

DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

SOCIAL LEARNING WITH PRIVATE AND COMMON VALUES

Jacob Goeree
University of Amsterdam

Thomas Palfrey
California Institute of Technology

Brian Rogers
California Institute of Technology



SOCIAL SCIENCE WORKING PAPER 1187

November 2003

Social Learning with Private and Common Values¹

Jacob K. Goeree, Thomas R. Palfrey, and Brian W. Rogers²

December 9, 2003

Abstract

We consider an environment where individuals sequentially choose among several actions. The payoff to an individual depends on her action choice, the state of the world, and an idiosyncratic, privately observed preference shock. Under weak conditions, as the number of individuals increases, the sequence of choices always reveals the state of the world. This contrasts with the familiar result for pure common-value environments where the state is *never* learned, resulting in herds or informational cascades. The medium run dynamics to convergence can be very complex and non-monotone: posterior beliefs may be concentrated on a wrong state for a long time, shifting suddenly to the correct state.

JEL classification numbers: D82, D83

Key words: social learning, information cascades, herd behavior

¹Financial support from the National Science Foundation NSF (SBR-0098400 and SES-0079301) and the Alfred P. Sloan Foundation is gratefully acknowledged. We thank Richard McKelvey posthumously for insights and conjectures about information aggregation that helped shape our thinking about the problem. We also acknowledge helpful comments from Kim Border, Tilman Börgers, Bogachen Celen, Luis Corchon, Matthew Jackson and seminar participants at University College London, UCLA, NYU, The University of Arizona, Universitat Autònoma de Barcelona, University of Edinburgh, 2003 annual meeting of ESA in Pittsburgh, the 2003 Malaga Workshop on Social Choice and Welfare Economics, the 2003 SAET meetings in Rhodos, and the 2003 ESSET meetings in Gerzensee.

²Goeree: CREED, Faculteit voor Economie en Econometrie, Universiteit van Amsterdam, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands. Palfrey and Rogers: Division of Humanities and Social Sciences, California Institute of Technology, Mail code 228-77, Pasadena, CA 91125, USA.

1. Introduction

Consider a traveller deciding at which of many restaurants to dine in Barcelona. Reliable rumor has it that a famous guitarist will perform after dinner at one restaurant for the exclusive benefit of the diners. The traveller has gathered some useful information about where the guitarist will likely perform, but this information is inconclusive. Before deciding, however, the traveller is able to observe the dining choices of others who have acquired possibly different information regarding the performance location, and have also observed previous diners' choices. How should the traveller decide, and how likely are restaurant goers to learn where the guitarist will perform?

Following Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992), the standard analysis of this situation assumes all restaurants are identical so that the only payoff relevant difference between restaurants is whether the guitarist will be there. In this case, the public information revealed by the choices of the first several individuals dominates the private information of subsequent decision-makers. Once this occurs, later choices are made independently of privately held information, creating an *informational cascade*. In a cascade, beliefs about the performance location become “stuck” since decisions do not reveal any privately held information. Furthermore, once a cascade starts, since individuals share identical preferences, they will all choose the restaurant indicated by the first few decisions, a phenomenon known as a *herd*. Since beliefs in a cascade are based on the imperfect information of a few initial decision-makers, restaurant goers never learn the true location and with positive probability they will *all* make the wrong choice. Cascades will occur despite the wealth of information in the economy, which, if aggregated, would almost surely reveal the location.¹

In this paper we show that this negative conclusion is not robust to relaxing the assumption that all restaurants are identical. Stated differently, restaurants are assumed to be homogeneous goods so that dining choices are driven entirely by a commonly-valued “vertical” or quality

¹Laboratory data offer some support for the theory that cascades will arise, at least in the short run. However, there is also evidence indicating that such cascades may not persist indefinitely. See, e.g., Anderson and Holt (1997) and Goeree, Palfrey, and Rogers (2003).

component, i.e. the performer. In reality, however, restaurants are also “horizontally” differentiated with some specializing in Catalan cuisine, others in Spanish, French, tapas, seafood, etc. In this more realistic setting, individuals’ choices may also (partly) reflect their idiosyncratic tastes and common beliefs do not necessarily translate into identical decisions.² It is difficult (if not impossible) to think of any decision-making environment where everyone’s preferences are perfectly aligned and only a single common-value element plays a role.³

We show that information will almost surely be revealed if there are sufficiently diverse idiosyncratic tastes. At first glance this result is paradoxical. The direct effect of introducing idiosyncratic tastes is to dilute the information content of observed actions, which would seem to hinder information aggregation. As we demonstrate, however, it is precisely this dilution of inferences from observations that allows observational learning to continue to the limit, rather than becoming stuck. In our example, if consumers have heterogeneous tastes over restaurants, the guitarist’s location becomes publicly known with arbitrarily high probability as the number of observed dining choices grows. No matter how strong the evidence becomes that the guitarist will be at a particular restaurant, there may still be individuals who choose a different restaurant. Because of this possibility, there is something to be learned from every decision. We provide a general set of conditions on payoff functions and on the information structure such that complete learning is achieved in the limit. Moreover, due to the idiosyncratic taste differences built into our model, different actions may be selected even when beliefs have (almost) converged. This contrasts sharply with the existing literature where the emphasis has been mostly on information cascades and herding.

Other authors have obtained positive convergence results under different assumptions. Lee (1993) provides conditions under which full learning must occur in a pure common-value model. His result relies on a sufficiently rich action space, so that actions can perfectly reveal signals.

²For example, if the common belief indicates the guitarist’s location is likely to be a steak house, vegetarians may opt to not follow the crowd.

³Even in stock markets, transactions are based not only on beliefs about earnings (the common-value component) but may also depend on investors’ risk attitudes, liquidity needs, portfolio balancing, and other idiosyncratic attributes.

Smith and Sorensen (2000) show that full learning occurs when private signals and beliefs are “unbounded,” i.e. when some signals are fully (or arbitrarily close to fully) revealing of the state.⁴

Smith and Sorensen (2000) also admit the possibility of heterogeneous tastes, but in a different manner than the present paper. They assume that preferences differ only along the vertical dimension of common values.⁵ In the context of the above example, this means that all restaurants are identical, except some diners like to hear guitar music, while others prefer to avoid it. Smith and Sorensen (2000) show that in such environments the process of social learning may stop due to “confounded learning.” The idea is that choices are sensitive to privately held information, but others cannot extract any information from them because they cannot distinguish whether the choices were due to signals or preferences.

The main difference between our setup and that of Smith and Sorensen (2000) can be illustrated by the following political science example. Consider a dynamic election campaign, say a primary campaign for President of the United States, where voters vote sequentially and observe past votes (or poll results). Suppose two candidates differ in two dimensions: competence and ideology. Voters are privately informed about their own ideological preferences, say an “ideal point” in the policy space. The environment considered by Smith and Sorensen corresponds to the case where either: (i) voters do not care about ideology, and receive imperfect information about the candidates’ competence, but some voters prefer a less competent candidate; or (ii) voters do not care about candidates’ competence, and receive imperfect information about the candidates’ ideology.⁶ In contrast, our approach integrates the private-value (ideology) and common-value (competence) elements in a way that includes both vertical and horizontal

⁴There is also a more distantly related literature on Bayesian learning in games. The paper in this literature closest to ours is Jackson and Kalai (1997), who identify sufficient conditions under which agents in a population are able to infer the true distribution over player types by observing a history of stage game strategies. Our environment does not satisfy the conditions for their results.

⁵A second restriction that Smith and Sorensen (2000) impose in the latter case is that the set of possible voters’ ideal points is finite. We address this later in the paper.

