Learning from Market Share when Consumers are Rationally Inattentive*

Andrew Caplin, John Leahy, and Filip Matějka
New York University, N.B.E.R. and CERGE-EI

September 2014

*We thank Alessandro Pavan, Chris Tonetti, and Xavier Vives for helpful discussions.

1 Introduction

Social learning is ubiquitous. Booksellers lure readers with the label "New York Times Best Seller." Restaurant goers consult Zagat and Yelp for recommendations. Vacationers consult TripAdvisor.com. In most such settings there are many choice options and agents are heterogeneous. Consumers want to choose goods that match their own preferences. This gives rise to a particularly interesting interplay between social and private learning. After all, if private learning was costless, social learning would serve little obvious purpose.

The themes of individual heterogeneity and multiple options also arise in applied work in such areas as technology adoption (Munshi [2003]), audience dynamics (Moretti [2010]), and market share dynamics (Sorenson [2006]). While applications involve heterogeneity, the theoretical literature remains largely focused on cases with limited consumer heterogeneity and few choice options, following on the early work of Bikhchandani, Hirshleifer, and Welch (1992), Caplin and Leahy (1994, 1998), Chamley and Gale (1994). Hence there is at present a considerable gap between theory and application.

We introduce a general model of social and private learning with unrestricted individual heterogeneity and an unrestricted choice set. A continuum of agents choose from a finite set of options. Agents differ in their type. Different types prefer different goods. Agents observe data on the distribution of past choices. This market share data is assumed to be freely available, as in Conlisk and Smallwood (1979), Becker (1981), and Caminal and Vives (1996).[1] In addition to this public information, individuals can engage in incremental private learning to learn their own type, albeit at some cost.

Our first results concern convergence properties of the model. We establish a general convergence result without placing any restriction on the form of the information cost

[1] These are among the few papers that explicitly model learning from market share. Smallwood and Conlisk (1979) study the dynamics of a market with non-rational consumers who use adaptive strategies in which the probability of purchasing a good depends on its market share. The idea is that consumers tend to imitate other consumers. Becker (1991) assumes that individual demand for a product depends on market demand. He justifies this reduced form as representing either learning or a preference for conformity. Caminal and Vives (1996) is the closest in spirit to our paper. They construct a model in which homogeneous consumers choose among products of heterogeneous quality. Consumers receive private signals on quality and observe market shares. They show that as time passes, market shares reveal true qualities.

There is another literature in which market share plays an indirect role. In this literature, agents meet other agents randomly and exchange information. Market share affects the types of agent that any individual is likely to meet. Ellison and Fudenberg (1995) ask whether word-of-mouth communication aggregates information in an environment in an environment with exogenously specified rules of behavior. Burnside, Eichenbaum, and Rebelo (2013) study asset bubbles in a model in which "optimistic" agents may "infect" other agents through bilateral meetings.

function. Over time, market shares converge to steady state levels, with the new entrants reproducing the same market shares after undertaking optimal private learning.

To get sharp results we then specialize to the case of Shannon costs, as in the rational inattention model (Sims [1998,2003], Matějka and McKay [2014], Caplin, Dean and Leahy [2014]). We show that in this case prior beliefs about the distribution of tastes have a significant effect on long run market shares. However this effect operates only through one channel, which is the set of options that is chosen in the limit. Given this limit set, one can identify market shares without any reference to initial beliefs: they are the market shares associated with agents knowing the true distribution of tastes, regardless of whether or not participants in the market are actually able to pin down the population distribution for sure. This "as if" result has implications for long run market efficiency. A natural question is whether or not the market achieves in the long run a Pareto optimum in light of the costs of information acquisition and the ability only to make inferences from aggregate market information. Our "as if" result implies that the only possible failure of optimality arises when prior beliefs about the population distribution result in what would in fact be popular options remaining unchosen. Choices are optimal given the steady state choice set. Note that this particular failure of optimality is hard to spot in an actual functioning market place, since unchosen options are by definition absent from the market-place.

These results reveal the importance of identifying which options are chosen in the long run. As in Matějka and Sims (2011), we find that it only a small subset of available goods is typically chosen in steady state. We also find that limit market shares can greatly exaggerate the market share of popular options as individual choice tends to conflate private and public preferences. One result of this is that majority types tend to do better than minority types. Minority types may find themselves choosing options preferred by the common types much to their detriment.

The interplay between private and social learning is important not only for successful predictions of market outcomes, but also for inference of agents' preferences from them. We show that while the probability of choosing any single option conflates an agents own preference with the preferences of others, type-specific market share data reveals the choice rankings for each type of agent. The reason is that individual learning skews choice probabilities in the direction of the individually optimal choices. With the Shannon cost function, a simple formula summarizes the skewing of choice and fully reveals the ranking of choices by type.

We present a series of applications. Our first application is to consumer choice. There is a long tradition in industrial organization of inferring preferences from market shares (McFadden [1974], Berry, Levensohn and Pakes [1995]). Such inference is not straight forward in our setting, since market share conflates the demands of many types of consumer and exaggerates the influence of popular types. We show, however, that there is one instance in which inference is relatively straight forward. If one has information on market share by type, we show that one can recover type-specific utilities (up to a constant that depends on the cost of learning) by comparing market shares across types. Each type is more likely to choose their type appropriate option and the strength of this tendency depends directly on the type-specific payoffs and costs of inference. Inference in this case requires the observer to have more precise data on market share than the individual decision makers. In this era of big data, however, it is not difficult to imagine situations in which this would be the case.

Our second application is to political choice. As in industrial organization, there is a long tradition in political economy of inferring voter preferences from vote shares and polling data. It is well understood that voters are only partially informed so that how voters actually vote may differ from how they would vote if they had full information. More informed voters vote differently than ill informed voters after controlling for observable characteristics such as age, race, education and party affiliation. Bartles [1996] and Delli Carpini and Keeter [1996] attempt to uncover the "true" distribution preferences by projecting the votes of better informed voters on less informed voters. In our model, this approach is conceptually correct so long as the more informed voters are in fact fully informed. In all other cases even the choices of the informed voters are biased towards the most popular choice. We can also model the effect of information on the outcome of an election. We show that the presence of informed voters generally reduces the welfare of uninformed voters, as informed voters only affect the outcome of the election if their preferences differ from those of uninformed voters, and thus when they skew the decision of the uninformed towards an undesirable outcome for them.

Our final application is to product regulation. Smallwood and Conlisk (1979) argue that in a model in which market share affects demand product regulation may reduce welfare. The argument is simple. In their model, product regulation is modelled as a minimum quality standard. Product regulation therefore improves the quality of low quality goods. This has the effect of increasing their market share. The resulting shift in demand to low quality goods reduces welfare.

Section 2 introduces the model, as well as the general convergence result for arbitrary information cost functions. Section 3 introduces the Shannon model. This is used to establish the "as if" result and the results on long run welfare. We then investigate the importance of unchosen options and present comparative statics for the two-choice two-type case. Section 4 presents applications. Section 5 concludes

2 Model

We model a dynamic market in which successive generations of agents make a single choice from a fixed set of available options. The generational structure of the model makes it best suited to choices that are made infrequently. Agents differ in their type and choices differ in their type specific payoffs. Agents do not know their own type. They observe the distribution of past choices, and from this distribution they infer the distribution of types in the economy. This forms their prior. Agents then privately gather further information on their type. This is costly, as there needs to be some impediment to private learning for social learning to influence behavior as it appears so often to do in practice. The limits on private learning imply that mistakes are made, so that an individual sometimes ends up with a choice that they like less than available alternatives. The precise pattern of these mistakes reflects the costs and benefits of assessing the unknown state of the world, which is the consumer's type. This pattern therefore contains information regarding the distribution of types. This is the reason that market share is informative, which in turn accounts for market dynamics.

