The Power of Knowledge in Games

Rohit Parikh[1,2], Çağıl Taşdemir[2], and Andreas Witzel[3]

[1] Brooklyn College of the City University of New York, rparikh@gc.cuny.edu
[2] Graduate Center of the City University of New York
[3] Courant Institute, New York University

3:42PM
June 21, 2011

Abstract: We propose a theory of the interaction between knowledge and games. Epistemic game theory is of course a well developed subject [4,5]. But there is also a need for a theory of how some agents can *affect* the outcome of a game by affecting the knowledge which other agents have and thereby affecting their actions.

We concentrate on games of incomplete or imperfect information, and study how conservative, moderate, or aggressive players might play such games. We provide models for the behavior of a knowledge manipulator who seeks to manipulate the knowledge states of active players in order to affect their moves and to maximize her own payoff even while she herself remains inactive.

1 Introduction

It is a commonplace that what we do depends on what we know[4]. And the *theory of mind* (Premack and Woodruff [15]) predicts that we also know that what others do will depend on what *they* know.

Bill can copy the answers from a fellow student's exambook if he knows that the teacher is not looking. And Betty can shoplift if she knows that the store does not have cameras, and that the guard is distracted.[5]

[4] or believe. We will use the word 'knowledge' neutrally, being well aware that actions often proceed from false beliefs.

[5] Deception also occurs among non-human primates. As [11] note, "... chimpanzees, one of humans' two closest primate relatives, sometimes attempt to actively conceal things from others. Specifically, when competing with a human in three novel tests, eight chimpanzees, from their first trials, chose to approach a contested food item via a route hidden from the human's view (sometimes using a circuitous path to do so)." Note that we are not claiming that chimps actually have what is called a *theory of mind*. Merely that some of their behavior seems deceptive.

But we can also proceed one level up. If Bill wants to copy from Jack's answerbook, he might ask Betty to *distract* the teacher, perhaps by asking his permission to go to the bathroom. If the store does not want Betty to shoplift then it might install TV cameras which record what happens in the store.

This second level of *arranging* for some level of knowledge or ignorance in others, in order to influence their actions, seems not to be sufficiently investigated formally.

Manipulation of knowledge can happen in two different contexts. One might manipulate knowledge for a particular purpose. For instance if Jack and Ann are going on a picnic and do not want Betty to come along, they may simply not reveal the existence of the picnic to her. That way they can avoid the situation where she says, "Oh, can I come too?" and have to either put up with her presence or else offend her by saying no.

But manipulation can also happen in a more general context. For instance a university may reveal the email address of a student to a professor who is teaching a course which the student is taking; and yet, not reveal the email address to another professor even though there is no *specific reason* why this knowledge would be harmful in some way.

1.1 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 can 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 paper, we will assume that the active players do not concern themselves with the motives of KM.

1.2 Defining rationality

Suppose an agent is in a situation of uncertainty where it has to choose between two moves L and R but does not know for sure what the outcome will be with either choice. How will the agent choose?

One option is the maxmin route. The agent can choose L if the *worst* possible outcome with L is better than the *worst* outcome with R. We will describe such an agent as *conservative*. However, an ambitious agent may choose R if the *best* outcome under R is better than the *best* outcome with L. We will describe such an agent as *aggressive*.

It is clear then that in the same situation, an aggressive agent with 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 under not buying. But those who do buy such tickets are clearly judging by the best outcome.

In most of this paper we assume that utilities are ordinal. In other words, between any two choices a, b, the agent may be neutral, prefer a or prefer b. Numbers 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, 3 or 1, 2, 4 or 1, 3, 4. It is also generally assumed that comparing utilities between different players makes no sense.

If utilities are cardinal and a subjective probability is available, we could also use expected value as a measure. However, in this work our utilities will be ordinal, and the notion of expected utility will not be available to us.

