LEARNING FROM NEIGHBORS

Venkatesh Bala[1]
and
Sanjeev Goyal[1]

First Version: October 1993
This Version: July 1996

Abstract

When payoffs from different actions are unknown, agents use their own past experience as well as the experience of their neighbors to guide their current decision making. This paper develops a general framework to study the relationship between the structure of information flows and the process of social learning.

We show that in a connected society, local learning ensures that all agents obtain the same utility, in the long run. We develop conditions under which this utility is the maximal attainable, i.e. optimal actions are adopted. This analysis identifies a structural property of information structures – local independence – which greatly facilitates social learning. Our analysis also suggests that there exists a negative relationship between the degree of social integration and the likelihood of diversity. Simulations of the model generate spatial and temporal patterns of adoption that are consistent with empirical work.

Key Words: Connected societies, conformism, social integration, diffusion, diversity, locally independent agents, Royal Family.

JEL Classification: D83, L15, O30, Q16, R10.

---

[1]Dept. of Economics, McGill University, Montreal, and Econometric Institute, Erasmus University, Rotterdam, respectively. E-mail addresses: bala@musicb.mcgill.ca and goyal@wke.few.eur.nl
The paper has benefitted greatly from the suggestions of the editor and two anonymous referees. We would like to thank Beth Allen, Larry Blume, David Easley, Glenn Ellison, Yossi Greenberg, David Hirshleifer, Hsueh-Ling Huynh, Boyan Jovanovic, Ehud Kalai, Nick Kiefer, Krishna Kumar, Bentley McLeod, Jean Francios Mertens, Charles Noussair, Heracles Polemarchakis, Roy Radner, Karl Shell, Jonathan Thomas and Ivo Welch for helpful comments. The paper has also benefitted from presentations at Cambridge, CORE (Louvain), Cornell, CRDE (Montreal), 1994 Decentralization Conference, Toronto, 1994 Econometric Society Meetings (Maastricht and Quebec City), Erasmus, Essex, Institute for Social and Economic Change, Bangalore, Indian Statistical Institute (Bangalore and New Delhi), McGill, NYU, Stoney Brook and Warwick. We also thank Gerrit Draisma, Melanie Raymond and Qian Xie for research assistance and acknowledge the financial support from Quebec's Fonds FCAR and SSHRC.

1. Introduction

In a wide variety of economic environments, individual agents are obliged to choose a course of action without being fully informed about the true payoff from the different options. As time goes by, they learn from their own past experience; moreover, since experiments are often expensive and time consuming, they also try and gather information from the experience of others, both through personal communication as well through magazines and professional journals.[2] In this paper we develop a general framework to understand how the structure of information links in a society affects the generation of information (via the actions individuals choose) as well as its social dissemination.

We consider a society with a large number of agents, each of whom faces a similar decision problem: choose an action at regular intervals without knowing the true payoffs from different actions. The action chosen generates a random reward and also provides information concerning the true payoffs. Before choosing an action, an agent uses her own past experience as well as the experience of a subset of the society, viz. her *neighbors*, to revise her beliefs about the true payoffs.[3] Given these beliefs, an agent chooses an optimal action. In this setting, we study the evolution of agents' beliefs, actions and utilities. Our interest is in the following types of questions:

- What features of an information structure facilitate/hinder the social adoption of an optimal action?
- What type of neighborhood structures sustain diversity/conformism?[4]

[2]Examples of such environments include consumers learning about different brands, farmers learning about the productivity of a new crop and doctors learning about the efficacy of new treatments. Empirical work has documented the importance of learning from 'others' in several contexts, such as the adoption of new crops (Ryan and Gross, 1943), the diffusion of patent drugs (Coleman, 1966), the choice of new agricultural techniques (Hagerstrand, 1969; Rogers, 1983), economic demography (Watkins, 1991) and the purchase of consumer products (Kotler, 1986).

[3]This experience includes the choice of action as well as the corresponding outcome. This is a natural formulation in the examples mentioned in footnote 1 above. We assume that agents use only the information available from the realizations of actions taken by their neighbors, and that they *do not attempt to make any inferences* from the choices of their neighbors per se. Thus, we suppose that agents are 'bounded Bayesians'. Our assumptions concerning individual decision rules are discussed in greater detail in section 2.3.

[4]Diversity refers to a situation in which different groups of agents choose different actions, while conformism describes the outcome with everyone choosing the same action.

- What are the spatial and temporal patterns of adoption when individuals learn from their neighbors? Are these patterns consistent with empirical observations?

Our analysis is restricted to connected societies[5] and we start by establishing an important property of such societies: *the limiting (expected) utilities of all agents are equal* (Theorem 3.2). The proof of this theorem uses the following arguments. We first establish that if an agent i takes some action x infinitely often then the limiting utility is equal to the true payoff from action x. Next, we consider two agents i and j and suppose that j is a neighbor of i. If agent j takes some action x' infinitely often then her limiting utility is equal to the true payoff from action x'. We then establish the following intuitive property: if i observes j then the true payoff from x must be at least as high the true payoff from x'. We note that this property of limiting utilities is transitive. The proof is completed by using the definition of connectedness along with this transitivity of limiting utilities.

Theorem 3.2 implies that, in the long run, agents cannot choose actions yielding different payoffs. This motivates the study of two related questions: one, do agents choose the optimal action, and two, can different actions with the same payoff survive, in the long run?

We first study the *complete learning* question. In this analysis we focus on large societies, i.e. societies with a countably infinite number of agents.[6] We begin with an example of incomplete learning. In this example, every agent has to choose between two actions, one whose payoff is known and a second action whose payoff is unknown; thus agents do not know which action is optimal. We suppose that agents are located on integer points of the real line and each agent observes the agent on either side of her. In addition, there exists a 'Royal Family', i.e. a small set of agents who are observed by everyone.[7] We suppose that the action with unknown payoff is the optimal action and

[5]A society is said to be connected, if for very pair of agents i and j, either j is a neighbor of i or there exist agents $\{i_1, \cdots, i_m\}$, such that i_1 is a neighbor of i, i_2 is a neighbor of i_1 and so on, until j is a neighbor of i_m. This is a very general class of societies. Familiar examples of connected societies are (a) agents located on points of a d-dimensional lattice in which every agent observes her immediate 2^d neighbors; (b) an organization tree where each person observes their immediate superior and subordinates; (c) agents located around a circle, observing their immediate neighbors and in addition observing a common set of agents who are sampled by a consumer magazine; (d) a group of agents with public observability.

[6]Using standard arguments it can be shown that learning is generally incomplete in finite agent societies.

[7]This structure corresponds to situations in which individuals have access to local as well as a common public source of information. Thus it is quite prevalent in everyday life. For example, such a structure arises naturally in the context of agriculture where individual farmers observe their neighboring farmers

also that initially everyone's prior beliefs favor the adoption of this action. Thus an infinite number of independent trials of this action are undertaken in the society; despite this, we show that there is a positive probability that the society will choose the suboptimal action eventually. This happens because, in our example, the Royal Family can generate sufficient negative information that can overwhelm any locally gathered positive information, thereby inducing all agents to switch to the action with known payoffs. This means that no further information is generated and thus the society is locked into an inferior choice. We can also show that in the absence of the Royal Family, the society will choose the optimal action in the long run. Thus, this example illustrates an interesting aspect of social learning: *more information links can make the society worse off*.

The example with incomplete learning also helps us identify alternative sets of conditions on the *structure of neighborhoods*, *the distribution of prior beliefs* and the *informativeness of actions*, that are sufficient for complete learning (Theorems 4.1-4.2 and Proposition 4.1). Our results on complete learning highlight the role of *locally independent* agents. The general argument proceeds as follows. First, given an agent i, we can choose a set of sample paths A_i having positive probability with the following properties: A_i depends only upon the realizations agent i observes, and moreover sample paths in A_i have a uniform upper bound on the amount of negative information concerning optimal actions.[8] The conditions in Theorem 4.1 ensure that the prior beliefs of some agents are sufficiently optimistic to overcome this negative information; thus an agent i with 'optimistic beliefs' will choose an optimal action forever on the set A_i.[9] We say that two agents i and i' are locally independent if they have non-overlapping neighborhoods, i.e. they observe different sets of agents. For two such agents the corresponding events A_i and $A_{i'}$ are independent. This implies that the probability that neither i nor i' tries an optimal action forever is bounded above by the product of the probabilities

but all the farmers observe a few large farmers and research laboratories. Another setting with this structure is a consumer goods market; individual consumers discuss purchase decisions with their colleagues and friends and all potential customers read one or two consumer magazines which report on some experiments/consumer experiences. A third example pertains to research activity; individual researchers typically keep abreast of developments in their own narrow area of specialization, and also try to keep informed about the work of the pioneers/intellectual leaders in their subject more broadly defined.

[8]This construction is possible in our model because agents do not make inferences from the choices of their neighbors, but only from the realizations of the choices.

[9]Our other complete learning results, Proposition 4.1 and Theorem 4.2, impose restrictions on the informativeness of actions to ensure that an agent will choose only optimal actions from a finite point onwards, with positive probability.

that neither A_i nor $A_{i'}$ occur. Hence the probability bound on the event that no one from a set of locally independent agents chooses an optimal action forever is exponentially decreasing in the number of such agents and in the limit equals 0. The final step in the argument invokes Theorem 3.2 to show that in a connected society the probability of a society choosing suboptimal actions in the long run is subject to the same upper bound.

We next study the *conformism vs. diversity* question: can different groups of agents in a connected society take different actions (having the same payoffs)? Diversity suggests that there exist boundaries, with agents on one side of the boundary choosing one action while agents on the other side choose a different action. Our analysis focuses on the sources of information of the boundary agents. We argue, with the help of an example, that diversity is easier to sustain when agents on each side of the boundary have more information links with agents who choose as they do, i.e. with agents on the same side of the boundary (Proposition 5.1). The proof of this result uses the law of iterated logarithm and is of independent technical interest. We also show that, in the context of this example, public observability implies conformism (Proposition 5.2). Taken together, the propositions also make the more general point that the structure of information flow in societies have a direct bearing on the likelihood of diversity.

We study the *temporal* and *spatial patterns* of diffusion by simulating the choices of a group of farmers trying to learn the true productivity of a new crop. We find that the temporal pattern (percentage of adopters vs. time) is described quite well by the logistic function, and that the rate of adoption is positively related to the profitability of the new crop. These results are consistent with empirical findings (Griliches, 1959; Feder, Just and Zilberman, 1985). We also observe that for different model specifications and parameter values the speed of convergence is fairly rapid. Finally, with regard to the spatial patterns we find that initially small groups of farmers adopt the new crops and then it slowly spreads as neighboring agents adopt it as well. Eventually these regions join up and the pace of diffusion accelerates. These findings match the empirically observed spatial patterns (see e.g., Hagerstrand (1969)).

Our paper is a contribution to the theory of social learning; we now discuss its relationship with some recent work by Ellison and Fudenberg (1993, 1995).[10] In Ellison and Fudenberg's models, agents use only currently available social information such as recent

[10]Some other papers in this area are An and Kiefer (1992) and Smallwood and Conlisk (1983).

popularity weighting and disregard historical data (including their own past experience) in making decisions. By contrast, in our model agents *do* use historical information; moreover, the bounded Bayesian decision rule they employ precludes the use of popularity weighting. Ellison and Fudenberg study the possibility of obtaining efficient outcomes and social diversity under different levels of popularity weighting and sample sizes. While our paper also studies efficiency and conformism, we focus on the role of prior beliefs and neighborhood structures. These differences suggest that our paper should be viewed as complementary to their work.

More generally, our paper should be seen as contributing to the theory of Bayesian learning. Research in this tradition has focused on cases where individual agents privately observe a signal and also have access to some central statistic. This central statistic varies depending on the model; in models of rational expectations learning, for instance, market prices are the central statistic, while in the recent work on herding/information cascades the actions of all previous agents are publicly observable.[11] The work on herding/information cascades considers situations where a sequence of individuals (who make one-shot decisions) learn from the actions of their predecessors. By contrast, in our framework, agents take actions repeatedly and learn from their own past experience as well as the experience of their neighbors. Our formulation of neighborhoods captures in a natural form the flow of information in such settings, thereby enabling us to explore the relationship between social structure and learning. Our paper can be viewed as integrating the two strands of literature dealing with social learning and Bayesian learning respectively.

Finally, our paper can also be regarded as studying the dynamics of technology adoption. Our example on incomplete learning in the presence of a Royal Family provides new insights about how the structure of information flows can generate 'lock-ins' into inferior technologies.[12] In this connection we would also like to mention the early work of Allen (1982a, 1982b) which explores the role of neighborhood influence on the invariant distribution of a process of technology adoption. Our paper extends her work by considering social learning in an explicit model of (Bayesian) individual decision making and learning.

[11] See Blume and Easley (1992) for a survey of the rational expectations learning literature; recent papers on multi-agent models of Bayesian learning include Aghion, Paz-Espinosa and Jullien (1993), Bala and Goyal (1994, 1995) and Bolton and Harris (1992). Banerjee (1992) and Bikchandani, Hirshleifer and Welch (1992) are the standard references on herding/information cascades.

[12] See e.g. Arthur(1989).

The rest of the paper is organized as follows. Section 2 describes the model. Section 3-5 presents our results while section 6 discusses simulations of spatial and temporal patterns of learning. Section 7 concludes.

2. The Model

2.1 Preliminaries: Let Θ be a finite set of possible states of the world, X be a finite set of actions and let Y be the space of outcomes. If the state of the world is $\theta \in \Theta$ and an agent chooses action $x \in X$, he observes outcome y with conditional density $\phi(y; x, \theta)$ and obtains reward $r(x, y)$. We make the following assumptions about Y, ϕ and $r(x, y)$.

