k)$.

Figure 9.a shows an example of a game tree, and Figure 9.b shows a history-based tree corresponding to 9.a. Here KM might send any of the two different signal pairs σ_1, σ_2 at α_1 and the two different signal pairs σ_3, σ_4 at β_1 . We should note that under the node function N , a corresponds to all of $\alpha_1, \alpha_2, \alpha_3$. Similarly for b and the three β 's. (i.e. $N(\alpha_1) = N(\alpha_2) = N(\alpha_3) = a$, and $N(\beta_1) = N(\beta_2) = N(\beta_3) = b$)

Pradeep Dubey [22] has remarked that if we see KM as an additional active player and interpreting her signals as moves, a knowledge based game can be understood as a conventional game of partial information with information sets. However, we should note that in our framework, information sets over nodes are defined not only for the player who is making the move but for all active players.



(9.a) Game tree τ_1



(9.b) A history-based tree $H(\tau_1)$ corresponding to τ_1

Figure 9: A game tree and a history-based tree corresponding to it

5.4 States of Knowledge

We described a model for representing a game with possibly complex signal functions.

Let Δ be the set of nodes of a history-based game tree $H(\tau)$. By abuse of language, we refer to A , the set of nodes of τ , as $N(\Delta)$. We stipulate that for

each element $a \in N(\Delta)$, a is also an atomic formula which is true precisely when the play is at the (physical) node a , i.e. it is true at node α of $H(\tau)$ iff $N(\alpha) = a$. Since there are several nodes α with the same $N(\alpha) = a$, a will be true at all of them, but knowledge states, before receiving a signal from KM and after receiving that signal will of course be different. When we talk about KM sending a signal, of course she sends such a signal at the first α such that $N(\alpha) = a$.

We created a formal language LK by closing under truth functions, operators K_1, K_2 and the operator C . (Here K_1 means that 1 knows, K_2 means that 2 knows, and C stands for common knowledge).

Then a perfect information game is simply a game where formulas of the form $a \rightarrow C(a)$ are true at all α such that $N(\alpha) = a$.

Consider an actual play of a game with two players 1 and 2 and where the formula $K_1(a) \wedge K_2(a)$ holds at some realized nodes α such that $N(\alpha) = a$, but for instance $K_1K_2(a)$ does not hold at those nodes. At those α 's, where they *have* received the signal, both players know what node it is but they do not know that the other knows. Now, in order for this to work, clearly the formula $a \rightarrow (K_1(a) \wedge K_2(a))$ cannot be common knowledge. This means that KM *had the option* of sending each player at each node a signal revealing that node, but KM also had other options so that a player receiving such a signal knows which node she is at but does not know what signal another player received and hence whether the other player knows which node he is

at. (We will see an actual example later.)

Then with $(K_1(a) \wedge K_2(a))$, both players know which node they are at. But if 1 had made a choice between L and R , 2 knows which choice 1 made, but 1 did not know that 2 *would* know. Thus 1 might well play differently. So it is not a perfect information game, strictly speaking. Yet we cannot indicate the ‘imperfection’ by indicating information sets, for both players know where they are and the information sets would be singletons. We illustrate this by means of the following example.¹²

Example 5.1 Figure 10 illustrates a Kripke structure M_α at node α . The set of states is $\{\alpha, \alpha', \alpha'', \beta, \gamma\}$. The actual state is α . We have a in states α, α' and α'' . We have $b \neq a$ in β , and $c \neq a$ in γ . Even though state α' will not be arrived at, at the actual state α , 1 considers it possible, and even though 1 knows that β will not be arrived at, she considers that 2 considers it possible, etc. In this case, both players know that they are at node a , but they do not know that the other player knows.

To represent such situations, we modify the knowledge requirement. We stipulate that with each node α is associated a Kripke structure M_α with two knowers 1 and 2. *Such a Kripke structure would represent a complex state of partial knowledge on the part of the players.*

Let M_α^- be the unpointed structure corresponding to M_α . We will assume

¹²Note that the Roman letters a, b etc refer to the nodes of the game tree, Greek letters α, β etc refer to the nodes of the corresponding history-based tree.

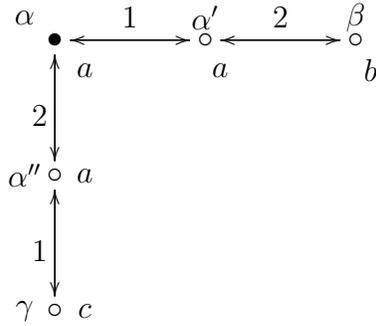


Figure 10: Kripke structure M_α associated with node α

that the map $\alpha \rightsquigarrow M_\alpha^-$ is common knowledge. In [54] terms this means that the *protocol*, or which plays and signals are *possible* is common knowledge. Each player will also know the block from her partition that includes the actual state. Thus if α is the actual state of the pointed Kripke structure M_α , then player i will know M_α^- as well as the set $\{\delta : \alpha \sim_i \delta\}$. In the example above, at state α , 1 will know that the actual state is in $\{\alpha, \alpha'\}$ and 2 will know that the actual state is in $\{\alpha, \alpha''\}$.

