Philosophical Studies manuscript No.
(will be inserted by the editor)

Local Interactions and the Dynamics of Rational Deliberation

J. McKenzie Alexander

Version submitted 3 September 2008

Abstract Whereas *The Stag Hunt and the Evolution of Social Structure* supplements *Evolution of the Social Contract* by revisiting some of the earlier work's strategic problems in a local interaction setting, no equivalent supplement exists for *The Dynamics of Rational Deliberation*. In this article, I develop a general framework for modeling the dynamics of rational deliberation in a local interaction setting. In doing so, I show that when local interactions are permitted, three interesting phenomena occur: (a) the attracting deliberative equilibria may fail to agree with any of the Nash equilibria of the underlying game, (b) deliberative dynamics which converged to the *same* deliberative outcome in *The Dynamics of Rational Deliberation* may lead to very different deliberative outcomes, and (c) Bayesian deliberation seems to be more likely to avoid nonstandard deliberative outcomes, contrary to the result reported in *The Dynamics of Rational Deliberation*, which argued in favour of the Brown-von Neumann-Nash dynamics.

Keywords Evolution, Rationality, Bayesianism, Brown-von Neumann-Nash dynamics, Social Network

1. Introduction

There are many models of rational deliberation. The best known model of the rational deliberator is that of *homo economicus*, the perfectly rational agent beloved by economists, who makes choices which explicitly maximise his expected utility given the information available. Another model of rational deliberation is that of the *virtually* rational agent (Pettit, 1995), who only explicitly acts to maximise his expected utility when the anticipated loss

Many thanks to Brian Skyrms, Brad Armendt, and Kevin Zollman for their comments on an earlier draft of this paper. I would also like to thank the LSE Choice Group, who heard a talk based on the formal model used in this paper, for many useful suggestions.

J. McKenzie Alexander
Department of Philosophy, Logic and Scientific Method
London School of Economics and Political Science
Houghton Street, London, WC2A 2AE, United Kingdom
Tel: +44 (0)207 955 6496
E-mail: jalex@lse.ac.uk

dips below a previously identified toleration threshold. If this does not happen, the virtually rational agent will carry on doing whatever it was he had been doing, where this is compatible with making choices according to criteria that have nothing to do with maximising behavior. Then there are all the models of *boundedly* rational deliberation (see, for example, Simon, 1957; Gigerenzer et al., 1999) where individuals make choices using heuristics. The justification for employing heuristics is that they take advantage of structural features of the choice problem to deliver approximately optimal results, while imposing a lesser cognitive burden upon the individual. Lastly, the work of Kahneman et al. (1982) makes one wonder whether people, really, achieve anything even remotely close to an approximately optimal outcome from their own point of view.

In *The Dynamics of Rational Deliberation*, Skyrms introduced a model of rational deliberation which falls into the middle of this hierarchy. I locate it there because, although individuals are not conceived of as an explicit maximisers, they modify their beliefs according to dynamical rules which "seek the good", where this means

> 1. the rule raises the probability of an act only if that act has utility greater than the status quo;
> 2. the rule raises the some of the probabilities of all acts with utility greater than that of the status quo (if any).
>
> (Skyrms, 1990, pg. 30)

Bayes's Rule seeks the good, in the above sense, as does the function used by Nash (1951) to prove his famous existence result regarding equilibria in games. Dynamical rules which seek the good sit above boundedly rational heuristics in the heirarchy because, whereas heuristics only guarantee approximately optimal results in suitable environments, dynamical rules that seek the good lead to an equilibrium which maximises expected utility.

Let us investigate the Skyrmsian dynamics of rational deliberation more closely. The precise scenario he concerns himself with is that where "two (or more) [. . .] deliberators are deliberating about what action to take in a noncooperative non-zero-sum matrix game" (Skyrms, 1990, pg. 32). An *action*, here, corresponds to the strategy that a player will adopt in the game. Because Skyrms wishes to model the deliberative process by which each player comes to choose an action, he assumes that each player initially exists in a state of indecision over what to do (otherwise no deliberative problem exists). A state of indecision is represented by a probability distribution over the set of possible actions.

As Skyrms noted, if we assume that the states of indecision are common knowledge, an iterative revision process becomes possible. A player, given his own state of indecision and his knowledge of the states of indecision of his opponents, can calculate, using his specified dynamical rule, what incremental revision to his own state of indecision would maximally increase his expected utility. (And, of course, given what he knows about his opponents, he can also calculate how *they* will incrementally revise their states of indecision as well.) Moreover, each of his fellow opponents also performs the same process themselves. This leads to a new state of indecision for each of the players, where these new states of indecision are also common knowledge, and the process begins again.

One should bear in mind an importance difference between these models and the more recent work of Skyrms. Both *Evolution of the Social Contract* and *The Stag Hunt and the Evolution of Social Structure* develop evolutionary models which are most naturally understood as models of phenotypic change in a population. In these more recent works, a population of individuals faced with a strategic problem (*e.g.*, divide-the-cake, the ultimatum game, or the stag hunt), experience differential rates of "success" in the strategic problem based upon the strategy they employ. These differential rates of success translate into phenotypic change between an earlier moment of time and a later moment of time.[1] However, in *The Dynamics of Rational Deliberation*, the concern lies with how two (or more) players involved in a strategic problem will modify their *internal* states of indecision (understood as a probability distribution over possible actions) under conditions where suitably strong common knowledge assumptions obtain.

Indeed, much of *The Dynamics of Rational Deliberation* concerns itself with investigating questions in such a framework. What can one say about the existence of equilibria? What about stability? And how do the results of these investigations point towards a theory of rational deliberation?

Yet a serious omission remains. Although Skyrms considers cases where two (or more) rational deliberators play a game, these games occur in the absence of a structured social context. The game played by the two deliberators is a deliberative island, the outcome of which fails to carry over and influence the outcomes of other processes of rational deliberation. In part, this omission affects all of rational choice theory, traditionally understood. The classic theory of games, and even social choice theory, typically assume the choices of persons are independent and occur in an unstructured environment. But this does not seem right, for surely any observations I can make about how you revise your beliefs when interacting with another give me some insight into how you are likely to revise your beliefs when you interact with me.[2] Moreover, when agents obtain information which they then use to change their beliefs, the specific source of that information, and the channels through which information flows, may matter as well.

