NBER WORKING PAPER SERIES

SOCIAL LEARNING AND SELECTIVE ATTENTION

Andrew Caplin
John Leahy
Filip Matějka

Working Paper 21001
http://www.nber.org/papers/w21001

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2015

We thank Sourav Bhattacharya, Nobuhiro Kiyotaki, Juan Pablo Nicolini, Alessandro Pavan, Avner Shaked, Chris Tonetti, Xavier Vives, and Krešo Žigič for helpful discussions. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Social Learning and Selective Attention
Andrew Caplin, John Leahy, and Filip Matějka
NBER Working Paper No. 21001
March 2015
JEL No. D83

ABSTRACT

Popularity is self reinforcing. The attention garnered by popular options propels further interest in them. Yet rather than blindly follow the crowd, most pay attention to how well these items match their tastes. We model this role of social learning in guiding selective attention and market dynamics. We confirm that attention focuses on options that quickly achieve popularity. Information externalities render the chosen set smaller than socially optimal. This rationalizes antitrust policies that encourage early experimentation. When attention costs are based on Shannon entropy, optimal policies are computable. With rich data, optimal choices can be identified for all consumer types.

Andrew Caplin
Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10012
and NBER
andrew.caplin@nyu.edu

John Leahy
Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10012
and NBER
john.leahy@nyu.edu

Filip Matějka
CERGE-EI
Filip.Matejka@cerge-ei.cz

1 Introduction

Popularity is to some extent self-reinforcing. The attention garnered by books, films, and restaurants that quickly achieve popularity raises further interest in them. Many readers pay particular attention to books that are New York Times Best Sellers, high grossing movies attract additional viewers (as in Moretti (2010)), and musicians regard rising to the top in the music charts as important in terms of attracting further fan interest. Indeed the value of such social signals may be higher than ever in the era of "choice overload" that the Internet heralds (Iyengar and Lepper (2000)).

Social learning theory considers the effect that early adaptors have on subsequent market participants (Banerjee (1992), Bikhchandani, Hirshleifer, and Welch (1992), Caplin and Leahy (1994, 1998), and Chamley and Gale (1994)). In these models popularity typically acts as one of several signals available to agents. The examples above, however, suggest that social learning not only provides direct information, but also determines the form and selectivity of subsequent private learning. Rather than blindly following the crowd, many agents work to sort out how well popular options match their tastes. Market shares guide their private choice of where to direct their attention, as these agents rationally choose to learn more about popular options than unpopular ones. This attentional channel is absent in current models of social learning, which either rule out additional private learning following observation of market share, or restrict such learning to be independent of market shares. These models are therefore silent on how social learning shapes private learning.

We develop a simple, dynamic model of market share that highlights interactions between social learning and private attentional choice. The key role of social learning in our model is to guide buyers not only on how much total attention to allocate to each decision, but also on precisely how to divide this attention up in an optimal manner. We show that allowing the agents the opportunity to gather information selectively after observing market behavior has a fundamental effect on individual learning and on the resulting market dynamics.

In technical terms, the key innovation in our model is introduction of a flexible learning technology. Given that our focus is on how social learning shapes subsequent private learning, we do not restrict the structure of the chosen private signals. Instead we specify a cost to information acquisition in which more attentive strategies are more costly as in rational inattention theory (Sims (2003)). Note that private costs of learning are essential to the social learning literature. If private learning was in fact costless, there would be no need for social learning.

As in Conlisk and Smallwood (1979), Becker (1981), and Caminal and Vives (1996), we consider a market in which new buyers observe all past market shares. However, unlike these earlier models, agents in our model have an unrestricted choice of additional information. Since such private learning is costly, its intensity and form depends on past market shares. Informational effort is allocated across alternatives optimally, and agents optimally balance the costs and benefits of additional private learning. While past popularity shapes private learning, individual differences are also evident in market behavior.

Our first result establishes convergence of market shares to a stable steady-state distribution. In this steady state, we show that attention is commonly restricted to only a small subset of available options, as in Matějka and Sims (2011). Products that do not receive significant market share in early periods may be ignored by future buyers, as their low market share encourages agents to direct their attention elsewhere. Such products might then easily be missed even by consumers who would very much prefer them to the best sellers. Costs of private learning also increase market concentration even among goods that are chosen. Long run market shares may greatly exaggerate the demand for the most popular options as individual choice conflates private and social preferences. This has implications for the distribution of benefits across types: learning from market share helps those who prefer popular items over less common types.

When costs of private information are based on Shannon entropy, we establish an "as if" result on long-run market shares. In the long run, market shares are those that would arise if agents were perfectly informed about the distribution of tastes, but were restricted to choose from the set of options with positive long-run market share. In fact, conditional on options with positive market share, social learning is optimal in the long run. This implies that the only long run failure of social optimality arises when potentially popular options remain unchosen. Note that such failures are hard to spot in practice, since unchosen options are absent from the market place.

