Part XI

THE DYNAMICS OF RATIONAL DELIBERATION

Dynamic Deliberation

William L. Harper

The University of Western Ontario

Brian Skyrms' investigation of dynamic deliberation began when he contrasted dynamical deliberation based on evidential decision theory with dynamical deliberation based on causal decision theory(Skyrms1982). According to Skyrms, counterintuitive features of deliberation dynamics based on evidential decision theory undercut attempts by Ellery Eells and Richard Jeffrey to use dynamical considerations to argue that evidential decision theory could be trusted to agree with the recommendations of causal decision theory in the examples that had been used to motivate causal decision theory.

Skyrms went on to investigate the light dynamic deliberation could throw on game theory and on the relation between decision and information. His well known book,*The Dynamics of Rational Deliberation,* reports the interesting results he had obtained by 1990. Since the book, Skyrms' has begun to investigate dynamical deliberation more closely tied to rational belief change and to induction. He has also begun to investigate interpretations of dynamic deliberation that can illuminate the new trend in game theory to interpret equilibria as more or less stable configurations that can be reached as states of indecision evolve by processes that are more like evolution than like deliberations of rational agents.

1. Basics

Skyrms models the state of indecision of an agent at a stage of deliberation as a probability assignment to the acts the agent is deliberating about. The utility of the status quo for such a state of indecision can be represented as the expected utility of the corresponding mixed act. At each stage of deliberation the agent calculates expected utility and then adjusts the state of indecision by a dynamical rule. The rule is required to:

(a) raise the probability of an act only if that act has utility greater than the status quo.

(b) raise the sum of the probabilities of all acts with utility greater than the status quo.

A rule which satisfies these conditions is said to *seek the good.*

PSA 1992, Volume 2, pp. 353-364

https://doi.org/10.1017/S0270864700009437 Published online by Cambridge University Press

Upon applying a dynamical rule to change her state of indecision an agent may also provide information feedback which changes the probabilities she attributes to the events on which the outcomes of her acts depend. A deliberational equilibrium for a given dynamical rule and information feedback system will be a fixed point in the deliberation process—a state which is not changed by additional stages of deliberation. Skyrms shows that for any continuous information feedback system a dynamical rule that seeks the good will lead to a deliberational equililibrium which will also count as a deliberational equilibrium for any other dynamical rule which seeks the good.

A deliberational equilibrium may be a decision to perform a specific act if it concentrates all the probability on that act, or it may correspond to a mixed act where the probabilities are distributed over several acts if such a state of indecision is a fixed point of deliberation. We can think of the default mixed act corresponding to such a state as the best estimate of the probabilities if the agent were forced to choose while in such a state. This interpretation corresponds to the assumption that an agent's deliberation is effective in changing her state of indecision in the way it specifies.[1]

Skyrms suggests that dynamic deliberation can illuminate the relation between game theory and individual rational choice theory. In a game the decisions of the other players are events on which the outcomes of an agent's acts depend. Skyrms supposes that the initial state of indecision and the dynamical rule of each player is common knowledge. This provides for updating by emulation as an information feedback, where at each stage the other players go through the calculations of a given player and use the resulting state of indecision as their best estimate of what that player will end up doing. Skyrms shows that under these assumptions the adaptive rules seek the good and each player is at a deliberational equilibrium if and only if the resulting acts count as a Nash equilibrium of the game. That is the act(mixed or pure) corresponding to the end point state of each player is a best reply to the combination of acts that correspond to the end point states of the other players.

Skyrms offers dynamic deliberation as a contribution toward illuminating the very important role Nash equilibria have played in game theory . To the extent that the strong common knowledge assumptions about the initial states of indecision together with the assumptions about the dynamical rules and feedback by emulation are reasonable the result can support the claim that solutions of games played by rational agents ought to be equilibria. In the 80s one of the most vibrant topics in game theory was the idea that refinements of the equilibrium concept would lead to solutions for many games with multiple non interchangeable equilibria by giving reasons to suppose that some of the equilibria would not count as rational. Skyrms shows that dynamical deliberation can lead to such refinements of the Nash equilibrium.[2]

In the book Skyrms focused, for the most part, on two dynamical rules: The Nash dynamic, which is based on a construction that John Nash used to prove that all finite games have equilibria, and the Darwin dynamic based on the evolutionary games studied by Maynard Smith.[3] I shall concentrate on the Darwin dynamic,[4] but will mention corresponding results for the Nash dynamic where appropriate.

