DEPARTMENT OF ECONOMICS

WORKING PAPER SERIES

203; -0:

McMASTER UNIVERSITY

Department of Economics
Kenneth Taylor Hall 426
1280 Main Street West
Hamilton, Ontario, Canada
L8S 4M4

<http://www.mcmaster.ca/economics/>

Consumer Learning and Firm Dynamics ^{*}

Zachary Mahone[†]

Filippo Rebessi[‡]

September 3, 2019

Abstract

We propose a general equilibrium model of industry where consumers learn about firms' unobserved product quality over time. Because consumers learn through purchase decisions, price setting is a crucial lever through which firms manipulate future demand. We map equilibrium policies to a range of empirical evidence on industry, firm, product and price dynamics. We then study how firms respond as consumer information varies. Specifically, we show that firms exacerbate information problems by constraining learning more aggressively in those markets where consumers are less informed. Developing an indicator of consumer information by product category, we find these are typically markets for consumer durables. Finally, the efficiency implications of this behavior and interaction with size-dependent policies are explored.

Keywords: Learning, Firm Dynamics, Product Quality, Welfare

JEL Classification: E23, D83, L11

^{*}We would like to thank participants at the Midwest Macro conference in Athens, Georgia, the macro bag lunch at the University of Toronto and the WIMP at McMaster University. Colleagues at McMaster University and Cal State East Bay gave important feedback on early drafts, particularly Bettina Brueggemann. We also thank Edouard Schaal very useful discussions in the early stages of development. All errors are ours.

[†]Corresponding Author. Department of Economics, McMaster University, Hamilton, Ontario, Canada, L8S 4M4. Email: mahonez@mcmaster.edu, Phone: 1-905-525-9140 ext. 23814

[‡]Department of Economics, California State University East Bay, Hayward, CA USA, 94542. Email: filippo.rebessi@csueastbay.edu

1 Introduction

In many industries, the quality of products consumers purchase is difficult to observe. A restaurant is hard to evaluate without trying. A new technology may turn out less useful than initially thought. This uncertainty requires consumers to form beliefs about product quality, which they may learn over time. For producers in these industries, firm dynamics are then intimately linked to consumer learning and the evolution of beliefs about product quality: which firms grow or shrink, enter or exit, depends on which products rise or fall in consumer favor. Welfare and efficiency depend not just on which firms are active but also which products they are producing and how those products are perceived. The relevance of this problem is made manifest in the explosion of consumer review websites and the prime importance placed on brand and reputation management by firms.¹ This paper introduces a framework to jointly study consumer learning and industry dynamics. The key interaction distinct from existing models is that because consumers learn from consumption decisions, pricing choices are a crucial lever through which firms influence future demand. As a result, changes in consumer information yield changes in pricing behavior, with rich implications for industry concentration, firm growth and price and product switching dynamics.

This paper explores firm and industry dynamics in an environment where product qualities are difficult to observe and both firms and consumers learn over time. We study a market where a large number of long-lived firms produce differentiated products of unknown quality. Firms are allowed to enter, exit and may choose to replace products. The fortunes of these firms are tied to the market belief about their product's quality, which is akin to a brand or reputation. A higher belief allows firms to charge higher prices and earn greater profits, and these beliefs fluctuate as consumers learn. At the same time, because learning takes place through consumption decisions, pricing choices are a crucial lever to influence future demand. The consumers buying these products enter the market only once and have limited private information about product quality. They make purchase decisions based on their own information and public beliefs. Because they enter sequentially, consumer learning requires information to flow across periods, from past consumers to the current one, and we must specify a framework to govern this. We adopt the social learning model of Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992) which assumes that private signals cannot be observed but may possibly be inferred depending on the publicly known belief, price and purchase behavior of others. This inference rule gives rise to the key interaction of the model: through pricing choices, firms can control whether private signals are revealed and hence whether consumer learning takes place. This is not a trivial problem for the firm since learning is risky. Public beliefs may deteriorate, forcing the firm to either exit or pursue costly product replacement.

¹See Keller and Lehmann (2006) for a review of the literature discussing the importance of brand management and existing research. An example of the expansion in online reviews can be seen in the exponential growth of cumulative Yelp reviews since inception in 2005. There are currently 184 million reviews as of March 31, 2019. <https://www.yelp.com/factsheet>

In the stationary equilibrium we study, all new products enter with the same initial public prior. Equilibrium pricing rules imply that firms constrain the public belief about their product to an interval around it. If the public belief falls sufficiently below the initial prior, firms stop the learning process and either exit or pursue costly product replacement. If it rises sufficiently above, learning also stops, as the current flow of profits outweighs the expected gain from allowing beliefs to evolve. The size of the resulting learning space between these two thresholds is endogenously determined and accompanied by a rich set of empirical implications, entirely driven by the product level learning process. As in other models of firm learning, our equilibrium features a skewed firm size distribution and declining market shares by cohort age of firms. It also generates conditional exit and growth rates by age and size of firms that are consistent with the data. A feature unique to our framework is that product switching and price dynamics are jointly driven by the learning process. We show that the model is able to capture empirical regularities along both dimensions. Bernard et al. (2010) find that the likelihood of product dropping declines in the age of the product and size of shipments. In the model, product age and shipment size are both positively correlated with public beliefs, which generates the declining hazard rates of product dropping. Argente and Yeh (2018) document that the frequency and size of price adjustments decline in product age. In the model, because older products are more likely to have no learning occurring (at the upper threshold of the learning space) or update less, equilibrium prices follow similar age dependent patterns.

We next turn to the role of consumer information. In the market we study, private information is equally limited across consumers (specifically, the precision of the noisy signal received by consumers is identical, while the draws are independent and identically distributed). However, it is reasonable to think that consumer information may vary across markets or over time in systematic ways. Across industries, as products are increasingly differentiated or consumers are in the market less often (in the case of durable goods, for example), consumers' private information about product quality likely declines. Alternatively, the development of the internet has yielded an explosion of sites in which regular consumers and experts review every kind of product imaginable. This likely means that in any given market consumers arrive better informed than in the past. We use the model to study how such differences in consumer's private information about unobserved product quality affect industry dynamics. We find that firms exacerbate information problems by constraining learning more aggressively in those markets where consumers arrive less informed. This is because less informed consumers rely more heavily on the public belief, giving a stronger incentive for firms to manipulate it and making learning relatively more risky. In equilibrium this implies a smaller learning space, and more muted firm, product and price dynamics which we explore in detail. We then present some evidence on the variation of consumer information across product markets by developing an indicator using product reviews. Consistent with the intuition presented above, markets with more limited information are dominated by durable goods and relatively opaque service industries like financial services.

Finally, we turn to the policy implications of consumer learning. As a benchmark for comparison, we begin with a planner’s problem. We numerically characterize the constrained efficient solution in which consumers must still learn through prices but firm policies are controlled by a planner. Compared to the planner’s solution, the competitive equilibrium features inefficiently high product replacement as profit seeking firms are too quick to discard a product with a relatively low public belief. This arises because firms internalize neither the quantity-quality tradeoff available to the planner nor do they recognize that upon exit a new firm, with a new product, will enter and that this may be socially desirable. The resulting planner policies are more patient, allowing beliefs to deteriorate further (with lower average product quality) while providing a larger quantity of output. We then compare the planner’s policy to the competitive outcomes in markets with a progressive corporate tax, which we view as a stylized approximation to size-dependent policies such as small business subsidies or employment regulations that bind on large firms. Our focus on size dependent policies is motivated by their widespread use (see Guner et al. (2008), for a similar motivation in a different model). While the relevance of this class of policies for output and productivity has been investigated in other settings, our model reveals a cost of these policies that, to our knowledge, is novel in the literature. Progressive taxation reduces the value of learning in equilibrium, shrinking the learning space and reducing selection and productivity in the market. Product replacement declines, as in the planner’s solution, but at the cost of a much less dynamic market. Welfare falls below the competitive equilibrium. Comparisons across markets show that the cost of size dependent policies is increasing as consumers are worse informed.

Existing models of firm dynamics have studied tradeoffs between static profits and information accumulation in the context of one-sided firm learning. For example, in the active learning environments of Bachmann and Moscarini (2012) or Argente and Yeh (2018), pricing choices affect the extent of information accumulation by the firm but have no impact on demand.² In our framework, the presence of consumer learning implies that if prices influence information accumulation, they also indirectly influence demand. This distinction is crucial, since firms use prices to manipulate information accumulation and future sales. Low-quality products may persist much longer in this environment than in a one-sided learning framework, as firms impede consumer learning through pricing strategies. High-quality products may not be fully revealed since learning is risky, weakening selection and lowering their market share.

In our framework, firms invest by setting prices that allow learning at the expense of current profits. However, alternative ways have been considered in which firms may invest in future demand. In models of customer acquisition, for example, consumers are unaware of or unmatched with firms, and costly effort can be exerted to contact them. Gourio and Rudanko (2014) study the implications

²A recent example of passive learning is Arkolakis et al. (2018), who bring monopolistic competition to the classic model of Jovanovic (1982) and calibrate it to Colombian data, finding a quantitative role for learning in explaining firm dynamics. Empirical evidence on the importance of demand learning for firms can be found in Foster et al. (2016) or, for active learning through pricing, in Gaur and Fischer (2005).

of costly customer acquisition in a model of consumer search. Dinlersoz and Yorukoglu (2012) develop a related model to study the impact of a decline in the cost of advertising on firm and industry dynamics.

While we ignore strategic interactions and explicitly model industry dynamics, the link between pricing and learning at the heart of our model is closest in spirit to several game theoretic studies. Bergemann and Valimaki (1997), for example, study two-sided learning in a duopoly. Related environments have been studied in Bergemann and Valimaki (1996), Bergemann and Valimaki (2000) or Villas-Boas (2004).