The possibility that potentially popular items may never be chosen provides a new rationale for antitrust policies that reduce initial market share disparities. Unlike many existing policies, which are typically based on considerations of lower prices, these policies aim at limiting initial market share of popular options to increase the quality of selected products. We show that information revelation and long run market efficiency may be improved by an appropriate handicapping scheme. This scheme involves more ex ante popular items being initially taxed and ex ante unpopular items being subsidized in a manner that moves prior popularities toward equality.

On the empirical front, our 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. 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 ways to generalize current methods of recovering information on preferences from data on market shares (McFadden (1974), Berry, Levinsohn and Pakes (1995)).

We discuss related literature in the next section. The model is introduced and the general convergence result established in Section 3. In Section 4 we consider the case in which costs of information acquisition depend on Shannon's mutual information and establish our "as if" result. A series of examples in Section 5 illustrate both the basic workings of the model, including the property that only a small number of options are typically chosen. We consider issues of welfare in section 6. Issues of inference are addressed in section 7. Section 8 concludes.

2 Related Literature

This paper contributes to the branch of the social learning literature that studies how agents learn from the actions of others. The first generation of such models (Banerjee (1992), Bikhchandani, Hirshleifer, and Welch (1992), Caplin and Leahy (1994, 1998), Chamley and Gale (1994)) focused on the externalities associated with learning from others: the incentive to wait for others to act and the possibility that the market fails to converge to the first best due to the suppression of private information.

While social learning theory has since progressed in many important directions, it remains largely focused on cases with few choice options, limited private information acquisition strategies, and limited consumer heterogeneity. The vast majority of papers fall into one of two modelling traditions. One such tradition involves binary choice, two payoff relevant states, and agents who receive idiosyncratic Bernoulli signals of the true state prior to observing choices of others. The other involves a continuum of actions choices, a continuous state space, a quadratic loss function, and normal signals from a pre-specified distribution. These cases are studied more for their analytic tractability than for their economic realism. These are the two classic settings in which agents' beliefs maintain the same functional form from period to period, a property that greatly simplifies the recursive solution of these models (Chamley, 2004).

Our model departs in two key respects from the prior social learning literature. First, we stress individual differences: consumers want to choose goods that match their own preferences about which the choices of others are only imperfectly informative. In the applied literature on social learning, the practical importance of individual differences is evident in such disparate areas as technology adoption (Munshi (2003)), audience dynamics (Moretti (2010)), and market share dynamics (Sorenson (2006)). Second, we apply rational inattention theory to stress the rich learning opportunities that may be open after receiving social information.

Our use of the rational inattention framework for private gathering of information connects us with the pioneering work of Sims (1998), and subsequent work by Sims (2003), Luo (2008), Woodford (2009), Mackowiak, Wiederholt (2009), Van Nieuwerburgh and Veldkamp (2010), Mondria (2010), Matějka (2010), Matějka and McKay (2015), Caplin, Dean and Leahy (2015). The departure from this literature is introduction of social learning. In our model successive generations of agents make choices from a fixed set of available options. Each generation costlessly observes the choices of past agents and these observations shape their prior. Given this prior, agents then gather additional information that tailors their choice to their own tastes.

We follow rational inattention theory and model the cost of private information acquisition as proportionate to Shannon's mutual information. This results in a surprisingly tractable model. This simplicity allows us to consider cases beyond the two-act-two-state-two-signal case and the continuous-act-continuous-state-normal-signal case, and to consider general action sets and general forms of heterogeneity in the individual payoffs to actions. That our model is tractable despite its richness derives from the unlimited flexibility of information acquisition that rational inattention theory embodies. We show that rationally inattentive agents choose to receive signals of a particular form. The desired form of signals "twists" the unconditional probability of choosing an option in the direction of options with higher payoffs to the agent. With the Shannon cost function, this twist has a logit form that connects with the empirical discrete choice literature. Since the twist depends only on beliefs about chosen alternatives and not on the realized states, one only needs to solve for the unconditional probability of choosing an option to solve for the state dependent probabilities. In a market with N options chosen, only $(N-1)$ degrees of freedom are needed to fully specify the choice behavior regardless of the number of states of the world.

Moreover, at the steady state the $(N-1)$ degrees of freedom coincide with the market share, and thus provide useful conditions for the choice behavior at the steady state.

Our approach draws from a small yet important prior literature on the feedback effects of market share on choice. 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 branch of the social learning 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.

As indicated in the introduction, our model has implications for inference. Interestingly, the political science literature has begun to grapple with the problem of 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 over candidates 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.

3 Model

We consider a dynamic market in which successive generations of agents make choices from a fixed set of available options. Agents differ in their type, and types differ in the payoffs that they receive from available options. 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.[1]

In addition to common prior beliefs, agents who enter the market in any period have access to two sources of information. First, 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.[2] The limits on private learning imply that mistakes are made, so that an individual sometimes ends up with a choice that they would like less than available alternatives.

