American Economic Journal: Microeconomics 2012, 4(1): 1–31
http://dx.doi.org/10.1257/mic.4.1.1

Observational Learning and Demand for Search Goods[†]

By Kenneth Hendricks, Alan Sorensen, and Thomas Wiseman*

We develop a model of herds in which consumers observe only the aggregate purchase history, not the complete ordered history of search actions. We show that the purchasing information changes the conditions under which herds can occur for both low- and high-quality products. Inferior products are certain to be ignored; high quality products may be ignored, but complete learning may also occur. We obtain closed form solutions for the probabilities of these events and conduct comparative statics. We test the model's predictions using data from an online music market created by Salganik, Dodds, and Watts (2006). (*JEL* D11, D12, L82)

When estimating demand in markets for horizontally differentiated goods, industrial organization economists typically assume that consumers are aware of all products and know their preferences for them. But in many markets the number of available products is very large, with many new products entering the market each week. As a result, consumers are often poorly informed about many products—especially new products—and have to spend time and resources learning about them before deciding whether (or which one) to purchase.[1] In such markets, demand depends not only upon consumer preferences, but also upon which products they choose to investigate.

A consumer's decision to learn about a product can be influenced by the choices of other consumers—an effect that is commonly referred to in the literature as observational learning. For example, in the markets for 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 sales rankings to choose which products to investigate.[2] In online markets for consumer durables, most sellers sort search results by sales ranks, effectively steering consumers toward the most popular products. Consumers'

* Hendricks: Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706 (e-mail: hendrick@ssc.wisc.edu); Sorensen: Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706 (e-mail: sorensen@ssc.wisc.edu); Wiseman: Department of Economics, University of Texas, 1 University Station, Austin, TX 78712 (e-mail: wiseman@austin.utexas.edu). We received very helpful comments from Daron Acemoglu, Nathan Hendren, Johannes Horner, Phil Reny, Lones Smith, Joel Sobel, four anonymous referees, and the participants at the microeconomics seminars at Brown University, Duke University, Hong Kong University, The Workshop on Industrial Organization at University of Santiago, Chile, and the "Theory in A2" conference at the University of Michigan.

[†] To comment on this article in the online discussion forum, or to view additional materials, visit the article page at http://dx.doi.org/10.1257/mic.4.1.1.

[1] In Nelson's (1970) classification, the products we are concerned with are "search goods": those for which a consumer can learn her preferences before purchasing.

[2] Sorensen (2007) examines the impact of bestseller lists on book sales.

tendency to check out products that others buy creates a feedback effect that can cause a product's success or lack of success to reinforce itself. The key questions in this kind of environment are: What is the likelihood that observational learning leads consumers to ignore superior products or waste time and resources learning about inferior products? How does the likelihood of these events depend upon product quality, the kind of information that consumers observe about the product, and what they observe about the decisions of other consumers?

Our objective in this paper is to address these questions both theoretically and empirically. We use a variant of the herding models introduced by Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992) and subsequently generalized by Smith and Sorensen (2000) to capture the salient features of observational learning in product markets. A large number of consumers with heterogeneous preferences arrive sequentially and decide whether or not to buy a product. Consumers do not know the quality of the product, which is either high or low, or their idiosyncratic preferences for it. The common quality implies that preferences are correlated across consumers, so purchases are informative. Each consumer observes the aggregate purchases of previous consumers and a private informative signal about her utility for the product. She uses this information to decide whether or not to investigate the product and learn more about her utility before making a purchasing decision. We refer to the action of investigating the product as "search." Consumers only purchase products that they search and like; but search is costly, so consumers do not search products that they believe they are unlikely to buy.[3] We show that the model has clear empirical predictions about the formation of herds. We then test those predictions using experimental data from an online music market created by Salganik, Dodds, and Watts (2006) (hereafter SDW).

The information structure of our model differs from that of the standard herding models in two important ways. First, consumers only observe aggregate purchases, and not the entire ordered history of purchases. As a result, the sequence of beliefs does not form a martingale, so we need to develop different proof techniques to establish convergence of beliefs. Second, the failure to search is not observed, and search is observed only if the consumer purchases. Thus, consumers cannot distinguish between consumers who did not search and consumers who searched but did not purchase because they did not like the product. As a result, each consumer faces a more complicated inference problem: does low sales mean that relatively few people like the product or that relatively few people have searched the product?

Our first main result is that the purchasing information changes the conditions under which herding can occur. For low-quality products, the possibility of a bad herd is eliminated: in the long-run, the likelihood of search for these products is certain to go to zero. However, for high-quality products bad herds are not eliminated: the likelihood of search goes to zero with positive probability. But, in contrast to the standard models, complete learning can occur despite our assumption that beliefs

[3] Since purchases in our model are always preceded by a search, the best real-world examples are products that are one-time purchases where the price is high relative to the search cost, such as music albums and consumer durables.

are bounded. When this is the case, the market share of the product converges to the share it would have obtained if consumers were able to observe product quality.

Our second main result is that the probability of a bad herd on high quality products has a closed form solution if the density of the private signal satisfies a certain boundedness condition. The condition ensures that the beliefs never "jump" into the cascade sets from outside of the set. We use the closed form solution to conduct comparative statics. Our most striking findings are that the probability of a bad herd is decreasing in product quality and independent of the precision of the private signal.

Hence, in product markets where observational learning is important, our model yields two testable implications. The first is that, in the long run, the probability that a consumer searches (buys) a high-quality product is a random variable with a two-point support. Either a bad herd will form, in which case the probability of search (buy) is essentially zero; or the market will converge to complete learning, in which case the probability of search (buy) approaches a positive constant. The second are the monotonicity results: the probability of a bad herd is decreasing in product quality and, conditional on a bad herd not forming, in the long run, the probability of search (buy) is increasing in product quality. We test these predictions using data from the SDW experiment. In the experiment, participants arrived sequentially at a music website and chose whether to listen to and download songs. Participants in treatment groups were shown which songs had been most heavily downloaded by previous participants, whereas in the control group no such information was given.

Two features of the experiment make it particularly well-suited for testing our model's predictions. First, in order to estimate the empirical model implied by the theory, an exogenous measure of product quality is needed. In the SDW experiment, such a measure is easily calculated from the control group, in which participants were given no information about the choices of other participants. The ratio of a song's downloads to listens in this group (which SDW called the "batting average") is a natural measure of song quality. Second, in order to test for bimodality of outcomes, it is necessary to observe long-run outcomes for multiple realizations of the stochastic learning process. In the SDW experiment we observe realizations of the learning process for 48 songs, and for each song there were 8 independent sequences of at least 700 participants.

The results from the SDW experiment confirm both of our model's main predictions. We show that the data strongly reject a model in which long-run outcomes are drawn from a single distribution in favor of a model in which outcomes are drawn from a mixture of two binomial distributions. And, using the batting average as a measure of quality, we estimate that the probability of a bad herd is indeed decreasing in song quality, ranging from 50 percent for the highest quality song to 96 percent for the lowest quality song.

The paper is organized as follows. In Section I, we briefly discuss related literature. In Section II, we describe our basic model. In Section III, we characterize the equilibrium dynamics and outcomes. Section IV examines the determinants of bad herds. Section V describes the model's empirical implications and tests them with data from the SDW experiment. Section VI concludes.

I. 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. Smith and Sorensen (2000) 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 (2006) and Acemoglu et al. (2009) 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 (2008) derives a similar result under a different condition on heterogeneity, for the case where preference types are publicly observed.

We focus on products where consumers learn their preferences through search, but there are other products where consumers learn through experimentation.[4] Orphanides and Zervos (1995) 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 (2011) 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 social learning relax the assumption that consumers observe the complete, ordered history of actions. Herrera and Horner (2009) study a model in which individuals arrive randomly but arrival times are not observed. Thus, each individual observes the sequence of individuals who acted but cannot observe who has failed to act. They find that the inability to track failures does not alter the conditions under which cascades occur but does affect the probability of bad herds. Banerjee and Fudenberg (2004) 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 (2008b) 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 (1991), does examine that situation.) Further, their proofs do not seem to apply to our setting. Monzón and Rapp (forthcoming) study agents who are uncertain about both their own place in the order of decision-makers and the places of those who are sampled. They also focus on cases other than complete, ordered histories. Celen and Kariv (2004a) assume that agents observe only their immediate predecessor's action, and show that beliefs and actions cycle indefinitely. Acemoglu et al. (2009) study an environment where each agent observes the choices of a random neighborhood of other agents, and provide conditions under which complete learning occurs when signals

[4] Smith and Sorensen (2008a) examine the link between observational learning and a repeat consumer's optimal experimentation.

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.

Burguet and Vives (2000) relax the assumption that the precision of private signals is exogenous. They study observational learning where agents choose how much to invest in increasing that precision. In their setting, agents have heterogeneous preferences, but their model of information acquisition can be viewed as a smoothed version of the discrete option in our model to learn one's preferences exactly. Vives (1997) also examines a model of sequential choice where some consumers are more informed than others about the quality of a product, but in that paper the precision of the signal is exogenous. Both papers quantify the welfare effects of social learning.

There is a large literature on non-Bayesian observational learning on networks, including Ellison and Fudenberg (1993, 1995), Bala and Goyal (2001), DeMarzo, Vayanos, and Zwiebel (2003), and Golub and Jackson (2010). Acemoglu et al. (2009) provides a thoughtful discussion of this branch of the literature.