Finally, the specific form of learning we adopt, the observational model of Bikhchandani et al. (1992) and Banerjee (1992), has been studied extensively in the context of consumer behavior. Much of the existing work comes from the industrial organization literature, with datasets ranging across a variety of markets. Zhang (2010) finds evidence of observational learning using data from the U.S. kidney market. Cai et al. (2009) conduct a randomized field experiment in a restaurant setting and show strong support for observational learning among consumers. Hendricks et al. (2012) study a model of observational learning and show its predictions consistent with data from an online music market. Chen et al. (2011) exploit a natural experiment to study word of mouth versus observational learning on Amazon markets, finding support for both types of social learning. While many of these papers build on partial equilibrium consumer models, there has been a growing extension of observational learning to other environments. Garcia and Shelegia (2018) incorporate it into a search framework. Perhaps closer to the present paper, observational learning has been introduced into general equilibrium models with production, although primarily focusing on investment dynamics. Loisel et al. (2012) study monetary policy in a world where entrepreneurs invest in a technology of uncertain productivity. Taschereau-Dumouchel and Schaal (2018) embed a similar idea into a real business cycle framework and investigate the model’s ability to generate technology-led boom and bust cycles. Our paper complements both literatures, yielding macro and micro predictions in a model of observational learning while incorporating dynamic pricing considerations.

After developing the model in Section 2, Section 3 investigates the model’s ability to replicate empirical patterns at the industry, firm, product and price levels. Section 4 turns to the role of consumer information. We first use numerical simulations to study the implications of changes in consumer information and highlight testable predictions. We then develop an empirical measure of information by product category to identify markets with limited consumer information. In Section 5 we study the constrained efficient benchmark and then compare the impact of a generic size-dependent policy on market outcomes. Section 6 concludes.

2 Consumer Model

In this section we describe an equilibrium industry model of consumer learning. Time is discrete, with infinite horizon. Firms are infinitely lived and produce one product at a given point in time, although they may have the option to switch products in a given period. Consumers are one-period lived and arrive sequentially. We begin by describing the information set at the beginning of each period and then discuss the consumer, firm and entrant problem before defining the stationary equilibrium. Because of the dynamic nature of the problem we will use a recursive form and drop all time subscripts.

2.1 Initial Information

At the beginning of any period, some mass I firms are in operation. Each firm produces a product that takes on one of two values $\nu \in \{\nu_l, \nu_h\}$ where ν is the unobserved, true quality. Market information about the quality of the product is summarized by a public belief σ , which is the probability that $\nu = \nu_h$. For each product, the consumer receives a private signal $x \in \{h, l\}$, where the precision of the signal is $\gamma > 0.5$. That is, the signal $x = j$ is received with probability γ when $\nu = \nu_j$. The resulting posterior $b(x, \sigma)$ is then updated according to Bayes' Rule, so that the private belief of a consumer is

$$b(x, \sigma) = \frac{Pr(x = j | \nu = \nu_h)}{Pr(x = j)} \quad (1)$$

As will become apparent below, a firm's only state variable is the public belief about their product quality σ . With atomistic firms, we assume that all producers with the same belief will act identically and construct our equilibrium over the distribution of beliefs rather than firms. This distribution is μ where $\int_0^1 \mu(\sigma) d\sigma = I$.

2.2 Static Consumer Problem

Each period a new consumer enters the market, observing public beliefs and receiving private signals. They then decide which products to purchase (including the option of a riskless consumption good $\tilde{C}$), and how much labor to provide.

At the beginning of each period, the consumer observes the distribution $\mu(\sigma)$, prices $p(\sigma)$ and receives their product specific signals $x \in \{h, l\}$. Let $\pi(x, \sigma)$ denote the probability of signal x given belief σ (and known precision γ). Under correct initial priors on new products, this will also be the fraction of products with belief σ that actually receive signal x . Consumers then make zero-one purchasing decisions for each producer. Within a given product group (x, σ) , the assumption that firms price identically implies the consumer is indifferent between products. A

consumer then randomizes across producers, purchasing a mass $\delta(x, \sigma, p) \in [0, \mu(\sigma)\pi(x, \sigma)]$ of the available products. Given this purchasing decision, the aggregate quality-adjusted consumption of these products is

$$C(\delta) = \int_0^1 \sum_x \delta(x, \sigma, p) [b(x, \sigma)\nu_h + (1 - b(x, \sigma))\nu_l] d\mu(\sigma) \quad (2)$$

Notice that while the value of a given product is unknown, there is no risk in the aggregate bundle since the average quality of products in group (x, σ) is known with certainty. Taking the distribution μ , prices p and prevailing wage w as given, consumers solve:

$$\begin{aligned} & \max_{\delta, L} U(C(\delta)) + \tilde{C} - V(L) \\ & \text{s.t.} \\ & wL + \int_0^1 \sum_x \Pi^*(x, \sigma) d\mu(\sigma) \geq \int_0^1 \sum_x \delta(x, \sigma, p) p(\sigma) d\sigma + \tilde{C} \\ & \delta(x, \sigma, p) \in [0, \mu(\sigma)\pi(x, \sigma)] \quad \forall (x, \sigma) \\ & L \leq \bar{L}, \quad \tilde{C} \geq 0 \end{aligned}$$

That is, consumers choose how much total labor to offer in the market, the set of products to buy, and the quantity of a riskless consumption good $\tilde{C}$ that enters linearly into preferences. These purchases are funded by the labor income earned and profits (Π^*) from owning the set of firms. The solution to this problem yields a purchasing rule over products $\delta^*(x, \sigma, p)$.

2.3 Belief Evolution

Before discussing the firm problem we now briefly describe how beliefs evolve. Firms and future consumers observe prices, sales and understand the purchasing rule $\delta^*(x, \sigma, p)$. In the frameworks of Banerjee (1992) and Bikhchandani et al. (1992), if $\delta^*(l, \sigma, p) \neq \delta^*(h, \sigma, p)$ then purchase decisions are sufficient to infer signals. In our general equilibrium framework however, ensuring budget balance requires us to permit that not all products in a set (x, σ) are purchased (i.e. interior values of δ). This potentially complicates the inference problem since sales decisions in principle depend on prices, signals and the budget constraint. For the moment, we assume purchase decisions depend only on prices and signals, hence are all-or-nothing for each product type (x, σ) , and we will show below that such a purchase rule is indeed consistent with equilibrium. Inference then yields the following updating rule

$$\tilde{\sigma}(x, \sigma, p, \delta) = \begin{cases} b(x, \sigma) & \text{if } \delta^*(l, \sigma, p) \neq \delta^*(h, \sigma, p) \\ \sigma & \text{o.w.} \end{cases} \quad (3)$$

That is, if different signals lead to different purchase rules for some product, then the signal can be inferred and public beliefs will update. Otherwise, they remain unchanged. This yields a post-inference distribution of beliefs

$$\tilde{\mu}(z) = \int_0^1 \mu(\sigma) \sum_x \pi(x, \sigma) I_{[\tilde{\sigma}(x, \sigma, p, \delta)=z]} d\sigma \quad \forall z \in [0, 1]$$

2.4 Firms

We now consider the problem of a firm producing a product with arbitrary belief σ . We assume all firms must hire one unit of labor as a fixed production cost each period. Writing the problem recursively, given the consumer purchasing rule and interest rate R the firm solves

$$J(\sigma) = \max_p \sum_x \pi(x, \sigma) \left[\hat{\delta}^*(x, \sigma, p)p - w + R^{-1} \tilde{J}(\tilde{\sigma}(\sigma, p, \delta)) \right]$$

where expectations are taken over the private signal received by the consumer and $\hat{\delta}(x, \sigma, p) = \frac{\delta(x, \sigma, p)}{\mu(\sigma)\pi(x, \sigma)}$ is adjusted as the probability a given product is sold, where we assume consumers randomize over products of type (x, σ) when some but not all are purchased. The continuation value for the firm $\tilde{J}$ incorporates product switching and exit choices. Existing products obsolesce independently at a fixed rate d_p and we assume new product ideas arrive with probability λ , fixed initial public prior $\bar{\sigma}$ and implementation cost c denominated in labor units.

$$\begin{aligned} \tilde{J}(\tilde{\sigma}) &= d_p \lambda \max\{J(\bar{\sigma}) - c, 0\} \\ &+ (1 - d_p) [\lambda \max\{J(\bar{\sigma}) - c, J(\tilde{\sigma}), 0\} + (1 - \lambda) \max\{J(\tilde{\sigma}), 0\}] \end{aligned}$$

The first term of the expression denotes the case in which the current product is obsolescent at the beginning of the period. With probability λ a firm obtains a new idea and may choose to implement it, or exit. With probability $1 - \lambda$ the firm has no new idea to implement and must exit, obtaining value zero. The second term describes the similar set of options available to a firm whose current product does not obsolesce independently.

In the firm problem, prices control both expected static profits and continuation values given the evolution of the public belief. The solution to the firm problem results in a pricing rule $p^*(\sigma)$, switching rule $r^*(\tilde{\sigma}) \in \{0, 1\}$ and exit rule $e^*(\tilde{\sigma}) \in \{0, 1\}$. It also yields the expected static profits used in the consumer problem

$$\Pi^*(x, \sigma) = \sum_x \pi(x, \sigma) \hat{\delta}^*(x, \sigma, p^*(\sigma)) p^*(\sigma) - w$$

2.5 Entrants and Updating the Distribution

At the end of each period new entrants may arrive by paying a fixed cost κ in labor units to enter. We assume they enter with a new product in hand. Free entry then implies

$$J(\bar{\sigma}) \leq \kappa$$

Let the associated entry flow be e_n . Then, given $(\tilde{\mu}, e_n)$ the belief distribution evolves, for any $\sigma \in [0, 1]$, according to

$$\mu'(\sigma) = \begin{cases} \tilde{\mu}(\sigma)(1 - r^*(\sigma))(1 - e^*(\sigma)) & \text{if } \sigma \neq \bar{\sigma} \\ \tilde{\mu}(\sigma)(1 - r^*(\sigma))(1 - e^*(\sigma)) + e_n + \int_0^1 \tilde{\mu}(\sigma)(1 - e^*(\sigma))r^*(\sigma)d\sigma & \text{if } \sigma = \bar{\sigma} \end{cases}$$

For all beliefs not equal to the entry value $\bar{\sigma}$, the mass is simply the post-inference mass adjusted for replacement and exit decisions. At the entry value we must additionally account for entering firms and existing firms that switch products. Figure 1 summarizes the timing of all the choices described above.

