Selective Observation in Social Learning

Nathan Larson
Centre for Economic Learning and Social Evolution
University College London
Gower Street, London WC1E 6BT
Email: n.larson@ucl.ac.uk

July 31, 2003

Abstract

Social learning is studied when heterogeneous agents can select whom to observe. An agent learns from the average action in her neighborhood. She may face a tradeoff between joining a large neighborhood where most agents have different preferences than she does or a small neighborhood where agents share her preferences. Typically, her choice will not be socially optimal, and the equilibrium outcome may involve inefficient pooling.

1 Introduction

The study of social learning formalizes the idea that people facing decisions under imperfect information may seek guidance from the way that other people they know have acted (and fared) in similar situations. Implicit elements of this story are a social structure indicating whom an agent observes (and who observes her) and an inference process describing how an agent incorporates her observations in choosing her optimal action. Most work on social learning has focused on the inference process, taking the social structure to be exogenous. However, when agents differ in their preferences, the value of social learning to an agent depends critically on choosing the right set of agents to observe. This paper looks at a model in which agents with diverse preferences must first choose a neighborhood and then choose an action based on their observation of how other agents in that neighborhood have acted (plus some private information). In choosing a neighborhood, agents will often face a new tradeoff between large samples with diffuse preferences and smaller, targeted ones. When large samples prevail, there will be pooling on a neighborhood, creating an obstacle to complete learning that is reminiscent of, but distinct from, herding on actions.

To introduce the main ideas, consider the decision problem of an agent who wants to go to the movies. There are only two films, X and Y, both of which are showing at the two theaters in town, the arthouse (A) theater and the megaplex (M). Filmgoer preferences are of two types, arty and mainstream. Both types generally agree on the elements that make a good movie but differ on their relative importance (good dialogue vs. car chases, for example). Each filmgoer first chooses which theater to drive to and then decides which movie to see on the basis of her (noisy) private signal and what people ahead of her in line seem to prefer. (But she doesn't mind waiting in line *per se*.) How should a filmgoer decide where to go?

If there were only one theater and only a single type of agent, then standard models of sequential social learning would fit this story quite well. However, when agents have diverse preferences and different choices about where to gather information, the picture becomes more complex. Once she has arrived at a theater, our filmgoer must assess how much information can be gleaned from the line's behavior. She must consider how much her predecessors have relied on their signals, as in the standard case, but also how similar their preferences are likely to be to hers. One step back, her belief about which theater is likely to be more informative for her will generally depend on her type and her beliefs about the composition of the two lines.

For example, suppose that it is well known that arty filmgoers are relatively rare and tend to frequent the arthouse theater, while mainstream filmgoers go to the megaplex. In deciding which theater to go to, an arty agent must trade off noise due to a small observation sample (at A) versus noise due differing preferences (at M). In general, this tradeoff can go in either direction; let us imagine that it favors M. Now consider the choice of the next arty agent. The line at A is no more informative than it would have been for her predecessor, while the line at M has become strictly more informative – so she will drive to M as will all subsequent arty agents. Thus, pooling on an information source is a possibility.

As is usual with social learning, agents do not account for the value that later agents derive from the private information embodied in their film choices. However, the sign of this externality is no longer straightforward; mainstream agents may wish their arty predecessors had used their private information less rather than more, for example. There is also a second externality that is new: the choice of an information source affects the information available to later filmgoers, but agents do not take this into account when deciding which theater to drive to. Pooling on M is an example of this – although pooling may reflect each agent's short run interest, in the long run it can lead to incomplete learning if the preferences of arty and mainstream filmgoers cannot be disentangled. In contrast, if the filmgoers were to segregate at different theaters, complete learning might obtain.[1]

The rest of this paper develops a simple model in which these issues can be explored. Like the example above, the model has two types of agents and two "information neighborhoods"; however, the action space will be continuous, not discrete. The neighborhoods could be thought of as physical locations, such as theaters or retail stores, or non-geographic communities like professional organizations or online chat sites. What is important is that neighborhoods are mutually exclusive – an agent can join only one – and that information flows within neighborhoods, but not between them. Agents join a neighborhood, observe the average action taken there, and then choose their own actions, which are then incorporated into the average observed by subsequent agents. Section 2 introduces this model and develops a recursive characterization of its social learning equilibria (Lemmas 1, 2, and 3). Section 3 provides a deeper exploration of equilibrium behavior. We show that if pooling starts, then it persists forever (Lemma 4) and provide conditions under which pooling either occurs in the first period or not at all (Proposition 1). Next we turn to long run learning. If the agents separate by type, each type eventually converges to its optimal action, but because of the underutilization of private information, this convergence is only at rate $\ln t$ (Proposition 2). On the other hand, when the agents pool, their actions eventually converge, but this limiting action is suboptimal for

[1] With the discrete actions in this example, we would also have to worry the possibility that complete learning fails due to herding. This will not be a concern in the model developed later, as the action spaces are continuous.

both types. Although the limiting action favors the preferences of the more populous type in absolute terms, it overweights the preferences of the minority type relative to their share of the population (Proposition 3). Finally, we consider welfare-improving interventions in a pooling equilibrium. A social planner can always improve time average utility by reducing noise in either public or private signals (Propositions 4 and 5). Furthermore, a policy of enforcing segregation for a limited time can suffice to divert the population into an separating equilibrium with complete learning (Proposition 6). However, complete learning can be a deceptive benchmark for efficiency because the rate of learning is quite slow. We present several examples demonstrating that for relatively modest social discount factors, pooling may actually lead to higher welfare than separation. Before turning to the model, this section concludes with a brief discussion of where our model fits in broader literature on social learning.

In its basic template of agents acting sequentially after observing a private signal and some of the actions of agents before them, our model builds on the seminal contributions of Banerjee [2] and Bikhchandani, Hirshleifer, and Welch (BHW) [3]. These papers developed the insight that with discrete actions and identical preferences it is possible for all agents past a certain point to rationally disregard their private signals and follow the lead of the agents before them. Much subsequent work has focused on clarifying and expanding the conditions under which these information cascades can occur. Along these lines, Smith and Sorensen [6] embellish BHW-style models with heterogeneous agents, while retaining an exogenous information structure. They show that contrarian types can tend to break up herds and also how pooling by different types can lead to incomplete learning, an effect they call confounded learning. In their model, confounded learning arises because agents have no way to restrict their attention to predessors with similar preferences to theirs. Our work shows that even when agents can restrict their attention in this way, they may rationally choose not to, so incomplete learning can be a persistent phenomenon. In assuming that different agents may have access to different observation samples, our work has parallels with many non-sequential models of social learning, including Ellison and Fudenberg [4], [5], and Bala and Goyal [1]; however, all of these papers assume observation samples to be exogenous. Finally, we follow Vives [7] (and depart from the work above) in focusing on normal learning with a continuous action space. The continuous action space is interesting in its own right and has the added benefit of always generating complete learning when agents are homogeneous, permitting the effects of pooling with multiple types to stand out in sharper relief.

2 The Model

There is an infinite sequence of cohorts of agents, indexed by the natural numbers. Each cohort consists of r Red agents and g Green agents, with $g < r$. An agent's type is private. The preferences of an agent of type $T \in \{R, G\}$ are given by

$$u_T(x) = -(\lambda_T - x)^2 \tag{1}$$

where $\lambda_T = \theta + \mu_T$ is composed of a term θ that is common to all agents and a type-specific parameter μ_T. Prior beliefs about θ are diffuse, while priors on μ_R, and μ_G are assumed to follow independent normal distributions with mean zero and variance σ_μ^2. Each agent i receives a (conditionally independent) signal s_i about its preference $\lambda_{T(i)}$ which is also assumed to be normally distributed: $s_i \sim N(\lambda_{T(i)}, \sigma_s^2)$. Occasionally, we will need to refer to i's private signal error $\varepsilon_i = s_i - \lambda_{T(i)} \sim N(0, \sigma_s^2)$.