(A.1) *Y is a non-empty, separable metric space. The distribution of outcomes*[13] *conditional on x and θ can be represented by the density $\phi(\cdot; x, \theta)$ with respect to a measure , defined on the Borel subsets of Y.*

(A.2) *For each $x \in X$, $r(x, \cdot)$ is bounded and measurable in Y.*

Agents do not know the true state of the world, and they enter with a prior belief in the set $\mathcal{D}(\Theta)$ of beliefs (probability distributions) over the state of nature :

$$\mathcal{D}(\Theta) \;=\; \{\mu = \{\mu(\theta)\}_{\theta\in\Theta} \mid \text{for all } \theta \in \Theta, \mu(\theta) \geq 0 \text{ and } \sum_{\theta\in\Theta} \mu(\theta) \;=\; 1\}. \tag{2.1}$$

Given belief μ an agent's one-period expected utility $u(x, \mu)$ from taking action x is :

$$u(x, \mu) = \sum_{\theta\in\Theta} \mu(\theta) \int_Y r(x, y)\phi(y; x, \theta)d, (y). \tag{2.2}$$

Note that $u(x, \cdot)$ is linear on $\mathcal{D}(\Theta)$ for every $x \in X$. We assume that individuals have the same preferences.[14] Let $G : \mathcal{D}(\Theta) \to X$ be the one-period optimality correspondence:

$$G(\mu) = \{x \in X \mid u(x, \mu) \geq u(x', \mu) \text{ for all } x' \in X\}, \quad \mu \in \mathcal{D}(\Theta). \tag{2.3}$$

Let δ_θ be the point mass belief on the state θ; then $G(\delta_\theta)$ denotes the set of expost optimal actions if the true state is $\theta \in \mathcal{D}(\Theta)$. (In the rest of the paper, we refer to expost optimal actions as 'optimal actions').

[13] In what follows, we shall use the words outcomes/realizations/observations interchangeably.

[14] We have also explored the learning process when agents have heterogeneous utilities. Our results on limiting utilities and learning carry over if for each group of agents of a given preference type, taken separately, connectedness obtains.

We now give two examples which are special cases of the above framework. The first example helps to clarify the basic structure, while the second example is the canonical bandit model (Berry and Fristedt, 1985) and illustrates the generality of our framework.

Example 2.1: There are two actions x_o and x_1 and two states, θ_o and θ_1. In state θ_1, action x_1 yields Bernoulli distributed payoffs with parameter $\pi \in (1/2, 1)$; in state θ_o payoffs from x_1 are Bernoulli distributed with parameter $1 - \pi$. Furthermore, in both states action x_o yields payoffs which are Bernoulli distributed with parameter 1/2. Hence, x_1 is the optimal action if the true state is θ_1 while x_o is the optimal action if θ_o is the true state. The belief of an agent is a number $\mu \in (0, 1)$, which represents the probability that the true state is θ_1.[15] In this example, the one period optimality correspondence is given by:

$$G(\mu) \;=\; \begin{cases} x_1 & \text{if } \mu \geq 1/2; \\ x_o & \text{if } \mu \leq 1/2. \end{cases} \tag{2.4}$$

Example 2.2: There is a finite set of actions X; each of the actions can be one of $s \geq 2$ quality levels or types. We suppose that the s quality types are labelled $\{q_1, \ldots, q_s\}$. If an action x is of quality q_m then it generates observations with a density $\phi_m(y)$ and a reward $r(y)$. The expected value of an action of type q_m is

$$V_m \equiv \int_Y r(y)\phi_m(y)\mathrm{d},\, (y). \tag{2.5}$$

Let the quality levels be strictly ordered according to ascending expected value, i.e. $V_1 < V_2 < \ldots < V_s$. This induces an ordering $\prec$ among quality levels where $q_j \prec q_k$ if and only if $V_j < V_k$. Any two distinct actions are independent. This implies that a belief μ can be written as

$$\mu \;=\; \{\{\mu(x;q)\} \mid \sum_{q \in \{q_1,\ldots,q_s\}} \mu(x;q) = 1,\; \mu(x;q) \geq 0,\; \forall x \text{ and } q\}. \tag{2.6}$$

In terms of the model presented earlier, a state $\theta \in \Theta$ is a specification of the quality types of the various actions. Let μ be the initial belief of an agent in the society. We assume as before that the belief is interior, i.e. for each x and each quality type q, $\mu(x;q) > 0$. Recall

[15] A natural interpretation of this example is to view action x_o as an established technology, whose payoff is known, and action x_1 as a new technology, whose payoff is uncertain.

that $u(x, \mu)$ gives the expected one-period utility of choosing x when the belief is μ. Thus, equation (2.2) can be rewritten as

$$u(x, \mu) \;=\; \sum_{q_j \in \{q_1, \ldots, q_s\}} \mu(x; q_j)\, V_j. \tag{2.7}$$

Finally, we also allow for an additional kind of action which is completely uninformative i.e. provides the same expected payoff in all states of nature.[16] The set of actions is thus given by $X = X_T \cup \{x_u\}$ where X_T is the set of actions each of which can be one of s types and x_u is the uninformative action.

2.2 The Social Structure: The set of agents is a non-empty set N which can be finite or countably infinite. For each $i \in N$, let $N(i)$ denote the set of *neighbors* of agent i. The statement 'j is in $N(i)$' is to be interpreted as saying that agent i has access to the entire past history of agent j's actions and outcomes. By contrast, if j is not a neighbor of i, then i does not observe any of j's actions or outcomes. Throughout this paper we shall suppose $i \in N(i)$ for every agent i. We also assume that the set $N(i)$ is a finite set for all $i \in N$. Let $N^{-1}(i) = \{j \in N \mid i \in N(j)\}$; the set $N^{-1}(i)$ is the set of all agents who observe agent i.[17] The 'Royal Family' is the set $R = \{j \in N \mid N^{-1}(j) = N\}$, i.e. those agents who are observed by everyone.

A society comprises of the set of agents and the neighborhoods of each of the agents. We shall say that a society is *connected* if, for every $i \in N$ and every other agent $j \in N$ there exists a sequence of agents $\{i_1, i_2, \ldots, i_m\}$ (depending upon i and j) such that $i_1 \in N(i)$, $i_2 \in N(i_1)$, and so on until $j \in N(i_m)$. The analysis in this paper is restricted to connected societies; we focus on such societies because all other types of societies can be analyzed as a collection of connected societies. In what follows, for expositional simplicity, we shall usually omit the term 'connected' while referring to societies.

2.3 The Dynamics of the Model: Time is discrete and is indexed by $t = 1, 2, \ldots$. At the beginning of period 1, each agent i has a prior belief $\mu_{i,1} \in \mathcal{D}(\Theta)$. We assume :

(A.3) For all $i \in N$, $\mu_{i,1} \in \text{Int}(\mathcal{D}(\Theta))$.

[16] In the context of crop choice, this corresponds to a case where the farmer decides not to plant any crop. Likewise, in situations where consumers are making brand choices this action is the 'no purchase' option.

[17] If the observation relation is symmetric, this set clearly coincides with $N(i)$. However, there are many sources of communication (e.g. radio, television, books, journals and gossip!) which do not possess symmetry. Our framework allows for asymmetric observational links.

where $\text{Int}(\mathcal{D}(\Theta))$ denotes the interior of the belief space. It is worth noting that we do not restrict the agents to have identical priors.

For each $i \in N$, let $g_i : \mathcal{D}(\Theta) \to X$ be a selection from the one-period optimality correspondence G of equation (2.3) above. In period 1, each agent i plays the action $g_i(\mu_{i,1})$ and observes the outcome. Agent i also observes the actions taken and outcomes obtained by the other agents in $N(i)$. In periods $t = 2$ and beyond, each agent i first computes her posterior belief $\mu_{i,t}$ based on the experiences of the agents in $N(i)$. In this regard, we assume that agents employ a "bounded Bayesian" learning algorithm. This algorithm specifies that agents modify their prior beliefs to posterior ones, using Bayes rule in conjunction with the information obtained from their own and their neighbors' experiences. However, they do not attempt to extract any information from the observed choices of their neighbors.[18] After forming her posterior $\mu_{i,t}$ in the manner described above, agent i then chooses the action $g_i(\mu_{i,t})$ which maximizes one-period expected utility, and the process continues in this manner. Thus, agents are being boundedly rational both in choosing their optimal action myopically given their beliefs and also in forming posterior beliefs.

The assumptions on individual decision making described above are not standard and we now discuss the motivation behind them. We are interested in analyzing the process by which individual agents make use of information gathered from their neighbors and the implications of this local learning for aggregate social outcomes. This suggests, first, that individual choice should not be arbitrary and, second, that there must be a well defined mechanism through which information from neighbors is incorporated in individual decisions. Given that the information observed by agents is partial, a model with fully rational agents would require that the learning problem of the agents be formulated as a dynamic Bayesian game of incomplete information. In such a formulation, the influence of neighborhood structure would interact with the incentives for strategic experimentation in addition to inducing a complex inference problem for agents. In this paper our concern is with the relationship between neighborhood structure and learning. To keep the model mathematically tractable and to allow us to focus on this relationship we have made certain simplifying assumptions on individual decision rules. Thus we assume that agents are myopic expected utility maximizers, which eliminates incentives for strategic

[18]This formulation thus rules out the use of popularity weighting and related measures in the learning process. Note, however, that Bayesian updating provides a relatively simple way for each agent to keep track of the information in past history. We also remark that while the bounded Bayesian learning rule employed here may not be efficient, it is consistent in our framework.

experimentation. Secondly, we also note that a Bayesian who knows the structure of the model should be able to incorporate information on the popularity of different choices among her neighbors in forming posterior beliefs. This possibility is precluded in our model, thus eliminating the inference problem and simplifying the belief revision process considerably.

We now briefly sketch the construction of the probability space since the notation is required for the results. Details are provided in Appendix A. For a fixed $\theta \in \Theta$ we define a probability triple $(\Omega, \mathcal{F}, \mathrm{P}^\theta)$, where Ω is the space containing sequences of realizations of actions of all agents over time, and P^θ is the probability measure induced over sample paths in Ω by the state $\theta \in \Theta$.

Let Θ be endowed with the discrete topology, and suppose $\mathcal{B}$ is the Borel σ-field on this space. For rectangles of the form $A \times H$ where $A \subset \Theta$ and H is a measurable subset of Ω, let $\mathrm{P}_i(A \times H)$ be given by

$$\mathrm{P}_i(A \times H) \;=\; \sum_{\theta \in A} \mu_{i,1}(\theta)\mathrm{P}^\theta(H). \tag{2.8}$$

for each agent $i \in N$. Each P_i extends uniquely to all of $\mathcal{B} \times \mathcal{F}$. Since every agent's prior belief lies in the interior of $\mathcal{D}(\Theta)$, the measures $\{\mathrm{P}_i\}$ are pairwise mutually absolutely continuous. All stochastic processes are defined on the measurable space $(\Theta \times \Omega, \mathcal{B} \times \mathcal{F})$. A typical sample path is of the form $\omega = (\theta, \omega')$ where θ is the state of nature and ω' is the infinite sequence of sample outcomes denoted by:

$$\omega' \;=\; ((y^x_{i,1})_{x \in X, i \in N}, (y^x_{i,2})_{x \in X, i \in N}, \ldots) \tag{2.9}$$

where $y^x_{i,t} \in Y^x_{i,t} \equiv Y$. Let $C_{i,t} \equiv g_i(\mu_{i,t})$ denote the action of agent i at time t, $Z_{i,t}$ the outcome of agent i's action at time t (i.e., the signal of her own action from the outcome space Y) and let $(Z_{j,t})_{j \in N(i)}$ be the set of outcomes of the neighbors of i at time t. Also let $U_{i,t}(\omega) = u(C_{i,t}, \mu_{i,t})$ be the expected utility of i with respect to her own action at time t.[19] The posterior belief of agent i in period $t+1$ is computed as follows:

$$\mu_{i,t+1}(\theta) \;=\; \frac{\prod_{j \in N(i)} \; \phi(Z_{j,t}; C_{j,t}, \theta) \; \mu_{i,t}(\theta)}{\sum_{\theta' \in \Theta} \prod_{j \in N(i)} \; \phi(Z_{j,t}; C_{j,t}, \theta') \; \mu_{i,t}(\theta')}. \tag{2.10}$$

[19]The outcomes of actions are projections of ω onto the respective coordinates. We assume that if agent i has chosen action x' for the t^{th} time on ω, he observes the coordinate $y^{x'}_{i,t}(\omega)$.

The σ-field of agent i's information at the beginning of time 1 is $\mathcal{F}_{i,1} \equiv \{\emptyset, \Theta \times \Omega\}$. For every $t \geq 2$, define $\mathcal{F}_{i,t}$ as the σ-field generated by the past history of agent i's observations of her neighbors' actions and outcomes, i.e. the random variables $(C_{j,1}, Z_{j,1})_{j \in N(i)}$, $\ldots, (C_{j,t-1}, Z_{j,t-1})_{j \in N(i)}$. Since by the rules of the process, agents only employ the information generated by their neighbors, the set classes $\{\mathcal{F}_{i,t}\}$ are the relevant σ-fields for our purposes. We shall denote by $\mathcal{F}_{i,\infty}$ the smallest σ-field containing all $\mathcal{F}_{i,t}$ for $t \geq 1$.

3. Aggregation of Information

In this section we establish that (roughly speaking) in a connected society every agent expects the same utility, in the long run. The first step in the study of the long run distribution of individual utilities consists of showing the convergence of a typical individual's beliefs and utilities. The following result shows that the sequence of posterior beliefs of a typical agent converges almost surely to a limit belief which is measurable with respect to the (direct) limit information of the agent.[20] This result is an immediate consequence of the Martingale Convergence Theorem.[21]

Theorem 3.1 $\triangleright$ *There exists $Q \in \mathcal{B} \times \mathcal{F}$ satisfying $\mathrm{P}_i(Q) = 1$ for all $i \in N$ and random vectors $\{\mu_{i,\infty}\}_{i \in N}$ such that*
(a) For each $i \in N$, $\mu_{i,\infty}$ is $\mathcal{F}_{i,\infty}$-measurable.
(b) $\omega \in Q \Rightarrow$ for all $i \in N$, $\mu_{i,t}(\omega) \to \mu_{i,\infty}(\omega)$.

In what follows, we restrict attention to a specific state of nature which is taken to be the true state. We shall denote this state by θ_1. Clearly, the set

$$Q^{\theta_1} \equiv \{\omega = (\theta, \omega') \in Q \mid \theta = \theta_1\}.$$

has P^{θ_1} probability 1. (Strictly speaking, the domain of definition of P^{θ_1} is the measurable subsets of Ω, not of $\Theta \times \Omega$. However, we can regard P^{θ_1} as the conditional probability induced by θ_1 on the product space, which is the same for all agents). Without loss of generality we assume that the strong law of large numbers holds on Q^{θ_1}. In what follows statements of the form "with probability one" are with respect to the measure P^{θ_1}.

[20] It is worth emphasizing that Theorem 3.1 does not preclude the possibility of limit beliefs being different across individual agents.
[21] Proofs of results in this section are given in Appendix B.