2. Noncredible Threats

Let us see what happens when the Darwin dynamics with updating by emulation is applied to represent the deliberations of rational agents in two person non-cooperative games. Our first game illustrates the idea of non-credible threats which Selten used to initiate what has become a major game theoretic industry of investigating refinements to the Nash equilibrium concept. Here is the game in normal form.

https://doi.org/10.1017/S0270864700009437 Published online by Cambridge University Press

	a	b
C	(.5,1)	(.5,1)
D	(0,0)	(1,.5)

Row chooser's pure strategies are C (for comply) and D (for defy) and has utilities represented by the numbers on the left.[5] The non-credible threat is for column chooser to play a (aggressive) in an attempt to force row chooser to play C. The other pure strategy is b (to back down) which lets Row chooser profit from playing D.

Let our deliberation start from the initial position (.5C ,.5D) and (.5a ,.5b). This initial position is assumed to be common knowledge. The Darwin dynamic is

$$P_{n+1}(X) = P_n(X)(U_n(X)/U_n(SQ))$$

where n is a stage of deliberation , X is a players strategy , $P_{n+1}(X)$ is the probability assigned to X at stage n+1 , $P_n(X)$ is the probability assigned to X at stage n , $U_n(X)$ is the utility of X at the nth stage of deliberation , and $U_n(SQ)$ is the utility of the status quo at the nth stage of deliberation. For row chooser the utility of the status quo at stage 0 is

$$U_0(SQ)= P_0(C)U_0(C) + P_0(D)U_0(D) = (.5)(.5) + (.5)U_0(D)$$

At stage 0 the probabilities are those specified in the initial position. The utility to row chooser of C is .5 whatever column chooser does. The utility of D depends on the probability assigned to column chooser's acts. At stage 0 these are each .5, so $U_0(D) = (.5)(0) + (.5)(1) = .5$. In this example the Darwin dynamic makes no change for row chooser at stage 0 , since $U_0(D)=U_0(C)=U_0(SQ)$.

Consider column chooser at stage 0. We have $U_0(a) = .5(1) + .5(0) = .5$ and $U_0(b) = .5(1) + .5(.5) =.75$,which makes $U_0(SQ) = .625$, since $U_0(SQ) = (.5)U_0(a) + (.5)U_0(b)$. Here the Darwin dynamic makes a change. $P_1(b) = P_0(b)(U_0(b)/U_0(SQ)) = .6$. Row chooser is already starting to move toward b. This dynamic leads from the initial state (.5C,.5D) (.5a,.5b) to the equilibrium (D,b). Here are the first few steps as probabilities of D and of b :

step 0 .5D .5b
step 1 .5D .6b
step 2 .545D .692b
step 3 .624D .782b
step 4 .722D .871b
step 5 .819D .94b
step 6 .895D .982b

The dynamics eventually settles on the equilibrium (D,b). The equilibrium (C,a) is not reached from this starting point. Indeed, in this game the Darwin dynamic will reach (D,b) from any starting point that does not already assign probability 1 to a. We see here that the Darwin dynamic can act as an equilibrium selection. This is an advantage in a game like this one where there is more than one equilibrium and a fairly compelling reason for selecting one as most rational.

In the extensive form the non-credibility of the threat can be more obvious. Suppose Row chooser gets to move first, and has the utilities at the top, so that the game tree is as follows. Row chooser is player 1, column chooser is player 2.

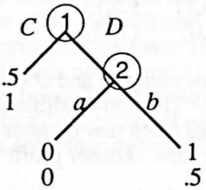

This makes player 2's threat reduce to a commitment to choose 0 over .5. Skyrms has provided a way of building extensive form information into the dynamics. The idea is that probabilities are evaluated relative to information sets governing choices. In this game player 2's utilities are $U_0(a)=0$ and $U_0(b)=.5$,so that $U_0(SQ) = .5(0) + .5(.5) = .25$. This makes the Darwin dynamics give a one step jump to *b*.[6]

$$P_1(b) = P_0(b)\ (U_0(b)/U_0(SQ)) = .5(.5/.25)=1$$

In this example the Darwin dynamics exploits the information provided by the extensive form to very quickly reach player 2's optimal choice which is to back down.

