

# DYNAMIC MODELS OF DELIBERATION AND THE THEORY OF GAMES

Brian Skyrms  
Department of Philosophy  
University of California, Irvine  
Irvine, Ca. 92717  
bskyrms@uci.bitnet

## ABSTRACT

Deliberation can be modeled as a dynamic process. Where deliberation generates new information relevant to the decision under consideration, a rational decision maker will (processing costs permitting) feed back that information and reconsider. A firm decision is reached at a fixed point of this process - a deliberational equilibrium. Although there may be many situations in which informational feedback may be neglected and an essentially static theory of deliberation will suffice, there are others in which informational feedback plays a crucial role. From the point of view of procedural rationality, computation itself generates new information. Taking this point of view seriously leads to dynamic models of deliberation within which one can embed the theory of non-cooperative games.

In the sort of strategic situations considered by the theory of games, each player's optimal act depends on the acts selected by the other players. Thus each player must not only calculate expected utilities according to her current probabilities, but must also think about such calculations that other players are making and the effect on the probabilities of their acts, and of the import of other players thinking at this level as well, and so forth. In one idealized version of this problem, the players are endowed with enough initial common knowledge and computational resources so that each can emulate the reasoning of the others. A conjecture that all players have reached a firm decision as to the optimal act is consistent just in case it corresponds to a deliberational equilibrium on the part of each player. Under these idealized conditions, such a joint deliberational equilibrium on the part of the players is just a Nash equilibrium of the game. This embedding of classical game theory in the theory of dynamic deliberation suggests some non-classical extensions of the theory of strategic rationality. (I) Once one has a genuine dynamics one can ask a very rich set of questions instead of just asking which points are equilibria. As one illustration we can notice that notions of dynamical accessibility and stability of equilibria provide natural refinements of the Nash equilibrium concept. (II) In a more realistic theory, the highly idealized assumptions of common knowledge required for the coincidence of Nash equilibrium and joint deliberational equilibrium should be relaxed.

## 1 PHILOSOPHICAL ORIENTATION

The point of view adopted here depends on the cotenability of two principles. The first is that the principle of expected utility maximization is the touchstone of the theory of rationality in strategic as well as non-strategic contexts. The second is that a theory of deliberation should conceive of rationality as procedural. The first principle stands in direct opposition to a widely held view that rational decisions in the sort of situations treated by the theory of games demand an entirely different standard of rationality than those of agents dealing only with "nature". In a unified theory of rational decision, other decisionmakers should be regarded as part of nature, and rational decision will consist in maximizing one's expected payoff relative to one's uncertainty about the state of nature. The relevance of such a Bayesian viewpoint for the theory of games has become more widely appreciated as a result of fundamental papers by Harsanyi (1967), Aumann (1987), Pearce (1984), Bernheim (1984), and Kreps and Wilson (1982a), (1982b). The question which forms the basic theme of these investigations can already be found in von Neumann and Morgenstern (1944): when is a combination of strategies on the part of all the players of a game consistent with expected utility maximization on the part of all the players? The Nash equilibrium concept gives a sufficient but not a necessary condition for this to be true. The framework of this literature, although Bayesian, is still static. The focus is on a solution which satisfies a set of conditions rather than on the procedures by which the players attempt to arrive at an optimal decision.

The importance of the procedural aspect of rationality has long been emphasized by Simon (1957), (1972), (1986). In accordance with this point of view, common knowledge of rationality in strategic situations is to be thought of as common knowledge of a rational deliberational procedure. In such a situation, computations must be conceived of as generating new information. As a result, probabilities can change as a result of pure thought. There are provocative discussions of such a possibility under the name of "dynamic probability" in the writings of Good. [see Good (1983) Ch. 10 for a sample.] As probabilities evolve, so do expected utilities. There is some tension between the ideal of a full Bayesian analysis based on expected utility and a procedural approach which takes into account limitations of computational resources of the deliberators. But I do not think that the tension is so great that expected utility cannot play the central role in the analysis of deliberation. Nevertheless, the tension is responsible for some of the more interesting aspects of a deliberational approach to strategic rationality.

