

Product Life Cycle, Learning, and Nominal Shocks

DAVID ARGENTE

Pennsylvania State University

and

CHEN YEH

Federal Reserve Bank of Richmond

First version received November 2018; Editorial decision October 2021; Accepted February 2022 (Eds.)

This article documents a new set of stylized facts on how pricing moments depend on product age and emphasizes how this heterogeneity is crucial for the amplification of nominal shocks to the real economy. Exploiting information from a unique panel containing billions of transactions in the US consumer goods sector, we show that our empirical findings are consistent with a narrative in which firms face demand uncertainty and learn through prices. Such a mechanism of active learning from prices can strongly influence an economy's aggregate price level and can thus be important for assessing the degree of monetary non-neutrality. To quantify this, we build a general equilibrium menu cost model with active learning and exogenous entry that features heterogeneity in pricing moments over the life cycle of products. Under this setup, firms engage in active learning to deal with uncertainty on their demand curves. Firms choose prices not only to maximize static profits but also to create signals to obtain valuable information on their demand. In the calibrated version of our model, the cumulative real effects of a nominal shock are approximately three times as large compared to a standard price-setting model. The main intuition behind this result is that active learning weakens the selection effect. Price changes are mainly determined by forces of active learning and, hence, become more orthogonal to aggregate shocks, which reduces the aggregate price flexibility of the economy.

Key words: Menu cost, Firm learning, Fixed costs, Nominal shocks.

JEL Codes: D4, E3, E5

1. INTRODUCTION

How large is the transmission of nominal shocks to the real economy? This remains an important question in macroeconomics because of the disconnect that exists between micro price-setting models, which typically predict small real effects (e.g. [Golosov and Lucas, 2007](#)), and macro estimates where the variation in output to nominal shocks is sizable at short horizons (e.g. [Shapiro and Watson, 1988](#)).¹ This article seeks to reconcile these views by focusing on the

1. [Shapiro and Watson \(1988\)](#) refer to these shocks as “demand” shocks, but we follow the convention in [Lucas \(2003\)](#), as done by [Nakamura and Steinsson \(2009\)](#), that these shocks capture more than just monetary shocks. In fact,

The editor in charge of this paper was Dirk Krueger.

age-dependence of several pricing moments. We argue that this heterogeneity is crucial for understanding the amplification of nominal shocks to the real economy.

To come to this conclusion, we document new stylized facts on the age-dependence of pricing moments and interpret them through the lens of a structural price-setting model with active learning. Using a panel containing billions of transactions in the US consumer good sector, we show that the empirical evidence supports a mechanism of active learning with prices. We argue that this narrative rationalizes our set of stylized facts comprehensively. Then, we build a price-setting model with active learning to quantify the impact of a nominal shock on the real economy. We find that the cumulative real effects of a nominal shock are approximately three times as large compared to a standard menu cost model. The main intuition behind this result is that active learning weakens the selection effect: firms have less of a need to adjust their prices in response to nominal shocks because they are more prone to adjust their prices due to idiosyncratic reasons (i.e. active learning) instead. This raises an economy's non-neutrality.²

We start by using a large panel of US products to document new facts about the distribution of both the duration and size of price changes over the product's life cycle; a dimension so far ignored by traditional price-setting models. We show that pricing moments, such as frequency and absolute size, strongly depend on the age of a product: entering products change their prices twice as often as the average product and the average size of these adjustments is 50% larger compared to the average price change.³ Furthermore, large price changes occur disproportionately more often in the early stages of a product's life cycle.

Previous studies from a select number of industries show that the introduction of a new product is associated with a significant amount of demand uncertainty from the firm's perspective. Empirical evidence from these few case studies indicate that firms vary their prices to obtain information about their demand curves.⁴ We argue that our stylized facts provide evidence from a broad line of product categories to support such a hypothesis. Although we cannot rule out the possibility that a combination of alternative mechanisms can account for our stylized facts, the active learning mechanism can *simultaneously* rationalize them and do so in a parsimonious way. In addition, the active learning narrative generates a unique prediction on the age-dependence of pricing moments. Under the assumption that more novel products bring more demand uncertainty with them, the negative relation between a product's pricing moments and its age should be more pronounced for more novel products. To confirm this relation, we construct a *newness* index that measures the novelty of a product. This novelty reflects observed characteristics such as, for example, brand, volume or packaging (as opposed to product age). We indeed find that more novel entering products adjust their prices more often and by larger amounts. Also, we exploit variation across space and the timing of product launches within retailers to find that retailers carry information obtained during the first launch forward to any subsequent launch of the same

he argues that these shocks also capture temporary non-neutrality due to real shocks. While [Smets and Wouters \(2007\)](#) show that monetary shocks on their own account for a relatively modest fraction of business cycle variations in output, the broader interpretation of nominal shocks can explain "a significant fraction of the short-run forecast variance in output." Other examples of this sort include [Justiniano, Primiceri and Tambalotti \(2010, 2013\)](#).

2. The term "selection" was introduced by [Golosov and Lucas \(2007\)](#) to indicate that firms that change prices after a nominal shock are those whose prices are in greatest need of adjustment. Given that the distribution of the size of price changes fully encodes this type of selection, a wide range of papers in recent years have taken advantage of micro data to match the size distribution of price changes, such as [Nakamura and Steinsson \(2008\)](#), [Midrigan \(2011\)](#), [Vavra \(2014\)](#), and [Karadi and Reiff \(2019\)](#).

3. The patterns at exit are quite different, as the frequency and absolute size of regular price adjustments stay mostly constant before exit. Nonetheless, the frequency and depth of sales increase significantly at exit. These findings are described in more detail in [Supplementary Appendix A1](#).

4. [Gaur and Fisher \(2005\)](#), for example, survey 32 US retailers and finds that 90% conduct price experiments to learn about their demand.

product. As a result, a framework where firms engage in active learning is a good starting point for rationalizing standard and new moments of the price change distribution.

Our empirical findings are the first in the price-setting literature that focus on heterogeneity in product age. While previous contributions have found heterogeneity in pricing moments along other dimensions and demonstrated its importance for nominal shock propagation (e.g. [Nakamura and Steinsson, 2009](#)), none of them have focused on a product's life cycle. Importantly, this is a source of heterogeneity *within* sectors. Our facts emphasize that this type of heterogeneity is a salient feature of the data and provide insights on the underlying reason *why* firms adjust prices.⁵ We argue that the age distribution of products at a given point in time is crucial for assessing the propagation of a nominal shock to the real economy. To do so, we build a general equilibrium menu cost model in the spirit of [Golosov and Lucas \(2007\)](#) to rationalize our stylized facts and to quantify the magnitude of this propagation. Based on our empirical findings, we let firms engage in active learning pricing strategies to deal with demand uncertainty. To model active learning, we follow the insights by [Mirman, Samuelson and Urbano \(1993\)](#) and [Bachmann and Moscarini \(2012\)](#).⁶ In particular, the latter contribution is useful for our approach in modelling active learning because it highlights how to do this in a concise way.

In our framework, a firm is uncertain about its demand elasticity upon its product's entry. Firm-specific demand shocks prevent a firm from directly inferring its type: a firm that sets a relatively high price and observes a low amount of sales cannot distinguish between the fact that its product has a high elasticity of substitution or that the realization of its demand shock was simply low.

To deal with this uncertainty, firms engage in a Bayesian learning process in which their beliefs about their elasticity can be updated after observing the amount of quantity they have sold.⁷ These sold quantities are the firm's only source of noisy signals about its elasticity. The key insight of active learning is that prices not only affect a firm's level of current profits, but they also impact the future set of signals a firm receives. Changes in its price alter the speed at which the firm learns about its elasticity of demand. Thus, the firm balances its incentives between maximizing static profits and altering its future information set in the most efficient way. Active learning thus features a trade-off between "current control" and "estimation." As a firm ages and learns more about its elasticity, its incentives for active learning decline and the dispersion of price changes decreases. This is consistent with evidence from other markets such as the newly deregulated market of frequency response services in the UK, a service required to keep electricity running smoothly. [Doraszelski, Lewis and Pakes \(2018\)](#) find that in response to uncertainty, firms experiment with their bids by adjusting them more frequently and by larger amounts and, over time, the market experienced an important reduction in the range of bids; adjustments of bids became less frequent and smaller.

In our framework, the age-dependence of pricing moments is completely generated by learning from the supply side. Consumers have standard preferences over a composite of two baskets

5. The theoretical literature on price-setting shows that different types of price changes have substantially different macroeconomic implications. For example, the way the timing of price changes is modelled, is crucial for its macroeconomic implications; be it endogenous (e.g. [Caplin and Spulber, 1987](#)) or exogenous (e.g. [Calvo, 1983](#)).

6. The focus of the latter paper is very different as they study how negative first moment aggregate shocks induce risky behaviour. In their model, when firms observe a string of poor sales, they become pessimistic about their own market power and contemplate exit. At that point, the returns to price experimentation increase as firms "gamble for resurrection."

7. Our work is related to the literature in optimal control problems with active learning that has been studied in many areas of economics since [Prescott \(1972\)](#). Its application to the theory of imperfect competition consists of relaxing the assumption that the monopolist knows the demand curve it faces. The first application of this concept can be found in [Rothschild \(1974\)](#). More recent examples can be found in [Wieland \(2000b\)](#), [Willems \(2017\)](#), and [Ilut, Valchev and Vincent \(2020\)](#).

of goods facing no frictions, hence we abstract from demand-side narratives (e.g. learning by customers or customer base). While we acknowledge the limitations of the demand side of our framework, we argue that our model of active learning is a parsimonious way of capturing salient facts on the age-dependence of pricing moments. Furthermore, we provide unique evidence from the newness of products and staggered product introduction across space and time that supports our mechanism of active learning by firms.⁸

We calibrate our framework to standard pricing moments and our newly found set of stylized facts regarding the product's life cycle and quantify the cumulative real output effect of a nominal spending shock. Our findings indicate that the output response is 2.8–3.3 times larger than in the benchmark model with no demand uncertainty and accounts for 7–11% of the US business cycle, which is about half of the measured variation predicted by [Shapiro and Watson \(1988\)](#).

The reasoning behind this result is 2-fold. The incentives for active learning dampen selection in the *size* of price changes; that is, pricing with active learning motives pushes firms away from the margin of price adjustment and lessens the mass of firms that adjust their prices due to a nominal shock. Since firms have an additional motive to change prices, they become less sensitive to changes in their marginal costs, which is the primary source of price gaps in most menu cost models.

In addition, the concept of a product's life cycle introduces an additional form of cross-sectional heterogeneity in the frequency of price adjustment. In an environment with active learning incentives combined with menu costs, uncertain firms are willing to adjust their prices more often to acquire information on their demand. These firms will most likely adjust their price several times before firms with sharper beliefs adjust their price once after a nominal shock. However, all those price changes by more uncertain firms, after the one in response to the nominal shock, have no effect on real output because these firms have already adjusted to the shock. Given that the model is calibrated to match the average frequency of price changes, firms that are more certain about their type significantly delay the adjustment of the aggregate price level after a nominal shock. These firms tend to be older and have a lower frequency of price adjustment on average. As a result, this delay reduces selection in the *timing* of price changes.⁹ Importantly, the lower selection effect in the size and timing of price changes arises endogenously in our model.¹⁰

The remainder of this article is organized as follows. In [Section 2](#), we present the data and the main empirical findings. [Section 3](#) discusses a simple two-period setup to highlight the intuition behind a firm's pricing strategy engaged in active learning. Moreover, we present a quantitative, general equilibrium menu cost model that is able to explain our stylized facts. [Section 4](#) calibrates our quantitative framework and evaluates it along a set of targeted and untargeted pricing moments. In [Section 5](#), we discuss our results on the propagation of nominal shocks and compare our results with other models used in the literature. [Section 6](#) provides additional empirical evidence

8. Importantly, we perform extensive robustness tests in [Section 7](#) showing that a customer base model with deep habits ([Ravn, Schmitt-Grohé and Uribe, 2006](#)) is not consistent with our stylized facts. It is worthwhile to note, however, that this model abstracts from strategic customer responses through the extensive margin as featured in recent macroeconomic models of the customer base (e.g. [Kleshchelski and Vincent, 2009](#); [Drozd and Nosal, 2012](#); [Gourio and Rudanko, 2014](#); [Paciello, Pozzi and Trachter, 2019](#); [Roldan and Gilbukh, 2021](#)).

9. Recent contributions highlighting the importance of selection in the timing of price changes are [Kiley \(2002\)](#), [Nakamura and Steinsson \(2009\)](#), [Sheedy \(2010\)](#), [Alvarez, Lippi and Paciello \(2011\)](#), and [Carvalho and Schwartzman \(2015\)](#).

10. Although this logic is similar to the one described in [Alvarez, Le Bihan and Lippi \(2016a\)](#) in relation to the response of real output to nominal shocks, our framework does not fit the class of models for which the kurtosis of the distribution of price changes is a sufficient statistic. This is because price gaps in our setup are endogenous since prices and beliefs are jointly determined. Then, conditional on adjustment, firms do not fully close their price gaps. We discuss this matter more extensively in [Section 5](#).

in favour of the active learning mechanism. Section 7 discusses alternative mechanisms that can rationalize our empirical findings and extends our quantitative framework along several dimensions to ensure that our results on monetary non-neutrality are robust. Section 8 concludes. The Appendix provides additional empirical findings and details on extensions of the model.

2. STYLIZED FACTS ON THE LIFE CYCLE OF US PRODUCTS

In this section, we use a large scanner data set to show a new set of stylized facts on pricing moments over a product's life cycle in the US economy. We begin by showing the importance of new products; both in terms of their count and revenues relative to the aggregate. Then, we develop a set of facts that show that pricing moments at the product level are considerably different across a product's age; in particular near entry. At entry, the frequency of regular price changes, the absolute size of regular price adjustments, and the cross-sectional standard deviation of regular price changes are higher. All of these moments approximately settle to their respective averages as the product matures. Furthermore, the fraction of large price changes, defined as those changes larger than two standard deviations after prices are demeaned in a given category and store, is considerably larger at the beginning of the product's life cycle.

2.1. Data

The life cycle patterns of products' prices have typically not been studied much as the requirements on the data are quite stringent. Doing so requires a large panel of products with information about their entry date and prices at a high sampling frequency. The Consumer Price Index (CPI) Research Data, for example, is only available at a monthly frequency and the age of products is unknown. All Entry-Level Items (ELIs) are added to the CPI basket long after their first appearance in the national market.

For this reason, we rely on the IRI Marketing data set which provides more than 10 years of data at the store-product-week level. The data are generated by point-of-sale systems: each retailer reports the total dollar value of its weekly sales and total units sold for each product. A product is identified by its Universal Product Code (UPC), a code consisting of 12 numerical digits that is uniquely assigned to each specific product and represents the finest level of disaggregation at the product level.

The data contain approximately 2.4 billion transactions from January 2001 to December 2011 which represents roughly 15% of household spending in the Consumer Expenditure Survey (CEX). Our sample contains approximately 170,000 products and 3,000 distinct stores across 43 metropolitan areas (MSA). The data cover 31 product categories and includes detailed information about each product such as its brand, volume, colour, flavour, and size.¹¹

Given the properties of the data, we can identify the first appearance of a UPC in a certain store by using the retail and product identifiers. We assume that if a UPC changes, some noticeable characteristic of the product has also changed. This is because it is rare that a meaningful quality change occurs without a change to its UPC. Considering each UPC as a product is, in fact, a broad

11. The product categories include Beer, Carbonated Beverages, Coffee, Cold Cereal, Deodorant, Diapers, Facial Tissue, Photography Supplies, Frankfurters, Frozen Dinners, Frozen Pizza, Household Cleaners, Cigarettes, Mustard & Ketchup, Mayonnaise, Laundry Detergent, Margarine & Butter, Milk, Paper Towels, Peanut Butter, Razors, Blades, Salty Snacks, Shampoo, Soup, Spaghetti Sauce, Sugar Substitutes, Toilet Tissue, Toothbrushes, Toothpaste, and Yogurt. The data set is discussed in more detail in Bronnenberg, Kruger and Mela (2008). See also Coibion, Gorodnichenko and Hong (2015), Alvarez *et al.* (2016a), Chevalier and Kashyap (2019), Stroebel and Vavra (2019), and Gagnon and López-Salido (2020) for applications of the data to related questions.

definition since it includes classically innovative products, which are “breakthrough” products that deliver innovation to an existing or new product category; line extension products, which are new products within an already existing category; and temporary products, which have a short life cycle and are typically seasonal. We find that product line extensions, such as flavour/form upgrades or novelty and seasonal items, are much more prevalent than the introduction of new brands.

Using UPC and retailer identifiers, we are able to determine at what week and store each product first appears. We define entering products as those that enter the US market after January 2002. Our data start from January 2001, thus an entering product is one that has no observable transactions in any store across the US for at least 1 year. This assumption avoids the inclusion of products with a left-censored age. In addition, we only consider products that entered the market before the first week of 2007. We impose this restriction for two reasons. First, the prices of products born during downturns can have substantially different patterns than those of products born in normal times.¹² More importantly, IRI Marketing undertook a substantial reorganization of its product categories and expanded their scope at the beginning of 2007. Thus, the data after this specific date might include some entering products that might not correspond to actual product introductions.¹³ By restricting our sample of entering products between January 2002 and January 2007, we avoid this reclassification bias.

Further, in order to minimize concerns of potential measurement error in the calculation of product-level entry and exit, our baseline sample excludes private-label products and only considers products that last at least 2 years in the market. We exclude private-label items from the main analysis because all private-label UPCs have the same brand identification so that the identity of the retailer cannot be recovered from the labelling information.¹⁴ Also, we exclude short-lived products in order to minimize the problem that some newly entering products get assigned UPCs of previously existing products as noted by [Chevalier, Kashyap and Rossi \(2003\)](#). Last, we drop promotional items or products with very little revenue to minimize biases due to measurement error.¹⁵

2.2. *The importance of new products*

[Broda and Weinstein \(2010\)](#) emphasize the importance of entering and exiting products for the aggregate performance of the US economy through aggregate price indices. [Argente, Lee and Moreira \(2018\)](#) show that product turnover in the US is substantial as one-third of all products are either created or destroyed in a given year and more than 20% of US products are aged less than 1 year. In this subsection, we sketch an identical picture in our sample to highlight the importance of entering products. We begin by using information on the number of

12. [Moreira \(2016\)](#), for example, provides evidence that the average business size across cohorts is significantly affected by aggregate economic conditions at inception.

13. More specifically, IRI undertook the following actions: (i) reorganization of private-label items (i.e. organic private labels are broken out for some categories), (ii) dropping of UPCs that have not moved in past years, (iii) collapse of UPCs into a main UPC to avoid clutter (i.e. products that came to a store as part of a special promotional code rather than with a standard UPC code), (iv) reorganization of categories (i.e. a category might have increased in scope and as a consequence experienced an increase in items), and (v) addition of UPCs that were introduced at the beginning of each stub. All of these are consistent with changes in the number of entering and exiting UPCs due to changes in the product stub rather than new product introductions or products being phased out.

14. A “private label” product is one that a retailer gets from a third party (who also produces it) but sells under its own brand name.

15. These restrictions are only a small part of our sample (approximately 2.5%). In the end, our results are robust to not excluding private-label items, short-lived products, and promotional items.

TABLE 1
Product entry and exit

	UPC 5-year	UPC 1-year	UPC×Store 5-year	UPC×Store 1-year
Entry	0.45	0.14	0.66	0.27
Entry (W)	0.29	0.07	0.47	0.15
Exit	0.42	0.13	0.61	0.25
Exit (W)	0.08	0.01	0.39	0.10

Notes: The table shows the statistics of the entry and exit rates for one and five year intervals. Entry and exit rates denoted with “W” are weighted by revenues. Columns (1) and (2) show the statistics at the UPC level whereas columns (3) and (4) show these at the UPC-store level.

new products, exiting products, and the total number of products in each category $k \in \mathcal{C}$ to define aggregate entry and exit rates at the product level:

$$n(t, s) = \frac{\sum_{k \in \mathcal{C}} N_k(t, s)}{\sum_{k \in \mathcal{C}} T_k(t)}$$

$$x(t, s) = \frac{\sum_{k \in \mathcal{C}} X_k(t, s)}{\sum_{k \in \mathcal{C}} T_k(s)}$$

where $N_k(t, s)$, $X_k(t, s)$, and $T_k(t)$ are the number of entering products, exiting products, and total products in period t relative to period s for category k . We define the entry rate in period t relative to s as the number of new products in period t relative to period s as a share of the total number of products with strictly positive sales in period t . A new product is one that records at least one transaction in any store in period t and that was not sold in any store in period s . We also report entry and exit rates that are weighted by revenue.

Using a scanner data set collected at the store level offers the advantage of observing, for the categories available, the entire universe of products for which a transaction is recorded in a given week. For this reason, we can distinguish between products entering the market and products being launched at each store, where our unit of observation is every UPC-store pair. We find a substantial degree of entry of products at both levels.

Table 1 reports the entry and exit rates for the case in which t and s are one and five years apart. It shows that 14% of the UPCs in the market and on average 27% of the products in each store entered in the last year. Approximately 45% of the products in the market entered in the last 5 years accounting for 29% of total expenditures.

At the store level, 66% of all products sold were first introduced by the store in the last 5 years and they account for about half of the total revenue of the store. Although the exit rate is very similar to the entry rate of products, it is lower than entry whenever weighted by revenue. This lower rate means that consumers spend more on new products than on products that are about to exit. The rate of product turnover indicates that at any point in time, there is a large amount of products being launched or being phased out.

Table 2 shows that the median duration of a product in a given store is slightly above 3 years.¹⁶ The large rates of revenue-weighted entry and revenue-weighted exit, both at the store and at the market level, along with a short product life cycle indicate that the pricing of entering products are relevant for determining the dynamics of aggregate prices.¹⁷

16. Since our data set ends the last week of 2011 and we are considering products that entered the last week of 2006 at the latest, right censoring is only an issue for products that last more than 261 weeks in the market.

17. We observe large rates of entry and exit (revenue-weighted) at the brand level as well. On average, 16% of brands are either created or destroyed in a given year which is consistent with Broda and Weinstein (2010) who find

TABLE 2
Distribution of duration by UPC \times Store

	(1)	(2)	(3)	(4)
	Unweighted		Revenue Weighted	
	Weeks since entry	Observations since entry	Weeks since entry	Observations since entry
1st percentile	1.0	1.0	12.4	7.7
25th percentile	37.4	16.3	108.4	73.5
50th percentile	96.3	47.1	183.7	131.8
75th percentile	209.5	122.0	280.6	208.3
99th percentile	450.7	369.7	466.3	405.1
Mean	134.0	83.1	198.9	148.9
Std. Dev.	118.9	90.6	117.9	100.0

Notes: The table shows the statistics of the distribution of durations of a UPC-store pair. In columns (1) and (2), we compute the duration of each UPC-store pair and aggregate them to the category level using equal weights. Categories are further aggregated using equal weights. In columns (3) and (4), we aggregate to the category level using revenue weights and aggregate across categories using equal weights. Weeks since entry refers to the number of weeks elapsed since the product was first observed. Observations since entry refers to the number of times a product is observed in our data set. A product is observed only if it records a transaction in a given week and store.

2.3. Empirical strategy

To study the price dynamics of products along their life cycle, we begin by computing the average retail price in a given week:

$$P_{mcjst} = \frac{\text{sales}_{mcjst}}{\text{units}_{mcjst}},$$

where m, c, j, s , and t index markets (at the MSA level), product categories, UPCs, stores, and time, respectively. A considerable advantage of the IRI Marketing data set is that it provides information on whether and when a product was on sale in a certain store (the so-called “sales flag”) that is absent in other scanner data sets. Since our goal is to study the speed of price adjustment following a nominal shock, we focus on studying the life cycle patterns of regular price changes given that retailers’ use of sales/promotions in our data do not vary with macroeconomic conditions.¹⁸ Nonetheless, the main stylized facts discussed below are robust to the inclusion of sales in the analysis.

We adopt the same conventions as Coibion *et al.* (2015) to distinguish between regular price changes and sales. A *regular* price change is defined as any change in price that is larger than one cent or 1% in absolute value. For prices larger than 5 dollars in value, this cut-off is 0.5%. To identify sales, we use the sales flag provided in the data, but our results are robust to applying the sales filter introduced by Nakamura and Steinsson (2008).¹⁹ The size of a price change is calculated as the log difference between the price levels in the current and the previous week.

that most creation does not come from new sizes or new flavours. Table A3 shows similar entry and exit patterns at the category level.

18. Using the same data set, Coibion *et al.* (2015) find that retailers’ use of sales does not vary with the unemployment rate. Anderson, Malin, Nakamura, Simester and Steinsson (2017) argue that sales prices are governed by sticky plans and they are planned in advance according to a “trade promotion calendar.” They also find that retailers do not respond to macroeconomics shocks by adjusting the size or frequency of sales. Figure A5 in Appendix A shows that neither the frequency or the size of sales/promotions are larger at entry.

19. Under this approach, a good is on sale if a price is reduced but returns to its same previous level within 4 weeks. Coibion *et al.* (2015) use two approaches to identify a price spell. The first treats missing values as interrupting price spells. In the second approach, missing values do not interrupt price spells if the price is the same before and after the periods with missing values. Since the incidence of sales from applying these two approaches does not significantly differ from the one identified by the sales flag provided in the IRI Marketing data set, we use the union of sales flags obtained

Thus, we get:

$$\Delta P_{mcjst} = \ln(P_{mcjst}) - \ln(P_{mcjst-1}).$$

Let $a \in \{1, \dots, A\}$ denote the number of weeks since entry (which we will define as the *age* of the product) where $a = 1$ denotes entry. To assess the movements of the pricing moments over the life cycle of a product, we adopt the following age-cohort-period model:

$$Y_{jstc} = \alpha + \sum_{a=1}^A \phi_a \cdot D_{js}^a + \theta_{js} + \tau_t + \gamma_c + \varepsilon_{jstc}, \quad (1)$$

where j , s , t , and c are the UPC, store, time period and cohort $c = t - a$ the product belongs to respectively. Y_{jstc} is the variable of interest (e.g. the price change indicator or the size of the price change). D_{js}^a is a dummy variable that takes the value of one if the product is in its a th week since entry. θ_{js} captures fixed effects at the UPC-store level whereas τ_t and γ_c denote time and cohort fixed effects, respectively. We are interested in the regression coefficients $\{\phi_a\}_{a=1}^A$ which capture age heterogeneity of the pricing moment of interest.

In our empirical specification, it is not possible to identify the heterogeneous effects of age conditional on a product's cohort and time period due to perfect collinearity. To resolve this, we follow Heckman and Robb (1985) who argue that "age, cohort, and time effects are proxy variables for underlying unobserved variables which are not themselves linearly dependent." We estimate equation 1 under the assumption that trends appear only in cohort effects. Time fixed effects are replaced with the seasonally adjusted unemployment rate at the MSA level to control for local cyclical economic variation.²⁰

2.4. Pricing moments over the life cycle of US products

We first use regression specification 1 with a price change indicator as the dependent variable. Figure 1 plots the frequency of regular price changes over the life cycle of a product.

The dots represent the estimates of the age fixed effects $\{\phi_a\}_{a=1}^A$ associated with specification 1.²¹ Newer products clearly see their prices being changed more often. This allows us to state our first stylized fact.

EMPIRICAL FACT 1. The average frequency of price adjustment declines with the product's age. The decline is most pronounced at the early stage of the product's life cycle as entering products change their prices twice as often as the average product.

from applying these two approaches and the flag provided in the IRI Marketing data to identify the incidence of sales. Our results are not sensitive to any of these choices.

20. Some examples of studies that use this proxy approach are Deaton and Paxson (1994), Gourinchas and Parker (2002), De Nardi *et al.* (2010), and Aguiar and Hurst (2013). We have also applied another normalization in which trends only appear in the period effects instead. In this case, the time fixed effects are included in the estimation of equation 1 and we use the local unemployment rate to represent cohort fixed effects. The idea behind this is that products introduced during (local) downturns might perform differently in terms of their pricing dynamics over the life cycle. This is reminiscent of Moreira (2016) who shows that the size and performance of businesses over their life cycle is heavily influenced by the state of the economy at times of business inception. Our baseline results are not sensitive to the chosen normalization.