Finally, recent field experiments have confirmed the empirical relevance of observational learning in which consumers draw inferences from the purchasing decisions of other consumers (e.g., Cai, Chen, and Fang 2009; and Tucker and Zhang 2007). The fact that observational learning can lead to herd behavior has been documented in laboratory experiments (e.g., Anderson and Holt 1997 and Çelen and Kariv 2004b).

II. The Basic Model

An infinite sequence of consumers indexed by t enter in exogenous order. Each consumer makes an irreversible decision on whether or not to purchase the 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. Let F_U denote the distribution of U. The 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 two 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 action that she can choose is to *Not Search* (N), in which case she does not purchase the product.[5] Let $a_t \in \{N, S\}$ 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

$$\int_{p-X}^{\infty} (X - p + u) dF_U(u) = [1 - F_U(p - X)] E[X - p + u | X - p + u \geq 0].$$

In what follows, it will be convenient to assume that F_U is exponential. Under this assumption, the expectation of the surplus $X - p + U$ conditional on that surplus being positive is independent of X ; call its value G. Hence, the payoff to a consumer from action S in state X is $[1 - F_U(p - X)]G - c$. To restrict attention to the economically interesting cases, we impose the following restrictions on the payoffs from search in each state:

(i) $[1 - F_U(p - H)]G - c > 0$ and

(ii) $[1 - F_U(p)]G - c < 0.$

Condition (i) states that, conditional on H, the expected payoff to search is positive. Condition (ii) states that the consumer's expected payoff to search is negative if she knows that the state is L—i.e., search costs are high enough that consumers do not want to search a product if they know its quality is low.

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 her utility for the product. Specifically, the private signal is informative about the probability of the event $\{V_t \geq p\}$, that is, the probability that the consumer likes the product enough to buy it. In that case, Smith and Sorensen (2000) have shown that there is no loss in generality in defining the private signal, σ, that a consumer receives as her *private belief* that $V_t \geq p$. That is, σ is the result of updating a half-half prior with the information in the private signal. Let ω_{1t} denote the event $\{V_t \geq p\}$ and let ω_{0t} denote the complementary event $\{V_t < p\}$. Conditional on event $\omega \in \{\omega_0, \omega_1\}$, the signals are identically and independently distributed across consumers and drawn from a distribution F_ω. We assume that F_{ω_0} and F_{ω_1} are continuous, mutually absolutely continuous, and differentiable with densities $f_{\omega 0}$ and $f_{\omega 1}$. Smith and Sorensen (2000) show that F_{ω_1} stochastically dominates $F_{\omega 0}$. Private beliefs are *bounded* if the convex hull of the common support of F_{ω_0} and F_{ω_1} consists of an interval $[\underline{\sigma}, \overline{\sigma}]$ where ½ $> \underline{\sigma} > 0$ and ½ $< \overline{\sigma} < 1$.

[5] We ignore the possibility of purchasing without search since our focus is on the search decision. The premise 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. Also, this assumption holds for the experiment we analyze in Section VI. Nevertheless, we discuss the implications of allowing consumers to purchase without search at the end of Section IV.

In addition to the private signal, consumer t also obtains information on the decisions by consumers 1 through $t - 1$. The space of possible t-period search histories is given by $H_t = \{0, 1\}^{t-1}$. A particular history is denoted by h_t. The initial history is defined as $h_1 = \varnothing$. However, consumers do not observe the entire ordered history of purchase decisions. Instead, they observe only the aggregate number of consumers who have purchased so far. Let $n_t \in \{0, 1, \ldots, t - 1\}$ denote the aggregate history observed by consumer t.

Given any history n_t, consumer t updates her beliefs about X using Bayes' rule. Let $\mu_t(n_t)$ represent her posterior belief that the state is H conditional on history n_t. Following Acemoglu et al. (2009), we call μ_t the *social belief* in period t. The social belief μ_t generates consumer t's *social probability* ρ_t that her state is ω_{1t} (that is, $V_t \geq p$)

$$\rho_t = \mu_t[1 - F_U(p - H)] + (1 - \mu_t)[1 - F_U(p)]. \tag{1}$$

Given social probability ρ_t and private signal σ_t, consumer t's *private probability* that her state is ω_{1t} is

$$r(\sigma_t, \rho_t) = \frac{\sigma_t \rho_t}{\sigma_t \rho_t + (1 - \sigma_t)(1 - \rho_t)}. \tag{2}$$

Her expected net payoff to action S is

$$r_t(\sigma_t, \rho_t)G - c. \tag{3}$$

Recall that if a consumer chooses N, then her payoff is zero. We look for a Bayesian equilibrium where everyone updates beliefs using Bayes' rule, knows the decision rules of all consumers, and knows the probability laws determining outcomes under those rules. In studying the dynamics of beliefs and actions, it is sometimes useful to follow Smith and Sorensen (2000) and work with the social likelihood ratio that the state is L versus H rather than social beliefs. Define the social likelihood ratio

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

and let l_0 denote the prior likelihood.

A *cascade* on action $a \in \{S, N\}$ has occurred if 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 τ if each consumer $t \geq \tau$ 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. A consumer may ignore her private signal in choosing S, but subsequent consumers may be dissuaded from searching if she decides not to purchase. 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.

III. 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 long-run behavior and, in particular, the likelihood of "inadequate learning." Inadequate learning (in the language of Aghion et al. 1991) occurs when the long-run market share of the product is different from what it would if the quality were known to all consumers.

We begin by defining thresholds. Let $\hat{r}$ represent the private probability at which a consumer is indifferent between S and N. From equation (3),

$$\hat{r} = \frac{c}{G}. \tag{4}$$

The restrictions on payoffs implies that $\hat{r} \in (0, 1)$. Using equations (2) and (4), we can then define the private signal at which a consumer is indifferent between S and N (assuming it is interior) as

$$\hat{\sigma}(l) = \frac{c[lF_U(p) + F_U(p - H)]}{c[(lF_U(p) + F_U(p - H)] + (G - c)[l(1 - F_U(p)) + 1 - F_U(p - H)]}. \tag{5}$$

Thus, given social likelihood ratio 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 (5), $\underline{l}$ is defined as the solution to $\hat{\sigma}(\underline{l}) = \underline{\sigma}$. Solving this equation for $\underline{l}$ yields

$$\underline{l} = \frac{\left[(G - c)\left(\frac{\underline{\sigma}}{1 - \underline{\sigma}}\right) + c\right][1 - F_U(p - H)] - c}{c - \left[(G - c)\left(\frac{\underline{\sigma}}{1 - \underline{\sigma}}\right) + c\right][1 - F_U(p)]}. \tag{6}$$

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 (5), $\bar{l}$ satisfies $\hat{\sigma}(\bar{l}) = \overline{\sigma}$. Solving this equation for $\bar{l}$ yields

$$\bar{l} = \frac{\left[(G - c)\left(\frac{\overline{\sigma}}{1 - \overline{\sigma}}\right) + c\right][1 - F_U(p - H)] - c}{c - \left[(G - c)\left(\frac{\overline{\sigma}}{1 - \overline{\sigma}}\right) + c\right][1 - F_U(p)]}. \tag{7}$$

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

$$[1 - F_U(p - X)][1 - F_{\omega_1}(\hat{\sigma}(l))] + F_U(p - X)[1 - F_{\omega_0}(\hat{\sigma}(l))];$$

and when $l > \bar{l}$, there is a cascade on *N*. Note that when the support of the private signal σ is large, either or both cascade regions may be empty. To focus on the empirically relevant case, we will assume throughout that $\bar{l} > 0$. That is, when a consumer is pessimistic enough about the quality of the product, then even the most positive private signal will not convince her to search.

We now characterize the dynamics of the social likelihood ratio. Consumer t does not know what history consumer $t - 1$ observed. 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 did not purchase. (The exceptions are when either all previous consumers searched ($n_t = t - 1$) or none of them did ($n_t = 0$). A consequence is that the sequence of social likelihood ratios does not form a martingale, and the proof techniques used in the standard herding model to establish convergence no longer apply. (See Acemoglu et al. 2009 and Smith and Sorensen 2008b).

Define $\pi_t(n, X)$ as the equilibrium probability that $n_t = n$ in state X and $\beta_t(n, X)$ as the probability that consumer t purchases in state X after observing $n_t = n$. Recall that $l_t(n)$ is the social likelihood when $n_t = n$. We recursively define $\pi_t(n, X)$ as follows. First, let

$$\pi_1(0, L) = \pi_1(0, H) = 1, \; l_1(0) = l_0$$

$$\beta_1(0, X) = [1 - F_U(p - X)]\big[1 - F_{\omega_1}\big(\hat{\sigma}\,(l_1(0))\big)\big].$$

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

$$\pi_t(n, X) = \begin{cases} \pi_{t-1}(0, X)[1 - \beta_{t-1}(0, X)] \text{ if } n = 0 \\ \pi_{t-1}(n, X)[1 - \beta_{t-1}(n, X)] + \pi_{t-1}(n - 1, X)\beta_{t-1}(n - 1, X) \text{ if } 1 \leq n \leq t - 1 \\ \pi_{t-1}(t - 2, X)\beta_{t-1}(t - 2, X) \text{ if } n = t - 1 \end{cases}$$

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

$$\beta_t(n, X) = [1 - F_U(p - X)]\big[1 - F_{\omega_1}\big(\sigma(l_t(n))\big)\big].$$

Under these dynamics, if the social likelihood ratio l_t converges over time, its limit l_∞ must either equal 0 or be at least $\bar{l}$. That is, either consumers learn for sure that the state is H or eventually consumers stop searching. The reason, simply, is that whenever there is a positive probability that a consumer searches, the probability of purchase varies with the state, and thus the outcome is informative. Note that, in contrast to the standard herding model, learning in our model continues even if a herd on S occurs.

