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

October 2014

Abstract

If private learning was costless, social learning would serve little purpose. Yet, for instance, booksellers lure readers with the label "New York Times Best Seller." We characterize the evolution of market share when agents freely observe past shares and also engage in costly private learning a la rational inattention. The model allows for straight-forward welfare analysis of steady state. The only inefficiency arises form the fact that some good options may have zero market share. Given the chosen options, the market shares are such that they maximize welfare. Even among chosen goods, long run market shares skew toward popular items, and thus increase the market power of the largest players. However, the products with large market shares help economize on the costly private learning by providing a reasonable default options. The model provides a new rationale for anti-trust policies or policies that encourage experimentation. It also has implications for inference from data on market shares as market shares refect a combination of preferences and learning. Finally, we show that an outside observer with rich data may, using a simple test, be better able to understand preferences than are decision makers themselves. Hence preferences and learning costs may be separately identified from suitably rich data.

*We thank Nobuhiro Kiyotaki, Juan Pablo Nicolini, Alessandro Pavan, Chris Tonetti, and Xavier Vives for helpful discussions.

1 Introduction

If private learning were costless, social learning would serve little obvious purpose. Yet social learning is in fact 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 settings with social learning there are many choice options and agents are heterogeneous. In all of the above examples, consumers want to choose goods that match their own preferences. The practical importance of individual differences for social learning is evident in such disparate areas as technology adoption (Munshi (2003)), audience dynamics (Moretti (2010)), and market share dynamics (Sorenson (2006)). Yet despite its practical importance, the interaction with private learning has been down-played in the burgeoning theoretical literature on social learning. Following the early work of Banerjee (1992), Bikhchandani, Hirshleifer, and Welch (1992), Caplin and Leahy (1994, 1998), and Chamley and Gale (1994), this literature remains largely focused on cases with limited consumer heterogeneity and few choice options.

We model the interplay between social and private learning while allowing for unrestricted heterogeneity in tastes and an unrestricted choice set. Heterogeneous agents costlessly observe past market shares, as in Conlisk and Smallwood (1979), Becker (1981), and Caminal and Vives (1996).[1] They can also engage in costly private learning. We provide a general characterization of the evolution of market share and establish convergence. Over time, market shares converge to steady state levels, with the new entrants reproducing the same market shares after undertaking optimal private learning.

When costs are based on Shannon entropy as in the rational inattention model (Sims

[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.

(1998, 2003)), we establish an "as if" result on long run market shares.[2] For goods that are chosen in the long run, market shares are those that would arise if agents knew for sure the true distribution of tastes. This implies that the only failure of optimality arises when prior beliefs result in potentially popular options remaining unchosen.[3]

Rational inattention coupled with social learning from market shares provides a surprisingly tractable model. Typically, choice probabilities in learning models have as many degrees of freedom as prior beliefs. This is implied by the Bayes law and some exogeneity of the form of signals the agent receives. On the other hand, it turns out that in rational inattention the number of degrees of freedom is for a choice from N options $(N-1)$ only. This is because rationally inattentive agents choose to receive signals of a particular form, which is largely independent of the form of prior beliefs. In our model, the $(N-1)$ degrees of freedom that are needed to fully specify the choice behavior coincide with the market shares in steady state.

As in Matějka and Sims (2011), we find that only a small subset of available goods is typically chosen in steady state. This provides a rationale for policies that reduce initial market share disparities. Information revelation and long run market efficiency is improved by a handicapping scheme in which more ex ante popular items are initially taxed and ex ante unpopular items subsidized in a manner that moves prior popularities toward equality. Even among goods that are chosen, we find that limit market shares can greatly exaggerate the market share of popular options as individual choice tends to conflate private and public preferences. This has implications for the distribution of benefits across types: it helps those who prefer popular items over less common types.

