THE UNIVERSITY OF CHICAGO

ESSAYS ON APPLIED MACROECONOMICS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE SOCIAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

KENNETH C. GRIFFIN DEPARTMENT OF ECONOMICS

BY

DAVID ORLANDO ARGENTE AMAYA

CHICAGO, ILLINOIS

JUNE 2018

A mi familia.

En la calle codo a codo somos mucho más que cinco...

El que es de aquí.
El que acaba de llegar
y ya es de aquí.
El que dice “ciudad” por decir
tú y yo y Pedro y Marta
y Francisco y Guadalupe.
El que lleva dos días sin luz
ni agua.
El que todavía respira.
El que levantó un puño
para pedir silencio.
Los que le hicieron caso.
Los que levantaron el puño.
Los que levantaron el puño
para escuchar
si alguien vivía.
Los que levantaron el puño para
escuchar si alguien
vivía y oyeron
un murmullo.
Los que no dejan de escuchar.

-Juan Villoro

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I would like to thank my advisor Fernando Alvarez for never giving up on me and for teaching me much more than just economics. I would also like to thank the rest of my committee – Robert Shimer, Erik Hurst, and Joseph Vavra – for pushing me hard and, at the same time, always finding time to meet me to discuss ideas that go much beyond what is included in this dissertation. You gave me all the support and encouragement I needed to complete this work, all the many imperfections it may have are my own. I thank Francesco Lippi for taking me under his wing and Robert Lucas Jr., Gary Becker, and Nancy Stokey for all the great conversations and advice they gave me. I also thank Gaby Silva for adopting me and taking care of me.

This dissertation is dedicated to my family without whom I would have never finished this adventure. To my father, David, for being my reference for courage and compassion. To my mother, Aida, the best person I know and who lived with me every step of this process. I would have never done it without you; all my accomplishments will always be because of you and for you. To Haydée, my sister, for always bringing me inspiration in the most difficult days. To my brother, Héctor, for believing in me even when I myself doubted. I hope I never disappoint you. And to my grandparents Arturo, Ursula, Paz, and Erlinda. Your story has built my character and your constant sacrifices brought me here. You will always have the most special place in my mind and my heart.

I thank Jorge, Bong Geun, Munseob, Chen, and Sara for never leaving me behind and my friends Elisa, Nuno, Claudia, Kerstin, Roberto, Big, Yi, Cristian, Aaron, Holly, Ursula, Pallavi, Juan, and Nick for turning the difficult path of obtaining this degree into something that I will greatly miss. Lastly, I would like to thank the University of Chicago and the Science and Technology Council of Mexico for giving me the chance and support to pursue my dreams.

ABSTRACT

This dissertation studies macroeconomic questions in the fields of monetary economics, economic growth, and income inequality through the lens of detailed micro data sets. In the first chapter of this dissertation (joint with Chen Yeh), we study the role of product entry and exit in propagating nominal shocks to the real economy. Toward that goal, we show that product turnover has an extensive role in the aggregate economy and that the frequency and size of price adjustments are negatively related to a product's age. We exploit information from product-level characteristics and the timing of products' launches to provide empirical support that these stylized facts can be rationalized by an active learning motive: firms with new products are faced with demand uncertainty and can optimally obtain valuable information by varying their prices. Building on the empirical findings, we construct a menu cost model with active learning and quantify the importance of age-dependent pricing moments for the propagation of nominal shocks. In the calibrated version of our model, the cumulative response of output to a nominal shock more than doubles compared to the standard menu cost model and this response is higher during economic booms.

In the second chapter (joint with Munseob Lee and Sara Moreira), we use detailed product- and firm-level data to study the sources of innovation and the patterns of productivity growth over the period from 2007 to 2013. We document several new facts on product reallocation. First, every quarter around 8 percent of products are reallocated in the economy, and the entry and exit of products are prevalent among different types of firms. Second, most reallocation of products occurs within the boundaries of the firm. The entries and exits of firms only make a small contribution in the overall creation and destruction of products. Third, product reallocation is strongly pro-cyclical and declined by more than 25 percent during the Great Recession. This cyclical pattern is almost entirely explained by a decline in within firm reallocation. Motivated by these facts, we study the causes and consequences of reallocation within incumbent firms. As predicted by Schumpeterian growth theories, the rate of product reallocation strongly depends on the innovation efforts of the firms and has important implications for revenue growth, improvements in products' quality, and productivity dynamics. Our estimates suggest that the decline in product reallocation through

these margins has contributed greatly to the slow growth experienced after the Great Recession.

Lastly, in the third and last chapter of this dissertation (joint with Munseob Lee), we construct income-specific price indices for the period from 2004 to 2010. We find substantial differences across income groups that arise during the Great Recession. The annual inflation in the cost of living of the highest quartile is on average 0.59 percentage points lower than that of the lowest quartile of the income distribution. We find that product quality substitution, a margin mostly available to richer households, is the main mechanism explaining the differences. Our evidence shows that not accounting for these differences could lead to significant biases in the calculation of inequality and poverty measures.

CHAPTER 1

PRODUCT LIFE CYCLE, LEARNING, AND NOMINAL SHOCKS (WITH CHEN YEH)

1.1 Introduction

A recent surge of research on product innovation has shown that product creation and destruction is extensive in the U.S. and has emphasized its importance for macroeconomic performance (e.g. Broda and Weinstein, 2010). In particular, new products are prevalent. Argente, Lee, and Moreira (2017a) find that one in three products are either created or destroyed in a given year and that more than 20 percent of the products are less than one year old. While turnover rates are clearly high, whether product creation and destruction are quantitatively important for the propagation of nominal shocks to the real economy is not clear.[1]

In this paper, we study the real effects of a nominal stimulus in an economy featuring high rates of product turnover and age-dependent pricing moments. To do so, 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 which price-setting models ignore. We show that pricing moments at the product level strongly depend on their age: entering products change their prices twice as often as average products and the average size of these adjustments is 50 percent larger compared to the average price change.[2]

We then explore the underlying reason why new products display such different pricing dynamics. The empirical evidence from a select number of industries shows that the introduction of a new product comes with a significant amount of demand uncertainty from the firm's perspective. These firms vary their prices to obtain valuable information about their demand curves.[3] We

1. Nominal shocks, as discussed in Lucas (2003), can affect real variables and relative prices in the short run but not in the long run. Shapiro and Watson (1988) refers to these shocks as "demand" shocks.

2. The patterns at exit are quite different as the frequency and absolute size of the price adjustments stay mostly constant before exit at the product level. Nonetheless, the frequency and depth of sales increase significantly at exit which indicates that firms attempt to liquidate their inventories before phasing out their products. These findings are described in more detail in the appendix.

3. Gaur and Fisher (2005), for example, surveys 32 US retailers and finds that 90 percent conduct price experiments to learn about their demand.

provide further evidence of this and extend the set of previous empirical findings in the literature to a broad line of product categories. Exploiting variation in the timing of product launch within retailers, we find that retailers carry forward information obtained during the first launch to any subsequent launch of the same product at a different location. As a result, retailers adjust prices less frequently and less aggressively after the first launch of their product. Furthermore, 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 unique *newness* index that measures the novelty of a product. We indeed find that more novel, entering products adjust their prices more often and by larger amounts.

Our set of stylized facts is the first in the price-setting literature that focuses on age-dependent pricing moments. While previous contributions have found heterogeneity in pricing moments along other dimensions (e.g. Nakamura and Steinsson, 2009), none of them has focused on a product's life cycle. Our empirical findings provide insights on the underlying reason why firms adjust prices. This is relevant because the theoretical literature on price-setting shows that different types of price changes have substantially different macroeconomic implications.[4] We argue that the age distribution of products at a given point in time is important for the propagation of a nominal shock to real output.

To support our argument, we build a quantitative menu cost model in the spirit of Golosov and Lucas (2007). The model encompasses demand uncertainty and active learning similar to that in Bachmann and Moscarini (2012).[5] Whenever a firm's product enters the market, it is unaware about its elasticity of demand. However, the firm does know that its elasticity is constant, and that it is either high or low. 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 whether the realization of its demand shock was simply low. To deal with this uncertainty, firms engage in a Bayesian

4. Especially if the timing of the price change is endogenous (e.g. Caplin and Spulber (1987)) or exogenous (e.g. Calvo (1983)).

5. Their focus is very different as they study how negative first moment aggregate shocks induce risky behavior. 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." However, we observe no evidence for differential pricing moments at a product's exit in the data.

learning process and form beliefs about their elasticity.[6] 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 deviating in order to affect its posterior beliefs in the most efficient way so that it acquires more valuable information about its type. This is also known as the trade-off between current control and estimation. As a firm ages and learns more about its elasticity, 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 in the UK. Doraszelski, Lewis, and Pakes (2016) find that in response to uncertainty, firms experiment with their bids by adjusting them more frequently and by larger amounts. As firms acquires more information about demand, the adjustment of the bids became less frequent and smaller.

Then, we calibrate this framework to standard pricing moments and our newly found pricing moments that vary over a product's life cycle, and quantify the cumulative real output effect of a nominal spending shock. Our findings indicate that the output response is 2.3 times larger than the benchmark model with no demand uncertainty. The reasoning behind this result is twofold. First, the learning incentives dampen the selection effect; 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.[7] Since firms have an additional motive to change prices, they are less sensitive to changes in their costs.[8]

Second, the concept of a product's life cycle introduces an additional form of cross-sectional

6. Our work is related to the literature in optimal control problems with active learning that have 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), Ilut, Valchev, and Vincent (2015) and Willems (2016).

7. 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), Karadi and Reiff (2016) and Vavra (2014).

8. Although this logic is similar to the one described in Alvarez, Bihan, and Lippi (2014) 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 firms might optimally choose to deviate from the static profit-maximizing price in absence of a menu cost in order to learn about their demand elasticity.

heterogeneity in the adjustment frequency. 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 price changes by more uncertain firms that occur 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 the 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 lower frequency of price adjustment on average. As a result, this delay reduces the selection in timing of price changes.[9]

The remainder of this paper is organized as follows: In Section 3.2, we present the data and the main empirical findings. In section 1.3, we set up a quantitative model of menu costs that is able to explain our stylized facts. In addition, we develop the relevant conditions for active learning and describe what learning regimes might occur. In Section 3.7, we discuss our results on the propagation of nominal shocks and compare our results with other models used in the literature. We also discuss the sensitivity of our results to matching other moments that vary over the product's life cycle such as sales profiles and exit rates. We also extend the model to include endogenous entry to quantify the response of output in periods of high and low product entry. Section 3.8 concludes. The appendix provides proofs to the propositions, additional empirical findings, and extensions of the model.

1.2 Stylized Facts of Product's Life Cycle

In this section, we use a large scanner data set to show a new set of stylized facts on 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 revenue relative to the aggregate. Then, we develop a set of facts that clearly show that pricing moments at the product level are considerably different across a product's age; in

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

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, is considerably larger at the beginning of the product's life cycle. We hypothesize that an active learning motive can rationalize these facts. The last part of this section provides additional evidence consistent with this interpretation by exploiting the variation in products' entry time across different stores and their characteristics across different product categories.

1.2.1 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 Dataset, 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 its first appearance in the national market.

For this reason, we rely on the IRI Marketing data set instead that provides more than ten years of data at the store-product-week level. The data is 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 contains approximately 2.4 billion transactions from January 2001 to December 2011 that represents roughly 15 percent 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 covers 31 product categories and includes detailed information about each product such as its brand, volume, color, flavor, and size.[10]

10. The product categories include Beer, Carbonated Beverages, Coffee, Cold Cereal, Deodorant, Diapers, Facial

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 very broad 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 flavor or 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 starts from January 2001, thus an entering product is one that has no observable transactions in any store across the US for at least one 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 reason. First, the prices of products born during downturns can have substantially different patterns than those of products born in normal times.[11] More importantly, IRI Symphony 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.[12] By restricting our sample of entering products between January 2002

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 dataset is discussed in more detail in Bronnenberg, Kruger, and Mela (2008). See also Alvarez, Bihan, and Lippi (2014), Chevalier and Kashyap (2014), Gagnon and Lopez-Salido (2014), Stroebel and Vavra (2014) and Coibion, Gorodnichenko, and Hong (2015a) for applications of the data to related questions.

11. Moreira (2015), for example, provides evidence that the average business size across cohorts is significantly affected by aggregate economic conditions at inception.

12. 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

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 two years in the market. We exclude private label items from the analysis because all private-label UPCs have the same brand identification so that the identity of the retailer cannot be recovered from the labeling information. We exclude short-lived products in order to minimize the problem that some UPCs might get discontinued only to have the same product appear with a new UPC as noted by Chevalier, Kashyap, and Rossi (2003). We also drop promotional items or products with very little revenue to minimize biases due to measurement error.

1.2.2 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. Argente, Lee, and Moreira (2017a) 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 percent of US products are aged less than one 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 new products, exiting products, and the total number of products in each category c to define aggregate entry and exit rates at the product level:

$$n(t,s) = \frac{\sum_i N_i(t,s)}{\sum_i T_i(t)}$$

$$x(t,s) = \frac{\sum_i X_i(t,s)}{\sum_i T_i(s)}$$

where $N_i(t,s)$, $X_i(t,s)$, and $T_i(t)$ are number of entrant products, exit products, and total products in period t relative to period s. 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

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.

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. Creation and destruction are the revenue weighted analogues of the entry and exit rates.

Using a scanner dataset 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$\times$Store pair. We find a substantial degree of entry of products at both levels.

Table 1.1 reports the entry and exit rates for the case in which t and s are one and five years apart. It shows that 15 percent of the UPCs in the market and on average 27 percent of the products in each store entered in the last year. Approximately 45 percent of the products in the market entered in the last five years accounting for 30 percent of total expenditures. At the store level, 66 percent of all products sold were first introduced by the store in the last five years and they account for more than half of the total revenue of the store. Although the exit rate is very similar to the entry rate of products, destruction is lower than creation. 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.

Panel A in table 1.2 shows that the median duration of a product in a given store is only slightly above three years.[13] The remarkably large rates of product creation and destruction, 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.

1.2.3 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:

13. Since our dataset 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.

$$P_{mcjst} = \frac{\texttt{sales}_{mcjst}}{\texttt{units}_{mcjst}}$$

where m, c, j, s, and t index markets (at the MSA level) are product categories, UPCs, stores and time respectively. A considerable advantage of the IRI Symphony 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 in our data do not vary with macroeconomic conditions.[14] Nonetheless, the main stylized facts discussed below are robust to the inclusion of sales in the analysis.

We adopt the same conventions as Coibion, Gorodnichenko, and Hong (2015a) 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 percent in absolute value. For prices larger than 5 dollars in value, this cut-off is 0.5 percent. 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).[15] The size of a price change is calculated as the log difference between the price levels in the current and the previous week. Thus, we get:

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

Let $a = 1, \ldots, A$ denote the number of weeks since entry (which we will define as the *age* of the

14. Using the same data set, Coibion, Gorodnichenko, and Hong (2015a) find that retailers' use of sales does not vary with the unemployment rate. Anderson, Malin, Nakamura, Simester, and Steinsson (2013) 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. In addition, we do not find evidence that retailers use sales to actively learn at entry as neither the frequency or the size of sales are larger. See figure A.15 in the online appendix for a more elaborate discussion on this issue.

15. Specifically, under this approach, a good is on sale if a price is reduced but returns to its same previous level within four weeks. Coibion, Gorodnichenko, and Hong (2015a) 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 dataset, we use the union of sales flags obtained 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.

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 empirical specification:

$$Y_{jsct} = \alpha + \sum_{a=1}^{A} \phi_a D_{js}^{a} + \theta_{js} + \tau_t + \gamma_c + \varepsilon_{jsct} \tag{1.1}$$

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 the 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 on 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 estimate equation 1.1 using two different normalizations. The first assumes that trends appear only in cohort effects. In this case, we replace the time fixed effects with the seasonally adjusted unemployment rate at the level of Metropolitan Statistical Area (MSA) to control for local cyclical economic variation.[16] The second assumes trends only appear in the period effects. In this case, the time fixed effects are included in the estimation of equation 1.1 and we use the local unemployment rate to represent the cohort fixed effects. Our baseline results use the latter approach, but our results are not sensitive to the chosen normalization.

1.2.4 Pricing Moments over the Life Cycle of US Products

We first use regression specification 1.1 with a price change indicator as the dependent variable. Figure 1.1 plots the frequency of regular price changes over the life cycle of a product. The dots represent the estimates of the age fixed effects $\{\phi_a\}_{a=1}^{A}$ computed with equation 1.1.[17] The newer

16. Some examples of studies that use this proxy approach are Deaton and Paxson (1994), Gourinchas and Parker (2002), Aguiar and Hurst (2013), and De Nardi, French, and Jones (2010).

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

products clearly have their prices changed more often. We 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.

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 percent. 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.[18] 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.

Our findings are consistent with those in Alvarez, Borovicková, and Shimer (2015). They test whether hazard rates of price changes depend on the age of the product once unobserved heterogeneity is taken into account. They find this is the case and argue that this statistical model is a reasonable representation of the data. Figure A.1 in the appendix decomposes these price changes into increases and decreases. We find similar results for both the frequency of price increases and decreases.

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 1.2 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 percent. Further, the dispersion of price changes as measured by the weekly cross-sectional standard deviation is almost 40 percent larger during the first four months after entry with respect to its level 12 months after the product is launched. Importantly, this fact holds for both price increases and decreases as shown in figures A.2 and A.3 in the appendix.[19]

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

19. Bachmann and Moscarini (2012) focus on the pricing moments of products at exit. In their work, at the end of the product life cycle, the returns to price experimentation increase as firms "gamble for resurrection". We found little support for this mechanism in the IRI Marketing data. Figure A.14 in the appendix shows that both the frequency and size of regular price changes stay mostly constant at exit. Nonetheless, figure A.14 shows that there is an increase in the frequency of sales during the last weeks of the life cycle of the product. We interpret this pattern as an increase in "clearance" sales.

This leads to our second stylized fact:

Empirical fact 2. The absolute size of price adjustment declines monotonically with 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. Our stylized facts remain robust whenever we use a linear specification in the age of a product. This is summarized in table 1.3. 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.

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, Bihan, and Lippi (2014) to minimize the issue of heterogeneity across products and stores. We define "cells" at the UPC-store level, say $(j, s) = i$, and standardize each price change at this level through $z_{it} = (\Delta p_{it} - \mu_i)/\sigma_i$ where μ_i and σ_i are the mean and standard deviation of price changes in cell i across time. Price changes equal to zero are disregarded.

Figure 1.3 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 to entry particularly during the first 20 weeks. About 40 percent 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.[20] Figure A.5 in the appendix shows that this occurs for price changes in both directions.[21] Our third stylized fact is then summarized by:

Empirical fact 3. Empirical fact 3: About 40 percent of large price changes occur during the first 20 weeks of the product's life cycle. A quarter of them are observed in the first three weeks of a product's entry.

20. This finding shows that idiosyncratic shocks from fat-tailed distributions need to be used with caution when used to generate large price changes. In standard menu cost models it is often assumed that these shocks arrive at a constant rate. This assumption includes the family of Poisson shocks used in Midrigan (2011) and Karadi and Reiff (2016) where the distribution of large price changes is independent of age. This is relevant as fat-tailed shocks have drastic implications on the degree of non-neutrality. Their presence reduces the selection effect after a monetary shock as the mass of firms responding to it is smaller.

21. Figure A.6 in the appendix shows that our results are not sensitive to our standardization since the fraction of non-standardized price changes larger than 30 percent shows the same pattern.

1.2.5 The Case for Active Learning

Our empirical findings so far have not shed any light on the economic mechanism that could rationalize them. In this section, we argue that active learning can rationalize our empirical findings, and we provide additional empirical evidence to support this claim. There is already a substantial body of literature that shows that firms actively use their prices to obtain information about their demand. This is widespread across the US economy according to empirical and anecdotal evidence from Gaur and Fisher (2005, US retailers), Einav, Kuchler, Levin, and Sundaresan (2015, online listings on eBay), and Campbell and Eden (2005, grocers).

Whenever firms face uncertainty over the demand of their product due to some aggregate (e.g., industry- or location-specific shocks) or idiosyncratic (e.g., consumer taste shocks) factors, they can gain information from large movements in a product's price. This is because these factors are less likely to be important whenever a firm's sales move due to large, deliberate variations in their price. At entry, a firm might not be aware of some of its demand characteristics. As a result, entering products should change their prices more often and by larger amounts.[22]

Further, active learning has two additional testable implications. First, if the same product is launched at different times or different locations by the same retailer, we should observe that the patterns of learning attenuate after the introduction. This is true when assuming demand characteristics are persistent across time and space. Under this scenario, retailers incorporate attained information from the introduction to later rounds in which the incentives for learning are dampened. Second, if firms are actively learning about their demand, then we should observe that active learning increases when the demand for the product being launched is more uncertain. This increase can occur if the product in question is more novel or innovative. The nature of the IRI

22. An alternative explanation could be penetration pricing. Under this scenario, firms increase long-run profits by launching a low-priced product to secure market share or a solid customer base. This tactic then results in higher future profits as the firm is able to benefit from the consumer's higher willingness to pay. In the appendix, we provide evidence that this strategy does not seem to be consistent with our empirical findings. We do not find evidence of lower entry prices (figure A.9) or rapid price increases at the beginning of the product's life cycle (figure A.8). In fact, figure A.7 shows that price increases and decreases are mostly balanced after entry. Conditional on having at least two price changes during the first 6 months, less than 15 percent of the products record only price increases and the probability of two consecutive price increases is less than half. To provide further evidence, in appendix A.6.3 we extend the standard price-setting model to include costumer base concerns and show that the implications of the model are not consistent with several of the stylized facts we have documented.

Marketing data allows us to test both of these conjectures.

Timing of Product Launch

The first conjecture consists of studying whether retailers carry forward any information obtained during the first launch to any subsequent launch of the same product at a different location. For this purpose, we take advantage of the variation in products' launch dates across different periods in time or locations. If information is held at the retailer level and they are uncertain about some demand characteristics, then retailers should adjust prices less frequently and less aggressively after the first launch of their product.[23]

To show this, 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 the completion of the first year since the UPC was first introduced in the national market. 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.

Figure 1.4 shows that after following the baseline regression specification 1.1 for each wave, 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. On average, at entry, the size of price changes is 7 percent larger than the mean for products in the first wave and only 4 percent larger for products in the second wave. The size of absolute price changes then converges back to the mean of its respective wave.

In the appendix, we show that these findings also hold if we condition on the fact that waves should occur across different cities and, importantly, these patterns occur for both price increases and decreases. These figures show that retailers obtain relevant information about the demand their product faces in the first wave. As a result, their incentives to learn at the time of entry in the second wave are then lower. Thus, there is less need for active learning, and we see that the age-dependence of products' pricing moments become attenuated.

23. We assume that markets across space are not completely independent. This assumption means that any information obtained on the local demand of some product in Chicago is at least somewhat informative for another city such as New York.

Newness Index

To strengthen our findings on learning, we confirm the hypothesis that the age-dependence of pricing moments is more pronounced for products that are more novel. Our definition of what constitutes a "new" product in our earlier empirical exercises is relatively broad. However, it is very well possible that product introductions vary significantly in their degree of novelty. A new product could be an entirely new brand, exist within an incumbent brand, or simply be an improvement or variation on an existing product (e.g., new color or flavor). In order to quantify the novelty of a product, we first compare the pricing behavior of retailers when launching a new brand relative to the rest of the products. The reasoning for this comparison is that brand extensions are more rare and constitute larger innovations relative to products already in the market. We do so using the following regression specification:

$$Y_{jsct} = \alpha + \gamma \text{age}_{jsct} + \phi \text{age}_{jsct} \times \text{NewBrand}_{jsct} + \theta_{js} + \tau_t + \gamma_c + \varepsilon_{jsct}. \tag{1.2}$$

where as before j, s, c, and t are the UPC, store, cohort and time period, respectively. The dependent variable Y_{jsct} either represents the price change indicator of the absolute size of price changes. NewBrand_{jsct} is an indicator that equals one if the product introduction is a new brand with a new volume. γ captures the heterogeneity of pricing moments in age, and ϕ measures the strength of this heterogeneity with respect to the novelty of the product's brand. Not surprisingly, the effect of age is negative and strongly significant as γ only summarizes the patterns described in section 1.2.4 with a linear function. More importantly, table 1.4 shows that the coefficient for the interaction term ϕ is also negative and significant and shows that newer products tend to adjust their prices more often and by larger amounts during the first six months of their life cycle.

In order to provide a more comprehensive measure of the novelty of each product, we construct a *newness index* that uses detailed information about the characteristics of each UPC provided in the IRI Marketing data set. The 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., novel).

We define a product j in category k as a vector of characteristics $V_{kj} = \left[v_{j1}, v_{j2}, .., v_{jN_k}\right]$ where N_k denotes the number of attributes we observe in category k in our data.[24] Then, if Ω_{kst} 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:

$$\mathrm{NI}^k_{jst} = \frac{1}{N_k} \sum_{i=1}^{N_k} \mathbb{1}[v_{ji} \notin \Omega_{kst}]. \tag{1.3}$$

For example, if a new product within the beer category enters with a flavor and a volume that has never been sold at the store before, its newness index is 2/9. We assume that each attribute is equally weighted in order to remain agnostic about the relative importance of each attribute to the degree of newness of a product.[25] To understand if stores launching more novel products are actively learning, we estimate:

$$Y^k_{jsct} = \alpha + \gamma \mathrm{age}_{jsct} + \phi \mathrm{age}_{jsct} \times \mathrm{NI}^k_{jst} + \theta_{js} + \tau_t + \gamma_c + \varepsilon_{jstc}. \tag{1.4}$$

where ϕ is our coefficient of interest because it summarizes the degree of heterogeneity in a product's age depending on its novelty. Table 1.5 shows that our index has substantial power to explain the price-setting patterns we observe at entry. The index confirms the second testable conjecture on learning by showing that more novel products change prices more actively at the beginning of their life cycle. This evidence provides additional support to the active learning hypothesis. The remainder of the paper takes these empirical facts as given and extends the standard price-setting

24. For example, the product category beer consists of $N_{beer} = 9$ attributes for each barcode: vendor, brand, volume, type (e.g., ale or lager), package (e.g., can or keg), flavor, size (e.g., bottle or six pack), calorie level (e.g., light or regular) and color.

25. 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 that might not describe a product in its entirety. On average, we observe ten product characteristics in each category.

models to include the concept of the product's life cycle through an active learning mechanism.

1.3 Quantitative Model

Our framework is a discrete time menu cost model in the tradition of Golosov and Lucas (2007) in which firms face uncertainty on their demand curves. A firm's type, that is, its elasticity of demand, is either high or low, but the firm does not know its type. As a result, a firm forms beliefs on its type. It can adjust its price to change the speed at which its beliefs get updated. Thus, a firm faces a trade-off between maximizing its static profits and gaining more information about its type when choosing its price. The active learning mechanism is based on Mirman, Samuelson, and Urbano (1993) and its implementation in general equilibrium is closely related to Bachmann and Moscarini (2012). However, we deviate from their framework by removing the firm's fixed costs of production. This removal eliminates the "gambling for resurrection" effect for which we cannot find evidence in our micro-level data. Furthermore, the moments in the data indicate that the frequency and absolute size of price changes converge from above to a fixed level over time. This convergence indicates that the incentives for price changes are *not* driven by active learning alone. We deal with these empirical regularities by adding a menu cost and firm-level idiosyncratic shocks. We choose to model active learning in the most transparent manner that is still consistent with the data. Nonetheless, we explore several extensions to our benchmark model to show its robustness (section 1.4.4) and to highlight further implications (section 1.4.3).

1.3.1 Households

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

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

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 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) where the first index $i \in \{1, 2\}$ denotes the good's *type*. Its *variety* within a group type is denoted by $k \in J_i$. Goods come in two types or baskets: J_1 and J_2 are the groups of *specialties* and *generics* respectively. A group J_i is characterized by its Lebesgue measure φ_i. Varieties within the first basket are hard to substitute with each other whereas generic varieties are mutually substitutable with a relatively high elasticity of substitution, thus we have $\sigma_2 > \sigma_1$. Each good is assumed to be produced by a single monopolistically competitive producer.[26] Each good is identified through the index pair (i, k). We assume that i is time-invariant whereas the consumer's variety-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):

$$c_t^i(k) = \alpha_t^i(k) \left(\frac{p_t^i(k)}{P_{it}} \right)^{-\sigma_i} \frac{\eta_i S_t}{P_{it}} \tag{1.5}$$

where, with some abuse of notation, the income shares for each basket are given by $\eta_1 = \eta$ and $\eta_2 = 1 - \eta$. $p_t^i(k)$ denotes the price of a good (i, k) in period t and the price index of group i is denoted by:

26. As a result, we will use the term "firm" and "product" interchangeably.

$$P_{it} = \left(\int_{k \in J_i} \alpha_t^i(k) \left(p_t^i(k) \right)^{1-\sigma_i} dk \right)^{1/1-\sigma_i} \tag{1.6}$$

These price indices satisfy the following equalities:

$$\begin{aligned} P_{1t}C_{1t} + P_{2t}C_{2t} &= S_t \\ P_{1t}^{\eta} P_{2t}^{1-\eta} &= P_t \\ P_t C_t &= S_t \end{aligned}$$

Thus, 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.[27] 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. Therefore, the period t budget constraint is characterized by:

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

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

27. In particular, we use the *Glivenko-Cantelli* theorem. The argument is identical to the one in Bachmann and Moscarini (2012).

problem are then:

$$q_{t,T} = \beta^{T-t} \left(\frac{C_T}{C_t} \right)^{-\theta} \frac{P_t}{P_T} \tag{1.7}$$

$$\frac{W_t}{P_t} = \omega L_t^{\chi} C_t^{\theta} \tag{1.8}$$

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

1.3.2 Firms

In contrast to most state-dependent pricing models, firms in our framework set prices under incomplete information. 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 are comprised of:

$$q_t^i(k) = \alpha_t^i(k) \left(\frac{p_t^i(k)}{P_{it}} \right)^{-\sigma_i} \frac{\eta_i S_t}{P_{it}}$$

Given that this demand specification is log-linearly separable, we obtain:

$$\begin{aligned} \log\left(q_t^i(k)\right) &= \log(\alpha_t^i(k)) - \sigma_i \log\left(\frac{p_t^i(k)}{P_{it}}\right) + \log(\eta_i S_t) - \log(P_{it}) \\ &= -\sigma_i \log(p_t^i(k)) + \log(S_t) + (\sigma_i - 1)\log(P_{it}) + \log(\eta_i) + \log(\alpha_t^i(k)) \end{aligned}$$

where we define $\log(\alpha_t^i(k)) = \varepsilon_t^i(k)$. 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:

$$q = -\sigma_i p + s + \mu_i + \varepsilon_k \tag{1.9}$$

A firm (i, k) does not know to which group it belongs to, that is, it does not know whether its type satisfies $\sigma_i = \sigma_1$ or $\sigma_i = \sigma_2$. Furthermore, it does not observe the realization of the (log) taste shock ε_k. From the firm's point of view, there is no longer a one-to-one mapping between quantities and prices. Whenever a firm sets a price p, demand q can be high for three reasons: (1) the variety belongs to a basket within which substitution is hard (and therefore the market power is high), that is, $\sigma_i = \sigma_1$, (2) the realization of the taste shock ε_k is high, or (3) the basket to which the firm belongs to only has a few competitors (i.e., low φ_i) so μ_i is high and the consumer spends much of his or her income on goods in basket i. For example, if a firm knows its elasticity of substitution σ_i but does not know the realization of the taste shock ε_k and hence takes expectations over it, then the optimal set price is characterized by $p^* = \frac{\sigma_i}{\sigma_i - 1} MC$ where MC denotes the firm's marginal cost. Thus, we immediately can deduce $P_i = \left(\frac{\sigma_i}{\sigma_i-1} MC\right) \varphi_i^{1/1-\sigma_i}$ and $\log(P_i) = \frac{1}{1-\sigma_i} \log(\varphi_i) + \log\left(\frac{\sigma_i}{\sigma_i-1} MC\right)$. Since $\sigma_i > 1$, $\log(P_i)$ strictly decreases with φ_i.

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. Our quantitative results rely on CES preferences that most of the price-setting models use. However, our results on firm-level learning do not rely on these specific type of preferences. The results will hold as long as the demand function is linearly separable after some uniform transformation.

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 a 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 $\exp(q) = \exp(-\sigma_i p + s + \mu_i + \varepsilon_k)$ conditional on the set price p.

3. The firm is contractually obliged to supply $\exp(q)$.

Let λ denote the firm's prior belief in a low elasticity of demand (i.e. $\sigma_i = \sigma_1$) and f 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 beliefs to the posteriors λ' according to Bayes' rule:

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

However, dynamic decision-making requires knowing the evolution of beliefs conditional on a given state i, where the true data generating process for sold quantities is $q = -\sigma_i p + s + \mu_i + \varepsilon_k$. The firm rationally anticipates that the price it sets will affect the informative quantity it will observe the period after. Prices in period t affect a firm's future beliefs in period $t + 1$ conditional on the true state being i. 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. The posterior belief, conditional on state i, is equal to:

$$b_i(\lambda, p, \varepsilon) = B(\lambda, p, -\sigma_i p + s + \mu_i + \varepsilon, s).$$

In the remainder of this paper, we assume that log demand shocks are normally distributed. Thus, we get $\varepsilon_k \sim N(0, \sigma_\varepsilon^2)$. After some algebra, we can derive that:

$$b_i(\lambda, p, \varepsilon) = \left(1 + \frac{1 - \lambda}{\lambda} \exp\left(\tfrac{1}{2}(-1)^{\mathbf{1}(i=2)} \left[\left(\frac{\varepsilon}{\sigma_\varepsilon}\right)^2 - \left(p \cdot \frac{\Delta\sigma}{\sigma_\varepsilon} - \frac{\Delta\mu}{\sigma_\varepsilon} + \frac{\varepsilon}{\sigma_\varepsilon}\right)^2\right]\right)\right)^{-1} \tag{1.10}$$

where $\Delta\sigma \equiv \sigma_2 - \sigma_1 > 0$. Expression 1.10 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, if the firm sets a price equal $\Delta\mu/\Delta\sigma$, then the posterior beliefs do not change regardless of the realization of the taste shock. As a result, at this price, beliefs are self-reinforcing.

This model of active learning with discrete types is parsimonious and is a sufficient ingredient for generating the empirical findings of section 1.2.4. Alternatively, we could work with a continuum of types such that a firm's beliefs are captured through a probability distribution. In appendix A.6.1, we show that it is unlikely that our quantitative results will be affected under this model because it shares many of the same features as the model with discrete types.

However, there are additional advantages to using discrete types. First, it simplifies the computational procedure significantly. The simplest case of active learning with a continuum of types that is still tractable features uncertainty about its demand elasticity only and Gaussian conjugates its priors. Even under this case, a firm's beliefs will consist of at least a pair, that is, a mean and a variance, which is more than the single state for the prior belief in our setup. As a result, our computational procedure suffers much less from the curse of dimensionality. Second, Kiefer and Nyarko (1989) show that under a continuum of types there are multiple limit beliefs that are outcomes of optimal policy but that do not coincide with the true parameter values. Under discrete types we show that only one limit belief exists that does not converge to the truth.

Two Period Model with Active Learning

To display the active learning mechanism as clearly as possible, we use a model with only two periods. If the firm is type i, then its profits are conditional on setting a price p equal to $\Pi^i(p)$. We impose that $\Pi^i(\cdot)$ is strictly concave, which is standard.

By construction, the firm only cares about maximizing myopic profits in the second period. These myopic profits are a linear combination of the concave functions $\Pi^1(p)$ and $\Pi^2(p)$ based on the prior belief λ, that is, $M(p;\lambda) = \lambda\Pi^1(p) + (1-\lambda)\Pi^2(p)$. Thus, we must have:

$$V_2(\lambda) \equiv \max_{p \in \mathcal{P}} M(p;\lambda)$$

Therefore, its maximizer $p^M(\lambda)$ is unique and monotonically increasing in the belief λ.[28] In the first period however, the firm must balance its incentives between obtaining higher myopic profits and sharpening its posterior beliefs to increase its continuation value.

28. A formal argument can be found in proposition 2 in the appendix.

$$V_1(\lambda_0) = \max_{p \in \mathcal{P}} \left\{ M(p;\lambda_0) + \beta \mathbb{E}_\varepsilon \Big[\lambda_0 V_2(b_1(\lambda_0, \log(p), \varepsilon)) + (1-\lambda_0) V_2(b_2(\lambda_0, \log(p), \varepsilon)) \Big] \right\}$$

where $b_i(\lambda_0, \log(p), \varepsilon)$ is defined as in (10) and we denote its maximizer as $p^*(\lambda_0)$. Further, the policy $p^M(\lambda)$ is the optimal pricing function whenever the firm is unable to affect its posterior beliefs. As a result, a firm actively learns with its price at the belief λ_0 if it deviates from this myopic price function. This deviation $|p^*(\lambda_0) - p^M(\lambda_0)|$ reflects the firm's incentive to gain information to increase the speed of learning at the expense of its current period profits.

Convexity of the value of information. A relatively large literature has established that the firm's active learning is formally captured by a continuation value that is convex in a firm's beliefs.[29] The following lemma establishes this feature.

LEMMA 1. The value function $V_2(\cdot)$ is convex and $\mathcal{C}^2$.

Proof. See appendix A1.1. □

The result is shown explicitly for the two period setup. However, it is generalizable to the infinite period framework. The convexity of $V_2(\cdot)$ is important because, to establish sufficient conditions for active learning, we follow Mirman, Samuelson, and Urbano (1993). In their proposition 1, the convexitity of $V_2(\cdot)$ is one of their two sufficient conditions. Informally, the second condition states that adjustments in prices must be capable of increasing the informativeness of a firm's sales.

Incentives for active learning. A firm sets its price to identify its elasticity of substitution by "separating" these two possible demand curves as far apart as possible. This separation indicates that the price at which the demand curves cross in expectation results in sales that are completely uninformative. Thus, we deduce that the expected demand curves cross if and only if $p = \frac{\Delta\mu}{\Delta\sigma}$.

29. For example, this argument can be found in Aghion, Bolton, Harris, and Jullien (1991) and Mirman, Samuelson, and Urbano (1993).

This intersecting price can be defined as the confounding price $\widehat{p}$. If the firm decides to choose its active learning policy $p^*(\lambda)$ to be equal to $\widehat{p}$, then there should be no benefits from active learning. This reasoning is formalized in the following proposition:

Proposition 1. Let $\widehat{p} = \frac{\Delta\mu}{\Delta\sigma} \in (p_2^*, p_1^*)$, then there exists a confounding belief $\widehat{\lambda}$ such that either one of the following two cases hold.

I. $p^*(\widehat{\lambda}) = p^M(\widehat{\lambda}) = \widehat{p}$.

II. $p^*(\lambda)$ is discontinuous at $\lambda = \widehat{\lambda}$.

Furthermore, the confounding belief $\widehat{\lambda}$ is unique up to $\lambda \in \{0, 1\}$ and strictly increasing (decreasing) in $\Delta\mu$ ($\Delta\sigma$).

Proof. See appendix A.2.2. □

Numerical example. In this example, we parameterize the profit function using CES demand curves with elasticities σ_1 and σ_2 and, for simplicity, we set $\lambda_0 = \widehat{\lambda}$. Figure 1.5 plots the firm's static profits, the continuation value, and the total payoff (which is the sum of the latter two) as a function of the firm's set price p. The dotted lines at the extremes of the figure depict the optimal prices p_2^* and p_1^* under perfect information. Given that $M(p; \lambda_0)$ is a weighted sum of strictly concave functions in p, it is strictly concave in p itself. By definition, it is 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 1.5 shows that $\mathcal{V}(\cdot; \lambda_0)$ is convex. It also shows that its minimum lies at the confounding price $\widehat{p}$. The reason is that a firm's sales become completely informative 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. The convexity of $\mathcal{V}(\cdot; \lambda_0)$ is also affected by the signal-to-noise ratio. For extremely large values of σ_ε, for instance, the optimal policy converges to the myopic policy. This is because there is no amount of variation in its price that the firm can use to induce an informative signal. As a result, the firm behaves as if its price does not affect its posterior beliefs, which is equivalent to behaving myopically.

A firm bases its pricing strategy by maximizing its total payoff, which is the sum of strictly concave and convex functions. 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] As mentioned above, the degree of active learning is captured by the difference between $p^*(\lambda_0)$ and $p^M(\lambda_0)$.[32] In appendix B.1, we show that the active learning mechanics with a continuum of types are identical in a two-period setting. Just as in this section, a firm faces the "current control-estimation" trade-off by maximizing a total payoff that consists of a strictly concave myopic profit function and a convex continuation value.

Active learning regimes. The gains from active learning are strongly related to the convexity of $\mathcal{V}(\cdot;\lambda_0)$. This, in turn, is determined by the prior belief, the signal-to-noise ratio, and the discount factor. A firm's prior belief determines how certain it is about its type. A firm has less incentive to engage in active learning as its belief moves closer to zero or one. The signal-to-noise ratio summarizes the sensitivity of a firm's posterior beliefs to price deviations relative to the confounding price. Thus, firms that face extremely large levels of noise will basically never receive an informative signal through their sales. As a result, they have no incentives to actively learn. Lastly, the discount factor indicates how much a firm values more information in future periods. The convexity of $\mathcal{V}(\cdot;\lambda_0)$ determines the shape of the total payoff function. In our previous numerical example, the total payoff function was double-peaked because $\mathcal{V}(\cdot;\lambda_0)$ was sufficiently

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 proposition 3 of the appendix, we derive a set of sufficient conditions to guarantee that $p^*(\lambda_0) \in [p_2^*, p_1^*]$ for all λ_0.