21. Specifically, we plot $\hat{\alpha} + \hat{\phi}_a$ for every $a \in \{1, \dots, A = 50\}$. $\hat{\alpha}$ is the unconditional average of the frequency of price changes of a product that has been in the market for 50 weeks.

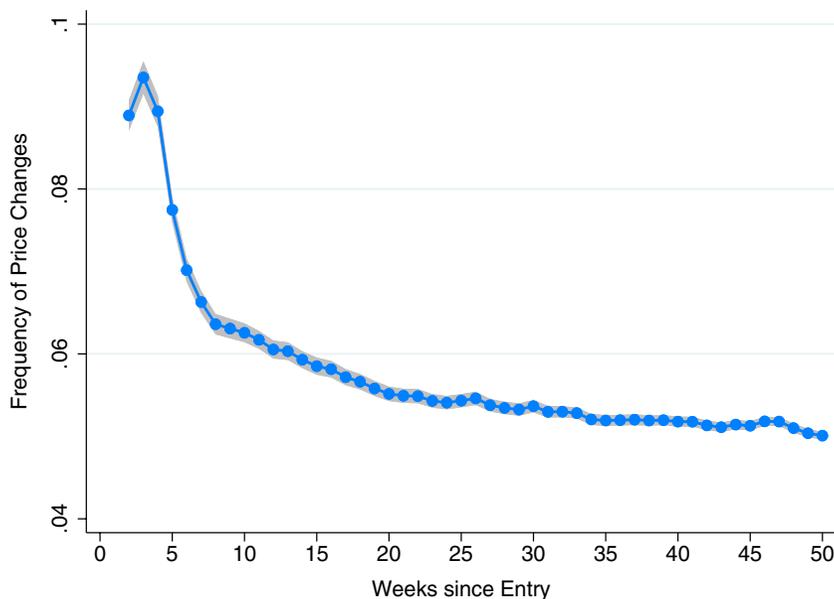


FIGURE 1

Frequency of price adjustment since entry

Notes: The graph plots the average weekly frequency of price adjustments of entering products. The y-axis denotes the probability that a product adjusts its price in a given week whereas the x-axis denotes the number of weeks the product has been observed in the data after entry. The graph plots the coefficients for the age fixed effects in equation 1 where we use the regular price change indicator as the dependent variable. Regression specification 1 is computed by controlling for UPC-store and time fixed effects while the local unemployment rate proxies for cohort fixed effects. The calculation uses approximately 130 million observations and 2.5 million UPC-store pairs. Standard errors are clustered at the store level.

The frequency of price adjustment is almost 4 percentage points higher at entry and takes approximately 20 weeks to settle to its average value of 5%. The magnitude of this significantly higher frequency is best reflected in the expected amount of time it takes for a product to change its price.²² If we maintain the frequency of adjustment at entry, then a price should change approximately every 12 weeks. This is *twice* as often relative to the average of 24 weeks that we observe in the data.

In order to study whether the magnitude of these price adjustments also changes over the life cycle of the product we use a similar approach but instead use the absolute size of price changes as the dependent variable. Figure 2 depicts our results.

During the first few months, the absolute value of price changes is much larger and almost 5 percentage points higher than the average which amounts to approximately 9%. Further, the dispersion of price changes as measured by the weekly cross-sectional standard deviation is almost 40% larger during the first 4 months after entry with respect to its level 12 months after the product is launched.²³ Importantly, this fact holds for both price increases and decreases.²⁴ This leads to our second stylized fact.

22. This is equal to $-1/\ln(1-f)$, where f denotes the frequency of price adjustment.

23. Our findings are consistent with those in Alvarez, Borovicková and Shimer (2016c) who find that the hazard rates of price changes depend on the age of the product once unobserved heterogeneity is taken into account. Figure 7 decomposes these price changes into increases and decreases. Qualitatively, we find similar results for both the frequency of price increases and decreases.

24. See Figure 7 in Section 4. Figure A4 in Appendix A shows that both the frequency and size of regular price changes stay mostly constant at exit. Figure A6 shows that there is an increase in the frequency of sales during the last weeks of the life cycle of the product.

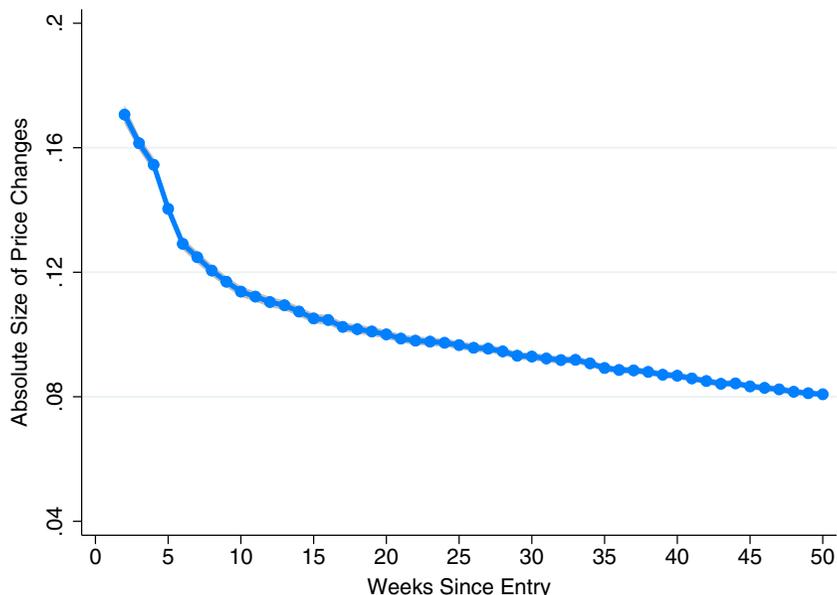


FIGURE 2

Absolute value of price changes since entry

The graph plots the average absolute size of price adjustments of entering products. The y-axis is the absolute value of the log price change in that week whereas the x-axis denotes the number of weeks since the product entered. The graph plots the coefficients for the age fixed effects in equation 1 where we use the absolute value of the log price change as the dependent variable. Regression specification 1 is computed controlling for UPC-store, time fixed effects while the local unemployment rate proxies for cohort fixed effects. The calculation uses approximately 5.8 million price changes and 2.5 million UPC-store pairs. Standard errors are clustered at the store level.

EMPIRICAL FACT 2. The absolute size of price adjustment declines monotonically with the product's age. The decline is most pronounced at the early stage of the product's life cycle as the average absolute size of entering products is almost twice as large as the average change.

Our baseline specification uses a non-parametric specification for the age of a product to allow for non-linearities of pricing moments in a product's age. However, our results stay robust when allowing for a richer set of time fixed effects (e.g. UPC-time effects or store-time effects).²⁵

Thus, we conclude that firms not only price more often but also in a more extreme fashion during the early stages of their products' life cycles. Importantly, the standard class of menu cost models is not capable of capturing the salient age-dependent features we observe in the data.²⁶

Next, we investigate whether very large price changes are more or less frequent as products get older. To do so, we follow the approach by Alvarez *et al.* (2016a) to minimize complications regarding heterogeneity across products and stores. We define "cells" at the UPC-store level, say (j, s) , and standardize each price change at this level through $z_{jst} \equiv (\Delta p_{jst} - \mu_{js}) / \sigma_{js}$, where μ_{js} and σ_{js} are the mean and standard deviation of price changes in cell (j, s) across time.

Figure 3 shows the distribution of regular price changes larger than two standard deviations as a function of the age of the product. We observe a sizeable share of large price changes close

25. These results are displayed in Figure A1 in Appendix A. Our conclusions remain unchanged whenever we impose linearity in the age of a product. These results are summarized in Table A1 in Appendix A. Figures A2 and A3 show that our results are very similar using alternative scanner data sets such as the Nielsen RMS.

26. To see this, consider the benchmark model of Golosov and Lucas (2007) with exogenous entry and exit. In this setting, the frequency of price adjustment rises with age if the entry price is chosen optimally. Furthermore, this benchmark makes no predictions on the absolute size of price changes over a product's life cycle.

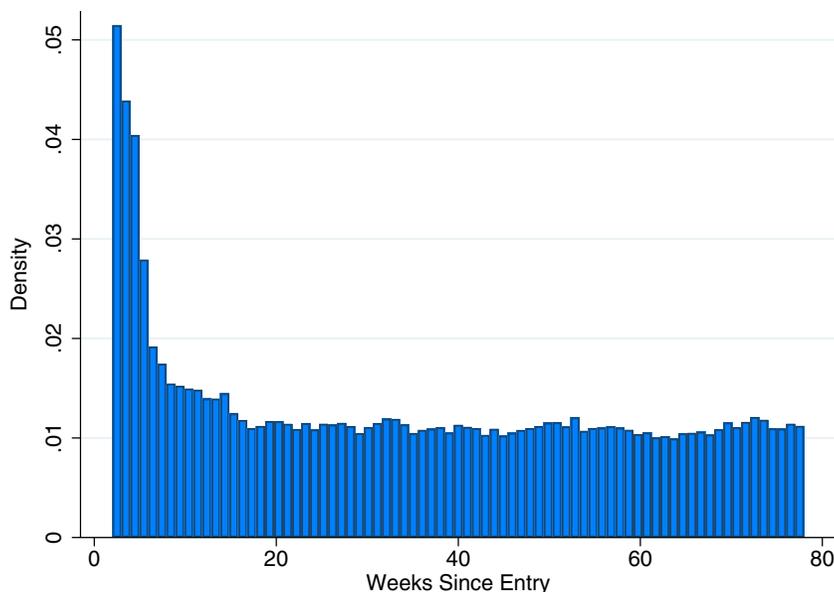


FIGURE 3

Fraction of price changes larger than two standard deviations

Notes: The figure shows the fraction of price changes larger than two standard deviations from the mean in a given category and store as a function of the age of the product. The products considered are those that last at least 2 years in the market.

to entry particularly during the first 20 weeks. About 40% of the price changes larger than two standard deviations of the product life cycle occur during these weeks. The distribution of large price changes is roughly uniform after that.

These results do not rely on the standardization of prices nor on our definition of “large” price changes.²⁷ Lastly, this pattern holds for both positive and negative large price changes; just like our previous set of stylized facts.²⁸ As a result, we summarize our third stylized fact as follows.

EMPIRICAL FACT 3. Large price changes mostly occur in the early stages of a product’s life cycle. 40% of large price changes occur during the first 20 weeks of a product’s life cycle whereas a quarter of them are already observed in the first 3 weeks of a product’s entry.

This finding implies that age-dependent pricing moments are not limited to “normally sized” price changes. This is important as our third empirical fact relates to the use of idiosyncratic shocks from fat-tailed distributions in menu cost models to generate large price changes. Previous work has assumed that these shocks arrive at a constant rate. This assumption includes the family of Poisson shocks used in [Midrigan \(2011\)](#). However, the prevalence of large price changes in the beginning of the life cycle contradicts this assumption. Fat-tailed shocks have substantial implications on the degree of non-neutrality as well. Their presence reduces the selection effect after a monetary shock as the mass of firms responding to it is smaller. This will prove to be important for understanding what economic mechanism could rationalize our set of stylized facts.

27. Figure [A8](#) in Appendix [A](#) shows our robustness exercise for non-standardized price changes larger than 30%.

28. This is displayed in Figure [A7](#) in Appendix [A](#).

3. A GENERAL EQUILIBRIUM MODEL OF ACTIVE LEARNING

In the following, we present a model of active learning that allows us to structurally interpret our empirical findings of the previous section. Our framework is a discrete-time menu cost model in the tradition of [Golosov and Lucas \(2007\)](#). The crucial assumption differentiating it from the former model is that firms face uncertainty on their demand curves. A firm does not know its elasticity of demand (or type), but forms beliefs and learns about its type over time. We adopt an approach in which firms *actively* learn about their type: firms can adjust their prices to change the speed at which its beliefs get updated. As a result, a firm faces an intertemporal trade-off between maximizing its static profits and gaining more information about its type when choosing its price. Hence, information at the firm level is endogenously generated. The learning mechanics of our framework are similar to those in [Bachmann and Moscarini \(2012\)](#).

We start with a simple two-period setup of active learning to highlight the aforementioned trade-off and how this impacts a firm's pricing strategy. Then, we present a quantitative, general equilibrium model that can rationalize our empirical facts and allows us to explore the aggregate implications of active learning in a menu cost model.

3.1. A two-period model of active learning

We start with a two-period setup to highlight the economic intuition behind pricing under active learning motives. Firms are characterized by a type $\sigma \in \{\sigma_1, \sigma_2\}$ that determines their profitability. In particular, the inverse demand curve is given by $q = D(p; \sigma) + \varepsilon$, where ε is a zero mean random variable drawn from a density $f(\cdot)$ that is continuous and satisfies the monotone likelihood ratio property (MLRP). A firm does not know its type σ but has some initial belief λ_0 over it. A firm cannot deduce after the first period what its type is due to the random variable ε , but it can update its belief through Bayes' rule.

We assume that firms have market power and can determine their own prices. After a firm sets its price p , it can observe how much quantity it sold q and update its beliefs. Thus, this means that a firm's informative signal (i.e. its quantity sold) directly depends upon its control variable (i.e. its price). As a result, the firm's posterior belief is an explicit function of its price p . These posterior beliefs follow Bayes' rule and are denoted by $\lambda' = B(\lambda, p, q)$. Conditional on some price p and type σ , profits are given by $\Pi(p; \sigma)$. We assume that $\Pi(\cdot; \sigma)$ is concave for all types σ which is a fairly weak assumption. In the end, we can write a firm's dynamic problem as:

$$v(\lambda_0) = \max_{p \in \mathcal{P}} M(p; \lambda_0) + \beta \cdot \mathbb{E}_\varepsilon \left[\lambda_0 \cdot V(B(\lambda_0, p, D(p; \sigma_1) + \varepsilon)) \right. \\ \left. + (1 - \lambda_0) \cdot V(B(\lambda_0, p, D(p; \sigma_2) + \varepsilon)) \right],$$

where

$$M(p; \lambda_0) = \lambda_0 \Pi(p; \sigma_1) + (1 - \lambda_0) \Pi(p; \sigma_2)$$

$$V(\lambda') = \max_{p' \in \mathcal{P}} M(p'; \lambda').$$

The firm's value $V(\lambda')$ in the last period is straightforward: a firm maximizes its expected, static or myopic profits $M(p; \lambda')$ given its posterior belief λ' in the second period. Let the maximizer of myopic profits be defined as the *myopic* price $p^M(\lambda) \equiv \max_{p' \in \mathcal{P}} \lambda \Pi(p'; \sigma_1) + (1 - \lambda) \Pi(p'; \sigma_2)$. The myopic price is unique for each belief λ since $\Pi(\cdot; \sigma)$ was assumed to be concave.

In the first period however, a firm internalizes the fact that the price it sets also affects its posterior belief. Thus, a firm must balance its incentives between obtaining higher myopic profits

and sharpening its posterior beliefs to increase its continuation value. The optimal price that strikes a balance between these two forces, given a prior belief λ_0 , is denoted by $p^*(\lambda_0)$. This price is the maximizer associated with the value function $v(\lambda_0)$.

As a result, we say that a firm actively learns with its price at the belief λ_0 if it deviates from the myopic price. This deviation $|p^*(\lambda_0) - p^M(\lambda_0)|$ reflects the firm's incentive to gain information for increasing the speed of learning at the expense of its current period profits. However, it is not clear ex ante that any price will lead to more information for the firm. In the following, we describe what prices are more informative from the firm's point of view and formally establish what conditions are sufficient for active learning.

INCENTIVES FOR ACTIVE LEARNING. Even though a firm's type σ is uncertain, it is aware that it can only take two values. As a result, there are only two possible *expected* demand curves: $D(p; \sigma_1)$ and $D(p; \sigma_2)$. We define the *confounding* price \hat{p} as that price at which the demand curves intersect in expectation, i.e., $D(\hat{p}; \sigma_1) = D(\hat{p}; \sigma_2)$. To learn about its own type, a firm's strategy is to induce variation in its quantity at a price that is as far as possible from the confounding price. This is because at the confounding price the firm does not obtain any relevant information in expectation. In fact, the posterior belief is equal to its prior at the confounding price, i.e., we get that $B(\lambda_0, \hat{p}, D(\hat{p}; \sigma) + \varepsilon) = \lambda_0$ for any ε, σ .

Therefore, an uncertain firm that wants to learn about its demand curve, can do so by "separating" the two possible demand curves from each other. We formalize this intuition by relying on the insights of [Aghion, Bolton, Harris and Jullien \(1991\)](#) and [Mirman et al. \(1993\)](#). The latter provides a set of sufficient conditions in order for a monopolistic firm to experiment. First, information must be valuable which is reflected by the strict convexity of $V(\cdot)$. A relatively large literature has established that a firm's motive for active learning can be captured by the convexity of its continuation value. In other words, the more convex is the function $V(\cdot)$ the more valuable information becomes. In our specific setup, it is straightforward to show that this is the case.²⁹ Second, prices must be able to affect the informativeness of quantities. In a similar setup to our two-period model, [Mirman et al. \(1993\)](#) show that the latter condition is satisfied whenever the slopes of the demand curves are not equal to each other at the myopic price, i.e., we have $\frac{dD(p, \sigma_1)}{dp} \Big|_{p=p^M(\lambda_0)} \neq \frac{dD(p, \sigma_2)}{dp} \Big|_{p=p^M(\lambda_0)}$. This condition can be satisfied by assuming, for example, iso-elastic demand curves.

Recall however that a firm's pricing decision is not purely determined by active learning motives. A firm also cares about its myopic profits in the first period. Then, the optimal pricing policy strikes a balance between a concave myopic profit function and a convex continuation value. This is also known as the "current control-estimation" trade-off in the experimentation literature. This emphasizes the importance of the convexity of $V(\cdot)$: a more convex continuation value stresses the importance of the value of information and steers the firm away from a concave profit function which reflects static, profit maximization incentives that are standard in most menu cost models. In order to shed some light on the active learning motives, below we provide a numerical example to illustrate how a firm's pricing incentives are influenced by the value of information.

NUMERICAL EXAMPLE. In this example, we parameterize the profit function through loglinear (or CES) demand curves and a constant marginal cost of production. Furthermore, we draw ε from a normal distribution with mean zero and variance σ_ε^2 . For simplicity, we set the prior belief at $\lambda_0 = \hat{\lambda}$. For convenience, we denote $\mathcal{V}(\cdot; \lambda_0) \equiv \mathbb{E}_\varepsilon [\lambda_0 \cdot V(B(\lambda_0, \cdot, D(\cdot; \sigma_1) + \varepsilon)) + (1 - \lambda_0) \cdot$

29. We prove this formally in [Supplementary Appendix BII.1](#).

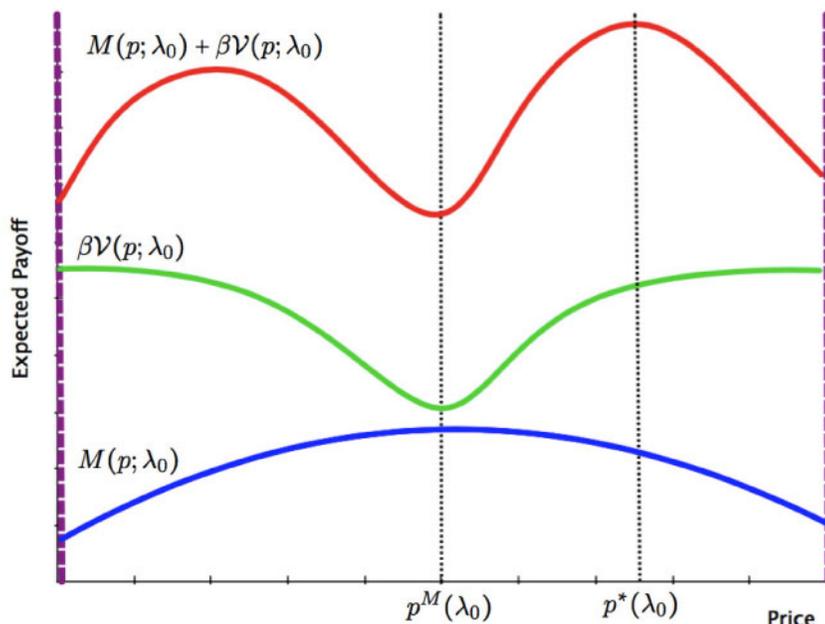


FIGURE 4

Example of the two-period model

Notes: The figure denotes static profits $M(p; \lambda_0)$, continuation value $\beta V(p; \lambda_0)$ and total payoff $M(p; \lambda_0) + \beta V(p; \lambda_0)$ at $\lambda_0 = \hat{\lambda}$. The dotted purple lines represent the optimal prices p_2^* and p_1^* under complete information. $p^M(\lambda_0)$ represents the myopic policy whereas $p^*(\lambda_0)$ represents the pricing policy under active learning.

$V(B(\lambda_0, \cdot, D(\cdot; \sigma_2) + \varepsilon))$. Figure 4 plots the firm's myopic profits \mathcal{M} , the continuation value \mathcal{V} , and the total payoff (which is simply the sum of the former and discounted latter) as a function of the firm's set price. The dotted lines at the extremes of the figure depict the optimal static prices with perfect information which satisfy $p_i^* = \max_{p \geq 0} \Pi(p; \sigma_i)$ for $i \in \{1, 2\}$. By construction, myopic

profits are maximized at $p^M(\lambda_0)$. The concavity of $M(p; \lambda_0)$ illustrates the costs of active learning as prices far away from $p^M(\lambda_0)$ represent profit losses in the first period.

Figure 4 shows that $\mathcal{V}(\cdot; \lambda_0)$ is convex which follows from the convexity of $V(\cdot)$. It also shows that its minimum lies at the confounding price \hat{p} . The reason is that a firm's sales become completely uninformative at the confounding price. In this case, small deviations from the confounding price lead to large gains. Thus, the benefits from active learning are strongly related to the convexity of $\mathcal{V}(\cdot; \lambda_0)$. For example, prior beliefs closer to zero and one lead to less convex continuation values. The reason is because the marginal benefit of information decreases for firms that are more certain about their type.

A firm bases its pricing strategy by maximizing its *total* payoff. In this example, the total payoff is double-peaked and its global maximum is at $p^*(\lambda_0)$.³⁰ The figure shows that the global maximum lies in the interior of $\mathcal{P} = [p_2^*, p_1^*]$ and, most importantly, the optimal pricing strategy deviates from its myopic counterpart.³¹

30. In general, the sum of concave and convex functions can have multiple peaks, however the results of our baseline framework always have either a single or a double-peaked continuation value.

31. In [Supplementary Appendix BII.4](#), we derive a set of sufficient conditions to guarantee that $p^*(\lambda_0) \in [p_2^*, p_1^*]$ for all λ_0 . [Supplementary Appendix BII.5](#) shows the different learning regimes, as consistent with the findings by [Keller and Rady \(1999\)](#), that could arise in our setup and describes the case with menu costs.

In the following, we bring these pricing mechanics of active learning to a quantitative, general equilibrium menu cost model. In Section 4, we show that such a model is well suited in matching our empirical results and as such provides a good justification for interpreting our facts through the lens of a pricing framework with active learning.

3.2. Quantitative model

3.2.1. Households. Households in the economy maximize their expected, discounted utility over aggregate consumption C_t and labour supply L_t that is characterized by:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\theta} - 1}{1-\theta} - \omega \frac{L_t^{1+\chi}}{1+\chi} \right],$$

where \mathbb{E}_t denotes the expectations operator conditional on information available to the household at time t . Households have CRRA preferences over an aggregate consumption good with risk aversion parameter θ and the level of disutility from working is denoted by ω . The inverse Frisch elasticity is given by χ and households discount by a factor β per period.

The aggregate consumption good C_t is a Cobb–Douglas composite of two Dixit–Stiglitz indices of differentiated goods:

$$C_t = C_{1t}^\eta C_{2t}^{1-\eta} \text{ with } C_{it} = \left[\int_{k \in J_i} \left(\alpha_t^i(k) \right)^{\frac{1}{\sigma_i}} \left(C_t^i(k) \right)^{\frac{\sigma_i-1}{\sigma_i}} dk \right]^{\frac{\sigma_i}{\sigma_i-1}}$$

There are two continua of differentiated goods consisting of perishable consumption units or services. A good is indexed by a pair (i, k) . The first index $i \in \{1, 2\}$ denotes the good's *basket* whereas its *variety* within a basket i is denoted by $k \in J_i$. Varieties within the first basket are hard to substitute with each other whereas other varieties are mutually substitutable with a relatively high elasticity of substitution, thus we set $\sigma_2 > \sigma_1$.

Each good is assumed to be produced by a single monopolistically competitive producer.³² Each good is identified through the index pair (i, k) . We assume that i is time-invariant. Consumer's good-specific preference shocks $\alpha_t^i(k)$ are drawn every period, independently over time, and within and across group types. Draws are the same for all consumers. Households' decisions are taken *after* observing these taste shocks.

Within each period, households choose how much to consume of each differentiated good to maximize the level of the aggregate consumption good C_t . For a given level of spending S_t , we obtain the following downward-sloping demand curve for each differentiated good (i, k) :

$$C_t^i(k) = \alpha_t^i(k) \left(\frac{P_t^i(k)}{P_{it}} \right)^{-\sigma_i} \frac{\eta_i S_t}{P_{it}}, \quad (2)$$

where we denote the income shares for each basket by $\eta_1 = \eta$ and $\eta_2 = 1 - \eta$. $p_t^i(k)$ denotes the price of a good (i, k) in period t . By construction, the price index for composite good C_{it} is given by:

$$P_{it} = \left(\int_{k \in J_i} \alpha_t^i(k) \left(P_t^i(k) \right)^{1-\sigma_i} dk \right)^{1/1-\sigma_i}. \quad (3)$$

32. We will use the term “firm,” “good,” and “product” interchangeably since there are no multi-product firms in our setup.

These price indices satisfy the following equalities $P_{1t}C_{1t} + P_{2t}C_{2t} = P_t C_t = S_t$. The aggregate price level P_t is such that $P_t C_t$ is the minimum amount of expenditure necessary to obtain C_t units of the aggregate consumption good. We assumed that the realization of taste shocks are independent across groups. Since there is a continuum of goods within each basket i , we can use a law of large numbers.³³ This implies:

$$P_{it} = \left(\int_{k \in J_i} P_t^i(k)^{1-\sigma_i} dk \right)^{1/1-\sigma_i},$$

where the law of large numbers gives us that $\int \alpha_t^i(k) dk \xrightarrow{P} \mathbb{E}(\alpha_t^i(k)) = 1$ since we normalize the expected value of the taste shocks to be equal to unity.

Households have access to a complete set of Arrow–Debreu securities.³⁴ Therefore, the period t budget constraint is characterized by:

$$P_t C_t + \mathbb{E}_t(q_{t,t+1} B_{t+1}) \leq B_t + W_t L_t + \sum_i \int_{k \in J_i} \Pi_t^i(k) dk,$$

where W_t denotes the nominal wage rate and $\Pi_t^i(k)$ is the profit that households receive from owning the firm producing good (i, k) . B_{t+1} denotes the state-contingent payoffs in period $t+1$ from purchasing assets in period t . These claims are priced in period t by the unique (stochastic) discount factor $q_{t,t+1}$. The first-order conditions of a household's intertemporal maximization problem are then:

$$\beta^{T-t} \mathbb{E}_t \left(\frac{C_T^{-\theta}}{P_T} \frac{1}{q_{t,T}} \right) = \frac{C_t^{-\theta}}{P_t} \quad (4)$$

$$\omega L_t^\chi C_t^\theta = \frac{W_t}{P_t}. \quad (5)$$

These equations describe the determination of asset prices and labour supply.

