A GUIDE TO KNOWLEDGE AND GAMES

Tommy C. Tan
Graduate School of Business
University of Chicago
Chicago, Illinois

Sergio R.D.C. Werlang
IMPA &
Getulio Vargas Foundation
Rio de Janeiro, Brazil

ASTRACT

This paper serves as an informal introduction to some of the main concepts in non-cooperative games as well as a guide to where more careful treatments can be found. We pay particular attention to the explicit and implicit role of knowledge and information in games. No attempt is made to be formal and exhaustive. The interested reader should refer to the excellent book of Luce and Raiffa (195) which covers the classical material exhaustively and to the almost encyclopaedic lecture note of Van Damme (1983) which covers the vast literature on refinements of solution concepts. Games of incomplete information games are dealt with in Harsanyi (1967-68) and Myerson (1985a) at a foundational level. Tan and Werlang (1984, 1986) provides an environment in which games and the knowledge of players in games are explicitly modelled and several solution concepts are derived from axioms placed on the knowledge of players. As this note was prepared in a short time, some of the notation used may appear to be cryptic and unexplained. The intention is to go through these notes at the conference to clarify the points. The notation however, is more or less standard in literature on games.

The Nash Equilibrium

There are n players in a game identified by the index i. A **normal form game** (see Luce and Raiffa or Van Damme) is given by $\{\pi_i, A_i\}$ for i=1,...,n where A_i is a finite set of actions for each player and π_i: $X A_i \rightarrow\rightarrow \mathbf{R}$ is the payoff function for player i. Let $s_i \in S_i$ be a mixed strategy for player i (S_i is the $|A_i|$- 1 dimensional simplex). The interpretation is that each player chooses a course of action from his alternatives represented by A_i , then the payoff to each player given the strategic choice of all the players is determined by the payoff function $\pi_i(a_1,...,a_n)$. This may appear to be a somewhat limited class of games in which all players choose strategies simultaneously once and only once and the outcome of the game is determined. One easily imagines many games in which the rules are far more complicated: e.g. poker where players bet, watch others bet, then bet again etc, or chess where players move sequentially and many times before the outcome of a game is determined. Harold Kuhn proved a classic result though that it is possible to model any game equivalently as a normal form game provided the sets A_i are made rich enough (a strategy in A_i may be a very complex computer program which specifies a move in any contingency which a player encounters in an actual sequential game). I have not yet worked out whether this equivalence theorem is still valid when one changes the common knowledge assumptions implicit in much of game theory. Given a normal form game, the classical solution concept or equilibrium concept is that of Nash.

A **Nash equilibrium** is $s = \{s_1,...,s_n\}$ such that given s_{-i}, s_i maximizes player i's expected payoff. The nash equilibrium is a profile of strategies (one strategy for each player) such that given the strategies of the other players in the profile, the best a player can do for himself is to play the strategy for him in the nash equilibrium profile as well. Every player is maximizes his own payoff by playing the nash equilibrium strategy given that the other players are playing their nash equilibrium strategies. Starting at a nash equilibrium, there is no reason for any of the players to move unilaterally away from it. In general, there may be many nash equilibria in a given game. Much of recent research in non-cooperative game theory has been to provide refinements of nash equilibrium (see Van Damme) which narrow the set of equilibria. We shall discuss some of these below.

The nash equilibrium is widely used in the literature and most of the recent refinements are subsets of nash equilibria. It is regarded by many to be a necessary condition for any solution concept. The strength of the nash concept is that it is consistent with it being common knowledge. That is, common knowledge that a particular nash equilibrium will be played does not lead to a contradicting situation in which one player wishes to deviate from playing his nash equilibrium strategy. Any other equilibrium concept which is not also a nash equilibrium leads to such a contradiction.