Figure 1: Timing

Finally, we define a stationary equilibrium of this economy.

Definition 2.1. A stationary equilibrium is a wage w , labor supply L , entry flow e_n , belief distribution $\mu(\sigma)$, pricing rule $p^*(\sigma)$, purchasing rule $\delta^*(x, \sigma, p)$, switching rule $r^*(\sigma)$ and exit rule $e^*(\sigma)$

1. L and δ^* solve the consumer problem.
2. p^*, r^*, e^* solves the individual firm problem for a given σ and inference rule.
3. Labor markets clear: $L = I + e_n\kappa + c\lambda \int_0^1 \tilde{\mu}(s)r^*(s)(1 - e^*(s))ds$.
4. The free entry condition binds.
5. Beliefs update following inference rule and distribution is stationary $\mu' = \mu$.

2.6 Equilibrium and its Properties

Before investigating the dynamics of the model we characterize several features of policy functions in the stationary equilibrium we study. These help convey the intuition of how the model works. Recall that the inference rule in Equation 3 requires that consumption decisions of products are never budget constrained. We first show that there is a unique consumption rule consistent with both equilibrium and inference.

Proposition 2.2. *There is a unique purchasing rule consistent with both the equilibrium defined above and inference rule from Equation 3, taking the form*

$$\delta(x, \sigma, p) = \begin{cases} 1 & \text{if } U'(C)v(x, \sigma, p) \geq 1 \\ 0 & \text{o.w.} \end{cases}$$

where $U'(C)v(x, \sigma, p)$ denotes the expected marginal utility of an additional unit of a product type (x, σ) , with quoted price $p(\sigma)$.

Further, under this purchasing rule $\tilde{C} = 0$ in equilibrium.

Proof. All proofs are in Appendix A. □

The proposition above shows that we can indeed study an equilibrium consistent with the form of inference we impose.³ Access to the risk free consumption good makes such a rule intuitive, and the ownership of firms by consumers allows for such a consumption rule to be essentially self-financed in equilibrium. This self-financing also implies that $\tilde{C} = 0$.

Properties of the firm policies and equilibrium prices can also be characterized.

Proposition 2.3. *In the stationary equilibrium associated with the purchase rule above,*

1. *Firms optimally chooses between one of two prices:*

$$\begin{aligned} p_h(\sigma) &= U_{C,h,\sigma} \rightarrow \text{learning} \\ p_l(\sigma) &= U_{C,l,\sigma} \rightarrow \text{no learning} \end{aligned}$$

2. *The equilibrium learning space has interior bounds and number of points N is finite.*
3. *The equilibrium pricing rule is non-monotonic in belief σ , and single peaked (increasing and then falling).*

³It turns out however that we cannot rule out other equilibria. Alternative equilibria may be possible in which binding budget constraints yield randomization across products of type (x, σ) . In these economies inference is complicated by the fact that a non-purchase may arise due to a bad signal or budget constraints. The inference problem in Hendricks et. al. (2012) has a similar flavor.

The two-point pricing problem is intuitive and gives a simple mapping between pricing choices and learning in the model. Because purchasing decisions are zero-one for consumers, there is a natural cutoff solution, associated with the highest willingness to pay conditional on good (p_h) or bad (p_l) signals. The choice between these two creates a potential tradeoff between profits and information, as in other active learning environments. The difference here is that more learning is not always better for the firm. This is precisely because in our two-sided learning environment, future consumers may actually learn they do not like the product, in which case the firm would have been better off shutting down learning in the first place. These tradeoffs yield an upper bound, above the initial prior $\bar{\sigma}$, for when firms allow learning. Once beliefs about their product are sufficiently good, it is profitable for firms to shut down learning and price low, $p_l(\sigma)$.⁴ The exit option yields a similar bound below the initial prior, such that when beliefs are sufficiently bad the firm will exit (or may shut down learning and wait for a new product). The result is a finite learning space on the support of beliefs. The dynamic pricing policy in the third part of the claim arises from these choices around learning. While the high and low prices in Proposition 2.3 are strictly increasing in σ , the resulting pricing policy from the firm maximization problem is not. This reflects the decision of firms to stop learning at sufficiently high belief σ . With N belief nodes in equilibrium, order them in increasing fashion such that N denotes the upper threshold of the learning space. To stop learning, firms with these products actually price below firms with products of belief σ_{N-1} , the second highest belief. The highest price charged in this economy then is for the second highest quality product, in which learning still occurs. This implies that given a time series of prices upon product introduction, the “long run” price at which it eventually settles should be strictly less than the highest price in the series.

An example of such a bounded space when $\bar{\sigma} = 0.5$ and $\gamma = 0.55$ is presented below in Figure 2, with the learning space illustrated by the bracketed gray points. The associated pricing policy would be to price high at the interior gray points and low at the bracketed gray points.

Figure 2: A Bounded Learning Space

2.6.1 The value of learning

For a firm, the value of learning may imply a tradeoff between static and dynamic profit maximization but this is not always true. In some cases the price that maximizes static profits and continuation values may coincide. To show the impact of learning explicitly, captured in the continuation value of the firm, we compare here learning spaces as the firm discount factor increases,

⁴In the language of Bikhchandani et al. (1992) and Banerjee (1992), this means that firms induce an information cascade, such that purchases are no longer informative (given the price and public belief).

Figure 3: Learning Space and Discount Factor

plotted in Figure 3.

We plot this learning space as an indicator function, where a value of one indicates learning (pricing high) and zero no learning (pricing low). This mapping generates an inverted U-shape. Notice that when $\beta = 0$, the implied pricing rule maximizes static profits alone and features regions of both high and low pricing. As we increase β the learning space expands, and this expansion is due entirely to the value of learning. In these regions where learning occurs only as β increases, the implication is that static profits are maximized by pricing low while the value of learning implies dynamic profits are maximized by pricing high. In the interior space, where learning occurs with $\beta = 0$, this tradeoff disappears.

3 Dynamics in the Stationary Equilibrium

The model developed above generates a rich set of predictions with empirical counterparts. In this section we discuss the performance of the model with respect to a range of empirical evidence at the industry, firm, product and price levels. We present a set of empirical patterns from the literature and argue that the model does a good job in qualitatively matching them.

3.1 Industry Dynamics

Empirical studies at the industry and firm level are widespread. At the risk of shortchanging many valuable papers, we summarize this literature, here and in the next subsection, mainly through the stylized facts [7-10] of Klette and Kortum (2004) (reproduced as Facts 1, 2 and 4-6). *Fact 3* relies on Dunne et al. (1989). We begin with considerations of industry.

Fact 1: The firm size distribution is skewed.

Fact 2: The market share of an entering cohort generally declines with age.

Fact 3: Across industries, there is a positive correlation between firm entry and exit rates.

Figure 4: Industry Dynamics in the Model

(a) Fact 1 - Skewed Distribution

(b) Fact 2 - Market Share by Age

(c) Fact 3 - Entry and Exit Correlation

Graphs of model simulations matching Facts 1-3 are reproduced in Figure 4. The skewed firm distribution of Fact 1 arises from the interaction of learning and the option of product switching or exit. In the model, firm size and belief about product quality move hand in hand. Hence, small firms are those producing products with lower public beliefs. Learning, along with a consistent flat prior such that half of all new products are truly low quality, generates a high rate of product churning or exit. This, along with exogenous obsolescence, yields a large mass of small firms in the stationary equilibrium and a fat right tail. Of course, the length of the tail depends on the number

of points in the support of the belief space, and a smaller support leads to a right-truncation in the distribution. Extending the model to include a heterogeneous firm productivity factor could smooth this out.

The decline in market share by age, Fact 2, arises from an interaction between survival and size. We compute market share by age as the share of total revenues accruing to firms of some age T . While true that older firms are larger (producing products with higher beliefs), there are also less of them in the stationary distribution. The latter effect dominates.

Finally, we turn to the positive correlation of entry and exit rates from Fact 3. We compute different industry equilibria by changing the precision of consumer signals, which affects rates of endogenous exit and entry. We then compute these statistics for each equilibrium and graph them in Figure 4c. Improving signal precision tends to increase exit and entry rates, as bad products are discarded more quickly and firms that cannot produce new ideas exit. However, the relative size of exiting firms also changes, which matters because new entrants must keep quality adjusted consumption constant in the stationary equilibrium. These general equilibrium effects yield a positive but somewhat noisy correlation between entry and exit.

3.2 Firm Dynamics

Following much of the focus on firm dynamics, we cite existing evidence regarding firm age and size, entirely reproduced from Klette and Kortum (2004):

Fact 4: (a) Small firms have a higher probability of exit and, (b) conditional on survival, grow faster than larger firms.

Fact 5: (a) Younger firms have a higher probability of exit and (b) conditional on survival, grow faster than older firms.

Fact 6: The variance of growth rates for small firms is higher.

Simulations of the model are presented in Figure 5 and statistics are computed at different intervals for robustness. Firm size is measured as the present discounted value of profits, which is increasing in the public belief on product quality. As a result, large firms are those that produce a product with high public beliefs. Exit and growth rates by firm size, Facts 4a and 4b, are then driven by prices, which in turn are driven by learning. Because exit (4a) is driven by a decline in the belief about product quality, larger firms are less likely to exit as it takes a longer sequence of publicly revealed negative signals before an exit decision is made. Growth rates by size (4b) are also driven by the learning environment. Smaller firms produce products with a public belief close to the flat prior. As a result, revealed signals are more informative and lead to larger changes in

Figure 5: Firm Dynamics in the Model

(a) Fact 4a - Exit by Size

(b) Fact 4b - Growth by Size

(c) Fact 5a - Exit by Age

(d) Fact 5b - Growth by Age

(e) Fact 6 - Growth Rate Variance

the public belief and subsequent revenues (prices). Conditional on survival, these will mostly be positive signals, yielding higher growth rates for smaller firms.