3.2.2. Firms. In contrast to most state-dependent pricing models, firms in our framework set prices under incomplete information. In the following, we incorporate the active learning mechanics of Section 3.1 into a richer, quantitative environment.³⁵ At any point in time t , a firm (i, k) can observe its total amount of sales $Q_t^i(k)$ after setting some price $P_t^i(k)$. Due to the Dixit–Stiglitz specification above, these sales satisfy the following loglinear form:

$$\log(Q_t^i(k)) = -\sigma_i \log(P_t^i(k)) + \log(S_t) + (\sigma_i - 1) \log(P_{it}) + \log(\eta_i) + \log(\alpha_t^i(k)),$$

33. In particular, we use the *Glivenko–Cantelli* theorem. The argument is identical to the one in [Bachmann and Moscarini \(2012\)](#).

34. Our main results in Sections 4 and 5 are based on a stationary environment, hence there is no need for households to trade these Arrow–Debreu securities in this case. However, we present our model in a general setting since we discuss the implementation of aggregate shocks in Appendix C.3 and [Supplementary Appendix CIII](#).

35. Note that the intuition behind active learning in the two-period model of Section 3.1 is identical to that in the quantitative framework.

where we define $\log(\alpha_i^i(k)) = \varepsilon_i^i(k)$. If the time index t is temporarily dropped and a monopolistically competitive firm producing good (i, k) is considered, then with some abuse of notation, we can rewrite the above sales equation as:

$$q = -\sigma_i p + s + \mu_i + \varepsilon_k. \quad (6)$$

From the firm's point of view, there is no longer a one-to-one mapping between quantities and prices as it cannot observe taste shocks ε_k and it does not know its elasticity of substitution. Whenever a firm sets a price p , demand q can be high for two reasons: (1) the variety belongs to a basket within which substitution is hard or (2) the realization of the taste shock ε_k is high.

A firm does observe the amount of sold quantities q of its product. It can use this information to learn about its elasticity of demand and update its type. As a result, a firm might want to deviate from the static profit maximizing price to learn more about the price elasticity of its corresponding basket.

Taste shocks are specific to each variety, but these are unobserved by the firm. Furthermore, a firm is unaware of its type i but uses observed sales as an informative signal to learn about its type. As a result, its pricing policy is independent of (i, k) , and we can drop this index without loss of generality for determining the optimal pricing strategy. Our setup imposes the following timing on the firm's pricing decisions and the consumer's realized demand shock for each period.

1. A firm decides on its price p before the realization of the demand shock ε_k .
2. The shock ε_k is realized and households decide to consume $q = -\sigma_i p + s + \mu_i + \varepsilon_k$ conditional on the set price p .
3. The firm is contractually obliged to supply $Q = \exp(q)$ and forms its posterior beliefs.

Let λ denote the firm's prior belief that its type is of a low elasticity of demand σ_1 and $f(\cdot)$ the probability distribution function of ε_k . Whenever a firm observes log sales q and aggregate income s , and sets some price p , it can update its prior belief to the posterior λ' according to Bayes' rule:

$$\begin{aligned} \lambda' &= B(\lambda, p, q, s) \\ &= \left[1 + \frac{1-\lambda}{\lambda} \frac{f(q + \sigma_2 p - \mu_2 - s)}{f(q + \sigma_1 p - \mu_1 - s)} \right]^{-1}. \end{aligned}$$

The firm rationally anticipates that the price it sets will affect the informative quantity it will observe the period after: prices set in period t affect a firm's future beliefs. As a result, the firm's motives are not solely rooted in the maximization of its static profits because a firm's pricing strategy can increase the value of its sales' informativeness. In the remainder of this article, we assume that log demand shocks are normally distributed with mean zero and variance σ_ε^2 . Then, a firm's posterior belief, conditional on its true type being i , is equal to:

$$\begin{aligned} b_i(\lambda, p, \varepsilon) &= B(\lambda, p, -\sigma_i p + s + \mu_i + \varepsilon, s) \\ &= \left(1 + \frac{1-\lambda}{\lambda} \exp \left(\frac{1}{2} (-1)^{1(i=2)} \left[\left(\frac{\varepsilon}{\sigma_\varepsilon} \right)^2 - \left(p \cdot \frac{\Delta\sigma}{\sigma_\varepsilon} - \frac{\Delta\mu}{\sigma_\varepsilon} + \frac{\varepsilon}{\sigma_\varepsilon} \right)^2 \right] \right) \right)^{-1}, \end{aligned}$$

where $\Delta\sigma \equiv \sigma_2 - \sigma_1 > 0$ and $\Delta\mu \equiv \mu_1 - \mu_2$. The expression above shows that the speed of learning or the rate at which posterior beliefs change is heavily influenced by the firm's pricing decision. In fact, the expression indicates that posterior beliefs are more responsive to prices whenever the signal-to-noise ratio $\Delta\sigma/\sigma_\varepsilon$ is high. Further, it can be seen that beliefs

are self-reinforcing at the confounding price $\widehat{p} = \Delta\mu/\Delta\sigma$ as we have $b_i(\lambda, \widehat{p}, \varepsilon) = \lambda$ for all ε and i .³⁶

3.2.3. Pricing. A firm has access to a linear production technology in labour. Its production function is given by:

$$y_t^i(k) = z_t^i(k) \ell_t^i(k),$$

where $y_t^i(k)$ denotes the output of some firm (i, k) in period t . Similarly, $\ell_t^i(k)$ is the quantity of labour the firm uses for production purposes in period t . Its idiosyncratic productivity is given by $z_t^i(k)$. Labour is supplied competitively at the nominal rate W_t . In addition, we assume that log productivity follows a mean-reverting process:

$$\log(z_{t+1}^i(k)) = \rho \cdot \log(z_t^i(k)) + \sigma_\zeta \zeta_{t+1}^i(k), \text{ where } \zeta_{t+1}^i(k) \sim N(0, 1).$$

A firm's ex-interim expected profits conditional on type i , price p , and idiosyncratic productivity z is equal to:

$$\begin{aligned} \Pi_t^i(p; z) &= \mathbb{E}_t \left[\left(p - \frac{W_t}{z} \right) \alpha_t^i(k) \left(\frac{p}{P_{it}} \right)^{-\sigma_i} \frac{\eta_i S_t}{P_{it}} \right] \\ &= \left(p - \frac{W_t}{z} \right) \left(\frac{p}{P_{it}} \right)^{-\sigma_i} \frac{\eta_i S_t}{P_{it}}, \end{aligned}$$

where the last equality follows from our previous normalization. Myopic profits are then formed by taking expectations with respect to a firm's current prior belief λ_t . Furthermore, firms are required to pay a fixed cost of ψ in units of labour to adjust their nominal price. Static profits are then equal to:

$$\lambda_t \Pi_t^1(p; z) + (1 - \lambda_t) \Pi_t^2(p; z) - \psi W_t \cdot \mathbf{1}(p \neq p_{t-1}^i(k)),$$

where $\mathbf{1}(\cdot)$ is an indicator function equal to unity whenever its argument holds true. Given these constraints, a firm chooses a path of prices $\{p_t^i(k)\}_{t \geq 0}$ to maximize its expected, discounted profits:

$$\mathbb{E} \left(\sum_{t=0}^{\infty} q_{t,t+1} \Pi_t(P_t^i(k); z_t^i(k)) \middle| \mathcal{F}_0 \right),$$

where the expectation is with respect to the path of future beliefs, demand, and productivity shocks. Its information set available at time 0 is denoted by the filtration \mathcal{F}_0 . Any firm makes its pricing decisions while taking aggregate prices, spending, and the wage rate as given. These variables are determined in general equilibrium. In the following, we focus on a stationary recursive equilibrium in which nominal aggregate spending trends at a constant rate and satisfies $\log(S_{t+1}) = \log(S_t) + \pi$. Thus, there is no aggregate uncertainty and aggregate variables are constant under this equilibrium. Profits are then discounted at the rate β .

Entering firms start out with a common prior λ_0 and some initial productivity draw. Only upon entering, firms can adjust their price without incurring the menu cost. In our model, firms have substantial incentives to learn their type at the beginning of their life cycle and do so by making

36. This is formally shown in [Supplementary Appendix BII.2](#).

large and frequent price adjustments. As the product matures the gains to obtaining additional information are extremely small and are not sufficient to offset the menu cost. Given that the frequency and absolute size of price changes at the later stages of the product life cycle are far from negligible, we capture the incentives for price changes at these later stages through standard state-contingent channels: idiosyncratic cost shocks and allowing for positive inflation levels.

Thus, a firm's dynamic programming problem is summarized by the following Bellman equation:

$$v(\lambda, z, p_{-1}) = \max \left\{ v^A(\lambda, z), v^{NA}(\lambda, z, p_{-1}) \right\},$$

where the value of adjusting and not adjusting are respectively given by:

$$\begin{aligned} v^A(\lambda, z) &= \max_{p \geq 0} \lambda \Pi^1(p; z) + (1 - \lambda) \lambda \Pi^2(p; z) - \psi \frac{W}{P} \\ &\quad + \beta \lambda \mathbb{E}_{\varepsilon, z'} \left[v \left(b_1(\lambda, \log(\frac{p}{1+\pi}), \varepsilon), z', \frac{p}{1+\pi} \right) \middle| z \right] \\ &\quad + \beta (1 - \lambda) \mathbb{E}_{\varepsilon, z'} \left[v \left(b_2(\lambda, \log(\frac{p}{1+\pi}), \varepsilon), z', \frac{p}{1+\pi} \right) \middle| z \right] \\ v^{NA}(\lambda, z, p_{-1}) &= \lambda \Pi^1(p_{-1}; z) + (1 - \lambda) \lambda \Pi^2(p_{-1}; z) \\ &\quad + \beta \lambda \mathbb{E}_{\varepsilon, z'} \left[v \left(b_1(\lambda, \log(\frac{p_{-1}}{1+\pi}), \varepsilon), z', \frac{p_{-1}}{1+\pi} \right) \middle| z \right] \\ &\quad + \beta (1 - \lambda) \mathbb{E}_{\varepsilon, z'} \left[v \left(b_2(\lambda, \log(\frac{p_{-1}}{1+\pi}), \varepsilon), z', \frac{p_{-1}}{1+\pi} \right) \middle| z \right]. \end{aligned}$$

We define the optimal pricing policy $p^*(\lambda, z)$ as the maximizer associated with the value function $v^A(\lambda, z)$. In a menu cost model without active learning, a price-setting firm only considers its static profits and its effect on the continuation value through the fact that changing prices is costly.³⁷ However, sold quantities are observable and informative from the firm's point of view. Thus, a firm can also affect its posterior beliefs through its price. This is highlighted by the posterior belief functions b_1 and b_2 in the firm's continuation value. As a result, the policy function $p^*(\lambda, z)$ reflects the optimal deviation from the myopic policy function that summarizes the balance between sacrificing static profits and increasing the rate at which it learns about its type.³⁸

3.2.4. Stationary recursive equilibrium. We close the model by clearing the labour market and imposing an exogenous structure for entry and exit. In particular, we assume that every firm starts out with the prior $\lambda_0 \in (0, 1)$ in the beginning of its life cycle, thus firms are *ex ante* homogeneous in this dimension. However, different realizations of the idiosyncratic taste and productivity shocks lead to *ex post* heterogeneity of a firm's beliefs in the cross-section.

37. This class of frameworks include standard price-setting models such as Barro (1972), Dixit (1991), Golosov and Lucas (2007), and Alvarez and Lippi (2014) for example.

38. Note that our framework is fundamentally different from most price-setting models with learning. In the framework by Baley and Blanco (2019), a firm is faced with uncertainty about its productivity. As a result, the problem can be formulated as a Kalman–Bucy filtering problem. Information however evolves exogenously: in their baseline case, these information flows are driven by Brownian motions and a Poisson shock. In contrast, our model considers firms that can explicitly affect their information set. As a result, the flow of information becomes an *endogenous* object.

Note that firms are also heterogeneous in their level of productivity. This dispersion in firms' beliefs and their idiosyncratic productivity is captured by the cross-sectional distribution $\varphi_i(\lambda, z)$ for firms of type i .

A firm's pricing policy is independent from i and k . To obtain price consistency in the aggregate, we must have:

$$P_i = \left(\int_{(\lambda, z)} p^*(\lambda, z)^{1-\sigma_i} d\varphi_i(\lambda, z) \right)^{\frac{1}{1-\sigma_i}}. \quad (7)$$

The market clearing condition for goods is explicitly incorporated in the firm's problem, thus the only remaining market to clear is the factor or labour market. Given the linear production technology, labour demand is simply characterized by:

$$L^d = \sum_i \int_{k \in J_i} \frac{C^i(k)}{z^i(k)} dk.$$

We target labour supply \bar{L} to be equal to one-third, then labour market clearing implies $L^d = \bar{L}$.

We focus on a stationary equilibrium in which any dying firm, which occurs at the rate δ , is immediately replaced by a new firm. Nature assigns a new firm to have an elasticity $\sigma = \sigma_1$ with probability λ_0 . This is common knowledge for all firms in the economy. Hence, it naturally follows that all firms have a prior belief at entry equal to λ_0 as well. This is identical to [Bachmann and Moscarini \(2012\)](#). We simplify the analysis by normalizing the measure of firms to unity in which a fraction γ_1 of firms belong to basket J_1 . This immediately implies that the measure of σ_2 -type firms is $1 - \gamma_1$.

Our assumptions on entry then implies that $\gamma_1 = \lambda_0$ is necessary for a balanced measure of in- and out-flows at the product level. Whenever nominal total spending grows deterministically at the rate π , there is no aggregate uncertainty. If W is the economy's numéraire, then we can define a stationary recursive equilibrium accordingly.

DEFINITION. A *stationary recursive equilibrium* with numéraire $W = 1$ is a tuple (P_1, P_2, S) and a pair of invariant distributions $(\varphi_1(\lambda, z), \varphi_2(\lambda, z))$ in which real variables are constant such that: (1) consumers maximize their utility, (2) firms adopt optimal pricing policies, (3) market for goods and labour clear, (4) aggregate prices are consistently aggregated, and (5) firms exit and enter at the rate δ , and the fraction of type 1 entering firms is λ_0 .

4. QUANTITATIVE PERFORMANCE

In this section, we calibrate our general equilibrium model of active learning and evaluate its quantitative performance. We discuss in detail how well our model is able to hit the moments of our calibration. For external validation, we also assess the model's performance vis-à-vis other micro moments that were not part of the calibration. The details of our computational procedure can be found in [Appendix B.1](#).

4.1. Calibration

The IRI Marketing data set is defined at the weekly level, so we set the model period at 1 week. There are four parameters that are externally calibrated, i.e., β , π , δ , and σ_ε . We do so because most of these parameters have direct, empirical counterparts. Since the model features a weekly frequency, the discount factor is set at $\beta = 0.96^{1/52}$ which reflects an annual interest rate of around

3.8%. Note that the discount factor β already incorporates the exogenous, weekly exit rate of $\delta = 0.4\%$ which comes directly from the IRI Marketing data at the UPC-store level.³⁹ We assume an annual inflation rate of 2%, thus we set $\pi = (1.02)^{1/52} - 1 = 0.038\%$ as the weekly rate of inflation.

The standard deviation of taste shocks σ_ε is disciplined through the variation of sold quantities conditional on no price change in the IRI Marketing data. To be more precise, we regress quantities on prices and a rich set of non-parametric controls and obtain its residuals.⁴⁰ For each store, we then calculate the standard deviation (over time) of these residuals. These standard deviations are then averaged across stores within each product category using store-level revenues as weights. Finally, we aggregate across product categories using equal weights. In the end, we calculate a value of 40% which is similar to the number reported in [Eichenbaum, Jaimovich and Rebelo \(2011\)](#).⁴¹

In the remainder of our analysis, we follow [Golosov and Lucas \(2007\)](#) by assuming log utility for consumption and linear disutility in labour (i.e. $\theta = 1$ and $\chi = 0$). The rest of the parameters are calibrated internally to match various micro-data moments. There are eight remaining parameters: the two elasticities of substitution σ_1 and σ_2 , the prior belief at entry λ_0 , the basket division of income η , the fixed menu cost ψ , the persistence and standard deviation of idiosyncratic productivity ρ and σ_ζ and the disutility of labour ω . These parameters are calibrated jointly and are selected to hit eight moments from the data. They include moments that are typically used in the menu cost literature, such as the frequency, fraction of positive price changes, and size of price changes. The disutility of labour ω is chosen such that aggregate labour supply equals one-third as in [Golosov and Lucas \(2007\)](#) and [Vavra \(2014\)](#). We also target the frequency and absolute size of price changes at different points of time in a product's life cycle to generate age dependence in pricing moments.⁴²

Even though our moments are jointly calibrated, implying that no parameter identifies exactly one specific moment of the data, we can still provide some intuition on which parameters are more informative for a certain set of calibration targets. As we showed in Section 2.4, pricing moments converge to their long-term averages after about 20 weeks. From this point onward, the data seem roughly consistent with a standard menu cost model in the spirit of [Golosov and Lucas \(2007\)](#). As a result, the menu cost ψ and idiosyncratic productivity parameters (ρ, σ_ζ) are most informative for the long-run average pricing moments of the whole sample. This is because firms learn about their type and, hence, their incentives to actively learn decrease over time as well.

In our framework, we generate age dependence in pricing moments through an active learning mechanism. Intuitively, the steepness of these pricing moments (as a function of a product's age) should be informed by those parameters that govern the speed of learning in our model. As

39. We have verified that our results on monetary non-neutrality in Section 5 are not sensitive to variations in entry and exit rates across product categories as observed in our data.

40. These controls include fixed effects at the UPC-store and time level. Our results are similar when we include fixed effects at the UPC-store, UPC-time and store-time level instead (see Table A2 in Appendix A.2).

41. Our findings are similar whenever we use no (revenue) weights in the first (second) stage of averaging. The revenue-weighted standard deviations of residuals do not vary much across product categories. Our estimates for σ_ε vary from 0.38 to 0.43 depending on the set of non-parametric controls. We set $\sigma_\varepsilon = 0.4$ as a reflection of the median estimate. We performed some sensitivity analysis with respect to the parameter σ_ε and found that our results on monetary non-neutrality in Section 5 do not vary much when σ_ε is allowed to vary within a relatively wide range between 0.3 and 0.5. The cross-product category standard deviation of demand shocks is about 0.05. Hence, our range covers two standard deviations below and above our baseline estimate of 0.4. Furthermore, the standard deviation of taste shocks in the data is approximately constant over the product life cycle. Therefore, the size and variation of demand shocks does not vary over the life cycle in our model.

42. In particular, we focus on Weeks 2 and 10 of the product life cycle, but our results are robust to picking other weeks in the early stages of a product's life cycle.

TABLE 3
Internally calibrated values of the model's parameters

Description	Parameter	Value
Elasticity of Substitution 1	σ_1	4.6
Elasticity of Substitution 2	σ_2	15.2
Prior belief at entry	λ_0	0.25
Basket division of income	η	0.15
Fixed cost	ψ	0.03
Productivity persistence	ρ	0.58
Productivity standard deviation	σ_ζ	0.04
Disutility of labour	ω	2.98

emphasized in Section 3.1, this is most accurately captured by the signal-to-noise ratio $\Delta\sigma/\sigma_\varepsilon = (\sigma_2 - \sigma_1)/\sigma_\varepsilon$ and the common prior belief at entry λ_0 . Whenever the signal-to-noise ratio is high, active learning is effective and firms learn about their types fast. Learning is more important whenever there is some level of uncertainty at the firm level (i.e. prior beliefs are not close to the extreme values of zero or unity).

Moreover, the common prior λ_0 also guides the model on the *sign* of price changes. To see why this is the case, recall we follow Bachmann and Moscarini (2012) in that λ_0 is also equal to the fraction of low elasticity firms. Whenever λ_0 is low, then there are also few firms with low elasticities. Hence, the majority of firms have high elasticities with prior beliefs that are aligned with their true type. On the other hand, a small fraction of firms has a low elasticity (who wish to set high prices in the long run), but their prior beliefs are aligned with high elasticities instead. In our environment with decreasing-gains learning, a firm will eventually discover its type (conditional on survival). This implies that if learning occurs relatively fast, as implied by the life cycle patterns of pricing moments in the data, then this set of low elasticity firms will change their prices aggressively upward in order to learn about their type. It is the pricing behaviour of this fraction of firms that will dominate the dynamics on the sign of price changes. This logic for low values of λ_0 gets reinforced whenever η is relatively low as well. Whenever η is low, true static profits for low elasticity firms are actually much lower than according to their initial beliefs. This implies it is more costly for these firms to not learn about their type and, hence, these firms have incentives to learn about their type faster.⁴³

Table 3 shows the model's best parameters in terms of fitting the selected moments, and Table 4 displays the resulting moments from the framework compared to the data.⁴⁴

The productivity parameters (ρ, σ_ζ) are in line with previous estimates in the menu cost literature. The specification for the menu cost ψ implies that the cost conditional on adjustment relative to revenues is around 2%, which is in line with the estimates in Zbaracki, Ritson, Levy, Dutta and Bergen (2004). The value of σ_1 is in the range of values typically used in the menu cost literature. The model requires a somewhat large σ_2 to induce a sufficient amount of active learning. Nonetheless, σ_2 is within the range of estimates in Broda and Weinstein (2010) who compute elasticities of substitution for a variety of products using data similar to ours.

43. A sensitivity analysis of our calibration can be found in Appendix B.2.

44. The model's parameters $\theta = (\sigma_1, \sigma_2, \eta, \lambda_0, \psi, \rho, \sigma_\zeta)$ are chosen to minimize the criterion function $\mathcal{L}(\mathbf{x}, \mathcal{D}; \theta) = \frac{1}{8} \sum_{i=1}^8 \left| \frac{x_i(\theta) - \mathcal{D}_i}{x_i(\theta)} \right|$, where $\mathcal{D} = \{\mathcal{D}_i\}_{i=1}^8$ and $\mathbf{x}(\theta) = \{x_i(\theta)\}_{i=1}^8$ denote the sets of targeted moments in the data and those generated by our quantitative framework, respectively. To ensure that the criterion function was minimized, we have tried different initializations for our calibration.

TABLE 4
Moments of price change distribution

Moment	Data (1)	Model with learning (2)	Full info (3)
Frequency Week 2	0.09	0.10	0.005
Frequency Week 10	0.06	0.06	0.05
Absolute size Week 2	0.17	0.17	0.05
Absolute size Week 10	0.11	0.10	0.07
Frequency	0.05	0.05	0.05
Fraction up	0.66	0.56	0.58
Average size	0.02	0.003	0.008
Not targeted			
Std. of price changes	0.11	0.11	0.08
75th Pct size price changes	0.10	0.11	0.09
90th Pct size price changes	0.18	0.13	0.11

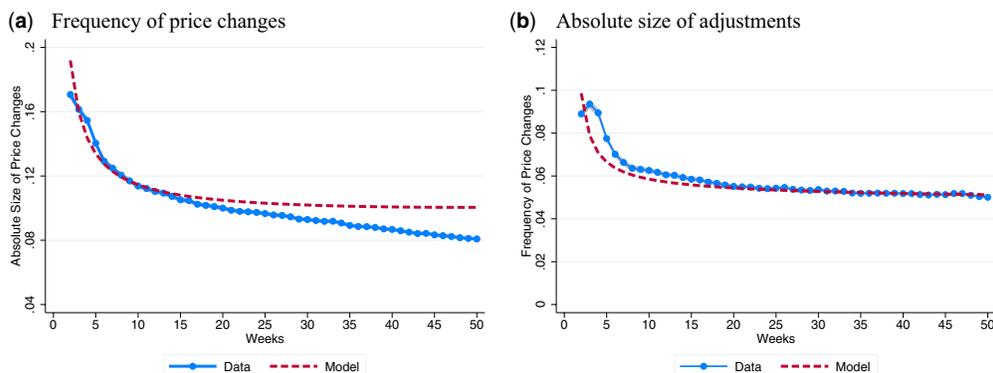


FIGURE 5

Frequency and absolute size of price changes at entry (model vs. data)

Notes: The figure shows the simulation results of the quantitative model and compares them with their data counterpart. We simulate a panel of 1,000 firms over 1,000 periods and compute both the predicted frequency of price adjustments and the absolute size of price changes over the life cycle of a product. The results for the frequency of price changes are shown in (a) and those for the absolute size of price changes are shown in (b).

4.2. Calibrated moments

Table 4 shows that our framework is able to capture pricing moments that are independent of a product's age. The model matches the overall frequency of adjustment and fraction of increasing price changes closely. This is not surprising since our quantitative framework is an extension of a standard menu cost model.

More importantly, our framework is able to endogenously generate age dependence in pricing moments. Table 4 indicates that we are able to match selected weeks of the frequency and absolute size of price changes. Our framework does not only capture the early stages of the product life cycle's pricing moments but also pricing moments along the *entire* life cycle. This is illustrated in Figure 5 shown below.

For comparison, column (3) of Table 4 displays the performance of our framework in a perfect information setting, i.e., active learning motives are absent but firms are heterogeneous in their elasticities. The frequency of price adjustment is slightly increasing over the product life cycle and there is no age dependence in the absolute size of price adjustment; both are characteristics

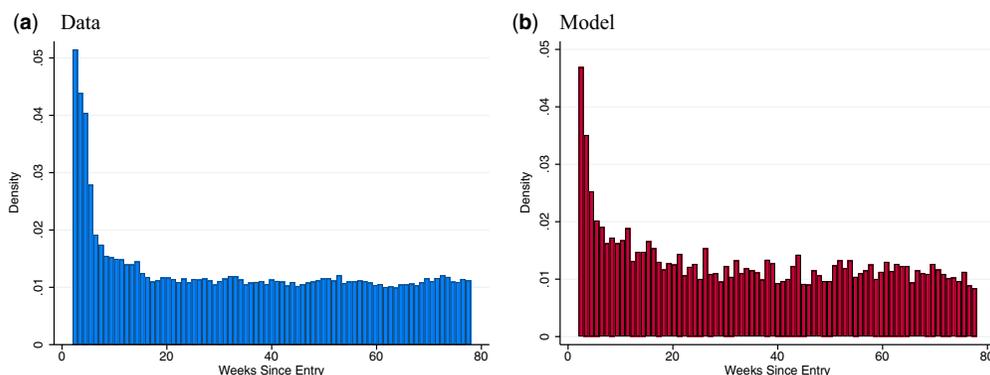


FIGURE 6

Fraction of price changes larger than two standard deviations (model vs. data)

Notes: (a) shows the fraction of price changes larger than two standard deviations from the mean in a given category and store as a function of the age of the product. The products considered are those that last at least two years in the market. (b) shows the results of the quantitative model after simulating a panel of 1,000 firms over 1,000 periods.

that are counterfactual. In this case, since firms choose their entry price optimally, in the presence of menu costs and without learning incentives, they are less likely to adjust their prices during the early stages of their life cycles.

4.3. *Untargeted moments*

We also evaluate the performance of our quantitative framework relative to data moments not targeted in our calibration. To emphasize the consistency of the active learning mechanism with our data, we validate our framework externally along several dimensions of the price change distribution. Our model is consistent with the prevalence of large price changes in the early stages of the product life cycle, as emphasized in our third stylized fact. This is illustrated in Figure 6. Active learning generates this feature endogenously. Intuitively, large price changes generate valuable information for younger products. Our calibration also matches the standard deviation of price changes and price changes in the 75th percentile in absolute value without explicitly targeting them. The model, however, somewhat underpredicts the prevalence of price changes in the 90th percentile of the size distribution.

The model is also consistent with the frequency of positive and negative price changes over a product's life cycle. Figure 7(a) shows that the model is able to capture the full age profile for price changes in both directions. Figure 7(b) shows that, in terms of the life cycle for the size of price changes, the model replicates the empirical patterns; particularly for positive price changes which induce the most curvature as a function of a product's age. Overall, our calibrated model is able to capture salient features of the data on the sign of price changes.

Lastly, we focus on the hazard rate of price changes. Alvarez, Borovicková and Shimer (2021), using the same data set as ours, estimate a decreasing hazard for regular price changes after correcting for unobserved heterogeneity.⁴⁵ Our framework of active learning also generates a hazard rate that is downward-sloping with a small hump at short durations; see Figure 8.

This is not obvious at first glance, and it is the result of several opposing forces. The presence of a menu cost typically results in upward-sloping hazard rates since firms are less likely to adjust

45. We explore and discuss the performance of our quantitative framework relative to other untargeted moments in [Supplementary Appendix DIV](#).

