

Observational Learning and Demand for Search Goods*

Ken Hendricks

University of Texas at Austin

Alan Sorensen

Stanford University & NBER

Tom Wiseman

University of Texas at Austin

September 22, 2009

Abstract

In many differentiated good markets like music, books, and movies, the choice set of available products is overwhelmingly large and growing as new products flow into the market each week. Consumers are not aware of or are poorly informed about many of the available products. They learn about products and their preferences for them from the search and purchasing decisions of other consumers and through their own costly search. We use a variant of the sequential search models of Banerjee [7], Bikhchandani, Hirshleifer, and Welch [11] and Smith and Sorensen [33] to study market demand in these kinds of markets. We characterize the equilibrium dynamics of demand under different assumptions about the kind of information that consumers observe about the decisions of other consumers. Our main result consists of showing that long-run behavior in herding models can be invariant to the learning dynamics. We exploit this result to study the determinants of the likelihood that high-quality products end up with no sales and of the likelihood that low quality products end up with significantly positive sales.

*Hendricks: Department of Economics, University of Texas, Austin, TX 78712. Sorensen: Graduate School of Business, Stanford University, Stanford, CA 94305. Wiseman: Department of Economics, University of Texas, Austin, TX 78712. We received very helpful comments from Daron Acemoglu, Nathan Hendren, Johannes Horner, Phil Reny, Lones Smith, and the participants at the microeconomics seminars at Duke University and Hong Kong University, Workshop on Industrial Organization at University of Santiago, Chile, and the “Theory in A2” conference at the University of Michigan.

1 Introduction

In estimating demand in horizontally differentiated good markets, industrial organization economists typically assume that consumers are aware of all products and know their preferences for them. But, in many of these markets, the number of available products is very large with new products entering the market at more or less constant rates. As a result, consumers are often unaware of or poorly informed about many products, especially new products, and have to spend time learning about them before making their purchase decisions. According to the marketing literature, consumers typically focus their search on a relatively small set of products, called consideration sets. Consumers learn their preferences for the products in these smaller sets and make their purchases from them. Market demand then depends not only upon consumers preferences, but also upon the process by which they learn about the product space. Furthermore, its impact on demand determines in part seller markups and the structure of the market.¹ For these reasons, it is important to understand how consumers learn about the product space.

Many consumers learn by observing the choices of other consumers or information related to those choices. For example, in markets like music, consumers tend to buy albums they hear on the radio, but they hear what others buy since playing time is largely determined by album sales. Similarly, in markets for books and movies, consumers frequently use product sales or sales rankings to determine which products to consider. In online markets, sellers often post number of visits or consumer reviews to induce consumers to check out their products. Search engines respond to queries by selecting and ranking a small number of sellers that they believe are most likely to meet the needs of consumers based on search and purchasing decisions of other consumers.² The feedback effect of observational learning can cause a product's success or lack of success to reinforce itself. The key questions then become: how does observational learning affect the efficiency of the match between products and consumers with heterogeneous preferences? More precisely, what is the likelihood that, in the long-run, consumers ignore superior products or waste time and resources learning about inferior products? How does the likelihood of these events depend upon the kind of information that consumers observe about the behavior of other consumers?

¹Varian's [37] classic model of sales illustrates this point in a vertically differentiated product space.

²Advertisers bid on the basis of conversions (that is, purchases by visitors to their website), and the search engine scores their bids in part on the basis of click-through-rates (the fraction of ad viewers who visit the advertiser's website).

Our main goal in this paper is to address these questions. To study them, we use a variant of the sequential learning models introduced by Banerjee [7] and Bikhchandani, Hirshleifer and Welch [11] and subsequently generalized by Smith and Sorensen [33]. A large number of consumers with heterogenous preferences sequentially choose which products to search prior to making their purchasing decisions. Consumers do not know the quality of the products, which is either high or low, or their idiosyncratic preferences for them. The common quality implies that preferences are correlated across consumers and hence their actions are informative. Prior to making their search decisions, consumers also observe a private, informative signal about their preferences. In our baseline model, they also observe the ordered history of search and/or purchasing decisions of previous consumers. Searching a product is costly but it informs them about their preferences for the product. For example, in online markets, search would involve submitting a product search query, clicking on the links of one or more online sellers, and spending time at their websites learning about their product. In the market for recorded music, search consists of listening to songs on the internet or at the listening post in the store. A positive search cost provides a rationale for consumers *not* to consider products that they believe they are unlikely to buy. The cost per product may be quite small, but the number of products is too large for a consumer to search every product. We further simplify the analysis by assuming that consumer preferences over products are additive and the private signals about the products are independently distributed. These assumptions imply that the learning dynamics for different products in the choice set can be analyzed separately.

We show that, in the long-run, two outcomes are possible for a product: either a herd forms on the “no search” action and product sales converge to zero, or it forms on the “search” action and product sales converge to the market share that would occur if product quality was costlessly known. However, the probability of these two outcomes depends upon the kind of information that is provided to consumers. If only the search history is observed, then both outcomes occur with positive probability for high and low quality products. This model corresponds to the standard herding model in which actions are observable and the market never learns the true product quality. Observational learning can generate “bad herds” by leading a population of consumers to search low quality products that are not worth the cost of search or to not search (and therefore not buy) high quality products. By contrast, when consumers observe the purchasing history, the market can learn the true quality of high quality products and the market shares of low quality products are certain

to converge to zero. The sales information does not eliminate the bad herd on high quality products - market shares can still go to zero - but it does eliminate the bad herd on low quality products.

The fact that observational learning can generate bad herds is not surprising. The more interesting question is how does the probability of a bad herd depend upon the information provided, as well as upon model parameters such as product quality, price, signal quality, and search costs. Such an analysis is absent from the literature on herding, primarily because that literature considers models in which beliefs can enter the cascade sets (i.e., the sets of beliefs where consumers ignore their private information in choosing an action) from outside of the set. In these cases, the distribution of limit beliefs (and hence the probability of a bad herd) depends upon the learning dynamics. To derive comparative statics, the key restriction that needs to be imposed is that the posterior belief that the product is high quality has to be monotone increasing in the prior belief following each outcome. We provide sufficient conditions (roughly, log concavity) on the joint distribution of signals and preferences for the monotonicity condition to hold. Given this condition, the belief is forever trapped between the *boundary* points of the cascade sets associated with the two actions. Beliefs converge asymptotically, and the support of the distribution of limit beliefs consists of only two points. The probability mass on the point associated with the cascade set on the “wrong” action is then the probability of a bad herd. This result leads to two striking conclusions. First, models with different learning dynamics and signal structures but with the same cascade sets have the same distribution over limit beliefs. For example, the learning model in which only purchasing information is observed has different dynamics depending upon the precision of the signal or whether consumers also observe the search history, but the probability of a bad herd is the same. Second, we are able to obtain closed form solutions for the probability of bad herds. We use these formulas to analyze the determinants of bad herds.

An important, and standard, assumption of the sequential learning model is that consumers observe the entire past history of consumer decisions. This assumption plays a critical role in the analysis since it implies that the sequence of likelihood ratios that the product is low quality rather than high forms a martingale and allows one to use the Martingale Convergence Theorem to establish convergence of beliefs. But markets typically report only the total number of searches or purchases, so it is important in our context to determine whether the comparative static results are robust to the way in which informa-

tion about the decision history is aggregated. We consider one natural extension, namely, that consumers observe the aggregate number of prior searches or purchases. We develop different proof techniques to show that, given the monotonicity condition, beliefs do indeed converge to same two possible limits as in the baseline case. Furthermore, the distribution of limit beliefs is the same as in the baseline case. Thus, the probability of a bad herd occurring when consumers only observe aggregate numbers of search or purchases is the same as when they observe the entire history of search or purchasing decisions.

In Nelson’s [25] classification our model applies to “search” goods (i.e., those for which the consumer can learn her preferences before purchasing) rather than “experience” goods (which must be consumed before preferences are learned). The important distinction for our analysis is not whether a consumer can learn her utility perfectly or only partially through search, but rather whether or not she potentially buys the product repeatedly, in which case she may experiment by sampling a product rather than paying the search cost. That is, our model fits an intermediate case between search and experience goods, as long as either the product is a one-time (at most) purchase, or the product’s price is high relative to the search cost. Smith and Sorensen [34] examine the link between observational learning and a repeat consumer’s optimal experimentation.

The paper is organized as follows. In Section 2, we briefly discuss related literature. In Section 3, we describe our basic model. In Section 4, we characterize the equilibrium dynamics and outcomes. Section 5 examines two extensions of the basic model that make it more relevant for applications. In Section 6 we present the comparative static results. Section 7 examines evidence of herding in the market for recorded music. Section 8 concludes.

2 Related Literature

Previous work on sequential choice with heterogeneous preferences has focused on the case where agents observe the private component of their preferences before acting, so that the role of search is greatly diminished. Smith and Sorensen [33] model a finite set of preference types and show that with bounded signals an incorrect herd may arise. They further demonstrate the possibility of “confounded learning,” in which eventually all agents follow their own signals but no further public information about the state is generated. In our model, confounded learning cannot arise because every consumer’s utility from buying the product is increasing in product quality. Goeree, Palfrey, and Rogers [21] and Acemoglu,

Dahleh, Lobel, and Ozdaglar [2] show that if there is sufficient heterogeneity of preferences (specifically, if each action is optimal in all states for some type), then public beliefs converge to the truth almost surely, even if signals are bounded. Wiseman [38] derives a similar result under a different condition on heterogeneity, for the case where preference types are publically observed.

Consumers with incomplete knowledge of their preferences can also learn through experimentation. Orphanides and Zervos [26] show that a potential drug user who attaches high probability to the state in which he is not susceptible to addiction may experiment and in fact become addicted. Ali [4] shows that a consumer who is pessimistic about his level of self-control may always choose to commit himself in advance, and thus never learn whether he has high or low self-control. Those settings are very different from ours, but the results that incomplete information about preferences can have long-run consequences are similar in spirit.

A number of recent papers on learning in social networks relax the assumption that consumers observe the complete, ordered history of actions. Banerjee and Fudenberg [8] show that with a continuum of agents and uniform sampling of previous actions, under weak conditions complete learning results. With countable agents, Smith and Sorensen [35] examine general sampling rules, although they do not focus on the case that we are interested in, where consumers observe complete but unordered histories. (A precursor of that work, Smith [31], does examine that situation.) Further, their proofs do not seem to apply to our setting, where search is available. Celen and Kariv [16] assume that agents observe only their immediate predecessor's action, and show that beliefs and actions cycle indefinitely. Acemoglu, Dahleh, Lobel, and Ozdaglar [1] study an environment where each agent observes the choice of a random neighborhood of other agents, and provide conditions under which complete learning occurs when signals are unbounded. They also show that in some cases, complete learning occurs even with bounded signals, a result similar to ours that comes from a very different model.

There is a large literature on non-Bayesian observational learning on networks, including Ellison and Fudenberg [19, 20], Bala and Goyal [6], DeMarzo, Vayanos, and Zwiebel [17], and Golub and Jackson [22]. Acemoglu, Dahleh, Lobel, and Ozdaglar [1] provides a thoughtful discussion of this branch of the literature.

3 The Basic Model

An infinite sequence of consumers indexed by t enter in exogenous order. The choice set consists of J products. Each consumer has to decide on which products to search and, given the results of search, which products to buy. Prior to making these decisions, she observes private signals on each product and the history of decisions taken by previous consumers. We simplify the consumer's decision problem by assuming that her preferences across products are additive and the private signals she receives on the products are independently distributed. Given these assumptions, the consumer's search and purchasing decisions on a product are independent of what she knows about the other products and her decisions on these products. Thus, in what follows, we can drop the product index and focus on the learning dynamics of an individual product.

Consumer t 's utility for the product is given by

$$V_t = X + U_t$$

where X denotes the mean utility or quality of the product and U_t is consumer t 's idiosyncratic preference shock. Here U_t is identically and independently distributed across consumers and independently across products. Let F_U denote the distribution of U . Each product has two possible quality levels: $X = H$ and $X = L$, where $H > L$. We will refer to H as the high quality state and L as the low quality state. We normalize $L = 0$. Consumer t does not know X or U_t . There is a common prior belief that assigns a probability μ_0 to the event that $X = H$. (We interpret the prior as the belief that consumers have after observing all public information about the product's quality, such as media reviews and advertisements.) For convenience, we assume that both states are equally likely. The price of the product is p . Consumers' utility is quasilinear in wealth, so consumer t 's net payoff from purchasing the product is $V_t - p$.

Consumer t has three available actions. Buying a product involves risk since the ex post payoff may be negative. She can reduce the likelihood of this event by choosing to *Search* (S) before making her purchasing decision. Search involves paying a cost c to obtain a private, informative signal about V_t , and then purchasing only if the expectation of V_t conditional on the signal exceeds p . For notational simplicity, it will be convenient to assume that the signal is perfectly informative and reveals V_t precisely. Note that search remains a valuable option for consumer t even if she has learned X . The other actions that she can choose are

to *Not Search and Not Buy* (N) and *Buy Without Search* (B). Let $a_t \in \{B, N, S\}$ denote the action chosen by consumer t .

Given the consumer's purchasing rule following search, the expected value of search conditional on state X is

$$w(X) = \int_{p-X}^{\infty} (X - p + u) dF_U(u).$$

Hence, the payoff to a consumer from action S in state X is $w(X) - c$. To restrict attention to the economically interesting cases, we impose the following restrictions on the payoffs from search in each state.

A1: (i) $w(H) - c > 0$; (ii) $w(0) - c < 0$; (iii) $0 > H - p$.