32. In our example, total expenditure S is constant across the two periods. Suppose that $S_1 \neq S_2$, then note that the incentives to engage in active learning increase if the firm expects demand to increase in the second period since the instantaneous profits are proportional to aggregate demand. In particular, if $\frac{S_1}{P_1} < \frac{S_2}{P_2}$, then the cost of active learning is lower in the first period. In the data, the demand for a new product increases the first few periods after entry which increases the incentives to actively learn early on the life cycle. More than 80% of the variation in the demand for a product can be explained by the extensive margin of demand.

convex but this might not always be the case.

The shape of the total payoff function determines the active learning regime.[33] In our setup, there are two qualitatively different regimes determined by the shape of $\mathcal{V}(p;\lambda_0)$: extreme and moderate active learning. Under extreme active learning, the total payoff function is double-peaked. As a result, the firm never chooses to price at the confounding price, and $p^*(\lambda)$ displays a discontinuity at $\lambda = \widehat{\lambda}$. Since the value of information is minimized at the confounding belief $\widehat{\lambda}$, the firm has the most incentive to change its price at this specific belief and deviates in a discontinuous fashion. But, under moderate active learning, the total payoff function is single-peaked and the policy function $p^*(\cdot)$ is continuous between p_2^* and p_1^*.

Figure 1.6 depicts the two active learning regimes. The thin gray line shows the myopic policy function $p^M(\lambda)$ that is monotonically increasing in λ whereas the purple line is the policy function $p^*(\lambda)$ under active learning. The figure shows that $p^M(\lambda)$ and $p^*(\lambda)$ are bounded from below and above by p_2^* and p_1^*, respectively, which proposition 3 in the appendix predicts. Under extreme active learning, the policy function shows a discontinuity at the confounding belief. The firm actively learns mostly near $\widehat{\lambda}$ as it tries to keep the informativeness of its observed sales as high as possible. It can only do this to a limited extend as otherwise the firm would lose too many static profits. With moderate active learning, the myopic policy coincides with the active learning policy at the confounding price $\widehat{p}$ as predicted by proposition 1. Once the firm updates its posterior closer to the boundaries (i.e., $\lambda \in \{0, 1\}$), the incentives for active learning decline again as the firm's information set converges to the complete information case. In this case, the myopic and active learning policies coincide at $\lambda \in \{0, 1\}$. Hence, the firm would never pay the opportunity costs (i.e., give up static profits) through active learning whenever its beliefs reach either zero or one.[34]

33. The terminology is borrowed from Keller and Rady (1999) who show that different active learning regimes could arise in a problem of a seller choosing quantities subject to a randomly changing state.

34. In the two-period model with menu costs, the firm must decide to either adjust its price or maintain it at the same level. Under perfect information, the firm follows a standard (s, S) policy and the region of inaction depends on the curvature of the profit functions and the menu cost. But, under demand uncertainty, the width of the inaction band also depends on the firm's prior belief and it is larger close to the confounding belief because the variance in the belief changes is higher. This variance induces a high option value of waiting that is reflected in the larger width of inaction (figure A.17). This inaction, in turn, reduces the adjustment frequency. On the other hand, higher uncertainty pushes the firm to adjust for a given region of inaction.

Dynamic Pricing Policies under Incomplete Information

Firms are ex ante identical but can generate heterogeneous ex-post pricing paths as different realizations of the log demand shocks induce differently updated posterior beliefs. A firm has access to a linear production technology in labor. 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 labor the firm uses for production purposes in period t. Its idiosyncratic productivity is given by $z_t^i(k)$. Labor is supplied competitively at the nominal rate W_t, then a firm's static profits, conditional on being type i, are equal to:

$$\left(p_t^i(k) - \frac{W_t}{z_t^i(k)}\right)\alpha_t^i(k)\left(\frac{p_t^i(k)}{P_{it}}\right)^{-\sigma_i}\frac{\eta_i S_t}{P_{it}}$$

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 chooses a price to determine the trade-off between maximizing current profits and obtaining more accurate information in the future about its elasticity of demand. Since a firm cannot observe the realization of its demand shock $\alpha_t^i(k)$ whenever it has to decide on its pricing policy, it has to take expectations over it. Due to our normalization $\mathbb{E}(\alpha_t^i(k)) = 1$, we obtain a firm's ex-interim expected profits conditional on type i, setting some price p, and having idiosyncratic productivity z:

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

A firm does not know its elasticity of demand, so it takes expectations over these ex-interim profits using their current prior belief λ_t. Furthermore, firms are required to pay a fixed cost of ψ in units of labor to adjust their nominal price. This results in a firm's (ex-ante) expected profits. Thus, we define:

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

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

$$\mathbb{E}_0 \sum_{t=0}^{\infty} q_{t,t+1} \Pi_t(p_t^i(k); z_t^i(k))$$

where the expectation is with respect to the path of future beliefs, demand, and productivity shocks. Any firm makes its pricing decisions while taking aggregate prices, spending, and the wage rate as given. These variables are determined in general equilibrium and are summarized by the aggregate state $\xi_t \equiv (P_{1t}, P_{2t}, W_t, S_t) \in \Xi$. In the following, we focus on a stationary equilibrium in which nominal aggregate spending trends are at a constant rate $\tilde{\pi} \geq 0$:

$$\log(S_{t+1}) = \log(S_t) + \tilde{\pi}$$

Thus, there is no aggregate uncertainty and the state ξ_t is constant in this stationary equilibrium. Profits are then discounted at the rate β.

Our framework is a hybrid version of standard state-dependent pricing models and frameworks that feature active learning, which in our setup only adds a state variable. Firms start out with a prior λ_0 and an initial productivity draw. They then choose their entry price optimally without paying a menu cost. In our model firms have substantial incentives to learn their type at the beginning of their life cycle by adjusting their prices to obtain information. As the product matures the gains to obtaining additional information are extremely small and they do not offset the menu cost. Given that the frequency and absolute size of price changes at these stages are nonnegligible, we capture the incentives for price changes at the later stages of a product's life cycle 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:

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

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

$$\begin{aligned}
V^A(\lambda, z) = \max_{p \geq 0} \, & (p - \tfrac{W}{z}) \left[\lambda\eta \frac{p^{-\sigma_1}}{P_1^{1-\sigma_1}} + (1-\lambda)(1-\eta) \frac{p^{-\sigma_2}}{P_2^{1-\sigma_2}} \right] \frac{S}{P} - \psi \frac{W}{P} \\
& + \beta\lambda \int_{z'} \int_{\varepsilon} V\left(b_1(\lambda, \log(\tfrac{p}{1+\bar{\pi}}), \varepsilon\right), z', \tfrac{p}{1+\bar{\pi}}) dF(\varepsilon) dG(z', z)
\end{aligned}$$

$$\begin{aligned}
& + \beta(1-\lambda) \int_{z'} \int_{\varepsilon} V\left(b_2(\lambda, \log(\tfrac{p}{1+\bar{\pi}}), \varepsilon\right), z', \tfrac{p}{1+\bar{\pi}}) dF(\varepsilon) dG(z', z) \\
V^{NA}(\lambda, z, p_{-1}) = & (p_{-1} - \tfrac{W}{z}) \left[\lambda\eta \frac{p_{-1}^{-\sigma_1}}{P_1^{1-\sigma_1}} + (1-\lambda)(1-\eta) \frac{p_{-1}^{-\sigma_2}}{P_2^{1-\sigma_2}} \right] \frac{S}{P} \\
& + \beta\lambda \int_{z'} \int_{\varepsilon} V\left(b_1(\lambda, \log(\tfrac{p_{-1}}{1+\bar{\pi}}), \varepsilon\right), z', \tfrac{p_{-1}}{1+\bar{\pi}}) dF(\varepsilon) dG(z', z) \\
& + \beta(1-\lambda) \int_{z'} \int_{\varepsilon} V\left(b_2(\lambda, \log(\tfrac{p_{-1}}{1+\bar{\pi}}), \varepsilon\right), z', \tfrac{p_{-1}}{1+\bar{\pi}}) dF(\varepsilon) dG(z', z)
\end{aligned}$$

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.[35] However, sales are observable and informative. 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

35. 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.

increasing the rate at which it learns about its type.[36]

1.3.3 Stationary Equilibrium

Aggregate Price Consistency

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 shocks lead to ex-post heterogeneity of a firm's prior belief λ in the cross-section. Furthermore, firms also become heterogeneous due to different realizations of the idiosyncratic productivity shock in the cross-section. Note there is not only cross-sectional dispersion in beliefs across firms of different types but also within groups. 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. We previously defined the aggregate price index as:

$$P_i = \left(\int_{k \in J_i} p^i(k)^{1-\sigma_i} dk\right)^{\frac{1}{1-\sigma_i}}$$

However, the optimal pricing policy $p^*(\lambda, z)$ is independent from i and k. To obtain price consistency in the aggregate, we thus require:

$$P_i = \left(\int p^*(\lambda, z)^{1-\sigma_i} d\varphi_i(\lambda, z)\right)^{\frac{1}{1-\sigma_i}} \tag{1.11}$$

Labor Market Clearing

The market clearing condition for goods is explicitly incorporated in the firm's problem, thus the only remaining factor market to clear is the labor market. In the remainder of our analysis, we use the log utility in consumption (i.e., $\theta = 1$) and an inverse Frisch elasticity of zero (i.e., $\chi = 0$).

36. Note that our framework is fundamentally different from most price-setting models with learning. In the framework by Baley and Blanco (2017), 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 flows are driven by Brownian motions and a Poisson shock. In contrast, our model considers firms that can explicitly affect their set of information. As a result, the flow of information becomes an endogenous object.

This restriction with separable, additive utility in consumption and leisure means that wages are proportional to aggregate spending.[37] This gives us:

$$W = \chi S$$

Total nominal spending S equals PC, where $P = P_1^{\eta} P_2^{1-\eta}$, and thus gives us an expression for the *real* wage rate:

$$\frac{W}{P} = \chi C$$

Given the linear production technology, labor demand is simply characterized by:

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

Labor market clearing means $L^d = L$ where L is equal to one third in our calibration exercises.

Equilibrium

We focus on a stationary equilibrium in which any dying firm is immediately replaced by a new firm. The latter is a type 1 firm with probability λ_0 that serves as its prior belief at entry. We simplify the analysis by normalizing the measure of firms to one and by organizing the industry composition as follows:

$$J_1 = [0, \gamma_1] \text{ and } J_2 = (\gamma_1, 1].$$

Our restrictions on entry then mean that $\gamma_1 = \lambda_0$, which guarantees a balanced measure of in- and out-flows at the product level. Whenever nominal total spending grows deterministically at the rate $\tilde{\pi}$, there is no aggregate uncertainty. If W is the economy's numéraire, then we can define

37. See Golosov and Lucas (2007) who use the same specification for consumer preferences. In their setup, this specification means that wages are proportional to the stock of money. Thus, it grows at the same rate as inflation. We have a similar proportionality rule as wages become proportional to total spending which grows at the rate of inflation $\tilde{\pi}$.

a stationary equilibrium.[38]

Definition 1 (Stationary equilibrium). *A stationary equilibrium is a tuple* (W, P_1, P_2, S) *and a pair of invariant distributions* $(\varphi_1(\lambda, z), \varphi_2(\lambda, z))$ *such that the real variables are constant. Thus:*

I. *Consumers maximize utility by consuming varieties* $c^i(k)$, $k \in J_i$, $i \in \{1, 2\}$.

II. *Firms maximize profits by adopting* $p^*(\lambda, z)$ *when adjusting prices.*

III. *Factor markets clear.*

IV. *Prices are consistently aggregated.*

V. *Firms die at the rate* δ *and enter the economy as a type 1 firm with probability* λ_0.

1.4 Propagation of Nominal Shocks

1.4.1 Calibration and Results

Because the IRI Symphony data is weekly, we set the model period at one week. As a result, the discount factor is set at $\beta = 0.96^{1/52}$ that reflects an interest rate of around 3.8% and incorporates the exogenous exit rate of 0.4%. This rate comes directly from the IRI Symphony data at the UPC-store level. The mean yearly growth rates of nominal and real GDP equal $g_n = 0.04$ and $g_r = 0.02$ respectively. Since there is no long-run real growth in the model economy, we set $\tilde{\pi} = (g_n - g_r)^{1/52} = 0.00038$ as the weekly rate of inflation. Furthermore, the standard deviation of the taste shock σ_ε equals 0.4, which matches the standard deviation of sold quantities conditional on no price change in the IRI Symphony, which is the value of 42 percent that is reported in Eichenbaum, Jaimovich, and Rebelo (2011). Lastly, the disutility of labor χ is chosen so that the aggregate employment is approximately $\frac{1}{3}$ because we normalize the amount of time available to the consumer to unity.

The remaining parameters are calibrated to match various micro-data moments. There are seven remaining parameters: 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, and the persistence and standard

38. In the appendix A.5, we describe the numerical algorithm to solve this framework computationally.

deviation of idiosyncratic productivity ρ and σ_ζ. These parameters are calibrated jointly and are selected to hit eight moments from the data: the average frequency of adjustment, the average size of increases, the average size of decreases, the fraction of price changes that are increases, the frequency of adjustment on the second and tenth week after entry, and the absolute size of the price changes during the second and tenth week after entry.

As is standard in the class of menu cost models, the fixed cost of adjustment ψ partially governs the average frequency and size of adjustments. The extent to which active learning is more present at the beginning of a product's life cycle is determined by the amount of information a firm has at entry summarized by the signal-to-noise ratio $(\sigma_2 - \sigma_1)/\sigma_\varepsilon$ and the prior belief at entry λ_0 that represents the fraction of firms facing the elasticity of substitution σ_1. We assume this parameter to be equal across all entering firms. As the incentives to actively learn decrease, price changes are mainly driven driven by idiosyncratic cost shocks. Thus, the parameters ρ and σ_ζ have a relatively large impact on the pricing moments at the later stage of the product's life cycle.

Table 1.6 shows the model's best parameters in terms of fitting moments, and table 1.7 displays the resulting moments from the framework compared to the data. The productivity parameters are in line with previous estimates in the menu cost literature. The model matches the frequency of adjustment and fraction of increasing price changes closely. The specification for the menu cost ψ means that the total adjustment costs in the economy represent approximately 0.7% of steady-state weekly revenues. The cost conditional on adjustment is around 1.4%, 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 enough active learning. Nonetheless, σ_2 is well within the estimates of Broda and Weinstein (2010) who compute elasticities of substitution for a variety of products using data similar to ours.

The model also performs well in replicating the life cycle patterns in the frequency of price adjustments and the absolute size of the adjustments. Figure 1.7 shows that in our simulations entering products are more likely to adjust prices and they do so by larger amounts. This is driven by 1) the size of the signal-to-noise ratio and 2) the overall level of the elasticities of substitution since firms with lower market power (i.e., high elasticities of substitution) have higher incentives to get their prices "right" as the opportunity costs of active learning (i.e., sacrificing static profits) are higher. The incentives to actively learn also affect the size distribution of price changes by generating large price changes endogenously. Our calibration matches the standard deviation of price changes and price changes in the 75th percentile in absolute value without explicitly targeting them. The model, however, underpredicts the prevalence of price changes in the 90th percentile of the size distribution. [39]

1.4.2 Implications

We perform a counterfactual experiment in which the log nominal output increases permanently by a size that is comparable to a one week doubling of the nominal output growth rate. We observe that on impact approximately 70 percent of the nominal shock goes into output. In a baseline menu cost model with full information, this value is around 60 percent.[40]

Alvarez, Lippi, and Passadore (2017) show that in a large class of continuous time, price-setting models, both state- and time-dependent, the effect of a small nominal shock is identical across these models. In discrete time, however, the response when the shock occurs conveys information about the aggregate price flexibility. To quantify this flexibility, we follow Caballero and Engel (2007) and compute the "Flexibility Index": an accounting relation that describes how inflation will respond when a small shock occurs and that is fully pinned down by the current distribution of the firms' desired price changes and the adjustment hazard.

39. The resulting hazard of price changes in our economy is downward-sloping with a small hump at short durations (figure 1.8). 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 as firms are less likely to adjust after they reset their prices. On the other hand, as firms learn, the probability of consecutive price changes is larger at entry since new information might cause 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 decrease in the slope of the hazard at long horizons. This is in contrast to the hazard rate in Bachmann and Moscarini (2012) that is completely flat at zero with a spike at 21 months. This is because of the presence of idiosyncratic cost shocks and due to the fact that in order to match the life cycle moments, the signal to noise ratio in our calibration is larger, which increases the incentives to learning.

40. The full information model reflects the Golosov-Lucas benchmark with two types of firms (those facing elasticity σ_1 or σ_2), but both groups of firms are fully aware of their type. This model is then calibrated to match the same moments as the model with active learning and features the same fraction of firms of each type. Since we are interested in the effect of active learning on real output, the full information Golosov-Lucas model with two types is the appropriate benchmark.

This index is valid in all models including our model with active learning. We begin by decomposing the price response when the shock occurs into intensive and extensive margins. If $x_t(\lambda) \equiv \ln\left(\frac{p_t}{p_t^*(\lambda)}\right)$ is the price gap, then it can be defined as the difference between a firm's current price and its desired price, that is, the price it chooses as a function of its beliefs conditional on adjustment. If the economy-wide distribution of price gaps is given by $f(x,\lambda)$, we assume that firms have an adjustment probability $\Lambda(x,\lambda)$ that is increasing in their price gap.[41] If there is some unexpected, positive shock $\triangle S > 0$ to firms' desired prices, the price response equals:

$$\lim_{\triangle S \to 0} \frac{\triangle \pi}{\triangle S} = \underbrace{\int \Lambda(x,\lambda) f(x,\lambda) dx d\lambda}_{=\text{ intensive}} + \underbrace{\int x(\lambda) \Lambda_x(x,\lambda) f(x,\lambda) dx d\lambda}_{=\text{ extensive}} \tag{1.12}$$

that can be seen as the sum of two components: intensive and extensive margins. The intensive margin measures the contribution to inflation of the firms' products whose prices would have adjusted without the aggregate shock taking place. These firms adjust to the aggregate shock by changing the size of their adjustment. Equation 1.12 shows that this margin equals the frequency of adjustment. The extensive margin captures the strength of the selection effect and measures the additional inflation contribution of firms whose decision to adjust is either triggered or canceled by the aggregate shock. This margin becomes naturally more relevant as the number of firms near the margin of adjustment increases (i.e., large $\Lambda_x(x,\lambda)$) or when the difference between adjusting and not adjusting is large (i.e., large $|x(\lambda)|$). Calibrated to the same frequency of price adjustment, the difference between any two models in quantifying the effect solely reflects the difference in the extensive margin.[42] The desire to actively learn from prices pushes firms away from the margin of adjustment; both lowering the mass of firms at the original bounds of inaction and substantially reducing the importance of the extensive margin. Even though this approach recovers only the price response to the shock, Berger and Vavra (2015) show that it is highly predictive of the overall price stickiness in the economy.

Furthermore, the half-life of the real response more than doubles in our framework with respect

41. To simplify the math, we assume here that a positive small shock $\triangle S$ does not affect firms' beliefs. However, our results do take this effect into consideration as we calculate the extensive and intensive margins numerically.

42. In the Calvo model, the extensive margin is zero as there is no selection effect.

to that of the Golosov-Lucas benchmark. This is because, by introducing the product's life cycle, we introduce cross-sectional heterogeneity in the frequency of price adjustments across firms of different ages endogenously. As a result, the coefficient of variation of price spells duration is 30 to 35 percent larger than in the full information benchmark.

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 the output because these firms have already adjusted to the shock. Given that the model is calibrated to match the average frequency of price changes, the fact that firms certain about their type have, on average, a lower frequency of price adjustment significantly delays the adjustment of the aggregate price level after a nominal shock. In other words, a higher level of cross-sectional heterogeneity in the duration of price spells reduces selection in timing after an unanticipated monetary shock as pointed out by Alvarez, Lippi, and Paciello (2011) and Carvalho and Schwartzman (2015).[43]

Quantitatively, the cumulative effects on real output are 2.3 times larger under demand uncertainty than in the model with full information (as shown in figure 1.9).[44] In our baseline setup with firms producing only one good ($n = 1$), no fat-tailed shocks and no random menu costs, our model with active learning has cumulative effects on real output that are about 2/5 of that in the Calvo framework. This magnitude is very comparable to the multi-product model by Alvarez and Lippi (2014) for the case of $n = 10$ or to the single product Golosov-Lucas model with random menu costs (also known as the "CalvoPlus" specification) where the fraction of free price adjustments is 80 percent ($l = 0.8$). Furthermore, the cumulative real effects in our framework are only 20 percent lower than the $n = \infty$ case in Alvarez, Bihan, and Lippi (2014). In contrast, the cumulative effects on real output in our benchmark model with the addition of random menu costs (which are added

43. 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. If, for example, the overall frequency of price adjustment in the economy is a convex function of the frequency of price changes of firms actively learning and those certain about the elasticity they face, heterogeneity in the cross-sectional distribution of firms will amplify monetary non-neutrality.

44. The area under the impulse response in the full information model is 1/6 of that in the Calvo model which is also found by Alvarez, Bihan, and Lippi (2014).

in order to generate small price changes) are almost 1.3 times as large as in the Golosov-Lucas model with a fraction of free price adjustments of $l = 0.8$.[45]

1.4.3 Nominal Shocks in Periods of High Product Entry

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. There are several reasons why focusing on the business cycle could be important. First, Argente, Lee, and Moreira (2017a) show that the entry rate of new products is highly procyclical. Second, previous contributions also show that the impact of nominal shocks on real output vary over the business cycle. For example, Vavra (2014) shows that monetary policy is less effective in stimulating real output during downturns. For the sake of brevity, we provide only a summary of the model in this section. A more detailed description can be found in appendix A.6.4.

Consumers and firms are identical as in the baseline framework. However, a firm's productivity now consists of two components: an idiosyncratic and an aggregate component denoted by z and Z respectively. Furthermore, the entry rate of products is endogenous which allows it to vary with the aggregate state of the economy. In the beginning of each period, a pool of entrants observes the aggregate state and decides if they want to become producers by paying a fixed entry cost (denoted in units of labor). Lastly, we assume that the level of aggregate productivity follows a two-state symmetric Markov chain. A boom and a bust are defined as a one standard deviation increase and decrease, respectively, from average aggregate productivity. The latter is normalized to unity and we calibrate the former by estimating a standard autoregressive process on Fernald's (2014) utilization-adjusted Total Factor Productivity series. Transition probabilities are then calibrated such that the average length of a cycle is about 35 months.

Despite its simplicity, this setup shows whether nominal shocks are amplified more during booms. In our extension, periods of high aggregate productivity mean periods of high product entry. The calibration of this framework shows that the real output effects of a nominal shock

45. We calibrate the fraction of free adjustments to match the fraction of small price changes defined as $|dp_{it}| < \frac{1}{2} mean|dp|$ which is the data is approximately 40%.

are 15 percent larger in booms than during busts. This is because as the entry rate of products increases, the average firm gets younger and a higher proportion of firms then engage in active learning. These firms are less likely to adjust their prices after a nominal shock as their incentives to change their prices in response to idiosyncratic cost shocks are lower.

Although the findings of this section are robust to aggregate shocks of ordinary size, our findings might differ for very large aggregate shocks. In our calibration, 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, weakens the selection effect. However, the possibility exists that with a sufficiently large number of entering firms, 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.

1.4.4 Robustness and Extensions

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 1.2.4 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. We extend our framework in two different ways to deal with this issue. First, we allow for an exogenous, age trend in the demand shocks similar to Foster, Haltiwanger, and Syverson (2016) to incorporate the fact that entering products' sales grow over time after starting at a relatively low level. Appendix B.3.1 shows the details of our 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 1.4.2 virtually unchanged.

Second, we extend the canonical price-setting model of Golosov and Lucas (2007) by incorporating customer retention concerns. A firm's current level of demand depends positively on (a fraction of) the level of demand in the previous period. As a result, a firm has incentives to set low prices in the beginning of their life cycle to attract customers and build up their customer base. Whenever this base has reached sufficiently high levels, a firm starts exercising its market power by raising its price. These incentives are also known as "investing" and "harvesting". In appendix A.6.3, we show that such a framework is not consistent with our stylized facts from section 1.2.4.

Another 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 1.2.4. In appendix A.6.2, 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 consistent with the data, our conclusions are not affected significantly.

1.5 Conclusion

The increasing availability of micro-level data sets has allowed researches to delve deeper into the mechanics of a firm's dynamic pricing behavior. Recent studies have found new insights into firms' pricing behavior 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 paper, we aim to fill this gap by developing the 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 through their pricing strategies and show that this model can rationalize standard price-setting moments and our set of stylized facts.

In this context, we develop sufficient conditions for active learning to occur and describe the different regimes that could arise in this setup. 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 support our observations in the data. In our model, relative to the full information benchmark, the real effects of nominal shocks are at least twice as large and persistent when measured by their cumulative effect on real output.

Our quantitative framework contains the minimal amount of ingredients to rationalize the key mechanisms and 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.

Tables and Figures

Table 1.1: 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
Creation	0.29	0.07	0.47	0.15
Exit	0.42	0.13	0.61	0.25
Destruction	0.08	0.01	0.39	0.10

Note: The table shows the statistics of the entry rate, exit rate, creation, and destruction for 1-year and 5-year intervals. Columns (1) and (2) show the statistics at the UPC level and columns (3) and (4) at the UPC × Store level.

Table 1.2: Distribution of Duration by UPC × Store

	(1)	(2)	(3)	(4)
	Unweighted		Revenue Weighted	
	Weeks since Entry	Observations since Entry	Weeks since Entry	Observations since Entry
1^{st} percentile	1.0	1.0	12.4	7.7
25^{th} percentile	37.4	16.3	108.4	73.5
50^{th} percentile	96.3	47.1	183.7	131.8
75^{th} percentile	209.5	122.0	280.6	208.3
99^{th} 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

Note: 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.

Table 1.3: Life cycle Properties of Selected Pricing Moments - First 6 months

Dependent Variable	(1)	(2)	(3)	(4)
	Equal Weights		Revenue Weights	
Frequency	-0.054***	-0.076***	-0.072***	-0.104***
	(0.001)	(0.002)	(0.003)	(0.004)
Frequency increases	-0.051***	-0.069***	-0.066***	-0.092***
	(0.001)	(0.002)	(0.002)	(0.003)
Frequency decreases	-0.002***	-0.006***	-0.006***	-0.012***
	(0.000)	(0.001)	(0.001)	(0.002)
Absolute size	-0.154***	-0.171***	-0.154***	-0.174***
	(0.002)	(0.005)	(0.004)	(0.008)
Size increases	-0.194***	-0.223***	-0.185***	-0.223***
	(0.003)	(0.006)	(0.005)	(0.010)
Size decreases	0.075***	0.078***	0.091***	0.071***
	(0.002)	(0.006)	(0.004)	(0.011)
UPC $\times$ Store FE	✓	✓	✓	✓
Time FE		✓		✓
Cohort controls		✓		✓

Note: The table reports the coefficients (in percent) from the OLS tests. The independent variable is the age of the product and the dependent variables are the moments defined in the table. The sample is the first 6 months (or 26 weeks) after the product was first launched. The controls include UPC $\times$ 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 that assume equal weights for each UPC $\times$ store. Columns (3) and (4) report the results with revenue weights. The standard errors are clustered at the store level. The ***, **, and * denote significance at 0.01, 0.05, and 0.10 levels respectively.

Table 1.4: New Brand - First 6 Months

	Equal Weights		Revenue Weighted	
	Frequency	Size	Frequency	Size
	(1)	(2)	(3)	(4)
Age	-0.075***	-0.171***	-0.104***	-0.177***
	(0.002)	(0.005)	(0.004)	(0.008)
Age×New Brand	-0.028***	-0.047***	-0.020**	-0.085***
	(0.005)	(0.014)	(0.008)	(0.018)
UPC × Store FE	✓	✓	✓	✓
Time FE	✓	✓	✓	✓
Cohort	✓	✓	✓	✓

Note: The table reports the estimates of equation 1.2. The independent variable is the age of the product interacted with an indicator that equals one if the brand and the volume of the product are new. The dependent variables are the frequency and absolute size of the price changes. The sample is the first 6 months (or 26 weeks) after the product was first launched. The controls include 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 that assume equal weights for each UPC × store. Columns (3) and (4) report the results with revenue weights. The standard errors are clustered at the store level. The ***, **, and * denote significance at 0.01, 0.05, and 0.10 levels respectively.

Table 1.5: Newness Index - First 6 months

	Equal Weights		Revenue Weights	
	Frequency	Size	Frequency	Size
	(1)	(2)	(3)	(4)
Age	-0.076***	-0.173***	-0.109***	-0.177***
	(0.004)	(0.006)	(0.010)	(0.012)
Age×Newness	-0.053*	-0.228***	0.038	-0.310***
	(0.031)	(0.053)	(0.058)	(0.101)
UPC × Store FE	✓	✓	✓	✓
Time FE	✓	✓	✓	✓
Cohort	✓	✓	✓	✓

Note: The table reports the estimates of equation 1.4. The independent variable is the age of the product interacted with the Newness index. The dependent variables are the frequency and absolute size of the price changes. The sample is the first 6 months (or 26 weeks) after the product was first launched. The controls include 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 that assume equal weights for each UPC × store. Columns (3) and (4) report the results with revenue weights. The standard errors are clustered at the store level. The ***, **, and * denote significance at 0.01, 0.05, and 0.10 levels respectively.

Table 1.6: Internally Calibrated Values of the Model's Parameters

Description	Parameter	Value
Elasticity of Substitution 1	σ_1	4.7
Elasticity of Substitution 2	σ_2	13.4
Prior Belief at Entry	λ_0	0.75
Basket Division of Income	η	0.36
Fixed Cost	ψ	0.02
Productivity Persistence	ρ	0.56
Productivity Standard Deviation	σ_ς	0.05

Table 1.7: Moments of Price Change Distribution

Moment	Data	Model with Learning	Full Info
Frequency Week 2	0.09	0.09	0.002
Frequency Week 10	0.06	0.05	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	0.58	0.58
Size Up	0.09	0.09	0.06
Size Down	0.07	0.11	0.07
		Not Targeted	
Std. of Price Changes	0.11	0.11	0.07
75th Pct Size Price Changes	0.10	0.12	0.07
90th Pct Size Price Changes	0.18	0.14	0.08

Figure 1.1: Frequency of Price Adjustment at Entry

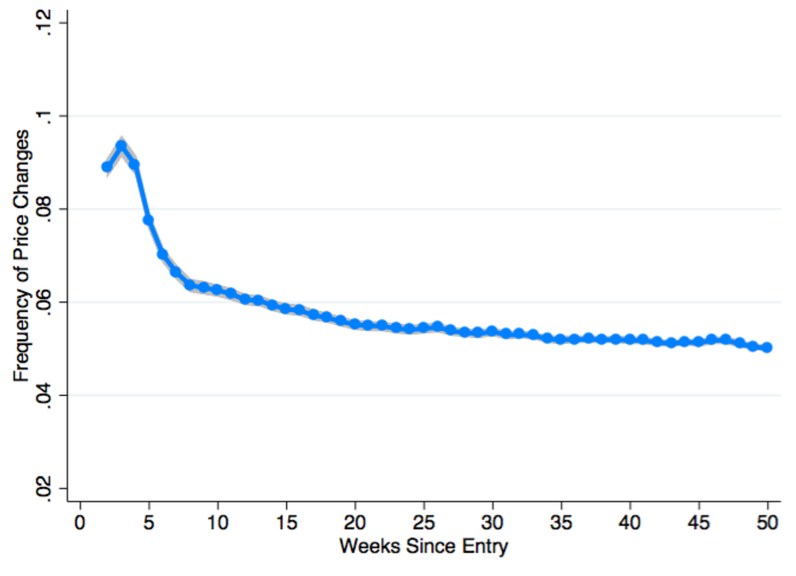

Note: The graph plots the average weekly frequency of price adjustments of entering products. The y-axis denotes the probability that the product adjusts its price in a given week and the x-axis denotes the number of weeks the product has been observed in the data after entry. The graph plots the coefficients for the age fixed effects of equation 1.1 where we use the regular price change indicator as the dependent variable. Equation 1.1 is computed by controlling for UPC-store and time fixed effects and the local unemployment rate to represent the cohort fixed effects. The calculation uses approximately 130 million observations and 2.5 million UPC $\times$ store pairs. The standard errors are clustered at the store level. The underlying data source is the Symphony IRI.

Figure 1.2: Absolute Value of Price Changes at Entry

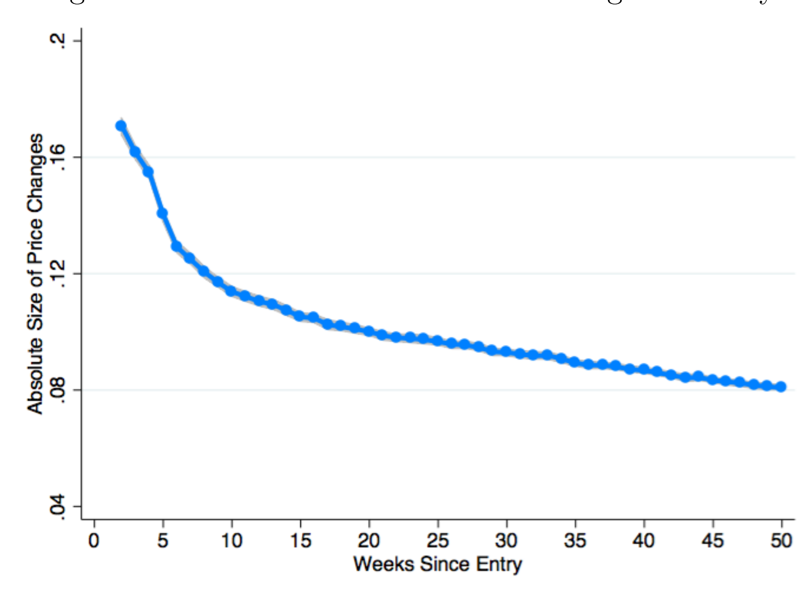

Note: The graph plots the average absolute size of price adjustments of entering products. The y-axis is the absolute value of the log price change in that week, and the x-axis denotes the number of weeks since the product entered. The graph plots the coefficients for the age fixed effects of equation 1.1 where we use the absolute value of the log price change as the dependent variable. Equation 1.1 is computed controlling for UPC-store, time fixed effects, and the local unemployment rate to represent the cohort fixed effects. The calculation uses approximately 5.8 million price changes and 2.5 million UPC $\times$ store pairs. The standard errors are clustered at the store level. The underlying data source is the Symphony IRI.

Figure 1.3: Fraction of Price Changes Larger than Two Standard Deviations

Note: 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 two years in the market. Source: IRI Symphony dataset

Figure 1.4: Pricing Moments by Waves

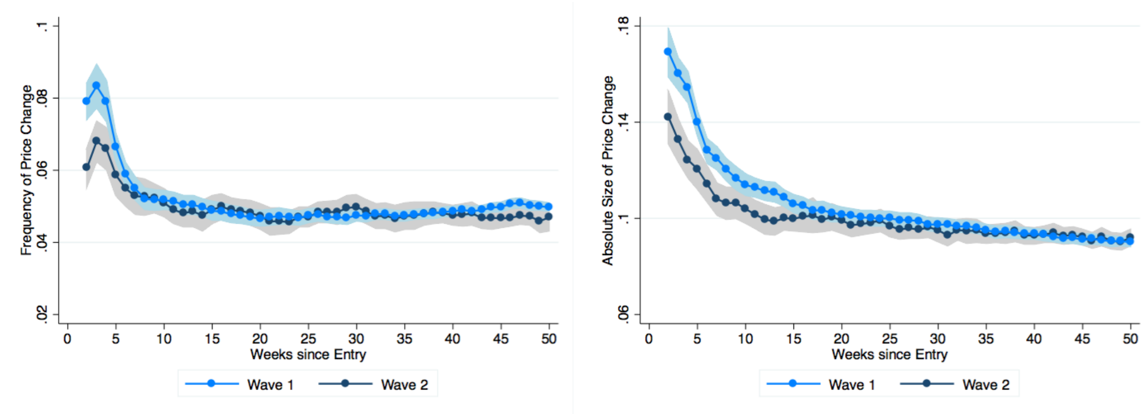

(a) Frequency of Price Changes (b) Absolute Size of Price Changes

Note: Panel (a) shows the probability of adjusting prices and panel (b) shows the absolute size of the price changes by waves. Wave 1 represents products that were launched at some location during a period in the first year since the product was introduced *nationally*. Wave 2 represents the same products when launched in different stores a year after their national entries. The graphs control for stores, product, and time fixed effects.

Figure 1.5: Numerical Example of the Two Period Model

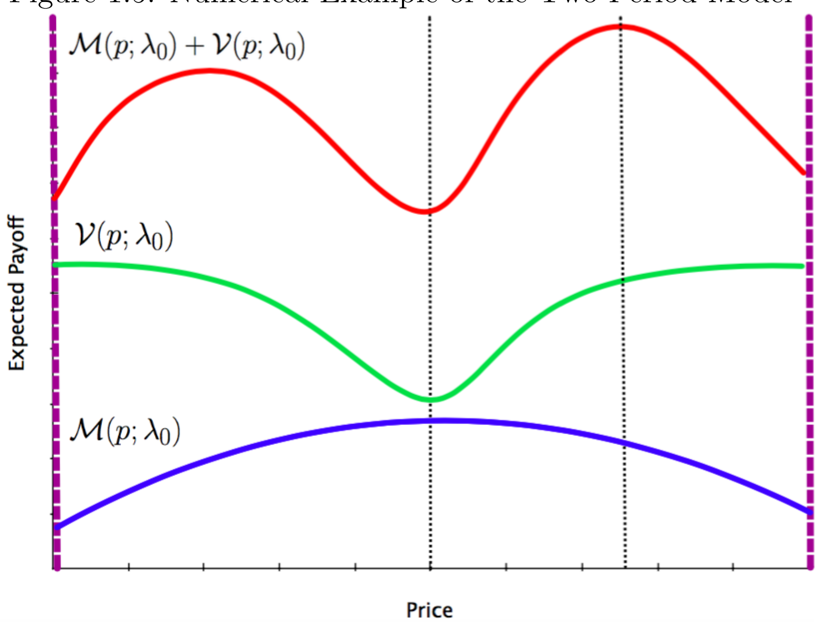

Note: Static profits $\mathcal{M}(p;\lambda_0)$, continuation value $\mathcal{V}(p;\lambda_0)$ and total payoff $\mathcal{M}(p;\lambda_0)+\mathcal{V}(p;\lambda_0)$ at $\lambda_0=\widehat{\lambda}$. The dotted purple lines represent the optimal prices P_2^* and P_1^*. $P^M(\lambda_0)$ represents the myopic policy and $P^*(\lambda_0)$ the policy under active learning.

Figure 1.6: Active Learning Regimes

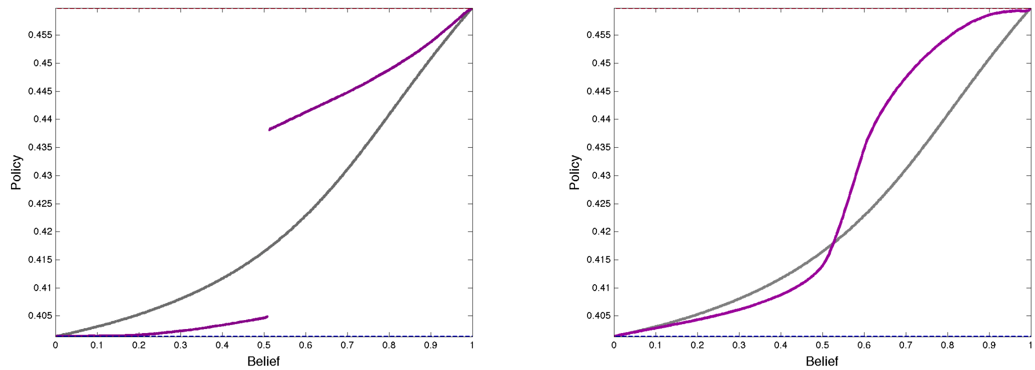

(a) Extreme Active Learning. (b) Moderate Active Learning.

Note: Panel (a) shows the extreme active learning regime and panel (b) its moderate counterpart. The gray line depicts the myopic policy $P^M(\lambda)$ and the purple lines the policy under active learning $P^*(\lambda)$. The dotted lines at the top and bottom of the panels indicate the optimal prices P_1^* and P_2^*.