2.1 Model structure

Time is discrete and indexed by $t \in \{0, 1, \dots\}$. There is a fixed finite set $A = \{1..N\}$ of options in a particular market. Each period a continuum of new agents enters the market. Upon entry each observes the fraction of agents that made each choice in each prior period, undertakes optimal private learning, and then make a once-off choice from A, at which point they exit the market never to return.

To capture heterogeneity, we assume that agents are of a finite number of distinct preference types $\omega \in \Omega$, and there is an underlying utility function,

$$u : \Omega \times A \to \mathbb{R}.$$

Let $g_{pop} \in \Delta(\Omega)$ denote density of preference types in the population and $g_{pop}(\omega)$ the share of type ω agents. g_{pop} is fixed and does not change over time. We normalize the total population of agents to 1. We place no restriction on the form of the heterogeneity in utility.

Buyers new to the market do not know their types which means that they do not know the utilities from selecting different options. The only information freely available to agents in period $t = 0$ is their common prior G^0, which comprises a probability measure over distributions in $\Delta(\Omega)$. Since Ω is finite, $\Delta(\Omega)$ is isomorphic to the $|\Omega| - 1$ dimensional simplex in $R^{|\Omega|}$. So that marginal distributions of G are well defined, we will assume that G has a continuous density on this simplex. It will be useful in what follows to define $\Gamma^0 \equiv \text{supp}(G^0) \subseteq \Delta(\Omega)$ as the set of possible distributions. We require that $g_{pop} \in \text{int}(\Gamma^0)$.[2]

Given G^0 and Γ^0, we can calculate agents' prior beliefs over preference types as the expected distribution of types:

$$\mu^0(\omega) = \frac{1}{G(\Gamma^0)} \int_{g \in \Gamma^0} g(\omega) dG$$

In addition to relying on the prior, each agent can process additional costly information about ω, and thus about the utilities of available options, by exploring preferences over the offered options in more detail. We do not specify this process: it might involve personal examination of a product such as test driving a car or a visit to a store; it might involve a detailed reading of the product reviews in Amazon.com or yelp.com; or it might involve discussions with friends, colleagues, or other people that the agent regards as similar to him or herself. At this point all that we need is that individual learning leads to a type-dependent choice function of the form

$$P(i, \omega) = \Pr\{i \in A | \omega \in \Omega\}$$

where the realized choices i are independent across agents of the same type ω. Caplin and Dean [2014] show that a broad class of learning models generate behavior of this type. Intuitively, if learning is expensive, $P(i, \omega)$ will not vary much across types ω and will be larger for choices i that appear desirable ex ante. As learning becomes less expensive, $P(i, \omega)$ will conform more and more closely with choices that are relatively desirable for type ω. We assume that all agents understand the mapping from priors μ to type-dependent

[2] $\text{supp}(G^0)$ is the set of $g \in \Delta(\Omega)$ such that every open neighborhood of g has positive measure. This assumption implies that the density of G is strictly positive at g_{pop}.

choice

$$P(\mu) = \{P(i, \omega|\mu) | i \in A, \omega \in \Omega, \mu \in \Delta(\Omega)\}.$$

In the next section we will place some structure on the $P(i, \omega)$'s by imposing a Shannon information cost function.

2.2 Recursive Learning From Market Share

Agents who enter the market in periods $t > 0$ can learn about their type in part by observing all past market shares. Realized market share provides information about the distribution of types in the economy. This form of learning from market share involves winnowing down the set Γ^0 by eliminating distributions that are inconsistent with any prior observed market shares as now specified recursively.

In period t, taking as given G and Γ^t, we can calculate agents' prior beliefs over preference types as the expected distribution of types conditional on the set of possible densities Γ^t:

$$\mu_{\Gamma_t}(\omega) = \frac{1}{G(\Gamma^t)} \int_{g \in \Gamma^t} g(\omega) dG$$

With $P(i, \omega|\mu_\Gamma)$ specified from the general rational inattention model outlined above, we can generate the realized market shares since the learning is conditionally independent across agents:

$$M(i|\mu_\Gamma, g_{pop}) = \sum_{\omega \in \Omega} g_{pop}(\omega) P(i, \omega|\mu_\Gamma). \tag{1}$$

Agents born in the ensuing period $t + 1$ observe period t aggregate market shares. These observed market shares eliminate from consideration some of the ex ante possible type distributions, in particular those in the elimination set $E(g_{pop}, \Gamma)$,[3]

$$E(g_{pop}, \Gamma) = \left\{ g \in \Gamma \,\middle|\, \sum_{\omega \in \Omega} g(\omega) P(i, \omega|\mu_\Gamma) \neq M(i|\mu_\Gamma, g_{pop}) \right\}.$$

Defining $E^t \equiv E(g_{pop}, \Gamma^t)$ as the period t elimination set enables updating on possible type distributions:

$$\Gamma^{t+1} = \Gamma^t \backslash E^t.$$

[3]One of the things that differentiates our approach from other models of learning from market share such as Smallwood and Conlisk (1979) is that our agents do not naively treat market share as the prior over acts but use market share to construct the prior over types.

Note that $g_{pop} \in \Gamma^t$ for all t because it is always consistent with observed behavior and can never be ruled out. Given Γ^{t+1} period $t+1$ proceeds in a manner similar to period t completing the recursion.

2.3 Orthogonality and Convergence

Given Γ^t, Γ^{t+1} has a simple structure. According to (1), all $g \in \Gamma^{t+1}$ were elements of Γ^t and generated the same market shares in period t as g_{pop}. It follows that $g \in \Gamma^{t+1}$ if $g \in \Gamma^t$ and for all $i \in A$

$$\sum_{\omega \in \Omega} [g(\omega) - g_{pop}(\omega)]\, P(i, \omega | \mu_{\Gamma^t}) = 0. \tag{2}$$

This orthogonality condition enables us to characterize updating precisely as a function of the chosen options $A(\mu_{\Gamma^t})$ and probabilities $P(i, \omega | \mu_{\Gamma^t})$.

A key question concerns whether or not population settles down. A set $\bar{\Gamma} \subseteq \Delta(\Omega)$ is a steady state of the model if $\Gamma^t = \bar{\Gamma}$ implies that $\Gamma^{t+1} = \bar{\Gamma}$. A steady state set of possible beliefs, $\bar{\Gamma}$, generates a steady state measure over possible distributions of beliefs $\bar{G}$, which is the distribution G conditioned on $\bar{\Gamma}$, as well a steady state prior $\bar{\mu}$, which is the expectation of $g \in G$ conditional on $\bar{\Gamma}$. This, in turn, pins down the steady state choices $P(i, \omega | \bar{\mu})$ and market shares $M(i | \bar{\mu}, g_{pop})$. The following proposition establishes the existence of a steady state $\bar{\Gamma}$ and that the market converges to steady state in a finite number of periods. The proofs of all propositions are contained in the appendix.

Proposition 1 There exists $\bar{\Gamma}$ such that $\Gamma^t \to \bar{\Gamma}$. Moreover $\Gamma^{|\Omega|} = \bar{\Gamma}$.