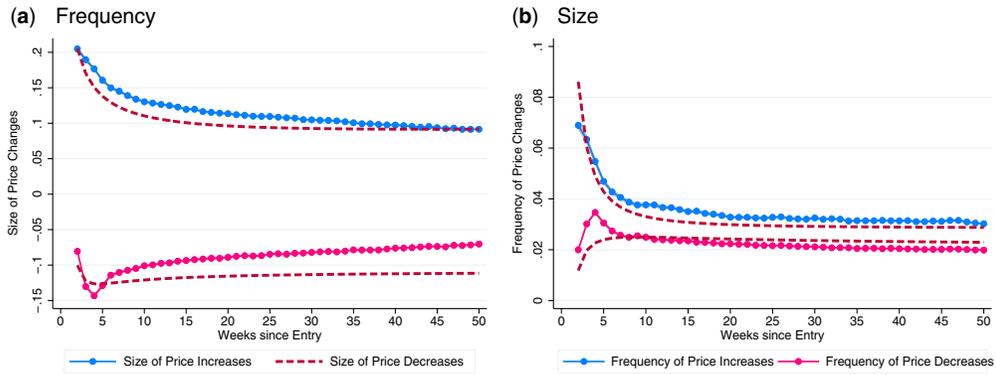


FIGURE 7

Frequency and size of price increases and decreases at entry (model vs. data)

Notes: (a) shows the frequency of increases and decreases in the data and those generated by the model. (b) shows the size of price increases and decreases in the data, and those generated by the model. We simulate a panel of 1,000 firms over 1,000 periods and compute both the predicted frequencies of adjustment and the sizes of price changes over the life cycle of a product.

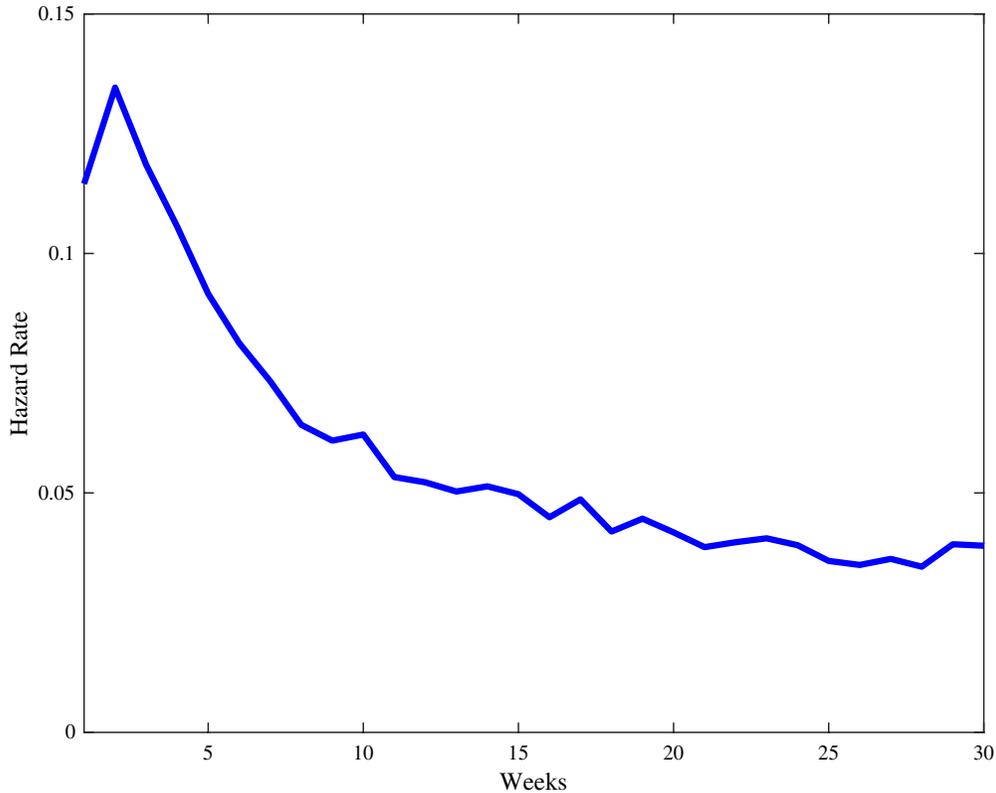


FIGURE 8

Hazard of price changes

Notes: The figure shows the hazard of price changes under the baseline calibration of the model. We simulate a panel of 1,000 firms over 1,000 periods and compute the hazard of price changes using the Kaplan–Meier method.

after they reset their prices. On the other hand, as firms learn, the probability of consecutive price changes is larger at entry since firms can leverage newly obtained information causing a new adjustment. This force, in addition to the fact that the variance of idiosyncratic shocks is large relative to the rate of inflation, contributes to the declining nature of the hazard at long horizons as in the data. This is relevant since [Carvalho and Schwartzman \(2015\)](#) point out that, all else equal, declining hazards of price changes are consistent with higher monetary non-neutrality.

5. PROPAGATION OF NOMINAL SHOCKS

We previously showed that a quantitative model of active learning is able to hit a set of standard pricing moments and can generate age dependence in pricing moments as consistent with our data. This does not only allow us to provide a structural interpretation of our stylized facts on the age dependence of pricing moments, but it also disciplines our framework which is necessary to run counterfactual experiments. In particular, we focus on assessing an economy's level of monetary non-neutrality.

To do so, we perform a counterfactual experiment in which log nominal output increases by a size that is comparable to a 1 week doubling of the nominal output growth rate. That is, the economy's nominal output growth rate is subjected to an unexpected, one-time shock after which it grows again according to the rate in the stationary equilibrium. This experiment follows the one in [Vavra \(2014\)](#) closely. In the following, we will compare the results of our active learning framework relative to a benchmark with perfect information. Our focus lies on the role of age-dependent pricing moments which are generated through an active learning mechanism in our setup. To investigate the importance of active learning, therefore, the proper benchmark is the perfect information setting with heterogeneous elasticities. The latter is calibrated to the same set of moments as in Section 4.1.

In the following, we will focus on the *cumulative* effects on real output after a monetary shock. In other words, our measure for monetary non-neutrality is defined by the cumulative area under the output impulse response function (IRF). We do so for two reasons. First, the efficacy of monetary policy is not always displayed on impact.⁴⁶ Second, a large body of papers in the menu cost literature evaluates and compares the level of monetary non-neutrality of different models along this dimension.

The results are displayed in Figure 9 below. Quantitatively, the cumulative effects on real output are 2.8–3.3 times larger under active learning than in the model with full information.⁴⁷ On impact, approximately 61.9% of the nominal shock goes into output. In a baseline menu cost model with full information, this value is around 57.6%. This implies that, on impact, the real effects of a nominal shock due to active learning increase by about 7.5%.⁴⁸ Importantly, the half-life of the real output response more than doubles with respect to that of the Golosov–Lucas model with heterogeneous firms.

We can also compare our results to other benchmarks in the literature. In our baseline setup with firms producing only one good (i.e. $n = 1$), no fat-tailed shocks and no random menu costs, our

46. See, for example, the discussion in Appendix E of [Alvarez, Lippi and Oskolkov \(2020\)](#). Moreover, in continuous-time models, [Alvarez, Lippi and Passadore \(2016b\)](#) show that the impact effect is exactly zero in time-dependent models. Furthermore, monetary shocks have no first-order effects on impact in state-dependent models whenever these monetary shocks are small and idiosyncratic productivity follows a diffusion process.

47. The cumulative area under the IRF in the full information model is approximately one-sixth of that in the Calvo model.

48. To put these results into perspective and compare with other studies, consider the framework by [Vavra \(2014\)](#) who highlights the importance of time-varying volatility for monetary non-neutrality. He shows that 60.7% of the nominal shock goes into output at the 90th percentile of the volatility distribution whereas it increases to 70.1% at the 10th percentile.

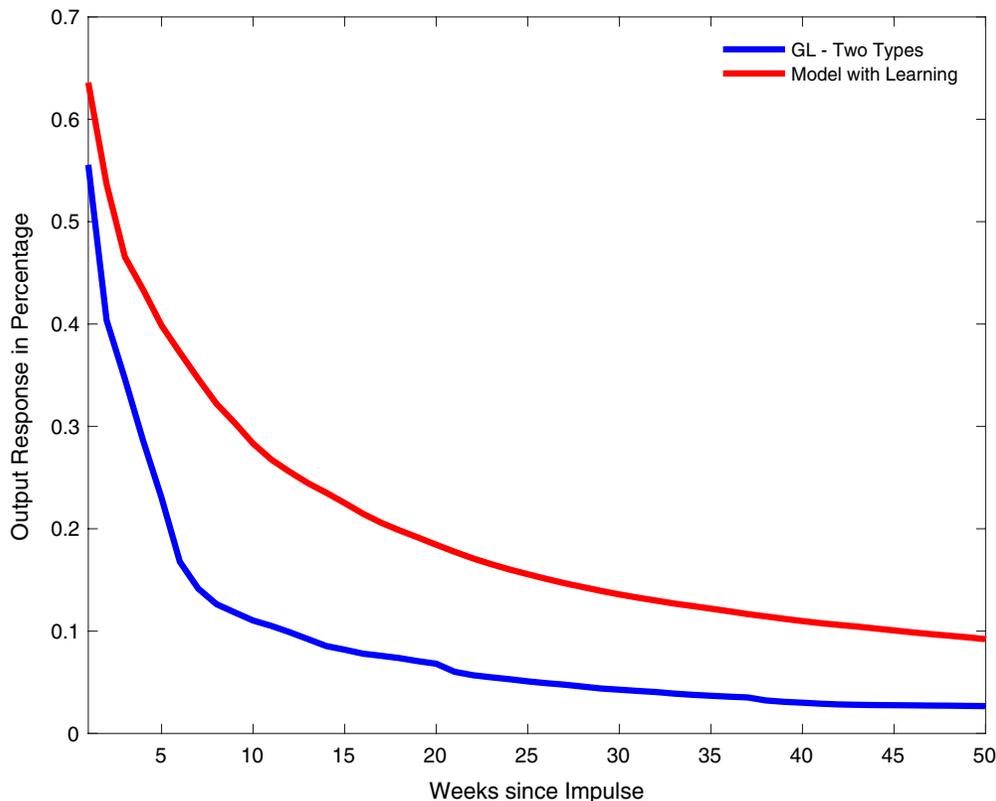


FIGURE 9

Real output response to nominal shock

Notes: The figure shows the response of log real output to an increase in the nominal output growth rate of size 0.00038. The output response is shown in the graph as a percentage of the nominal shock. The blue line depicts the output response in the setup of Golosov and Lucas (2007) with two different types of firms (i.e. σ_1 and σ_2), and the red line is the response in a price-setting model with active learning. All models are calibrated to match the same moments and feature the same fraction of firms of each type.

model with active learning has cumulative effects on real output that are about half of that in a Calvo (1983) framework. This magnitude is comparable to the multi-product setup in Alvarez and Lippi (2014) with firms having $n = \infty$ products in their portfolio. Our results also compare well to a single-product Golosov–Lucas model with random menu costs (also known as the “CalvoPlus” specification) where the fraction of free price adjustments is about 85%.⁴⁹

To put our results in perspective, recall that Shapiro and Watson (1988) attribute 28% of the variation in output at short horizons to what they called “demand” shocks. We follow Lucas (2003) in interpreting them as “nominal” shocks. A model with active learning alone can account for close to a third of this variation; that is, our model generates fluctuations of real output that can account for roughly 7–11% of the US business cycle.

Our experiment indicates that active learning is important in order to understand the impact of monetary policy. The intuition behind this result can be explained through the selection effect,

49. The cumulative effects on real output in our benchmark model with the addition of random menu costs (which are added to generate small price changes) are as large as in the Golosov–Lucas model with a fraction of free price adjustments of 90%. We calibrate the fraction of free adjustments to match the fraction of small price changes defined as $|dp| < \frac{1}{2} \text{mean}(|dp|)$ which in the data are approximately 40%.

where firms adjusting their prices after an aggregate shock are exactly those whose prices display the largest misalignment. [Goloso and Lucas \(2007\)](#) is an extreme example when it comes to selection since price changes in the steady state equilibrium are concentrated at the adjustment barriers. In our framework, the desire to actively learn through prices pushes firms away from the margin of adjustment: price changes are more orthogonal to the nominal shock. This lowers the mass of firms at the original bounds of inaction and, therefore, substantially reduces selection in *size*. This intuition is reflected in the size of the variance of price changes in our model, which is larger than in the baseline model and is not explicitly targeted in our calibration.

This is not the only type of selection that is reduced in our framework. Actively learning firms have vastly higher frequencies of price changes. These firms will most likely adjust their price several times before firms with sharper beliefs after a nominal shock. However, all price changes after the first one made by firms actively learning have no consequence on real output because these firms have already adjusted to the shock. Recall that firms with sharper posterior beliefs have, on average, a lower frequency of price adjustment. Given that the model is calibrated to match the average frequency of price changes, the previous observation implies that the adjustment of the aggregate price level after a nominal shock is significantly delayed. In other words, a higher level of cross-sectional heterogeneity in the duration of price spells reduces selection in the *timing* of price changes after an unanticipated monetary shock (see [Alvarez et al., 2011](#); [Carvalho and Schwartzman, 2015](#)).⁵⁰ As a result, the interquartile range of the duration of price spells is 45–50% larger than in the full information benchmark. Previous work highlights the importance of heterogeneity across sectors. Our model is different in that active learning produces heterogeneity in the frequency of price changes *within* sectors because the joint determination of prices and beliefs generates heterogeneity in the age of each product. This heterogeneity matters for the propagation of nominal shocks since the cumulative real output effect of a nominal spending shock depends on the joint distribution for price gaps and beliefs.

[Alvarez et al. \(2016a\)](#) point out that the kurtosis of the price change distribution captures both of these types of selection in a broad class of state- and time-dependent models. As a result, it can be interpreted as a sufficient statistic for the degree of monetary non-neutrality in an economy. It is worth noting that our framework is *not* nested within this broad class of models. A crucial assumption in their encompassing framework is that a firm's price gap evolves exogenously. Under this assumption, firms close their gaps to a common point (i.e. zero) and this allows for a characterization involving the kurtosis of the price change distribution. In our setup, price gaps evolve *endogenously* because prices and beliefs are jointly determined.⁵¹ Also, firms do not close their price gap conditional on adjustment whenever this price gap is the difference between a firm's current price relative to its full information (or myopic) price.⁵² As emphasized in our

50. [Nakamura and Steinsson \(2009\)](#) illustrate this concept within the context of a simple Calvo model. In that framework, the degree of monetary-non-neutrality is convex in the frequency of price changes across sectors. A straightforward argument based on Jensen's inequality then implies that heterogeneity in the cross-sectional distribution of the frequency of price changes will amplify monetary non-neutrality.

51. In [Supplementary Appendix BII.5](#), we show a simple example with active learning and menu costs in which the boundaries of inaction are dependent on a firm's beliefs.

52. The standard practice in the price-setting literature is that the price gap is defined as a firm's current price relative to the price it would set in the absence of menu costs. The latter is also referred to as a "micro" target. [Baley and Blanco \(2022\)](#) define price gaps relative to steady state values. Conditional on a firm's survival, a firm's full information price can also be considered as its steady state average since firms eventually learn their type. Using their terminology, [Baley and Blanco \(2022\)](#) note that "specifying a micro target is irrelevant for the study of impulse responses centred around steady state: the micro target cancels out as it enters symmetrically the impulse response and the steady state." They emphasize that it is only necessary to specify the *relative* position of a firm's price in the overall distribution; not its absolute level.

simplified model in Section 3.1, firms trade off static profits with obtaining more information and sharpening their posterior beliefs implying that prices are not equal to their static optima.⁵³

6. MORE EVIDENCE ON THE ACTIVE LEARNING MECHANISM

In the following, we provide additional empirical evidence that supports our narrative of active learning. First, we document that firms incorporate attained information from the introduction of a product to later rounds; products that are introduced locally at least 1 year after their national introduction have an attenuated age profile in the average frequency and absolute size of price adjustment. Second, we show that learning increases when the demand for the product being launched is more uncertain (e.g. more novel). More novel products adjust their prices more often and by larger amounts at the early stage of the product's life cycle.

6.1. Variation across space and the timing of products

We first divide every UPC-store pair into two different *waves*. A UPC-store pair belongs to the first wave if it was launched by a retailer before 1 year has passed since the UPC was introduced nationally. Then, a UPC-store pair belongs to the second wave if it was introduced by the same retailer at least one year after the product was first launched at the national level. For each wave, we use an identical empirical strategy as in Section 2. That is, we run the following regression for each wave:

$$Y_{jsct} = \alpha + \sum_{a=1}^A \phi_a \cdot D_{js}^a + \theta_{js} + \tau_t + \gamma_c + \varepsilon_{jsct}.$$

Figure 10 shows that products in the first wave have a higher frequency of adjustment at entry than those in the second wave. The figure also shows the same patterns for the absolute size of price changes. At entry, the size of price changes is 7% larger than the mean for products in the first wave and only 4% larger for products in the second wave. The size of absolute price changes then converges back to the mean of its respective wave. The fraction of new products that is introduced in a specific MSA or by a particular retailer is extremely similar across the two waves. This leads us to believe that the attenuation in pricing moments we observe in the second wave is not due to selection.⁵⁴ Our results remain robust whenever waves are required to occur across different cities. Last, these patterns occur for both price increases and decreases.⁵⁵

These figures, in addition to our other empirical findings, provide suggestive evidence of a learning mechanism. Retailers obtain relevant information about the demand of their product in the first wave and incorporate this knowledge in the second wave. In addition, retailers' incentives to learn about their demand diminish over time as valuable information is collected.

53. The model by Alvarez and Lippi (2020) is another example of a model not encompassed by the sufficient statistic approach of Alvarez *et al.* (2016a); precisely because firms also do not close their price gap upon adjustment in that environment.

54. See Figures AIX (MSA) and Figure AX (retailer) of the Supplementary Appendix. This result conditions on the fact that products are introduced by the same retailers across the two waves. Our findings are slightly reinforced whenever this restriction is relaxed. Consistent with Gilbert (2017), in our data retailers mostly coordinate the introduction of new products across all cities in which they operate. In our data, 82% of new products are introduced in all of a retailer's markets within a 1 year period. We use the remaining variation to define the "second wave" introduction of a product. Our results, however, are similar if we define the length of a wave to be shorter than 1 year.

55. See Figures AVII (different retailers), AVIII (different cities), A.2, and AVI (price increases and decreases) in the Supplementary Appendix.

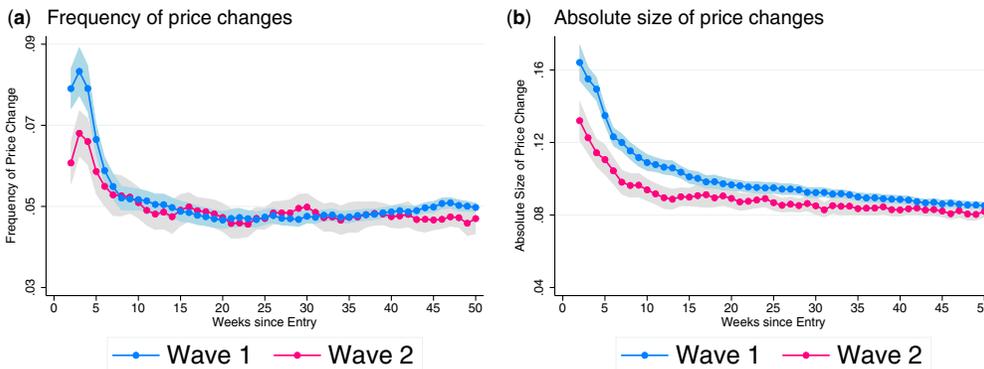


FIGURE 10
Pricing moments by waves (same retailer)

Notes: (a) shows the probability of adjusting prices and (b) shows the absolute size of price changes by waves. Wave 1 contains products that were launched at some location during a period in the first year since the product was introduced *nationally*. Wave 2 contains the same products when launched in different stores a year after their national entries within the same retailer. The results control for fixed effects at the store, product, and time level. The underlying source is the IRI Marketing data.

Under such a mechanism, retailers should make larger price changes and do so more frequently in the early stage of their products’ life cycles. This is exactly what we documented above.

6.2. *Newness index*

Under a setup with demand uncertainty, a firm is required to learn more about the demand for its product that has more novel features (from the consumer’s point of view). As a result, we should expect the age dependence of pricing moments to become more pronounced for more novel products. To test this hypothesis, we construct a measure for the “novelty” of a product that uses detailed information about the characteristics of each UPC which we label as the *newness index*.

This index counts the number of new and unique attributes a product has at the time of its introduction relative to all of the other products ever sold by a store within the same category. Our measure assigns a higher value to products with more unknown features to the store. Our aim is to capture the novelty of a product from the store’s perspective in order to study whether its pricing patterns differ when it sells a product whose demand parameters are more uncertain (i.e. more novel).

We define a product j in category k as a vector of characteristics $V_j^k = (v_{j1}^k, v_{j2}^k, \dots, v_{jN_k}^k)'$ where N_k denotes the number of attributes we observe in category k . Let Ω_{st}^k contain the set of product characteristics for each product ever sold in category k at store s at time t , then the *newness index* of a product j in category k , launched at time t , and in store s is defined as follows:

$$NI_{jst}^k = \frac{1}{N_k} \sum_{i=1}^{N_k} \mathbb{1}(v_{ji}^k \notin \Omega_{st}^k). \tag{8}$$

We assume that each attribute is equally weighted in order to remain agnostic about the relative importance of each attribute for the degree of newness of a product.⁵⁶ We then assess whether

56. For example, the product category “beer” consists of $N_{\text{beer}}=9$ attributes for each barcode: vendor, brand, volume, type (e.g. ale or lager), package (e.g. can or keg), flavour, size (e.g. bottle or six pack), calorie level (e.g. light or

TABLE 5
Newness index—first 6 months

	Equal weights		Revenue weights	
	Frequency (1)	Size (2)	Frequency (3)	Size (4)
Age	-0.074*** (0.003)	-0.164*** (0.005)	-0.106*** (0.005)	-0.168*** (0.009)
Age × Newness	-0.052*** (0.011)	-0.220*** (0.025)	0.0404 (0.024)	-0.290*** (0.040)
UPC × Store FE	✓	✓	✓	✓
Time FE	✓	✓	✓	✓
Cohort	✓	✓	✓	✓

Notes: The table reports the estimates for γ and ϕ of equation 9. The independent variables are the age of the product and the age of the product interacted with the newness index. The dependent variables are the frequency and absolute size of the price changes multiplied times 100. The sample is the first 6 months (or 26 weeks) after the product was first launched. Each regression specification includes UPC-store fixed effects, time fixed effects, and cohort controls that are approximated by the local unemployment rate in the city and month the product was launched. Columns (1) and (2) report the coefficients in which each UPC-store observation is equally weighted. Columns (3) and (4) report the results with revenue weights. Standard errors are clustered at the store level. The ***, **, and * denote significance at 1, 5, and 10% levels, respectively.

the age dependence of pricing moments is stronger for those products that are newer according to our measure. To do so, we run the following regression:

$$Y_{jsct}^k = \alpha + \gamma \cdot \text{age}_{jsct} + \phi \cdot \text{age}_{jsct} \times \text{NI}_{jst}^k + \theta_{js} + \tau_t + \gamma_c + \varepsilon_{jstc}, \quad (9)$$

where ϕ is our coefficient of interest: it reflects how the age heterogeneity of a product's pricing moments depends on its novelty.

Table 5 shows that our index has substantial power in explaining the price-setting patterns we observe at entry. In Appendix A.2, Table A4 shows that our conclusions remain unchanged whenever we measure the newness of a product by whether the introduced product is associated with a new brand. Hence, our exercise confirms the hypothesis that the incentives for active learning increase as the novelty of a product (and, hence, its uncertainty) is higher. We interpret this finding as additional evidence in favour of the active learning narrative.

7. ROBUSTNESS AND EXTENSIONS

In this section, we show some robustness exercises along two dimensions. First, we argue that our model of active learning outperforms other mechanisms (such as the customer base) that potentially generate age-dependent pricing moments. Second, we show that our findings on the importance of active learning for monetary non-neutrality are robust along several, other dimensions.

regular), and colour. If a new product within the category $k = \text{beer}$ enters with a flavour and a volume that has never been sold at store s before, its newness index is $(1+1)/N_{\text{beer}} = 2/9$. Our index should only be considered as an approximation of the novelty of an item given that it relies heavily on the number of attributes provided by the data which might not describe a product in its entirety. On average, we observe ten product characteristics in each category. Figure AXI in the Supplementary Appendix shows how our newness index varies by product category.

7.1. *Customer base*

We argued in Section 4 that our model of active learning provided a good fit to the micro data; in particular, those moments related to the life cycle of US products. However, another narrative that could generate age dependence in a firm’s pricing moments is the customer base. Under this mechanism, future demand directly depends on the level of sold quantity in the past. This generates “investing-harvesting” incentives in which a firm slowly builds up its customer base (or customer capital) by initially pricing low and sets high prices whenever its base reaches a certain level. In this section, we will briefly discuss how a basic menu cost model extended with the customer base falls short of explaining some key moments in the data.

We extend the canonical model of Golosov and Lucas (2007) with a customer base by incorporating “deep habits” as in Ravn *et al.* (2006) and Gilchrist, Schoenle, Sim and Zakrajšek (2017). The full details of this extension can be found in Appendix C.1, but we discuss its qualitative features and its quantitative implications in this section. The main intuition behind the mechanics of the customer base can be summarized by the following system of equations:

$$c_{it} = \left(\frac{p_{it}}{P_t}\right)^{-\sigma} (b_{it-1})^{\eta(1-\sigma)} C_t \tag{10}$$

$$b_{it} = (1 - \delta^C)b_{it-1} + \delta^C c_{it}. \tag{11}$$

Under the customer base, current demand directly depends on the past level of the customer base (parameterized by $\eta < 0$). In our specification, the level of the customer base acts as a demand shifter. Customer capital in period t evolves as a convex combination with weight $\delta^C \in [0, 1]$ between customer capital in period $t - 1$ and period t ’s sold quantity. This immediately implies that pricing decisions are dynamic (even in the absence of menu costs): a firm decides on its price by making the optimal trade-off between losses in static profits and gains in its future customer capital. We embed this version of the customer base in a standard menu cost model à la Golosov and Lucas (2007). To provide the main intuition behind a firm’s pricing decision in a model of the customer base, suppose there are no menu costs for the moment. Then, a firm’s optimal pricing strategy can be characterized as follows:

$$p^{CB}(b, z) = \frac{\sigma}{\sigma - 1} \left(\frac{w}{z} - \beta \mathbb{E}_{z'} \left[v_b(b, z') \delta^C | z \right] \right), \tag{12}$$

where $v_b(b, z)$ is the marginal value of increasing the firm’s customer base when its current base and productivity are valued at b and z , respectively.⁵⁷ Given the concavity of $v(\cdot, z)$, a firm starts off with a low price in the beginning of its life cycle and gradually increases its price as its customer base grows; a firm has incentives to set a low price upon entry to attract customers (at the expense of static profits) and sets high prices once it has established a solid clientele.