Figure 1.7: Model vs Data

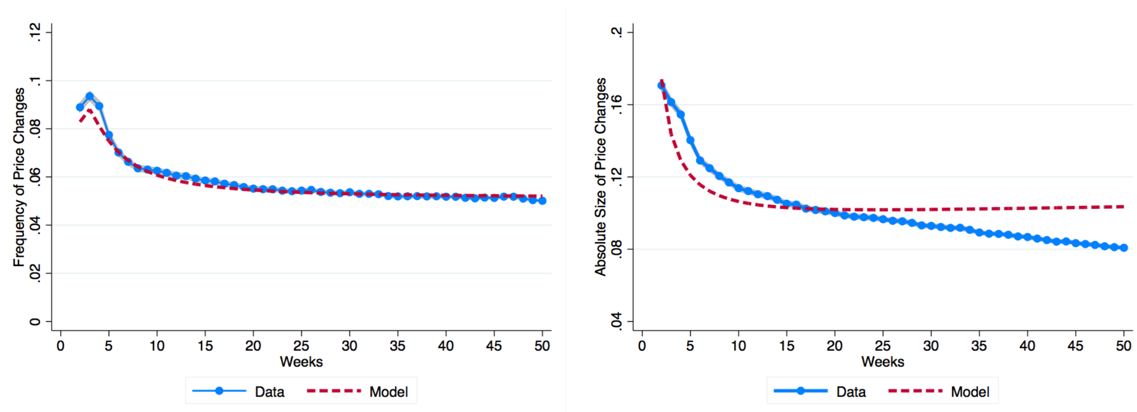

(a) Frequency of Price Changes (b) Absolute Size of Adjustments

Note: The figure shows the results of the model and compares them with the data. 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 the price changes over the life cycle of a product. The results of the frequency of price changes are shown in panel (a) and those of the absolute size of price changes are shown in panel (b).

Figure 1.8: Hazard of Price Change (Model)

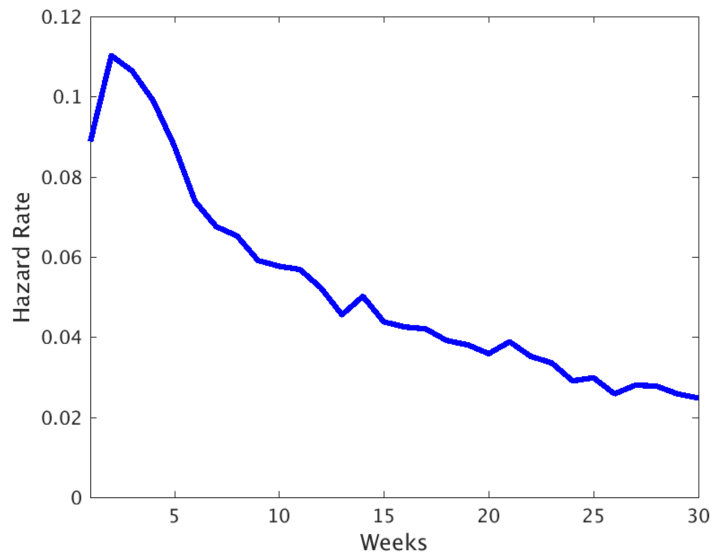

Note: The figure hazard of price changes generated by the model with active learning under our baseline calibration. It shows that the hazard is downward-sloping with a small hump at short durations. This is a consequence of the presence of a menu cost, which typically results in upward-sloping hazard rates, and the learning mechanism since new information at entry increases the likelihood of consecutive price changes.

Figure 1.9: Real Output Response to Nominal Shock

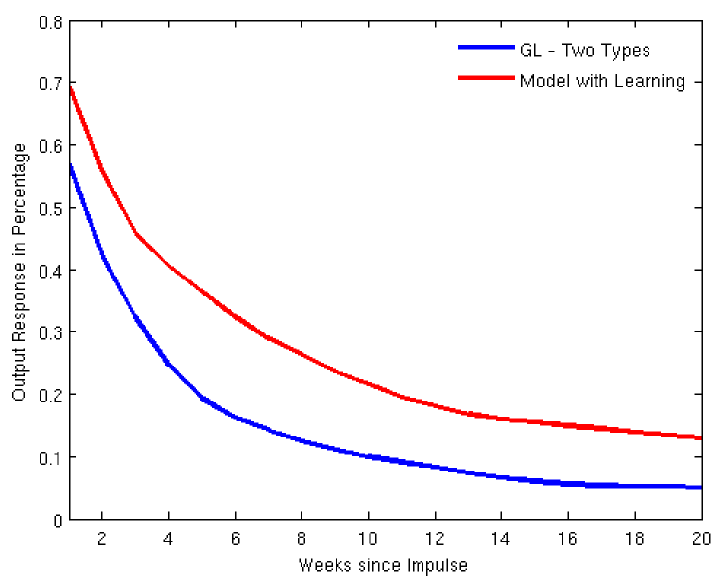

Note: The figure shows the response of the log real output to a 0.00038 increase in the nominal output growth rate. The output response is shown in the graph as a percent of the nominal shock. The red line depicts the output response in Golosov and Lucas (2007) with two different types of firms (i.e., σ_1 and $sigma_2$) and the blue line is the response in a price-setting model with active learning. Both models are calibrated to match the same moments and feature the same fraction of firms of each type.

CHAPTER 2

INNOVATION AND PRODUCT REALLOCATION IN THE GREAT RECESSION (WITH MUNSEOB LEE AND SARA MOREIRA)

2.1 Introduction

For decades, economists have identified product entry and exit as one of the key mechanism through which product innovation translates into economic growth (Aghion and Howitt, 1992; Grossman and Helpman, 1991; Aghion, Akcigit, and Howitt, 2014). But despite the important theoretical implications of product innovation, little is known empirically about about the process of the creation and destruction of a product, and how this process differs across different types of firms. In this paper, we study product reallocation across and within producers and how it evolved during the Great Recession. What is the role of product reallocation on output growth and quality improvements in the recent decade? How sensitive is innovation by new firms, small incumbents, and large incumbents to changes in aggregate economic conditions? New evidence on these questions will shed light on how resources are allocated to their best use within an economy and inform the recent debate on the sources of productivity slowdown in the US (Davis and Haltiwanger, 2014; Decker et al., 2014).

We begin by assessing the magnitude of product creation and destruction in the consumer goods sector over the period from 2007Q1 to 2013Q4. We use detailed product- and firm-level data at the barcode level and find that new products are systematically displacing existing products in the market. In our data set, a 12-digit number called the Universal Product Code (UPC) uniquely identifies each product, which is the finest level of disaggregation at the product level. Under this definition, firms reallocate more than 8 percent of the products in the economy every quarter. In this setting reallocation results from both the introduction of new products and the destruction of existing products. This is particularly relevant for large and well-diversified firms that sell products in several product categories. Consistent with several theories of creative destruction, we find that firms expanding, as well as firms contracting, contribute to the overall destruction of products. This source of dynamism in the US economy occurs within the boundaries of the firm, and as the result of the entries and exits of new firms. We find that most product reallocation is made by surviving incumbent firm that add or drop products in their portfolios.

After establishing the magnitude and pervasiveness of the reallocation of products, we evaluate the evolution during and after the Great Recession. We find that product reallocation is strongly pro-cyclical; the quarterly reallocation rate declined by more than 25 percent during the Great Recession. To better understand the sources of the cyclicality in the reallocation rate we decompose it in a within and a between firms component. We find that the cyclical pattern is overwhelmingly a consequence of within firm reallocation. In particular, most of the decline in reallocation within firms resulted from the decline in the creation of products during the recession.

In the second part of the paper we provide evidence that the decline in dynamism in the product market affected the economic growth and recovery after the Great Recession. Schumpeterian growth models have traditionally linked the speed of product reallocation to the innovation efforts of firms and to subsequent gains in productivity. To uncover the causes and consequences of the reallocation slowdown, we begin by establishing that the speed of product reallocation is strongly related to the innovation efforts of the firms as captured by their expenditures on research and development. This is consistent with theories featuring creative destruction where new and better varieties replace obsolete ones.

We then establish the relation between product reallocation and several innovation outputs such as revenue growth, improvements in products quality, and productivity growth. To do so, we follow Akcigit and Kerr (2010) and distinguish between two different types of innovation from the perspective of the firms: incremental innovations and extensions. Incremental innovations represent new products within the existing product lines of the firms, where they can use their capabilities and resources and benefit from economies of scale or scope. Extensions represent products outside the main business line of the firm. They are less common than incremental innovations because they represent larger innovations, which are likely to be more costly to develop. We find that incremental innovations have an immediate large impact on revenue. Extensions, on the other hand, are in general more innovative new products launched with higher average quality and have a higher impact on the total factor productivity (TFP) of the firm. In a similar way, we divide product exits into two types: products that are more likely to be terminated due to creative destruction (replaced by new products within the same product category) and those that were phased out due to the scaling down of firms' operations (products without replacement). Consistent with Schumpeterian theories[1], exits due to creative destruction are correlated with gains in TFP. Overall, we find that firms that have higher reallocation rates grow faster, launch products with higher average quality, and experience larger gains in productivity. Our evidence indicates that the decline in reallocation during the recession can explain around 15 percent of the drop in aggregate productivity in this period and had substantial implications for economic growth in the years that followed.

For most of our analysis, we rely primarily on the Nielsen Retail Measurement Services (RMS) scanner data set. It consists of more than 100 billion observations of weekly prices, quantities, and store information of approximately 1.4 million products identified at the UPC level. We combine the information on prices with the weight and volume of the product to compute unit values in order to approximate the quality of each product. In addition, we identify the firm owning each UPC by obtaining information from GS1, the single official source of barcodes in the United States. Our combined data set provides the revenue, price, quantity, and the quality for each product in a firm's portfolio and allows us to study how the within and between margins of product entry and exit evolve over time. Furthermore, we complement these data with measures of TFP and research and development expenses from Compustat. To the best of our knowledge, our paper is the first one to link the product-level information available in the Nielsen RMS with firm-level observables available in other data sets.

Our paper contributes to several active research areas. Despite the vast theoretical implications of product reallocation, the empirical analysis on the aggregate behavior of product reallocation lags far behind its theoretical counterpart due to data limitations. The literature on reallocation has focused on the input markets by using establishment and labor market data (Davis and Haltiwanger, 1992; Foster, Haltiwanger, and Krizan, 2001, 2006). By contrast, we focus on the reallocation in output markets. Importantly, we study the relative contribution of incumbents to

1. See Aghion, Akcigit, and Howitt (2014) for more detail.

the aggregate reallocation rate without inferring it from their job flow information.

Few papers have studied the degree of product reallocation directly. Bernard, Redding, and Schott (2010) study the extent of product switching within firms by using production classification codes (five-digit SIC codes), and Bernard and Okubo (2016) studies the role of product adding and dropping within Japanese manufacturing firms by using six-digit products according to the Japanese Standard Industrial Classification. Given the level of aggregation of their data, several firms could produce the same product. We substantially improve on this data by measuring products at a much finer level by using scanner data. With these data, we can explore the dynamics of each firms' unique portfolio of products as opposed to studying the dynamics of their product lines.

Our work is also closely related to Broda and Weinstein (2010) who study the patterns of product entry and exit using a similar data set to ours. But, they collect data from consumers rather than stores. Collecting data 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 rather than the products consumed by a sample of households. Therefore, with our data set, we can cover less frequently consumed goods and can provide less noisy measures of the entry and exit of products. Our paper builds on their work by examining between and within firm reallocation separately and by examining the contribution of each of these components during the Great Recession. Moreover, we examine the reallocation patterns of firms by subdividing them into several different dimensions: according to their size, their level of diversification (i.e., firms selling in a single product category versus firms selling in multiple product categories) and whether they are expanding or contracting at a given point in time.

Furthermore, by studying the connection between reallocation and different measures of innovation, our work links studies on reallocation that focus mainly on moving resources from less to more efficient uses to enhance productivity growth to the parallel literature on innovation (Klette and Kortum, 2004; Lentz and Mortensen, 2008; Akcigit and Kerr, 2010; Acemoglu et al., 2013; Garcia-Macia, Hsieh, and Klenow, 2016). Although we examine only the retail sector of the economy, to the best of our knowledge, our paper is the first to empirically establish the relation between product entry and exit and the innovation activities of a firm. In particular, we can empirically test and validate several predictions of Schumpeterian growth models with our matched data set; predictions that have been hard to examine in the past due to data availability issues.

Our work is also related to the literature on firm dynamics that studies the propagation of aggregate shocks after large contractions in output (Caballero and Hammour, 1994; Moreira, 2016). We find that both the reallocation rate and the entry rate of products suffered a persistent decline after the Great Recession. The decline in product creation had consequences in terms of revenue for the firms in the short run. But, more importantly, this missing generation of products, in the spirit of Gourio and Siemer (2014), combined with the evidence we provide on the relation between reallocation and productivity growth, can have substantial implications for the slow recovery experienced by the US economy in the years following the Great Recession.

Further, our paper complements the growing literature on how business formation and product creation amplifies business cycles (Chatterjee and Cooper, 1993; Jaimovich and Floetotto, 2008; Bilbiie, Ghironi, and Melitz, 2012; Minniti and Turino, 2013). Our estimations can be used to discipline the parameters these models use to replicate the number of firms and products at different stages in the business cycle. More importantly, the evidence we present emphasizes the endogenous interaction between the innovation efforts of firms, their product scope, and outcomes such as revenue and productivity. Our work highlights the importance of multiproduct firms in business cycle modeling and the role of firms heterogeneity in understanding the degree to which macroeconomic shocks propagate in the economy.

The rest of the paper is organized as follows. Section 2.2 presents the data and describes our procedure to link our product-level data set with the firm-level information available in Compustat. In section 2.3, we define reallocation and provide several decompositions to explore the relative contributions of the between and within margins. In this section we also provide an interpretation of the magnitudes of the reallocation rate we observe and describe its evolution during the Great Recession. In section 2.5 we examine the possible determinants of relocation. We examine its relation with R&D and define incremental innovations and extensions along with exits due to creative destruction and terminations due to firms scaling down their operations. Section 2.6 tests and validates the predictions of the models involving creative destruction and shows the relations between reallocation and revenue growth, quality improvements, and productivity dynamics. Section 2.7 concludes. We include several robustness tests and additional empirical findings in the appendix.

2.2 Data Description

2.2.1 Baseline Product-Level Dataset

We rely primarily on the Nielsen Retail Measurement Services (RMS) scanner data set that is provided by the Kilts-Nielsen Data Center at the University of Chicago Booth School of Business. The RMS consists of more than 100 billion unique observations at the week × store × UPC level. Each individual store reports weekly prices and quantities of every UPC code that had any sales volume during that week.[2]

The data is generated by point-of-sale systems and contains approximately 40,000 distinct stores from 90 retail chains across 371 MSAs and 2,500 counties between January 2006 and December 2014. The data set comprises around 12 billion transactions per year worth, on average, \$220 billion. Over our sample period the total sales across all retail establishments are worth approximately \$2 trillion and represent 53% of all sales in grocery stores, 55% in drug stores, 32% in mass merchandisers, 2% in convenience stores, and 1% in liquor stores.

The baseline data consist of approximately 1.64 million distinct products identified by UPC. The data is organized into 1,070 detailed product modules that are aggregated into 114 product groups that are then grouped into 10 major departments.[3] For example, a 31-ounce bag of Tide

2. In comparison to other scanner data sets collected at the store level, the RMS covers a much wider range of products and stores. Table G.I in the appendix shows that in comparison to the IRI Symphony data set, a similar data set widely used in the academic literature, the RMS covers 14 times more products in a given year. In terms of revenue the RMS represents roughly 2 percent of total household consumption whereas the IRI Symphony is 30 times smaller. In comparison to scanner data sets collected at the household level, the RMS also has a wider range of products because it reflects the universe of transactions for the categories it covers as opposed to the purchases of a sample of households. The Nielsen Homescan, for example, that contains information on the purchases of 40,000-60,000 US households covers less than 60% of the products the RMS covers in a given year.

3. The ten major departments are: Health and Beauty aids, Dry Grocery (e.g., baby food, canned vegetables), Frozen Foods, Dairy, Deli, Packaged Meat, Fresh Produce, Non-Food Grocery, Alcohol, and General Merchandise).

Pods has UPC 037000930389 and is produced by Procter & Gamble and is mapped to product module "Detergent-Packaged" in product group "Detergent", which belongs to the "Non-Food Grocery" department. Each UPC contains information on the brand, size, packaging, and a rich set of product features. We use the weight and the volume of the product to compute unit values.

Our data set combines all sales at the national and quarterly level, although we also conduct some exercises at the annual frequency given that some firm-level observables are only available at that frequency. For each product j in quarter t, we define revenue r_{jt} as the total revenue across all stores and weeks in the quarter. Likewise, quantity q_{jt} is defined as total quantities sold across all stores and weeks in the quarter. Price p_{jt} is defined by the ratio of revenue to quantity, which is equivalent to the quantity weighted average price.

A critical part of our analysis is the identification of entries and exits. For each product we use the panel structure to identify the entry and exit periods. In addition, we follow Broda and Weinstein (2010) and Argente and Yeh (2017) and use the UPC as the main product identifier. This is because it is rare that a meaningful quality change occurs without resulting in a UPC change. A concern that can arise from this assumption is that a new UPC might not always represent a new product. For instance, Chevalier, Kashyap, and Rossi (2003) notes that some UPCs might get discontinued only to have the same product appear with a new UPC. This is not a concern in our data set because Nielsen detects these UPCs and assigns them their prior UPC.

We define entry as the first quarter of sales of a product and exit as the quarter after we last observe a product being sold. To study the patterns in the entry and exit rates, we use information for all products in the period from 2007Q1 to 2013Q4, that include cohorts born from 2007Q1 to 2013Q4 and cohorts born before that period, from whom we cannot determine the cohort and age.[4] In addition, given that our estimates of products entries and exits might be affected by the entries and exits of stores in the sample, we consider only a balanced sample of stores during our sample period.

4. Note that we excluded the first four and last four quarters of the sample. Because we define entry as the first quarter of sales of a product and exit as the first quarter after we last observe a product being sold, we could identify an abnormally high entry in the first quarters and abnormally high exit in the last quarters. Our procedure ensures that we only classify a product as entering if it was not observed for at least a full year before, and a product as exiting if we no longer observe it for at least a full year past exit.

In order to minimize concerns of potential measurement error in the calculation of a products entry and exit, our baseline sample excludes private label products, considers products with at least one transaction per quarter after entering, and excludes products in the Alcohol and General Merchandise departments. We exclude private label goods because, in order to protect the identity of the retailer, Nielsen alters the UPCs associated with private label goods. As a result, multiple private label items are mapped to a single UPC that makes it difficult to interpret the entry and exit patterns of these items since it is not possible to determine the producer of these goods. We consider products without missing quarters to rule out the possibility that our results are driven by seasonal products, promotional items, or products with very small revenue. And, finally, we exclude the two departments for which the coverage in our data is smaller and less likely to be representative.

Our final sample is described in Table 2.1. On average, more than 222,000 distinct UPCs are present in our sample each year. Most products have revenue of less than $500 per year but almost 2.5% of the products make more than $50 million. A product module contains approximately 242 products, a product group 2,486 products, and a department 25,688 products on average. The table shows that these numbers remain very stable before, during, and after the Great Recession.

Nonetheless, all of the results that follow are robust to using the full sample of products that are available in the RMS. We present these results in Appendix G. Lastly, Appendix G also includes results where, instead of using the barcode as the main product identifier, we identify products using a broader definition using the product attributes provided by Nielsen as in Kaplan and Menzio (2015). Under this alternative definition, a good is the same if it shares the same observable features, the same size, and the same brand, but may have different UPCs. We use this definition to minimize the concern that new products, when identified by their UPCs, represent only marginal innovations from the perspective of the firms. Under this definition, each new entry represents at least a new product line for the firm. Appendix G shows that the results we describe below on the aggregate reallocation rate and on the impact of product reallocation on several innovation outputs remain very similar under this specification.

2.2.2 Matching Firm and Products

We link firms and products with information obtained from GS1 US, the single official source of UPCs. In order to obtain a UPC, firms must first obtain a GS1 company prefix. The prefix is a five- to ten-digit number that identifies firms and their products in over 100 countries where the GS1 is present. The number of digits in a company prefix indicates different capacities for firms to create UPCs. For example, a ten-digit prefix allows firms to create ten unique UPCs, and a six-digit prefix allows them to create up to 100,000 unique UPCs. Although the majority of firms own a single prefix, it is not rare to find that some own several. Small firms, for example, often obtain a larger prefix first, which is usually cheaper, before expanding and requesting a shorter prefix.[5] Larger firms, on the other hand, usually own several company prefixes due to past mergers and acquisitions. For example, Procter & Gamble owns the prefixes of firms it acquired such as Old Spice, Folgers, and Gillette. For consistency, in what follows we perform the analysis at the parent company level.

Given that the GS1 US data contains all of the company prefixes generated in the US, we combine these prefixes with the UPC codes in the RMS.[6] With this data set, we can compute the revenue, price, quantity, and quality of each product in a firm's portfolio to study how the within and between margins of product creation and destruction evolve over time.

Table 2.1 describes the characteristics of the firms in our data. We have a yearly average of 22,356 firms with slightly more firms present after the recession. Similar to the size distribution of products, the size distribution of firms is fat-tailed.[7] In addition, most firms are well diversified: 26% of the firms own a single product, 26% are multi-product firms that belong to a single module, 13% are multi-module firms that belong to a single product group, 19% sell in multiple product groups but in a single department, and 16% are multi-department firms.

5. Previous studies, including Broda and Weinstein (2010), have assumed that the first six digits of the UPC identify the manufacturing firm. This assumption is valid for 93% of the products in our sample.

6. Less than 5 percent of the UPCs belong to prefixes not generated in the US. We were not able to find a firm identifier for those products.

7. Table G.II in the appendix shows the top 20 firms in terms of revenue in our data. The top 10 firms alone account for approximately 27% of the total revenue.

2.2.3 Matching Nielsen RMS and Compustat

For the later analysis, we obtain firm-level characteristics from Compustat. To combine the Nielsen data with the Compustat database, we match the names in the GS1 to those in Compustat using a string matching algorithm that is described in Schoenle (2017). After applying the algorithm, we match 479 publicly traded firms over our sample period.[8] Our matched sample represents 22% of the total sales in Compustat and 45% of the total revenue in the RMS. Approximately 21% of the total number of products in the data belong to publicly traded firms. We describe in detail the construction of firm-level variables in Appendix D and the summary statistics in Table G.III in the appendix.

2.3 Reallocation of Products

2.3.1 Measurement of Reallocation

In this section we document several new stylized facts on the level and evolution of product creation, destruction, and reallocation in the U.S. consumer goods sector. We start with a description of the measures that we use to identify the aggregate levels and cyclical patterns of product reallocation. Most products' entries and exits do not necessarily translate into entry and exit of firms because the majority of products are produced by multi-product firms (Table 2.1). In order to study the degree of heterogeneity in this sector, we also compute the firm-specific reallocation rates for products.

Aggregate reallocation: To capture the importance of product entry and exit, we use information on the number of new products, the number of exits of products, and the total number of products for each firm i over time t, and define the aggregate entry and exit rates as follows:

$$n_t = \frac{\sum_i N_{it}}{\sum_i T_{it}} \tag{2.1}$$

8. A few public firms in our sample are conglomerates combining more than one independent corporation. For the later analysis, we combine their information into a single firm to perform our reduced-form analysis at the public firm level.

$$x_t = \frac{\sum_i X_{it}}{\sum_i T_{it-1}} \tag{2.2}$$

where N_{it}, X_{it}, and T_{it} are the numbers of entering products, exiting products, and total products, respectively. The entry rate is defined as the number of new products in period t as a share of the total number of products in period t. The exit rate is defined as the number of products that exited in period t (i.e., the last time we observe a transaction was in $t-1$) as a share of the total number of products in period $t-1$.[9]

Two relations link these concepts: the net growth rate of the stock of available products equals the entry rate minus the exit rate; the overall change in the portfolio of products available to consumers can be captured by the sum of the entry and exit rates. We refer to this last concept as the product reallocation rate, in particular:

$$r_t = n_t + x_t \tag{2.3}$$

With this measure we can measure the extent of the changes in the status of a product in our data, either from the entry or the exit margin.

Average within firm reallocation: Using information on the numbers of entering products, exiting products, and total products by each firm i over time t, we can define the average reallocation of products by firms as the (unweighted) mean entry and exit rates across all firms as follows:

$$\bar{n}_t = \frac{1}{\gamma_t} \sum_{i=1} n_{it} \tag{2.4}$$

$$\bar{x}_t = \frac{1}{\gamma_{t-1}} \sum_{i=1} x_{it} \tag{2.5}$$

where $n_{it} = \frac{N_{it}}{T_{it}}$, $x_{it} = \frac{X_{it}}{T_{it-1}}$, and γ_t is the number of firms active in t. The average reallocation rate of firms is then defined as:

$$\bar{r}_t = \bar{n}_t + \bar{x}_t \tag{2.6}$$

9. The main advantage of assigning a product exit to the quarter following the last observed transaction of a product is that we can define relative changes in the stock of products as the difference between entry and exit rate.

Aggregated and average within firm reallocation: The aggregate level of reallocation and the average level within can be easily related following an Olley and Pakes decomposition. The aggregate reallocation rate is composed of the average reallocation and a component that measures the covariance between the market share and reallocation rates:

$$r_t = \bar{r}_t + \sum_{i \in \Gamma_t} (r_{it} - \bar{r}_t)(t_{it} - \bar{t}_t) \tag{2.7}$$

where $t_{it} = \frac{T_{it}}{\sum_i T_{it}}$ measures the product share of firm i at t, $t_{it} \geq 0$ and sums to one; and Γ_t is the set of active firms in t. The second component of the decomposition captures whether firms with more products are more likely to be those with high or low reallocation rates.

2.3.2 Magnitude and Heterogeneity of Product Reallocation

The rates of aggregate product creation and destruction are remarkably large. Table 2.2 shows that, on average, 8 percent of all products are reallocated every quarter in the period from 2007 to 2013. This amount means that about one in three products are either destroyed or created over a typical 12-month interval. This fact highlights the fluidity in the consumer goods sector.

The level of reallocation depends on the product definition. In our baseline sample, products are defined at the UPC level for a set of consumer goods industries that excludes generics, alcohol, and general merchandise. In the alternative sample, where both generics and general merchandise are included, we observe an average quarterly reallocation rate of 7.6 percent, which is very close to the 7.9 percent that we observe in the baseline sample. The alternative sample, for the same universe of goods, but for the more coarse definition of product (as defined at the level of different module and brand), we still find an average quarterly reallocation rate of 4.7 percent. This percentage means that while some creations and destructions of products might involve small changes in their characteristics, a big share of reallocation happens with the creation and destruction of new brands.

Our measures of product reallocation can be compared with measures of reallocation at the production unit level and input level. Using data from Business Dynamics Statistics, we compute analogous measures of firm reallocation using information on the entries and exits of establishments.

We find that, during the same period, the entries and exits of establishments are about 20 percent of total establishments over a one-year period. The reallocation of establishments weighted by employment is about 9 per cent per year. In our dataset, we observe entry and exit of firms of about 17 percent over a one-year period (Table 2.1), which is similar to the whole economy's reallocation of firms. Foster, Grim, and Haltiwanger (2016) find that the evolution of job reallocation, computed as defined by Davis, Haltiwanger, and Schuh (1996), points to an average level of about 13 percent a quarter over the period from 2006 to 2012, for a total of around 150 million jobs.

Over the 2007 to 2013 period, the quarterly entry rate of products was 4.3 percent, and the quarterly exit rate was 3.6 percent. These rates mean that over a typical 12-month interval, about one in five new products are created in these sectors, and about one in six are no longer available (Table 2.2). Overall, while the growth rate of products in the consumer goods sector increased almost 1 percent per quarter over this period, both the entry and exit margins are important in explaining the changes in the portfolio of products available to consumers.

To better understand the sources of reallocation, we examine the degree of heterogeneity in the firm-specific reallocation of products. Table 2.3 shows the average quarterly reallocation, entry, and exit rates for the period from 2007 to 2013. On average, firms in the consumer goods industry add or drop about 10.8 percent of the products in their portfolios. The fact that this rate is larger than the aggregate reallocation rate means a negative covariance exists between reallocation rates and product shares, that is, firms that produce a lower number of products have, on average, higher reallocation rates. This negative covariance is driven entirely by the entries and exits of firms. Most products are produced by multi-product firms, and thus the entries and exits of firms only account for a small share of product reallocation (only 1 out of 20 products are created and destroyed by entering or exiting firms).

Over the period from 2007 to 2013, the average firm-specific quarterly entry rate of products was 5.5 percent, and the average quarterly exit rate was 5.6 percent. We classify firms by their net creation of products (expanding, contracting, and unchanged), and access to the market (entering, exiting, and incumbent).[10] Overall, most additions of products are made by expanding firms and

10. Appendix B provides details on the disaggregation.

most product destructions are made by contracting firms. Table 2.3 shows that expanding firms add around one product out of every three (one out of every four if we exclude entering firms). As expected, firms that are reducing the total number of products on net are adding products at a smaller rate (only about 1.3 percent, on average). Expanding firms destroy products at a rate of about 2 percent, while firms that are destroying products on net phase out about 33.6 percent of their products (23.4 percent when we exclude exiting firms).

There is substantial heterogeneity in the size of firms that produce consumer goods. We classify firms by their quartile of revenue, and we measure the contribution of each group to the aggregate reallocation. The average reallocation rates among incumbent firms by revenue quartile are slightly larger among high revenue firms, which hold several products on average, and thus an overwhelmingly large share of products created or destroyed every quarter originate in firms in the top quartile of the distribution of revenue.

Another important source of heterogeneity in this industry is the degree of diversification of products between firms (Table 2.1). Single-product firms have higher rates of product reallocation because they are also more likely to be entering or exiting firms. When we exclude single-product firms, diversified firms (in particular, multi-department firms) have slightly larger average rates of reallocation, and thus diversified firms make a higher contribution to the aggregate reallocation of products (Table 2.3).

2.3.3 Evolution of Product Reallocation in the Great Recession

After examining the sources of heterogeneity in the product reallocation rates, we analyze the evolution of our measures of product reallocation over the business cycle. The main takeaway from this analysis is that the reallocation of products in the period under analysis is very procyclical. The share of products that were created or destroyed was approximately 9.4 percent on average during 2007, dropping to about 7.0 percent on average during 2010, and recovering to 7.8 percent three years later (Table 2.2).

A significant fraction of this cyclical component is explained by the variation in the number of new products that firms created during the Great Recession. The quarterly entry rate declines from around 4.7 percent to about 3.7 percent in the period from 2007 to 2010, followed by a full recovery by 2013. The aggregate exit rate trends downwards during this period and the deviations from trend are also pro-cyclical. The aggregate quarterly exit rate varies from 4.6 percent to 3.3 percent from 2007 to 2010, followed by a very tepid decline until the end of 2013.

This evolution contrasts with the evidence in Broda and Weinstein (2010). Their period of analysis includes the 2001 recession and they find that the aggregate creation of products is pro-cyclical, while the aggregate product destruction is countercyclical, although the magnitude of the latter is quantitatively less important. This pattern indicated that product reallocation was only slightly pro-cyclical. We find the same strong pro-cyclicality in the entries of new products but we do not find any evidence of counter-cyclicality in the exit rate. Our findings differ from those in Aghion et al. (2017) who report countercyclical product churn after using data from the US Census of Manufacturers. Because the US Census of Manufacturers is only available in years ending in 2 and 7, their measure can only capture a period of recession followed by a recovery and is unlikely to capture product dynamics occurring during the Great Recession. Our findings of a strong decline in product reallocation during the Great Recession and the subsequent slow recovery are similar to the evolution of job creation and destruction documented in Foster, Haltiwanger, and Krizan (2001). In the Great Recession, job creation fell by as much or more than the increase in job destruction. In this respect, the Great Recession was not a time of increased reallocation. These patterns also contrast with the responsiveness of job creation and destruction in prior recessions. In prior recessions, periods of economic contraction had a sharp increase in job destruction and a mild decrease in job creation.[11]

The aggregate cyclical patterns of the product reallocation rates are pervasive across different types of firms. We find that during our period of analysis, the strong decline in the reallocation rates during the Great Recession is present across all types of firms. Nonetheless, we also find some evidence of systematic heterogeneity as some firms are more procyclical than others. The decline in average reallocation in 2008 and 2009 was larger among firms that reduced their stock of products; such decline results from decreases in both entry rates and exit rates (Table 2.2).

11. As highlighted in Davis, Haltiwanger, and Schuh (1996), the greater responsiveness of job destruction relative to job creation in these earlier cyclical downturns means that recessions are times of increased reallocation.

When we sort all firms based on quartiles of revenue, we find that all quartiles experience a decline in product reallocation during the Great Recession that is mostly explained by the evolution of the rate at which firms create products. The decline in reallocation was particularly large among low revenue firms and resulted from the decline in both entry rates and in exit rates. We also find, however, that after the Great Recession the product reallocation rates of the lower quartile of revenue show a greater rebound. The cyclical evolution of the product reallocation rates for both diversified and undiversified firms is similar over the period, and does not exhibit substantial differences.

2.4 Decomposition of Reallocation

Next, we apply decomposition methods to shed further light on the evolution of our product reallocation measures and explain what economic forces drive the evolution of this rate. The literature that examines the aggregate productivity in the economy has developed decomposition methods to investigate the sources of productivity change. Aggregate productivity is typically computed as a weighted average of productivity at the producer level. Because the productivity levels of producers are heterogenous, aggregate productivity changes over time can reflect both shifts in the distribution of producer-level productivity and changes in the composition of firms. In turn, changes in the composition of firms in the economy can result not only from changes in market shares among surviving firms, but also from the entry of new producers and the exit of old ones. These three sources of changes in the composition are often named the effect of reallocation of producers in the economy.

We borrow from this literature and apply these methods to our setting. Our goal is to decompose changes in the aggregate rate of product reallocation between changes in the re-allocative behavior of firms and changes in the distribution of firms. The idea is that product reallocation can evolve both because the incumbent firms change their behavior or because firms enter and exit markets. In our case, incumbent firms can increase the rate at which they add or destroy products, while their share of products varies over time, that is, firms that reallocate more might be gaining or losing overall market share. We use these methods to identify the main sources that explain the decline in reallocation during the Great Recession and in the post-recession period.

2.4.1 Decomposing Changes in Reallocation: Accounting for Entry and Exit of firms

Using equation 2.7, we can decompose the changes in reallocation between quarter t and a reference quarter 0, $\Delta r_{t,0} = r_t - r_0$, as follows:

$$\Delta r_{t,0} = \bar{r}_t - \bar{r}_0 + \sum_{i \in \Gamma_t} (r_{it} - \bar{r}_t)(t_{it} - \bar{t}_t) - \sum_{i \in \Gamma_0} (r_{i0} - \bar{r}_0)(t_{i0} - \bar{t}_0)$$

where $t_{ik} = \frac{T_{ik}}{\sum_i T_{ik}}$ measures the product share of firm i in quarter k, $t_{ik} \geq 0$ and sums to 1, and Γ_k is the set of active firms in k, $k = t, 0$. After simplifying the notation, we express this decomposition in the following components:

$$\Delta r_{t,0} = \Delta \bar{r}_{t,0} + \Delta \sum_{i \in \Gamma_t} (r_{it} - \bar{r}_t)(t_{it} - \bar{t}_t)) \tag{2.8}$$

The first component represents changes in the average reallocation rate within firm, and the second component is the adjustment by differences in size across firms. It is a natural way to capture changes in the first moment of the distribution of entry rates, and changes in market share reallocation via changes in the covariance. Thus, the evolution in reallocation rates of products can come from changes in the average within firm reallocation rate, and changes in the distribution of products across firms that reallocate more or less intensively.

Melitz and Polanec (2015) proposed an extension of the Olley and Pakes decomposition to accommodate entry and exit of firms, such that we can separately obtain the contribution of continuing, entering and exiting firms. The underlying idea is that we can write the change in reallocation rates as:

$$\Delta r_{t,0} = \bar{r}_t^{C_{t,0}} - \bar{r}_0^{C_{t,0}} + \sum_{i \in C_{t,0}} (r_{it} - \bar{r}_t)(t_{it} - \bar{t}_t) - \sum_{i \in C_{t,0}} (r_{i0} - \bar{r}_0)(t_{i0} - \bar{t}_0)$$

$$+\sum_{i\in EN_{t,0}} t_{it}(\sum_{i\in EN_{t,0}} \frac{t_{it}}{\sum_{i\in EN_{t,0}} t_{it}} r_{it} - \sum_{i\in C_{t,0}} \frac{t_{it}}{\sum_{i\in C_{t,0}} t_{it}} r_{it})$$

$$-\sum_{i\in EX_{t,0}} t_{i0}(\sum_{i\in EX_{t,0}} \frac{t_{i0}}{\sum_{i\in EX_{t,0}} t_{i0}} r_{i0} - \sum_{i\in C_{t,0}} \frac{t_{i0}}{\sum_{i\in C_{t,0}} t_{i0}} r_{i0}) \tag{2.9}$$

where the contribution of each firm to the aggregate change in the reallocation rate is separated into three categories for continuing $C_{t,0}$, entering $EN_{t,0}$ and exiting $EX_{t,0}$ firms. The first terms of the decomposition apply the Olley and Pakes decomposition to the the subset of surviving firms, that is decomposed between the change in the average reallocation rate among continuing firms and the change in the covariance between the product share and the reallocation rate. The latter two terms measure the contribution of entry and exit of firms to the aggregate change in the reallocation rates. The entry component is defined as the weighted average difference between the reallocation rate of entrants and reallocation rate of continuers. The exit component is defined as the weighted average difference between the reallocation rate of exit firms and reallocation rate of continuers.

An alternative approach to identify the importance of the different margins that can potentially generate changes in the aggregate product reallocation is to explore the equality $r_t = \sum_{i\in\Gamma_t} r_{it}t_{it}$, and we can write the changes as

$$\Delta r_{t,0} = \sum_{i\in C_{t,0}} (r_{it}t_{it} - r_{i0}t_{i0}) + \sum_{i\in EN_{t,0}} r_{it}t_{it} - \sum_{i\in EX_{t,0}} r_{i0}t_{i0}$$

For continuing firms, we can further disentangle between the sum of the changes in the reallocation rate, holding firms' shares of the product market constant (within-firm component), and the percentage sum of shares changes holding all firms' entry constant (between-firm component). The decomposition will be then:

$$\Delta r_{t,0} = \sum_{i\in C_{t,0}} t_{i0}(r_{it} - r_{i0}) + \sum_{i\in C_{t,0}} r_{it}(t_{it} - t_{i0}) + \sum_{i\in EN_{t,0}} r_{it}t_{it} - \sum_{i\in EX_{t,0}} r_{i0}t_{i0} \tag{2.10}$$

For continuing firms, the first component captures changes in the reallocation rate within them, while the second captures the contribution of changes in product shares between them. Under this decomposition, entry (exit) of firms has a positive (negative) contribution. In order to address this, Griliches and Regev (1995) redefines the decomposition above such that the average aggregate reallocation rate is the reference $\bar{r}_{0,t} = \frac{r_0+r_t}{2}$. The decomposition is then given by

$$\Delta r_{t,0} = \sum_{i\in C_{t,0}} t_{it}(r_{it} - \bar{r}_{0,t}) - \sum_{i\in C_{t,0}} t_{0,t}(r_{i0} - \bar{r}_{t0}) + \sum_{i\in EN_{t,0}} t_{it}(r_{it} - \bar{r}_{0,t}) - \sum_{i\in EX_{t,0}} t_{i0}(r_{i0} - \bar{r}_{0,t})$$

And we can split the contribution of continuing firms between within and between components as follows

$$\Delta r_{t,0} = \sum_{i\in C_{t,0}} \bar{t}_{i,0t}(r_{it} - r_{i0}) + \sum_{i\in C_{t,0}} (\bar{r}_{i,0t} - \bar{r}_{0,t})(t_{it} - t_{i0}) + \sum_{i\in EN_{t,0}} t_{it}(r_{it} - \bar{r}_{0,t}) - \sum_{i\in EX_{t,0}} t_{i0}(r_{i0} - \bar{r}_{0,t}) \tag{2.11}$$

where $\bar{r}_{i,0t} = \frac{r_i0+r_it}{2}$ and $\bar{t}_{i,0t} = \frac{t_i0+t_it}{2}$. The contribution of the within-firm component among surviving firms is now weighted by the average product share of each firm, while the between-firm contribution is weighted by the average reallocation rate. The main advantage of this last decomposition is that the contribution of entrants can now be negative, and the contribution of exits can be positive.

Foster, Haltiwanger, and Krizan (2001) proposes a slightly modified decomposition, where the reference level is period 0 instead of a time varying average. This approach facilitates comparisons across different time periods. The third contribution of the surviving firms is the cross-firm component, that captures the covariance between the change in the share of products and the change in entry rate. The decomposition is then given by

$$\Delta r_{t,0} = \sum_{i\in C_{t,0}} \bar{t}_{i0}(r_{it} - r_{i0}) + \sum_{i\in C_{t,0}} (r_{i0} - r_0)(t_{it} - t_{i0}) + \sum_{i\in C_{t,0}} (r_{it} - r_{i0})(t_{it} - t_{i0})$$

$$+ \sum_{i\in EN_{t,0}} t_{it}(r_{it} - r_0) - \sum_{i\in EX_{t,0}} t_{i0}(r_{i0} - r_0) \tag{2.12}$$

Similar to the decomposition above, the contribution of entry and exit can be negative or positive, depending on how the reallocation rates among entrants and exiters compare with the reallocation rate in the baseline period 0.

