

Inductive Learning, Knowledge Asymmetries and Convention

Peter Vanderschraaf
Department of Philosophy
University of California, Irvine
Irvine, California 92717

[T]his may properly enough be call'd a convention or agreement betwixt us, tho' without the interposition of a promise; since the actions of each of us have a reference to those of the other, and are perform'd upon the supposition, that something is to be perform'd on the other part. Two men, who pull the oars of a boat, do it by an agreement or convention, tho' they have never given promises to each other. . . . In like manner are languages gradually establish'd by human conventions without any promise. In like manner do gold and silver become the common measures of exchange, and are esteem'd sufficient payment for what is of a hundred times their value.

David Hume, *A Treatise of Human Nature*, p. 490.

Introduction

In the classic passage quoted above, Hume observes that people are able to coordinate many of their activities without explicitly agreeing to coordinate, and he resolves this apparent paradox by explaining coordination in terms of individual expectations. If a social coordination problem is “solved” in this way, that is, each agent acts so as to coordinate and expects others to do the same, the agents involved are said to follow a *convention*.

Convention goes far in explaining how persons coordinate many of their activities without continually negotiating over *how* they will coordinate, but the notion raises fundamental questions: (1) *Why do certain particular conventions persist over time?*, and (2) *How does a particular convention arise in the first place?* Beginning with the seminal work of Schelling (1960) and Lewis (1969), a number of authors have addressed these questions by applying the theory of *noncooperative games* developed in von Neumann and Morgenstern's classic *Theory of Games and Economic Behavior*. There are two common threads in the game-theoretic literature on convention. First, a convention is usually defined as a certain type of *equilibrium* of a noncooperative game, which helps account for why conventions remain stable over time. Second, the emergence of particular conventions is often explained as the result of a *dynamical adjustment process*, in which the behavior of agents engaged in a repeated strategic situation converges to an equilibrium corresponding to a convention.

This paper presents a new game theoretic definition of convention, and applies the theory of *inductive deliberation* (Skyrms 1991) to give an account of the emergence of convention. Lewis (1969) gives a widely accepted definition of convention as a *coordination equilibrium* of a noncooperative game that satisfies *common knowledge of a mutual expectations criterion* (MEC) because it is *salient*, that is, the equilibrium is somehow conspicuous to all of the agents involved.¹ I propose an alternate definition of convention as a *correlated equilibrium* (Aumann 1974, 1987) satisfying a *public intentions criterion* (PIC), and argue that this definition is more satisfactory than Lewis' definition.² A convention is defined as a function from a space of “states of the world”, which formalizes salience, to strategy combinations of a noncooperative game which meet

¹Lewis attributes the intuitive notion of salience to Schelling (1960). See Lewis (1969), p. 35 and Schelling (1960), pp. 55-56.

the PIC, so that the system is at equilibrium. I argue that many conventions correspond to correlated equilibria that are not Nash equilibria.

To account for the emergence of correlated equilibria corresponding to conventions, I model agents as a dynamical system of *inductive deliberators*. Choosing which correlated equilibrium of many to follow as a convention is an instance of the more general problem of equilibrium selection in game theory. One way to address this problem is to have the players adjust their beliefs about each other recursively with a common *inductive rule*. This process of *inductive deliberation* (Skyrms 1991, Vanderschraaf and Skyrms 1993) enables players to reach an equilibrium from a state of initial indecision. This paper applies the theory of inductive deliberation specifically to the problem of explaining the emergence of convention. The approach in this paper differs from other dynamical explanations of convention in the literature in that players receive external signals at each stage of deliberation, and they do not assume that strategy choices are stochastically independent.³ Such deliberators can learn to play a correlated equilibrium under the right conditions. In particular, *asymmetries* in the deliberators' knowledge of their situation can lead to their converging to a correlated equilibrium corresponding to a convention.

§1. Convention as Correlated Equilibrium

Lewis (1969) defines a convention as a state in which: (1) agents engaged in a game play a *coordination equilibrium*, and (2) their preferences to conform with this coordination equilibrium are *common knowledge*. A coordination equilibrium of a noncooperative game is a strategy combination such that *all* players would be strictly worse off if any player were to deviate from this equilibrium.⁴ For instance, in the 2-player "Battle of the Sexes" game with payoff structure given in Figure 1, there are two coordination equilibria in pure strategies, namely (A1, A1) and (A2, A2).⁵

Figure 1. Battle of the Sexes

		Kay	
		A1	A2
Amie	A1	$(\sqrt{2}, 1)$	(0, 0)
	A2	(0, 0)	$(1, \sqrt{2})$

This game models a situation in which there are at least two possible conventions, corresponding to the

²Shubik (1982) and Skyrms (1990) suggest that conventions can be defined as correlated equilibria, but neither gives an explicit definition. See Shubik (1982), p. 249, and Skyrms (1990), pp. 52-61.

³For other dynamical explanations of the emergence of conventions, see Crawford and Haller (1990), Young (1993) and Kandori, Mailath and Rob (1993). Note that all of these papers model a convention as a Nash equilibrium, which I argue is too restrictive.

⁴Lewis (1969), p. 14.

⁵This game is called "Battle of the Sexes" because in some of its early interpretations the two players were a wife who wants to attend the opera and a husband who wants to attend a prize fight!

coordination equilibria $(A1, A1)$ and $(A2, A2)$. At either of these equilibria, the desires of the two players are perfectly coordinated, that is, each player has an incentive to play her end of the equilibrium, and she wants her opponent to play her end of the equilibrium as well.

Lewis' definition of convention requires not only that agents play their ends of a coordination equilibrium, but also that the agents have *common knowledge* that they all have a decisive reason to conform with the equilibrium. A proposition p is common knowledge for a set of agents if and only if

- (1) Each agent k knows that p ,
- (2) Each agent i knows that each agent k knows that p , each agent j knows that each agent i knows that each agent k knows that p , and so on.

For Lewis, a coordination equilibrium is a convention only if the players have common knowledge of a *mutual expectations criterion*:

MEC Each agent has a decisive reason to conform to his part of the convention *given that she expects the other agents conform to their parts*.⁶

Lewis adds this common knowledge requirement to his definition of convention because he contends, rightly I think, that agents who follow a convention must know that all are acting in order to achieve coordination. Lewis' common knowledge restriction is meant to rule out cases in which agents coordinate by accident or as the result of false beliefs regarding their opponents. Hence, Lewis defines a convention as follows:

A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if it is true that, and it is common knowledge in P that, in any instance of S among the members of P ,

- (1) everyone conforms to R ;
- (2) everyone expects everyone else to conform to R ;
- (3) everyone prefers to conform to R on condition that the others do, since S is a coordination equilibrium and uniform conformity to R is a coordination equilibrium in S .⁷

Lewis argues that the agents engaged in a coordination problem will conform with a particular coordination equilibrium because this equilibrium is *salient*, that is, the equilibrium somehow stands out so that all expect each other to coordinate on this equilibrium. Lewis points out that salience need not necessarily arise from pre-game communication. A variety of factors, including the description of the game, environmental clues and precedent, can result in a coordination equilibrium being salient. For instance, the coordination equilibrium $(A1, A1)$ in Battle of the Sexes might be salient to the two players, because they have coordinated on $(A1, A1)$ sometime in the past.

I propose an alternate definition of convention that differs from Lewis' definition in two fundamental respects. First, I contend that satisfaction of a *public intentions criterion* (PIC) reflects the nature of a convention better than common knowledge of the MEC. Players will generally have common knowledge of the MEC at any strict equilibrium of a noncooperative game, but only coordination equilibria satisfy the following:

⁶Lewis (1969), p. 25. The term "mutual expectations criterion" is mine, not Lewis'.

⁷Lewis (1969), p. 58.

PIC Each agent will desire that his choice of pure strategy is common knowledge among all agents engaged in the game.

Second, I contend that many conventions do not correspond to Nash equilibria, a fact that Lewis does not discuss in his work, but which is consistent with his general point of view. To illustrate both these points, suppose that in Battle of the Sexes, Amie and Kay correlate their strategies with the toss of a fair coin, so that both play *A1* if the coin lands heads-up (“*H*”) and both play *A2* if the coin lands tails-up (“*T*”). Then the coordinated strategy combinations (*A1,A1*) and (*A2,A2*) are each played with probability $\frac{1}{2}$, and the uncoordinated combinations are played with probability zero. If Amie computes her conditional expected payoffs, she will note that

$$E(u_1(A1) | H) = \sqrt{2} > 0 = E(u_1(A2) | H) , \text{ and}$$

$$E(u_1(A2) | T) = 1 > 0 = E(u_1(A1) | T)$$

so she will not want to defect from the correlated strategy of playing *A1* if *H* and *A2* if *T*. Moreover, Amie will want Kay to know that she intends to follow the correlated strategy so that Kay will want to conform. Likewise, Kay will want to adhere to her end of this correlated strategy combination, and she will want Amie to know this so that Amie will do the same. Hence this strategy combination is a *correlated equilibrium* (Aumann 1974, 1987) which satisfies the PIC. This correlated equilibrium is also *salient* to the players because their expectations are correlated with the same piece of external information, namely the coin toss, and should therefore be regarded as a convention.

To see that common knowledge of the MEC is not by itself sufficient to characterize convention, note that in the “Chicken” game with payoff matrix given in Figure 2, the strict Nash equilibrium (*A1,A2*) can satisfy common knowledge of the MEC, but it fails the PIC.