Thus the class of knowledge situations we can consider is more general than perfect information games or games whose imperfection can be indicated simply by information sets.

We define an *extended knowledge-based game* (or KB-game) as a history-based game supplemented by such a function M_α . As we noted, a perfect information game is (can be seen as) a special case of such a KB-game. For in that case, for each α , with $N(\alpha) = a$, the structure M_α has a single state satisfying a and no other states are accessible to any player.

Can *every* knowledge-based game arise from a history-based game? Not so, because in a knowledge-based game we do not account for the fact that players retain knowledge of where they were at some previous stage, a necessary consequence of the fact that a player who knows his local history also knows his previous local histories. If a player knew that he *was* at some node $a = N(\alpha)$ then he must know now that he is at some node b which is a descendent of a . Thus if $b = N(\beta)$, there are obvious limitations on what M_β^- could be given what M_α^- was.¹³

However, at least in **one-shot** situations, KM has enormous power to create states of knowledge:

Theorem 5.1 *Any knowledge situation represented by a finite Kripke structure M can be created in a single signaling step.*

Proof: The knowledge manipulator (KM) picks a world w in M and sends player i the signal (M, X_i^w) where $X_i^w = \{v | wR_iv\}$ and R_i is the accessibility relation of player i . This tells us the local history of i . The global history H_w is $(M, w, X_1^w, \dots, X_n^w)$. An atomic formula p holds at H_w iff it holds at w . We now show by induction on the complexity of the formula F that $M, w \models F$ iff $H_w \models F$ where we use the [54] history-based semantics to define $H_w \models F$

We have already noticed that atomic formulas behave correctly (by stipulation) and truth functions are clear. Consider $F = K_i(B)$.

¹³The situation would be different if, like the inhabitants of Gulliver's *Laputa*, the players were forgetful and had to be reminded all the time of past events.

- $M, w \models K_i(B)$ iff
- $(\forall v)(wR_iv \rightarrow M, v \models B)$ iff
- (IH) $(\forall v)(wR_iv \rightarrow H_v \models B)$ iff
- $(\forall H_v)(\lceil_i(H_v) = \lceil_i(H_w) \rightarrow H_v \models B)$
- iff $H_w \models K_i(B)$

We recall that \lceil_i is the projection for i of a global history, and two histories H_w, H_v here have the same projection for i iff $X_i^w = X_i^v$ iff wR_iv □

5.5 Using Knowledge Manipulation

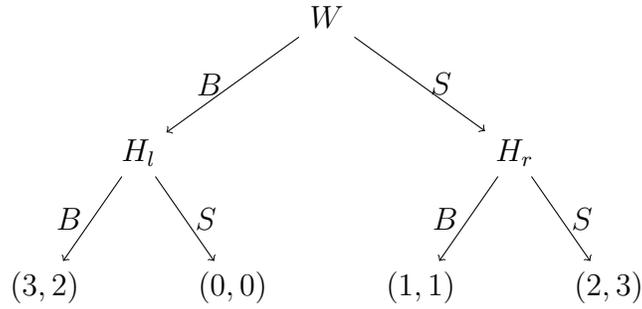


Figure 11: Bach-Stravinsky game in extensive-form with generic payoffs

Example 5.2 We consider now the question of how KM can create the various knowledge scenarios of Example 4.2. KM is capable of creating these situations by means of signals, as well as another one we did not mention

where the husband does not know the wife's move, and the wife does not know that he will not.

The game is shown again in Figure 11. We assume that the wife moves first and the husband after. Thus the payoffs are in the form (wp, hp) where wp refers to the wife's payoff and hp to the husband's.

- Case 1. The husband does not know the wife's move and she knows this.
- Case 2. The husband *will* know what node he is at and the wife knows this.
- Case 3. We will also consider a third case where the husband does not know the wife's move but the wife does not know that he will not.

For case 1, KM can associate the signal tuple (l, a) with the wife's left move (i.e. KM sends the wife an l , and the husband an a if the wife moves left), and the signal tuple (r, a) with the right (i.e. KM sends the wife an r , and the husband an a if the wife moves right). In this case after the wife moves and KM sends the signals, the wife knows (if she did not already) which node they are at, but the husband will not, all he gets is the uninformative a in both cases.

For case 2, KM can associate the tuple (l, l) with the wife's left move, and (r, r) with the right. After the signals are sent, both players will know which node they are at, and this fact will be common knowledge. (Note that it is

also possible to have the signal function such that the husband knows what node he is at, the wife knows that he knows, but the husband doesn't know that the wife knows he knows.)

Finally, for case 3, if KM wants the wife to be in doubt whether the husband knows, she could associate two signals $\{(l, l), (l, a)\}$ with the wife's left move, and two other signals $\{(r, r), (r, a)\}$ with her right move. Then if the wife chose left and receives an l , she will not know if the husband got an l or the neutral a . If KM sends (l, l) then the husband will know what node he is at, but will also know that his wife did not know whether he *would* know.

We have not indicated KM's utilities above. They could appear as a third component of the payoff function. When the game finishes, all three players including KM receive their payoffs and so KM has an interest in seeing to it that the game is played in a certain way. She can do this, to a limited extent, by influencing the structures M_a .