2.4.2 Results

Table 2.4 reports the results of the decompositions. We apply them to changes in the aggregate entry, exit, and reallocation rates. In particular, we report the decomposition for the Great Recession by adding the cumulative one-quarter changes between 2007Q1 and 2009Q4 and for the period following the Great Recession by adding the cumulative one-quarter changes between 2010Q1 and 2012Q4.[12]

First, we present the results for the method developed by Olley and Pakes (1996). This method does not accommodate firm entry and exit but is used as a reference and baseline for the other methods. During the Great Recession, the weighted average reallocation rate declines by around 3.6 percentage points, and is decomposed into a change of -5.3 percentage points in the first moment of firms' reallocation distribution (the unweighted mean), and an increase of 1.6 percentage points in the joint distribution of reallocation and market shares (the covariance between reallocation and product shares). This means that the Great Recession saw a substantial decline in the average reallocation rates, and that firms reallocating more were increasing their product share relative to firms reallocating less. In the post-recession period, the aggregate reallocation increased by 1.3 percentage points as a result of a recovery of 2.5 percentage points in the average reallocation of firms, and a 1.2 percentage point decline in the covariance between product shares and reallocation.

Next, we implement the methodology developed by Melitz and Polanec (2015) to further understand the contributions of the entries and exits of firms to product reallocation rates. The results show that the decline in the average product reallocation rate during the Great Recession was partially offset by a 0.9 percentage points increase in the reallocation rate from net entry of firms (which in its turn is further decomposed into 3.1 percentage points stemming from the entry of products from entering firms, and -2.2 percentage points explained by the exit of products from exiting firms). In the period following the Great Recession, net entry contributed 1.3 percentage points to the recovery in the aggregate product reallocation rates.[13] The contribution

12. The two periods correspond to a 12-quarter (3-year) overall change. We select these particular dates to match the overall evolution that we observe for the aggregate reallocation

13. It is worth pointing out that the sign of the contribution of entry is always positive and the size of exit is always negative, given that the reallocation rates of entering and exiting firms are by definition equal to 1, while for surviving firms is closer to the level of 0.1.

of the net entry of firms in the recession and post-recovery periods is positive and very similar, which indicates that the distinct evolution of the reallocation in those periods was mainly driven by surviving incumbents firms, which seems to be the group that was more dynamic in adjusting their re-allocative behavior.

When we adapt the standard within and between decompositions to the product reallocation rate, we obtain similar results for the impact of entry and exit of firms. The Griliches and Regev (1995) decomposition shows that the decline in reallocation in the recession period results in -4.0 percentage points decline in the rate of product reallocation within surviving firms and -0.4 percentage points from variation between surviving firms. This decomposition indicates that there is almost not between effect, that is, the market share of high reallocation firms is very similar. In the post-recession period the within component is -1.8 percentage points, while the between amounts to 1.2 percentage points. Comparing the results for the two periods shows that both components recovered. The Foster, Haltiwanger, and Krizan (2001) decomposition assigns a larger negative contribution to the within component (-8.0 percentage points), a larger negative component to the between firms reallocation (-4.4 percentage points), and a sizable positive cross effect (8.0 percentage points). This decomposition allows a clear counterfactual exercise where changes in reallocation rates are calculated holding constant the product shares at their initial levels. The above findings suggest that the smaller effect of within and between firms variation in explaining the decline in reallocation can result from the cross-term, i.e. the relation between the change in shares and the change in reallocation rates.

Overall, the findings from this section show that the aggregate reallocation rate is largely explained by the decisions by incumbents firms to create and destroy products, followed by the contributions of entering and exiting firms. Moreover, the decompositions show that the decline in aggregate reallocation in the recession resulted largely from declines in reallocation within surviving firms. Further, Table 2.4 shows that these conclusions are robust to the choice of decomposition method.

These results motivate us to better understand the consequences of reallocation within incumbent firms. We interpret these empirical facts as evidence that some of the variation in the productivity and growth within surviving firms that Foster, Haltiwanger, and Syverson (2008) find is related to how they manage heterogenous multi-product portfolios that are comprised of winners (high revenue and high productivity products) and losers (low revenue and low productivity).

2.5 Product Reallocation and Innovation

What does product reallocation represent? In most models of creative destruction, output reallocation plays an important role in determining productivity dynamics. These models emphasize that adopting new products inherently involves the destruction of the old ones and that the pace at which this destruction takes place depends crucially on the innovation activities of the firm.[14] In this section we establish that there is a positive relationship between product reallocation and innovation.

2.5.1 Exploring Heterogeneous Types of Entry and Exit

The results of the previous section do not distinguish products being added or destroyed in what regards to how innovative they are. When we observe an entry of a new UPC, it might be a good that is very similar to others that the firm already has in its portfolio, or a good that is truly unique and innovative. As discussed in Section 2.2 defining a product as a unique UPC can cause some measurement concerns. In our data set, small changes in packing or volume likely result in a new bar code.[15] This type of new product is not what researchers have in mind when developing models of the effect of innovation in reallocation. We address this issue in two ways. First, we distinguish between two different types of innovation – incremental and extensions– and we examine their evolution in the recession and post-recession periods. Second, we show that the results reported in the previous sections do not qualitatively change when we consider coarser definitions of products.

14. For example, in Aghion and Howitt (1992) firms get monopoly rents for their innovations until the next innovation arrives. In this case, the incentives for investing in innovation are substantial.

15. Broda and Weinstein (2010) also acknowledge that the their measures of product creation and destruction include changes in characteristics that might be secondary and use information on the UPC's characteristics to show that only a small part results from changes in size and flavor. We follow an alternative approach that fits better with our overall goal.

Under the first approach, our goal is to distinguish between the entry of a new product within the main product line and a new product that is beyond the main product line of the firm. New products that constitute only marginal changes in the stock of existing products, such as changes in volume and other minor characteristics of the products, are unlikely to involve a lot of resources when developed and to have an insignificant impact on the outcomes of the firm. By contrast, new products that are not within the core business of the firm are likely to involve substantial changes in the production technology with sizable consequences to the outcomes of the firm. We implement a distinction between types of product by using the classification system in the Nielsen data set. In particular, we classify a new product at t as an improvement if the firm already has other products of that type, that is, if the firm at $t-1$ already produces goods in the same module as the product being created. We classify a new product as an extension if it is in a new module for the firm.

We apply the same principle to classify exits by type. Exits are classified as improvements if the firm maintains operations in that module. Exits are classified as extensions if they correspond to a cessation of activity in that module. The distinction shows if some products are terminated due to creative destruction (replaced by new products within the same product category) and those that are phased out due to the scaling down of the firms operations (products without replacement).

Table 2.5 presents the decomposition of aggregated and average entry and exit rates by type over the period under analysis. For comparison, we also show the share of entering and exiting products introduced by entering and exiting firms. As expected, most changes in UPCs occur within the same product module: around 80 percent of entering and exiting UPCs. Product extensions correspond to 11 percent of all entries, and exit extensions correspond to 14 percent of all exits. Over the period under analysis, we observe that product extensions and shut downs of product lines are only slightly more cyclical than the entrants and exits of within firm's product lines. We interpret this as evidence that over the business cycle, firms change the rate at which they make marginal changes in their stock of products, as well as the introduction of new product lines.

2.5.2 Research and Development Expenses

In order to further understand the relation between firms' innovation activities on the reallocation of their products, we use the various measures of product creation and destruction described in the previous section along with information on R&D expenses available in Compustat. This measure is particularly relevant because, as it is defined by Compustat, it encompasses all planned search aimed at the discovery of new knowledge that could lead to new products or the improvement of the existing ones.

Given that our main interest is to explore the determinants of reallocation within incumbent firms, we focus on studying firms present in every period in our sample. Our specification is the following:

$$r_{f,t+1} = \alpha + \beta \mathrm{R\&D}_{f,t} + \Gamma X_{f,t} + \mu_f + \lambda_t + \epsilon_{f,t} \tag{2.13}$$

where $r_{f,t}$ represents the reallocation rate of firm f in year t. R&D represents the ratio of research and development expenses to total sales for firm f at time t. Our main focus is on β that captures the direct impact of R&D on product reallocation. $X_{f,t}$ is a vector of firm-level controls that vary over time. All our specifications include year fixed effects, λ_t, to control for possible trends and firm fixed effects, μ_f, to control for other types of heterogeneity.

Table 2.6 shows the results. Column (1) shows that R&D expenditures have a large positive impact on reallocation after controlling for firm size because larger firms tend to engage in more R&D activities. Next, we add a wide range of controls to disentangle the effect of R&D from potentially confounding firm-level factors. In column (2) we include the price cost margin and in column (3) a control for firm idiosyncratic volatility. Our results do not vary under these specifications or if measures of financial constraints are included (column (4)). This is not surprising given that even without any time varying control the inclusion of both firm and time fixed effects account for most of the possible variation. In all cases the point estimates are large and statistically significant; a hypothetical increase in R&D expenditures relative to sales of 1 percentage point increases the reallocation rate by 0.6-0.9 percentage points. This is equivalent to an increase of close to 10% in the reallocation rate.[16]

2.6 Reallocation and Growth of Firm

Our findings so far strongly suggest that the innovation efforts of the firms are associated with higher reallocation rates. The second key prediction of Schumpeterian growth models to test is whether increases in the reallocation rates of products lead to larger growth rates for firms and to improvements in the products they produce.

2.6.1 Reallocation and Revenue Growth

We first confirm the prediction on revenue growth by estimating the following equation in the data:

$$\text{Revenue}_{f,t+1} = \alpha + \beta\, r_{f,t} + \Gamma X_{f,t} + \mu_f + \lambda_t + \epsilon_{f,t} \tag{2.14}$$

where $\text{Revenue}_{f,t}$ is the sum of the revenue of all products in the firm's portfolio at time t. As before, all our specifications include both firm and time fixed effects and consider only a balanced sample of firms. Furthermore, given that in order to run this specification we only require the information available in the RMS, equation 2.14 is estimated using quarterly data.

Column (1) in table 2.7 shows that β, our coefficient of interest, is both economically and statistically significant. This is after controlling for revenue in the previous period. The table also shows that, not surprisingly, most of the revenue growth due to reallocation of products comes from the entry margin. The exit rate on the other hand is negatively related to the revenue growth in the next quarter.

At the entry margin, the entry of products in the module where a firm operated before, Column (3) in the table, and the entry of products in a new module, Column (4), are associated with revenue growth by similar magnitudes. At the exit margin, closing down a product module completely, Column (7), is more strongly correlated with revenue growth, compared to destroying products in

16. In the Appendix E, we test a placebo specification by using future R&D expenditures instead of past R&D expenditures in predicting the change in the reallocation rate, and rule out a concern about confounding factors, such as time-varying firm-level shocks.

the module they keep operating in (related to the idea of creative destruction), Column (6).

2.6.2 Reallocation and Quality Improvements

A similar analysis can be done to explore whether higher reallocation rates lead to increases in the average quality of firms' portfolios. Several growth models, such as those in Klette and Kortum (2004) and Lentz and Mortensen (2008), predict that higher quality versions of a product are the outcome of the innovation activities of the firms. To study these predictions, we use three different measures to approximate product quality: relative prices, a percentile-based measure, and a perceived quality measure.

Benchmark Quality To measure a products average quarterly quality at the firm level, we use prices to approximate quality as in Argente and Lee (2016). This measure is similar to those used in the international trade literature where, if a sector or firm in a country is able to export a large volume at a high price, then it must be producing high-quality goods (Hummels and Klenow, 2005; Hallak and Schott, 2011; Kugler and Verhoogen, 2012). As a benchmark measure, we represent quality with the average relative price of the UPC-level good within each product category. First, we measure the log-difference between the price of good j and the median price for category c in quarter t.

$$R_{jt}^{\text{benchmark}} = \log \frac{P_{jt}}{\bar{P}_{ct}}$$

where $R_{jt}^{\text{benchmark}}$ is the relative price, and $\bar{P}_{ct}$ is the median price of category c. Therefore, if the price of a high quality type of milk, say organic milk, is much higher than the median price of milk, then $R_{jt}^{\text{benchmark}}$ is positive and high.

We then calculate the firm-level average quality by combining information on product-level quality and on the product portfolio of each firm. The average product quality of firm f is:

$$Q_{ft}^{\text{benchmark}} = \sum_{jf} \omega_{jft} R_{jt}^{\text{benchmark}}$$

where ω_{jft} is a revenue weight. $Q_{ft}^{\text{benchmark}}$ captures how far the prices of the products produced by firm f are from the median price level in each of their categories at time t.

Percentile-based Quality Measure Since the cross-sectional distribution of prices in a given category is fat-tailed, a cardinal measure of quality might be noisy and thus problematic. For this reason, we also define a percentile-based measure of quality as follows:

$$R_{jt}^{\text{percentile}} = \frac{1}{N_{ct}}(n_{jct} - \frac{1}{2})$$

where N_{jt} is the total number of products in category c at quarter t and n_{jt} is an ordinal rank of product j. $Q_{ft}^{\text{percentile}}$ can be defined as before by aggregating the product portfolio of each firm by using revenue weights.

Using these quality measures, we now test the association between our measure of product reallocation and improvement in the average quality of the product at the firm level. We use the following specification:

$$Q_{f,t+1} = \alpha + \beta\, r_{f,t} + \Gamma X_{f,t} + \mu_f + \lambda_t + \epsilon_{f,t} \tag{2.15}$$

where f is the firm, and t is the quarter. Our main focus is on β that captures the direct impact of reallocation on firm-level average quality in the next quarter. $X_{f,t}$ is a vector of firm-level controls, μ_f represents firm fixed effects, and λ_t represents time fixed effects. By construction, the benchmark quality measure is centered at zero and the percentile-based quality measure is center at 50. Table 2.8 reports the relation between our reallocation measures and our benchmark quality measure, $Q_{f,t+1}^{\text{benchmark}}$. We again keep a balanced panel of firms to investigate the importance of reallocation among surviving firms. An increase in reallocation is associated with quality improvements in the following quarter. This correlation is mainly driven by the entry margin of products. Large firms tend to produce higher quality products on average. Furthermore, as shown in columns (3) and (4), quality improves more for product extensions beyond the module than for incremental innovations.

2.6.3 Reallocation and Productivity

The remaining central implication of models with creative destruction to be tested is whether the reallocation of products is a major source of productivity growth. This prediction has been hard to examine directly in the data given the lack of availability of data sets combining both product- and firm-level information.[17]

We begin by computing total factor productivity in the Compustat data relying on the methodology developed by İmrohoroğlu and Tüzel (2014). We then regress the natural logarithm of TFP on the annual reallocation rate as follows:

$$TFP_{f,t+1} = \alpha + \beta\, r_{f,t} + \Gamma X_{f,t} + \mu_f + \lambda_t + \epsilon_{f,t} \tag{2.16}$$

where as before f is the firm, and t is the year. Our main focus is once again on β which captures the direct impact of reallocation on firm's productivity. Table 2.9 reports our results. Column (1) shows that both variables are strongly correlated even after controlling for the size of the firm. This is important because in general larger firms have higher productivity and, as we have shown, they also have higher rates of reallocation. Column (2) includes controls for market power by including the price-to-cost margin as control. Column (3) includes the standard deviation of sales to control for the possibility that firms with faster sales growth have higher rates of product entry and exit. Lastly, in column (4) we control for differences in financial constraints across firms by including the Kaplan-Zingales index. Our estimates of β remain similar across specifications and show that, on average, an increase of 1 percentage point in reallocation increases TFP around 0.35%.[18]

When we examine the contribution of entries and exits separately, we find that the contribution of reallocation to TFP mainly comes from exits. But, interestingly, when we explore improvements and extensions separately, we find that product extensions, products that are more likely to involve larger innovations, have a positive and significant contribution to TFP. On the other hand, exits

17. This question has been, nonetheless, explored in other contexts such as the reallocation of establishments (Bartelsman and Doms, 2000; Foster, Haltiwanger, and Krizan, 2001; Foster, Grim, and Haltiwanger, 2016). In both cases, a large share of productivity growth can be explained by the reallocation of resources.

18. In the Appendix E, we test a placebo specification using past TFP instead of future TFP as an outcome variable.

within the main module of the firm, those that are more likely to come from replacing outdated products for better products, contribute positively and significantly to TFP.

2.6.4 Discussion

The importance of product reallocation has been central in models of creative destruction for decades but, as the theoretical literature evolved, lack of data availability made many of its central implications hard to test. Our calculations validate many predictions of these models. First, we find that faster innovation-led growth is associated with higher rates of reallocation of products (Aghion, Akcigit, and Howitt, 2014). Product reallocation is positively related to R&D and to revenue growth. Second, although both entrants and incumbents innovate (Bartelsman and Doms, 2000), most growth appears to come from incumbents improving on existing varieties (Garcia-Macia, Hsieh, and Klenow, 2016; Acemoglu and Cao, 2015). Third, small firms and new entrants have a comparative advantage in achieving major innovations or, as we called them within the context of the consumer goods sector, extensions (Akcigit and Kerr, 2010). And, lastly, a more innovative firm has higher levels of productivity (Lentz and Mortensen, 2008).

Implications to Labor Reallocation

In the absence of direct measures of product creation and destruction, many authors have recently tried to infer the sources of growth indirectly from the patterns of job flows. For example, in the quality ladder model with R&D activities and labor decisions, the degree of creative destruction is closely related to the employment growth rates in the economy (Klette and Kortum, 2004; Lentz and Mortensen, 2008; Garcia-Macia, Hsieh, and Klenow, 2016; Atkeson and Burstein, 2017). Using information from our Nielsen-Compustat matched data, we test whether the relation between product reallocation and employment growth holds in the data. Our evidence is suggestive of this relation but not conclusive. We find a positive but not significant correlation between product reallocation and contemporaneous firm-level employment growth rates. Table 2.10 shows that for entry rates this relation is positive but insignificant, with a slightly higher correlation for incremental innovations. By contrast, the correlation is much weaker for exit rates. Exit rates of incremental innovations are positively associated with employment growth while the correlation is negative for radical innovations. Overall, the direction of these correlations shows some support for the use of indirect inference to understand the sources of innovation but many other concerns remain. For example, in the presence of adjustment costs or wage rigidities, the contemporaneous employment growth might not be the proper statistic to identify the degree of product-level reallocation activities.

Implications to Aggregate Productivity

How much of the decrease in aggregate productivity can be attributed to changes in the product reallocation? Our baseline estimate in column 1 of Table 2.9 shows that total factor productivity increases by approximately 0.35 percent for every 1 percentage point increase in reallocation. Given that the reallocation rate decreased by 3.8 percentage points during the recession and that TFP declined almost 5% from 2007 to 2010 in our data, product reallocation can explain around 20 to 25% of the total decline in total factor productivity.[19] This evidence suggests that a significant drop in aggregate productivity was driven by firms slowing down their innovation activities during this period. This, in turn, decreased the dynamism in which they replaced older products with improved products decreasing the pace of quality improvements and ultimately economic growth.

Business Cycle Modeling

Our work highlights the importance of studying the role of product creation and destruction in propagating business cycle fluctuations. There is a substantial amount of literature that studies how business formation affects business cycle dynamics (e.g., Chatterjee and Cooper (1993) and Jaimovich and Floetotto (2008)) and a growing body of work that emphasizes the endogenous determination of the number of products over the cycle (e.g. Bilbiie, Ghironi, and Melitz (2012) and Minniti and Turino (2013)). Our results emphasize the importance of studying the role of multi-product firms in the amplification of shocks. Our estimations can be used to discipline the

19. The interpretation of our results should consider the fact that they were computed using a sample of large publicly traded firms. Moreover, although within firm reallocation is by far the most important component of the overall reallocation rate, our estimates in section 2.6.3 ignore the contribution of firm entry and exit.

parameters governing the number of producers and products within each firm at different stages of the business cycle. More importantly, the fact that the reallocation rate differs substantially across different types of firms has significant implications for business cycle modeling. In traditional business cycle models, firms are homogeneous and, even in models with multi-product firms, there are no differences in the amount of products they produce or in the rate at which they introduce them to the market. We show that larger and more diversified firms launch and phase out products more often (on average) but we also provide evidence that the reallocation rate of smaller and less diversified firms is more sensitive to aggregate conditions (see tables 2.1 and 2.3). These shifts in the distribution of sales over the cycle could potentially be an important source of amplification, and they highlight the importance of introducing firm-level heterogeneity to these models. Lastly, considerably more work needs to be done to understand the potential links between business cycles and innovation-based growth theory. Standard business cycle models do not address the determinants of product variety within firms and that changes in the product scope of firms occur exogenously. Given the strong correlation we find between R&D and reallocations and the correlation between product turnover and changes in TFP, our work shows that modeling the endogenous interaction between the innovation efforts of the firm and its product scope could substantially improve our understanding of business cycle fluctuations.

2.7 Conclusion

In this paper, we describe the extent of product innovation and reallocation in the consumer goods sector over the period from 2007 to 2013. We find a 25 percent decline in product reallocation during the Great Recession, and investigate the impact of this drop on firm-level outcomes such as revenue, product quality, and total productivity. The analysis provides several findings. First, product reallocation is strongly pro-cyclical and the cyclical pattern is overwhelmingly a consequence of within firm reallocation. Second, the rate of product reallocation is strongly related to the innovation efforts of the firms. Third, firms that have higher reallocation rates grow faster, launch higher quality goods, and experience larger gains in productivity.

Given that higher reallocation activities lead firms to grow both quantitatively and qualitatively, the fact that its pace suffered an important drop had substantial implications for aggregate growth in this period. More importantly, the fact that the reallocation rate took so long to return to its pre-recession level suggests it was an important factor in the slow recovery the economy experienced after the Great Recession. Our findings also show that industrial and innovation policies aimed at increasing economic growth should contemplate the relative importance of the product-mix decisions. This is particularly relevant for incumbent firms as they account for the majority of the decline in dynamism in the retail sector that ultimately led to important declines in total factor productivity.

Tables and Figures

Table 2.1: Summary Statistics of Products and Firms

The table reports summary statistics of products and firms included in the baseline sample. The variables are defined at the quarter level and grouped as averages for the period 2007Q1–2013Q4, and for the subperiods 2007Q1–2007Q4, 2010Q1–2010Q4, and 2013Q1–2013Q4. "Entrants" refers to the average share of products and firms that are identified for the first time in the data set in each quarter. "Exits" refers to the average share of products and firms that are identified for the last time in the data set in each quarter. "Continuers expanding (contracting)" refers to firms that had products in the previous quarter and are increasing (decreasing) the number of products. The diversification statistics report the average number of products and the share of firms within each categories. The revenue is the total sales (in dollars) across all stores and weeks of the quarter, deflated by the Consumer Price Index for All Urban Consumers. The revenue presents the share of products and firms in each revenue interval and is computed using all surviving products and countinuing firms.

			2007-2013	2007	2010	2013
Average # products			222105	211101	214001	252189
Share products	Entrants		0.043	0.047	0.037	0.047
	Exits		0.036	0.046	0.033	0.030
# products	Per module		242	234	232	272
	Per group		2486	2356	2391	2836
	Per dpt.		25688	24348	24709	29304
Share products	[0,500[		0.610	0.626	0.605	0.615
	[500,5000[		0.230	0.223	0.232	0.228
	[5000,50000[		0.136	0.128	0.138	0.135
	>=50000		0.024	0.023	0.025	0.022
Average # firms			12861	13074	12361	13319
Share firms	Entrants		0.021	0.025	0.017	0.024
	Exits		0.020	0.023	0.018	0.019
	Continuers Expanding		0.122	0.118	0.110	0.141
	Continuers Contracting		0.128	0.145	0.125	0.118
Share firms	Single Product	Share firms	0.262	0.280	0.265	0.252
	Multi-Product & Single Module	Share firms	0.259	0.256	0.261	0.257
		Average # products	5.7	5.6	5.6	5.7
	Multi-Module & Single Group	Share firms	0.126	0.121	0.126	0.128
		Average # products	12.9	12.5	13.2	12.9
	Multi-Group & Single Dpt.	Share firms	0.192	0.187	0.191	0.195
		Average # products	23.5	22.4	24.4	23.8
	Multi-Dpt.	Share firms	0.161	0.155	0.157	0.168
		Average # products	61.6	56.9	59.4	65.4
Share continuing	$[0,10^4[$		0.462	0.480	0.455	0.462
	$[10^4,10^5[$		0.262	0.256	0.265	0.264
	$[10^5,10^6[$		0.182	0.177	0.185	0.180
	$>=10^6$		0.093	0.087	0.095	0.094

Table 2.2: Aggregate and Average Within Entry, Exit, and Reallocation Rates

The table reports the aggregate and average entry, exit, and reallocation rates, as defined in Section 2.3.1, for the baseline sample. The entry, exit, and reallocation rates are computed at the quarter level, seasonally adjusted, and then grouped as averages for the periods 2007Q1–2013Q4, and for the subperiods 2007Q1–2007Q4, 2010Q1–2010Q4, and 2013Q1–2013Q4.

		All 2007-2013	(1) 2007	(2) 2010	(3) 2013	Change (2)/(1)-1	 (3)/(2)-1
Aggregate Rates	Reallocation	0.079	0.094	0.070	0.078	-25%	11%
	Entry	0.043	0.047	0.037	0.047	-21%	25%
	Exit	0.036	0.046	0.033	0.031	-30%	-5%
Average Within Rates	Reallocation	0.108	0.125	0.095	0.113	-24%	19%
	Entry	0.055	0.061	0.047	0.065	-22%	37%
	Exit	0.056	0.068	0.050	0.051	-27%	3%

Table 2.3: Average Within Reallocation Rates by Types of Firms

The table reports the average entry, exit, and reallocation firm-specific rates by types of firms, as defined in Section 2.3.1, for the baseline sample. The entry, exit, and reallocation rates are computed for different sets of firms at the quarter level, seasonally adjusted, and then grouped as averages for the period 2007Q1–2013Q4, and for the subperiods 2007Q1–2007Q4, 2010Q1–2010Q4, and 2013Q1–2013Q4.

	All 2007-2013	(1) 2007	(2) 2010	(3) 2013	Change (2)/(1)-1	(3)-(2)-1
Reallocation						
Expanding Entrant	1.000	1.000	1.000	1.000	0%	0%
Expanding Incumbent	0.256	0.267	0.250	0.261	-6%	4%
Contracting Exit	1.000	1.000	1.000	1.000	0%	0%
Contracting Incumbent	0.246	0.276	0.235	0.235	-15%	0%
Unchanged Incumbent	0.006	0.009	0.004	0.006	-51%	41%
Entry						
Expanding Entrant	1.000	1.000	1.000	1.000	0%	0%
Expanding Incumbent	0.237	0.243	0.234	0.244	-4%	4%
Contracting Incumbent	0.013	0.018	0.011	0.013	-39%	20%
Unchanged Incumbent	0.003	0.004	0.002	0.003	-51%	41%
Exit						
Expanding Incumbent	0.019	0.024	0.017	0.018	-31%	6%
Contracting Exit	1.000	1.000	1.000	1.000	0%	0%
Contracting Incumbent	0.234	0.259	0.225	0.223	-13%	-1%
Unchanged Incumbent	0.003	0.004	0.002	0.003	-51%	41%
Reallocation						
Q1 revenue	0.173	0.209	0.150	0.183	-28%	22%
Q2 revenue	0.098	0.096	0.089	0.108	-7%	22%
Q3 revenue	0.086	0.091	0.078	0.092	-14%	18%
Q4 revenue	0.075	0.084	0.069	0.075	-18%	9%
Entry						
Q1 revenue	0.084	0.090	0.072	0.107	-20%	48%
Q2 revenue	0.049	0.050	0.042	0.061	-16%	45%
Q3 revenue	0.046	0.050	0.039	0.054	-22%	37%
Q4 revenue	0.043	0.049	0.039	0.044	-20%	11%
Exit						
Q1 revenue	0.102	0.135	0.088	0.089	-35%	2%
Q2 revenue	0.051	0.048	0.049	0.050	2%	3%
Q3 revenue	0.041	0.042	0.040	0.040	-5%	1%
Q4 revenue	0.032	0.035	0.030	0.032	-15%	7%
Reallocation						
Single-product	0.112	0.137	0.098	0.110	-28%	13%
Single-Module	0.070	0.078	0.060	0.075	-23%	25%
Single-Group	0.081	0.095	0.072	0.080	-24%	11%
Single-Department	0.072	0.085	0.065	0.072	-24%	11%
Multi-department	0.080	0.095	0.069	0.080	-28%	17%
Entry						
Single-product	0.027	0.034	0.021	0.033	-39%	61%
Single-Module	0.024	0.026	0.020	0.030	-23%	51%
Single-Group	0.031	0.033	0.028	0.034	-16%	21%
Single-Department	0.029	0.033	0.026	0.034	-22%	31%
Multi-department	0.035	0.039	0.028	0.039	-26%	39%
Exit						
Single-product	0.087	0.107	0.079	0.080	-26%	1%
Single-Module	0.046	0.053	0.041	0.046	-23%	12%
Single-Group	0.051	0.062	0.045	0.047	-28%	5%
Single-Department	0.043	0.053	0.039	0.039	-26%	-1%
Multi-department	0.045	0.057	0.040	0.041	-29%	1%

Table 2.4: Decomposition

The table reports decomposition exercises on the change in the aggregate entry, exit, and reallocation rates by types of firms, as defined in Section 2.4. We decompose the first differences of the aggregate entry and exit rates. The decomposed series are seasonally adjust and then summed over the periods 2007Q1–2009Q4, and 2010Q1–2012Q4. The decomposition of the reallocation rate is computed by adding the subcomponents of the entry and exit decompositions.

		Within (+)	Between (+)	Cross (+)	Entry (+)	Exit (-)	Change
OP - Non-dynamic							
Entry Rate							
	07Q1 - 09Q4	-2.9	1.0	-	-	-	-1.9
	10Q1 - 12Q4	2.2	-1.0	-	-	-	1.2
Exit Rate							
	07Q1 - 09Q4	-2.4	0.7	-	-	-	-1.7
	10Q1 - 12Q4	-0.3	-0.2	-	-	-	-0.5
Reallocation Rate							
	07Q1 - 09Q4	-5.3	1.6	-	-	-	-3.6
	10Q1 - 12Q4	1.8	-1.2	-	-	-	0.7
OP - Dynamic							
Entry Rate							
	07Q1 - 09Q4	-23.4	18.7	-	3.1	0.2	-1.9
	10Q1 - 12Q4	-19.2	17.5	-	3.0	0.1	1.2
Exit Rate							
	07Q1 - 09Q4	19.5	-19.1	-	0.2	2.2	-1.7
	10Q1 - 12Q4	18.4	-17.3	-	0.1	1.7	-0.5
Reallocation Rate							
	07Q1 - 09Q4	-4.0	-0.4	-	3.1	2.2	-3.6
	10Q1 - 12Q4	-0.8	0.2	-	3.0	1.7	0.7
GR							
Entry Rate							
	07Q1 - 09Q4	-8.1	3.3		3.0	0.2	-1.9
	10Q1 - 12Q4	-5.6	3.9		3.0	0.1	1.2
Exit Rate							
	07Q1 - 09Q4	4.1	-3.7		0.2	2.2	-1.7
	10Q1 - 12Q4	3.8	-2.7		0.1	1.7	-0.5
Reallocation Rate							
	07Q1 - 09Q4	-4.0	-0.4	-	3.0	2.2	-3.6
	10Q1 - 12Q4	-1.8	1.2	-	3.0	1.7	0.7
FHK							
Entry Rate							
	07Q1 - 09Q4	-10.4	1.0	4.7	3.0	0.2	-1.9
	10Q1 - 12Q4	-8.1	1.3	5.1	3.0	0.1	1.2
Exit Rate							
	07Q1 - 09Q4	2.4	-5.4	3.3	0.2	2.2	-1.7
	10Q1 - 12Q4	2.6	-3.9	2.5	0.1	1.7	-0.5
Reallocation Rate							
	07Q1 - 09Q4	-8.0	-4.4	8.0	3.0	2.2	-3.6
	10Q1 - 12Q4	-5.5	-2.6	7.5	3.0	1.7	0.7

Table 2.5: Summary Statistics for Aggregate and Within Entry and Exit Rates by Types

The table reports aggregate and average within firm entry and firm exit rates by different types of products. “Entry Improvement” rate is defined as the share of total products created within the firms’ previous portfolio of modules. “Entry Extensions” rate is the share of total products created outside a firm’s previous portfolio of modules. “Entry firm” refers to the ratio of new products by new firms relative to the total products. “Exit Improvement” rate is the share of total products eliminated within firm’s previous portfolio of modules. “Exit Extensions” rate is the share of total products exiting that eliminated modules from a firm’s portfolio. ”Exit firm” refers to the share of firms that exit the market. The rates are computed at the quarter level, seasonally adjusted, and then grouped as averages for the period 2007Q1–2013Q4, and for the subperiods 2007Q1–2007Q4, 2010Q1–2010Q4, and 2013Q1–2013Q4.

		All	(1)	(2)	(3)	Change	
		2007-2013	2007	2010	2013	(2)/(1)-1	(3)-(2)-1
i. Aggregate Rates							
Entry							
	Improvement	0.036	0.039	0.031	0.038	-20%	22%
	Extension	0.005	0.005	0.004	0.005	-26%	40%
	Firm	0.003	0.003	0.002	0.003	-27%	36%
	ALL	0.043	0.047	0.037	0.047	-21%	25%
Exit							
	Improvement	0.030	0.038	0.027	0.025	-29%	-7%
	Extension	0.005	0.006	0.004	0.004	-35%	-3%
	Firm	0.002	0.002	0.001	0.002	-29%	11%
	ALL	0.036	0.046	0.033	0.031	-30%	-5%
ii. Average Within Rates							
Entry							
	Improvement	0.024	0.025	0.022	0.028	-15%	30%
	Extension	0.009	0.010	0.008	0.012	-22%	50%
	Firm	0.022	0.025	0.018	0.025	-30%	41%
	ALL	0.055	0.061	0.047	0.065	-22%	37%
Exit							
	Improvement	0.020	0.023	0.018	0.018	-21%	2%
	Extension	0.016	0.022	0.014	0.013	-35%	-6%
	Firm	0.020	0.024	0.018	0.020	-24%	10%
	ALL	0.056	0.068	0.050	0.051	-27%	3%

Table 2.6: Reallocation Activities and R&D Expenses

The table reports the coefficients of OLS regressions with revenue weights. The dependent variable reallocation at t, defined as the product entry rate plus the product exit rate, at the firm level as defined in the main text. The main independent variable is the ratio of R&D expenses to total sales at $t-1$. The construction of the rest of the control variables is described in Appendix D. Other controls include firm fixed-effects and year fixed-effects. Revenue is winsorized at the 1% level. The sample used include the set of publicly traded firms that are found in the Nielsen RMS. Standard errors are presented in parentheses. ***, **, and * represent statistical significance at 1%, 5%, and 10% levels, respectively.

Dep. var.: $r_{f,t+1}$	(1)	(2)	(3)	(4)
R&D	0.668*	0.625*	0.873**	0.872**
	(0.363)	(0.363)	(0.392)	(0.392)
Size	0.006	0.014	0.014	0.015
	(0.032)	(0.033)	(0.035)	(0.035)
Price Cost Margin		0.344	0.405*	0.399*
		(0.210)	(0.224)	(0.226)
Std. Sale			-0.169	-0.161
			(0.127)	(0.127)
Kaplan-Zingales				0.000
				(0.001)
Observations	661	661	599	595
R-squared	0.563	0.565	0.576	0.579
Year Effects	Yes	Yes	Yes	Yes
Firm Effects	Yes	Yes	Yes	Yes

Table 2.7: Reallocation Activities and Revenue Growth

The table reports the coefficients of OLS regressions with revenue weights. The dependent variable is the revenue growth in the next quarter. The reallocation rate at t of firm f, $r_{f,t}$, is defined as the product entry rate plus the product exit rate at the firm level as defined in the main text. Revenue is winsorized at the 1% level. Standard errors are presented in parentheses. ***, **, and * represent statistical significance at 1%, 5%, and 10% levels, respectively.

Dep. var.: $\text{Revenue}_{f,t+1}$	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$r_{f,t}$	0.2869***						
	(0.011)						
$n_{f,t}$		0.5281***					
		(0.014)					
$n_{f,t}$ (in module)			0.5445***				
			(0.017)				
$n_{f,t}$ (beyond module)				0.5452***			
				(0.028)			
$x_{f,t}$					-0.7333***		
					(0.017)		
$x_{f,t}$ (in module)						-0.0396*	
						(0.023)	
$x_{f,t}$ (beyond module)							-1.7489***
							(0.027)
$\text{Revenue}_{f,t}$	0.8007***	0.7976***	0.7972***	0.7989***	0.7554***	0.7565***	0.7598***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	242,660	242,803	242,803	242,803	242,711	242,711	242,711
R-squared	0.970	0.970	0.970	0.970	0.967	0.967	0.967
Year Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 2.8: Reallocation Activities and Benchmark Quality Improvement

The table reports the coefficients of OLS regressions with revenue weights. The dependent variable is the improvement in benchmark quality in the next quarter. Reallocation rate at t of firm f, $r_{f,t}$, is defined as the product entry rate plus the product exit rate at the firm level as defined in the main text. Revenue is winsorized at the 1% level. Standard errors are presented in parentheses. ***, **, and * represent statistical significance at 1%, 5%, and 10% levels, respectively.

Dep. var.: $Q^{\text{benchmark}}_{f,t+1}$	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$r_{f,t}$	0.0255***						
	(0.004)						
$n_{f,t}$		0.0383***					
		(0.005)					
$n_{f,t}$ (in module)			0.0295***				
			(0.006)				
$n_{f,t}$ (beyond module)				0.0617***			
				(0.009)			
$x_{f,t}$					0.0144**		
					(0.006)		
$x_{f,t}$ (in module)						-0.0043	
						(0.007)	
$x_{f,t}$ (beyond module)							0.0446***
							(0.009)
Revenue$_{f,t}$	0.0064***	0.0059***	0.0059***	0.0060***	0.0064***	0.0064***	0.0064***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	242,537	242,679	242,679	242,679	242,588	242,588	242,588
R-squared	0.925	0.924	0.924	0.924	0.924	0.924	0.924
Year Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 2.9: Reallocation Activities and Firm-Level Productivity

The table reports the coefficients of OLS regressions with revenue weights. The dependent variable is the natural logarithm of the total factor productivity at the firm-level at $t+1$. Reallocation at t is defined as the product entry rate plus the product exit rate at the firm level as defined in the main text. The construction of the control variables is described in Appendix D. Other controls include firm fixed-effects and year fixed-effects. Revenue is winsorized at the 1% level. The sample used include the set of publicly traded firms that are found in the Nielsen RMS. Standard errors are presented in parentheses. ***, **, and *, represent statistical significance at 1%, 5%, and 10% levels, respectively.

$TFP_{f,t+1}$	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$r_{f,t}$	0.36***	0.37***	0.35***	0.35***						
	(0.10)	(0.10)	(0.10)	(0.10)						
$n_{f,t}$					0.026					
					(0.11)					
$n_{f,t}$ (in module)						-0.06				
						(0.12)				
$n_{f,t}$ (beyond)							1.89***			
							(0.55)			
$x_{f,t}$								0.91***		
								(0.17)		
$x_{f,t}$ (in module)									0.96***	
									(0.18)	
$x_{f,t}$ (beyond)										0.24
										(0.72)
Size	0.17***	0.15***	0.19***	0.19***	0.22***	0.22***	0.22***	0.20***	0.20***	0.22***
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.033)
Price Cost Margin		-0.30	-0.40*	-0.38*	-0.77***	-0.76***	-0.76***	-0.31	-0.32	-0.76***
		(0.19)	(0.22)	(0.22)	(0.21)	(0.21)	(0.21)	(0.22)	(0.22)	(0.21)
Std. Sale			-0.28**	-0.27**	-0.24**	-0.24**	-0.24**	-0.28**	-0.29**	-0.24**
			(0.11)	(0.12)	(0.11)	(0.11)	(0.11)	(0.11)	(0.11)	(0.11)
Kaplan-Zingales				0.00	-0.00	-0.00	-0.00	0.00	0.00	-0.00
				(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Observations	834	834	777	773	865	865	865	773	773	865
R-squared	0.85	0.85	0.86	0.86	0.84	0.84	0.85	0.86	0.87	0.84
Year Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 2.10: Correlation between Reallocation Activities and Employment Growth Rates

The table shows the correlation between the reallocation activities and the employment growth rates from the Nielsen-Compustat matched database. P-values are presented in parentheses. ***, **, and * represent statistical significance at 1%, 5%, and 10% levels, respectively. Each observation is at the year-firm level. The number of observations is 981.

	$r_{f,t}$	$n_{f,t}$	$n_{f,t}$ (in module)	$n_{f,t}$ (beyond)	$x_{f,t}$	$x_{f,t}$ (in module)	$x_{f,t}$ (beyond)
Correlation w/ emp. growth	0.0255 (0.4095)	0.0342 (0.2697)	0.0240 (0.4379)	0.0261 (0.3998)	0.0099 (0.7490)	0.0524* (0.0904)	-0.0292 (0.3463)

CHAPTER 3

COST OF LIVING INEQUALITY DURING THE GREAT RECESSION (WITH MUNSEOB LEE)

3.1 Introduction

For decades, economists have tried to study the patterns of poverty and inequality by focusing only on the disparities in nominal income or nominal wages. The measures of real income inequality that account for differences in the cost of living are arguably better measures of the differences in the standard of living. However, the literature largely ignores this dimension due to the difficulty in calculating cost of living indices at the income-group level or at finer levels of disaggregation. If the cost of living differs across income groups, then estimates of poverty and inequality could be misleading. This problem could be worse during downturns when the well-known upward bias in the consumer price index (CPI) is larger.[1] In this paper, we calculate income-specific price indices and show that high-income households experienced lower inflation in their cost of living than low-income households during the Great Recession. We argue that high-income households were able to cope with the recession by adjusting their shopping behavior unlike low-income households.