## 2 DYNAMIC DELIBERATION BASED ON ADAPTIVE RULES

Rational deliberators aim at maximizing expected utility. But in a situation where it is expected that new information may change your probabilities of the states of nature during deliberation, you may expect that your current expected utilities will be revised. Upon calculating expected utility the act which comes out on top may be more likely the act that looks best in the end than it was before that calculation, but it is by no means certain to be the optimal act at the end of deliberation. In such situation, then, calculations of expected utility should change the probabilities of the decisionmaker's acts in a way consistent with the decisionmaker's aiming at maximizing expected utility, but in a way which falls short of giving acts with non-maximal current expected utility probability zero. How should the decisionmaker's calculations of expected utility change the probabilities of her acts?

In principle, the decision maker might carry out an elaborate Bayesian preposterior analysis at each stage of deliberation but here we will be interested in modeling the case in which that is not feasible. Alternatively, the decisionmaker may modify probabilities of her acts according to some simple adaptive rule which reflects the fact that she is aiming for optimality. Let us say that a rule **seeks the good** just in case:

- (1) it raises the probability of an act only if that act has utility greater than that of the status quo, and
- (2) it raises the sum of the probabilities of all acts with utility greater than that of the status quo (if any).

[The expected utility of the status quo is just the average of the expected utilities of the pure acts, weighted by the decisionmaker's probabilities that she will do them.]

All dynamical rules which seek the good have the same fixed points. These are the states in which the utility of the status quo is maximal.

As a concrete example of such a rule we can take the function that Nash (1951) used to prove the existence of equilibria for finite non-zero sum games. Define the **Covetability** of an act in a state of indecision,  $p$ , as the difference in expected utility between the act and the state of indecision if the act is preferable to the state of indecision, and as zero if the state of indecision is preferable to the act  $COV(A) = \text{MAX} [U(A)-U(p), 0]$ . Then the Nash Map takes the decision maker from state of indecision  $p$  to state of indecision  $p'$  where each component  $p_i$  of  $p$  is changed to:

$$p'_i = [p_i + COV(A_i)] / [1 + \sum_i COV(A_i)]$$

Here a bold revision is hedged by averaging with the status quo. We can get a whole family of NASH MAPS by allowing different weights for the average:

$$p'_i = [k p_i + \text{COV}(A_i)] / [k + \sum_i \text{COV}(A_i)]$$

The constant  $k$  ( $k > 0$ ) is an index of caution. The higher  $k$  is, the more slowly the decision maker moves in the direction of acts which look more attractive than the status quo. In continuous time, one has the corresponding NASH FLOWS:

$$dp(A)/dt = [\text{Cov}(A) - p(A)\sum_j \text{Cov}(A_j)] / [k + \sum_j \text{Cov}(A_j)]$$

Let us model the deliberational situation in an abstract and fairly general way. A Bayesian has to choose between a finite number of acts:  $A_1 \dots A_n$ . Calculation takes time for her; although its cost is negligible. We assume that she is certain that deliberation will end and she will choose some act (perhaps a mixed one) at that time. Her state of indecision will be a probability vector assigning probabilities to each of the  $n$  acts, which sum to one. These are to be interpreted as her probabilities now that she will do the act in question at the end of deliberation. A state of indecision,  $P$ , carries with it an expected utility; the expectation according to the probability vector  $P = \langle p_1 \dots p_n \rangle$  of the expected utilities of the acts  $A_1 \dots A_n$ . The expected utility of a state of indecision is thus computed just as that of the corresponding mixed act. Indeed, the adoption of a mixed strategy can be thought of as a way to turn the state of indecision for its constituent pure acts to stone. We will call a mixed act corresponding to a state of indecision, its default mixed act. The decisionmaker's calculation of expected utility and subsequent application of the dynamical rule constitutes new information which may affect the expected utilities of the pure acts by affecting the probabilities of the states of nature which together with the act determine the payoff. In the typical game theoretical contexts they consist of the possible actions of the opposing players. For simplicity, we will assume here a finite number of states of nature.