In addition to conservative and aggressive players, we can also consider moderate players who try to find the middle way. The general issue is that a player in uncertainty is choosing between two sets (or sequences) of payoffs. The payoff with L is say, $a_1 > a_2 > ... > a_k$ and with R it is $b_1 > b_2 > ... > b_m$. A conservative player chooses L over R iff a_k is preferred to b_m. An aggressive player chooses R over L iff b_1 is better than a_1. More generally, let a player use a function f to represent a sequence of outcomes by a single element. A conservative player uses the minimum, an aggressive player uses the maximum, and a moderate player uses (say) the median.

Such points of view are often taken into account by stockbrokers advising people on investments. A younger investor may prefer a stock with a high potential growth but significant risk. An investor close to retirement age may, on the contrary prefer a stock with less growth but also less risk. A middle aged investor may accept a moderate amount of risk.

The function f should satisfy some rationality conditions.

Definition 1. *A choice function f is* suitable *if it satisfies the following two conditions:*

1. *If X is a final segment of Y, then $f(X) \geq f(Y)$ and if X is an initial segment of Y then $f(Y) \geq f(X)$.*
2. *(Dubey) If sequences X and Y overlap, but all elements in $X - Y$ are higher than all elements of $Y - X$, then $f(X)$ is higher than $f(Y)$.*
3. *If sequences X and Y are in an order preserving one-one correspondence g then $g(f(X)) = f(Y)$.*

Lemma 11 *The minimum, the median and the maximum are all suitable functions in the sense above (and the corresponding notions of f-rationality are equivalent to being conservative, moderate, and aggressive respectively).*

Note that an SCF need not satisfy Nash's IIA condition that if $a = f(X)$, $Y \subseteq X$, and $a \in Y$ then $a = f(Y)$. It so happens that both the maximum and the minimum do satisfy this condition, but not the median. 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 have cardinal utilities and a subjective probability.

Definition 2. *Given an SCF f, An f-rational agent is an agent who, when uncertain between sets X and Y of alternatives, always picks X if $f(X) > f(Y)$.*

It is easily seen that if all payoffs in X are higher than those in Y then an f-rational player will choose X over Y. Thus all three kinds of players, conservative, moderate and aggressive will never pick a strictly dominated strategy.[6]

[6] By a *strictly dominated* strategy we will mean a strategy which is dominated by another pure strategy. See appendix for details.

1.3 An Example

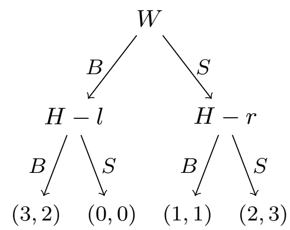

Fig. 1

In figure 1, we assume that the wife moves first and the husband after. 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) Husband does not know 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 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).

2) Finally if the husband *will* know what node he is at (and the wife knows this), then the wife will choose B, the husband will also choose B and they will end up at (3,2).

1.4 Example 2

Artemov [2] is concerned with rationality in the presence of uncertainty.[7] A rational player for him is one who makes a decision based on the highest *guaranteed* payoff, subject to the player's knowledge. In other words he describes as rational the kind of player we have chosen to call conservative.

[7] His utilities are also ordinal.

Artemov shows (his theorem 1) that a rational player *in his sense* will follow the backward induction solution *even in the absence of common knowledge of rationality*. Thus Artemov generalizes Aumann's result, replacing common knowledge of rationality by plain rationality.[8]

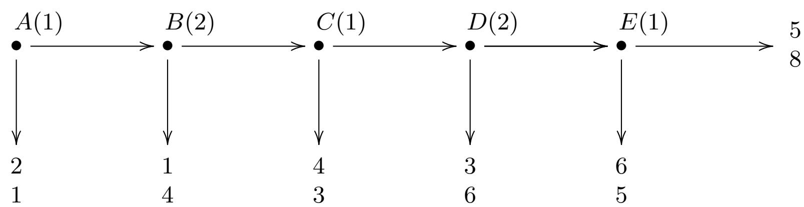

Fig. 2: Centipede game

Now consider a *moderate* player playing this game. If he had been conservative and used backward induction, he would go down at once and get a certain payoff. But since he is a moderate, he will see that the median from going across is much higher. If both players are moderate players, and agnostic about the rest of the game, then the game will continue for quite a while, with both players going across and earning much larger payoffs. Thus our notion of a moderate player shows the rationality of the common pattern seen in ordinary behaviour where players play across for quite a while.[9]