The recent literature shows that households' shopping behavior changed during the Great Recession. Households changed the quality of the items they bought, the stores where these purchases were made, or simply took advantage of sales and coupons. However, the extent to which households can adjust along such margins can vary with income. Because households have different margins within which to adjust their shopping behavior, they face heterogeneous inflation rates. Nevertheless, the conventional price indices do not capture these changes in shopping behavior. This is because the US Bureau of Labor Statistics (BLS) does not collect information on quantities, so it updates the expenditure weights infrequently. Further, the BLS only collects information on a limited variety of products within a category, therefore the Consumer Price Index (CPI) does not accurately reflect substitution within a category. Taking advantage of a large scanner database of

1. See, e.g., Chevalier and Kashyap (2014), Coibion, Gorodnichenko, and Hong (2015b), Kryvtsov and Vincent (2014), and Handbury, Watanabe, and Weinstein (2013).

consumers, and relying on the method developed by Feenstra (1994), we calculate price indices for different income groups. We allow five different components in the price index vary with income: i) the set of goods consumed, ii) expenditure shares for each of the goods, iii) within-category elasticity of substitution, iv) perceived quality of each product, and iv) the actual prices paid. By allowing these components to differ, our price indices can incorporate changes in shopping behavior.

From our income-specific price indices, we find substantial differences across income groups during the Great Recession. In our preferred specification, the annual inflation for the cost of living of the highest quartile of the income distribution was 0.59 percentage points lower than that of the lowest quartile from 2004 to 2010. This gap is substantially large: the annual inflation rates of the lowest and highest quartile were 1.39% and 0.80% respectively.

To see if the gap comes from different margins of shopping behavior across income groups, we investigate three major ways consumers can adjust when they suffer income declines. The first margin is shopping intensity. Shopping intensity is the frequency at which consumers purchase and the types of purchases, such as sale items, coupon usage, generic products, and bulk items. The second is store switching: the reallocation of expenditures toward lower priced retailers. The third is quality substitution: the reallocation of expenditures from high to low quality brands within the same product category. We study the patterns of the average price paid of each income group in a regression that controls for each of these three margins, and we distinguish which mechanism causes the difference. We find that quality substitution can explain an important part of the differences across income groups. That is, high-income households are better able to pay lower prices for the same category of goods by shifting their expenditures to less expensive brands, while low-income households have fewer options for quality substitution during downturns. This extra margin of adjustment for high-income households allowed them to lower their cost of living relative to low-income households during the Great Recession.

We also provide direct evidence that the responses of the high- and low-income households differ when they suffer an income decline. By combining retail-level scanner data and household-level panel data, we are able to measure the degree of quality substitution for each household. With this measure, we explore how households change their shopping behavior when local economic conditions deteriorate. We find that only high-income households take advantage of the quality substitution margin. And, even though high-income households respond along other margins, such as store switching, the reallocation of expenditures within product categories explains most of the gap we observe.

Further, in order to rule out the possibility that our result arises from changes in the pricing behavior of retailers or manufactures, we construct measures of posted prices and show that retailers do not price discriminate in favor of high-income consumers when economic conditions worsen. Second, we provide evidence that, in the retail sector, consumers purchase goods from the same manufacturers; we do not find a clear link between certain types of manufacturers and consumers of different income groups. Therefore, we rule out the possibility that supply shocks to the retailers or manufacturers account for the differences in the prices paid by consumers.

Our study has implications on the measurement of real variables such as income and wealth as well as other economic indicators such as the poverty rate and inequality. Our findings indicate that there could be a systematic bias in the measurement of real variables during a recession. This bias could greatly affect, for example, the allocation of social transfers such as food stamps and other means-tested programs.

Our work contributes to various strands of the literature. First, this study relates to the literature that focuses on the differences in the cost of living across households. The studies by Michael (1979), Hagemann (1982), and Garner, Johnson, and Kokoski (1996) focus on computing group-specific inflation rates. More recently, Hobijn and Lagakos (2005) explore inflation inequality across households in the United States and find that the average difference in inflation between poor and non-poor people is less than 0.1%. However, these studies only allow the shares of expenditures of different categories to vary across groups. In our study, by taking advantage of a vast amount of scanner data, we allow the share of expenditures of different products within categories to vary across income groups as in Broda and Romalis (2009). They find that from 1994 to 2005, the inflation for poorer consumers was lower than the inflation for richer consumers and argue that half of the increase in conventional inequality measures was due to a bias caused by ignoring the variation in consumption behavior across income groups. This study builds on Broda and Romalis (2009) by computing income-specific elasticities of substitution within product categories to allow for non-homotheticicities across groups.[2] Our findings show that the reallocation of expenditures within categories accounts for most of the differences across income groups.

In subsequent work, Kaplan and Schulhofer-Wohl (2016) and Jaravel (2016) extend the sample period of our study and confirms the inflation disparities we document. Kaplan and Schulhofer-Wohl (2016) finds that low-income households experience higher inflation, even after controlling for all other demographics. Jaravel (2016) argues that increases in the relative demand for products consumed by high-income households led firms to introduced more new products to cater such households. As a result, continuing products decreased their prices due to an increased competitive pressure. His mechanism refers to a long-term trend whereas we focus on the Great Recession. As a result, it is not the main driver of the disparities we observe given that the crisis was characterized by a dramatic drop in demand, particularly for goods purchased by high-income households, and a significant decrease in the product entry rate which is highly cyclical.[3] In fact, the magnitude of the variety correction term in our price indices is relatively small and can explain at most 15% of the inflation disparities we observe.

Our study is also related to the extensive literature that examines the heterogeneity of shopping behavior across income groups. Aguiar and Hurst (2007) find that low-income households shop more intensively and typically pay lower prices for identical products. Griffith, Leibtag, Leicester, and Nevo (2009) use a national representative sample of UK households in 2006 and find that savings from sales and buying in bulk are of a similar order of magnitude to those from purchasing

2. Given the similarity of our methodologies, the differences in our results are likely due to the sample periods examined and the fact that the more recent data is richer. Nielsen does not longer supply the data from 1994-2003 so we are unable to replicate their exercise using our methodology. Our data set is much richer and allows us to compute both the bias correction term and the expenditure weights for each income group every quarter. Broda and Romalis (2009) are unable to compute an exact price index (EPI) for most of the periods in their sample due to lack of household level information. They examine two different basket of goods, food and non-durables. For food, which represent 50-70% of expenditures in their sample, they are unable to compute the EPI for half of the periods and construct a Paasche index instead. Due to the same limitation, they use extrapolations of the bias correction term for that part of the sample. For non-durables, they only have household level information in one quarter (2003q4). As a result, they construct a Paasche index for most of their sample periods and use the same bias correction term for all groups.

3. See Broda and Weinstein (2010) and Argente, Lee, and Moreira (2017b) for evidence on the procyclicality of product entry.

generics and store switching. They also find that low-income households mostly use generic and bulk purchases to save. Our findings are based on the United States during the Great Recession and show the importance of quality substitution within product categories as a margin of adjustment during economic downturns. These findings provide evidence that the ability of consumers to use this margin increases with income.

Further, this study contributes to the growing literature on how shopping behavior changes over the business cycle. The studies by Nevo and Wong (2015) and Stroebel and Vavra (2016) show that households lower their grocery bill during economic downturns by increasing their shopping intensity. McKenzie and Schargrodsky (2011) show that an increase in the activity of searching for shopping was one of the most prevalent adjustment mechanisms used by consumers to cope with the 2002 Argentinean financial crisis. Coibion, Gorodnichenko, and Hong (2015b) find that the inflation in effective prices declines significantly with higher unemployment, while little change occurs in the posted price's inflation. This difference, they argue, reflects the reallocation of household expenditures across retailers particularly by households at the highest income quintile. Although, store switching contributed to the decline of effective prices during the Great Recession, we find that the main mechanism that explains inequality in the inflation of the cost of living is quality substitution. Burstein, Eichenbaum, and Rebelo (2005) highlight this mechanism by showing that a large measurement error emerged in the CPI after households substituted low quality goods in the aftermath of the Argentina's devaluation. Further, Jaimovich, Rebelo, and Wong (2015) show that higher quality goods lost market share during the Great Recession due to consumers trading down in quality.

The structure of the paper is as follows. Section 3.2 presents the data used for the empirical analysis. In Section 3.3 we construct the income-specific price indices. Section 3.4 has a description of our reduced form approach to analyze the implication of shopping behavior on the inequality of the cost of living. In this section we also provide cross-sectional evidence that quality substitution is the most important channel driving the differences across income groups. Section 3.5 presents direct evidence of the changes in shopping behavior. In Section 3.6 we rule out the possibility that the pricing behavior of retailers and firms drives our result. In Section 3.7 we discuss the implications of our results, and Section 3.8 concludes.

3.2 Data

To construct income-group specific price indices, we need to observe not only the households' demographic information but also the details on their purchases. The Homescan data set, which is provided by the Kilts-Nielsen Data Center at the University of Chicago Booth School of Business, is the best available source for this information. We use this data set to construct our price indices. One limitation of the Homescan data set is that it does not contain all of the transactions that take place in the market. Given that we can only observe transactions made by households in the sample, in order to directly measure the degree of quality substitution, we combine the Homescan data set with the Nielsen Retail Measurement Services (RMS) scanner data set. The Nielsen RMS collects its scanner data at the store level and contains all of the transactions made in a given store-week for each product category covered. We therefore also use the Nielsen RMS scanner data to study changes in the prices posted by retailers. In order to focus on the periods around the Great Recession, we use the Homescan data set from 2004 to 2010 and the Nielsen RMS scanner data set from 2006 to 2010. Further, in order to approximate the changes in local economic conditions, we use the BLS's monthly unemployment data available at the city and county level. Below we discuss each data source in detail and provide information on how we assembled the data.

3.2.1 Nielsen Homescan Data Set

The Nielsen Homescan tracks the shopping behavior of 40,000 to 60,000 households every year in 54 Metropolitan Statistical Areas (MSA). Each panelist uses in-home scanners to record their purchases.[4] A 12-digit universal product code (UPC) identifies the items the panelists purchase. The data contain around 1.6 million distinct UPCs grouped into 1,235 product modules that range from food to beauty aids to computer software. Our data cover around 40% of all of the

4. Nielsen offers a variety of incentives to join and stay active such as monthly prize drawings, gift points, and regular sweepstakes. The incentives are designed to be non-biasing (i.e., Nielsen does not provide account-specific coupons out of concern for the potential effect on the natural purchase selection of outlets and products).

expenditures on goods in the CPI.[5]

Our data represent well the categories and products it covers. To check the consistency of the data, Figure C.2 in Appendix **??** plots a price index constructed using the Nielsen Homescan Panel data and the BLS's Food at home CPI for all urban consumers. We use a procedure that mimics the construction of the CPI and we pick product groups in our data that match those used in the construction of the CPI-U Food at home. As shown in the figure, our index closely matches the overall patterns of the CPI-U Food at home.

For each UPC, the data contain information on the brand, size, packaging, and a rich set of product features. If the panelist purchases the good at a store covered by Nielsen, the price is set automatically to the average price of the good at the store during the week when the purchase was made. If not, the panelist directly enters the price. Nielsen reports detailed transaction information for each product purchased (e.g., UPC code, quantity, price, deals, and coupons). We combine this information with the weight and volume of the product to compute unit values. Because our data lack other measures of quality, we follow the industrial organization and the international trade literature and approximate the quality of a product with its unit value.

The data also contain information about each purchasing trip the panelist makes, such as the retailer, the location, and the date of the transaction. Further, the data have demographic variables such as age, education, annual income, marital status, and employment that are updated annually based on surveys sent to the panelists. The surveys are sent in Q4 of each year and the variables are implemented in the first week of January of the following year. Nielsen provided 16 income bins top-coded at \$100,000 up to 2005. After 2006, it has provided 20 income bins top-coded at \$200,000. Nielsen asks panelists to report their combined total household annual income as of year-end of the previous calendar year. Nielsen believes panelists are actually reporting their "annualized" estimated income as of the time of the survey and not referring to the previous year's tax returns. Self-reported annual income is likely to be the total labor income. Nielsen constructs projection weights that make the sample representative of the US population that we use in all

5. Table C.1 in Apppendix **??** depicts the distribution of UPCs and expenditures across different product groups as defined by Nielsen. The table shows that our sample includes a wide set of goods including a few durable goods (e.g., cameras, flashlights, and cookware). Unfortunately, our data do not include other important durables or goods purchased online. Our results should be interpreted accordingly.

our calculations.[6]

We restrict our sample to households whose head is between 25 and 64 years of age because the reported annual income might not accurately represent the households' income sources after retirement. Furthermore, a retiree's low opportunity cost of time has a direct bearing on the total cost of consumption, which makes market expenditures a poor proxy for actual consumption (Aguiar and Hurst, 2005, 2007).[7] Each year we divide the households into four groups according to their annual income: (i) less than $25,000, (ii) $25,000 to $50,000, (iii) $50,000 to $100,000, and (iv) more than $100,000. These groups roughly approximate the quartiles for the cross-sectional distribution of household income which were very stable during our sample period. The median household income increased from $44,344 (2004) to $50,303 (2008) and then dropped to $49,276 (2010) during the Great Recession.[8] On average, approximately 85% of the households remain in the same income bin the following year.[9]

3.2.2 Nielsen RMS (Retail Measurement Services) Scanner Data Set

The RMS consists of nearly 76 billion observations of weekly prices, quantities, and store information generated by point-of-sale systems. The data contain about 40,000 individual stores across 371 MSAs between January 2006 and December 2011. The data consist of approximately 1.4 million distinct products. Approximately 97% of the sales in the data come from food, drug, and mass merchandising stores. The data represent roughly 30% of the total US expenditures on food and

6. Nielsen has a comprehensive program of dropping and replacing panelists that do not perform to minimum reporting standards. Currently, Nielsen retains about 80% of its active panel each year. Nielsen uses a stratified sampling design to ensure that the panel is demographically balanced.

7. Figure C.1 in Appendix **??** shows the age profile of the households' average annual income in the Homescan data. As the figure shows, it mimics the standard lifecycle of earnings.

8. Einav, Leibtag, and Nevo (2010) find that income is not systematically correlated with the quality of recording in the Homescan data. Furthermore, to check how representative are the households in our data, we compare the expenditures in the Homescan with those of the food-at-home category in the CEX. For each income group, the expenditures-to-income ratio in both dataset are remarkably similar. For the poorest households, the ratio is 0.25 in the Homescan and 0.22 in the CEX. For the highest income households the ratio is 0.04 in both datasets.

9. Our results are qualitatively similar if we use lagged income instead. We use annualized income reported as of the time of the last survey to avoid dropping one year of data for each household (each household stays in the sample 3.2 years on average from 2004 to 2010) and to be able to construct our price indices from 2004 onward. Our results are also robust to using size-adjusted income as in Handbury (2013) where we compute each household's log size-adjusted income by subtracting its size fixed effect from its log income and adding back the single-member-household size fixed effect. These results are available on request.

beverages and roughly 2% of the total household consumption (Beraja, Hurst, and Ospina, 2016). Using the unique store identifier, the barcode of each product, and the date of the purchase, we match the information on household purchases contained in the Homescan with transactions at the store level in the RMS. We provide more details of the final sample in Section 3.5.

3.3 Income-Specific Price Indices

Our analysis begins by calculating the income-specific, cost-of-living indices with the Homescan. Within each product category, the number of products we observe is much larger than the number of products the BLS uses to calculate the CPI. Nielsen sorts its UPCs into over 1,000 narrowly defined product modules (e.g., skim milk and ground coffee). Our cost of living measures use these modules so that we can aggregate them to construct a cost of living index for each income group.

The availability of both prices and quantities in our data set provides the necessary inputs to calculate expenditure weights each period. Therefore, we are able to compute a price index that allows consumer choices to vary in response to a relative price change. In this way, our price index does not suffer from the substitution bias pointed out by Boskin (1996) and Boskin, Dulberger, Gordon, Griliches, and Jorgenson (1998). This is important because this bias is likely to increase if the shopping effort of the household increases. Following the literature, we present our results using a Sato-Vartia price index that allows the expenditure weights to vary over time and is theoretically justified if one assumes a CES utility function. Constructing a price index using a CES functional form has important advantages as well as some disadvantages. The CES is the preferred specification of many macroeconomic models. It has the attractive feature that the preferences of a representative consumer can be derived from aggregating heterogeneous consumers in a random utility model in which they only demand their preferred good (see Anderson, De Palma, and Thisse (1992) for more details). And, although it is not a superlative index, it approximates indices like the Tornqvist and the Fisher index extremely well. In addition, by estimating a price index for each income group our approach allows for an underlying non-homothetic preference structure that is approximated by nonsymmetrical CES utility functions. In this context, consumers in each income group can be located on different points of the same Engel curve as in Broda and Romalis

(2009).

For robustness, we also construct exact price indices for different income groups following the methodology developed by Feenstra and Reinsdorf (2000). They show that a Divisia price index can be used to generate an exact price index for different income groups whose preferences generate the Almost Ideal Demand System (AIDS) developed in Deaton and Muellbauer (1980). The AIDS demand system provides a flexible description of non-homothetic consumer tastes and has been previously used to study the differences in inflation across income groups (e.g. Nakamura, Steinsson, and Liu (2014) for the case of China). Our results under this specification are qualitatively similar to those reported in section 3.3.3 and are presented in table 3.2.

3.3.1 Preferences

We take advantage of the following features of our data set to construct the income-specific price indices. First, following the work of Broda and Romalis (2009), we allow the set of UPCs and the expenditure weights of our price indices to be income specific. We also allow product entry and exit to affect our cost of living indices, following the method developed by Feenstra (1994), and we permit the effect to vary across income groups. To this end, we compute income-specific elasticities of substitution within product modules to allow for non-homotheticities across income groups. Further, we use the actual price paid by the household, as opposed to the posted price, to calculate the price indices. This method relaxes the standard assumption that each income group faces the same price increases for a particular UPC. As a result, our price indices capture differences in the consumers' shopping behavior.[10]

We specify the upper level utility function of income group I as:

$$U_t^I = \left(\sum_{g \in G} (C_{gt}^I)^{\frac{\sigma - 1}{\sigma}} \right)^{\frac{\sigma}{\sigma - 1}} \tag{3.1}$$

where the product modules are indexed by g, σ is the elasticity of substitution across modules,

10. As Chevalier and Kashyap (2014) show, some consumers chase discounts, and thus, the actual prices paid are substantially lower than the posted prices. They argue that price indices that rely on posted prices overstate the price level experienced by consumers.

and G is the set of all product modules. The set G is fixed over time and is equal across income groups.[11] We model the lower tier as:

$$C_{gt}^{I} = \left(\sum_{u \in U_g^I} (d_{ugt}^{I} C_{ugt}^{I})^{\frac{\sigma_g^I - 1}{\sigma_g^I}} \right)^{\frac{\sigma_g^I}{\sigma_g^I - 1}} \tag{3.2}$$

where C_{ugt}^{I} is the total quantity of UPC u that is consumed in product category g by income-group I at time t, σ_g^I is the income-specific elasticity of substitution within product category g, and U_g^I is the set of all possible UPCs within a product category g that group I consumes. The set of existing UPCs in period t is a subset of this set (i.e., $U_{gt}^I \subset U_g^I$) and can vary over time. We define the set of consumed UPCs throughout the whole time period as $U_g^I = U_{gt}^I \cap U_{gs}^I$. The parameter d_{ugt}^I plays a crucial role in the analysis because it captures the perceived quality of each UPC at the income-group level. Furthermore, an increase in d_{ugt}^I raises the demand for UPC u and also lowers the unit costs. The remarkable feature of the exact price index (EPI) developed by Diewert (1976) is that the index does not depend on the unknown quality parameter as long as it is constant over time (i.e., $d_{ug,t-1}^I = d_{ugt}^I$). This is an assumption we make in the construction of our price indices.

If the set of UPCs available for each group is fixed over time, Sato (1976) and Vartia (1976) derive the EPI in the case of any multilevel CES utility function as:

$$\prod_{g \in G} \left\{ \prod_{u \in U_g^I} \left(\frac{p_{ugt}^I}{p_{ugt-1}^I} \right)^{\omega_{ugt}^I} \right\}^{\omega_{gt}^I} \tag{3.3}$$

This is the geometric mean of the price changes in each UPC u that belong to the set U_g^I, where the weights are ideal log-change weights that are computed as follows:

$$s_{ugt}^I = \frac{p_{ugt}^I C_{ugt}^I}{\sum_{u \in U_g^I} p_{ugt}^I C_{ugt}^I}$$

11. Quality weights to the aggregate consumption good are redundant given that we later assign a quality weight to each individual UPC.

and ω_{ugt}^{I} is defined similarly:

$$\omega_{ugt}^{I} = \frac{\dfrac{s_{ugt}^{I} - s_{ug,t-1}^{I}}{\ln s_{ugt}^{I} - \ln s_{ug,t-1}^{I}}}{\sum_{u \in U_g^I} \dfrac{s_{ugt}^{I} - s_{ug,t-1}^{I}}{\ln s_{ugt}^{I} - \ln s_{ug,t-1}^{I}}}$$

The weights capture how likely consumers are to substitute within and across categories, and they are income-group specific. The products that are highly substitutable can receive a much smaller weight than the products whose share barely moves after a price change; these products get a weight close to their average expenditure share.[12] In this specification of the price index, the prices vary by income group to capture the differences in consumers' shopping behavior.

Because this index ignores new and disappearing product varieties, it is often referred as a "conventional" exact price index (CEPI). In order to take into account increases in quality due to the entry and exit of UPCs, Feenstra (1994) generalize the EPI to allow for different, but overlapping, sets of goods in the two periods. If there is a set of UPCs available in both periods, and their quality parameters are constant, then we can derive different cost of living indices by income group by allowing for product creation and destruction as follows:

$$\prod_{g \in G} \left\{ \prod_{u \in U_g^I} \left(\frac{p_{ugt}^{I}}{p_{ugt-1}^{I}} \right)^{\omega_{ugt}^{I}} \times \left(\frac{\lambda_{gt}^{I}}{\lambda_{gt-1}^{I}} \right)^{\frac{1}{\sigma_g^I - 1}} \right\}^{\omega_{gt}^{I}} \tag{3.4}$$

where

$$\lambda_{gt}^{I} = \frac{\sum_{u \in U_g^I} p_{ugt}^{I} C_{ugt}^{I}}{\sum_{u \in U_{gt}^I} p_{ugt}^{I} C_{ugt}^{I}}$$

The cost of living index for income group I is now adjusted for a new goods bias between periods t and $t-1$. As Broda and Weinstein (2010) show, this correction is important because of the cyclical nature of the bias. The magnitude of the bias depends on the ratio of the share of common goods and the elasticity of substitution. The ratio of the share of common products

12. The limit of $\frac{s_{ugt}^{I} - s_{ug,t-1}^{I}}{\ln s_{ugt}^{I} - \ln s_{ug,t-1}^{I}}$ when $s_{ugt}^{I} \to s_{ug,t-1}^{I}$ is $s_{ug,t-1}^{I}$.

within a category in period t relative to the share of common products in period $t-1$ within category g is $\lambda^I_{gt}/\lambda^I_{gt-1}$. This term captures the importance of quality shifts due to the creation and destruction of UPCs in category g for income group I. The within category elasticity of substitution σ^I_g indicates the effect that these shifts have on the price index and varies by income. As the elasticity of substitution rises (i.e., highly substitutable UPCs), a given movement in the share of common goods over time has a smaller effect on the price index.[13]

3.3.2 Estimating Elasticities by Income Group

In order to estimate the elasticity of substitution within a product category σ^I_g, we rely on the method developed by Feenstra (1994) and extended by Broda and Weinstein (2006) and Broda and Weinstein (2010). The procedure consists of estimating a demand and supply equation for each UPC by using only the information on prices and quantities. In order to overcome the standard endogeneity problem, we use information on the joint distribution of supply and demand parameters provided by the data. We start by writing supply and demand equations. We define $\tilde{p}^I_{ugt} = p^I_{ugt}/d^I_{ugt}$ as the quality-adjusted price. Under a CES demand system, unobserved quality-adjusted prices can be recovered from observed purchases. That is, the share of consumption of UPC u within category g can be written as a function of the quality-adjusted price as follows:

$$s^I_{ugt} = \left(\frac{\tilde{p}^I_{ugt}}{P^I_{gt}}\right)^{1-\sigma^I_g} \tag{3.5}$$

where P^I_{gt} is the minimum unit cost function of the subutility function in equation 3.6 and is given by the following expression:

$$P^I_{gt} = \left(\sum_{u\in U^I_{gt}} (\tilde{p}^I_{ugt})^{\sigma^I_g}\right)^{\frac{1}{\sigma^I_g}}$$

13. We do not allow the entry and exit of a product category. Because the definition of a product category in the Nielsen data is disaggregated, we find some product categories where no single transaction is made in a quarter. We keep balanced product categories over time, requiring more than ten UPC's covered by each income groups every quarter. Our baseline results are robust to various cutoffs. In addition, we exclude the cigarette category because its price mainly depends on state taxes instead of a market mechanism.

Equation 3.5 shows that, under CES preferences, UPCs with higher quality have lower quality-adjusted prices and higher market share. Using equation 3.5 and the supply curve of each UPC we derive a system of differenced supply and demand equations to estimate the demand elasticities.[14] For each income group we define:

$$\Delta_g^k ln\, s_{ugt} = -(\sigma_g - 1)\Delta_g^k\, ln\, p_{ugt} + \epsilon_{ugt}^{k_g} \tag{3.6}$$

$$\Delta_g^k ln\, p_{ugt} = \frac{\omega_g}{1+\omega_g}\Delta_g^k\, ln\, s_{ugt} + \delta_{ugt}^{k_g} \tag{3.7}$$

Equation 3.6 and 3.8 are the demand and supply equations of UPC u in product category g differenced with respect to time and with respect to a benchmark UPC of the same product category. The k^{th} good corresponds to the largest selling UPC in each product category. The k-differencing removes any category level shocks (e.g., seasonal shifts, advertising, assembly line shocks) from the data. As a result, we assume that $\epsilon_{ugt}^{k_g}$ and $\delta_{ugt}^{k_g}$ are uncorrelated (i.e., $\mathbb{E}_t(\epsilon_{ugt}^{k_g}\delta_{ugt}^{k_g}) = 0$). In addition, we assume that the elasticities of supply and demand are the same for each UPC within a product category. Feenstra (1994) shows that under these two assumptions and because of the panel structure of our data set, we can define a set of moment conditions for each product category to identify the within product category demand elasticities:

$$G(\beta_g) = \mathbb{E}_t(\nu_{ug}(\beta_g)) \quad \forall\, u, g \tag{3.8}$$

where $\nu_{ug} = \epsilon_{ugt}\delta_{ugt}$ and $\beta_g =$ is a vector containing σ_g and ω_g. Having as many moment conditions as UPCs in a product category, $\hat{\beta}_g$ is estimated using the GMM under the assumption that the relative variances of the demand and supply shocks differ across barcodes.[15] To obtain estimates of the elasticities for each income group, we first distinguish between the purchases made by each income group in each product category. Then, we follow Broda and Weinstein (2010) and implement this procedure by using quarterly data and considering only the elasticities that are

14. To simplify the notation, in what follows we omit the superscript I.

15. See Feenstra (1994) and Broda and Weinstein (2006) for a more detailed description of the procedure and the derivation of the main equations.

economically feasible (i.e., $\sigma_g > 0$ and $\omega_g > 0 \; \forall \, g$).[16]

Table 3.1 shows the estimated distribution elasticities by income group. The median within-category elasticity is around 17 for all income groups. This number is unsurprisingly high given that it describes the substitutability of a UPC within a finely defined product category. The table shows that higher-income consumers are slightly less price elastic but the differences in the distribution of within-category elasticities across income groups are not substantial.

3.3.3 Results

Figure 3.1 plots the cost of living indices by income group from 2004 to 2010. The figure shows that the indices for all income groups track each other closely but drastically vary during the Great Recession. Over our sample period, the average annual inflation in the cost of living of the highest quartile is 0.80% while that of the lowest quartile is 1.39%. Table 3.2 shows that this gap appears mainly during the Great Recession. During the three years prior to the recession there is virtually no difference between the average annual inflation of the highest quartile and the lowest quartile of the income distribution. Remarkably, this difference is on average 0.62 percentage points in the periods following the start of the recession.[17]

The differences we observe across income groups could come from any of the income-specific components in our EPI. To understand which component of the EPI is driving the gap, we first study the correction term for the new goods bias. This term has two components: the ratio of the share of common goods across periods (i.e., $\lambda^I_{gt}/\lambda^I_{gt-1}$) and the elasticity of substitution (i.e., σ^I_g). In our sample period, the share of common goods across periods does not differ significantly across income groups. Furthermore, Table 3.1 shows that there are no substantial differences in the elasticity of substitution across the groups either. Therefore, the difference we observe comes, almost entirely, from the CEPI; removing the bias correction term does not reduce the gap across

16. If the procedure renders imaginary estimates or estimates of the wrong sign, we use a grid search to evaluate the GMM objective function. The objective function is evaluated at intervals that are approximately 5% apart. Overall we find feasible elasticites for around 80% of the product categories we consider.

17. The gap is approximately 0.5 percentage points when we exclude categories in the fresh produce and the dry grocery department indicating that our results are not driven by the Global Food crisis of 2007-2008.

income groups as shown in Panel B in Figure 3.2 and in table 3.3.[18] Further, given that the estimated elasticities of substitution only appear in the bias correction term, the variation in this parameter across income groups cannot explain the differences we observe.[19] Without the bias correction term, the annual inflation rate for all income groups is higher because in this specification the price indices do not account for the benefits from the introduction of new products. In this case, the difference between the lowest and the highest income groups is 0.48 percentage points, only slightly lower than the benchmark gap. This difference indicates that the entry and exit of products benefit high-income households disproportionately. However, this channel only explains about 15% of the differences in the EPI.

Next, we examine the components of the CEPI. The gaps we observe across groups could be driven by the differences in the prices paid for the same product due to differences in shopping intensity, the differences in the consumers' ability to substitute across product categories ω^I_{gt}, or by the differences in the consumers' ability to substitute across goods within product categories ω^I_{ugt}. To rule out the possibility that the differences in the prices they pay for the same product drives the differences we observe across groups, we compute our price indices using the posted prices by retailers instead of the effective prices paid by consumers. Panel C in Figure 3.2 shows that the differences persist and are of a similar magnitude as in our benchmark index. To explore whether the gaps come from the consumers' readiness to substitute across product categories, we compute our price index by fixing ω^I_{gt} for each income group to its baseline value. Panel D shows that the differences we observe across income groups prevail; table 3.3 shows that the difference between the top and bottom quartiles is on average 0.52 percentage points. Collectively, given that the disparities across groups persist after controlling for each component of the EPI, the results of this analysis show that the consumers' ability to substitute within product modules (i.e. changes in

18. A concern might be that the transactions we observe do no accurately reflect the choice set of an income group. The possibility exists, for instance, that our sample under-represents the availability of low quality products to high-income households in the pre-recession period. We minimize this potential bias by requiring only one transaction by any household in a given group and period to capture the presence of a specific variety in the choice set of that group. Thus, most of the products high-income households switch to during the recession were available before the crisis and, if not, their presence is captured by our variety correction term.

19. The differences we observe in the price indices across income groups are qualitatively the same when we compute them using a common elasticity of substitution (we use 11.5) or the same elasticities of substitution for every income group and product category.

ω^I_{ugt}) is what drives the main difference across the income groups

Nonetheless, the price indices developed in this section are only indicative of what happened during the Great Recession. The main disadvantage to this approach is that it does not control for certain household characteristics that could be more prominent in certain income groups. Similarly, people of different income groups are concentrated in different geographic regions and our indices do not take that into account. Further, this approach does not rule out the possibility that the differences in the cost of living's inflation do not arise from the pricing behavior of firms and/or retailers instead of the changes in the household behavior. For example, the case could be that particular firms or stores where wealthy people purchase lower their prices more. For this reason, in the next section, we use transaction level data to assess the impact of different shopping activities on the average prices paid by households of different income groups controlling for different household and geographic observables as well as store characteristics.

3.4 Reduced Form Evidence on Shopping Behavior

In this section we examine the contribution of each shopping activity on the prices paid by households of different income groups. Given the richness of our data, we are able to include demographic and geographic controls as well as various fixed effects to examine the contribution of each margin. For simplicity, we focus on comparing the shopping behavior of the highest quartile of the income distribution versus the lowest quartile.[20]

Our empirical specification is the following:

$$\log P^h_{i,j,m,t} = \lambda + \delta_t + \delta_t \times D^h + \beta_j + \alpha_h + \gamma_m + \epsilon^h_{i,j,m,t} \tag{3.9}$$

where $P^h_{i,j,m,t}$ is the unit value paid by household h for product i of category j in market m at time t. The δ_t denotes the time fixed effects, and D^h is an indicator that equals one if household h belongs to the highest quartile of the income distribution and zero if it belongs to the lowest quartile. The β_j denotes the product category fixed effects, the α_h denotes the household

20. All our results are robust to conducting a similar exercise using information of all income groups at the same time. These results are available on request.

fixed effects, and the γ_m is the county level fixed effects. In this context, δ_t represents the mean percent deviations from the price of the base period for the lowest quartile after controlling for demographics, product categories, and regions; $\delta_t \times D^h$ represents the difference of the mean percent deviations across income groups. We cluster the standard errors at the product category level.

Figure 3.4 plots the time fixed effects using the first quarter of 2004 as the base period. The figure shows that the mean percent deviations follow a similar pattern as the price indices in the previous section. Panel A in Figure 3.3 plots the interaction term. If the mean percent deviations of the two income groups are the same, then the line in the graph should be completely flat. Instead, the line starts flat and then decreases at the beginning of the Great Recession. This setup shows the mechanisms that households use to increase their shopping effort during bad times and whether they use them differently.

We first control for shopping intensity by adding, at the household level, the mean units purchased per category in a given time period (buying bulk), the fraction of items purchased on sales, the fraction of items purchased with coupons, the fraction of expenditures in generic brands, and the number of shopping trips a household makes in a given quarter.[21] The inclusion of these variables controls for the possible heterogeneity biases that could arise from shoppers who search intensely for the lowest price and who purchase the same good in the same store for less.[22]

As depicted in Panel B, the difference between the highest and the lowest income quartiles barely reduces. This finding shows that the differences in shopping intensity do not explain the differences in the price indices. However, that does not mean that households are not increasing their shopping intensity during the recession, in fact they do. But the question is what drives the differences in the inflation that they face. Next, we control for the reallocation of expenditures across retailers by introducing store fixed effects to the benchmark regression.[23] These controls for possible biases in the retailer's heterogeneity arise when stores systematically charge different

21. Table C.2 in Appendix **??** presents the descriptive statistics for these variables.

22. This procedure is similar to the one used by Handbury and Weinstein (2014) to control for different types of heterogeneity biases common in the construction of price indices.

23. For about 50% of the transactions the store code is missing. Instead we approximate a store as a retailer×county combination.

prices for the same good. By including these controls we eliminate the possibility that households could be reducing their relative prices by reallocating their purchases across retailers. Panel C shows that the difference in the average prices paid reduces slightly but the standard errors widen. This finding indicates that this is not the main mechanism behind the gap in the price indices.

To control for quality substitution, we add the UPC fixed effect to the benchmark regression. Without the UPC fixed effects, consumers could be paying lower average prices by reallocating their expenditures across brands within the same product category. The inclusion of the UPC fixed effects eliminates this channel. Panel D shows that at the end of the recession, the difference between the inflation rates faced by the highest and lowest income quartiles vanishes. This result shows that the differences we observe in price indices comes from the consumers' ability to substitute lower for higher quality brands within a product category. The graph shows that their ability to substitute increases with household income during economic downturns and that it provides an additional margin of adjustment to higher income households. In Figure C.3 in Appendix **??**, we show the same calculations by adding controls cumulatively to the benchmark. Panel D in Figure C.3 shows that even if we control for shopping intensity and the store fixed effects at the same time, the differences among the prices paid by consumers of different income groups only disappear when we include the UPC fixed effects. The results strongly indicate that the quality substitution channel is the most important margin for explaining the difference in the inflation in the cost of living across consumers of different income groups during the Great Recession.

This result also occurs when we analyze the cross-sectional relationship between the relative prices paid and household income.[24] Figure 3.5 shows how much more or less households pay per unit for products within a category. In the figure, the relative prices are measured in a regression of the log unit price paid against income category dummies and product category, region, chain, quarter and household fixed effects. Each dot represents how much more households in each income category pay per unit for products within a category than households in the lowest income category of between $5,000 and $8,000. The figure shows that a distinct upward slope exists; high-

24. This relationship is also known as quality Engel curve. The curves are computed under the premise that across households, at a point in time, those paying higher unit prices are buying higher quality goods (typically richer households). See Bils and Klenow (2001) for a detailed description.

income households pay around 8%-10% more for products in the same category than low-income households.[25] As income declines, households reduce the relative price paid by switching to cheaper products within a category. This relation is true when we split the sample for the years before and after the recession, particularly for high levels of income. Nonetheless, the figure shows that after 2008 the slope distinctively flattens for lower income groups, particularly for households below the median of the income distribution. That is, in the aftermath of the recession, the substitution margin is no longer available for low-income households. And, consistent with our previous findings, the ability of high-income households to take advantage of this margin is significantly higher.

3.5 Direct Evidence on Shopping Behavior

In this section, we focus on providing a direct measure of quality substitution and calculating the extent to which it changes when economic conditions deteriorate. Because we observe the share of expenditures of each household in each product, we construct a household-specific measure of the average quality of goods they purchase. We then explore the sensitivity of this measure to changes in local economic conditions. Here we show that quality substitution is a margin households use not only when aggregate conditions worsen but also when local conditions deteriorate. We also show that high income households are better able to use this margin partly due to the fact that most of the purchases of low-income households already come from products ranked low in the quality distribution. Finally, we explore other margins such shopping intensity and store switching. Our results suggest that quality substitution is the most important margin of adjustment during economic downturns.

25. This is consistent with the findings of Broda, Leibtag, and Weinstein (2009), Handbury (2013), and Faber (2014). In particular, Handbury (2013) shows that high-income households also pay slightly higher prices for the same UPC (same regression as above but controlling for product fixed effects instead of category fixed effects) but in this case the slope of the log unit price paid is much smaller. Thus, high-income households are paying more mainly because they purchase more expensive products within a product category and not because they pay more for the same UPC.

3.5.1 Combining Retail and Household Level Data

In order to study the extent of quality substitution and store switching, we construct a "ranking" of products and stores based on their relative prices. To do this, we use the RMS for 2006 to 2010 to obtain a universe of products sold by the retailers in our sample. Using information on 76 billion transactions at the store level, we calculate the average relative prices of products to approximate their quality. We consider a product with higher relative prices to be higher on the quality ladder. For example, within the milk category, organic whole milk is higher on the quality ladder than regular whole milk given that it is sold, on average, at a higher unit price. Then, we combine this information with the Homescan to construct a household-specific estimate of product quality in a given period of time. Our final sample consists of 54,558 stores that 83,564 households across the United States visited within 49 continental states to purchase 1,258,962 products within 1,074 product modules.

3.5.2 Definitions of Shopping Behavior Variables

Using this data set, we focus on the three ways consumers can save after they suffer income shocks: increasing shopping intensity, switching to low quality stores, and substituting down to lower quality products within a category.

Shopping Intensity: To approximate shopping intensity, we use the number of shopping trips and the fraction of items purchased on sale.[26] A shopping trip is defined by the date and location of the transaction (Aguiar and Hurst (2007)). That is, transactions at two stores on the same day are counted as two trips. Similarly, two transactions at the same store on two different days are counted as two trips. By making more shopping trips, households could exploit store and manufacturer discounts more frequently.[27] In addition, the Homescan data defines a transaction

26. We perform the same analysis with coupon usage, purchases of generic products, and buying in bulk and find similar results. This is consistent with the results by Petev, Pistaferri, and Eksten, who argue that these activities have similar business-cycle patterns.

27. We do not have information on the length of each shopping trip. Aguiar and Hurst (2007) supplements the consumer panel data using the American Time Use Survey (ATUS). Since the time spent at shopping in ATUS has remained relatively stable over our sample period (Petev, Pistaferri, and Eksten), we only focus on the number of

as being on sale if the household records that the item purchased involved a deal. We measure the fraction of items purchased on sale by a given household in a given quarter.

Store Switching To assess the reallocation of expenditures across stores, we first construct the stores' quarterly relative prices as Coibion, Gorodnichenko, and Hong (2015b) suggest. First, for each UPC-level good j in store s and county m, we calculate the log-difference between the price of good j in store s and the median price for good j across all of the stores in a given county and quarter.[28]

$$R_{jst} = \log \frac{P_{jst}}{\bar{P}_{jt}}$$

where R_{jst} is the relative price and $\bar{P}_{jt}$ is the median price.