3.1 Model structure

Time is discrete and indexed by $t \in \{0, 1, \dots\}$. There is a fixed finite set $A = \{1, \dots, N\}$ of options in a particular market. Each period a continuum of 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 makes a once-off choice from the set 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,

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

which specifies the payoff of each option to each type. Let $\Delta(\Omega)$ denote the set of densities over types and let $g^* \in \Delta(\Omega)$ denote the true density of preference types in the population. $g^*(\omega)$ is then the share of type ω agents. g^* is fixed and does not change over time. We normalize the total population of agents to 1, so that g^* is a probability density. We place no restriction on the form of the utility function and how it varies across goods and types.

Buyers new to the market do not know their type 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

[1] 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.

[2] 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.

distributions in $\Delta(\Omega)$. 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 $g^* \in \text{int}(\Gamma_0)$.[3] Given G and Γ_0, we can calculate time zero agents' prior beliefs over preference types as the expected distribution of types,

$$\mu_0(\omega) = \int g(\omega) dG.$$

In addition to relying on the prior, each agent can obtain additional information about ω, and thus about the utilities of available options, by exploring the offered options in more detail. This 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 in yelp.com; or it might involve discussions with friends, colleagues, or other people that the agent regards as similar to him or herself.

We do not place structure on the process of private learning. Rather we allow agents to choose an arbitrary information structure subject to a cost of choosing more informative structures. This means that agents can process information in any way they like. They may choose to learn more about the utility they derive from several of the available options, or work to compare particular pairs of products, etc. The only structural assumption we make is that more Blackwell informative information structures are more costly.

This model can be formulated in a compact way, since the distribution of option choices conditional on type,

$$P(i, \omega|\mu) = \Pr\{i \in A|\omega \in \Omega, \mu\}, \tag{1}$$

describes both the information acquisition and the selection of options. To understand this formulation, note that an information strategy is described by a joint distribution of types and signals, which defines how likely are agents of each type to receive each signal (see Laffont (1989)). In terms of its impact on action choice, one can characterize each signal by the corresponding posterior belief that it induces. Upon acquiring each particular signal and identifying the corresponding posterior belief, the agent chooses among available options to maximize expected utility. Therefore, signals can be associated with corresponding optimal choices. Since more Blackwell informative information structures are more costly, no option will be chosen from two distinct signals. Hence a strategy of information acquisition is

[3] supp(G) is the set of $g \in \Delta(\Omega)$ such that every open neighborhood of g has positive measure. With $g^* \in \text{int}(\Gamma_0)$, the continuity of G ensures that the density of G is strictly positive at g^*.

equivalent to a type-dependent stochastic choice function $P(i,\omega|\mu)$. Each such function specifies the probability of observing information that prompts the choice of action i when the agent's true type is ω and the prior is μ. We assume that this private learning is idiosyncratic so that the realized choices i are independent across agents of the same type ω.

Note that the choice probabilities in (1) depend on the agents' true types. This does not imply that agents know their types. In fact, conditional on choosing option i, Bayes' rule implies that the agents' posterior beliefs regarding their own types are,

$$\gamma(\omega|i,\mu) = \frac{P(i,\omega|\mu)\mu(\omega)}{\sum_j P(j,\omega|\mu)\mu(\omega)},$$

where $\gamma(\omega|i,\mu)$ is the probability that the agent is type ω conditional on choice of option i given the prior μ. What information acquisition does is to improve agents' understanding of their true types and thereby shift behavior in the direction of choices that on average are better for them.

3.2 Recursive Learning From Market Share

Agents who enter the market in periods $t > 0$ can learn about their type in part by observing past market shares. Since choice probabilities depend on agents' types, realized market shares may provide information about the distribution of types in the economy. As we shall see, 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 all densities in the support of Γ_0 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 $P(i,\omega|\mu_{\Gamma_t})$. These choice probabilities, along with the true density of types g^*, generate the realized market shares

$M(i|\mu_{\Gamma_t}, g^*)$. We assume that learning is conditionally independent across agents, so that:

$$M(i|\mu_{\Gamma_t}, g^*) = \sum_{\omega \in \Omega} g^*(\omega) P(i, \omega|\mu_{\Gamma_t}). \tag{2}$$

Agents born in period $t+1$ observe period t aggregate market shares and eliminate distributions that are inconsistent with observed market shares.[4] Γ_{t+1} includes all densities in Γ_t that generate $M(i|\mu_{\Gamma_t}, g^*)$,

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

Note that if $g^* \in \Gamma_t$ then $g^* \in \Gamma_{t+1}$ as it trivially satisfies this condition. Given Γ_{t+1} period $t+1$ proceeds in a manner similar to period t completing the recursion.