We calibrate the customer base model to the same set of moments as in Section 4.1, evaluate its performance vis-à-vis several micro moments and explicitly verify whether other stark predictions of the customer base model are present in the data. Figure C1 in Appendix C.1 indicates that there is age dependence in the frequency of price changes, but this is not true for the absolute size of price changes. However, it also indicates that the frequency of price changes is *increasing* over a product’s life cycle. These observations contradict our first two stylized facts. Given that there is no age dependence in the absolute size of price adjustments, it

57. Note that a firm adopts the standard, myopic CES price whenever it is fully myopic (i.e. $\beta = 0$) or the customer base is static and does not evolve as a function of previously sold quantities (i.e. $\delta^C = 0$).

is not surprising that large price changes also do not vary over the product life cycle. Hence, the customer base is also not consistent with our third stylized fact. These results are not surprising since our calibration finds values for η and δ^C that are close to zero. This implies that the incentives for building up a customer base are weak in the data.⁵⁸ Note that our findings are consistent with a recent empirical literature that has not found strong evidence in favour of the customer base narrative.⁵⁹

7.2. Cross-sectional learning

In our quantitative model, there is no distinction between firms and products. However, retailers sell multiple products and each of them can be sold across multiple locations. If demand curves are somewhat correlated across locations, then a firm can learn faster by experimenting with its prices, not only over time, but also by exploiting cross-sectional variation across locations. To understand why cross-sectional active learning can speed up the learning process, consider the following example.⁶⁰ Suppose that demand curves for each product i in location j are iso-elastic and, for simplicity, assume that elasticities are common across locations:

$$q_{ijt} = s - \sigma_i p_{ijt} + \varepsilon_{ijt}.$$

If demand shocks are perfectly correlated across locations, i.e., $\varepsilon_{ijt} = \varepsilon_{it}$ for all j , then a firm can learn its demand within one period by exploiting price and quantity variation across two stores only. In the other extreme of independent shocks across N stores, a firm still has $N - 1$ more observations than in the single location case. This implies that this firm also has $N - 1$ more opportunities to experiment. A possible concern is that the results in our baseline framework are significantly diminished when allowing for active learning in the cross-section. In the following, we argue that our baseline with single-location firms is nevertheless a good benchmark since some key predictions of the multi-location setup are not supported by our data.

Recall that firms have incentives to learn about their type quickly since active learning is costly. This is because active learning price policies involve losses in static profits as emphasized in our simple model in Section 3. To learn in the most efficient way, a firm would like to maximize the variation in prices across its stores at the early stages of its life cycle. If cross-sectional learning would be important, therefore, we would expect that the price variation across stores of a given retailer's product is declining as a function of a product's age. However, this is counterfactual with the within-retailer standard deviation of prices over the life cycle of US products. We show that the standard deviation of prices within the average retailer is relatively small and constant over the product's life cycle (see Figure A9 in Appendix A.2).⁶¹ This finding is in line with the evidence on uniform pricing by Della Vigna and Gentzkow (2019) and Hitsch, Hortacsu and Lin

58. In fact, our calibrated model of the customer base behaves very similarly to a standard menu cost model à la Golosov and Lucas (2007). In Supplementary Appendix C.1, we consider an alternative calibration based on Foster, Haltiwanger and Syverson (2016). Even though the fit of the model worsens, this calibrated version of the customer base does generate frequencies of price adjustments that are decreasing over the product's life cycle. However, this calibration also features no age dependence in the (absolute) size of price adjustments, overestimates the importance of positive price changes, and underestimates that of negative price changes.

59. This includes work by Fitzgerald, Haller and Yedid-Levi (2016) using customs data and Argente, Fitzgerald, Moreira and Priolo (2021), who find no evidence that markups change systematically with firm/brand/product age in a market using a variety of consumer-level, wholesale-level, and retail-level data sets. We find similar results in our data: prices are trending downward over their life cycle (see Appendix C.1).

60. We thank an anonymous referee for providing this example.

61. Our findings are robust to using different measures of dispersion, e.g., 90-10 differential and IQR.

(2019) who document that “prices and promotions are substantially more similar within stores that belong to the same chain than across stores that belong to different chains.”

Lastly, the motivating example illustrated that cross-sectional learning has more bite in an environment in which price elasticities for a given retailer are common across locations. In fact, if price elasticities were independent of each other, then our framework would be a good description of the data since sales from one location are completely uninformative for making inferences about the demand curve of another location. While it is unclear *ex ante* which scenario is more likely, we believe current evidence lends more support to the latter case. For example, DellaVigna and Gentzkow (2019) show that estimates of within-chain variation in price elasticities are substantial despite wide variation in consumer demographics and competition.

7.3. *Endogenous entry over the cycle*

Our baseline framework reflects a stationary environment in which the number of entrants is constant over time. To investigate whether cyclical changes in the extensive margin of products play an important role in the amplification of nominal shocks, we construct a dynamic version of the model. For the sake of brevity, a full description of the model can be found in Appendix C.3. In this extension, periods of high aggregate productivity imply periods of high product entry. The calibration of this framework shows that the real output effects of a nominal shock are 15% larger in booms than during busts. As the entry rate of products increases, the average firm gets younger and a higher proportion of firms then engages in active learning. Monetary non-neutrality then increases as firms are less likely to adjust their prices after a nominal shock due to active learning. In our calibration, for aggregate shocks of ordinary size, the size distribution of price changes plays a large role in determining the degree to which shocks get propagated during booms. Further, the kurtosis of the distribution increases which, in turn, is indicative of a weakening of the selection effect. However, for very large aggregate shocks, the possibility exists that the average frequency of price adjustments increases, which can offset this effect. An extreme example of this effect is whenever all firms in the economy are replaced every period. In this case, prices are close to fully flexible and the effects on real output are small.

7.4. *Learning with a continuum of types*

Our baseline framework in Section 3 features the simplest form of active learning with firms varying their price as a control. Even though a firm is only uncertain about its demand elasticity and its type can only be high or low, our menu cost model with active learning is already consistent with the life cycle patterns of Section 2. Nevertheless, we show in Appendix C.4 that the key patterns and incentives for active learning are preserved when we model it in a more elaborate form. In this case, there are a continuum of types and firms learn about the slope *and* intercept of their demand function. This particular setup of learning, which is taken from Wieland (2000a), preserves the intuition that we have presented in our two-period model. Optimal pricing under active learning remains a choice between striking a balance between concave, static profits and a convex, continuation value which is identical to our quantitative model.

7.5. *Age trend in demand shocks*

Our benchmark model can capture many features of the data including standard pricing moments and those related to the product’s life cycle as Section 2.4 showed. However, there might be other features of the data concerning entering products that could affect our conclusions. A possible source of concern lies in the fact that entering products do not immediately feature high

quantities of sales. In fact, it might require some time to build up sales for new products (e.g. building up a customer base). Our baseline framework does not reflect a gradual buildup of sales for entering products. Thus, we could be overestimating the importance of new products that in turn affects our results on the propagation of nominal shocks. To deal with this issue, we allow for an exogenous age trend in demand shocks. Appendix C.2 shows the details of this implementation. In this specification, younger products contribute less to aggregate output, but their incentives to actively learn are higher given the prospects of higher sales in the future. These two forces contribute in different directions when measuring the response of real output to a nominal shock, which leaves our results discussed in Section 5 virtually unchanged.

7.6. *Age-dependent exit rates*

A possible concern could be our assumption of a constant rate of exit. Younger products are more likely to exit the product market, so our assumption of age-independent exit rates could potentially bias our results on the propagation of nominal shocks. This is because the composition of products is biased toward younger products that experience a higher frequency and absolute size of price adjustment as discussed in Section 2.4. In Appendix C.5, we show that the product hazard function as a function of age is downward sloping in our data. However, the slope of the hazard function with respect to age is relatively small. Whenever we extend our framework by exogenously incorporating age-dependent exit rates that are consistent with the data, our conclusions are not affected significantly.

8. CONCLUSION

The increasing availability of micro-level data sets has allowed researchers to delve deeper into the mechanics of a firm's dynamic pricing behaviour. Recent studies have found new insights into firms' pricing behaviour along several dimensions. Although there is substantial anecdotal evidence that firms choose different pricing strategies over the life cycle of their products, the degree of price heterogeneity along this dimension and its aggregate implications have remained largely unexplored.

In this article, we aim to fill this gap by developing salient facts on the evolution of products' pricing moments over their life cycle and by providing a structural interpretation for them. We construct a quantitative framework in which firms that face uncertainty about their demand curves can actively learn from their pricing strategies and show that this model can rationalize standard price-setting moments and a new set of stylized facts.

We then investigate the implications of active learning incentives for the propagation of nominal shocks. The calibration of our model can be interpreted as a hybrid between standard menu cost models and active learning models. It delivers the life cycle facts that we documented in the data and is consistent with other salient facts such as the sign of price changes over a product's life cycle and the hazard rate of price changes. In our model, relative to the full information benchmark, the real effects of nominal shocks are at least three times as large when measured by their cumulative effect on real output.

We believe that our quantitative framework contains the minimal amount of ingredients to rationalize our empirical findings. Nonetheless, our model could be extended to cover more complicated mechanisms. We have briefly explored several of them, but we leave the full economic implications of these extensions for future research.

Acknowledgments. We thank Fernando Alvarez, Erik Hurst, Francesco Lippi, Robert Shimer, and Joseph Vavra for their advice and support. We are grateful to Treb Allen, Bong Geun Choi, Steve Davis, Guido Menzio, Elisa Giannone, Mikhail Golosov, Veronica Guerrieri, Greg Kaplan, Oleksiy Kryvtsov, Anthony Landry, Munseob Lee, Thomas Lubik, Robert Lucas Jr., Virgiliu Midrigan, Sara Moreira, Giuseppe Moscarini, Jón Steinsson, Nancy Stokey, Harald Uhlig,

Gianluca Violante, and Alex Wolman. David Argente gratefully acknowledges the hospitality of the Bank of Mexico and the Einaudi Institute for Economics and Finance where part of this article was completed. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System. All estimates and analyses in this article based on Information Resources Inc. data are by the authors and not by Information Resources Inc.

Data Availability Statement

The data underlying this article were provided by Information Resources Inc. under licence. Details on how to obtain access and supplementary material for this article including the code are available on Zenodo at <https://doi.org/10.5281/zenodo.5885183>.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online. And the replication packages are available at <https://dx.doi.org/10.5281/zenodo.5885183>.

APPENDIX

A. DATA

A.1. *Robustness exercises for stylized facts*

A.1.1. **Linear specifications for product age.**

$$Y_{jsct} = \alpha + \phi \cdot \text{age}_{jsct} + \theta_{js} + \tau_t + \gamma_c + \varepsilon_{jsct} \tag{A.1}$$

TABLE A1
Life cycle properties of selected pricing moments—first 6 months

Dependent variable	(1)	(2)	(3)	(4)
	Equal weights		Revenue weights	
Frequency	-0.054*** (0.001)	-0.076*** (0.002)	-0.072*** (0.003)	-0.105*** (0.004)
Frequency increases	-.051*** (0.001)	-0.069*** (0.002)	-0.066*** (0.002)	-0.092*** (0.003)
Frequency decreases	-0.002*** (0.000)	-0.006*** (0.001)	-0.006*** (0.001)	-0.012*** (0.002)
Absolute size	-0.155*** (0.002)	-0.172*** (0.005)	-0.155*** (0.004)	-0.174*** (0.008)
Size increases	-0.194*** (0.003)	-0.223*** (0.006)	-0.186*** (0.005)	-0.223*** (0.010)
Size decreases	0.075*** (0.002)	0.078*** (0.006)	0.091*** (0.004)	0.071*** (0.011)
UPC × Store FE	✓	✓	✓	✓
Time FE		✓		✓
Cohort controls		✓		✓

Notes: The table reports the coefficients (in percent) from OLS regressions. The independent variable is the age of the product and the dependent variables are the moments defined in the table. The dependent variables are multiplied times 100. The sample is the first 6 months (or 26 weeks) after the product was first launched. Each regression specification includes UPC-store fixed effects, time fixed effects, and cohort controls that are approximated by the local unemployment rate in the city and month the product was launched. Columns (1) and (2) report the coefficients in which each UPC-store observation is equally weighted. Columns (3) and (4) report the results with revenue weights. Standard errors are clustered at the store level. The ***, **, and * denote significance at the 1, 5, and 10% levels, respectively.

A.1.2. Age-dependent pricing moments—fixed effects interacted with time.

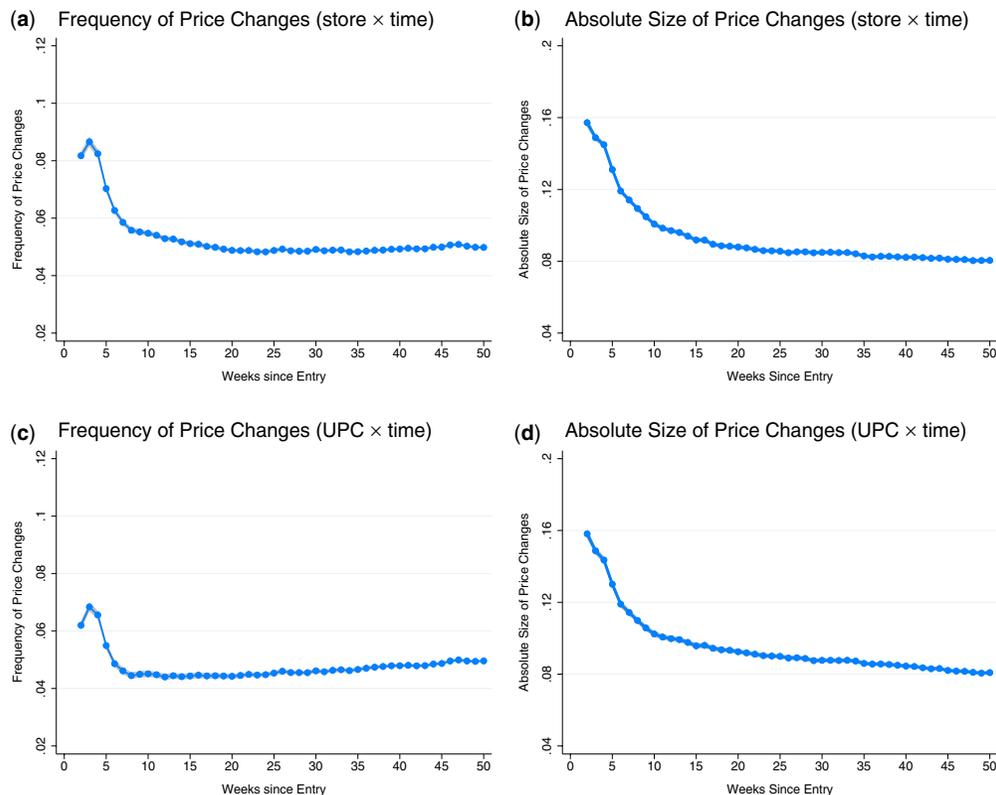


FIGURE A1

Frequency and absolute value of price adjustments at entry

Notes: The graph plots the average weekly frequency of price adjustments ((a) and (c)) and the average absolute size of price adjustments ((b) and (d)) of entering products. The y-axis denotes the probability (absolute size) of price adjustments in a given week and the x-axis denotes the number of weeks the product has been observed in the data since entry. The graph plots the coefficients for the age fixed effects of equation 1 where we use the regular price change indicator as the dependent variable. Equation 1 is computed by controlling for UPC-store effects and the local unemployment rate to represent the cohort fixed effects. (a) and (b) Control for store-time effects and (c) and (d) for good-time effects. The calculation uses approximately 130 million observations and 2.5 million UPC-store pairs. Standard errors are clustered at the store level. The underlying source is the IRI Marketing data.

A.1.3. Age-dependent pricing moments—all price changes in IRI marketing and Nielsen RMS.

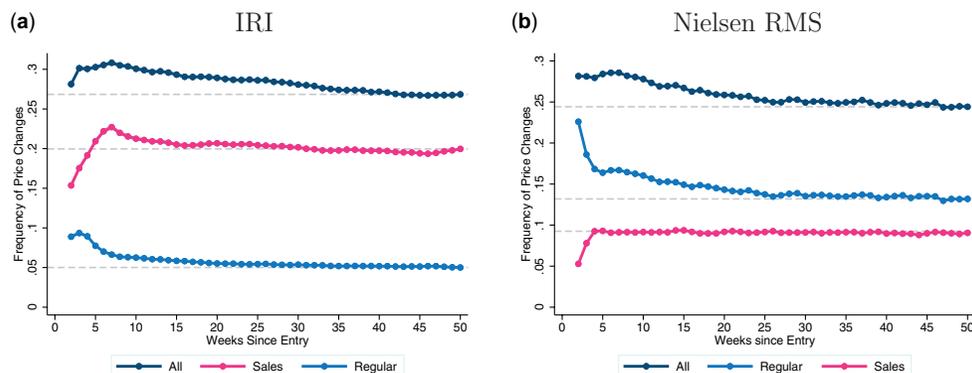


FIGURE A2
Frequency of all price changes

Notes: The figure shows the average weekly frequency of all price changes as a function of the number of weeks since a product entered. (a) shows the frequency of regular prices changes, the frequency of sales, and the frequency of all price changes in the IRI Marketing data. (b) shows the same variables computed using the Nielsen RMS data for the city of Chicago. Since the Nielsen RMS data do not provide a sales flag, we use the sales filters developed in Nakamura and Steinsson (2008). The graph plots the coefficients for the age fixed effects in equation 1 where we use the price change indicator as the dependent variable. Equation 1 is computed by controlling for the store, UPC and time-fixed effects whereas the local unemployment rate proxies for cohort fixed effects.

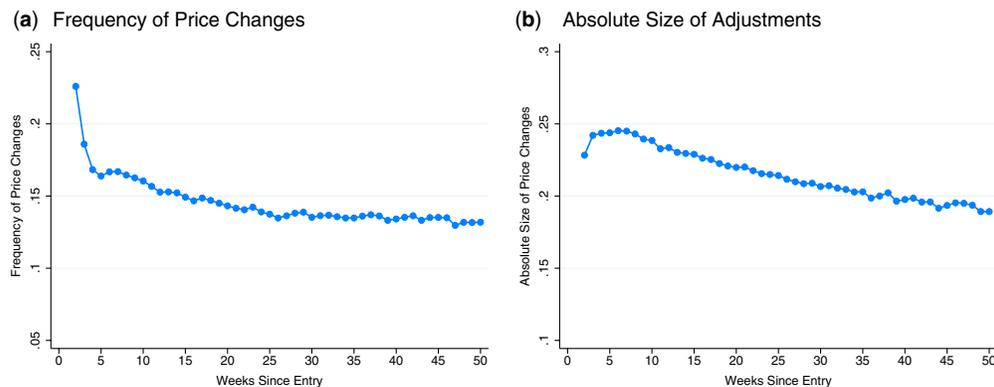


FIGURE A3
Frequency and absolute size of price changes at entry (Nielsen data)

Notes: The figure shows the average weekly frequency and absolute size of price changes as a function of the number of weeks since a product entered. (a) The frequency of regular prices changes whereas (b) shows the absolute size of regular price changes. Since the Nielsen RMS data do not provide a sales flag, we use the sales filters developed in Nakamura and Steinsson (2008). The graphs plot the coefficients for the age fixed effects in equation 1 where we use the price change indicator and the absolute value of the log price change as the dependent variables. Equation 1 is computed by controlling for the store, UPC and time fixed effects whereas the local unemployment rate proxies for cohort fixed effects.

A.1.4. Price changes at exit.

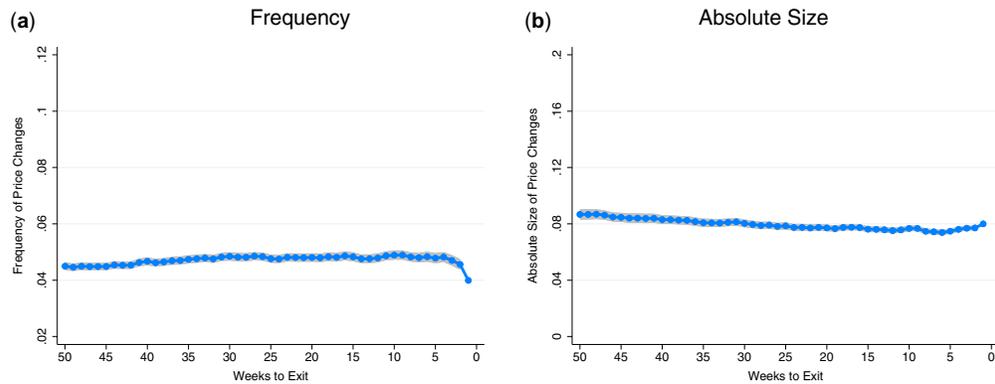


FIGURE A4

Frequency and size of price changes at exit

Notes: (a) plots the frequency of regular price changes at exit. (b) plots the absolute size of regular price changes at exit. The x -axis denotes the number of weeks a product has left in the market before exiting. The graph plots the coefficients for the age fixed effects in the regression where we use the regular price change indicator and absolute value of the log price change as dependent variables. The estimates control for store, UPC, time fixed effects, and the local unemployment rate represents the cohort fixed effects. (a) The frequency of price changes stays mostly constant and decreases only around 1 percentage point near exit. The calculation uses approximately 5.8 million price changes and 2.5 million store-UPC pairs. Standard errors are clustered at the store level. The underlying source is the IRI marketing data.

A.1.5. Sales/promotions over the product’s life cycle.

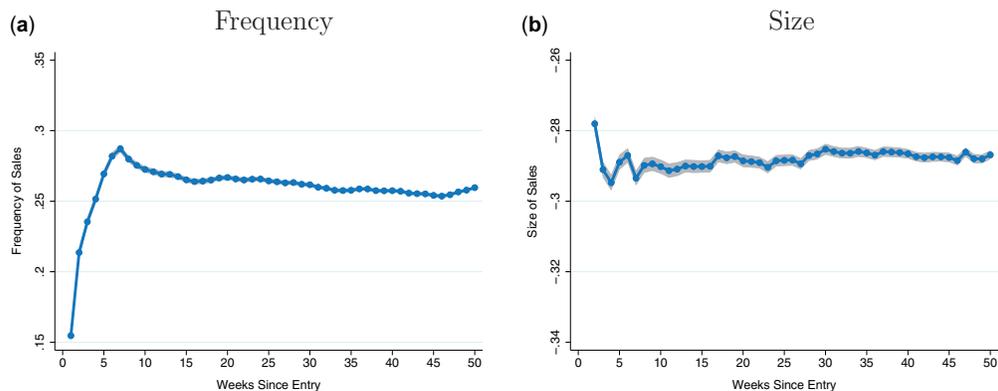


FIGURE A5
Frequency and size of sales at entry

Notes: (a) plots the frequency of sales changes at entry. (b) plots the size of the sales at entry. The x-axis denotes the number of weeks a product has been on the market. The graph plots the age fixed effects coefficients for the regression where we use the sales indicator (provided by the data) and the size of the sales (in logs) as dependent variables. The estimates control for store, UPC, time fixed effects, and the local unemployment rate represents the cohort fixed effects. (a) shows that the probability that a product is on sale is lower at entry. Similarly, the size of sales stays mostly constant during the first year after the product is launched. Standard errors are clustered at the store level. The underlying source is the IRI Marketing data.

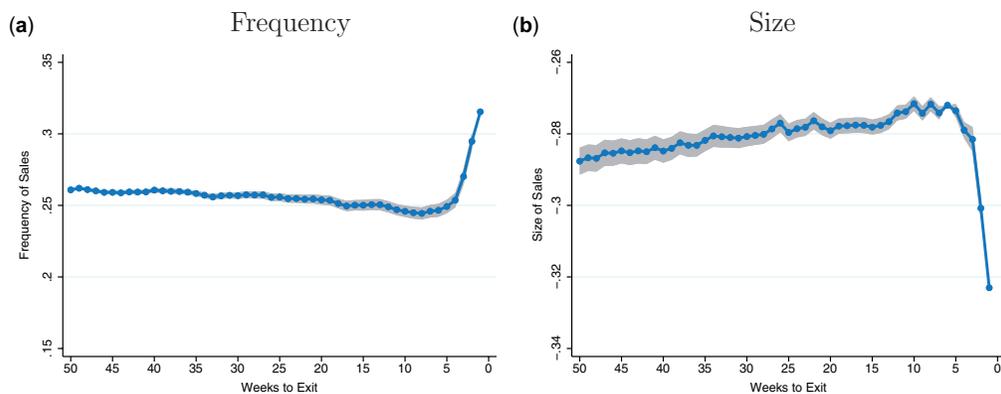


FIGURE A6
Frequency and size of sales at exit

Notes: (a) plots the frequency of sales at exit. (b) plots the absolute size of the sales at exit. The x-axis denotes the number of weeks a product has left in the market before exiting. The graph plots the coefficients for the age fixed effects in the regression where we use the sales indicator (provided by the data) and the size of the sales (in logs) as dependent variables. The estimates control for store, UPC, time fixed effects, and the local unemployment rate represents the cohort fixed effects. The figure shows that at exit, products are more likely to be on sale and the size of these discounts are larger. Standard errors are clustered at the store level. The underlying source is the IRI Marketing data.

A.1.6. Large price changes.

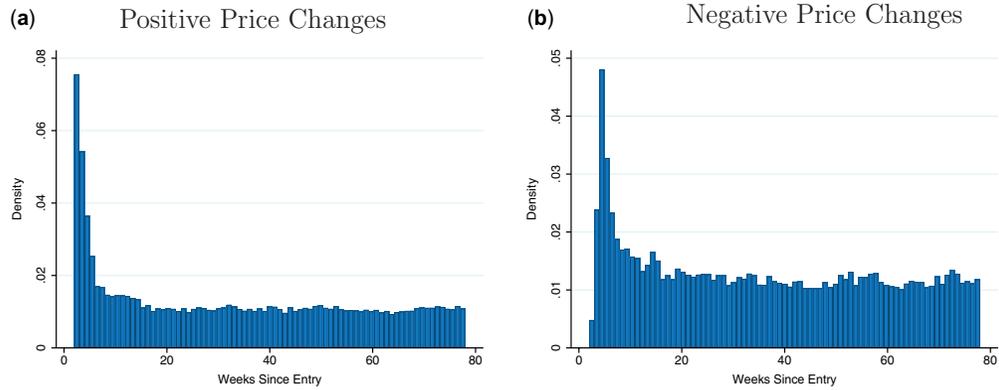


FIGURE A7

Fraction of price changes larger than two std. (positive and negative)

Notes: The figure shows the fraction of price changes larger than two standard deviations from the mean in a given category and store as a function of the age of the product. (a) shows the distribution of large price increases and (b) does this for the distribution of price decreases. The products considered are those that last at least two years in the market. The underlying source is the IRI Marketing data.

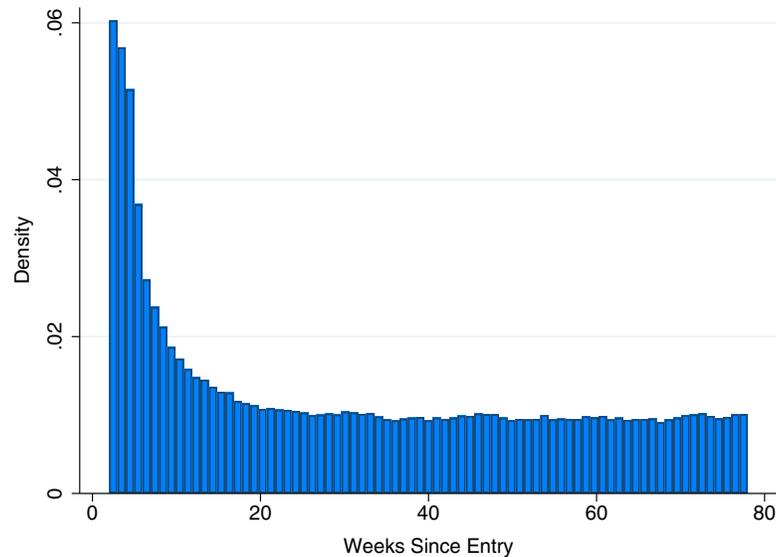


FIGURE A8

Fraction of price changes larger than 30%

Notes: The figure shows the fraction of price changes larger than 30% in a given category and city as a function of the age of the product. The products considered are those that last at least 2 years in the market. The underlying source is the IRI Marketing data.