We impose the following restriction on the distribution of private signals:

ASSUMPTION 1: *The density of the private signal conditional on the event* $\{V_t \geq p\}, f_{w_1}$, *is bounded above by* $2\frac{F_U(p)F_U(p-H)}{F_U(p)-F_U(p-H)}$.

Assumption 1 imposes a bound on the density of the private belief conditional on the event that the product is worth purchasing. For example, if 40 percent of consumers like the high quality product enough to buy it and 10 percent like the low quality product enough to buy it, then the upper bound on the density is 3.6. In the uniform case, the bound implies that the distance between the upper and lower bounds of the support of F_{ω_1} has to be at least 0.28. Thus, in this case, Assumption 1 is violated if signals are uninformative (i.e., the bounds, $\underline{\sigma}$ and $\overline{\sigma}$, are close to ½).

The role of Assumption 1 is to rule out the possibility of the social likelihood ratio "jumping" into the cascade set for N from outside.

LEMMA 1: *If* $l_0 \in (0, \bar{l})$, *then* $l_t(n) \in (0, \bar{l})$ *for all* t *and all* $n < t$.

Lemma 1 establishes that the social likelihood ratios stay in the learning region between 0 and $\bar{l}$ as long as the prior l_0 lies in that range. Suppose, contrary to Assumption 1, that the distribution of private signals has a mass point at $\overline{\sigma}$. If the social likelihood ratio l_t is just below $\bar{l}$, then with probability bounded away from zero consumer t will receive a private signal $\overline{\sigma}$ and search. If she does not purchase, then the social likelihood ratio jumps strictly above $\bar{l}$. A similar jump can result with continuous distributions that have big spikes in density, like a normal or exponential distribution with very low variance. In the standard herding model, Smith and Sorensen (2000) and Herrera and Horner (2009) show, respectively, that either log-concavity of the density of the log of the likelihood ratio of the private belief or the increasing hazard and failure ratios property for the distribution of private signals in the two states implies the conclusion of Lemma 1. We note that Assumption 1 is distinct from both of these conditions.

The next two lemmas demonstrate that $\bar{l}$ is a stable fixed point of the social likelihood ratio, and that l_t must converge to either 0 or $\bar{l}$.

LEMMA 2: *If* $l_t \leq \bar{l}$ *is within* ε *of* $\bar{l}$ *infinitely often for all* $\varepsilon > 0$, *then* l_t *converges almost surely to* $\bar{l}$.

LEMMA 3: *If the prior* $l_0 \in (0, \bar{l})$, *then* l_t *converges almost surely to a random variable* l_∞ *with support in the set* $\{0, \bar{l}\}$.

It is worth emphasizing that our convergence results rely upon Assumption 1. By contrast, if consumers can observe the entire ordered purchase history, then l_t converges to a random variable l_∞ even when it can jump into the cascade set from outside. The reason is that the sequence of likelihood ratios form a martingale and convergence follows from the Martingale Convergence Theorem. The likelihood ratios in our model do not have the martingale property. Hence, the proof of Lemma 3 is novel and may be of independent interest, so we include a sketch of the intuition in the text. The idea is to show that if l_t does not converge to 0, then it must approach $\bar{l}$ infinitely often. In that case, because $\bar{l}$ is a stable fixed point, l_t must eventually converge to $\bar{l}$. Suppose, on the other hand, that l_t eventually has support equal to $[a, b]$, with $0 < a \leq b < \bar{l}$. Then, in either state, infinitely often there occur t and n_t such that $l_t(n_t)$ is very close to b. But, if n_t occurs with positive probability in state H, then αn_t (or lower) must occur with the same probability in state L, where

$$\alpha = \frac{1 - F_U(p)}{1 - F_U(p - H)} < 1,$$

is the relative probability of purchase after search in the two states. And then αn_t (or lower) must also have positive probability in state H, or else the corresponding l_t would be larger than any finite b. Iterating, $\alpha^k n_t$ (or lower) must occur infinitely often in both states for any positive integer k. But arbitrarily low fractions of purchases correspond to arbitrarily high belief in state L, contradicting the assumption that b is an upper bound on the social likelihood ratio in the long-run. Thus, $[a, b]$ cannot be the eventual support of l_t. This argument can be generalized to rule out a nonconnected support whose upper bound is strictly below $\bar{l}$.

Together, Lemmas 1 and 3 show that the social likelihood ratio is trapped forever between the stationary points 0 and $\bar{l}$, converging asymptotically to one of those endpoints. We can now state the main technical result of this section. Lemma 4 gives the probabilities, conditional on quality, that the social likelihood ratio converges to each of its two possible limits.

LEMMA 4: *Suppose that* $l_0 \in (0, \bar{l})$. *Then*

$$\Pr\{l_\infty = \bar{l} | L\} = 1$$

and

$$\Pr\{l_\infty = \bar{l} | H\} = \frac{l_0}{\bar{l}} = l_0 \frac{c - \left[(G - c)\left(\frac{\overline{\sigma}}{1 - \overline{\sigma}}\right) + c\right](1 - F_U(p))}{\left[(G - c)\left(\frac{\overline{\sigma}}{1 - \overline{\sigma}}\right) + c\right](1 - F_U(p - H)) - c}.$$

The first probability follows from the fact that Bayesian updating almost surely does not assign probability 1 to the wrong state: $l_\infty = 0$ with probability 0 in state L. The Dominated Convergence Theorem, together with the the fact that the time-0 expectation of l_t is equal to the prior l_0 for all t, yields the second probability. Thus, in state L, eventually no more consumers search. Intuitively, if a positive fraction of consumers continued to search, then the frequency of purchases (i.e., $1 - F_U(p)$ versus $1 - F_U(p - H)$) would reveal that the state is L.

But, by assumption, searching is not optimal in state L. Thus, in state L, the social likelihood ratio must converge to $\bar{l}$, where search stops. Learning is *incomplete* (the belief assigned to L is less than 1), but it is *adequate*: eventually all consumers choose the optimal action for state L.

When the state is H, on the other hand, either continued search eventually reveals the true state, or a string of "no purchases" discourages consumers, the social likelihood drifts to $\bar{l}$, and a herd on N starts. In the first case, learning is complete, but in the second case, learning is both incomplete and inadequate. The second case is a "bad herd": the product is high quality, but consumers stop searching or buying it. (Note that, in the first case, complete learning results when $\underline{\sigma} > 0$: unbounded signals are not necessary for complete learning, in contrast to Smith and Sorensen 2000.)

Along those lines, we can use Lemma 4 to describe the long-run outcomes of our model. When the limiting social likelihood ratio $l_\infty = \bar{l}$, a herd on N eventually arises, and so the long-run fraction of consumers who purchase falls to 0. When $l_\infty = 0$, consumers who get an optimistic-enough private signal $\left(\text{above } \hat{\sigma}(0)\right)$ search, and they purchase if their taste shock also is high enough. (Note that we have not ruled out the possibility that $\hat{\sigma}(0) < \underline{\sigma}$, in which case all consumers, regardless of their private signal, will search if the state is known to be H.)

PROPOSITION 5: *Suppose that $l_0 \in (0, \bar{l})$. Then (i) in state L, almost surely a herd on N occurs and the fraction of consumers who purchase λ_t converges to* 0*; and (ii) in state H, with probability $l_0/\bar{l}$ a herd on N occurs and λ_t converges to* 0*; otherwise, λ_t converges to* $[1 - F_U(p - H)][1 - F_{\omega_1}(\hat{\sigma}(0))]$.

In the standard herding model, individuals observe the actions of previous arrivals which, in our context, corresponds to whether or not they search. In a previous version of this paper, we considered the case where consumers observe previous search decisions rather than purchase decisions. Using standard herding arguments, we showed that the likelihood ratio converges to either $\underline{l}$ or $\bar{l}$ in both states. If it converges to $\underline{l}$, a herd on S occurs. In this case, complete learning does not occur in state H and inadequate learning can occur in state L, that is, consumers search products that are not worth searching. The probability of a herd on N is lower in both states. Thus, the ability of consumers to observe purchasing information changes the conditions under which herds occur. It eliminates the bad herd in state L (i.e., everyone searching an inferior product). It increases the probability of a bad herd in state H (i.e., everyone ignoring a superior product), but it reduces costs of search by eliminating the herd on S. Instead, consumers learn the state and only consumers with strong signals search.

We conclude this section with a discussion of the role of the assumptions of our model. We have made two important simplifying assumptions. First, we assume that the distribution of the idiosyncratic preference shock F_U is exponential, so that the expected surplus from purchase $V_t - p$, conditional on the surplus being positive, is the same in both states. Second, we assume that the distribution of consumers' private signals F_ω depends only on whether or not $V_t \geq p$ and not otherwise on the value of V_t. These assumptions give us a closed-form solution for the search threshold $\hat{\sigma}$, which in turn gives us closed forms for $\bar{l}$ and the probabilities of long-run outcomes in Proposition 5. They also give us a clean condition to ensure that the social likelihood ratio does not jump into the cascade set (Assumption 1). In a previous version of this paper, we discuss how to relax the two assumptions and obtain the same qualitative results.