Bernheim (1982) and Pearce (1982) questioned the necessity of the nash property. Pearce provides a persuasive attack on many of the standard arguments in favor of the nash concept. Tan and Werlang (1984, 1986) investigated the common knowledge justification. Why after all should a solution concept be common knowledge to begin with? What are the behavioral and informational assumptions which would lead to players either behaving as though a nash equilibrium were common knowledge or to playing a nash strategy even if it were not common knowledge? Pearce proposed instead a concept called rationalisable equilibria which is obtained by successively and iteratively eliminating dominated strategies. The elimination of dominated strategies requires reasoning that a player is rational. Iteratively eliminating dominated strategies requires reasoning that other players are rational, that other players know all players are rational etc. That is to say, rationality is common knowledge. Tan and Werlang (1984, 1986) provides a formal model in which the statement "bayesian rationality (in the sense of Savage) is common knowledge in a game" is explicitly formalised and the behavioral implications of this axiom derived. Not surprisingly, it implied that players must play Pearce's rationalisable equilibrium -- not necessarily a nash equilibrium. Tan and Werlang (1986) also investigates other sets of axioms which lead to nash behavior in limited classes of games and provides counter examples in more general contexts. It appears that more than what is commonly assumed to be in the common knowledge base is required to generate nash behavior.

The question of how one arrives at a nash equilibrium if one does not assume that it is common knowledge to begin with is still an open one. Learning models have been studied in game theory and in economics but the convergence of the learning process to nash equilibrium requires strong assumptions. See Samuelson in this conference volume for some discussions.

The remainder of this note will be devoted to some of the refinements of nash equilibrium which have become very prominent in the literature. The fact that there may be in general many different nash equilibria in a given game poses several problems. Firstly, how do players choose which equilibrium to play even if it is common knowledge that they would play a nash equilibrium. Somehow, not only the concept of equilibrium is common knowledge but also a particular equilibrium must be common knowledge to the players. A focal point equilibrium has been suggested by some. Secondly, if all nash equilibria are admitted by the game theorist, his prediction as to the outcome of a particular game may not be very enlightening -- alternatively, if his theory does not have a narrow set of predictions, it is easily tested or rejected scientifically. Thirdly, there are many games in which the set of nash equilibria includes some rather "unreasonable" outcomes. It is this third which motivates many of the refinements of nash equilibrium and which we now turn our attention to.

Refinements of Nash Equilibrium

Van Damme contains formal treatments of most of the equilibrium concepts raised here. Reny in this conference volume discusses the role of knowledge in extensive form games. Extensive form games are representations of games which closely follow the actual play of a game: e.g. our poker game or chess game would be represented as a sequential game tree starting at a unique node (the first move of a game) which would have branches representing choices available to the first player. At the end of each of these branches, there are either more nodes which represent the decision points for another player (it is now the 2nd player's turn to move after the first player moves) with more branches now representing the choices to the 2nd player, or payoff numbers at the end of a branch representing the end of the game. Information sets are also defined in an extensive form representation to specify what each player observes or knows when he moves. For example, the second player may not get to see what the first player has chosen when the second player has to make a choice and the game tree or extensive form has to reflect this lack of information. See Luce and Raiffa for a detailed description of game trees.

Not all Nash equilibria are sensible: Game I

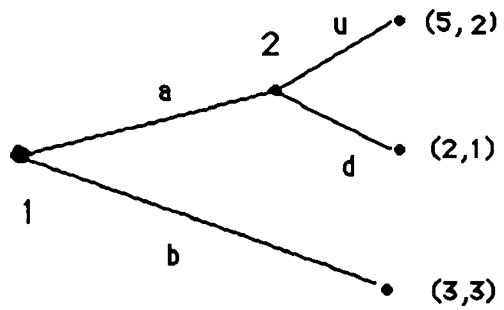

(a, u) and (b, d)

are Nash equilibria
but d is an incredible
threat by 2.

Think of game 1 as a game of extortion. Player 2 is a bank robber with a bomb threatening player 1 (the bank). Player 1 who moves first has two choices, a which is to not to pay extortion money or b which is to pay. Player 2, the bank robber, has two choices if player 1 chooses not to pay up. Player 2 can play u which represents not blowing up the bank or d which is to blow up the bank (and presumably himself). Presumably, if player 1 plays b, the bank robber just runs away and the game ends immediately. The numbers represents the payoffs to each player for each choice of strategies. The left number is the payoff to player 1 and the right the payoff to 2. The numbers are not important by themselves and only the relative sizes of the payoffs are important. The bank prefers not paying and not having the bank blown up to paying and not having the bank blown up and both of these over not paying and having the bank blown up. The robber prefers being paid and running away to not being paid and not blowing the bank up and both of these to not being paid and blowing the bank and himself up. Two of the nash equilibria in this game are (a,u) and (b,d). (a,u) is the equilibrium in which player 1 chooses not to pay and player 2 chooses not to blow up the bank. The reader should verify that this pair is indeed an equilibrium. (b,d) is the equilibrium of interest here: many would argue that this in an example of the indiscriminate nature of nash equilibrium to permit such unreasonable predictions. Here, the bank robber threatens to blow up the bank if the bank does not pay up. Given this choice of strategy by the robber, the best the bank can do is to pay up. It is therefore a Nash equilibrium.