⁶Confounded learning may occur in case (i) because votes for a particular candidate may have been cast by voters who value (in)competence, and in case (ii) because votes may have been cast by leftists who believe the candidate’s ideology is left, or rightist voters who believe the opposite.

differentiation in preferences. For this example, voters' *private* values would be determined by how close their ideal points are to the candidates' known ideology, as in standard voting models. In addition, voters receive imperfect information about the *common* value component, competence. *Ceteris paribus*, all voters prefer a more competent candidate, but some voters may vote for the less competent candidate for ideological reasons.

This approach permits a tractable analysis when there is an arbitrary number of states, actions, and signals. We thus extend the previous literature, which has focused mainly on the case of two states. The inclusion of multiple states allows a deeper understanding of the dynamics of the underlying stochastic belief process. Convergence of beliefs may seem to imply trivial or monotonic dynamics. Indeed, we show that, *on average*, the weight public beliefs assign to the true state rises. Yet it is not the case that any specific trajectory of beliefs will necessarily follow a monotone dynamic. We construct examples where beliefs are likely to first drift toward an incorrect state and then suddenly shift to converge upon the true state.

The paper is organized as follows. The next section presents the model. The evolution of public beliefs and optimal choice behavior are characterized in Section 3 and Section 4 respectively. Section 5 develops the main convergence result. Section 6 contains a more detailed analysis of the dynamic processes of beliefs and actions. In section 7 we connect our results to the previous literature and discuss generalizations. Section 8 concludes. Most proofs are sketched in the body of the paper, with complete proofs given in an Appendix.

2. The Model

There is a countable set $\mathcal{T} = \{1, 2, \dots\}$ of agents who choose, in sequence, one of several actions. For each $t \in \mathcal{T}$, let $a^t \in \mathcal{A}$ denote agent t 's chosen action where $\mathcal{A} = \{1, \dots, A\}$ is the set of $A > 1$ available actions. Possible states of the world are elements of the set $\mathcal{K} = \{1, \dots, K\}$ where $K > 1$. Agents do not know the state of the world but have common prior beliefs that the state is k with probability P_k^0 . We assume $P_k^0 > 0$ for all k so that all states are possible *a priori*.

2.1. Signals

Each agent receives one conditionally independent private signal about the state. The finite set of signals is denoted $\mathcal{S} = \{1, \dots, S\}$ where $S > 1$. In state k , agent t receives signal $s^t \in \mathcal{S}$ with probability $q(s^t|k)$. This defines a matrix, Q , with elements $q_{sk} \equiv q(s|k)$ for $s \in \mathcal{S}$ and $k \in \mathcal{K}$. We say that signals are *informative* when the probability distributions of signals differ across states, i.e. no two columns of Q are the same. Pairs of states for which this does not hold are obviously indistinguishable.

Definition 1. *Signals are informative when $k \neq k' \in \mathcal{K}$ implies $q_{sk} \neq q_{sk'}$ for some $s \in \mathcal{S}$.*

Furthermore, we say that signals are *bounded* if no single signal can reveal the state of the world: $\text{Prob}(k|s) \equiv \frac{q_{sk}P_k^0}{\sum_{k'=1}^K q_{sk'}P_{k'}^0} < 1$ for all $k \in \mathcal{K}$, $s \in \mathcal{S}$, and all interior P^0 . This inequality holds when at least two elements in every row of Q are strictly positive.

Definition 2. *Signals are bounded if for all $s \in \mathcal{S}$ there exist $k \neq k' \in \mathcal{K}$ such that $q_{sk} > 0$ and $q_{sk'} > 0$.*

Denote the rank of Q by r . The number of rows of Q generally exceeds r , as $2 \leq r \leq \min(K, S)$.⁷ If so, the number of signals can be reduced by combining different elements of \mathcal{S} into r “independent” signals. Pick $r - 1$ signals $\tilde{s}_1, \dots, \tilde{s}_{r-1} \in \mathcal{S}$ that correspond to $r - 1$ linearly independent rows of Q . Let \tilde{s}_r correspond to the union of all other elements of \mathcal{S} . In other words, one can think of receiving an \tilde{s}_r signal as being equivalent to *not* receiving any of the $\tilde{s}_1, \dots, \tilde{s}_{r-1}$ signals. Define the reduced signal set $\tilde{\mathcal{S}} = \cup_{i=1}^r \tilde{s}_i$ and the corresponding matrix of signal probabilities \tilde{Q} with elements $\tilde{q}_{\tilde{s}k} \equiv q(\tilde{s}|k)$ for $\tilde{s} = 1, \dots, r$, $k = 1, \dots, K$. Thus, in a sense, $\tilde{\mathcal{S}}$ contains the same independent information as \mathcal{S} . This construction motivates the following definition.

Definition 3. *Signals are non-redundant if $r = S$.*

⁷Under Assumption 1 the rank of Q is at least 2, since if $r = 1$ all columns of Q would have to be multiples of the first column. Since all columns add up to 1, however, they would have to be identical, which violates the informative-signal assumption below.

When signals are non-redundant we necessarily have $S \leq K$ since $S = r \leq \min(K, S)$.

Assumption 1. *Signals are bounded, informative, and non-redundant.*

The assumption that signals are bounded is not needed for our convergence results but is made to rule out trivial cases where some signals are fully revealing.⁸

2.2. Payoffs

Individual payoffs have two components, a private-value component and a common-value component. The latter depends on the action taken by the agent and the state of the world. If agent t 's action is a and the state of the world is k , the common-value component is given by Π_{ak} , where Π is the common-value payoff matrix. Without loss of generality we choose units such that $0 < \Pi_{ak} < 1$ for all $a \in \mathcal{A}$ and $k \in \mathcal{K}$. Furthermore, we say that Π is *admissible* if $\sum_{k=1}^K x_k = 0$ and $\sum_{k=1}^K (\Pi_{ak} - \Pi_{a'k}) x_k = 0$ for all $a \neq a' \in \mathcal{A}$ imply $x_k = 0$ for all $k \in \mathcal{K}$.

Assumption 2. *Π is admissible with elements strictly between 0 and 1.*

Note that admissibility requires that the gain (or loss) from switching from action a to a' varies across states. In particular, all columns of Π have to be distinct, i.e. states are (common-value) payoff-distinguishable. Admissibility holds generically only when there are at least as many actions as states; we explore the implications of relaxing Assumption 2 in Section 7.2.

Private-value components are assigned by nature as *iid* draws from a commonly known distribution of action-specific payoff disturbances. We assume the distribution of private values has a joint density, denoted by $f(v^t) = f(v_1^t, \dots, v_A^t)$ with corresponding distribution function $F(v^t)$. Let $\text{supp}(f)$ denote the (closure of the) set of points where $f(\cdot)$ is strictly positive.

Assumption 3. $\text{supp}(f) \supseteq [0, 1]^A$.

⁸The case of unbounded signals is somewhat more complicated with a continuous signal space. See Smith and Sorensen (2000).

By choosing action $a \in \mathcal{A}$, agent t receives a private value v_a^t . Hence, in state k , agent t 's payoff of choosing action a is

$$u^t(a|v^t, k) = v_a^t + \Pi_{ak}. \quad (2.1)$$

Summarizing the informational setup, each agent t has a multidimensional type, $\theta^t \in \Theta^t$, consisting of A private-value components and a single common-value signal. The prior state probabilities, P^0 , the conditional signal distributions, Q , the joint distribution of private-values, F , and the common-value payoff matrix, Π , are assumed to be common knowledge among the agents. Moreover, agent t observes the ordered sequence of action choices of all her predecessors, $h^t = \{c^1, \dots, c^{t-1}\} \in \mathcal{H}^t$, but not their types. Agent t 's strategy is a mapping from $\mathcal{H}^t \times \Theta^t$ into the A -dimensional probability simplex, specifying for every history $h^t \in \mathcal{H}^t$ and every type realization $\theta^t \in \Theta^t$, the probabilities with which each action is chosen.