Structure of the Game

The sequence of actions is as follows. Agents live for a single period. In period t, each agent i in cohort t must choose to join one of two neighborhoods, 1 or 2. At this stage, the only information available to her is the history of the size of each neighborhood in periods 1 through $t-1$, h_1^{t-1} and h_2^{t-1}. After joining a neighborhood n, the agent observes the average action $\bar{x}_n^{t-1}$ taken in her neighborhood in the past (that is, over all previous cohorts, 0 through $t-1$) as well as her private signal s_i. She then chooses an action x_i to try to maximize her payoff as given by (1). We make one final assumption: there is a seed cohort (which could be thought of as cohort 0) containing r_0 Red agents, all assigned to neighborhood 1, and g_0 Green agents, all assigned to neighborhood 2, and all of this is common knowledge. These agents have no history to observe; they must rely on their private signals alone.

Equilibrium

The strategy of an agent consists of a neighborhood choice function and an action function. The neighborhood choice function $n_T^t(h_1^{t-1}, h_2^{t-1}) \in \{1, 2\}$ specifies which neighborhood to join as a function of the history of each neighborhood in period $t-1$ and the agent's type T. The action function $x_T^t(s, \bar{x}, n, h_1^{t-1}, h_2^{t-1})$ specifies the agent's action as a function of her private signal, the average action she sees, the identity of the neighborhood she has joined, the neighborhood histories, and her type. In the sequel, we will suppress arguments whenever doing so should not cause confusion.

Definition 1 *A Social Learning Equilibrium (SLE) is a collection of strategies and beliefs for each agent such that her beliefs are obtained via Bayes' rule whenever possible and her strategy maximizes the expectation of (1) given her beliefs.*

Beliefs and Social Learning

In order to begin to characterize equilibria of the model, we will present a few lemmas capturing some observations about optimal behavior and beliefs.

Lemma 1 *Suppose that a Red agent believes that the average action in her neighborhood to be given by*

$$\bar{x} = \rho_R \lambda_R + \rho_G \lambda_G + \upsilon$$

where $\rho_R + \rho_G = 1$, $\upsilon \sim N(0, \sigma_\upsilon^2)$, and υ is uncorrelated the other stochastic elements in the model. Then her optimal action x is given by

$$x = \alpha s + (1 - \alpha)\bar{x}$$

where

$$\alpha = \frac{\tau_s}{\tau_s + \tau_{\bar{x}}}$$

and

$$\begin{aligned} \tau_s &= 1/\sigma_s^2 \\ \tau_{\bar{x}} &= (2\rho_G^2\sigma_\mu^2 + \sigma_\upsilon^2)^{-1} \end{aligned}$$

If she is Green, we have the same result, replacing $\tau_{\bar{x}}$ with

$$\tau_{\bar{x}} = (2\rho_R^2\sigma_\mu^2 + \sigma_\upsilon^2)^{-1}$$

Proof. Appendix. ■

To paraphrase, if an agent believes the average action she observes to be a weighted average of her type's target and the other type's target, plus noise, then her optimal action is a weighted average of her private signal and this average action. Note that the weight that an agent places on the average action is higher when the weight it places on her target is higher. Now consider the agent's neighborhood choice.

Lemma 2 *Under the conditions of Lemma 1, the agent's expected payoff is $-(\tau_s + \tau_{\bar{x}})^{-1} = -\alpha\sigma_s^2$. Thus, if the agent's beliefs about each neighborhood take the form in Lemma 1, she should choose the neighborhood with the higher value of $\tau_{\bar{x}}$.*

Proof. Omitted. ■

Finally, let us demonstrate that the hypothesized form of beliefs is justified. Suppose the agents in cohort t share beliefs about the two neighborhoods that take the form described in Lemma 1. Let us write υ_n^{t-1} and ρ_{nT}^{t-1} for the beliefs about neighborhood n, and τ_{nT}^{t-1} for the precision of the average action in n as a signal about λ_T. If $\tau_{1T}^{t-1} > \tau_{2T}^{t-1}$ ($\tau_{1T}^{t-1} < \tau_{2T}^{t-1}$) for both types, then all of the agents in cohort t will choose neighborhood 1 (2). In this case, we will say that the agents *pool* on neighborhood 1 (2). Otherwise, we will say that the agents *separate*. Now consider the beliefs of cohort $t+1$. By assumption, they observe the same neighborhood histories as cohort t agents, so their beliefs conditional on h_1^{t-1} and h_2^{t-1} are the same. Furthermore, they can anticipate the neighborhood choices that the cohort t agents will make and the weights that they will give to their private signals in making action choices. Thus, cohort $t+1$ agents will believe $\bar{x}_n^t$ to be a weighted average of $\bar{x}_n^{t-1}$ and the private signals of cohort t agents. It is not difficult to see that this weighted average will preserve the form described in Lemma 1. The details of this mapping from cohort t to cohort $t+1$ beliefs are laid out below. In anticipation of the equilibrium outcomes, only certain cases are considered. Note one more bit of notation: we will write D_n^t for the total number of agents who have chosen neighborhood n in all cohorts up to and including cohort t.

Lemma 3 *(Updating of beliefs)*

Let r_n^t and g_n^t be the number of Red and Green agents who in cohort t who choose neighborhood n, following Lemma 2. SLE beliefs take the form specified by Lemma 1 and can be *characterized by the following recursive relationship for all* $t \geq 1$:

$$\begin{aligned}
\rho_{nR}^{t} &= \frac{K_n^t \rho_{nR}^{t-1} + r_n^t \alpha_{nR}}{D_n^t} \\
\rho_{nG}^{t} &= \frac{K_n^t \rho_{nG}^{t-1} + g_n^t \alpha_{nG}}{D_n^t} \\
\sigma^2_{\upsilon_n^t} &= (K_n^{t^2} \sigma^2_{\upsilon_n^{t-1}} + (r_n^t \alpha_{nR}^{t^2} + g_n^t \alpha_{nG}^{t^2})\sigma_s^2)/D_n^{t^2}
\end{aligned}$$

where

$$\begin{aligned}
K_n^t &= D_n^t - r_n^t \alpha_{nR}^t - g_n^t \alpha_{nG}^t \\
D_n^t &= D_n^{t-1} + r_n^t + g_n^t
\end{aligned}$$

and α_{nR}^t *and* α_{nG}^t *are the weights used by cohort* t *agents according to Lemma 1, and*

$$\begin{aligned}
\rho_{1R}^0 &= \rho_{2G}^0 = 1 \\
\sigma^2_{\upsilon_1^0} &= \sigma_s^2 / r_0 \\
\sigma^2_{\upsilon_2^0} &= \sigma_s^2 / g_0
\end{aligned}$$

Proof. (Appendix) ■

These three lemmas fully characterize a social learning equilibrium of the model. Although we will not prove it, the fact that this equilibrium is generically unique should be clear. The only possible source of multiplicity would be indifference over which neighborhood to join, and with our assumption of initial separation, this indifference can only occur for zero-measure subsets of parameters.

The critical tradeoff in the model is illustrated by the expression for the precision of a neighborhood's average action as a signal about λ_T, $\tau_{nT} = (2(1-\rho_{nT})^2\sigma_\mu^2 + \sigma_{\upsilon_n}^2)^{-1}$. As υ_n is a weighted average of past private signal errors in neighborhood n, its variance will tend to be lower in in the neighborhood that has been more popular in the past. This reflects the benefit of learning in the neighborhood that offers a larger sample. However, there is also noise due to the fact that the sample that an agent observes may represent the actions of many agents who do not share her preferences. The smaller is ρ_{nT}, the less the information of type T agents is reflected in neighborhood n's average action.

3 Results

The model permits us to address some natural questions about the long run outcomes of social learning. We start by examining the circumstances under which agents separate or pool over neighborhoods in the long run.

As an opening shot, we can show that once a neighborhood is abandoned by one cohort, later cohorts never return to it. The basic logic here is quite simple: the signal provided by an abandoned neighborhood does not change from period to period, while the other neighborhood's signal tends to improve as its sample grows. The story is a bit more complicated than this, as pooling may temporarily reduce the value of the signal for one of the types at the unabandoned neighborhood. However, in equilibrium, this effect never makes it worthwhile to switch to the neighborhood that was abandoned. Let us say that the seed cohort is *consistent* with the later cohorts if it preserves the same ratio of types: $r_0 = kr$ and $g_0 = kg$ for some $k > 0$. We say that the seed cohort is *consistent and full strength* if $k \geq 1$, that is, if it is as least as populous as later cohorts. Then we have the following.