We return to the case where KM wants the wife to be in doubt whether the husband knows. Recall that she could associate two signals $\{(l, l), (l, a)\}$ with the wife's left move and two other signals $\{(r, r), (r, a)\}$ with her right move.

Thus KM could have two moves for each of the wife's moves. After her move B , KM could have a move corresponding to the signal pair (l, l) and another move corresponding to the signal pair (l, a) . Similarly after her S move, KM could have a move corresponding to the signal pair (r, r) and another move

corresponding to the signal pair (r, a) . This gives us four nodes corresponding to the moves by the wife and KM, and let us denote them in the natural way as LL, LR, RL and RR .

The nodes LL, LR are indistinguishable for the wife and similarly RL and RR . She knows what she moved, but does not know what the husband got. The husband cannot distinguish between LR and RR , because in the signal description he got an a in either case. But the other two, LL and RL are singletons for him. If he gets an l or an r he knows how the wife moved. This construction is illustrated in Figure 12.

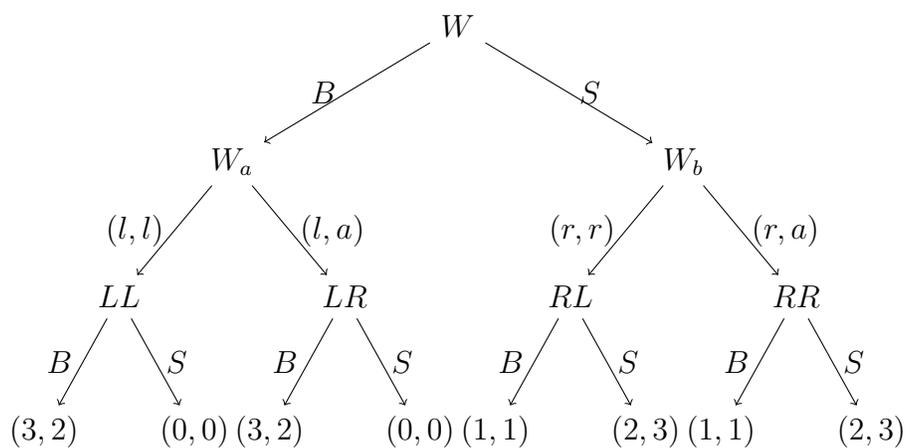


Figure 12: Bach-Stravinsky game with KM's signals as moves

5.5.1 An Example of Manipulation

Let us consider the game in Figure 13, first without KM's moves included. Player 1 has three moves: A , B and C , and player 2 has two moves: L and R . The payoffs at the leaf nodes are for player 1 and player 2 respectively.

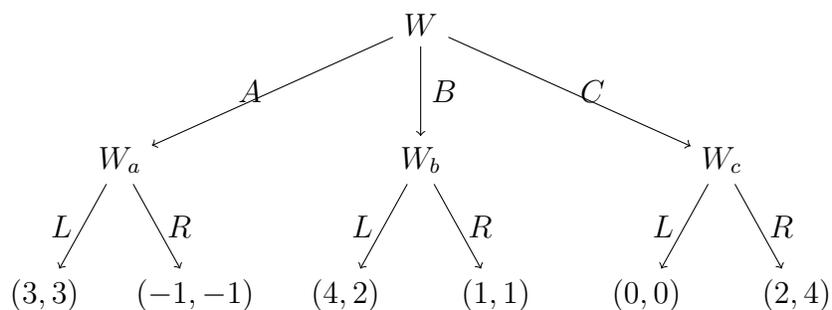


Figure 13: Another game tree

KM may choose to make this game a perfect information game by using the signal tuples shown in Figure 14. That is, KM sends the signal tuple (a, a) if player 1 moves A , the tuple (b, b) if she plays B , and the tuple (c, c) if she plays C . In this case, the backward induction solution will apply: Knowing that player 2 will know what node he will be at, player 1 will play B ; KM will send the signal pair (b, b) making player 1's move common knowledge, and player 2 will play L . Payoffs for the active players will be 4 and 2 respectively.

KM may also choose to use the signal structure in Figure 15. In this version of the game, it is common knowledge that if player 1 plays A , this fact will be common knowledge. (Both players associate the signal a only with the move

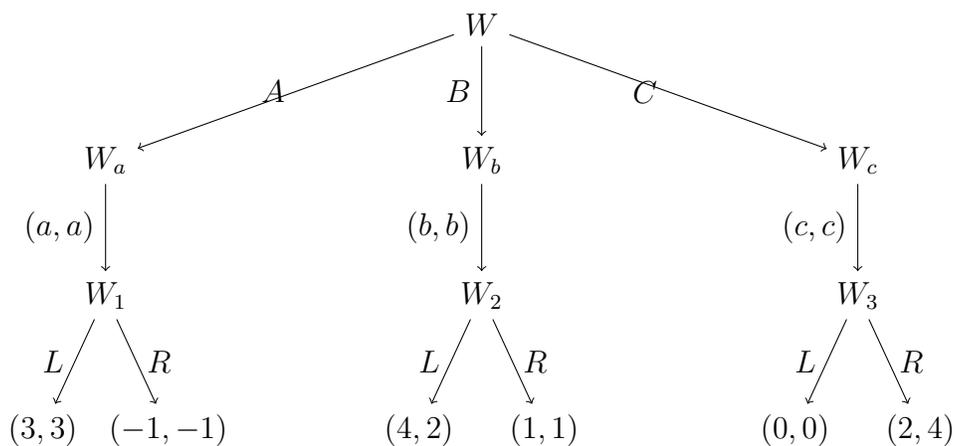


Figure 14: KM turns the game into a perfect information game

A). If player 1 chooses to play B or C , KM will leave player 2 in uncertainty by sending him an x .