The rational inattention model involves a non-standard information asymmetry. An outside observer with access to suitably rich data on market shares may be better able to understand preferences than are decision makers themselves. Availability of such enriched data is the rule rather than the exception in the era of big data. When learning costs are based on the reduction in Shannon entropy, a simple test reveals optimal choices by type. This suggests possible methods of generalizing current methods of recover information on preferences from data on market shares (McFadden (1974), Berry, Levinsohn and Pakes (1995)).

Interestingly, the political science literature has begun to grapple with the problem of

[2] Other papers using this form of cost are Woodford (2009), Mackowiak, Wiederholt (2009), Van Nieuwerburgh and Veldkamp (2010), Matějka (2010), Matějka and McKay (2014), Caplin, Dean and Leahy (2014)

[3] 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.

inferring preferences when there is incomplete information processing. Bartles (1996) and Delli Carpini and Keeter (1996) show that more informed voters vote differently than ill informed voters after controlling for observable characteristics such as age, race, education and party affiliation. They attempt to uncover the "true" distribution of 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.

The model is introduced in section 2 in which we also establish the general convergence result. We specialize in section 3 to the case of Shannon costs and establish our "as if" result. A series of examples in section 4 illustrates both this result and the small number of goods that is typically chosen. We consider issues of welfare in section 5. Issues of inference are addressed in section 6. Section 7 concludes.

2 Model

We consider a dynamic market in which successive generations of agents each make a single choice from a fixed set of available options. Agents differ in their type, and types differ in the payoffs that they receive from the choices that they make. If agents knew their type, they would simply choose the option best suited to their type. The problem is that agents do not know their type.[4] Agents have access to two sources of information about their type. Fist, they observe the distribution of past choices. These choices provide information about the distribution of types in the economy. This forms their prior. Second, agents privately gather further information on their type. Private information acquisition is costly.[5] 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. If the distribution of choices differs by type, then the pattern of choices in the economy will provide information about the distribution of types. This is the reason that market share is informative, which in turn accounts for market dynamics.

[4]While we model a market in which agents are uncertain about their type, it would be equivalent to assume that agents know their type but do not know the match between their type and the available options and do not know how many people are like them.

[5]There needs to be some impediment to private learning for social learning to influence behavior as it appears so often to do in practice. If private learning were costless there would be no need for social learning.

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,

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

which specifies the payoff of each option to each type. 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, so that g_{pop} is a probability density. 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, 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 $\Gamma_0 \equiv \text{supp}(G) \subseteq \Delta(\Omega)$ as the set of possible distributions. We require that $g_{pop} \in \text{int}(\Gamma_0)$.[6]

Given G 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

[6] supp(G) 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}.

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,

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

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, $P(\omega, i)$ will not vary much across types ω and will be larger for choices i that appear desirable ex ante. As learning becomes less expensive, $P(\omega, i)$ 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 choice. In the next section we will place some structure on the $P(\omega, i|\mu)$'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 shares provide 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 observed market shares.

We now describe this process for an arbitrary period t. Let $\Gamma_t \subseteq \Gamma_0$ denote the set densities that are consistent with all market shares observed in periods prior to t. Given Γ_t, we can calculate agents' prior beliefs over preference types as the expected distribution of types conditional Γ_t,

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

The prior μ_{Γ_t} determines the type dependent choice probabilities $P(\omega, i|\mu_{\Gamma_t})$, which along with the true density of types g_{pop}, generates the realized market shares $M(i|\mu_{\Gamma_t}, g_{pop})$. Given that learning is conditionally independent across agents:

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

Agents born in period $t+1$ observe period t aggregate market shares and eliminate distributions that are inconsistent with observed market shares.[7] Γ_{t+1} includes all densities in Γ_t that generate $M(i|\mu_{\Gamma_t}, g_{pop})$,

$$\Gamma_{t+1} = \left\{ g \in \Gamma_t \left| \sum_{\omega \in \Omega} g(\omega) P(\omega, i|\mu_{\Gamma_t}) = M(i|\mu_{\Gamma_t}, g_{pop}) \right. \right\}.$$