A.2. *Additional empirical exercises*

A.2.1. **Standard deviation of demand shocks by product category.**

TABLE A2
Standard Deviation of Demand Shocks by Category (σ_ε)

Category			
Beer	0.434	0.430	0.420
Blades	0.391	0.377	0.362
Carbonated beverages	0.489	0.457	0.438
Cigarettes	0.423	0.405	0.386
Coffee	0.435	0.412	0.399
Cold cereal	0.429	0.402	0.383
Deodorant	0.368	0.364	0.358
Diapers	0.388	0.376	0.361
Facial tissue	0.513	0.472	0.437
Frozen dinner	0.495	0.478	0.462
Frozen pizza	0.531	0.509	0.482
Household cleaners	0.387	0.373	0.361
Frankfurters	0.481	0.439	0.408
Laundry detergent	0.437	0.422	0.408
Margarine and butter	0.374	0.343	0.318
Mayonnaise	0.419	0.377	0.346
Milk	0.364	0.315	0.291
Mustard and ketchup	0.441	0.411	0.385
Paper towels	0.520	0.452	0.424
Peanut butter	0.400	0.377	0.348
Photography supplies	0.442	0.423	0.371
Razors	0.357	0.346	0.309
Salty snacks	0.460	0.412	0.399
Shampoo	0.359	0.353	0.347
Soup	0.465	0.438	0.415
Spaghetti sauce	0.456	0.442	0.415
Sugar substitutes	0.407	0.379	0.353
Toilet tissue	0.479	0.427	0.403
Toothbrushes	0.395	0.388	0.380
Toothpaste	0.434	0.426	0.415
Yogurt	0.414	0.382	0.357
UPC×Store	✓	✓	✓
Time	✓		
UPC×Time		✓	
UPC×Time and Store×Time			✓

Notes: The table shows the standard deviation of demand shocks by product category. Conditional on no price change, we regress quantities on a set of fixed effects. Then, we calculate the standard deviation (over time) of the obtained residuals for each store. For each product category, these standard deviations at the store level are then aggregated using revenue weights. The underlying source is the IRI Marketing data.

A.2.2. Entry and exit rates by product category.

TABLE A3
Product entry and exit by category

Category	UPC				UPC×Store			
	Entry	Entry (W)	Exit	Exit (W)	Entry	Entry (W)	Exit	Exit (W)
Beer	0.113	0.020	0.103	0.008	0.239	0.107	0.200	0.064
Blades	0.139	0.122	0.122	0.005	0.303	0.202	0.276	0.091
Carbonated beverages	0.116	0.032	0.108	0.002	0.298	0.153	0.247	0.086
Cigarettes	0.171	0.010	0.134	0.001	0.187	0.071	0.205	0.065
Coffee	0.126	0.032	0.111	0.006	0.256	0.108	0.205	0.067
Cold cereal	0.179	0.049	0.157	0.002	0.264	0.120	0.226	0.076
Deodorant	0.160	0.105	0.132	0.014	0.288	0.188	0.276	0.111
Diapers	0.216	0.183	0.193	0.065	0.389	0.274	0.364	0.236
Facial tissue	0.231	0.108	0.159	0.013	0.300	0.169	0.293	0.165
Frozen dinner	0.150	0.073	0.144	0.059	0.294	0.159	0.267	0.133
Frozen pizza	0.124	0.057	0.101	0.005	0.265	0.128	0.223	0.075
Household cleaners	0.152	0.058	0.135	0.006	0.271	0.152	0.265	0.100
Frankfurters	0.078	0.009	0.082	0.004	0.219	0.097	0.206	0.070
Laundry detergent	0.180	0.076	0.135	0.010	0.297	0.154	0.267	0.108
Margarine and butter	0.081	0.059	0.106	0.010	0.206	0.114	0.199	0.059
Mayonnaise	0.110	0.086	0.095	0.002	0.234	0.141	0.195	0.076
Milk	0.084	0.020	0.097	0.024	0.274	0.100	0.237	0.079
Mustard and ketchup	0.086	0.030	0.119	0.002	0.209	0.092	0.196	0.055
Paper towels	0.189	0.082	0.145	0.014	0.318	0.206	0.318	0.150
Peanut butter	0.099	0.031	0.076	0.000	0.220	0.095	0.184	0.070
Photography supplies	0.087	0.028	0.203	0.007	0.197	0.118	0.291	0.099
Razors	0.231	0.378	0.119	0.001	0.395	0.427	0.322	0.132
Salty snacks	0.158	0.101	0.165	0.015	0.348	0.184	0.332	0.157
Shampoo	0.151	0.126	0.145	0.012	0.336	0.229	0.296	0.133
Soup	0.097	0.043	0.091	0.002	0.227	0.096	0.182	0.062
Spaghetti sauce	0.090	0.034	0.073	0.005	0.225	0.099	0.192	0.065
Sugar substitutes	0.104	0.030	0.080	0.003	0.208	0.096	0.170	0.052
Toilet tissue	0.221	0.147	0.156	0.006	0.322	0.221	0.292	0.133
Toothbrushes	0.134	0.076	0.148	0.006	0.272	0.173	0.277	0.119
Toothpaste	0.185	0.073	0.138	0.005	0.307	0.161	0.269	0.107
Yogurt	0.147	0.117	0.122	0.013	0.322	0.205	0.265	0.090

Notes: The table shows the statistics for entry and exit rates at different levels of aggregation for one year intervals. Columns with “W” indicate that rates are weighted by revenues. Columns (1)–(4) show the statistics at the UPC level whereas columns (5)–(8) show these at the UPC-store level.

A.2.3. Novelty of product: new brand.

TABLE A4
New brand—first 6 months

	Equal weights		Revenue weighted	
	Frequency (1)	Size (2)	Frequency (3)	Size (4)
Age	-0.076*** (0.003)	-0.171*** (0.005)	-0.104*** (0.005)	-0.173*** (0.009)
Age × New brand	-0.028*** (0.005)	-0.047*** (0.014)	-0.020* (0.008)	-0.086*** (0.019)
UPC × Store FE	✓	✓	✓	✓
Time FE	✓	✓	✓	✓
Cohort	✓	✓	✓	✓

Notes: The independent variable is the age of the product interacted with an indicator that equals one if the brand and volume of the product are new. The dependent variables are the frequency and absolute size of price changes multiplied times 100. The sample is the first 6 months (or 26 weeks) after the product was first launched. Each regression specification includes UPC-store fixed effects, time fixed effects, and cohort controls that are approximated by the local unemployment rate in the city and month the product was launched. Columns (1) and (2) report the coefficients in which each UPC-store observation is equally weighted. Columns (3) and (4) report the results with revenue weights. Standard errors are clustered at the store level. The ***, **, and * denote significance at 1, 5, and 10% levels, respectively.

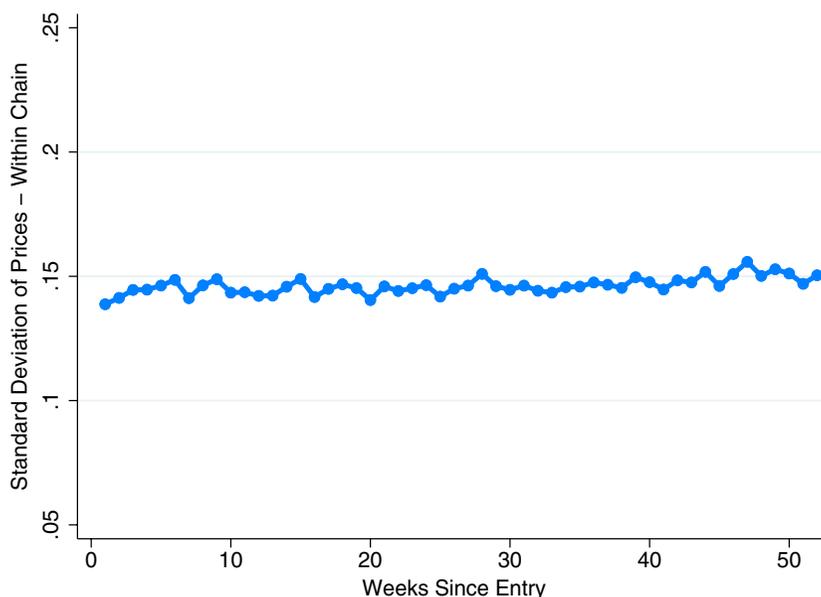


FIGURE A9
Within-chain dispersion of prices

Notes: The figure shows the within-chain standard deviation of prices over the life cycle of US products. The y-axis is the standard deviation of prices within retailers, and the x-axis denotes the number of weeks the retailer’s product has been observed in the data after it entered the market. The underlying source is the IRI Marketing data.

A.2.4. Within-chain dispersion of prices.

B. MODEL

B.1. Computational details

ALGORITHM: PSEUDO-CODE

We assume (log) consumer taste shocks are drawn from a normal distribution, i.e., we have $\varepsilon \sim N(0, \sigma_\varepsilon^2)$. We implement integrals over taste shocks computationally using Gaussian quadrature methods. With some abuse of notation, let the weights and nodes of this approximation method be denoted by $\{\omega_j^{GH}\}_{j=1}^{M^G}$ and $\{\zeta_j^{GH}\}_{j=1}^{M^G}$, respectively.⁶² The quadrature weights and nodes are chosen “optimally.” Specifically, the nodes $\{\zeta_j^{GH}\}_{j=1}^{M^G}$ are the roots of the Hermite polynomial $H_M(\zeta)$ which is defined as $H_M(\zeta) = \int \frac{M^G!}{2\pi i} \exp(-t^2 + 2t\zeta) t^{-(M^G+1)} dt$ and the weights are equal to:

$$\omega_i^{GH} = \frac{2^{M^G-1} M^G! \sqrt{\pi}}{(M^G!)^2 H_{M^G-1}(\zeta_j^{GH})^2}.$$

We discretize firms’ persistent productivity processes with the procedure from Tauchen (1986). This results in a symmetric M^T -dimensional transition matrix whose elements are denoted by $\{T(i, j)\}_{i, j}$.

I (initialization). Set P_1^0, P_2^0, ω^0 and convergence criteria $\Delta_\varepsilon > 0$. Let the counter k be equal to 0.

II (out). Given k , set P_1^k, P_2^k , and ω^k .

III (in). Set $L = \frac{1}{3}$ and calculate aggregate profits Π^k using:

$$L^\lambda = \frac{1}{\omega^k} \frac{1}{L + \Pi^k}$$

Then, set aggregate income equal to $S^k = L + \Pi^k$. Define the aggregate state as $\mathcal{A}^k = (P_1^k, P_2^k, S^k)$ and set the aggregate price level as $P^k = (P_1^k)^\eta (P_2^k)^{1-\eta}$.

IV. Using value function iteration over a finite grid, solve the firm’s problem by obtaining the value function $v(\lambda, z, p_{-1})$:

$$v(\lambda, z, p_{-1}) = \max \left\{ v^A(\lambda, z), v^{NA}(\lambda, z, p_{-1}) \right\}$$

with

$$v^A(\lambda, z) = \max_{p \geq 0} \lambda \Pi^1(p; z) + (1 - \lambda) \lambda \Pi^2(p; z) - \psi \frac{1}{p^k}$$

$$+ \beta \lambda \mathbb{E}_{\varepsilon, z'} \left[v \left(b_1(\lambda, \log(\frac{p}{1+\pi}), \varepsilon), z', \frac{p}{1+\pi} \right) \middle| z \right]$$

$$+ \beta (1 - \lambda) \mathbb{E}_{\varepsilon, z'} \left[v \left(b_2(\lambda, \log(\frac{p}{1+\pi}), \varepsilon), z', \frac{p}{1+\pi} \right) \middle| z \right]$$

$$v^{NA}(\lambda, z, p_{-1}) = \lambda \Pi^1(p_{-1}; z) + (1 - \lambda) \lambda \Pi^2(p_{-1}; z)$$

$$+ \beta \lambda \mathbb{E}_{\varepsilon, z'} \left[v \left(b_1(\lambda, \log(\frac{p_{-1}}{1+\pi}), \varepsilon), z', \frac{p_{-1}}{1+\pi} \right) \middle| z \right]$$

$$+ \beta (1 - \lambda) \mathbb{E}_{\varepsilon, z'} \left[v \left(b_2(\lambda, \log(\frac{p_{-1}}{1+\pi}), \varepsilon), z', \frac{p_{-1}}{1+\pi} \right) \middle| z \right],$$

where integrals over ε are approximated by Gaussian quadrature methods with $M^G = 21$ and integrals over firm-level productivity $z' | z$ are calculated using the methods from Tauchen (1986) with $M^T = 23$. Furthermore, we have $M^P = 1501$ grid points for prices and $M^\ell = 51$ grid points for beliefs.

IV. Store the optimal pricing policy function $p^*(\lambda, z; \mathcal{A}^k)$. We explicitly verify that this maximizer associated with $v^A(\lambda, z)$ is single-valued and the global maximizer over the interval $[p_2^*, p_1^*]$.

62. The superscript stands for “Gaussian-Hermite” quadrature. This is useful to approximate functions of the form $f(x) = \exp(-x^2)$ which includes the family of normal distributions.

V.a. SIMULATION. Simulate a panel of $T=1,000$ weeks and $N=1,000$ firms who use the policy function $p^*(\lambda, z; \mathcal{A}^k)$. The first 50 periods are burn-in periods.⁶³

SIMULATION INITIALIZATION. The initial distributions of cross-sectional beliefs and productivity $\varphi_{i,0}(\lambda, z; \mathcal{A}^k)$ for $i=1,2$ are independent. Furthermore, beliefs are degenerate at λ_0 and productivity is set at its stationary distribution. For each firm $n \in \{1, 2, \dots, N\}$, assign it to be a firm of type $\sigma_n = \sigma_1$ with probability λ_0 . Set time counter t to zero.

V.b. Given a firm's belief $\lambda_{n,t}$, let firm n set price $p^*(\lambda_{n,t}, z_{n,t}; \mathcal{A}^k)$. Generate log sales by drawing log demand shocks $\varepsilon_{n,t} \sim N(0, \sigma_\varepsilon^2)$ through:

$$\log(q_{n,t}) = -\sigma_n \log\left(p^*(\lambda_{n,t}, z_{n,t}; \mathcal{A}^k)\right) + \mu_i + \log(S^k) + \varepsilon_{n,t},$$

where $\mu_i = (\sigma_i - 1)\log(P_i^k) + \log(\eta_i)$. Update the firm n 's posterior to:

$$\lambda_{n,t+1} = B(\lambda_{n,t}, p^*(\lambda_{n,t}, z_{n,t}; \mathcal{A}^k), q_{n,t}, S^k).$$

Apply exogenous death shocks δ for each firm. If a firm exits, then replace it by a new firm which is assigned to be a type σ_1 firm with probability λ_0 . Its prior becomes λ_0 . Furthermore, its idiosyncratic productivity is drawn from the unconditional distribution. A firm is allowed to change its price in the first period upon entry without incurring the menu cost.

V.c. Calculate $\varphi_{i,t+1}(\lambda, z; \mathcal{A}^k)$ for each $i=1,2$. Stop the simulation when the distribution of beliefs and productivity settles in both measures of active firms or when the number of simulation periods exceed some upper bound $T > 1$, i.e., $\sup_{\lambda, z} |\varphi_{i,t+1}(\lambda, z; \mathcal{A}^k) - \varphi_{i,t}(\lambda, z; \mathcal{A}^k)| < \Delta_\varphi$ for $i \in \{1, 2\}$ and/or $t = T$. Otherwise, set $t := t + 1$ and repeat step **V.b.**

VI. Calculate \bar{P}_i^{temp} with the *simulated* density $\tilde{\Phi}_i(\lambda, z; \mathcal{A}^k)$:

$$\bar{P}_i^{\text{temp}} = \left(\sum_{\lambda, z} p^*(\lambda, z; \mathcal{A}^k)^{1-\sigma_i} \tilde{\Phi}_i(\lambda, z; \mathcal{A}^k) \right)^{\frac{1}{1-\sigma_i}},$$

where $\tilde{\Phi}_i(\lambda, z; \mathcal{A}^k)$ is the *empirical* cross-sectional probability distribution function of beliefs. Then, calculate aggregate labour and label it as L^{temp} . If we have:

$$\max \left\{ \left| \frac{P_1^k - \bar{P}_1^{\text{temp}}}{0.5(P_1^k + \bar{P}_1^{\text{temp}})} \right|, \left| \frac{P_2^k - \bar{P}_2^{\text{temp}}}{0.5(P_2^k + \bar{P}_2^{\text{temp}})} \right| \right\} < \Delta_\varepsilon$$

then, stop. Otherwise, set $P_1^{k+1} = \bar{P}_1^{\text{temp}}$ and $P_2^{k+1} = \bar{P}_2^{\text{temp}}$. Let $\omega^{k+1} > \omega^k$ if and only if $L - L^{\text{temp}} < 0$. Update the counter to $k := k + 1$ and repeat step **II**.

Finally, we simulate the model $M^S = 50$ times and perform our counterfactual experiment for each simulation. The reported output impulse response function in Section 5 is the average impulse response function across these simulations.

B.2. Elements of identification

In this section, we confirm our intuition of Section 4.1 and verify that those parameters governing the signal-to-noise ratio and prior beliefs are most informative for pricing moments at the early stages of a product's life cycle. To do so, we set all parameters to their estimated values (reported in Table 3) and then vary each of them separately while holding the others constant. We focus on the parameters σ_2 and λ_0 and show that these parameters are likely more informative for active learning.

In this case, we show the frequency and absolute size of price adjustment in Week 2 since they summarize the slope of the life cycle patterns for both moments.⁶⁴ Figure B1 reports comparative statics with respect to σ_2 . The value of this parameter in our baseline calibration is represented by the vertical red line. Given that both σ_1 (vertical black line) and σ_ε are held constant, the horizontal axis also represents the signal-to-noise ratio. By construction, it is zero when $\sigma_1 = \sigma_2$. Figure B1(a) and (b) shows that as σ_2 approaches σ_1 the life cycle patterns of both the frequency and the absolute

63. Quantitatively, our results do not change much whenever we simulate $N = 10,000$ firms instead.

64. The comparative statics for the frequency of price adjustment and the absolute size of price changes in Week 10 are very similar.

size of price adjustment flatten. This is because, for a given σ_ε and σ_1 , the gains from actively learning decrease as the signal-to-noise ratio is reduced. Similarly, on the other extreme, if σ_2 is very far apart from σ_1 the demand curves are far enough from each other such that the gains from active learning are also reduced. Recall from our intuition in Section 3.1 that a firm is basically able to learn its type through a single price change whenever demand curves are separated out a lot; that is, whenever σ_2 is much higher than σ_1 . This is reflected by the flattening of the life cycle profile of the frequency for large values of σ_2 . Our estimated value of σ_2 lies between these two extremes showing that in our model active learning motives are important and somewhat persistent.

Figure B1(c) and (d) considers the same moments as before but vary the common prior λ_0 instead. Both panels show that as λ_0 approaches unity, the incentives for active learning decrease and the life cycle patterns of both moments flatten. Figure B1(c) shows that the life cycle patterns of the frequency of adjustment informs us on the value of the initial belief. A value of λ_0 closer to 0.35 increases the steepness of the life cycle profile for the frequency of adjustment reflecting larger incentives for learning. This is because beliefs closer to 0.35 are closer the confounding belief $\hat{\lambda}$. This is the level for λ_0 at which the continuation value of gaining more information is minimized. As a result, any subsequent price change away from the confounding price $\hat{p} = p^*(\hat{\lambda})$ will generate a lot of information on the margin and, hence, incentives for active learning are maximized at this point.

Figure B2 reports more conventional comparative statics of menu cost models. It shows, for example, that the menu cost ψ affects both the frequency of price adjustment and the absolute size of price adjustments in Week 10 in opposite directions (see Figure B2(a) and (b)). It also shows that the standard deviation of idiosyncratic productivity σ_ζ are positively associated to the frequency of price adjustment.

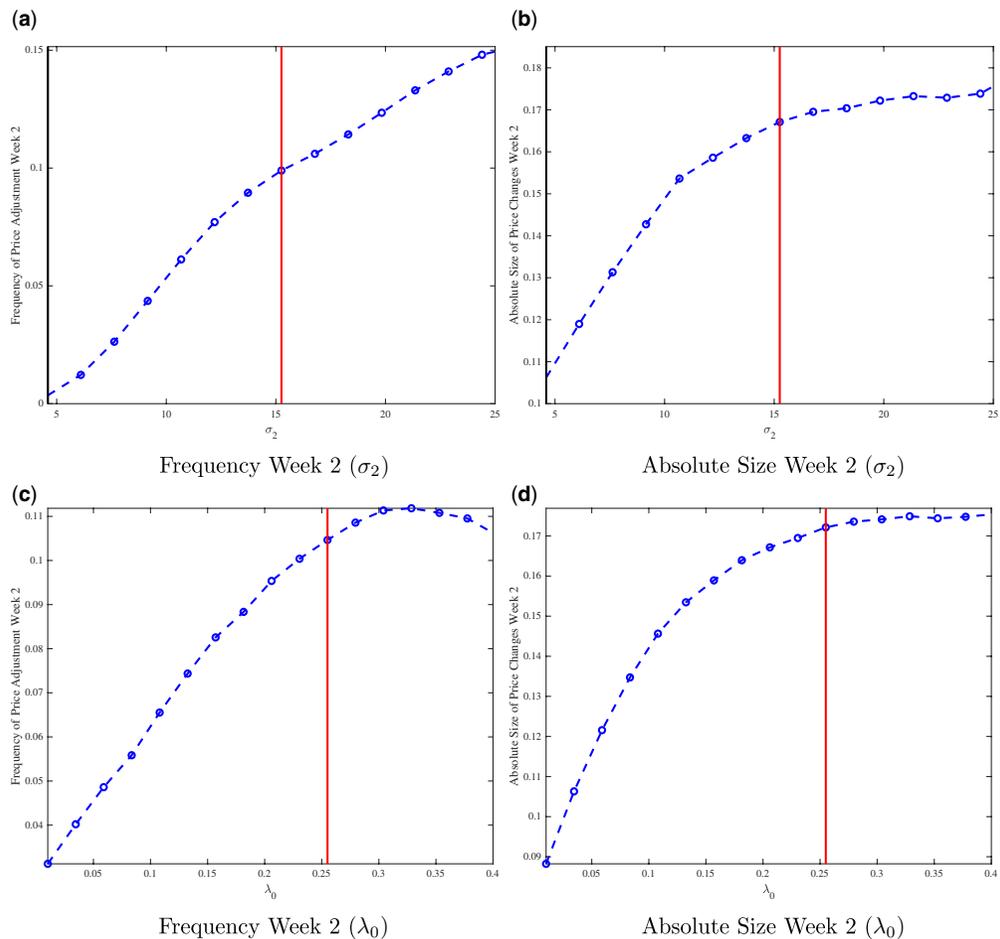


FIGURE B1

Elements of identification

Notes: The graphs plot the relationship between the estimated parameters and several moments relevant for active learning. We set all parameters to their baseline estimates reported in Table 3 and represented by the vertical red line. Then, we move each parameter around its estimated value holding the others constant. In (a) and (c), the moment reported is the frequency of price adjustment in Week 2. In (b) and (d), the moment reported is the absolute size of price adjustment in Week 2. The parameter in (a) and (b) is σ_2 and the black vertical line indicates the value of σ_1 . The parameter in (c) and (d) is λ_0 .

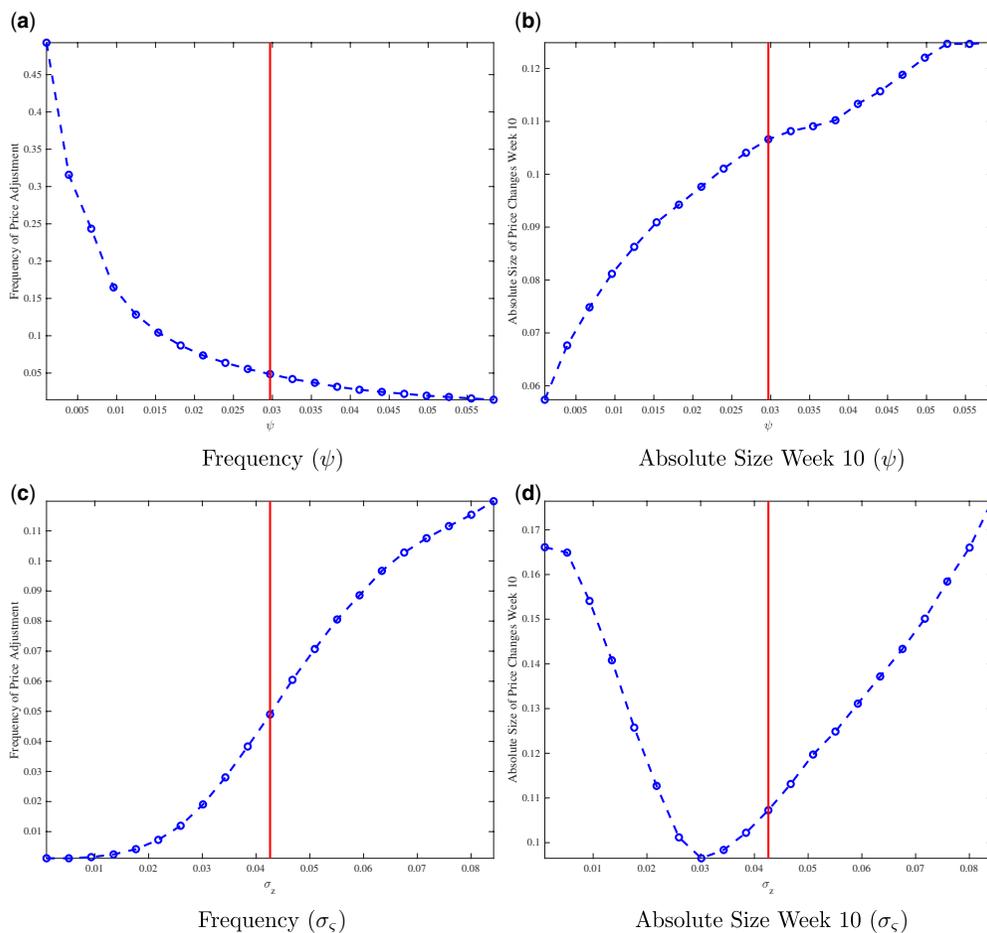


FIGURE B2
Elements of identification

Notes: The graphs plot the relationship between estimated parameters and several pricing moments that are relevant for active learning. We set all parameters to their baseline estimates reported in Table 3. The value for each parameter in the baseline calibration is represented by vertical red lines. Then, we move each parameter around its estimated value while holding the others constant at their baseline values. The moment reported in (a) and (c) is the frequency of price adjustment. The moment reported in (b) and (d) is the average size of price increases. The parameter in (a) and (b) is the menu cost ψ . The parameter in (c) and (d) is the standard deviation of the idiosyncratic cost shock σ_ζ .

C. EXTENSIONS

C.1. Customer base

In this section, we extend the canonical framework of [Golosov and Lucas \(2007\)](#) by adding a customer base in the most parsimonious way as possible. We do so by modelling the customer base through “deep habits” as in [Ravn *et al.* \(2006\)](#) and [Gilchrist *et al.* \(2017\)](#). We show that such a model can rationalize the fact that product-level sales are dependent on age, but its pricing implications are not consistent with the documented life cycle patterns of Section 2.