We note that under those two assumptions, our model is isomorphic (for any fixed parameter values) to the following variation. Each consumer's valuation V_t is either high ($\overline{v} > p$) or low (normalized to zero). The i.i.d. probability that $V_t = \overline{v}$ depends on the quality of the good: it is P_h for a high-quality product, and $P_l < P_h$ for a low-quality product. Parameter values are such that $P_h\overline{v} - p > c > P_l\overline{v} - p$; that is, the expected gain from search is positive in state H and negative in state L. The rest of the model is the same. In particular, consumers' private signals are informative about whether or not $V_t - p > 0$; that is, about whether $V_t = \overline{v}$ or $V_t = 0$. It is straightforward to see the equivalence between that model and our baseline model: simply replace G with $\overline{v} - p$, $F_U(p - H)$ with $1 - P_h$, and $F_U(p)$ with $1 - P_l$.

That equivalence, however, breaks down when we consider comparative statics (in the next section). For example, consider the effect of raising the price p on search and purchase dynamics. In our baseline model, that increase has two effects. First, consumers are (all else equal) less likely to search, because the expected surplus is lower. Second, conditional on search, consumers are less likely to buy: only consumers with valuations greater than p purchase. In the alternative, two-valuation model, only the first effect is present. Conditional on search, demand is a stairstep function, so quantity is (locally) invariant to price. Thus, our baseline model may be more appropriate than the alternative model for examining optimal pricing, or for examining environments where prices vary (such as the Amie Street online music store, where, until the site was purchased by Amazon.com, the price of a song increased with the number of times that it had been downloaded).

A third assumption is that the private signal is informative about utility and not, as in the standard herding model, only about product quality. When the signal is about X only, then the probability of search in state H is equal to one whenever l_t is less than $\underline{l}$. In our model, this probability can be less than one even when the state is known to be H because of selection effects. Consumers who search are more likely to like the product because the signal is informative about U as well as X. We suspect that this kind of selection effect is present in any realistic model of choice with private signals and costly search. It is certainly present in the SDW experiment where 40 percent of the participants choose not to search any songs. We have considered the polar cases in which the private signal is informative only about X or U and obtained similar qualitative results.

Finally, we have assumed that consumers must search before they can purchase. If instead consumers were allowed to purchase without search, then it would be possible for bad herds to arise even for low-quality products—i.e., in the limit, consumers could all end up buying a low-quality product. Thus, our model intentionally focuses on the opposite kind of bad herd, in which consumers end up ignoring high-quality products.

IV. Determinants of the Bad Herd

In this section, we analyze the determinants of inadequate learning. This kind of analysis has largely been absent in the literature on herding models, mainly because it considers models in which the likelihood ratio may "jump" into the cascade sets (e.g., with discrete signals). In these cases, comparative static results are difficult, if not impossible, to obtain without making strong parametric assumptions on the distribution of preferences and, in particular, on the distribution of the private signal. The reason is simple: changes in model parameters cause changes in the decision rules of consumers, 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 can reverse the direct impact of the parameter change on decision rules. For example, an increase in product quality makes consumers more likely to purchase and search, but it is not difficult to construct examples in which the likelihood of paths ending in the cascade set for N is higher. Pastine and Pastine (2006) obtain similar "perverse" results when they study the effect of changing the accuracy of signals on the probability of incorrect herds, and Burguet and Vives (2000) find that increasing the noise in public information can improve welfare. By restricting the likelihood ratio to lie within the learning region, we are able to obtain a closed form solution for the probability of a bad herd.

A striking implication of Lemma 4 is that this probability does not depend upon the specific functional form of F_ω, the distribution of the private signal. The only property of this distribution that matters (other than the upper bound on the density) is $\overline{\sigma}$, the upper bound of its support.[6] This observation also applies to signals about the utility of consumers who have purchased the product, such as consumer reviews and other forms of word-of-mouth communication. A number of empirical studies (e.g., Chevalier and Mayzlin 2006) have tried to measure the impact of word-of-mouth communication on consumer decisions. Our model suggests that, while word-of-mouth communication may have an impact on short-run behavior, it has no impact on long-run behavior and, in particular, on the probability of a bad herd. Similarly, the details of what consumers observe about previous purchase decisions may also have no long-run impact. In a previous version of this paper, we showed that, if consumers observed the entire ordered history of previous purchases (and the density of the log of the likelihood ratio of private signals is log-concave), then the likelihood ratio converges to either 0 or $\bar{l}$ and the probability of these events is the same as in the above model.

The next proposition describes how the probability of a bad herd changes with the parameters of the model. Differentiating the probability $l_0/\bar{l}$ in Proposition 5 yields the following results.

[6] Note that, as $\overline{\sigma}$ approaches 1, $\bar{l}$ goes to infinity, and the probability of a bad herd in H goes to zero. Thus, in our model, unbounded beliefs lead to complete learning. (See Smith and Sorensen 2000.)

PROPOSITION 6: *Suppose the state is* H *and* $l_0 \in (0, \bar{l})$. *Then the probability a herd on* N *occurs and the fraction of consumers who purchase* λ_t *converges to* 0 *is strictly decreasing in* H, *increasing in* c, *and increasing in* p.

Proposition 6 yields several intuitively plausible results. An increase in search costs increases the likelihood that long-run sales of a high quality product are zero and hence reduces its expected long-run sales. An increase in price also increases the probability of zero long-run sales, while an increase in the quality level decreases the probability. Intuitively, the first two effects reduce the expected net value of search for any pair of social likelihood ratio and private signal, while the third raises the value of search.

Rosen (1981) 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. This effect may also help explain why prices of products like albums, books, and videos do not vary with quality. In our model, 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. It explains why a seller of high quality products may want to keep prices low, at least initially, to encourage a positive herd on its product.

A number of papers (e.g., Brynjolfsson, Hu, and Simester 2006) 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 7: *Suppose the state is* H *and* $l_0 \in (0, \bar{l})$. *Then the impact of an increase in* c *on the probability that a herd on* N *occurs and the fraction of consumers* λ_t *converges to* 0 *is smaller* (*in absolute value*) *for higher quality products.*

The result follows from differentiating the probability $l_0/\bar{l}$ with respect to c and then H. Proposition 7 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.

V. Empirical Model

As described in the previous section, the central empirical predictions of our model relate to the likelihood of bad herds. In particular, our model predicts that long-run search probabilities have a two-point support: ultimately the market will either learn the product's true quality, in which case it will be searched with probability $1 - F_{\omega_1}(\hat{\sigma}(0))$, or the product will be ignored. Proposition 6 states that the likelihood of the latter outcome (the bad herd) should be declining in product quality.

A clever online experiment conducted by Salganik, Dodds, and Watts (2006) provides a nice opportunity to test these predictions. 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, participants were shown information about the downloads of previous participants in their world: song listings included the total number of previous downloads, and the song list was sorted by download rank. In the control world, the songs were shown in a random order, with no information about previous participants' listens or downloads. A song's download rate or "batting average" in this world provided a natural measure of its quality since it simply reflects the probability that participants choose to download a song conditional on listening to it. The eight treatment worlds operated independently of one another, so that the researchers could observe eight separate realizations of the stochastic learning process for each song.[7]

The design of the SDW 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.[8] 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. As in our model, downloading (purchasing) was only possible if the participant had first listened to (searched) the song. The signal that a participant observes is the title of the song and name of the artist, which may signal whether the song's style or genre is one that the participant likes. Since this is mostly about the participant's idiosyncratic taste, U, the signal is private. The participants assigned to treatment worlds were shown the aggregate number of downloads by previous participants, so 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) and the cost of downloading (i.e., the opportunity cost of time and disk space) were apparently large enough to matter, because on average participants listened to fewer than four songs and downloaded fewer than two.

To test for bimodality of long-run outcomes, it is obviously important that we observe multiple learning processes. For each song in the SDW experiment, we observe eight separate sequences of participants (since there were eight independent treatment worlds). By itself, this would not be enough to perform any meaningful statistical analysis. But the experiment generated data for 48 different songs, and since the songs' qualities can be estimated cleanly from the control world, we can control for quality differences and pool data across songs. Hence, we effectively observe 48×8 realizations of the learning process, and the heterogeneity in quality across songs allows us to test our model's monotonicity predictions.

[7] See Salganik (2007) for a much more thorough description of the methods and results of the experiment.

[8] They 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.

In pooling the data across songs, we are treating the sequences of listens and downloads for different songs in the same treatment world as realizations of independent stochastic processes. For independence to hold, we must assume that participants' preferences are additive. This is a strong assumption since it ignores substitution effects that may arise when participants make choices from a menu of songs. The fact that many participants 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.

Since our model's predictions are based on the assumption that the social likelihood ratio has converged, we analyze the songs' listening rates among participants who arrived late (i.e., toward the end of the sequence) in the treatment worlds. The sample consists of the last 200 participants of the roughly 700 participants in each treatment world.[9] Let L_{jwt} be an indicator variable equal to 1 if participant t listened to song j in world w and let Y_{jw} denote the number of listens for song j among the last N participants in world w. If the social likelihood ratio has converged to zero (that is, the song is known to be high quality), then X is known, and the L_{jwt} are independently and identically distributed. If instead the social likelihood has converged to $\bar{l}$, then in theory the probability of search should be zero and we should observe L_{jwt} to be a sequence of zeroes. In the data, few songs get exactly zero listens, so in estimating the model below we allow the search probability (which we denote p_0) to be a small positive number.