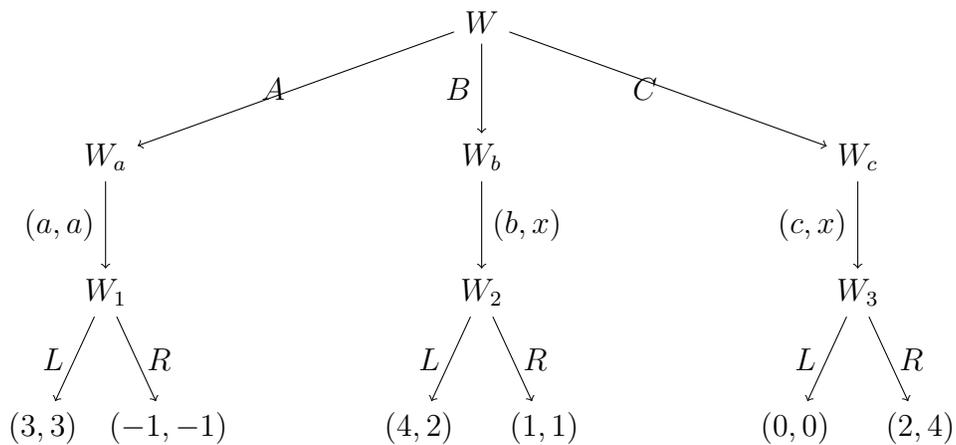


Figure 15: KM gives away some information but leaves the game as an imperfect information game

So how might the active players play this game? We will assume that both

players are conservative and that this is common knowledge.

We can start with player 1's reasoning:

- If she (player 1) plays A , KM will send (a, a) , and player 2, knowing what node he is at, will play L . The payoffs will be $(3, 3)$.
- If she plays B , KM will send (b, x) , player 2 will not know if he is at W_2 or W_3 . Playing L brings him a minimum of 0 whereas playing R brings him a minimum of 1. He will play R , and the payoffs will be $(1, 1)$.
- If she plays C , KM will send (c, x) , player 2 will again not know where he is. With the same reasoning he will play R , and the payoffs will be $(2, 4)$.

So if player 1 is a conservative player, she will choose to play A , which secures her a payoff of 3, higher than the secured payoffs with B and C .

In the original game, B dominates A for player 1 in the usual game theoretic sense, however in this case it makes sense for player 1 not to eliminate the move A without seeing KM's signal structure first.

Given that both players are conservative, and that this is commonly known, both players can reduce the game to the form in Figure 16 after learning KM's signal structure. In this reduced form, the move A is no longer dominated for player 1.

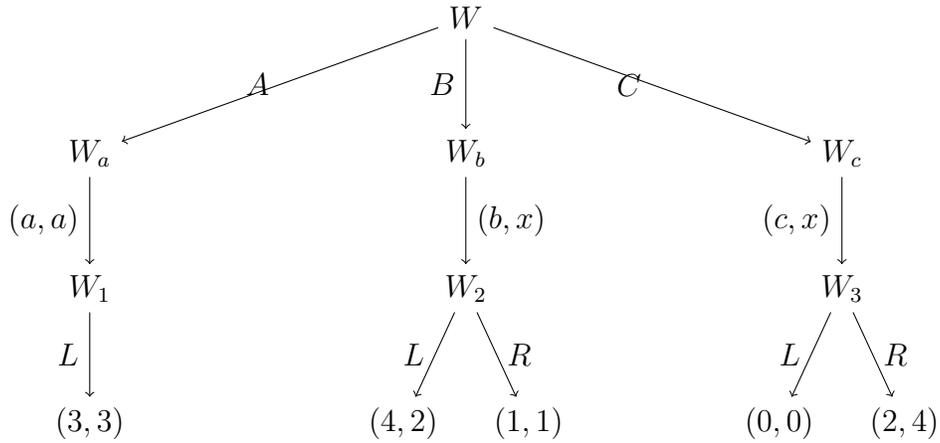


Figure 16: The reduced game

5.5.2 Can KM Always Create Uncertainty?

Consider again the game in Figure 12. In this extended game, if player 1 simulates player 2's reasoning, she gets the following:

- If she (player 1) plays B and KM sends (l, l) , player 2 will play B . The payoffs will be $(3, 2)$.
- If she plays B and KM sends (l, a) , player 2 will play B in order to secure himself a payoff of 1. However, the actual payoffs will be $(3, 2)$.
- If she plays S and KM sends (r, r) , player 2 will play S . The payoffs will be $(2, 3)$.
- If she plays S and KM sends (r, a) , player 2 will play B in order to secure himself a payoff of 1. The payoffs will be $(1, 1)$.