Condition (i) of Assumption A1 states that, conditional on H , the payoff to search is positive. Condition (ii) states that the consumer's payoff to search is negative if she knows that the state is L . It implies that the consumer's optimal action in state L is N . Condition (iii) states that buying without search is dominated by not searching, even when consumers know that the state is H . It implies that consumers always choose either S or N . The motivation for imposing the dominance condition is that it allows us to better focus on the search decision. The implicit assumption is that the heterogeneity in consumer preferences is sufficiently important relative to search costs that most consumers do not want to buy a product without first checking it out, but search costs are high enough that they do not want to search a product that they know is low quality.

Consumer t 's action generates a purchasing outcome $b_t \in \{0, 1\}$. Here $b_t = 0$ is the outcome in which consumer t does not purchase the good, and $b_t = 1$ is the outcome in which consumer t purchases the product. Outcome 0 occurs if consumer t chooses N or if she chooses S and obtains a realization of V_t such that her net payoff from purchase is negative. Outcome 1 arises if consumer t chooses S and obtains a realization of V_t such that her net payoff from purchase is positive.

Before taking her action, consumer t observes a private signal about the quality of the product. (In a later section we consider the case where the private signal is informative about both X and U_t .) Smith and Sorensen [33] have shown that there is no loss in generality in defining the private signal, σ , that a consumer receives as her *private belief* that the state is H . That is, σ is the result of updating a half-half prior with the information in the private

signal. Conditional on the state, the signals are identically and independently distributed across consumers and drawn from a distribution F_X , $X = H, L$. We assume that F_L and F_H are continuous, mutually absolutely continuous, and differentiable with densities f_L and f_H . Under this assumption, the unconditional distribution of σ is $F = (F_L + F_H)/2$ with density f . The conditional signal densities are given by $f_L(\sigma) = 2(1 - \sigma)f(\sigma)$ and $f_H(\sigma) = 2\sigma f(\sigma)$.³ Smith and Sorensen (2000) show that defining the private signal in this way implies that the joint distribution of X and σ possesses the monotone likelihood ratio property. As is well known, this property implies that the conditional distributions F_L and F_H as well as their hazard and reverse hazard rates are ordered. Private beliefs are *bounded* if the convex hull of the common support of F_L and F_H consists of an interval $[\underline{\sigma}, \bar{\sigma}]$ where $\frac{1}{2} > \underline{\sigma} > 0$ and $\frac{1}{2} < \bar{\sigma} < 1$.

In addition to the private signal, consumer t also obtains information on the decisions by consumers 1 through $t - 1$. In the basic model, we assume that, if consumers observe the prior consumers' search or purchasing decisions, then they observe the entire ordered history. We consider three possible information structures. One structure is that consumer t observes only the search decisions of prior consumers; a second is that she observes only the purchasing decisions of prior consumers; and the third is that she observes both decisions. In the first two structures, the space of possible t -period search histories is given by $\Omega_t = \{0, 1\}^{t-1}$ and in the last structure, the space of t -period histories is $\Omega_t = \{0, 1\}^{t-1} \times \{0, 1\}^{t-1}$. A particular history is denoted by ω_t . The initial history is defined as $\omega_1 = \emptyset$. Note that the private signals are not informative about U . We relax this assumption in a later section. We also relax the strong assumption that consumers observe the entire ordered history ω . In most markets, consumers know only the aggregate number of consumers who have searched (e.g., online markets where sellers post number of visits to their product webpages) or purchased (e.g. offline markets which report product sales or sales rankings).

Given any history ω_t , consumer t updates her beliefs about X using Bayes' rule. Let $\mu_t(\omega_t)$ represent her posterior belief that the state is H conditional on history ω_t . Since ω_t is publicly observable, μ_t is also the *public* belief in period t . Given public belief μ_t and

³Thus, the ratio

$$\frac{f_L(\sigma)}{f_H(\sigma)} = \frac{1 - \sigma}{\sigma},$$

a property that we exploit repeatedly in the analysis.

private signal σ_t , consumer t 's *private* belief that the state is H is

$$r(\sigma_t, \mu_t) = \frac{\sigma_t \mu_t}{\sigma_t \mu_t + (1 - \sigma_t)(1 - \mu_t)}. \quad (1)$$

In studying the dynamics of beliefs and actions, we follow Smith and Sorensen [33] and work with the public likelihood ratio that the state is L versus H rather than public beliefs.

Define

$$l_t = \frac{1 - \mu_t}{\mu_t}.$$

and let l_0 denote the prior likelihood ratio. Using this transformation of variables in equation (1), consumer t 's private belief that the state is H becomes $r(\sigma_t, l_t)$. Her expected net payoff to S is

$$W(S; \sigma_t, l_t) = r(\sigma_t, l_t)w(H) + (1 - r(\sigma_t, l_t))w(0) - c. \quad (2)$$

Recall that L is normalized to zero. Therefore, if the consumer chooses N , her payoff is zero. We look for a Bayesian equilibrium where everyone computes posterior beliefs using Bayes' rule, knows the decision rules of all consumers and knows the probability laws determining outcomes under those rules.

A *cascade* on action $a \in \{S, N\}$ occurs when a consumer chooses a regardless of the realization of her private signal σ . Because of the distinction between actions and outcomes, we have to be careful in defining a herd. We say that a *herd* on action a occurs at time n if each consumer $t \geq n$ chooses action a . Note that while a cascade on N implies a herd on N , a cascade on S does not imply a herd on S if consumers observe purchases. The outcome for a consumer who chooses S depends not only on X (which is common across consumers) but also on the realization of the idiosyncratic component U . In fact, a herd on S precludes the event that all future outcomes are the same (almost surely) - if the outcome does not vary with the realization of U , then it is not worthwhile paying c to search.

A related concept is outcome convergence. Let $\lambda_t \in [0, 1]$ be the fraction of the first $t - 1$ consumers whose outcome was 1 (purchase). Outcome convergence is the event that λ_t converges to some limit $\lambda \in [0, 1]$. A herd implies outcome convergence. A herd on N leads to $\lambda = 0$; and a herd on S leads to $\lambda = 1 - F_U(p - X)$ in state H and $\lambda = 1 - F_U(p)$ in state L .

How does our model compare to other herding models? When consumers observe only the search history, our model is similar to the standard herding model in which actions are

observable but payoffs are not. Search is “good news” about the quality of the product because it implies that the consumer obtained a high private signal. When consumers observe the purchasing history, they observe binary signals about payoffs of the actions taken, which can lead to different market dynamics and outcomes. In particular, purchasing decisions provide information not only about the private signals that consumers have observed but also about their utility. Purchase is very good news because it implies that the consumer obtained a high private signal and a high realization of utility, both of which are more likely when the state is H . But when consumer t does not purchase the product and search is not observable, the news is mixed: subsequent consumers do not know whether it is because she had a low private signal and chose N or because she had a high private signal, chose S , and obtained a realization on V_t such that her net payoff from purchase was negative. When consumers observe both decisions, they obtain two binary signals about the state and are able to distinguish the positive news of search from the negative news of no purchase.

4 Equilibrium Dynamics and Outcomes

In this section we characterize the equilibrium dynamics and outcomes of the model. The main question that we are interested in exploring is the effect of the different information structures on long-run behavior and, in particular, the likelihood of “bad herds.” A bad herd occurs in state H when there is a herd on N and market share of the product goes to zero; in state L , a bad herd occurs when there is a herd on S and the market share of the product converges to $1 - F_U(p)$.

We begin by defining thresholds. Let \hat{r} represent the private belief at which a consumer is indifferent between S and N . From equation (2),

$$\hat{r} = \frac{c - w(0)}{w(H) - w(0)}. \quad (3)$$

Assumption A1 implies that $\hat{r} \in (0, 1)$. Using equations (1) and (3), we can then define the private signal at which a consumer is indifferent between S and N (assuming it is interior) as

$$\hat{\sigma}(l) = \frac{(c - w(0))l}{w(H) - c + (c - w(0))l}. \quad (4)$$

Thus, given l , the consumer’s optimal action is to choose S if $\sigma \geq \hat{\sigma}$ and to choose N if $\sigma < \hat{\sigma}$. We will refer to $\hat{\sigma}$ as the *search threshold*.

Next we define the cascade regions. Let \underline{l} denote the largest value of the public likelihood ratio such that a consumer is certain to choose S . Using equation (4), \underline{l} is defined as the solution to $\widehat{\sigma}(\underline{l}) = \underline{\sigma}$. Solving this equation for \underline{l} yields

$$\underline{l} = \frac{\underline{\sigma}(w(H) - c)}{(1 - \underline{\sigma})(c - w(0))}. \quad (5)$$

Let \bar{l} denote the lowest value of the public likelihood ratio such that a consumer is certain to choose N . Once again, using equation (4), \bar{l} satisfies $\widehat{\sigma}(\bar{l}) = \bar{\sigma}$. Solving this equation for \bar{l} yields

$$\bar{l} = \frac{\bar{\sigma}(w(H) - c)}{(1 - \bar{\sigma})(c - w(0))}. \quad (6)$$

Thus, we can partition the values of the public likelihood ratio into three intervals. When $l < \underline{l}$, there is a cascade on S ; when $\underline{l} \leq l \leq \bar{l}$, the consumer searches with probability $1 - F(\widehat{\sigma}(l))$ and does not search with probability $F(\widehat{\sigma}(l))$; and when $l > \bar{l}$, there is a cascade on N .

We now characterize the dynamics of the public likelihood ratio for the different models. We consider first the model in which consumers only observe the search history. Suppose $\underline{l} < l_t < \bar{l}$. Then the probability that consumer t searches the product in state X is

$$\Pr\{a_t = S | X, l_t\} = 1 - F_X(\widehat{\sigma}(l_t)).$$

Using Bayes' rule, the public likelihood ratio in period $t + 1$ is given by

$$l_{t+1}(a_t) = \begin{cases} \psi_S(l_t)l_t = \left[\frac{1 - F_L(\widehat{\sigma}(l_t))}{1 - F_H(\widehat{\sigma}(l_t))} \right] l_t & \text{if } a_t = S \\ \psi_N(l_t)l_t = \left[\frac{F_L(\widehat{\sigma}(l_t))}{F_H(\widehat{\sigma}(l_t))} \right] l_t & \text{if } a_t = N \end{cases} \quad (7)$$

When $l_t < \underline{l}$, $F_L(\widehat{\sigma}) = F_H(\widehat{\sigma}) = 0$ and $l_{t+1} = l_t$; when $l_t > \bar{l}$, $F_L(\widehat{\sigma}) = F_H(\widehat{\sigma}) = 1$ and $l_{t+1} = l_t$. The properties of the conditional distributions F_H and F_L imply the following Lemma.

Lemma 1 (i) ψ_N is continuous on (\underline{l}, ∞) with $\lim_{l \downarrow \underline{l}} \psi_N(l) = \frac{1 - \underline{\sigma}}{\underline{\sigma}} > 1$, $\psi_N(l) > 1$ on (\underline{l}, \bar{l}) and $\psi_N(l) = 1$ on $[\bar{l}, \infty)$; (ii) $\psi_S(l)$ is continuous on $[0, \bar{l})$ with $\lim_{l \uparrow \bar{l}} \psi_S(l) = \frac{1 - \bar{\sigma}}{\bar{\sigma}} < 1$, $\psi_S(l) = 1$ on $[0, \underline{l}]$, and $\psi_S(l) < 1$ on (\underline{l}, \bar{l}) ;

For our purposes, a particular useful property is that the posterior likelihood ratio following each action and/or outcome is monotone increasing in the prior likelihood ratio. Smith and Sorensen [34] find that their herding model has this feature if the density of the private belief log-likelihood ratio is log-concave. We impose this condition for our models as well. For any private belief σ , let γ denote the natural log of the corresponding likelihood ratio:

$$\gamma = \ln \left(\frac{1 - \sigma}{\sigma} \right).$$

Denote the unconditional distribution of γ by F^γ , with density f^γ .

A2: The density of the log likelihood ratio, f^γ , is strictly log-concave.

As Smith and Sorensen [34] note, all commonly-used continuous distribution have log-concave densities. Further, if the density of *private beliefs* is uniform, normal with variance at least twice the mean) or exponential with parameter below 2 (all with support in $(0, 1)$ and expectation 0.5), then straightforward calculation shows that the density of the *log-likelihood ratio* will be log-concave. Assumption A2 will be violated if there is a discrete set of possible private beliefs. Suppose, for example, that the distribution of private beliefs has a mass point at $\bar{\sigma}$. If consumer t has a public likelihood ratio l_t just above \bar{l} , then she will never search, and so $l_{t+1} = l_t$. If l_t is just below \bar{l} , on the other hand, then with probability bounded away from zero she will receive private belief $\bar{\sigma}$ and search. In that case, if she does not purchase, then the public likelihood ratio rises strictly above \bar{l} . A similar nonmonotonicity can result with continuous distributions that have big spikes in density, like a normal or exponential distribution with very low variance.

Lemma 2 *Assumption A2 implies that $\psi_N(l_t)l_t$ and $\psi_S(l_t)l_t$ are strictly increasing in l_t .*

Figure 1 illustrates the learning dynamics for a system that satisfies Assumption A2. Recall that the cascade set for S is the interval $[0, \underline{l}]$ and the cascade set for N is the interval $[\bar{l}, \infty)$. By Lemma 1, $\psi_N(l_t)l_t$ is continuous, lies above the diagonal on (\underline{l}, \bar{l}) , and intersects the diagonal at \bar{l} . Lemma 1 also implies that $\psi_S(l_t)l_t$ is continuous, is equal to the diagonal below \underline{l} , and lies below the diagonal between \underline{l} and \bar{l} . Lemma 2 then implies that both of these functions are strictly increasing on the interval (\underline{l}, \bar{l}) . Active dynamics occur only when the prior likelihood $l_0 \in (\underline{l}, \bar{l})$.