We then compute the average relative price for a store across the set of UPC products available in the county. The average relative price (quality) of a store is

$$Q_{st}^{\text{store}} = \sum_{j} \omega_{jsmt} R_{jst}$$

where ω_{jst} is a revenue weight. The Q_{st}^{store} captures how far a store's average price level is from the median price level in a given county and quarter. We construct an average relative price (quality) for the stores at which household h shopped in quarter t as follows:

$$Q_{ht}^{\text{store}} = \sum_{s} \psi_{hst} Q_{st}^{\text{store}}$$

where ψ_{hst} is an expenditure weight. The Q_{ht}^{store} represents the average quality of a store where household h consumes. Given this measure, we assess the store switching phenomenon at the household level.

Product Quality Substitution Similar to the average relative price (quality) of the store, we

shopping trips.

28. To simplify the notation, we omit the subscript m for county.

construct the average relative price (quality) of the UPC-level good within each product category. First, we measure the log-difference between the price of good j in store s and the median price for category c in store s.

$$R_{jst} = \log \frac{P_{jst}}{\bar{P}_{cst}}$$

where R_{jst} is the relative price and $\bar{P}_{cst}$ is the median price of category c in store s. Therefore, if the price of a high quality type of milk, say organic milk, is much higher than the median price of milk in that store, then R_{jst} is positive and high. We compute the average relative price for the UPC-level goods across the set of stores within a county. The average relative price (quality) of a good is

$$Q_{jt}^{\text{product}} = \sum_{s \in \Omega} \omega_{jst} R_{jst}$$

where ω_{jst} is a revenue weight. The Q_{jt}^{product} captures how far a product's average price level is from the median price level of a category in a quarter. We construct an average relative price for the products that household h buys in quarter t as follows:

$$Q_{ht}^{\text{product}} = \sum_{j} \psi_{hjt} Q_{jt}^{\text{product}}$$

where ψ_{hst} is an expenditure weight. A low (high) value of Q_{ht}^{product} indicates that household h buys relatively low (high) quality products in quarter t.

Figure 3.6 shows the distribution of product quality for households belonging to the lowest and highest quartiles of the income distribution before the Great Recession. Almost 65% of the wealthier households purchased goods that were, on average, above the median quality within a product category. On the other hand, around 60% of low-income households consumed products below the median quality. We observe, nonetheless, substantial heterogeneity across the households within an income group; some wealthy households bought lower-than-average quality products, and vice versa. Figure 3.7 depicts how these distributions evolved after the recession. It shows that wealthier households, particularly those that on average purchased very high quality goods, were able to switch to lower quality goods in the aftermath of the recession. The figure also shows that the distribution of low income households stayed mostly stable with a small increase in the fraction of households purchasing goods belonging to the lowest quartile of the product quality distribution. This finding suggests that low-income households were less able to use product substitution as a margin of adjustment during this period.

3.5.3 Change in Shopping Behavior after an Income Shock

We use our household-specific estimates of product and store quality, along with our shopping intensity measures, to assess the changes in shopping behavior after an income shock. We represent an income shock by the change in the unemployment rate of the county where a household lives.[29] We use the following empirical specification:

$$Y_{ht} = \lambda + \theta \times \mathrm{UR}_{ct} + \alpha_h + \delta_t + \epsilon_{ht} \tag{3.10}$$

where Y_{ht} represents the variables for the shopping behavior of household h in quarter t, and UR_{ct} is the unemployment rate of county c where a household lives. The α_h and δ_t are household and time fixed effects. Because the error term is likely to be serially and cross-sectionally correlated, we use Driscoll and Kraay (1998) standard errors. When all fixed effects are added, the estimates of θ assess the strength of the correlation between the changes in the local unemployment rate and the households' shopping behavior. While aggregate shocks could lead to simultaneous movements in households' shopping behavior and local economic conditions, we control for time fixed effects that eliminate this issue. Therefore, by exploiting only the cross-sectional variation, our aim is to show that quality substitution is also a relevant margin of adjustment when local unemployment rises.

The results are presented in Table 3.4. The table shows that households in a region with a higher

29. We decide not to use the annual income reported by households to represent income shocks for several reasons. First, income data is top-coded. Thus, we cannot capture income changes for the wealthier households. And, second, the average number of years that panelists stay in the sample is only three years and their self-reported income does not change much over the sample period. So, we only observe, on average, less than two income changes per household.

unemployment rate reallocate their expenditures to stores with lower average relative prices. This is consistent with Coibion, Gorodnichenko, and Hong (2015b) who find that households reallocate their purchases toward low-price stores when economic conditions worsen.[30] Furthermore, the average product quality is also negatively correlated with the unemployment rate. A one percentage point increase in the unemployment rate is associated with a 0.07 percentage point decrease in the average product quality of the households' consumption bundles. This association indicates that households that face a negative income shock reallocate their expenditures to lower quality products within the same product categories. This finding is consistent with recent work by Jaimovich, Rebelo, and Wong (2015) and Dubé, Hitsch, and Rossi (2015). They find consumers trade down in the quality of the goods and services they consume during the recession.[31] Further, the coefficients for the variables that measure shopping intensity are positive and significant. Households are likely to make more shopping trips and buy more items on sale when the local unemployment rate increases.

Next, we examine which types of households are most likely to adjust their shopping behavior. Here again we focus on comparing the behavior of the top quartile versus the lowest quartile for simplicity. To do so, we use the following empirical specification:

$$Y_{ht} = \lambda + \theta_1 \times \text{UR}_{ct} + \theta_2 \times \text{UR}_{ct} \times D_h + \alpha_h + \delta_t + \epsilon_{ht} \tag{3.11}$$

where Y_{ht} and UR_{ct} are specified as before and D_h is a dummy variable that equals one if the household belongs to the highest quartile and zero if it belongs to the lowest quartile of the income distribution. The θ_2 indicates whether the response of households to changes in the local unemployment rate varies with their income level. Table 3.5 shows that θ_2 is significantly positive in all margins of shopping behavior. This indicates that higher income households adjust the quality of the products they purchase, the stores at which they buy, and their shopping intensity more

30. Our analysis is complementary to Coibion, Gorodnichenko, and Hong (2015b) who use household information from only two cities (Eau Claire, WI and Pittsfield, MA) to assess the extent of store switching. While we use a data set with a much wider geographical coverage, our results are consistent with their findings.

31. Jaimovich, Rebelo, and Wong (2015) also approximates the quality of goods and services by their relative price. Their evidence of quality substitution is from the Yelp! website matched with US Census of Retail Trade and Compustat. Dubé, Hitsch, and Rossi (2015) finds a negative effect of income on private-label shares.

with respect to low-income households. This adjustment could reflect the fact that high-income households more easily devote additional resources to shopping.[32] For the low-income households, only the quality substitution of stores is significant. The substitution of product quality is almost zero and not significant. In other words, the product quality mechanism is only present for high income households. A one percentage point increase in the unemployment rate is associated with a 0.16 percentage point decrease in the average product quality of households for this group. This is relevant given that on average the unemployment rate increased almost 5 percentage points during the recession. Our results show that, because poor households purchase a larger share of their products at the bottom of the distribution of product quality, they lack the substitution margin when they face an income drop. Higher income households, on the other hand, are able to take advantage of this margin to mitigate income shocks.

3.6 Supply Side Explanations

The possibility exists that the differences in the cost of living's inflation that we observe stem from changes in the pricing behavior of retailers or firms rather than from the adjustment of the consumers' behavior. This could be the case, for instance, if retailers reduce the prices of products targeted to high-income consumers more or if these consumers source their consumption from firms that decrease their prices more on average. In this section we explore whether our results are driven by this possibility.

3.6.1 Posted vs Effective Inflation

A growing body of literature studies the cyclicality in both the prices posted by retailers as well as the effective prices actually paid by consumers See, e.g., the recent work of Kaplan and Menzio (2015), Stroebel and Vavra (2016), Chevalier and Kashyap (2014), and Beraja, Hurst, and Ospina (2016). This task was in fact very difficult until recently because it required the availability of data for both quantities and prices. In a recent study, Coibion, Gorodnichenko, and Hong

32. Our results are consistent with the findings of Stroebel and Vavra (2016) who also show heterogeneous responses of different consumers after wealth shocks. Using a similar approach, they find that renters are less likely to adjust their shopping behavior after changes in housing wealth than homeowners.

(2015b) show that consistent with the idea of significant consumer reallocation of expenditures in response to economic conditions, effective price inflation is indeed more cyclically sensitive than the inflation in posted prices. Although they argue that the discrepancy between the cyclical changes reflects consumers reallocating their expenditures across stores, they also provide evidence of the importance of the reallocation of expenditures across different goods within a category. That is, provided some dispersion of prices across goods within the same category in any given period, deteriorating local economic conditions should lead price-sensitive consumers to reallocate some of their consumption to goods with a lower unit price. In fact, Coibion, Gorodnichenko, and Hong (2015b) show that cross-good substitution is quantitatively more important than store switching in explaining the cyclical sensitivity of effective inflation. They construct two measures of effective inflation. In the first, they fix the composition of the consumption basket but allow consumers to switch stores in order to focus on across-store substitution. In the second, they fix the store weights but allow consumers to substitute goods in the basket in order to study the effect of cross-good substitution on effective inflation. They report that the second measure lowers the inflation of effective prices relative to the inflation in posted prices at least 0.1 percentage points more than the first. Under the second measure, a 2% increase in the local unemployment rate lowers the inflation rate of effective prices by 0.3% to 0.4% at an annual rate below the inflation rate of posted prices. This result shows that not only effective prices, precisely the object that our price indices aim to capture, are significantly more flexible and sensitive to economic conditions than the underlying prices set by retailers or firms, but also that cross-good substitution plays a large role in explaining this sensitivity. This finding indicates that product substitution is more likely to drive the inequality in the cost of living that we observe among consumers than the pricing behavior of retailers.

3.6.2 Price Discrimination

An alternative explanation could be retailers' second degree price discrimination motive that induces them to charge different prices to different consumers. This could be the case if they try, for instance, to minimize the possibility of consumers switching to less preferred goods to save money.

Chevalier and Kashyap (2014) argue that retailers charge different prices to "active" consumers, who are more likely to substitute across products in narrowly defined categories, and "passive" consumers, who are not.[33] If the retailers forecast that high-income consumers are more likely to switch brands given that they are more likely to purchase high-priced products, then they might lower the prices of more expensive products to limit the degree of switching among these consumers. If this is the case, then our results could be driven by price discrimination rather than consumers substituting across brands to smooth their consumption. In order to explore this hypothesis we divide the products into two types, high-priced and low-priced, which follows the procedure developed by Chevalier and Kashyap (2014). We calculate the total amount spent in each week at each store divided by the total volume sold in that week. We call this the benchmark price per volume for that store in that week for each category. We define the high-priced products as those products that are on average 25% above the benchmark price and low-priced products as those that are on average 25% below the benchmark price. This separation is consistent with Figure 3.5 and is meant to distinguish products mainly consumed by high-income households, such as premium brands, from those mainly consumed by low-income households.[34] Then, we construct the posted price inflation for each type in order to check whether the inflation of high-priced products is more sensitive to changes in local economic conditions than that of low-price products. To determine the posted price's inflation for each type we follow the method developed by Coibion, Gorodnichenko, and Hong (2015b). We first calculate the average retail price in a given week by using the RMS data as follows:

$$P_{msctj} = \frac{TR_{msctj}}{TQ_{msctj}}$$

where m, s, c, t, and j are the market, store, product category, time, and UPC respectively. We then compute the inflation rates of the UPC-specific monthly posted price as $log(\frac{P_{mscjt}}{P_{mscj,t-1}})$. We then aggregate at the market, product category, product type, and month level using three weight-

33. In their model, the seller optimally sets prices to encourage deal-driven brand switching only by those consumers who get the least disutility from switching brands. Their empirical results show that many consumers get a low disutility from switching brands.

34. Faber and Fally (2017) provide cross-sectional evidence that wealthier households source their consumption from brands with higher unit values and larger market shares.

ing schemes: i) equal weights, ii) expenditure shares for each market and year (market-specific), and iii) cross-market expenditure shares for each year (common). We also add the monthly inflation rates to the annual inflation rates: $\bar{\pi}^{post}_{mct}$.[35] By comparing the differential response of the inflation rates of high-priced versus low-priced products, we can approximate the degree to which the retailers charge differential prices to customers in response to changes in economic conditions. Therefore, we use the following specification:

$$\bar{Y}_{mcqt} = \beta_1 \mathrm{UR}_{mt} + \beta_2 \mathrm{UR}_{mt} \times D^{H}_{mct} + \lambda_t + \theta_{mcq} + \varepsilon_{mct} \tag{3.12}$$

where $\bar{Y}_{mcqt}$ refers to the posted inflation of the products in market m, product category c, type q, and time period t; UR_{mt} is the local unemployment rate; and D^{H}_{mct} is a dummy variable equal to one if the products in category c and market m are high-priced. Table 3.6 presents the results. We do not find any evidence that the prices of high-priced products decrease more than those of low-priced products when economic conditions deteriorate. This is true under the three weighting schemes and is robust to the addition of time fixed effects. The specifications in columns 4-6 are particularly relevant because one might interpret this estimate as causal. This is because we are able to control for aggregate shocks that simultaneously move prices and local economic conditions and due to the fact that most of the goods in our sample are not produced locally. In Appendix **??**, we show that our results are robust to different definitions of high-priced products (e.g., 15% above and below the benchmark price) and to alternative methods to construct the posted price inflation such as the one suggested by Gagnon, Lopez-Salido, and Sockin (2015). Importantly, our main conclusion does not change: There is no evidence that when economic conditions deteriorate, retailers lower the prices of high-priced products more than the prices of low-priced products to limit the degree of brand switching by higher income consumers.

35. Our measure of posted inflation is a geometric price index similar to the one constructed by the BLS with the exception that we update expenditure weights once a year whereas the BLS updates its weights every five years.

3.6.3 Firm's Price-Setting Behavior

There could also be the concern that the actual price setters are not the retailers but the firms, and a supply shock to the firms during the recession might affect consumers differently. For example, Gilchrist, Schoenle, Sim, and Zakrajšek (2017) use confidential price data on producers to show that during the crisis firms with relatively weak balance sheets increased prices while firms with strong balance sheets lowered prices. If the firms with strong balance sheets are mainly selling to high-income consumers, our results could be driven by changes in the firms' price-setting behavior. In order to explore whether a clear link between firms of certain characteristics and consumers of different income levels exists, we first identify the manufacturer of each UPC-level product in the Homescan data by using the first six to nine digits of the barcode.[36] Then, in order to link the manufacturer to the firm that owns each product, we merge the Homescan data with information obtained from GS1, the single official source of barcodes in the United States. The GS1 data contains the list of all US barcodes along with information on the firms that own them, such as their location and their country of origin.[37] Once we identify the firms producing each product, we compute the total expenditures of the consumers in our sample at the firm level.[38] Table C.3 in Appendix **??** shows the share of total expenditures by firm. The table shows that the top 15 firms account for more than 60% of the total expenditures in the Homescan data. This amount is due to the fact that larger firms own several barcode-level products along a diverse set of product categories. For example, Procter and Gamble owns more than 9,000 barcodes that belong to many product categories such as baby care products, detergents, abrasive cleaners, shampoos, and cookies. This pattern also holds if we consider the expenditure shares of each quartile of the income distribution separately. The lowest quartile spends on average 63% of its total expenditures on products owned by the top 15 firms whereas the top quartile spends around 61%. That is, most of the products purchased by consumers belong to the same set of firms. The table also shows that

36. See Figure C.4 for examples of a 6-digit and a 9-digit prefix.

37. GS1 is the official source for firm prefixes used to create barcodes in over 100 countries. The barcodes created by GS1 are unique and used in a variety of industries including retail. The starting fee for a UPC barcode is $250 (see www.gs1us.org for more details).

38. The Homescan data has less than 5% of the non-US barcodes. We were not able to find a firm identifier for these products so we dropped them in this part of the analysis.

the share of expenditures of each income group in each of the firms is similar. For example, the share of expenditures in Nestle products by the highest and lowest income quartiles are 8.36% and 8.29% respectively.[39] Importantly, this pattern is very similar for the rest of the approximately 450 firms that we are able to identify. The data thus shows that most consumers buy from the same set of firms and that, across all income groups, the fraction of purchases on each of these firms is very similar. Therefore, at least in this sector, there is no clear match between specific firms and consumers of different income groups.

3.7 Implications

Our results have important implications for the measurement of inflation. The presence of product substitution indicates that the substitution bias that arises from using a fixed basket to compute price indices when people substitute away from expensive items, is higher during recessions and is mainly driven by substitution within product categories. This bias suggests that statistical agencies should attempt to track effective rather than posted prices particularly during bad times. In particular, an important effort needs to be made to study the dynamics of prices at a finer level of product detail. Although we acknowledge that this study might impose a large burden on statistical agencies, alternative methods to do so, such as the one proposed by Chevalier and Kashyap (2014), could be a good starting point. In addition, our study shows that given the large differences in the cost of living's inflation across income groups, these agencies should explore alternatives to produce income-specific price indices. These indices could be particularly relevant to understand how these differentials evolve over time; a very likely outcome given the differences in the basket of goods consumed by different income groups, their different shopping behavior patterns, and the fact that the relative price distribution is not stationary.

Our results also affect the interpretation of the studies that estimate both consumption and income inequality. In recent years, a great deal of the literature has compared the trend of consump-

39. The one exception is Wal-Mart Stores Inc. that sells mostly generic items. In this case, the share of expenditures by the higher income groups is lower than that of the lower income groups. This is mainly due to the fact that the share of expenditures of higher income households in Wal-Mart is significantly lower than that of the other groups.

tion inequality to that of income inequality.[40] During the Great Recession, unemployment rose quickly, asset prices fell, and many households suffered large changes in wealth. Thus, accurately measuring the patterns of income and consumption before and after the recession is particularly important.

Many studies associate the Great Recession with a decline in consumption inequality. Attanasio and Pistaferri (2014) find this result when using data from the Panel Study of Income and Dynamics (PSID). Meyer and Sullivan (2013) use data from the Current Population Survey and the Consumer Expenditure Survey to show that, between 2008 and 2011, income inequality rose sharply while consumption inequality fell drastically. Using the Nielsen Homescan data, we conduct a similar exercise and compute the total annual expenditures by income quartile. We then compute the ratio of the highest quartile to the lowest quartile (i.e., 75/25) to examine its trend over time.

Figure 3.8 shows that the ratio of expenditures in our data (dotted line in the figure) follows a similar declining trajectory as the consumption inequality described by Attanasio and Pistaferri (2014) and Meyer and Sullivan (2013). The figure also shows that, when we adjust the total expenditures of each quartile using our income-specific price indices, the trajectory of the ratio flattens. Although we calculate our price indices by using only the information on the goods available in the retail sector, this example illustrates that the patterns of consumption and income inequality could drastically change if one accounts for the differences in the cost of living's inflation. In particular, given that both the real income and the capacity to spend of low-income households eroded more drastically, official measures could be understating the increases in inequality during the Great Recession.

The welfare implications of our findings are much harder to quantify. This is because the quality of the products purchased cannot remain constant as consumers switch products within product categories. Our price indices summarize the utility implications for each type of consumer only if consumers experience low disutility by substituting between alternative brands within a product category. According to Chevalier and Kashyap (2014) this is the case for many consumers and

40. The literature is still divided between studies that have found that the rise in both have been similar (Aguiar and Bils, 2015; Attanasio, Hurst, and Pistaferri, 2015), while other studies have argued that consumption inequality has risen less than income inequality in recent decades (Krueger and Perri, 2006; Heathcote, Perri, and Violante, 2010; Fisher, Johnson, and Smeeding, 2014; Meyer and Sullivan, 2013).

only those with the least disutility actually switch. Nonetheless, this assumption is a limitation of our analysis.[41]

Finally, the mis-measurement in the cost of living has important implications for determining the eligibility and allocation of public assistance programs. This is because for a substantial amount of programs, the government adjusts the benefit levels and eligibility guidelines annually for inflation (e.g., food stamps). Our study indicates that in order to improve the allocation of public resources, particularly during downturns, the price adjustment for each group should use their own market basket, rather than using common price indices for everyone. More importantly, given that inflation measures are used to adjust poverty thresholds, poverty rates and, in turn, government funds to poor areas could be affected if there are large discrepancies between the CPI and a price index that correctly measures the experience of poor consumers.

3.8 Conclusion

We find substantial differences in the cost of living's inflation across income groups during the Great Recession. We argue that these differences are largely due to the ability of high-income households to adjust their shopping behavior to mitigate negative income shocks. Nonetheless, the gap in the inflation persists beyond the Great Recession. Further research is necessary to find the extent of the cyclicality of this gap. In addition, our study also points toward the need for statistical agencies to produce income-specific price indices. While the production of these indices is likely to be costly, an accurate measurement of these indices will have, undoubtedly, important implications for the measurement of other real variables.

41. This issue is present in other studies examining the evolution of the cost of living of consumers adjusting their shopping behavior. For example, as consumers switch retailers to reduce the prices they pay for the same product as in Coibion, Gorodnichenko, and Hong (2015b), store switching results in clear welfare increases only if the store where the products are purchased does not enter utility.

Tables and Figures

Table 3.1: Estimated Elasticities of Substitution by Income Group

Percentile	<25k	25k-50k	50k-100k	>100k
5	4.7	4.9	4.7	4.9
25	9.9	9.1	8.8	9.1
Median	19.0	18.7	16.7	16.5
75	48.5	45.3	46.6	44.8
95	233.1	221.6	196.4	196.4

Note: This table shows the distribution of the within-category elasticity of substitution by income group for the balanced product category. The estimation procedure follows Feenstra (1994) and Broda and Weinstein (2010)

Table 3.2: Average Annual Inflation Rates Before and After the Recession

Pre-2007					
	<25k	25k-50k	50k-100k	>100k	Difference
Benchmark EPI	0.67	0.40	0.46	0.69	-0.02
Divisia	0.77	0.44	0.39	0.70	0.07
Post-2007					
	<25k	25k-50k	50k-100k	>100k	Difference
Benchmark EPI	-0.10	-0.54	-0.82	-0.73	0.62
Divisia	0.62	-0.08	-0.67	-0.69	1.32

Note: This table shows the income-specific average annual inflation rates for the pre-2007 and post-2007 periods. It presents two different specifications: (i) benchmark exact price index, (ii) Divisia price index. The last column indicates the difference between the average annual inflation of the lowest income group (<25k) and the highest income group (>100k).

Table 3.3: Average Annual Inflation Rates by Income Group 2004-2010

	<25k	25k-50k	50k-100k	>100k	Difference
Benchmark EPI	1.39	0.97	0.75	0.80	0.59
No Bias Correction	1.66	1.35	1.14	1.18	0.48
Common Price	1.40	1.02	0.83	0.76	0.64
Fixed Category Weights	1.50	1.11	0.90	0.98	0.52

Note: This table shows the income-specific average annual inflation rates for the period 2004-20010 for four different specifications: (i) benchmark exact price index, (ii) no bias correction term, (iii) common (posted) prices, and (iv) fixed expenditure weights across product modules. The last column indicates the difference between the average annual inflation of the lowest income group (<25k) and the highest income group (>100k).

Table 3.4: Shopping Behavior as a Function of Local Unemployment Rate

Dependent variable:	Store Quality	Product Quality	Trips	Sales
Unemployment Rate	-0.0908***	-0.0670***	0.0923*	0.0751*
	(0.0182)	(0.0204)	(0.0511)	(0.0444)
Time Fixed Effect		✓		
Household Fixed Effect		✓		
# of Observations		208,769		

Note: This table reports the estimates of specification (3.10). The coefficients for store quality and product quality are the percentage change in quality as represented by the median price when the local unemployment rate increases by one percentage point. The coefficient of the variable trips (sales) is the number of additional shopping trips (additional fraction of items purchased on sale) after a 1% increase in the local unemployment rate. The correction of the standard errors follows Driscoll and Kraay (1998).

Table 3.5: Heterogeneous Shopping Behavior as a Function of Local Unemployment Rate

Dependent var:	Store Quality	Product Quality	Trips	Sales
UR_{ct}	-0.0749***	0.0124	-0.0116	-0.0286
	(0.0167)	(0.0256)	(0.0483)	(0.0388)
$\text{UR}_{ct} \times D_h$	-0.0321***	-0.1602***	0.2096***	0.2091***
	(0.0041)	(0.0360)	(0.0287)	(0.0332)
Time FE		✓		
Household FE		✓		
# of Observations		208,769		

Note: This table reports the estimates from specification (3.11). The coefficients for store quality and product quality are the percentage change in quality as represented by the median price when the local unemployment rate increases by one percentage point. The coefficient of the variable trips (sales) is the number of additional shopping trips (the additional fraction of items purchased on sale) after a 1% increase in the local unemployment rate. The correction of the standard errors follows Driscoll and Kraay (1998).

Table 3.6: Posted Inflation and Local Demand Shocks by Product Type

	Equal (1)	Market (2)	Common (3)
UR	-0.038***	-0.061***	-0.069***
	(0.007)	(0.010)	(0.010)
UR×High Priced	0.047***	0.108***	0.122***
	(0.005)	(0.007)	(0.007)
Market ×Category×Type FE		✓	
Month FE		✓	

Notes: We use market×category×quality-type fixed effects and month fixed effects. The number of observations is 10,909,370. Driscoll and Kraay (1998) standard errors are in parenthesis. The ***, **, and * denote significance at 0.01, 0.05, and 0.10 levels respectively.

Figure 3.1: Exact Price Index by Income Group

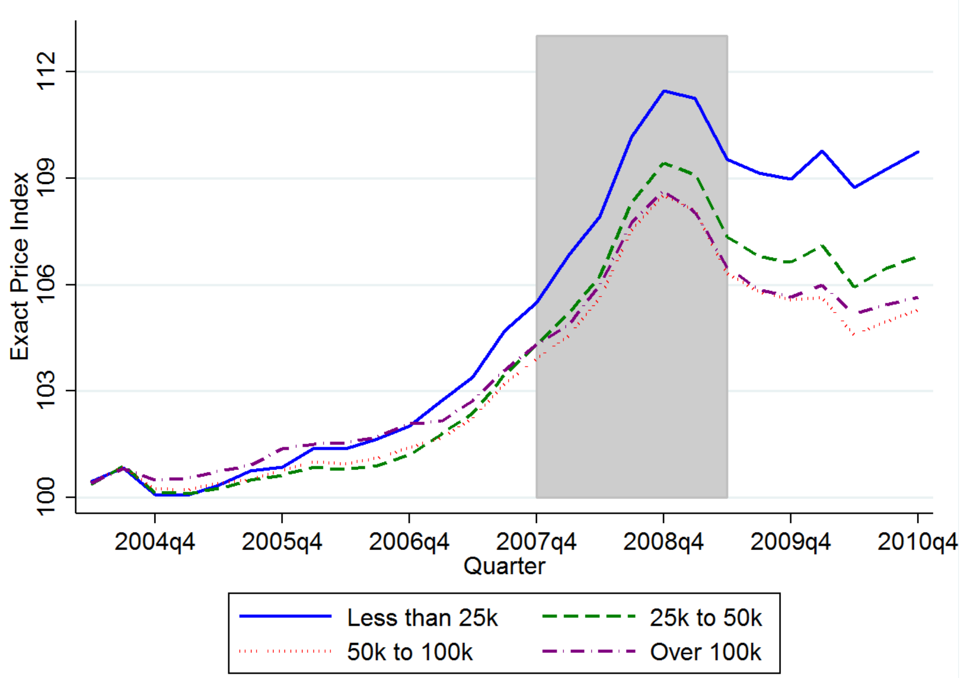

Note: This graph plots the exact price index for each income group from the first quarter of 2004 to the last quarter of 2010. The shaded areas indicate periods designated as recessions by the NBER.

Figure 3.2: Price Indices by Income Group after Controlling for Each Component

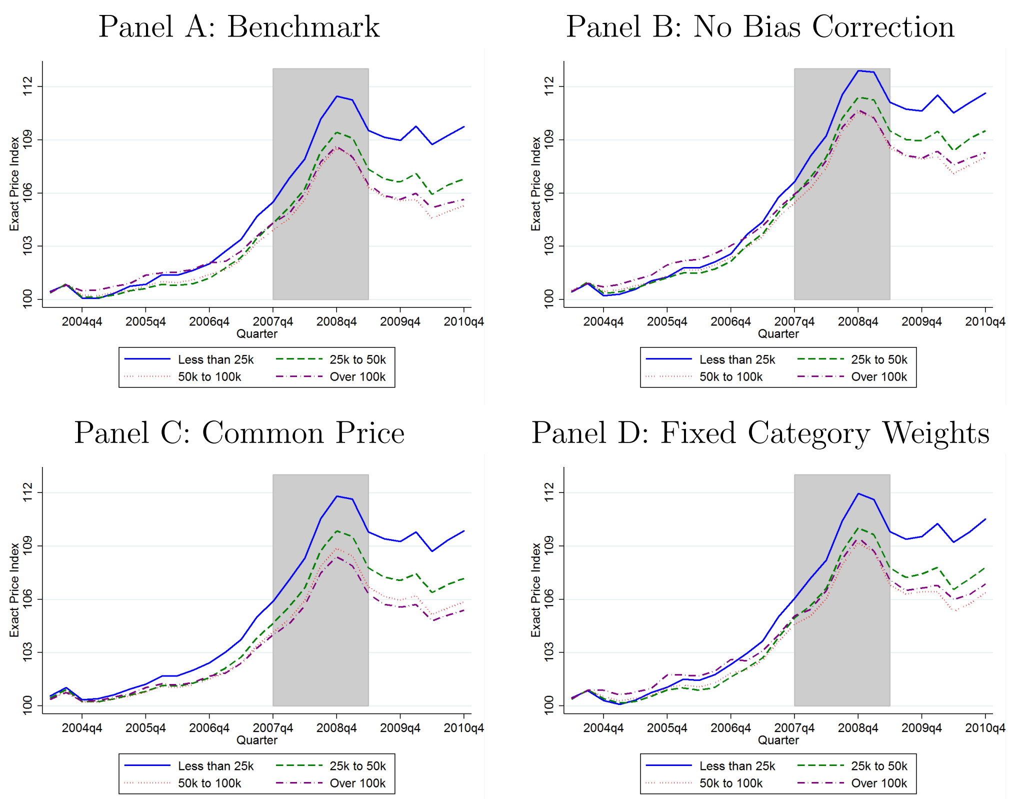

Note: This graph plots the exact price index (EPI) for each income group from the first quarter of 2004 to the last quarter of 2010. The shaded areas indicate periods designated as recessions by the NBER. The y-axis indicates the price index by income group after controlling for each component. Panel A of Figure 3.1 shows the benchmark EPI. Panel B reports the results for the conventional exact price index (CEPI) in which the bias correction term is dropped. Panel C depicts the results after using common prices across income groups instead of different effective prices. Panel D reports the results when expenditures at the product group level cannot be reallocated.

Figure 3.3: Difference between the Price Index of the Highest vs. the Lowest Quartile after Controlling for their Shopping Behavior

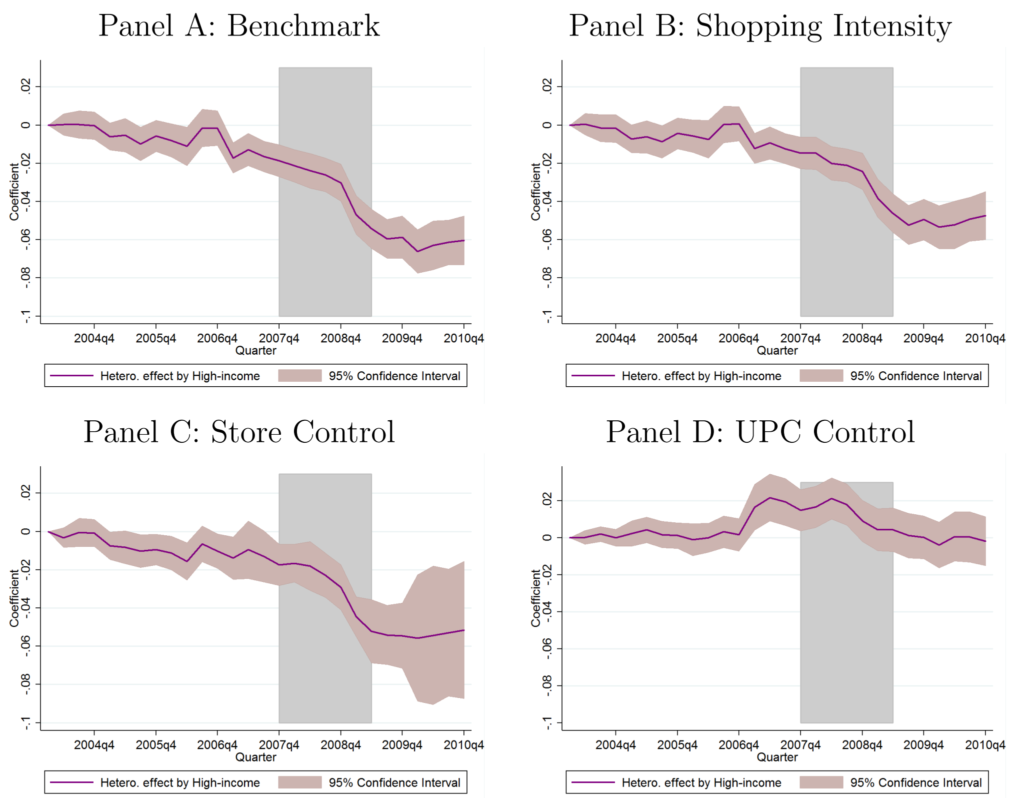

Note: This graph plots the coefficients of the interaction term in equation (6). The y-axis indicates the mean percent deviations of the prices paid by the highest quartile with respect to the lowest quartile of the income distribution. Panel A is the benchmark specification in equation (6). Panel B reports the results after controlling for shopping intensity: the fraction of goods bought on sale, the fraction of goods bought with coupons, the share of expenditures on generics, the number of trips in a given quarter, and the mean number of units purchased in a given quarter and product category. Panel C depicts the results after introducing the store fixed effects to the benchmark. Panel D reports the results obtained after introducing the UPC fixed effects. The standard errors are clustered at the product category level. The pink area indicates the 95% confidence interval. The shaded areas indicate the periods designated as recessions by the NBER. The number of observations is 103,940,344 and the R-squareds are 0.81 (Panel A), 0.82 (Panel B), 0.86 (Panel C), and 0.97 (Panel D) respectively.

Figure 3.4: Time Fixed Effects - Price Index of the Lowest Quartile

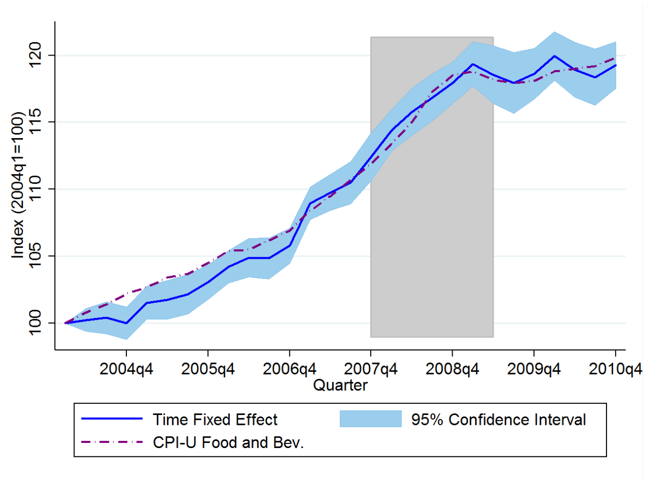

fixed effects.png

Note: This graph plots the coefficients of the time fixed effects in equation (6). The units of the y-axis are the mean percent deviations from the price of the base period (2004q1) for the lowest quartile of the income distribution. The standard errors are clustered at the product category level. The blue area indicates the 95% confidence interval. The shaded areas indicate periods designated as recessions by the NBER.

Figure 3.5: Relative Price Paid by Household Income

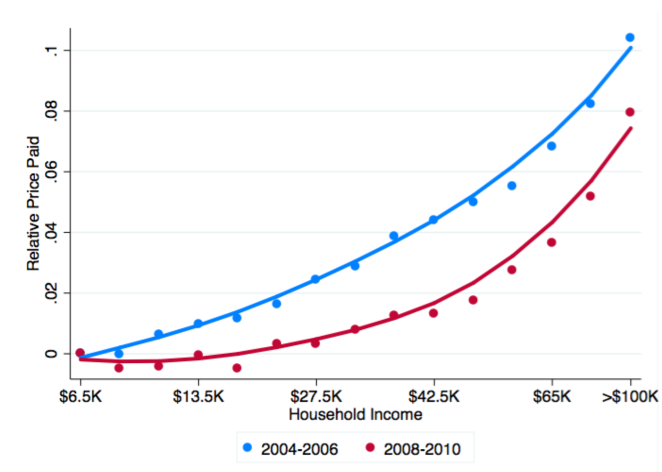

Note: The relative prices are measured in a regression of the log unit price paid against income category dummies and product category, region, chain, quarter and household fixed effects. Each dot represents how much more each income category pays per unit for products with respect to households earning between $5,000 and $8,000. The blue dots represent the cross-sectional relation between 2004 and 2006. The red dots cover from 2008 to 2010.

Figure 3.6: Product Quality Distribution for 2006-2007

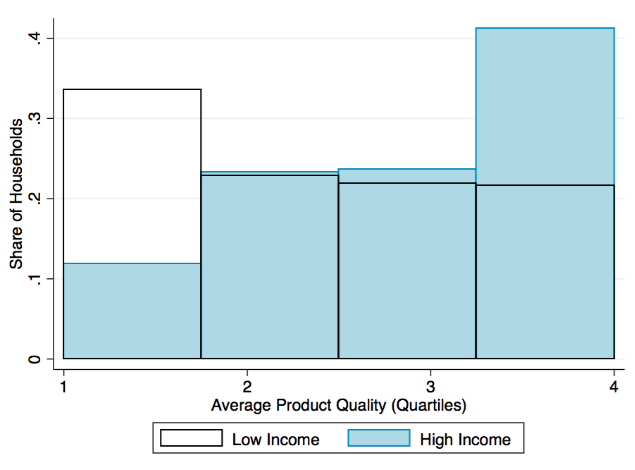

Note: This graph shows the distribution of the average relative prices paid by the highest and the lowest quartiles of the income distribution for 2006 to 2007. The average product quality is divided into quartiles. The height of the bar represents the fraction of households belonging to a given income group in each of the quartiles of the quality distribution. The source is the Nielsen Homescan.

Figure 3.7: Product Quality Distribution Before and After the Recession

Panel A: Highest Income Quartile

Panel B: Lowest Income Quartile

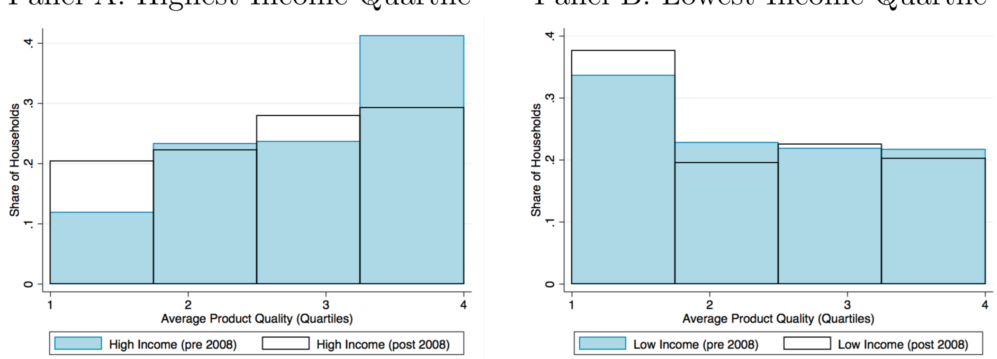

Note: This graph shows the distribution of the average relative prices paid by the highest and the lowest quartiles of the income distribution before (green bars) and after (white bars) the Great Recession. The average product quality is divided into quartiles. The height of the bar represents the fraction of households belonging to a given income group in each of the quartiles of the quality distribution. The source is the Nielsen Homescan.

Figure 3.8: Highest to Lowest Quartile Expenditure Ratio

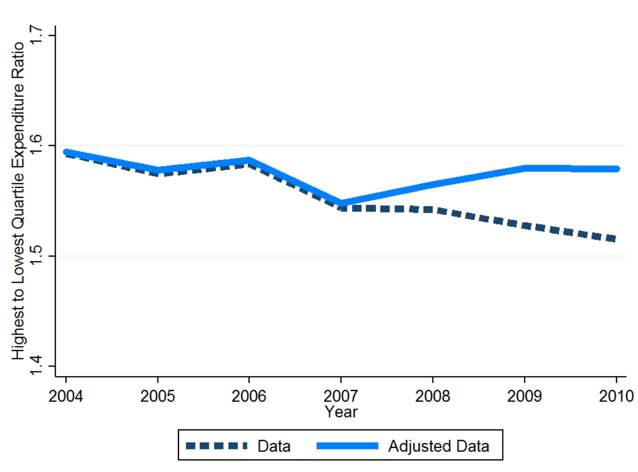

Note: This graph shows the change in the expenditure ratio of the highest to the lowest quartile over time. The dotted line is the calculated expenditure ratio data from Nielsen and the solid line is the ratio adjusted by income-specific price indices.

REFERENCES

Acemoglu, Daron, Ufuk Akcigit, Nicholas Bloom, and William R Kerr. 2013. "Innovation, Reallocation and Growth." *NBER Working Paper No. 18993* .