Taken together, the two possible outcomes for the social likelihood ratio imply that, in the long-run, Y_{jw} is a weighted average of two binomial random variables:

(8)

$$\Pr\{Y_{jw} = n\} = q(x_j)\binom{N}{n} p_0^n(1 - p_0)^{N-n} + (1 - q(x_j))\binom{N}{n} p(x_j)^n(1 - p(x_j))^{N-n},$$

where x_j is a measure of song j's quality, $q(x_j)$ is the probability that the market has converged to a bad herd, and $p(x_j)$ is the probability that a participant would choose to listen to song j if its quality x_j is known (i.e., the market has learned the quality of the song).

For purposes of estimation, we parameterize the probability functions as

$$q(x_j) = \frac{\exp\{\alpha_0 + \alpha_1 x_j\}}{1 + \exp\{\alpha_0 + \alpha_1 x_j\}} \qquad \text{and} \qquad p(x_j) = \frac{\exp\{\beta_0 + \beta_1 x_j\}}{1 + \exp\{\beta_0 + \beta_1 x_j\}}.$$

[9] We found that after the 500th participant, the listening probabilities had mostly "converged," in the sense that they changed very little from one cohort of 100 participants to the next. The results we report below are largely unchanged if we look only at the last 100 participants (instead of 200).

Table 1—Estimation Results: Test of Bimodality of Outcomes

	Listens		Downloads	
Parameter	(1)	(2)	(3)	(4)
α_0	3.796	—	3.492	—
	(1.151)		(2.160)	
α_1	−6.386	—	−6.729	—
	(2.684)		(4.383)	
β_0	−3.891	−4.462	−5.506	−6.109
	(0.718)	(0.435)	(1.279)	(0.455)
β_1	5.341	4.723	6.446	6.478
	(1.732)	(1.159)	(2.636)	(1.242)
Observations	384	384	384	384
Avg. log-likelihood	−4.403	−7.608	−3.699	−5.139

Notes: Estimates in columns 1 and 2 are based on the number of listens among the last 200 participants in the treatment worlds. Estimates in columns 3 and 4 are based on the number of downloads among the last 300 participants. An observation is a song/world pair. Standard errors (corrected for clustering on songs) are reported in parentheses.

We measure x_j, song quality, as the ratio of downloads to listens for song j in the control world. We set the minimum listening rate, p_0, to 0.043, which is the average listening rate for songs falling in the bottom quarter of the list.

While the SDW experiment generated data on both the listen (search) and download (purchase) decisions, in real-world markets it is much more common to observe purchases only. The test we described above can be applied to purchase outcomes in the same way as to search outcomes. Under the assumption that consumers have learned song quality, the long-run sequence of listen decisions is i.i.d.—and the same argument applies to download decisions. So the number of downloads is also a weighted average of two binomial random variables, and an equation analogous to equation (8) can be derived for it. Hence, our model's main prediction can be tested by looking for bimodality in either search outcomes or purchase outcomes.

The first column of Table 1 reports maximum likelihood estimates of the model in equation (8). The second column reports estimates of the same model with q, the probability of a bad herd, constrained to be zero—so the number of listens is a binomial random variable based on a single success probability, $p(x)$. The third and fourth columns report analogous estimates based on download decisions—i.e., with the number of downloads used as the dependent variable instead of number of listens.[10] The number of observations is 384, since we observe listen counts (Y_{jw}) for 48 songs, and 8 treatment worlds for each song.

[10] For the estimates based on downloads, we set p_0 to 0.014, the average download rate for songs falling in the bottom quartile of the list. We also set $N = 300$, because the relative infrequency of downloads makes the download counts noisier than the listen counts.

The results clearly confirm the predictions of the model. First, whether we base the test on listens or on downloads, the data easily reject the constrained model.[11] The reason is that the constrained model cannot fit the bimodal outcomes of high-quality songs. Even though songs with batting averages in the top quartile had listening rates above 20 percent in many worlds, in a third of cases the listening rates for these songs were below 4 percent. For example, the song with the highest control-world batting average had listening rates between 18–30 percent in 5 of the 8 treatment worlds, but in two other worlds this same song's listening rates were only 2.5 percent and 3.5 percent.

Second, the estimates confirm the comparative static prediction of Proposition 6: the estimate of α_1 is negative and statistically significant, indicating that the probability of a bad herd (q) is a decreasing function of quality. For the estimates based on listening decisions, the predicted values of q lie between 49.7 percent (for the best song) and 95.6 percent (for the worst song). The estimates for p are also consistent with the prediction of our model: the estimate of β_1 is positive and statistically significant, indicating that the probability of listening is an increasing function of quality. It lies between 3.6 percent and 33.1 percent. The estimates based on download decisions (in columns 3 and 4 of the table) are quantitatively very similar to the estimates based on listens, although with larger standard errors. This imprecision reflects the fact that downloads are more infrequent, so that download counts among the late-arriving participants are noisier than listen counts.

The estimates imply that if bad herds were eliminated so that song qualities were always learned in the long run—i.e., if q were zero for all songs—then expected listening rates would increase by an average of 4.2 percentage points. The difference would be negligible for low-quality songs, and as high as 16 percentage points for the high-quality songs.[12]

Alternative Models.—In our model, consumers are influenced by others' purchases only because those purchases affect the perceived benefits of search. One plausible alternative is a model with "social preferences," in which herds form because individuals prefer to consume the same products as others. Our model implies that download information should influence participants' listening decisions, but not their decisions about whether to download a song conditional on listening to it (i.e., the song's batting average). In contrast, a social preferences model would predict that batting averages are influenced by download information.

Let D_{jwt} be an indicator equal to 1 if participant t of world w downloaded song j. We can ask whether L_{jwt} and D_{jwt} are influenced by song j's download share among participants $1, \ldots, t-1$ in world w. Table 2 reports the results from probit regressions in which L_{jwt} and D_{jwt} are assumed to depend on song j's current

[11] The p-value of the likelihood ratio test is essentially zero in both cases.

[12] This mirrors the analysis of Hendricks and Sorensen (2009), who estimate the counterfactual distribution of music album sales in a world where consumers knew their preferences for every album (instead of being unaware of most albums).

Table 2—Probit Regressions

	Listen (1)	Download (2)
Download share	1.074	−0.222
	(0.011)	(0.061)
Control world Listening share	0.414	
	(0.036)	
Control world Cond. prob. download		0.783
		(0.040)
Observations	264,288	20,194

Note: Probit regressions; marginal effects reported; standard errors in parentheses.

download share (i.e., the fraction of prior participants who chose to download the song j). In the latter regression, we are interested in whether download information affects the probability of download *conditional* on listening, so we use only the observations of D_{jwt} for which $L_{jwt} = 1$. 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. The coefficient on download share in this regression is positive and highly significant, indicating that participants' decisions to listen to a song were influenced by information about previous participants' downloads.

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_{jwt} 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. As explained above, this conditional probability provides an independent measure of the song's relative quality.[13] When included, it forces the coefficient on download share to be identified from time variation in the download 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

[13] Unconditional 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 percent for the highest-quality song, and only 11 percent for the lowest-quality song.

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.

Another model that could plausibly explain the patterns in the SDW data is one in which listening decisions are determined by songs' list positions. Participants in the experiment may have been generally inclined to listen to songs shown at the top of the list, irrespective of any information that was posted about previous downloads. Since songs in the treatment worlds were ordered by download rank, the data would then conflate a list-position effect with the effect of download information.

SDW ran a separate experiment in which songs in the treatment worlds were randomly ordered (with download information still displayed). According to our model, if social beliefs have converged, the listening decisions of late participants should depend only upon their private signals and download information. However, we find that list position was also influential in this alternate experiment. As a result, low-quality songs got more listens than they otherwise would have because they were sometimes listed at the top, and high-quality songs got fewer listens because they were sometimes listed at the bottom. Consequently, the top songs (in terms of downloads) had long-run listening probabilities in the 15–25 percent range, as opposed to 25–45 percent in the experiment where songs were ranked by downloads. This "flattening" of the data makes it more difficult to detect the bimodality of outcomes. If we estimate the binomial-mixture model described above using data from the experiment with random ordering, we get coefficients of the same signs—$q(x)$ and $p(x)$ are estimated to be decreasing and increasing functions of x, respectively—but the estimates of the coefficients in $q(x)$ are statistically imprecise.

To some extent, this result calls into question whether our results are driven purely by herd effects resulting from observational learning. However, one plausible interpretation is that social beliefs in this alternate experiment had not yet converged, so that the underlying assumption of the test is not satisfied. Indeed, listen rates in the alternate experiment did not settle down like they did in the main experiment. For example, in the main experiment, the listen rates among participants 501+ hardly changed relative to the listen rates among participants 301–500. The simple correlation between the songs' listen rates for those two cohorts is 0.949. By contrast, the listen rates in the alternate experiment appeared to still be in flux. The correlation between listen rates for cohorts 301–500 and 500+ is only 0.640.

Table 3—Probit Regressions: Experiment with Songs Ordered Randomly

	Listen (first 200) (1)	Listen (last 200) (2)	Download (3)
Download share	0.457	0.959	0.147
	(0.033)	(0.050)	(0.134)
List position	−0.00151	−0.00098	0.00041
	(0.000076)	(0.000064)	(0.000215)
Control world	0.642	0.509	
Listening share	(0.060)	(0.054)	
Control world			0.849
Cond. prob. download			(0.034)
Observations	76,800	76,800	21,940

Note: Probit regressions from experiment in which songs were randomly ordered (instead of sorted by download rank); marginal effects reported; standard errors in parentheses.