Next we consider the model in which consumers only observe the purchase history. Suppose $\underline{l} < l_t < \bar{l}$. Then the probability that consumer t buys the product in state X is

$$\Pr\{b_t = 1|X, l_t\} = (1 - F_X(\hat{\sigma}(l_t)))(1 - F_U(X - p)).$$

It is the probability that consumer t searches in state X times the probability that she gets a realization of U that lies above the purchasing threshold $p - X$. The probability that consumer t does not buy the product in state X is

$$\Pr\{b_t = 0|X, l_t\} = F_X(\hat{\sigma}(l_t)) + (1 - F_X(\hat{\sigma}(l_t)))F_U(p - X).$$

The first term is the probability of the event that consumer t does not search in state X ; the second term is the probability of the event that consumer t searches in state X and gets a value of U that lies below the purchasing threshold. Using Bayes' rule, the public likelihood ratio in period $t + 1$ is given by

$$l_{t+1}(b_t) = \begin{cases} \psi_1(l_t)l_t = \left[\frac{(1 - F_L(\hat{\sigma}(l_t)))(1 - F_U(p))}{(1 - F_H(\hat{\sigma}(l_t)))(1 - F_U(p - H))} \right] l_t & \text{if } b_t = 1 \\ \psi_0(l_t)l_t = \left[\frac{F_L(\hat{\sigma}(l_t)) + (1 - F_L(\hat{\sigma}(l_t)))F_U(p)}{F_H(\hat{\sigma}(l_t)) + (1 - F_H(\hat{\sigma}(l_t)))F_U(p - H)} \right] l_t & \text{if } b_t = 0 \end{cases} \quad (8)$$

In this case, when $l_t < \underline{l}$, $F_L(\hat{\sigma}) = F_H(\hat{\sigma}) = 0$, and the dynamic system reduces to a pair of *linear*, first-order difference equations. When $l_t > \bar{l}$, $F_L(\hat{\sigma}) = F_H(\hat{\sigma}) = 1$ and $l_{t+1} = l_t$. Lemma 3 follows from the properties of F_L and F_H .

Lemma 3 (i) ψ_0 is continuous, $\psi_0(l) > 1$ on $[0, \bar{l})$, and $\psi_0(l) = 1$ on $[\bar{l}, \infty)$; (ii) $\psi_1(l) < 1$ and continuous on $[0, \bar{l})$ with

$$\lim_{l \uparrow \bar{l}} \psi_1(l) = \left(\frac{1 - \bar{\sigma}}{\bar{\sigma}} \right) \left(\frac{1 - F_U(p)}{1 - F_U(p - H)} \right).$$

Log concavity is *not* sufficient to ensure that the model in which consumers observe only purchasing histories possesses the monotonicity property. We need to impose an additional restriction on the distribution of U [and σ] to ensure monotonicity.

A3: Either (i) the unconditional density of the private signal f is bounded above by

$$\left[\frac{F_U(p)F_U(p - H)}{1 - F_U(p - H)} \right]$$

or (ii)

$$\left[\frac{F_U(p)}{F_U(p-H)} \right] \leq \left[\frac{F_L(\sigma)}{F_H(\sigma)} \right] \Big/ \left[\frac{1-F_L(\sigma)}{1-F_H(\sigma)} \right] \quad \forall \sigma \in [\underline{\sigma}, \bar{\sigma}].$$

The first condition of Assumption A3 bounds the change in the probability of search that results from a fixed change in the public belief. The second condition ensures that, at any public belief, a decision not to search is more suggestive of state L than deciding to search but then not buying. Note that the right-hand side of the inequality is strictly greater than one.

Lemma 4 *Assumption A2 implies that $\psi_1(l_t)l_t$ is strictly increasing in l_t and Assumption A2 and A3 imply that $\psi_0(l_t)l_t$ is strictly increasing in l_t .*

Figure 2 illustrates a dynamic system that satisfies Assumptions A2 and A3. Applying Lemma 3, $\psi_0(l_t)l_t$ intersects the diagonal at 0 and at \bar{l} and lies everywhere above the diagonal between these two values. It is linear on the cascade set for S and, by Lemma 4, strictly increasing on the interval (\underline{l}, \bar{l}) . Lemma 3 states that $\psi_1(l_t)l_t$ intersects the diagonal at 0 and lies below the diagonal for $l_t \in (\underline{l}, \bar{l})$. Monotonicity follows from Lemma 4. In this model, active dynamics occur when the prior is such that $l_0 \in (0, \bar{l})$.

Finally, we consider the model in which consumers observe both histories. Suppose $\underline{l} < l_t < \bar{l}$. Then the public likelihood ratio in period $t+1$ is given by

$$l_{t+1}(a_t, b_t) = \begin{cases} \psi_1(l_t)l_t = \left[\frac{(1-F_L(\hat{\sigma}(l_t)))(1-F_U(p))}{(1-F_H(\hat{\sigma}(l_t)))(1-F_U(p-H))} \right] l_t & \text{if } a_t = S, b_t = 1 \\ \psi_{S0}(l_t)l_t = \left[\frac{(1-F_L(\hat{\sigma}(l_t)))F_U(p)}{(1-F_H(\hat{\sigma}(l_t)))F_U(p-H)} \right] l_t & \text{if } a_t = S, b_t = 0 \\ \psi_N(l_t)l_t = \left[\frac{F_L(\hat{\sigma}(l_t))}{F_H(\hat{\sigma}(l_t))} \right] l_t & \text{if } a_t = N, b_t = 0 \end{cases} \quad . \quad (9)$$

In this case, when $l_t < \underline{l}$, $F_L(\hat{\sigma}) = F_H(\hat{\sigma}) = 0$, and the dynamic system reduces to the same pair of *linear*, first-order difference equations as in the previous model. As in the other two models, when $l_t > \bar{l}$, $F_L(\hat{\sigma}) = F_H(\hat{\sigma}) = 1$ and $l_{t+1} = l_t$. The updating equation following no search is the same as in the model in which search is observable. The updating equation conditional on purchase is the same as in the model in which purchases are observable. The new feature is the updating equation when consumer observe search and no purchase.

Lemma 5 ψ_{S0} is continuous and strictly increasing on (\underline{l}, \bar{l}) , and

$$\lim_{l \downarrow \bar{l}} \psi_{S0}(l) = \left(\frac{1 - \bar{\sigma}}{\bar{\sigma}} \right) \left(\frac{F_U(p)}{F_U(p - H)} \right).$$

The model possesses the monotonicity property if Assumption A2 is satisfied. However, this condition is not sufficient to eliminate the possibility that the likelihood function can enter the cascade set for N because ψ_{S0} can lie everywhere above the diagonal. If this is the case, then the likelihood ratio can “jump” into the cascade set for N when l_t is sufficiently close to \bar{l} and consumer t searches but does not purchase. Condition (ii) of Assumption A3 rules out this situation. Figure 3 illustrates a dynamic system that satisfies Assumptions A2 and A3.

Lemma 3 of Smith and Sorensen [33] establishes the following:

Lemma 6 *Conditional on state H , the public likelihood ratio is a martingale. It converges to a random variable with support in $[0, \infty)$, so fully wrong learning has probability zero. Conditional on state L , the inverse public likelihood ratio is a martingale. It converges to a random variable with support in $[0, \infty)$.*

The martingale property rules out convergence to nonstationary limit beliefs such as cycles or to incorrect point beliefs.

We can adapt Smith and Sorensen’s [33] arguments to derive the following characterization of long-run behavior when consumers observe only the search history.

Proposition 7 *Suppose consumers observe only the search history, Assumption A2 holds, and $\underline{l} < l_0 < \bar{l}$. Then (a) outcome convergence occurs almost surely. In state H , $\lambda = 1 - F_U(p - H)$ with positive probability and $\lambda = 0$ with positive probability; in state L , $\lambda = 1 - F_U(p)$ with positive probability and $\lambda = 0$ with positive probability. (b) Limit beliefs in both states converge either to \underline{l} or \bar{l} and learning is incomplete. Beliefs can never enter the cascade set from outside. (c) Actions in state H converge almost surely to S when $\lambda = 1 - F_U(p - H)$; otherwise they converge almost surely to N ; actions in state L converge almost surely to S when $\lambda = 1 - F_U(p)$; otherwise they converge almost surely to N .*

The key feature of the dynamics of the model is that the public likelihood ratio never enters the cascade sets from outside. For any $l_0 \in (\underline{l}, \bar{l})$, the dynamics are forever trapped between the stationary points \underline{l} and \bar{l} and convergence is always asymptotic. We illustrate

this point in Figure 1. A sequence of “search” actions causes the likelihood ratio to decrease in ever smaller increments towards \underline{l} ; on the other hand, a sequence of “no search” actions causes the likelihood ratio to increase in ever smaller increments to \bar{l} . A herd always eventually starts. In both states, the herd forms with positive probability on either N or on S , and in neither event do consumers learn the state.

A similar argument leads to a characterization of long-run behavior when consumers observe only the purchasing history or the search and purchasing histories.

Proposition 8 *Suppose that $0 < l_0 < \bar{l}$. Further suppose either that consumers observe the purchasing history and Assumptions A2 and A3 hold, or that consumers observe the search and purchasing histories and Assumptions A2 and A3(ii) hold. (a) Outcome convergence occurs almost surely. In state H , $\lambda = 1 - F_U(p - H)$ with positive probability and $\lambda = 0$ with positive probability; in state L , $\lambda = 0$. (b) In state H , beliefs converge to the truth when $\lambda = 1 - F_U(p - H)$; otherwise, the limit belief converges to \bar{l} and learning is incomplete. Beliefs can never enter the cascade set for N from outside. (c) Actions in state H converge almost surely to S when $\lambda = 1 - F_U(p - H)$; otherwise they converge almost surely to N .*

For any $l_0 < \bar{l}$, the dynamics in this model are forever trapped between the stationary points 0 and \bar{l} . Figure 2 illustrates the dynamics when consumers observe only the purchasing history. A sequence of “buy” outcomes causes the likelihood ratio to decrease in ever smaller increments towards 0; a sequence of “no buy” outcomes causes the likelihood ratio to increase in ever smaller increments to \bar{l} . A herd always eventually starts. In state H , the herd can form with positive probability on either N or on S but, if it forms on S , then the market eventually learns the true state from the frequency of purchases. In state L , the herd can only form on N . Intuitively, if a herd formed on S , then the frequency of purchases (i.e., $1 - F_U(p)$ vs. $1 - F_U(p - H)$) reveals that the state is L , but by assumption it is not optimal to search in state L . Figure 3 shows how the dynamics change when consumers also observe the search history, but the long-run behavior is the same since the stationary points are the same.

We can now state the main result of this section.

Proposition 9 *Suppose that $0 < l_0 < \bar{l}$. (i) Suppose consumers observe only the search*

history and Assumption A2 hold. Then

$$\begin{aligned}\Pr\{l_\infty = \bar{l}|H\} &= \frac{l_0 - \underline{l}}{\bar{l} - \underline{l}} \\ \Pr\{l_\infty = \bar{l}|L\} &= \frac{\underline{l}^{-1} - l_0^{-1}}{\underline{l}^{-1} - \bar{l}^{-1}}.\end{aligned}$$

(ii) Suppose consumers observe only the purchasing history and Assumptions A2 and A3 hold or they observe the search and purchasing histories and Assumptions A2 and A3(ii) hold. Then

$$\begin{aligned}\Pr\{l_\infty = \bar{l}|L\} &= 1 \\ \Pr\{l_\infty = \bar{l}|H\} &= \frac{l_0}{\bar{l}}.\end{aligned}$$

Proposition 9 provides formulas for computing the probability of bad herds. The proof follows from the fact the the stochastic process for the likelihood ratio $\langle l_t \rangle$ is a bounded martingale conditional on state H (and the inverse likelihood ratio is a martingale conditional on state L). The Dominated Convergence Theorem implies $E[l_\infty|H] = l_0$. The restriction that the public likelihood ratio cannot enter the cascade sets from outside implies that the support of the distribution of l_∞ consists of only two points: $\{\underline{l}, \bar{l}\}$ when consumers observe only the search history and $\{0, \bar{l}\}$ when they observe the purchasing history. Combining these two results yields the formulas for bad herds.

The importance of Proposition 9 is that it permits an analysis of the determinants of bad herds. This kind of analysis has been absent in the literature on herding models. The main reason is that the literature typically considers models in which the likelihood ratio “jumps” into the cascade sets (e.g., with discrete signals). In these cases, the support of l_∞ consists of nondegenerate subsets of the cascade sets and convergence typically occurs in finite time. It is not difficult to construct examples of sequences in which the number of consumers who search or purchase is actually higher at higher prices or lower quality. Pastine and Pastine [27] obtain similar “perverse” results when they study the effect of changing the accuracy of signals on the probability of incorrect herds in a herding model where beliefs can enter the cascade set from outside. The intuition is simple: changes in model parameters cause changes in the purchasing and search rules, which in turn affect the informativeness of the decisions observed by consumers. Consumers take these changes into account when they update their beliefs about the state. These equilibrium effects make it

difficult, if not impossible, to predict the impact of parameter changes on the distribution of the number of searches or purchases. The martingale property can be used to bound the probability of a bad herd, but it is impossible to predict how this bound will vary with the parameters of the model.