3. Public and Private Beliefs

From the definition of payoffs, it is clear that agent t cares about the history only to the extent that it is informative about the state of the world. Let $P_k^t \equiv \text{Prob}(k|h^t)$ denote the *public belief* that the state is k given the history of choices $h^t = \{a^1, \dots, a^{t-1}\}$. Since $\mathcal{H}^1 = \emptyset$, public beliefs in period 1 coincide with prior beliefs: $P_k^1 = P_k^0 > 0$ for all k . The analysis below is greatly simplified by the observation that, for any given strategy profile, public beliefs in period $t + 1$ are completely determined by agent t 's choice and public beliefs in period t . In other words, the public belief P_k^t serves as a sufficient statistic for the history of choices, h^t .

After observing her private signal s^t , agent t updates her *private belief* that the state is k from P_k^t to $p_k^t(s^t|P^t) \equiv \text{Prob}(k|P^t, s^t)$. Bayes' rule implies

$$p_k^t(s^t|P^t) = \frac{q_{s^t k} P_k^t}{\sum_{k'=1}^K q_{s^t k'} P_{k'}^t}. \quad (3.1)$$

4. Optimal Choice Behavior: Cutpoints

Agent t 's optimal action is a if it yields the highest expected payoff, i.e. if for all $a' \neq a$

$$v_a^t + \sum_{k=1}^K \Pi_{ak} p_k^t(s^t|P^t) > v_{a'}^t + \sum_{k=1}^K \Pi_{a'k} p_k^t(s^t|P^t).$$

To derive agent t 's choice probabilities, we define for each signal s^t and each pair of actions, $a, a' \in \mathcal{A}$ the *cutpoint* $v_{a,a'}^t(s^t|P^t)$:⁹

$$\bar{v}_{a,a'}^t(s^t|P^t) \equiv \sum_{k=1}^K (\Pi_{a'k} - \Pi_{ak}) p_k^t(s^t|P^t). \quad (4.1)$$

Therefore, agent t 's optimal action is a if and only if the difference between t 's private value disturbances for a and any other action a' exceeds the cutpoint defined by (4.1). Hence, conditional on public beliefs at t , P^t and private signal s^t , agent t 's choice probabilities are given by:

$$C_a^t(s^t|P^t) \equiv \text{Prob}(a|P^t, s^t) = \int_{v_a = \max_{a' \in \mathcal{A}} (v_{a'} + \bar{v}_{a,a'}^t(s^t|P^t))} dF(v), \quad (4.2)$$

for each $a \in \mathcal{A}$. We first establish that, under Assumptions 2 and 3, all actions have a positive chance of being selected.

Lemma 1. *The choice probabilities $C_a^t(s^t|P^t)$ are strictly positive and strictly decreasing in the cutpoints $\bar{v}_{a,a'}^t(s^t|P^t)$ for all $a' \neq a \in \mathcal{A}$ and $s^t \in \mathcal{S}$.*

Proof. See Appendix.

5. Convergence

Here we investigate the evolution of agents' beliefs and corresponding choices. We define state-dependent transition probabilities $T_{ka}^t(P^t) \equiv \text{Prob}(a|P^t, k)$ for $k \in \mathcal{K}$ and $a \in \mathcal{A}$, which

⁹Note that $\bar{v}_{a,a}^t(s^t|P^t) = 0$, $\bar{v}_{a,a'}^t(s^t|P^t) = -\bar{v}_{a',a}^t(s^t|P^t)$, and $\bar{v}_{a,a'}^t(s^t|P^t) + \bar{v}_{a',a''}^t(s^t|P^t) = \bar{v}_{a,a''}^t(s^t|P^t)$ for all $a, a', a'' \in \mathcal{A}$.

are the probabilities of observing action a when the state is k , given the history h^t . Note that

$$T_{ka}^t(P^t) = \sum_{s=1}^S q_{sk} C_a^t(s|P^t), \quad (5.1)$$

and $\sum_{a=1}^A T_{ka}^t(P^t) = 1$. Furthermore, let $P_k^{t+1}(a) \equiv \text{Prob}(k|P^t, a)$ denote the updated public belief that the state is state k if agent t chooses action a following history h^t . By Bayes' rule

$$P_k^{t+1}(a) = \frac{T_{ka}^t(P^t)P_k^t}{\sum_{k'=1}^K T_{k'a}^t(P^t)P_{k'}^t}. \quad (5.2)$$

We next establish that no alternative is ruled out in finite time.

Lemma 2. $P_k^t > 0$ for all $k \in \mathcal{K}$ and $t \in \mathcal{T}$.

Proof. See Appendix.

Since all actions have a strictly positive chance of being chosen (Lemma 1), if, for some t , $P_k^{t+1}(a) = P_k^t$, for all $k \in \mathcal{K}$, $a \in \mathcal{A}$, then $P_k^\tau(a) = P_k^t$, for all $k \in \mathcal{K}$, $a \in \mathcal{A}$ and for all $\tau > t$. If this happens, then we say the process of learning stops at time t .

Definition 4. *Learning stops at t if and only if $P_k^{t+1}(a) = P_k^t$, for all $k \in \mathcal{K}$, $a \in \mathcal{A}$.*

Since $P_k^t > 0$ by Lemma 2, equation (5.2) implies that learning stops if and only if $T_{ka}^t(P^t) = \sum_{k'=1}^K T_{k'a}^t(P^t)P_{k'}^t$, for all $k \in \mathcal{K}$, $a \in \mathcal{A}$, which in turn holds if and only if:

$$T_{ka}^t(P^t) = T_{1a}^t(P^t), \quad \forall k \in \mathcal{K}, a \in \mathcal{A}. \quad (5.3)$$

The next three Lemmas establish that as long as beliefs are non-degenerate, learning does not stop. Together with Lemma 2 this implies that learning does not stop in finite time. To prove this we show an intermediate result that learning never stops unless, for some public belief P_k^t , choice probabilities are the same regardless of the signal observed by agent t .

Definition 5. *Choice probabilities are independent of signals at P^t if and only if $C_a^t(s|P^t) =$*

$C_a^t(s'|P^t)$ for all $s, s' \in \mathcal{S}$ and for all $a \in \mathcal{A}$.

Lemma 3. *Learning stops at t if and only if choice probabilities are independent of signals at P^t .*

Proof. (Only if) Using (5.1) condition (5.3) can be rewritten as

$$\sum_{s=1}^S C_a^t(s|P^t)(q_{sk} - q_{s1}) = 0, \forall k \in \mathcal{K}, a \in \mathcal{A}.$$

The S -dimensional vector $e = (1, 1, \dots, 1)$ is a null (left) eigenvector of the matrix $Q_1 = q_{sk} - q_{s1}$, with $k \neq 1$ and $s = 1, \dots, S$. Assumption 1 implies that Q_1 has rank at least $S - 1$, so e is the unique null eigenvector. Hence $C_a^t(s|P^t) = C_a^t(s'|P^t)$ for all $s, s' \in \mathcal{S}$ and for all $a \in \mathcal{A}$.

(If) Suppose at P^t choice probabilities are independent of signals, so $C_a^t(s|P^t) = C_a^t(s'|P^t) = \bar{C}$ for all $s, s' \in \mathcal{S}$ and for all $a \in \mathcal{A}$. Then $T_{ka}^t = \bar{C}$ for all $k \in \mathcal{K}$ and substituting into (5.2) gives $P_k^{t+1}(a) = P_k^t$ for all $k \in \mathcal{K}$ and $a \in \mathcal{A}$, so learning stops at t . *Q.E.D.*

Definition 6. *Cutpoints are independent of signals if and only if $\bar{v}_{a,a'}^t(s|P^t) = \bar{v}_{a,a'}^t(s'|P^t)$ for all $s, s' \in \mathcal{S}$ and for all $a \neq a' \in \mathcal{A}$.*

Lemma 4. *Choice probabilities are independent of signals at P^t if and only if cutpoints are independent of signals.*

Proof. Here we illustrate the proof for the specific case $A = K = 3$ and $S = 2$. The proof for the general case can be found in the Appendix.