Figure 2. Chicken

		Kay	
		<i>A1</i>	<i>A2</i>
Amie	<i>A1</i>	(6,6)	(2,7)
	<i>A2</i>	(7,2)	(0,0)

Chicken is a game of *conflicting interests*, since if a player chooses to play *A1*, she would prefer that her opponent sacrifice her own interests by also playing *A1* rather than maximize expected utility by playing *A2*. Suppose that Amie elects to “play it safe” by choosing *A1*. If Kay knows Amie’s choice, then Kay will play the aggressive strategy *A2*, and if their respective strategy choices are common knowledge, then Amie and Kay are at the equilibrium (*A1,A2*) and common knowledge of the MEC is satisfied. However, if Kay does not know Amie’s choice, then Amie can serve her own best interests by appearing to be aggressive, in order to deceive Kay into playing *A1* in hopes of gaining the higher payoff associated with the nonequilibrium point (*A1,A1*).

So this equilibrium does not satisfy the PIC.

I now give a formal definition of convention, motivated by the discussion above. I first give the definitions of a game in strategic form and of Aumann's correlated equilibrium concept (Aumann 1974, 1987) to fix notation.

Definition 1.1. A game Γ is an ordered triple (N, S, \mathbf{u}) consisting of the following elements:

- (1) A finite set $N = \{1, 2, \dots, n\}$, called the *set of players*.
- (2) For each player $k \in N$, there is a finite set $S_k = \{A_{k1}, A_{k2}, \dots, A_{kn_k}\}$, called the *alternative pure strategies* for player k . The Cartesian product $S = S_1 \times \dots \times S_n$ is called the *pure strategy set* for the game Γ .
- (3) A map $\mathbf{u}: S \rightarrow \mathbb{R}^n$, called the *payoff function* on the pure strategy set. At each strategy combination $\mathbf{A} = (A_{1j_1}, \dots, A_{nj_n}) \in S$, player k 's payoff is given by the k th component of the value of \mathbf{u} , that is, player k 's payoff u_k at \mathbf{A} is determined by

$$u_k(\mathbf{A}) = I_k \circ \mathbf{u}(A_{1j_1}, \dots, A_{nj_n}),$$

where $I_k(\mathbf{x})$ projects $\mathbf{x} \in \mathbb{R}^n$ onto its k th component. \square

In Aumann's model, the players are engaged in a game $\Gamma = (N, S, \mathbf{u})$, and each player has a personal *information partition* \mathfrak{H}_k of a finite probability space Ω . The elementary events $\omega \in \Omega$ are called the *states of the world*, and at each ω , every player k knows that the element $H_{kj} \in \mathfrak{H}_k$ such that $\omega \in H_{kj}$ has occurred, but does not in general know which ω has occurred. Intuitively, H_{kj} represents k 's private information regarding the states of the world. A function $f: \Omega \rightarrow S$ defines a system of *exogenously correlated strategy n -tuples*, that is, for each $\omega \in \Omega$, the players select a strategy combination $f(\omega) = (f_1(\omega), \dots, f_n(\omega)) \in S$ correlated with the state of the world ω . Informally, f is a correlated equilibrium if at each state of the world $\omega \in \Omega$, $f_k(\omega)$ is optimal for each player k , in the sense that $f_k(\omega)$ maximizes k 's expected payoff given k 's private information and k 's beliefs regarding his opponents. In the following definition and in the sequel, ' E ' denotes expectation, $p_k(\cdot)$ denotes player k 's subjective probability distribution and the subscript ' $-k$ ' indicates the result of removing the k th component of an n -tuple or an n -fold Cartesian product. For instance,

$$f_{-k}(\omega) = (f_1(\omega), \dots, f_{k-1}(\omega), f_{k+1}(\omega), \dots, f_n(\omega))$$

that is, $f_{-k}(\omega)$ denotes the strategy combinations of player k 's opponents determined by f at $\omega \in \Omega$.

Definition 1.2. Given $\Gamma = (N, S, \mathbf{u})$, Ω , and the information partitions \mathfrak{H}_k of Ω as defined above, $f: \Omega \rightarrow S$ is a *correlated equilibrium* if and only if, for each $k \in N$,

- (i) f_k is an \mathfrak{H}_k -measurable function, that is, for each $H_{kj} \in \mathfrak{H}_k$, $f_k(\omega')$ is constant for each $\omega' \in H_{kj}$, and
- (ii) For each $\omega \in \Omega$,

$$(1.2.1) \quad E(u_k \circ f \mid \mathfrak{H}_k)(\omega) \geq E(u_k \circ (f_{-k}, g_k) \mid \mathfrak{H}_k)(\omega)$$

for any \mathfrak{H}_k -measurable function $g_k: \Omega \rightarrow S_k$.

The correlated equilibrium f is *strict* iff the inequalities of (1.2.1) are all strict. \square^8

The measurability restriction on f_k formalizes the intuition that at each state of the world, player k knows

which action he will perform. Note that Aumann's correlated equilibrium concept generalizes the Nash equilibrium concept. Players are at a *Nash equilibrium* $f = (f_1, \dots, f_n)$ if each player k pegs his strategy f_k on a randomizing device Ω_k that is *probabilistically independent* of the randomizing devices of his opponents and such that $f_k(\omega)$ is optimal for k given the strategies of k 's opponents. Hence f is formally a correlated equilibrium where $\Omega = \Omega_1 \times \dots \times \Omega_n$.

In the following definition of convention, the agents refer to a common information partition of the states of the world. While each agent k has a private information partition \mathfrak{H}_k of Ω , there is a partition of Ω , namely the intersection $\mathfrak{H} = \bigcap_{k \in N} \mathfrak{H}_k$, of the states of the world such that for each $\omega \in \Omega$, all of the agents will know which cell $H(\omega) \in \mathfrak{H}$ occurs.⁹ The agents' expected utilities in Definition 1.3 are conditional on their common partition \mathfrak{H} , reflecting the intuition that conventions rely upon information that is public to all.

Definition 1.3. Given $\Gamma = (N, S, \mathbf{u})$, Ω , and the partition \mathfrak{H} of Ω of cells that are common knowledge among the players, a function $f: \Omega \rightarrow S$ is a *convention* if and only if for each $\omega \in \Omega$, and for each $k \in N$, f_k is \mathfrak{H} -measurable and

$$E(u_k \circ f \mid \mathfrak{H})(\omega) > E(u_k \circ (f -_j, g_j) \mid \mathfrak{H})(\omega)$$

for each $j \in N$ and for any \mathfrak{H} -measurable function $g_j: \Omega \rightarrow S_j$. \square

In words, if any player j deviates unilaterally from the convention f , then every player $k \in N$, including j , is worse off. Note that if f is a convention, then f is clearly a strict correlated equilibrium. Moreover, this definition of convention satisfies the PIC. Since the agents in a convention refer to a common partition, at each $\omega \in \Omega$, a given player k 's opponents all know exactly which strategy k will play, and k will want his opponents to know this. If any of k 's opponents mistakenly believed that k would play a strategy $g_k(\omega) \neq f_k(\omega)$ at some $\omega \in \Omega$, then this opponent might be tempted to deviate from the convention, which would leave k strictly worse off. On the other hand, if k 's opponent's all know that k will play $f_k(\omega)$ at every $\omega \in \Omega$, then they will all have a decisive reason to conform with $f(\omega)$, with the result that k is strictly better off.

Unlike Lewis' definition of convention, Definition 1.3 formally incorporates the notion of salience. Indeed, throughout his work Lewis leaves the notion of salience somewhat imprecise. This is perhaps not surprising, since salience can arise from so many different factors. On the other hand, salience is built into the correlated equilibrium definition of convention, since the correlated equilibrium which characterizes a specific

⁸As Brandenburger and Dekel (1987) point out, this definition of correlated equilibrium requires the conditional expectations to be defined and condition (1.2.1) to be met even if some players assign zero probabilities to some of the cells of their information partitions. Following their lead, I require that for every player $k \in N$ and for each $H_{kj} \in \mathfrak{H}_k$,

(i) $p_k(\cdot \mid H_{kj})$ is a probability measure on Ω , and

(ii) $p_k(H_{kj} \mid H_{kj}) = 1$

for each $H_{kj} \in \mathfrak{H}_k$. Of course, these conditions follow immediately from the conventional definition of conditional probability if $p_k(H_{kj}) > 0$, but the point is that imposing (i) and (ii) on every cell of every player's information partition extends the definition of conditional probability to events of zero probability.

⁹Since the intersection of any collection of partitions of Ω is again a partition, \mathfrak{H} is a partition of Ω which contains (by definition) cells common to every partition \mathfrak{H}_k , $k \in N$.

convention is a function of “states of the world” on which the agents coordinate their actions. In other words, a convention f as defined in Definition 1.3 is salient to the players because they correlate (via f) their actions and expectations with various pieces of information at their disposal, which are formalized as elements of a partition of an event space Ω . For instance, in Battle of the Sexes, if Amie and Kay follow the convention of playing $(A1, A1)$ or $(A2, A2)$ depending on the result of the fair coin toss, then one can construct the event space $\Omega = \{\omega_1, \omega_2\}$, where ω_1 denotes “ H ” and ω_2 denotes “ T ”. The convention

$$(1.4) \quad f(\omega) = \begin{cases} (A1, A1) & \text{if } \omega = \omega_1 \\ (A2, A2) & \text{if } \omega = \omega_2 \end{cases}$$

is salient for the two players in virtue of how their beliefs regarding their own and their opponents’ choices are correlated with the elements of Ω . In general, the event space Ω in the correlated equilibrium definition of convention makes explicit *what* the agents who follow a convention correlate their expectations on. But it raises the following question: How do the agents come to correlate on the elements of Ω in any particular way? I will address this question in §2.