We next show the convergence of utilities of a typical individual in the society. For each agent i, given $\omega \in Q^{\theta_1}$, let $X^i(\omega)$ be the set of actions which are chosen infinitely often on the sample path. We shall refer to $X^i(\omega)$ as the set of limiting actions (of agent i) on ω. Given that every individual is a myopic optimizer, it seems natural that the set of limiting actions should be optimal with respect to the limiting beliefs. This is true, as part (a) of the following result shows. This result immediately implies that each agent's one period expected payoff converges as well. Recall that $U_{i,t}(\omega) \equiv u(C_{i,t}(\omega), \mu_{i,t}(\omega))$.

Lemma 3.1 ▷ *Suppose* $\omega \in Q^{\theta_1}$.
(a) If $x' \in X^i(\omega)$ *then* $x' \in \operatorname{argmax}_{x \in X} u(x, \mu_{i,\infty}(\omega))$.
(b) There exists a real number $U_{i,\infty}(\omega)$ *such that* $\{U_{i,t}(\omega)\} \to U_{i,\infty}(\omega)$. *Furthermore,* $U_{i,\infty}(\omega) = u(x', \mu_{i,\infty}(\omega))$ *where* x' *is any member of* $X^i(\omega)$.

We now examine the distribution of these limiting utilities and actions in the society. Our analysis is summarized in the following result.

Theorem 3.2 ▷ *Suppose that the society is connected. Then* $U_{i,\infty}(\omega) = U_{j,\infty}(\omega)$ *for all agents* i *and* j *in* N, *with probability 1.*

The proof of this result employs the following arguments. On a fixed sample path, consider two agents i and j and suppose $i \in N(j)$. We show that if x' is an action taken infinitely often by j then j's long run expected utility $U_{j,\infty}$ will be $u(x', \delta_{\theta_1})$. Likewise, if i chooses x infinitely often, then $U_{i,\infty} = u(x, \delta_{\theta_1})$. Furthermore, the assumption that j observes i is shown to imply that $u(x', \delta_{\theta_1}) \geq u(x, \delta_{\theta_1})$. Thus, $U_{j,\infty} \geq U_{i,\infty}$. Connectedness of the society now yields the result.

4. Long Run Social Learning

In this section we study the optimality of long run actions.[22] Our analysis suggests that the distribution of prior beliefs, the structure of neighborhoods and the informativeness of

[22]Suppose that θ_1 is the true state of the world. We shall say that long run actions of agent i are optimal on ω if $X^i(\omega) \subset G(\delta_{\theta_1})$. Likewise, beliefs of an agent i will be said to converge to the truth along ω if $\mu_{i,\infty}(\omega) = \delta_{\theta_1}$. We shall say that complete learning (or learning with probability 1) obtains if $\mathrm{P}^{\theta_1}(\cap_{i \in N}\{X^i(\omega) \subset G(\delta_{\theta_1})\}) = 1$.

actions all play an important role in determining whether or not long run actions will be optimal.[23]

We start by noting that social learning will typically be incomplete in finite societies.[24] This motivates the study of learning in societies with *infinitely* many agents. We begin our analysis with two observations. The first observation concerns the importance of the initial distribution of priors. It is easy to see (with the help of Example 2.1) that even in a large (infinite agent) society, no learning will occur if all agents start out with prior beliefs that lead them to choose the uninformative action. *Thus for learning to occur some restrictions on the distribution of prior beliefs are necessary.* Our second observation is that even when beliefs are favorable, the social structure of information flows may preclude learning. The following example illustrates this point and also helps us derive sufficient conditions for complete learning subsequently.

Example 4.1: Consider the setting of Example 2.1. Suppose that the true state is θ_1 and that the society has an infinite number of agents. Assume that the prior beliefs of agents satisfy the following condition:

$$\inf_{i \in N} \mu_{i,1} > \frac{1}{2}, \quad \sup_{i \in N} < \frac{1}{1+p^2}. \tag{4.1}$$

where $p = (1-\pi)/\pi \in (0,1)$. The above assumption implies that in period 1 all agents will choose the optimal (and informative) action x_1. We suppose that society N is given by the one dimensional integer lattice. For $i \in N$, the set of neighbors is assumed to be $N(i) = \{i-1, i, i+1\} \cup R$, where $R = \{1,2,3,4,5\}$ constitute the Royal Family. We

[23] An action $x \in X$ is said to be fully informative if, for all θ, θ' in Θ such that $\theta \neq \theta'$ we have :

$$\int_Y |\phi(y; x, \theta) - \phi(y; x, \theta')| d, (y) > 0.$$

Thus a completely informative action x is statistically capable of distinguishing between any two distinct states in the long run. Let X_I denote the set of completely informative actions. We shall say that an action x_u is uninformative if $\phi(\cdot; x_u, \theta)$ is independent of θ. It is worth noting that if the set of uninformative actions X_U is non-empty, then there is no essential loss of generality in assuming that it consists of a single element x_u.

[24] To see why this is true consider the set up of Example 2.1. Suppose that the true state is θ_1 and that the society is finite. Prior beliefs of agents are then represented by a number $\mu_{i,1}$ which is the probability that true state is θ_1. Let $\inf_{j \in N} \mu_{j,1} > 1/2$ and focus on the agent with the highest value of $\mu_{i,1}$. Standard arguments imply that there exists a finite sequence of T realizations of 0, such that this agent would switch to action x_o. Now consider the set of sample paths on which all agents get realizations of 0 for the first T periods. The probability of this set is positive given that realizations are independent and the number of agents finite. The argument is completed by observing that on any sample path in this set every agent will choose the sub-optimal (and uninformative) action x_o after time period T.

now note the possibility of incomplete learning: *there is a strictly positive probability that every agent will choose the suboptimal action x_o for all $t \geq 2$.* The proof is as follows: Define $\bar{Q} = \{Z_{j,1} = 0, \ \forall \ j \in R\}$; by construction, $\mathrm{P}^{\theta_1}(\bar{Q}) = (1-\pi)^5 > 0$. We show that if $\omega \in \bar{Q}$, then $C_{i,t}(\omega) = x_o$ for all $t \geq 2$, for $i \in N$. Note that on $\omega \in \bar{Q}$, an agent $i \in N$ observes 5 'negative' realizations from the Royal Family, while the maximum number of 'positive' realizations that can be observed locally is 3. Thus there is a minimum amount of residual negative information. Since $\mu_{i,1} < 1/(1+p^2)$ this negative residual information is sufficient to push the posterior belief $\mu_{i,2}$ below 1/2, making agent i choose the uninformative action. The argument is completed by noting that i has been chosen arbitrarily.[25]

In the example above, one reason for incomplete learning is that the prior beliefs of agents are not very dispersed. This allows a 'little' bad experience of a few people to convince everyone to switch to the uninformative action. This aspect of the example motivates a study of connected societies with dispersed prior beliefs. A second source of incomplete learning is the structure of individual neighborhoods. This structure allows the negative experience of a few people (the 'Royal Family') to overwhelm the locally gathered positive information of everyone in the society. This insight motivates the study of structures where the observability level is bounded. It is also worth noting that the example 'works' because the negative information generated by the Royal Family exceeds any positive information that a local neighborhood can generate. This suggests that if an agent (or a group of them) is able to generate an arbitrarily 'large' amount of positive information with non-zero probability, then complete learning may obtain.

In the rest of this section, we develop alternative sets of restrictions on these three factors – the distribution of prior beliefs, the structure of neighborhoods and the informativeness of actions – that are sufficient for complete learning (Theorems 4.1-4.2 and Proposition 4.1). While our analysis generates several useful insights, we are aware that our results do not provide a full characterization of the conditions under which social learning occurs.

[25] We briefly report simulations which assess the impact of the Royal Family on the probability of incomplete learning. The simulations have been carried out for a society with $|N| = 100$ for values of $|R|$ between 1 and 10, and where the uncertainty is described by a Bernoulli or Normal distribution. The results indicate that the probability of incomplete learning is not insignificant and can be as high as 0.2. The simulations also suggest that for low values of $|R|$, there is a roughly positive relationship between the probability of incomplete learning and the size of the Royal Family.

Our results exploit the notion of *locally independent agents.* We shall say that agents i and i' are locally independent if they have non-overlapping neighborhoods, i.e. $N(i) \cap N(i') = \emptyset$. A *pairwise locally independent group of agents* is a subset of N such that any two agents in the group are locally independent. We begin with Theorem 4.1, a result which focuses on the distribution of prior beliefs. It requires the following lemma.

Lemma 4.1 ▷ *Let K be a given positive integer and suppose i is an agent satisfying $|N(i)| \leq K$. There exists $d \in (0,1)$ and $\lambda > 0$ such that if $\mu_{i,1}(\theta_1) \geq d$ then $\mathrm{P}^{\theta_1}(\cap_{t \geq 1}\{C_{i,t}(\omega) \in G(\delta_{\theta_1})\}) \geq \lambda$.*

Proof: Let $\mu \in \mathcal{D}(\Theta)$. We start by noting that there exists a number $\hat{d} \in (0,1)$ such that $\mu(\theta_1) \geq \hat{d} \Rightarrow G(\mu) \subset G(\delta_{\theta_1})$. This follows since utility is continuous with respect to beliefs and the number of actions is finite.

Suppose next that agent i only chooses action x for $t-1$ periods, and observes a sequence $\{y^x_{i,\tau}\}_{\tau=1}^{t-1}$, where each $y^x_{i,\tau} \in Y^x_{i,t} \equiv Y$. The agent's information about state $\theta \neq \theta_1$ based upon her observations can be summarized by the *product likelihood ratio* $r^{x,\theta}_{i,t}$, defined as:

$$r^{x,\theta}_{i,t}(\omega) = \frac{\prod_{\tau=1}^{t-1} \phi(y^x_{i,\tau}(\omega); x, \theta)}{\prod_{\tau=1}^{t-1} \phi(y^x_{i,\tau}(\omega); x, \theta_1)}. \tag{4.2}$$

(If $t = 1$, we follow the convention that $r^{x,\theta}_{i,t} = 1$). It follows from an application of the law of large numbers that $r^{x,\theta}_{i,t} \to \bar{r}^{x,\theta}_i$ where $\bar{r}^{x,\theta}_i < \infty$, almost surely (see e.g., DeGroot, 1970; p. 201-204). Since this is true for all $\theta \neq \theta_1$ and all $x \in X$, there exists σ and a set A^σ_i of sample paths defined as:

$$A^\sigma_i \;=\; \prod_{x \in X}\Big\{\max_{\theta \in \Theta \setminus \theta_1} \sup_{t \geq 1} r^{x,\theta}_{i,t} \leq \sigma\Big\} \times \prod_{t=1}^{\infty} \prod_{j' \in N \setminus i} \prod_{x \in X} Y^x_{j',t}. \tag{4.3}$$

such that $\mathrm{P}^{\theta_1}(A^\sigma_i) \geq \delta$, for some number $\delta > 0$. It follows from our convention that $\sigma \geq 1$. Intuitively, on a sample path $\omega \in A^\sigma_i$, the maximum amount of "negative information" about state θ_1 vis-a-vis state θ that i can obtain from her own actions is bounded above by $\sigma^{|X|}$. We now consider each agent $j \in N(i)$ other than i. Since observations of individual agents are identically distributed (conditional on θ_1), it follows that for each neighbor $j \in N(i) \setminus i$, there exists a similarly defined set A^σ_j with $\mathrm{P}^{\theta_1}(A^\sigma_j) = \mathrm{P}^{\theta_1}(A^\sigma_i) = \delta > 0$. (This is done by just replacing i by j everywhere in equation (4.3)). Define the set $A_i = \cap_{j \in N(i)} A_j^\sigma$. Using the independence of individual observations, it follows that:

$$\mathrm{P}^{\theta_1}(A_i) = \delta^{|N(i)|} \geq \delta^K > 0. \tag{4.4}$$

The weak inequality holds since, by assumption, $|N(i)| \leq K$. Define $\lambda = \delta^K > 0$. Note that individual i's posterior belief about state θ_1 at time t can be written as:

$$\mu_{i,t}(\theta_1)(\omega) = \frac{\mu_{i,1}(\theta_1)(\omega)}{\mu_{i,1}(\theta_1)(\omega) + \sum_{\theta \neq \theta_1} \prod_{j \in N(i)} \prod_{x \in X} r_{j,t}^{x,\theta}(\omega) \mu_{i,1}(\theta)(\omega)}. \tag{4.5}$$

where $r_{j,t}^{x,\theta}(\omega)$ now refers to the product likelihood ratio along the sample path when the actions $\{C_{j,\tau}\}$ are chosen. We have :

$$\mu_{i,t}(\theta_1)(\omega) \geq \frac{\mu_{i,1}(\theta_1)(\omega)}{\mu_{i,1}(\theta_1)(\omega) + \sum_{\theta \neq \theta_1} \sigma^{K|X|} \mu_{i,1}(\theta)(\omega)} \tag{4.6}$$

by construction of the set A_i.[26] Since the expression on the right side of (4.6) is independent of t, it is evident that there will exist a value of $d \in (0,1)$ such that if $\mu_{i,1}(\theta_1) \geq d$ and $\omega \in A_i$, then $\mu_{i,t}(\omega)(\theta_1) \geq \hat{d}$ for all $t \geq 1$. The lemma follows. $\square$

In the context of Example 2.1 it follows from the theory of random walks that σ can be chosen to be 1. Since $\hat{d}$ can be any number greater than the cutoff $\mu = 1/2$, equation (4.6) implies that d can be chosen to equal $\hat{d}$. More generally, d like its counterpart λ will depend upon K. In the rest of this section, we shall suppose K is a fixed number.

Now consider the collection of agents $i \in N$ who have at most K neighbors each and such that $\mu_{i,1}(\theta_1) \geq d$, where d is as given by Lemma 4.1. Let $N_{K,d}$ be a maximal group of pairwise locally independent agents chosen from this collection, i.e., a subset of the above collection which has the highest cardinality.[27] We can show the following result:

[26]On this set of sample paths, irrespective of the choice of actions by $j \in N(i)$ upto time $t-1$ the corresponding $r_{j,t}^{x,\theta}$ will be bounded above by σ. See the discussion following equation (2.9) in Section 2.