1.5 Comparison with previous work

Two relevant sources are the book by Chwe [7] and the recent paper by Artemov [2]. Chwe's book is largely concerned with the manipulation of beliefs through some form of advertising. An advertiser may seek to create the common belief that everyone is drinking beer X and so the viewer of the TV show should also drink beer X. However, Chwe's treatment is largely non-technical and does not bring in game theoretic techniques for the most part.

[8] Artemov's argument applies only to the tree. For other games Artemov's solution could diverge from the backward induction solution.

[9] We are assuming here that the players will be agnostic about the actions of the other player rather than carry out the elaborate backward induction argument.

Artemov does mention a case (section 5.2 of [2]) where revealing true information changes the behavior of the players.

What is novel in our present work is that we make knowledge manipulation the central aspect of our considerations and we do bring in some technical considerations.

Thus while we acknowledge a debt to Chwe and Artemov, we are carrying the ideas considerably further.

Other work like that of Brandenburger et al [6] is also relevant but unlike us they rely on cardinal utilities. They also do not speak about actual manipulation of behavior by limiting knowledge.

Finally, Agotnes et al have written a very interesting paper about the power which agents have over other agents who want some knowledge. Suppose A knows $P \rightarrow R$, B knows $Q \rightarrow R$ and C knows $P \wedge Q$. Then, if the interest is in knowing that R is true, C has the most power since *either* of the pair A and C or the pair B and C could derive R; but A and B together could not.

2 Game Theory

Let us consider a game tree for two *(The number two has no special significance and is only used to simplify notation.)* players with a set X of nodes, divided into X_1, the nodes where player 1 moves, X_2 where player 2 moves, and T the set of terminal nodes so that X is the disjoint union of X_1, X_2, 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 [3,2].

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 with the example in figure 1, 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 the KM, then KM will seek to manipulate the active players' knowledge in such a way as to maximize her own payoff.

2.1 States of Knowledge

We now describe a model for representing a game with possibly complex knowledge situations. We will use Kripke models for that. X is the set of nodes of the game tree.

Let us stipulate that for each element $A \in X$, A is also an atomic formula which is true precisely when the play is at node A. We create a formal language L 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 for all nodes A.

But now consider a game with two players 1 and 2 and where the formula $A \rightarrow (K_1(A) \wedge K_2(A))$ holds at all nodes A, but for instance $K_1 K_2(A)$ does not hold at node A. At each node, both players know what node it is but they do not know that the other knows.

With $(K_1(A) \wedge K_2(A))$, both players know which node they are at. But if 1 makes a choice between L and R, 2 knows which choice 1 made, but 1 does not know that 2 *will* know, then 1 might well play differently. So it is not a perfect information game, strictly speaking. Yet we cannot indicate the 'imperfection' by indicating an information set.

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

We assume that the map $A \rightsquigarrow M_A$ is common knowledge.[10] To fix thoughts, we also assume that common knowledge of temperaments (conservative, moderate or aggressive) exists. Each player plays according to his own temperament subject to what he believes about the choice situation he will be in.[11] Thus the class of knowledge situations we can

[10] We of course mean the unpointed Kripke structure M_A, since an agent who knows also what the real world is would know everything.

[11] Thus it is even open in our model to consider players who have not carried out certain deductions which they were entitled to carry out. They choose according to their belief.

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 an extended game supplemented by such a function M_A. As we noted, a perfect information game is a special case of such a KB-game. For in that case, for each A, the structure M_A has a single state satisfying A and no other states are accessible to any player.

2.2 Creating knowledge States

How would the KM create the structure M_A? One way that KM can create such structures is, at each node she sends signals to the players – the signal function s is common knowledge, and based on the signal received by the player he can infer something about the node he is at.