To investigate this explanation further, we ran probit regressions analogous to those shown in Table 2, but including list position (1–48) as an additional explanatory variable. Table 3 reports results from these regressions. A song's position on the screen has a significant influence on whether the participant listens to it, but as more download information accumulates, the relative importance of list position diminishes. Columns 1 and 2 compare results for the first 200 participants to arrive versus the last 200 participants. The coefficient on download share more than doubles for the late-arriving participants, and the magnitude of the list position coefficient shrinks.[14] The third column shows that neither download information nor list position has a significant influence on the probability of download conditional on listening (i.e., the batting average).

The first two columns suggest that, in the long run, the download information would swamp the list position effects. The third column rules out any feedback effect of list position on listening decisions through download information (i.e., list position is a pure framing effect). Together, they imply that list position may have little impact on long-run outcomes. In other words, had there been 7,000 participants in each treatment world instead of 700, say, we would have obtained results very close to those reported in Table 1. Under this interpretation, the primary consequence of sorting songs by download rank was to amplify the effects of the download information and make the learning processes converge more quickly.

VI. 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 that they have checked out, and doing so is costly. Consumers would prefer

[14] The difference between the download share coefficients is highly statistically significant. Equality of the list position coefficients can be rejected at a 10 percent significance level (the p-value is 0.06).

not to pay this cost if they believe they are unlikely to buy the product. The 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 search the product. We have shown that beliefs converge to one of two possible limits and characterized the probabilities of these events. The probabilities have closed form solutions that can be used to study the determinants of inadequate learning.

The experimental data of Salganik, Dodds, and Watts (2006) confirm our model's central empirical prediction, which is that long-run search probabilities should be bimodal. The test we employ is to estimate a model in which the number of searches (listens) is a weighted sum of two binomial random variables with different success probabilities. Although some features of this test are specific to the SDW experiment, the basic principle could be applied in other real-world markets. What is critical is the availability of an independent measure of quality. A key advantage of the SDW experiment is that the control world yields such a measure. In real-world markets, quality could be measured based on user ratings (e.g., user reviews of books on Amazon or of movies on Netflix). The point is that if quality can be measured independently of the product's observed market success, then with data on many products it is possible to estimate the frequency with which high-quality products end up with near-zero market shares (i.e., bad herds). One could also then test whether this frequency depends on quality, search costs, or prices in the ways that our model predicts.

We have provided an equilibrium analysis of consumer learning in product markets when prices are fixed and sellers report aggregates sales. In practice, the seller can take other actions to increase the likelihood that consumers will investigate its product. One interesting open question is how a monopoly seller, assuming it does not know the quality of its product, can use dynamic pricing to manipulate learning. Another issue is the competition among sellers for the attention of consumers. For example, Eliaz and Spiegler (2008) develop a static model in which competing firms use advertising to influence the decision of consumers to check out their products. Display advertising clearly plays a role in directing consumer search, and incorporating advertising into a dynamic model like ours—with Bayesian learning and meaningful search costs—with multiple sellers should be a promising direction for future research.

Appendix

PROOF OF LEMMA 1:

Using the definition of $\pi_t(n, X)$, we can rewrite the social likelihood ration $l_t(n)$ for non-extreme values of n (greater than 0 and less than $t - 1$) after some manipulation as

$$l_t(n) \;=\; q\,\psi_0\big(l_{t-1}(n)\big)\,l_{t-1}(n) + (1-q)\psi_1\big(l_t(n-1)\big)\,l_{t-1}(n-1),$$

where

$$q = \frac{\pi_{t-1}(n,H)[1-\beta_{t-1}(n,H)]}{\pi_{t-1}(n,H)[1-\beta_{t-1}(n,H)] + \pi_{t-1}(n-1,H)\beta_{t-1}(n-1,H)} \tag{A1}$$

$$\psi_0(l)l = \left[\frac{1-[1-F_U(p)][1-F_{\omega_1}(\hat{\sigma}(l))]}{1-[1-F_U(p-H)][1-F_{\omega_1}(\hat{\sigma}(l))]}\right] l$$

$$\psi_1(l)l = \frac{[1-F_U(p)]}{[1-F_U(p-H)]}\, l.$$

That is, the social likelihood 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 $n_t = n_t - 1$ and consumer $t-1$ did purchase. The weights are the relative probabilities of these events in state H. When $n_t = 0$ or $n_t = t-1$, then $q = 1$ or $q = 0$, respectively. It is easily checked that the fixed points of the dynamics under $\psi_0(l)l$ and $\psi_1(l)l$ are $l = 0$ and $l \geq \bar{l}$. Therefore, to prove the lemma, it is sufficient to show that $\psi_0(l)l$ and $\psi_1(l)l$ are increasing in l. If l_t is between 0 and $\bar{l}$, then so will be $\psi_0(l_t)l_t$ and $\psi_1(l_t)l_t$ and any weighted average of $\psi_0(l_t)l_t$ and $\psi_1(l_t)l_t$.

It is clear from the definition that $\psi_1(l)l$ is strictly increasing in l. For $\psi_0(l)l$, differentiating equation (A1) yields

$$\begin{aligned}
[\psi_0(l)l]' &= \psi_0(l) + \psi_0'(l)l \\
&= \psi_0(l) - \frac{[F_U(p) - F_U(p-H)]\, f_{\omega_1}(\hat{\sigma}(l))\hat{\sigma}'(l)}{\left[1-(1-F_U(p-H))(1-F_{\omega_1}(\hat{\sigma}(l)))\right]^2}\, l \\
&= \frac{\left[1-(1-F_U(p-H))(1-F_{\omega_1}(\hat{\sigma}(l)))\right][1-(1-F_U(p))][1-F_{\omega_1}(\hat{\sigma}(l))]}{\left[1-(1-F_U(p-H))(1-F_{\omega_1}(\hat{\sigma}(l)))\right]^2} \\
&\quad - \frac{[F_U(p) - F_U(p-H)]\, f_{\omega_1}(\hat{\sigma}(l))\hat{\sigma}'(l)}{\left[1-(1-F_U(p-H))(1-F_{\omega_1}(\hat{\sigma}(l)))\right]^2}\, l \\
&= \frac{[F_U(p-H) + F_{\omega_1}(\hat{\sigma}(l))(1-F_U(p))][F_U(p) + F_{\omega_1}(\hat{\sigma}(l))(1-F_U(p))]}{\left[1-(1-F_U(p-H))(1-F_{\omega_1}(\hat{\sigma}(l)))\right]^2} \\
&\quad - \frac{[F_U(p) - F_U(p-H)]\, f_{\omega_1}(\hat{\sigma}(l))\hat{\sigma}'(l)}{\left[1-(1-F_U(p-H))(1-F_{\omega_1}(\hat{\sigma}(l)))\right]^2}\, l \\
&\geq \frac{F_U(p-H)F_U(p) - [F_U(p) - F_U(p-H)]\, f_{\omega_1}(\hat{\sigma}(l))\hat{\sigma}'(l)l}{\left[1-(1-F_U(p-H))(1-F_{\omega_1}(\hat{\sigma}(l)))\right]^2}.
\end{aligned}$$

Thus, $\psi_0(l)l$ is increasing in l if, for all l,

$$\text{(A2)} \qquad f_{\omega_1}\big(\hat{\sigma}(l)\big)\,\hat{\sigma}'(l)l\big[F_U(p - H) - F_U(p)\big] + F_U(p)F_U(p - H) \geq 0.$$

Next, we bound the magnitude of the term $\hat{\sigma}'(l)l$. To simplify the algebra, define $\Gamma \equiv \frac{1-\rho}{\rho}$ as the inverse likelihood ratio of the social probability. Expressing the indifference threshold $\hat{\sigma}$ in terms of Γ, we obtain

$$\hat{\sigma}(\Gamma) = \frac{\Gamma}{\Gamma + \hat{\Gamma}},$$

where $\hat{\Gamma} \equiv \dfrac{G - c}{c}$. Then

$$\frac{d\hat{\sigma}(\Gamma)}{d\Gamma} = \frac{\hat{\Gamma}}{[\Gamma + \hat{\Gamma}]^2}.$$

Thus,

$$\begin{aligned}
\hat{\sigma}'(l)l &= \frac{d\hat{\sigma}(\Gamma)}{d\Gamma}\,\frac{d\Gamma}{d\rho}\,\frac{d\rho}{d\mu}\,\frac{d\mu}{dl}\,l \\
&= \frac{\hat{\Gamma}}{[\Gamma + \hat{\Gamma}]^2}\left(\frac{-1}{\rho^2}\right)[F_U(p) - F_U(p - H)](-\mu^2)l \\
&\leq \frac{1}{2\Gamma\rho^2}\,[F_U(p) - F_U(p - H)]\mu^2\,l \\
&= \frac{1}{2(1 - \rho)\rho}\,[F_U(p) - F_U(p - H)]\mu(1 - \mu) \\
&= \frac{1}{2}\left(\frac{[F_U(p) - F_U(p - H)]\mu}{(1 - F_U(p)) + (F_U(p) - F_U(p - H))\mu}\right) \\
&\qquad \times \left(\frac{(1 - \mu)}{F_U(p) - \big(F_U(p) - F_U(p - H)\big)\mu}\right) \\
&\leq \frac{1}{2}\left(\frac{F_U(p) - F_U(p - H)}{F_U(p) - F_U(p - H) + F_U(p - H)\big(1 - F_U(p)\big)}\right) \\
&\leq \frac{1}{2}.
\end{aligned}$$

The second inequality follows from the fact that the second term is increasing in μ and the third is decreasing in μ. Substituting the bound into expression (A2) implies that $\psi_0(l)l$ is increasing in l if

$$0 \le \frac{1}{2} f_{\omega_1}(\hat{\sigma}(l))[F_U(p-H) - F_U(p)] + F_U(p)F_U(p-H)$$
$$\Leftrightarrow f_{\omega_1}(\hat{\sigma}(l)) \le 2\frac{F_U(p)F_U(p-H)}{F_U(p) - F_U(p-H)},$$

the condition ensured by Assumption 1.