Note that if $g_{pop} \in \Gamma_t$ then $g_{pop} \in \Gamma_{t+1}$ as it trivially satisfies this condition. 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 $i \in A$

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

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

A key question concerns whether or not market shares settle 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 $P(\omega, i|\bar{\mu})$ 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 $g \in \Delta(\Omega)$ can be represented 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_t^i in $R^{|\Omega|}$. Given that A is finite there are $|A|$ such hyperplanes. Γ_{t+1} is equal

[7] 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.

to the intersection of Γ_t and these $|A|$ hyperplanes,

$$\Gamma_{t+1} = \Gamma_t \cap (\cap_{i \in A} X_t^i) = \Gamma_0 \cap (\cap_{i \in A} X_0^i) \cap \ldots (\cap_{i \in A} X_t^i).$$

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 (recall that we have normalized the total population to one). In particular,

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

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, $\sum_{\omega \in \Omega} g(\omega) P(\omega, i|\bar{\mu}) = M(i|\bar{\mu})$ for all $g \in \bar{\Gamma}$.

Proposition 2 The steady state market shares $M(i|\bar{\mu})$ are equal to the expected choice probabilities,

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

In order to say more about the behavioral and welfare properties of the model we need to place some structure on the $P(\omega, i)$'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) in deriving the type-dependent stochastic choice map $P(\omega, i|\mu)$. The agent is a Bayesian expected utility maximizer who is rationally inattentive. The agents' information problem is that the agent does not know his or her type. Given the prior $\mu \in \Delta(\Omega)$, we can think of the agent as choosing an information revelation strategy which we model as a random mapping ϕ from possible types into signals or information states. Inverting the signal provides information on the agent's type. Since in the end choices will depend on the agents beliefs, we can without loss of generality associate these signals with the choices $i \in A$ so that $\phi : \Omega \to \Delta(A)$.[8] This set up leads immediately to type-dependent stochastic choice: $P(\omega, i|\mu)$ is the probability of signal i in state ω.

We will assume that more informative information revelation strategies ϕ are more costly and we will measure the information content of any information revelation strategy by the expected entropy of the corresponding posteriors.[9] More precise signals lead to more informative posteriors and are therefore more costly. If $P(\omega, i|\mu)$ is the probability of signal i in state ω, then Bayes rule implies that $\gamma(\omega, i) = P(\omega, i|\mu)\mu(\omega)/P(i|\mu)$ is the posterior probability of state ω conditional on signal i. The cost of an information strategy is proportionate to the reduction in entropy:

$$\lambda \left[\sum_{i \in A} P(i|\mu) \sum_{\omega \in \Omega} \gamma(\omega, i) \ln \gamma(\omega, i) - \sum_{\omega \in \Omega} \mu(\omega) \ln \mu(\omega) \right]$$

The first term in brackets is the expected entropy of the posteriors. The second term is the entropy of the prior. λ is the marginal cost of entropy reduction. Using Bayes rule we can rewrite this cost as:

$$\lambda \left[\sum_{\omega \in \Omega} \mu(\omega) \left(\sum_{i \in A} P(\omega, i|\mu) \ln P(\omega, i|\mu) \right) - \sum_{i \in A} P(i|\mu) \ln P(i|\mu) \right]$$

[8] In the end choice will be a function of the signals. If the desired choice as a funciton of the signals is random, one can redefine the signals as the result of the randomization. It is easy to show that it is not informationally efficient to have two signals associated with the same choice. An act that is unchosen can be associated with a probability zero signal.

[9] Let $P = \{p_1, \ldots, p_n\}$ denote a probabability density. The entropy of P is defined as $-\sum_{i=1}^{n} p_i \ln p_i$ with the convention that $p_i \ln p_i = 0$ when $p_i = 0$. This measure is concave and maximized at $p_i = 1/n$.