[27]It is worth noting that there may be many such maximal groups of agents. For instance let N be the set of integers, with $N(i) = \{i-1, i, i+1\}$ for all $i \in N$. Fix $K = 3$ and suppose all agents $i \in N$ satisfy $\mu_{i,1}(\theta_1) \geq d$. Then the sets of agents $\{0, 3, 6, 9, 12, \ldots\}$, and $\{\ldots, -6, -2, 2, 6, 10, \ldots\}$ are just two of infinitely many possible candidates for $N_{K,d}$. Note that for an infinite society with a Royal Family the maximal group is at most a singleton.

Theorem 4.1 $\triangleright$ *Consider an infinite agent society which is connected and let $d \in (0,1)$ and $\lambda > 0$ be as defined above. Then*

$$\mathrm{P}^{\theta_1}(\cup_{i\in N}\{X^i(w) \not\subset G(\delta_{\theta_1})\}) \leq (1-\lambda)^{|N_{K,d}|}. \tag{4.7}$$

In particular, if $|N_{K,d}| = \infty$ then complete learning obtains.

We note the following observations concerning the above result before providing its proof. Firstly, if $|N_{K,d}| = \infty$ and $G(\delta_{\theta_1}) \subset X_I$, where X_I is the set of fully informative actions, then for all $i \in N$, $\mu_{i,\infty}(\omega) = \delta_{\theta_1}$, with probability 1: in other words, agents' beliefs converge almost surely to the truth. We also remark that in Example 4.1, if the Royal Family is removed, so that only the local neighborhoods remain, i.e., for all $i \in N$, $N(i) = \{i-1, i, i+1\}$, then complete learning obtains. This follows directly from the proof of Theorem 4.1 and shows that in some situations additional information links can actually lead to a lowering of long run social welfare! Finally, we note that the uniform upper bound on the number of neighbors of the locally independent agents can be relaxed to allow for the number of neighbors to increase at a sufficiently "slow" rate. Lemma 4.1 is the first step in the proof of Theorem 4.1. The remaining steps are now given:

Step 2: Since agents in the set $N_{K,d}$ are pairwise locally independent, the probability of every agent $i \in N_{K,d}$ trying a sub-optimal action is given by

$$\mathrm{P}^{\theta_1}(\cap_{i\in N_{K,d}} A_i^c) \leq (1-\lambda)^{|N_{K,d}|}. \tag{4.8}$$

Step 3: We use connectedness of the society and the following general argument to show that the probability of incomplete social learning is bounded above by the same expression. Consider some $\omega \in Q^{\theta_1}$. Suppose there is $i(\omega) \in N$ such that $C_{i(\omega),t} \in G(\delta_{\theta_1})$ all but finitely often. Since the set of actions X is finite, this implies the existence of $x' \in G(\delta_{\theta_1})$ such that on ω we have $x' \in X^{i(\omega)}(\omega)$. Remark B.1 (in the appendix) implies $U_{i(\omega),\infty}(\omega) = u(x', \delta_{\theta_1})$. This remark further implies that if $x \neq x'$ also lies in $X^{i(\omega)}(\omega)$, then $u(x, \delta_{\theta_1}) = u(x', \delta_{\theta_1})$. Hence $x \in G(\delta_{\theta_1})$ as well, so that $X^{i(\omega)}(\omega) \subset G(\delta_{\theta_1})$. Connectedness (and hence Theorem 3.2) implies that $U_{j,\infty}(\omega) = U_{i(\omega),\infty}(\omega)$ for all $j \in N$. Hence $U_{j,\infty}(\omega) = u(x', \delta_{\theta_1})$. Furthermore, using the same remark again, $U_{j,\infty}(\omega) = u(x, \delta_{\theta_1})$ for all $x \in X^j(\omega)$. Hence $X^j(\omega) \subset G(\delta_{\theta_1})$ as well. Thus $\cup_{i\in N_{K,d}} A_i \subset \cap_{j\in N}\{X^j(\omega) \subset G(\delta_{\theta_1})\}$ and (4.8) above implies $\mathrm{P}^{\theta_1}(\cup_{i\in N}\{X^i(\omega) \not\subset G(\delta_{\theta_1})\}) \leq (1-\lambda)^{|N_{K,d}|}$ as required. $\square$

We briefly discuss some simulations of Example 2.1. We suppose the agents in N are arranged in a circle with $N(i) = \{i-1, i, i+1\}$ (no Royal Family). Figure 1 displays the probability of incomplete learning as a function of societal size $|N|$ assuming the payoffs in Example 2.1 are Bernoulli distributed, while Figure 2 concerns the Normal case. We note that the probability decays quite rapidly with the size of the society; furthermore, a regression of the log incomplete learning probability on $|N|$ yields a very good fit (the R^2 values are all above 0.99 and between 0.94 and 0.98 in the Bernoulli and Normal cases respectively), and suggests that the bound in (4.7) is tight.

Theorem 4.1 allows for fairly general neighborhood structures[28] and imposes no restrictions on the informativeness of actions. The presence of a Royal Family, however, means that $|N_{K,d}| \leq 1$ so that the result no longer ensures complete learning. We now propose some conditions that are sufficient for learning in societies with a Royal Family. We start by noting that when the society contains a Royal Family, incomplete learning can arise even if beliefs of agents are 'highly optimistic', so long as they are bounded away from point mass on the truth.[29] This motivates a stronger restriction on beliefs, which we refer to as the heterogenous priors assumption.

> **(H)** The distribution of prior beliefs is *heterogeneous* if for every $\theta \in \Theta$, and for any open neighborhood around δ_θ, there exists an agent whose prior belief lies in that neighborhood.[30]

We now show that *in any infinite connected society with* $\sup_{i \in N} |N(i)| \leq K < \infty$ *and which satisfies (H), the probability of incomplete learning is less than* $1-\lambda$ *for every* $\lambda \in (0,1)$. *In particular, this is also true in the presence of a Royal Family* R *where* $|R| \leq K$. The intuition behind this claim is as follows: from arguments slightly extending those in Lemma 4.1,

[28] In particular, agents outside $N_{K,d}$ can have any (finite) number of neighbors. Theorem 4.1 also allows for complete learning in some societies having agents observed by infinitely many other agents. For instance in Example 4.1, if $N(i) = \{i-1, i, i+1\} \cup R$ for all $i \geq 0$ and $N(i) = \{i-1, i, i+1\}$ for all $i < 0$, complete learning occurs.

[29] To see why this is true, consider Example 4.1 again but suppose that when x_1 is chosen, the outcome is distributed according to the negative exponential distribution $\phi(y; x, \theta_k) = \mathbf{1}_{\{y<0\}} \theta_k^{-1} \exp(y/\theta_k)$ for $k = 0{,}1$. Assume $\theta_o > \theta_1 > 0$. Choose $u(x_o)$ to lie between $u(x_1, \delta_{\theta_o}) = -\theta_o$ and $u(x_1, \delta_{\theta_1}) = -\theta_1$. Suppose θ_1 is the true state and that all agents i have beliefs of at most $\hat{\mu} < 1$. It is not difficult to see that no matter how close $\hat{\mu}$ is to 1, there is a strictly positive probability of a large negative shock from the Royal Family which is enough to push all agents' beliefs to a level where they will only choose the suboptimal action x_o. Note that this can happen even if $|R| = 1$.

[30] Heterogeneity of beliefs may be interpreted as saying that the truth must lie in the support of the distribution of prior beliefs across agents. Since the true state is unknown, this requirement leads naturally to the formulation above, where for any θ, δ_θ lies in the support of the distribution of prior beliefs.

for every $\lambda \in (0,1)$ there is a corresponding σ and a set A_i with $\mathrm{P}^{\theta_1}(A_i) \geq \lambda$ such that the amount of negative information that agent i can get about state θ_1 from her own as well as her neighbor's actions is bounded above by $\sigma^{|X||K|}$. Using equation (4.6) we can now establish the existence of a $d \in (0,1)$ such that if $\mu_{i,1}(\theta_1) \geq d$ and $\omega \in A_i$, then the posterior $\mu_{i,t}(\theta_1) \geq \hat{d}$ for all time periods. It follows from (H) that there exists some individual i whose priors satisfy $\mu_{i,1}(\theta_1) \geq d$. Thus agent i will try an optimal action forever with probability at least λ. The argument now follows along the lines of Steps 2 and 3 of Theorem 4.1.

Condition (H) imposes strong restrictions on the distribution of prior beliefs[31] and it is therefore useful to consider alternative conditions under which complete learning may obtain. In our discussion following Example 4.1, we noted that if an agent is able to generate arbitrarily large amounts of positive information with positive probability, then locally independent groups of agents may be able to overcome negative information from the Royal Family, ensuring complete learning. This suggests the following condition on the informativeness of actions.

> **(UPI)** An action $x \in X$ generates unbounded positive information concerning the true state θ_1, if for every $\alpha \in (0,1)$, there is a set $B^{x,\alpha} \subset Y$ with $\int_{B^{x,\alpha}} \phi(y;x,\theta_1)d, \ > 0$ such that
>
> $$y \in B^{x,\alpha} \Rightarrow \max_{\theta \in \Theta \backslash \theta_1} \frac{\phi(y;x,\theta)}{\phi(y;x,\theta_1)} \leq \alpha. \tag{4.9}$$

In the presence of a Royal Family two agents cannot be locally independent. We thus consider the following generalization of local independence: two individuals $i \notin R$ and $i' \notin R$ are called *quasi-locally independent* if $N(i) \cap N(i') = R$. For more than two agents the corresponding condition is that of *pairwise quasi-local independence*. Recall that $\hat{d} \in (0,1)$ is such that if $\mu(\theta_1) \geq \hat{d}$ then $G(\mu) \subset G(\delta_{\theta_1})$. Let $\hat{N}_{K,\hat{d}}$ be a maximal group of pairwise quasi-locally independent agents with at most K neighbors each and having prior beliefs which satisfy $\mu_{i,1}(\theta_1) \geq \hat{d}$. We now have:

Proposition 4.1 $\triangleright$ *Consider a connected society with* $|\hat{N}_{K,\hat{d}}| = \infty$. *Suppose each* $x \in G(\delta_{\theta_1})$ *satisfies condition (UPI). Then complete learning obtains.*

[31]In particular, a single agent with a very positive prior on the truth is responsible for virtually all of the learning carried out by the society.

We note that since $R \subset N(i)$ for $i \in \hat{N}_{K,\hat{d}}$, the fact that $|R| \leq K$ is implicit in Proposition 4.1. The proof follows along the lines of Theorem 4.1 and may be found in Appendix C. The main difference arises in the case where $|R| > 0$. The argument for this case proceeds by contradiction. If learning is incomplete then there must exist a set which has positive probability, on which the negative information generated by the Royal Family is bounded above by some number and yet learning is incomplete. We use a construction similar to step 2 in Theorem 4.1 to establish that if there are an infinite number of quasi-locally independent agents then at least one of them will get sufficiently positive information in their first trial with an optimal action to offset this negative information. Thus at least one agent will try an optimal action forever on any sample path of this set. This observation taken along with the connectedness of society yields a contradiction and completes the proof.[32]

In the results described so far, social learning relies on the set of locally independent agents who each try optimal actions with positive probability from the first period onwards. We now examine the possibilities of complete learning when agents do not necessarily start with prior beliefs that favor optimal actions. In this setting, the likelihood of social learning is sensitive to the nature of information generated by non-optimal actions across agents, both regarding the payoffs of these actions themselves as well as the payoffs of optimal actions.[33]

We provide two alternative sets of sufficient conditions on the informativeness of actions. One set of conditions applies when realizations from an action convey no payoff relevant information concerning any other action. The second set of conditions deal with the complementary situation when realizations on an action can reveal information about other actions. The conditions are used in Theorem 4.2.

To state the result we need to introduce additional concepts. First note that $x \in X$ induces an ordered partition of the states denoted by $\Theta_1(x) \prec_x \Theta_2(x) \prec_x \ldots \prec_x \Theta_{s(x)}(x)$ such that

[32] To continue Example 4.1 further, suppose that when x_1 is chosen the outcome is distributed according to the normal, exponential, Poisson or geometric (but not, of course, the negative exponential!) distributions. Then the above result applies. It also holds more generally: for example if the density functions take one of the above forms and $x \in G(\delta_{\theta_1}) \Rightarrow u(x, \delta_{\theta_1}) > u(x, \delta_\theta)$ for all $\theta \neq \theta_1$ then complete learning obtains.

[33] In this context it is also worth noting that Proposition 4.1 also holds if an infinite number of quasi-locally independent agents have priors that lead them to try sub-optimal actions provided that these actions satisfy the condition (UPI).

(a) for each $k = 1, \ldots, s(x)$, the expected payoff $u(x, \delta_{\theta_k})$ is constant for all $\theta_k \in \Theta_k(x)$.
(b) if $\Theta_m(x) \prec_x \Theta_k(x)$ then $u(x, \delta_{\theta_m}) < u(x, \delta_{\theta_k})$ for $\theta_m \in \Theta_m(x)$ and $\theta_k \in \Theta_k(x)$.

For $x \in X$, let $k(x)$ denote the payoff equivalence set of states of nature which contains θ_1, i.e. $\theta_1 \in \Theta_{k(x)}(x)$. Also let $\Theta(x)^{++} \equiv \cup_{m>k(x)}\Theta_m(x)$, $\Theta(x)^{+} \equiv \cup_{m\geq k(x)}\Theta_m(x)$ and $\Theta(x)^{-} \equiv \cup_{m<k(x)}\Theta_m(x)$. The first set of assumptions on informativeness of actions are given by condition (I) stated below:

> **(Ia)** For $x,\ x' \in X$, where $x' \neq x$, if action x' is chosen and $y \in Y$ is observed, then for any $\mu \in \mathcal{D}(\Theta)$ the posterior belief $\mu(\Theta_m(x))' = \mu(\Theta_m(x))$ for each $m = 1, \ldots, s(x)$.
>
> **(Ib)** There exists $x_1 \in G(\delta_{\theta_1})$ such that if x_1 is chosen and $y \in Y$ is observed, then $\phi(y; x_1, \theta)/\phi(y; x_1, \theta_1) = 1$ for all $\theta \in \Theta_{k(x_1)}(x_1)$.
>
> **(Ic)** For x_1 as above, there exists a set $B^{x_1} \subset Y$ satisfying $\int_{B^{x_1}} \phi(y; x_1, \theta_1)d,\ (y) > 0$ and $\alpha \in (0, 1)$ such that
>
> $$y \in B^{x_1} \Rightarrow \frac{\phi(y; x, \theta)}{\phi(y; x, \theta_1)} \leq \alpha < 1. \tag{4.10}$$
>
> for all $\theta \in \Theta(x_1)^{-}$.