Survival and growth by firm age, Facts 5a and 5b, tell a similar story but also show the distinction between age and size in the model. While size is a function exclusively of product belief, age is not, due to the product switching option. There is however a positive correlation between age and size, so that older firms are more likely to have better products, and hence exit and grow less for the same reasons discussed above. The peak occurring at early ages is common in learning models, as exit will not happen immediately but rather after enough negative signals have been received, leading to an initial spike in exit and subsequent slower decline. Growth rates initially demonstrate a rapid decline, which slows with age.

Finally, the variance of firm growth rates by size, Fact 6, echoes the intuition from Figure 5b. Larger firms enjoy a higher public belief about their product quality. The result is that an additional public signal carries relatively less information, yielding smaller changes in the public belief and lower variance in growth rates.

3.3 Product Dynamics

As existing work has already shown the ability of learning to generate some of the industry and firm dynamics explored above⁵, it is perhaps not surprising that our model is also capable of doing so. To our knowledge however, ours is the only equilibrium model of industry in which product pricing and switching choices are jointly driven by learning at the product level. We turn to model predictions for these two choices now. Much of the existing evidence on product switching at the firm level comes from one of two sources, private scanner databases on consumer purchases or confidential Census micro data. Bernard et al. (2010), using the latter, report the following two “stylized” facts from their data.

Fact 7: The likelihood of dropping a product declines in product age.

Fact 8: The likelihood of dropping a product declines in size of shipment.

We show the ability of the model to match these facts in Figure 6, where measured shipment size equals the firm’s posted price. Again, we compute these patterns over different intervals, particularly as the frequency of data in the literature varies widely (5 years in the Census of Manufactures to daily or more in scanner data). The positive correlation between age and shipment size arising in the model suggests it should successfully reproduce either both or neither of Facts 7 and 8. Older products, having survived learning, enjoy higher beliefs on average and command a higher price, yielding larger shipments. Higher beliefs imply the likelihood of product dropping

⁵This was part of the motivation of Jovanovic (1982), for example.

Figure 6: Product Dynamics in the Model

declines with product age as a longer sequence of negative shocks is needed to induce the firm to switch products. This yields Fact 7. That higher beliefs yields larger shipments generates the same implications for Fact 8.⁶

3.4 Price Dynamics

Turning to price dynamics, we are aware of only one recent paper measuring age-dependent pricing for a wide range of products.⁷ Relying on scanner data from grocery stores, Argente and Yeh (2018) report the following empirical findings,

Fact 9: The frequency of price adjustment declines with product age.

Fact 10: The size of price adjustment declines with product age.

We present evidence of the model's ability to capture these facts in Figure 7. We report price changes only at one-period intervals since the evidence is very high frequency, although these patterns are also robust. Again, learning drives both of these results. Recall that price adjustment occurs in the model as long as products lie in the learning part of the belief space (off the upper and lower thresholds). The decline in the frequency of price adjustment, then, is a statement about the share of products of age T that are still in the learning space. Through selection, older products

⁶Notice that, as we increase the time horizon to four periods, the measured product dropping shows a sharp increase. This is due to the fact that we measure product dropping conditional on firm survival. Firms are more likely to exit at lower shipment sizes.

⁷Papers in the industrial organization literature, such as Copeland and Shapiro (2016), discuss pricing dynamics but in a restricted market and one in which innovations are occurring rapidly. Our model is not designed to capture pricing and product switching due to technical change.

Figure 7: Price Dynamics in the Model

(a) Fact 9 - Frequency of Adjustment by Age

(b) Fact 10 - Size of Adjustment by Age

enjoy higher beliefs by the public and as a result are more likely to be at the upper threshold of the belief space, where prices do not change, generating Fact 9. Further, because additional signals become somewhat less informative as beliefs become stronger, older products with higher public beliefs should also see smaller absolute sizes of price changes, generating Fact 10.

3.5 Relating to models of multi-product firms

In the model developed above, firms may choose to replace their existing product with a new one (i.e. product churning) if beliefs about product quality are sufficiently low. As a result, our model features multi-product firms across periods, although firms are single-product within a period.

In all of our simulations, a majority of firms are multi-product as measured over a relatively short number of periods.⁸ The corollary of this at the product level is that product churning rates within an industry are quite high. Further, by construction, product churning will happen almost entirely within the firm. These patterns are all supported by existing empirical evidence.

A variety of datasets have shown a large role for multi product firms in production. Using the ACNielsen Home Scanner database, Broda and Weinstein (2010) find that 60% of sales in their data come from multi product firms. For Bernard et al. (2010), using the Census of Manufactures, this number is 87%. Product churning is similarly widespread, as Broda and Weinstein (2010) find that 50% of products did not exist four years prior (these new products accounting for 37% of expenditures). Bernard et al. (2010) find 30-35% of firm output are products added in the last five years. Using Nielsen Retail scanner data, Argente et al. (2018) find that 1/3 of all products are created and destroyed in a given year. Finally, the concentration of product churning within, rather

⁸We define firms as multi-product if they have switched products at least once in some number (3-5) periods

than across, firms, comes through clearly in the data. Broda and Weinstein (2010) show that in a four year window, 82% and 87% of product creation and destruction, respectively, occur within the firm. Among multi product firms Bernard et al. (2010) find that 75% of output is produced by those firms that churn products, as opposed to simply add, drop or neither.

While definitions and datasets vary across these papers, the main message is clear. Multi-product firms are prevalent, if not dominant, in most industries and product churning is both ubiquitous and concentrated within the firm. This is consistent with our model of learning at the product level.

4 The Role of Consumer Information

A key driver of the stationary equilibrium analyzed above is the extent of consumer information, represented by the strength of the private signal consumers receive for each product. We argue that changes in this value likely derive from industry characteristics or technological advances over time. Across industries, examples of markets where consumers may be less informed are those that are less commodity focused, with greater product variety, or produce durable goods that are purchased only occasionally. Over time, the amount of information that consumers arrive with has increased in many markets thanks to the spread of consumer review websites. Analysis of a single stationary equilibrium has already shown that firms will manipulate how consumers learn through their pricing decisions, and intuition suggests that the motive for such manipulation depends strongly on how well informed consumers are to begin with.⁹ As pricing rules change, so will industry dynamics and welfare. In this section we study how firms respond to changes in consumer information and the implications for industry dynamics. We then construct a measure of consumer information across products markets to identify where these dynamics may be most relevant.

4.1 Consumer Information and Learning Across Markets

We begin by asking how firm decisions and learning change as consumers are variably informed i.e. the precision γ varies. Because changes in γ lead to general equilibrium effects on the distribution which are hard to characterize, we pursue a numerical approach in our analysis. In Figure 8 we vary consumer information and plot the learning space, with the same interpretation as in Figure 3.

The immediate result we see is that as consumers are better informed the learning space widens. Conversely, as consumers become worse informed, learning is more constrained. The intuition for this is that in markets where consumers are less informed, they rely more heavily on the public signal. This, in turn, gives a stronger incentive for firms to manipulate the signal and in equilibrium results in learning being riskier. We can see the logic of this in Figure 9a, in which we plot the value

⁹As one example, if quality were perfectly observable there would be no incentive to limit learning.

Figure 8: Consumer Information and Learning

Figure 9: Decomposing Changes in Learning

(a) Learning

(b) Static v. Dynamic Decomposition

functions of pricing high and low for two different levels of consumer information. As consumers are better informed, two things occur. First, the value functions flatten, which is consistent with the idea that as consumers rely less on public information, the relative value to firms of having a higher public signal is diminished. Second, the relative value of learning increases, as pricing high dominates in a larger region of the belief space. Columns one and two of Table 1 show the widening thresholds of the learning space.

It is also true of course that as γ increases, the static pricing rule ($\beta = 0$ from Figure 3) will alone imply a larger region for learning (that is, pure static profit maximization would yield more learning, ignoring any dynamic considerations). In Figure 9b we show that this only partially accounts for the growth of the learning space. Moving from the less informed benchmark (solid line) to the better informed, static market (dot-dashed line) we see the learning space expands

Table 1: Learning Spaces, Wages and Selection

γ	σ_{min}	σ_{max}	Wage (real)	Quality	Productivity
0.70	0.15	0.84	1.60	0.73	3.21
0.75	0.10	0.96	1.38	0.77	3.72
0.80	0.06	0.98	1.28	0.83	4.25

slightly and shifts right. The majority of the expansion in the learning space however is attributed to the changing value of consumer learning to firm profits (dashed line).

Table 1 summarizes the impacts of these changing thresholds of the learning space on productivity, selection and wages. While we focus on these aggregate moments here, Appendix B reproduces the stylized facts of Section 3 and provides other summary statistics that offer a more detailed comparison of stationary equilibria with different degrees of consumer information.

We see clearly that the bounds of the learning space increase as consumers are better informed. Consequentially, we would expect to see stronger selection in unobserved product quality, since firms allow quality to be identified with greater certainty. Indeed, we see that average product quality rises as consumers are better informed, along with productivity. It is worth noting here that as the learning space expands, two opposing forces are being exerted on productivity. First, the larger learning space means that, on average, more products are in the learning space. This is good for future consumers, but means that a larger share of firms will not produce in a given period, putting downward pressure on productivity. This is dominated however by the fact that, among the products produced, quality is higher, hence overall output per worker is rising in γ . The change in real wage may seem surprising, as wages are higher in less informed markets, but this highlights general equilibrium effects in the model. In less informed markets, quality-adjusted consumption is lower, raising marginal utility. This higher marginal utility implies higher prices and firm profits. Market clearing and free entry then combine to imply that wages must rise to bring profits back down.