3.3 Orthogonality and Convergence

Given Γ_t, Γ_{t+1} has a simple structure. According to (2), all $g \in \Gamma_{t+1}$ were elements of Γ_t and generated the same market shares in period t as g^*. It follows that $g \in \Gamma_{t+1}$ if $g \in \Gamma_t$ and for all $i \in A$

$$\sum_{\omega \in \Omega} \left[g(\omega) - g^*(\omega) \right] P(i, \omega|\mu_{\Gamma_t}) = 0. \tag{3}$$

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

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 $P(i, \omega|\mu)$ and market shares $M(i|\bar{\mu}, g^*)$. 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 a set $\bar{\Gamma} \subseteq \Delta(\Omega)$ such that $\Gamma_t \to \bar{\Gamma}$. Moreover $\Gamma_{|\Omega|} = \bar{\Gamma}$.

[4] One feature 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.

The idea behind the proof in the appendix 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 $i \in A$ that is chosen in period t, there is an orthogonality condition (3) with non-zero $P(i, \omega|\mu_{\Gamma_t})$. Each such orthogonality condition defines a $|\Omega - 1|$ dimensional hyperplane X_t^i in $R^{|\Omega|}$. Given that A is finite there are at most $|A|$ such hyperplanes. Γ_{t+1} is equal to the intersection of Γ_t and these hyperplanes,

$$\Gamma_{t+1} = \Gamma_t \cap (\cap_{i \in A_t} X_t^i),$$

where $A_t \subseteq A$ is the set of options that are chosen with positive probability by some type. 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 chosen 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 $\Gamma_t = g^*$. Otherwise learning is incomplete. In general complete learning cannot be guaranteed as we show in Section 5.

What drives convergence is a disconnect between the expected probability of choosing an 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,

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

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

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

$$M(i|\bar{\mu}) = P(i|\bar{\mu}). \tag{4}$$

In order to say more about the behavioral and welfare properties of the model we need to place some structure on type-dependent stochastic choice. In the next section we relate the cost of information acquisition to the mutual information between act choice and the agent's type as in Sims (1998, 2003).

4 Rational Inattention and State Dependent Stochastic Demand

We follow Matějka and McKay (2014) and Caplin, Dean and Leahy (2015) in deriving the optimal type-dependent stochastic choice map $P(i, \omega|\mu)$ when the cost of private learning is proportionate to the mutual information between the choice i and the type ω, as in Sims (2003). Mutual information is a measure of the amount of information flow; it equals the reduction of agent's uncertainty about the type ω:[5]

$$I(i; \omega) = \sum_{\omega \in \Omega} \mu(\omega) \left(\sum_{i \in A} P(i, \omega|\mu) \ln P(i, \omega|\mu) \right) - \sum_{i \in A} P(i|\mu) \ln P(i|\mu) \tag{5}$$

where $P(i|\mu) \equiv \sum_{\omega \in \Omega} \mu(\omega) P(i|\omega, \mu)$ is the average probability of choosing i conditional on the prior μ. If the agent does not learn about type ω, then the cost is zero; according to (5), it costs nothing to choose an information structure that yields a strategy $P(i, \omega|\mu)$ that is independent of the type ω. Learning more about the type is more costly. Given the concavity of $\ln P(i, \omega|\mu)$, it is increasingly costly to make $P(i, \omega|\mu)$ type-contingent. This specific entropy-based cost of information allows for tractability. Moreover, it can be derived from not unreasonable axioms (Csiszár (2008)) and can be microfounded using fairly general assumptions on the technology of information acquisition (Cover and Thomas, 2006).

The agent maximizes:

[5]Entropy, $E\{\ln X\}$, is a measure of the information content of a single random variable, Mutual information $I(Y; X)$ measures the reduction in uncertainty about X from the observation of Y. To see this note 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. Substituting, (5) becomes:

$$I(P; \mu) = \sum_{i \in A} P(i|\mu) \sum_{\omega \in \Omega} \gamma(\omega, i) \ln \gamma(\omega, i) - \sum_{\omega \in \Omega} \mu(\omega) \ln \mu(\omega)$$

The first term is the expected entropy of the posteriors, whereas the second term is the entropy of the prior. Hence I measures the expected reduction in entropy of beliefs over ω.

$$
\begin{aligned}
V(A,\mu) &= \max_{\{P(\omega,i)\}_{i\in A,\omega\in\Omega}} \sum_{\omega\in\Omega}\mu(\omega)\left(\sum_{i\in A} P(i,\omega|\mu)u(i,\omega)\right) \\
&\quad -\lambda\left[\sum_{\omega\in\Omega}\mu(\omega)\left(\sum_{i\in A} P(i,\omega|\mu)\ln P(i,\omega|\mu)\right) - \sum_{i\in A} P(i|\mu)\ln P(i|\mu)\right].
\end{aligned} \tag{6}
$$

The first term on the right-hand side is the expected utility of the strategy $\{P(i,\omega|\mu)\}_{i\in A,\omega\in\Omega}$. The second term is the cost of information acquisition which we take as proportionate to the mutual information $I(i;\omega)$. $\lambda > 0$ is the marginal cost of information.