Taking all of these things into account requires consideration of the socially networked nature of society, and requires one to investigate the effect of local interactions upon rational deliberation. However, doing so reveals some important differences between the social network case and the ordinary two-person case. In what follows, I will argue for three points:

1. Allowing for local interactions in the dynamics of rational deliberation breaks the link between convergent points of the deliberative dynamics and Nash equilibrium points of the underlying game.
2. Whereas Skyrms (1990) observes that "All dynamical rules which seek the good have the same fixed points: those states in which the expected utility of the status quo is maximal", this fails to be true, in general, in the local interaction framework.
3. Whereas *The Dynamics of Rational Deliberation* identified some reasons for preferring Nash dynamics over the Bayesian dynamics (for example, "[Bayesian] deliberation can lead to an equilibrium that is not only improper but also imperfect"), the effect of local interactions reveal reasons for preferring the Bayesian dynamics over the Nash dynamics, in some cases.

2. A local interaction model of rational deliberation.

In this section, I present a general framework for modeling the dynamics of rational deliberation on social networks. I then consider the functional form of the dynamical rule used by

[1] I am deliberately being ambiguous as to the nature of the phenotypic change. Most of the models used in Skyrms (1996, 2003) can be understood as models of either biological evolution or cultural evolution, and so the phenotypic change may occur as a result of strategic learning or as a result of differential rates of reproduction.

[2] Of course, there are a variety of reasons why this type of inference might go awry. You might believe that the person you are currently interacting with is out to take advantage of you, but you believe that I will not do so. My point is that, in the absence of reasons of this kind, the inference is plausible. At least to some extent.

people to modify their state of indecision, arguing that considerations of character stability and bounded rationality suggest a linear pooling model. The section ends with some further assumptions which shall be made for sake of simplicity in subsequent sections.

Let $G = (V, E)$ be a graph, where $V = \{A_0, \ldots, A_n\}$ is the set of agents (also called the population) and E the set of edges connecting the agents. For simplicity, we assume that G is a connected graph.[3] Each edge represents a local interaction between A_0 and his neighbors, and corresponds to a pairing for which a game will be played. For example, in figure 1, the agent A_0 is connected to three other agents and hence will play three games during each iteration.

Suppose we have a two-player nonzero sum game of N strategies with a common payoff structure for all members of the population and assume that this payoff structure is common knowledge. Following the general approach of *The Dynamics of Rational Deliberation*, also assume that at a given time t, each player A_i exists in a certain state of indecision regarding what to do. This state of indecision corresponds to a probability distribution $\mathbf{p}_i(t) = \langle p_i^1(t), \ldots, p_i^N(t) \rangle$ over the various actions $\{1, \ldots, N\}$ available to the agent at t. If two players A_i and A_j are connected by an edge, $\mathbf{p}_i(t)$ and $\mathbf{p}_j(t)$ are common knowledge between A_i and A_j.

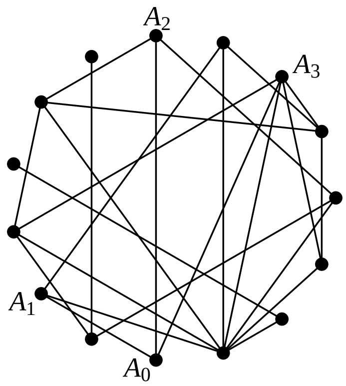

Fig. 1 A simple social network.

This scenario already changes significantly the deliberative framework from that of *The Dynamics of Rational Deliberation*. There the set of two (or more) players all participate in the same *common* game, where each player's action was instrumental in determining the payoff to all. Here, though, we have multiple games being played side-by-side, and the deliberative outcome between two players may have some, little, or no influence whatsoever on the deliberative outcome of the rest of society. The exact nature of the influence depends upon the specification of the dynamics.

[3] Informally, a connected graph has the property that, for any two vertices, there exist a path which links them. For directed graphs, where each edge (a_i, a_j) is understood as "pointing" from a_i to a_j, one can distinguish two senses of connectedness. A strongly connected graph is a directed graph such that, for any two vertices, there exists a path which links them (where walking the path respects the directionality of the edges). A weakly connected graph is a graph where, although any two vertices are linked by a path, one may need to disregard the directionality of the edges when walking the path. Clearly all strongly connected graphs are weakly connected, but not vice versa. Since I will ultimately consider directed graphs, I must note that by "connected", I mean "weakly connected".

What are natural dynamics to use for these local interaction models of rational deliberation? To answer this, consider the functional form of the dynamical rule R used by a player to adjust her state of indecision. A particular player, say A_i, is connected to a number of other players as specified by the graph. Call these other players the *neighbors* of A_i, denoted by $\eta_i = \{i_1, \ldots, i_k\}$. Each one of A_i's neighbors has her own state of indecision over actions, $\mathbf{p}_{i_1}(t), \ldots, \mathbf{p}_{i_k}(t)$, and it is assumed that A_i knows these states of indecision. Furthermore, each of $A_{i_1}, \ldots, A_{i_k}$ known A_i's state of indecision.[4]

The dynamical rule R_i used by A_i specifies how A_i changes her state of indecision given the state of indecision of all of her neighbors. In short:

$$\mathbf{p}_i(t+1) = R_i\big(\mathbf{p}_i(t); \mathbf{p}_{i_1}(t), \ldots, \mathbf{p}_{i_k}(t)\big).$$

In *The Dynamics of Rational Deliberation*, Skyrms primarily considered the case where a player has exactly one neighbor. In that case, the dynamical rule would be of the form $\mathbf{p}_i(t+1) = R_i\big(\mathbf{p}_i(t); \mathbf{p}_{i_1}(t)\big)$

Note that we assume each player has a *single* state of indecision which governs her interactions with all of her neighbors. This may strike some as controversial. Some may think it more natural to allow players to conditionalise their act upon the identity of the neighbor they are playing against. This makes sense if the identity of one's neighbor is known in advance, and a player can take that information into account when choosing. However, this is not always possible. The model proposed here can be thought of one where individuals "play the field" determined by their neighbors.[5]

There are other reasons to think that people do not always conditionalise their action upon the person they are interacting with. For one, this imposes a greater memory requirement, as a person must remember her conditional states of indecisions for every one of her neighbors. Although this is a trivial memory requirement, in the grand scheme of things it will eventually become burdensome: given the number of interactions we have in our daily lives, and the different types of games we play, even someone without a vast number of friends and acquaintences would accumulate well over a hundred conditional states of indecision to remember.[6]

Secondly, rampant conditionalisation seems a phenomenologically false description of how we deliberate about our social interactions. I do not consciously formulate a different state of indecision when I begin to interact with my plumber than when I interact with my baker, butcher, or bourgeois neighbor. Perhaps it would be more rational if I did, but I don't, at least not consciously, and it would be a very odd thing, indeed, if I did so unconsciously!