4.1.1 Testable Predictions

As noted above, Appendix B lays out a detailed comparison across markets as consumer information varies. In this subsection we highlight a number of the most relevant, measurable, conclusions from these comparisons. Ideally, our discussion would contain a more direct test of these predictions against empirical evidence. To our knowledge however, the remarkable growth in high quality microdata over the last few decades has so far not yielded broad panel datasets at the

product-price level that are suitable to this purpose.¹⁰ Nevertheless, as guidance for future research on consumer learning and in the hopes that future data availability provides better measurement opportunities, we briefly discuss several predictions of the model as consumer information improves.

As consumers become individually better informed through their private signals, two things happen simultaneously. First, where learning occurs, it happens more quickly, as each purchase decision conveys more information. The result is that consumers converge more rapidly on which products are good and which are bad. This leads to faster exit, of products and firms, at the bottom of the distribution and greater dominance of large firms at the top. Second, as average quality rises, the marginal value of product quality declines. In spite of faster learning, this actually flattens positive price growth for new products. Jointly, these lead to four predictions.

1. Prediction 1: Faster growth of successful products and firms. Quicker exit of unsuccessful ones. There is a larger gap between the winners and losers.
2. Prediction 2: Greater right-skewness of the firm size distribution. Increased dominance of large firms.
3. Prediction 3: The likelihood of new product exit increases.
4. Prediction 4: For new products, price growth is flatter while price declines are steeper.

The general idea conveyed by the model, that improved consumer information leads to greater dominance of winners in the market, is consonant with the growing literature studying market concentration in the United States. In particular, it may provide an alternative interpretation of the winner-take-all story of “superstar” firms (Dorn et al. (2017)) or ICT price declines (Van Reenen (2018)). Our exercises suggest that if technological change has improved the ability and speed with which consumers identify the best (perceived) firms and products, this may also contribute to the observed growth in market concentration.

4.2 Which markets?

The numerical experiments point to a potentially important role for consumer information in determining consumer learning, welfare and productivity. But exactly which markets are those with limited information? Intuition suggests they should be markets that are visited less often and where quality is difficult to observe. In this section however we provide some more direct empirical evidence. Our approach is to construct an indicator of consumer information using product reviews. While our model does not explicitly include a channel for consumer reviews, they

¹⁰The literature on multi product firms typically relies on one of two sources. Census data will provide a direct mapping from products to producers but product categories are quite broad and prices are not available. The other source is detailed scanner data, such as the Nielsen panels. These contain high quality product-price level data but are largely restricted to grocery store products, making cross industry comparison difficult. Further, because product data is at the point of sale, product adding and dropping is observed at the store, rather than firm level.

Table 2: Summary Statistics from Consumer Reports Reviews, 2016

No. Product Categories	35
Avg. No. Reviews	2.85
Median No. Reviews	2.0
Avg. log price	4.31
Median log price	3.91

(a) A total of 190 reviews were counted in 2016. Collapsed so that a month-product pair counts once, 103 separate observations remain with 35 separate 3-digit product categories.

are a form of social learning with the virtue of being observable, and therefore measurable. The resulting indicator, while imperfect, does give some sense of the product markets where consumers may be worse informed and, consequentially, the dynamics we study may be most relevant.

The intuition for our indicator is simple - consumers with sufficient private information will rely less on public, or “market”, information. For products where consumers are well informed, we would expect less reliance on, and demand for, these reviews. In principle we would like the existence of reviews to reflect demand for information rather than a supply effect (it is not clear what many angry Amazon reviews tell us about demand for information by other consumers). To address this issue we turn to “Consumer Reports”, a subscription magazine providing professional reviews of consumer products. As a non-profit company supporting itself through subscription fees, we assume the choice of products reviewed at least partially reflects the degree of consumer demand for information across product categories.

We proceed as follows. For the year 2016 we extract the entire list of product reviews from their “5 Year Index”. Because product reviews are typically by category and mention several brands, for each product category we record the category, average price, and date of review. An example would be “Fitness Trackers” with an average price of \$150.00, reviewed July, 2016.¹¹ We next map each product category into the North American Product Classification System (NAPCS), a product-based system of classification complementary to the better known North American Industry Classification System (NAICS). For example, “Fitness Trackers” map into three-digit code “362 - Communication and audio and visual equipment”, five-digit code “36211 - Telephone and data communications equipment”, and seven-digit code “3621132 - Mobile telephones and other wireless communications equipment”. Among the “Illustrative examples” of this category is a “smart fitness tracker watch”.

Having mapped product reviews into NAPCS codes, we use a simple count of reviews at the three-digit product level as an inverse indicator of consumer information. That is, product categories with a higher review count identify product markets where consumers are less well informed. We briefly summarize the resulting dataset in Table 2.

¹¹See Appendix C for further details on Consumer Reports and the construction of this dataset.

The five product categories with the highest counts are *Passenger Cars and Light Trucks*, *Financial Services (except insurance)*, *Appliances*, *Communication and Audio and Video Equipment* and *Miscellaneous Goods and Supplies*. Consistent with intuition, we see three large consumer durables industries (automobiles, appliances and consumer electronics) along with a service sector that is both opaque and infrequently sampled. Financial services reviews covered themes such as student debt and home ownership financing, products the average consumer purchases only so many times. These are all examples of goods or services whose quality is hard to observe prior to purchase and where consumers are not regularly in the market.

5 Size Dependent Policies and Efficiency

Before concluding, we want consider the policy implications of consumer learning. Intuition suggests that in our model, with consumer learning driven by profit motives, policies that alter the payoff to firm growth will indirectly effect the market’s accumulation of information, selection and productivity. Following a number of papers (Garicano et al. (2016), Guner et al. (2008), Gollin (1995)), we restrict our focus to size-dependent regulations faced by producers. As a benchmark for welfare, we begin by numerically characterizing the constrained efficient policies of a planner who controls all firm decisions. We then consider the impact of a progressive corporate tax, meant to mimic a number of existing size-dependent policies, and analyze their impact on welfare through the relevant margins of the planner’s problem.

5.1 Planner’s problem

In the decentralized environment, firms face a trade-off between profits and learning. When they set prices, firms may prevent consumers from inferring product quality, giving rise to a bounded learning space. The amount of information on product quality arising in the competitive equilibrium is unlikely to be efficient, as firms do not internalize the impact of their learning choices on the stationary distribution and aggregate quantities. We examine the extent of this inefficiency with a numerical characterization of the allocations chosen by a social planner subject to the same information constraints faced by agents and firms.

The social planner chooses firms’ pricing, replacement and exit policies to maximize the representative consumer’s welfare in the economy’s stationary equilibrium. Learning takes place as in the competitive equilibrium: when a high price is selected, signals are inferred from purchasing decisions and the public belief on product quality updates in bayesian fashion. With low pricing,

no learning occurs. Formally, the planner solves the following problem

$$\begin{aligned}
& \max_{\mathbb{P}=\{\hat{r}(\bar{\sigma}), \hat{e}(\bar{\sigma}), \hat{p}(\sigma)\}} U(C) + v(1 - L) \\
& \text{subject to} \\
& C = \int_0^1 \sum_x \delta(x, s; \mathbb{P}) [b(x, s)\nu_h + (1 - b(x, s))\nu_l] ds \\
& L = \int_0^1 \mu(s; \mathbb{P}) ds + \kappa e_n(\mathbb{P}) + c\lambda \int_0^1 \tilde{\mu}(s, \mathbb{P}) \hat{r}(s)(1 - \hat{e}(s)) ds \\
& \tilde{\sigma}(\sigma, \hat{p}, \delta) = \begin{cases} b(x, \sigma) & \text{if } \delta(l, \sigma, \hat{p}) \neq \delta(h, \sigma, \hat{p}) \\ \sigma & \text{o.w.} \end{cases} \tag{4} \\
& \tilde{\mu}(z) = \int_0^1 \mu(s) \sum_x \pi(x, s) I_{[\tilde{\sigma}(s, \hat{p}, \delta)=z]} ds \quad \forall z \in [0, 1] \\
& \mu'(x, \mathbb{P}) = \mu(x, \mathbb{P}) \\
& \delta(x, s; \mathbb{P}) = \begin{cases} 1 & \text{if } \hat{p}(\sigma) \geq U_{C,x,\sigma} \\ 0 & \text{o.w.} \end{cases}
\end{aligned}$$

where $e_n(\mathbb{P})$ denotes the entry flow consistent with stationary equilibrium, given $\mathbb{P}$. It is worth re-emphasizing that the planner's policy $\mathbb{P}$ belongs to the same set of choices taken by firms in the competitive equilibrium. $\mathbb{P}$ induces a purchasing rule for the consumer δ . Jointly, δ and $\mathbb{P}$ determine the stationary distribution of firms over beliefs on product quality. We solve the problem assuming $U(C) + v(1 - L) = \frac{C^{1-\chi}}{1-\chi} + \log(1 - L)$, with $\chi = 2$.

The results reported in Table 3 make clear that the decentralized allocation is indeed inefficient. Relative to the Social Planner (last column), firms in the undistorted competitive equilibrium ($\tau = 0\%$) rely more heavily on replacement and choose to stop learning too quickly. These differences arise from two sources. First, there is a quality-quantity tradeoff introduced by the existence of product quality variation in the model. Firms, by construction, do not face this. For an individual firm, profits are highest when their product enjoys a high public belief. In the competitive equilibrium then, firms work to reach a high public belief as fast as possible. As a result, in the competitive equilibrium products are replaced (given the option) anytime the public belief on product quality falls below the flat prior of $\bar{\sigma} = 0.5$ (see Figure 10). This generates high average product quality and generates an extremely thick right tail in the firm distribution. The Social Planner however recognizes that chasing high average quality is a very costly way to achieve aggregate consumption C . Instead, they choose to allow for lower average quality, reduced selection, and simply provide more - there are more operating firms in the planner's economy.