Matějka and McKay (2014) and Caplin, Dean and Leahy (2015) show that the resulting pattern of state dependent stochastic choice is of the form,

$$
P(i,\omega|\mu) = \frac{P(i|\mu)\exp(u(i,\omega)/\lambda)}{\sum_{j\in A} P(j|\mu)\exp(u(\omega,j)/\lambda)}, \tag{7}
$$

where,

$$
\sum_{\omega\in\Omega}\mu(\omega)\left\{\frac{\exp(u(i,\omega)/\lambda)}{\sum_{j\in A} P(j|\mu)\exp(u(\omega,j)/\lambda)}\right\} \le 1 \qquad \forall i, \tag{8}
$$

with equality if $P(i|\mu) > 0$. Note that once one recovers the $\{P(j|\mu)\}_{j\in A}$ from (8), the $P(i,\omega|\mu)$ follow directly from (7). All type dependent choices $P(i,\omega|\mu)$ are determined by the average choice probabilities $P(i|\mu)$, the payoffs $u(i,\omega)$ and the cost of information λ.

The above solution indicates precisely how optimal policy "twists" the average choice probabilities $P(i|\mu)$ in the direction of those choices that yield higher utility to type ω agents. The twisting takes a logit form which depends only on parameters of the utility function. The higher is $u(i,\omega)$ relative to $u(j,\omega)$ the more likely is option i chosen in state ω conditional on $P(i|\mu)$. This effect is stronger when λ is lower and learning is easier.

Using (7), we can rearrange (6) to express the expected utility in terms of $\{P(i|\mu)\}$ only:

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

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

Axiom 1.

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

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

The solution to (6) can be calcuated by applying a simple Blahut-Arimoto algorithm which involves guessing a function $P^0(i|\mu)$, using (7) to obtain $P^0(i,\omega|\mu)$, deriving $P^1(i|\mu) = \sum_{\omega \in \Omega} \mu(\omega) P^0(i,\omega|\mu)$, and iterating (Cover and Thomas (2006)).

4.1 An "As If" Result

Given the form of type-dependent stochastic choice in (7) we can show that in steady state agents behave as if they know the true distribution of types g^* and are choosing from the steady state set of options. This is the case 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 (2) and the optimal policies (7), we have:

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

Dividing both sides by $P(i|\bar{\mu})$ yields,

$$\sum_{\omega \in \Omega} g^*(\omega) \left\{ \frac{\exp(u(i,\omega)/\lambda)}{\sum_{j \in A} P(j|\bar{\mu}) \exp(u(\omega,j)/\lambda)} \right\} = 1. \tag{9}$$

Equation (9), however, is simply a restatement of the necessary and sufficient conditions (8) for an optimal policy over the choice set $\bar{A}$ given the prior g^*. It follows that an agent with a prior equal g^* and an option set equal to $\bar{A}$, would also choose $P(i|\bar{\mu})$. Agents therefore act as if they know the true distribution of types.

Corollary 1 In the steady state, $P(i|\bar{\mu})$ satisfy the necessary and sufficient conditions for optimal choice if the prior were g^* 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 $\bar{A}$, while the set of choices itself may not be optimal. We will discuss welfare in Section 6.

5 Solving the Model

5.1 The Two-by-Two Case

While much of our interest is in selective attention in large choice sets, we first solve fully the simplest case with two types, ν and η, and two choices, a and b. As in the theory of international trade, the two-by-two case reduces the substitution possibilities among options, allowing for a clear illustration of the underlying forces at work. Suppose that type ν prefers option a while type η prefers option b,

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

Since according to (7) choices depend only on $u(i,\omega) - u(\omega,j)$, normalizing the value of option b to zero for both types is without loss of generality.

Solving (8) assuming that both options are chosen yields the following probability of choosing option a given the prior μ:

$$\tilde{P}(a|\mu) = \frac{\mu(\nu)}{1-\exp(u(a,\eta)/\lambda)} - \frac{1-\mu(\nu)}{\exp(u(a,\nu)/\lambda)-1} \tag{10}$$

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 then follow directly form (7). 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, the market converges to steady state in either period zero or period 1. 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 in subsequent periods. If both options are chosen in period 0, (7) implies that type ν are more likely to choose option a. Since these probabilities are known, period zero market share perfectly reveals the share of type ν agents. In this case, the steady state is reached in period 1.