(Only if) Suppose, in contradiction, cutpoints are not independent of signals: $\bar{v}_{1,2}^t(1|P^t) > \bar{v}_{1,2}^t(2|P^t)$. By Lemma 1, $C_1^t(1|P^t) = C_1^t(2|P^t)$ then implies $\bar{v}_{1,3}^t(1|P^t) < \bar{v}_{1,3}^t(2|P^t)$. Since $\bar{v}_{2,3}^t(s|P^t) = \bar{v}_{1,3}^t(s|P^t) - \bar{v}_{1,2}^t(s|P^t)$ for all s we have $\bar{v}_{2,3}^t(1|P^t) < \bar{v}_{2,3}^t(2|P^t)$. Furthermore, since $\bar{v}_{2,1}^t(s|P^t) = -\bar{v}_{1,2}^t(s|P^t)$ for all s we have $\bar{v}_{2,1}^t(1|P^t) < \bar{v}_{2,1}^t(2|P^t)$. But recall from Lemma 1 that choice probabilities are strictly decreasing in cutpoints, so $C_2^t(1|P^t) > C_2^t(2|P^t)$, the desired contradiction. A similar reasoning rules out $\bar{v}_{1,2}^t(1|P^t) < \bar{v}_{1,2}^t(2|P^t)$, so $\bar{v}_{1,2}^t(1|P^t) = \bar{v}_{1,2}^t(2|P^t)$. Finally, applying the same steps shows $\bar{v}_{1,3}^t(1|P^t) = \bar{v}_{1,3}^t(2|P^t)$ and $\bar{v}_{2,3}^t(1|P^t) = \bar{v}_{2,3}^t(2|P^t)$.

(If) Suppose cut-point are independent of signals. Then by (4.2), the choice probabilities are independent of signals. *Q.E.D.*

Definition 7. *Public beliefs are degenerate at t if $P_k^t = 1$ for some k .*

Lemma 5. *Cutpoints are independent of signals if and only if public beliefs are degenerate.*

Proof. (Only if) Independence of cutpoints with respect to signals implies:

$$\sum_{k=1}^K (\Pi_{a'k} - \Pi_{ak}) p_k^t(s|P^t) = \sum_{k=1}^K (\Pi_{a'k} - \Pi_{ak}) p_k^t(s'|P^t), \quad \forall a, a' \in \mathcal{A}, s, s' \in \mathcal{S},$$

or, equivalently,

$$\sum_{k=1}^K (\Pi_{a'k} - \Pi_{ak}) (p_k^t(s|P^t) - p_k^t(s'|P^t)) = 0, \quad \forall a, a' \in \mathcal{A}, s, s' \in \mathcal{S}.$$

Since $\sum_{k=1}^K (p_k^t(s|P^t) - p_k^t(s'|P^t)) = 0$, Assumption 2 implies that $p_k^t(s|P^t) = p_k^t(s'|P^t)$ for all $k \in \mathcal{K}$ and $s, s' \in \mathcal{S}$ so $p_k^t(s|P^t) = P^t$ for all $k \in \mathcal{K}$ and $s \in \mathcal{S}$. Using (3.1) this can be written as

$$P_k^t q_{sk} = P_k^t \sum_{k'=1}^K q_{sk'} P_{k'}^t. \quad \forall k \in \mathcal{K}, s \in \mathcal{S}. \quad (5.4)$$

Let $\Omega \subseteq \mathcal{K}$ denote the set of states for which the public belief is strictly positive, i.e. $P_k^t > 0$ for all $k \in \Omega$. From (5.4) we conclude that q_{sk} is the same for all $s \in \mathcal{S}$ and all $k \in \Omega$, which violates the informative-signal assumption unless Ω contains only a single element.

(If) Suppose public beliefs are degenerate, then private beliefs, and hence cutpoints, are independent of signals. *Q.E.D.*

Theorem 1. *Learning stops if and only if public beliefs are degenerate.*

Proof. Immediate from Lemmas 3 through 5. *Q.E.D.*

We next show that the learning process in fact converges, and that, in the limit, the public beliefs put all mass on the true state k . In what follows we assume, without loss of generality,

that the true state is $k = 1$. We define the (public) likelihood ratio against the true state $\ell_1^t \equiv (1 - P_1^t)/P_1^t$.

Lemma 6. *The likelihood ratio ℓ_1^t defines a martingale process conditional on state 1.*

Proof. Define $\ell_1^t(a) = -1 + 1/P_1^t(a)$ for $a \in \mathcal{A}$. Conditional on $k = 1$, the transition probabilities are given by $T_{1a}^t(P^t)$ so

$$E(\ell_1^{t+1} | \ell_1^t, k = 1) = \sum_{a=1}^A T_{1a}^t(P^t) \ell_1^t(a) = \sum_{a=1}^A T_{1a}^t(P^t) \left(\frac{1}{P_1^t(a)} - 1 \right). \quad (5.5)$$

Using (5.2) condition (5.5) can be worked out as

$$E(\ell_1^{t+1} | \ell_1^t, k = 1) = -1 + \sum_{k=1}^K \sum_{a=1}^A \frac{T_{ka}^t(P^t) P_k^t}{P_1^t} = -1 + \frac{1}{P_1^t} = \ell_1^t. \quad (5.6)$$

Hence, ℓ_1^t defines a martingale conditional on $k = 1$. *Q.E.D.*

By the Martingale Convergence Theorem (Doob, 1953) there exists a limit random variable to which ℓ_1^t converges almost surely. Hence, $P_1^t = (\ell_1^t + 1)^{-1}$ also converges almost surely to a limit random variable. We are now in position to state our main result.

Theorem 2. *Under Assumptions 1-3, public beliefs converge to the correct state almost surely.*

A brief sketch of the proof follows (see the Appendix for details). By Fatou's lemma the martingale property implies $\lim_{t \rightarrow \infty} E(\ell_1^t) \leq E(\ell_1^0) = \ell_1^0$. By assumption $P_1^0 \neq 0$, so ℓ_1^0 is finite. Hence $\lim_{t \rightarrow \infty} E(\ell_1^t)$ is finite, which implies that public beliefs cannot converge to an incorrect state. By Lemmas 3 to 5, public beliefs also cannot converge to a non-degenerate distribution over states, so they must converge to the true state with probability one.

6. Dynamics

The martingale property of the conditional likelihood ratio against the true state implies that the belief for the true state obeys a sub-martingale. In other words, the expected change

in beliefs for the true state is always non-negative.

Lemma 7. *The public belief P_1^t defines a sub-martingale process conditional on state 1.*

Proof. Since $P_1^{t+1}(a) = (\ell_1^{t+1}(a) + 1)^{-1}$ for all $a \in \mathcal{A}$, the public belief is a strictly convex transformation of the likelihood ratio against the true state. Hence

$$\begin{aligned} E(P_1^{t+1} | P_1^t, k = 1) &= \sum_{a=1}^A T_{1a}^t(P^t) P_1^{t+1}(a) = \sum_{a=1}^A T_{1a}^t(P^t) (\ell_1^{t+1}(a) + 1)^{-1} \\ &\geq \left(\sum_{a=1}^A T_{1a}^t(P^t) (\ell_1^{t+1}(a) + 1) \right)^{-1} = (\ell_1^{t+1} + 1)^{-1} = P_1^t, \end{aligned} \quad (6.1)$$

applying Jensen's inequality and Lemma 6 in the last line.

Q.E.D.

Thus the public beliefs converge to the true state almost surely (Theorem 2) and they always increase in expectation. Together these results may seem to suggest that the belief process tends to follow a monotone dynamic, converging smoothly to the correct beliefs from any non-degenerate prior. However, this is generally not the case. That is, there is a wide range of possible paths that beliefs may follow, many of which are non-monotonic and exhibit sudden jumps.