Condition (I) can be best understood in terms of the canonical bandit model of Example 2.2. Condition (Ia) requires that there be no essential information flows across actions, i.e. actions are independent of each other. Condition (Ib) says that the action x_1 is incapable of distinguishing between states which are payoff equivalent for it: in the bandit model, payoff equivalent states for x_1 correspond to states where the quality types of actions other than x_1 vary. As actions are independent, x_1 will not be able to distinguish between these states. Condition (Ic) requires that x_1 should be capable of generating a minimum amount of negative information concerning payoff inferior states. In the bandit model if the conditional density functions $\{\phi(\cdot)\}$ have the standard monotone likelihood ratio property (MLRP), then (Ic) holds.

We now impose some restrictions on beliefs. Let $x_1 \in G(\delta_{\theta_1})$ be as above. By definition, it must be the case that $u(x_1, \delta_{\theta_1}) > \max_{x \in X \setminus G(\delta_{\theta_1})} u(x, \delta_{\theta_1})$. Hence we can find $\xi \in (0, 1)$ and $\epsilon > 0$ such that

$$\xi u(x_1, \delta_{\theta_1}) + (1 - \xi)u(x_1, \delta_{\theta_L}) \geq \max_{x \in X \setminus G(\delta_{\theta_1})} u(x, \delta_{\theta_1}) + \epsilon \equiv u_{\min}. \tag{4.11}$$

where $\theta_L \in \Theta_1(x_1)$. Recall that $\Theta_{k(x_1)}(x_1)$ is the set of states payoff equivalent to state θ_1 for action x_1. Consider the collection of agents $i \in N$, who have at most K neighbors each and such that $\mu_{i,1}(\Theta_{k(x_1)}(x_1)) \geq \xi$ for each i. Let $N_{K,\xi}$ be a maximal group of pairwise locally independent agents chosen from this collection. The restriction on the belief of an agent $i \in N_{K,\xi}$ ensures that i will choose x_1 at least once; however, it does not preclude suboptimal actions from being chosen at the outset.

We next consider the class of situations where actions potentially provide information on states which are payoff relevant for other actions. Recall that X_I is the set of fully informative actions. Assume that $X = X_I \cup \{x_u\}$ and let $x_1 \in G(\delta_{\theta_1})$ be given. The case where $x_u \in G(\delta_{\theta_1})$ is trivial and thus there is no loss of generality in assuming that $x_1 \in X_I$. We now state the alternative conditions on the informativeness of actions.

(**Ia***) For each $x \in X \backslash x_u$, there exists a set $B^x \subset Y$ satisfying $\int_{B^x} \phi(y; x, \theta_1) d, \ (y) > 0$ such that if action x is chosen and $y \in B^x$ is observed, then for any $\mu \in \mathcal{D}(\Theta)$ the posterior belief $\mu(\Theta(x_1)^+)' \geq \mu(\Theta(x_1)^+)$.

(**Ib***) For each $x \in X \backslash x_u$ and for B^x as in (Ia*) above there exists $\alpha(x) \in (0,1)$ such that $y \in B^x$ implies $\max_{\theta \in \Theta \backslash \theta_1} \phi(y; x, \theta)/\phi(y; x, \theta_1) \leq \alpha(x)$.

Condition (Ia*) requires that for each informative action x, if $y \in B^x$ is observed then this does not yield negative information concerning payoffs of the optimal action x_1. Condition (Ib*) requires that all informative actions should be capable of generating a certain minimum amount of positive information concerning the true state. We note that condition (I*) is always satisfied when $|\Theta| = 2$.

As before, fix $\epsilon > 0$ and $\xi^* \in (0,1)$ such that $\xi^* u(x_1, \delta_{\theta_1}) + (1 - \xi^*) u(x_1, \delta_{\theta_L}) \geq u(x_u) + \epsilon$ where $\theta_L \in \Theta_1(x_1)$. Let N_{K,ξ^*} be a maximal collection of locally independent agents having at most K neighbors each and whose beliefs satisfy $\mu_{i,1}(\theta_1) \geq \xi^*$. We can now state the following theorem.

Theorem 4.2 ▷ *Consider an infinite agent society which is connected. (a) Suppose that actions satisfy condition* (I). *Then there exists* $\lambda \in (0,1]$ *such that*

$$\mathrm{P}^{\theta_1}\Big(\bigcup_{i \in N} \{X^i(\omega) \not\subset G(\delta_{\theta_1})\}\Big) \leq (1 - \lambda)^{|N_{K,\xi}|}. \tag{4.12}$$

In particular if $|N_{K,\xi}| = \infty$ *then complete learning obtains. (b) The above conclusions continue to hold if condition* (I) *is replaced by condition* (I*) *and* $N_{K,\xi}$ *by* N_{K,ξ^*} *everywhere.*

We briefly intuition behind Theorem 4.2.[34] The basic difference from the earlier results lies in the construction of the set A_i. In part (a) we show that for a sample path in A_i, agent i will observe a critical number of trials T with the optimal action x_1. By virtue of (Ic) this is sufficient to ensure that the agent will choose only the optimal action from some finite time onwards.

We briefly discuss how condition (I*) is used in part (b). As in Lemma 4.1, we can isolate a set of sample paths A_i, on which the amount of negative information obtained by agent $j \in N(i)$ concerning θ_1 is uniformly bounded above by σ. Recall that $\hat{d}$ is a number such that $\mu(\theta_1) \geq \hat{d}$ implies $G(\mu) \subset G(\delta_{\theta_1})$. Let $\alpha = \max_{x \in X \setminus x_u} \alpha(x)$; since $\alpha < 1$, we can choose T to satisfy $\xi^*/(\xi^* + \alpha^T \sigma^{K|X|}(1-\xi^*)) \geq \hat{d}$. Define A_j^σ as follows:

$$A_j^\sigma = \prod_{x \in X \setminus x_u} \{\prod_{t=1}^{T} B_{j,t}^x \times \{\max_{\theta \in \Theta \setminus \theta_1} \sup_{t \geq T+1} r_i^{x,\theta}(T+1,t) \leq \sigma\}\} \times \prod_{t=1}^{\infty} \prod_{j \in N(i)} \prod_{x \in X} Y_{j,t}^x \qquad (4.13)$$

where $B_{j,t}^x = B^x$ for $j \in N(i)$ and all $t \geq 1$. Let $A_i = \cap_{j \in N(i)} A_j^\sigma$; familiar arguments can be used to establish that $\mathrm{P}^{\theta_1}(A_i) = \lambda > 0$. Using condition (Ia*) we next show that along sample paths in A_i, the choices $C_{i,t} \neq x_u$, for all $t \leq T$. This guarantees that agent i tries an informative action long enough and generates positive information that is sufficient to offset any subsequent negative information concerning state θ_1. The rest of the proof is standard.

The discussion so far has focused on the optimality of long run actions: we now summarize our findings on the distribution of limit beliefs. Recall that the beliefs of every agent converge almost surely (Theorem 3.1). An issue of importance is whether agents learn the truth, i.e., if limit beliefs place point mass on the true state. In general, even in cases where long run actions are optimal, there is no guarantee that beliefs will converge to the truth. This is because the support of the limiting beliefs distribution depends crucially on the informativeness of the optimal actions.[35] However, as we remarked after Theorem 4.1, if an agent chooses optimal actions in the long run and these actions are fully informative

[34]The proof of part (a) is given in Appendix C. The proof of part (b) is similar, and is omitted.

[35]It is not difficult to construct instances of Example 2.2 (the canonical bandit model) where beliefs fail almost surely to place point mass on the truth despite long run actions being optimal.

about the true state then the agent will learn the truth. Lastly, we note that in Proposition 5.1 below, with positive probability some agents in a society learn the truth state while simultaneouly other agents do not.

5. Conformism *vs.* Diversity

This section studies the relationship between the social structure of information flow and the likelihood of diversity in a connected society.[36] We begin by noting an implication of Theorem 3.2: in a connected society, two actions with different expected payoffs (conditional on the true state of the world) will be chosen in the long run, with probability zero. This observation leads us to examine the possibility of two equally attractive actions surviving in the long run.

To analyze issues of conformity/diversity in general is quite a difficult problem. We therefore study these phenomena in the context of an example. Our analysis takes as fixed an individual decision problem and a certain distribution of priors. We then vary the structure of neighborhoods and look at how the probability of diversity varies. We use a combination of analytical and simulation methods. Our findings suggest that probability of diversity is inversely related to the *degree of social integration*, as elaborated below.

Example 5.1: We consider the following special case of Example 2.2. There are two actions x_1 and x_2 each of two possible types, High (H) and Low (L). An action of High type yields outcomes (and rewards) $y = 1$ and 0 with probabilities π and $1 - \pi$ respectively, where $\pi \in (1/2, 1)$ is a parameter. An action of the Low type also yields outcomes of 1 and 0 but with probabilities $1 - \pi$ and π. Thus there are four states of nature, labelled as (H,H), (H,L), (L,H) and (L,L). The types of actions are *independent* of each other and the beliefs of an agent are given by a pair $(\mu, \nu) \in [0, 1]^2$ where μ is the probability that x_1 is a High type and ν the probability that x_2 is a High type. Since agents are assumed to choose actions to maximize single-period expected utility, it is clear by symmetry that an agent will choose x_1 if $\mu > \nu$, x_2 if $\mu < \nu$ and arbitrarily if $\mu = \nu$. Furthermore, denoting the outcome (1 or 0) from choosing x_1 by Z, the posterior belief is

[36]We shall say that a society exhibits conformism if all agents choose the same action in the long run, with probability 1. Correspondingly, a society exhibits diversity if different agents choose different actions in the long run, with positive probability.

(μ', ν) where

$$\mu' = \frac{\mu\pi^Z(1-\pi)^{1-Z}}{\mu\pi^Z(1-\pi)^{1-Z} + (1-\mu)(1-\pi)^Z\pi^{1-Z}} \equiv \frac{\mu}{\mu + (1-\mu)p^{2Z-1}}. \quad (5.1)$$

where $p \equiv (1-\pi)/\pi \in (0,1)$. A similar expression may be derived for ν if x_2 is chosen instead of x_1. Equation (5.1) implies that if an agent observes $n \geq 0$ independent trials with x_1, and $m \geq 0$ trials with x_2, then x_1 will be chosen in the current period if

$$\frac{\mu}{\mu + (1-\mu)p^{2S_n - n}} > \frac{\nu}{\nu + (1-\nu)p^{2R_m - m}} \Rightarrow p^{2(S_n - R_m)+m-n} < \frac{\mu(1-\nu)}{\nu(1-\mu)}. \quad (5.2)$$

where S_n denotes the number of successes in n trials with x_1 and R_m the number of successes in m trials with x_2. (We adopt the convention that $S_n = 0$ if $n = 0$ and $R_m = 0$ if $m = 0$).

We next specify the distribution of priors. We shall suppose there are $2k+2$ agents labelled as $1, \ldots, k$, α, β and $k+1, \ldots, 2k$. We shall refer to agents 1 to k and α collectively as the group N_1, and the rest of the agents as group N_2. Agents in N_1 have the same prior belief (μ_1, ν_1) with $\mu_1 > \nu_1$; thus they choose x_1 in the first period. Agents in the N_2 all have the common belief given by (μ_2, ν_2) with $\nu_2 > \mu_2$; thus they choose x_2 in the first period. Finally, we shall suppose that the (unknown) true state of nature is $\theta_1 = (H, H)$.

We consider the following general information structure. All observation relations are assumed to be symmetric. There is public observability within Group N_1 (i.e. all agents in the group observe each other) and likewise within Group N_2. The observation links across the two groups is specified by a *degree of integration* parameter $\eta \in \{1, \ldots, k+1\}$. If the degree of integration is η this means that agents $\{1, \ldots, \eta - 1, \alpha\}$ observe (and are observed by) agents $\{\beta, k+1, \ldots, k+\eta-1\}$.[37] Note that for any k, the case $\eta = k+1$ corresponds to full public observability while $\eta = 1$ is the minimum required for a society to be connected. Figure 3 shows a society with $k = 3$ (i.e. 8 agents) when $\eta = 1$, 2 and $k+1$. The first result derives conditions on the degree of integration that are sufficient for diversity to obtain.

[37]This structure is motivated by examples of societies with well defined sub-groups – based on ethnic, linguistic, spatial or professional considerations – where interaction within a sub-group is very regular and at a high level but only a few members of a sub-group interact with members from another sub-group.

Proposition 5.1 ▷ *Suppose the society is as described above. If $k - \eta + 1 > \eta/(2\pi - 1)$ then there is a strictly positive probability that all agents in group N_1 will choose x_1 forever and all agents in group N_2 will choose x_2 forever.*

The idea underlying the proof[38] can be explained by considering the case of $\eta = 1$: First, applying standard arguments we identify a set of sample paths, which has positive probability, and on which each agent, acting in isolation, will choose his initial action forever. We then consider agent α's behavior when the society is linked in the manner described above. In each period, agent α observes the other members of N_1 choosing x_1 and also sees β, who is choosing x_2. If $k > 1/(2\pi - 1)$ then, with the help of the Law of the Iterated Logarithm, we can show that agent α will receive positive information about x_1 at a "faster" rate than positive information about x_2 and so will choose x_1 forever. Furthermore, all $i \in N_1 \backslash \alpha$ will choose x_1 forever since they only observe positive information on x_1 and no information on x_2. By symmetry, agent β will choose x_2 forever as will all members of N_2. Thus diversity occurs for the following reason: the agents at the "boundary" maintain their original choices because they are *more highly connected to agents who choose as they do as compared to those who choose differently.* If the information forthcoming from their own experience and experience of like-minded neighbors' actions is sufficiently positive it 'overcomes' the positive information concerning the other action coming from across the boundary, and the agents maintain their original decision.

It is worth noting some additional features of this result. As the agents α and β observe actions x_1 and x_2 infinitely often, they will learn the true state in the long run. However, the remaining agents observe only either x_1 or x_2 and therefore do not learn the true state. Thus it is possible for some agents in a connected society to learn the truth while others simultaneously do not. Secondly, the example described above exhibits "path-dependency" since there exists another set of sample paths having positive probability such that one action will become extinct in finite time. The proof of this statement is not difficult and is omitted.