5 Extensions

In this section we consider two extensions of the basic model. For convenience, we restrict the analysis to the model in which consumers observe only the purchasing history. In the first extension, we relax the assumption that signals are informative about the state but not about tastes. In practice, the signal that consumers obtain is likely to provide information about both components of utility. For example, when a consumer hears a song on the radio, he may recognize the song to be high quality but not like it because he does not like the genre. Similarly, when a consumer searches for product by submitting a query on a search engine, he observes not only the positions of advertisers on the page but also their text ad, which are informative signals about the idiosyncratic component of the match. In the second extension, we relax the assumption that consumers observe all past purchases. In most applications, consumers observe only aggregate statistics of the purchasing history such as total sales.

5.1 Private Signals of Utility

Suppose each consumer t gets a (noisy) private signal Z of his utility V_t for the product rather than a signal only of the product's quality X . In particular, let

$$Z_t = V_t + \varepsilon_t$$

where ε_t is independent of X and U_t and is i.i.d. across consumers. The distribution of ε is denoted F_ε . For simplicity, suppose that both F_ε and F_U have smooth densities, and that the support of the sum

$$\eta = u + \varepsilon$$

is the real line (so that no signal perfectly reveals the state). Let F_η denote the distribution of η ; it is the convolution of F_ε and F_U . The expected value of search conditional on state

X after signal z is observed is then defined as:

$$w(X, z) = \int_{p-X}^{\infty} (X - p + u) dF_{U|\eta=z-X}(u);$$

$F_{U|\eta=c}$ is the distribution of the taste component conditional on $\eta = c$:

$$F_{U|\eta=c} = \int_{-\infty}^{\infty} f_U(c - \varepsilon) f_{\varepsilon}(\varepsilon) d\varepsilon.$$

We need to impose restrictions on payoffs in order to focus on the interesting cases. In particular, we assume that (i) a consumer who knows that the product is high quality is willing to search the product when she has a high enough signal but never buys without searching and that (ii) a consumer who knows that the product is low quality always chooses N regardless of the realization of his signal. The assumption, which is analogous to Assumption A1 in our baseline case, is formalized as Assumption A1':

A1': (i) $w(H, z) - c > 0$ for high enough z and $H - p + E[U|X = H, z] < 0$ for all z ;
 (iii) $w(L, z) - c < 0$ and $0 > L - p + E[U|X = L, z]$ for all z .

We also need to order the distributions of the signal. We assume that the distribution of the private signal Z conditional on quality H first-order stochastically dominates the distribution conditional on L , so Z is informative about quality. Let $\sigma(z)$ denote the *belief component* of the signal. It is defined as

$$\sigma(z) = \frac{f_{\eta}(z - H)}{f_{\eta}(z) + f_{\eta}(z - H)},$$

the ratio of the density of z in state H to the total density of z . We assume in addition that the distributions of signals in the two states satisfy the monotone likelihood ratio property, so that $\sigma(z)$ is an increasing function, and that the support of $\sigma(z)$ is bounded:

A4: (i) $\frac{f_{\eta}(z - H)}{f_{\eta}(z)}$ is increasing in z ; and (ii) $\sigma(z) \in [\underline{\sigma}, \bar{\sigma}]$ for all $z \in \mathfrak{R}$.

With these assumptions in place, we can proceed to describe the dynamics. A consumer t who starts with public belief μ_t and observes private signal z_t updates his belief to

$$r(\sigma(z_t), l_t) = \frac{\sigma(z_t)}{\sigma(z_t) + (1 - \sigma(z_t))l_t}.$$

The consumer's expected net payoff from search is then

$$W(l_t, z_t) = r(\sigma(z_t), l_t)w(H, z_t) + (1 - r(\sigma(z_t), l_t))w(0, z_t) - c.$$

Note that $W(l, z)$ is decreasing in l and increasing in z , so for each public likelihood ratio l_t consumer t follows a cutoff strategy, choosing to search if z_t exceeds a threshold \hat{z} that is increasing in l_t . Let

$$\alpha(X, z) = \frac{\int_{u=p-X}^{\infty} [1 - F_\varepsilon(z - X - u)] f_U(u) du}{1 - F_\eta(z - X)}$$

denote the probability in state X that a consumer's utility exceeds the price p conditional on the event that his signal exceeds s . From Bayes' rule, then, the updated public likelihood ratio in period $t + 1$ after a purchase by consumer t ($b_t = 1$) or no purchase ($b_t = 0$) is, respectively,

$$l_{t+1}(l_t) = \begin{cases} \varphi_1(l_t)l_t & \text{if } b_t = 1 \\ \varphi_0(l_t)l_t & \text{if } b_t = 0 \end{cases}$$

where the functions φ_0 and φ_1 are defined as

$$\begin{aligned} \varphi_0(l_t) &= \frac{\alpha(0, \hat{z}(l_t))[1 - F_\eta(\hat{z}(l_t))]}{\alpha(H, \hat{z}(l_t))[1 - F_\eta(\hat{z}(l_t) - H)]} \\ \varphi_1(l_t) &= \frac{[1 - \alpha(0, \hat{z}(l_t))][1 - F_\eta(\hat{z}(l_t))] + F_\eta(\hat{z}(l_t))}{[1 - \alpha(H, \hat{z}(l_t))][1 - F_\eta(\hat{z}(l_t) - H)] + F_\eta(\hat{z}(l_t) - H)}. \end{aligned}$$

As in the baseline case, the public likelihood ratio is a martingale, conditional on H . The fixed points of the dynamics are $l_t = 0$ and $l_t \geq \bar{l}$ where here \bar{l} is the largest value of l_t such that $\hat{z}(l_t)$ is finite. It follows from Assumption A1'(ii) and A4 that \bar{l} is finite. In fact, we can obtain a closed form solution for \bar{l} using arguments similar to those in the basic model:

$$\bar{l} = \frac{\bar{\sigma}(\hat{w}(H) - c)}{(1 - \bar{\sigma})(c - \hat{w}(0))} \quad (10)$$

where $\hat{w}(X) = \lim_{z \rightarrow \infty} w(X, z)$.

As in the baseline case, if consumer t purchases the product, future consumers know that he received a signal whose value was above his cutoff, and that he learned that his utility exceeded the price p . If he does not purchase, then future consumers know that either his signal was below the cutoff, or his utility was less than p . The complication relative to the baseline case is that, given quality, the probability of purchase after search is not

independent of the private signal (since the signal is now correlated with the idiosyncratic taste component U). The correlation introduces a selection effect into the model: consumers who search are more likely to like the product and buy it.

The correlation also leads to less-than-clean sufficient conditions for the monotonicity property that is needed for the comparative statics. We need to establish that both $\varphi_0(l_t)l_t$ and $\varphi_1(l_t)l_t$ are increasing in l_t . As an analog to Assumption A3, we assume that the density of the log of the likelihood ratio $(1 - \sigma(z))/\sigma(z)$ is log-concave and that either (i) not searching is a more negative signal than searching but not buying, or (ii) the density of the belief component $\sigma(z)$ is bounded above.

A3': (i) The density of the log of the likelihood ratio, $(1 - \sigma(z))/\sigma(z)$, is strictly log-concave; and (ii)

$$\frac{1 - \alpha(0, z)}{1 - \alpha(H, z)} \leq \left[\frac{F_\eta(z)}{F_\eta(z - H)} \right] \bigg/ \left[\frac{1 - F_\eta(z)}{1 - F_\eta(z - H)} \right] \text{ for all } z \in \mathfrak{R}.$$

To cope with the complications created by the dependence of the probability of purchase after search on the search probability, we introduce an additional assumption:

$$\text{A5: } \frac{1 - \alpha(0, z)}{1 - \alpha(H, z)} \geq \frac{\frac{\partial \alpha(0, z)}{\partial z}}{\frac{\partial \alpha(H, z)}{\partial z}} \geq \frac{\alpha(0, z)}{\alpha(H, z)}.$$

Assumption A5 has the flavor of a monotone likelihood ratio property for the probability of purchase conditional on search in the two states.

Lemma 10 *Assumptions A3' and A5 imply that $\varphi_0(l_t)l_t$ and $\varphi_1(l_t)l_t$ are increasing in l_t .*

Having established the monotonicity of the dynamics of the public likelihood, we can extend the results of the baseline case concerning long-run behavior to the case of private signals of both common and idiosyncratic components.

Proposition 11 *If Assumptions A3', A4 and A5 hold, then l_t converges to \bar{l} with probability 1 in state L and with probability l_0/\bar{l} in state H.*

The proof of Proposition 11 proceeds along the same lines as the proofs of Propositions 8 and 9, and therefore is omitted. This concludes our analysis of the model in which consumers get private signals of utility and observe past purchases. A similar analysis applies when consumers observe only searches or both searches and purchases.

5.2 Aggregate Histories

Suppose each consumer t observes only her place in the sequence of consumers (that is, t) and the aggregate number of previous purchases, denoted n_t . In other words, consumer t observes the unordered history of purchase decisions rather than, as in the baseline case, the ordered history ω_t . Otherwise the model is the same as before. We want to examine whether the long-run behavior of the model when only the aggregate history is observed is the same as when the purchasing history is observed.

The main difficulty in analyzing a model in which aggregate histories are observed is that, in general, consumer t does not know what history consumer $t - 1$ observed. In our model, she knows only that either $n_{t-1} = n_t - 1$ and consumer $t - 1$ searched and purchased or $n_{t-1} = n_t$ and consumer $t - 1$ did not purchase the product. (The exceptions are when either all previous consumers purchased ($n_t = t$) or none of them did ($n_t = 0$)). One consequence is that the player t 's belief after observing the period- t aggregate history but before receiving her private signal can no longer be called a public belief; instead, we follow Acemoglu, Dhaleh, Lobel, and Ozdaglar [1] and call it the *social belief* because it is based on the decisions of other consumers. A more important consequence is that the sequence of likelihood ratios no longer form a martingale, and some of the proof techniques used in the baseline case no longer work. (See for example, Acemoglu, Dhaleh, Lobel, and Ozdaglar [1] and Smith and Sorensen [35].) The difficult new step is to show that the social likelihood ratios converge almost surely either to 0 or to \bar{l} , the same two possible limits as in the baseline case. The rest of the analysis is roughly the same as before.

We begin by introducing some new notation. Define $\pi_t(n, X)$ as the probability that $n_t = n$ in state X ; $l_t(n)$ as consumer t 's social likelihood ratio (probability of state L divided by the probability of state H) after observing $n_t = n$; and $\beta_t(n, X)$ as the probability that consumer t purchases in state X after observing $n_t = n$. We recursively define $\pi_t(n, X)$ as follows. First, let

$$\begin{aligned}\pi_1(0, L) &= \pi_1(0, H) = 1, \\ l_1(0) &= l_0, \\ \beta_1(0, X) &= [1 - F_U(p - X)][1 - F_X(\hat{\sigma}(l_1(0))).\end{aligned}$$

That is, consumer 1 necessarily observes zero previous purchases and so her belief equals the prior, and her probability of purchase is the product of the probability of getting a high

enough private signal to induce search and the probability of purchase after search. Next, given $\pi_{t-1}(n_{t-1}, X)$, $l_{t-1}(n_{t-1})$, and $\beta_{t-1}(n_{t-1}, X)$, define

$$\pi_t(n, X) = \begin{cases} \pi_{t-1}(0, X)(1 - \beta_{t-1}(0, X)) & \text{if } n = 0 \\ \pi_{t-1}(n, X)[1 - \beta_{t-1}(n, X)] + \pi_{t-1}(n-1, X)\beta_{t-1}(n-1, X) & \text{if } 1 \leq n \leq t-1 \\ \pi_{t-1}(t-1, X)\beta_{t-1}(t-1, X) & \text{if } n = t \end{cases}$$

$$l_t(n) = \frac{\pi_t(n, L)}{\pi_t(n, H)} l_0,$$

$$\beta_t(n, X) = [1 - F_U(p - X)][1 - F_X(\hat{\sigma}(l_t(n)))].$$

Using the definition of $\pi_t(n, X)$, we can rewrite the social likelihood ratio $l_t(n)$ for non-extreme values of n (above 0 and less than t) after some manipulation as

$$l_t(n) = \frac{\pi_{t-1}(n, L)[1 - \beta_{t-1}(n, L)] + \pi_{t-1}(n-1, L)\beta_{t-1}(n-1, L)}{\pi_{t-1}(n, H)[1 - \beta_{t-1}(n, H)] + \pi_{t-1}(n-1, H)\beta_{t-1}(n-1, H)} l_0 \quad (11)$$

$$\begin{aligned} &= l_{t-1}(n) \frac{[1 - \beta_{t-1}(n, L)]}{[1 - \beta_{t-1}(n, H)]} \\ &\quad \left(\frac{\pi_{t-1}(n, L)[1 - \beta_{t-1}(n, L)]}{\pi_{t-1}(n, H)[1 - \beta_{t-1}(n, H)] + \pi_{t-1}(n-1, H)\beta_{t-1}(n-1, H)} \right) \\ &\quad + l_{t-1}(n-1) \frac{\beta_{t-1}(n, L)}{\beta_{t-1}(n, H)} \\ &\quad \left(\frac{\pi_{t-1}(n, L)\beta_{t-1}(n, L)}{\pi_{t-1}(n, H)[1 - \beta_{t-1}(n, H)] + \pi_{t-1}(n-1, H)\beta_{t-1}(n-1, H)} \right) \\ &= \psi_1(l_{t-1}(n))l_{t-1}(n) \left(\frac{\pi_{t-1}(n, L)[1 - \beta_{t-1}(n, L)]}{\pi_{t-1}(n, H)[1 - \beta_{t-1}(n, H)] + \pi_{t-1}(n-1, H)\beta_{t-1}(n-1, H)} \right) \\ &\quad \psi_0(l_{t-1}(n-1))l_{t-1}(n-1) \\ &\quad \left(\frac{\pi_{t-1}(n, L)\beta_{t-1}(n, L)}{\pi_{t-1}(n, H)[1 - \beta_{t-1}(n, H)] + \pi_{t-1}(n-1, H)\beta_{t-1}(n-1, H)} \right) \end{aligned} \quad (12)$$

Thus, the social likelihood ratio is a weighted average of the ratios that consumer t would have if she knew that $n_{t-1} = n_t$ and consumer $t-1$ did not purchase, and if she knew that $n_{t-1} = n_t - 1$ and consumer $t-1$ did purchase. (The weights are the relative probabilities of those events in state H .) An immediate consequence is that social likelihood ratio cannot jump into the cascade sets if Assumptions A2 and A3 are satisfied.