The decision maker's personal state is then, for our purposes, determined by two things: her state of indecision and the probabilities that she assigns to states of nature. Her personal state space is the product space of her space of indecision and her space of states of nature. Deliberation defines a dynamics on this space. We could model the dynamics as either discrete or continuous, but for the moment, we will focus on discrete dynamics. We assume a dynamical function,  $\phi$ , which maps a personal state  $\langle x, y \rangle$  into a new personal state  $\langle x', y' \rangle$  in one unit of time. The dynamical function,  $\phi$ , has two associated rules: (I) the adaptive dynamical rule,  $D$ , which maps  $\langle x, y \rangle$  onto  $x'$  and (II) the informational feedback process,  $I$ , which maps  $\langle x, y \rangle$  onto  $y'$  (where  $\langle x', y' \rangle = \phi \langle x, y \rangle$ .)

A personal state  $\langle x, y \rangle$  is a deliberational equilibrium of the dynamics,  $\phi$ , if and only if  $\phi \langle x, y \rangle = \langle x, y \rangle$ . If  $D$  and  $I$  are continuous, then  $\phi$  is continuous and it follows from the Brouwer fixed point theorem that a deliberational equilibrium exists. Let  $N$  be the Nash dynamics for some  $k > 0$ . Then if the informational feedback process,  $I$ , is continuous, the dynamical function  $\langle N, I \rangle$  is continuous, and has a deliberational equilibrium. Then, since  $N$  seeks the good, for any continuous informational feedback process,  $I$ ,  $\langle N, I \rangle$  has a deliberational equilibrium  $\langle x, y \rangle$  whose corresponding mixed act maximizes expected utility in state  $\langle x, y \rangle$ . This is a point from which process  $I$  does not move  $y$  and process  $N$  does not move  $x$ . But if process  $N$  does not move  $x$ , then no other process which seeks the good will either. Therefore:

**I: If  $D$  seeks the good and  $I$  is continuous, then there is a deliberational equilibrium,  $\langle x, y \rangle$  for  $\langle D, I \rangle$ . If  $D'$  also seeks the good then  $\langle x, y \rangle$  is also a deliberational equilibrium for  $\langle D', I \rangle$ . The default mixed act corresponding to  $x$  maximizes expected utility at  $\langle x, y \rangle$ .**

Now let us consider a situation in which two (or more) Bayesian deliberators are deliberating about what action to take in a finite non-cooperative non-zero sum matrix game. We assume here that each player has only one choice to make, and that the choices are causally independent in that there is no way for one player's decision to influence the decisions of the other players. Then, from the point of view of decision theory, for each player the decisions of the other players constitute the relevant state of the world which, together with her decision, determines the consequence in accordance with the payoff matrix. [If there are more than two players, we assume here that each player calculates the probability of combinations of acts of other players by taking the product of the individual probabilities.]

Suppose, in addition, that each player has an adaptive rule,  $D$ , which seeks the good (they need not have the same rule) and that what kind of Bayesian deliberator each player is, is common knowledge. Suppose also, that each player's initial state of indecision is common knowledge, and that other player's take a given player's state of indecision as their own best estimate of what that player will ultimately do. Then initially, there is a probability assignment to all the acts for all the players which is shared by all the players and is common knowledge.

Under these strong assumptions of common knowledge, an interesting informational feedback process becomes available. Starting from the initial position, player 1 calculates expected utility and moves by her adaptive rule to a new state of indecision. She knows that the other players are Bayesian deliberators who have just carried out a similar process. And she knows their initial states of indecision and their updating rules. So she can simply go through their calculations to see their new states of indecision and update her probabilities of their acts accordingly. We will call this sort of informational feedback process Updating by Emulation. Suppose that all the players update by emulation. Then, in this ideal case, the new state is common knowledge as well and the process can be repeated. Thus, the joint state of all players is common knowledge at all times.