The Darwin dynamic does not do such an efficient job for player 1. Here are the probabilities of *D* in the first few stages of deliberation, .5 , .667 , .8 , .889 , .941 , .969 , .983 , .991 , .995 . By stage 8 player 1 still has not settled on choosing *D*. This failure to move player 1 more quickly to *D* is some what at odds with the motivation Skyrms provides for rational agents to leave room for additional stages of deliberation.

> Let us suppose that one deliberates by calculating expected utility. In the simplest cases, deliberation is trivial; one calculates expected utility and maximizes. But in more interesting cases, the very process of deliberation may generate information that is relevant to the evaluation of the expected utilities. Then,processing costs permitting, a Bayesian deliberator will feed back that information, modify his probabilities of states of the world, and recalculate expected utilities in light of the new knowledge.(1990,p28)

In the normal form game when player 1 changes the probabilities assigned to *C* and *D* this can provide information (in the feedback by emulation) which changes the probabilities assigned to player 2's playing *a* or *b*. In the extensive form example, however, player 2's choice of *b* is fixed after the first stage, so that player 1's later stages of deliberation are not providing any feedback about what player 2 may be expected to do.[7] This suggests that Darwin would need to be supplemented by some rule for shortcutting in such examples if it is to model dynamic deliberation appropriate to rational deliberators.

3. An Incorrect Recommendation

In the non credible threat example the Darwin dynamics gives what is intuitively the correct answer from almost any starting point. I have criticized its failure to shortcut stages of deliberation when the motivation Skyrms has offered for having additional stages is missing, but we have not yet seen any grounds for supposing that it would lead to a recommendation that would not be rational. I shall argue that the game we are about to consider does raise such worries. In order to make this case I shall contrast the dynamic deliberation approach with a different approach based on a reconstruction of a suggestion by Von Neumann and Morgenstern made possible by applying causal decision theory to game theory.

Here is the game in extensive form.[8] Player 1 has pure strategies *A*, *L* and *R*, and has the utilities at the top. Player 2 chooses at information set 2, which leaves open whether Player 1 has played *L* or *R*.

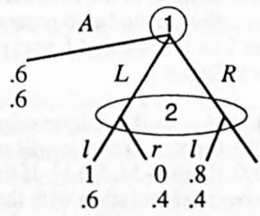

Let us suppose, as before, that the initial states of indecision distribute probabilities equally over the alternative pure strategies. In this case we have (1/3 *A* , 1/3 *L* , 1/3 *R*) and (1/2 *l* , 1/2 *r*). Consider player 1. We have $U_0(A) = .6$, while $U_0(L) = .5(1) + .5(0) = .5$ and $U_0(R) = .5(.8) + .5(.2) = .5$ as well. $U_0(SQ)$ for player 1 is (1/3)(.6) + (1/3)(.5) + (1/3)(.5) = .5333. The Darwin dynamic gives $P_1(A)$=.375 and gives $P_1(L) = .3125 = P_1(R)$. An important thing to notice here is that player 1 moves toward *A*, but keeps dividing the rest of the probability evenly between *L* and *R*.

Consider player 2. We have $U_0(l) = (1/3)(.6) + (1/3)(.4)$ and $U_0(r) = (1/3)(.6) + (1/3)(.4) + (1/3)(.6)$, so that *l* and *r* have equal utility as long as the probabilities assigned to *L* and *R* are equal. So long as these probabilities are equal the utilities of *l*, *r* and the status quo will be equal for player 2 so that player 2 will stay at the initial state of indecision which assigned .5*r* and .5*l*. The result of applying the Darwin dynamics to this game from this starting point will be that player 2 stays at (.5*r*,.5*l*), while player 1 converges to *A*. The pair (*A* ,(.5*r*,.5*l*)) is a Nash equilibrium which, I shall argue, is not a solution which rational players should opt for in this game.