Lemma 4 *Suppose the seed cohort is consistent. In a SLE, suppose that no agent chooses some neighborhood for a cohort t. Then no agent chooses that neighborhood for any $t' > t$.*

Proof. Let t be the first period in which the agents pool. Then in cohorts 1 through $t-1$, Red (Green) agents choose neighborhood 1 (2).[2] The first step is to show that the cumulative private signal error is lower in every period for the neighborhood with the larger sample: $\sigma^2_{\upsilon_1^z} < \sigma^2_{\upsilon_2^z}$ for all $z < t$. This straightforward exercise is left to Lemma 5 in the appendix. Now observe that

$$\begin{aligned}
\tau^{t-1}_{1R} &= 1/\sigma^2_{\upsilon_1^{t-1}} \\
\tau^{t-1}_{1G} &= (2\sigma^2_\mu + \sigma^2_{\upsilon_1^{t-1}})^{-1} \\
\tau^{t-1}_{2R} &= (2\sigma^2_\mu + \sigma^2_{\upsilon_2^{t-1}})^{-1} \\
\tau^{t-1}_{2G} &= 1/\sigma^2_{\upsilon_2^{t-1}}
\end{aligned}$$

Period t pooling cannot be on neighborhood 2, as $\sigma^2_{\upsilon_1^{t-1}} < \sigma^2_{\upsilon_2^{t-1}}$ and $\sigma^2_\mu > 0$ imply $\tau^{t-1}_{1R} > \tau^{t-1}_{2R}$. Suppose then that pooling is on 1, so we must have $2\sigma^2_\mu + \sigma^2_{\upsilon_1^{t-1}} < \sigma^2_{\upsilon_2^{t-1}}$. Suppose toward a contradiction that $t' > t$ is the first period in which some agent returns to 2. Suppose this agent is Green. Then, because there has been pooling on 1 in the meantime, $1 > \rho^{t'-1}_{1G} > 0$. Furthermore, because $\sigma^2_{\upsilon_1^z}$ is monotonically decreasing, we have $\sigma^2_{\upsilon_1^{t'=1}} \leq \sigma^2_{\upsilon_1^{t-1}}$. Thus, $\tau^{t'-1}_{1G} = (2(1-\rho^{t'-1}_{1G})^2\sigma^2_\mu + \sigma^2_{\upsilon_1^{t'=1}})^{-1} > \tau^{t-1}_{1G}$. On the other hand, with no agents settling in 2, we have $\tau^{t'-1}_{1G} = \tau^{t-1}_{2G}$. Thus, if Greens preferred 1 in period t, they will still prefer 1 in period t', so it cannot be a Green agent who returns to 2. Suppose instead that it is a Red agent. Then we have $\tau^{t'-1}_{2R} = \tau^{t-1}_{2R} = (2\sigma^2_\mu + \sigma^2_{\upsilon_2^{t-1}})^{-1}$. On the other hand, $\tau^{t'-1}_{1R} = (2(\rho^{t'-1}_{1G})^2\sigma^2_\mu + \sigma^2_{\upsilon_1^{t'=1}})^{-1}$. Because $\rho^{t'-1}_{1G} < 1$ and $\sigma^2_{\upsilon_1^{t'=1}} \leq \sigma^2_{\upsilon_1^{t-1}}$, $\tau^{t'-1}_{1R} > \tau^{t'-1}_{2R}$ so Reds cannot switch to neighborhood 2 either. Consequently, if pooling on neighborhood 1 starts, no agent will ever return to neighborhood 2. ■

Lemma 4 means that equilibrium outcomes are easy to characterize. Either both neighborhoods are always occupied and separated by type, or at some point pooling on the "majority" neighborhood begins. Requiring the seed cohorts to be consistent simplifies the proof, but a similar result could be obtained without this assumption. However, consistency is necessary for the following result showing that pooling occurs immediately, if at all.

Proposition 1 *If the seed cohort is consistent and full strength, then the SLE entails immediate*

[2] We have ignored the possibility that there is a "flip-flop" in some period $t' < t$, with cohort t' Reds choosing 2 and Greens choosing 1. This cannot happen; if Greens prefer 1 (for example), even though its history at t' contains only Reds, then Reds must strictly prefer 1 as well.

pooling (i.e., in cohort 1) on neighborhood 1 if

$$\frac{\sigma_\mu^2}{\sigma_s^2} < \frac{1}{2k}(\frac{1}{g} - \frac{1}{r})$$

and separation otherwise.

Proof. Appealing to the proof of Lemma 4, pooling occurs in the first period for which $\sigma^2_{\upsilon_2^{t-1}} - \sigma^2_{\upsilon_1^{t-1}} > 2\sigma_\mu^2$ (as this is the condition for Green agents to choose 1). Lemma 6, in the appendix, shows that if the seed cohort is consistent and full strength, then $\sigma^2_{\upsilon_2^{t-1}} - \sigma^2_{\upsilon_1^{t-1}}$ is decreasing in t; this implies that any pooling must be immediate. The condition for pooling follows from earlier expressions for $\sigma^2_{\upsilon_1^0}$ and $\sigma^2_{\upsilon_2^0}$. ■

The comparative statics of the pooling condition are relatively straightforward and intuitive. If preference differences between the two types are substantial (large σ_μ^2), then it will be optimal to learn from one's own type. On the other hand, if idiosyncratic error matters (large σ_s^2) and the majority smooths out significantly more idiosyncratic variation ($r >> g$), then even minority types will find it optimal to learn in the minority neighborhood. The intuition behind the consistency condition is also straightforward. If neighborhood 1 is smaller than neighborhood 2 at date 0 but larger in all subsequent cohorts, the public signal in neighborhood 1 may leapfrog the signal in neighborhood 2 at some $t > 1$, in which case there may be delayed pooling. Consistency ensures that any advantage to being in neighborhood 1 is present from the start. The necessity of the full strength condition is less clear; it is helpful in proving Lemma 6 but perhaps could be relaxed.

Next we look at whether learning will be complete in the long run and sketch some implications for welfare. By complete learning, we really mean complete type-specific learning; that is, the estimate of a type T agent should converge in distribution to a point mass at λ_T. It should come as no surprise that in a separating equilibrium, complete learning prevails. Informally, as long as the average action in a neighborhood is only finitely precise, agents will continue to incorporate their private signals with a strictly positive weight, driving convergence to the true value of λ_T. However, it is also clear that complete learning can never arise in a pooling equilibrium, as the average action in the active neighborhood can never reveal both λ_R and λ_G with certainty. In what follows, we formalize these points and characterize the limiting information revealed in a pooling equilibrium.

Proposition 2 *In a separating SLE, the average action in neighborhood 1 (2) converges to λ_R (λ_G) at rate $\ln t$. That is, complete learning obtains.*

As noted above, it is not surprising that complete learning is obtained. The fact over-reliance on public information results in a convergence rate for σ_υ^2 that is slower than the standard rate of t is also to be expected following the work of Vives. However, while Vives predicts convergence at rate $t^{1/3}$, in our model convergence is at the substantially slower rate of $\ln t$. The difference can be attributed to the fact that in Vives' model, agents observe a noisy signal of the action of the most recent cohort, while here the signal incorporates the actions of all earlier agents.

Proposition 3 *In a pooling SLE, the average action converges to* $\bar{x}^{\infty} = \rho_R^{\infty}\lambda_R + \rho_G^{\infty}\lambda_G$*, where* ρ_R^{∞} *and* ρ_G^{∞} *are defined by*

$$\begin{aligned} \frac{g}{r} &= \left(\frac{\rho_G^{\infty}}{\rho_R^{\infty}}\right)^3 \frac{\sigma_s^2 + 2\rho_R^{\infty^2}\sigma_\mu^2}{\sigma_s^2 + 2\rho_G^{\infty^2}\sigma_\mu^2} \\ &= \frac{\rho_G^{\infty}}{\rho_R^{\infty}}\frac{\alpha_R^{\infty}}{\alpha_G^{\infty}} \end{aligned}$$

Proof. The proof proceeds through a sequence of steps. First we show that under pooling, the weights that agents place on their own signals cannot all go to zero. This means that private signal error vanishes from $\bar{x}$ as $t \to \infty$. Third, as long as own-signal weights do not explode, ρ_R^t and ρ_G^t must converge. This in turn implies the convergence of α_R^t and α_G^t. Finally, given the laws of motion for ρ_R^t, ρ_G^t, α_R^t, and α_G^t, when they converge, they must satisfy the expression above.