A combination of strategies in a non-cooperative game is a **Nash equilibrium** just in case if all players know the others' strategies, each player's strategy maximizes her expected utility. It follows immediately from the foregoing that in our idealized model Nash equilibrium coincides with deliberational equilibrium:

**II: In a game played by dynamic deliberators with a common prior probability assignment, an adaptive rule which Seeks the Good, and Updating by Emulation, with common knowledge of all the foregoing, each player is at a deliberational equilibrium at a state of the system if and only if the assignment of the default mixed acts to each player constitutes a Nash equilibrium of the game.**

### 3 DYNAMIC DELIBERATION BASED ON INDUCTIVE RULES

In game theoretic situations there is an alternative way to conceptualize deliberational dynamics according to which the dynamics is driven by inductive rules. Given other players probabilities (for acts of players other than themselves), a player can calculate their expected utilities for their options just as they do. On some form of the hypothesis that the players are optimizers, each player makes an inductive inference about the eventual play of other players, and updates her probabilities accordingly. Given sufficient common knowledge, these calculations can be emulated. Then each player again knows the probabilities that other players have (for acts of players other than themselves) and thus can calculate their expected utilities, starting the cycle anew. On the inductive dynamics, no player needs to have subjective probabilities for her own acts. She only needs to have probabilities for other players' acts, and to think about other players' probabilities on her acts. An equilibrium in inductive dynamics is thus an equilibrium in degrees of belief. This is of special interest in connection with the theory of games because of a growing conviction among game theorists that the most viable interpretation of mixed equilibria is as equilibria in beliefs.

As a simple illustration let us consider 2 by 2 two person games, where the players rely on Laplace's rule of succession, treating each round of deliberation as a virtual trial. And suppose that it is common knowledge between them that they are such Laplacian deliberators. Then each can emulate the other's calculations and discover the other's current expected utilities and utilities at each stage of deliberation.

Laplace's rule is that given  $n$  instantiations of a given act,  $A_i$  in  $N$  trials, the probability of an instantiation on a new trial is:

$$\text{Laplace: } \text{pr}(A_i) = (n + 1)/(N + 2)$$

So initially our players each assign the other probability 1/2 for each act. Each player then calculates the other player's expected utilities, identifies the act with highest expected utility and counts that act as exemplified on the first trial to get an updated probability over the other player's acts. Each player can now emulate the other player's calculation in this process to find the other player's updated probability. This process is then repeated, generating a trajectory in the joint belief space of the players. This space can be represented as the unit square with the  $y$  axis measuring Column's degree of belief that Row will play 2 and the  $x$  axis representing Row's degree of belief that Column will play 2.

If a player has more than two acts the natural generalization of Laplace's rule is a rule used in Carnap's inductive logic. [Historical details are omitted here.] Suppose that column has  $m$  possible acts,  $A_1 \dots A_m$ . After  $N$  trials in which act  $A_i$  has been chosen  $n$  times:

$$\text{Carnap: } \text{pr}(A_i) = (n + 1)/(N + m)$$

If two or more acts are tied for maximum expected utility on a round, then they are each counted as having a "fractional success" with the fractions being proportion to the acts current probabilities and adding to one.

If the game has more than two players, we will assume as before that each player takes the product measure in computing the probabilities on combinations of acts by other players. [The assumption here is made for simplicity, not on principle. Ultimately it should be relaxed.] Under this assumption (and assuming again sufficient common knowledge to allow updating by emulation) the appropriate joint belief space is just the product of the belief simplices for each player, and deliberation generates an orbit in this space.

There is a clear sense in which deliberational equilibrium corresponds to Nash equilibrium for Carnap deliberators:

**III. If Carnap deliberators are at a point in the joint belief space, then they will stay at that point just in case that point is a Nash equilibrium in beliefs of the game.**

If the point is not a Nash equilibrium then some pure strategy used with positive probability by some player does not maximize expected utility. Then under the Carnap rule the probability of that strategy is diminished. Conversely, inspection of the Carnap rule show that if the point does not move then for every player every pure strategy used with positive probability must maximize expected utility for that player. Then the point is a Nash equilibrium. ||

But this result leaves open some questions. The Carnap rule specifies the initial beliefs of the players. They are built into the inductive rule. Carnap deliberation may not be able to reach some points in the joint belief space. In particular, in every game probabilities of one or zero are never reached by a finite number of stages of deliberation. This raises the question of **accessibility**. We will say that a point in the joint belief space is **accessible** under Carnap deliberation if Carnap deliberation converges to it as the number of stages of deliberation becomes arbitrarily great. Accessible points will be discussed in the next section.