It is considered unreasonable because " the bank should know that the robber is rational and if the bank refuses to pay, the robber faced with blowing or not blowing himself up would rationally choose not to."

Hence the bank reasoning thus should not pay. Notice this argument presumes that the bank knows that the robber is rational. Backward induction arguments of this sort requires common knowledge (or at least high levels of knowledge depending on how many steps there are in the argument) assumptions about rationality. See Reny and also Bichierri in this volume on inconsistencies in the maintained hypothesis of rationality being common knowledge at every stage of a game. Reny in particular points out that this line of argument cannot be formalised consistently in the extensive form to eliminate unreasonable equilibria. "Rationality is common knowledge" leads to logical difficulties in a game tree.

In order to eliminate these unreasonable equilibria, several refinements of nash equilibrium have been proposed. These generally have not been directly related to common knowledge arguments but their motivations frequently derive from the "rationality is common knowledge" hypothesis.

A **subgame** is a collection of branches of a game such that they start from the same node and the branches and the node together form a game tree by itself. See the picture above, player 2's decision node as well as his moves form a subgame of the game.

A **subgame perfect equilibrium** (Selten 1965) is a nash equilibrium such that the strategies when restricted to any subgame, remain a nash equilibrium of the subgame. Hence, the equilibrium of game I using the incredible threat by player 2 is not a subgame perfect equilibrium. Player 2 can improve his payoff in the subgame by playing u instead of d so that it is not the best he can do in the subgame.This equilibrium concept is the earliest refinement. It clearly works well in game 1 and seems to address some of the "unreasonableness" of nash. Subgame perfection is essentially a backward induction argument, using rationality of players at each stage of the game to decide what is a good choice and then rolling backwards. Hence players who move early in the game assume that the players in the remainder of the game are rational and would respond rationally to earlier moves by themselves.

Binmore (1985) and Reny point out that some of the counter factual reasoning in a backward induction argument are highly questionable. In particular, some stages of a game are reached only if players earlier in a game have behaved irrationally. Yet, the backward induction argument (as embodied in the subgame perfect nash equilibrium) requires that players still continue to reason that the players should behave as though rationality were common knowledge for the rest of the game. Binmore discusses the Rosenthal example and Reny discusses the repeated prisoners dilemma in this context. In both games, backward induction - or even just nash equilibrium - requires non cooperation at every stage of the game and the players end up with low payoffs. However, if one player can act irrationally early in the game by cooperating and persuading the other player that he is irrational and likely to continue cooperating, the resultant outcome may give higher payoffs than the nash equilibrium.

Despite some of these recent criticisms of backward induction it is still widely used and in fact the formal definition of subgame perfection is not efficacious in eliminating all "unreasonable equilibrium".

Not all subgame perfect equilibria are sensible: Game II

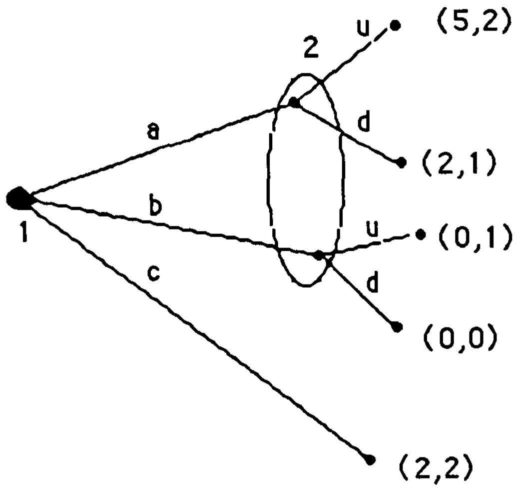

(c,d) and (a,u) are both subgame perfect even though this is just the same game as I with a strictly dominated strategy added for player 1. Trouble here is that there are no subgames.