Definition 3. *A* game tree with knowledge function *is a standard extensive-form game tree with nodes A along with a set of signals Σ and a function $s : A \to P(\Sigma^n)$ where n is the number of players and P stands for the power set. We extend s to sequences $\in A^*$ in the obvious way. The associated* protocol *(see [14]) $H(A)$ consists of all sequences $(a_1, \sigma_1), m_1, (a_2, \sigma_2) \ldots, m_{k-1}, (a_k, \sigma_k) \in (A \times \Sigma)^*$ such that $a_1, \ldots, a_k$ is a path in the game tree starting at the root, for all $i < k$, a_{i+1} is a child of a_i resulting from the move m_i, and $\sigma_i \in s(a_i)$ for each $1 \leq i \leq k$. We define a* valuation function *$V : H(A) \to 2^A$ by setting $V((a_1, \sigma_1), \ldots, (a_k, \sigma_k)) := \{a_k\}$ (re-using the nodes as propositions). Further, we assume an* observability function *on Σ for each player which gives rise to synchronous epistemic accessibility relations in the usual way. Thus for each $\sigma_i = (s_1, ..., s_n)$, player j observes s_j and moreover the player observes all the m_i which were his own moves.*

Pradeep Dubey [8] has pointed out that by including 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. For details see the appendix.

2.3 Example 1 revisited

We consider now the question of how KM can create these various knowledge scenarios of example 1.

KM is capable of creating all these three situations by means of signals, as well as the one we did not mention where the husband does not know but the wife does not know that he will not.

For case 1a), $s(H-l) = (l, a)$ and $s(H-r) = (r, a)$. The wife knows (if she did not already) which node they are at, but the husband will not.

For case 2, $s(H-l) = (l, l)$ and $s(H-r) = (r, r)$. Both will know which node they are at.

Finally if KM wants the wife to be in doubt whether the husband knows, he could make $s(H-l) = \{(l, l), (l, a)\}$ and $s(H-r) = \{(r, r), (r, a)\}$. 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 does send (l, l) then the husband will know, 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.

The Kripke structures which arise this way will be special in three ways. In the first place it will be common knowledge that wherever the players are, they are all at *some* node of the game tree (but they may not know the actual node). Secondly, (assuming perfect recall) if a player was uncertain among nodes A and B, and only these, then she will know in the future that she must be at some node below one of A and B. Finally, if she herself performed an action α when she was so uncertain, then whatever node she is at now will have be below either the α successor of A or the α successor of B.

2.4 Predicting the play

Can the KM *always* predict how a game will be played in a less than perfect information state which he has brought about? This is indeed true in a decision theoretic situation if the temperament of the player is common knowledge[12] For instance a conservative agent faced with uncertainty will choose the least risky alternative. And since we assume that no two out-

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

comes have the same value, the least risky alternative will always be well defined and *known to the KM.*

With two person games, there may not be a unique way that the players will play in a game with imperfect information and so the KM may not be able to predict *how* they will indeed play. In particular the reasoning process of the players can be order-dependent, for consider figure 3 below.

With perfect information and CK of (conventional) rationality, the backward induction solution applies. In figure 3, 2 would choose right at B and left at A. The resulting payoffs for 1 are 4 with left, and 3 with right. He chooses L and so does 2 so they get (4,4).

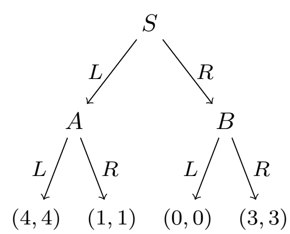

Fig. 3

But now suppose that when it is his turn to play, 2 (who is conservative) does not know whether he is at node A or B. Then he will choose Right which gives him one of {1,3}, safer than {4,0}, which he would get with Left. 1 will anticipate this and choose Right. So they end up at (3,3).

However, 2 might start his reasoning by trying to figure out 1's move. 1 will get one of {4,1} if she plays Left, and one of {0,3} if she plays Right. So she will play Left. 2 will anticipate this and will play Left. So they end up at (4,4).

Clearly the KM (whose payoffs we have not included) cannot count on any particular play.