The two-by-two case has well-behaved comparative statics. It is immediate from (10) and (7) that an increase in either $u(a,\nu)$ or $u(a,\eta)$ will increase the probability that each type will choose option a. In each case, the increase in the payoff to act a to type ω increases the probability that type ω chooses a, which then increases the average probability that a is chosen and thereby increases the probability that the other type chooses a as well. In a similar fashion an increase in $\mu(\nu)$ will cause the probability of choice a to rise. Finally, as learning costs rise, the influence of the ex ante optimal choice grows. When learning is costless, each agent chooses the option that is best for them and the market share of each choice is then equal to the proportion of agents who prefer that choice. As learning costs rise, agents tend to rely more and more on their prior. Eventually, λ rises so high that $P(a|\mu)$ hits either zero or one and only the ex ante optimal option is taken. Since $M(a|\bar{\mu}, g^*) = P(a|\bar{\mu})$ in steady state, these comparative statics apply to the steady state market share as well.

5.2 How Many Goods are Chosen?

In the general case, we have seen that 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 $\Omega = A = \{1, .., N\}$. Each agent would like to choose the option matched to their type $i = \omega$. The payoffs are:

$$\exp(u(i,\omega)/\lambda) = \begin{cases} x(1+\delta) & \text{if } i = \omega; \\ x & \text{if } i \neq \omega; \end{cases} \tag{11}$$

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.

We consider a market that has settled to steady state and order goods according to their perceived likelihood in steady state beliefs $\bar{\mu}$, with lower indexed types perceived as more likely

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

We leave 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.

The necessary and sufficient conditions for the rational inattention model enable us to fully characterize demand for any prior (see Caplin, Dean, and Leahy (2015) and hence is steady state. The key is to identify K, the highest index action that is chosen in steady state. If $\bar{\mu}(N) > \frac{1}{N+\delta}$. then $K = N$. If $\bar{\mu}(N) < \frac{1}{N+\delta}$, then $K < N$ is the unique integer such that,

$$\bar{\mu}_K > \frac{\sum_{\omega=1}^{K} \bar{\mu}(\omega)}{K+\delta} \geq \bar{\mu}(K+1). \tag{12}$$

Having identified chosen actions, the associated choice probabilities satisfy,

$$P(i|\bar{\mu}) = \frac{\bar{\mu}(i)(K+\delta) - \sum_{\omega=1}^{K} \bar{\mu}(\omega)}{\delta \sum_{\omega=1}^{K} \bar{\mu}(\omega)} > 0 \tag{13}$$

for $i \leq K$, with $P(i|\bar{\mu}) = 0$ for $i > K$. Note that 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.

One phenomenon that arises in this setting is that conditions for all actions to be chosen are highly restrictive. If δ is small or there are many choices, it takes only a small deviation from uniformity for $\mu(N) < \frac{1}{N+\delta}$, resulting in some actions being unchosen. This opens the door for inefficient learning. For example if $N = 2$ and $\bar{\mu}(2) < \frac{1}{1+\delta}$, then $P(1|\bar{\mu}) = 1$ and there is no way for the market to learn how many $\omega = 2$ types there are in the population.

A second important phenomenon is that information on market share tends to exaggerate demand for "popular" choices. Consider (13) 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 $i, j \leq K$,

$$P(i|\bar{\mu}) - P(j|\bar{\mu}) = (\bar{\mu}(i) - \bar{\mu}(j)) \left[\frac{(K+\delta)}{\delta \sum_{\omega=1}^{K} \mu(\omega)} \right]. \tag{14}$$

Note that the term in square brackets is strictly greater than one (the denominator is no greater than δ while the numerator is strictly larger than delta). Choice is skewed towards the options with higher prior probability of success. In fact, the unconditional probability of the most likely popular choice $P(1|\bar{\mu})$ is easily seen to be greater than the prior probability $\bar{\mu}(1)$, and the probability of the least popular choice $P(N|\bar{\mu})$ is less than the share of type N agents $\bar{\mu}(N)$. The following numerical example illustrates this skewing of choice probabilities.

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

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

In this case $K = 4$,

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

The existence of any additional options beyond these most likely five are therefore irrelevant. Equation (13) implies

$$P(1|\bar{\mu}) - P(2|\bar{\mu}) = P(2|\bar{\mu}) - P(3|\bar{\mu}) = P(3|\bar{\mu}) - P(4|\bar{\mu}) = 0.01 \left[\frac{5}{0.34} \right] = \frac{5}{34}.$$

Hence,

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

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

6 Welfare and Policy

6.1 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 (6) that is common across types. The corresponding value function $V(A, \mu)$ therefore provides a universal measure of subjective wellbeing. We can therefore analyze the perceived effect of any policy by studying the response of $V(A, \mu)$. A look at (6) 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.[6]

There is another interesting notion of welfare. Since our agents may hold incorrect

[6] The optimal policies (7) 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.