In game II, player 2 has an information set around choice a and b of player 1 reflecting the fact that player 2 is unable to tell whether player 1 has chosen a or b (although he is able to distinguish between c and these two). Consequently, there is no proper subgame in this game and all nash equilibria are subgame perfect by default. However, (c,d) is as unreasonable here as it was in the equivalent equilibrium in game 1. Even though player 2 is unable to tell between a or b, player 2 has a dominant strategy of playing u since it gives him a higher payoff regardless of whether a or b were played. Strategy d in (c,d) equilibrium acts as before like an empty threat.

A **sequential equilibrium (Kreps and Wilson 1982)** is a profile of strategies, one for each player, and a system of beliefs with the following properties: each player has a belief over the nodes at each information set. Beginning at any information set, given a player's belief there and given the strategies of the other players, his own strategy for the rest of the game still maximizes his expected payoff. Beliefs must be obtained from the beliefs at earlier information sets and the equilibrium strategies by Bayes rule whenever it makes sense. This takes care of the unreasonable equilibrium in game II since regardless of player 2's belief, up dominates down, so down will never be part of an equilibrium strategy for 2.

Probability measures in the form of priors and posteriors have long been used to represent knowledge (or equivalently to an economist beliefs). A prior probability measure represents the initial beliefs or knowledge of a player in a game. A posterior represented an updated belief or updated knowledge base given the new information received. Game theorists run into the same problems as other researchers in modelling how knowledge is updated. In particular, one of the most controversial topics is how to update beliefs or the knowledge base when new evidence which contradicts one's prior (existing knowledge base) is encountered. See Cho and Kreps (1987), McClennan (1986a,b) and Myerson in this conference volume. We shall run into this problem in one of the examples below. As one might suspect by now even sequential equilibrium eliminates all "unreasonable equilibria".

Not all sequential equilibria are sensible: Game III

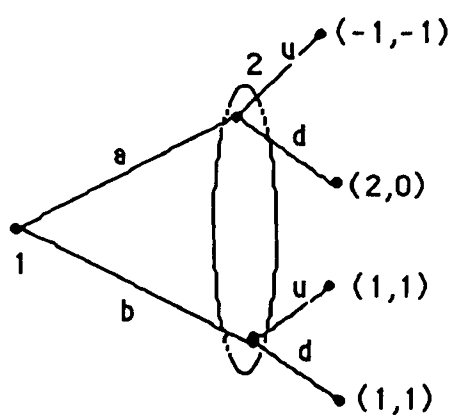

(b,u) is a sequential equib. with 2 believing that 1plays b with probability one. Notice that u is weakly dominated by d for 2.

Up is weakly dominated by down for 2. U is no better than d and if player 1 plays a it is strictly worse for player 2. Yet, (b,u) is and equilibrium. This equilibrium is supported by player 2 believing that player 1 plays b with probability one and in fact player one will do so since player 2 is playing d. The reasoning which goes on in player 1 to reject the "threat" by player 2 to play down is as follows: when it is his turn to move, player 2 should play u since it does no worse than d and in fact strictly better if for what ever reason I played a instead. When it is 2's turn to play, there is no reason for him to play d at all. Therefore I should play a if he reasons that 2 will play u. Notice that this is beyond assuming common knowledge of rationality bayesian rationality since (b, u) with player 2 having a mass point belief at b is perfectly bayesian rational. Some measure of caution or of fear of mistakes is being assumed in the players motivations and it is further assumed to be common knowledge. Selten's notion of perfection addresses this point

ε-perfect equilibrium : s_i is a mixed strategy and s_{ik} is the probability that i will play $a_{ik} \in A_i$ under s_i . s is an ε perfect equilibrium if given ε, (i) s_i is a completely mixed strategy (it's in the interior of the simplex - every strategy is played with positive probability) and (ii) if $\pi_i(s|a_{ik}) < \pi_i(s|a_{il})$ then $s_{ik} < \varepsilon$ for every k,l . If k is not a best response, then play k with probability less than ε. Let $s(\varepsilon)$ be an ε perfect equilibrium depending on ε.