We should note that the signaling function is intended to see to it that player 2 will not know which node he is at if he gets the signal a . But there is an issue: Player 2 himself is also able to perform all this reasoning. So if he *does* get the signal a , will he not be able to deduce that he must be at node LR after all?

It might not be so. Player 2 can start his reasoning by checking player 1's payoffs first. Since with the given signal structure player 1 cannot know if player 2 will know which node he is at, she (player 1) might choose to play S which secures her a payoff of 1 as opposed to the secured payoff of 0 with the move B .

So when player 2 does get the signal a from KM, he indeed cannot know which node he is at.

However, KM is not always able to create uncertainty. One such case is when both players have a move that dominates their all other moves. Consider the Prisoner's Dilemma game in extensive-form in Figure 17:

In this game, if both players are conservative (or aggressive) they will play D . Then even if KM tries to create uncertainty for player 2 by sending him an x in both W_c and W_d , player 2 will know that player 1 must have played D and therefore he must be at node W_2 .

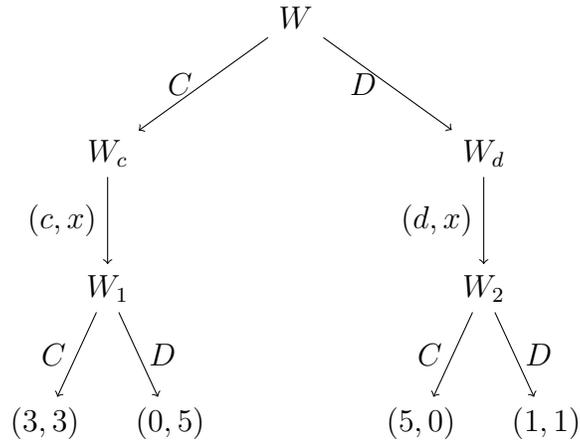


Figure 17: Prisoner's Dilemma in extensive-form with signals

5.5.3 Predicting the Play

Can KM *always* predict how a game will be played in a less than perfect information state which she has brought about? This is indeed true in a decision theoretic situation if the temperament of the player is known to KM.¹⁴ For instance a conservative agent faced with uncertainty will choose the least risky alternative. And since we assume that no two outcomes have the same value, the least risky alternative will always be well defined and *known to KM*.

In two-person games with imperfect information, there may not be a unique way that the players will play, therefore KM may not be able to predict *how* they will indeed play. In particular the reasoning process of the players can

¹⁴By decision theoretic we mean that there is only one agent *apart* from KM, who has a decision theoretic problem to solve.

be order-dependent, as we have seen in Section 5.5.2.

Theorem 5.2 *If player 2 does not know player 1's payoffs but player 1 does know player 2's payoffs, then (given their temperaments) there is a unique solution to the game.*

Proof: Although our examples do take this form, we do not assume that player 2 always plays after 1. For instance player 2 could start the game¹⁵ or the game may be over more than two stages in which case we would have both agents playing after each other.

At any particular node, an f -rational player 2 has a set of nodes X which he *might* be at. He considers all possible strategies s of player 1 which are compatible with their presently being at a node in X . For each such s he considers various strategies s' which he himself could play and the payoff $p(s, s')$ to himself of s, s' . Then he chooses that s' for which $f(\{p(s, s') | s \in X\})$ is highest. This defines the strategy s' of 2 as a function of the node. Player 1 can simulate player 2's reasoning and plays so as to maximize her own payoff. This yields a unique outcome.

Note that since player 2 does not know player 1's payoffs, he is not able now to think of a proper response to player 1's choice - he has no idea what it is. So there is no 'cycle of reasoning'. □

More generally, with two or more players, if the players are linearly ordered

¹⁵It is worth noting that unlike chess, in the Japanese game of *Go* the black player starts.

so that no player knows the payoffs of any player above him then there is a unique solution.

5.5.4 A Note on Normal Form Games: Correlated Equilibrium

Unlike the rest of the chapter, we do assume cardinal utilities here. Given that assumption, it makes sense to talk of mixed strategies. In a correlated equilibrium, players play co-ordinated strategies [8, 6]. Consider for instance the Bach-Stravinsky game in Figure 18, where the wife is the row player and the husband is the column player.

	Bach	Stravinsky
Bach	2, 1	0, 0
Stravinsky	0, 0	1, 2

Figure 18: Bach-Stravinsky game in normal form

The players can play mixed strategies [47]. $2/3$ Bach and $1/3$ Stravinsky for the wife, and $1/3$ Bach (B) and $2/3$ Stravinsky (S) for the husband. This yields an expected payoff of $2/3$ for each.

However, if they can co-ordinate, with (B, B) half of the time, and (S, S) half of the time, then they will get an average payoff of $3/2$.

How is this correlation to be achieved? Suppose there is a co-ordinator (knowledge manipulator) who tosses a coin and when it is heads, she in-

forms both that it is B and when it is tails she informs both that it is S. This works, *if the manipulator is honest*.