Theorem 21 *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.*

More generally, with 2 or more players, if the players are linearly ordered so that no player knows the payoffs of any player above him then there is a unique solution.

Future work: In the setup we investigated, there is only one knowledge manipulator who, moreover, is trusted by the other players. But we can consider variants.

One possibility is where the manipulator is, well, 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 [10].

Another possibility to consider is that while the KM is presumed honest, every player is both an actor and an informer. This case would be investigated by enriching the purely informational structure of [14] and augmenting it with actions.

Acknowledgement: We thank Sergei Artemov, Pradeep Dubey, and Johan van Benthem for comments.

References

1. Thomas Agotnes, Wiebe van der Hoek and Michael Wooldridge, Scientia Potentia Est, *Proc. 10th Int. Conf. on Autonomous Agents and Multiagent Systems*, Tumer, Yolum, Sonenberg and Stone (eds.) May 2-6, 2011, Taipei, Taiwan.
2. Sergei Artemov, Rational Decisions in non-probabilistic settings, technical report TR-2009012, CUNY Ph.D. program in computer science, 2009.
3. R. Aumann. Backward Induction and Common Knowledge of Rationality, *Games and Economic Behavior*, 1995.
4. Battigalli, Pierpaolo and Bonanno, Giacomo, 1999. "Recent results on belief, knowledge and the epistemic foundations of game theory, " Research in Economics, Elsevier, vol. 53(2), pages 149-225, June
5. Adam Brandenburger, "Epistemic Game Theory: An Overview" (pre-publication version) and "Epistemic Game Theory: Complete Information" (pre-publication version) in The New Palgrave Dictionary of Economics, 2nd edition, edited by Steven Durlauf and Lawrence Blume, Palgrave Macmillan 2008
6. Adam Brandenburger, Amanda Friedenberg and Jerome Keisler, Admissibility in Games, *Econometrica* **76** (2) 307-352 (2008)
7. Michael Chwe, *Rational Ritual: Culture, Coordination and Common Knowledge*, Princeton University Press, 2001.
8. Pradeep Dubey, personal communication, March 18, 2011.
9. Jan van Eijck, and Rineke Verbrugge, *Discourses on Social Software*, Amsterdam University Press, (Paperback - Apr 1, 2010)
10. J. Farrell and M. Rabin, Cheap talk, *J. Economic Perspectives*, **10** (3) 1996, 113-118.
11. Brian Hare, Josep Call, and Michael Tomasello, Chimpanzees deceive a human competitor by hiding, *Cognition*, Volume 101, Issue 3, October 2006, Pages 495-514

12. Martin Osborne and Ariel Rubinstein, *A Course in Game Theory*, MIT Press (1994)
13. Rohit Parikh, Social Software, *Synthese*, **132**, Sep 2002, 187-211.
14. Parikh and Ramanujam, A Knowledge based Semantics of Messages, *J. Logic, Language and Information* **12** (2003), 453-467
15. Premack, D. G. and Woodruff, G. (1978). Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, **1**, 515-526 (1978.)
16. Tommy Tan and Sergio Werlang, The Bayesian foundations of solution concepts of games, *J. Economic Theory*, **45** (1988) 370-391.

3 Appendix

Proof of lemma 11 It is obvious that the median, the maximum and the minimum are preserved by isomorphism. we check the Dubey property just for the median. Suppose that X and Y overlap so that X is $a_1 > a_2 > ... > a_k > b_1 > ....b_m$ and Y is $b_1 > b_2 > ...b_m > c_1 > ... > c_p$. $X - Y$ 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. Then $i + k = m - i + 1$ and $j = p + m - j + 1$. Thus we get $2i = m + 1 - k$ and $2j = p + m + 1$. Thus $i < j$ and $b_i > b_j$. $\square$

Sketch of proof of theorem 31: We do not assume that player 2 always plays after 1. For instance the game may be over more than two stages.

At any particular node, 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 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 $\min\{p(s, s')|s \in X\}$ 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. $\square$

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'.