To illustrate this point, consider the following example where $A = S = K = 2$. We can define the signal technology and private value distributions such that, with high probability, beliefs initially tend towards the incorrect state. Specifically, the information structure is given by $P_1^0 = \frac{1}{2}$, $q_{1,1} = \frac{4}{5}$, and $q_{2,2} = \frac{19}{20}$. Let the common-value payoff matrix be the identity matrix, i.e. $\Pi_{1,1} = \Pi_{2,2} = 1$ and $\Pi_{1,2} = \Pi_{2,1} = 0$, and the private values be normally distributed with the same variance but different means: $v_1^t \sim N(-1.25, .4)$ and $v_2^t \sim N(0, .4)$. In this case, 66% of all individuals choose action 2 even when they believe the true state is 1. In other words, no matter how strong the evidence for state 1, a majority of agents will nevertheless choose action 2 for idiosyncratic reasons.

An important feature of the dynamics in this example is that the change in posterior beliefs after observing either choice is highly asymmetric. Specifically, after observing a choice for 2,




Figure 1: *The asymmetry in private valuations causes most choices to be for alternative 2 (right panel). The few choices for alternative 1, however, are sufficient to eventually tilt beliefs to the correct alternative (left panel).*

the evidence for 2 is only slightly stronger, since successors rationally realize that the decision was likely driven by private value considerations, and not by information regarding the common value. Thus even after observing many choices for 2, the effective sample of “2 signals” inferred from the history is comparatively small. In contrast, after a choice for 1 is observed, the evidence for 1 increases substantially since agents know that it is unlikely that the choice was based on idiosyncratic tastes; instead, it is very likely that $s^t = 1$. This situation is illustrated in Figure 1, which presents a simulation of ten trials of 250 periods each using the above specifications. The belief for state 1 is plotted in the left panel and the corresponding cumulative choice frequencies for alternative 1 in the right panel. Note that many of the belief paths are non-monotonic and exhibit sudden shifts to the correct state following a choice for alternative 1. This example demonstrates that individuals are capable of learning the true state even when few of them actually choose it, as can be seen in the right hand panel of Figure 1, where the cumulative frequencies for alternative 1 are all less than 34%.

When we consider more general settings, where there are several states and alternatives, it is in fact *not* necessary that some individuals choose the superior alternative in order for learning to converge to the correct beliefs. Consider the following situation where $A = S = K = 3$. As

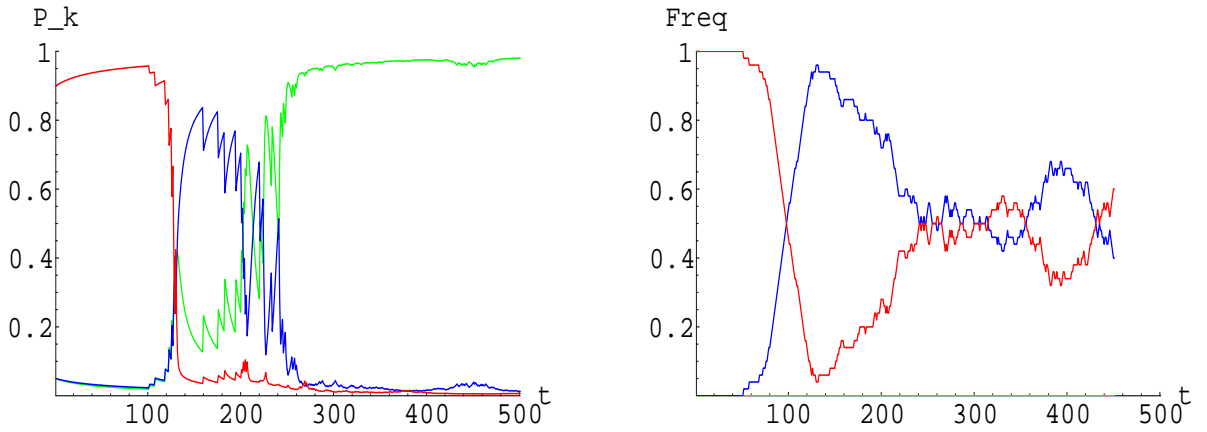


Figure 2: *Prior beliefs favor state 3, then shift suddenly to 2, finally shifting and converging to state 1 (left panel). Initial choices are for 3 but after a spurt of 2 choices, choice frequencies are roughly equal for 2 and 3 (right panel).*

before, prior beliefs are uniform, but the signal technology is now given by

$$Q = \begin{pmatrix} .7 & .15 & .1 \\ .2 & .7 & .2 \\ .1 & .15 & .7 \end{pmatrix}.$$

The common value payoff matrix is again given by the identity matrix: $\Pi = I$, so that actions correspond directly to states. Private values are drawn independently across alternatives with $F_a(v_a^t) = (1 + \exp(-\lambda(v_a^t + \mu_a)))^{-1}$ where $\mu_1 = 2$, $\mu_2 = \mu_3 = 0$, and the scale parameter is $\lambda = 5$. Thus the distribution of private values for the true state 1 is shifted down by 2 units and with a high value of λ it is very unlikely that individuals ever choose action 1. Instead there will be only choices for alternatives 2 and 3. However, from the frequencies of these choices, individuals learn the relative frequencies with which signals for alternatives 2 and 3 are received. Since these relative frequencies differ across all three states, individuals still are able to learn that the true state is 1.

Figure 2 depicts a simulation of public beliefs (left panel) along with a moving average of corresponding choice frequencies (right panel) for 500 periods, where green denotes the true state 1, blue denotes 2, and red corresponds to 3. Prior beliefs put mass 90% on state 3,

with the remaining 10% split equally between states 1 and 2. The right panel shows that the first 100 choices or so are for 3, so that beliefs for state 3 increase from the initial prior. But once some individuals draw high private valuations for alternative 2, they break the herd and beliefs shift very quickly to state 2. However, the continuing presence of choices for alternative 3 causes beliefs to eventually converge upon state 1, even though no single 1-choice is made. Note that after beliefs have (almost) converged, choice frequencies are approximately 50% for alternatives 2 and 3, since conditional on state 1, the common value payoff from each of these alternatives are the same (zero) and the private values are identically distributed. Thus while beliefs converge upon the true state, a herd never arises.

7. Extensions

In this section we explore the consequences of relaxing some key assumptions of the model presented in Section 2. We first discuss how the *iid* assumption of private valuations can be relaxed. We also show why the assumption of full support is essential for complete learning. Second we turn to the common-value payoffs, and show through an example the role of admissibility. We sketch a proof that complete learning still holds generically when the admissibility condition is not met.

7.1. Private Valuations

We first discuss two ways in which one can relax the assumptions on how agents receive their private values for the various alternatives in \mathcal{A} .

7.1.1 Private Values: Full Support

Recall the full support Assumption 3, which requires the support of private values for each alternative to contain the unit interval. In order to connect our results to previous findings, we shall be mainly interested in the alternative case of a finite number of private values, which

of course violates the full support condition. Our main finding is that without full support, beliefs fail to converge upon the true state. However, even in the absence of complete learning, it is never possible for inferences to cause beliefs to converge to an incorrect state (see Smith and Sorensen, 2000).¹⁰

Lemma 8. *For any distribution $F(\cdot)$, if $P_1^0 > 0$, then beliefs a.s. cannot converge to an incorrect state.*

We next discuss two examples of private-value distributions for which “pathological outcomes” arise, i.e. for which learning is incomplete. First, consider the case where the distribution of private values has a single mass point at $v^t = (0, \dots, 0)$, so that choices are driven purely by the common-value component. This case generalizes the basic model discussed in Bikhchandani, Hirshleifer, and Welch (1992).¹¹ With bounded signals there always exist choice histories where the resulting public beliefs are strong enough to outweigh any private information, creating an informational cascade. The set of beliefs at which a cascade occurs, i.e. the *cascade set*, is a connected set containing the boundaries of the beliefs simplex, and has open interior.¹² An example for the case of two states is given in the left panel of Figure 3¹³ where the expected change in public beliefs is shown for all possible values of the public belief. For the case of two states it can be shown that learning stops (i.e. condition (5.3)) if and only if the expected change in beliefs is zero, and whenever beliefs are not in a cascade set the expected change to the true state is strictly positive. As t becomes large, public beliefs eventually enter one of the

¹⁰This result is again a direct consequence of the martingale property of the likelihood ratio (since it holds for any distribution of values) and Fatou’s Lemma, since if beliefs were to converge upon an incorrect state, the likelihood ratio against the true state would explode. Note that even in such cases, however, beliefs must still converge by the Martingale Convergence Theorem, and since it is not possible for beliefs to converge upon any incorrect state, beliefs must settle upon some non-degenerate point that places mass on multiple states.