Using Lemma 1 and the fact that $\rho_R^t + \rho_G^t = 1$, one can show that $\max\{\alpha_R^t, \alpha_G^t\} \geq \sigma_\mu^2/(\sigma_\mu^2 + 2\sigma_s^2)$. This means that $\sigma_{\upsilon^t}^2 \to 0$. The proof, which is similar to one for the separating case, is omitted.

Next, observe that

$$\begin{aligned} \rho_R^t - \rho_R^{t-1} &= (K^t/D^t - 1)\rho_R^{t-1} + r\alpha_R^t/D^t \\ &= (r\alpha_R^t - \rho_R^{t-1}(r\alpha_R^t - g\alpha_G^t))/D^t \end{aligned}$$

The terms in the numerator are uniformly bounded, and the total number of agents grows without bound, $D^t \to \infty$, so $\rho_R^t - \rho_R^{t-1} \to 0$. Then, by completeness of the unit interval, ρ_R^t converges to some ρ_R^{∞} (and similarly for ρ_G^t).

Again using Lemma 1, the convergence of ρ_R^t, ρ_G^t, and $\sigma_{\upsilon^t}^2$ implies

$$\alpha_T^t \to \alpha_T^{\infty} = \frac{2(1-\rho_T^{\infty})^2\sigma_\mu^2}{2(1-\rho_T^{\infty})^2\sigma_\mu^2 + \sigma_s^2}$$

Now return to the expression for $\rho_R^t - \rho_R^{t-1}$ above. We have shown that the numerator converges to some constant m_R. Then, because D^t grows at rate t, for t large enough, successive increments are close to proportional to m_R/t. If $m_R \neq 0$, the sum of these terms diverges, contradicting the convergence of ρ_R^t, so we must have $m_R = 0$. We proceed similarly for $\rho_G^t - \rho_G^{t-1}$ and substitute for α_R^{∞} and α_G^{∞} to get the result in the proposition. ■

Some special cases may help to illustrate this result. First, suppose that the difference in preferences between types is large relative to the error in individual signals (σ_μ^2/σ_s^2 large). Then $\rho_R^{\infty} \to r/(r+g)$ and $\rho_R^{\infty} \to g/(r+g)$; that is, the weight of each type's target action in the average action of the group converges to the type's population share. This results from the fact that agents rely entirely on their own signals, even in the limit as t grows large.[3] At the other extreme, when individual errors are large relative to the difference in preferences, $\rho_R^{\infty} \to r^{1/3}/(r^{1/3} + g^{1/3})$ and $\rho_R^{\infty} \to g^{1/3}/(r^{1/3} + g^{1/3})$. In other words, the limiting average action places more weight on the preferences of the minority type than its population share

[3] Of course, under the assumption of initial separation, a pooling equilibrium is unlikely to arise in these circumstances.

would indicate, although this weight is still less than one-half. This "bias" is due the fact that minority types place more weight on their private signals in the limit than majority types do because the average pooled action is less informative for them. This bias can be substantial; for example, if Green agents represent 11% of the population, the weight of their preferences on the average action will be 33%. Of course, the significance of this bias depends on whether a small σ_μ^2/σ_s^2 ratio is driven by small preference differences or noisy individual signals. In the latter case, the effect of pooling on welfare may be important.

3.1 Welfare

We will consider the perspective of a social planner whose welfare function is the undiscounted time average payoff across all agents. This is equivalent to the limiting average payoff of cohort t, as t goes to infinity. Under this definition, separating equilibria trivially match the social optimum, but pooling equilibria will be inefficient.

Let us write W_{SLE} for social welfare in a pooling equilibrium. Using Lemma 2, we have

$$W_{SLE} = -\frac{r\alpha_R^\infty + g\alpha_G^\infty}{r+g}\sigma_s^2 \tag{2}$$

We will look at a number of possible policy interventions and ask whether each could improve social welfare. These include interventions that affect welfare in a pooling equilibrium as well as interventions that shift the population from a pooling to a separating equilibrium. First, the former case.

One obvious channel through which a social planner might influence learning is *via* public information. For example, if the social planner could identify Red and Green agents and observe their actions, she could release the average action taken by each type. This would generate complete learning and a limiting welfare loss of 0. Of course, to assume that the agents' types are observable is, in a sense, to assume away the social learning problem. Alternatively, imagine a social planner whose only role is to publish the average action in neighborhood 1. Could she ever improve welfare by injecting noise into her announcement? The question arises because the pooling equilibrium puts too much weight, from a welfare perspective, on the minority target. Adding noise to the announced average action will induce both types to put more weight on their private signals, but this effect will tend to be greater for the majority type (as it is starting from a lower private signal weight). This should shift the average action toward the majority's target, which in principle could improve welfare. However, it turns out that this is not a possibility:

Proposition 4 *The social planner cannot improve welfare in a pooling equilibrium by adding noise to the average action.*

Proof. The problem is already evident in (2): adding noise can only improve welfare if the ultimate effect of this is to reduce α_R^∞ or reduce α_G^∞, but in fact, they must both increase. Suppose the social planner adds $\omega^t \sim N(0,\gamma)$ to the average action before announcing it. Let α_T and ρ_T (α'_T and ρ'_T) be the limiting weights without (with) the intervention. Then we have

$$\alpha'_T = \frac{2(1-\rho'_T)^2\sigma_\mu^2 + \gamma}{2(1-\rho'_T)^2\sigma_\mu^2 + \gamma + \sigma_s^2}$$

If $\rho'_T \leq \rho_T$, then $\alpha'_T > \alpha_T$. This must be true for one of the types; suppose wlog it is true for G. Then by the second part of Proposition 3, which still applies, we have

$$\alpha'_R = \frac{g}{r}\frac{\rho'_R}{\rho'_G}\alpha'_G$$

Then, $\alpha'_G > \alpha_G$ and $\rho'_G \leq \rho_G$ imply $\alpha'_R > \alpha_R$. Thus, adding noise increases both α_R^∞ and α_G^∞, which by (2) decreases welfare. ■

On the other hand, if the social planner can make the agents' private signals more precise, it should always do so. The benefits here come through two channels. There is the direct effect that can be observed in (2) and there is also an indirect effect: when private signals improve, the average action moves closer to the majority's target, and this is welfare-improving.

Proposition 5 *The social planner can improve welfare in a pooling equilibrium by making private signals more precise (reducing σ_s^2).*

Proof. Appendix. ■

It may also be possible for the social planner to steer the population into an efficient separating equilibrium through appropriate an appropriate policy. Here we suppose that the planner can enforce separation for a limited period of time before allowing agents to choose neighborhoods freely.[4] This is modeled in reduced form with the assumption that the planner can set the the seed populations in each neighborhood to $r_0 = rk$ and $g_0 = gk$ for some k reflecting the length of time for which separation can be imposed. We ask whether there are any values of k for which the planner can guide the population to an efficient separating equilibrium.

Proposition 6 *There exists some $\bar{k}$ such that imposing $k \geq \bar{k}$ periods of separation suffices to generate a separating and efficient SLE.*

Proof. By taking $\bar{k}$ large enough, we can ensure that each neighborhood's error is closer to its target action than the expected preference difference between types. To be precise, take $\bar{k}$ such that $\sigma_s^2/\bar{k}r < \sigma_s^2/\bar{k}g < 2\sigma_\mu^2$. Then assuming separation through cohort t (for arbitrary t), $\tau_{2G}^t > 1/2\sigma_\mu^2 > \tau_{1G}^t$ and $\tau_{1R}^t > 1/2\sigma_\mu^2 > \tau_{2R}^t$, so cohort $t+1$ separates as well. ■

Given the slow rate at which agents learn, both in separating and pooling equilibria, it is also worth asking how welfare compares along the learning path. To illustrate the importance of this question, consider a separated Red neighborhood with a cohort size $r = 1$ (and $r_0 = 1$). The expected utility of cohorts 1, 1000, and 1,000,000 respectively are -0.5, -0.0757, and -0.0380. Looking only at the fact that expected utility rises eventually to 0 is to miss a large part of the picture. Here, the analysis becomes intractable, and we resort to numerical examples.