Lemma 12 *Assumptions A2 and A3 imply that if $l_0 \in (0, \bar{l})$, then $l_t(n) \in (0, \bar{l})$ for all t and all $n \leq t$.*

Thus, social likelihood ratios stay between 0 and \bar{l} , as long as the prior l_0 lies in that range. The next two lemmas demonstrate that \bar{l} is a stable fixed point of the social likelihood ratio and that l_t cannot converge to any value other than 0 and \bar{l} .

Lemma 13 *Assumption A2 and A3 imply that if $l_t \leq \bar{l}$ is within ε of \bar{l} infinitely often for all $\varepsilon > 0$, then l_t converges to \bar{l} .*

Lemma 14 *Assumption A2 and A3 imply that if $l_0 \in (0, \bar{l})$, then l_t converges almost surely, either to 0 or to \bar{l} .*

The proof of Lemma 13 is novel and of independent interest so we include a sketch of the proof in the text. Suppose that almost surely l_t eventually has support equal to the interval $[l^*, l^{**}]$, with $0 < l^* \leq l^{**} < \bar{l}$. Then for T large enough, the probability that $l_t \in [l^*, l^{**}]$ for all $t > T$ is close to 1. Then some consumer t with positive probability observes n^{**} such that $l_t(n^{**})$ is very close to l^{**} , and any lower n would result in $l_t(n) > l^{**}$. If consumer t does not purchase (an event that has probability at least $F_U(p - X)$), then consumer $t + 1$ also observes n^{**} . Since she assigns very high probability to the event $\{n_t = n^{**}\}$ rather than $\{n_t = n^{**} - 1\}$, her social likelihood ratio $l_{t+1}(n^{**})$ is very close to

$$\psi_0(l_t(n^{**}))l_t(n^{**}) \approx \psi_0(l^{**})l^{**}.$$

Since $l^{**} < \bar{l}$, $\psi_0(l^{**})$ is strictly greater than 1, and so $l_{t+1}(n^{**}) > l^{**}$. Thus, $[l^*, l^{**}]$ cannot be the eventual support of l_t . That argument can be generalized to show that unless l_t converges to 0, it must approach arbitrarily close to \bar{l} infinitely often.

To extend the results of the base line model on long-run behavior, we need to show that the probabilities of the two limits in the two different states are unchanged.

Proposition 15 *If Assumptions A2 and A3 hold, then l_t converges to \bar{l} with probability 1 in state L and with probability l_0/\bar{l} in state H.*

This concludes our analysis of the model in which consumers observe only aggregate sales. A similar analysis applies when consumers observe only aggregate searches.

6 Comparative Statics

Our analysis yields two interesting results on long-run behavior. First, as long as the likelihood ratio cannot enter the cascade sets from outside, long-run behavior is determined

by two boundary points on the cascade sets. The learning dynamics can differ depending upon the type of information provided or the way in which it is provided but, if the cascades sets are the same, then they possess the same long-run behavior. For example, the learning model in which only purchasing information is observed has different dynamics depending upon whether consumers observe the entire history or only aggregate sales or whether they also observe the search history, but the probability of a bad herd is the same.

Second, the cascade sets can depend upon the kind of information that is provided. If search actions are observed but purchasing outcomes are not, then boundary point on the cascade set for S is bounded away from 0, as a result, the market never learns the true quality of products. Bad herds can arise in both states. By contrast, when consumers observe purchasing decisions, the cascade set on S is degenerate at 0. In this case, the market can learn the true quality of high quality products and the market shares of low quality products are certain to converge to zero. The purchasing information does not eliminate the bad herd on high quality products - they can still fail - but it does eliminate the bad herd on low quality products. This result is driven in part by our assumption that consumers never purchase without buying. If we relax Assumptiopn A1(iii), then the cascade sets are on N and B (there is no cascade set on S) and the boundary points on the cascade set for B is bounded away from 0. In this case, complete learning does not occur in state H and a herd on B can form with positive probability in both states.

The implication of our analysis is that sellers who want to influence market outcomes should focus on factors that determine the boundary points of the cascade sets (e.g., prices and search costs) rather than factors that affect the learning dynamics (e.g., precision of the private signal). We analysis the determinants of bad herds when consumers observe information on purchasing decisions of prior consumers. The analysis for the case where consumers observe information only about the search history is similar and yields similar results.

Differentiating \bar{l} with respect to H , c and p yields the following results:

Proposition 16 *Suppose the state is H . Then the probability of a limit cascade on N is (a) strictly decreasing in H ; (b) increasing and convex in c ; (c) increasing in p .*

The proposition yields several predictions. An increase in search costs increases the likelihood that long-run sales of a high quality product is zero and hence reduces its expected long-run sales. The impact of the increase is larger at higher cost levels. An increase in price

also increases the probability of zero long-run sales but the sign of the second derivative of \bar{l} with respect to p depends upon the distribution of the private signal. This issue is important because, if \bar{l} is convex, then the model could explain why search goods like music albums and books sell for the same prices even when the distributors know that the quality of their products vary.

Rosen [28] argued that the reward function to quality is convex because, in equilibrium, more talented artists can sell more units at higher unit prices. Our results suggest another source of convexity: in the long-run, consumers are more likely to learn about higher quality products and are more likely to buy them. Thus, small differences in product quality can lead to large differences in expected sales even when prices do not vary with quality. Our model can also explain why prices of products like albums, books and videos do not vary with quality. A small increase in price can have a disproportionate effect on expected sales since it decreases market share *and* increases the probability of a “bad” herd.

A number of papers (e.g., [13]) have argued that the decline in search costs due to the Internet has disproportionately increased sales of niche products and reduced the concentration of sales. The next proposition provides support for this claim.

Proposition 17 *Suppose the state is H . Then the impact of an increase in c on the probability of a limit cascade on N is smaller (in absolute value) for higher quality products.*

The results follows from differentiating \bar{l} with respect to H and c . Proposition 17 implies that a decrease in search costs (due for example to Internet technologies) has a larger impact on long-run sales of niche products (i.e., medium quality products) than on high quality products.

A striking feature of the model is that long-run behavior does not depend upon the precision of the private signal about product quality or tastes. The only property of F that matters for the probability of a bad herd (aside from log concavity) is $\bar{\sigma}$, the upper bound of its support. As $\bar{\sigma}$ goes to 1, \bar{l} approaches infinity and the probability of a bad herd goes to zero. Thus, in our model, unbounded beliefs lead to complete learning. (See Smith and Sorensen [33].)

Acemoglu, Dhaleh, Lobel, and Ozdaglar [2] investigate the issue of whether heterogeneous preferences make social learning more or less likely to occur. They show in the context of their model that heterogeneity makes social learning more likely. We obtain a similar result in our model but for different reasons.

Proposition 18 *Suppose the state is H . Then a mean-preserving spread in F_U reduces the probability of a limit cascade on N .*

The argument is straightforward. A mean-preserving spread in the distribution of tastes increases the value of search in both states. As a result, \bar{l} increases and the probability of limit cascade on N falls.

When sales converge to zero, a seller has an incentive to invest in a signal (e.g., advertisements) that can change public beliefs and lead the market to correct its mistake.

Proposition 19 *The introduction of a sufficiently positive, public signal after beliefs have converged (a) leads to an increase in short-run sales and, with positive probability, a herd on S if the initial herd is on N ; (b) has no effect on sales if the initial herd is on S .*

7 Empirical Evidence

Like most models of herd behavior, our model describes a simple form of observational learning in which consumers draw inferences from search or purchase decisions of other consumers. Recent field experiments have confirmed the empirical relevance of this kind of learning (e.g., Cai, Chen, and Fang [14] and Tucker and Zhang [36]), and the fact that this kind of observational learning can lead to herd behavior has been documented in laboratory experiments (e.g., Anderson and Holt [5] and Çelen and Kariv [15]). To highlight the relevance of our particular model, in this section we briefly describe two recent studies whose results relate directly to our model's predictions.

Hendricks and Sorensen [23] use detailed album sales data for 355 music artists who released at least two albums to document that many albums flop when they are first released, but then later succeed if the artist releases another album that is a hit. While a change in preferences could explain this effect, the authors show that the spillover more likely results from learning: the success of the new album generates new information that leads many consumers to learn about the previous album and consider buying it. In other words, some debut albums suffer from a “bad herd,” but a new hit album can cause consumer beliefs about the debut album to change enough that many consumers will consider buying it. If enough of these consumers buy the album, it can attract enough attention to make the market for the album re-converge to a higher level of sales. This result is consistent with Proposition 18, which states that a sufficiently positive public signal can “overturn” a bad

herd. Using estimates of the discovery probability, they compute the “lost” sales of debut albums from consumers not knowing about the album during its release period. They find that the difference between counterfactual and actual sales is largest for the moderately successful artists. Most consumers know about the albums of the top artists and, even though most consumers are unaware of the sub-par albums, these albums’ qualities are sufficiently low that sales would be minimal even if everyone were fully informed. This finding is consistent with Proposition 16 and implies that sales would have been substantially less skewed if consumers were more fully informed about the albums.

A clever online experiment conducted by Salganik, Dodds, and Watts [30] (hereafter SDW) generated data that are directly informative about the dynamics of learning. In the experiment, thousands of subjects were recruited to participate in artificial online music markets. Participants arrived sequentially and were presented with a list of 48 songs, which they could listen to, rate, and then download (for free) if they so chose. In real time, each participant was randomly assigned to one of nine “worlds.” In the treatment worlds, of which there were eight, songs were listed by download rank: the first song listed was the one with the most downloads by previous participants in that same world, the second song listed had the second most downloads, and so on. The listing also included the total number of downloads of each song. In the control world, the 48 songs were shown in a random order, with no information about previous participants’ listening, rating, or downloading behavior. The eight treatment worlds operated independently of one another, so that the researchers could observe eight separate realizations of the stochastic process.⁴

The central finding of the SDW study is that listening and downloading probabilities were more variable in the treatment worlds than in the control worlds. In the control world, the fraction of participants who listened to a given song was roughly equal across songs, ranging from 6% to 11%. In the treatment worlds, listening probabilities varied widely, with some songs in some worlds being listened to by more than 40% of participants. The higher variance of listening probabilities naturally translated to a higher variance in downloading probabilities. Hence, downloads were much more “skewed” in the treatment worlds than in the control world. Downloads were also more unpredictable in the treatment worlds. As SDW put it, “the best songs rarely did poorly, and the worst rarely did well, but any other result was possible.”

⁴See Salganik [29] for a much more thorough description of the methods and results of the experiment.

The data from the SDW experiment can be used to examine some of the specific predictions of our model, because the design of the experiment matches the assumptions of our model remarkably well. As in our model, the products in these experiments were search goods. The songs were carefully screened to ensure that they would be unknown to the participants.⁵ Choosing whether to sample a song is analogous to the decision of whether to search in our model, and downloading a song (after listening to it) is analogous to the purchase decision. Since participants assigned to treatment worlds were shown the aggregate number of downloads by previous participants, the information they received is essentially the same as in our model with consumers observing aggregate purchases. The cost of search in the experiment (i.e., the opportunity cost of the time spent listening to a song) was apparently large enough to matter: Table 1, which summarizes the behavior of participants in the experiment, shows that most participants listened to very few songs. On average, participants listened to fewer than four songs and downloaded fewer than two. The median number of listens was 1 and the median number of downloads was zero. It was not possible in the experiment to download a song without first listening to it, just as in our model a consumer cannot buy the product without first searching.

Recall that our single-product model is a suitable representation of demand for these songs if participants' preferences were additive, because in that case the model would apply individually to each song. The fact that some participants listened to and downloaded several different songs suggests that additivity is at least plausibly correct. We do not test this assumption, however, so it should be viewed as a caveat to the discussion below.

An important assumption that we *can* test is that the treatment effects in the SDW study reflect effects on consumers' information, not their preferences. In our model, consumers are influenced by others' purchases only insofar as those purchases affect the decision to search. There are no social effects in the traditional sense: preferences are unaffected by previous consumers' purchases. In the SDW experiment, this means that download information shown in the treatment worlds should influence participants' listening decisions, but not their decisions about whether to download a song conditional on listening to it. (SDW refer to the probability that a participant downloaded a song conditional on listening to it as the song's "batting average.") In fact this appears to be the case. Participants in the treatment

⁵They were obtained from the music website *purevolume.com*, a website where aspiring bands can create homepages and post music for download. Bands that had played too many concerts or received too many hits on their homepages were excluded.

worlds were roughly 10 times more likely to listen to the top-ranked (i.e., most downloaded) song than any song ranked below 30, but high download ranks did not appear to increase songs’ batting averages.