The idea behind the proof is that since Ω is finite each $g \in \Delta(\Omega)$ exists as a point in the $|\Omega - 1|$ dimensional simplex in $R^{|\Omega|}$. For each choice $i \in A$, there is an orthogonality condition (2). Each orthogonality condition defines a $|\Omega - 1|$ dimensional hyperplane X_i^t in $R^{|\Omega|}$. Given that A is finite there are $|A|$ such hyperplanes. Γ^{t+1} is equal to the intersection of Γ^t and these $|A|$ hyperplanes::

$$\Gamma^{t+1} = \Gamma^t \cap (\cap_{i \in A} X_i^t) = \Gamma^0 \cap (\cap_{i \in A} X_i^0) \cap \ldots (\cap_{i \in A} X_i^t)$$

Each additional orthogonality condition either reduces the dimension of Γ^{t+1} relative to Γ^t or does not. If none of the period t conditions reduce the dimension of Γ^t, then a steady state has been reached and the model has converged. The finite dimension of Ω guarantees that the model converges in a finite number of periods. In fact, if we have as many options as types and given the prior, the vectors $P(i|\omega, \mu^0)$ are independent, convergence will be immediate. If the dimension of Γ^t falls to zero, then $\Gamma^t = g_{pop}$. Otherwise learning is incomplete. In general complete learning cannot be guaranteed as we show using the Shannon cost function in section 4.

What drives convergence is a disconnect between the expected probability of choosing and option and the observed market shares. To see this note that, in steady state, all distributions $g \in \bar{\Gamma}$ must give rise to the observed market shares. Otherwise it would be possible to eliminate some of them and further reduce $\bar{\Gamma}$. It follows that the observed market share of each good i is equal to the expected probability of choosing good i given the steady state prior prior (recall that we have normalized the total population to one). In particular, for all $i \in \bar{A}$

$$P(i|\bar{\mu}) \equiv \sum_{\omega\in\Omega} \bar{\mu}(\omega) P(i|\omega, \bar{\mu}) = \left[\frac{1}{G(\bar{\Gamma})} \int_{g\in\bar{\Gamma}} \sum_{\omega\in\Omega} g(\omega) P(i, \omega|\bar{\mu}) dG \right] = M(i|\bar{\mu})$$

where the second equality follows from the definition of $\bar{\mu}$ and Fubini's theorem and the last equality follows from the steady state orthogonality condition, $\sum_{\omega\in\Omega} g(\omega) P(i, \omega|\bar{\mu}) = M(i|\bar{\mu})$ for all $g \in \bar{\Gamma}$.

Proposition 2 Given a steady state set of goods with positive market share $\bar{A} \subseteq A$, the steady state market shares $M(i)$ are equal to the expected choice probabilities:

$$M(i|\mu) = P(i|\mu) \tag{3}$$

In order to say more about the behavioral and welfare properties of the model we need to place some structure on the $P(i, \omega)$'s. In the next section we model the cost of information acquisition as a function of the reduction in entropy as in Sims [1998,2003].

3 Rational Inattention and State Dependent Stochastic Demand

We follow Matějka and McKay (2014) and Caplin, Dean and Leahy (2014) to describe the derivation of the type-dependent stochastic choice map $P(i, \omega|\mu)$. The agent is a Bayesian expected utility maximizer who is rationally inattentive. Given the prior $\mu \in \Delta(\Omega)$, the agent chooses $\{P(i,\omega)\}_{i\in A,\omega\in\Omega}$, the probability of taking each option i conditional being each type ω, in order to maximize expected utility. The agent can gather information at a cost and thereby make the $P(i,\omega)$ more responsive to his or her true type ω. There is a cost to increasing the correlation between $P(i,\omega)$ and the state which on the reduction in entropy. The agent maximizes:

$$\begin{aligned} V(\mu, A) &= \max_{\{P(i,\omega)\}_{i\in A,\omega\in\Omega}} \sum_{\omega\in\Omega} \mu(\omega) \left(\sum_{i\in A} P(i,\omega) u(\omega, i) \right) \\ &\quad -\lambda \left[\sum_{\omega\in\Omega} \mu(\omega) \left(\sum_{i\in A} P(i,\omega) \ln P(i,\omega) \right) - \sum_{i\in A} P(i) \ln P(i) \right] \end{aligned} \tag{4}$$

The first term on the right-hand side is the expected utility of the strategy $\{P(i,\omega)\}_{i\in A,\omega\in\Omega}$. The second term is cost of information acquisition. $P(i)$ is the probability of option i conditional on the prior μ :

$$P(i) = \sum_{\omega\in\Omega} \mu(\omega) P(i,\omega)$$

The second term in square brackets is therefore the entropy of a strategy that is independent of type. It costs nothing to follow this strategy for all ω. Costs are incurred if one chooses $P(i,\omega)$ that are more highly correlated with ω. The first term in square brackets is the entropy of such a type-contingent strategy. λ is the marginal cost of entropy reduction.[4]

Matějka and McKay (2014) and Caplin, Dean and Leahy (2014) characterize the resulting pattern of state dependent stochastic choice

$$P(i,\omega|\mu) = \frac{P(i|\mu)\exp(u(\omega,i)/\lambda)}{\sum_{j\in A} P(j|\mu)\exp(u(\omega,j)/\lambda)}, \tag{5}$$

[4] To see that this is a cost of information acquisition, note that Bayes rule implies that $\gamma^i(\omega)$, the posterior probability of state ω conditional on the choice of action, is equal to $P(i,\omega)\mu(\omega)/P(i)$. Hence the term in square brackets can rewritten as

$$\sum_{i\in A} P(i) \sum_{\omega\in\Omega} \gamma^i(\omega) \ln \gamma^i(\omega) - \sum_{\omega\in\Omega} \mu(\omega) \ln \mu(\omega)$$

In this formulation the second term is the entropy of the prior, whereas the first term is the expected entropy of the posterior.

where for each option $i \in A$, the $\{P(j)\}_{j \in A}$ satisfy

$$\sum_{\omega \in \Omega} \mu(\omega) \left\{ \frac{\exp(u(\omega, i)/\lambda)}{\sum_{j \in A} P(j|\mu) \exp(u(\omega, j)/\lambda)} \right\} \leq 1 \qquad \forall i, \tag{6}$$

with equality if $P(i) > 0$. The inequalities (6) are the necessary and sufficient conditions for a solution to (4). Note that once one recovers the $\{P(j)\}_{j \in A}$ from (6), the $P(i, \omega)$ follow directly from (5).

Given the concavity of the problem, the optimal choice exists. The solution is unique if the vectors $\exp(u(\omega, i)/\lambda) \in \mathbb{R}_{++}^{|\Omega|}$ for $i \in A$ are affinely independent. We assume that this is the case.

Axiom 1.