Example 1

Suppose $r = 10$, $g = 1$, $k = 1$, $\sigma_s^2 = 1$, and $\sigma_\mu^2 = 0.1$. In this case, the SLE will be a pooling equilibrium. The table below provides two measures of short run welfare for each type. The first is the time average payoff per agent for the first 1,000,000 cohorts. The second is a discounted sum of payoffs, weighted by cohort size, with discount rate $\delta = 0.01$. The first two columns for each type present welfare under a (hypothetical) separating scenario and in the pooling SLE. The third column presents the percentage welfare gain under pooling.

[4] This, of course, requires that the planner can identify types, at least temporarily.

	Red			Green		
	Separating (s)	Pooling (p)	$\frac{p-s}{\lvert s \rvert}$	Separating (s)	Pooling (p)	$\frac{p-s}{\lvert s \rvert}$
Time Avg. Utility	−0.0290	−0.0283	2.2%	−0.0426	−.1358	−219%
Discounted Utility	−56.9	−52.8	7.2%	−14.7	−17.3	−16.4%

In this case, pooling provides modest benefits to the majority type over any reasonable time horizon. In contrast, pooling hurts the minority type rather substantially, particularly later cohorts who bear almost all of the losses. That these losses occur in the SLE can be traced to the short-run incentives of early Green cohorts.

Example 2

Suppose $r = g = 1$, $k = 1$, $\sigma_s^2 = 1$, and $\sigma_\mu^2 = 0.01$. In this case, the SLE will be separating. However, the preference difference between the types is small enough that if they were to pool initially, both types would benefit from the larger sample size, and these benefits would persist for a relatively long time. To illustrate, we consider the same two measures of welfare as in Example 1. (The Green type is omitted, as payoffs for the two types are identical.)

	Red		
	Separating (s)	Pooling (p)	$\frac{p-s}{\lvert s \rvert}$
Time Avg. Utility	−0.0426	−0.0421	1.1%
Discounted Utility	−14.7	−12.1	18.0%

It is an indication of the slow speed of learning that the utility of the one millionth cohort is higher under pooling than under separation, even though the former converges to −0.005 and the latter to 0. Any welfare function with even a modest amount of discounting will recommend intervention to encourage pooling in this example.

4 Discussion

1. **Observing average vs. individual, actions**

   Sequential learning models often assume that agents observe the full sequence of preceding actions. What are the consequences of replacing this unrealistic assumption with the equally stylized assumption that an agent only observes the average preceding action? First, we remove the possibility for an agent in a pooling neighborhood to infer additional information about her type's target from the distribution of actions. In many markets, the premise that agents do not have or process highly detailed information about the distribution of past actions is probably reasonable. Even if agents could mine these distributions, they would be running a race against time – with each successive period, actions by the two types differ less as each type relies more on the public history. Whether complete learning could prevail in this situation is an open question.

   A second implication of observing averages is that agents cannot distinguish the actions of more recent cohorts from those who acted earlier; we discuss this in the next section.

2. **The speed of learning**

   The slow learning result of Proposition 2 depends on two factors. The first is that the reliance of new cohorts on their private signals shrinks commensurately as the precision of the public statistic increases. The second is that the contribution of new cohorts to the public statistic shrinks at rate t. Boosting the influence of new cohorts through either of these channels will tend to speed up learning. For example, in the model of Vives, cohorts are observed separately, but with noise. Due to the first factor above, the signal to noise ratio in the observed actions of later cohorts deteriorates, and consequently, the contribution of new cohorts to the public statistic diminishes, but at rate slower than t. The net result is learning at rate $t^{1/3}$. Going further, if the aggregate statistic places a constant weight on the actions of the newest cohorts, then learning at rate t is restored.

3. **Benefits of diversity**

   The idea that a diversity of preferences can benefit a group by discouraging it from swinging to extremes is a familiar one. There is a natural extension of this idea to sequential learning, namely that neighborhoods including several different types of agents are less susceptible to herding on the "wrong" action because it is more difficult to find histories for which all of the types ignore their signals. Our formulation rules out this scenario in order to focus on the negative implications of pooling, but a more complete welfare analysis would take both factors into account.

5 Appendix

Proof of Lemma 1

Proof. We can write $\bar{x}$ as $\bar{x} = \lambda_R + \rho_G(\mu_G - \mu_R) + \upsilon$. Consider that the agent has two conditionally independent, normally distributed signals regarding the true value of λ_R. It is a standard result that in order to minimize a quadratic loss function, her optimal action is to take a linear combination of the two signals with weights proportional to their precisions. The precision of her private signal is just τ_s, and the variance of $\rho_G(\mu_G - \mu_R) + \upsilon$ is $2\rho_G^2\sigma_\mu^2 + \sigma_\upsilon^2$, hence the result. ■

Proof of Lemma 3

Proof. By assumption, the seed cohort is separated, so $\rho_{1R}^0 = \rho_{2G}^0 = 1$. Agents in the seed cohort place full weight on their private signals, so $\sigma^2_{\upsilon_1^0} = \sigma_s^2/r_0$ and $\sigma^2_{\upsilon_2^0} = \sigma_s^2/g_0$. Now suppose for arbitrary t that cohort t's beliefs about the average action in the two neighborhoods can be summarized by ρ_{nT}^{t-1} and $\sigma^2_{\upsilon_n^{t-1}}$, $n \in \{1,2\}$, $T \in \{R,G\}$. Then, apply Lemmas 1 and 2. All agents of the same type have the same optimal neighborhood choice, so there are four possible neighborhood assignments for the cohort t agents: $(r_1^t, g_1^t) \in \{(0,0),(0,g),(r,0),(r,g)\}$ (with the complement of agents in neighborhood 2). In any of these cases, the optimal own signal weight for a cohort t agent of type T in neighborhood n is α_{nT}^t as given by Lemma 1. Since the period $t-1$ average action in neighborhood n was given by

$$\bar{x}_n^{t-1} = \rho_{nR}^{t-1}\lambda_R + \rho_{nG}^{t-1}\lambda_G + \upsilon_n^{t-1}$$

the action of an arbitrary cohort t, type T agent in this neighborhood is

$$\alpha_{nT}^t(\lambda_T + \varepsilon_i^t) + (1-\alpha_{nT}^t)(\rho_{nR}^{t-1}\lambda_R + \rho_{nG}^{t-1}\lambda_G + \upsilon_n^{t-1})$$

The average action after period t is composed of r_n^t such actions with $T = R$, g_n^t such actions with $T = G$, and D_n^{t-1} earlier actions averaging $\bar{x}_n^{t-1}$. Thus the new average action is

$$\begin{aligned}
\bar{x}_n^t &= \frac{\begin{array}{r} D_n^{t-1}\bar{x}_n^{t-1} + (r_n^t(1-\alpha_{nR}^t) + g_n^t(1-\alpha_{nG}^t))\bar{x}_n^{t-1} \\ + r_n^t\alpha_{nR}^t\lambda_R + g_n^t\alpha_{nG}^t\lambda_G + \alpha_{nR}^t\Sigma_{i=1}^{r_n^t}\varepsilon_i^t + \alpha_{nG}^t\Sigma_{j=1}^{g_n^t}\varepsilon_j^t \end{array}}{D_n^{t-1} + r_n^t + g_n^t} \\
&= \frac{K_n^t}{D_n^t}\bar{x}_n^{t-1} + \frac{r_n^t\alpha_{nR}^t}{D_n^t}\lambda_R + \frac{g_n^t\alpha_{nG}^t}{D_n^t}\lambda_G + \frac{\alpha_{nR}^t\Sigma_{i=1}^{r_n^t}\varepsilon_i^t + \alpha_{nG}^t\Sigma_{j=1}^{g_n^t}\varepsilon_j^t}{D_n^t} \\
&= \frac{K_n^t\rho_{nR}^{t-1} + r_n^t\alpha_{nR}^t}{D_n^t}\lambda_R + \frac{K_n^t\rho_{nG}^{t-1} + g_n^t\alpha_{nG}^t}{D_n^t}\lambda_G + \upsilon_n^t
\end{aligned}$$

where

$$\upsilon_n^t = \frac{K_n^t}{D_n^t}\upsilon_n^{t-1} + \frac{\alpha_{nR}^t\Sigma_{i=1}^{r_n^t}\varepsilon_i^t + \alpha_{nG}^t\Sigma_{j=1}^{g_n^t}\varepsilon_j^t}{D_n^t}$$