PROOF OF LEMMA 2:

In state X, the probability of the transition from n to $n+1$, $\beta_t(n, X)$, shrinks continously to 0 as $l_t(n)$ approaches $\bar{l}$. Also, as $l_t(n)$ approaches $\bar{l}$, 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 (2000) "Rest of Proof of Theorem 4" (page 397) applies directly: $\bar{l}$ is a stable fixed point.

PROOF OF LEMMA 3:

Lemma 1 establishes that l_t stays in the set $(0, \bar{l})$. Define the random variable y_T as

$$y_T \equiv \sup\{l_t : t \ge T\},$$

and the random variable y as

$$y \equiv \lim_{T\to\infty} \sup y_T.$$

(That is, for any $\varepsilon > 0$, l_t exceeds $y + \varepsilon$ only a finite number of times, but it exceeds $y - \varepsilon$ infinitely often. If l_t converges to z, then $y = z$. If l_t does not converge, then y is the upper bound of its eventual support.) If $y = \bar{l}$, then l_t approaches $\bar{l}$ arbitrarily closely infinitely often, and Lemma 2 implies that l_t converges to $\bar{l}$. If $y = 0$, then l_t converges to 0, since l_t is non-negative. Thus, if y equals 0 or $\bar{l}$, then we have the desired result.

To finish the proof, we suppose that $y \in (0, \bar{l})$ and derive a contradiction. Choose $\varepsilon > 0$, and let $\tau(\varepsilon)$ denote the last time at which $l_t > y + \varepsilon$. Let $A_t(\varepsilon)$ denote the event that $t > \tau(\varepsilon)$ and $l_t(n_t) > y - \varepsilon$. By definition of y, this event must occur infinitely often. Further, it must occur infinitely often in both states (since the likelihood ratio y assigns positive probability to both states). Thus,

$$\lim_{s\to\infty} \inf \sum_{t=s}^{\infty} \Pr[A_t(\varepsilon) | X] > 0 \text{ for } X \in \{L, H\}.$$

But, for any aggregate number of purchases $n_t = n$ that occurs with probability π through t periods in state H, we claim that the event

$$n_t \leq \left\lfloor \frac{1 - F_U(p)}{1 - F_U(p - H)} n \right\rfloor + 1 \equiv \nu(n),$$

has probability at least π in state L. The proof is by induction. Fix any infinite sequence of independent draws from F_{ω_1}. In period 1, the ratio of the probabilities of purchase in state H versus state L is

$$\frac{[1 - F_U(p)][1 - F_{\omega_1}(\hat{\sigma}(l_0))]}{[1 - F_U(p - H)][1 - F_{\omega_1}(\hat{\sigma}(l_0))]} = \frac{1 - F_U(p)}{1 - F_U(p - H)}.$$

Further, since $u - p \geq 0$ implies $H + u - p \geq 0$, for any realization of u, purchase in state L implies purchase in state H. Now choose any period $t > 1$ and any $n_H, n_L \in \{0, \dots, t - 1\}$ such that $n_H \geq n_L$. The ratio of the probability of purchase in state H given $n_t = n_H$ to the probability of purchase in state L given $n_t = n_L$ is

$$\frac{[1 - F_U(p)][1 - F_{\omega_1}(\hat{\sigma}(l_t(n_L)))]}{[1 - F_U(p - H)][1 - F_{\omega_1}(\hat{\sigma}(l_t(n_H)))]} \leq \frac{1 - F_U(p)}{1 - F_U(p - H)}.$$

The inequality follows because (i) $[1 - F_{\omega_1}(\hat{\sigma}(l_t(n_L)))]$ is decreasing in l_t and (ii) $l_t(n)$ is decreasing in n.

Property (ii) is intuitive, but not quite trivial.[15] For $t = 2$, the result is immediate. So pick any $t > 2$, and assume that $l_k(n)$ is decreasing in n for all $k < t$ and any prior between 0 and $\bar{l}$. Consider any two ordered histories through period $t - 1$, h and h', with the property that the number of purchases n in history h is greater than the number $n - 1$ in h'. Let $s \in \{0, \dots, t - 2\}$ be the largest number such that the first s periods of both histories contain the same number of purchases; denote that number as $n(s)$. Let $l_t(n, s, n(s))$ and $l_t(n - 1, s, n(s))$ be the social likelihood ratios that consumer t would have if (contrary to fact) she observed s and $n(s)$ in addition to n and $n - 1$, respectively. We will show that $l_t(n) < l_t(n - 1)$ by showing that $l_t(n, s, n(s)) < l_{t+1}(n - 1, s, n(s))$ for any s and $n(s)$. If $s > 0$, then $l_t(n, s, n(s))$ and $l_t(n - 1, s, n(s))$ are the likelihood ratios that result from observing a $((t - 1) - s)$-period history containing $n - n(s)$ purchases and $n - 1 - n(s)$ purchases respectively, starting from an initial likelihood ratio $l_{s+1}(n(s))$. By hypothesis, then, $l_t(n, s, n(s)) < l_t(n - 1, s, n(s))$, since $s + 1 < t$. If $s = 0$ (so that histories h and h' are identical except that the first consumer purchased in h and not in h'), the argument is similar. In that case, $l_t(n, s, n(s))$ and $l_t(n - 1, s, n(s))$ are the likelihood ratios that result from observing a $(t - 2)$-period history containing $n - 1$ purchases, starting from the initial likelihood rations $l_2(1)$ and $l_2(0)$ respectively. Since $l_2(1) < l_2(0)$ by the induction hypothesis, and the social likelihood ratio is increasing in the prior likelihood ratio (by a proof analogous to the argument that $\psi_0(l)l$ and $\psi_1(l)l$ are increasing in the proof of Lemma 1), the result follows.

[15] Callander and Horner (2009), in fact, show that property (ii) can fail if Assumption 1 does not hold.

Thus, since $n_t \in A_t$ occurs infinitely often in state H, $n_t' \leq \nu(n_t)$ occurs infinitely often in state L. But then $n_t' \leq \nu(n_t)$ must also occur infinitely often in state H, since otherwise $l_t(n_t')$ would be larger than any finite y. Iterating, $n_t' \leq \nu(\nu(n_t))$ must also occur infinitely often in both states, as must $n_r' \leq \nu(\nu(\nu(n_t)))$, and so on. For large n (above $[1 - F_U(p - H)]/[F_U(p - H) - F_U(p)]$), $\nu(n)$ is strictly less than n. Thus, histories in which an aribitrarily low fraction of consumers have purchased must occur infinitely often in both states. But, at such histories, the social likelihood ratio assigns weight close to 1 to state L, contradicting the definition of y. Thus, y cannot lie in $(0, \bar{l})$.

PROOF OF LEMMA 4:

Lemma 3 establishes that l_t converges almost surely to a $l_\infty \in \{0, \bar{l}\}$. In state L, consumer's beliefs cannot converge to something completely wrong with positive probability, and so l_t must converge almost surely to $\bar{l}$. Next, suppose that the state is H. Lemma 1 ensures that l_t stays between 0 and $\bar{l}$, so the Dominated Convergence Theorem implies that

$$\lim_{t\to\infty} E[l_t] = E[l_\infty].$$

For any $t > 0$, the expectation (at time 0) of l_t is equal to the prior l_0, since

$$\begin{aligned} E[l_t(n)] &= \sum_{n=0}^{t-1} l_t(n)\pi_t(n,H) \\ &= \sum_{n=0}^{t-1} \left[\frac{\pi_t(n,L)}{\pi_t(n,H)}\, l_0\right] \pi_t(n,H) \\ &= \sum_{n=0}^{t-1} \pi_t(n,L) l_0 = l_0. \end{aligned}$$

Thus, $E[l_\infty] = l_0$. Since l_∞ is almost surely either 0 or $\bar{l}$, convergence to $\bar{l}$ must have probability $l_0/\bar{l}$.

PROOF OF PROPOSITION 5:

The consumers' decision rule in equation (5) is continuous in the social likelihood ratio l_t, so convergence of l_t implies convergence in actions. The proposition then follows immediately from Lemma 4.

PROOF OF PROPOSITION 6:

From Proposition 5, the probability a herd on N occurs and the fraction of consumers who purchase λ_t converges to 0 is

$$\text{(A3)} \qquad \Pr\{l_\infty = \bar{l}\} = \frac{l_0}{\bar{l}} = l_0 \frac{c - \varphi(c)[1 - F_U(p)]}{\varphi(c)[1 - F_U(p - H)] - c},$$

where $\varphi(c) \equiv (G - c)\frac{\bar{\sigma}}{1-\bar{\sigma}} + c > 0$. Note that $\bar{l} > 0$ implies that both the numerator and the denominator in equation (A3) are positive, facts that will be used repeatedly below.