Lastly, and most importantly, conditionalisation of our states of indecision seems to run afoul of the perceived constancy of our character. Mill refers to this fact in the following well-known passage from *A System of Logic*: "[G]iven the motives which are present to an individual's mind, and given likewise the character and disposition of the individual, the

[4] Note, though, that A_i's state of indecision is *not* necessarily common knowledge amongst *all* of his neighbors. For example, A_{i_1} need not know that A_{i_2} knows that A_i's state of indecision is $\mathbf{p}_i(t)$. Common knowledge exists for the states of indecision of two players joined by an edge.

[5] Furthermore, note that if A_i were capable of conditionalising her state of indecision upon the person she was interacting with, then we would simply have a case where the pairwise dynamics investigated in *The Dynamics of Rational Deliberation* apply to the game played along each edge in G, modifying the conditionalised states of indecision.

[6] Why so many? Because the states of indecision will depend upon the nature of the *game* one is playing. Even if you only tend to interact with a few people, say ten, on a regular basis, you will play a number of different games with each person; each of these games will likely require a different response, and so you will have a different state of indecision for each deliberative problem, even though each deliberate problem is against the *same* person.

manner in which he will act might be unerringly inferred." That is, a person's character provides a basis or foundation, one reasonably stable and enduring over time, upon which we can predict future behavior, even in novel circumstances.[7]

Note that this is not to say that we don't conditionalise our actions at least some times and in some cases. Even the most trusting of individuals would be well advised to look over his shoulder when walking through Southwark late at night. And we do certainly assign a different default degree of belief to any statement uttered by a used car salesman or an insurance broker than statements uttered by our friends. My point is simply that these conditionalised states of indecision mark important deviations from our general state which, in some rough sense, we might call our character.

Given these considerations, I propose that the dynamical rule R_i which A_i uses to update his state of indecision takes on the following form: A_i first calculates, for each pairwise interaction, what his new state of indecision *would* be if that particular pairwise interaction was the only interaction he had. Let us call these the *notional* pairwise refinements of a player's distribution. Denote these notional pairwise refinements for player A_i by $\mathbf{p}'_{i,i_1}(t+1),\ldots,\mathbf{p}'_{i,i_k}(t+1)$, where $\mathbf{p}'_{a,b}(t+1)$ represents the incremental refinement of player a's state of indecision given his knowledge about player b's state of indecision. Player A_i then *pools* these notional pairwise refinements to determine his new state of indecision by forming the weighted average:

$$\mathbf{p}_i(t+1) = \sum_{j=1}^{k} w_{ii_j} \cdot \mathbf{p}'_{i,i_j}(t+1).$$

The weights w_{ii_j} are exogenous parameters given by the edges of the social network, and provide a measure of "social importance" for that particular interaction. In the case where all interactions are taken to be equally important, then the weights are all the same for player i.[8]

This is a linear pooling method for probability aggregation. The reason why it is reasonable to suppose that individuals use such a method is that it is the only aggregation method which satifies the following constraints (see Lehrer and Wagner, 1981):

1. The aggregate probability player i assigns to strategy j at $t+1$ depends only upon the probability player i assigns to strategy j in all of his notional pairwise refinements at $t+1$.
2. If player i assigns zero probability to strategy j in each of his notional pairwise refinements at $t+1$, then player i assigns zero probability to strategy j in the aggregate.

Both of these are plausible when the aggregate probability is viewed as a person's general tendency to act a certain way in his dealings with other people.

What dynamical rule will players use when calculating the notional pairwise refinements of their state of indecision? In *The Dynamics of Rational Deliberation*, Skyrms considered two types of dynamical rules: the first derived from the mapping used by Nash to prove his famous existence theorem (which had been studied earlier by Brown and von Neumann

[7] Yet, having said that, I must acknowledge that the immediately following sentence in *A System of Logic* challenges this interpretation. Mill subsequently writes "if we knew the person thoroughly, and knew all the inducements which are acting upon him, we could foretell his conduct with as much certainty as we can predict any physical event." The fact that Mill includes the additional qualifier "and knew all the inducements which are acting upon him" makes his claim ambiguous. Since each interaction counts as a different "inducement" acting upon a person, it would be consistent with what Mill writes to say that a person could, in principle, act differently whenever he interacts with a different person.

[8] This does not mean that the weights will be the same *across* players, though, since there is an implicit dependence upon the number of edges a player is incident upon. If A_i has three neighbors and A_j has five, then $w_{ii_1} = \cdots = w_{ii_3} = \frac{1}{3}$, but $w_{jj_1} = \cdots = w_{jj_5} = \frac{1}{5}$.

		Them	
		Don't Swerve	Swerve
You	Don't Swerve	$(-10,-10)$	$(5,-5)$
	Swerve	$(-5,5)$	$(0,0)$

Fig. 2 The game of Chicken.

(1950)); the second derived from Bayes' Rule.[9] We shall consider what happens under both of these rules.

The above discussion obscures one technical issue which needs to be made explicit. Skyrms (1990) assumed that players were assigned the role of Row or Column as a permanent feature of the deliberative problem. When we consider the case of local interactions, more than one way exists for handling the assignment of roles to players, and none suggests itself as necessarily right. In a social network like that of figure 1, when should A_0 be considered as playing Row and when should A_0 play Column? One could require the underlying graph to be directed, in which case the presence of the edge (A_i, A_j) would assign A_i the role of Row and A_j the role of Column. Alternatively, we might think of the deliberative process as one where both A_i and A_j are *uncertain* as to whether they will be assigned Row or Column, and they calculate the expected utility of the status quo taking this into account.

In the following, I assume that the underlying graph is directed. The presence of an edge $A \to B$ will mean that, insofar as the pairwise interaction of A and B is concerned, A always plays as Row and B always plays as Column.[10] Finally, I also assume that a player uses equal weights when he pools his notional pairwise refinements, and that the interactions of Row and Column are symmetric with respect to the available strategies (although not necessarily the payoffs).

3. Chicken

Consider the game of Chicken defined by the payoff matrix of figure 2. In the ordinary two-player setting, both Nash and Bayesian deliberators beginning at symmetric starting points with a sufficiently high index of caution[11] will converge to the mixed-strategy equilibrium; if the Nash deliberators have asymmetric starting points, the population will converge to the equilibrium in pure strategies where the player who initially assigned greater probability to Don't Swerve winds up assigning probability 1 to Don't Swerve. Bayesian deliberators beginning at a completely mixed asymmetric state will also likewise converge to one of the two equilibria in pure strategies (Skyrms, 1990), although the vector field determining their revision trajectories is bent differently from that of Nash deliberators (see figure 3).