An equilibrium s is **perfect** (Selten 1975) if $\exists\ \varepsilon^m \to 0,\ \exists\ s(\varepsilon^m) \to s$. This knocks out the (b,u) equilibrium in game III because for any positive probability that up will be played by 1, player two should minimize the probability that he would play up since that is dominated by down. Selten has recently presented a paper (which we have only heard about) which defends this equilibrium concept as being justified by a very simple theory of counter factuals: that of making mistakes.

Not all perfect equilibria are sensible: Game IV.

(c,d) is perfect because you can make player 1 play the strategy b with much higher probability than a, but with both close to zero and c with probability close to one. That is, a and b are both mistakes but one mistake is played with an order of magnitude higher than another. Then 2's response is to play down with probability close to one. Notice though that a strictly dominates b for player 1.

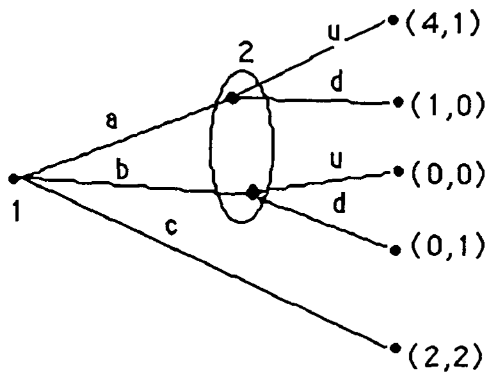

(c,d) is perfect since d is not dominated for 2. Though b is dominated for 1 by both a and c.And d only makes sense if the mistake b occurs with higher probability than mistake a.

ε proper equilibrium : Let $\varepsilon \in (0,1)$. s is an ε proper equilibrium if (i) s_i is completely mixed and (ii) if $\pi_i(s|a_{ik}) < \pi_i(s|a_{il})$ then $s_{ik} < \varepsilon . s_{il}$. If k is a costlier mistake than l, then play k with an order smaller probability. Let s(p,ε) be an ε proper equilibrium depending on ε.

proper equilibrium (Myerson 1978): s is a proper equilibrium if $\exists\ \varepsilon^m \rightarrow 0$ and

$\exists\ s(p,\varepsilon^m) \rightarrow s$. Gets rid of (c,d) because b is a costlier mistake than a for player 1. In this case, b is played with an order less probability. Hence, 2 plays up with as high a probability as possible.

Not all proper equilibria are sensible: Game V

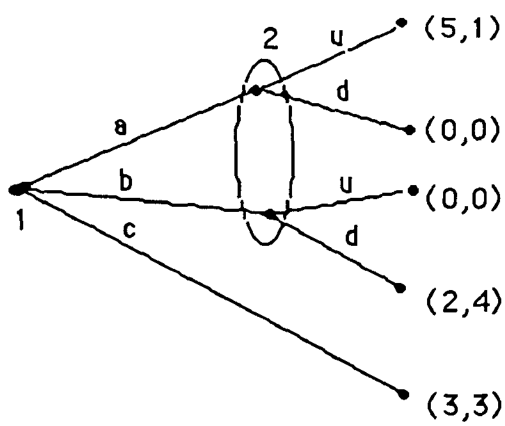

(a,u) and (c,d) are both proper even though b is strictly dominated by c and d only makes sense if b occurs with higher prob. than a by means of mistakes This is permitted in the proper equilibrium because if d is played with high prob. then b is better for 1 than a he plays a with an order lower probabilty. In which case, d witl high prob. is best for two.