Assumption *The vectors* $\exp(u(\omega, i)/\lambda) \in \mathbb{R}_{++}^{|\Omega|}$ *for* $i \in A$ *are affinely independent:*

$$\sum_{i \in A} \alpha(i) \exp(u(\omega, i)/\lambda) = 0 \Longrightarrow \alpha(i) \equiv 0.$$

3.1 An "As If" Result

Given the form of type-dependent stochastic choice in (5) we can show that in steady state agents behave as if they know the true distribution of types g_{pop} and are choosing from the steady state set of options even though in fact they might be quite uncertain which distribuiton of types is in fact generating the observed market shares. Let $\bar{A} \subseteq A$ denote the set of options with positive market shares in steady state. Recall that Proposition 2 states that steady state choice probabilities are equal to market shares. Using the definition of market share (1) and the optimal policies (5), we have:

$$P(i|\bar{\mu}) = M(i|\bar{\mu}) = \sum_{\omega \in \Omega} g_{pop}(\omega) P(i, \omega|\bar{\mu}) = \sum_{\omega \in \Omega} g_{pop}(\omega) \left\{ \frac{P(i|\bar{\mu}) \exp(u(\omega, i)/\lambda)}{\sum_{j \in A} P(j|\bar{\mu}) \exp(u(\omega, j)/\lambda)} \right\}$$

or dividing both sides by $P(i|\bar{\mu})$,

$$\sum_{\omega \in \Omega} g_{pop}(\omega) \left\{ \frac{\exp(u(\omega, i)/\lambda)}{\sum_{j \in A} P(j|\bar{\mu}) \exp(u(\omega, j)/\lambda)} \right\} = 1 \tag{7}$$

Equation (7), however, is simply a statement of the necessary and sufficient conditions (6) for an optiomal policy over the choice set $\bar{A}$ given the prior g_{pop}. It follows that the $P(i|\bar{\mu})$ are optimal for the prior g_{pop} and the option set $\bar{A}$. Agents act as if they know the true distribution of types.

Corollary 3 In the steady state, $P(j|\bar{\mu})$ satisfy the necessary and sufficient conditions for optimal choice if the prior were g_{pop}. The choice behavior is thus given by the solution to (5)-(4) for $\mu = g_{pop}$ and $A = \bar{A}$.

Several comments are in order. First, the fact that agents act as if they know the true distribution of types in steady state, greatly simplifies the analysis of the model and limits the range of steady state behavior. If one knows the steady state choice set $\bar{A}$, one can always assume that agents know the true distribution of types. Second, market inefficiency takes a very limited form. Individual choice is optimal given the observed set of choices $\bar{A}$. The set of choices $\bar{A}$, however, may not be optimal. It may be the case that the left-hand side of (7) is greater than one for goods in $A/\bar{A}$, in which case the necessary and sufficient conditions (6) would be violated on the larger choice set. In such cases agents would be better off choosing some of the unchosen options.

3.2 Welfare

In our model there are agents of different types who often choose options that they would prefer not to take if they had more information. This would normally make complicate welfare calculations (See Bernheim and Rangel [2009]). Our agents, however, solve a well defined maximization problem (4). $V(\mu, A)$ therefore provides a measure of subjective well being. We can therefore analyze the perceived effect of any policy by studying the response of $V(\mu, A)$. A look at (4) shows that there is a sense in which our agents make interpersonal comparisons of utility. They must imagine the payoff of each option to each type in order to learn optimally about their type.[5]

There is another potential notion of welfare. Since our agents may hold incorrect beliefs,

[5] The optimal policies (5) depend only on the differences in utility across options, $u(i, \omega) - u(j, \omega)$. Agents therefore must be able to evaluate the utility of each type up to an additive constant. Since this constant is a fixed effect tied to the type and does not affect choice it can be ignored in most policy experiments.

a social planner who knew the true population distribution would want to calculate

$$\begin{aligned}\tilde{V}(\mu, g_{pop}, A) &= \max_{\{P(i,\omega)\}_{i\in A,\omega\in\Omega}} \sum_{\omega\in\Omega} g_{pop}(\omega)\left(\sum_{i\in A} P(i,\omega)u(\omega,i)\right) \\ &\quad -\lambda\left[\sum_{\omega\in\Omega}\mu(\omega)\left(\sum_{i\in A} P(i,\omega)\ln P(i,\omega)\right) - \sum_{i\in A} P(i)\ln P(i)\right]\end{aligned}$$

Note that here we retain μ in the information cost as we interpret the learning cost as subjective.[6] Since in steady state choice is made as if the prior were g_{pop}, steady state policy maximizes both V and $\tilde{V}$ on the observed steady state choice set $\bar{A}$. The social planner, however, wish to expand that choice set. Also outside of steady state the social planner would like to nudge the agents in the direction of policies that are optimal for types more common than believed.

3.3 How Many Goods are Chosen?

Market inefficiency occurs only when there are few options chosen. We now present a simple example that illustrates a phenomenon noted by Matejka and Sims (2011) in which few options are indeed chosen. We remove all heterogeneity beyond the distribution of types and we consider the steady state of a class of symmetric models with $\Omega = A = \{1, .., M\}$. Each agent would like to choose the option matched to their type $i = \omega$. The payoffs are:

$$\exp(u^i_\omega/\lambda) = \begin{cases} x(1+\delta) & \text{if } i = \omega; \\ x & \text{if } i \neq \omega; \end{cases} \tag{8}$$

with $x > 0$ and $\delta \geq 0$. Note that $1 + \delta = \exp\left(\frac{u^i_i - u^j_i}{\lambda}\right)$ so that increases in the utility differential or reductions in learning costs are associated with increases in δ. We order

[6] The expected probability of choosing option i, $P(i|\mu)$, may differ from the realized frequency with which the option is actually taken, $M(i|\mu)$, which raises the question of what the agent is actually choosing in (4) and what exactly the information cost represents. Our interpretion of the maximization problem (4) is the following. The individual chooses $\{P(i,\omega|\mu)\}_{i\in A,\omega\in\Omega}$ to maximize the expected payoff net of costs of entropy reduction. The cost of entropy reduction is subjective. The agent chooses the strategy of gathering information that for a given expected payoff minimizes the expected cost of information given the belief about his type. The realized cost can be different since it depends on the true type. The cost is proportional to the expected loss of entropy in moving the agent's prior μ to the posterior $\gamma^i(\omega) = P(i,\omega|\mu)\mu(\omega)/P(i|\mu)$ where the latter follows directly from Bayes rule.

goods according to perceived likelihood with lower indexed types perceived as more likely

$$\mu_\omega \geq \mu_{\omega+1}.$$

The next proposition characterizes the solution to this simple model.

Proposition 4 If $\mu_M > \frac{1}{M+\delta}$ define $K = M$. If $\mu_M < \frac{1}{M+\delta}$, then define $K < M$ as the unique integer such that,

$$\mu_K > \frac{\sum_{\omega=1}^{K} \mu_\omega}{K+\delta} \geq \mu_{K+1} \tag{9}$$

Then the unique solution involves,

$$P^i = \frac{\mu_i(K+\delta) - \sum_{\omega=1}^{K} \mu_\omega}{\delta \sum_{\omega=1}^{K} \mu_\omega} > 0 \tag{10}$$

for $i \leq K$, with $P^i = 0$ for $i > K$.

The proposition characterizes choice as a function of the prior μ and the payoff to the correct option δ. Note x merely scales utility without affecting behavior as choice depends only on δ through $\frac{u_i^i - u_i^j}{\lambda}$, a large value of $\frac{u_i^i - u_i^j}{\lambda}$ being associated with a large value of delta.

Two main phenomenon arise in this setting. The first is that if $\mu_M < \frac{1}{M+\delta}$, $|A| < M$. If δ is small or there are many choices, it takes only a small deviation from uniformity for this condition to hold. This opens the door for inefficient learning. For example if $M = 2$ and $\mu_2 < \frac{1}{1+\delta}$, then $P^1 = 1$ and there is no way for the market to learn how many $\omega = 2$ types there actually are.