Can inductive deliberational dynamics be consistent with adaptive deliberational dynamics? Well, it could be the case that each player could find no better adaptive rule than to look at virtual trials and apply Carnap to herself. If this were the case, then the adaptive and inductive deliberational dynamics would be two descriptions of the same process.

#### 4 ACCESSIBLE POINTS

Since Carnap deliberation cannot start at a pure strategy and cannot reach one in a finite number of steps, the question is pressing whether there is a connection with the Nash equilibrium concept relevant to combinations of pure strategies. The following proposition gives an affirmative answer:



**IV. Accessible points under Carnap deliberation (pure or mixed) are Nash equilibria.**

Suppose not. Then at the point,  $p$ , some pure act,  $A$ , which does not maximize expected utility gets positive probability. By continuity of expected utility of the acts as a function of the probabilities, there is some neighborhood of  $p$  throughout which  $A$  does not maximize expected utility. Inspection of the Carnap rule shows that if deliberation converges to  $p$ , then the probability of  $A$  must - contrary to hypothesis - converge to zero.  
 ||

Not every Nash equilibrium is accessible under the Carnap dynamics. What equilibria are accessible depends in part on the structure of the particular game in question, and on the a priori probabilities built into the Carnapian inductive logic. But there are some properties that accessible equilibria have in general under the Carnap dynamics, and which remain properties of accessible equilibria under various Bayesian generalizations of the Carnap dynamics.

Not all Nash equilibria are created equal. Consider the following game:

Example 1

	C1	C2
R1	1,1	-2,-2
R2	0,0	0,0

There is a Nash equilibrium in beliefs with Row believing with probability one that Column will do C2 and Column believing with probability one that Row will do R2. Doing R2 maximizes expected utility for row and doing either act has maximal expected utility for column.

Column's act C2 is weakly dominated by C1. That is, for some state of the world (i.e. R1), C1 gives column greater payoff than C2 and for all states of the world, C1 gives column at least as great a payoff as C2. Acts which are not weakly dominated are called admissible. Since C2 is inadmissible, if column is not absolutely certain about row choosing R2 - if column assigns each of Row's choices positive probability - then C2 would no longer maximize expected utility for column.

It is often assumed that a rational player will never choose an inadmissible act. However this conclusion is not a consequence of the postulate that players maximize expected utility. If Column's probability of R2 is 1, then his choice of the inadmissible act C2 does maximize expected utility. It is of some interest then, to see if there is a connection between accessibility and admissibility. For the Carnap dynamics one can show:

**V. Accessible points under the Carnap dynamics only give positive probability to weakly admissible acts.**

Suppose not. Then there is a weakly dominated act, A, which has positive probability at the accessible point, p. Since A is weakly dominated it does not maximize expected utility at any completely mixed point. Inspection of the rules shows that if deliberation converges to p, then probability of A must converge to zero. ||

Readers familiar with Selten's (1975) notion of perfection can note that it is an immediate corollary that *in two person games, equilibria that are accessible under Carnap deliberation are perfect*. Due to limitations of space I cannot develop the concept of perfection here. Connections between accessibility, admissibility and perfection are discussed in more detail in Skyrms (forthcoming c).

Carnap's inductive rule may strike one as rather special. If so, it is of interest that that propositions III, IV and V continue to hold good when we generalize the Carnap dynamics to a wide class of Bayesian models. Suppose we model a process generating our evidence as being like sampling without replacement from an urn of unknown composition. The natural conjugate prior is the Dirichlet distribution. If there are m possible outcomes then the Dirichlet distribution is characterized by positive parameters  $\alpha_1 \dots \alpha_m$ . Take  $\beta = \alpha_1 + \dots + \alpha_m$ . The if n is the number of occurrences of  $A_i$  in N trials we get the general rule:

$$\text{(Dirichlet) } \Pr(A_i) = (n + \alpha_i) / (N + \beta)$$

If the  $\alpha_i$  are all 1 we get Carnap. [In his posthumously published system, Carnap in fact moved to this parametric Bayesian rule.] Other choices of the parameters allow any starting point in the interior of the joint belief space of the players. It is easy to see that the reasoning used to establish propositions III-V holds good when (Dirichlet) is substituted for (Carnap). The results can be further generalized, using mixtures of Dirichlet priors to approximate arbitrary subjective priors. [See Diaconis and Ylvisaker (1984)]