Under the setup with deep habits, the aggregate consumption good C_t consists of a continuum of monopolistically competitive goods and is constructed as follows:

$$C_t = \left[\int_0^1 \left(\frac{c_{it}}{b_{it-1}^\eta} \right)^{\frac{\sigma-1}{\sigma}} di \right]^{\frac{\sigma}{\sigma-1}},$$

where b_{it} is the habit stock associated with good i at time t and we have $\sigma > 1$ and $\eta < 0$. The former denotes the elasticity of substitution between goods whereas the latter indicates the relative importance of habits for current consumption. The good-specific habit stock is assumed to be external: consumers take this level of stock as given. In addition to being more tractable, the assumption of external habits avoids the time-inconsistency problem of a firm setting its price associated with good-specific internal habits ([Nakamura and Steinsson, 2011](#)). Following [Ravn *et al.* \(2006\)](#) and [Gilchrist *et al.* \(2017\)](#), we impose an exogenous law of motion for the external habit:

$$b_{it} = (1 - \delta^C) b_{it-1} + \delta^C c_{it},$$

where δ^C denotes the depreciation rate of the customer base. Given the fact that consumers take the stocks of external habits $\{b_{it}\}_{i \in [0,1]}$ as given at time t , its good-specific demand can be derived as:

$$c_{it} = \left(\frac{p_{it}}{P_t} \right)^{-\sigma} (b_{it-1})^{\eta(1-\sigma)} C_t.$$

The CES price index, adjusted for external habits, is denoted by:

$$P_t = \left(\int_0^1 (p_{it} b_{it-1}^\eta)^{1-\sigma} di \right)^{\frac{1}{1-\sigma}}.$$

Identical to the baseline framework, we assume that each firm produces with a labour only production function that features constant returns to scale. Given a consumer’s demand for good c_{it} , we can derive that a monopolistically competitive firm i ’s real profits are equal to:

$$\Pi_{it}(p_{it}; b_{it-1}) = \left(p_{it} - \frac{W_t}{z_{it}} \right) \left(\frac{p_{it}}{P_t} \right)^{-\sigma} (b_{it-1})^{\eta(1-\sigma)} \frac{C_t}{P_t}.$$

In the absence of menu costs, a firm’s dynamic pricing problem can be characterized as:

$$\begin{aligned} v(b_{-1}, z) &= \max_{p \geq 0} \Pi(p; b_{-1}) + \beta \mathbb{E}_{z'} [v(b, z') | z] \\ \text{s.t. } b &= (1 - \delta^C) b_{-1} + \delta^C c(p) \\ c(p) &= \left(\frac{p}{P} \right)^{-\sigma} (b_{-1})^{\eta(1-\sigma)} C. \end{aligned}$$

The optimal pricing policy is characterized by its first-order condition which satisfies:

$$c(p) + \left(p - \frac{W}{z} \right) \frac{\partial c(p)}{\partial p} = -\beta \mathbb{E}_{z'} \left[v_b(b, z') \delta^C \frac{\partial c(p)}{\partial p} \right].$$

Given that consumer demand has a constant price elasticity, we can rearrange this equation as:

$$p^{\text{CB}}(b_{-1}, z) = \frac{\sigma}{\sigma-1} \left(\frac{W}{z} - \beta \mathbb{E}_{z'} [v_b(b, z') \delta^C | z] \right).$$

Hence, a firm’s optimal price under a customer base is always dominated by its myopic pricing policy since we have $v_b(\cdot, z) \geq 0$ for all z . In the quantitative extension, we assume that firms are faced with a nominal rigidity in the form of a menu cost (denoted in units of labour) and allow for inflation. A firm’s dynamic programming problem is then summarized by the following Bellman equation:

$$v(b_{-1}, z, p_{-1}) = \max \left\{ v^A(b_{-1}, z), v^{\text{NA}}(b_{-1}, z, p_{-1}) \right\},$$

TABLE C1
Internally Calibrated Values of the Model's Parameters (Customer Base)

Description	Parameter	Value
Elasticity of substitution	σ	5.7
Elasticity of customer capital	η	-0.08
Customer capital depreciation	δ^C	0.09
Fixed cost	ψ	0.10
Productivity persistence	ρ	0.90
Productivity standard deviation	σ_ζ	0.03

where the values of adjusting and not adjusting are given by:

$$\begin{aligned}
 v^A(b_{-1}, z) &= \max_{p \geq 0} \left(\frac{p}{P} - \frac{W}{zP} \right) \left(\frac{p}{P} \right)^{-\sigma} b_{-1}^{\eta(1-\sigma)} C - \psi \frac{W}{P} \\
 &\quad + \beta \mathbb{E}_{z'} \left[v \left(b, z', \frac{p}{1+\pi} \right) \middle| z \right] \\
 \text{s.t. } b &= B(b_{-1}, p) = (1 - \delta^C) b_{-1} + \delta^C \left[\left(\frac{p}{P} \right)^{-\sigma} b_{-1}^{\eta(1-\sigma)} C \right] \\
 v^{\text{NA}}(b_{-1}, z, p) &= \left(\frac{p}{P} - \frac{W}{zP} \right) \left(\frac{p}{P} \right)^{-\sigma} b_{-1}^{\eta(1-\sigma)} C - \psi \frac{W}{P} \\
 &\quad + \beta \mathbb{E}_{z'} \left[v \left(b, z', \frac{p}{1+\pi} \right) \middle| z \right] \\
 \text{s.t. } b &= B(b_{-1}, p) = (1 - \delta^C) b_{-1} + \delta^C \left[\left(\frac{p}{P} \right)^{-\sigma} b_{-1}^{\eta(1-\sigma)} C \right].
 \end{aligned}$$

The two crucial parameters that govern the customer base are η and δ^C . As mentioned before, η determines how important a build-up customer base is for current demand whereas δ^C determines the speed at which a customer base accumulates.

Then, we calibrate our model of the customer base to the same moments as mentioned in Section 4.1 and verify whether such a framework is consistent with our stylized facts. The calibrated parameters can be found below.

Our calibration finds low values for η and δ^C which implies that the incentives for the customer base are weak in the data. As a result, this calibrated model essentially behaves as a standard menu cost model à la Golosov and Lucas (2007). As can be seen from the figure below, there is no age dependence in the absolute size of price adjustments. Even though there is age dependence for the frequency of price changes, it increases (rather than decreases) over a product's life cycle. This is not surprising since, upon entry, firms are allowed to change their prices without incurring a menu cost. Thus, a firm has little to no incentives to change its price directly afterward. After this period, the frequency of price changes does not display any age dependence: price changes are then purely determined by the stochastic process for idiosyncratic productivity.

Since there is no age dependence in the absolute size of price changes, there should also be no age dependence in the fraction of large price changes over a product's life cycle.⁶⁵ Our intuition is confirmed by the figure below.

Lastly, the customer base model does predict that positive price changes are relatively more important than negative ones. However, its patterns for the frequency and size of price adjustment are counterfactual. This is illustrated in the figure below.

C.2. Demand shocks with age-dependent trend

In Section 2, we showed that entering products play a substantial role in the aggregate economy. Approximately 45% of products in the US market entered in the last 5 years and they account for about 30% of total expenditures. However,

65. Our pricing formula (12) implies that price changes in the customer base setup are roughly constant over the life cycle whenever the customer base increases in a convex fashion. This is because of the concavity of $v(\cdot, z)$.

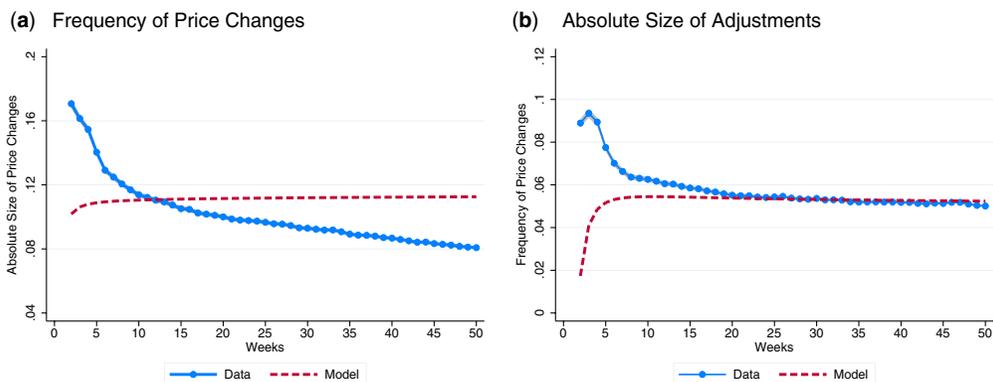


FIGURE C1

Frequency and absolute size of price changes at entry (customer base)

Notes: The figure shows the simulation results of the quantitative model with a customer base and compares them with their data counterparts. We simulate a panel of 1,000 firms over 1,000 periods and compute both the predicted frequency of price adjustments and the absolute size of price changes over the life cycle of a product. The results for the frequency of price changes are shown in (a) and those for the absolute size of price changes are shown in (b).

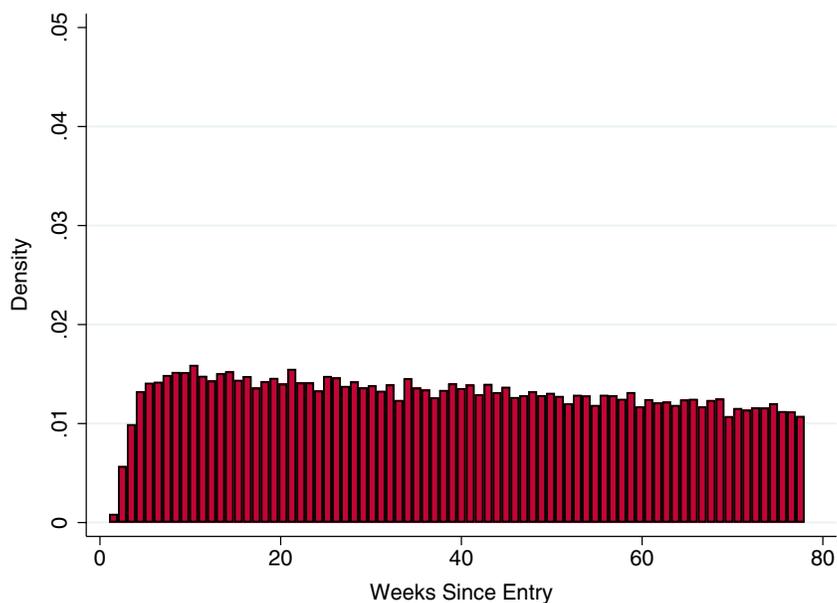


FIGURE C2

Fraction of price changes larger than two standard deviations (customer base)

Notes: The figure shows the fraction of price changes larger than two standard deviations from the mean in a given category and store as a function of the age of the product. The products considered are those that last at least 2 years in the market. The underlying source is the IRI Marketing data.

entering products do not immediately reach these high levels of sales because, for example, they need to build up customer bases. [Paciello et al. \(2019\)](#) argue that the pricing dynamics of firms are heavily influenced by customer retention concerns that are relatively more important for entering products. As a result, our quantitative results could be biased whenever we do not take into account that product sales require some time to be built up.

The easiest, albeit mechanical, way of incorporating the fact that entering products' sales grow over time starting from a relatively low level is through an age trend in demand shocks. We add the following adjustment to the baseline

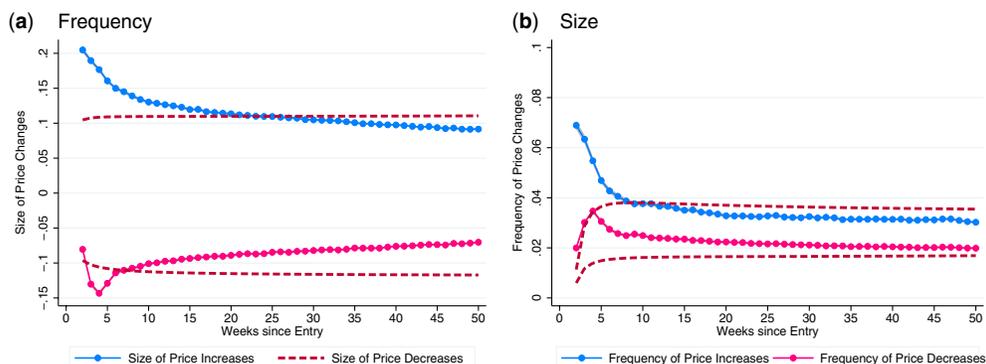


FIGURE C3

Frequency and size of price increases and decreases at entry (customer base)

Notes: (a) shows the frequency of price increases and decreases in the data and those generated by the calibrated customer base model. (b) shows the size of price increases and decreases in the data, and those generated by the calibrated customer base model. We simulate a panel of 1,000 firms over 1,000 periods and compute both the predicted frequencies of adjustment and the sizes of price changes over the life cycle of a product.

TABLE C2
Moments of Price Change Distribution (Customer Base)

Moment	Data (1)	Model with learning (2)	Customer base (3)
Frequency Week 2	0.09	0.10	0.002
Frequency Week 10	0.06	0.06	0.06
Absolute size Week 2	0.17	0.17	0.10
Absolute size Week 10	0.11	0.11	0.11
Frequency	0.05	0.05	0.05
Fraction up	0.66	0.56	0.68
Average size	0.02	0.003	0.04
		Not Targeted	
Std. of price changes	0.11	0.11	0.10
75th Pct size price changes	0.10	0.11	0.13
90th Pct size price changes	0.18	0.13	0.14

framework. Assuming that the realization of taste shocks are independent across all groups and over time, we specify demand shocks to depend on a product's age through:

$$\alpha_t^i(k, a) = \omega^i(k, a) \alpha_t^i(k).$$

Then, the price index becomes:

$$P_{it} = \left(\int_{k \in J_i} \int_a \omega^i(k, a) p_i^i(k)^{1-\sigma_i} da dk \right)^{1/1-\sigma_i}.$$

Recall that we are trying to capture the fact that entering products' sales start at a relatively low level and grow asymptotically towards some steady rate in a concave fashion. This steady rate is achieved fairly quickly in our data set and occurs within the first 3 months after the product enters the market. As a result, we use the following functional form:

$$\omega^i(k, a) = t \cdot a^\nu,$$

where we calibrate the initial level t and age-dependent slope ν to match the age profile of sales. Even though younger firms contribute less to output under this specification, their incentives to actively learn are higher given the prospects of higher sales in the future. In other words, the opportunity cost of poor sales at entry is low relative to future potential sales. These two forces contribute in different directions when measuring the response of real output to a nominal shock. Nonetheless, we find that the effects on real output remain quantitatively similar to those obtained in our benchmark model.

C.3. Endogenous entry over the cycle

Our baseline framework reflects a stationary environment in which the number of entrants is constant over time. Even though our baseline framework is successful in replicating the stylized facts that we documented in Section 2, it does not capture whether the magnitude of the nominal shock amplification varies over the business cycle. In this section, we construct a fully dynamic version of our model to investigate whether cyclical changes in the extensive margin of products play an important role in the amplification of nominal shocks.

Following work by Lee and Mukoyama (2018), aggregate productivity shocks are the source of aggregate fluctuations. In addition, the entry rate of products is endogenous, which allows it to vary with the aggregate state of the economy. In this section, we present the necessary ingredients to allow for a procyclical entry rate in our model.

Consumers and firms are identical as in the baseline framework. However, a firm's productivity now consists of two components: an idiosyncratic one as described in Section 3 and an aggregate component $Z_t \in \{Z_L, Z_H\}$. Aggregate productivity Z_t follows a symmetric, two state Markov chain. The transition matrix between high and low aggregate productivity is then characterized by:

$$\begin{bmatrix} \vartheta & 1 - \vartheta \\ 1 - \vartheta & \vartheta \end{bmatrix}$$

The average duration of a state is given by:

$$\sum_{\tau=1}^{\infty} \tau (1 - \vartheta) \vartheta^{\tau-1} = \frac{1}{1 - \vartheta}.$$

Recall that a price-adjusting incumbent firm has a three-dimensional idiosyncratic state. We denote this state by $\mathbf{v}_t^i(k) = (\lambda_t^i(k), z_t^i(k), p_{t-1}^i(k))$. Then, a firm's *ex ante* expected profits, excluding its menu cost, can be written as a function of the idiosyncratic state $\mathbf{v}_t^i(k)$ and the aggregate state $\xi_t \equiv (S_t, P_{1t}, P_{2t}, Z_t)$:

$$\Pi_t(p; \mathbf{v}_t^i(k), \xi_t).$$

A firm chooses a path of prices $\{p_t^i(k)\}_{t \geq 0}$ to maximize its expected, discounted profits. The firm's problem in Bellman form is then equal to:

$$v(\mathbf{v}, p_{-1}; \xi) = \max \left\{ v^A(\mathbf{v}; \xi), v^{NA}(\mathbf{v}; \xi) \right\},$$

where the value of adjusting and not adjusting are respectively given by:

$$\begin{aligned} v^A(\mathbf{v}; \xi) &= \max_{p \geq 0} \Pi(p; \mathbf{v}, \xi) - W \cdot \psi \\ &+ \mathbb{E}_{\xi', \varepsilon, z'} \left[\mathbf{q}(\xi, \xi') \lambda v \left(b_1 \left(\lambda, \log \left(\frac{p}{1+\pi} \right), \varepsilon \right), z', \frac{p}{1+\pi}; \xi' \right) \right. \\ &\left. + \mathbf{q}(\xi, \xi') (1 - \lambda) v \left(b_2 \left(\lambda, \log \left(\frac{p}{1+\pi} \right), \varepsilon \right), z', \frac{p}{1+\pi}; \xi' \right) \right] \xi, z \\ v^{NA}(\mathbf{v}; \xi) &= \Pi(p_{-1}; \mathbf{v}, \xi) \\ &+ \mathbb{E}_{\xi', \varepsilon, z'} \left[\mathbf{q}(\xi, \xi') \lambda v \left(b_1 \left(\lambda, \log \left(\frac{p_{-1}}{1+\pi} \right), \varepsilon \right), z', \frac{p_{-1}}{1+\pi}; \xi' \right) \right. \\ &\left. + \mathbf{q}(\xi, \xi') (1 - \lambda) v \left(b_2 \left(\lambda, \log \left(\frac{p_{-1}}{1+\pi} \right), \varepsilon \right), z', \frac{p_{-1}}{1+\pi}; \xi' \right) \right] \xi, z \end{aligned}$$

where the stochastic discount factor is given by $\mathbf{q}(\xi, \xi') = \beta \frac{u'(C')}{u'(C)}$.

There is a pool of potential entrants. In the beginning of a period, everyone observes the aggregate state ξ_t . Furthermore, every potential entrant is endowed with an idiosyncratic productivity z drawn from the exogenous distribution H . If a potential entrant wants to become a producer, she needs to pay a fixed entry cost c_E , which is denoted in units of labour. At entry only, we assume that the entrant is allowed to choose its price without incurring the menu cost. The value of becoming a producer then becomes:

$$v_t^E(z_t; \xi_t) = v_t^A(\lambda_0, z_t, 1; \xi_t) + W_t \cdot \psi.$$

This structure indicates that only those potential entrants with sufficiently high values for z_t can actually enter the product market. In fact, there is a threshold value z_t^* that is defined by the free entry condition:

$$v_t^E(z_t^*; \xi_t) = W_t \cdot c_E$$

$$= \omega S_t \cdot c_E$$

such that potential entrants become producers if and only if their drawn level of productivity satisfies $z_t \geq z_t^*$.

To analyse the model in general equilibrium, we need to consider an environment in which consumers and firms engage in optimal behaviour while the markets for goods and labour clear. Optimization behaviour is apparent from the representative consumer's first-order conditions and firms' value functions. The market for goods clears by construction because we plug the optimal consumer demand into the firm's optimization problem. As a result, we only need to clear the labour market.

Let $\varphi_t(\lambda, z, p_{-1})$ denote the labour demand of a firm with idiosyncratic state \mathbf{v} . Assuming the mass of potential entry in each period is one, its distribution at period t is denoted by $\mu_t(\mathbf{v})$. Then, the quantity of labour demanded by incumbent producers is:

$$L_t^{d,p} = N_t \cdot \int_{\mathbf{v}} \varphi_t(\mathbf{v}) d\mu_t(\mathbf{v}),$$

where N_t denotes the actual mass of potential entrants in period t . Furthermore, labour is used for the costs of entry. Thus, total labour demand L_t^d can be characterized as:

$$L_t^d = L_t^{d,p} + N_t \cdot (1 - H(z_t^*)) \cdot c_E.$$

The above equation characterizes the optimal labour demand. Then, the market for labour clears when the labour supply equals labour demand.

In the following, we describe how we calibrated the exogenous, aggregate productivity process. Following the RBC literature, we assume that the level of aggregate productivity can be well approximated by an autoregressive process. We use Fernald's (2014) quarterly utilization-adjusted time series for TFP and detrend it with the HP filter (using a smoothing parameter of 1,600). Then, we run a linear regression of this detrended series on its lagged counterpart and calculate the standard deviation of the residuals. These residuals are then interpreted as shocks to aggregate productivity. We find a value of 0.009 at the quarterly level. Converting this to the weekly level, we obtain $\sigma_Z = \frac{0.009}{\sqrt{12}} \simeq 0.0026$. We normalize the trend for aggregate productivity to unity and define a boom or bust as a one standard deviation increase or decrease from the trend respectively. As a result, we obtain $Z_H = 1.0026$ and $Z_L = 0.9974$.⁶⁶ Furthermore, we assume that the average duration of a boom or bust is 35 months, which is approximately 140 weeks. This indicates a value of the transition probability $\vartheta = 1 - \frac{1}{140} \simeq 0.99286$.

In terms of computation, it is conventional to adopt a Krusell–Smith procedure for extensions with aggregate shocks; see Vavra (2014) for a pricing application. Under this procedure, agents use a forecasting rule that contains first (and possibly higher-order) moments of the aggregate price indices P_{1t} and P_{2t} rather than whole distributions. When we computed this extension, firms' pricing policies are based on the stationary price indexes for two reasons. First, this procedure approximates pricing policy functions in a dynamic environment fairly well. In fact, several papers argue that aggregate shocks do not have a large impact in terms of non-neutrality in models without active learning (see e.g. Section 6 in Golosov and Lucas, 2007; p. 1178 in Midrigan, 2011; pp. 992–993 in Nakamura and Steinsson, 2009). Specifically, Alvarez and Lippi (2014, p. 7 of Technical Appendix) show that the impact of aggregate shocks on the aggregate price level impulse response can be well approximated by imposing steady-state pricing policy functions (rather than solving for the exact solution). Numerically, it can be shown that for small shocks this is also true in our model. Second, a full-blown procedure in the spirit of Krusell and Smith Jr. (1998) would increase a firm's number of state variables significantly. This is because firms need to conjecture law of motions for the two aggregate price indexes. Even in the simplest computational application (when these law of motions only contain first moments), this would imply that firms have five state variables to keep track of. Higher-order approximations would increase it by even more which might render the computational procedure infeasible. In Supplementary Appendix CIII, we do lay out the details on how such a model could be solved.

Furthermore, there are intuitive arguments why the previous logic on the irrelevance of aggregate shocks for monetary non-neutrality also applies in our framework. For instance, we can allow for nominal income to follow a random walk with drift in logs as in Vavra (2014). This is a convenient example since shocks to nominal income are isomorphic to, for example, aggregate productivity shocks. Note that we need to keep track of two aggregate states (one for each aggregate price index) and an updating rule that possibly requires higher-order moments to be sufficiently precise. However, nominal spending S_t (or s_t in logs) does not appear in a firm's posterior belief functions, i.e., $b_1(\lambda, p, \varepsilon)$ and $b_2(\lambda, p, \varepsilon)$. In fact, we have:

$$b_i(\lambda, p, \varepsilon) = \left(1 + \frac{1-\lambda}{\lambda} \exp \left(\frac{1}{2} (-1)^{1(i=2)} \left[\left(\frac{\varepsilon}{\sigma_\varepsilon} \right)^2 - \left(p \cdot \frac{\Delta\sigma}{\sigma_\varepsilon} - \frac{\Delta\mu}{\sigma_\varepsilon} + \frac{\varepsilon}{\sigma_\varepsilon} \right)^2 \right] \right) \right)^{-1}.$$

66. Vavra (2014) performs a similar exercise with real output per hours worked and finds a standard deviation for aggregate productivity shocks of 0.006 at the monthly level. This means a value of $\sigma_Z = \frac{0.006}{\sqrt{4}} \simeq 0.003$ which is similar to what we find above.

Recall that the key component behind the active learning mechanism is that a firm can affect its posterior beliefs with its prices. As shown in the above however, these beliefs are independent of S_t . As a result, there are no first-order effects of S_t on firms' pricing policies under active learning.

C.4. Bayesian learning with a continuum of types

Our baseline framework in Section 3 features the simplest form of active learning with firms varying their price as a control. Even though a firm is only uncertain about its demand elasticity and its type can only be high or low, our menu cost model with active learning is already consistent with the life cycle patterns that we showed in Section 2. Nevertheless, we show that the key patterns and incentives for active learning are preserved when we use a more elaborate form of learning.

Consider a monopolistically competitive producer with constant marginal costs c who is faced with a linear demand curve of the following form:

$$q = \alpha - \sigma p + \varepsilon,$$

where the demand shock satisfies $\varepsilon \sim N(0, \sigma_\varepsilon^2)$. There are two key differences in this framework compared to the baseline. First, the firm faces uncertainty about the intercept α and the slope σ of its demand curve. Second, the pair (α, σ) is now part of a continuous parameter space. A firm's prior belief is thus specified by a probability density function on (α, σ) over \mathbb{R}^2 . This is denoted by $f(\alpha, \sigma | \theta, \mathcal{Q})$, where \mathcal{Q} denotes its information set that consists of the history of previous realized sales. θ parameterizes the distribution f . We specify the firm's initial prior over (α, σ) to be a multivariate normal distribution that is parameterized by the mean vector $(a, s)'$ and variance-covariance matrix Σ . The latter is symmetric and satisfies:

$$\Sigma = \begin{pmatrix} v_a & v_{as} \\ v_{as} & v_s \end{pmatrix}.$$

Therefore, we have $\theta = (a, s, \text{vec}(\Sigma))' = (a, s, v_a, v_{as}, v_s)'$.

Given a prior distribution $f(\alpha, \sigma | \theta_t, \mathcal{Q}_{t-1})$ at time t , where $\mathcal{Q}_{t-1} = \{q_1, q_2, \dots, q_{t-1}\}$ and after observing a realized sales value of q_t , a firm will update its prior to a posterior distribution according to Bayes' rule:

$$\begin{aligned} f(\alpha, \sigma | \theta_{t+1}, \mathcal{Q}_t) &= \frac{f(q_t | \alpha, \sigma, \theta_t, \mathcal{Q}_{t-1}) \cdot f(\alpha, \sigma | \theta_t, \mathcal{Q}_{t-1})}{f(q_t | \theta_t, \mathcal{Q}_{t-1})} \\ &\propto f(q_t | \alpha, \sigma, \theta_t, \mathcal{Q}_{t-1}) \cdot f(\alpha, \sigma | \theta_t, \mathcal{Q}_{t-1}). \end{aligned}$$

The family of Gaussian distributions is conjugate to itself with respect to a Gaussian likelihood function, so this means that the posterior function must be of the multivariate normal form as well. A standard application of the Kalman filter shows that:

$$\begin{pmatrix} a \\ s \end{pmatrix}_{t+1} = \begin{pmatrix} a \\ s \end{pmatrix}_t + \frac{\Sigma_t X_t}{X_t' \Sigma_t X_t + \sigma_\varepsilon^2} \left(q_t - X_t' \begin{pmatrix} a \\ s \end{pmatrix}_t \right) \quad (\text{K1})$$

$$\Sigma_{t+1} = \Sigma_t - \frac{\Sigma_t X_t X_t' \Sigma_t}{X_t' \Sigma_t X_t + \sigma_\varepsilon^2}, \quad (\text{K2})$$

where $X_t = (1, -p_t)'$. A more direct derivation can be found in Zellner (1971). The function that changes the parameters from the prior distribution into their posterior counterparts, as a function of observed sales and a chosen price, is denoted by $B: \Theta \times \mathcal{P} \times \mathbb{R}_+ \rightarrow \Theta$. Thus, the above system of equations can be compactly written as $\theta_{t+1} = B(\theta_t, p, q)$. The *ex ante* expected profits are defined as:

$$\Pi(p; \theta) = \int_{\varepsilon \in \mathbb{R}} \int_{\alpha, \sigma \in \mathbb{R}^2} (p - c)(\alpha - \sigma p + \varepsilon) f(\alpha, \sigma | \theta, \mathcal{Q}_{-1}) q(\varepsilon; \sigma_\varepsilon^2) d(\alpha, \sigma) d\varepsilon$$

with $q(\cdot; \sigma_\varepsilon^2)$ being a normal distribution with a mean of zero and a variance of σ_ε^2 . Then, a firm's Bellman equation can be written as:

$$V(\theta) = \max_{p \in \mathcal{P}} \left\{ \Pi(p; \theta) + \beta \int_{\varepsilon \in \mathbb{R}} \int_{\alpha, \sigma \in \mathbb{R}^2} V(B(\theta, p, \alpha - \sigma p + \varepsilon)) f(\alpha, \sigma | \theta, \mathcal{Q}_{-1}) q(\varepsilon; \sigma_\varepsilon^2) d(\alpha, \sigma) d\varepsilon \right\}.$$

Under this setup, a firm chooses its optimal price by trading off two forces. To maximize its current profits, a firm chooses a price that maximizes myopic profits $\Pi(p; \theta)$. For a given prior $\theta = (a, s, \text{vec}(\Sigma))'$, we can derive that the optimal myopic price equals:

$$p^{\text{my}}(\theta) = \arg \max_{p \in \mathcal{P}} \Pi(p; \theta)$$

$$= \frac{a+sc}{2s}.$$

However, a firm's price will affect its sales. The observed amount of sales in the future serves as an useful signal for the firm to update its prior beliefs. A firm internalizes this signal and thus needs to take into account how its price will affect its posterior beliefs. This consideration is also known as the trade-off between current control and estimation. More importantly, these incentives do not necessarily align with each other. The reasoning is as follows: for moderate beliefs, a firm prefers to choose a price that is not too extreme in order to maximize the myopic profits.⁶⁷ However, large deviations in a firm's price are more likely to result in large deviations in a firm's future sales which in turn means that its signals become more volatile and are thus more informative. In the end, a firm needs to strike a balance between maximizing strictly concave myopic profits and a convex continuation value.