¹¹We allow for asymmetries in the signal technology and common value payoffs which BHW do not explicitly consider.

¹²We follow Smith and Sorensen (2000) in using this term.

¹³The cascade sets are the regions where the expected change in beliefs vanish. In each panel, $A = S = K = 2$, $q_{1,1} = q_{2,2} = \frac{3}{5}$, and common values are given by the identity matrix. Private value distributions are given by a single mass point at $v^t = (0, 0)$ (upper left), mass points with equal weight at $v^t \in \{(-\frac{4}{3}, 0), (-\frac{2}{3}, 0), (0, 0), (\frac{2}{3}, 0), (\frac{4}{3}, 0)\}$ (upper right), the uniform distribution on $[-\frac{4}{3}, \frac{4}{3}]$ (lower left), and the uniform distribution on $[-\frac{2}{3}, \frac{2}{3}]$ (lower right).




Figure 3: *Expected change in beliefs $E(P_1^{t+1} - P_1^t | P_1^t)$ as a function of P_1^t for the case of two possible states and no private values (upper left panel), a finite number of private values (upper right), continuous private values with full support (lower left), and continuous values without full support (lower right).*

cascade sets with probability one, after which learning stops.

Next consider the case of multiple mass points in the unit square. As in the case of a single mass point, the probability of eventually reaching an interior belief at which a cascade occurs is one. The principle difference arising from introducing more private values to the support is that the cascade set has more regions. In general, with finite support, the cascade set will consist of the union of a finite collection of separated regions (see the upper right panel of Figure 3). Which region beliefs eventually settle into depends on prior beliefs and the specific history of signals and choices.

The intuition behind the failure of complete learning in the above examples may be generalized in a straightforward manner.

Lemma 9. *Beliefs a.s. cannot converge to the correct state if $F(\cdot)$ has finite support.*

To understand this result note that when the support of $F(\cdot)$ is finite, its graph will be a step function with “flat” regions almost everywhere. The expression for the choice probabilities in (4.2) reduces to a finite sum, which is invariant with respect to small variations in the cutpoints $\bar{v}_{a,a'}^t(s^t|P^t)$. Now suppose public beliefs P^t converge to a degenerate point. Then private beliefs $p^t(s^t|P^t)$ must also converge to the same point following any signal s^t , which is a direct consequence of Bayes’ Law. Thus the difference $p^t(s|P^t) - p^t(s'|P^t)$ is small for any pair of signals $s, s' \in \mathcal{S}$, which in turn implies that the cutpoints in (4.1) are nearly independent of signals when beliefs are close to degenerate. Hence, choice probabilities become invariant with respect to signals and choices no longer reveal any privately held information. The fact that learning must then stop follows directly from Lemma 3.

In contrast, with a continuum of types and full support, the process of social learning never stops (see the lower left panel of Figure 3) and thus necessarily reveals the true state. With a continuum of types and incomplete support it is possible to get into a cascade. The lower right panel in Figure 3 illustrates this when private values are uniform on $[-\frac{2}{3}, \frac{2}{3}]$ and the common-value payoff matrix is the identity.

7.1.2 Private Values: IID across Agents

We have assumed that private values are distributed according to a joint distribution $f(\cdot)$ that is independent and identically distributed across agents. Thus while we allowed for correlation across alternatives, we prohibited any correlation in private values across agents. This is in fact not necessary for our results. All of the results still hold if we allow private values to have agent-specific distributions $f^t(\cdot)$ provided that there is some $\epsilon > 0$ such that $\text{supp}(f^t) \supseteq [-\epsilon, 1 + \epsilon]$ for all $t \in \mathcal{T}$.

Relaxing independence of private values across agents presents a more complex argument, but we believe our results are also robust to this change. Suppose there was some positive correlation in private values across agents. Then after observing an action a^t all agents revise

their beliefs in such a way that they believe the subsequent agent is more likely to have a high value of v_a^{t+1} . So if agent $t + 1$ chooses a also, then beliefs on the state are updated less than in the case of independence. This illustrates that correlation may affect the rate of convergence, but should not change the limit results.

7.2. Common Value Payoffs: Admissibility

Recall the definition of admissibility required by Assumption 2. Without this assumption, our results hold generically, but not universally. To see this, first note that admissibility implies $A \geq K$. Consider the following example where $A = S = 2$ and $K = 3$, so admissibility is violated. Let common values and signals be given by

$$\Pi = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & \frac{1}{2} \end{pmatrix}, \quad Q = \begin{pmatrix} \frac{6}{18} & \frac{8}{18} & \frac{13}{18} \\ \frac{12}{18} & \frac{10}{18} & \frac{5}{18} \end{pmatrix}.$$

Notice that Π is not admissible. Furthermore, if public beliefs ever reach $P^t = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, for instance, then learning stops in the sense of Definition 4; i.e., $P_k^{t+1}(a) = P_k^t = \frac{1}{3}$ for each action $a = 1, 2$ and each state $k = 1, 2, 3$. To see this, consider the cut-points $\bar{v}_{21}(s|P^t)$, which, for $s = 1$, equals $1(\frac{2}{9}) - 1(\frac{8}{27}) + \frac{1}{2}(\frac{13}{27}) = \frac{1}{6}$, and for $s = 2$, equals $1(\frac{4}{9}) - 1(\frac{10}{27}) + \frac{1}{2}(\frac{5}{27}) = \frac{1}{6}$. Since the cut-points are independent of signals, Lemma 4 shows that choice probabilities are also invariant to signals. The intuition is that the cut-points measure the marginal amount of idiosyncratic utility required to make a particular action better than another. When these values are constant across signals, then so are choice probabilities. This implies that if initial beliefs are $P^0 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, agents can never learn anything about the true state by observing the actions of others. However, beliefs can be perturbed slightly from $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ such that learning will continue. Indeed, the set of beliefs from which learning continues to the limit is open and dense (and full Lebesgue measure) in the belief simplex.

This is illustrated for the example in Figure 4, which displays the simplex of beliefs over the three states. The upper vertex corresponds to beliefs concentrated on state 1, the lower left




Figure 4: *The belief simplex for the example of non-admissible payoffs. The darkest curve is the set of absorbing beliefs, with the preimage of that set given by the two adjacent curves, and higher-order preimages shown as progressively thinner curves.*

vertex state 2, and the lower right vertex state 3. The *absorbing* set of beliefs at which learning stops is the dark curve connecting the lower two vertices, passing through $P^t = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.¹⁴

More generally, for any common value matrix Π with distinct columns and any non-redundant signal distribution matrix Q , the set of beliefs at which learning stops is at most a lower-dimensional subset of the K -simplex. The same is true for the “pre-image” of such a public belief vector P^t under the belief updating process. Working backwards from t to initial beliefs P^0 in an inductive fashion, the set of initial beliefs that could potentially lead to beliefs P^t in finite time is a countable union of sets of measure zero, and so is itself a set of measure zero.