Second, information on market share tends to exaggerate demand for "popular" choices. Consider (10) and consider the differences in choice probabilities among options that are chosen. The key observation is that this difference is strictly proportionate to the difference in prior probabilities. Given options $i, j \leq K$,

$$P^i - P^j = (\mu_i - \mu_j) \left[\frac{(K+\delta)}{\delta \sum_{\omega=1}^{K} \mu_\omega} \right].$$

Note that the denominator in the term in square brackets is no higher than δ while the numerator is strictly larger than delta. This term is therefore strictly greater than one. Choice is skewed towards the options with higher prior probability of success. In fact, the unconditional probability of the most likely popular choice P^1 is easily seen to be greater than the the prior probability μ_1, and the probability of the least popular choice P^M is less than the share of type M agents μ_M. The following numerical example illustrates the skewing of choice due to priors

Example Suppose that $\delta = 1$ and that the five most probable states have the following probabilities

$$(\mu_1, \mu_2, \mu_3, \mu_4, \mu_5) = \left(\frac{10}{100}, \frac{9}{100}, \frac{8}{100}, \frac{7}{100}, \frac{6}{100}\right).$$

Since

$$\frac{\sum_{\omega=1}^{4} \mu_\omega}{4+\delta} = \frac{0.34}{5} \in \left(\frac{7}{100}, \frac{6}{100}\right)$$

the cutoff condition in the proposition implies that $K = 4$. The existence of any additional options beyond these most likely five are therefore irrelevant. The first difference condition above implies

$$P^1 - P^2 = P^2 - P^3 = P^3 - P^4 = 0.01 \left[\frac{5}{0.34}\right] = \frac{5}{34}.$$

This together with the requirement that the probabilities sum to one implies

$$(P^1, P^2, P^3, P^4) = (\frac{16}{34}, \frac{11}{34}, \frac{6}{34}, \frac{1}{34}).$$

This illustrates the great twist in favor of the likely more popular option.

The form of type-specific stochastic demand in equation (5) implies that agents of a given type are more likely to choose their preferred choice than are agents in general. Hence not only are more commonly desired choices proportionately more likely to be chosen, but more common types are more likely to make these common choices than the average type. This skewing of choice has obvious welfare implications. Common types tend to do better than uncommon types. For example, we can calculate type specific demand in the example above. Type 1 chooses good one 64% of the time. Type 2 chooses correctly 49% of the time, while types three and four choose correctly 30% and 6% respectively. The remaining 66% never choose the correct option. They would be better off choosing randomly.

3.4 The Two by Two Case

As in trade theory, the two-by-two case reduces the substitution possibilities among options, allowing for a clear illustration of the fundamental forces operating in the model. Suppose that there are two types ν and η and two choices a and b. Suppose that,

$$u(a,\nu) > u(b,\nu) = 0 = u(b,\eta) > u(a,\eta)$$

so that type ν prefers option a and type η prefers option b. Since according to (5) choices depend only on $u(\omega,i) - u(\omega,j)$, normalizing the value of option b to zero for both types is without loss of generality.

Solving (6) assuming that both options are chosen yields the following probability of choosing option a given the prior μ:

$$\tilde{P}(a|\mu) = \frac{\mu(\nu)}{1 - \exp(u(a,\eta)/\lambda)} - \frac{1 - \mu(\nu)}{\exp(u(a,\nu)/\lambda) - 1} \tag{11}$$

If $\tilde{P}(a|\mu) \in [0,1]$ then this equation gives the true choice probabilities and $P(a|\mu) = \tilde{P}(a|\mu)$. The type-dependent choice probabilities follow directly form (5). If instead $\tilde{P}(a|\mu) > 1$, then only option a is chosen and $P(a|\mu) = 1$, while if $\tilde{P}(a|\mu) < 0$, only option b is chosen and $P(a|\mu) = 0$. Hence prior beliefs determine whether or not both options are chosen.

In the two-by-two case, convergence to steady state occurs in one or two periods. If prior beliefs are such that only one option is chosen in period zero, agents learn nothing about the population from market share and no learning takes place. The period 0 choices repeat themselves. If both options are chosen in period 0, (5) implies that type ν are more likely to choose a. Since these probabilities are known, market share perfectly reveals the population distribution and the steady state is reached in period 1.

The model has well-behaved comparative statics. It is immediate from (11) that an increase in $u(i,\omega)$ will increase the probability of choosing i in both states and an increase in $\mu(\nu)$ will cause the probability of choice a to rise. When learning is costless, each agent chooses the option that is best for them. The market share of each choice is then equal to the proportion of agents who prefer that choice. As learning costs rise, the influence of the ex ante optimal choice grows. If $\mu_t u(a,\nu) > -(1-\mu_t)u(a,\eta)$ then option a is more desirable ex ante, and $P(a|\mu)$ rises with λ. Eventually, λ rises so high that $P(a|\mu)$ hits either zero or one and only the ex ante optimal option is taken. Since $M(a|\mu,g) = P(a|\mu)$ in steady state, these comparative statics apply to the steady state market share as well.

4 Applications

4.1 Type specific demand is very informative

It follows immediately from the type-specific choice probabilities (5) that individual choice probabilities of an agent of type ω skew average choice probabilities in the direction of choices preferred by agents of type ω. Optimal private learning leads to choices that are correlated with one's type. The converse is that we may use type-specific choice probabilities to infer agents' preferences. The following proposition captures the idea that because agents tend to choose options that they prefer, market shares will be very informative about preferences.

Proposition 5 In the steady state,

$$\frac{u(\omega,i)-u(\omega,j)}{\lambda}=\log\left(\frac{M(i,\omega)}{M(j,\omega)}\bigg/\frac{M(i)}{M(j)}\right). \tag{12}$$

The proposition follows directly from type-specific stochastic choice (5) and the observation that in steady state the unconditional choice probabilities $P(i)$ are equal to the market shares $M(i)$.

The proposition implies that an outside observer can infer preferences from detailed market share data. Learning from market share skews choice in the direction of popular choices, and for this reason popular choices tend to be popular for all types. That being said, optimal private learning also skews choices in the direction of individual payoffs. One can infer whether an agent of type ω prefers choice i to good j by comparing the frequency by which agents of type ω choose these goods to the average frequency of purchase in the population. Note that since choice depends only on $\frac{u(\omega,i)-u(\omega,j)}{\lambda}$ we can normalize the payoff to one choice to zero for all types ω. Normalizing $u(\omega,j)$ to zero:

$$\frac{u(\omega,i)}{\lambda}=\log\left(\frac{M(i,\omega)}{M(j,\omega)}\bigg/\frac{M(i)}{M(j)}\right).$$

which identifies utility up to the learning cost. Notice that this inference does not depend on knowledge of the agents' beliefs or of the initial prior G. Moreover, the observer does not even need to know the true distribution of types in the population g_{pop}.

It is interesting to note that the term in brackets on the right-hand side of (12) very similar to Balassa's (1965) measure of "revealed" comparative advantage. Belassa measures of comparative advantage as the ratio of the share of country a's exports of good i in the total exports of country a to the share of world exports of good i to total world exports

$$\frac{\frac{\text{exports}_{ai}}{\sum_j \text{exports}_{aj}}}{\frac{\sum_n \text{exports}_{ni}}{\sum_n \sum_j \text{exports}_{nj}}}$$

A country has a comparative advantage in good i, if it exports relatively more of good i than the average country. In our setting, an agent prefers option i, if the agent takes that option relatively more often than does the average agent. In both cases the presence of the average in the denominator controls for common forces that tend to raise exports in all countries or increase the probability of an option across all agents.

4.2 Consumer demand

In this application we associate the choices A with differentiated products and the types Ω with groups of heterogenous consumers. We let the utility of each choice depend on the value of the good and its market price

$$u^i_\omega = \delta^i_\omega - \alpha p_i$$

Note that we assume that agents see and understand prices. The inference problem is one of figuring out which good to purchase.