By way of contrast, we now consider the case of a fully integrated society, i.e. with $\eta = k + 1$, for which the above proposition does not apply. Our analysis is summarized by the following result.

[38] A proof of Proposition 5.1 is presented in Appendix D.

Proposition 5.2 ▷ *Let the society be as described above and suppose that $\eta = k + 1$, so that there is full public observability. Then conformity occurs, i.e. with probability 1, all agents choose the same action forever after a finite time.*

The intuition[39] for this result is as follows: Expressions (5.1)-(5.2) describe the relationship between information generated by trials, summarized by $L \equiv 2(S_n - R_m) - n + m$, and the priors of agents. In particular, they imply that for 'high' values of L, both groups of agents will choose action x_1, for 'low' values of L both groups of agents will choose action x_2, while for intermediate values of L, group 1 will choose action x_1 while group 2 chooses action x_2. The process L can be represented as a stationary Markov Process; the proof then follows from the observation that for high (low) values of L, the process is a random walk with a positive drift (negative drift), while for intermediate values it is a symmetric random walk.

We have so far derived sufficient conditions on the degree of integration for diversity to obtain and also shown that when a society is fully integrated then conformity obtains. We now examine how the probability of diversity varies in relation to the parameter of integration, with the help of simulations. As a check on the robustness of our analysis, we also consider the case of actions with normally distributed outcomes. These simulations are summarized in Figures 4 and 5, respectively. Our simulations were carried out for a society with $k = 9$ (total number of agents is 20). The figures suggest that there exists an inverse relationship between the degree of integration and the likelihood of diversity. The intuition for this is that as η increases, the information set of agents on the 'boundary' becomes more alike, making different optimal choices by them less likely. Thus the agents who make the 'across-group' observations (such as α and β) choose the same actions. This allows information about both actions to flow into a sub-group thereby precipitating a breakdown in the 'information barrier' separating N_1 and N_2 and eventually leading to social conformism.The simulations also suggest that for fixed η, the likelihood of diversity increases as π and the value of the mean (in the normal distribution case, keeping standard deviation constant) decreases. Lowering π or the value of the mean in the normal case both have the effect of decreasing the informativeness of action x_1, which allows the influence of the initial differences in priors of the two groups N_1 and N_2 to survive more easily.

[39]A proof of this result is available from the authors upon request.

6. Temporal and Spatial Patterns of Learning

While the results of Sections 3-5 characterize the long run outcomes in our framework, they do not tell us much about the temporal and spatial evolution of social learning. In this section we discuss simulations of our framework to get some idea about these issues. In particular, we wish to compare the results of our simulations with the findings of the extensive empirical literature on diffusion as a means of validating our theoretical paradigm.

We assume the following social structure: The set of farmers N is arranged in a $k \times k$ grid, with each farmer owning a single plot of land. In our simulations we take $k = 20$, so that we have a total of 400 agents. Each farmer i observes the actions and payoffs (observations) of her surrounding 8 neighbors. We perform simulations under different specifications, which are special cases of Example 2.2. We now summarize our findings.[40]

Temporal Patterns: In the first simulation, we assume that there are two crops, one of which (Crop 0) has a known payoff of 1/2, while the other (Crop 1) represents a new, unknown technology. Crop 1 can be of quality level $q_1 = 0.45$ or $q_2 = 0.55$; if the crop is of quality q_k for $k = 1, 2$ then its payoff is Bernoulli-distributed with parameter q_k. We suppose that the true quality of Crop 1 is q_2, and so it is better than Crop 0. We also assume that the farmers' beliefs at the beginning of period 1 are heterogeneous, with about 1% of the farmers having a prior above 1/2 and therefore experimenting with the new crop.

The diffusion curve of a typical simulation is given in Figure 6a. As can be seen, the logistic curve fitted from the data matches the adoption curve quite well. The R^2 is 0.987, which is in the same range as obtained by Griliches (1957) in his study of the diffusion of hybrid corn. We also report a simulation where the new technology is more profitable than in the earlier case (we chose $q_1 = 0.43$ and $q_2 = 0.57$ as the quality levels). The adoption curve for this simulation is given in the Figure 6b. As can be seen, the logistic still provides a good fit ($R^2 = 0.99$); however, the adoption rate is far higher, as it takes approximately half the time for the population to adopt compared to the earlier case. This is consistent with the result of Griliches, who found that the adoption rate was strongly positively linked to profitability. Finally, we also note that both adoption curves exhibit small downward

[40] In our simulations, the opposite edges of the rectangular grid are identified with each other to ensure that all farmers have 8 neighbors, including those living along an edge.

fluctuations, an empirical phenomenon which has been discussed by Rogers (1983). As a check on the robustness of these patterns, we also ran simulations of a two crop model in which the returns of the new crop were normally distributed, with unknown means. Two typical simulations are plotted in Figures 7a and 7b. The R^2 values are 0.988 and 0.982 respectively. These figures corroborate the findings that emerged from the Bernoulli case.[41]

Spatial Patterns: We begin with a simulation of the two crop model discussed above when $q_1 = 0.45$ and $q_2 = 0.55$. Figure 8 shows the spatial evolution of such a simulation. Initially, there are only 3 farmers who experiment with the new crop. By $t = 25$, one farmer has dropped out due to bad experiences with the new crop. However, a group of agents around the other two farmers have chosen the new crop as well. By $t = 50$ the two clusters are almost in contact with each other, after which the adoption rate increases rapidly. (At $t = 50$, the proportion of adopters is about 0.15, while at $t = 100$ it has almost tripled to 0.41). By $t = 200$ adoption is nearly complete. This pattern of spatial diffusion is consistent with empirical evidence (Hagerstrand, 1969; Rogers 1983).

We next consider a model which allows for the possibility that some of the actions may be payoff equivalent. This setting is interesting as it allows us to examine whether learning from neighbors can generate diversity in social structures other than the class explored in Section 5. In our example, there are a total of 4 crops. There are 3 quality types, with $q_1 = 0.45$, $q_2 = 0.55$ and $q_3 = 0.60$. As before, Crop 0 has a known payoff of 1/2. In our simulations, we suppose that crops 1 and crop 3 are the most profitable, having quality type q_3, while crop 2 is of type q_2. The ex-post ranking of crops in increasing order of profitability is therefore $\{0, 2, (1, 3)\}$. Finally, farmers have heterogeneous prior beliefs, which makes the initial choice of crops random.

The results of a typical simulation are presented in Figure 9. This simulation can be summarized as follows : (a) the pattern of crop choice begins to display features of clustering very quickly. (b) Over time, less profitable crops get replaced by more profitable ones. (c) In the long run, only the most profitable ones survive, with agents growing the same crop being linked together reflecting local conformism.

[41] All four figures also reveal high positive serial correlation of the residuals from the logistic fit. Given the local learning structure of our model, this is intuitive, and suggests a (reduced-form) test of the hypothesis of neighborhood learning.

7. Concluding Remarks

When payoffs from different actions are unknown, agents use their own past experience as well as the experience of their colleagues, friends and acquaintances as a guide for current decisions. We model these information flows across agents in terms of neighborhoods of individual observation. Our analysis suggests that the structure of these neighborhoods has important implications for the likelihood of adoption of new technologies, the coexistence of different practices as well as the temporal and spatial patterns of diffusion in a society. Our conclusions raise an important question: what types of information structures are likely to occur/emerge in societies?

Appendix A

We begin with a formal construction of the probability space, $(\Omega, \mathcal{F}, \mathrm{P}^\theta)$. Fix $\theta \in \Theta$. For each $i \in N$, $x \in X$ and $t = 1, 2, \ldots$ let $Y_{i,t}^x \equiv Y$. For each $t = 1, 2, \ldots$ let $\Omega_t = \prod_{i \in N} \prod_{x \in X} Y_{i,t}^x$ be the space of the t^{th} outcomes across all agents and all actions. Ω_t is endowed with the product topology. Let $H_t \subset \Omega_t$ be of the form

$$H_t = \prod_{i \in N} \prod_{x \in X} H_{i,t}^x \tag{A.1}$$

where $H_{i,t}^x$ is a Borel subset of $Y_{i,t}^x$ for each $i \in N$ and $x \in X$. (If the number of agents n is countably infinite, $H_{i,t}^x \equiv Y_{i,t}^x$ for all but a finite set of i's). Define the probability P_t^θ of the set H_t as :

$$\mathrm{P}_t^\theta(H_t) = \prod_{i \in N} \prod_{x \in X} \int_{H_{i,t}^x} \phi(y; x, \theta)\mathrm{d}, (y). \tag{A.2}$$

P_t^θ extends uniquely to the σ-field on Ω_t generated by sets of the form H_t. Let $\Omega = \prod_{t=1}^\infty \Omega_t$. For cylinder sets $H \subset \Omega$ of the form

$$H = \prod_{t=1}^{T} H_t \times \prod_{t=T+1}^{\infty} \Omega_t. \tag{A.3}$$

let $\mathrm{P}^\theta(H)$ be defined as $\mathrm{P}^\theta(H) = \prod_{t=1}^T \mathrm{P}_t^\theta(H_t)$. Let $\mathcal{F}$ be the σ-field on Ω generated by sets of the type given by (A.3). P^θ extends uniquely to the sets in $\mathcal{F}$. This completes the construction of the probability space $(\Omega, \mathcal{F}, \mathrm{P}^\theta)$.

Let Θ be endowed with the discrete topology, and suppose $\mathcal{B}$ is the Borel σ-field on this space. For rectangles of the form $A \times H$ where $A \subset \Theta$ and H is a measurable subset of Ω, let $\mathrm{P}_i(A \times H)$ be given by

$$\mathrm{P}_i(A \times H) \;=\; \sum_{\theta \in A} \mu_{i,1}(\theta)\mathrm{P}^{\theta}(H). \qquad (A.4)$$

for each agent $i \in N$. Each P_i extends uniquely to all of $\mathcal{B} \times \mathcal{F}$. Since every agent's prior belief lies in the interior of $\mathcal{D}(\Theta)$, the measures $\{\mathrm{P}_i\}$ are pairwise mutually absolutely continuous.

Appendix B

Proof of Theorem 3.1: For each $\theta \in \Theta$, the belief $\mu_{i,t}(\theta)$ of agent i at the beginning of time t can be regarded as a version of the conditional expectation $E[\mathbf{1}_{\{\theta\}\times\Omega} \| \mathcal{F}_{i,t}]$ where the expectation is with respect to the measure P_i. Since this sequence of random variables is a uniformly bounded martingale (see Easley and Kiefer, 1988) with respect to the increasing sequence of σ-fields $\{\mathcal{F}_{i,t}\}$ the Martingale Convergence Theorem applies, so that $\mu_{i,t}$ converges almost surely to the $\mathcal{F}_{i,\infty}$-measurable limit belief $\mu_{i,\infty}$. Let Q_i be the set of sample paths where agent i's beliefs converge, where $P_i(Q_i) = 1$. Since the measures are pairwise mutually absolutely continuous and the set of agents N is at most countable, the set $Q = \cap_{i \in N} Q_i$ also has P_i measure 1 for each i. $\square$

Proof of Lemma 3.1: Let $x \in X$. Since $x' \in X^i(\omega)$ there exists a subsequence $\{t_k\}$ such that $u(x', \mu_{i,t_k}(\omega)) \geq u(x, \mu_{i,t_k}(\omega))$. Taking limits and using the continuity of u on the set $\mathcal{D}(\Theta)$, we get $u(x', \mu_{i,\infty}(\omega)) \geq u(x, \mu_{i,\infty}(\omega))$. Since x is arbitrary, this proves statement (a). Statement (b) follows from the maximum theorem and part (a). $\square$

Let supp(μ) to denote the support of a probability distribution μ. We have:

Lemma 3.2 $\triangleright$ *Suppose* $i \in N(j)$ *and* $\omega \in Q^{\theta_1}$. *If, for some* $\theta \neq \theta_1$, $\theta \in \mathrm{supp}(\mu_{j,\infty}(\omega))$ *then* $u(x, \delta_\theta) = u(x, \delta_{\theta_1})$ *for all* $x \in X^i(\omega) \cup X^j(\omega)$.

Proof: Suppose the conditions of Lemma 3.2 hold but $u(x, \delta_{\theta_1}) \neq u(x, \delta_\theta)$ for some $x \in X^i(\omega) \cup X^j(\omega)$. Then, by definition, we have

$$\int_Y |\phi(y; x, \theta_1) - \phi(y; x, \theta)| \mathrm{d}, (y) > 0. \qquad (B.1)$$

Since x is chosen infinitely often either by agent i or by j (or both), and agent j observes agent i, the law of large numbers ensures that $\mu_{j,\infty}(\theta)(\omega) = 0$, so that θ is not in the support of $\mu_{j,\infty}(\omega)$. This contradiction establishes the result. $\square$

Remark B.1: Since $i \in N(i)$ for every $i \in N$, the above lemma implies that for every $x \in X^i(\omega)$, $u(x, \cdot)$ is constant on the set

$$\{\mu \mid \text{supp}(\mu) \subset \text{supp}(\mu_{i,\infty}(\omega))\}.$$

In particular, $U_{i,\infty}(\omega) \equiv u(x, \mu_{i,\infty}(\omega)) = u(x, \delta_{\theta_1})$ for each $x \in X^i(\omega)$.

Lemma 3.3 $\triangleright$ *Suppose* $\omega \in Q^{\theta_1}$. *If* $i \in N(j)$, *then* $U_{j,\infty}(\omega) \geq U_{i,\infty}(\omega)$.