§2. Dynamical Explanations of Conventions

How do agents come to regard one of many correlated equilibria as salient, so that this equilibrium becomes a convention? I address this question by modeling the agents as *inductive deliberators* who modify their beliefs about one another recursively. One way to explain how players select a particular correlated equilibrium is to exploit the common knowledge assumptions that justify the correlated equilibrium concept, and to have the players modify their beliefs, quantified as subjective probability distributions, with a common *inductive rule*. Aumann’s (1974, 1987) correlated equilibrium concept implies that the following are common knowledge if the players are to *know* that they are at correlated equilibrium:

- (C1) Each player knows the payoff structure of the game Γ , the partitions of Ω and the function $f: \Omega \rightarrow S$,
- (C2) Each player is *Bayesian rational*, that is, she acts so as to maximize her expected utility, and
- (C3) Each player knows the beliefs of her opponents, as well as her own beliefs.

Given these common knowledge assumptions, if the players update their beliefs inductively, then even if their beliefs are not initially such that they will play an equilibrium, the sequence of updated beliefs can eventually converge to an equilibrium. One well-known inductive rule for belief updating is the *Dirichlet rule*: If $n_{A_{-k}}$ is the number of times that k ’s opponents have played the strategy combination A_{-k} over the first n rounds of deliberation, then k ’s probability that the opponents play A_{-k} at the n th stage of deliberation is

$$p_k^n(A_{-k}) = \frac{n_{A_{-k}} + \lambda_k p_k^0(A_{-k})}{n + \lambda_k}$$

where $\lambda_k > 0$ and $p_k^0(A_{-k})$ is k ’s *prior probability* that A_{-k} is played. The Dirichlet rule is called *inductive* because it is an instance of the rule for inductive logic proposed by Carnap (1980). Note that after n rounds of play, k ’s probability that the opponents will play A_{-k} is a *mixture* of k ’s prior and the relative frequency of A_{-k} during the first n rounds. The parameter λ_k adjusts the weighting placed on the effect of the initial rounds. This method of updating probabilities generalizes the *method of fictitious play* introduced by Brown (1951) as a means of computing mixed Nash equilibria. In the present context, the players use the inductive

rule to *learn* which strategy they ought to play given what their opponents play. The sequence $(\mathbf{p}^m) = ((p_1^m, \dots, p_n^m))$ determined by the successive applications of the Dirichlet rule forms a *dynamical system*, and consequently one may refer to Dirichlet deliberation as *Dirichlet dynamics*. A *deliberational equilibrium* of a game is a fixed point of this dynamical process.

There are two natural ways to interpret inductive deliberation:

- (1) *Actual Play Interpretation*: A fixed set of players play the game repeatedly, and update probabilities as functions of the frequencies of strategy combinations they observe. Each player chooses a strategy that maximizes expected utility at each round of deliberation, and updates his probabilities according to the Dirichlet rule.
- (2) *Representative Interpretation*: The n players correspond to n distinct populations. At each stage of deliberation, one representative from each population is selected to play a round of the game. Each representative plays the strategy that maximizes expected utility given what the representatives in past plays have done, where probabilities for the current round are determined by the Dirichlet rule, and then drops out.¹⁰

The model of Dirichlet deliberation can be generalized in a variety of ways.¹¹ Vanderschraaf and Skyrms (1993) extend the Dirichlet model by introducing external signals that players observe at each round of deliberation, and by not requiring that players' strategies be stochastically independent. In this model of *conditional Dirichlet deliberation*, each player k has an information partition \mathfrak{H}_k of a finite probability space Ω , each element of which occurs with positive probability.¹² At each round of deliberation, a state of the world $\omega^t \in \Omega$ occurs, and k learns that an event $H_{kj}(\omega^t) \in \mathfrak{H}_k$ has occurred. Ω is the external event space that underpins Aumann's (1974, 1987) correlated equilibrium concept. However, the players are not necessarily settled on a correlated equilibrium to begin with.

Definition 2.1. Let $\Gamma = (N, S, U)$, Ω , and the partitions \mathfrak{H}_k of Ω be given, and let an elementary event $\omega^m \in \Omega$ occur at each round m of deliberation. The players $1, \dots, n$ are *conditional Dirichlet deliberators* iff their conditional probabilities $p_k^m(\cdot | \mathfrak{H}_k)$, $k \in N$, at each stage of deliberation are defined recursively as follows:

- (1) Each player k has a *prior conditional distribution* $p_k^0(\cdot | \mathfrak{H}_k)$ for the act combinations of his opponents in S_{-k} and a value $\lambda_{H_{kj}} > 0$ for each $H_{kj} \in \mathfrak{H}_k$.

¹⁰There is a third interpretation of this process as *fictitious play*: The players mentally simulate a sequence of successive plays by computing one another's expected utilities at each stage of deliberation, noting which strategy maximizes expected utility for each player, and updating probabilities according to the Dirichlet rule. However, the extensions of Dirichlet dynamics introduced in §3 do not admit of a plausible fictitious play interpretation. See Vanderschraaf (forthcoming) for details.

¹¹See Monderer and Sela (1993) for an overview of some of the generalizations of Dirichlet dynamics which do not involve correlated strategies.

¹²This restriction is added to avoid mathematical trivialities. To justify it, one might argue that in a deliberational context, a deliberator will only want to condition on events that can occur with positive probability, since these are the only events that with positive probability can give the deliberator information during the sequence of deliberation. However, it is in principle possible to extend the model of deliberation to include partition cells of measure zero, as Brandenburger and Dekel (1987) do for Aumann's correlated equilibrium concept, but this would make the formalism somewhat more complicated without adding substantively to the deliberational model in any way I can see.

(2) At each m th round of deliberation, each player k plays a strategy $A_{kj}^m \in S_k$ such that

$$E(u_k(A_{kj}^m) | \mathfrak{H}_k)(\omega^m) \geq E(u_k(A_{kl}) | \mathfrak{H}_k)(\omega^m) \text{ for } j \neq l.$$

For each $H_{kj} \in \mathfrak{H}_k$ and each $A_{-k} \in S_{-k}$,

$$n_{H_{kj}} = \sum_{t=1}^m 1_{H_{kj}}(\omega^t) \text{ and } n_{A_{-k} \wedge H_{kj}} = \sum_{t=1}^m 1_{A_{-k} \wedge H_{kj}}(\omega^t)^{13}$$

that is, $n_{H_{kj}}$ is the number of times that H_{kj} has occurred during the first m rounds of deliberation, and $n_{A_{-k} \wedge H_{kj}}$ is the number of times that A_{-k} and H_{kj} have both occurred during the first m rounds of deliberation. At the $m+1$ st round of deliberation, player k updates his probability distribution $p_k^m(\cdot | \mathfrak{H}_k)$ according to the rule

$$p_k^{m+1}(A_{-k} | H_{kj}) = p_k^m(A_{-k} | H_{kj})$$

if H_{kj} did not occur at the $m+1$ st round of deliberation, and, if H_{kj} occurred at the $m+1$ st round of deliberation, then

$$p_k^{m+1}(A_{-k} | H_{kj}) = \frac{n_{A_{-k} \wedge H_{kj}} + \theta + \lambda_{H_{kj}} p_k^0(A_{-k} | H_{kj})}{n_{H_{kj}} + 1 + \lambda_{H_{kj}}}$$

$$\text{where } \theta = \begin{cases} 1 & \text{if the opponents play } A_{-k} \text{ at round } m+1 \\ 0 & \text{otherwise} \end{cases}$$

The sequence of updated distributions is denoted by

$$\mathfrak{P}^m = (p_1^m(\cdot | \mathfrak{H}_1), \dots, p_n^m(\cdot | \mathfrak{H}_n)). \quad \square$$

For illustration, consider Battle of the Sexes once more. Suppose that Kay and Amie are to play Battle of the Sexes repeatedly, but are not initially at equilibrium. They are not allowed to communicate, though they both observe the toss of a fair coin at each round of play. They cannot explicitly agree to play the convention $(A1, A1)$ if H and $(A2, A2)$ if T , but if they are conditional Dirichlet deliberators with the appropriate priors, they can *learn* to play this convention. In this case, $\Omega = \{H, T\}$, and $\mathfrak{H}_1 = \mathfrak{H}_2 = \{\{H\}, \{T\}\}$. For instance, if their priors are:

$$p_1^0(\text{Kay plays } A1 | H) = p_2^0(\text{Amie plays } A1 | H) = .51,$$

$$p_1^0(\text{Kay plays } A2 | H) = p_2^0(\text{Amie plays } A2 | H) = .49,$$

$$p_1^0(\text{Kay plays } A1 | T) = p_2^0(\text{Amie plays } A1 | T) = .49, \text{ and}$$

$$p_1^0(\text{Kay plays } A2 | T) = p_2^0(\text{Amie plays } A2 | T) = .51.$$

and $\lambda_{1H} = \lambda_{1T} = \lambda_{2H} = \lambda_{2T} = 100$, then conditional Dirichlet dynamics converges almost surely to the distribution that characterizes the correlated equilibrium convention defined by (1.4).¹⁴

Why would deliberators who do not communicate with one another condition their updated

¹³Following the usual convention, the indicator function $1_E: \Omega \rightarrow \{0, 1\}$ is defined as

$$1_E(\omega) = \begin{cases} 1 & \text{if } \omega \in E \\ 0 & \text{if } \omega \notin E \end{cases}$$

¹⁴Conditional Dirichlet deliberation converges to the convention $(A1, A1)$ if H and $(A2, A2)$ if T unless either H or T occurs only finitely many times in the sequence of coin flips, in which case one of the subsequences of conditional deliberation will stop short of its attractor. By the Borel-Cantelli lemma, with probability one H and T both occur infinitely often in a sequence of flips of a fair coin, so the conditional deliberation converges to the correlated equilibrium with probability one.

probabilities on information from the “states of the world”? Deliberators who condition on their information partitions are simply using all of the information at their disposal when they update. If the information a deliberator receives from the external event turns out to be statistically independent of the results of play, the updated probability will be the same as if she had updated unconditionally. On the other hand, if the deliberator notices a correlation between her opponents’ strategies and her signals, she can incorporate this correlation into his deliberations with a rule like the conditional Dirichlet rule.