Differentiating equation (A3) with respect to H yields

$$\text{(A4)} \qquad \frac{\partial \Pr\{l_\infty = \bar{l}\}}{\partial H} = l_0 \frac{(c - \varphi(c)[1 - F_U(p)])\varphi(c) f_U(p - H)}{(\varphi(c)[1 - F_U(p - H)] - c)^2}$$

$$= l_0 \Pr\{l_\infty = \bar{l}\} \frac{\varphi(c) f_U(p - H)}{\varphi(c)[1 - F_U(p - H)] - c} > 0.$$

Differentiating equation (A3) with respect to c yields

$$\frac{\partial \Pr\{l_\infty = \bar{l}\}}{\partial c} = \left(\frac{l_0}{(\varphi(c)[1 - F_U(p - H)] - c)^2}\right)$$

$$\times \left\{(\varphi(c)[1 - F_U(p - H)] - c)[1 - \varphi'(c)[1 - F_U(p)]\right.$$

$$\left. + [c - \varphi(c)(1 - F_U(p))][1 - \varphi'(c)(1 - F_U(p - H))]\right\}$$

$$= l_0 \frac{[F_U(p) - F_U(p - H)][\varphi(c) - c\varphi'(c)]}{(\varphi(c)[1 - F_U(p - H)] - c)^2} > 0,$$

where the last inequality follows from the fact that $\varphi'(c) = 1 - \frac{\bar{\sigma}}{1-\bar{\sigma}} < 0$.

Finally, differentiating equation (A3) with respect to p yields

$$\frac{\partial \Pr\{l_\infty = \bar{l}\}}{\partial p} = \left(\frac{l_0}{(\varphi(c)[1 - F_U(p - H)] - c)^2}\right)$$

$$\times \left\{[\varphi(c)[(1 - F_U(p - H))] - c]\varphi(c)\, f_U(p)\right.$$

$$\left. + [c - \varphi(c)(1 - F_U(p))]\,\varphi(c)\, f_U(p - H)\right\} > 0.$$

PROOF OF PROPOSITION 7:

Differentiating equation (A4) with respect to H yields

$$\frac{\partial^2 \Pr\{l_\infty = \bar{l}\}}{\partial c \partial H} = l_0[\varphi(c) - c\varphi'(c)]\frac{\partial}{\partial H}\left\{\frac{F_U(p) - F_U(p - H)}{(\varphi(c)[1 - F_U(p - H)] - c)^2}\right\}$$

$$= \frac{l_0[\varphi(c) - c\varphi'(c)]}{(\varphi(c)[1 - F_U(p - H)] - c)^4}$$

$$\times \{(\varphi(c)[1 - F_U(p - H)] - c)^2 f_U(p - H)$$

$$- (F_U(p) - F_U(p - H))2(\varphi(c)[1 - F_U(p - H)] - c)$$

$$\times \varphi(c) f_U(p - H)\}$$

$$= \frac{l_0[\varphi(c) - c\varphi'(c)] f_U(p - H)}{(\varphi(c)[1 - F_U(p - H)] - c)^3}$$

$$\times \{(\varphi(c)[1 - F_U(p - H)] - c)$$

$$- (F_U(p) - F_U(p - H))2(\varphi(c)\}$$

$$= \frac{l_0[\varphi(c) - c\varphi'(c)] f_U(p - H)}{(\varphi(c)[1 - F_U(p - H)] - c)^3}$$

$$\times \{(\varphi(c)[1 - F_U(p - H)] - c)$$

$$- \varphi(c)[F_U(p) - F_U(p - H)]\} < 0,$$

where the last inequality follows because $\varphi(c) > 0$, $\varphi'(c) < 0$, and both the numerator and the denominator of equation (A3) are positive.

REFERENCES

Acemoglu, Daron, Munther Dahleh, Ilan Lobel, and Asuman Ozdaglar. 2009. "Bayesian Learning in Social Networks." Unpublished. http://inber.org/papers/w14040.

Aghion, Philippe, Patrick Bolton, Christopher Harris, and Bruno Jullien. 1991. "Optimal Learning by Experimentation." *Review of Economic Studies*, 58(4): 621–54.

Ali, S. Nageeb. 2011. "Learning Self-Control." *Quarterly Journal of Economics*, 126(2): 857–93.

Anderson, Lisa R., and Charles A. Holt. 1997. "Information Cascades in the Laboratory." *American Economic Review*, 87(5): 847–62.

Bala, Venkatesh, and Sanjeev Goyal. 2001. "Conformism and Diversity under Social Learning." *Economic Theory*, 17(1): 101–20.

Banerjee, Abhijit V. 1992. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics*, 107(3): 797–817.

Banerjee, Abhijit, and Drew Fudenberg. 2004. "Word-of-Mouth Learning." *Games and Economic Behavior*, 46(1): 1–22.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change in Informational Cascades." *Journal of Political Economy*, 100(5): 992–1026.

Brynjolfsson, Erik, Yu Jeffrey Hu, and Duncan Simester. 2011. "Goodbye Pareto Principle, Hello Long Tail: The Effect of Search Costs on the Concentration of Product Sales." *Management Science*. 57(8): 1373–1386.

Burguet, Roberto, and Xavier Vives. 2000. "Social Learning and Costly Information Acquisition." *Economic Theory*, 15(1): 185–205.

Cai, Hongbin, Yuyu Chen, and Hanming Fang. 2009. "Observational Learning: Evidence from a Randomized Natural Field Experiment." *American Economic Review*, 99(3): 864–82.

Callander, Steven, and Johannes Horner. 2009. "The Wisdom of the Minority." *Journal of Economic Theory*, 144(4): 1421–39

Çelen, Bogachan, and Shachar Kariv. 2004a. "Distinguishing Informational Cascades from Herd Behavior in the Laboratory." *American Economic Review*, 94(3): 484–98.

Çelen, Bogachan, and Shachar Kariv. 2004b. "Observational Learning under Imperfect Information." *Games and Economic Behavior*, 47(1): 72–86.

Chevalier, Judith A., and Dina Mayzlin. 2006. "The Effect of Word of Mouth on Sales: Online Book Reviews." *Journal of Marketing Research*, 43(3): 345–54.

DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel. 2003. "Persuasion Bias, Social Influence, and Unidimensional Opinions." *Quarterly Journal of Economics*, 118(3): 909–68.

Eliaz, Kfir, and Ran Spiegler. 2011. "Consideration Sets and Competitive Marketing." *Review of Economic Studies*, 78(1): 235–62.

Ellison, Glenn, and Drew Fudenberg. 1993. "Rules of Thumb for Social Learning." *Journal of Political Economy*, 101(4): 612–43.

Ellison, Glenn, and Drew Fudenberg. 1995. "Word-of-Mouth Communication and Social Learning." *Quarterly Journal of Economics*, 110(1): 93–125.

Goeree, Jacob K., Thomas Palfrey, and Brian W. Rogers. 2006. "Social Learning with Private and Common Values." *Economic Theory*, 28(2): 245–64.

Golub, Benjamin, and Matthew O. Jackson. 2010. "Naive Learning in Social Networks and the Wisdom of Crowds." *American Economic Journal: Microeconomics*, 2(1): 112–49.

Hendricks, Ken, and Alan Sorensen. 2009. "Information and the Skewness of Music Sales." *Journal of Political Economy*, 117(2): 324–69.

Hendricks, Kenneth, Alan Sorensen, and Thomas Wiseman. 2012. "Observational Learning and Demand for Search Goods: Dataset." *American Economic Journal: Microeconomics*. http://dx.doi.org/10.1257/mic.4.1.1.

Herrera, Helios, and Johannes Horner. 2009. "Biased Social Learning." Yale University Cowles Foundation Discussion Paper 1738.

Monzón, Ignacio, and Michael Rapp. 2011. "Observational Learning with Position Uncertainty." http://carloalberto.org/assets/working-papers/no.206.pdf.

Nelson, Phillip. 1970. "Information and Consumer Behavior." *Journal of Political Economy*, 78(2): 311–29.

Orphanides, Athanasios, and David Zervos. 1995. "Rational Addiction with Learning and Regret." *Journal of Political Economy*, 103(4): 739–58.

Pastine, Ivan, and Tuvana Pastine, 2006. "Signal Accuracy and Informational Cascades." Centre for Economic Policy and Research (CEPR) Discussion Paper Series 5219.

Rosen, Sherwin. 1981. "The Economics of Superstars." *American Economic Review*, 71(5): 845–58.

Salganik, Matthew J. 2007. "Success and Failure in Cultural Markets." PhD diss. Columbia University.

Salganik, Matthew J., Peter Sheridan Dodds, and Duncan J. Watts. 2006. "Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market." *Science*, 311(5762): 854–56.

Smith, Lones. 1991. "Error Persistence, and Experimental versus Observational Learning." Unpublished.

Smith, Lones, and Peter Sorensen. 2000. "Pathological Outcomes of Observational Learning." *Econometrica*, 68(2): 371–98.

Smith, Lones, and Peter Sorensen. 2008a. "Informational Herding and Optimal Experimentation." Unpublished.

Smith, Lones, and Peter Sorensen. 2008b. "Rational Social Learning by Random Sampling." http://lonessmith.com/sites/default/files/rational.pdf.

Sorensen, Alan T. 2007. "Bestseller Lists and Product Variety." *Journal of Industrial Economics*, 55(4): 715–38.

Tucker, Catherine, and Juanjuan Zhang. 2007. "Long Tail or Steep Tail? A Field Investigation into How Online Popularity Information Affects the Distribution of Consumer Choices." MIT Sloan School of Management Working Paper 4655-07.

Vives, Xavier. 1997. "Learning from Others: A Welfare Analysis." *Games and Economic Behavior*, 20(2): 177–200.

Wiseman, Thomas. 2008. "Disagreement Leads to Complete Learning: Sequential Choice with Continuous Types." *Economics Letters*, 100(1): 53–55.