One might, however, be tempted to overgeneralize. Consider the adaptive dynamics that underlies the game theory of Maynard Smith (1982):

$$(\text{Darwin}) \text{ NEWPr}(A_i) = \text{OldPr}(A_i) [U(A_i)/U(\text{Status Quo})]$$

In example 1, dynamic deliberation according to this rule can start at a completely mixed point and converge to an equilibrium where  $\text{Pr}(R2) = 1$  and  $2/3 < \text{Pr}(C2) < 1$ . Such an equilibrium is imperfect and has Column using the inadmissible strategy, C2, with positive probability. The discussion of generalized evolutionary game theory in Samuelson (1988) is highly relevant to this point.

## 5 IMPRECISE PROBABILITIES

Our idealized model of games played by Bayesian deliberators makes the unrealistic assumption that at the onset of deliberation precise states of indecision of the players are common knowledge. Precise states of indecision are then maintained throughout deliberation. It is of interest to weaken this assumption as a move in the direction of greater realism and for the theoretical interest of the interaction of the dynamics and the imprecision of the beliefs. In this section, for simplicity, I will fix the dynamics as the Nash dynamics.

There are various ways in which imprecise states of indecision might be modeled. Here we will look at the computationally simplest alternative. Instead of taking a player's state of indecision to be a probability measure over his space of final actions, we will take it to be a convex set of probability measures. We focus here on the simplest case of 2 person games, where each player has only two possible actions, and a player's state of indecision is given by a closed interval. If, for example, Row's probability of act 2 is to lie in the interval between Row's upper probability of act 2 = .7 and Row's lower probability of act 2 = .6, then the extreme probability measure corresponding to Row's upper probability of act 2 is  $\text{Pr}(A2) = .7, \text{Pr}(A1) = .3$  and the extreme measure corresponding to Row's lower probability of act 2 is  $\text{Pr}(A2) = .6, \text{Pr}(A1) = .4$ . The convex set in question is composed of all probability measures over the space  $[A1, A2]$  which can be gotten by a weighted average of the extreme measures.

How should Row calculate expected utilities given Column's probability interval? He should have a set of expected utilities, one corresponding to each possible point probability consistent with Column's probability interval. Because of the nature of the expectation however, Row need only compute the expected utilities relative to the endpoints of Column's interval with assurance that the other point utilities lie inbetween.

How should Row modify his probability sets in the light of new expected utility sets. He should have new probability sets corresponding to every point gotten by applying his dynamical law to a point chosen from the expected utility set and a point chosen from his old probability set. But for the Nash dynamics, and a large class of reasonable dynamical laws to which it belongs, it is a consequence of the form of the dynamical law that if old  $Pr_2(A)$  is in the interval from old  $Pr_1(A)$  to old  $Pr_3(A)$  and old  $U_2(A)$  is between old  $U_1(A)$  and old  $U_3(A)$  then new  $Pr_2(A)$  is between new  $Pr_1(A)$  and new  $Pr_3(A)$ . It is a consequence of these observations, that Row can achieve the results of point deliberation on every pair consisting of one point from his interval and one from column's interval, by performing four point computations on pairs consisting of one endpoint from his interval and one from columns interval. The new maximum and minimum probabilities of A among the four possibilities form the endpoints of his new probability for A. Here is the relevant subroutine:

```

SUB INTDYNAMICS
  CALL NASH(UPrR2,UPrC2) 'Call Nash dynamics applied to upper probabilities of R2 and C2
  Y1=OutPrR2' Set Y1 equal to the output for Pr(R2)
  X1=OutPrC2
  CALL NASH(UPrR2,LPrC2)
  Y2=OutPrR2
  X2=OutPrC2
  CALL NASH(LPrR2,UPrC2)
  Y3=OutPrR2
  X3=OutPrC2
  CALL NASH(LPrR2,LPrC2)
  Y4=OutPrR2
  X4=OutPrC2
  UPrR2=FNMAX(Y1,Y2,Y3,Y4)
  LPrR2=FNMIN(Y1,Y2,Y3,Y4)
  UPrC2=FNMAX(X1,X2,X3,X4)
  LPrC2=FNMIN(X1,X2,X3,X4)
END SUB

```

The general points made above continue to hold good *mutatis mutandis* for numbers of acts greater than two, with intervals being generalized to convex sets of probability measures and endpoints being generalized to extreme points. With regard to computational tractability, deliberational dynamics, as sofar developed, has a certain affinity for convex set representations of imprecise probabilities.