In the special case in which deliberators all have the same information partition \mathfrak{H} , one can prove some fairly general convergence results for conditional Dirichlet dynamics. Such results are especially relevant to dynamical accounts of the emergence of conventions, since we define convention in terms of information from the *public* information that the various agents gain from their common partition of the states of the world. The key idea underlying these results is that when deliberators condition on a common partition \mathfrak{H} , then the resulting dynamical process can be regarded as a union of sequences of unconditional deliberation.

Definition 2.2. Given a set conditional Dirichlet deliberators who condition on a common information partition \mathfrak{H} of the states of the world, the H_k -subsequence of deliberation (\mathfrak{P}^{m_k}) is the subsequence of updated conditional distributions (\mathfrak{P}^m) of the deliberators such that $\omega^{m_k} \in H_k \in \mathfrak{H}$. \square

If deliberators condition on a public information partition \mathfrak{H} , then trivially the distribution of any H_k -subsequence with conditional priors $p_k^0(\cdot | H_k)$, $k \in N$, is the same as the distributions of a sequence of unconditional Dirichlet deliberation with priors $p_k^0(\cdot)$. This observation enables us to take advantage of the following result, proved in Monderer and Shapley (1993):

Proposition 2.1 (Monderer and Shapley). If Γ is a game with identical payoff functions for each player $k \in N$, then unconditional Dirichlet deliberation converges to a Nash equilibrium.

Proposition 2.2. If Γ is a game with identical payoff functions for each player, and deliberators condition on a common information partition \mathfrak{H} of Γ , then conditional Dirichlet deliberation converges almost surely to a correlated equilibrium.

PROOF. By Proposition 2.1, each H_k -subsequence of conditional Dirichlet deliberation converges to a Nash equilibrium of Γ if H_k occurs infinitely often in the sequence of deliberation. By the Borel-Cantelli lemma, this occurs with probability one. Hence, conditional Dirichlet deliberation converges almost surely to a function $f: \Omega \rightarrow S$ such that f is constant on each $H_k \in \mathfrak{H}$ and for each $\omega \in H_k \in \mathfrak{H}$, $f(\omega)$ is a Nash equilibrium, so that f satisfies (1.2.1). \square

In a game with identical payoffs, when the limit of each H_k -subsequence of conditional Dirichlet deliberation is a strict Nash equilibrium, then conditional Dirichlet deliberation converges almost surely to a convention. However, it is possible for the dynamics to converge to a mixed equilibrium, which is not a convention since at a mixed equilibrium the players are indifferent to their alternative pure strategies.

On the other hand, a convention is always an *absorbing state* of conditional Dirichlet dynamics, in the sense that if conditional Dirichlet deliberators coordinate just once at each state of the world, then they will always coordinate at each state of the world, with the result that their beliefs will converge to the

correlated equilibrium corresponding to a convention. This result relies upon the following fundamental result regarding unconditional deliberation:

Proposition 2.3 (Absorption Theorem). If Dirichlet deliberators play a strict Nash equilibrium A^* at any round of deliberation, they will play A^* at all future rounds of deliberation.¹⁵

Proposition 2.4. A convention $f: \Omega \rightarrow S$ is an absorbing state of conditional Dirichlet dynamics in the case where deliberators condition on a common partition \mathfrak{H} . That is, if conditional Dirichlet deliberators all condition on \mathfrak{H} , and if for each $\omega \in \Omega$, there is a value t^ω such that the deliberators play $f(\omega)$ at round t^ω , then for all rounds $m \geq \max\{t^\omega \mid \omega \in \Omega\}$, the deliberators will always play the convention $f(\omega)$.

PROOF. Let $\omega \in \Omega$, and let $H_k \in \mathfrak{H}$ be such that $\omega \in H_k$. By hypothesis, there is a round t^ω at which the deliberators play $f(\omega)$. Since f is a convention, $f(\omega)$ is a strict Nash equilibrium, so by the Absorption Theorem the H_k -subsequence of deliberation remains on $f(\omega)$ for each round $m_k \geq t^\omega$ of deliberation such that $\omega^{m_k} \in H_k$. Since $\omega \in \Omega$ was chosen arbitrarily, and there are finitely many ω , for all $t \geq \max\{t^\omega \mid \omega \in \Omega\}$, the deliberators will play $f(\omega)$. \square

§3. Deliberation With Knowledge Asymmetries

Conditional Dirichlet deliberators can converge to the equilibrium of a convention given the right priors, but in many cases, deliberators may fail to converge to a convention. Often this failure is the result of perfect symmetry in the deliberators' beliefs and knowledge, which results in their converging to a mixed equilibrium rather than to a convention. In particular, deliberators might fail to converge to a convention if their priors are "uncorrelated", that is, if the priors satisfy probabilistic independence. For example, if Amie and Kay play Battle of the Sexes as conditional Dirichlet deliberators as in the example in §2, but with uniform conditional priors, then each subsequence of conditional Dirichlet deliberation converges to the mixed equilibrium defined by the distributions

$$p_1(\text{Kay plays } A1) = p_2(\text{Amie plays } A2) = \frac{1}{1 + \sqrt{2}},$$

at which each player is indifferent to her own alternative strategies. Hence, the PIC is not satisfied at this equilibrium, so the limit of this particular sequence of Dirichlet deliberation is not a convention.

Nevertheless, inductive deliberators starting with uncorrelated priors can converge to a convention if they have *asymmetric knowledge* of the history of deliberation. One can introduce asymmetries in the players' knowledge in at least two ways: (1) If each player recalls only a random number of the most recent m past plays, rather than the full sequence of past plays, then *random recall Dirichlet deliberation* (RRC-Dirichlet deliberation) introduces a knowledge asymmetry by making it possible for deliberators to have unequal memories of the past history of plays. (2) If each player periodically fails to observe an element of his private information partition \mathfrak{H}_k , then *variable state conditional Dirichlet deliberation* (VSC-Dirichlet deliberation) introduces a knowledge asymmetry by occasionally giving deliberators unequal knowledge of the states of the world which determine salience. In both of these extensions of the conditional Dirichlet model, I show by

¹⁵The result I call the Absorption Theorem is well-known among game theorists, though I have not found a proof in the literature. For a proof, see Vanderschraaf (forthcoming).

example that correlation in deliberators' beliefs can emerge spontaneously, with the result that they can converge to a convention even if their initial beliefs are uncorrelated.

I first consider RRC-Dirichlet dynamics. In this setting, each player k will remember which strategy profiles have been played for a certain number η_k^t of the most recent stages of play, where η_k^t is a random variable. As before, I use the Battle of the Sexes game to illustrate the dynamics. Suppose that Kay and Amie begin conditional Dirichlet deliberation with the common partition $\mathfrak{H} = \{\{H\}, \{T\}\}$ of the event space $\Omega = \{H, T\}$ and with uniform priors, that both have complete recall for the first 16 rounds of deliberation, but that afterwards they start to forget some of the early history of play. One possible sequence of the first 18 rounds of play is summarized in Table 1. On the 17th round of deliberation, the coin comes up H and Amie can recall 7 of the most recent rounds at which she observed H , while Kay can only recall 6 of the most recent rounds in which she observed H . Given their distributions updated on the rounds they each *remember*, Amie and Kay both maximize expected utility by playing $A1$. On the 18th round of deliberation, the coin comes up T , and at this round Kay can recall the 7 most recent rounds at which she observed T , while Amie can only recall the 6 most recent rounds at which she observed T . Given the updated distributions based on what each player remembers at this round, they coordinate on $(A2, A2)$. Moreover, in subsequent rounds of deliberation the two players will continue to coordinate, so long as they can recall at least one round of previous play. This is because, given the parameters of the dynamics in this example, if the players can recall that they have coordinated at the most recent round of play, then their distributions will be in the basin of attraction of the conditional distribution vector corresponding to $(A1, A1)$ given H or $(A2, A2)$ given T . Consequently, deliberation with random recall in this instance converges almost surely to the convention $f: \Omega \rightarrow S$ defined by (1.4).

The RRC-Dirichlet dynamics is similar in spirit to the *stochastic adaptive dynamics* developed in Young (1993) and in Kandori, Mailath and Rob (1993), but also differs from stochastic adaptive dynamics in several important respects. Stochastic adaptive deliberators play their best response given their probabilities at each round of deliberation, same as Dirichlet deliberators. They also have imperfect histories of past plays, as in RRC-Dirichlet dynamics. However, while RRC-Dirichlet deliberators each remember a number of the *most recent* past plays which varies at each round of deliberation, stochastic adaptive deliberators update their probabilities from *random samples* of past opponents' plays, which need not include the most recent plays. Also, stochastic adaptive deliberators all sample the same number of past opponents' plays, though they need not sample the same past plays, while RRC-Dirichlet deliberators each recall a random number of the most recent past plays which can differ from the random numbers of the most recent past plays their opponents recall. Finally, the models of stochastic adaptive dynamics developed in the works cited do not incorporate external signals, unlike RRC-Dirichlet dynamics. Due to this last difference, the research on stochastic adaptive dynamics has focused on identifying sufficient conditions under which the dynamics converge to Nash equilibrium, while I am introducing RRC-Dirichlet dynamics as one means of explaining the emergence of correlated equilibrium.