What happens in a socially structured context? The simplest interesting case consists of the three-player directed cycle given by $1 \to 2 \to 3 \to 1$. We can investigate this by simulation, starting the population in a variety of randomly chosen initial conditions and then

[9] An explicit definition of both of these dynamical rules can be found in appendix A.

[10] No particular assumption will be made regarding what constitutes a "natural" or "reasonable" assignment of directions to edges. Quite often the same player will play as both Row and Column, although in games with different people. This represents the fact that the role we play in a game is, quite often, a matter of historical accident. Although I play as Row in an interaction with you, now, it could have been the case that I played as Column, with you playing as Row, instead. In this model, the fact that a given player permanently takes on the role of Row, in some interactions, and the role of Column, in other interactions, should be seen as a consequence of these "historical accidents" determined by social and other causal forces which lie outside the scope of the model.

[11] See the definition of the dynamical rules in appendix A.

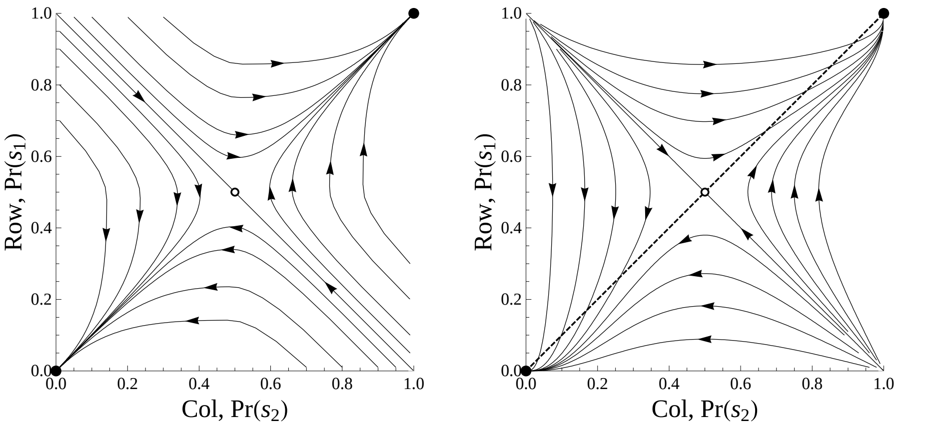

(a) Nash dynamics, $k = 25$ (b) Bayesian dynamics, $k = 100$

Fig. 3 Flow diagrams for two-player Chicken. The x-axis represents the probability assigned to Swerve by Column, the y-axis represents the probability assigned to Don't Swerve by Row. (This allows one to interpret the flow diagram by visualizing the payoff matrix superimposed over the trajectories.) The dotted line in the Bayesian case illustrates the asymmetric bending of the vector field.

stepping the model forward until it converges (if it does). When we do this with Nash deliberators using an index of caution of 1, we find that the group settles into one of two types of attractors. The first type is a cycle of period two, where all players alternate between the distribution $\langle 0.276393, 0.723607\rangle$ and the distribution $\langle 0.723607, 0.276393\rangle$.[12] The second type is a fixed point of the dynamics, where one player adopts the distribution $\langle 0,1\rangle$ and the rest $\langle 0.614508, 0.385492\rangle$ or one player adopts $\langle 1,0\rangle$ and the rest $\langle 0.385492, 0.614508\rangle$. Figure 4 illustrates the second type of attractor. Note that *neither* type corresponds in any way to the standard Nash equilibria for the game of Chicken.

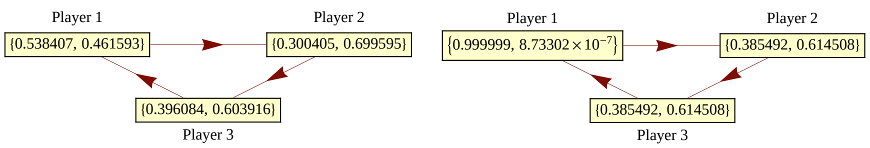

(a) Initial conditions (b) $t = 1{,}000{,}000$

Fig. 4 The game of Chicken played on a three-person directed cycle with Nash deliberators having an index of caution of 1. Probabilities shown as (Don't Swerve, Swerve).

The first type of attractor is not a consequence of the local interaction framework. One can prove that the probability distributions of the first type of equilibrium are, under the Nash dynamics, a cycle of length two in the ordinary two-person case. (See appendix B for a proof.) In our model, since each player revises his distribution by determining the notional revisions for each of his pairwise interaction and then averaging, cycles of length two in the ordinary two-person case will remain cycles of length two in our social network setting.

[12] The probabilities are rounded to six decimal places.

What of the second type of attractor? Here, again, we can prove that they are indeed a fixed point of the dynamics (see appendix C). The important lesson this example demonstrates is that just because the deliberative dynamics converge does not mean that the probability distributions of the players will be a Nash equilibrium of the underlying game.

If we run one thousand simulations on a cycle of length three, using Nash deliberators having an index of caution of 25, we find the state where one person's state of indecision is $\langle 1,0\rangle$ and the other two have states of indecision of $\langle 0.354245, 0.645755\rangle$ occurring 499 times.[13] The converse state, with probabilities of $\langle 0,1\rangle$ and $\langle 0.645755, 0.354245\rangle$, respectively, occurs 477 times.[14]

Very different behavior occurs when we consider Bayesian deliberators. Out of one thousand simulations using Bayesian deliberators having an index of caution of 100,[15] *all* of them converged to a state where one player assigned probability 1 to Don't Swerve and the other two assigned probability 1 to Swerve. It's also very easy to predict what the final convergent state will be: the player who initially assigns the greatest prior probability, however small, to Don't Swerve will, in the limit, assign probability 1 to Don't Swerve. The other two will assign probability 1 to Swerve.

Given our pooling method for how people aggregate probabilities from their pairwise revisions, it seems that the probability distribution a person arrives at will depend crucially on the network topology. What happens on a cycle containing four players? Simulations show that in this case each individual ultimately converges to playing a pure strategy, where the assignment of pure strategies to individuals is such that each pairwise interaction corresponds to one of the ordinary pure strategy Nash equilibria of the game. (See figure 5 for an illustration.) Here, individuals learn to coordinate their behavior so that people alternate between swerving and not swerving in a way which ensures that no collisions occur. This happens regardless of whether people are Nash or Bayesian deliberators.[16]

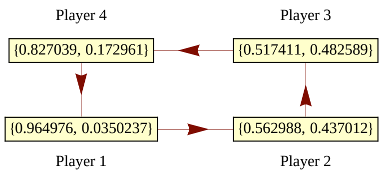

(a) Initial conditions

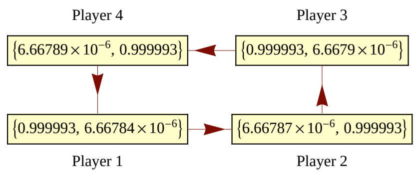

(b) $t = 750{,}000$

Fig. 5 Chicken played on a cycle of length four with Nash deliberators using an index of caution of 25. The probabilities appearing here are not fixed and are continuing to move towards either $\langle 1,0\rangle$ or $\langle 0,1\rangle$.