Returning to the case of two players each of which must choose between two acts, a state of indecision in the interval valued sense is now represented as a rectangle in the old space of indecision - the product of row and column's intervals. Points are considered degenerate intervals, and point states of indecision are special cases of rectangles of indecision. The area of a rectangle of indecision need not be preserved by deliberational dynamics. For example, players may start out with non-degenerate interval-valued probabilities, and be carried by deliberation to point probabilities. One might call such a process *elicitation of point probabilities through deliberation*. This process is illustrated in the case of a game with elements of both competition and coordination in figure 1. Here, in the game of Chicken, we have the orbit of  $[\cdot, 1], [\cdot, 1]$  converging to  $0, 0$  and that of

[.6,.9],[.6,.9] converging to 1,1.




figure 1: Interval Deliberation in Chicken

		Chicken	
		C1	C2
R2	-	-5,5	0,0
R1	-	-10,-10	5,-5

In the case of Chicken, the behavior of intervals under the Nash dynamics is just what one might expect from an investigation of the point dynamics of this game. One can predict the orbit of a rectangle of indecision by applying the point dynamics to its four corners. The relation between the point dynamics and the interval dynamics is not, however, always so simple. This can be illustrated by the simple game of Matching Pennies:

		Matching Pennies	
		C1	C2
R2	-	1,-1	-1,1
R1	-	-1,1	1,-1

In this game there is a unique mixed equilibrium in which each player gives each act equal probability. In the Nash point dynamics this point is a strongly stable equilibrium which has as its basin of attraction the whole space of indecision. In the Nash interval dynamics, any starting point in which both players have non-degenerate intervals, no matter how small, leads to a state where both players have maximal  $[0,1]$  intervals.

These two examples certainly do not exhaust the field, but perhaps they are sufficient to allow a modest moral. The questions posed by imprecise beliefs cannot be adequately addressed in isolation from questions about the dynamics of deliberation.

### Acknowledgement

This paper sketches parts of an area treated more fully in forthcoming publications [Skyrms (forthcoming a,b,c)]. The relevant research was partially supported by the National Science Foundation Grant SES-8721469 and by the John Simon Guggenheim Foundation.

## References

- Aumann, R. J. 1987. "Correlated Equilibrium as an Expression of Bayesian Rationality" *Econometrica* 55:1-18.
- Bernheim, B. D. 1984. "Rationalizable Strategic Behavior" *Econometrica* 52:1007-1028.
- Binmore, K. 1987. "Modeling Rational Players I and II" *Economics and Philosophy* 3, 179-214; 4, 9-55.
- Brown, G. W. 1951. "Iterative Solutions of Games by Fictitious Play" in *Activity Analysis of Production and Allocation* (Cowles Commission Monograph) New York: Wiley. 374-376.
- Carnap, R. 1950. *Logical Foundations of Probability* Chicago: University of Chicago Press.
- Carnap, R. 1952. *The Continuum of Inductive Methods* Chicago: University of Chicago Press.
- Carnap, R. 1971. "A Basic System of Inductive Logic, Part 1" in *Studies in Inductive Logic and Probability* vol. I ed. R. Carnap and R. C. Jeffrey. Berkeley: University of California Press.
- Carnap, R. 1980. "A Basic System of Inductive Logic, Part 2" in *Studies in Inductive Logic and Probability* vol. II ed. R. C. Jeffrey. Berkeley: University of California Press.
- Diaconis, P. and Ylvisaker, D. 1984. "Quantifying Prior Opinion" in *Bayesian Statistics 2* ed. J. M. Bernardo et al. Amsterdam: North Holland.
- Eells, E. 1984. "Metatrickles and the Dynamics of Deliberation" *Theory and Decision* 17: 71-95.
- Good, I. J. 1965. *The Estimation of Probabilities: An Essay on Modern Bayesian Methods* Cambridge, Mass.: MIT Press.
- Good, I. J. 1983. *Good Thinking: The Foundations of Probability and Its Applications* Minneapolis: University of Minnesota Press.
- Harper, W. 1988. "Causal Decision Theory and Game Theory: A Classic Argument for Equilibrium Solutions, a Defense of Weak Equilibria, and a New Problem for the Normal Form Representation" in *Causation in Decision, Belief Change, and Statistics* ed. Harper and Skyrms. Dordrecht: Kluwer. 25-48.
- Harsanyi, J. C. 1967. "Games with Incomplete Information Played by Bayesian Players" parts I,II,III *Management Science* 14:159-183, 320-334, 486-502.
- Harsanyi, J. C. 1973. "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed Strategy Equilibrium Points" *International Journal of Game Theory* 2:1-23.
- Harsanyi, J. C. and Selten, R. 1988. *A General Theory of Equilibrium Selection in Games* Cambridge, Mass.: MIT Press.
- Kreps, D. and Wilson, R. 1982a. "Sequential Equilibria" *Econometrica* 50:863-894.
- Kreps, D. and Wilson, R. 1982b. "Reputation and Incomplete Information" *Journal of Economic Theory* 27:253-279.