The second distinction between the firm and planner is that the planner recognizes two ways to introduce new products, through new firms and the replacing of existing firms' products. Individual

Figure 10: Stationary Distributions and Replacement: Competitive Equilibrium vs. Social Planner

firms do not internalize that upon exit a new firm and new product will enter, which may be desirable for overall welfare. In Table 3 we see that even though product replacement rates are lower in the planner's economy, the average lifespan of firms and products is shorter. This reflects the increase churn of producers.

In general, these results suggest that in markets where consumer learning plays a role there is too little learning, too much product replacement at the expense of firm entry and exit and too few firms in operation (too small a market).

5.2 Taxes, Subsidies and Learning

Having characterized the constrained efficient policies, we now turn to examining the impact of size-dependent policies in the model. In particular, we are interested in how such policies affect welfare through the margins important to the planner - the thresholds of the learning space, product and firm churning, and market size. Rather than picking a particular size-dependent policy, we adopt a simple progressive tax function as a proxy for a variety of possible policies (i.e. small business subsidies or employment regulations that bind on larger firms)

$$T(\pi) = \pi^{(1-\tau)} \quad (5)$$

We assume that firms with negative profits do not owe any taxes, and tax revenue is returned to consumers in the same period as a lump sum. Our focus here is solely on how learning and associated firm dynamics are affected by progressivity, which is increasing in τ . Conceptually, the

Table 3: Progressive Taxation and Social Planner Allocations, $\gamma = 0.7$

τ	0%	5%	10%	20%	30%	Social Planner
Average Tax	0%	4.99%	3.89%	-7.53%	-13.14%	
Welfare	-3.612	-3.618	-3.729	-3.970	-3.971	-3.575
C	0.380	0.377	0.367	0.347	0.346	0.416
L	0.625	0.620	0.635	0.664	0.663	0.674
σ_{min}	0.155	0.155	0.300	0.300	0.300	0.155
σ_{max}	0.844	0.844	0.700	0.500	0.500	0.927
Mass of Operating Firms	0.515	0.512	0.535	0.578	0.578	0.571
Entry (% of Op. firms)	13.125	13.119	11.659	10.696	10.692	14.770
Replacement (% of Op. firms)	10.856	10.855	10.706	3.801	3.806	6.520
Average Quality	0.736	0.736	0.686	0.600	0.600	0.729
Average Productivity	3.220	3.220	2.710	2.156	2.156	3.255
Average Product lifespan	5.706	5.664	5.997	7.013	6.745	5.459
Average Firm lifespan	8.885	8.686	10.163	9.879	9.695	7.703

impact of progressive taxes is unclear. When learning, taxes are only paid in the event of a high signal (when sales occur) whereas taxes are always paid if pricing low. This is a force for increased learning. On the other hand, the value of learning derives from the profits of a higher public belief. As progressive taxes limit this gain, the incentive to allow learning declines.

In Table 3 we report results for increasing levels of the progressivity parameter τ ¹². As progressivity increases, it is noticeable that churning declines both in the product replacement and firm exit margins. At the same time, the learning space shrinks. By contrast, while the planner also reduces product replacement, it compensates by increasing firm exit rates and widening the thresholds of the learning space. While true that the mass of operating firms increases and approaches the efficient level, it is a mere consequence of the sizeable decrease in average product quality. As progressivity increases, firms shut down learning earlier and average firm productivity and product quality fall. All else equal, this causes marginal utility (and prices) to rise, incentivizing firms to enter and increase the size of the market. In essence, a market with progressive taxes broadly mimics a market with less informed consumers. The learning space is smaller and the market is less dynamic, both at the firm and product level. Relative to the more patient policies of the planner, our illustration suggests that size dependent policies lead firms to limit consumer learning more aggressively and are ill-suited to correct the inefficiency arising in our competitive environment.

A final dimension that we have yet to consider is the interaction between these size dependent policies and consumer information. We have seen that progressive taxation limits learning through the profit motive of firms, but we also know that the learning motive changes with consumer information. It is therefore reasonable to think that the impact of these policies will depend on the market in which they are implemented. We compare two progressivity scenarios ($\tau = 0\%$ and

¹²When the average tax level is negative, firms are on average being subsidized.

Table 4: Learning Spaces, Wages and Selection

γ	τ	σ_{min}	σ_{max}	Wage (real)	Quality	Productivity
0.70	0%	0.155	0.844	1.603	0.736	3.220
	20%	0.300	0.500	2.102	0.600	2.156
0.75	0%	0.100	0.964	1.387	0.773	3.726
	20%	0.100	0.900	1.435	0.765	3.530
0.80	0%	0.058	0.984	1.284	0.833	4.253
	20%	0.058	0.941	1.310	0.825	4.072

$\tau = 20\%$) for increasing levels of γ . The results are reported in Table 4. While in all economies progressive taxation works to reduce learning and associated selection and productivity, the clear message of the numerical experiments is that these kinds of progressive policies are particularly harmful in markets where consumers are poorly informed. In Table 4 it is noticeable how the learning space shrinks dramatically with progressive taxation when $\gamma = 0.7$, while in the economies with increased signal precision the change is relatively modest. Average product quality and firm productivity fall respectively by 18% and 26% when $\gamma = 0.7$, while the impact of progressive taxes is barely felt in the other cases. The relatively small changes for better informed markets does not mean taxes have no effect, however. Figure 11 illustrates the impact of τ on the stationary distribution of firms for the $\gamma = 0.8$ economy. First, the learning space shrinks with taxation. Everything else equal, this implies lower average product quality, since the fraction of high quality products at the top of the firms' distribution is smaller (from 98.4% to 94.1%). At the same time, however, the relative mass of firms at the top of the distribution is increasing, which raises average quality. In the stationary equilibrium, these opposing effects mostly compensate each other (the measured changes for average quality and productivity are 0.96% and 4%). Overall, the results confirm the logic behind our findings about consumer learning and information. When consumers depend heavily on public information, so do firm profits, which means that taxation has a much larger impact on incentives to learn in those markets. In markets where consumers are already well informed, taxation has a much more limited effect, if any.

6 Conclusion

Product quality is difficult to observe in many markets, requiring consumers to form beliefs about the quality of their consumption choices. The relevance of this information problem, to both consumers and firms, is clear from the explosion of consumer review websites and the importance (and value) that firms place on brand management. This paper develops an equilibrium model of industry and consumer learning to study the feedback effects between learning and industry dynamics. While stylized, the model is able to qualitatively reproduce empirical observations at the

Figure 11: Taxation and Firm Distribution

industry, firm, product and price levels of aggregation. We then show that consumer information can have important consequences for the shape and efficiency of industry dynamics. The key interaction driving this is that, since consumers learn through consumption decisions, firms can manipulate learning with their pricing policies. We find that firms optimally restrict learning in equilibrium, and do so more aggressively in markets where consumers are less informed. This has important implications for market dynamics and policy. Finally, while product level panel data across industries is quite limited, we highlight testable predictions of the model in the hopes that future improvements in data availability will allow for more precise quantification of the effects consumer learning on industry outcomes and welfare.

References

- Argente, D., M. Lee, and S. Moreira (2018, March). How do firms grow? the life cycle of products matters. *Working Paper*.
- Argente, D. and C. Yeh (2018, November). Product life cycle, learning and nominal shocks. *Working Paper*.
- Arkolakis, C., T. Papageorgiou, and O. A. Timoshenko (2018, January). Firm learning and growth. *Review of Economic Dynamics* 27, 146–168.
- Bachmann, R. and G. Moscarini (2012). Business cycles and endogenous uncertainty. *Working Paper*.
- Banerjee, A. V. (1992, August). A Simple Model of Herd Behavior. *The Quarterly Journal of Economics* 107(3), 797–817.
- Bergemann, D. and J. Valimaki (1996). Learning and strategic pricing. *Econometrica* 64(5), 1125–1149.
- Bergemann, D. and J. Valimaki (1997). Market diffusion with two-sided learning. *RAND Journal of Economics* 28(4), 773–795.
- Bergemann, D. and J. Valimaki (2000). Experimentation in markets. *Review of Economic Studies* 67, 213–234.
- Bernard, A. B., S. J. Redding, and P. K. Schott (2010). Multiple-product firms and product switching. *American Economic Review* 100(1), 70–97.
- Bikhchandani, S., D. Hirshleifer, and I. Welch (1992, October). A Theory of Fads, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy* 100(5), 992–1026.
- Broda, C. and D. E. Weinstein (2010, June). Product creation and destruction: Evidence and price implications. *American Economic Review* 100(3), 691–723.
- Cai, H., Y. Chen, and H. Fang (2009). Observational learning: Evidence from a randomized natural field experiment. *American Economic Review* 99(3), 864–882.
- Chen, Y., Q. Wang, and J. Xie (2011). Online social interactions: A natural experiment on word of mouth versus observational learning. *Journal of Marketing Research* XLVIII, 238–254.
- Copeland, A. and A. H. Shapiro (2016, July). Price setting and rapid technology adoption: The case of the pc industry. *The Review of Economics and Statistics* 98(3), 601–616.
- Dinlersoz, E. M. and M. Yorukoglu (2012). Information and industry dynamics. *American Economic Review* 102(2), 884–913.
- Dorn, D., L. F. Katz, C. Patterson, and J. V. Reenen (2017). The fall of the labor share and the rise of superstar firms. *NBER WP*.
- Dunne, T., M. J. Roberts, and L. Samuelson (1989). The growth and failure of u.s. manufacturing plants. *Quarterly Journal of Economics* 104(4), 67–98.