To show this balance, we work out a numerical two-period version of the above framework. We denote $\theta_t = (a_t, s_t, v_{a,t}, v_{as,t}, v_{s,t})$. There are only two periods, thus in the second and last period, we must have:

$$\begin{aligned} V_2(\theta_2) &= \max_{p \in \mathcal{P}} \Pi(p; \theta_2) \\ &= \max_{p \in \mathcal{P}} (p-c)(a_2 - s_2 p) \\ &= (p^{\text{my}}(\theta_2) - c)(a_2 - s_2 p^{\text{my}}(\theta_2)) \\ &= \frac{(a_2 - s_2 c)(3a_2 - s_2 c)}{4s_2}. \end{aligned}$$

By backward induction, we obtain:

$$\begin{aligned} V_1(\theta_1) &= \max_{p \in \mathcal{P}} (p-c)(a_1 - s_1 p) + \beta \int_{\varepsilon \in \mathbb{R}} \int_{\alpha, \beta \in \mathbb{R}^2} \frac{(a_2 - s_2 c)(3a_2 - s_2 c)}{4s_2} f(\alpha, \sigma | \theta_1, q_1) g(\varepsilon; \sigma_\varepsilon^2) d(\alpha, \sigma) d\varepsilon \\ \text{s.t. } a_2 &= a_1 + \frac{(v_{a,1} + v_{as,1} p)(\alpha - \sigma p + \varepsilon - a - sp)}{v_{a,1} + 2v_{as,1} p + v_{s,1} p^2 + \sigma_\varepsilon^2} \\ s_2 &= s_1 + \frac{(v_{as,1} + v_{s,1} p)(\alpha - \sigma p + \varepsilon - a - sp)}{v_{a,1} + 2v_{as,1} p + v_{s,1} p^2 + \sigma_\varepsilon^2}. \end{aligned}$$

In the following example, we initialize the prior through $\theta_1 = (5, -1.2, 0.5, -0.1, 2)'$ and normalize $\sigma_\varepsilon^2 = 1$. The graph below reflects the reasoning we just described.

We only consider prices for which quantities are non-negative in expectation (with respect to the demand shock ε and the prior distribution), thus we define $\mathcal{P} = [0, \frac{a_1}{s_1}]$ in this example. The blue, dashed myopic profits are concave as expected. If the firm ignores its incentives for estimation (i.e. changing its price to affect its posterior beliefs), then it is optimal to set $p^{\text{my}}(\theta_1) = 2.58$. However, setting a more extreme price delivers a more informative signal in the second period. This is reflected in the convex shape of the continuation value (depicted in green). In the end, a rational firm balances the trade-off between control and estimation. As a result, it maximizes the sum of myopic profits and its continuation value, which is depicted in the red line. The maximum of this function is obtained at $p^*(\theta_1) = 3.5$. By definition, the firm engages in active learning as $p^{\text{my}}(\theta_1) \neq p^*(\theta_1)$.

Note that the mechanics of active learning in this example with a continuum of types is identical to the two-period example we illustrated in Section 3.1. In the baseline setup, a firm also faces the trade-off between current control and estimation through concave myopic profits and a strictly convex continuation value. As a result, our results on the propagation of nominal shocks in Section 5 should be robust to a more complicated version of active learning.

In fact, the incentives for active learning are stronger under a setup with a continuum of types. To understand this argument, we rely on the insights of Kiefer and Nyarko (1989). They show that under a setup with a linear demand curve all limiting beliefs and policy pairs $(\bar{\theta}, \bar{p})$ must satisfy a set of three properties that we outline below:

$$\bar{\theta} = B(\bar{\theta}, \bar{p}, \alpha - \sigma \bar{p} + \varepsilon) \quad (\text{B1})$$

$$\Pi(\bar{p}, \bar{\theta}) = \max_{p \in \mathcal{P}} \Pi(p; \bar{\theta}) \quad (\text{B2})$$

$$\mathbb{E}(\alpha | \bar{\theta}) - \mathbb{E}(\sigma | \bar{\theta}) \bar{p} = \alpha - \sigma \bar{p}. \quad (\text{B3})$$

Equation B1 is also known as belief invariance and follows directly from the definition of a limiting belief. In the limit (if one exist), beliefs converge to a constant vector that is defined as the fixed point of the function B conditional on \bar{p} . If

67. Further, this price must exist and is unique since the profit function is strictly concave in p for every pair (α, σ) .

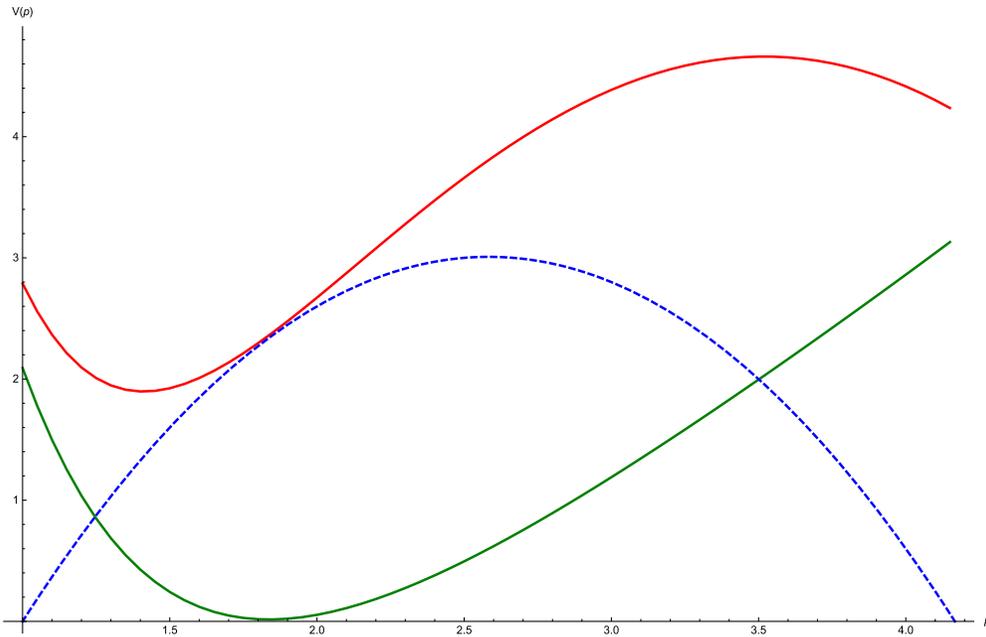


FIGURE C4

Numerical example of the two-period model: continuum of types

Notes: The figure shows the static profits (blue), the continuation value (green), and the total payoff (red) of the two-period model with a continuum of types. The y-axis represents the total payoff whereas the x-axis displays the price.

beliefs do not change in the limit, then there are no incentives to actively learn. As a result, the optimal policy must be the myopic one conditional on the limiting beliefs $\bar{\theta}$ as described in equation B2. Kiefer and Nyarko (1989) refer to this policy as one-period optimization. Further, if prices are forever held at \bar{p} , then a firm will at least infer the true amount of sales associated at the price \bar{p} . Equation B3 is also known as the mean prediction property.

The solution $(\bar{\theta}, \bar{p})$ that satisfies B1, B2, and B3 contains the correct limit belief but is in general not unique. Wieland (2000a) shows that any solution that contains incorrect limit beliefs must satisfy the following three properties:

- Perfect correlation.** $\frac{\bar{v}_{as}^2}{\bar{v}_a \bar{v}_s} = 1$.
- Uncertainty.** $\bar{v}_a, \bar{v}_s > 0$.
- Limit actions.** $\bar{p} = -\frac{\bar{v}_{as}}{\bar{v}_s} = -\frac{\bar{v}_a}{\bar{v}_{as}}$.

As a result, there is a set of incorrect, confounding beliefs under the continuum of types case. Recall from Section 3 that a firm does not learn anything under the confounding belief and thus avoids setting prices that are equal to the confounding price. This is reflected by the discontinuity in policy function under the extreme active learning regime. Under a continuum of types, there are a multitude of such points. Thus, firms vary their prices more due to active learning in this case.

Another advantage of restricting our attention to active learning in which there are only two types, (μ_1, σ_1) and (μ_2, σ_2) with $\sigma_2 > \sigma_1$, is that equations B1, B2, and B3 can be used to show that there exists only one limit belief that does not converge to the truth (i.e. $\lambda \notin \{0, 1\}$). This incorrect limit belief is equal to:

$$\bar{\lambda} = \frac{\sigma_2 \Delta \sigma c - \mu_2 (\sigma_1 + \sigma_2) + 2 \mu_1 \sigma_2}{\Delta \sigma (\Delta \sigma c - \Delta \mu)}$$

where c denotes a firm’s marginal cost of production.

C.5. Age-dependent exit rates

In our baseline framework, the exit rate is fixed at $\delta > 0$, which applies for each product. However, just like for firms (Caves, 1998), younger products are more likely to exit the market. Our assumption of a constant exit rate that is independent of the product’s age could potentially bias our results on the propagation of nominal shocks. This is because the composition of products is biased towards younger products that experience a higher frequency and absolute size of price adjustments.

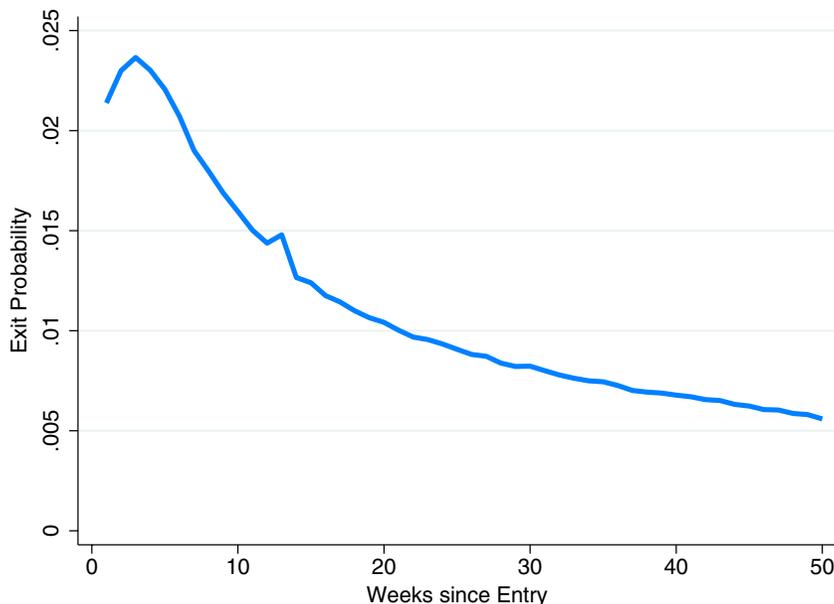


FIGURE C5
Exit probability by age

Notes: The graph plots the average exit probability of a UPC-store pair as a function of its age. We first compute the probability of exit for each category in the IRI Marketing data set. We then aggregate across categories using equal weights. The y-axis denotes the average exit probability in a given week whereas the x-axis denotes the number of weeks the product has been observed in the data since it entered the market.

In this section, we show that the assumption of a constant exit rate does not significantly bias the results generated by the baseline framework.

First, we show in the IRI Marketing data that product-level exit declines in age. Second, we propose an extension of our baseline framework to incorporate this observation. Third, we calibrate the extended framework and recalculate the real effects of a nominal shock.

In the first exercise, we compute the fraction of products that exit the market by age for each year and product category. Then, for each age bin, we average across years and product categories and plot these average exit rates by product age. The result can be found in the figure below that shows that exit rates indeed slope downward with age.

An alternative way of showing this fact is by estimating the product hazard function. There are multiple ways of doing this estimation. We allow for a high degree of flexibility in the hazard rate by estimating it parametrically through the Weibull distribution. The hazard function is then given by $h(a) = \lambda p a^{p-1}$ where the scale parameter is denoted by λ . The shape parameter p indicates whether the hazard rate varies with age. A value $p < 1$ means that the exit rates decline with a product's age whereas $p = 1$ and $\lambda = \delta$ corresponds to the exponential hazard function that we assume in our baseline framework. Thus, the Weibull distribution is flexible in that it allows for age-varying hazard rates. Another advantage of the Weibull specification is that it is straightforward to calibrate. Let T be a random variable that denotes the product's duration, then the following equalities hold whenever $T \sim WEI(\lambda, p)$:

$$\mathbb{E}(T) = \lambda^{-1/p} \Gamma(1 + p^{-1})$$

$$V(T) = \lambda^{-2/p} \left[\Gamma(1 + 2/p) - \Gamma(1 + p^{-1})^2 \right],$$

where $\Gamma(\cdot)$ denotes the gamma function. The IRI data shows that the average amount of weeks that a product lasts in the market is 81.25 weeks. Furthermore, its variance is given by 8,326.5. Then, we can form a system of two equations in the pair of unknowns (λ, p) . Its solution is given by $\hat{\lambda} = 48.13$ and $\hat{p} = 0.8922 < 1$. Note that our calibrated value for p is not too far away from unity. This value for p implies that product exit rates do not depend too strongly on age.

Note that the previous method relies on the structure of the Weibull distribution. Alternatively, we perform a non-parametric exercise. It is fairly difficult to obtain hazard functions non-parametrically, but we can still infer whether product-level exit rates depend on age in a non-parametric fashion. Recall that survival and hazard functions are related through the following identity:

$$-\log(S(a)) = \int_0^a h(\tau) d\tau.$$

A concave, increasing cumulative hazard function then indicates that hazard rates decline with the product's age. This is useful as it is straightforward to obtain the survival function non-parametrically through the Kaplan–Meier estimator $\widehat{S}(a)$.

To capture age-dependent exit rates in our framework, we extend the baseline model of Section 3 by assuming that product-level exit rates depend on age as follows:

$$\delta(a) = \delta_0 \exp(-\delta_1 a).$$

The parameters δ_0 and δ_1 can then be chosen accordingly to match our observations from the above graph. This estimation can be done by running a linear regression of average exit rates (in natural logs) on age. The estimated intercept and slope of this regression then correspond to $\widehat{\delta}_0$ and $\widehat{\delta}_1$. We could also match the observed (cumulative) hazard rate function. In this case, we would add the two parameters δ_0 and δ_1 to our calibration.

REFERENCES

- AGHION, P., BOLTON, P., HARRIS, C. and JULLIEN, B. (1991), “Optimal Learning by Experimentation”, *Review of Economic Studies*, **58**, 621–654.
- AGUIAR, M. and HURST, E. (2013), “Deconstructing Life Cycle Expenditure”, *Journal of Political Economy*, **121**, 437–492.
- ALVAREZ, E., LIPPI, F. and PACIELLO, L. (2011), “Optimal Price Setting With Observation and Menu Costs”, *Quarterly Journal of Economics*, **126**, 1909–1960.
- ALVAREZ, F. and LIPPI, F. (2014), “Price Setting With Menu Cost for Multiproduct Firms”, *Econometrica*, **82**, 89–135.
- _____, and _____, (2020), “Temporary Price Changes, Inflation Regimes, and the Propagation of Monetary Shocks”, *American Economic Journal: Macroeconomics*, **12**, 104–152.
- _____, LE BIHAN, B. and LIPPI, F. (2016a), “The Real Effects of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach”, *American Economic Review*, **106**, 2817–2851.
- _____, LIPPI, F. and OSKOLKOV, A. (2020), “The Macroeconomics of Sticky Prices with Generalized Hazard Functions” (Working Paper No. w27434, NBER).
- _____, _____, and PASSADORE, J. (2016b), “Are State-and Time-Dependent Models Really Different?”, *NBER Macroeconomics Annual*, **31**, 379–457.
- _____, BOROVICKOVÁ, K. and SHIMER, R. (2016c), “The Proportional Hazard Model: Estimation and Testing using Price Change and Labor Market Data” (Working Paper No. 22188, NBER).
- _____, _____, and _____, (2021), “Consistent Evidence on Duration Dependence of Price Changes”.
- ANDERSON, E., MALIN, B. A., NAKAMURA, E., SIMESTER, D. and STEINSSON, J. (2017), “Informational Rigidities and the Stickiness of Temporary Sales”, *Journal of Monetary Economics*, **90**, 64–83.
- ARGENTE, D., FITZGERALD, D., MOREIRA, S. and PRIOLO, A. (2021), “How Do Firms Build Market Share?”.
- _____, LEE, M. and MOREIRA, S. (2018), “Innovation and Product Reallocation in the Great Recession”, *Journal of Monetary Economics*, **93**, 1–20.
- BACHMANN, R. and MOSCARINI, G. (2012), “Business Cycles and Endogenous Uncertainty”.
- BALEY, I. and BLANCO, A. (2022), “Aggregate Dynamics in Lumpy Economies”, *Econometrica*, forthcoming.
- _____, and BLANCO, J. A. (2019), “Firm Uncertainty Cycles and the Propagation of Nominal Shocks”, *American Economic Journal: Macroeconomics*, **11**, 276–337.
- BARRO, R. J. (1972), “A Theory of Monopolistic Price Adjustment”, *Review of Economic Studies*, **39**, 17–26.
- BRODA, C. and WEINSTEIN, D. E. (2010), “Product Creation and Destruction: Evidence and Price Implications”, *American Economic Review*, **100**, 691–723.
- BRONNENBERG, B. J., KRUGER, M. W. and MELA, C. F. (2008), “Database Paper—The IRI Marketing Data Set”, *Marketing Science*, **27**, 745–748.
- CALVO, G. A. (1983), “Staggered Prices in a Utility-Maximizing Framework”, *Journal of Monetary Economics*, **12**, 383–398.
- CAPLIN, A. S. and SPULBER, D. F. (1987), “Menu Costs and the Neutrality of Money”, *Quarterly Journal of Economics*, **102**, 703–725.
- CARVALHO, C. and SCHWARTZMAN, F. (2015), “Selection and Monetary Non-Neutrality in Time-Dependent Pricing Models”, *Journal of Monetary Economics*, **76**, 141–156.
- CAVES, R. E. (1998), “Industrial Organization and New Findings on the Turnover and Mobility of Firms”, *Journal of Economic Literature*, **36**, 1947–1982.
- CHEVALIER, J. A., KASHYAP, A. K. and ROSSI, P. E. (2003), “Why Don't Prices Rise During Periods of Peak Demand? Evidence from Scanner Data”, *American Economic Review*, **93**, 15–37.
- _____, and _____, (2019), “Best Prices: Price Discrimination and Consumer Substitution”, *American Economic Journal: Economic Policy*, **11**, 126–159.
- COIBION, O., GORODNICHENKO, Y. and HONG, G. H. (2015), “The Cyclicalities of Sales, Regular and Effective Prices: Business Cycle and Policy Implications”, *American Economic Review*, **105**, 993–1029.

- DEATON, A. and PAXSON, C. (1994), “Intertemporal Choice and Inequality”, *Journal of Political Economy*, **102**, 437–467.
- DELLAVIGNA, S. and GENTZKOW, M. (2019), “Uniform Pricing in U.S. Retail Chains”, *Quarterly Journal of Economics*, **134**, 2011–2084.
- DE NARDI, M., FRENCH, E. and JONES, J. B. (2010), “Why Do the Elderly Save? The Role of Medical Expenses”, *Journal of Political Economy*, **118**, 39–75.
- DIXIT, A. (1991), “Analytical Approximations in Models of Hysteresis”, *Review of Economic Studies*, **58**, 141–151.
- DORASZELSKI, U., LEWIS, G. and PAKES, A. (2018), “Just Starting Out: Learning and Equilibrium in a New Market”, *American Economic Review*, **108**, 565–615.
- DROZD, L. and NOSAL, J. (2012), “Understanding International Prices: Customers as Capital”, *American Economic Review*, **102**, 364–395.
- EICHENBAUM, M., JAIMOVICH, N. and REBELO, S. (2011), “Reference Prices, Costs, and Nominal Rigidities”, *American Economic Review*, **101**, 234–262.
- FERNALD, J. G. (2014), “A Quarterly, Utilization-Adjusted Series on Total Factor Productivity” (San Francisco: Federal Reserve Bank of San Francisco).
- FITZGERALD, D., HALLER, S. and YEDID-LEVI, Y. (2016), “How Exporters Grow” (Working Paper No. w21935, NBER).
- FOSTER, L., HALTIWANGER, J. and SYVERSON, C. (2016), “The Slow Growth of New Plants: Learning about Demand?”, *Economica*, **83**, 91–129.
- GAGNON, E. and LÓPEZ-SALIDO, D. (2020), “Small Price Responses to Large Demand Shocks”, *Journal of the European Economic Association*, **18**, 792–828.
- GAUR, V. and FISHER, M. L. (2005), “In-Store Experiments to Determine the Impact of Price on Sales”, *Production and Operations Management*, **14**, 377–387.
- GILBERT, L. (2017), “Gains from Product Variety and the Local Business Cycle”, 2018 Meeting Papers 46, Society for Economic Dynamics.
- GILCHRIST, S., SCHOENLE, R., SIM, J. and ZAKRAJŠEK, E. (2017), “Inflation Dynamics during the Financial Crisis”, *American Economic Review*, **107**, 785–823.
- GOLOSOV, M. and LUCAS, R. E. (2007), “Menu Costs and Phillips Curves”, *Journal of Political Economy*, **115**, 171–199.
- GOURINCHAS, P.-O. and PARKER, J. A. (2002), “Consumption Over the Life Cycle”, *Econometrica*, **70**, 47–89.
- GOURIO, F. and RUDANKO, L. (2014), “Customer Capital”, *Review of Economic Studies*, **81**, 1102–1136.
- HECKMAN, J. and ROBB, R. (1985), “Using Longitudinal Data to Estimate Age, Period and Cohort Effects in Earnings Equations”, in Mason, W. M. and Fienberg, S. E. (eds) *Cohort Analysis in Social Research*, chapter 5. (Springer) 137–150.
- HITSCH, G., HORTACSU, A. and LIN, X. (2019), “Prices and Promotions in U.S. Retail Markets: Evidence from Big Data” (Technical Report).
- ILUT, C., VALCHEV, R. and VINCENT, N. (2020), “Paralyzed by Fear: Rigid and Discrete Pricing under Demand Uncertainty”, *Econometrica*, **88**, 1899–1938.
- JUSTINIANO, A., PRIMICERI, G. and TAMBALOTTI, A. (2010), “Investment Shocks and Business Cycles”, *Journal of Monetary Economics*, **57**, 132–145.
- _____, _____, and _____, (2013), “Is There a Trade-Off Between Inflation and Output Stabilization?”, *American Economic Journal: Macroeconomics*, **5**, 1–31.
- KARADI, P. and REIFF, A. (2019), “Menu Costs, Aggregate Fluctuations, and Large Shocks”, *American Economic Journal: Macroeconomics*, **11**, 111–146.
- KELLER, G. and RADY, S. (1999), “Optimal Experimentation in a Changing Environment”, *Review of Economic Studies*, **66**, 475–507.
- KIEFER, N. M. and NYARKO, Y. (1989), “Optimal Control of an Unknown Linear Process with Learning”, *International Economic Review*, **30**, 571–586.
- KILEY, M. T. (2002), “Partial Adjustment and Staggered Price Setting”, *Journal of Money, Credit, and Banking*, **34**, 283–298.
- KLESHCHELSKI, I. and VINCENT, N. (2009), “Market share and price rigidity”, *Journal of Monetary Economics*, **56**, 344–352.
- KRUSELL, P. and SMITH, A. A. Jr. (1998), “Income and Wealth Heterogeneity in the Macroeconomy”, *Journal of Political Economy*, **106**, 867–896.
- LEE, Y. and MUKOYAMA, T. (2018), “A Model of Entry, Exit, and Plant-level Dynamics Over the Business Cycle”, *Journal of Economic Dynamics and Control*, **96**, 1–25.
- LUCAS, R. E. (2003), “Macroeconomic Priorities”, *American Economic Review*, **93**, 1–14.
- MIDRIGAN, V. (2011), “Menu Costs, Multiproduct Firms, and Aggregate Fluctuations”, *Econometrica*, **79**, 1139–1180.
- MIRMAN, L. J., SAMUELSON, L. and URBANO, A. (1993), “Monopoly Experimentation”, *International Economic Review*, **34**, 549–563.
- MOREIRA, S. (2016), “Firm Dynamics, Persistent Effects of Entry Conditions, and Business Cycles” Persistent Effects of Entry Conditions, and Business Cycles (October 1, 2016).
- NAKAMURA, E. and STEINSSON, J. (2008), “Five Facts about Prices: A Reevaluation of Menu Cost Models”, *Quarterly Journal of Economics*, **123**, 1415–1464.

- _____ and _____, (2009), “Monetary Non-Neutrality in a Multi-sector Menu Cost Model”, *Quarterly Journal of Economics*, **125**, 961–1013.
- _____ and _____, (2011), “Price Setting in Forward-Looking Customer Markets”, *Journal of Monetary Economics*, **58**, 220–233.
- PACIELLO, L., POZZI, A. and TRACHTER, N. (2019), “Price Dynamics with Customer Markets”, *International Economic Review*, **60**, 413–446.
- PRESCOTT, E. C. (1972), “The Multi-Period Control Problem Under Uncertainty”, *Econometrica*, **40**, 1043–1058.
- RAVN, M., SCHMITT-GROHÉ, S. and URIBE, M. (2006), “Deep Habits”, *Review of Economic Studies*, **73**, 195–218.
- ROLDAN, P. and GILBUKH, S. (2021), “Firm dynamics and pricing under customer capital accumulation”, *Journal of Monetary Economics*, **118**, 99–119.
- ROTHSCHILD, M. (1974), “A Two-Armed Bandit Theory of Market Pricing”, *Journal of Economic Theory*, **9**, 185–202.
- SHAPIRO, M. D. and WATSON, M. W. (1988), “Sources of Business Cycle Fluctuations”, *NBER Macroeconomics Annual*, **3**, 111–148.
- SHEEDY, K. D. (2010), “Intrinsic Inflation Persistence”, *Journal of Monetary Economics*, **57**, 1049–1061.
- SMETS, F. and WOUTERS, R. (2007), “Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach”, *American Economic Review*, **97**, 586–606.
- STROEBEL, J. and VAVRA, J. (2019), “House Prices, Local Demand, and Retail Prices”, *Journal of Political Economy*, 2019, **127**, 1391–1436.
- TAUCHEN, G. (1986), “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions”, *Economic Letters*, **20**, 177–181.
- VAVRA, J. (2014), “Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation”, *Quarterly Journal of Economics*, **129**, 215–258.
- WIELAND, V. (2000), “Learning by Doing and the Value of Optimal Experimentation”, *Journal of Economic Dynamics and Control*, **24**, 501–534.
- _____ (2000), “Monetary Policy, Parameter Uncertainty and Optimal Learning”, *Journal of Monetary Economics*, **46**, 199–228.
- WILLEMS, T. (2017), “Actively Learning by Pricing: A Model of an Experimenting Seller”, *Economic Journal*, **127**, 2216–2239.
- ZBARACKI, M., RITSON, M., LEVY, D., DUTTA, S. and BERGEN, M. (2004), “Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets”, *Review of Economics and Statistics*, **86**, 514–533.
- ZELLNER, A. (1971), *An Introduction to Bayesian Inference in Econometrics* (New York: Wiley).