Recall $P(i|\mu) \equiv \sum_{\omega\in\Omega} \mu(\omega)P(i|\omega,\mu)$ is the average probability of choosing i. In this interpretation, it costs nothing to choose a strategy $P(\omega, i|\mu)$ that is independent of type ω and it is increasingly costly to make $P(\omega, i|\mu)$ type-contingent.

Given this cost, the agent maximizes:

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

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

Matějka and McKay (2014) show that the resulting pattern of state dependent stochastic choice is of the form,

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

According to (5), the $P(\omega, i|\mu)$ are completely determined by the average choice probabilities $P(i|\mu)$, the payoffs $u(\omega, i)$ and the cost of information λ. The optimal policy "twists" the average choice probabilities in the direction of the choices i that yield higher utility to type ω. The twisting takes a logit form.

Caplin, Dean and Leahy (2014) introduce complementary slackness conditions (6) that characterize the $P(i|\mu)$. For each option $i \in A$, the $\{P(j|\mu)\}_{j\in A}$ must satisfy.

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

with equality if $P(i) > 0$. Note that once one recovers the $\{P(j)\}_{j\in A}$ from (6), the $P(\omega, i|\mu)$ follow directly from (5). Using (5), we can rearrange (4) to express the expected utility in terms of $\{P(i|\mu)\}$ only:

$$V(\mu, A) = \max_{\{P(i|\mu)\}_{i\in A}} \sum_{\omega\in\Omega} \mu(\omega) \log \left(\sum_{i\in A} P(i|\mu) \exp(u(\omega, i)/\lambda \right).$$

Given the concavity of the problem, 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. *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 they might be quite uncertain which distribution 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:

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

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 optimal 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, $P(i|\bar{\mu})$ satisfy the necessary and sufficient conditions for optimal choice if the prior were g_{pop} and the choice set were $\bar{A}$.

The "as if" result helps us out in two ways. 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. We will use this in the next section. Second, the result implies that market inefficiency takes a very limited form. Individual choice is optimal given the observed set of choices $\bar{A}$, but the set of choices $\bar{A}$, however, may not be optimal. We will discuss welfare in Section 5.

4 Solving the Model

4.1 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 underlying forces at work. 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:

$$\tilde{P}(a|\mu) = \frac{\mu(\nu)}{1 - \exp(u(a,\eta)/\lambda)} - \frac{1 - \mu(\nu)}{\exp(u(a,\nu)/\lambda) - 1} \tag{8}$$

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 (8) that an increase in $u(\omega, i)$ 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. 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.

4.2 How Many Goods are Chosen?

Market performance depends on how many options are chosen. A simple example illustrates a phenomenon noted by Matějka and Sims (2011) whereby it is optimal to entirely ignore options that are ex ante unlikely to be best. We remove all heterogeneity beyond the distribution of types and we consider the steady state of a class of symmetric models with $\Omega = A = \{1, .., S\}$. Each agent would like to choose the option matched to their type $i = \omega$. The payoffs are:

$$\exp(u(\omega, i)/\lambda) = \begin{cases} x(1+\delta) & \text{if } i = \omega; \\ x & \text{if } i \neq \omega; \end{cases} \tag{9}$$

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

The next proposition characterizes the solution to this simple model for a given structure of long run population beliefs. We order goods according to perceived likelihood in steady state beliefs $\bar{\mu}$, with lower indexed types perceived as more likely

$$\bar{\mu}(\omega) \geq \bar{\mu}(\omega + 1).$$

Proposition 4 If $\bar{\mu}(S) > \frac{1}{S+\delta}$ define $K = S$. If $\bar{\mu}(S) < \frac{1}{S+\delta}$, then define $K < S$ as the unique integer such that,

$$\bar{\mu}_K > \frac{\sum_{\omega=1}^{K} \bar{\mu}(\omega)}{K + \delta} \geq \bar{\mu}(K + 1) \tag{10}$$