Proof: We shall show that if $x' \in X^j(\omega)$, then $u(x', \delta_{\theta_1}) \geq u(x, \delta_{\theta_1})$, for all $x \in X^i(\omega)$. This will suffice for the proof since from Lemma 3.2 and Remark B.1 we have

$$U_{j,\infty}(\omega) \equiv u(x', \mu_{j,\infty}(\omega)) = u(x', \delta_{\theta_1}) = u(x', \delta_{\theta}) \text{ for all } \theta \in \text{supp}(\mu_{j,\infty}). \qquad (B.2)$$

and

$$U_{i,\infty}(\omega) \equiv u(x, \mu_{i,\infty}(\omega)) = u(x, \delta_{\theta_1}) = u(x, \delta_{\theta}) \text{ for all } \theta \in \text{supp}(\mu_{j,\infty}). \qquad (B.3)$$

There are two cases: if $\mu_{j,\infty}(\omega) = \delta_{\theta_1}$ the result follows trivially from Lemma 3.1. In the second case, suppose that $\theta \neq \theta_1$ also lies in the support of $\mu_{j,\infty}(\omega)$. We now proceed by contradiction. Assume that $u(x', \delta_{\theta_1}) < u(x, \delta_{\theta_1})$. Since $\theta \neq \theta_1$ lies in the support of $\mu_{j,\infty}(\omega)$, Lemma 3.2 above together with the facts that $x' \in X^j(\omega)$ and $x \in X^i(\omega)$ implies that $u(x', \mu_{j,\infty}(\omega)) < u(x, \mu_{j,\infty}(\omega))$. However this contradicts Lemma 3.1 above and hence $u(x', \delta_{\theta_1}) \geq u(x, \delta_{\theta_1})$. $\square$

Proof of Theorem 3.2: If i and j are two agents in N, then either $i \in N(j)$ or there exist agents $j_1, \ldots, j_m$ such that $j_1 \in N(j)$, $j_2 \in N(j_1)$ and so on until $i \in N(j_m)$. In the first case, Lemma 3.3 applies directly to show that $U_{j,\infty}(\omega) \geq U_{i,\infty}(\omega)$ while in the latter case the same is true by transitivity. The result follows by interchanging the roles of i and j. $\square$

Appendix C

Let $i \in N$. If agent i were to choose $x \in X$ between period t and $t' - 1$ and observe the corresponding outcomes $\{y^x_{i,n}\}_{n=t}^{t'-1}$, the product likelihood ratio of state θ with respect to θ_1 at the beginning of time t' would be:

$$r_i^{x,\theta}(t, t') = \prod_{n=t}^{t'-1} \frac{\phi(y^x_{i,n}; x, \theta)}{\phi(y^x_{i,n}; x, \theta_1)}. \qquad (C.1)$$

By convention we assume that $r_i^{x,\theta}(t,t') = 1$ if $t = t'$. Moreover, if $t = 1$ we write $r_i^{x,\theta}(1,t')$ simply as $r_{i,t'}^{x,\theta}$.

Proof of Proposition 4.1 (Sketch): Let $j \in N$. For $\alpha \in (0,1)$ and $x \in G(\delta_{\theta_1})$ let $B_{j,1}^{x,\alpha}$ be the set $B^{x,\alpha}$ whose existence is assumed in condition (UPI). Using arguments analogous to Lemma 4.1, we can establish that there exists a $\sigma \geq 1$, an $\alpha \in (0,1)$ such that $\alpha\sigma^{K|X|} < 1$ and a set A_j^σ defined as:

$$A_j^\sigma = \prod_{x \in G(\delta_{\theta_1})} B_{j,1}^{x,\alpha} \times \{\max_{x \in G(\delta_{\theta_1})} \sup_{\tau \geq 2} r_j^{x,\theta}(2,\tau) \leq \sigma\} \times \{\max_{x \in X \setminus G(\delta_{\theta_1})} \sup_{t \geq 1} r_{j,t}^{x,\theta} \leq \sigma\} \times \prod_{t=1}^{\infty} \prod_{j' \in N(j)} \prod_{x \in X} Y_{j',t}^x. \tag{C.2}$$

such that $\mathrm{P}^{\theta_1}(A_j^\sigma) = \delta > 0$ (by using the assumption that $x \in G(\delta_{\theta_1})$ satisfies the (UPI) property). Fix $i \in \hat{N}_{K,\hat{d}}$. Define $A_i = \cap_{j \in N(i)} A_j^\sigma$. Clearly $\mathrm{P}^{\theta_1}(A_i) \geq \delta^K > 0$. Note that since agent i is assumed to have a belief $\mu_{i,1}(\theta_1) \geq \hat{d}$, she will choose an action $x \in G(\delta_{\theta_1})$; by construction of the set A_i, she will observe an outcome $y \in B^{x,\alpha}$. As such a y provides sufficiently strong positive information concerning state θ_1, agent i's posterior belief will be very close to the truth, as in equation (4.6). The proof for the case of $|R| = 0$ now follows along the lines of steps 2 and 3 of Theorem 4.1, and is omitted.

The case $|R| > 0$*:* Let $\hat{Q} = \{\cup_{i \in N} X^i(\omega) \not\subset G(\delta_{\theta_1})\}$. We shall assume $\mathrm{P}^{\theta_1}(\hat{Q}) > 0$ initially. Clearly, there exists $\sigma \geq 1$ (without loss of generality having the same value as above) such that $\mathrm{P}^{\theta_1}(\hat{Q} \cap A_R^\sigma) > 0$, where A_R^σ is the set

$$A_R^\sigma = \bigcap_{j \in R} \{\max_{\theta \in \Theta \setminus \theta_1} \max_{x \in X} \sup_{t \geq 1} r_{j,t}^{x,\theta} \leq \sigma\} \times \prod_{t=1}^{\infty} \prod_{j \in N \setminus R} \prod_{x \in X} Y_{j,t}^x. \tag{C.3}$$

For $i \in \hat{N}_{K,\hat{d}}$ consider the set A_i constructed as above, but excluding all $j \in N(i)$ who are members of R. The probability of A_i conditional on A_R^σ satisfies

$$\mathrm{P}^{\theta_1}(A_i \mid A_R^\sigma) = \frac{\prod_{j \in N(i) \setminus R} \mathrm{P}^{\theta_1}(A_j^\sigma) \times \mathrm{P}^{\theta_1}(A_R^\sigma)}{\mathrm{P}^{\theta_1}(A_R^\sigma)} \geq \delta^K > 0. \tag{C.4}$$

Using (C.4) we can establish the analog of Step 2 in Theorem 4.1 above, i.e., $\mathrm{P}^{\theta_1}(\cap_{i \in \hat{N}_{K,\hat{d}}} A_i^c \mid A_R^\sigma) \leq \lim_{|\hat{N}_{K,\hat{d}}| \to \infty} (1 - \delta^K)^{|\hat{N}_{K,\hat{d}}|} = 0$. Note that for $\omega \in A_i \cap A_R^\sigma$, as $\mu_{i,1} \geq \hat{d}$, our construction ensures that $G(\mu_{i,t}) \subset G(\delta_{\theta_1})$ for all $t \geq 1$. Thus on the set $A_i \cap A_R^\sigma$, agent i will always choose an action in $G(\delta_{\theta_1})$, so that we can employ the arguments in Step 3 in Theorem 4.1 above to ensure that $X^j(\omega) \subset G(\delta_{\theta_1})$ for all agents $j \in N$ on the set A_R^σ. It follows that $P^{\theta_1}(\cup_{i\in N}\{X^i(\omega) \not\subset G(\delta_{\theta_1})\} \mid A_R^\sigma) = 0$. However, this implies $P^{\theta_1}(\hat{Q}\cap A_R^\sigma) = P(\cup_{i\in N}\{X^i(\omega) \not\subset G(\delta_{\theta_1})\}\cap A_R^\sigma) = 0$, which contradicts our earlier statement that $P^{\theta_1}(\hat{Q}\cap A_R^\sigma) > 0$. □

Proof of Theorem 4.2: We suppose for simplicity that $G(\delta_{\theta_1})$ is a singleton. The arguments presented below extend easily to cover the case where there are multiple optimal actions. We first establish the following lemma:

Lemma 4.2 ▷ *Suppose (Ia)-(Ic) hold. Let $\mu \in \mathcal{D}(\Theta)$ satisfy $\mu(\Theta_{k(x_1)}(x_1)) \geq \xi$. (a) If action x_1 is chosen t times, and outcomes $y_1 \in B^{x_1}$,...,$y_t \in B^{x_1}$ are observed, then the posterior belief $\mu(\Theta(x_1)^+)' \geq \xi$. (b) The conclusion in (a) is unaffected if an action $x \in X\backslash x_1$ has also been chosen and $y \in Y$ is observed.*

The proof exploits condition (Ib) and involves some straightforward calculations. We omit it due to space constraints. Lemma 4.2 is useful since if $\mu \in \mathcal{D}(\Theta)$ satisfies $\mu(\Theta(x_1)^+) \geq \xi$ then $u(x_1,\mu) \geq u_{\min}$.

Proof (Theorem 4.2) Let $j \in N$. Arguments analogous to those used in Lemma 4.1 establish that there exists a real number $\sigma \geq 1$ such that

$$P^{\theta_1}(\sup_{t'>t}\max_{\theta\in\Theta(x_1)^-} r_j^{x_1,\theta}(t,t') \leq \sigma) \;=\; \delta > 0. \qquad (C.5)$$

Choose T to satisfy $\alpha^T\sigma^K < 1$, where $\alpha \in (0,1)$ is the number assumed in condition (Ic). Let A_j^σ be defined as:

$$A_j^\sigma \;=\; \prod_{t=1}^{T} B_{j,t}^{x_1}\times\{\sup_{t'>T}\max_{\theta\in\Theta(x_1)^-} r_j^{x_1,\theta}(T{+}1,t') \leq \sigma\}\times\prod_{x\in X\backslash x_1}\prod_{t=1}^{\infty} Y_{j,t}^{x}\times\prod_{j'\in N\backslash j}\prod_{x\in X}\prod_{t=1}^{\infty} Y_{j',t}^{x}. \quad (C.6)$$

where we have written B^{x_1} as $B_{j,t}^{x_1}$ to avoid confusion. Fix $i \in N_{K,\xi}$. Let $A_i = \cap_{j\in N(i)}A_j^\sigma$. By construction $P^{\theta_1}(A_i) \;=\; \delta^{|N(i)|} \geq \delta^K \;=\; \lambda > 0$.

We claim that if $\omega \in A_i$ then agent i will choose the optimal action x_1, for all time periods after some finite point. The first step is to show that agent i will observe at least T trials of action x_1. We begin by showing it is tried at least once by some agent $j \in N(i)$. The proof is by contradiction. Suppose not. This implies, in particular, agent i observes infinitely many trials of some action $x \in X\backslash x_1$. Since x is suboptimal, the strong law of large numbers will ensure that $\lim_{t\to\infty}\mu_{i,t}(\theta) = 0$ for all states θ where $u(x,\delta_\theta) > u(x,\delta_{\theta_1})$. Choose $\bar{\epsilon} > 0$ such that $u_{\min} - \bar{\epsilon} > \max_{x\in X\backslash x_1} u(x,\delta_{\theta_1})$. The above argument implies that at a finite time t', agent i's expected utility $u(x,\mu_{i,t'}) \leq u_{\min} - \bar{\epsilon}$.

Since x_1 has not been chosen and the choice of other actions does not affect i's beliefs concerning $\Theta_{k(x_1)}(x_1)$, we have $\mu_{i,t'}(\Theta_{k(x_1)}(x_1)) \geq \xi$. By the observation following Lemma 4.2 this implies $u(x_1, \mu_{i,t'}) \geq u_{\min}$, which implies that x_1 would be preferable to x at the time of the next choice of x by agent i. Thus action x_1 must be tried by agent i at some time $\hat{t}$, and this contradicts our original supposition.

We now make the following observation. Suppose that at time t each agent $j \in N(i)$ has chosen action x_1 for $0 \leq t_j \leq T$ periods. Hence upto time t, for each $j \in N(i)$ agent i observes the outcomes $y_{j,1}^{x_1} \in B_{j,1}^{x_1}, \ldots, y_{j,t_j}^{x_1} \in B_{j,t_j}^{x_1}$. It follows from Lemma 4.2 that agent i's posterior belief $\mu_{i,t}(\Theta(x_1)^+) \geq \xi$. Note by Lemma 4.2(b) that the possibility that agents $j \in N(i)$ may have also chosen actions in $X \backslash x_1$ does not alter the conclusion. The same argument used above can be repeated in conjunction with this observation to show that agent i must observe at least T choices of action x_1 by agents $j \in N(i)$.

Let $t(T)$ be the time when agent i has observed a trial of x_1 for the T^{th} time. Let $\hat{\mu}_{i,t'} \in \mathcal{D}(\Theta)$ be agent i's belief after incorporating all information about actions $x \in X \backslash x_1$ upto time $t' \geq t(T)$. We get

$$\mu_{i,t'}(\Theta(x_1)^-) = \frac{\sum_{\theta \in \Theta(x_1)^-} \hat{\mu}_{i,t'}(\theta) \prod_{j \in N(i)} r_j^{x_1,\theta}(1, t_j)}{\hat{\mu}_{i,t'}(\Theta_{k(x_1)}(x_1)) + \sum_{\theta \in \Theta(x_1)^{++}} \hat{\mu}_{i,t'}(\theta) \prod_{j \in N(i)} r_j^{x_1,\theta}(1, t_j) + \sum_{\theta \in \Theta(x_1)^-} \hat{\mu}_{i,t'}(\theta) \prod_{j \in N(i)} r_j^{x_1,\theta}(1, t_j)} \tag{C.7}$$

Since $\omega \in A_i$ by assumption we have $\prod_{j \in N(i)} r_j^{x_1,\theta}(1, t_j) \leq \alpha^T \sigma^K < 1$ for all $\theta \in \Theta(x_1)^-$. This is because, by construction of the set A_i, for the first T observation of x_1 by agent i, the product likelihood ratio $r^{x_1,\theta}$ for any $\theta \in \Theta(x_1)^-$ is at most α^T, and in all subsequent trials for each agent $j \in N(i)$ the product likelihood ratio is at most σ. However, by (Ia) we have $\hat{\mu}_{i,t'}(\Theta_{k(x_1)}(x_1)) = \mu_{i,1}(\Theta_{k(x_1)}(x_1)) \geq \xi$ and $\hat{\mu}_{i,t'}(\Theta(x_1)^-) \leq 1 - \xi$. Thus $\sum_{\theta \in \Theta(x_1)^-} \mu_{i,t'}(\theta) \leq \alpha^T \sigma^K (1 - \xi) < (1 - \xi)$. It follows from (C.7) that $\mu_{i,t'}(\Theta(x_1)^-) < (1 - \xi)/(\xi + (1 - \xi)) = 1 - \xi$. Thus $\mu_{i,t'}(\Theta(x_1)^+) \geq \xi$ and hence $u(x_1, \mu_{i,t'}) \geq u_{\min}$. As t' is arbitrary, this means that agent i's belief on ω will henceforth never fall below $u_{\min}$. As all suboptimal actions will fall below $u_{\min} - \bar{\epsilon}$ in finite time, agent i must choose action x_1 from some finite time onwards. The rest of the proof now proceeds as in steps 2 and 3 of Theorem 4.1. □

Appendix D

Proof of Proposition 5.1: Let $N_1' = \{j \in N_1 | N(i) = N_1\}$ denote the set of agents who are linked only to members of their own sub-group, and let $N_1'' = \{i \in N_1 | N(i) \cap N_2 \neq \phi\}$ denote the set of agents who are also linked to members of the other sub-group. Similarly define $N_2' = \{i \in N_2 | N(i) = N_2\}$ and $N_2'' = \{i \in N_2 | N(i) \cap N_1 \neq \phi\}$. Note that $|N_1''| = |N_2''| = \eta$; also note that $\theta_1 = (H, H)$.