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

$$\tilde{V}(A,\mu,g^*) = \max_{\{P(\omega,i)\}_{i\in A,\omega\in\Omega}} \sum_{\omega\in\Omega} g^*(\omega)\left(\sum_{i\in A} P(i,\omega|\mu)u(i,\omega)\right)$$
$$-\lambda\left[\sum_{\omega\in\Omega}\mu(\omega)\left(\sum_{i\in A} P(i,\omega|\mu)\ln P(i,\omega|\mu)\right) - \sum_{i\in A} P(i|\mu)\ln P(i|\mu)\right]$$

Note that here we retain μ in the information cost as we interpret the learning cost as subjective.[7] However, since in steady state choice is made as if the prior were g^*, 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 implies the following proposition.

Proposition 3 In steady state, given a set of chosen options $\bar{A}$, the agent maximizes the following expected utility.

$$V(\bar{A},g^*) = \max_{\{P(i|\mu)\}_{i\in\bar{A}}} \sum_{\omega\in\Omega} g^*(\omega)\log\left(\sum_{i\in\bar{A}} P(i|g^*)\exp(u(i,\omega)/\lambda)\right), \tag{15}$$

which also equals social welfare.

The immediate implication is:

Corollary 2 In steady state, given $\bar{A}$, the market shares $M(i|\bar{\mu})$ are such that they maximize (15).

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, it follows directly that an expansion of the set of chosen options weakly cannot reduce long run social welfare.

[7] 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 (6) and what exactly the information cost represents. Our interpretation of the maximization problem (6) 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. Given prior beliefs, the agent chooses the strategy of gathering information that for a given expected payoff minimizes the expected cost of information. The realized cost can be different since it depends on the true type. The expected 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.

Corollary 3 Let $\bar{A}, \bar{B}$ be two sets of chosen options such that $\bar{A} \subset \bar{B}$. Then

$$V(\bar{A}, g^*) \leq V(\bar{B}, g^*).$$

6.2 Handicapping Policies

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

One such setting involves antitrust policies. Typical arguments for such policies are based on the level of prices. When competition is low, the firms tend to charge higher prices. By contrast, our model provides a rationale for policies that address the nature and quality of chosen options more than their prices. Specifically, a policy that limits market shares of the market leaders provides room for new entrants whose quality would be tested by heterogeneous buyers. Selection decisions of agents testing these new entrants would generate positive information externalities for future generations of agents. An appropriate antitrust policy would incentivize experimentation to expand the set of chosen products, and thereby increase long run welfare.

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. Direct computation reveals this to be so in example 1. In this example, 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 $g^*(\omega)$. At this point all taxes and subsidies could be removed and the market would settle to the social optimum. Note that this would involve choice only of the most popular goods according to $\bar{\mu}(\omega) = g^*(\omega)$, 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.

6.3 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. They 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 $u(i,\omega)$ 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 $u(i,\omega)$ will typically increase the expected utility net of information cost of all market participants in steady state. Unlike in Smallwood and Conlisk (1979), welfare in (15) obviously achieves a higher value after such a change to $u(i,\omega)$ in case all previously chosen goods are still chosen. Again, full consideration of the role of regulation is beyond the scope of the paper.

6.4 Heterogeneity and Welfare

The form of type-specific stochastic demand in equation (7) 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. By way of illustration, we can calculate type specific demands in example 1 above. Direct computations shows that type 1 chooses correctly 64% of the time, type 2 does so 49% of the time, while types three and four choose correctly 30% and

6% of the time respectively. The remaining types (66% of the population) never choose the correct option. They would be better off choosing randomly.

7 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-dependent market shares $M(\omega, i|\bar{\mu})$. Because agents tend to choose options that they prefer, these type-dependent market shares will be very informative about agents' preferences.

The following proposition shows how to use type-dependent market shares to recover preferences. A simple ratio test allows one to infer agents' preferences and reveals optimal choices by type.

Proposition 4 In the steady state,

$$\frac{u(i,\omega)-u(j,\omega)}{\lambda} = \log\left(\frac{M(\omega,i|\bar{\mu})}{M(\omega,j|\bar{\mu})} \Big/ \frac{M(i|\bar{\mu})}{M(j|\bar{\mu})}\right). \qquad (16)$$

The proposition follows directly from type-specific stochastic choice (7) and the observation that in steady state the unconditional choice probabilities $P(i|\bar{\mu})$ are equal to the market shares $M(i|\bar{\mu})$.

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. Even if an agent of type ω chooses option i very rarely and option j quite frequently, if they choose i relatively more frequently than does the average agent, one can infer that they in fact prefer option i to option j.

Note that since choice depends only on $\frac{u(i,\omega)-u(\omega,j)}{\lambda}$ we can normalize the payoff to one choice to zero for all types ω. Normalizing $u(\omega, j)$ to zero:

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

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^*.