I now give the definition of RRC-Dirichlet deliberation. Note that the priors are defined as functions of *initial prior states*, which deliberators can eventually forget. In an actual play interpretation of the dynamics, it seems plausible to suppose that if deliberators can forget part of the past history of play, they will not be able to remember indefinitely the the states that defined their priors before they began play.¹⁶

Definition 3.1. Let $\Gamma = (N, S, U)$, Ω , and the partitions \mathfrak{H}_k of Ω be given, and let an elementary event $\omega^t \in \Omega$ occur at each round of deliberation. The players are *RRC-Dirichlet deliberators* iff their conditional probabilities $p_k^m(\cdot | \mathfrak{H}_k)$, $k \in N$, at each stage of deliberation are defined recursively as follows:

- (1) k associates an *initial weight* $\gamma_{iA_{-k} \wedge H_{kj}} \geq 0$ with each of $\tau_{A_{-k} \wedge H_{kj}}$ *initial states*. k 's initial weight for the conjunction $A_{-k} \wedge H_{kj}$, $\gamma_{A_{-k} \wedge H_{kj}}^0$ is defined by

$$\gamma_{A_{-k} \wedge H_{kj}}^0 = \sum_{i=1}^{\tau_{A_{-k} \wedge H_{kj}}} \gamma_{iA_{-k} \wedge H_{kj}}.$$

k 's conditional prior distribution $p^0(\cdot | \mathfrak{H}_k)$ is defined by

$$p_k^0(A_{-k} | H_{kj}) = \frac{\gamma_{A_{-k} \wedge H_{kj}}^0}{\sum_{B_{-k} \in S_{-k}} \gamma_{B_{-k} \wedge H_{kj}}^0}, H_{kj} \in \mathfrak{H}_k$$

- (2) For round t of deliberation and for each player k , let η_k^t be a random variable with range $\{1, 2, \dots, M_k^t\}$, $1 \leq M_k^t \leq t$, and let τ_k^t be a random variable with range $\{1, 2, \dots, N_k^t\}$, $1 \leq N_k^t \leq t$. For each $H_{kj} \in \mathfrak{H}_k$, and each $A_{-k} \in S_{-k}$

$$\kappa_{H_{kj}}^t = \sum_{i=t-\eta_k^t}^t 1_{H_{kj}}(\omega^i) \text{ and } \kappa_{A_{-k} \wedge H_{kj}}^t = \sum_{i=t-\eta_k^t}^t 1_{A_{-k} \wedge H_{kj}}(\omega^i)$$

that is, $\kappa_{H_{kj}}^t$ and $\kappa_{A_{-k} \wedge H_{kj}}^t$ are the number of times that H_{kj} has occurred and the number of times that A_{-k} and H_{kj} have both occurred, respectively, during the most recent η_k^t rounds rounds of deliberation. At the $t+1$ st round of deliberation, player k updates his probability distribution $p_k^t(\cdot | \mathfrak{H}_k)$ according to the rule

$$p_k^{t+1}(A_{-k} | H_{kj}) = p_k^t(A_{-k} | H_{kj})$$

if H_{kj} did not occur at the $t+1$ st round of deliberation, and, if H_{kj} occurred at the $t+1$ st round of deliberation, then

$$p_k^{t+1}(A_{-k} | H_{kj}) = \frac{\kappa_{A_{-k} \wedge H_{kj}}^t + \theta + \gamma_{A_{-k} \wedge H_{kj}}^t}{\kappa_{H_{kj}}^t + 1 + \sum \gamma_{B_{-k} \wedge H_{kj}}^t}, \gamma_{B_{-k} \wedge H_{kj}}^t = \sum_{i=t-\tau_k^t+1}^{\tau_{B_{-k} \wedge H_{kj}}} \gamma_{iB_{-k} \wedge H_{kj}}$$

that is, $\gamma_{B_{-k} \wedge H_{kj}}^t$ is the sum of the initial initial weights that k can recall at time t , and θ is defined as in Definition 2.1. \square

In words, Definition 3.1 says that at time t , deliberator k can recall the results of the η_k^t most recent rounds of play and, if $t - \tau_k^t < \tau_{B_{-k} \wedge H_{kj}}$, then k can also recall $\tau_{B_{-k} \wedge H_{kj}} - t - \tau_k^t$ of the initial states that define k 's priors. If one regards η_k^t and τ_k^t as a degenerate random variables, then conditional Dirichlet deliberation

¹⁶Thanks to Vince Crawford for clearing me up on this point.

with full memory is a special case of RRC-Dirichlet deliberation, where $\eta_k^t = \tau_k^t = t$ and $\gamma_{\mathbf{B}-k}^t \wedge H_{kj} = \gamma_{\mathbf{B}-k}^0 \wedge H_{kj}$ for all $k \in N$ and all $t \geq 1$.

One can interpret this kind of dynamics plausibly either as a sequence of actual plays by individuals with imperfect memories, or in representative terms. Under an actual play interpretation in which players update their probabilities by applying the conditional Dirichlet rule, one can easily imagine that they might not be able to remember the entire history of plays, but that they are able to remember what their opponents did at a number of the most recent plays. If one interprets the deliberators to be representatives of populations who play the game once and then drop out, then it makes sense to suppose that each representative can observe some of the sequence of plays immediately preceding her turn at the game, but not necessarily the entire history of the game or the same set of recent past plays the representatives she will oppose observe.

I now turn to the VSC-Dirichlet model, which allows for the possibility that at some rounds of deliberation, a deliberator might have a more coarse information partition of Ω than at other times. One can regard this intermittent “coarsening” of deliberators’ partitions as a form of “noise” in the deliberators’ information, which can lead to exogenous correlation in strategies. Once again, to motivate this type of dynamics I use the Battle of the Sexes example. In this scenario, Kay and Amie receive information from an external event space Ω , such that

$$\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$$

where the ω_k ’s are equiprobable. For concreteness, one can interpret the points of Ω as the outcomes of two independent flips of a fair coin, such that $\omega_1 = HH$, $\omega_2 = HT$, $\omega_3 = TH$, $\omega_4 = TT$. If Amie and Kay are Dirichlet deliberators who are always able to observe the elementary event $\omega_k \in \Omega$, then their common information partition is $\mathfrak{K} = \{\{\omega_1\}, \{\omega_2\}, \{\omega_3\}, \{\omega_4\}\}$. Starting with uniform conditional priors, conditional Dirichlet deliberation converges almost surely to the mixed Nash equilibrium defined by

$$p_1(\text{Kay plays } A1 \mid \{\omega_j\}) = p_2(\text{Amie plays } A2 \mid \{\omega_j\}) = \frac{1}{1 + \sqrt{2}}, \quad 1 \leq j \leq 4.$$

However, suppose instead that at each round of deliberation there is a positive probability that Amie will not observe the outcome of the second coin toss, and there is likewise a positive probability that Kay will not gain information from the second coin toss. In this setting, the partition H_k^t for each deliberator can vary with respect to time. At each round of deliberation, Amie and Kay each observe an event $H_k^t \in \mathfrak{K}_k^t \subseteq \mathfrak{F}$, where \mathfrak{F} is the σ -algebra generated by \mathfrak{K} . They each update all of their conditional probabilities $p_k^t(\cdot \mid F_k)$, $F_k \in \mathfrak{F}$, and each plays a strategy that maximizes her conditional expectation given $H_k^t \in \mathfrak{K}_k^t$. One possible early sequence of VSC-Dirichlet deliberation with these partitions is summarized in Table 2. In this example, the deliberators converge to the correlated equilibrium $f: \Omega \rightarrow S$ defined by

$$f(\omega) = \begin{cases} (A1, A1) & \text{if } \omega \in \{\omega_1, \omega_2\} \\ (A2, A2) & \text{if } \omega \in \{\omega_3, \omega_4\} \end{cases}.$$

The leading idea behind VSC-Dirichlet deliberation is that the periodic coarsenings of the deliberators’ information partitions can introduce an asymmetry in their knowledge of the past history of deliberation. As in conditional Dirichlet deliberation with fixed information partitions, each deliberator in

VSC-Dirichlet deliberation uses all of the information she has at any given time in computing expected utilities. However, at certain times the information the deliberator receives about the states of the world can be less detailed than at other times, and in particular less detailed than the information her opponents receive, which is formalized by the deliberator's information partition being more coarse than at other times.

In the following definition, the key point to keep in mind is that in VSC-Dirichlet dynamics, deliberators update several different conditional probabilities simultaneously, namely the probabilities conditional on the events $F_k^t \in \mathcal{F}_k$, $k \in N$, that the deliberators observe at round t , as well as the conditional probabilities for each event $F_{kj} \in \mathcal{F}_k$ such that $F_k^t \subseteq F_{kj}$. This is because, when k has observed F_k^t at round t , then if $F_k^t \subseteq F_{kj}$, then k knows that F_{kj} has occurred at round t as well.