	a	b
u	5,1	0,0
m	0,0	2,4
d	3,3	3,3

The reasoning by which player 1 should conclude that he should play a instead of c is very interesting. Begin from the sequential equilibrium (c,d). d is supported by player 2 having a posterior at his information set which places a high probability that player 1 played b instead of a. Player 1 as a result should play c instead as this is better than playing a or b given that player 2 will play d. As a consequence, player 2's information set will not be reached since 1 plays c. The fact that 2's information set is reached with probability zero allows 2 to have any posterior at that information set as bayes rule places no restrictions on conditioning at probability zero events. Hence a posterior with a mass point at b is perfectly consistent with this belief (or knowledge) revision rule. However, 1 knows that 2 knows that 1 is rational (since rationality is common knowledge). 1 also knows that 2 knows that c is strictly dominated by c. So if 2 believes that his information set were reached through a rational choice by 1, then 2 must conclude that since 1 had a choice of c which is better than b but c was not chosen, then 1 must have rationally chosen a instead. Therefore it is inconsistent for 2 to believe that 1 chose b and 2 should therefore place a mass point belief on a instead and therefore play u. In this case 1 should play a instead of c. This sequence of reasoning therefore demonstrates that (c,d) is inconsistent with rationality being common knowledge in the game tree and eliminates the unreasonable equilibrium. The argument we have given is embodied in Kreps (1985) and Cho and Kreps (1987). It essentially is an argument which rejects certain belief updating rules as being inconsistent with the common knowledge of rationality -- a maintained hypothesis. J. Stiglitz in a private communication notices however that the line of reasoning which eliminates certain belief rules cannot be themselves assumed to be common knowledge as they then lend themselves to manipulation.

Strictly perfect equilibria: The definition of perfect equilibrium can be interpreted as slight perturbations of the strategies. An equilibrium is perfect <u>if there exists slight perturbations</u> of the strategies such that there is an equilibrium for each perturbed game converging to the perfect equilibrium as the perturbations converge to zero. A strictly perfect equilibrium is similar, except that it survives all perturbations.

Gets rid of (c, down) since one particular perturbation could require player one to play a with higher probability than b. In this case, up is better for 2 than down. So that (c, down) doesn't survive perturbations in this direction. It was perfect because there, we chose the probability to be higher on b than on a.

Not bad, but alas, it may not always exist: Game VI

	a	b	c
u	1,1	1,0	0,0
d	1,1	0,0	1,0

a is always the best for 2, so it is always part of the equilibrium. If the tremble towards b is higher than the tremble towards c, then the only u is the best response for 1. But if the trembles are the other way around for 2, then only d is a best response for 1. So no pair of strategies survives all trembles. The Kohlberg and Mertens (1985) notion of Stable component is essentially a generalization of the strictly perfect notion so that existence is assured always. The Stable component for this game is a for player 2, and all the mixed strategies for player 1. Notice that all nearby games have equilibria near this stable component.

Literature Survey

A survey on the economics literature on knowledge prior to March 1986 may be found in Tan and Werlang (1986). Much interest continues to be focused on economic models in which some economic agents who have private information have observable actions through which other uninformed agents try to infer the private information (e.g. Arbitrageurs who observe the actions of insiders (who know more about the future profits of companies they manage) in order to trade on better information). Bayesian priors are typically used to represent the knowledge of economic agents in these models and Bayes rule represents "learning" through observation of actions over time (see Tan and Werlang (1985, 1986) and Reny (1986)).

In these models, Nash equilibrium is unsatisfactory as a solution concept because there is typically a continuum of equilibria thus rendering the models devoid of predictions. Cho and Kreps (1987), Banks and Sobel (1985), McClennan (1986a) represent attempts at refining the Nash equilibrium concept for such games by placing additional restrictions beyond those implied by Savage (1954) on the updating of beliefs at events which occur with probability zero. Such considerations are not important in statistical decision theory since probability measures are chosen by Nature. In contrast, in economic situations, probabilities of events are chosen by the agents themselves as part of their actions and agents typically take into consideration the learning (updating of priors) of the other agents when they choose their actions. Observing prior probability zero events may occur because other agents are attempting to influence your beliefs and knowledge.

This line of inquiry has resulted in recent attempts to modify the Savage (1954) axioms so that the resultant priors and decision theory better capture the knowledge of economic agents. Bewley (1986) modifies the completeness axiom so that the agents in his world have a set of priors rather than a unique one as in Savage. Myerson (1986b), Blume (1986) and Brandenburger and Dekel (1986) modify the axioms in such a way that agents have a lexicographic ordering of priors: lower levels of priors are called into expected utility calculations at the probability zero events of the higher level priors. These high order priors correspond to beliefs about the world in which one places the highest confidence and the lower order priors correspond to alternative beliefs about the world in case the higher orders have been contradicted by evidence. Applications of these approaches have been made to game theory by the respective authors.