Consider yourself player 1 facing this game. You do not yet know what the solution to this game is, but, initially at least, you believe that there is a rational solution which each of you will end up doing your parts to achieve. You are now in a position to apply an indirect argument suggested by Von Neumann and Morgenstern to test your strategies as candidates for your part of such a solution.[9] Under the assumption that some strategy X is your part of the solution you can assume that the other player will figure this out and will make a best reply to it. let P be a best reply prior just in case P(Y/X)=0 unless Y is a best reply to X.

If on the assumption you were going to do X another act X' would be evaluated as better than X , then it would not be rational to commit youself to doing X. We can define ratifiability as a requirement on any act worthy of choice.[10]

X is ratifiable iff $U_X(X) \geq U_X(X')$ for all X',

where $U_X(X')$ is your evaluation on the assumption that you will do X of the ulility you would expect if you were to do X' instead.[11] Now we have the apparatus to apply the Von Neumann indirect argument as a test. Any strategy will fail this test unless some best reply prior will ratify it.

https://doi.org/10.1017/S0270864700009437 Published online by Cambridge University Press

Consider strategy R for player 1. The unique best reply is for player 2 to play r. This makes $U_R(R) = .2$ for any best reply prior, but $U_R(A) = .6$ whatever player 2 does, so R fails the test. Consider L. $U_L(L) = 1$, which is at least as high as $U_L(X)$ for any alternative X, so L passes. In order to evaluate A we need to consider $U_A(L)$. If player 1 were to do L instead of A this would put player 2 into the information set, which player 2 would not get to if player 1 did A. The evaluation of UA(L) depends on what player 2 can be expected to assume at the information set. The only strategy of player 1 which meets the information set and also passes the test is the pure strategy L.[12] Player 1 can expect player 2 to be sure that L was played, therefore $U_A(L) = 1 > .6 = U_A(A)$, which makes A unratifiable.

I have argued that this game has (L,l) as the unique solution that rational players would reach. We have seen that the Darwin dynamic will not reach this solution from the starting point (1/3 A, 1/3 L , 1/3 R) and (.5l,.5r).[13] If the dynamical deliberation approach to game theory is to recover this solution with the Darwin (or Nash) dynamics, then some way must be found to argue that this initial position we started from is not one that rational agents would use. Such a critique of starting points, however, would be more like the classical approach to game theory where one uses reasoning about the game to generate one's probabilities about what the other player will end up doing rather than starting out with an initial assignment of such probabilities as given.

4. Information Feedback

In the book Skyrms explores interesting generalizations of Good's theorem that Bayesian deliberators will evaluate an opportunity to learn before choosing as optimal so long as the learning is cost free.[14] This helps motivate dynamic deliberation so long as later stages of deliberation can be cost free sources of information. Skyrms also provides an interesting chapter on dynamic coherence which ends with the following passage:

> The question of how learning should be modeled in dynamic deliberation is a delicate question, because it is computation, not perception, that generates the new information. This is the realm of what I. J. Good (1950,1983) called "dynamic probability". The evidence could be modeled as data for conditioning in a very big space, but this does not square well with the emphasis of the approach on real-time computation and procedural rationality. It is possible—on the basis of such considerations—to doubt whether any sort of Bayesian approach to these situations is possible. Such worries should be eased by coherence results for quite general black-box learning models. The qualitative features of generalized coherent updating can serve as a touchstone for dynamic deliberation, without imposing unrealistic excess structure on the model. (1990 p. 125)

Let us explore the information feedback in the examples we have been considering.

The immediate impact of an agents calculation at a stage of deliberation governed by the Darwin (or the Nash) rule is to change the assignments of probability to her own acts. This suggests that the information about the other players acts to be fed back ought to fit Jeffrey's rule for generalizing conditionalization where the input is the change in the probability assignments over the partition of the agent's own acts.[15] Consider steps 2 and 3 of the dynamic deliberation in example 1. The output of stage 2 was to assign .545D and .692b. If the assignment to b is to agree with what player 1 could learn about player 2 by using Jeffrey's rule to update on the input consisting of the new assignments to 1's own acts C and D then we have

$$.692 = .455(y) + .545(x),$$

where y and x are respectively P(b|C) and P(b|D), which according to Jeffrey's rule should remain fixed.[16] From stage 3 we have

$$.782 = .376(y) + .624(x).$$

On the assumption that y and x do not change these give two equations for x which together yield x=1.2 which is impossible.