There is a long tradition in industrial organization of using market shares to infer utility parameters. Prominent examples include McFadden (1974) and Berry, Levensohn and Pakes (1995). Such inference is not straight forward in our setting, since market share conflates the demands of many types of consumer and exaggerates the influence of popular types, thereby biasing the results.

One case in which inference is straight forward is when one can identify the types. One could imagine a large agent who has access to detailed market data and is able to put together type-specific market shares. Such a situation is not far fetched. Agents in the model are learning optimally given limited resources. They focus their attention on matters that concern them directly. A large agent, such as the government, Google, a market research firm such as J. D. Power and Associates or Consumer Reports, or an academic researcher, may have greater incentives to learn and potentially greater resources available for learning. It is very likely that they have greater information on the details of demand and their distribution across agents.

Inference is relatively straight forward given type specific market share. Supposing that the value of good i to an agent of type ω depends on a vector of product characteristics X^i, so that $\delta^i_\omega = X^i\beta_\omega$, Proposition 5 suggests a regression that projects market shares onto product characteristics:

$$\log\left(\frac{M(i,\omega)}{M(j,\omega)}\Big/\frac{M(i)}{M(j)}\right) = (X^i - X^j)\frac{\beta_\omega}{\lambda} - \frac{\alpha}{\lambda}(p_i - p_j) + \xi \tag{13}$$

where ξ is a regression error reflecting measurement error or omitted factors.[7]

The regression is also similar to the standard logit model. There are a number of differences. In logit estimation one normally assumes that good j in the regression is the outside option and that the value of the outside option is zero. This normalization is without loss of generality as only differences in utility matter in the logit model. These assumptions imply that $p_j = 0$ and $\delta^i_\omega = X^i\beta_\omega = 0$ in the logit model, so that (13) depends only on good i. If one had data on type-specific market participation, one could identify the one of the choices with the outside option. Moreover, choice according to (13) depends only on $\frac{u(\omega,i)-u(\omega,j)}{\lambda}$; one may therefore normalize the utility of one option for each type of agent to zero without affecting behavior.

The main difference between (13) and the logit model is that the values of characteristics and the sensitivity of market share to price reflect a combination of utility parameters, β_ω and ω, and the cost of learning, λ. This will not matter in situations in which learning costs are stable. In other cases, changes in learning costs will look like changes in tastes.

Another difference between the logit model and our model of type-specific demand with social learning is the effect of changes in prices on market shares. It is well known that in the logit model the effect of price on market share is completely captured by market share itself:

$$\frac{dM_{Logit}(i)}{dp_i} = -\alpha M_{Logit}(i)[1 - M_{Logit}(i)] \text{ and } \frac{dM_{Logit}(i)}{dp_j} = -\alpha M_{Logit}(i)M_{Logit}(j)]$$

In our setting, both social and private learning alter this relationship. In steady state:

$$\frac{p_i}{M(i,\omega)}\frac{dM(i,\omega)}{dp_i} = -\frac{\alpha}{\lambda}p_i[1-M(i,\omega)]+\frac{p_i}{P(i|g_{pop})}\frac{dP(i|g_{pop})}{dp_i}-\sum_{j\in A}M(j,\omega)\frac{p_i}{P(j|g_{pop})}\frac{dP(j|g_{pop})}{dp_i}$$

[7]While straight forward in principle this regression suffers from all of the endogeneity issues that generally plague demand estimation.

The first term reflects the direct effect of prices on market share. Higher learning costs tend to dampen this effect. The remaining terms refect the effect of prices on the unconditional choice probabilitites. As a rise in p_i tends to reduce $P(i|g_{pop})$ and increase $P(j|g_{pop})$. These terms tend to increase the elasticity of demand relative to the logit benchmark. Out of steady state there is an additional channel by which p_i may affect demand as the change in market shares may provide additional information on the possible distributions of types, thereby further affecting $P(i|\mu)$ through μ.

4.3 Political Choice

In this application, we associate the options A with candidates or positions on policy issues and the types ω with groups of voters. Voters tend to know something about their type. They know, for example, if they are Democrat or Republican and polling information is often broken down by broad categories. We therefore partition the type space into subsets Ω^k with $\Omega = \cup_k \Omega^k$. Agents know which subset they belong to but not their type within the subset. Agents of type $\omega \in \Omega^k$ have a prior μ^k over types in Ω^k that they derive from past vote share or polling data. For each $\omega \in \Omega$ and each candidate or policy position $i \in A$, there is a payoff $u(i, \omega)$ to having that candidate elected or that policy enacted. Our behavioral assumption is that agents vote and respond to surveys in order to maximize this expected payoff. We ignore issues related to strategic voting or to the misrepresentation of preferences. With these assumptions, the model translates into the political setting.

Note that it does not change anything substantive in the model or the solution to allow the set of choices A to change from election to election as the set of candidates changes. Choice in any given period t depends on the current priors μ_t^k and the current choice set A_t. The only difference is that changing the choice set may delay convergence, as beliefs by become stuck in some region for a while until the choice set changes in a way that reveals more information. We could also delay convergence by allowing the distribution of types to change over time as the preferences of voters changes.

One new phenomenon that arises in this setting is that we can study the outcome of the election. One implication of the model is that popular candidates tend to win elections: if there are more type ω voters then candidates preferred by type ω voters tend to do better. Slightly more striking is the role of the prior. Given the form of type-specific stochastic choice (5), $P(i, \omega)$ is increasing in $P(i)$, so that agents of each type are more likely to vote for a candidate if they believe that others are more likely to be voting for that candidate. Perceived popularity is as important as popularity.

It is common practice to infer political preferences from opinion polls and vote share. It is also widely recognized that this practice is problematic. Bartels (1996) and Delli Carpini and Keeter (1996) show that the expressed opinions and voting behavior of informed voters differs greatly form those of uninformed voters even after controlling for observable differences such as age, gender and education. These authors distinguish between choice and true preference, associating true preference with a hypothetical choice under full information . They then attempt to reconstruct the distribution of true preferences by projecting the choice behavior of informed voters of each observable type on the set of uninformed voters. This exercise sometimes shifts the results with the true preference favoring a different candidate or policy proposal than the poll (See also Althaus (1998)). The underlying assumption is that the informed are fully informed and that there is no bias in their voting behavior or opinions. If the informed have $\lambda = 0$, then according to our model this assumption would be justified and, absent other measurement issues, this approach would recover g_{pop}. If $\lambda > 0$, however, then vote shares within a class Ω^k will tend to exaggerate preferences for candidates that are popular within that class even among the better informed. For example, if Ω^k were the set of types who identify themselves as Democrats and we assume that a majority of Democrats prefer the Democratic candidate, the Democratic candidates vote share among Democrats should exceed the fraction of Democrats who would prefer the candidate under full information. The flip side is the minority types within the class Ω^k do badly; their preferred candidates tend to receive vote shares that understate their support under full information.

In many settings introducing more informed agents improves market performance. In this political science application, introducing more informed voters can only make uninformed agents worse off. The following example indicates that all else equal outcomes of elections tend to reflect the preferences of types with low learning costs, as the choices of these types tend to correlate more closely with their preferences.