Then the unique solution involves,

$$P^i = \frac{\bar{\mu}(i)(K+\delta) - \sum_{\omega=1}^{K} \bar{\mu}(\omega)}{\delta \sum_{\omega=1}^{K} \bar{\mu}(\omega)} > 0 \tag{11}$$

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

The proposition characterizes choice as a function of the steady state beliefs $\bar{\mu}$ 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 phenomena arise in this setting. The first is that if $\mu(S) < \frac{1}{S+\delta}$, $|A| < S$. 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 $S = 2$ and $\bar{\mu}(2) < \frac{1}{1+\delta}$, 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 (11) 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 $i, j \leq K$,

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

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 $\bar{\mu}(1)$, and the probability of the least popular choice P^S is less than the share of type S agents $\bar{\mu}(S)$. The following numerical example illustrates this skewing of choice probabilities.

Example Suppose that $\delta = 1$ and that in steady state the five most probable states satisfy,

$$(\bar{\mu}(1), \bar{\mu}(2), \bar{\mu}(3), \bar{\mu}(4), \bar{\mu}(5)) = \left(\frac{10}{100}, \frac{9}{100}, \frac{8}{100}, \frac{7}{100}, \frac{6}{100} \right).$$

Proposition 4 implies that $K = 4$,

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

The existence of any additional options beyond these most likely five are therefore irrelevant. The first difference condition implies,

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

Hence,

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

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

The above example leaves unspecified the process of converging to steady state beliefs. In practice this depends intricately on the nature of initial beliefs, which priors ultimately determine the evolution of market shares.

5 Welfare and Policy

5.1 Social 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 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.[10]

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

[10]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

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

Note that here we retain μ in the information cost as we interpret the learning cost as subjective.[11] However, 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 $\bar{A}$. This together with expression above for the objective imply the following proposition.

Proposition 5 In steady state, given a set of chosen options $\bar{A}$, the agent maximizes the following expected utility.

$$V(g_{pop}, \bar{A}) = \max_{\{P(i|\mu)\}_{i\in\bar{A}}} \sum_{\omega\in\Omega} g_{pop}(\omega) \log \left(\sum_{i\in\bar{A}} P(i|g_{pop}) \exp(u(\omega, i)/\lambda\right), \tag{12}$$

which also equals social welfare.

The immediate implication is:

Corollary 3 In steady state, given $\bar{A}$, the market shares $M(i|\bar{\mu})$ are such that they maximize (12).

This statement holds since in steady state $M(i|\bar{\mu}) = P(i|\bar{\mu})$, and $P(i|\bar{\mu})$ maximize the expected utility and thus also welfare. Similarly, since market shares maximize welfare given the set of chosen options, then the following statement is also an immediate implication.

Corollary 4 Let $\bar{A}, \bar{B}$ be two sets of chosen options such that $\bar{A} \subset \bar{B}$, then

$$V(g_{pop}, \bar{A}) \leq V(g_{pop}, \bar{B}),$$

[11] 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 interpretation of the maximization problem (4) is the following. The individual chooses $\{P(\omega, i|\mu)\}_{i\in A,\omega\in\Omega}$ 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 $\gamma^i(\omega) = P(\omega, i|\mu)\mu(\omega)/P(i|\mu)$ where the latter follows directly from Bayes rule.

i.e. an expansion of the set of chosen options weakly increase welfare.

5.2 Handicapping Policies

Given the set of chosen options and the constraints on publicly available information, market shares are efficient. Hence the only inefficiency comes from the possibility that in steady state some good options may not be selected at all. We consider now the potential role for the social planner in improving the choice set.

One application of our model is thus to the antitrust policies. Typically the argument for such policies is based on the level of prices. When the competition is low, then the few firms in the market charge higher prices than if there were many firms. Our model provides a rationale for why such policies can be of interest that is instead based on addressing the quality of chosen products. A policy that limits market shares of the market leaders provides room for new entrants whose quality would be tested by the agents. Selection decisions of agents testing these entrants would then generate positive information externality to future generations of agents. A good antitrust policy would allow for experimentation of products, which could expand the set of the chosen products, and increase welfare in steady state.