- Maynard Smith, J. 1982. *Evolution and the Theory of Games* Cambridge: Cambridge University Press.
- Nash, J. 1951. "Non-Cooperative Games" *Annals of Mathematics* 54:286-295.
- Pearce, D. G. 1984. "Rationalizable Strategic Behavior and the Problem of Perfection" *Econometrica* 52:1029-1050.
- Samuleson, L. 1988. "Evolutionary Foundations for Solution Concepts for Finite, Two-Player, Normal-Form Games" in *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge* ed. M. Vardi. Los Altos, California: Morgan Kaufmann. 211-226.
- Selten, R. 1975. "Reexamination of the Perfectness Concept of Equilibrium in Extensive Games" *International Journal of Game Theory* 4:25-55.
- Simon, H. 1957. *Models of Man* New York: Wiley.
- Simon, H. 1972. "Theories of Bounded Rationality" In *Decision and Organization* ed. C.B.McGuire and R. Radner. Amsterdam: North Holland.
- Simon, H. 1986. "Rationality in Psychology and Economics" In *Rational Choice* ed. Hogarth and Rader. Chicago: University of Chicago Press.
- Skyrms, B. 1984. *Pragmatics and Empiricism* New Haven: Yale University Press.
- Skyrms, B. 1986. "Deliberational Equilibria" *Topoi* 5:59-67.
- Skyrms, B. 1988. "Deliberational Dynamics and the Foundations of Bayesian Game Theory" In *Epistemology [Philosophical Perspectives v.2]* ed. J.E.Tomberlin. Northridge: Ridgeview.
- Skyrms, B. 1989. "Correlated Equilibria and the Dynamics of Rational Deliberation" *Erkenntnis* 31: 347-364.
- Skyrms, B. forthcoming a. *The Dynamics of Rational Deliberation* Cambridge, Mass: Harvard University Press.
- Skyrms, B. forthcoming b. "Ratifiability and the Logic of Decision" in *Philosophy of the Human Sciences (Midwest Studies in Philosophy vol. 15)* University of Minnesota Press: Minneapolis, Minnesota.
- Skyrms, B. forthcoming c. "Inductive Deliberation, Admissible Acts and Perfect Equilibrium" in *Essays in the Foundations of Rational Decision* ed. M. Bacharach and S. Hurley. Oxford: Blackwells.
- Smith, C. A. B. 1966. "Consistency in Statistical Inference and Decision" *Journal of the Royal Statistical Society, Series B* 23: 1-37.
- Spohn, W. 1982. "How to Make Sense of Game Theory" in *Studies in Contemporary Economics vol2:Philosophy of Economics* ed. W. Stegmuller et. al. Heidelberg and New York: Springer Verlag.
- von Neumann, J. and Morgenstern, O. 1947. *Theory of Games and Economic Behavior* Princeton: Princeton University Press.