Thus quite generally the set of prior beliefs P^0 from which learning could stop at a non-degenerate outcome has measure zero, and our convergence result in Theorem 2 holds generically in this sense. When admissibility is satisfied, however, we get the stronger universal convergence result, where *all* priors necessarily lead to fully correct learning in the limit. Furthermore, in the

¹⁴The upper vertex is also an absorbing belief and is marked as a separate point.

case $A \geq K$, Π satisfies the admissibility condition generically, so that almost all common-value technologies are admissible.

An exceptional case arises if Π contains several columns that differ only by a constant,¹⁵ and full learning will generally not occur in this case. We conjecture, but have not proven, the following result: Let $\Pi = (\pi_1, \dots, \pi_K)$, where each π_k is an A -dimensional column vector of payoffs in state k . Let $\bar{\mathcal{K}} \subset \mathcal{K}$ be a set of states for which $\pi_k = \pi_{k'} + c_{kk'}$ for all $k, k' \in \bar{\mathcal{K}}$ and some constant $c_{kk'}$, and such that all $\pi_k \in \mathcal{K} \setminus \bar{\mathcal{K}}$ can not be written as additive shifts of each other or the elements of $\bar{\mathcal{K}}$. If the true state is an element of $\bar{\mathcal{K}}$, then we conjecture that generically beliefs converge to a point that puts mass one on $\bar{\mathcal{K}}$. If the true state is not in $\bar{\mathcal{K}}$, we conjecture the convergence result converges obtains as above.

8. Relaxing full rationality

A natural question to ask is to what extent the convergence results rely on perfectly rational Bayesian decision-making. We consider three possibilities: errors in beliefs, errors in recording or monitoring, and decision errors.¹⁶ In all three extensions, full learning still obtains. Errors in belief updating are modelled as in Kahnemann and Tversky (1973), which presented experimental evidence that some individuals will overweight their private signal relative to the prior, a judgement bias that has become known as the *base rate fallacy*. All of the information aggregation results of the previous sections continue to hold even in the presence of this fallacies (or the opposite, under-weighting). Of additional interest is the fact that the learning process is actually faster as agents' behavior reflects the base rate fallacy. The second type of irrational behavior we consider is that agents sometimes "tremble" and accidentally choose the wrong alternative with some probability (which may vary over time). Alternatively, one could suppose there is imperfect monitoring, and agents' choices are observed with error. This addition to the model does not change the results. The third type of irrational behavior is

¹⁵The constant could of course be zero, in which case these columns of Π would be identical.

¹⁶See Goeree, Palfrey, and Rogers (2003) for details.

that these trembles are payoff related. That is, agents are better at avoiding high cost errors than low cost errors. This is the basic idea behind quantal response equilibrium (McKelvey and Palfrey, 1995, 1998). In a quantal response equilibrium with payoff responsive errors, the convergence result holds, *even in the pure common values case*. The reason is that quantal response equilibrium with payoff responsive errors corresponds to the Bayesian equilibrium of the basic model with private values disturbances.

9. Conclusion

In this paper we establish some general conditions for positive results about convergence of beliefs in social learning models when preferences have both a common value and a private value component. Under weak assumptions on the information structure and preferences, neither herds nor cascades occur. Learning does not stop in finite time, and as a result information is fully aggregated in the limit. Observed behavior is asymptotically fully revealing in the sense that the public beliefs about the state converge to a degenerate random variable concentrated on the true state. The key assumptions required are that signals are informative and that the support of the distribution of private values is sufficiently rich. If signals were not informative, then full revelation of the state could not possibly occur, even if signals were public. If the distribution of private values were not sufficiently rich, then at some point learning could stop, because cutpoints for optimal decision rules reach a boundary (or gap) in the support, so that choices are not informative of signals.

The main reason for the difference between these new results and past negative results is the private value component of payoffs. In a related paper, Smith and Sorensen (2000) consider a model with a finite number of common-value payoff types, which results in incomplete or confounded learning. The key feature of the full-support private-values environment, where learning is complete, is that there is positive probability of every action conditional on any signal and any non-degenerate prior.

Intuition suggests that our result could be extended to allow for much more general joint distributions of signals and private values, which could include some positive or negative correlation, provided the joint distribution satisfies an absolute continuity condition. While the actual proof technique we employ does not extend in a direct way to such environments, we conjecture that our limit result will continue to hold quite generally as long as signals are informative, and the distribution of idiosyncratic payoffs satisfies a condition of “observationally full support,” i.e. for any action and any set of beliefs there exist private shocks for which that action would be optimal.

A. Appendix

Proof of Lemma 1. Since elements of Π lie strictly between 0 and 1, the cutpoints $\bar{v}_{a,a'}^t(s^t|P^t)$ lie strictly between -1 and 1 for all $a, a' \in \mathcal{A}$. Hence, by Assumption 3, for each $a \in \mathcal{A}$ there is a set $I(a) \subseteq \text{supp}(f)$ of positive measure such that if $v_a \in I(a)$ then $v_a - \bar{v}_{a,a'}^t(s^t|P^t) > 0$ for all $a' \neq a \in \mathcal{A}$. Hence, $C_a^t(s^t|P^t) > 0$. Moreover, the cumulative distribution $F(\cdot)$ is strictly increasing in all its arguments on $I(a)$, so choice probabilities are strictly decreasing in cutpoints. *Q.E.D.*

Proof of Lemma 2. The proof is by induction. Recall that $P_k^1 = P_k^0 > 0$ for all $k \in \mathcal{K}$, by assumption. Lemma 1 ensures $C_a^1(s^1|P^1) > 0$ for all $a \in \mathcal{A}$ and $s^1 \in \mathcal{S}$, so $T_{ka}^1(P^1) > 0$ for all $k \in \mathcal{K}$ and $a \in \mathcal{A}$ and (5.2) yields $P_k^1(a) > 0$ for all $k \in \mathcal{K}$ and $a \in \mathcal{A}$, so $P_k^2 > 0$ for all $k \in \mathcal{K}$. By a similar argument, for $t > 1$, $P_k^t > 0$ for all $k \in \mathcal{K}$ implies $P_k^{t+1} > 0$ for all $k \in \mathcal{K}$. *Q.E.D.*

Proof of Lemma 4. First consider the case $A > 3$. Suppose $\bar{v}_{a,a'}^t(s|P^t) > \bar{v}_{a,a'}^t(s'|P^t)$ for some actions $a \neq a'$ and some signals $s \neq s'$ (if no such cutpoints exist we are done). Relabel actions such that action a becomes action 1 and action a' becomes action 2. Lemma 1 and $C_1^t(s|P^t) = C_1^t(s'|P^t)$ together imply that $\bar{v}_{1,a''}^t(s|P^t) < \bar{v}_{1,a''}^t(s'|P^t)$ for some $a'' > 2$. Again relabel states such that $a'' = 3$. So we have:

$$\bar{v}_{1,2}^t(s|P^t) > \bar{v}_{1,2}^t(s'|P^t), \quad \bar{v}_{1,3}^t(s|P^t) < \bar{v}_{1,3}^t(s'|P^t). \quad (\text{A.1})$$

From this we derive constraints on other cutpoints by an induction argument.

Claim: *If*

$$\begin{cases} (-1)^a \bar{v}_{a,a'}^t(s|P^t) < (-1)^a \bar{v}_{a,a'}^t(s'|P^t) \\ (-1)^a \bar{v}_{a,a+2}^t(s|P^t) > (-1)^a \bar{v}_{a,a+2}^t(s'|P^t) \end{cases} \quad \text{for } a < a^*, \quad a' \leq a+1, \quad a' \neq a$$

for $2 \leq a^* \leq A-2$, then

$$\begin{cases} (-1)^a \bar{v}_{a,a'}^t(s|P^t) < (-1)^a \bar{v}_{a,a'}^t(s'|P^t) \\ (-1)^a \bar{v}_{a,a+2}^t(s|P^t) > (-1)^a \bar{v}_{a,a+2}^t(s'|P^t) \end{cases} \quad \text{for } a \leq a^*, \quad a' \leq a+1, \quad a' \neq a \quad (\text{A.2})$$

for $2 \leq a^* \leq A - 2$.