Definition 3.2. Let $\Gamma = (N, S, U)$, Ω , and the partitions \mathcal{H}_k of Ω be given, and let an elementary event $\omega \in \Omega$ occur at each round of deliberation. Let \mathcal{F}_k denote the σ -algebra generated by the events in \mathcal{H}_k . The players $1, \dots, n$ are *VSC-Dirichlet deliberators* iff they update their conditional probabilities recursively as follows:

- (1) At the start of deliberation, each player k assigns a value $\gamma_{A_{-k} \wedge H_{kj}} \geq 0$ to each $A_{-k} \in S_{-k}$ and each $H_{kj} \in \mathcal{H}_k$. For every $H_{kj} \in \mathcal{H}_k$, $\gamma_{B_{-k} \wedge H_{kj}} > 0$ for at least one $B_{-k} \in S_{-k}$. For each $F_{kj} \in \mathcal{F}_k$, define

$$\gamma_{A_{-k} \wedge F_{kj}} = \sum_{H_{kj} \subseteq F_{kj}} \gamma_{A_{-k} \wedge H_{kj}}$$

The prior conditional distribution $p^0(\cdot | \mathcal{F}_k)$ that each player has for the act combinations of his opponents in S_{-k} is defined by

$$p_k^0(A_{-k} | F_{kj}) = \frac{\gamma_{A_{-k} \wedge F_{kj}}}{\sum_{B_{-k} \in S_{-k}} \gamma_{B_{-k} \wedge F_{kj}}}, \quad F_{kj} \in \mathcal{F}_k$$

- (2) At each round of deliberation $t \geq 1$, each deliberator k observes an event $F_k^t \in \mathcal{F}_k$, and plays a strategy $A_{kj} \in S_k$ such that

$$E(u_k(A_{kj}) | F_k^t) \geq E(u_k(A_{kl}) | F_k^t) \text{ for all } A_{kl} \in S_k.$$

For each $F_{kj} \in \mathcal{F}_k$ such that $F_k^t \subseteq F_{kj}$, k knows that F_{kj} has occurred at round t as well. For each $F_{kj} \in \mathcal{F}_k$, $n_{F_{kj}}$ is the number of times that k knows that F_{kj} has occurred during the first t rounds of deliberation. For each $F_{kj} \in \mathcal{F}_k$ and each $A_{-k} \in S_{-k}$, $n_{A_{-k} \wedge F_{kj}}$ is the number of times that k knows that A_{-k} and F_{kj} have both occurred during the first t rounds of deliberation. At the $t+1$ st round of deliberation, player k updates her probability distribution $p_k^n(\cdot | \mathcal{F}_k)$ according to the rule

$$p_k^{t+1}(A_{-k} | F_{kj}) = p_k^t(A_{-k} | F_{kj}) \text{ if } F_{kj}^{t+1} \not\subseteq F_{kj}, \text{ and}$$

$$p_k^{t+1}(A_{-k} | F_{kj}) = \frac{n_{A_{-k} \wedge F_{kj}} + \theta + \gamma_{A_{-k} \wedge F_{kj}}}{n_{F_{kj}} + 1 + \sum \gamma_{B_{-k} \wedge F_{kj}}} \text{ if } F_{kj}^{t+1} \subseteq F_{kj}, \text{ where } \theta \text{ is defined as in}$$

Definition 2.1. \square

As with RRC-Dirichlet dynamics, I believe that VSC-Dirichlet dynamics admits of plausible actual play and representative interpretations. Allowing the partitions to be sporadically variable reflects the

intuition that deliberators may sometimes fail to observe all of the information the signals would ordinarily provide them, perhaps because they are somewhat distracted from time to time or because some sort of “noise” interferes with the signal.

In general, convergence results for Dirichlet dynamics are very hard to prove, and it is known that the dynamics does not converge in all cases.¹⁷ In the remainder of this section, I will present a convergence result for the RRC-Dirichlet dynamics. Given certain appropriate bounds on the deliberators’ memories, in the special case of 2×2 games with two Nash equilibrium conventions, RRC-deliberation converges almost surely to a correlated equilibrium convention.

Every strict Nash equilibrium of a game is such that, if a Dirichlet deliberator k ’s opponents play their end of this equilibrium often enough, then k will want to play his end of the equilibrium. More precisely, if A^* is a strict Nash equilibrium of a game Γ , then for each player k , there is a value $\beta_{kA_k^*}$ such that if k ’s opponents play A_{-k}^* at least $\beta_{kA_k^*}$ times consecutively, then $E(u_k(A_k^*)) > E(u_k(A_{kj}))$ for all $A_{kj} \neq A_k^*$, so that A_k^* becomes k ’s best response given k ’s beliefs. $\beta_{kA_k^*}$ is called player k ’s *transition frequency* for A_k^* . Since the number of players and the number of the game’s strict Nash equilibria are finite, one can set

$$\beta_{A^*} = \max\{\beta_{kA_k^*} \mid k \in N\} \text{ and } \beta_{\Gamma} = \max\{\beta_{A^*} \mid A^* \text{ is a strict Nash equilibrium of } \Gamma\},$$

called the *absorption frequencies* of the equilibrium A^* and of the strict equilibria of Γ , respectively.

Lemma 3.1. Let Γ be a 2×2 game Γ with two strict Nash equilibria, let the players be RRC-Dirichlet deliberators with a common partition \mathfrak{H} of Ω , and let $H_k \in \mathfrak{H}$ be given. If each deliberator k recalls at least $\min\{t, \beta_{\Gamma}\}$ rounds, and at most $2\beta_{\Gamma}$ rounds, of the H_k -subsequence of deliberation at every round, then for each round $t \geq \beta$ of the H_k -subsequence of RRC-Dirichlet deliberation there is a positive probability, bounded from below, that the deliberators will coordinate on a strict equilibrium of Γ .

PROOF. Without loss of generality, let $(A_1, A_1), (A_2, A_2)$ be the strict Nash equilibria of this 2×2 game, and consider only rounds of deliberation in which $\omega^t \in H_k$ occurs. Note that, for any $N \geq 1$, the H_k -subsequence of plays $(A^t)_{t=N}^{N+\beta}$ determined by Dirichlet deliberation with full memory is fixed, and must satisfy one of the following:

- (1) For some value η , $1 \leq \eta \leq \beta_{\Gamma}$, the deliberators coordinate on either (A_1, A_1) or (A_2, A_2) on round $N + \eta$,
- (2) The deliberators play the nonequilibrium strategy combinations $(A_1, A_2), (A_2, A_1)$ at each round of deliberation between $N + 1$ and $N + \beta_{\Gamma}$ inclusive, and switch from one nonequilibrium strategy combination to the other at least once.

If, for instance, the deliberators miscoordinate on (A_2, A_1) for the first $\beta_{1A_1} - 1 \leq \beta_{2A_2} - 1$ rounds of deliberation after round N , then on the $N + \beta_{1A_1}$ th round, Player 1 will switch strategies to A_1 , in which case either they will coordinate on (A_1, A_1) , or player 2 will also switch, so that they will miscoordinate on

¹⁷See Monderer and Shapley (1993) for a summary of most of the known convergence theorems for fictitious play dynamics, all of which generalize straightforwardly to Dirichlet dynamics. Shapley (1964) gives an example of a game in which the probability distributions of fictitious play deliberators cycle around the Nash equilibrium without ever reaching it. Richards (1993) proves that for a certain class of noncooperative games, the fictitious play dynamics will be *chaotic*.

(A1, A2).

Now apply the above reasoning with $N = \beta_\Gamma$. If (1) is the case, then at round t there is a positive probability that $\eta_1^t = \eta_2^t = \beta_\Gamma + \eta$, so that players coordinate on round t . If (2) is the case, then if τ , $\beta_\Gamma < \tau \leq 2\beta_\Gamma$ is a round at which deliberators with full memory would switch from one nonequilibrium strategy combination to the other, then at round t there is a positive probability that $\eta_1 = \beta_\Gamma + \tau$ and $\eta_2 = \beta_\Gamma + \tau - 1$. Player 1 recalls enough rounds of deliberation to switch her strategy, say from A2 to A1. However, Player 2 does not recall enough rounds of deliberation to switch his strategy from A1 to A2. This is because at the $\beta_\Gamma + \tau - 1$ st round of deliberation with full memory, the frequency of Player 1's playing A1 over the sequence $(A_1^{\beta_\Gamma}, \dots, A_1^{\beta_\Gamma + \tau - 1})$ is great enough to make A1 Player 2's best response, so that if Player 2 can only recall the most recent $\beta_\Gamma + \tau - 1$ rounds at round $\beta_\Gamma + \tau$, then Player 2 will compute the relative frequency of Player 1's playing A1 over the sequence $(A_1^{\beta_\Gamma + 1}, \dots, A_1^{\beta_\Gamma + \tau - 1}, A1)$, which must be at least as high as the frequency of A1's over the sequence $(A_1^{\beta_\Gamma}, \dots, A_1^{\beta_\Gamma + \tau - 1})$. Hence, if $\eta_2 = \beta_\Gamma + \tau - 1$, then A1 is Player 2's best reply, so the players will coordinate on (A1, A1). Hence, we see that if (2) is the case, then there is a positive probability that the deliberators will coordinate.

Hence, in either case, there is a positive probability p' that at round t , the deliberators will coordinate on either (A1, A1) or (A2, A2), and $p' \geq p$ for some $p > 0$, since the range of η_1 and η_2 is finite. \square

The proof of the following lemma is inspired by the proof of Theorem 1 in Young (1993).

Lemma 3.2. If there is a value T such that for every $t \geq T$, there is a positive probability bounded from below that RRC-Dirichlet deliberators with the information and memory restrictions of Lemma 3.1 coordinate on a strict equilibrium of Γ when $\omega^t \in H_k \in \mathfrak{H}$, then the H_k -subsequence of RRC-Dirichlet dynamics converges almost surely to a strict equilibrium of Γ .