More formally, letting L_{jt} be an indicator variable equal to 1 if participant t listened to song j , and D_{jt} an indicator equal to 1 if participant t downloaded song j (conditional on listening to it), we can ask whether L_{jt} and D_{jt} are influenced by song j ’s download share among participants $1, \dots, t - 1$. Table 2 reports the results from probit regressions in which L_{jt} and D_{jt} are assumed to depend on song j ’s current download share (i.e., song j ’s share of total downloads by previous participants). Because some song titles and/or artist names might be more appealing than others on average, we include the song’s listening share from the independent world as a control in the listening regression. In the listening probit regression, the coefficient on download share is positive and highly significant, indicating that participants’ decisions to listen to a song were influenced by information about previous participants’ downloads.⁶

By contrast, conditional download probabilities did *not* appear to be higher for top-ranked songs. This is a trickier issue, however, because a probit regression of D_{jt} on download share involves an obvious reflection problem. (Songs with the most downloads will naturally have higher average download probabilities.) Fortunately, the conditional download probability from the control world is a natural control variable. These conditional probabilities provide the best measure of the songs’ relative qualities.⁷ When included, it forces the coefficient on download share to be identified from time variation in the download

⁶In interpreting this result, one might argue that it partially reflects a framing effect. Because people generally tend to choose the first item when selecting from a list, the apparent influence of download information could be conflated with the impact of list position itself. Indeed, even in the control world, where songs were ordered randomly for each participant, participants were much more likely to listen to the first listed song. However, the download information can be shown to have an effect above and beyond the effect that comes from list position. SDW ran separate experiments in which download information was provided but songs were still randomly ordered; in these experiments, the provided information still had a substantial impact on listening probabilities (albeit not as large as in the experiments we analyze here).

⁷Unconditional download probabilities do not accurately reflect song quality because they conflate the probability of listening (which was highly variable across songs, and across worlds for a given song) with the conditional probability of downloading. On the other hand, conditional on listening to a song, the probability of downloading is clearly higher for songs with greater appeal. Measured this way, song quality varied substantially across songs: the conditional download probability was nearly 60% for the highest-quality song, and only 11% for the lowest-quality song.

share relative to what it “ought” to be (as indicated by its download probability in the treatment world). Estimates of this model are reported in the second column of Table 2. The coefficient on download share is actually negative, suggesting that participants were slightly less likely to download top-ranked songs (conditional on listening to them). Taken together, the estimates imply that the information provided in the treatment worlds primarily affected participants’ listening decisions, not their “preferences” (or at least not positively).

The negative impact of downloads on the download probability is an intriguing result. It is likely due to a selection effect. In the control world, participants who chose to listen to a particular song may have done so because something in the song’s title appealed to them. In other words, the title is an informative signal about the idiosyncratic component of their preferences and as a result, they are more likely to download the song than a randomly selected participant would be. When download information is shown, this selection effect is not as strong. Top-ranked songs are listened to by a wider selection of participants, not just those who liked the titles. Since the selection of listeners for top-ranked songs is less favorably inclined, the fraction who choose to download these songs after listening to them is lower. Conversely, lower-ranked songs are listened to by a narrower selection of participants, those who really liked the title. Since the selection of listeners for these songs are more favorably inclined, the fraction who choose to download is higher. Thus, the results reported in Table 2 accord well with the assumptions of our extended model in which the private signal is informative about the idiosyncratic component of preferences as well as the common component. We suspect that the selection effect will be present in any realistic model of choice with private signals and costly search.

The central predictions of our model relate to the long-run outcomes of the sequential decision-making process. To describe the long-run outcomes in the SDW experiment, we analyze the behavior of the last 100 participants in each of the eight treatment worlds. Table 3 reports the distribution of listening probabilities among these last 100 participants for three groups of songs: high, medium, and low quality (as indicated by their batting averages in the control world). An important result in our theoretical analysis is that when consumers observe purchase histories, bad herds on low-quality products are eliminated—i.e., if the state is L , then a herd on S (search) is not possible. This appears to hold in the SDW experiment: listening probabilities for the lowest-quality songs almost always converged to rates lower than 5%, with only 3% of songs obtaining listening probabilities

above 10%.⁸ For high-quality products, our model predicts that bad herds *are* possible: a herd may form on S , in which case the product’s market share converges to its “true” share, or a herd may form on N , in which case the market never learns that the product is good. The results shown in Table 3 appear to be consistent with this prediction. Outcomes for the highest quality songs were highly unpredictable. In 53% of the cases, the listening probabilities were less than 5%, but in 13% of the cases they exceeded 20%.

A fundamental question about observational learning is whether it increases consumer welfare (in expectation). Although the stochastic nature of the learning process implies that the market can sometimes converge to the “wrong” outcome, in general one might expect the provision of information about previous consumer’s decisions to make search more efficient. To test this in the experimental data, we compare the listen rates and download rates for three different groups of participants: (1) those who were randomly assigned to the control world; (2) those who were assigned to a treatment world and were among the *first* 100 participants to arrive; (3) those who were assigned to a treatment world and were among the *last* 100 participants to arrive. Listening and downloading behavior in group 2 should be similar to group 1, since not much information has yet accumulated. However, we should expect search to be noticeably more efficient for the third group, since they observe substantial information on previous participants’ downloads, and that information should in most cases guide them toward higher quality songs.

As a crude way of gauging the efficiency of participants’ searches, Table 4 reports the average of the number of listens and the average of the number of downloads of participants in each of the three groups. The listening rates and download rates for the first 100 participants look roughly similar to those in the control world. But the last 100 participants have lower listening rates and substantially higher download rates, increasing from .316 to .372. This difference is statistically significant ($p=0.022$). It suggests that search indeed became more efficient over time in the treatment worlds. The impact of information provision on the efficiency of search is an important issue in the study of directed search (e.g., internet search engines); we plan to explore the issue more fully in future research.

⁸Recall that roughly 2% of the participants had essentially zero search costs and listened to all 48 songs. Consequently, in contrast to our model, we would not expect to see listening probabilities converge to zero in the experiment.

8 Conclusion

We have studied a simple choice problem in which consumers have to decide whether or not to consider a product of unknown utility. Consumers only purchase products in their consideration sets, but including a product in the set involves a small search cost. Consumers would prefer not to pay this cost if they believe they are unlikely to buy the product. The purchasing decisions of other consumers influence their beliefs about the gains from search. A poor purchasing record can feed on itself and lead consumers to wrongfully omit high quality products from their consideration sets. On the other hand, a good purchasing record can also feed on itself and lead consumers to include the product in their consideration sets. In this case, the market learns the quality of the product and the product obtains its “true” market share. The experimental study by SDW provides evidence on the feedback mechanism and shows how it can affect outcomes. The results are largely consistent with our model.

In a multiple choice environment, the assumption that preferences are additive across products is quite strong. In many markets, consumers may only want one product. Substitution effects implies that the learning dynamics have to be studied at the level of the choice set and not on a product-by-product basis. We hope to explore the impact of the interactions in learning across products on long-run behavior in subsequent research.

We have largely ignored the issue of how firms may want to increase the likelihood that consumers will include their products in their consideration sets. Eliaz and Spiegler [18] explore this issue in a static model in which firms employ costly marketing devices such as advertising to influence the formation of consumer consideration sets. They use the model to study whether firms can profitably exploit the bounded rationality of consumers and the impact of their rational behavior on product variety and marketing. More generally, we think advertising plays an important role in directing consumer search and believe that exploring how firms try to use this information channel to influence purchase decisions is a promising direction for future research.

References

- [1] Acemoglu, D., Dahleh, M., Lobel, I. and Ozdaglar, A. “Bayesian Learning in Social Networks,” 2009, working paper.
- [2] Acemoglu, D., Dahleh, M., Lobel, I. and Ozdaglar, A. “Heterogeneity and Social Networks,” 2008, slides for MIT Theory talk.
- [3] Adler, M. “Stardom and Talent,” *American Economic Review*, 75, 1985, pp. 208-12.
- [4] Ali, N. “Learning Self-Control,” 2009, working paper.
- [5] Anderson, L. and Holt, C. “Information Cascades in the Laboratory,” *American Economic Review* 87, 1997, pp. 847-862.
- [6] Bala V. and Goyal, S. “Conformism and Diversity Under Social Learning,” *Economic Theory*, 17, 2001, pp. 101-120.
- [7] Banerjee, A. “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, 107, 1992, pp. 797-817.
- [8] Banerjee, A. and Fudenberg, D. “Word-of-Mouth Learning,” *Games and Economic Behavior*, 46, 2004, pp. 1-22.
- [9] Becker, G., Grossman, M., and Murphy, K. “An Empirical Analysis of Cigarette Addiction,” *American Economic Review*, 84, 1994, pp. 396-418.
- [10] Becker, G., and Murphy, K. *Social Economics: Market Behavior in a Social Environment*, Harvard University Press, 2000.
- [11] Bikhchandani, S., Hirshleifer, D. and Welch, I. “A Theory of Fads, Fashions, Custom, and Cultural Change as Informational Cascades,” *Journal of Political Economy*, 100, 1992, pp. 992-1026.
- [12] Brock, W. and Durlauf, S. “Interactions-Based Models,” in *Handbook of Econometrics*, vol. 5., J. Heckman and E. Leamer eds., Amsterdam: North-Holland, 2001.
- [13] Brynjolfsson, E., Hu, Y., and Simester, D. “Goodbye Pareto Principle, Hello Long Tail: The Effect of Search Costs on the Concentration of Product Sales,” 2006, working paper.

- [14] Cai, H., Chen, Y., and Fang, H. "Observational Learning: Evidence from a Randomized Natural Field Experiment," *American Economic Review*, 99, 2009, pp. 864-882.
- [15] Çelen, B. and Kariv, S. "Distinguishing Informational Cascades from Herd Behavior in the Laboratory," *American Economic Review*, 94, 2004, pp. 484-498.
- [16] Çelen, B. and Kariv, S. "Observational Learning Under Imperfect Information," *Games and Economic Behavior*, 47, 2004, pp. 72-86.
- [17] Demarzo, P.M., Vayanos, D., and Zwiebel, J. "Persuasion Bias, Social Influence, and Unidimensional Opinions," *Quarterly Journal of Economics*, 118, 2003, pp. 909-968.
- [18] Eliaz, K. and Spiegler, R. "Consideration Sets and Competitive Marketing," 2008, working paper.
- [19] Ellison, G. and Fudenberg, D. "Rules of Thumb for Social Learning," *Journal of Political Economy*, 101, 1993, pp. 612-643.
- [20] Ellison, G. and Fudenberg, D. "Word-of-mouth Communication and Social Learning," *Quarterly Journal of Economics*, 110, 1995, pp. 93-126.
- [21] Goeree, J., Palfrey, T., and Rogers, B. "Social Learning with Private and Common Values," *Economic Theory*, 28, 2006, pp. 245-264.
- [22] Golub, B. and Jackson, M. "Naive Learning in Social Networks: Convergence, Influence, and the Wisdom of Crowds," *American Economic Journal: Microeconomics*, forthcoming.
- [23] Hendricks, K. and Sorensen, A. "Information and Skewness of Music Sales," *Journal of Political Economy*, 117, 2009, pp. 324-369.
- [24] Manski, C. "Identification of Exogenous Social Effects: The Reflection Problem," *Review of Economic Studies*, 1993, 60, pp. 31-42.
- [25] Nelson, P. "Information and Consumer Behavior," *Journal of Political Economy*, 78, 1970, pp. 311-329.
- [26] Orphanides, A. and Zervos, D. "Rational Addiction with Learning and Regret," *Journal of Political Economy*, 103, 1995, pp. 739-758.

- [27] Pastine, I. and T. Pastine, “Signal Accuracy and Informational Cascades” CEPR Discussion Paper Series #5219, 2006.
- [28] Rosen, S., “The Economics of Superstars,” *American Economic Review*, 71, 1981, pp. 845-58.
- [29] Salganik, M., “Success and Failure in Cultural Markets,” Ph.D. thesis, 2007, Department of Sociology, Columbia University, New York.
- [30] Salganik, M., Dodds, P., and Watts, D. “Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market,” *Science*, 311, 2006, pp. 854-856.
- [31] Smith, L. “Error Persistence, and Experimental versus Observational Learning,” 1991, working paper.
- [32] Sorensen, A., “Bestseller Lists and Product Variety: The Case of Book Sales,” forthcoming, *Journal of Industrial Economics*.
- [33] Smith, L. and Sorensen, P. “Pathological Outcomes of Observational Learning,” *Econometrica*, 68, 2000, pp. 371-398.
- [34] Smith, L. and Sorensen, P. “Informational Herding and Optimal Experimentation,” 2008, working paper.
- [35] Smith, L. and Sorensen, P. “Rational Learning with Random Sampling,” 2008, working paper.
- [36] Tucker, C. and Zhang, J. “Long Tail or Steep Tail? A Field Investigation into How Online Popularity Affects the Distribution of Consumer Choices”, 2007, working paper, .
- [37] Varian, H. “A Model of Sales,” *American Economic Review*, 70, 1980, pp. 651-659.
- [38] Wiseman, T. “Disagreement Leads to Complete Learning: Sequential Choice with Continuous Types,” *Economics Letters*, 100, 2008, pp. 53-55.