The result is trivial for $a < a^*$ so the only cases to consider are $a = a^*$ and $a' \leq a^* + 1$, $a' \neq a^*$.

To show the result for $a = a^*$ and $a' = a^* + 1$ note that

$$\begin{aligned}
(-1)^{a^*} \bar{v}_{a^*, a^*+1}^t(s|P^t) &= (-1)^{a^*-1} (\bar{v}_{a^*-1, a^*}^t(s|P^t) - \bar{v}_{a^*-1, a^*+1}^t(s|P^t)) \\
&< (-1)^{a^*-1} (\bar{v}_{a^*-1, a^*}^t(s'|P^t) - \bar{v}_{a^*-1, a^*+1}^t(s'|P^t)) \\
&= (-1)^{a^*} \bar{v}_{a^*, a^*+1}^t(s'|P^t),
\end{aligned} \tag{A.3}$$

where the inequality follows from the induction hypothesis. For $a = a^*$ and $a' = a^* - 1$

$$\begin{aligned}
(-1)^{a^*} \bar{v}_{a^*, a^*-1}^t(s|P^t) &= (-1)^{a^*-1} \bar{v}_{a^*-1, a^*}^t(s|P^t) \\
&< (-1)^{a^*-1} \bar{v}_{a^*-1, a^*}^t(s'|P^t) = (-1)^{a^*} \bar{v}_{a^*, a^*-1}^t(s'|P^t).
\end{aligned} \tag{A.4}$$

Finally, the proof for $a = a^*$ and $a' < a^* - 1$ follows since

$$\begin{aligned}
(-1)^{a^*} \bar{v}_{a^*, a'}^t(s|P^t) &= (-1)^{a^*-2} (\bar{v}_{a^*-2, a'}^t(s|P^t) - \bar{v}_{a^*-2, a^*}^t(s|P^t)) \\
&< (-1)^{a^*-2} (\bar{v}_{a^*-2, a'}^t(s'|P^t) - \bar{v}_{a^*-2, a^*}^t(s'|P^t)) \\
&= (-1)^{a^*} \bar{v}_{a^*, a'}^t(s'|P^t).
\end{aligned} \tag{A.5}$$

This proves the top line of (A.2). Lemma 1 together with $C_{a^*}^t(s|P^t) = C_{a^*}^t(s'|P^t)$ then implies $(-1)^{a^*} \bar{v}_{a^*, a''}^t(s|P^t) > (-1)^{a^*} \bar{v}_{a^*, a''}^t(s'|P^t)$ for some $a'' > a^* + 1$. We can relabel states such that $a'' = a^* + 2$, which proves the bottom line of (A.2), thus verifying the claim.

Inequalities (A.1) together with the claim imply that (A.2) holds for all $2 \leq a^* \leq A - 2$. Furthermore, by repeating steps (A.3), (A.4), and (A.5) for the case $a^* = A - 1$ we can conclude:

$$(-1)^{A-1} \bar{v}_{A-1, a}^t(s|P^t) < (-1)^{A-1} \bar{v}_{A-1, a}^t(s'|P^t) \text{ for } a \leq A, a \neq A - 1. \tag{A.6}$$

Before finishing the proof of the claim we extend (A.6) to the cases $A = 2$ and $A = 3$. For $A = 2$, condition (A.6) simply states that the single cut-point $\bar{v}_{1,2}^t$ differs for the two possible signals.

For $A = 3$, condition (A.6) can be derived from (A.1). Recall that $\bar{v}_{2,1}^t(s|P^t) = -\bar{v}_{1,2}^t(s|P^t)$ and $\bar{v}_{2,3}^t(s|P^t) = \bar{v}_{2,1}^t(s|P^t) + \bar{v}_{1,3}^t(s|P^t)$ for all signals s , so (A.1) implies $\bar{v}_{2,1}^t(s|P^t) < \bar{v}_{1,2}^t(s'|P^t)$ and $\bar{v}_{2,3}^t(s|P^t) < \bar{v}_{2,3}^t(s'|P^t)$. Hence (A.6) holds for all $A \geq 2$.

The desired contradiction now follows as Lemma 1 implies that $C_{A-1}^t(s|P^t) \neq C_{A-1}^t(s'|P^t)$ since for all $a \neq A-1$ the cutpoints $\bar{v}_{A-1,a}^t(s|P^t)$ are greater (less) than $\bar{v}_{A-1,a}^t(s'|P^t)$ when A is even (odd) while $\bar{v}_{A-1,A-1}^t(s|P^t) = \bar{v}_{A-1,A-1}^t(s'|P^t) = 0$. To summarize, independence of choice probabilities with respect to signals implies that $\bar{v}_{a,a'}^t(s|P^t) = \bar{v}_{a,a'}^t(s'|P^t)$ for all $a, a' \in \mathcal{A}$ and $s, s' \in \mathcal{S}$.

Finally the reverse implication is clear from (4.2). *Q.E.D.*

Proof of Theorem 2. Lemma 6 shows that ℓ_1^t is a martingale. By the Martingale Convergence Theorem (see Doob, 1953), there exists a limit random variable to which ℓ_1^t converges almost surely. Since ℓ_1^t satisfies the martingale property, Fatou's lemma implies $\lim_{t \rightarrow \infty} E(\ell_1^t) \leq E(\ell_1^0) = \ell_1^0$, which is finite by assumption. Thus it is impossible for beliefs to converge upon an incorrect state with positive probability, as this would imply an infinite likelihood ratio against the true state. Smith and Sorensen (2000) (Theorems B.1 and B.2, p. 393) show that any public belief ℓ^* to which the belief process may converge with positive probability must also be a fixed point of the learning process in the sense that $\ell^t = \ell^*$ implies $\ell^{t+1} = \ell^*$. The reason is that if $\ell^t \rightarrow \ell^*$ then since the updating process is continuous in beliefs, if ℓ^* were not a fixed point of the learning process then almost surely beliefs eventually become bounded away from ℓ^* , contradicting the assumption that it is a limit point with positive probability. Lemmas 2 through 4 establish that the only possible such fixed points occur where beliefs are degenerate. Thus with probability one, $\ell_1^t \rightarrow 0$, i.e., beliefs converge to the correct state almost surely. *Q.E.D.*

References

- Anderson, L. and C. Holt, (1997) "Information Cascades in the Laboratory," *American Economic Review*, **187**: 797-817.
- Banerjee, A. (1992) "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, **107**, August, 797-817.
- Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992) "A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades," *Journal of Political Economy*, **100**, October, 992-1026.
- Goeree, J., T. Palfrey, and B. Rogers (2003), "Informational Cascades in the Laboratory? A re-examination," in preparation.
- Jackson, M. and Kalai, E. (1997) "Social Learning in Recurring Games," *Games and Economic Behavior*, **21**, 102-134.
- Kahneman, D. and A. Tversky (1973) "On the Psychology of Prediction," *Psychological Review*, **80**:237 - 251.
- Lee, I. H. (1993) "On the Convergence of Informational Cascades," *Journal of Economic Theory*, **61**, December, 395-411.
- McKelvey, R., and T. Palfrey (1995), "Quantal Response Equilibria in Normal Form Games," *Games and Economic Behavior*, **10**(July):6-38.
- McKelvey, R., and T. Palfrey (1998), "Quantal Response Equilibria in Extensive Form Games," *Experimental Economics*, **1**:9-41.

Smith, L., and Sorensen, P. (2000) “Pathological Outcomes of Observational Learning,” *Econometrica*, **68**, March, 371-398.