Because υ_n^{t-1} and the ε_i^t's and ε_j^t's are uncorrelated and distributed $N(0, \sigma^2_{\upsilon_n^{t-1}})$ and $N(0,\sigma_s^2)$ respectively, υ_n^t is normally distributed with zero mean and variance

$$\sigma^2_{\upsilon_n^t} = \left(\frac{K_n^t}{D_n^t}\right)^2 \sigma^2_{\upsilon_n^{t-1}} + r_n^t\left(\frac{\alpha_{nR}^t}{D_n^t}\right)^2\sigma_s^2 + g_n^t\left(\frac{\alpha_{nG}^t}{D_n^t}\right)^2\sigma_s^2$$

Thus $\bar{x}_n^t$ can be written in the form $\bar{x}_n^t = \rho_{nR}^t \lambda_R + \rho_{nG}^t \lambda_G + \upsilon_n^t$, where ρ_{nR}^t, ρ_{nG}^t, and $\sigma^2_{\upsilon_n^t}$ are specified by the recursive relationships stated in the Lemma whenever $\bar{x}_n^{t-1}$ also takes this form. Because this format is valid for $t = 0$, it is therefore valid for all t. ■

Lemma 5 *If the seed cohorts are consistent, then under separation public information is more precise in the Red neighborhood in every period, that is, $\sigma^2_{\upsilon_1^t} < \sigma^2_{\upsilon_2^t}$ for all $t \geq 1$.*

Proof. Consider an arbitrary neighborhood under separation in which the seed population is $n_0 > 0$ and the cohort size is n. (Ultimately we will look at $(n_0, n) = (r_0, r)$ or $(n_0, n) = (g_0, g)$. Let $k = n_0/n$. Write $\sigma_t^2 = \sigma^2_{\upsilon^t}$ for the variance of the public information in this neighborhood and let κ_t be defined by $\sigma_t^2 = \kappa_t \sigma_s^2$. We will use Lemmas 1 and 3 to characterize a recursive relationship for κ_t. Lemma 1 gives us $\alpha_t = \frac{\kappa_{t-1}}{1+\kappa_{t-1}}$, and Lemma 3 yields $D_t = n(t+k)$ and hence $\kappa_t = h(\kappa_{t-1}; n, \tilde{t})$, $\tilde{t}$where $\tilde{t} = t + k$ and

$$h(\kappa; n, \tilde{t}) = \left(1 - \frac{\kappa}{(\kappa+1)\,\tilde{t}}\right)^2 \kappa + \frac{1}{n} \frac{\kappa^2}{(\kappa+1)^2\,\tilde{t}^2} \tag{3}$$

First, we want to show that, holding k and n constant, h is monotonic: a larger κ_{t-1} implies a larger κ_t. Differentiating h and conveniently grouping terms yields

$$h_\kappa(\kappa; n, \tilde{t}) = \frac{n\left(\tilde{t}-1\right)^2 \kappa^3 + 3n\left(\tilde{t}-1\right)^2 \kappa^2 + 2\kappa + 3n\tilde{t}(\tilde{t}-1)\kappa + n\tilde{t}(\tilde{t}-\kappa)}{n\,(1+\kappa)^3\,\tilde{t}}$$

Since $0 < \kappa_t \leq 1$ and $k > 0$, all of the terms in the numerator are strictly positive for $t \geq 1$, so $h_\kappa(\kappa; n, \tilde{t}) > 0$. Next, observe that $n' > n \Rightarrow h(\kappa; n', \tilde{t}) < h(\kappa; n, \tilde{t})$. Finally, consider two separated neighborhoods with consistent seed cohorts $r_0 = kr$ in neighborhood 1 and $g_0 = kg$ in neighborhood 2. Suppose $\kappa_{1,t-1} < \kappa_{2,t-1}$ for some $t \geq 1$. Then

$$\begin{aligned} \kappa_{1t} = h(\kappa_{1,t-1}; r, t+k) &< h(\kappa_{2,t-1}; r, t+k) \\ &< h(\kappa_{2,t-1}; g, t+k) = \kappa_{2t} \end{aligned}$$

Furthermore, we have $\kappa_{1,0} = 1/r_0 = 1/kr < 1/kg = 1/g_0 = \kappa_{2,0}$. By induction, $\kappa_{1t} < \kappa_{2t}$ for all $t \geq 0$, and consequently, $\sigma^2_{\upsilon_1^t} < \sigma^2_{\upsilon_2^t}$ for all $t \geq 0$. ■

Lemma 6 *If the seed cohorts are consistent and full strength, then under separation, $\sigma^2_{\upsilon_2^t} - \sigma^2_{\upsilon_1^t}$ is strictly decreasing in t.*

Proof. Equivalently, we must show that

$$\begin{aligned} \kappa_{1,t-1} - \kappa_{1,t} &< \kappa_{2,t-1} - \kappa_{2,t}, \text{ or} \\ (1 - \kappa_{1,t}/\kappa_{1,t-1})\kappa_{1,t-1} &< (1 - \kappa_{2,t}/\kappa_{2,t-1})\kappa_{2,t-1} \end{aligned}$$

holds for all $t \geq 1$. Since we know that $\kappa_{1,t} < \kappa_{2,t}$, it suffices to prove the stronger claim that

$$\kappa_{1,t}/\kappa_{1,t-1} > \kappa_{2,t}/\kappa_{2,t-1}$$

for all $t \geq 1$. (Variance declines proportionately more slowly in neighborhood 1, after its initial advantage.) Complications arise because new signals arrive faster in neighborhood 1; we must show that this is outweighed by Red agents' greater reliance on public information. The first step is to show that $\tilde{t}(\kappa_{2,t}-\kappa_{1,t})$ is increasing. Loosely, this says that the greater incentive to use public information in neighborhood 1 does not decrease too quickly over time. Then we show that the importance of new information declines about one order of t faster than this.

Define $b_{n,t}=\tilde{t}\kappa_{n,t}$, so we must show that $\Delta b_t = b_{2,t}-b_{1,t}$ is increasing. Substituting in (3) and simplifying, we have

$$\begin{aligned} b_{t+1} &= \frac{\tilde{t}+1}{\tilde{t}} b_t \left(1-\frac{b_t}{(b_t+\tilde{t})(\tilde{t}+1)}\right)^2 + \frac{1}{n}\frac{b_t^2}{(b_t+\tilde{t})^2(\tilde{t}+1)} \\ &= \frac{\left(b_t+\tilde{t}+1\right)^2 \tilde{t} + b_t/n}{\left(\tilde{t}+1\right)\left(b_t+\tilde{t}\right)^2} b_t \end{aligned}$$

With some algebra, we have

$$\begin{aligned} b_{t+1}-b_t &= b_t \frac{(\tilde{t}-b_t)(b_t+\tilde{t})+\tilde{t}+b_t/n}{(b_t+\tilde{t})^2(\tilde{t}+1)} \\ &\equiv C_t(b_t,n) \end{aligned}$$

so $\Delta b_{t+1} = \Delta b_t + C_t(b_{2,t},g) - C_t(b_{1,t},r)$. If $C_t(b_{2,t},g)-C_t(b_{1,t},r)>0$, we are done. Since $C_t(b_t,n)$ is decreasing in n, it suffices to show $C_t(b_{2,t},r)-C_t(b_{1,t},r)>0$. We will show this for $r\geq 2$. (By assumption, $r>g\geq 1$.) Differentiate C_t :

$$\frac{dC_t(b_t,n)}{db_t} = \frac{\left(\tilde{t}^3 - b_t(b_t^2+3b_t\tilde{t}+\tilde{t}^2)\right)+\tilde{t}(\tilde{t}-b_t)+2\tilde{t}b_t/n}{\left(b_t+\tilde{t}\right)^3\left(\tilde{t}+1\right)}$$

The second and third terms in the numerator are positive ($\kappa_t \leq 1 \Rightarrow b_t \leq \tilde{t}$). The first term is positive whenever $\kappa_t < \sqrt{2}-1 \approx 0.41$. But if $n\geq 2$ and $k\geq 1$, then $\kappa_1 \leq \frac{13}{36}\approx 0.36$. (If $n=2$ and $k=1$, then $\kappa_1=\frac{13}{36}$ and straightforward algebra reveals that κ_1 is decreasing in n and in k for $n\geq 2$ and $k\geq 1$.) Then because κ_t is decreasing, $\kappa_t \leq \frac{13}{36}$ for all $t\geq 1$. Thus, if $r\geq 2$, $\Delta b_{t+1} > \Delta b_t$ for all $t\geq 1$. For $t=0$, use the fact that $\kappa_0=\frac{1}{kn}$ and $\tilde{t}=k$ to get

$$b_1 - b_0 = \frac{kn}{(1+kn)^2}$$

which is decreasing in n for $kn>1$. Thus we have $b_{2,1}-b_{2,0} > b_{1,1}-b_{1,0}$, so $\Delta b_1 > \Delta b_0$. Hence Δb_t is increasing for all $t\geq 0$.