- Foster, L., J. Haltiwanger, and C. Syverson (2016, January). The slow growth of new plants: Learning about demand? *Economica* 83(329), 91–129.
- Garcia, D. and S. Shelegia (2018). Consumer search with observational learning. *RAND Journal of Economics* 49(1), 224–253.
- Garicano, L., C. Lelarge, and J. V. Reenen (2016). Firm size distortions and the productivity distribution: Evidence from france. *American Economic Review* 106(11), 3439–79.
- Gaur, V. and M. L. Fischer (2005). In-store experiments to determine the impact of price on sales. *Production and Operations Management* 14(4), 377–387.
- Gollin, D. (1995). Do taxes on large firms impede growth? *Bulletins* (7488). University of Minnesota, Economic Development Center.
- Gourio, F. and L. Rudanko (2014). Customer capital. *Review of Economic Studies* 81, 1102–1136.
- Guner, N., G. Ventura, and Y. Xu (2008). Macroeconomic implications of size-dependent policies. *Review of Economic Dynamics* 11(4), 721–744.
- Hendricks, K., A. Sorensen, and T. Wiseman (2012). Observational learning and demand for search goods. *American Economic Journal: Microeconomics* 4(1), 1–31.
- Jovanovic, B. (1982). Selection and the evolution of industry. *Econometrica* 50(3), 649–670.
- Keller, K. L. and D. R. Lehmann (2006). Brands and branding: Research findings and future priorities. *Marketing Science* 25(6), 740–759.
- Klette, T. J. and S. Kortum (2004). Innovating firms and aggregate innovation. *Journal of Political Economy* 112(5), 986–1018.
- Loisel, O., A. Pommeret, and F. Portier (2012). Monetary Policy and Herd Behavior: Leaning Against the Wind in Asset Prices Booms and Busts. *Working Paper*.
- Taschereau-Dumouchel, M. and E. Schaal (2018). Herd-driven Business Cycles. *Working Paper*.
- Van Reenen, J. (2018). Increasing differences between firms: market power and the macro economy. *CEPR WP*.
- Villas-Boas, J. M. (2004). Consumer learning, brand loyalty and competition. *Marketing Science* 23(1).
- Zhang, J. (2010). The sound of silence: Observational learning in the u.s. kidney market. *Marketing Science* 29(2), 315–335.

A Equilibrium Characterization

In this appendix we include all characterizations of equilibrium policies as well as a brief discussion of how we compute the equilibrium.

We start by characterizing the consumption policies consistent with the inference rule assumed.

Proposition A.1. Define $v(x, \sigma, p) = \frac{b(x, \sigma)v_h + (1 - b(x, \sigma))v_l}{p}$. Then a consumer purchasing rule

$$\delta(x, \sigma, p) = \begin{cases} 1 & \text{if } U'(C)v(x, \sigma, p) \geq 1 \\ 0 & \text{o.w.} \end{cases}$$

is the only purchasing rule consistent with equilibrium and the inference equation 3. Further $\tilde{C} = 0$.

Proof. The purchasing rule determines that consumption of some good occurs as long as the marginal utility per dollar of another unit of type (x, σ) is greater than one, the cost of marginal utility from the riskless consumption good $\tilde{C}$. To be consistent with equilibrium, this rule must be utility maximizing and budget feasible. The first is obvious under the assumption that these purchases are not budget constrained. The second is equally clear from the assumption that consumers own the firms, so that all profits and losses pass to the budget of consumers. In equilibrium, purchases are essentially self financed, which also implies that $\tilde{C} = 0$.

Finally, note that no other purchase rule is consistent with both equilibrium and inference. Any purchase rule where products are not consumed when satisfying the weak inequality above violate the assumption of our inference rule. Any purchase rule buying products that violate the weak inequality cannot be utility maximizing given the access to risk free consumption $\tilde{C}$. $\square$

The next three propositions jointly yield the firm policy characterizations in the paper. We first show that the purchasing rule above leads to a two-point pricing problem.

Proposition A.2. Conditional on belief σ , firms optimally choose between two prices.

Proof. This result is straightforward given the purchasing rule described above. Recall the firm problem

$$J(\sigma) = \max_p \sum_x \pi(x, \sigma) \left[\hat{\delta}^*(x, \sigma, p)p - w + R^{-1} \tilde{J}(\tilde{\sigma}(\sigma, p, \delta)) \right]$$

Using the purchase rule above, we know that $\hat{\delta} \in \{0, 1\}$. Let $U_{c,j,\sigma}$ denote $U'(C) [b(j, \sigma)v_h + (1 - b(j, \sigma))v_l]$. Then we can divide the set of all possible prices into three sets:

- (I) $p \leq U_{c,l,\sigma}$
- (II) $p < U_{c,l,\sigma} \leq U_{c,h,\sigma}$
- (III) $p > U_{c,h,\sigma}$

In set (I), $\hat{\delta}^*(x, \sigma, p) = 1 \forall x$ and beliefs do not evolve (the continuation value is constant). Hence $p = U_{c,l,\sigma}$ is the optimal price within that set as it maximizes current profits. In set (II), $\hat{\delta}^*(x, \sigma, p) = 1$ only for $x = h$. Beliefs do evolve, but in the same way for all cases, so $p = U_{c,h,\sigma}$ is optimal within set (II). In set (III), no purchasing occurs and beliefs do not evolve. This set is strictly dominated

by any positive price in set (I), since current profits are positive and continuation values are the same. Hence firms optimally choose between one of two prices:

$$p = \begin{cases} U_{c,l,\sigma} & \text{no experimentation} \\ U_{c,h,\sigma} & \text{experimentation} \end{cases}$$

□

Given the experimentation problem described above, we next show that this leads to a learning space bounded on $(0,1)$.

Proposition A.3. *The belief space has interior bounds $[\underline{\sigma}, \bar{\sigma}] \subset (0,1)$ and number of points N is finite.*

Proof. Denote the expected values of pricing high or low today for a firm with product of belief σ to be $J^h(\sigma)$ and $J^l(\sigma)$, respectively. If learning does not happen, the claim is trivially true. Assume the learning space is non-degenerate, hence $J^h(\bar{\sigma}) > J^l(\bar{\sigma})$.

Given the fixed operating cost of $w > 0$, for sufficiently low σ firms will exit (as the price they can charge falls to zero) hence a lower bound $\underline{\sigma}$ exists. Note that depending on the parameter space firms may stop learning earlier, but exit is sufficient to guarantee the existence of a lower bound.

Turning to the upper bound, suppose that learning never stops, i.e. firms always price high. We can bound the value of J^h above using:

$$\bar{J}(\sigma) = \pi(h, \sigma)p(b(h, \sigma)) - w + \frac{\beta(p(1) - w)}{1 - \beta}$$

where the first term are expected static profits from pricing high and the second term is the discounted value of jumping to a belief $\sigma = 1$ next period. Clearly $\bar{J}(\sigma) > J^h(\sigma)$. The value of pricing low is simply

$$J^l(\sigma) = \frac{p(b(l, \sigma)) - w}{1 - \beta}$$

We will show then that there exists a σ sufficiently high such that $J^l(\sigma) > \bar{J}(\sigma) > J^h(\sigma)$, which implies that, for contradiction, pricing low is optimal. Essentially, we are searching for a level of belief such that the forgone static profits are not compensated by additional improvements in the public belief about product quality. Using the expressions for prices, eliminating production costs and multiplicative marginal utility (which are identical) and simplifying, we require

$$[b(l, \sigma)\nu_h + (1 - b(l, \sigma))\nu_l] + \frac{\beta[b(l, \sigma)\nu_h + (1 - b(l, \sigma))\nu_l]}{1 - \beta} > [\sigma\gamma\nu_h + (1 - \sigma)(1 - \gamma)\nu_l] + \frac{\beta\nu_h}{1 - \beta}$$

$$[b(l, \sigma) - \sigma\gamma]\nu_h + [(1 - b(l, \sigma)) - (1 - \sigma)(1 - \gamma)]\nu_l > \frac{\beta}{1 - \beta}(1 - b(l, \sigma))[\nu_h + \nu_l]$$

$$\frac{b(l, \sigma)}{1 - b(l, \sigma)} \left(1 - \frac{\gamma^2}{\pi(h, b(l, \sigma))}\right) \nu_h + \left(1 - \frac{(1 - \gamma)^2}{\pi(h, b(l, \sigma))}\right) \nu_l > \frac{\beta}{1 - \beta}(\nu_h + \nu_l)$$

where $\pi(h, b(l, \sigma)) = [b(l, \sigma)\gamma + (1 - b(l, \sigma))(1 - \gamma)]$. The right hand side is a positive constant. It is easy to see the left hand side goes to infinity as $\sigma \rightarrow 1$, hence there exists a $\sigma < 1$ such that for any values of β, ν_h, ν_l the value of pricing low must be greater than the value of pricing high,

contradicting our initial supposition. The resulting set of points in the belief space, N , is finite given the common initial prior. $\square$

Finally, we turn to the shape of the pricing function, which is non-monotonic.

Proposition A.4. *The equilibrium pricing rule is non-monotonic in belief σ , and single peaked (increasing and then falling).*

Proof. This is a direct implication of the bounded learning space. Clearly prices conditional on high or low are strictly increasing in σ since

$$p^j(\sigma) = U'(C) [b(j, \sigma)v_h + (1 - b(j, \sigma))v_l]$$

Assuming the learning space is non-degenerate with N points, let $\sigma_0 = \underline{\sigma}$ and $\sigma_N = \bar{\sigma}$.

At the lower threshold, $p^l(\sigma_0) < p^h(\sigma_1) < \dots < p^h(\sigma_{N-1})$, hence prices clearly rise monotonically over this space.

At the upper threshold however, notice that $p_l(\sigma_N) < p_h(\sigma_{N-1})$. Firms with the highest quality products do not charge the highest price, but rather a slightly lower one since the value of learning no longer compensates the forgone expected profits. This yields the non-monotonic, single peaked shape of equilibrium prices. $\square$

A.1 Solving for the Equilibrium

For a given entry cost κ we construct an equilibrium in the following way.