The Darwin dynamic with feedback by emulation does not satisfy the requirements for Jeffrey's generalization when the origin of the shift is to be the assignments to the partition of the agent's pure acts. Skyrms suggests that the coherence results for black box learning models can ease worries about information feedback in dynamical deliberation. In order to apply this suggestion to the stages of Darwin deliberation with feedback by emulation it would appear that at each stage the output probabilities about the other player's acts have to be treated as the appropriate result for a black box that takes the new assignments to the agents own acts as input. Such more fine grained specification as is suggested by Jeffrey's rule does not seem possible.[17]

Skyrms has begun to study dynamic deliberation based on rules where the agents treat each round of deliberation as an inductive trial(Skyrms 1991). Here is his characterization of how the players update their starting points.

> 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.(p227)

Here is his characterisation of updating by Carnap's rule C.

> Suppose that Column has m possible acts, $A_1,...,A_m$. After N trials in which act A_i has been chosen n times,
>
> $$pr(A_i) = (n + 1)/(N + m)$$
>
> This natural generalization of Laplace's rule sets the acts as equiprobable on no evidence and updates using a simple frequency count.(p229)

He proposes to award ties with fractional successes proportional to the tied strategies current probabilities and summing to one. (p227note3 , p229note4)

Let us apply this to the our game from section 3. The equiprobable starting point is the same one we used to try out the Darwin dynamic. A wins on the first trial for player 1 and l ties with r for player 2. This inductive dynamic leads right to the same bad equilibrium (A , (.5l , .5r)) that the Nash and Darwin dynamics got to from the same starting point.

Skyrms suggests that one may want to consider decision-makers who do not start with such symmetry based priors(p230). The most general version of such inductive models he discusses is based Dirichlet priors.

If there are m possible outcomes then the Dirichlet distribution is characterized by positive parameters $\alpha_1,...,\alpha_m$. Take $\beta = \alpha_1 + ... + \alpha_m$. Then if n is the number of occurrences of A_i in N trials we get the general rule

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

If the α_i are all unity we get (*C*); if they are all equal we get Carnap's continuum (*C C*). (p231)

To get to the rational equilibrium (*L* , *l*) in this game with an inductive rule we need a starting point that makes P(L)>P(*R*) or that makes P(l) > .6. Even with inductive deliberation the dynamical approach needs to be supplemented with some way of putting appropriate constraints of what can count as reasonable initial assignments of probabilities to what rational agents can be expected to do in this game.

Skyrms offers some general results about inductive deliberation that could be exploited to suggest a fairly interesting method for two person normal form games. An accessible point is one that can be reached as a limit by inductive players starting from some completely mixed initial state.

(1) Accessible points are Nash equilibria.(p235)

This suggests that an agent who is in a game with such a deliberator can ignore non-Nash strategies.[18] On the left is the Kreps Wilson game in normal form.[19] On the right is the result of deleting pure strategy *R* which is not part of any equilibrium strategy .

	l	*r*
A	(1,1)	(1,1)
L	(3,1)	(-2,0)
R	(2,0)	(-1,1)

	l	*r*
A	(1,1)	(1,1)
L	(3,1)	(-2,0)

Note that once *R* is deleted, *l* weakly dominates *r*. Skyrms gives another general result which can be applied to the reduced game. An admissible act is one which is not weakly dominated.

(2) Accessible points give probability only to admissible acts.(p. 236)

This supports dropping *r*, which leaves the good equilibrium (*L*,*l*) as the only reasonable solution.[20]

5. Evolutionary Game Theory

Skyrms' results about chaos in even quite simple evolutionary games are relevant to the recent trend ,among some economists,to treat equilibria in games as relatively stable configurations reached by processes more akin to evolution than to deliberation by rational agents. One response would be to take up, for games modelling interactions among human decision makers, the lessons Skyrms draws for evolutionary games. These would include embracing concepts from Chaos Theory such as that of strange attractor as of potentially more interest than more traditional game theoretic concepts such as that of equilibrium. It may,however, be instructive to investigate the extent to which the chaotic results are due to the extremely limited rationality of the dynamical deliberations at work.