If the true population is known, it is generally the case that welfare increases with the set of considered choices. Hence for long run purposes an ideal policy would be to induce full learning in the first period from choice on an unrestricted choice set. In certain cases, just such a policy of improving long run market performance by increasing knowledge and expanding the initial set of chosen options is available. This is so in the context of our simple running example. In that case, the policy maker can make it equally likely in the first period that all options are most preferred by appropriate use of tax and subsidy schemes. This would result in all options that are preferred by a positive mass of agents being chosen with strictly positive probability. While possibly raising attention costs for the first generation and diminishing the quality of their choices, this policy would benefit all future generations since the actual market shares would then reveal $g_{pop}(\omega)$. At this point all taxes and subsidies could be removed and the market would settle to the full information optimum. Note that this would involve choice only of the most popular goods according to $\bar{\mu}(\omega) = g_{pop}(\omega)$, precisely as in our general solution.

The example above is not general. If there are fewer goods than types, then there will not be full revelation of the population distribution based on a single prior. How best to induce experimentation in this general case is an open question. It may for example involve inducing a dynamic and state dependent system handicapping policy that aims sequentially to uncover remaining aspects of uncertainty. Another issue to be borne in mind in the full solution is the appropriate rate of discount as between the early group who are induced to experiment and the later groups who benefit from their incremental policy-induced experimentation. To fully address optimal handicapping policies is beyond the scope of this paper.

5.3 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, the increase in any $u(\omega, i)$ will typically increase the expected utility of all market participants in steady state. This is certainly the case provided all previously chosen goods are still chosen, since welfare in (12) can obviously achieve a higher value. Again, full consideration of the role of regulation is beyond the scope of the paper.

5.4 Heterogeneity and Welfare

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.

6 Enriched Choice Data and Inference

The rational inattention model introduces a non-standard information asymmetry. An outside observer with access to suitably rich data on market shares may be better able to understand preferences than are decision makers themselves. Agents in the model are learning optimally given their limited resources. They focus their attention on matters that concern them directly, but their powers are limited. One could imagine that a large agent, such as the government, Google, Amazon, or a market research firm such as J. D. Power and Associates or Consumer Reports, might have greater access to large amounts of detailed choice data as well as greater incentives process this information. Governments and research firms collect a broad range of statistics. Google sees the search behavior of a large fraction of agents. Amazon directly observes consumer choice. In this era of big data such an agent might be able to put together detailed market data and might be able to learn type-specific market shares.

The following proposition captures the idea that because agents tend to choose options that they prefer, market shares will be very informative about preferences. A simple ratio test allows one to infer agents' preferences and reveals optimal choices by type.

Proposition 6 In the steady state,

$$\frac{u(\omega,i)-u(\omega,j)}{\lambda} = \log\left(\frac{M(\omega,i)}{M(\omega,j)} \Big/ \frac{M(i)}{M(j)}\right). \tag{13}$$

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 $M(i)$. 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.

Proposition 6 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:

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

which identifies utility up to the learning cost (and the utility of option j). 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}.

In most other models, e.g. with optimal deterministic choice, type-specific choice would not be very informative as it would reveal the most preferred option only. Here, on the other hand, the choice is probabilistic with probabilities reflecting the preferences. Type-dependent choices reveal not only the most preferred option, but the entire ranking of observed choices.

It is interesting to note that the term in brackets on the right-hand side of (13), which emerges from the model, is 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.