Appendix

Proof of Lemma 1: The continuity of ψ_N and ψ_S follows from the continuity of F_L , F_H , and $\hat{\sigma}$. Since F_L and F_H satisfy the monotone likelihood ratio property, F_H first order stochastically dominates F_L , so $\psi_N(l) > 1$ and $\psi_S(l) < 1$ for $l \in (l, \bar{l})$. For $l \geq \bar{l}$, $F_H(\hat{\sigma}(l)) = F_L(\hat{\sigma}(l)) = 1$, and so $\psi_N(l) = 1$. Finally,

$$\begin{aligned} \lim_{l \uparrow \bar{l}} \psi_S(l) &= \lim_{l \uparrow \bar{l}} \left[\frac{1 - F_L(\hat{\sigma}(l))}{1 - F_H(\hat{\sigma}(l))} \right] \\ &= \lim_{l \uparrow \bar{l}} \left[\frac{f_L(\hat{\sigma}(l))\hat{\sigma}'(l)}{f_H(\hat{\sigma}(l))\hat{\sigma}'(l)} \right] \\ &= \frac{f_L(\hat{\sigma}(\bar{l}))}{f_H(\hat{\sigma}(\bar{l}))} \end{aligned}$$

The last equality follows from the fact that

$$\hat{\sigma}'(l) = \frac{(w(H) - c)(c - w(0))}{[w(H) - c + ((c - w(0))l)]^2} \neq 0.$$

Recall that $f_L(\sigma) = 2(1 - \sigma)f$ and $f_H(\sigma) = 2\sigma f$. Therefore,

$$\begin{aligned} f_L(\hat{\sigma}(\bar{l}))/f_H(\hat{\sigma}(\bar{l})) &= \frac{(1 - \hat{\sigma}(\bar{l}))}{\hat{\sigma}(\bar{l})} \\ &= \frac{(1 - \bar{\sigma})}{\bar{\sigma}} \end{aligned}$$

A similar argument establishes that

$$\begin{aligned} \lim_{l \downarrow l} \psi_N(l) &= \lim_{l \downarrow l} \left[\frac{F_L(\hat{\sigma}(l))}{F_H(\hat{\sigma}(l))} \right] \\ &= \lim_{l \downarrow l} \left[\frac{f_L(\hat{\sigma}(l))\hat{\sigma}'(l)}{f_H(\hat{\sigma}(l))\hat{\sigma}'(l)} \right] \\ &= \frac{f_L(\hat{\sigma}(\bar{l}))}{f_H(\hat{\sigma}(\bar{l}))} \\ &= \frac{1 - \bar{\sigma}}{\bar{\sigma}} \end{aligned}$$

Q.E.D.

Proof of Lemma 2: Smith and Sorensen's [34] Lemma 6 establishes that if f^γ is log-concave, then both

$$\left[\frac{F_L(\hat{\sigma}(l_t))}{F_H(\hat{\sigma}(l_t))} \right] l_t$$

and

$$\left[\frac{1 - F_L(\widehat{\sigma}(l_t))}{1 - F_H(\widehat{\sigma}(l_t))} \right] l_t$$

are increasing in l_t . (That is, the posterior likelihood ratio after observing only whether or not consumer t searches increases with l_t .) Q.E.D.

Proof of Lemma 3: The continuity of ψ_0 and ψ_1 follows from the continuity of F_L , F_H , and $\widehat{\sigma}$. Since F_L and F_H satisfy the monotone likelihood ratio property, F_H first order stochastically dominates F_L , so $\psi_0(l) > 1$ and $\psi_1(l) < 1$ for $l \in [0, \bar{l}]$. For $l \geq \bar{l}$, $F_H(\widehat{\sigma}(l)) = F_L(\widehat{\sigma}(l)) = 1$, and so $\psi_0(l) = 1$. Finally,

$$\begin{aligned} \lim_{l \uparrow \bar{l}} \psi_1(l) &= \frac{1 - F_U(p)}{1 - F_U(p - H)} \lim_{l \uparrow \bar{l}} \left[\frac{1 - F_L(\widehat{\sigma}(l))}{1 - F_H(\widehat{\sigma}(l))} \right] \\ &= \left(\frac{1 - F_U(p)}{1 - F_U(p - H)} \right) \left(\frac{1 - \bar{\sigma}}{\bar{\sigma}} \right) \end{aligned}$$

where the last equality follows from Lemma 1. Q.E.D.

Proof of Lemma 4: The posterior likelihood ratio after observing $b_t = 1$ is

$$\psi_1(l_t)l_t = \frac{1 - F_U(p)}{1 - F_U(p - H)} \left[\frac{1 - F_L(\widehat{\sigma}(l_t))}{1 - F_H(\widehat{\sigma}(l_t))} \right] l_t$$

Thus, $\psi_1(l_t)l_t$ is a scalar multiple of $\psi_S(l_t)l_t$ which is increasing in l_t by Lemma 2.

The posterior likelihood ratio after observing $b_t = 0$ is

$$\psi_0(l_t)l_t = \left[\frac{F_L(\widehat{\sigma}(l_t)) + (1 - F_L(\widehat{\sigma}(l_t)))F_U(p)}{F_H(\widehat{\sigma}(l_t)) + (1 - F_H(\widehat{\sigma}(l_t)))F_U(p - H)} \right] l_t. \quad (\text{A1})$$

First, suppose that the second part of Assumption A3 holds. Defining

$$\rho(l_t) = \frac{F_L(\widehat{\sigma}(l_t))}{F_H(\widehat{\sigma}(l_t)) + (1 - F_H(\widehat{\sigma}(l_t)))F_U(p - H)} \in (0, 1)$$

as the probability in state H that consumer t did not search conditional on his not purchasing, we can rewrite equation (A1) as

$$\psi_0(l_t)l_t = \rho(l_t) \left[\frac{F_L(\widehat{\sigma}(l_t))l_t}{F_H(\widehat{\sigma}(l_t))} \right] + (1 - \rho(l_t)) \left[\frac{(1 - F_L(\widehat{\sigma}(l_t)))F_U(p)l_t}{(1 - F_H(\widehat{\sigma}(l_t)))F_U(p - H)} \right]. \quad (\text{A2})$$

We know that the terms in brackets are increasing in l_t . Thus, if we can show that

$$(i) \quad \frac{F_L(\widehat{\sigma}(l_t))}{F_H(\widehat{\sigma}(l_t))} \geq \frac{(1 - F_L(\widehat{\sigma}(l_t)))F_U(p)}{(1 - F_H(\widehat{\sigma}(l_t)))F_U(p - H)}$$

and that (ii) $\rho(l_t)$ is increasing in l_t , then we have established that $\psi_0(l_t)l_t$ is increasing in l_t . The second part of Assumption A3 implies claim (i). To check claim (ii), differentiating $\rho(l_t)$ at $l_t \geq \underline{l}$ and simplifying yields

$$\begin{aligned} \rho'(l_t) &= \frac{F_U(p-H)f_H(\widehat{\sigma}(l_t))\widehat{\sigma}'(l_t)}{[F_H(\widehat{\sigma}(l_t)) + [1 - F_H(\widehat{\sigma}(l_t))]F_U(p-H)]^2} \\ &> 0 \end{aligned}$$

because

$$\widehat{\sigma}'(l_t) = \frac{(w(H) - c)(c - w(L))}{[w(H) - c + (c - w(L))l_t]^2} > 0.$$

At $l_t < \underline{l}$ (so that the probability of search is one), the value of $\rho(l_t)$ is constant at 0 (and thus weakly increasing).

Next suppose that the first part of Assumption A3 holds. For notational simplicity, define

$$\begin{aligned} a &= 1 - F_U(p), \\ b &= 1 - F_U(p - H) \end{aligned}$$

and

$$\widehat{l}(l_t) = \frac{1 - \widehat{\sigma}(l_t)}{\widehat{\sigma}(l_t)},$$

so equation (A1) can be rewritten as

$$\psi_0(l_t)l_t = \left[\frac{1 - a + aF_L(\widehat{\sigma}(l_t))}{1 - b + bF_U(\widehat{\sigma}(l_t))} \right] l_t. \quad (\text{A3})$$

Differentiating equation (A3) with respect to $l_t \geq \underline{l}$,

$$\begin{aligned} &\left[\frac{1 - a + aF_L(\widehat{\sigma}(l_t))}{1 - b + bF_U(\widehat{\sigma}(l_t))} \right] + \\ &l_t \left[\frac{1 - b + bF_H(\widehat{\sigma}(l_t))af_L(\widehat{\sigma}(l_t))\widehat{\sigma}'(l_t) - [1 - a + aF_L(\widehat{\sigma}(l_t))]bf_H(\widehat{\sigma}(l_t))\widehat{\sigma}'(l_t)}{[1 - b + bF_H(\widehat{\sigma}(l_t))]^2} \right] \\ &= \frac{1}{[1 - b + bF_H(\widehat{\sigma}(l_t))]^2} [1 - a + aF_L(\widehat{\sigma}(l_t))][1 - b + F_H(\widehat{\sigma}(l_t))] \\ &\quad + \widehat{\sigma}'(l_t)l_t [(1 - b + bF_H(\widehat{\sigma}(l_t)))af_L(\widehat{\sigma}(l_t)) - (1 - a + aF_L(\widehat{\sigma}(l_t)))bf_H(\widehat{\sigma}(l_t))]. \end{aligned}$$

Substituting $f_L(\sigma) = 2(1 - \sigma)f(\sigma)$ and $f_H(\sigma) = 2\sigma f(\sigma)$, and

$$\widehat{\sigma}'(l_t) = \frac{\widehat{l}(l_t)}{[\widehat{l}(l_t) + l_t]^2}$$

then yields

$$\begin{aligned}
&= \frac{1}{[1 - b + bF_H(\hat{\sigma}(l_t))]^2} [1 - a + aF_L(\hat{\sigma}(l_t))] [1 - b + F_H(\hat{\sigma}(l_t))] \\
&\quad + \frac{\hat{l}(l_t)}{[\hat{l}(l_t) + l_t]^2} 2f(\hat{\sigma}(l_t)) [(1 - b + bF_H(\hat{\sigma}(l_t)))a(1 - \hat{\sigma}(l_t)) - (1 - a + aF_L(\hat{\sigma}(l_t)))b\hat{\sigma}(l_t)] \\
&\geq \frac{1}{[1 - b + bF_H(\hat{\sigma}(l_t))]^2} [(1 - a)(1 - b) - bf(\hat{\sigma}(l_t))].
\end{aligned}$$

The first part of Assumption A3 ensures that the last expression is positive. If $l_t < \underline{l}$, then $\hat{\sigma}'(l_t) = 0$, and the derivative in equation (A3) reduces to

$$\frac{1 - a + aF_L(\hat{\sigma}(l_t))}{1 - b + bF_H(\hat{\sigma}(l_t))} > 0.$$

Q.E.D.

Proof of Lemma 5: By definition,

$$\psi_{S0}(l_t) = \left(\frac{1 - F_U(p)}{1 - F_U(p - H)} \right) \psi_S(l_t).$$

Since ψ_{S0} is proportional to ψ_S , it inherits the properties of ψ_S established in Lemmas 1 and 2. Q.E.D.

Proof of Proposition 7: Lemma 6 shows that the public likelihood ratio converges almost surely to a random variable l_∞ . The only fixed points of the Markov process on the public likelihood ratio (and thus the only possible values of l_∞) are $l \leq \underline{l}$ and $l \geq \bar{l}$. Given belief convergence and the monotonicity of the Markov process (Lemma 2), Smith and Sorensen's [34] Lemma 12a establishes that the public likelihood ratio cannot enter the cascade sets $l \leq \underline{l}$ and $l \geq \bar{l}$ from outside, and so if $\underline{l} < l_0 < \bar{l}$, then either $l_\infty = \underline{l}$ or $l_\infty = \bar{l}$. In state H , the argument of Smith and Sorensen's [33] Theorem 1d shows that the events $\{l_\infty = \underline{l}\}$ and $\{l_\infty = \bar{l}\}$ both have positive probability. In the former case, actions converge to S , and $\lambda = 1 - F(p - H)$; in the latter case, actions converge to N , and $\lambda = 0$. A similar argument establishes that the events $\{l_\infty = \underline{l}\}$ and $\{l_\infty = \bar{l}\}$ have positive probability in state L . In the former case, actions converge to S and $\lambda = 1 - F(p - H)$ and in the latter case, actions converge to N and $\lambda = 0$. Learning is always incomplete. Q.E.D.

Proof of Proposition 8: Lemma 6 shows that the public likelihood ratio converges almost surely to a random variable l_∞ . The only fixed points of the Markov process on the public

likelihood ratio (and thus the only possible values of l_∞) are $l = 0$ and $l \geq \bar{l}$. Given belief convergence and the monotonicity of the Markov process (Lemma 4), Smith and Sorensen's [34] Lemma 12a establishes that the public likelihood ratio cannot enter the cascade set $l \geq \bar{l}$ from outside. (In the case where consumers observe both searches and purchases, the same argument applies, because Assumption A3(ii) implies that $\psi_{S0}(l_t) \leq \psi_N(l_t)$.) Thus, if $l_0 < \bar{l}$, then either $l_\infty = 0$ or $l_\infty = \bar{l}$. Because l_∞ almost surely cannot be fully wrong, in state L , $l_\infty = \bar{l}$ with probability one. Consequently, in state L actions converge to N and so $\lambda = 0$. In state H , the argument of Smith and Sorensen's [33] Theorem 1d shows that the events $\{l_\infty = 0\}$ and $\{l_\infty = \bar{l}\}$ both have positive probability. In the former case, learning is complete, actions converge to S , and $\lambda = 1 - F_U(p - H)$. In the latter case, learning is incomplete, actions converge to N , and $\lambda = 0$. Q.E.D.