Example Suppose that there are two general types of agent Ω^1 and Ω^2 and that within each subset there are two types of voter ω_1 and ω_2. Suppose that type ω_1 agents would prefer candidate i if they had full information, while type ω_2 agents would prefer candidate j. So that mere numbers do not determine the result of the election. Suppose that the two groups are symmetric: half of the population is in subgroup Ω^1 and half in subgroup Ω^2. Suppose further that the realization of the types within the groups are symmetric: within each group a fraction $\psi > 1/2$ are type ω_1 or a fraction ψ are type ω_2. These two type distributions occur with equal probability and are independent across the subgroups Ω^i. What distinguishes the subgroups is the agents' learning costs: members of Ω^1 have lower learning costs than members of Ω^2. It is easy to see that a reduction in λ, increases the vote share of the candidate preferred by the subgroup. Let V denote the vote share of the candidate preferred by the majority of the subgroup, then

$$\frac{dV}{d\lambda} \sim -(2\psi - 1)(u - v)V(1 - V)$$

where u is the payoff to the preferred candidate and v is the utility to the other candidate. In this case, it the candidate preferred by the group with the lower learning cost always wins the election.

Note that since choice depends only on $\frac{u(\omega,i)-u(\omega,j)}{\lambda}$ we could replace low learning costs in the example with high intensity of preference. With costly learning those who care most about the result of an election will tend to learn the most, and their choices will correlate more closely with their preferences.

4.4 Regulation

Smallwood and Conlisk (1979) suggest the possibility that product market regulation in the form of minimum standards can reduce welfare. The idea is that in a model in which agents learn from market share improving minimum quality may raise the market share of low quality products thereby reducing average quality in the market place in steady state. Smallwood and Conlisk model minimum standards as an increase in the reliability of products and model learning as a mechanical feedback between market share and product choice. In our model the closest analogy to the Smallwood and Conlisk thought experiment would be an increase in the minimum $u(\omega, i)$ across products and agents. While such an increase will tend to increase the market share of good i for all agents, it can be shown that the increase in any $u(\omega, i)$ always weakly increases the expected utility of all market participants in steady state.

Proposition 1. *For all $i \in A$ and $\omega \in \Omega$, an increase in $u(i, \omega)$ raises $P(i)$ and $V(g_{pop}, A)$.*

If there is a role for regulation, it is in managing the externality that arises from social learning. In our model this externality shows up in the convergence to steady state and in the steady state set of chosen options. The role for policy is most likely in influencing the latter. In steady state, agents act as if they understand the true population and are choosing from $\bar{A}$. There is no guarantee that $\bar{A}$ contains the optimal options. It may therefore make sense for a regulator to encourage experimentation.

5 Conclusion

We have presented a model of social learning with heterogeneous agents who choose among heterogeneous options. We have shown that optimal choice probabilities are skewed in the direction of the most popular options and that relative demands may be more informative about preferences that aggregate demand.

There are several directions one might want to extend the model. One would be to alter the timing. Our agents make a single choice at a single date. Allowing agents to choose the date at which they make the choice would introduce a desire to wait for others to reveal information. Allowing for repeat purchases would allow agents to learn from their own experience and introduce a motive for experimentation. Another extension would be to introduce a taste for conformity as models of beauty contest. In this case the payoff would depend not only on the agent's own type but on the actions of others. Agents would use their knowledge of the distribution of types to predict the actions of others and there would be an additional effect of market share on outcomes. Another direction would be to introduce supply side considerations. How do firms position themselves in such markets? What are their incentives to advertise or to introduce new goods? How does market share influence profitability? All of these directions present interesting possibilities for future work.

References

[1] Althaus, Scott L. (1998),"Information Effects in Collective Preferences," *The American Political Science Review* 92, 545-558.

[2] Balassa, Bela (1965), "Trade Liberalization and Revealed Comparative Advantage," *The Manchester School* 33, 99-123.

[3] Bartles, Larry (1996), "Uninformed Votes: Information Effects in Presidential Elections," *American Journal of Political Science* 40, 194-203.

[4] Becker, Gary (1981), "A Note on Restaurant Pricing and other Examples of Social Influences on Price," *Journal of Political Economy* 99, 1109-1116.

[5] Bernheim, Douglas, and Antonio Rangel (2009), "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics," *Quarterly Journal of Economics* 124, 52-104.

[6] Berry, Steven, James Levinson and Ariel Pakes (1995), "Automobile Prices in Market Equilibrium," *Econometrica* 63, 841-890.

[7] Bikhchandani, Sushil, David Hirshleifer, Ivo Welch (1992), "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy* 100, 992-1026.

[8] Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo (2013), "Understanding Booms and Busts in Housing Markets," Northwestern University Working Paper.

[9] Caminal, Raomon, and Xavier Vives (1996), "Why Market Shares Matter: An Information Based Theory," *Rand Journal of Economics* 27, 221-239.

[10] Caplin, Andrew, Mark Dean (2014), "Rational Inattention and State-Dependent Stochastic Choice," NYU working paper.

[11] Caplin, Andrew, Mark Dean and John Leahy (2014), "Behavioral Implications of Rational Inattention," NYU working paper.

[12] Caplin, Andrew and John Leahy (1994), "Business as Usual, Market Crashes, and Wisdom after the Fact," *American Economic Review* 84, 548-565.

[13] Caplin, Andrew and John Leahy (1998), “Miracle on Sixth Avenue: Information Externalities and Search,” *Economic Journal* 108, 60-74

[14] Chamley, Christophe and Douglas Gale (1994), “Information Revelation and Strategic Delay in a Model of Investment,” *Econometrica* 62, 1065-1085.

[15] Cover, Thomas and Joy Thomas (2006), *Elements of Information Theory, Second Edition*, Hoboken, NJ: John Wiley & Sons.

[16] Delli Carpini, Michael X, and Scott Keeter (1996), *What Americans Know about Politics and Why It Matters*, New Haven, CT: Yale University Press.

[17] Ellison, Glenn, and Drew Fudenberg (1995), “Word-of-Mouth Communication and Social Learning,” *Quarterly Journal of Economics* 110, 93-125.

[18] McFadden, Daniel (1974), “Conditional Logit Analysis of Qualitative Choice Behavior,” in P. Zarembka (ed.), *Frontiers in Econometrics*, 105-142, Academic Press: New York, 1974.

[19] Matějka, Filip, and Alisdair McKay (2014), “Rational Inattention to Discrete Choices: a New Foundation for the Multinomial Logit Model,” *American Economic Review*, forthcoming.

[20] Matějka, Filip, and Christopher A. Sims (2011), “Discrete Actions in Information-Constrained Tracking Problems,” CERGE-EI Working Paper No. 441.

[21] Moretti, Enrico (2010), “Social Learning and Peer Effects in Consumption: Evidence from Movie Sales,” *Review of Economic Studies* 78, 356–393.

[22] Munshi, Kaivan (2003), “Social Learning in a Heterogeneous Population: Technology Diffusion in the Indian Green Revolution,” *Review of Economic Studies* 73, 175-203.

[23] Sims, Christopher A. (1998), “Stickiness,” *Carnegie-Rochester Conference Series on Public Policy* 49, 317–356.

[24] Sims, Christopher A. (2003), “Implications of Rational Inattention,” *Journal of Monetary Economics* 50, 665–690.

[25] Smallwood, Dennis, and John Conlisk (1979), “Product Quality in Markets where Consumers are Imperfectly Informed,” *Quarterly Journal of Economics* 93, 1-23.

[26] Sorenson, Alan (2006), “Social Learning and Health Plan Choice,” RAND Journal of Economics 37, 929-945.

A Proofs.