https://doi.org/10.1017/S0270864700009437 Published online by Cambridge University Press

A beginning of such investigations is to note that when the Carnap inductive dynamic is applied to the game of Skyrms' example 7 with the parameter *a* set at 5, the value at which chaos starts to show up with the Taylor - Jonker dynamics from evolutionary game theory, the inductive dynamic goes right to an equilibrium which is completely missed by the evolutionary dynamic.[21] Even this relatively simple minded inductive rule is enough to avoid the Chaos faced by the evolutionary dynamic. This suggests that more work is needed before the significance of results about evolutionary games for games modelling interactions among human agents will become clear.

Notes

[1]Skyrms expresses some discomfort with the classical interpretation of mixed acts in game theory, which requires that an agent turn her choice over to a random device(1991 p225). One feature he has found attractive about dynamics based on inductive learning rules is that mixed strategy equilibria can be quite naturally interpreted as equilibria of belief along the lines suggested in Aumann 1987.

[2]Skyrms also has interesting ideas about how dynamic deliberation can illuminate what happens when the classical common knowledge constraints are weakened in repeated games. His discussion of good habits contributes to the growing literature on conditions under which rational agents can cooperate in iterated prisoner's dilemmas.

[3]See Skyrms 1990 p. 30,31 for the Nash dynamic, p. 37,38 for the Darwin dynamic.

[4]Skyrms(1990 pp. 36,37)argues that the Darwin dynamic is more consistent with Bayesian updating because it does not change assignments of probability 1 or 0.

[5]This is a version of a game motivated by Puccini's opera *Gianni Schicchi* in Harper 1991 pp. 268,269. The Darwin dynamic requires non-negative utilities. Skyrms suggests that one can normalize to a scale such as [0,1]in transforming games with other utilities. The renormalization of the version of this game in Harper 1991 would have utilities (.75,1), (0,0), and (1,.75) instead of the (.5,1), (0,0) and (1,.5) used here.

[6]This also happens with the utilities (.75,1), (0,0) and (1,.75) that correspond to the version of the game in Harper 1991.

[7]Once the probability of *b* has reached 1 , player 1 can use the general result that Darwin dynamics will not change assignments of 1 or 0 to know that later stages of deliberation will not provide any information about what player 2 may be expected to do.

[8]This is a game from Kreps and Wilson 1982 that was discussed in Harper 1991. The utilities have been transformed to the interval [0,1] in accordance with Skyrms' recommentation for applying the Darwin dynamic.

[9]For a more detailed account of this type of analysis of the game, see Harper 1991. See Harper 1989 for an account of how this test explicates Von Neumann and Morgenstern's indirect argument, for references to Von Neumann and Morgenstern , and for a defense of their conclusion that solutions should be equilibria.

[10]Richard Jeffrey (1983) introduced the idea of ratifiability in one of his attempts

to make evidential decision theory yield the same recommendations as causal decision theory. See Harper (1986) for an account of some of the advantages of formulating ratifiability with the resources made possible by causal decision theory.

[11]Causal decision theory allows the following definition.

$U_X(X') = \Sigma_i P((X' \Box\!\!\rightarrow B_i) | X) U(X', B_i)$, where the B_i's are the outcome determining events , $(X' \Box\!\!\rightarrow B_i)$ is the subjunctive conditional

If I were to do X' it would be the case that B_i,

and $U(X',B_i)$ is the utility of doing X' when event B_i obtains. Game trees provide information that fixes the evaluation of these subjuctive conditionals (Harper 1991).

[12]Any mixed strategy which assigns positive probability to *R* will also fail to be best reply ratifiable. This is shown by a slight variation on a calculation given on p..287 of Harper 1991.

[13]The Nash dynamic also goes to the wrong equilibrium $(A,(.5l,.5r))$ from this starting point.

[14]1990 pp. 87-106. This chapter also reports Skyrm's discovery of an anticipation of Good's theorem in a manuscript of Frank Ramsey's.