We now provide definitions for the way in which a KM can create appropriate Kripke structures.

Intuitively, at each node a, the KM chooses and sends an n-tuple of signals $(s_1, ..., s_n) \in f(a)$. Player j observes only s_j but can infer something about the signals received by the other players. Moreover, he observes his own moves. Based on what he has seen, he can infer a set of possible sequences compatible with what he has seen, and what he knows is what is true in all these possible sequences [14].

As Dubey has pointed out, the kinds of structures we defined above can be replaced by traditional imperfect information games with information sets, *provided that* the knowledge player's signals are treated as actual moves arising *within the game tree* rather than outside it.

Consider the case from figure 1 where KM wants the wife to be in doubt whether the husband knows, she could make $s(H-l) = \{(l,l),(l,a)\}$ and $s(H-r) = \{(r,r),(r,a)\}$. 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, but will also know that his wife did not know whether he *would* know.

Thus KM could have two moves for each of the wife's moves. After her move L, she could have an L move corresponding to the signal pair (l,l) and an R move corresponding to the signal pair (l,a). Similarly after her R move, he could have an L move corresponding to the signal pair (r,r) and an R 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. (See figure 4).

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.

Definition 4. *A* D-tree *is a standard extensive-form game tree as above, but with the choices given by s interleaved after each move, and information sets added as follows: For each player i, at each depth in the tree, any two nodes share an information set iff their parents share an i-information set and*

(i) they both result from the same action by i himself, or
(ii) they result from two s-actions which lead to the same observation for i, or
(iii) they result from some other player's actions.

Additionally, we define a valuation function *which assigns a unique proposition to all nodes generated by f-actions from the same parent. The*

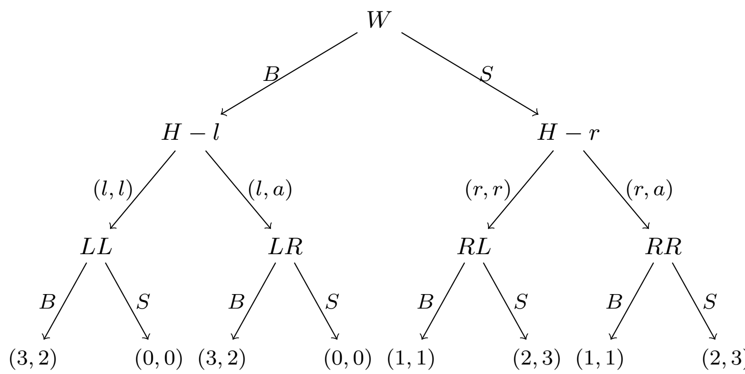

Fig. 4

knowledge situation *after n moves is the horizontal slice of this tree at depth $2n$.*

Theorem 31 *The knowledge situations in a game tree with knowledge function are isomporphic to the ones in the corresponding D-tree (modulo renaming of propositions).*

Proof. intuition: both constructions boil down to taking the product of a "normal" move and the signals that can be sent along with it, and in both constructions the indistinguishabilities are wired according to the observability of the signal part.

Theorem 32 *Any knowledge situation can be created in a single signaling step.*

Proof. Intuition: Take the Kripke structure representing the knowledge situation and create an edge from a unique (new) root node to each possible world. Label each edge with tuples $(\sigma_1, \ldots, \sigma_n)$ of signals, one for each player i, such that any two edge labels coincide in σ_i iff the worlds they lead to are indistinguishable to i. Define the observability function for player i as the restriction of a given tuple to its ith component.

Postscript on dominated strategies: Suppose that an agent believes he is facing various scenarios $s_1, ..., s_n$ but does not know which one. For each of these he has payoffs l_i from playing L and r_i from playing R and in each case $l_i > r_i$. then it is an easy one step argument that the set $\{l_1, ..., l_n\}$ has higher maximum, minimum and median values than the set $\{r_1, ..., r_n\}$. Thus whether a player is conservative, aggressive, or moderate, he will not choose R. Only the case of a moderate player requires a very short argument which we leave to the reader.