Proof of Proposition 1:

There are a finite number of types $\omega \in \Omega$. Hence each g may be represented by a point in the $|\Omega| - 1$ dimensional simplex in $R^{|\Omega|}$. Γ^0 is a subset of this simplex. Hence Γ^0 is the subset of an $|\Omega| - 1$ dimensional hyperplane in $R^{|\Omega|}$. This plane is the plane through g_{pop} that is orthogonal to the unit vector

$$(g - g_{pop}) \cdot 1 = 0$$

Define E^0 as the subspace generated by the unit vector.

Consider period t, with Γ^t and E^t. Choice in period t givers rise to a set of type specific choice probabilities $P(\omega, i)$. Let $Z^i \in R^{|\Omega|}$ denote the vector with $Z^i_\omega = P(\omega, i)$. Now the orthogonality conditions can be written as

$$(g - g_{pop}) \cdot Z^i = 0 \qquad \forall i \text{ such that } P(\omega, i) > 0.$$

Each orthogonality condition defines a $|\Omega| - 1$ dimensional hyperplance $R^{|\Omega|}$.

There are two possibilities in period t. First, all the Z^i lie in E^t. In this case there are no new restrictions placed on the set of possible distributions. $\Gamma^{t+1} = \Gamma^t$. Learning stops and the the market has converged. Alternatively, there exists $Z^i \notin E^t$. E^{t+1} is now the space generated by E^t and the $Z^i \notin E^t$. The dimensionality of E^{t+1} is strictly greater than E^t. Γ^{t+1} is the subset of Γ^t that is orthogonal to all vectors in E^{t+1}. The dimension of Γ^{t+1} is therefore strictly less than that of Γ^t. Note, by construction, $g_{pop} \in \Gamma^{t+1}$ if $g_{pop} \in \Gamma^t$.

As $|\Omega|$ is finite Γ^t converges in a finite number of periods.■

Proof of Proposition 2 and Corollary 3:

Let us first define the following mapping $f : \Delta(\Omega) \to \mathbb{R}$:

$$f(g, i, t) = \sum_{\omega \in \Omega} \frac{g(\omega) \exp(u(\omega, i)/\lambda)}{\sum_{j \in A} P^t(j) \exp(u(\omega, j)/\lambda)}. \tag{14}$$

Equations (5) and (1) imply that for all i such that $M^t(i) > 0$, then $P^t(i)$ is also positive and the following holds:

$$f(g_{pop}, i, t) = \frac{M^t(i)}{P^t(i)}. \tag{15}$$

Similarly, (5) together with the fact that the sum of probabilities in a distribution equals

1 implies:

$$f(\mu^t, i, t) = 1. \tag{16}$$

The agent in period $t+1$ knows M^t as well as P^t, and thus the agent does not deem possible those population distributions g that do not satisfy $f(g, i, t) = M^t(i)/P^t(i)$ for some i such that $M^t(i) > 0$,

$$\Gamma^{t+1} = \{g \in \Gamma^t; f(g, i, t) = M^t(i)/P^t(i)\}.$$

Now, if $f(g, i, t) = 1$ for all $g \in \Gamma^t$ and all i such that $M^t(i) > 0$, then $g_{pop} \in \Gamma^t$ implies $f(g_{pop}, i, t) = 1$ and thus also $M^t(i) = P^t(i)$. In this case $f(g, i, t) = M^t(i)/P^t(i)$ for all $g \in \Gamma^t$ so that $\Gamma^{t+1} = \Gamma^t$ and we have converged to a steady state $\bar{\Gamma}$.

If, on the other hand, there exist $g \in \Gamma^t$ and i such that $M^t(i) > 0$ for which $f(g, i, t) \neq 1$, then since $f(\mu^t, i, t) = 1$, $f(g, i, t)$ is linear in g and μ^t is the population distribution conditional on Γ^t, then there must exist $g' \in \Gamma^t$ for which $f(g', i, t) \neq M^t(i)/P^t(i)$ whatever $M^t(i)/P^t(i)$ is. Such g' then does not belong to Γ^{t+1}. Hence $\Gamma^{t+1} \subset \Gamma^t$.

The set of possible population distributions thus shrinks in every period, or reaches a steady state, where $M(i) = P(i)$. Γ therefore converges pointwise in $\Delta(\Omega)$.

Finally, in steady state $f(g, i, t) = 1$ for all $g \in \bar{\Gamma}$ and all i such that $M(i) > 0$. Let $S = \{i \in A | M(i) > 0\}$. Since $g_{pop} \in \bar{\Gamma}$, $f(g_{pop}, i, t) = 1$ for $i \in S$. But then

$$\sum_{\omega \in \Omega} \frac{g_{pop}(\omega) \exp(u(\omega, i)/\lambda)}{\sum_{j \in A} P(j) \exp(u(\omega, j)/\lambda)} = 1, \qquad \forall i \in S,$$

which means that P satisfies the necessary and sufficient conditions for optimality for the prior equal to g_{pop} and an option set S. ■

Proof of Proposition 4

In this case the necessary and sufficient conditions for an optimum are

$$\sum_{\omega \in \Omega} \mu(\omega) \left\{ \frac{\exp(u(\omega, i)/\lambda)}{\sum_{j \in A} P(j) \exp(u(\omega, j)/\lambda)} \right\} \leq 1 \text{ for all } i$$

where A is the set of goods chosen. Substituting the assumed payoffs, if good $i \in A$

$$\frac{x(1+\delta)\mu_i}{P^i x(1+\delta) + x\,(1-P^i)} + \sum_{\omega \in A \setminus i} \frac{x\mu_\omega}{P^\omega x(1+\delta) + x\,(1-P^\omega)} + \sum_{\omega \in M \setminus A} \mu_\omega = 1$$

or

$$\sum_{\omega\in A}\frac{\mu_\omega}{1+\delta P^\omega}+\frac{\mu_i\delta}{1+\delta P^i}=\sum_{\omega\in A}\mu_\omega$$

Define $X^\omega=\frac{\mu_\omega}{1+\delta P^\omega}$, then

$$\sum_{\omega\in A}X^\omega+X^i\delta=\sum_{\omega\in A}\mu_\omega$$

and the solution is symmetric

$$X^\omega=\frac{\sum_{\omega\in A}\mu_\omega}{|A|+\delta}$$

Substituting the definition of X^ω and rearranging

$$\frac{\mu_\omega}{1+\delta P^\omega}=\frac{\sum_{\omega\in A}\mu_\omega}{|A|+\delta}$$

or

$$P^\omega=\frac{\mu_\omega\left(|A|+\delta\right)}{\delta\sum_{\omega\in A}\mu_\omega}-\frac{1}{\delta}$$

Given the monotonicity of μ_ω, $P^\omega>P^{\omega+1}$

$$P^\omega=\frac{\mu_\omega\left(K+\delta\right)}{\delta\sum_{\omega=1}^{K}\mu_\omega}-\frac{1}{\delta}$$

The inequalities follow from the requirement that $P^K>0$ and $P^{K+1}<0$.■

Proof of Proposition 5:

According to (5)

$$P(i,\omega)=\frac{P(i)\exp(u(\omega,i)/\lambda)}{\sum_{j\in A}P(j)\exp(u(\omega,j)/\lambda)}$$

It follows that given, $P(i),P(j)>0$

$$\frac{P(i,\omega)}{P(j,\omega)}=\frac{P(i)\exp(u(\omega,i)/\lambda)}{P(j)\exp(u(\omega,j)/\lambda)}$$

So that

$$\frac{P(i,\omega)/P(i)}{P(j,\omega)/P(j)}=\frac{\exp(u(\omega,i)/\lambda)}{\exp(u(\omega,j)/\lambda)}$$

In steady state $P(i)=M(i)$ and the result follows.■