References

Armbruster, W. and W. Böge (1979), "Bayesian Game Theory", in GameTheory and Related Topics, ed. O. Moeschlin and D. Pallaschke,Amsterdam : North-Holland, 17-28.

Aumann, R.(1974), "Subjectivity and Correlation in Randomized Strategies",Journal of Mathematical Economics, Vol. 1, 67-96.

________ (1976), "Aggreeing to Disagree", The Annals of Statistics, Vol. 4,No. 6, 1236-1239.

________ (1985), "Correlated Equilibrium as an Expression of Bayesian Rationality", Econometrica, June.

Banks, J. and J. Sobel (1985): "Equilibrium Selection in Signalling Games"Discussion Paper 85-9, University of California, San Diego.

Basu, K. (1985), "Strategic Irrationality in Extensive Games", mimeo, School of Social Science, The Institute for Advanced Study, Princeton

Bernheim, D. (1984), "Rationalizable Strategic Behavior", Econometrica, Vol. 52, No. 4, July, 1007-1028.

________ (1985): "Rationalisable Characterizations of Rational Choice in Strategic Environments," mimeo, Stanford University.

Bewley, Truman (1986): "Knightian Decision Theory," Cowles Foundation Discussion Paper , Yale University.

Bichierri, C.(1987) see this conference volume

Billingsley, P. (1968), Convergence of Probability Measures, New York : John Wiley and Sons.

Binmore, K. G. (1984), "Equilibria in Extensive Games", The Economic Journal 95,51-59.

________ (1985), "Modelling Rational Players", mimeo, Mathematics Department, London School of Economics, July.

Blume, L. (1986): "Lexicographic Refinements of Nash Equilibrium," mimeo, University of Michigan

Böge, W. (1974),"Gedanken über die Angewandte Mathematik", in Mathematiker über die Mathematik, M. Otte (ed.). Heidelberg: Springer.

Böge, W. and T. H. Eisele (1979), "On Solutions of Bayesian Games", International Journal of Game Theory, Vol. 8, Issue 4, 193-215.

Brandenburger, A. and E. Dekel (1985): " Rationalisability and Correlated Equilibria," mimeo, Harvard University.(also Econometrica forthcoming)

________ (1985a), "Common Knowledge with Probability 1", Research Paper 796R, Graduate School of Business, Stanford University, September.

________ (1985b), "Hierarchies of Beliefs and Common Knowledge", mimeo, Department of Economics,Stanford University, July.

________ (1986),"On an Axiomatic Approach to Refinements of Nash Equilibrium," Economic Theory Discussion Paper 104, University of Cambridge.

Cho, I. K. (1985), "A Refinement of the Sequential Equilibrium Concept", mimeo,Princeton University, May, forthcoming Econometrica.

Cho, I.K. and D. Kreps (1987): "Signalling and Strategic Stability," Quarterly Journal of Economics forthcoming.

Dellacherie, C. and P. A. Meyer (1978), Probabilities and Potential A, Mathematical Studies No. 29. New York: North-Holland.

Fagin R., J. Y. Halpern and M. Y. Vardi (1984), "A Model-Theoretic Analysis of Knowledge :Preliminary Report", Proc. 25th IEEE Symposium on Foundations of Computer Science, West Palm Beach, Florida, 268-278.

Farrel, J. (1985), "Credible Neologisms in Games of Communication", Economics Working Paper 85-4, GTE Laboratories, June.

Grossman, S. and M. Perry (1987), "Perfect Sequential Equilibrium", Journal of Economic Theory forthcoming.

Harsanyi, J.C. (1967-1968), "Games with Incomplete Information Played by 'Bayesian' Players", Parts I - III, Management Science 14, 159-182, 320-334, 486-502.

________ (1975),"The Tracing Procedure: A Bayesian Approach to Defining a Solution for n-Person Non-cooperative Games", International Journal of Game Theory 4, 61-95.