We can consider yet another variation. Let us say that the manipulator is not 100% honest and that this fact is known to the active players. Moreover, suppose that the manipulator has her own payoff of 1 at cell (B, S), and 0 in every other cell as shown in Figure 19, and that all of this is also known to the active players.

	Bach	Stravinsky
Bach	2, 1, 0	0, 0, 1
Stravinsky	0, 0, 0	1, 2, 0

Figure 19: Bach-Stravinsky game with knowledge manipulator

Then she has the incentive to deceive and occasionally send a B to the wife and an S to the husband.

Is this deception risk free? Only if she does not deceive too much. If she signals (B, B) with probability $2/9 + \epsilon$, (S, S) with probability $2/9 + \epsilon$ and (B, S) with probability $5/9 - 2\epsilon$ then she has an expected value of $5/9 - 2\epsilon$ and the wife and husband are *still better off* than if they had played the mixed strategy Nash equilibrium. Their expected payoff will be strictly more than $2/3$. In that case, even though the wife and the husband know the manipulator is not 100% honest, they would still trust her signals.

5.6 Related Work

A very interesting approach loosely related to ours is described by Kamien, Tauman and Zamir [41] who refer to our knowledge manipulator as the *maven*. They describe the following puzzle.

A card is randomly picked of one of two colors red and black. Player A is asked to guess the color. After A announces his choice, B is asked to make a guess. The utilities are 5 for the one who guesses correctly if the other does not, and 0 for the one who didn't guess correctly; 2 for each if both guess correctly; and 0 for both if neither guesses correctly. It is easy to see that in the absence of information A should guess randomly and B should guess a different color. The expected value is 2.5 for each.

Suppose however that an informed person KM, the maven, shows the card to A and B knows that A has seen the card. Then A 's best move is to announce the correct color and B 's best move is to announce the same. The expected utility for both (and hence for A) has dropped from 2.5 to 2.

It appears that the knowledge has harmed A , contrary to the intuition that the more one knows, the better off one is. But note that A is actually harmed by the fact that B knows that A has learned. In a subsequent paper, Neyman [49] shows that if A 's knowledge is increased while leaving everything else the same, then A is indeed better off.

Other work like that of Brandenburger et al [17] is also relevant but unlike

us they rely on cardinal utilities. They also do not speak about actual manipulation of behavior by limiting knowledge. Other relevant references are [14, 59, 21].

5.7 Discussion

Our framework allows KM to make public as well as certain types of private announcements. A player can learn new information without the others knowing what information she has learned. However, with the signal structure being common knowledge, KM cannot make private announcements where a player gets information without the other players being aware of the fact that she received *some* information. We would expect that if KM is allowed to manipulate the game with private announcements of the second kind, she can achieve more (in terms of manipulating the game) since the other players would not be able to reason about an announcement they are unaware of.

In the setup we investigated, there is only one knowledge manipulator who, moreover, is trusted by the other players. But we can also consider variants. One possibility is where the manipulator is manipulative. Her payoff function is known to other players, and they are aware that they cannot fully trust her. This is the direction of cheap talk [26].

We can also consider the case where every player is both an actor and an informer. This case could be investigated by enriching the purely informational

structure of [54] and augmenting it with actions.

References

- [1] Itay Arieli and Robert J. Aumann. The Logic of Backward Induction. Discussion Paper Series, The Federmann Center for the Study of Rationality, the Hebrew University, Jerusalem, <http://EconPapers.repec.org/RePEc:huj:dispap:dp652>. (2013).
- [2] Sergei Artemov. Knowledge-Based Rational Decisions. Technical Report TR-2009011, CUNY Ph.D. Program in Computer Science, (2009).
- [3] Sergei Artemov, Rational decisions in non-probabilistic settings, Technical Report TR-2009012, CUNY Ph.D. Program in Computer Science, (2009).
- [4] Sergei Artemov. Robust knowledge and rationality. Technical Report TR-2010010, CUNY Ph.D. Program in Computer Science, (2010).
- [5] Geir B. Asheim. On the epistemic foundation for backward induction. *Mathematical Social Sciences*, 44, 121-144, (2002).
- [6] Robert J. Aumann. Subjectivity and Correlation in Randomized Strategies. *Journal of Mathematical Economics*, 1 (1), 67-96, (1974).
- [7] Robert J. Aumann. Agreeing to disagree. *Annals of Statistics*, 4, 1236-1239, (1976).
- [8] Robert J. Aumann. Correlated Equilibrium as an Expression of Bayesian Rationality. *Econometrica*, 55 (1), 1-18, (1987).
- [9] Robert J. Aumann. Backward induction and common knowledge of rationality. *Games and Economic Behavior*, 8, 6–19, (1995).
- [10] Alexandru Baltag, Sonja Smets, Jonathan Alexander Zvesper. Keep “hoping” for rationality: A solution to the backward induction paradox. *Synthese*, 169, 301-333 (Knowledge, Rationality and Action, 705-737), (2009).
- [11] Salvador Barbera, Walter Bossert and Prasanta K. Pattanaik. Ranking sets of objects, University of Montreal, Cahier 2001-2002, papyrus.bib.umontreal.ca