PROOF. By hypothesis, for each $t \geq T$, there is a positive probability $p \geq p^* > 0$ that the following occur together:

- (1) $\omega^t \in H_k$ occurs and random recall deliberators coordinate on coordinate on a strict equilibrium A^* of Γ .
- (2) For each of the next $\beta_\Gamma + 1$ -rounds, $\omega^t \in H_k$,
- (3) From rounds t through rounds $t + \beta_\Gamma$, each deliberator recalls the previous $\eta_k^t + m$ rounds of deliberation, $1 \leq m \leq \beta_\Gamma$, and
- (4) On the $t + \beta_\Gamma + 1$ st round, they forget all but the most recent β_Γ rounds.

If (1)-(3) occur, then by the Absorption Theorem the deliberators play A^* at each of the $t + m$ rounds, $0 \leq m \leq \beta_\Gamma$. If (4) also occurs, then they will forget having ever played anything other than A^* when H_k occurs over β_Γ rounds, which "locks" the deliberators onto A^* when H_k occurs from then on.

Now the probability that, for any given run of $\beta_\Gamma + 2$ rounds of deliberation after T , the deliberators do *not* coordinate on a strict equilibrium and then lock onto this equilibrium permanently when H_k occurs as described in the previous paragraph is at most $1 - p^*$. Hence the probability that the H_k -subsequence fails to converge to a strict equilibrium over $r(\beta_\Gamma + 2)$ rounds of deliberation after round T is at most $(1 - p^*)^r \rightarrow 0$ as $r \rightarrow \infty$. \square

Proposition 3.3. In a 2×2 game Γ with two Nash equilibrium conventions, if the players are RRC-Dirichlet deliberators with the information and memory restrictions of Lemma 3.1, then deliberation converges almost surely to a convention.

PROOF. Let H_j be a cell of \mathcal{H} . Since the Nash equilibrium conventions are strict equilibria, the hypotheses of Lemma 3.1 and Lemma 3.2 are satisfied. Hence, the subsequence (f^{t_u}) of RRC-Dirichlet deliberation defined by t_u such that $\omega^{t_u} \in H_j$ converges almost surely to one of the strict Nash equilibria of Γ . Without loss of generality, let this equilibrium be $f(\omega_j) = (A_i, A_j)$ for $\omega_j \in H_j$. Then, $(f^{t_u}(\omega^{t_u})) \rightarrow f(\omega^{t_u}) = (A_i, A_j)$ almost surely if $\omega^t \in H_j$ as $t \rightarrow \infty$, and by hypothesis

$$(1) \quad E(u_k \circ f | \mathcal{H})(\omega^t) > E(u_j \circ (g_j, f_{-j}) | \mathcal{H})(\omega^t)$$

Since $H_j \in \mathcal{H}$ was chosen arbitrarily, the entire system (f^t) converges almost surely to a correlated equilibrium f satisfying (1), which by definition is a convention. \square

Proposition 3.3 is analogous to Miyasawa's (1961) theorem, which guarantees that fictitious play converges to a Nash equilibrium in an arbitrary 2×2 game. However, the almost sure limit of RRC-Dirichlet deliberation in a 2×2 game with two Nash equilibrium conventions is a strict correlated equilibrium, while the limit of fictitious play in an arbitrary 2×2 game need not be a strict equilibrium.

One may regard also Proposition 3.3 as complimentary to Theorem 1 in Young (1993). Young proves that, given certain restrictions on the deliberators' memories, stochastic adaptive dynamics converges almost surely to a strict Nash equilibrium if the game is *weakly acyclic*. The class of weakly acyclic games is much more general than the class of 2×2 games with two conventions. One natural extension of Young's model would be to introduce external signals, as in RRC-Dirichlet deliberation, in such a way that stochastic adaptive deliberators can converge to conventions corresponding to a correlated equilibria. Such a generalization of stochastic adaptive dynamics would yield a more general convergence theorem than Proposition 3.3,¹⁸ but at a certain price. The only plausible interpretation of the stochastic adaptive dynamics, and the only interpretation Young gives, is the representative interpretation. An actual play interpretation of stochastic adaptive dynamics would rely upon the counterintuitive assumption that an individual's memory at any given time is a random sample of the events she has experienced. Moreover, the proof of Young's result depends crucially upon the possibility that at certain rounds of deliberation, the deliberators recall information from disjoint samples of previous plays. Since RRC-Dirichlet deliberators all recall subsets of the most recent plays, there is always some overlap in their samples, and consequently convergence results for the RRC-Dirichlet dynamics are much harder to prove than analogous convergence results for stochastic adaptive dynamics.

Acknowledgements

I thank Vince Crawford and Brian Skyrms for their many helpful comments on earlier versions of this essay.

¹⁸For the proof of such a result, see Vanderschraaf (forthcoming).

Table 1. Battle of the Sexes Played By RRC-Dirichlet Deliberators*

t	ω^t	A_1	A_2	$\kappa_{H_{1j}}^t$	$\kappa_{H_{2j}}^t$	$p_1^t(A_2 H_{1j})$	$p_2^t(A_2 H_{2j})$
1	H	A_1	A_2	0	0	$\frac{1}{2}$	$\frac{1}{2}$
2	T	A_1	A_2	0	0	$\frac{1}{2}$	$\frac{1}{2}$
3	H	A_2	A_1	1	1	$\frac{1+1}{1+2} = \frac{2}{3}$	$\frac{0+1}{1+2} = \frac{1}{3}$
4	T	A_2	A_1	1	1	$\frac{1+1}{1+2} = \frac{2}{3}$	$\frac{0+1}{1+2} = \frac{1}{3}$
5	H	A_1	A_2	2	2	$\frac{1+1}{2+2} = \frac{1}{2}$	$\frac{1+1}{2+2} = \frac{1}{2}$
6	T	A_1	A_2	2	2	$\frac{1+1}{2+2} = \frac{1}{2}$	$\frac{1+1}{2+2} = \frac{1}{2}$
7	H	A_2	A_1	3	3	$\frac{2+1}{3+2} = \frac{3}{5}$	$\frac{1+1}{3+2} = \frac{2}{5}$
8	T	A_2	A_1	3	3	$\frac{2+1}{3+2} = \frac{3}{5}$	$\frac{1+1}{3+2} = \frac{2}{5}$
9	H	A_1	A_2	4	4	$\frac{2+1}{4+2} = \frac{1}{2}$	$\frac{2+1}{4+2} = \frac{1}{2}$
10	T	A_1	A_2	4	4	$\frac{2+1}{4+2} = \frac{1}{2}$	$\frac{2+1}{4+2} = \frac{1}{2}$
11	H	A_1	A_2	5	5	$\frac{3+1}{5+2} = \frac{4}{7}$	$\frac{2+1}{5+2} = \frac{3}{7}$
12	T	A_1	A_2	5	5	$\frac{3+1}{5+2} = \frac{4}{7}$	$\frac{2+1}{5+2} = \frac{3}{7}$
13	H	A_2	A_1	6	6	$\frac{4+1}{6+2} = \frac{5}{8}$	$\frac{2+1}{6+2} = \frac{3}{8}$
14	T	A_2	A_1	6	6	$\frac{4+1}{6+2} = \frac{5}{8}$	$\frac{2+1}{6+2} = \frac{3}{8}$
15	H	A_1	A_2	7	7	$\frac{4+1}{7+2} = \frac{5}{9}$	$\frac{3+1}{7+2} = \frac{4}{9}$
16	T	A_1	A_2	7	7	$\frac{4+1}{7+2} = \frac{5}{9}$	$\frac{3+1}{7+2} = \frac{4}{9}$
17	H	A_1	A_1	7	6	$\frac{4+1}{7+2} = \frac{5}{9}$	$\frac{2+1}{6+2} = \frac{1}{3}$
18	T	A_2	A_2	6	7	$\frac{4+1}{6+2} = \frac{5}{8}$	$\frac{3+1}{7+2} = \frac{4}{9}$

* A_k denotes the strategy that player k plays at round t . $\kappa_{H_{kj}}^t$ denotes the number of rounds of play at which H_{kj} has occurred that player k recalls at round t .