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 (16), which emerges from the model, is very similar to Balassa's (1965) measure of "revealed" comparative advantage. Belassa measures 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 to option k if the ratio of individual to market choice for that agent is higher for option i than for option k. 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 exaggerate the influence of popular types, thereby biasing inference. Inference is more straightforward from the type-dependent demands $M(\omega, i)$. To capture this in our setting, we associate the choices A with differentiated products and the types Ω with groups of heterogeneous consumers. We then let the utility of each product depend on the value of the product and its market price,

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

Note that we assume that agents see and understand prices. Their inference problem is one of figuring out which product is right for them. To match prominent specifications in the differentiated product setting, we suppose that the value of product 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:

$$\log\left(\frac{M(\omega, i)}{M(\omega, j)} \bigg/ \frac{M(i)}{M(j)}\right) = (X^i - X^j)\frac{\beta_\omega}{\lambda} - \frac{\alpha}{\lambda}(p_i - p_j) + \xi. \tag{17}$$

where ξ is a regression error reflecting measurement error or omitted factors.[8] The regression is similar to the standard logit model of McFadden (1974).[9] If one had data on type-specific market participation, one could identify one of the products with the outside option. Alternatively, one may normalize the utility of one product for each type of agent to zero without affecting behavior.

The main difference between (17) 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 lies in the effect of changes in prices on market shares. In the logit model

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

[9] Note the model provides an alternative interpretation of stochastic choice data. Stochastic choice of one specific type in data is typically taken to refect unobserved heterogeneity. In our model such patterns may instead be driven by randomness in the outcome of attentional effort and the resulting pattern of choice mistakes.

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,

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

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 average choice probabilities. As a rise in p_i tends to reduce $P(i|\bar{\mu})$ and increase $P(j|\bar{\mu})$, 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 from 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 assumptions are 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 informed agents in fact choose their most preferred option and, absent other measurement issues, this approach would recover true preferences. If $\lambda > 0$ that above suggests that richer methods of inference are required, as observed choice tends to be skewed in the direction of choices that are perceived to be popular.

8 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.

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

[11] Caplin, Andrew, Mark Dean and John Leahy (2015), "Rational Inattention, Optimal Consideration Sets, and Stochastic Choice," 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 (2004), *Rational Herds*, Cambridge: Cambridge University Press.

[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] Csiszár, I. (2008), "Axiomatic characterizations of information measures," *Entropy*, 10(3), 261-273.

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

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

[20] Iyengar, Sheena S., and Mark R. Lepper (2000), "When choice is demotivating: Can one desire too much of a good thing?,"*Journal of personality and social psychology* 79.6, 995.

[21] Laffont, Jean-Jacques (1989), *The Economics of Uncertainty and Information*, Cambridge, MA: MIT Press.

[22] Luo, Yulei (2008), "Consumption dynamics under information processing constraints," *Review of Economic Dynamics*, 11(2):366-385.

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

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

[25] Matějka, Filip, and Alisdair McKay (2015), "Rational Inattention to Discrete Choices: a New Foundation for the Multinomial Logit Model," *American Economic Review*, 105(1),272-98.

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

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

[28] Mondria, Jordi (2010), "Portfolio choice, attention allocation, and price comovement, *Journal of Economic Theory*, 145(5):1837-1864.

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

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

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

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

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

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

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

[36] 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^* that is orthogonal to the unit vector

$$(g - g^*) \cdot 1 = 0$$

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

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

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, $g^* \in \Gamma_{t+1}$ if $g^* \in \Gamma_t$.

As $|\Omega|$ is finite Γ_t converges in a finite number of periods.■

Proof of Corollary 1:

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

$$f(g, i, t) = \sum_{\omega \in \Omega} \frac{g(\omega) \exp(u(i, \omega)/\lambda)}{\sum_{j \in A} P^t(j) \exp(u(\omega, j)/\lambda)}. \tag{18}$$

Equations (7) and (2) imply that for all i such that $M^t(i) > 0$, then $P^t(i)$ is also positive and the following holds:

$$f(g^*, i, t) = \frac{M^t(i)}{P^t(i)}. \tag{19}$$

Similarly, (7) together with the fact that the sum of probabilities in a distribution equals

1 implies:

$$f(\mu^t, i, t) = 1. \tag{20}$$

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 $g^* \in \Gamma_t$ implies $f(g^*, i, t) = 1$ 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 $N = \{i \in A | M(i) > 0\}$. Since $g^* \in \bar{\Gamma}$, $f(g^*, i, t) = 1$ for $i \in N$. But then

$$\sum_{\omega \in \Omega} \frac{g^*(\omega) \exp(u(i, \omega)/\lambda)}{\sum_{j \in A} P(j) \exp(u(\omega, j)/\lambda)} = 1, \qquad \forall i \in N,$$

which means that P satisfies the necessary and sufficient conditions for optimality for the prior equal to g^* and an option set N. ■

Proof of Proposition 4:

According to (7)

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

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

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

So that

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

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