- [12] Pierpaolo Battigalli and Giacomo Bonanno. Recent results on belief, knowledge and the epistemic foundations of game theory, *Research in Economics* Elsevier vol. 53 (2), 149-225, (1999).
- [13] Pierpaolo Battigalli and Marciano Siniscalchi. Strong belief and forward induction reasoning, *Journal of Economic Theory* 106 (2), 356-391, (2002).
- [14] Johan van Benthem. Games in Dynamic-Epistemic logic, *Bulletin of Economic Research*, 53 (4), 219-248, (2001).
- [15] Walter Bossert, Prasanta K. Pattanaik, Yongsheng Xu. Choice under complete uncertainty: axiomatic characterizations of some decision rules, *Economic Theory* 16 (2), 295-312, (2000).
- [16] Adam Brandenburger, "Epistemic Game Theory: An Overview" (pre-publication version) and "Epistemic Game Theory: Complete Information" (pre-publication version), *The New Palgrave Dictionary of Economics, 2nd edition*, Palgrave Macmillan, (2008).
- [17] Adam Brandenburger, Amanda Friedenberg and Jerome Keisler. Admissibility in Games, *Econometrica*, 76 (2), 307-352, (2008).
- [18] Michael Chwe. *Rational Ritual: Culture, Coordination and Common Knowledge*, Princeton University Press, (2001).
- [19] Robin P. Cubitt and Robert Sugden. Common Knowledge, Salience and Convention: A Reconstruction of David Lewis' Game Theory, *Economics and Philosophy*, 19 (2), 175-210 (2003).
- [20] Walter Dean and Rohit Parikh. The Logic of Campaigning, in M. Banerjee and A. Seth (Eds.): *ICLA 2011*, LNAI 6521, 38-49. Springer, Heidelberg (2011).
- [21] Hans van Ditmarsch. Knowledge games, *Bulletin of Economic Research*, 53 (4), 249-273, (2001).
- [22] Pradeep Dubey, personal communication, March 18, 2011.
- [23] Jan van Eijck, and Rineke Verbrugge. *Discourses on Social Software*, Amsterdam University Press, (Paperback - Apr 1, 2010).

- [24] Daniel Ellsberg. Risk, Ambiguity, and the Savage Axioms, *Quarterly Journal of Economics* 75 (4), 643–669, (1961).
- [25] Ron Fagin, Joseph Halpern, Yoram Moses and Moshe Vardi. *Reasoning about Knowledge*, MIT press (1995).
- [26] Joseph Farrell and Matthew Rabin. Cheap talk, *J. Economic Perspectives*, 10 (3), 113-118, (1996).
- [27] Yossi Feinberg. Subjective reasoning: Dynamic games. *Games and Economic Behavior*, 52, 54-93, (2005).
- [28] Bruno de Finetti. Foresight: its logical laws, its subjective sources, in *Studies in Subjective Probability*, ed. H. Kyburg and H. Smokler, John Wiley and Sons 1964, (originally published in French in 1937).
- [29] Peter Gärdenfors. Manipulation of social choice functions, *J Econ Theory*, 13 (2), 217-228, (1976).
- [30] John Geanakoplos. Game theory without partitions, and applications to speculation and consensus. *Journal of Economic Theory*, (1989).
- [31] John Geanakoplos and Herakles Polemarchakis. We can't disagree forever. *Journal of Economic Theory*, 28, 192-200, (1982).
- [32] John Geanakoplos and James K. Sebenius. Don't Bet on It: Contingent Agreements with Asymmetric Information, *Journal of the American Statistical Association*, 78, 424-426, (1983).
- [33] Jelle Gerbrandy. Communication strategies in games, *Journal of Applied Non-Classical Logics*, 16, 1-15, (2006).
- [34] Sujata Ghosh, Ben Meijering and Rineke Verbrugge. Empirical reasoning in games: Logic meets cognition. *Proceedings Third Logics for Resource Bounded Agents Workshop*, 15-34, (2010).
- [35] Joseph Halpern. Set-theoretic completeness for epistemic and conditional logic. *Annals of Mathematics and Artificial Intelligence* 26, 1-27, (1999).
- [36] Joseph Halpern. Substantive rationality and backward induction. *Games and Economic Behavior*, 37, 425–435, (2001).