This result makes sense. Suppose that, in a population of four Nash deliberators on a cycle, one player moves his state of indecision so that it is very close to $\langle 1,0\rangle$. As the two players adjacent to him adjust their states of indecision to shift probability away from Don't

[13] The probabilities existing in the convergent state differ here from the case discussed previously due to the higher value of the index of caution.

[14] The remaining simulations did not converge to six decimal places within 1,000,000 iterations.

[15] The reason for the higher value of the index of caution for Bayesian deliberation is simply due to the fact that the alternate mathematical form requires larger values to slow down the rate at which the state of indecision changes.

[16] Provided that the index of caution is sufficiently high. If $k < \frac{5}{4}$, Nash deliberators may arrive at one of the periodic orbits of the form identified earlier.

Swerve, they won't settle upon the $\langle 0.385492, 0.614508\rangle$ distribution,[17] the fixed point distribution we found earlier, because those two players don't interact with each other. Instead they are connected to a fourth player who now has an incentive to move *his* probability distribution towards $\langle 1,0\rangle$, as that is compatible with his neighbors' gradual adjustment towards $\langle 0,1\rangle$.

Four players connected via a ring move towards this stable convergent state extremely rapidly, even if very little asymmetry exists in their initial distributions. One can easily show that if the state of indecision for all four players is the mixed strategy $\langle \frac{1}{2}, \frac{1}{2}\rangle$, that population state is a fixed point of the dynamics for both Nash and Bayesian deliberators. However, if three players have the state of indecision $\langle \frac{1}{2}, \frac{1}{2}\rangle$ and one player deviates from this by a very small amount, say 10^{-6}, this small asymmetry is enough for both the Nash and the Bayesian deliberational dynamics to converge to the state where, as we go around the ring, players adopt the strategies Swerve, Don't Swerve, Swerve, and Don't Swerve.

Although a ring with an even number of players *may* converge to a state where players alternate between Don't Swerve and Swerve, this needn't happen. Figure 6 illustrates how nonstandard distributions similar to those appearing on a cycle of three players can also appear on a cycle of eight players. The explanation for this can be determined by considering again the final configuration of figure 4: although players two and three both interact with an individual who has adopted the distribution $\langle 1,0\rangle$, it isn't necessary for players two and three to interact with the *same* individual. The nonstandard distribution would equally well be stable if they both happened to interact with the same *type* of individual.

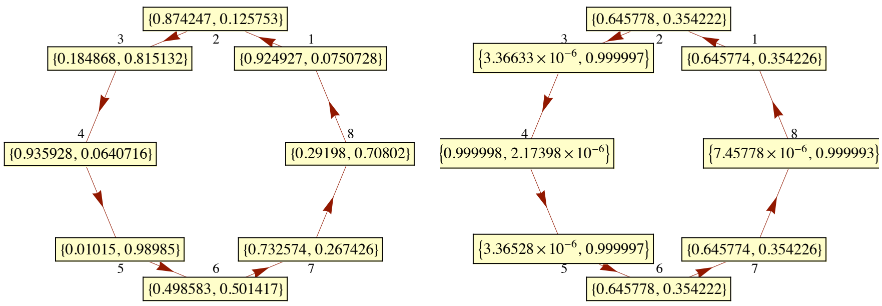

(a) Initial conditions

(b) $t = 2,300,000$

Fig. 6 Chicken played on a cycle of eight with Nash deliberators having an index of caution of 25.

And, indeed, that is what happens in figure 6. Player 1 adopts the nonstandard distribution $\langle 0.614508, 0.385492\rangle$ (as does player 2), interacting with player 8 whose distribution is effectively $\langle 0,1\rangle$. Although player 2 does not interact with 8, he does interact with 3, whose distribution also effectively equals $\langle 0,1\rangle$. The distributions of 3 and 8 are sufficiently close to being the same that the outcome produced is a stable nonstandard deliberational equilibrium.

Given this configuration, we know from our earlier proof that players 1 and 2 will not switch their strategies as long as players 8 and 3 do not. Moreover, we know that players 8 and 3 will not switch strategies because they are both connected to players who assign probability 1 to Don't Swerve. Lastly, notice that the configuration of players 4, 5, 6, and 7

[17] Assuming that the index of caution is 25.

		Billy	
		Boxing	Ballet
Maggie	Boxing	$(2,1)$	$(0,0)$
	Ballet	$(0,0)$	$(1,2)$

Fig. 7 The game of Battle of the Sexes.

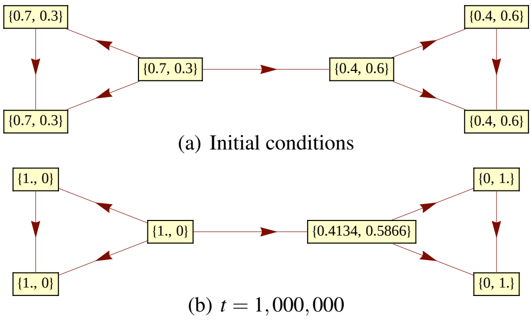

(a) Initial conditions

(b) $t = 1,000,000$

Fig. 8 Battle of the Sexes played by Nash deliberators ($k = 25$) on two cycles connected by a bridge edge (values rounded to the nearest 10^{-4}).

is just the opposite of that of players 1, 2, 3, and 8; hence the overall population state is stable under the Nash dynamics.

4. Battle of the Sexes

Turning now from anti-coordination to coordination games, consider the game of Battle of the Sexes as defined by the payoff matrix in figure 7. Simulations of the deliberational dynamics for both Nash and Bayesian deliberators on a cycle of length three reveal that, on that simple network, the population will coordinate on either All Go Boxing or All Go To The Ballet. It is also straightforward to predict which of these two outcomes will come about: look at the total aggregate probability assigned to Boxing and Ballet in the states of indecision for the population. Whichever activity has more probability assigned to it will be the activity the population converges upon.