Table 2. Battle of the Sexes Played By VSC-Dirichlet Deliberators*

t	ω^t	H_1^t	H_2^t	A_1	A_2	$p_1^t(A_1 H_1^t)$	$p_1^t(A_2 H_1^t)$	$p_2^t(A_1 H_1^t)$	$p_2^t(A_2 H_1^t)$
1	ω_1	$\{\omega_1\}$	$\{\omega_1\}$	A1	A2	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$
2	ω_4	$\{\omega_4\}$	$\{\omega_4\}$	A1	A2	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$
3	ω_1	$\{\omega_1\}$	$\{\omega_1\}$	A2	A1	$\frac{0+1}{1+2} = \frac{1}{3}$	$\frac{1+1}{1+2} = \frac{2}{3}$	$\frac{1+1}{1+2} = \frac{2}{3}$	$\frac{0+1}{1+2} = \frac{1}{3}$
4	ω_4	$\{\omega_4\}$	$\{\omega_4\}$	A2	A1	$\frac{0+1}{1+2} = \frac{1}{3}$	$\frac{1+1}{1+2} = \frac{2}{3}$	$\frac{1+1}{1+2} = \frac{2}{3}$	$\frac{0+1}{1+2} = \frac{1}{3}$
5	ω_1	$\{\omega_1\}$	$\{\omega_1\}$	A1	A2	$\frac{1+1}{2+2} = \frac{1}{2}$	$\frac{1+1}{2+2} = \frac{1}{2}$	$\frac{1+1}{2+2} = \frac{1}{2}$	$\frac{1+1}{2+2} = \frac{1}{2}$
6	ω_4	$\{\omega_4\}$	$\{\omega_4\}$	A1	A2	$\frac{1+1}{2+2} = \frac{1}{2}$	$\frac{1+1}{2+2} = \frac{1}{2}$	$\frac{1+1}{2+2} = \frac{1}{2}$	$\frac{1+1}{2+2} = \frac{1}{2}$
7	ω_2	$\{\omega_2\}$	$\{\omega_2\}$	A1	A2	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$
8	ω_1	$\{\omega_1\}$	$\{\omega_1, \omega_2\}$	A2	A1	$\frac{1+1}{3+2} = \frac{2}{5}$	$\frac{2+1}{3+2} = \frac{3}{5}$	$\frac{3+2}{4+4} = \frac{5}{8}$	$\frac{1+2}{4+4} = \frac{3}{8}$
9	ω_2	$\{\omega_1, \omega_2\}$	$\{\omega_2\}$	A1	A1	$\frac{2+2}{5+4} = \frac{4}{9}$	$\frac{3+2}{5+4} = \frac{5}{9}$	$\frac{1+1}{1+2} = \frac{2}{3}$	$\frac{0+1}{1+2} = \frac{1}{3}$
10	ω_2	$\{\omega_2\}$	$\{\omega_2\}$	A2	A1	$\frac{0+1}{1+2} = \frac{1}{3}$	$\frac{1+1}{1+2} = \frac{2}{3}$	$\frac{2+1}{2+2} = \frac{3}{4}$	$\frac{0+1}{2+2} = \frac{1}{4}$
11	ω_3	$\{\omega_3\}$	$\{\omega_3\}$	A1	A2	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$
12	ω_1	$\{\omega_1\}$	$\{\omega_1\}$	A1	A1	$\frac{2+1}{4+2} = \frac{1}{2}$	$\frac{2+1}{4+2} = \frac{1}{2}$	$\frac{2+1}{3+2} = \frac{3}{5}$	$\frac{1+1}{3+2} = \frac{2}{5}$
13	ω_4	$\{\omega_3, \omega_4\}$	$\{\omega_4\}$	A2	A1	$\frac{1+2}{4+4} = \frac{3}{8}$	$\frac{3+2}{4+4} = \frac{5}{8}$	$\frac{2+1}{3+2} = \frac{3}{5}$	$\frac{1+1}{3+2} = \frac{2}{5}$
14	ω_3	$\{\omega_3\}$	$\{\omega_4, \omega_3\}$	A2	A2	$\frac{0+1}{1+2} = \frac{1}{3}$	$\frac{1+1}{1+2} = \frac{2}{3}$	$\frac{3+2}{5+4} = \frac{5}{9}$	$\frac{2+2}{5+4} = \frac{4}{9}$
15	ω_3	$\{\omega_3\}$	$\{\omega_3\}$	A2	A1	$\frac{0+1}{2+2} = \frac{1}{4}$	$\frac{2+1}{2+2} = \frac{3}{4}$	$\frac{1+1}{1+2} = \frac{2}{3}$	$\frac{0+1}{1+2} = \frac{1}{3}$
16	ω_2	$\{\omega_2\}$	$\{\omega_2\}$	A1	A1	$\frac{1+1}{2+2} = \frac{1}{2}$	$\frac{1+1}{2+2} = \frac{1}{2}$	$\frac{2+1}{3+2} = \frac{3}{5}$	$\frac{1+1}{3+2} = \frac{2}{5}$
17	ω_1	$\{\omega_1\}$	$\{\omega_1\}$	A1	A1	$\frac{3+1}{5+2} = \frac{4}{7}$	$\frac{2+1}{5+2} = \frac{3}{7}$	$\frac{3+1}{4+2} = \frac{2}{3}$	$\frac{1+1}{4+2} = \frac{1}{3}$
18	ω_4	$\{\omega_4\}$	$\{\omega_4\}$	A2	A2	$\frac{1+1}{3+2} = \frac{2}{5}$	$\frac{2+1}{3+2} = \frac{3}{5}$	$\frac{2+1}{4+2} = \frac{1}{2}$	$\frac{2+1}{4+2} = \frac{1}{2}$
19	ω_2	$\{\omega_2\}$	$\{\omega_2\}$	A1	A1	$\frac{2+1}{3+2} = \frac{3}{5}$	$\frac{1+1}{3+2} = \frac{2}{5}$	$\frac{3+1}{4+2} = \frac{2}{3}$	$\frac{1+1}{4+2} = \frac{1}{3}$
20	ω_3	$\{\omega_3\}$	$\{\omega_3\}$	A2	A2	$\frac{1+1}{3+2} = \frac{2}{5}$	$\frac{2+1}{3+2} = \frac{3}{5}$	$\frac{1+1}{2+2} = \frac{1}{2}$	$\frac{1+1}{2+2} = \frac{1}{2}$
21	ω_4	$\{\omega_4\}$	$\{\omega_4\}$	A2	A2	$\frac{1+1}{4+2} = \frac{1}{3}$	$\frac{3+1}{4+2} = \frac{2}{3}$	$\frac{2+1}{5+2} = \frac{3}{7}$	$\frac{3+1}{5+2} = \frac{4}{7}$
22	ω_3	$\{\omega_3\}$	$\{\omega_3\}$	A2	A2	$\frac{1+1}{4+2} = \frac{1}{3}$	$\frac{3+1}{4+2} = \frac{2}{3}$	$\frac{1+1}{3+2} = \frac{2}{5}$	$\frac{2+1}{3+2} = \frac{3}{5}$

* A_1 is the strategy that Amie plays at round t , A_2 is the strategy that Kay plays at round t , H_1^t is the event in \mathfrak{F} that Amie observes at round t , and H_2^t is the event in \mathfrak{F} that Kay observes at round t .

References

- Aumann, Robert. 1974. "Subjectivity and Correlation in Randomized Strategies." *Journal of Mathematical Economics* 1: 67-96.
- Aumann, Robert. 1987. "Correlated Equilibrium as an Expression of Bayesian Rationality." *Econometrica* 55: 1-18.
- Brandenburger, Adam and Dekel, Eddie. 1987. "Rationalizability and Correlated Equilibria", *Econometrica* 55: 1391-1402.
- Carnap, Rudolf. 1980. "A Basic System of Inductive Logic, Part 2", in *Studies in Inductive Logic and Probability*, ed. Richard C. Jeffrey. Berkeley: University of California Press: 7-155.
- G. W. Brown. 1951. "Iterative Solutions of Games by Fictitious Play", in *Activity Analysis of Production and Allocation*, ed. T. C. Koopmans. New York: Wiley: 374-376.
- Crawford, Vincent P. and Haller, Hans. 1990. "Learning to Cooperate: Optimal Play in Repeated Coordination Games." *Econometrica* 58: 571-595.
- Hume, David. 1888. *A Treatise of Human Nature*. ed. L. A. Selby-Bigge. Oxford: Clarendon Press.
- Kandori, Michihiro, Mailath, George J. and Rob, Rafael. 1993. "Learning, Mutation, and Long Run Equilibria in Games." *Econometrica* 61,1: 29-56.
- Lewis, David. 1969. *Convention: A Philosophical Study*. Cambridge, Massachusetts: Harvard University Press.
- Luce, R. Duncan and Raiffa, Howard, 1957. *Games and Decisions: Introduction and Critical Survey*, Wiley Publications, New York.
- Miyasawa, K. 1961. "On the Convergence of the Learning Process in a 2×2 Non-Zero-Sum Two Person Game", Economic Research Program, Princeton University, Research Memorandum No. 33.
- Monderer, Dov and Shapley, Lloyd. 1993. "Fictitious Play Property For Games With Identical Interests", working paper, The Technion, Haiffa 32000 Israel and Department of Economics and Department of Mathematics, University of California, Los Angeles.
- Monderer, Dov and Sela, Aner. 1993. "Fictitious Play and No-Cycling Conditions", working paper, The Technion, Haiffa 32000 Israel.
- Richards, Diana. 1993. "Nonconvergence in Full-History Bayesian Decision Rules", working paper, Department of Political Science, University of Minnesota.
- Schelling, Thomas. 1960. *The Strategy of Conflict*. Cambridge, Massachusetts: Harvard University Press.
- Shapley, Lloyd. 1964. "Some Topics in Two-Person Games" in *Advances in Game Theory*, ed. M. Dresher, L. S. Shapley and A. W. Tucker. Princeton: Princeton University Press: 1-28.
- Shubik, Martin. 1982. *Game Theory in the Social Sciences*, The MIT Press, Cambridge, Massachusetts.
- Skyrms, Brian. 1990. *The Dynamics of Rational Deliberation*. Cambridge, Massachusetts: Harvard University Press.
- Skyrms, Brian. 1991. "Inductive Deliberation, Admissible Acts, and Perfect Equilibrium", in *Essays in the Foundations of Rational Decision*, ed. Michael Bacharach and Susan Hurley. Blackwell: Oxford: 220-241.
- Vanderschraaf, Peter and Skyrms, Brian. 1993. "Deliberational Correlated Equilibria." Technical Report Series, Report No. MBS 93-07, Irvine Research Unit in Mathematical Behavioral Sciences.
- Vanderschraaf, Peter. Forthcoming. *A Study in Inductive Deliberation*. Doctoral dissertation, University of California, Irvine.
- von Neumann, John and Morgenstern, Oskar. 1944. *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.
- Young, H. Peyton. 1993. "The Evolution of Conventions." *Econometrica*. 61: 57-84.