1. Propose an equilibrium wage w .
2. Compute the firm problem $J(\sigma, C, w)$, where J is continuous and increasing in σ, C and continuous and decreasing in w .
3. Solve for $C(w)$ such that the free entry condition holds (where $C(w)$ is defined in the consumer problem).
4. The optimal pricing, replacement and exit decisions from the previous step define a transition matrix over the N points in the belief space. Vary the entry flow of new firms until $C(w)$ is obtained, yielding $e_n(w)$.
5. Compute aggregate labor demand and supply. Vary the wage guess w until the labor market clears.

B Model Dynamics with Variable Consumer Information

B.1 Full Simulations for Variable Consumer Information

To get a sense of how variation in consumer information affects the market dynamics discussed in the paper, we reproduce Facts 1-10 in Figure B.1 for economies with differing degrees of signal precision γ .

B.2 Comparison of Stationary Equilibria for Variable Consumer Information

γ	0.7	0.75	0.8
Growth Rate (Expected 4-periods growth)	36.01%	40.32%	56.22%
Exit Rate (% exit within 2 periods)	24.45%	24.84%	30.96%
Age (Average periods)	8.39	8.18	7.14
Concentration (HH Index)	219	231	221
Churning (% of products alive in 4-periods)	59.86%	60.03%	53.38%
Price Change Frequency (probability within 2 periods)	50.46%	62.33%	55.80%
Price Change Size (average size)	1.97%	1.52%	1.09%
Average Product Lifespan	5.70	5.47	5.62
Average Firm Lifespan	8.88	8.55	7.43

The first statistic we compute is the expected four period growth rate of firm value before entering the market. The measure is strictly increasing in γ . The intuition is natural: increasing γ implies that high quality products have a greater chance of growth (up to the upper threshold of the learning space). While the opposite is true for low quality products, it is worth pointing out that replacement and firms' exit mitigate the impact of negative returns on the expected growth rate. In other words, higher γ is associated to more frequent positive shocks among surviving products and firms. The three economies share a very similar level of $J(\bar{\sigma})$, since free entry requires $J(\bar{\sigma}) = w\kappa$, and the equilibrium wage w varies little as γ changes. The absolute size of their potential change in value, however, varies substantially with γ (the largest potential increase occurs in the $\gamma_{0.75}$ economy). The expected growth rate depends both on the frequency of good and bad shocks, and on the average change in firm value associated with them. Overall, in the computed examples, the frequency component dominates the size component, resulting in the observed strictly increasing pattern.

Exit rates are also increasing in γ , since as signal precision increases it is more likely that low-

quality products receive bad signals. After a sequence of bad signals, firms reach the bottom of the belief distribution where it is optimal to exit the market. In other words, higher γ improves selection.

The latter effect is also reflected in the average firm age, which is decreasing in γ . It is however not obvious that signal precision should decrease average firm age, since high quality products are also more likely to reach the upper threshold of the learning space, where they will be alive until either the firm or the product is exogenously destroyed. One result we have shown in the paper is that increasing γ implies widening the learning space. With a larger learning space, a longer sequence of positive shocks may be required to reach the upper threshold where learning is shut-down. In other words, with a lower γ , there is a larger mass of bad-quality products that reach the top of the belief distribution, raising average age. Such firms are less likely to have a long lifespan with higher signal precision.

Several statistics (average product lifespan, churning, price change frequency, concentration) show a non-monotonic pattern. Some intuition can be grasped from looking at the firms' stationary distributions. The fraction of products in the highest belief point is respectively 46%, 35% and 41% in the $\gamma_{0.7}$, $\gamma_{0.75}$ and $\gamma_{0.8}$ economy. At the same time, the total mass of operating firms decreases from 0.5144, to 0.5103 and 0.4917. It is also worth noting that, as we go from $\gamma_{0.7}$ to $\gamma_{0.75}$, the number of belief points in the learning space increases from 5 to 6, while it is constant going from $\gamma_{0.75}$ to $\gamma_{0.8}$. There seem to be two patterns developing at the same time. On one hand, increasing γ coincides with greater selection and firm exit. On the other, a more precise signal gives firms stronger incentives to learn their product quality, which causes the learning space to widen. At $\gamma_{0.75}$ the dominant force seems to be the latter, while at $\gamma_{0.8}$ the former prevails.

These considerations allow us to understand the joint pattern of churning, frequency of price changes, and average product lifespan. When the selection force prevails, churning rises, as more products receiving bad signals are eliminated through exit and replacement. The surviving products have higher expected quality, and are more likely to reach the top of the belief's distribution, where their expected lifespan is the highest, and where prices are constant.

The pattern in market concentration, measured by the HH index, is also consistent with these dynamics. As we have shown, firms at the highest belief of the distribution sell with certainty, but also charge a lower price than competitors at lower belief points. Conditional on selling, these firms have a relatively lower individual market share, and they tend to have a heavier weight in the stationary distribution when the selection force dominates.

C Data Appendix

C.1 Indicator of Consumer Information

On the Consumer Reports (CR) website, we downloaded the most recent “5-Year Index”, which is available freely. This lists all product reviews published in the magazine for a given year. We chose 2016 as a representative year. In the index, entries are listed alphabetically, typically by product type (i.e. “bicycle helmets”). To start, we created a list of reviews by product category (using the description from the CR 5-Year Index). In general this was straightforward but some issues arose. For example, some reviews appeared double listed (entries “Casual restaurants, Nov 16, 38” and then “Restaurants, casual, Nov 16, 38” both identify a review of casual restaurants in the November 2016 (i.e. 16) issue on page 38.) An annual subscription to CR was purchased,

allowing us to crosscheck these double entries and ensure appropriate counting. A second, more difficult issue, was categorizing reviews. The CR monthly magazine is not just a set of product reviews. It includes industry guides, industry advocacy and gift lists as well. While these all convey some amount of information, we restricted focus only to what could plausibly be considered product reviews. While some service products were included (for example, Pet Insurance, May 16, page 12) they were not easily priced and as such were excluded from the final analysis. Service reviews accounted for a relatively small share of overall reviews so we do not believe selection is a serious problem. Finally, we also found that multiple reviews of a product category could be listed in one magazine issue. Most notably, each automobile review entered as a separate listing in the index as opposed to the majority of other reviews which mentioned several products. To limit the impact of this, we collapsed all reviews by month, so that any product category with at least one mention in a month received just one count. Other measures would likely be better, such as total words expended on a given product category, but were infeasible.

Having constructed a list of all reviews including product category, date of publication and average price (obtained through online versions of the original article) we next map each product review into its North American Product Classification System (NAPCS) code. These contain, seven, six and three digit codes which we manually mapped to products. Statistics Canada provides a fairly comprehensive description of the NAPCS along with ample product examples and in most cases the mappings were easily made. In the few cases where they were more difficult, the mapping was left at the highest level (3-digit) of aggregation and in some cases excluded entirely.

Finally, we collapsed products into 3-digit NAPCS codes and counted the number of reviews. A full list of these reviews is presented in Table C.1.

In Figure C.2 we average prices within three-digit NAPCS categories and plot log prices against review counts. A clear positive correlation emerges, even if we eliminate the highest-count product category. This is consistent with the idea that consumers will rely more on public information for products they do not regularly purchase, which are likely to be more expensive and more durable. Of course, it may also be the case that consumers will demand more information for products on which they spend more, holding constant their private information, simply because of the opportunity cost of buying low quality goods. This likely plays some role, however this logic would imply we should see the most reviews for the largest consumer expenditure categories. Data from the Bureau of Labor Statistics shows that for 2017, in spite of Food (at and away from home) representing 12.97% of average annual household expenditures we see few reviews in those product categories as compared to Vehicle Purchases (6.75%) or Entertainment (5.33%). See <https://www.bls.gov/news.release/cesan.nr0.html>.

Figure B.1: Product and Price Dynamics in Economies with Variable Consumer Information

(a) Firm Distribution is Skewed

(b) Product Dropping by Age

(c) Product Dropping by Shipment Size

(d) Frequency of Adjustment by Age

(e) Size of Adjustment by Age

Figure C.1: Consumer Reports Product Reviews, 2016
3-Digit NAPCS, Full List

Category	3-digit NAPCS	No. Reviews	Log Price
Other Information Technology	751	1	-
Fertilizers, Pesticides and Other Chemical Products	272	1	1.95
Prepared Meals and Beverages	841	1	3.91
Postal and Courier Services	514	1	4.03
Asphalt (except natural) and Asphalt Products	262	1	3.00
Meat Products	172	1	2.30
Other General Purpose Machinery and Equipment, n.e.c.	345	1	3.40
Carbonated and Non-Carbonated Drinks ...	193	1	0.70
Fresh Fruits, Nuts and Vegetables, and Pulse Crops	114	1	1.61
Plastic and Foam Building and Construction Materials	464	1	5.62
Other Crop Products	115	1	2.48
Carpets, Textiles Furnishings and Other Textile Products	232	2	3.912
Other Insurance Services	763	2	-
Telecommunications, Broadcasting Distribution and ...	723	2	3.91
Coffee and Tea	191	2	2.71
Electronic and Electrical Parts	371	2	4.74
Motor Vehicle Engines and Motor Vehicle Parts	413	2	6.13
Heating, Cooling and Air Purification Equipment	344	2	5.62
Furniture and Fixtures	391	2	6.32
Computers and Computer Peripherals	361	2	6.48
Software and Software Licensing	482	2	3.56
Plastic and Rubber Finished Products	284	3	3.22
Medical, Measuring and Other...	363	3	6.07
Clothing, Footware and Accessories	231	3	4.42
Health Services	811	3	4.38
Cleaning Products and Toiletries	274	3	2.74
Retail Services (Except Commissions)	561	4	3.26
Other Food Products	183	4	3.59
Agricultural, Lawn and Garden Machinery and Equipment	331	4	5.99
Pharmaceutical and Medicinal Products	273	4	2.71
Miscellaneous Goods and Supplies	475	5	5.58
Communication, and Audio and Video Equipment	362	5	6.54
Appliances	382	7	6.69
Financial Services (Except Insurance)	761	8	-
Passenger Cars and Light Trucks	411	13	10.50

Figure C.2: Prices and Consumer Information