The predictive success of that rule depends on the topology, though. Consider a graph defined by the sequence of edges $1 \to 2 \to 3 \to 1$, $4 \to 5 \to 6 \to 4$ and $1 \to 4$ (see figure 8). If each player in the left "lobe" of graph has the state of indecision $\langle .7, .3\rangle$ and each player in the right "lobe" has the state of indecision $\langle .4, .6\rangle$, the total aggregate probability assigned to Boxing in the population is 3.3, with Ballet receiving an aggregate of 2.7. Yet if each player is a Nash deliberator with an index of caution $k = 25$, the left lobe converges to Boxing and most of the right lobe converges to Ballet. This makes sense, given the topology, but it also shows that the predictive rule which works on a simple cycle fails to work here.

It's worth investigating what happens on more realistic and complex social networks. Consider, then, the following sequence of simulations: for one thousand trials, generate a random directed graph g_i consisting of twenty vertices, and a random assignment of states of indecision $\langle \mathbf{p}_{i_1}, \ldots, \mathbf{p}_{i_{20}} \rangle$ to each vertex in g_i.[18] Then, for each of these initial conditions, calculate the state resulting from stepping the model forward $1,000,000$ iterations under the Nash dynamics. Then do the same thing, except under Bayesian dynamics.

[18] The random graphs were generated using the following procedure: each of the 190 possible edges had a 20% chance of being included. If the resulting graph was connected, it was used; if the resulting graph was disconnected, it was thrown out and a new candidate was generated. The same graph was used for the i^{th} trial for both Nash and Bayes deliberators, as well as the same initial conditions.

	Nash	Bayesian
All Go Boxing:	492	482
All Go To Ballet:	507	512
Nonstandard:	1	6

Fig. 9 Comparison of outcomes for Nash and Bayes deliberators on one thousand random networks, beginning with the same initial conditions.

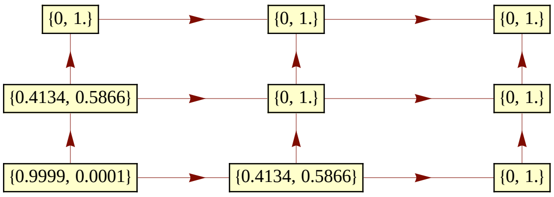

Fig. 10 Nonstandard deliberational outcome for Battle of the Sexes played on a 3×3 lattice with Nash deliberators $k = 25$. The states of indecision here were obtained after $1,000,000$ iterations. (Probabilities rounded to nearest 0.0001.)

	Nash	Bayesian
All Go Boxing:	123	280
All Go To Ballet:	149	276
Nonstandard:	728	444

Fig. 11 Comparison of outcomes for Nash and Bayes deliberators playing Battle of the Sexes on a 3×3 lattice.

Figure 9 tabulates the outcome of this sequence of simulation. Given that figure 8(b) shows that a connected graph may fail to converge to a state where complete agreement exists on the activity, it is somewhat surprising that an entire population of Nash or Bayes deliberators manages to coordinate upon either Boxing or Ballet the vast majority of the time. What the statistics don't reveal, though, is the extent to which the deliberational dynamics can lead to divergent outcomes. When both Nash and Bayes deliberators converged to a state of All Go Boxing or All Go to Ballet, they *disagreed* on the deliberational equilibrium 70 times.

It is difficult to draw a general lesson concerning coordinated action by rational deliberators in a networked environment from the examples discussed so far. We have seen that the deliberational dynamics can converge to a globally coordinated outcome, instead of some nonstandard outcome, on both regular structures (like a cycle of length three) or on random structures (like the one thousand randomly generated graphs). However, there are graphs which generally discourage both Nash or Bayesian deliberators from globally coordinating on an outcome. One example of this is a 3×3 lattice.

Simulations show that, out of one thousand trials run on a 3×3 lattice, the majority evolve to a nonstandard deliberational equilibria consisting of some people who opt for Boxing, some who opt for Ballet, and some whose state of indecision is a mix between the two. Figure 10 illustrates one of the nonstandard deliberational outcomes, and figure 11 lists the frequencies with which Nash and Bayesian deliberators arrived at various outcomes. The point to note is that, although both Nash and Bayesian deliberators are capable of globally coordinating their behavior on *random connected* graphs with considerable frequency, they seem to be *incapable* of globally coordinating their behavior with anything like the same degree of frequency on a social structure like a lattice.

5. Conclusion

I said that I would argue for three claims which reveal how allowing for local interactions in the dynamics of rational deliberation can overturn three claims appearing in *The Dynamics of Rational Deliberation*. Let us now retrace why those three claims are true. Regarding the first, we have seen repeatedly that the convergent state of the deliberational dynamics in the local interaction framework defined here may not agree with any of the traditional Nash equilibrium points of the underlying two-player game.

As for the second claim, we have also seen how, in the game of Battle of the Sexes, a population of Nash deliberators, each of whom began in the *same* state of indecision as their counterpart in a population of Bayesian deliberators might nevertheless arrive at a *different* deliberational outcome. Because these deliberational outcomes can be very different in kind, this demonstrates that it is false that all dynamical rules which seek the good have the same fixed points (which is true in the ordinary two-person case). Another way to see this is to reflect on the difference between the Nash and Bayesian dynamics on a cycle of length three for the game of Chicken: the state $\langle 1,0\rangle$, $\langle 0,1\rangle$ and $\langle 0,1\rangle$ is a fixed point (and an attractor) of the Bayesian dynamics, but a population of Nash deliberators with those states of indecision will quickly evolve to the state indicated in figure 4(b).

Lastly, note that the examples here seem to indicate that Bayesian deliberators are better able to avoid the nonstandard deliberational equilibria which Nash deliberators are prone to. In Chicken on a cycle of length three, Bayesian deliberators evolve to a state where one person plays Don't Swerve and the others Swerve. Nash deliberators, on the other hand, always arrive at a state where collisions are possible. On 3×3 lattices, Bayesian deliberators are more likely to achieve global coordination than Nash deliberators. So, although it is true that Bayesian deliberation "can lead to an equilibrium that is not only improper but also imperfect", Bayesian deliberation also seems more capable of achieving an intuitively reasonable equilibrium in this local interaction setting. The extent to which this is true, though, needs to be examined more carefully.

Clearly more work remains to be done concerning the local interactions of rational deliberation. Some of the results I cite here as the result of simulation should readily admit a rigorous proof. And it would be nice if some general properties could be discovered enabling us to predict which deliberational outcomes will come about given the initial states of indecision and the underlying network structure. Although it is now nearly twenty years since the publication of *The Dynamics of Rational Deliberation*, it is remarkable how much remains to be extracted from it.