Proof of Proposition 9: (i) In state H , $\langle l_t \rangle$ is a martingale bounded below by \underline{l} and above by \bar{l} . The Dominated Convergence Theorem then implies that $E[l_\infty | H] = l_0$. Since either $l_\infty = \underline{l}$ or $l_\infty = \bar{l}$, we have $l_0 = \Pr\{l_\infty = \bar{l} | H\}\bar{l} + (1 - \Pr\{l_\infty = \bar{l} | H\})\underline{l}$, and the result follows. In state L , $\langle 1/l_t \rangle$ is a martingale bounded above by $1/\underline{l}$ and below by $1/\bar{l}$, and so a similar argument yields the result. (ii) Now, in state H , $\langle l_t \rangle$ is a martingale bounded below by 0 and above by \bar{l} , and so $\Pr\{l_\infty = \bar{l} | H\} = l_0/\bar{l}$. In state L , since with probability one either $l_\infty = \underline{l}$ or $l_\infty = \bar{l}$, Lemma 6 implies that $l_\infty = \underline{l}$ almost surely. (Note that $\langle 1/l_t \rangle$ is not bounded above in this case, and so the Dominated Convergence Theorem does not apply.) Q.E.D.

Proof of Lemma 10: First, consider $\varphi_1(l_t)l_t$. As in the proof of Lemma 4, Assumption A3' and Smith and Sorensen's [33] Lemma 6 imply that

$$\frac{1 - F_\eta(\hat{z}(l_t))}{1 - F_\eta(\hat{z}(l_t) - H)} l_t$$

is increasing in l_t . Thus, it is sufficient to show in addition that the ratio

$$\frac{\alpha(L, \hat{z}(l_t))}{\alpha(H, \hat{z}(l_t))}$$

is increasing in l_t ; Assumption A5 implies that it is. Similarly, Assumption A5 implies that

$$\frac{1 - \alpha(L, \hat{z}(l_t))}{1 - \alpha(H, \hat{z}(l_t))}$$

is increasing in l_t , so Lemma 3's proof that $\varphi_0(l_t)l_t$ is increasing in l_t goes through here as well. Q.E.D.

Proof of Lemma 12: The proof is inductive. Since $l_1(0) = l_0$, the conclusion holds trivially for $t = 1$. Next, suppose that the conclusion holds for $t - 1$. By equation (11), $l_t(n)$ is a weighted average of $\psi_{11}(l_{t-1}(n))l_{t-1}(n)$ and $\psi_0(l_{t-1}(n-1))l_{t-1}(n-1)$. Lemma 3 ensures that both values lie in $(0, \bar{l})$. Q.E.D.

Proof of Lemma 13: In either state X , the probability of the transition from n to $n + 1$, $\beta_t(n, X)$, shrinks continuously to 0 as $l_t(n)$ approaches \bar{l} . Further, the slope of $l_{t+1}(n)$ with respect to $l_t(n)$,

$$\frac{l_{t+1}(n) - l_{t+1}(n')}{l_t(n) - l_t(n')} < 1,$$

because (i) $l_{t+1}(n) > l_t(n)$ for $l_t(n) < \bar{l}$ (no purchase is bad news) and (ii) $l_{t+1}(n) = l_t(n)$ if $l_t(n) = \bar{l}$ (\bar{l} is a fixed point). Therefore, although l_t is not a martingale, the argument of Smith and Sorensen's [33] "Rest of Proof of Theorem 4" (page 397) applies directly: \bar{l} is a stable fixed point. Q.E.D.

Proof of Lemma 14: Lemma 12 establishes that l_t stays in the set $(0, \bar{l})$. Define the random variable Y_t as

$$Y_T \equiv \sup\{l_t : t \geq T\},$$

and the random variable Y as

$$Y \equiv \lim_{T \rightarrow \infty} \sup Y_T.$$

(That is, l_t exceeds Y only a finite number of times, but is arbitrarily close to Y infinitely often. If l_t converges to c , then $Y = c$. If l_t does not converge, then Y is the upper bound of its eventual support.) Let y^* be the upper bound of the support of Y conditional on l_t not converging to \bar{l} . (If l_t converges to \bar{l} almost surely, then y^* is not defined, but in that case the claim of the lemma is satisfied.) Lemma 11 implies that y^* must be strictly less than \bar{l} : if l_t approaches \bar{l} arbitrarily closely infinitely often, then l_t converges to \bar{l} .

If $y^* = 0$, then we have established the desired result. To finish the proof, we suppose that $y^* > 0$ and derive a contradiction. In the event that $Y = y^*$, for any $\varepsilon > 0$ and $T > 0$, there must almost surely be a period $t > T$ at which $l_t > y^* - \varepsilon$, and for any $l > y^*$ there must be a last period $\tau(l)$ in which the social likelihood ratio exceeds l . Since $\lim_{t \rightarrow \infty} \Pr\{\tau > t | Y = y^*\} = 0$, then for any $\varepsilon > 0$ and $T > 0$ there exists $t^* > T$ and

$n^*(t^*)$ such that

- (i) consumer t^* observes $n_{t^*} = n^*(t^*)$ with positive probability
- (ii) $l_{t^*}(n^*(t^*)) \in [y^* - \varepsilon, y^*]$,
- (iii) $l_{t^*}(n^*(t^*) - 1) > y^*$, and
- (iv) $\frac{\pi_{t^*}(n^*(t^*), X)}{\pi_{t^*}(n^*(t^*) - 1, X)} \geq 1 - \varepsilon$ for either state X

That is, $n^*(t^*)$ is the smallest number of purchases that consumer t can observe without being more pessimistic than is likely when Y_T is converging to y^* . For large enough t , $\tau(l_{t^*}(n^*(t^*) - 1))$ is very likely less than t conditional on l_t being less than y^* . The aggregate history $n^*(t^*)$ has positive probability when $Y = y^*$, but $n^*(t^*) - 1$ is very unlikely for any realization of Y .

If consumer t^* observes $n_{t^*} = n^*(t)$, then with probability bounded below by $F_U(p - X)$ she does not purchase, and consumer $t^* + 1$ observes $n_{t^*+1} = n^*(t)$. By condition (iv) above, consumer $t^* + 1$ assigns probability near 1 to consumer t^* 's having observed $n^*(t)$ rather than $n^*(t) - 1$, and so by equation (12), her social likelihood ratio $l_{t^*+1}(n^*(t^*))$ will be close to

$$l_{t^*}(n^*(t^*)) \frac{[1 - \beta_{t^*}(n^*(t^*), L)]}{[1 - \beta_{t^*}(n^*(t^*), H)]}.$$

Because $l_t \leq y^* < \bar{l}$,

$$\frac{[1 - \beta_{t^*}(n^*(t^*), L)]}{[1 - \beta_{t^*}(n^*(t^*), H)]} > 1,$$

and so $l_{t^*+1}(n^*(t^*))$ is strictly greater than y^* . Thus $y^* > 0$ is not possible. Q.E.D.

Proof of Proposition 15: In state L , consumer's beliefs cannot converge to something completely wrong with positive probability, and so l_t must converge almost surely to \bar{l} . Suppose next that the state is H . For any $t > 0$, the expectation at time 0 of l_t is equal to the prior, l_0 since

$$\begin{aligned} E_0[l_t(n)] &= \sum_{n=0}^t l_t(n) \pi_t(n, H) \\ &= \sum_{n=0}^t \left[\frac{\pi_t(n, L)}{\pi_t(n, H)} l_0 \right] \pi_t(n, H) \\ &= \sum_{n=0}^t \pi_t(n, L) l_0 \\ &= l_0. \end{aligned}$$

Since l_t almost surely converges either to 0 or to \bar{l} , convergence to \bar{l} must have probability l_0/\bar{l} . Q.E.D.

Proof of Propositions 16 and 17: A limit cascade on N is the event that $l_\infty = \bar{l}$, which has probability

$$\Pr\{l_\infty = \bar{l}\} = \frac{l_0}{\bar{l}} = \frac{(c - w(0))(1 - \underline{\sigma})}{(w(H) - c)\underline{\sigma}} l_0 = l_0 \frac{(1 - \underline{\sigma}) \left(c - \int_p^\infty (-p + u) dF_U(u) \right)}{\underline{\sigma} \left(\int_{p-H}^\infty (H - p + u) dF_U(u) - c \right)}.$$

Differentiating with respect to H , c , and p yields, respectively,

$$\frac{\partial \Pr\{l_\infty = \bar{l}\}}{\partial H} = \frac{-(1 - F_U(p - H))(c - w(0))(1 - \underline{\sigma})}{(w(H) - c)^2 \underline{\sigma}} l_0 < 0,$$

$$\frac{\partial \Pr\{l_\infty = \bar{l}\}}{\partial c} = \frac{(w(H) - w(0))(1 - \underline{\sigma})}{(w(H) - c)^2 \underline{\sigma}} l_0 > 0,$$

and

$$\frac{\partial \Pr\{l_\infty = \bar{l}\}}{\partial p} = \frac{[(1 - F_U(p))(w(H) - c) + (1 - F_U(p - H))(c - w(0))](1 - \underline{\sigma})}{(w(H) - c)^2 \underline{\sigma}} l_0 > 0.$$

Further differentiating $\frac{\partial \Pr\{l_\infty = \bar{l}\}}{\partial c}$ with respect to c and H yields, respectively,

$$\frac{\partial^2 \Pr\{l_\infty = \bar{l}\}}{(\partial c)^2} = \frac{2(w(H) - c)(w(H) - w(0))(1 - \underline{\sigma})}{(w(H) - c)^3 \underline{\sigma}} l_0 > 0$$

and

$$\frac{\partial \Pr\{l_\infty = \bar{l}\}}{\partial c \partial H} = \frac{-(1 - F_U(p - H))[(c - w(0)) + (w(H) - w(0))](1 - \underline{\sigma})}{(w(H) - c)^3 \underline{\sigma}} l_0 < 0.$$

Q.E.D.


Proof of Proposition 18: As in the previous proof, a limit cascade on N has probability

$$\Pr\{l_\infty = \bar{l}\} = \frac{(c - w(0))(1 - \underline{\sigma})}{(w(H) - c)\underline{\sigma}} l_0.$$

The expected payoff to search in state X , $w(X)$, can be written as $w(X) = E_U[\max\{0, X - p + u\}]$. Since the max function is convex, a mean-preserving spread in F_U increases $w(X)$. Since $\Pr\{l_\infty = \bar{l}\}$ is decreasing in both $w(0)$ and $w(H)$, the result follows. Q.E.D.


Proof of Proposition 19: If the initial herd is on S (corresponding to $l_\infty = 0$), then subsequent buyers will choose action S (and then purchase with probability $1 - F_U(p - H)$) with or without the positive public signal. If the initial herd is on N (corresponding to $l_\infty = \bar{l}$ and $\lambda = 0$), then a positive public signal pushes the public likelihood ratio below \bar{l} . The public likelihood ratio will reconverge, to \bar{l} in state L but with positive probability to 0 (and positive long-run sales) in state H . Thus, in state H , a positive public signal raises expected sales when the initial herd is on N . Q.E.D.

Figure 1: Only search history is visible




The figure shows the dynamics of the public likelihood ratio when consumers can only observe the sequence of prior consumers' search decisions. The constant k is equal to $\frac{w(H)-c}{c-w(0)}$.

Figure 2: Only purchase history is visible



The figure shows the dynamics of the public likelihood ratio when consumers can only observe the sequence of prior consumers' purchase decisions. The constant k is equal to $\frac{w(H)-c}{c-w(0)}$.

Figure 3: Both search and purchase histories are visible



The figure shows the dynamics of the public likelihood ratio when consumers can observe the sequence of prior consumers' search *and* purchase decisions. The constant k is equal to $\frac{w(H)-c}{c-w(0)}$.

Table 1: Summary statistics: participants

	Treatment Worlds ($N = 5,746$)	Control World ($N = 1,446$)	Overall ($N = 7,192$)
Number of listens:			
Mean	3.52	3.90	3.60
Std. Dev.	7.13	8.38	7.40
Min	0	0	0
Median	1	1	1
Max	48	48	48
Number of downloads:			
Mean	1.41	1.51	1.43
Std. Dev.	4.35	5.05	4.49
Min	0	0	0
Median	0	0	0
Max	48	48	48

Table 2: Probit regressions

	(1)	(2)
	Listen	Download
	N = 274,512	N = 20,091
Download share	1.266 (.013)	-0.192 (.074)
Control world listening share	0.681 (.035)	
Control world cond. prob. download		0.754 (.039)

The table reports marginal effects from probit regressions, with standard errors in parentheses.

Table 3: Stochastic mapping of quality to outcomes

	Percent of cases with listening probabilities equal to:				
	0%-5%	6%-10%	11%-15%	16%-20%	>20%
Best 16 songs	52.34	16.41	10.16	7.81	13.28
Middle 16 songs	63.28	17.19	14.06	2.34	3.12
Worst 16 songs	81.25	15.62	2.34	0.00	0.78

The table reports the fraction of cases in which listening probabilities converged to the given ranges. The “converged” probabilities are calculated as the fraction of the last 100 participants who listened to the song. So, for example, for the best 16 songs (as measured by conditional download probabilities in the control world), in 52.34% of cases only 5 or fewer of the last 100 participants listened to the song.

Table 4: Search efficiency

	Number of participants	Listens per participant	Downloads per listen
Control world	1,446	3.902	0.301
Treatment worlds: First 100	800	3.740	0.316
Treatment worlds: Last 100	800	3.524	0.372

The “First 100” are participants who were among the first 100 to arrive in their randomly assigned world; the “Last 100” were among the last 100 to arrive. Since there were 8 separate treatment worlds, there are 800 participants in each of these categories. The differences between the listening rates for the three groups are not statistically significant. The difference in download rates for control world participants vs. first 100 treatment world participants is not statistically significant, but the difference in download rates for the first 100 vs. last 100 (0.372 vs. 0.316) is statistically significant at the 5% level (p -value for a two-sided t-test is 0.023).