Acemoglu, Daron and Dan Cao. 2015. "Innovation by entrants and incumbents." *Journal of Economic Theory* 157:255–294.

Aghion, Philippe, Ufuk Akcigit, and Peter Howitt. 2014. "What Do We Learn from Schumpeterian Growth Theory?" In *Handbook of Economic Growth*, vol. 2, edited by Philippe Aghion and Steven N Durlauf, chap. 1. Oxford: North Holland, 515–563.

Aghion, Philippe, Nicholas Bloom, Brian Lucking, Raffaella Sadun, and John Van Reenen. 2017. "Turbulence, Firm Decentralization and Growth in Bad Times." *NBER Working Paper No. 23354* .

Aghion, Philippe, Patrick Bolton, Christopher Harris, and Bruno Jullien. 1991. "Optimal learning by experimentation." *Review of Economic Studies* 58 (4):621–654.

Aghion, Philippe and Peter Howitt. 1992. "A Model of Growth through Creative Destruction." *Econometrica* 60 (2):323–351.

Aguiar, Mark and Mark Bils. 2015. "Has Consumption Inequality Mirrored Income Inequality?" *American Economic Review* 105 (9):2725–2756.

Aguiar, Mark and Erik Hurst. 2005. "Consumption versus Expenditure." *Journal of Political Economy* 113 (5):919–948.

———. 2007. "Life-Cycle Prices and Production." *American Economic Review* 97 (5):1533–1559.

———. 2013. "Deconstructing life cycle expenditure." *Journal of Political Economy* 121 (3):437–492.

Akcigit, Ufuk and William R Kerr. 2010. "Growth Through Heterogeneous Innovations." *NBER Working Paper No. 16443* .

Alvarez, Fernando, Hervé Le Bihan, and Francesco Lippi. 2014. "Small and large price changes and the propagation of monetary shocks." *NBER Working Paper No. 20155* .

Alvarez, Fernando, Katarına Borovicková, and Robert Shimer. 2015. "The proportional hazard model: Estimation and testing using price change and labor market data." *University of Chicago Mimeo* .

Alvarez, Fernando and Francesco Lippi. 2014. "Price setting with menu cost for multiproduct firms." *Econometrica* 82 (1):89–135.

Alvarez, Fernando, Francesco Lippi, and Juan Passadore. 2017. "Are State-and Time-Dependent Models Really Different?" *NBER Macroeconomics Annual* 31 (1):379–457.

Alvarez, Fernando E, Francesco Lippi, and Luigi Paciello. 2011. "Optimal Price Setting With Observation and Menu Costs*." *The Quarterly journal of economics* 126 (4):1909–1960.

Anderson, Eric, Benjamin A Malin, Emi Nakamura, Duncan Simester, and Jón Steinsson. 2013. "Informational rigidities and the stickiness of temporary sales." *NBER Working Paper No. 19350* .

Anderson, Simon P, Andre De Palma, and Jacques François Thisse. 1992. *Discrete choice theory of product differentiation.* MIT press.

Argente, David and Munseob Lee. 2016. "Cost of Living Inequality during the Great Recession." .

Argente, David, Munseob Lee, and Sara Moreira. 2017a. "Innovation and product reallocation in the great recession." *Journal of Monetary Economics* .

———. 2017b. "Innovation and Product Reallocation in the Great Recession." .

Argente, David and Chen Yeh. 2017. "Product's Life Cycle, Learning, and Nominal Shocks." .

Atkeson, Andy and Ariel Burstein. 2017. "The Aggregate Implications of Innovative Investment in the Garcia-Macia, Hsieh, and Klenow Model." .

Attanasio, Orazio, Erik Hurst, and Luigi Pistaferri. 2015. *The Evolution of Income, Consumption, and Leisure Inequality in the US, 1980-2010.* University of Chicago Press.

Attanasio, Orazio and Luigi Pistaferri. 2014. "Consumption inequality over the last half century: some evidence using the new PSID consumption measure." *American Economic Review* 104 (5):122–126.

Bachmann, Rüdiger and Giuseppe Moscarini. 2012. "Business Cycles and Endogenous Uncertainty." .

Baley, Isaac and Julio A Blanco. 2017. "Firm Uncertainty Cycles and the Propagation of Nominal Shocks." .

Barro, Robert J. 1972. "A theory of monopolistic price adjustment." *Review of Economic Studies* :17–26.

Bartelsman, Eric J and Mark Doms. 2000. "Understanding productivity: Lessons from longitudinal microdata." *Journal of Economic Literature* 38 (3):569–594.

Beraja, Martin, Erik Hurst, and Juan Ospina. 2016. "The Aggregate Implications of Regional Business Cycles." *NBER Working Paper No. 21956* .

Berger, David and Joseph Vavra. 2015. "Dynamics of the US price distribution." .

Bernard, Andrew B and Toshihiro Okubo. 2016. "Product switching and the business cycle." *NBER Working Paper No. 22649* .

Bernard, Andrew B, Stephen J Redding, and Peter K Schott. 2010. "Multiple-product firms and product switching." *American Economic Review* 100 (1):70–97.

Bilbiie, Florin O, Fabio Ghironi, and Marc J Melitz. 2012. "Endogenous Entry, Product Variety, and Business Cycles." *Journal of Political Economy* 120 (2):304–345.

Bils, Mark and Peter J Klenow. 2001. "Quantifying quality growth." *American Economic Review* :1006–1030.

Boskin, Michael J. 1996. *Toward a More Accurate Measure of the Cost of Living: Final Report to the Senate Finance Committee*. Advisory Commission to Study the Consumer Price Index.

Boskin, Michael J, Ellen R Dulberger, Robert J Gordon, Zvi Griliches, and Dale W Jorgenson. 1998. "Consumer prices, the consumer price index, and the cost of living." *Journal of Economic Perspectives* 12 (1):3–26.

Broda, Christian, Ephraim Leibtag, and David E Weinstein. 2009. "The Role of Prices in Measuring the Poor's Living Standards." *Journal of Economic Perspectives* 23 (2):77–97.

Broda, Christian and John Romalis. 2009. "The Welfare Implications of Rising Price Dispersion." .

Broda, Christian and David E Weinstein. 2006. "Globalization and the Gains From Variety." *Quarterly Journal of Economics* 121 (2):541–585.

———. 2010. "Product Creation and Destruction: Evidence and Price Implications." *American Economic Review* :691–723.

Bronnenberg, Bart J, Michael W Kruger, and Carl F Mela. 2008. "Database paper-The IRI marketing data set." *Marketing Science* 27 (4):745–748.

Burstein, Ariel, Martin Eichenbaum, and Sergio Rebelo. 2005. "Large Devaluations and the Real Exchange Rate." *Journal of Political Economy* 113 (4):742–784.

Caballero, Ricardo J and Eduardo MRA Engel. 2007. "Price stickiness in Ss models: New interpretations of old results." *Journal of Monetary Economics* 54:100–121.

Caballero, Ricardo J and Mohamad L Hammour. 1994. "The Cleansing Effect of Recessions." *American Economic Review* 84 (5):1350–68.

Calvo, Guillermo A. 1983. "Staggered prices in a utility-maximizing framework." *Journal of Monetary Economics* 12 (3):383–398.

Campbell, Jeffrey R and Benjamin Eden. 2005. "Rigid prices: evidence from US scanner data." .

Caplin, Andrew S and Daniel F Spulber. 1987. "Menu Costs and the Neutrality of Money." *The Quarterly Journal of Economics* 102 (4):703–725.

Carvalho, Carlos and Felipe Schwartzman. 2015. "Selection and monetary non-neutrality in time-dependent pricing models." *Journal of Monetary Economics* 76:141–156.

Caves, Richard E. 1998. "Industrial organization and new findings on the turnover and mobility of firms." *Journal of Economic Literature* 36 (4):1947–1982.

Chatterjee, Satyajit and Russell Cooper. 1993. "Entry and Exit, Product Variety and the Business Cycle." *NBER Working Paper No. 4562* .

Chevalier, Judith A and Anil K Kashyap. 2014. "Best Prices: Price Discrimination and Consumer Substitution." .

Chevalier, Judith A, Anil K Kashyap, and Peter E Rossi. 2003. "Why Don't Prices Rise During Periods of Peak Demand? Evidence from Scanner Data." *American Economic Review* 93 (1):15–37.

Coibion, Olivier, Yuriy Gorodnichenko, and Gee Hee Hong. 2015a. "The cyclicality of sales, regular and effective prices: Business cycle and policy implications." *American Economic Review* .

———. 2015b. "The Cyclicality of Sales, Regular and Effective Prices: Business Cycle and Policy Implications." *American Economic Review* 105 (3):993–1029.

Corrado, Carol A and Charles R Hulten. 2010. "How Do You Measure a" Technological Revolution"?" *The American Economic Review* 100 (2):99–104.

Davis, Steven J and John Haltiwanger. 1992. "Gross Job Creation, Gross Job Destruction, and Employment Reallocation." *Quarterly Journal of Economics* 107 (3):819–863.

———. 2014. "Labor Market Fluidity and Economic Performance." *NBER Working Paper No. 20479* .

Davis, Steven J, John C Haltiwanger, and Scott Schuh. 1996. "Job creation and destruction." *MIT Press Books* .

De Nardi, Mariacristina, Eric French, and John B Jones. 2010. "Why Do the Elderly Save? The Role of Medical Expenses." *Journal of Political Economy* 118 (1):39–75.

Deaton, Angus and John Muellbauer. 1980. "An Almost Ideal Demand System." *American Economic Review* 70 (3):312–326.

Deaton, Angus and Christina Paxson. 1994. "Intertemporal Choice and Inequality." *Journal of Political Economy* 102 (3):437–467.

Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda. 2014. "The Role of Entrepreneurship in U.S. Job Creation and Economic Dynamism." *Journal of Economic Perspectives* 28 (3):3–24.

Diewert, W Erwin. 1976. "Exact and Superlative Index Numbers." *Journal of Econometrics* 4 (2):115–145.

Dixit, Avinash. 1991. "Analytical approximations in models of hysteresis." *Review of Economic Studies* 58 (1):141–151.

Doraszelski, Ulrich, Gregory Lewis, and Ariel Pakes. 2016. "Just starting out: Learning and equilibrium in a new market." .

Driscoll, John C and Aart C Kraay. 1998. "Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data." *Review of Economics and Statistics* 80 (4):549–560.

Dubé, Jean-Pierre, Günter J Hitsch, and Peter E Rossi. 2015. "Income and Wealth Effects on Private-Label Demand: Evidence From the Great Recession." *NBER Working Paper No. 21446* .

Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo. 2011. "Reference prices, costs, and nominal rigidities." *American Economic Review* 101 (1):234–262.

Einav, Liran, Theresa Kuchler, Jonathan Levin, and Neel Sundaresan. 2015. "Assessing sale strategies in online markets using matched listings." *American Economic Journal: Microeconomics* 7 (2):215–247.

Einav, Liran, Ephraim Leibtag, and Aviv Nevo. 2010. "Recording discrepancies in Nielsen Homescan data: Are they present and do they matter?" *Quantitative Marketing and Economics* 8 (2):207–239.

Faber, Benjamin. 2014. "Trade Liberalization, the Price of Quality, and Inequality: Evidence from Mexican Store Prices." .

Faber, Benjamin and Thibault Fally. 2017. "Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data." *NBER Working Paper No. 23101* .

Feenstra, Robert C. 1994. "New Product Varieties and the Measurement of International Prices." *American Economic Review* 84 (1):157–177.

Feenstra, Robert C and Marshall B Reinsdorf. 2000. "An Exact Price Index for the Almost Ideal Demand System." *Economics Letters* 66 (2):159–162.

Fernald, John G. 2014. "A quarterly, utilization-adjusted series on total factor productivity." Federal Reserve Bank of San Francisco.

Fisher, Jonathan, David S Johnson, and Timothy M Smeeding. 2014. "Inequality of Income and Consumption in the U.S.: Measuring the Trends in Inequality from 1984 to 2011 for the Same Individuals." *Review of Income and Wealth* 61 (4):630–650.

Foster, Lucia, Cheryl Grim, and John Haltiwanger. 2016. "Reallocation in the Great Recession: Cleansing or Not?" *Journal of Labor Economics* 34 (S1):S293–S331.

Foster, Lucia, John Haltiwanger, and Cornell J Krizan. 2001. "Aggregate Productivity Growth. Lessons from Microeconomic Evidence." In *New Developments in Productivity Analysis*. University of Chicago Press, 303–372.

———. 2006. "Market Selection, Reallocation, and Restructuring in the U.S. Retail Trade Sector in the 1990s." *Review of Economics and Statistics* 88 (4):748–758.

Foster, Lucia, John Haltiwanger, and Chad Syverson. 2008. "Reallocation, firm turnover, and efficiency: Selection on productivity or profitability?" *American Economic Review* 98 (1):394–425.

———. 2016. "The slow growth of new plants: Learning about demand?" *Economica* 83 (329):91–129.

Gagnon, Etienne and David Lopez-Salido. 2014. "Small price responses to large demand shocks." .

Gagnon, Etienne, David Lopez-Salido, and Jason A Sockin. 2015. "The Cyclicality of Sales, Regular and Effective Prices: Comment." *FEDS Working Paper No. 2015-052* .

Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J Klenow. 2016. "How Destructive is Innovation?" *NBER Working Paper No. 22953* .

Garner, Thesia I, David S Johnson, and Mary F Kokoski. 1996. "An Experimental Consumer Price Index for the Poor." *Monthly Labor Review* 119:32.

Gaur, Vishal and Marshall L Fisher. 2005. "In-store experiments to determine the impact of price on sales." *Production and Operations Management* 14 (4):377–387.

Gilchrist, Simon, Raphael Schoenle, Jae Sim, and Egon Zakrajšek. 2017. "Inflation dynamics during the financial crisis." *The American Economic Review* 107 (3):785–823.

Golosov, Mikhail and Robert E Lucas. 2007. "Menu Costs and Phillips Curves." *Journal of Political Economy* 115 (2).

Gorodnichenko, Yuriy and Michael Weber. 2016. "Are sticky prices costly? Evidence from the stock market." *American Economic Review* 106 (1):165–199.

Gourinchas, Pierre-Olivier and Jonathan A Parker. 2002. "Consumption over the life cycle." *Econometrica* 70 (1):47–89.

Gourio, Francois and Leena Rudanko. 2014. "Customer capital." *Review of Economic Studies* 81 (3):1102–1136.

Gourio, Francois and Michael Siemer. 2014. "A Missing Generation of Firms? Aggregate Effects of the Decline in New Business Formation." .

Griffith, Rachel, Ephraim Leibtag, Andrew Leicester, and Aviv Nevo. 2009. "Consumer Shopping Behavior: How Much Do Consumers Save?" *Journal of Economic Perspectives* 23 (2):99–120.

Griliches, Zvi and Haim Regev. 1995. "Firm Productivity in Israeli Industry 1979–1988." *Journal of Econometrics* 65 (1):175–203.

Grossman, Gene M and Elhanan Helpman. 1991. "Quality ladders in the theory of growth." *The Review of Economic Studies* 58 (1):43–61.

Hagemann, Robert P. 1982. "The Variability of Inflation Rates across Household Types." *Journal of Money, Credit and Banking* 14 (4):494–510.

Hallak, Juan Carlos and Peter K Schott. 2011. "Estimating Cross-Country Differences in Product Quality." *Quarterly Journal of Economics* 126 (1):417–474.

Handbury, Jessie. 2013. "Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across US Cities." .

Handbury, Jessie, Tsutomu Watanabe, and David E Weinstein. 2013. "How Much Do Official Price Indexes Tell Us About Inflation?" *NBER Working Paper No. 19504* .

Handbury, Jessie and David E Weinstein. 2014. "Goods Prices and Availability in Cities." *Review of Economic Studies* 82:258–296.

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L Violante. 2010. "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967–2006." *Review of Economic Dynamics* 13 (1):15–51.

Hobijn, Bart and David Lagakos. 2005. "Inflation Inequality in the United States." *Review of Income and Wealth* 51 (4):581–606.

Hummels, David and Peter J Klenow. 2005. "The Variety and Quality of a Nation's Exports." *American Economic Review* 95 (3):704–723.

Ilut, Cosmin, Rosen Valchev, and Nicolas Vincent. 2015. "Paralyzed by Fear: Rigid and Discrete Pricing under Demand Uncertainty." .

İmrohoroğlu, Ayşe and Şelale Tüzel. 2014. "Firm-level productivity, risk, and return." *Management Science* 60 (8):2073–2090.

Jaimovich, Nir and Max Floetotto. 2008. "Firm Dynamics, Markup Variations, and the business cycle." *Journal of Monetary Economics* 55 (7):1238–1252.

Jaimovich, Nir, Sergio Rebelo, and Arlene Wong. 2015. "Trading Down and the Business Cycle." *NBER Working Paper No. 21539* .

Jaravel, Xavier. 2016. "The Unequal Gains from Product Innovations." .

Kaplan, Greg and Guido Menzio. 2015. "The morphology of price dispersion." *International Economic Review* 56 (4):1165–1206.

Kaplan, Greg and Sam Schulhofer-Wohl. 2016. "Inflation at the Household Level." *NBER Working Paper No. 22331* .

Kaplan, Steven N and Luigi Zingales. 1997. "Do investment-cash flow sensitivities provide useful measures of financing constraints?" *Quarterly Journal of Economics* 112 (1):169–215.

Karadi, Peter and Adam Reiff. 2016. "Menu Costs, Aggregate Fluctuations, and Large Shocks." .

Keller, Godfrey and Sven Rady. 1999. "Optimal experimentation in a changing environment." *Review of Economic Studies* 66 (3):475–507.

Kiefer, Nicholas M and Yaw Nyarko. 1989. "Optimal control of an unknown linear process with learning." *International Economic Review* 30 (3):571–86.

Kiley, Michael T. 2002. "Partial adjustment and staggered price setting." *Journal of Money, Credit, and Banking* 34 (2):283–298.

Klette, Tor Jakob and Samuel Kortum. 2004. "Innovating Firms and Aggregate Innovation." *Journal of Political Economy* 112 (5).

Krueger, Dirk and Fabrizio Perri. 2006. "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory." *Review of Economic Studies* 73 (1):163–193.

Kryvtsov, Oleksiy and Nicolas Vincent. 2014. "On the Importance of Sales for Aggregate Price Flexibility." .

Kugler, Maurice and Eric Verhoogen. 2012. "Prices, Plant Size, and Product Quality." *Review of Economic Studies* 79 (1):307–339.

Lee, Yoonsoo and T Mukoyama. 2015. "A Model of Entry, Exit, and Plant-level Dynamics Over the Business Cycle." .

Lentz, Rasmus and Dale T Mortensen. 2008. "An empirical model of growth through product innovation." *Econometrica* 76 (6):1317–1373.

Lucas, Robert E. 2003. "Macroeconomic Priorities." *American Economic Review* 93:114.

McKenzie, David and Ernesto Schargrodsky. 2011. "Buying Less, but Shopping More: The Use of Nonmarket Labor during a Crisis." *Economia* 11 (2):1–43.

Melitz, Marc J and Sašo Polanec. 2015. "Dynamic Olley-Pakes Productivity Decomposition with Entry and Exit." *The RAND Journal of Economics* 46 (2):362–375.

Meyer, Bruce D and James X Sullivan. 2013. "Consumption and Income Inequality and the Great Recession." *American Economic Review* 103 (3):178–83.

Michael, Robert T. 1979. "Variation Across Households in the Rate of Inflation." *Journal of Money, Credit and Banking* 11 (1):32–46.

Midrigan, Virgiliu. 2011. "Menu costs, multiproduct firms, and aggregate fluctuations." *Econometrica* 79 (4):1139–1180.

Minniti, Antonio and Francesco Turino. 2013. "Multi-product Firms and Business Cycle Dynamics." *European Economic Review* 57:75–97.

Mirman, Leonard J, Larry Samuelson, and Amparo Urbano. 1993. "Monopoly experimentation." *International Economic Review* 34 (3):549–563.

Moreira, Sara. 2015. "Firm dynamics, persistent effects of entry conditions, and business cycles." Tech. rep., mimeo.

———. 2016. "Firm Dynamics, Persistent Effects of Entry Conditions, and Business Cycles." .

Nakamura, Emi and Jón Steinsson. 2008. "Five facts about prices: A reevaluation of menu cost models." *Quarterly Journal of Economics* 123 (4):1415–1464.

Nakamura, Emi and Jon Steinsson. 2009. "Monetary Non-Neutrality in a Multi-Sector Menu Cost Model." *Quarterly Journal of Economics* 125 (3):961–1013.

Nakamura, Emi and Jón Steinsson. 2011. "Price setting in forward-looking customer markets." *Journal of Monetary Economics* 58 (3):220–233.

Nakamura, Emi, Jón Steinsson, and Miao Liu. 2014. "Are Chinese Growth and Inflation Too Smooth? Evidence from Engel Curves." *NBER Working Paper No. 19893* .

Nevo, Aviv and Arlene Wong. 2015. "The Elasticity of Substitution Between Time and Market Goods: Evidence from the Great Recession." *NBER Working Paper No. 21318* .

Olley, G. Steven and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64 (6):1263–1297.

Paciello, Luigi, Andrea Pozzi, and Nicholas Trachter. 2014. "Price dynamics with customer markets." .

Peters, Ryan H and Lucian A Taylor. 2017. "Intangible Capital and the Investment-q Relation." *Journal of Financial Economics* 123 (2):251–272.

Petev, Ivaylo, Luigi Pistaferri, and Itay Saporta Eksten. ???? "Consumption and the Great Recession: An analysis of trends, perceptions, and distributional effects."

Prescott, Edward C. 1972. "The multi-period control problem under uncertainty." *Econometrica* :1043–1058.

Rothschild, Michael. 1974. "A two-armed bandit theory of market pricing." *Journal of Economic Theory* 9 (2):185–202.

Sato, Kazuo. 1976. "The Ideal Log-Change Index Number." *Review of Economics and Statistics* 58 (2):223–228.

Schoenle, Raphael. 2017. "International menu costs and price dynamics." *Review of International Economics* 25 (3):578–606.

Shapiro, Matthew D and Mark W Watson. 1988. "Sources of business cycle fluctuations." *NBER Macroeconomics annual* 3:111–148.

Sheedy, Kevin D. 2010. "Intrinsic inflation persistence." *Journal of Monetary Economics* 57 (8):1049–1061.

Stroebel, Johannes and Joseph Vavra. 2014. "House prices, local demand, and retail prices." *NBER Working Paper No. 20710* .

———. 2016. "House Prices, Local Demand, and Retail Prices." *NBER Working Paper No. 20710* .

Vartia, Yrjö O. 1976. "Ideal Log-Change Index Numbers." *Scandinavian Journal of Statistics* 3 (3):121–126.

Vavra, Joseph. 2014. "Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation." *Quarterly Journal of Economics* 129 (1):215–258.

Wieland, Volker. 2000a. "Learning by doing and the value of optimal experimentation." *Journal of Economic Dynamics and Control* 24 (4):501–534.

———. 2000b. "Monetary policy, parameter uncertainty and optimal learning." *Journal of Monetary Economics* 46 (1):199–228.

Willems, Tim. 2016. "Actively learning by pricing: a model of an experimenting seller." *Economic Journal* .

Zbaracki, Mark J, Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen. 2004. "Managerial and customer costs of price adjustment: direct evidence from industrial markets." *Review of Economics and Statistics* 86 (2):514–533.

Zellner, Arnold. 1971. *An introduction to Bayesian inference in econometrics.*

APPENDIX A

PRODUCT LIFE CYCLE, LEARNING, AND NOMINAL SHOCKS

A.1 Tables and Figures

Figure A.1: Frequency of Price Increases and Decreases at Entry

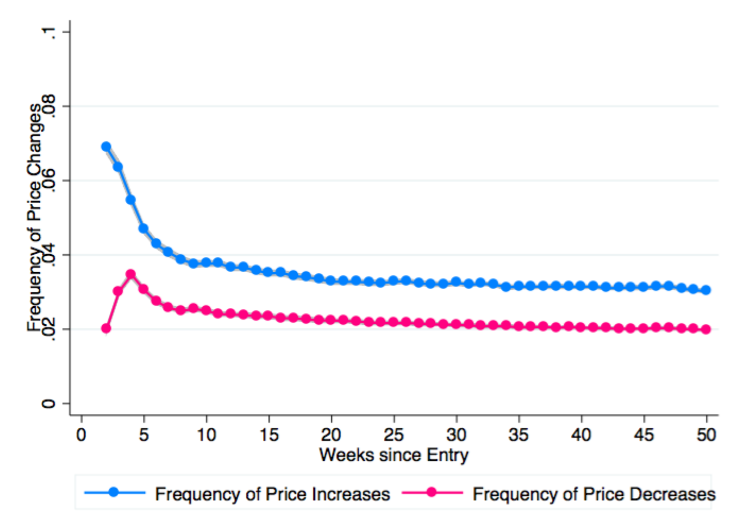

Note: The graph plots the average weekly frequency of price adjustments of products entering the market. The y-axis denotes the probability that the product adjusts prices in a given week, and the x-axis denotes the number of weeks the product has been observed in the data after it entered the market. The graph plots the age fixed effects where we use a regular price change indicator as the dependent variable that controls for the store, UPC, time fixed effects, and the local unemployment rate represents the cohort fixed effects. The blue line indicates the frequency of positive price adjustments and the red line the frequency of negative price adjustments. The calculation uses approximately 130 million observations and 2.5 million stores$\times$ UPC pairs. The data source is the Symphony IRI data set.

Figure A.2: Absolute Value of Price Increases and Decreases at Entry

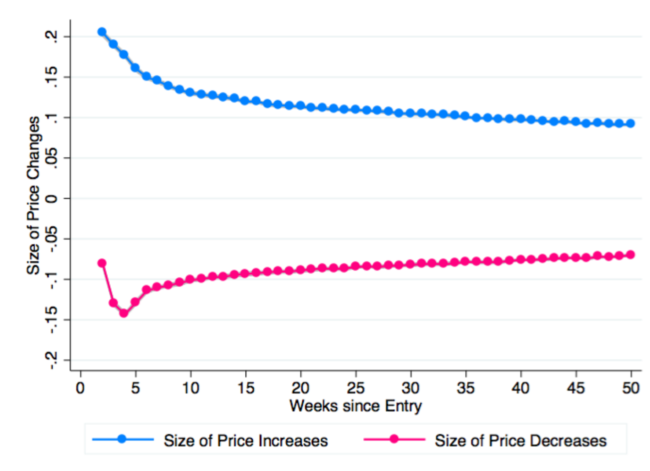

Note: The graph plots the average size of the price adjustments of products that enter the market. The y-axis is the value of the log price change in that week, and the x-axis denotes the number of weeks the product has been observed in the data after it entered the market. The graph plots the age fixed effects where we use the log price change as dependent variable that controls for the store, UPC, time fixed effects, and the local unemployment rate represents the cohort fixed effects. The blue line indicates the average size of positive price adjustments and the red line the average size of negative price adjustments. The calculation uses approximately 5.8 million price changes and 2.5 million stores× UPC pairs. The data source is the Symphony IRI data set.

Figure A.3: Distributions of Price Changes

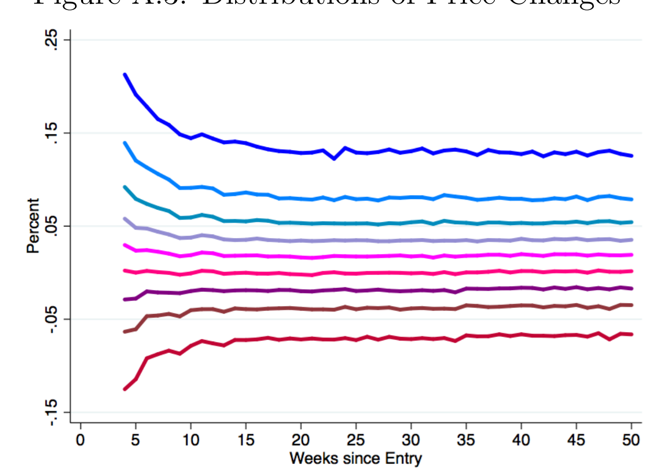

Note: The graph shows the distribution of regular price changes. The percentiles plotted are the 10th, 25th, 33rd, 50th, 66th, 75th, and 90th respectively. The y-axis is the value of the log price change in that week, and the x-axis denotes the number of weeks the product has been observed in the data after it entered the market. The calculation uses approximately 5.8 million price changes and 2.5 million stores× UPC pairs. The data source is the Symphony IRI data set.

Figure A.4: Frequency of All Price Changes

Panel A: IRI Symphony Panel B: Nielsen RMS

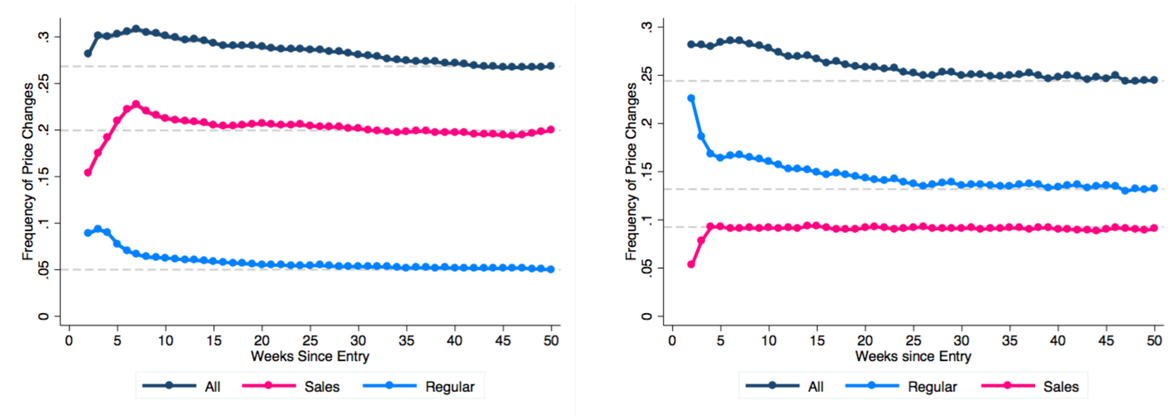

Note: The figure shows the average weekly frequency of all price changes as a function of the number of weeks after a product enters. Panel A shows the frequency of regular prices changes, the frequency of sales, and the frequency of all price changes in the IRI Symphony data. Panel B shows the same 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 fixed effects' coefficients for equation 1.1 where we use the price change indicator as the dependent variable. Equation 1.1 is computed by controlling for the store, UPC and time fixed effects, and the local unemployment rate represents the cohort fixed effects.

Figure A.5: Fraction of Price Changes Larger than Two Std. (Positive and Negative)

Panel A: Positive Price Changes Panel B: Negative Price Changes

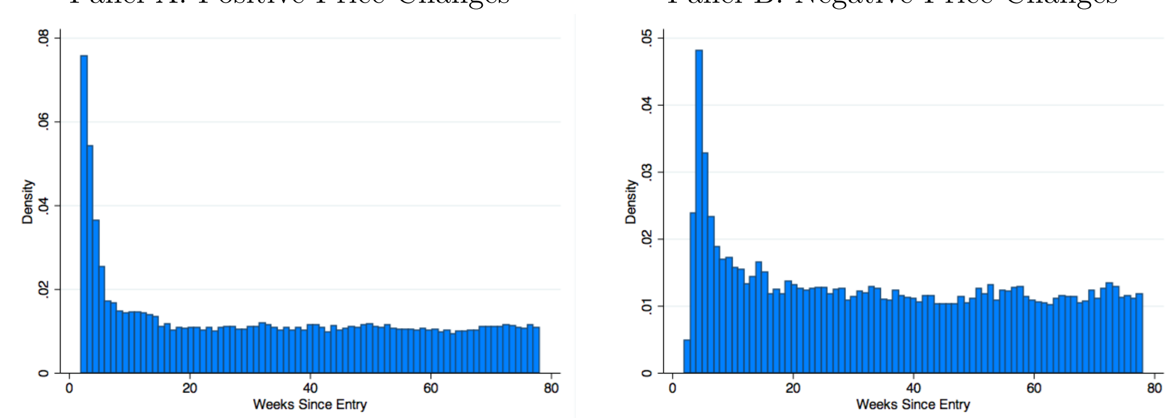

Note: The figure shows the fraction of price changes larger than two standard deviations from the mean in a given category and store as a function of the age of the product. Panel A shows the distribution of large price increases and Panel B the distribution of price decreases. The products considered are those that last at least two years in the market. Source: IRI Symphony dataset.

Figure A.6: Fraction of Price Changes Larger than 30%

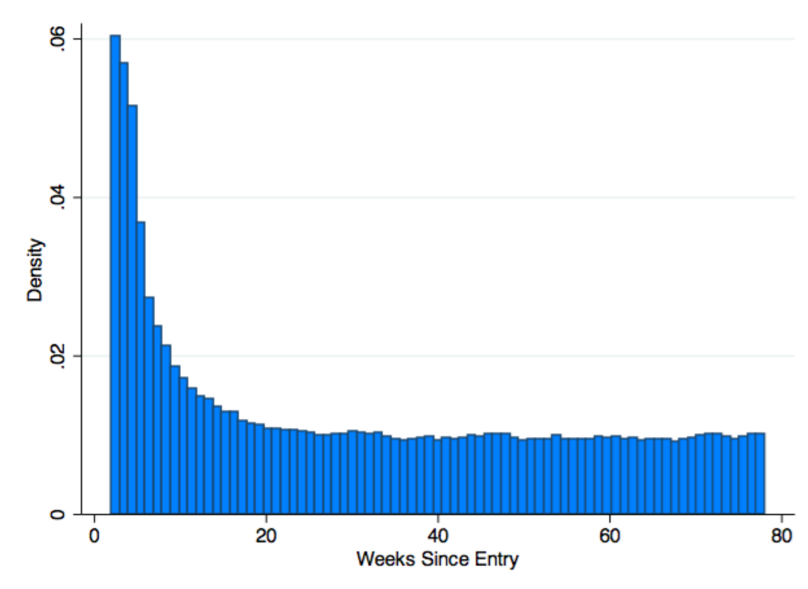

Note: The figure shows the fraction of price changes larger than 30% in a given category and city as a function of the age of the product. The products considered are those that last at least two years in the market. Source: IRI Symphony data set

Figure A.7: Direction of Price Changes Conditional on Adjustment

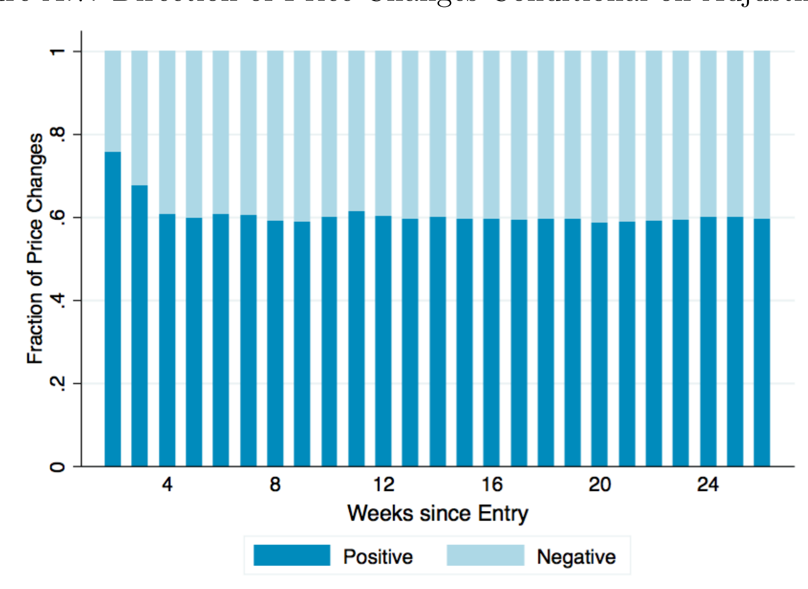

Note: The graphs plots the share of price increases and price decreases conditional on adjustment. It considers the first six months after entry. The data source is the Symphony IRI data set.

Figure A.8: Price Index for New Products

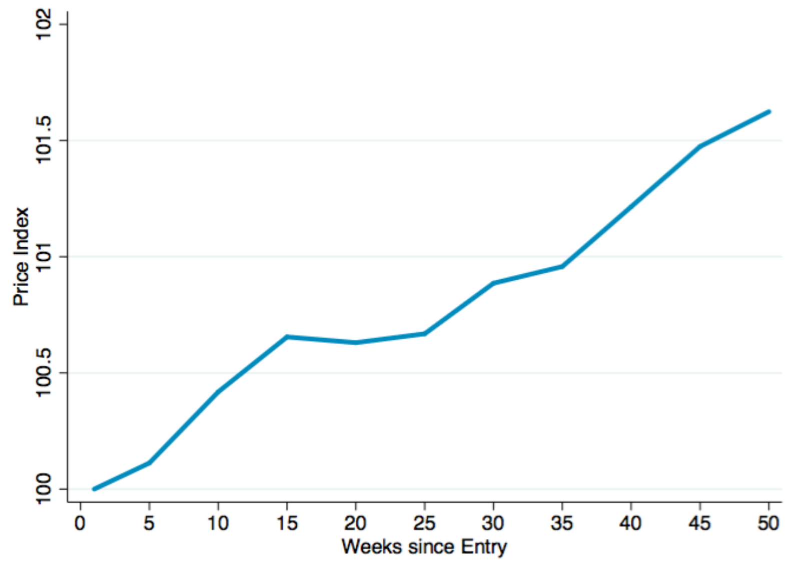

Note: The graph plots a geometric price index for new products. It considers the first year after entry. The expenditure weights are at the UPC level and based on the first year of sales of each product. The data source is the Symphony IRI data set.

Figure A.9: Distribution of Entry Prices

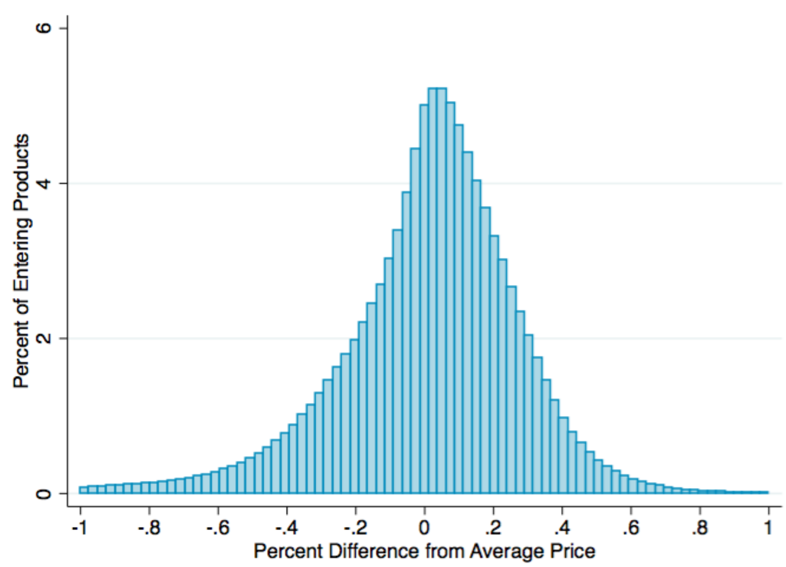

Note: The graph plots the percent difference between the entry price of all new products in our sample with respect to products of the same size, within the same category, at the store in which they were launched. The data source is the Symphony IRI dataset.

Figure A.10: Frequency of Price Adjustment (Positive and Negative)

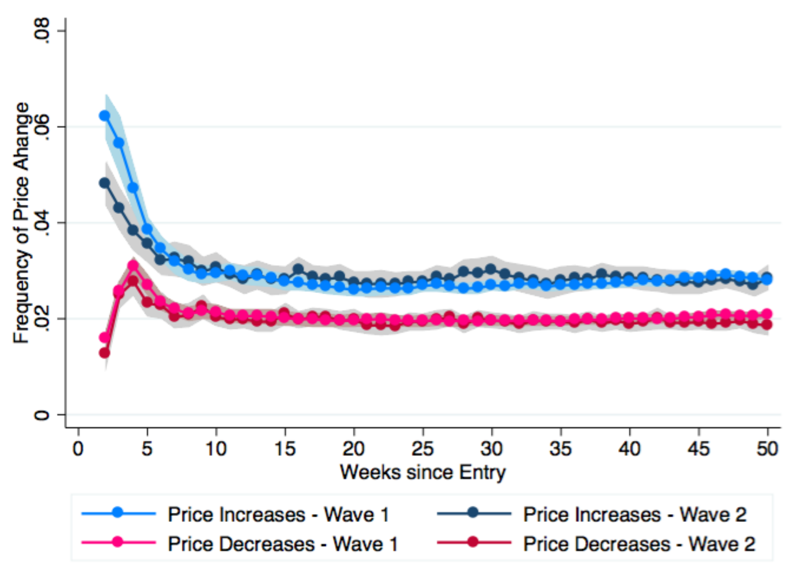

Note: The figure shows the probability of a price adjustment with respect to the mean for both price increases and decreases. Wave 1 represents products that were launched during the first year after the product was introduced. Wave 2 represents the same products when launched in different stores a year later. The graphs control for stores, time, and products fixed effects.