A. Definition of the Nash and Bayes dynamics.

Let $M=\langle r_{ij},c_{ij}\rangle|_{i,j=1}^{n}$ be the payoff matrix for a two-player game with n strategies and let $\mathbf{p}_{\mathrm{col}}(t)$ and $\mathbf{p}_{\mathrm{row}}(t)$ be the states of indecision for Row and Column, respectively. In these states of indecision, $p^i_{\mathrm{col}}(t)$ and $p^i_{\mathrm{row}}(t)$ denote the probability assigned by Column and Row to action i in their state of indecision at time t. Thus the expected utility for action i for Row at time t is

$$\underset{\mathrm{row}}{\mathrm{EU}}(i,t)=\sum_{j=1}^{n} r_{ij}\cdot p^j_{\mathrm{col}}(t)$$

(with the expected utility for Column defined *mutatis mutandis*). The expected utility of the *status quo* for Row at time t equals

$$\underset{\mathrm{row}}{\mathrm{ESQ}}(t)=\sum_{i=1}^{n} p^i_{\mathrm{row}}(t)\cdot\underset{\mathrm{row}}{\mathrm{EU}}(i,t).$$

Finally, let the *covetability* of act j for Row at time t be

$$\max\left(0,\underset{\mathrm{row}}{\mathrm{EU}}(i,t)-\underset{\mathrm{row}}{\mathrm{ESQ}}(t)\right).$$

Given this, the Nash dynamics (also known as the Brown-von Neumann-Nash dynamics) states that individuals revise their state of indecision according to the rule

$$p^i(t+1) = \frac{k \cdot p^i(t) + \mathrm{Cov}(i)}{k + \sum_{j=1}^{n} \mathrm{Cov}(j)}$$

where $k > 0$ is an index of caution specifying how quickly an agent revises his or her belief in a single iteration.

The Bayesian dynamics[19] takes a slightly different form:

$$p^i(t+1) = p^i(t) + \frac{1}{k} \cdot p^i(t) \cdot \frac{\mathrm{EU}(i,t) - \mathrm{ESQ}(t)}{\mathrm{ESQ}(t)}$$

where, again, $k > 0$ is an index of caution. (In order to ensure that the Bayes dynamics results in a probability distribution, the payoff matrix needs to be normalized so that the lowest payoff is 0 and the greatest payoff is 1.)

B. Chicken, cycle of period 2.

One can easily prove that the numerically obtained probabilities do form a cycle of period two under the dynamical rules described above, for the case where $k = 1$. Suppose that each player adopts the distribution $\mathbf{v} = \langle v_1, v_2 \rangle$ for the game of Chicken, where we list the strategies in the order Don't Swerve and Swerve. We want to find values of v_1 and v_2 which, under the Nash dynamics, have the property that $\langle v_1', v_2' \rangle = \langle v_2, v_1 \rangle$.

First, note that the expected utility of the status quo for both players is $-10v_1^2$ and the expected utility of the actions Don't Swerve and Swerve, are, respectively, $-10v_1 + 5v_2$ and $-5v_1$. Now, the covetability of each action is simply the difference between the expected utility of the action and the status quo, if this value is greater than zero, and zero otherwise. If we assume that the action Don't Swerve has a nonzero covetability, then this value is $-10v_1 + 10v_1^2 + 5v_2$. If we assume that *only* Don't Swerve has a nonzero covetability, then the new values of v_1 and v_2 under the Nash dynamics are

$$v_1' = \frac{v_1 + (-10v_1 + 10v_1^2 + 5v_2)}{1 + (-10v_1 + 10v_1^2 + 5v_2)}$$

and

$$v_2' = \frac{v_2}{1 + (-10v_1 + 10v_1^2 + 5v_2)}.$$

The system of equations given by $\langle v_1', v_2' \rangle = \langle v_2, v_1 \rangle$ has four solutions, of which only three can be probability distributions over actions.[20] The first solution is the standard mixed strategy Nash equilibrium $\langle v_1, v_2 \rangle = \left\langle \frac{1}{2}, \frac{1}{2} \right\rangle$; however, this solution must be discarded as well as it does not satisfy the assumption that Don't Swerve has a nonzero covetability. The remaining solutions are $\langle v_1, v_2 \rangle = \left\langle \frac{1}{10}(5 - \sqrt{5}), \frac{1}{10}(5 + \sqrt{5}) \right\rangle$ and $\langle v_1, v_2 \rangle = \left\langle \frac{1}{10}(5 + \sqrt{5}), \frac{1}{10}(5 - \sqrt{5}) \right\rangle$. Of these, only the first satisfies the assumption that Don't Swerve has a nonzero covetability and that Swerve has a zero covetability.

All that remains is to verify that $\langle v_2', v_1' \rangle = \langle v_1, v_2 \rangle$. This can be done by straightforward calculation or by writing down a set of equations similar to the above and solving it. Note that the numeric approximations of v_1 and v_2 are 0.276393 and 0.723607, respectively, which agree with the values found by simulation.

If we consider the generalized Nash dynamics with index of caution $k > 0$, where a distribution is revised according to the rule

$$v_i' = \frac{kv_i + \mathrm{Cov}(v_i)}{k + \sum_j \mathrm{Cov}(v_j)},$$

[19] For a discussion of why these dynamics warrant the label "Bayesian", see Skyrms (1990, pp. 36–37, 165–166).

[20] The trivial solution $\langle v_1, v_2 \rangle = \langle 0, 0 \rangle$ has to be discarded for this reason.

a similar analysis shows that cycles of length two exist for $\frac{5}{4} \geq k > 0$. These cycles are given by

$$v_1 = \frac{1}{10}\left(5 - \sqrt{5}\sqrt{5-4k}\right) \qquad v_1' = \frac{1}{10}\left(5 + \sqrt{5}\sqrt{5-4k}\right)$$
$$v_2 = \frac{1}{10}\left(5 + \sqrt{5}\sqrt{5-4k}\right) \qquad v_2' = \frac{1}{10}\left(5 - \sqrt{5}\sqrt{5-4k}\right).$$

C. The analytic solution for the fixed-point in Chicken on a cycle of length three.

Suppose that player one follows the distribution $\langle 1,0\rangle$. We need to show that there exist a distribution $\langle v_1, v_2\rangle$ which, when adopted by players two and three, produces a stable state for all three players given the dynamics. In order for player one to remain at $\langle 1,0\rangle$, it must be the case that the notional revisions obtained for each of his pairwise interactions with two and three do not assign any probability to Swerve. If they did, then the result of averaging the two notional distributions would cause player one to move away from the distribution $\langle 1,0\rangle$.