- [37] Joseph Halpern. A computer scientist looks at game theory. *Games and Economic Behavior*, 45 (1), 114–132, (2003).
- [38] Joseph Y. Halpern and Moshe Y. Vardi. The complexity of reasoning about knowledge and time, *Proceedings of the eighteenth annual ACM symposium on Theory of computing*, 304-315, (1986).
- [39] Sergiu Hart. Games in Extensive and Strategic Forms, *Handbook of Game Theory with Economic Applications*, Volume 1, Chapter 2, 19-40, (1992).
- [40] Jaakko Hintikka. *Knowledge and Belief*, Cornell University Press, (1962).
- [41] Morton I. Kamien, Yair Tauman and Shmuel Zamir. On the value of information in a strategic conflict. *Games and Economic Behavior*, 2, 129-153, (1990).
- [42] Yakar Kannai and Bezalel Peleg. A note on the extension of an order on a set to the power set, *J. Economic Theory*, 32, 172-175, (1984).
- [43] Aleks Knoks and Eric Pacuit. Deliberating between Backward and Forward Induction Reasoning: First Steps. TARK 2015, Fifteenth conference on Theoretical Aspects of Rationality and Knowledge, June 4-6, 2015, Carnegie Mellon University, Pittsburgh, USA.
- [44] Saul Kripke. Semantical Considerations on Modal Logic, *Acta Philosophica Fennica*, 16, 83-94, (1963).
- [45] David Lewis. *Convention: A philosophical study*, Cambridge: Harvard University Press, (1969).
- [46] Richard D. McKelvey and Thomas R. Palfrey. An experimental study of the centipede game, *Econometrica*, 60, 803-836, (1992).
- [47] John F. Nash, Jr. Equilibrium Points in n-Person Games, *Proceedings of the National Academy of Sciences of the United States of America*, Vol. 36, No. 1, 48-49, (1950).
- [48] John F. Nash, Jr. The Bargaining problem, *Econometrica*, Vol. 18, No. 2, 155-162, (1950).

- [49] Abraham Neyman. The positive value of information, *Games and Economic Behavior*, 3, 350-355, (1991).
- [50] Martin Osborne and Ariel Rubinstein. *A Course in Game Theory*, MIT Press, (1994).
- [51] Eric Pacuit and Olivier Roy. “Epistemic Foundations of Game Theory”. The Stanford Encyclopedia of Philosophy (Spring 2015 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/spr2015/entries/epistemic-game/>>.
- [52] Rohit Parikh. Social Software, *Synthese*, 132, 187-211, (2012).
- [53] Rohit Parikh and Paul Krasucki. Communication, consensus and knowledge. *Journal of Economic Theory*, 52, 178-189, (1990).
- [54] Rohit Parikh and R. Ramanujam. A Knowledge based Semantics of Messages, *J. Logic, Language and Information* 12, 453-467, (2003).
- [55] Rohit Parikh and Çağıl Taşdemir. The Strategy of Campaigning, in the *Lecture Notes in Computer Science, volume 9300, Fields of Logic and Computation II, Essays Dedicated to Yuri Gurevich on the Occasion of His 75th Birthday*, 253-260, (2015).
- [56] Rohit Parikh, Çağıl Taşdemir and Andreas Witzel. The power of knowledge in games, http://ceur-ws.org/Vol-751/reasoningminds2011_submission_3.pdf
- [57] Rohit Parikh, Çağıl Taşdemir and Andreas Witzel. Choice and Uncertainty in Games, in the *Lecture Notes in Computer Science, volume 7230, Logic and Program Semantics*, 244-255, (2012).
- [58] Rohit Parikh, Çağıl Taşdemir and Andreas Witzel. The power of knowledge in games, *International Game Theory Review* 15, 1340030, (2013).
- [59] Andres Perea. *Epistemic game theory: Reasoning and choice*, Cambridge University Press, (2012).
- [60] Jan Plaza. Logic of Public Communication, in *4th International Symposium on Methodologies for Intelligent Systems*, 201-216, (1989).

- [61] Frank P. Ramsey. Truth and probability, in *The foundations of Mathematics and Other Logical essays*, ed. R. Brathwaite, Kegan Paul, (1931).
- [62] Paramesh Ray, Independence of irrelevant alternatives, *Econometrica*, 41 (5) 987-991, (1973).
- [63] Robert van Rooy. Quality and Quantity of Information Exchange, *J. Logic, Language and Information* 12, 423-451, (2003).
- [64] Robert Rosenthal. Games of Perfect Information, Predatory Pricing, and the Chain Store. *Journal of Economic Theory*, 25, 92-100, (1981).
- [65] Ariel Rubinstein. *Modeling Bounded Rationality*, MIT Press, (1998).
- [66] Ariel Rubinstein and Yuval Salant. A Model of choice from lists, *Theoretical Economics*, 1, 3-17, (2006).
- [67] Dov Samet. Hypothetical knowledge and games with perfect information. *Games and Economic Behavior*, 17, 230-251 (1996).
- [68] Leonard Savage, *The Foundations of Statistics*, Second revised edition, Dover 1972.
- [69] Reinhard Selten. A reexamination of the perfectness concept for equilibrium points in extensive games, *International Journal of Game Theory*, 4, 25-55, (1975).
- [70] Brian Skyrms. *Signals: Evolution, Learning and Information*, Oxford U. Press 2010.
- [71] Robert Stalnaker. Belief revision in games: forward and backward induction. *Mathematical Social Sciences*, 36, 31–56 (1998).
- [72] Çağıl Taşdemir. On Tolerance Analysis of Games with Belief Revision. in the *Lecture Notes in Computer Science, volume 7734, Logical Foundations of Computer Science*, 350-358, (2013).
- [73] Amos Tversky and Daniel Kahneman. The Framing of Decisions and the Psychology of Choice, *Science*, 211, 453-458, (1981).