[15]Skyrms has provided dynamic Dutch book arguments to defend Jeffrey's rule when the input from a learning experience is correctly represented by changing the probability assignments over a given partition (pp. 119-121,and Skyrms 1987)

[16]Jeffrey's rule assumes that an agents conditional beliefs on elements of a partition do not change when the new assignments over that partition count as the origin for the shift. Where the origin of a shift is the probabilities assigned to a partition $A_1,..,A_n$ the new probability to be assigned to a proposition X is

$$P_1(X) = \Sigma_i P_1(A_i) P_0(X|A_i),$$

where P_0 is the agent's old conditional probability. Here is the result of applying Jeffrey's rule to represent player 1's stage 3 assignment to player 2's playing *b* as the out come of a Jeffrey rule shift where the input is the new assignments to player 1's own pure strategies *C* and *D* generated by the Darwin dynamic calculation at that stage.

$$P_3(b) = P_3(C)P_2(b|C) + P_3(D)P_2(b|D).$$

The P_3's are player 1's beliefs after the stage three calculation and the P_2 are the conditional beliefs before that stage.

[17]This suggests that for such dynamic deliberation the new probability assigned to your acts is treated as a single default mixed strategy that you are considering directly, rather than as a new assignment over the partition of your pure acts.

[18]In games where the only strategic information is that provided by the normal form, the strategies which survive the best reply ratifiability version of the vonNeumann Morgenstern test are exactly the strategies which are part of some equilibrium (Harper1991, p. 268)

[19]Here we are using the original Kreps Wilson utilities. The extensive form version we used above had the utilities transformed to the interval [0,1] in order to satisfy Skyrms' recommendation for applying the Darwin dynamic.

[20]This procedure of first dropping all non-Nash strategies and then deleting weakly dominated strategies was discussed in Harper1991. In more complicated games the deletion of weakly dominated strategies is done in stages where at each stage all weakly dominated strategies are deleted. This avoids the notorious problem of order dependence when the weakly dominated strategies are deleted one at a time. Somewhat surprisingly it turns out that this procedure for normal form analysis, which gave the correct result in the normal form game used by Myerson to motivate proper equilibria , also gives the correct result in all the other games treated in Harper1991 including the ones where the analysis given turned on the extensive form representation.

[21]Skyrms was kind enough to run the Carnap dynamic on this game with his computer when I asked him what the inductive dynamics would do in games where he had produced chaos with the Taylor - Jonker dynamic.

References

Aumann, R. (1987), " Correlated Equilibrium as an Expression of Bayesian Rationality", *Econometrica* 55, 1-18.

Bacharach, M. and Hurley, S. (eds.) (1991), *Foundations of Decision Theory. Oxford:Basil Blackwell.*

Harper, W. (1986), "Ratifiability and Causal Decision Theory: Comments on Eells and Seidenfeld,"in *PSA 1984* Volume 2.

Harper, W. (1989), "Decisions, Games and Equilibrium Solutons," pp. 344-362 in *PSA 1988* Volume2.

_ _ _ _ _ . (1991), "Ratifiability and RefineMents(in Two-person Noncooperative Games)"pp. 263-293 in Bacharach and Hurley (eds.).

Jeffrey, R. (1983), *The Logic of Decision*, 2nd ed. Chicago: University of Chicago Press.

Nash,John (1951), "Non-cooperative Games," *Annals of Mathematics* 54, 286-95.

Selten, R. (1965), "*Spieltheoretische Behandlung eines Oligopolmodells mit Naschfragetragheit*," *Zeitschrift fur die gesamkte Staatswissenschft* 121, 301-24, 667-89.

Skyrms, B. (1982), "Causal Decision Theory", *Journal of Philosophy 79,695 - 711.*

_ _ _ _ _ _. (1987), "Dynamic Cohrence and Probability Kinematics."*Philosophy of Science 54:1-20.*

_ _ _ _ _ _. (1990), *The Dynamics of Rational Deliberation*. Cambridge, MA:Harvard University Press.

______. (1991), "Inductive Deliberation, Admissible Acts, and Perfect Equilibrium", pp. 220-241 in Bacharach, Michael and Hurley, Susan eds.

______. (1993), "Chaos and the Explanatory Significance of Equilibrium: Strange Attractors in Evolutionary Game Theory," This volume.

https://doi.org/10.1017/S0270864700009437 Published online by Cambridge University Press