What this means is that the covetability of Don't Swerve and Swerve, for player one, must be zero when both player two and three adopt $\langle v_1, v_2\rangle$. Now, the expected utility of the status quo is $-10v_1 + 5v_2$ for player one, and the expected utility of the actions Don't Swerve and Swerve are $-10v_1 + 5v_2$ and $-5v_1$, respectively. The covetability of Don't Swerve is thus zero, and the covetability of Swerve will be nonzero exactly when $v_1 > v_2$. Requiring that player one's distribution remains unchanged forces, then, that $v_1 \leq v_2$.[21]

Consider now the notional revision generated for player two in his interaction with player one. (Recall that player one is Row and player two is Column.) The expected utility of the status quo for player two is $-10v_1 - 5v_2$. What are the covetabilities of Don't Swerve and Swerve? The expected utilities of each of these actions, given that player one currently follows $\langle 1,0\rangle$, are -10 and -5, respectively. From this we see that Don't Swerve has a nonzero covetability for player two if and only if $-10 + 10v_1 + 5v_2 > 0$. Given that $v_1 + v_2 = 1$, it is impossible for this inequality to be satisfied and so the covetability of Don't Swerve for player two is zero.

If we assume that Swerve has a nonzero covetability for player two, then his notional revision for this particular interaction will be as follows (we use η^i_{ab} to denote the probability assigned to action i when player a calculates the notional revision for his pairwise interaction with b):

$$\eta^1_{21} = \frac{v_1}{1 + (-5 + 10v_1 + 5v_2)}$$
$$\eta^2_{21} = \frac{v_2 + (-5 + 10v_1 + 5v_2)}{1 + (-5 + 10v_1 + 5v_2)}.$$

From this, it follows that $\eta^1_{21} < v_1$ and $\eta^2_{21} > v_2$.

Now consider what happens when player two interacts with player three. When both players use the distribution $\langle v_1, v_2\rangle$, the expected utility of the status quo for each is $-10v_1^2$. What is the covetability of Don't Swerve and Swerve for player two? Well, in order for $\langle v_1, v_2\rangle$ to be a fixed point under the dynamics, it must be the case that when player two calculates his notional revision for his interaction with player three, he finds that $\eta^1_{23} > v_1$ and $\eta^2_{23} < v_2$, because otherwise it would be impossible for the average of the notional revisions $\left\langle \eta^1_{21}, \eta^2_{21}\right\rangle$ and $\left\langle \eta^1_{23}, \eta^2_{23}\right\rangle$ to equal $\langle v_1, v_2\rangle$. This means that player two must find that the strategy Don't Swerve has a nonzero covetability when he considers his interaction with player three. And, consequently, that the strategy Swerve has a covetability of zero for his interaction with player three. Hence player two will calculate his notional revision in this case as follows:

$$\eta^1_{23} = \frac{v_1 + (-10v_1 + 10v_1^2 + 5v_2)}{1 + (-10v_1 + 10v_1^2 + 5v_2)}$$
$$\eta^2_{23} = \frac{v_2}{1 + (-10v_1 + 10v_1^2 + 5v_2)}.$$

One then simply needs to solve the following three equations:

$$\frac{\eta^1_{21} + \eta^1_{23}}{2} = v_1 \qquad \frac{\eta^1_{21} + \eta^2_{23}}{2} = v_2 \qquad v_1 + v_2 = 1$$

[21] Which, of course, agrees with the simulation results.

and check which solutions satisfy the assumptions made along the way. Of the four solutions, three occur in the complex plane. The remaining solution in real values is approximately $\langle v_1, v_2 \rangle = \langle 0.385492, 0.614508 \rangle$, which agrees with the simulation results. Although of little use, it is worth seeing the complexity of the analytic solution:

$$v_1 = \frac{19}{40} - \frac{1}{40}\sqrt{\frac{1}{3}\left(243 + 20\sqrt[3]{87\left(18-\sqrt{237}\right)} + 20\sqrt[3]{87\left(18+\sqrt{237}\right)}\right)} + \frac{1}{20}\sqrt{\frac{1}{6}\left(243 - 10\sqrt[3]{87\left(18-\sqrt{237}\right)} - 10\sqrt[3]{87\left(18+\sqrt{237}\right)} + \frac{2163}{\sqrt{\frac{1}{3}\left(243+20\sqrt[3]{87\,(18-\sqrt{237})}+20\sqrt[3]{87\,(18+\sqrt{237})}\right)}}\right)}$$

and

$$v_2 = \frac{21}{40} + \frac{1}{40}\sqrt{\frac{1}{3}\left(243 + 20\sqrt[3]{87\left(18-\sqrt{237}\right)} + 20\sqrt[3]{87\left(18+\sqrt{237}\right)}\right)} - \frac{1}{20}\sqrt{\frac{1}{6}\left(243 - 10\sqrt[3]{87\left(18-\sqrt{237}\right)} - 10\sqrt[3]{87\left(18+\sqrt{237}\right)} + \frac{2163}{\sqrt{\frac{1}{3}\left(243+20\sqrt[3]{87\,(18-\sqrt{237})}+20\sqrt[3]{87\,(18+\sqrt{237})}\right)}}\right)}.$$

References

Brown, G. and J. von Neumann (1950). "Solutions of Games by Differential Equations". In *Contributions to the Theory of Games*, Princeton University Press, number 24 in Annals of Mathematical Studies.

Gigerenzer, Gerd, Peter M. Todd, and the ABC Research Group (1999). *Simple Heuristics That Make Us Smart*. Oxford University Press.

Kahneman, D., P. Slovic, and A. Tversky (1982). *Judgement Under Uncertainty: Heuristics and Biases*. Cambridge University Press.

Lehrer, K. and C. Wagner (1981). *Rational Consensus in Science and Society*. Dordrecht:Reidel.

Nash, John (1951). "Non-cooperative games". *Annals of Mathematics* 54, 2: 286–295.

Pettit, Philip (1995). "The Virtual Reality of *homo economicus*". *Monist* 78: 308–329.

Simon, Herbert A. (1957). *Models of Man: Social and Rational*. John Wiley & Sons, Inc.

Skyrms, Brian (1990). *The Dynamics of Rational Deliberation*. Harvard University Press.

— (1996). *Evolution of the Social Contract*. Cambridge University Press.

— (2003). *The Stag Hunt and the Evolution of Social Structure*. Cambridge University Press.