For any agent $i \in N$, let S_n^i (R_m^i) denote the number of successes when agent i chooses action x_1 (x_2) for n (m) time periods. Note that since $\mu_1 > \nu_1$, every $i \in N_1$ chooses x_1 in $t = 1$. Assume that every agent $j \in N_1''$ chooses x_1 in each of the periods $t = 1, \ldots, \bar{n}$, where $\bar{n}$ is some positive integer. A sufficient condition for $i \in N_1'$, to continue to choose only x_1 in every period $t = 1, \ldots, \bar{n}$ then follows from (5.2):

$$p^{2(\sum_{j \in N_1''} S_t^j + \sum_{j \in N_1'} S_t^j) - (k+1)t} < \frac{\mu_1(1-\nu_1)}{\nu_1(1-\mu_1)}. \qquad (D.1)$$

for every $t = 1, \ldots, \bar{n}$. Note that $\mu_1 > \nu_1$ implies that the right hand side of (D.1) is strictly greater than one, and that $0 < p < 1$. Thus a simpler and stronger sufficient condition than (D.1) is that for every $t = 1, \ldots, \bar{n}$

$$p^{2(\sum_{j \in N_1''} S_t^j + \sum_{j \in N_1'} S_t^j) - (k+1)t} < 1 \Rightarrow \sum_{j \in N_1''} S_t^j + \sum_{j \in N_1'} S_t^j \geq \frac{(k+1)t}{2}. \qquad (D.2)$$

assuming as before that agent every $j \in N_1''$ chooses x_1 in each time period up to $\bar{n}$. Since (H,H) is the true state, for each agent i and time t the random variable $2S_t^i - t$ is the sum of independent and identically distributed random variables of the form $2Z - 1$ where Z has a Bernoulli distribution with parameter $1/2 < \pi < 1$. Since $E[Z] = 2\pi - 1 > 0$, it follows from the standard theory of random walks that $\mathrm{P}^{\theta_1}(\{2S_t^i - t \geq 0 \text{ for all } t = 0, 1, \ldots\}) > 0$. Fix $\epsilon > 0$ and for $t \geq 3$ define $\psi(\epsilon, t) = (1+\epsilon)\sqrt{2t\pi(1-\pi)\log\log t}$. For $t \geq 3$, let E_t^i be the event $E_t^i = \{S_n^i \in [n\pi - \psi(\epsilon, n), n\pi + \psi(\epsilon, n)] \text{ for all } n \geq t\}$.

Note that the sets E_t^i are increasing in t. The law of the iterated logarithm (see Billingsley 1986) implies that $\mathrm{P}^{\theta_1}(\bigcup_{t=3}^{\infty} E_t^i) = 1.$. It follows from the previous equations that there is some T such that $\mathrm{P}^{\theta_1}(\{2S_t^i - t \geq 0 \text{ for all } t = 0, 1, \ldots\} \cap E_T^i) > 0$. Let E^i denote the above event, i.e. $E^i = \{2S_t^i - t \geq 0 \text{ for all } t = 0, 1, \ldots\} \cap E_T^i$. The set E^i is such that agent i's observations will satisfy $S_t^i \geq t/2$ for every t and also from T onwards, S_t^i will be within the bounds prescribed by the law of the iterated logarithm. Note that given that the previous equation holds for T it also holds from $T' > T$ since the sets $\{E_t^i\}$ are increasing. Thus we shall assume, without loss of generality, that T also satisfies the condition

$$T \geq \inf \Big\{ t \mid ((k-\eta+1)(2\pi-1)-\eta)\frac{t}{2} - (k-\eta+1)\psi(\epsilon,t) > 0 \Big\}. \qquad (D.3)$$

Since $k-\eta+1 > \eta/(2\pi-1)$ by assumption, the term $(k-\eta+1)(2\pi-1)-\eta$ is positive. As the term $(k-\eta+1)\psi(\epsilon,t)$ in (D.3) is of the order of $t^{1/2}\log\log t$ it is eventually dominated by $((k-\eta+1)(2\pi-1)-\eta)t/2$. Hence, T in (D.3) will be finite. In addition, for all $t \geq T$ the expression $((k-\eta+1)(2\pi-1)-\eta)t/2 - (k-\eta+1)\psi(\epsilon,t) > 0$ as well.

We now consider the situation of agent $i \in N_1''$ given that the rest of N_1 choose only x_1 up to time $\bar{n}$. We shall assume that agent $j \in N_2''$ choose only x_2 in each of the $\bar{n}$ periods, and later isolate a set of sample paths where this will in fact be true. Recall that R_t^j denotes the number of successes that agent β obtains in t trials from using x_2. Using (5.2) and the arguments underlying (D.1) and (D.2) again, a sufficient condition for agent $i \in N_1''$ to choose x_1 for each $t = 1,\ldots,\bar{n}$ is

$$\sum_{j\in N_1''} S_t^j + \sum_{j\in N_1'} S_t^j \geq \sum_{j\in N_2''} R_t^j + \frac{(k-\eta+1)t}{2}. \qquad (D.4)$$

We now choose the set of sample paths for agent $j \in N_1''$ which satisfy the following requirements:

$$\bar{E}^j = \{S_T^j = T,\ S_t^j \geq T + (t-T)/2,\ \text{for all } t > T\}. \qquad (D.5)$$

Thus, for sample paths in $\bar{E}^j$, agent $j \in N_1''$ gets T successes in the first T trials with x_1 and subsequently gets a success rate of at least 50 percent in the remaining periods. Clearly $\mathrm{P}^{\theta_1}(\bar{E}^j) > 0$.

Next, we consider an agent $j \in N_2$. By definition of the set N_2, it follows that all $j \in N_2$ choose action x_2 in $t = 1$. Assuming that $j \in N_2''$ chooses x_2 in all periods up to $\bar{n}$, using arguments as above, it follows that a sufficient condition for each agent $i \in N_2'$ to choose x_2 for each $t = 1,\ldots,\bar{n}$ is :

$$\sum_{j\in N_2''} R_t^j + \sum_{j\in N_2'} R_t^j \geq \frac{(k+1)t}{2}. \qquad (D.6)$$

For agents $i \in N_2'$ consider the sample paths $\{F_n^i\}$ defined as $F_t^i = \{R_n^i \in [n\pi - \psi(\epsilon,n), n\pi + \psi(\epsilon,n)]$ for all $n \geq t\}$. By symmetry with the situation of N_1, it follows that the event F^i defined as $F^i = \{2R_t^j - t \geq 0 \text{ for all } t \geq 0\} \cap F_T^i$, has strictly positive probability. Lastly, we define the set $\bar{F}^j$ corresponding to $\bar{E}^j$ where (D.5) will be true for $j \in N_2^{''}$ in place of $j \in N_1^{''}$ and x_2 instead of x_1. We now consider the set of sample paths

$$\prod_{j \in N_1'} E^j \times \prod_{j \in N_1^{''}} \bar{E}^j \times \prod_{j \in N_2^{''}} \bar{F}^j \times \prod_{j \in N_2'} F^j. \qquad (D.7)$$

By construction, the above event has strictly positive probability. We claim that on this set of sample paths every $i \in N_1$ chooses x_1 and every $j \in N_2$ chooses x_2 forever. We first consider the situation for $t = 1, \ldots, T$. For any $i \in N_1'$, condition (D.2) is clearly satisfied, since on our choice of sample paths, $\sum_{j \in N_1'} S_t^j \geq (k - \eta + 1)t/2$ and $\sum_{j \in N_1^{''}} S_t^j = \eta t \geq t/2$ so that their sum is at least $(k+1)t/2$. Next consider an agent $i \in N_1^{''}$. Since $\sum_{j \in N_1^{''}} S_t^j + \sum_{j \in N_1'} S_t^j \geq \eta t + (k - \eta + 1)t/2$ and the right hand side of (D.4) equals this value (recall that $R_T^j = T$ for $j \in N_2^{''}$ on this set of sample paths) equation (D.4) also holds, ensuring that $j \in N_1^{''}$ will choose x_1 upto time T. By symmetry, we can show that any agent $j \in N_2$ chooses x_2 upto time T.

We now consider the position after time T. In this case, for each $i \in N_1'$, we have S_t^i in the interval $t\pi \pm \psi(\epsilon, t)$. Furthermore, for $j \in N_1^{''}$, S_t^j is at least equal to $T + (t-T)/2$ while R_t^j is at most equal to t. From (D.4) we have therefore

$$\begin{aligned} \sum_{j \in N_1^{''}} S_t^j + \sum_{j \in N_1'} S_t^j - \Big\{ \sum_{j \in N_2^{''}} R_t^j + \frac{(k-\eta+1)t}{2} \Big\} \geq \\ \geq \eta \Big\{ T + \frac{(t-T)}{2} \Big\} + (k - \eta + 1)\Big\{ t\pi - \psi(\epsilon, t) \Big\} - \Big\{ \eta t + \frac{(k-\eta+1)t}{2} \Big\} \\ = \frac{\eta T}{2} + \Big\{ (k-\eta+1)(2\pi - 1) - \eta \Big\} \frac{t}{2} - (k - \eta + 1)\psi(\epsilon, t). \qquad \text{(D.8)} \end{aligned}$$

By our choice of T in (D.3) the last expression above is always non-negative. It follows that provided that every $i \in N_1'$ chooses x_1, and $j \in N_2^{''}$ always chooses x_2, any agent $j \in N_1^{''}$ will always choose x_1. We next show agents in N_1' will continue to choose x_1 after time T given that an agent $j \in N_1^{''}$ chooses x_1 forever. This requires condition (D.2) to hold. However,

$$\sum_{j \in N_1^{''}} S_t^j + \sum_{j \in N_1'} S_t^j \geq \eta \Big\{ T + \frac{t-T}{2} \Big\} + \frac{(k-\eta+1)t}{2} \geq \frac{(k+1)t}{2}. \qquad (D.9)$$

so (D.2) continues to be satisfied. By symmetry of the agents N_2 to the agents N_1, analogous arguments establish that every agent in N_2 will always choose x_2.

□

References

Aghion, P., P. Bolton, C. Harris, and B. Jullien (1992), "Optimal Learning by Experimentation", *Review of Economic Studies*, 58, 621-654.

Aghion, P., Paz-Espinosa, M., and Jullien, B. (1993), "Dynamic Duopoly with Learning through Market Experimentation", *Economic Theory*, 3, p.517-540.

Allen, B. (1982a), "Some Stochastic Processes of Interdependent demand and technological diffusion of an innovation exhibiting externalities among adopters", *International Economic Review*, 23, p.595-607.

Allen, B. (1982b), "A Stochastic Interactive Model for the Diffusion of Information", *Journal of Mathematical Sociology*, 8, p.265-281.

An, M., and Kiefer, N. (1992), "Local Externalities and Societal Adoption of Technologies", CAE Working Paper 92-06, Cornell University, Ithaca.

Arthur, B (1989), "Competing Technologies, Increasing Returns, and Lock-in by Historical Events", *Economic Journal*, 99, p.116-131.

Bala, V., and Goyal, S. (1995), "A Theory of Learning with Heterogeneous Agents", *International Economic Review*, 36, 2, 303-323.

Bala, V. and Goyal, S. (1994), "The Birth of a New Market", *Economic Journal*, 104, 282-290.

Banerjee, A. (1992), "A Simple Model of Herd Behavior", *Quarterly Journal of Economics*, CVII, 3, 797-817.

Berry, D. A., and Fristedt, B. (1985), *Bandit Problems: Sequential Allocation of Experiments*, Chapman and Hall, New York.

Blume, L. and Easley, D. (1992), "Rational Expectations and Rational Learning", mimeo., Cornell University.

Bikhchandani, S., Hirshleifer, D., Welch, I. (1992), " A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades", *Journal of Political Economy*, 100, 5, 992-1023.

Billingsley, P. (1985), *Probability and Measure*, Wiley, New York.

Bolton, P., and Harris, C. (1992), " Strategic Experimentation", mimeo, Oxford University.

Chung, K.L. (1974), *A Course in Probability Theory*, 2nd Edition, Academic Press, Orlando, Florida.

Coleman, J. (1966), *Medical Innovation: A Diffusion Study*, Bobbs-Merrill, New York.

DeGroot, M. (1970), *Optimal Statistical Decisions*, McGraw Hill, New York.

Easley, D. and Kiefer, N. (1988), "Controlling a Stochastic Process with Unknown Parameters", *Econometrica*, 56, p.1045-1064.

Ellison, G. and Fudenberg, D. (1993), "Rules of Thumb for Social Learning", *Journal of Political Economy*, 101, 612-643.

Ellison, G. and Fudenberg, D. (1995), "Word of Mouth Communication and Social Learning", *Quarterly Journal of Economics*, 110, 1, 93-126.

Feder, G., Just, R.E. and Zilberman, D. (1985), "Adoption of Agricultural Innovations in Developing Countries: A Survey", *Economic Development and Cultural Change*, p.255-298.

Griliches, Z.(1957), "Hybrid Corn: An Exploration in the economics of technical change", *Econometrica*, 25, p.501-522.

Hagerstrand, T. (1969), "Innovation Diffusion as a Spatial Process", University of Chicago Press.

Kotler, P. (1986), *The Principles of Marketing*, 3rd Edition, Prentice-Hall, New York.

Rogers, E. M. (1983), *Diffusion of Innovations*, Third Edition, Free Press, New York.

Ryan, B., and Gross, N. (1943), "The Diffusion of Hybrid Seed Corn in Two Iowa Communities", *Rural Sociology*, 8, 15-24.

Vives, X. (1992), " Learning from Others", Mimeo, CSIS, Barcelona.

Watkins, S. C. (1991), *Provinces into Nations: Demographic Integration in Western Europe, 1870-1960*, Princeton University Press, New Jersey.