2. We want to show that $\kappa_{1,t}-\kappa_{1,t+1} < \kappa_{2,t}-\kappa_{2,t+1}$ for all $t\geq 0$. Expand 3 to get

$$\kappa_t - \kappa_{t+1} = \frac{\kappa_t}{(\tilde{t}+1)^2}\left(2(\tilde{t}+1)\alpha_t - \alpha_t^2 - \frac{1}{n}\frac{\alpha_t}{1+\kappa_t}\right)$$

where $\alpha_t = \kappa_t/(1+\kappa_t)$ as before. Because $\kappa_{1,t}<\kappa_{2,t}$, it suffices to show that the second term in this expression is smaller in neighborhood 1. That is, we need

$$L_t \equiv 2(\tilde{t}+1)(\alpha_{2,t}-\alpha_{1,t}) - (\alpha_{2,t}^2-\alpha_{1,t}^2) - (\frac{1}{g}\frac{\alpha_{2,t}}{1+\kappa_{2,t}} - \frac{1}{r}\frac{\alpha_{1,t}}{1+\kappa_{1,t}}) > 0 \tag{4}$$

for all $t \geq 0$. First, observe that for $t = 0$, we have $\kappa_t = \frac{1}{kn}$, so the final term in L_t becomes $k(\alpha_{2,0}^2 - \alpha_{1,0}^2)$, so for $t = 0$ we need the following to hold.

$$\begin{aligned} 2(k+1)(\alpha_{2,0} - \alpha_{1,0}) - (1+k)(\alpha_{2,0}^2 - \alpha_{1,0}^2) &> 0 \text{ or} \\ (k+1)(2 - (\alpha_{2,0} + \alpha_{1,0}))(\alpha_{2,0} - \alpha_{1,0}) &> 0 \end{aligned}$$

But $\alpha_{2,t} > \alpha_{1,t}$ whenever $\kappa_{2,t} > \kappa_{1,t}$ and $\alpha_{n,t} < 1$, so this inequality holds, and $\kappa_{1,t} - \kappa_{1,t+1} < \kappa_{2,t} - \kappa_{2,t+1}$ holds for $t = 0$. Now suppose that the relationship holds for all $t \leq t' - 1$ and evaluate L_t at t'. We can write the last term in brackets as

$$k(\kappa_{2,0}\beta_{2,t'} - \kappa_{1,0}\beta_{1,t'})$$

where $\beta_{n,t'} = \kappa_{n,t}/(1+\kappa_{n,t})^2$. Note that $\beta_{n,t}$ is increasing in $\kappa_{n,t}$ when $\kappa_{n,t} \leq 1$, as is the case here. Manipulate this to get

$$\begin{aligned} & k(\alpha_{2,0}^2 - \alpha_{1,0}^2) + k(\kappa_{2,0}(\beta_{2,t'} - \beta_{2,0}) + \kappa_{1,0}(\beta_{1,0} - \beta_{1,t'})) \\ < \; & k(\alpha_{2,0}^2 - \alpha_{1,0}^2) + k\kappa_{2,0}((\beta_{2,t'} - \beta_{1,t'}) - (\beta_{2,0} - \beta_{1,0})) \\ < \; & k(\alpha_{2,0}^2 - \alpha_{1,0}^2) + k\kappa_{2,0}(\beta_{2,t'} - \beta_{1,t'}) \\ < \; & k(\alpha_{2,0}^2 - \alpha_{1,0}^2) + k\kappa_{2,0}(\alpha_{2,t'} - \alpha_{1,t'}) \end{aligned}$$

The second and third steps follow because $\kappa_{2,0} > \kappa_{1,0}$ and because $\beta_{1,0} - \beta_{1,t'} > 0$ (monotonicity of $\beta_{n,t}$ and $\kappa_{n,t}$ decreasing in t). The last step uses the easily derived fact that $x > y \Rightarrow \frac{x}{1+x} - \frac{y}{1+y} > \frac{x}{(1+x)^2} - \frac{y}{(1+y)^2}$. Thus we have

$$\begin{aligned} L_t &> 2(\tilde{t}+1)(\alpha_{2,t'} - \alpha_{1,t'}) - (\alpha_{2,t'}^2 - \alpha_{1,t'}^2) - k(\alpha_{2,0}^2 - \alpha_{1,0}^2) - k\kappa_{2,0}(\alpha_{2,t'} - \alpha_{1,t'}) \\ &= \left(2k + 2t' + 2 - (\alpha_{2,t'} + \alpha_{1,t'}) - k\kappa_{2,0}\right)(\alpha_{2,t'} - \alpha_{1,t'}) - k(\alpha_{2,0} + \alpha_{1,0})(\alpha_{2,0} - \alpha_{1,0}) \\ &> \left(k + 2t' + 1\right)(\alpha_{2,t'} - \alpha_{1,t'}) - k(\alpha_{2,0} + \alpha_{1,0})(\alpha_{2,0} - \alpha_{1,0}) \end{aligned}$$

The last step uses the fact that $\kappa_{n,t} \leq 1$ and hence $\alpha_{n,t} \leq \frac{1}{2}$. There are two cases to consider. Suppose $\alpha_{2,t'} - \alpha_{1,t'} > \alpha_{2,0} - \alpha_{1,0}$. Then,

$$\begin{aligned} L_{t'} &> \left(k + 2t' + 1 - k(\alpha_{2,0} + \alpha_{1,0})\right)(\alpha_{2,0} - \alpha_{1,0}) \\ &> (2t'+1)(\alpha_{2,0} - \alpha_{1,0}) > 0 \end{aligned}$$

Now suppose $\alpha_{2,t'} - \alpha_{1,t'} < \alpha_{2,0} - \alpha_{1,0}. > 0$. We have

$$\begin{aligned} L_{t'} &> \left(k + 2t' + 1\right)(\alpha_{2,t'} - \alpha_{1,t'}) - k(\alpha_{2,0} - \alpha_{1,0}) \\ &= \Delta d_{t'} - \Delta d_0 + (t'+1)(\alpha_{2,t'} - \alpha_{1,t'}) \end{aligned}$$

where we define $d_{n,t} = \tilde{t}\alpha_{n,t}$ and $\Delta d_t = \tilde{t}(\alpha_{2,t} - \alpha_{1,t})$ by analogy to b_t. Observe that $\Delta d_t = \frac{\Delta b_t}{(1+\kappa_{2,t})(1+\kappa_{1,t})}$. Thus Δb_t increasing and $\kappa_{n,t}$ decreasing in t imply that Δd_t is increasing in t. But then, $\Delta d_{t'} - \Delta d_0 > 0$ in the expression above, so we have $L_t > 0$ as we claimed. This shows that $\kappa_{1,t} - \kappa_{1,t+1} < \kappa_{2,t} - \kappa_{2,t+1}$ for all $t \geq 0$, completing the proof.

■

Proof of Proposition 2

Proof. Consider an arbitrary separated neighborhood. We must show that $\sigma^2_{\upsilon^t_n} \to 0$ at rate $1/\ln t$, or equivalently, that $\kappa_t \to 0$ at rate $1/\ln t$, where $\kappa_{n,t}$ is defined by $\sigma^2_{\upsilon^t_n} \equiv \kappa_{n,t}\sigma^2_s$ and we drop the neighborhood subscript to avoid clutter. Let n_0 and n be the size of cohorts 0 and $t > 0$ respectively in this neighborhood, and let $k = n_0/n$. Lemma 5 shows that κ_t is defined by

$$\kappa_{t+1} = (1 - \alpha_t/(1 + k + t))^2\kappa_t + \frac{1}{n}\left(\frac{\alpha_t}{1 + k + t}\right)^2 \tag{5}$$

where the own-signal weight α_t is given by $\alpha_t = \kappa_t/(1 + \kappa_t)$. First we will show that κ_t is strictly decreasing and thus must converge to some κ_∞. Then we compare the sequence of finite differences $\kappa_{t+1} - \kappa_t$ to a family of differential equations in order to bound κ_t above and below by continuous functions that go to zero at rate $1/\ln t$.