Harsanyi,J. C. and R. Selten (1980-84), "A Noncooperative Solution Concept with Cooperative Applications" (draft chapters 1-5), Bielefeld discussion papers.

Hildenbrand, W. (1974), Core and Equilibria in a Large Economy, Princeton : Princeton University Press.

Hintikka, J. (1962), Knowledge and Belief, Ithaca : Cornell University Press.

Kalai, E. and D. Samet (1984), "Persistent Equilibria", International Journal of Game Theory 13, 129-144.

Kohlberg, E. and J. F. Mertens (1985), "On the Strategic Stability of Equilibria", Econometrica, forthcoming.

Kreps, D. (1985), "Signalling Games and Stable Equilibria", Stanford University, mimeo

Kreps, D. and R. Wilson (1982), "Sequential Equilibria", Econometrica 50, 863-894.

Kripke, S. (1963), "Semantical Analysis of Modal Logic", Zeitschrift für Mathematische Logik und Grundlagen der Mathematik 9, 67-96.

Luce, D. and H. Raiffa (1957), Games and Decisions, New York : Wiley.

McClennan, A. (1986a): "Justifiable Beliefs in Sequential Equilibrium," *Econometrica*, 53, 1337-1351

________ (1986b): "Consistent Conditional Systems in Noncooperative Game Theory," Discussion Paper 5217-86, Mathematical Sciences Research Institute, Berkeley.

Mertens, J. F. and S. Zamir (1985), "Formulation of Bayesian Analysis for Games with Incomplete Information", International Journal of Game Theory 14, 1-29.1337-1351

________ (1987) "On the Strategic Stability of Equilibria" Econometrica

Milgrom, P. and R. Weber (1979), "Topologies on Information and Related Strategies in Games with Incomplete Information", in Game Theory and Topics, ed. O. Moeschlin and D. Pallaschke, North-Holland, Amsterdam.

Moulin, H. (1984), "Dominance Solvability and Cournot Stability", Mathematical Social Sciences 7, 83-102.

Myerson, R. (1978), "Refinements of the Nash Equilibrium Concept", International Journal of Game Theory 7, 73-80.

________ (1985a), "Bayesian Equilibrium and Incentive Compatibility : An Introduction", in Social Goals and Social Organization L. Hurwicz, D. Smeidler and S. Sonnenschein eds., 229-259, Cambridge University Press, Cambridge.

________ (1985b), "Acceptable and Predominant Correlated Equilibria", Center for Mathematical Studies, Discussion Paper No. 591, Northwestern University.

________ (1986a): "Multistage Games with Communication," Econometrica, 54, 323-258.

________ (1986b): "Axiomatic Foundations of Bayesian Decision Theory," Discussion Paper 671, J.L. Kellogg Graduate School of Management, Northwestern University.

Nielsen, L. T. (1984), "Common Knowledge, Communication, and Convergence of Beliefs", Mathematical Social Sciences 8, 1-14.

Reny, P. (1986): "Rationality, Common Knowledge and the Theory of Games," mimeo, Princeton University. (see also present conference volume)

Rubinstein, A.(1986): "Finite Automata Play the Repeated Prisoner's Dilemma," Journal of Economic Theory, 39, 83-96.

Samuelson, L. (1987) see this conference volume.

Savage, L. (1954): The Foundations of Statistics, New York, Wiley.

Selten, R. (1965), "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit", Zeitschrift für Gesamte Staatswissenschaft 121.

________ (1975), "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games", International Journal of Game Theory, Vol. 4, 25-55.

Tan, T. and S. Werlang (1984): "The Bayesian Foundations of Rationalisable Strategic Behavior and Nash Equilibrium Behavior," mimeo, Princeton University.

________ (1985): "On Aumann's Notion of Common Knowledge -- An Alternative Approach," Discussion Paper in Economics and Econometrics 85-26, University of Chicago.

________ (1986), "The Bayesian Foundations of Solution Concepts of Games", Working paper #86-34, Graduate School of Business, University of Chicago. Also Journal of Economic Theory forthcoming.

van Damme, E. (1983), Refinements of the Nash Equilibrium Concept, Lecture Notes in Mathematical Economics and Social Systems No. 219, Berlin : Springer-Verlag.