The above characterization relates to a long tradition in industrial organization of using market shares to infer utility parameters. Prominent examples include McFadden (1974) and Berry, Levinsohn and Pakes (1995). This literature normally takes as its starting point observation of the aggregate market shares $M(i)$. Inference from market share is not straight forward in our setting, since the market shares conflate the demands of many types of consumer and exaggerates the influence of popular types, thereby biasing inference. Inference is more straightforward from the type-dependent demands $M(\omega, i)$. To capture this model in our setting, we associate the choices A with differentiated products and the types Ω with groups of heterogenous consumers. We then let the utility of each choice depend on the value of the good and its market price,

$$u(\omega, i) = \delta^i_\omega - \alpha p_i.$$

Note that we assume that agents see and understand prices. Their inference problem is one of figuring out which good to purchase. To match prominent specifications in the differentiated product setting, we suppose that the value of good i to an agent of type ω depends on a vector of product characteristics X^i, so that $\delta^i_\omega = X^i\beta_\omega$. Consider now a regression that projects type-specific market shares onto product characteristics:

$$\log\left(\frac{M(\omega,i)}{M(\omega,j)}\bigg/\frac{M(i)}{M(j)}\right) = (X^i - X^j)\frac{\beta_\omega}{\lambda} - \frac{\alpha}{\lambda}(p_i - p_j) + \xi. \tag{14}$$

where ξ is a regression error reflecting measurement error or omitted factors.[12] The regression is similar to the standard logit model of McFadden (1974).[13] If one had data on type-specific market participation, one could identify one of the choices with the outside option. Alternatively, one may therefore normalize the utility of one option for each type of agent to zero without affecting behavior.

The main difference between (14) 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 the cost of learning, λ. This will not matter much in situations in which learning costs are stable. In other cases, however, 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:

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

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

[13]Note the model provides an alternative interpretation of data. Dispersed choice of one specific type in data is typically interpreted as a proof of unobserved heterogeneity.In our model such choice can in fact arise for one particular type of given preferences, since the dispersion in choice is driven by noise in attention, i.e. by mistakes in the choice.

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

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

The first term reflects the direct effect of prices on market share. Higher learning costs tend to dampen this effect. The remaining terms reflect the effect of prices on the unconditional choice probabilities. 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.

As in industrial organization, it is common practice to infer political preferences from opinion polls and vote share. In this setting the influence of differences in information are increasingly under investigation. 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 $\lambda > 0$ that above suggests that richer methods of inference are required.

7 Conclusion

We characterize the evolution of market share when agents freely observe past shares and also engage in costly private learning. Our characterization of steady state behavior in particular opens the doors to analysis of market behavior, policy, and to issues of inference from suitably rich data. Generalizations of the model may also be of interest.

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] Banerjee, Abhijit (1992), "A Simple Model of Herd Behavior," *Quarterly Journal of Economics* 107, 797-817.

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

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

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

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

[8] 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.

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

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

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

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

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

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

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

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

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

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

[19] Maćkowiak, Bartosz and Mirko Wiederholt (2009), "Optimal sticky prices under rational inattention," *The American Economic Review* 99, 769-803.

[20] Matějka, Filip (2010), "Rationally inattentive seller: Sales and discrete pricing, " CERGE-EI Working Paper No. 408.

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

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

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

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

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

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

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

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

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

[30] Van Nieuwerburgh, Stijn, and Laura Veldkamp (2010), “Information Acquisition and Under-Eiversification,” *The Review of Economic Studies* 77, 779–805.

[31] Woodford, Michael (2009), “Information-Constrained State-Dependent Pricing,” *Journal of Monetary Economics* 56, 100-124.

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 Corollary 3:

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

$$f(g, i, t) = \sum_{\omega \in \Omega} \frac{g(\omega) \exp(u(\omega, i)/\lambda)}{\sum_{j \in A} P^t(j) \exp(u(\omega, j)/\lambda)}. \tag{15}$$

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:

$$f(g_{pop}, i, t) = \frac{M^t(i)}{P^t(i)}. \tag{16}$$

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

1 implies:

$$f(\mu^t, i, t) = 1. \tag{17}$$

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 6:

According to (5)

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

where we have suppressed the prior μ to simplify notation. It follows that given, $P(i),P(j)>0$

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

So that

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

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