Using the (5), we have

$$\begin{aligned}
\kappa_{t+1} &< (1 - \alpha_t/(1 + k + t))^2\kappa + \frac{1}{n}\left(\frac{\alpha_t}{1 + k + t}\right)\kappa_t \\
&\leq (1 - \alpha_t/(1 + k + t))\kappa + \frac{1}{n}\left(\frac{\alpha_t}{1 + k + t}\right)\kappa_t \\
&\leq (1 - \alpha_t/(1 + k + t))\kappa + \left(\frac{\alpha_t}{1 + k + t}\right)\kappa_t \\
&= \kappa_t
\end{aligned}$$

The first line uses $\alpha_t < \kappa_t$ and $1 + k + t \geq 1$, and the third line uses $n \geq 1$. Then, $\kappa_{t+1} < \kappa_t$ and $\kappa_t \geq 0\ \forall\, t$ imply that κ_t converges.

Now we find the upper bound. Select an arbitrary t^*. For all $t \geq t^*$ we have

$$\begin{aligned}
\kappa_t - \kappa_{t+1} &= \frac{\alpha_t}{1 + k + t}\left(2\kappa_t - \left(\frac{\alpha_t}{1 + k + t}\right)(\kappa_t + \frac{1}{n})\right) \\
&> \frac{\alpha_t}{1 + k + t}\left(2\kappa_t - \left(\frac{\alpha_t}{1 + k + t}\right)(\kappa_t + 1)\right) \\
&= \frac{\alpha_t\kappa_t}{1 + k + t}\left(2 - \frac{1}{1 + k + t}\right) \\
&> \frac{\alpha_t\kappa_t}{1 + k + t} \\
&= \frac{\kappa_t^2}{1 + \kappa_t}\frac{1}{1 + k + t} \\
&> \frac{\kappa_t^2}{1 + \kappa_0}\frac{1}{1 + k + t} \\
&> M_1\kappa_t^2/t
\end{aligned}$$

where the constant M_1 in the final line is chosen so that $M_1 < \frac{1}{1+\kappa_0}\frac{t^*}{1+k+t^*}$. Consider the family of solutions $y_s(t) = \frac{\kappa_s}{1+M_1\kappa_s \ln t/s}$ to the differential equation $\dot{y} = -M_1 y^2/t$ parameterized by the initial condition (s, κ_s). It is easy to see that y_s is decreasing in t, and therefore that $M_1 y_s^2/t$

is decreasing in t, so y_s is convex. Then, for each $t' \geq t^*$, we have

$$\begin{aligned} \kappa_{t'+1} - \kappa_{t'} &< -M_1 \kappa_{t'}^2 / t' \\ &< \int_{t'}^{t'+1} \dot{y}_{t'} \, dt \\ &= y_{t'}(t'+1) - \kappa_{t'} \end{aligned}$$

and hence

$$\kappa_{t'+1} < y_{t'}(t'+1)$$

But this means that

$$y_{t'+1}(t'+1) < y_{t'}(t'+1)$$

for all $t' \geq t^*$. Since differential equation trajectories cannot cross, we have $y_{t+1} < y_t$ for all $t' \geq t^*$. Putting the pieces together, we have $\kappa_t < y_{t^*}(t) = \frac{\kappa_{t^*}}{1 + M_1 \kappa_{t^*} \ln t / t^*}$ for all $t \geq t^*$. (Recall that our choice of t^* was arbitrary.) Thus, κ_t goes to zero at least as fast as $1/\ln t$.

Now we find a lower bound for κ_t in order to show that it vanishes no faster than $1/\ln t$. Proceeding in a manner similar to that above, we have

$$\begin{aligned} \kappa_t - \kappa_{t+1} &= \frac{\alpha_t}{1+k+t} \left(2\kappa_t - \left(\frac{\alpha_t}{1+k+t} \right) (\kappa_t + \frac{1}{n}) \right) \\ &< \frac{2\alpha_t \kappa_t}{1+k+t} \\ &< \frac{2\kappa_t^2}{t} \end{aligned}$$

We choose an arbitrary $t^* \geq 9$. Let $z_s(t)$ be the family of solutions to $\dot{y} = -M_2 y^2 / t$ with initial condition (s, κ_s). The idea is that the z_s should decline faster than κ_t; in order to achieve this, M_2 must be taken large enough to compensate for the convexity of z_s. We choose $M_2 = 10$. In this case, for $u \in [0,1]$ and $s \geq t^*$ we have

$$\begin{aligned} \dot{z}_s(s+u) &= -\frac{10}{s+u} \left(\frac{\kappa_s}{1 + 10 \left(\ln(s+u)/s \right) \kappa_s} \right)^2 \\ &< -\frac{10\kappa_s^2}{s+1} \left(\frac{1}{1 + 10 \ln(10/9)} \right)^2 \\ &< -\frac{20}{9} \frac{\kappa_s^2}{s+1} \\ &< -\frac{2\kappa_s^2}{s} \\ &< \kappa_{s+1} - \kappa_s \end{aligned}$$

From here we proceed as for the upper bound to show that $\kappa_{t'+1} > z_{t'}(t'+1)$ and $z_{t'+1} > z_{t'}$ for all $t' \geq t^*$. Consequently, we have $\kappa_t > z_{t^*}(t) = \frac{\kappa_{t^*}}{1 + 10 \kappa_{t^*} \ln t / t^*}$, so $\kappa_t \to 0$ at rate no faster than $1/\ln t$. ■

Proof of Proposition 5

Proof. We simply differentiate (2) with respect to σ_s^2. Using the definitions of α_T^∞ (and dropping the superscript ∞) we have

$$\frac{dW_{SLE}}{d\sigma_s^2} = -\frac{1}{r+g}\left(r\alpha_R^2 + g\alpha_G^2\right) + \left(\frac{\partial W_{SLE}}{\partial \rho_R}\frac{d\rho_R}{d\sigma_s^2} + \frac{\partial W_{SLE}}{\partial \rho_G}\frac{d\rho_G}{d\sigma_s^2}\right)$$

We must show that the second term is negative. First note that $d\rho_R/d\sigma_s^2 = -d\rho_G/d\sigma_s^2 < 0$, using Proposition 3 and the fact that $\rho_R > 1/2$. Thus it suffices to show that $\frac{\partial W_{SLE}}{\partial \rho_R} - \frac{\partial W_{SLE}}{\partial \rho_G} > 0$, or equivalently, that $Z \equiv r\partial\alpha_R/\partial\rho_G - g\partial\alpha_G/\partial\rho_R > 0$. Using the definition of α_R, we have

$$\begin{aligned}\frac{d\alpha_R}{d\rho_G} &= \frac{2}{\rho_G}\alpha_R(1-\alpha_R) \\ \frac{d\alpha_G}{d\rho_R} &= \frac{2}{\rho_R}\alpha_G(1-\alpha_G)\end{aligned}$$

Using Proposition 3, we can write $Z = ((1-\alpha_R)/\rho_G^2 - (1-\alpha_G)/\rho_R^2)\kappa$, where $\kappa = r\alpha_R\rho_G = g\alpha_G\rho_R > 0$. Using the fact that $\rho_R > 1/2$ and $\alpha_R < \alpha_G$, we have $Z > 0$, as claimed. ■

References

[1] Venkatesh Bala and Sanjeev Goyal. Conformism and diversity under social learning. *Economic Theory*, 17(1):101–120, January 2001.

[2] Abhijit V. Banerjee. A simple model of herd behavior. *Quarterly Journal of Economics*, 107(3):797–817, August 1992.

[3] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change in informational cascades. *Journal of Political Economy*, 100(5):992–1026, October 1992.

[4] Glenn Ellison and Drew Fudenberg. Rules of thumb for social learning. *Journal of Political Economy*, 101(4):612–43, August 1993.

[5] Glenn Ellison and Drew Fudenberg. Word-of-mouth communication and social learning. *Quarterly Journal of Economics*, 110(1):93–125, February 1995.

[6] Lones Smith and Peter Sorensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, March 2000.

[7] Xavier Vives. How fast do rational agents learn? *Review of Economic Studies*, 60(2):329–47, April 1993.