Figure A.11: Size of Price Adjustments (Positive and Negative)

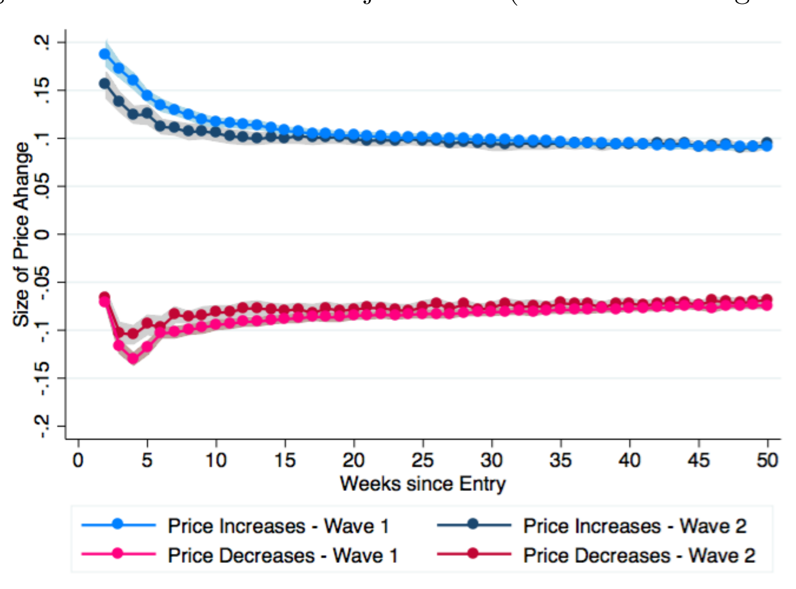

Note: The figure shows the size of price changes with respect to the mean for both price increases and decreases. Wave 1 represents products that were launched during the first year after the product was introduced. Wave 2 represents the same products when launched in different stores a year later. The graphs control for stores, time, and products fixed effects.

Figure A.12: Pricing Moments by Waves in Different Cities

Panel A: Positive Price Changes Panel B: Negative Price Changes

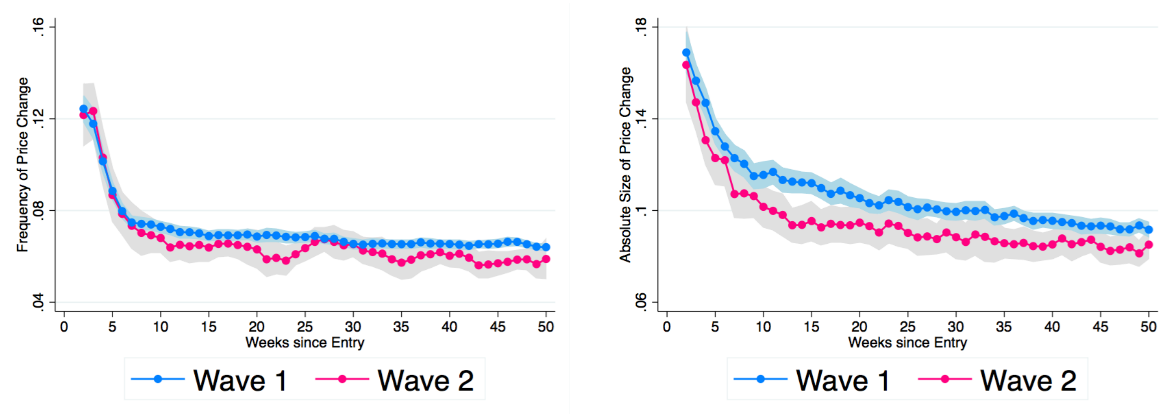

Note: The figure shows the probability of adjusting prices and the sizes of the adjustments by waves. Wave 1 represents products that were launched during the first year after the product was introduced. Wave 2 represents the same products when launched in different stores (located in different cities) a year later. Panel A shows the frequency of price adjustments and Panel B the absolute size of the price changes. The graphs control for stores, time, and products fixed effects.

Figure A.13: Fraction of Products Launched by Wave

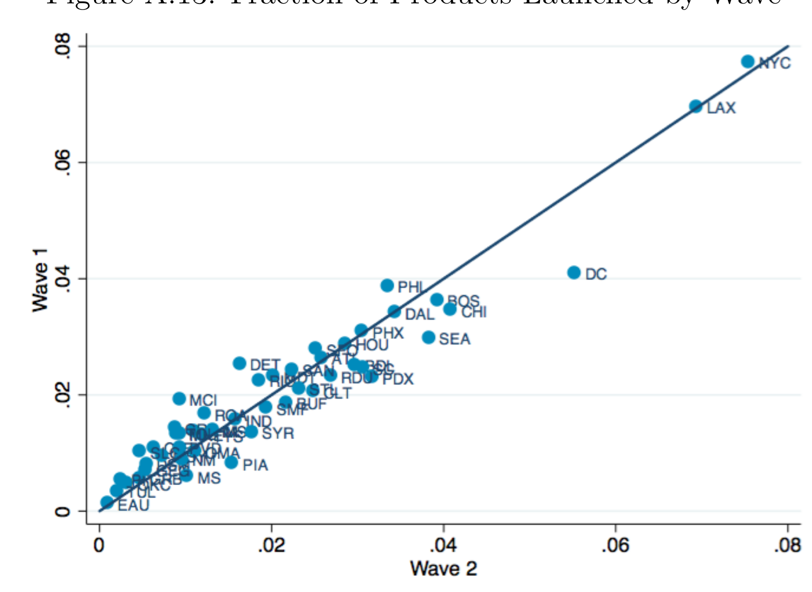

Note: The figure shows the fraction of products launched in each wave by MSA. Wave 1 represents products that were launched during the first year after the product was introduced. Wave 2 represents the same products when launched in different stores (located in different cities) a year later. The 45 degree line represents when the same fraction of new products launched in wave 1 and wave 2 for a given city.

Figure A.14: Frequency and Size of Price Changes at Exit

Panel A: Frequency — Panel B: Absolute Size

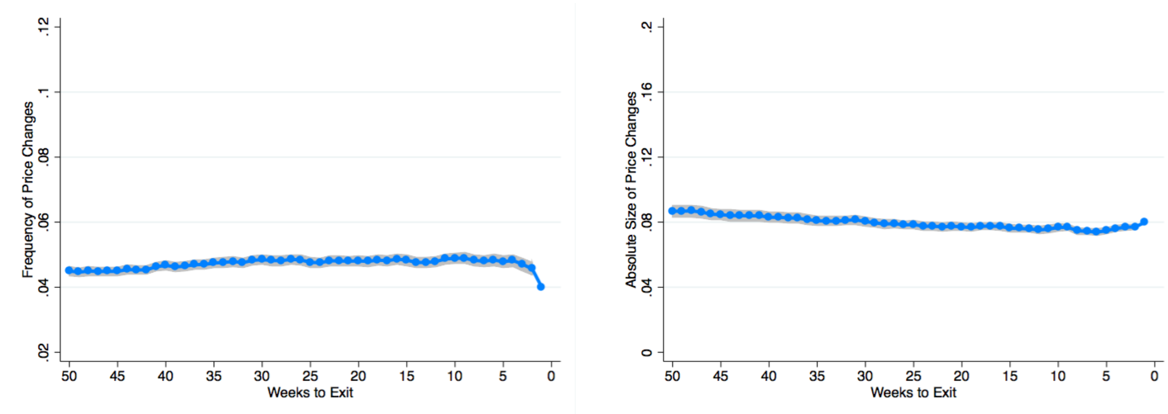

Note: Panel A plots the frequency of regular price changes at exit. Panel B plots the absolute size of regular price changes at exit. The x-axis denotes the number of weeks a product has left in the market before exiting. The graph plots the coefficients for the age fixed effects in the regression where 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. Panel A shows that the frequency of price changes stays mostly constant and decreases only around 1 percentage point near exit. Panel B shows that the absolute value of price changes stays close to its average value (around 10%) during the last weeks of the product. The calculation uses approximately 5.8 million price changes and 2.5 million stores× UPC pairs. The standard errors are clustered at the store level. The data source is the Symphony IRI data set.

Figure A.15: Frequency and Size of Sales at Entry

Panel A: Frequency — Panel B: Size

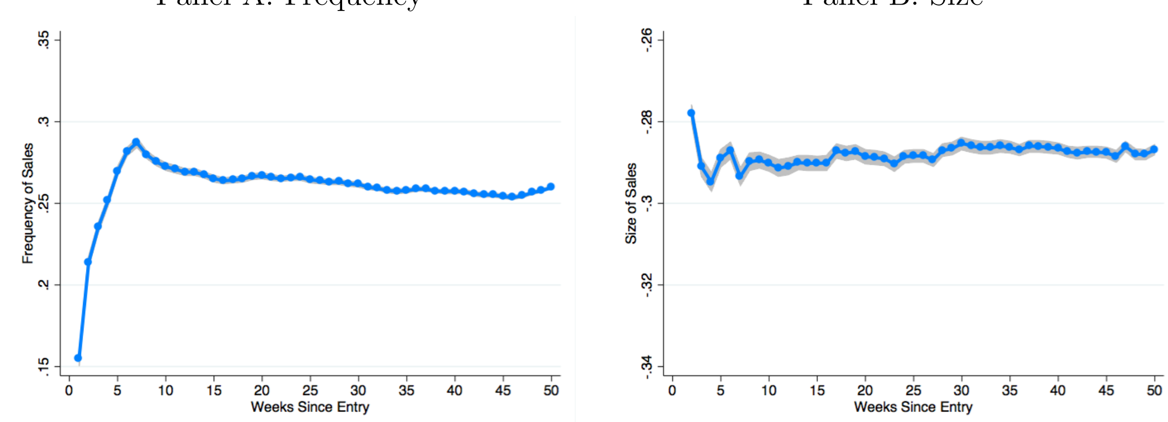

Note: Panel A plots the frequency of sales changes at entry. Panel B plots the size of the sales at entry. The x-axis denotes the number of weeks a product has been on the market. The graph plots the age fixed effects coefficients for the regression where we use the sales indicator (provided by the data) 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. Panel 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. The standard errors are clustered at the store level. The data source is the Symphony IRI data set.

Figure A.16: Frequency and Size of Sales at Exit

Panel A: Frequency   Panel B: Size

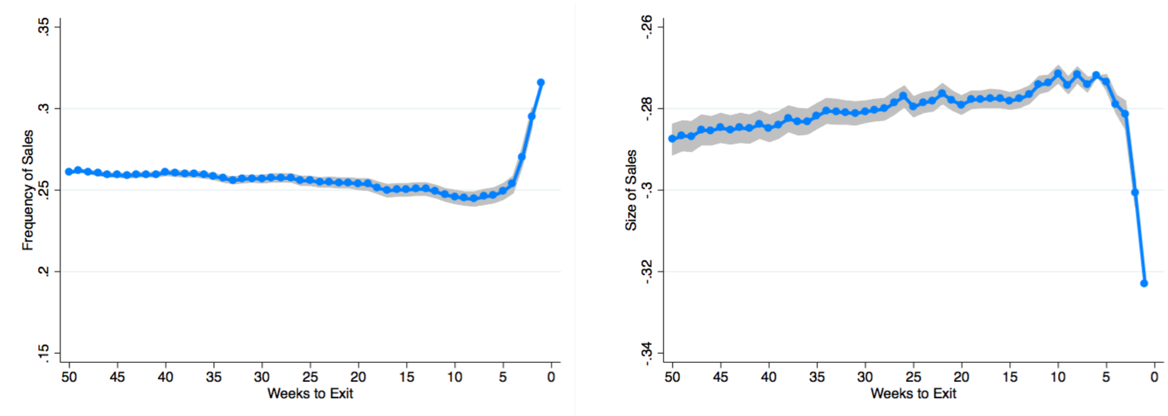

Note: Panel A plots the frequency of sales at exit. Panel B plots the absolute size of the sales at exit. The x-axis denotes the number of weeks a product has left in the market before exiting. The graph plots the coefficients for the age fixed effects in the regression where we use the the sales indicator (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. This finding indicates that firms might be attempting to liquidate their inventory before phasing out their product permanently, and they do so by offering extra discounts (”clearance” sales”). The standard errors are clustered at the store level. The data source is the Symphony IRI data set.

Figure A.17: Inaction Region: Two-Period Model with Menu Costs

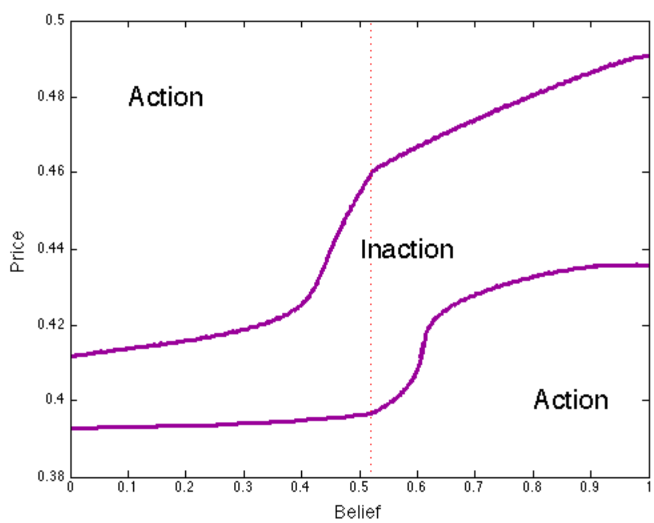

Note: The figure shows the region of inaction for the two-period model with menu costs. The y-axis denotes the previous price and the x-axis the belief. The region inside the purple lines is the region of inaction, and the dotted red line indicates the confounding belief.

A.2 Proof of Lemmas

A.2.1 Proof of Lemma 1

Proof. Recall that the value function $V_2(\lambda)$ is given by:

$$V_2(\lambda) = \max_{p \in \mathcal{P}} \left\{ \left(p - \frac{W}{z}\right) \left(\lambda\eta \frac{p^{-\sigma_1}}{\overline{P}_1^{1-\sigma_1}} + (1-\lambda)(1-\eta)\frac{p^{-\sigma_2}}{\overline{P}_2^{1-\sigma_2}}\right) \frac{S}{P} \right\}$$

The function $f(\lambda, p) \equiv \left(p - \frac{W}{z}\right) \left(\lambda\eta \frac{p^{-\sigma_1}}{\overline{P}_1^{1-\sigma_1}} + (1-\lambda)(1-\eta)\frac{p^{-\sigma_2}}{\overline{P}_2^{1-\sigma_2}}\right) \frac{S}{P}$ is continuous in $(\lambda, p) \in [0,1] \times \mathcal{P}$. The set $\mathcal{P} = [P_2^*, P_1^*]$ is furthermore compact. Then, the *Theorem of the Maximum* states that $V_2(\cdot)$ is continuous on $[0,1]$.

Convexity in λ follows almost directly. Fix an arbitrary $\alpha \in [0,1]$ and $\lambda, \lambda' \in [0,1]$. Let the convex combination $\tilde{\lambda}$ be defined as $\alpha\lambda + (1-\alpha)\lambda'$ and define the myopic policy function:

$$P^M(\lambda) = \arg\max_{p \in \mathcal{P}} \left\{ \left(p - \frac{W}{z}\right) \left(\lambda\eta \frac{p^{-\sigma_1}}{\overline{P}_1^{1-\sigma_1}} + (1-\lambda)(1-\eta)\frac{p^{-\sigma_2}}{\overline{P}_2^{1-\sigma_2}}\right) \frac{S}{P} \right\}$$

Then, we get:

$$\begin{aligned}
V_2(\tilde{\lambda}) &= \alpha \left(P^M(\tilde{\lambda}) - \frac{W}{z}\right) \left(\lambda\eta \frac{P^M(\tilde{\lambda})^{-\sigma_1}}{\overline{P}_1^{1-\sigma_1}} + (1-\lambda)(1-\eta)\frac{P^M(\tilde{\lambda})^{-\sigma_2}}{\overline{P}_2^{1-\sigma_2}}\right) \frac{S}{P} \\
&\quad + (1-\alpha) \left(P^M(\tilde{\lambda}) - \frac{W}{z}\right) \left(\lambda'\eta \frac{P^M(\tilde{\lambda})^{-\sigma_1}}{\overline{P}_1^{1-\sigma_1}} + (1-\lambda')(1-\eta)\frac{P^M(\tilde{\lambda})^{-\sigma_2}}{\overline{P}_2^{1-\sigma_2}}\right) \frac{S}{P} \\
&\leq \alpha V_2(\lambda) + (1-\alpha) V_2(\lambda')
\end{aligned}$$

Therefore, we showed $V_2(\alpha\lambda + (1-\alpha)\lambda') \leq \alpha V_2(\lambda) + (1-\alpha)V_2(\lambda')$ which is equivalent to $V_2(\cdot)$ being convex. □

A.3 Proof of Propositions

Proof of Proposition 1

Proof. Note that this proposition holds for the infinite period model as well. Suppose it is optimal for the firm to choose $P^*(\widehat{\lambda}) = \widehat{P} \in \text{int}(\mathcal{P})$ for some $\widehat{\lambda} \in (0,1)$. We show that a firm's continuation value is equal to zero whenever it chooses its price equal to $\widehat{P}$. Given some price P and prior belief λ_0, a firm's continuation value is defined as:

$$\beta\mathcal{V}(P;\lambda_0) \equiv \beta\Big(\lambda_0\mathbb{E}_\varepsilon\left[V(b_1(\lambda_0,\log(P),\varepsilon))\right] + (1-\lambda_0)\mathbb{E}_\varepsilon\left[V(b_2(\lambda_0,\log(P),\varepsilon))\right]\Big)$$

Recall that a firm faces a trade-off between maximizing current period expected profits and the value of information (through sharpening its posterior belief). The latter is captured by $\mathcal{V}(P;\lambda_0)$. As a result, a firm's marginal benefits are defined as:

$$\begin{aligned}&\lambda_0\mathbb{E}_\varepsilon\left[V'(b_1(\lambda_0,\log(P),\varepsilon))\frac{\partial b_1(\lambda_0,\log(P),\varepsilon)}{\partial\log(P)}\frac{1}{P}\right]\\ +&(1-\lambda_0)\mathbb{E}_\varepsilon\left[V'(b_2(\lambda_0,\log(P),\varepsilon))\frac{\partial b_2(\lambda_0,\log(P),\varepsilon)}{\partial\log(P)}\frac{1}{P}\right]\end{aligned}$$

Therefore, a firm's posterior belief at the confounding price $\widehat{P}$ equals its prior belief, that is, we have:

$$\begin{aligned}b_1(\lambda_0,\log(P),\varepsilon)\Big|_{P=\widehat{P}} &= b_2(\lambda_0,\log(P),\varepsilon)\Big|_{P=\widehat{P}}\\ &= \left(1+\frac{1-\lambda_0}{\lambda_0}\right)^{-1}\\ &= \lambda_0\end{aligned}$$

for *all* $\varepsilon \in \mathbb{R}$. Also, the expected change in a firm's posterior belief at $P = \widehat{P}$ is exactly equal to zero as:

$$\begin{aligned}\mathbb{E}_\varepsilon\left[\frac{\partial b_i(\lambda_0,\log(P),\varepsilon)}{\partial\log(P)}\Big|_{P=\widehat{P}}\right] &= \mathbb{E}_\varepsilon\left[\tfrac{\Delta\sigma}{\sigma_\varepsilon^2}(1-\lambda_0)\lambda_0(-1)^{\mathbf{1}(i=2)}\varepsilon\right]\\ &= 0\end{aligned}$$

for $i \in \{1,2\}$ as $\mathbb{E}_\varepsilon[\varepsilon] = 0$. Therefore, a firm's expected marginal benefit at $P = \widehat{P}$ reduces to:

$$V'(\lambda_0)\widehat{P}^{-1}\mathbb{E}_\varepsilon\left[\lambda_0\frac{\partial b_1(\lambda_0,\log(P),\varepsilon)}{\partial\log(P)}\Big|_{P=\widehat{P}} + (1-\lambda_0)\frac{\partial b_2(\lambda_0,\log(P),\varepsilon)}{\partial\log(P)}\Big|_{P=\widehat{P}}\right] = 0$$

If it is optimal for a firm to choose $P^*(\widehat{\lambda}) = \widehat{P}$, then it must be equal to $P^M(\widehat{\lambda})$ as there are no gains from active learning. Recall that $P^M(0) = P_2^*$, $P^M(1) = P_1^*$ and $P^M(\cdot)$ is strictly increasing and continuous in proposition 1. Therefore, the confounding price $\widehat{P} \in \mathcal{P}$ is guaranteed to exist. Furthermore, proposition 1 and the *Intermediate Value Theorem* state that there must be some $\widehat{\lambda}$ such that $P^M(\widehat{\lambda}) = \widehat{P}$.

By construction, we have $\widehat{P} \equiv \frac{\Delta\mu}{\Delta\sigma}$. Proposition 1 states that $P^M(\cdot)$ is strictly increasing. As a result, we derive that the confounding belief is strictly increasing (decreasing) in $\Delta\mu$ ($\Delta\sigma$) as we must have $P^M(\widehat{\lambda}) = \frac{\Delta\mu}{\Delta\sigma}$. □

Proof of Proposition 2

PROPOSITION 2. The myopic policy function $P^M(\cdot)$ is strictly increasing and $\mathcal{C}$.

Proof. The *Theorem of the Maximum* states that $P^M(\lambda)$ is a non-empty, compact-valued, and upper hemi-continuous correspondence. However, the objective function is a weighted average of strictly concave functions, thus it is strictly concave itself. As a result, $P^M(\lambda)$ must be single-valued. This value means that $P^M(\lambda)$ is not only upper hemi-continuous but continuous.

Appendix A.1 of Bachmann and Moscarini (2012) shows that $\frac{dP^M(\lambda)}{d\lambda} > 0$ if and only if $P^M(\lambda) > P^M(0) = P_2^*$ for $\lambda > 0$. By construction, this holds for $\lambda = 1$ as $P_1^* > P_2^*$ as $\sigma_2 > \sigma_1$. Thus, the inequality must hold as well for large enough λ through continuity of $P^M(\cdot)$.

Suppose by way of contradiction that for some $\lambda' > 0$, we have $P^M(\lambda') = P^M(0)$ instead. Then for some small $\Delta > 0$, we must either have $P^M(\lambda' - \Delta) > P^M(0)$, $P^M(\lambda' - \Delta) = P^M(0)$ or $P^M(\lambda' - \Delta) < P^M(0)$. The first case indicates that $\frac{dP^M(\lambda)}{d\lambda} < 0$, which contradicts the equivalence from Bachmann and Moscarini (2012). The second case states that $P^M(\lambda' - \Delta) = P^M(0)$ over an open interval of small strictly positive values of Δ. However, this value cannot be true as the expected profit function is strictly concave. Whenever $P^M(\lambda' - \Delta) < P^M(0)$, then we must have $\frac{dP^M(\lambda)}{d\lambda}\big|_{\lambda=\ell} < 0$ for all $\ell \in (0, \lambda')$. But, this means that for all $\ell \in (0, \lambda')$, we have $P^M(\ell) < P^M(0)$ but we assumed that $\lim_{\lambda \downarrow 0} P^M(\lambda) = P^M(\lambda')$. Therefore, $P^M(\lambda)$ must display a discontinuity at $\lambda = 0$. This is the desired contradiction as we showed that $P^M(\cdot)$ is continuous. Thus, $P^M(\lambda') > P^M(0)$ must hold for all $\lambda' > 0$ and $\frac{dP^M(\lambda)}{d\lambda} > 0$ follows. □

A.4 Additional Theoretical Results

LEMMA 2. The marginal expected change in a firm's posterior belief is bounded by its absolute value, that is,

$$\left|\mathbb{E}_\varepsilon\left[\frac{\partial b_i(\lambda_0,\log(p),\varepsilon)}{\partial\log(p)}\;\middle|\;\varepsilon\in\mathcal{F}\right]\right|\leq\frac{\Delta\sigma}{\sigma_\varepsilon^2}\left(\int_{\varepsilon\in\mathcal{F}}|\log(p)\Delta\sigma-\Delta\mu+\varepsilon|dF(\varepsilon)\right)$$

where the sign of $\log(p)\Delta\sigma-\Delta\mu+\varepsilon$ is constant for all $\varepsilon\in\mathcal{F}\subseteq\mathbb{R}$.

Proof. Let $x\equiv\log(P)\Delta\sigma-\Delta\mu$ and ε is contained in some set $\mathcal{F}\subseteq\mathbb{R}$. By construction of the ex post belief function $b_i(\lambda,\log(P),\varepsilon)$, we obtain:

$$\left|\mathbb{E}_\varepsilon\left[\frac{\partial b_i(\lambda_0,\log(P),\varepsilon)}{\partial\log(P)}\;\middle|\;\varepsilon\in\mathcal{F}\right]\right|=\left|\int_{\varepsilon\in\mathcal{F}}\frac{\exp\left(\frac{(\varepsilon+x)^2+\varepsilon^2}{2\sigma_\varepsilon^2}\right)\Delta\sigma(x+\varepsilon)(1-\lambda_0)\lambda_0}{\left(\exp\left(\frac{\varepsilon^2}{2\sigma_\varepsilon^2}\right)(1-\lambda_0)\sigma_\varepsilon+\exp\left(\frac{(x+\varepsilon)^2}{2\sigma_\varepsilon^2}\right)\lambda_0\sigma_\varepsilon\right)^2}dF(\varepsilon)\right|$$

$$\begin{aligned}&=\frac{\Delta\sigma}{\sigma_\varepsilon^2}\left|\int_{\varepsilon\in\mathcal{F}}\left(\frac{\exp\left(\frac{(\varepsilon+x)^2}{2\sigma_\varepsilon^2}\right)\lambda_0}{\exp\left(\frac{\varepsilon^2}{2\sigma_\varepsilon^2}\right)(1-\lambda_0)+\exp\left(\frac{(\varepsilon+x)^2}{2\sigma_\varepsilon^2}\right)\lambda_0}\right)\times\right.\\&\left.\left(\frac{\exp\left(\frac{\varepsilon^2}{2\sigma_\varepsilon^2}\right)(1-\lambda_0)}{\exp\left(\frac{\varepsilon^2}{2\sigma_\varepsilon^2}\right)(1-\lambda_0)+\exp\left(\frac{(\varepsilon+x)^2}{2\sigma_\varepsilon^2}\right)\lambda_0}\right)(x+\varepsilon)dF(\varepsilon)\right|\\&\leq\tfrac{\Delta\sigma}{\sigma_\varepsilon^2}\left(\int_{\varepsilon\in\mathcal{F}}|x+\varepsilon|dF(\varepsilon)\right)\end{aligned}$$

where the last inequality follows as the bracketed terms in the second equality are bounded by $[0,1]$ and the sign of $x+\varepsilon$ remains constant on the set $\mathcal{F}$ by assumption. This is exactly what we wanted to show. □

PROPOSITION 3. Whenever $V_2'(1)$ is small enough, then the firm's active learning policy has an interior solution, that is, $P^*(\lambda_0)\in(P_2^*,P_1^*)$ for all $\lambda_0\in(0,1)$.

Proof. We show the case for $\lambda_0\geq\frac{1}{2}$. The case for $\lambda_0<\frac{1}{2}$ follows a very similar process. We derive sufficient conditions such that $P^*(\lambda_0)\in\text{int}(\mathcal{P})$ for all $\lambda_0\in(0,1)$. This is equivalent to finding sufficient conditions such that a firm's expected marginal benefits strictly dominate its cost counterpart for $P=P_2^*$

and vice versa for $P = P_1^*$.

By construction of the ex post belief functions $b_i(\lambda_0, \log(P), \varepsilon)$, we can derive the following equality:

$$\frac{\partial b_2(\lambda_0, \log(P), \varepsilon)}{\partial \log(P)} = -\frac{\partial b_1(\lambda_0, \log(P), \varepsilon)}{\partial \log(P)} \tilde{\beta}(\varepsilon, \lambda_0, x)$$

where the function $\tilde{\beta}(\cdot)$ is characterized by:

$$\tilde{\beta}(\varepsilon, \lambda_0, x) = \left(\frac{\lambda_0 \exp\left(\frac{(x+\varepsilon)^2}{2\sigma_\varepsilon^2}\right) + (1-\lambda_0)\exp\left(\frac{\varepsilon^2}{2\sigma_\varepsilon^2}\right)}{(1-\lambda_0)\exp\left(\frac{(x+\varepsilon)^2}{2\sigma_\varepsilon^2}\right) + \lambda_0 \exp\left(\frac{\varepsilon^2}{2\sigma_\varepsilon^2}\right)} \right)^2$$

We show that $\tilde{\beta}'(\cdot, \lambda_0, x)$ is strictly increasing if and only if $x(2\lambda_0 - 1)$ is strictly positive. Furthermore, it satisfies $\lim_{\varepsilon \to +\infty} \tilde{\beta}(\varepsilon, \lambda_0, x) = \left(\frac{1-\lambda_0}{\lambda_0}\right)^2$ and $\lim_{\varepsilon \to -\infty} \tilde{\beta}(\varepsilon, \lambda_0, x) = \left(\frac{\lambda_0}{1-\lambda_0}\right)^2$. Let $x_{\min} = \log(P_2^*)\Delta\sigma - \Delta\mu < 0$ and $x_{\max} = \log(P_1^*)\Delta\sigma - \Delta\mu > 0$, then we need to show that:

$$\eta\lambda_0 \Pi_1'(P_2^*) + \frac{\beta}{P_2^*}\mathbb{E}_\varepsilon\Bigg[\Big(\lambda_0 V_2'(b_1(\lambda_0, \log(P_2^*), \varepsilon)) - (1-\lambda_0)V_2'(b_2(\lambda_0, \log(P_2^*), \varepsilon))\tilde{\beta}(\varepsilon, \lambda_0, x_{\min})\Big) \frac{\partial b_1(\lambda_0, \log(P), \varepsilon)}{\partial \log(P)}\Bigg|_{P=P_2^*}\Bigg] > 0 \qquad (\mathbf{A1})$$

$$(1-\eta)(1-\lambda_0)\Pi_2'(P_1^*) + \frac{\beta}{P_1^*}\mathbb{E}_\varepsilon\Bigg[\Big(\lambda_0 V_2'(b_1(\lambda_0, \log(P_1^*), \varepsilon)) - (1-\lambda_0)V_2'(b_2(\lambda_0, \log(P_1^*), \varepsilon))\tilde{\beta}(\varepsilon, \lambda_0, x_{\max})\Big) \frac{\partial b_1(\lambda_0, \log(P), \varepsilon)}{\partial \log(P)}\Bigg|_{P=P_1^*}\Bigg] < 0 \qquad (\mathbf{A2})$$

that are the first order conditions with respect to P in period 1 evaluate at $P = P_2^*$ and $P = P_1^*$. We start by finding a sufficient condition for the first inequality **A1**. To do this, we define the function $g(\varepsilon, \lambda_0) = \lambda_0 V_2'(b_1(\lambda_0, P_2^*, \varepsilon)) - (1-\lambda_0)V_2'(b_2(\lambda_0, P_2^*, \varepsilon))\tilde{\beta}(\varepsilon, \lambda_0, x_{\min})$. For $x = x_{\min} < 0$, and we show that $g(\cdot, \lambda_0)$ is monotonically decreasing. Furthermore, it satisfies $g(-x, \frac{1}{2}) < 0$ and $g(-x, 1) > 0$.

Whenever λ_0 is relatively close to $\frac{1}{2}$, we show that $\exists \varepsilon(\lambda_0) < -x_{\min}$ such that $g(\varepsilon, \lambda_0) > 0$ for all $\varepsilon < \varepsilon(\lambda_0)$ as $g(-x_{\min}, \frac{1}{2}) < 0$.[1] Furthermore, we derive that $\frac{\partial b_1(\lambda_0, \log(P), \varepsilon)}{\partial \log(P)}\Big|_{P=P_2^*} > 0$ if and only if $\varepsilon > -x_{\min}$. Now,

1. Whenever λ_0 is close to one, then the steps of the proof are similar. Instead, we have that $\varepsilon(\lambda_0)$ is greater than $-x_{\min}$ though.

denote $E_1 \equiv (-\infty, \varepsilon(\lambda_0))$, $E_2 \equiv (\varepsilon(\lambda_0), -x_{\min})$ and $E_3 \equiv (-x_{\min}, +\infty)$. By construction, it must be that $E_1 \cup E_2 \cup E_3 = \mathbb{R}$.

The observations above show that:

$$g(\varepsilon, \lambda_0) \frac{\partial b_1(\lambda_0, \log(P), \varepsilon)}{\partial \log(P)}\bigg|_{P=P_2^*} > 0$$

for $\varepsilon \in E_2$. Thus, it is sufficient to show:

$$\eta\lambda_0\Pi_1'(P_2^*) + \frac{\beta}{P_2^*}\mathbb{E}_\varepsilon\left[g(\varepsilon, \lambda_0)\frac{\partial b_1(\lambda_0, \log(P), \varepsilon)}{\partial \log(P)}\bigg|_{P=P_2^*}\,\bigg|\,\varepsilon \in E_1 \cup E_3\right] > 0$$

Let $\xi_2 \equiv \max_{\varepsilon \in E_3} \tilde{\beta}(\varepsilon) = \tilde{\beta}(-x_{\min})$, then observe the following strain of inequalities:

$$\begin{aligned}
\eta\lambda_0\Pi_1'(P_2^*) + \tfrac{\beta}{P_2^*}V_2'(1)\tfrac{\Delta\sigma}{\sigma_\varepsilon^2}\big(x_{\min} + \mathbb{E}_\varepsilon\left[\varepsilon|\varepsilon \leq \varepsilon(\lambda_0)\right] - \xi_2\mathbb{E}_\varepsilon\left[\varepsilon|\varepsilon \geq -x_{\min}\right]\big) <& \\
\eta\lambda_0\Pi_1'(P_2^*) + \tfrac{\beta}{P_2^*}V_2'(1)\tfrac{\Delta\sigma}{\sigma_\varepsilon^2}\big(x_{\min}F(\varepsilon(\lambda_0)) + \mathbb{E}_\varepsilon\left[\varepsilon|\varepsilon \leq \varepsilon(\lambda_0)\right]& \\
-\xi_2 x_{\min}\left(1 - F(-x_{\min})\right) - \xi_2\mathbb{E}_\varepsilon\left[\varepsilon|\varepsilon \geq -x_{\min}\right]\big) =& \\
\eta\lambda_0\Pi_1'(P_2^*) + \tfrac{\beta}{P_2^*}V_2'(1)\tfrac{\Delta\sigma}{\sigma_\varepsilon^2}\left(\mathbb{E}_\varepsilon\left[x+\varepsilon|\varepsilon \in E_1\right] - \xi_2\mathbb{E}_\varepsilon\left[x+\varepsilon|\varepsilon \in E_3\right]\right) \leq& \\
\eta\lambda_0\Pi_1'(P_2^*) + \tfrac{\beta}{P_2^*}V_2'(1)\Bigg(\mathbb{E}_\varepsilon\left[\tfrac{\partial b_1(\lambda_0,\log(P),\varepsilon)}{\partial \log(p)}\bigg|_{P=P_2^*}\,\bigg|\,\varepsilon \in E_1\right]& \\
-\xi_2\mathbb{E}_\varepsilon\left[\tfrac{\partial b_1(\lambda_0,\log(P),\varepsilon)}{\partial \log(p)}\bigg|_{P=P_2^*}\,\bigg|\,\varepsilon \in E_3\right]\Bigg) <& \\
\eta\lambda_0\Pi_1'(P_2^*) + \frac{\beta}{P_2^*}\mathbb{E}_\varepsilon\left[g(\varepsilon, \lambda_0)\frac{\partial b_1(\lambda_0, \log(P), \varepsilon)}{\partial \log(P)}\bigg|_{P=P_2^*}\,\bigg|\,\varepsilon \in E_1 \cup E_3\right]&
\end{aligned}$$

where the weak inequality follows from lemma 2 and the last strict inequality from the fact that $V_2'(1) > V_2'(\lambda')$ for any $\lambda' < 1$. This means that we are done whenever we can show:

$$\eta\lambda_0\Pi_1'(P_2^*) + \tfrac{\beta}{P_2^*}V_2'(1)\tfrac{\Delta\sigma}{\sigma_\varepsilon^2}\big(x_{\min} + \mathbb{E}_\varepsilon\left[\varepsilon|\varepsilon \leq \varepsilon(\lambda_0)\right] - \mathbb{E}_\varepsilon\left[\varepsilon|\varepsilon \geq -x_{\min}\right]\big) > 0$$

Recall that $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$. Therefore, we can use standard truncation formulas for our conditional expectations. These formulas give:

$$\mathbb{E}_\varepsilon\left[\varepsilon|\varepsilon\leq\varepsilon(\lambda_0)\right]-\xi_2\mathbb{E}_\varepsilon\left[\varepsilon|\varepsilon\geq -x_{\min}\right]=-\frac{\varphi\left(\frac{\varepsilon(\lambda_0)}{\sigma_\varepsilon}\right)}{\Phi\left(\frac{\varepsilon(\lambda_0)}{\sigma_\varepsilon}\right)}-\xi_2\frac{\varphi\left(\frac{-x_{\min}}{\sigma_\varepsilon}\right)}{1-\Phi\left(\frac{-x_{\min}}{\sigma_\varepsilon}\right)}$$

$$>-\varphi(0)\left[\frac{1}{\Phi\left(\frac{\varepsilon(\lambda_0)}{\sigma_\varepsilon}\right)}+\frac{\xi_2}{1-\Phi\left(\frac{-x_{\min}}{\sigma_\varepsilon}\right)}\right]$$

$$>-\varphi(0)\left[\frac{1+\xi_2}{1-\Phi\left(\frac{-x_{\min}}{\sigma_\varepsilon}\right)}\right]$$

Then, we can frame our first sufficient condition as:

$$\eta\lambda_0\Pi_1'(P_2^*)+\tfrac{\beta}{P_2^*}V_2'(1)\tfrac{\Delta\sigma}{\sigma_\varepsilon^2}\left[x_{\min}-\varphi(0)\tfrac{1+\xi_2}{1-\Phi\left(\frac{-x}{\sigma_\varepsilon}\right)}\right] \tag{\textbf{B1}}$$

In a similar fashion, we derive a sufficient condition for **A2**. Define $h(\varepsilon,\lambda_0)=\lambda_0V_2'(b_1(\lambda_0,P_1^*,\varepsilon))-(1-\lambda_0)V_2'(b_2(\lambda_0,P_1^*,\varepsilon))\tilde{\beta}(\varepsilon,\lambda_0,x_{\max})$, then we have $h(\cdot,\lambda_0)$ that is monotonically increasing. Once again, we can show that $\exists\varepsilon(\lambda_0)>0$ such that $h(\varepsilon,\lambda_0)$ if and only if $\varepsilon>\varepsilon(\lambda_0)$. By straightforward algebra, we can deduce that $\left.\frac{\partial b_1(\lambda_0,\log(P),\varepsilon)}{\partial\log(P)}\right|_{P=P_1^*}>0$ if and only if $\varepsilon>-x_{\max}$. Then, denote $E_1\equiv(-\infty,-x_{\max})$, $E_2\equiv(-x_{\max},\varepsilon(\lambda_0))$ and $E_3\equiv(\varepsilon(\lambda_0),+\infty)$ which satisfies $E_1\cup E_2\cup E_3=\mathbb{R}$.

With similar reasoning as before, the following condition is sufficient for **A2** to hold:

$$(1-\eta)(1-\lambda_0)\Pi_2'(P_1^*)+\frac{\beta}{P_1^*}\mathbb{E}_\varepsilon\left[\left(\lambda_0V_2'(b_1(\lambda_0,\log(P_1^*),\varepsilon))-\right.\right.$$
$$\left.\left.(1-\lambda_0)V_2'(b_2(\lambda_0,\log(P_1^*),\varepsilon))\tilde{\beta}(\varepsilon,\lambda_0,x_{\max})\right)\left.\frac{\partial b_1(\lambda_0,\log(P),\varepsilon)}{\partial\log(P)}\right|_{P=P_1^*}\middle|\varepsilon\in E_1\cup E_3\right]>0$$

Let $\xi_1\equiv\max_{\varepsilon\in E_1}\tilde{\beta}(\varepsilon)=\left(\frac{\lambda_0}{1-\lambda_0}\right)^2$, then we derive a similar chain of inequalities as before:

$$(1-\eta)(1-\lambda_0)\Pi_2'(P_1^*)+\tfrac{\beta}{P_1^*}V_2'(1)\tfrac{\Delta\sigma}{\sigma_\varepsilon^2}\left(\xi_1x_{\max}+\xi_1\mathbb{E}_\varepsilon\left[\varepsilon|\varepsilon\geq\varepsilon(\lambda_0)\right]-\mathbb{E}_\varepsilon\left[\varepsilon|\varepsilon\leq -x_{\max}\right]\right)>$$
$$(1-\eta)(1-\lambda_0)\Pi_2'(P_1^*)+\tfrac{\beta}{P_1^*}V_2'(1)\tfrac{\Delta\sigma}{\sigma_\varepsilon^2}\left(\xi_1x_{\max}\left[1-F(\varepsilon(\lambda_0))\right]+\xi_1\mathbb{E}_\varepsilon\left[\varepsilon|\varepsilon\geq\varepsilon(\lambda_0)\right]\right.$$
$$\left.-x_{\max}F(-x_{\max})-\mathbb{E}_\varepsilon\left[\varepsilon|\varepsilon\leq -x_{\max}\right]\right)$$
$$(1-\eta)(1-\lambda_0)\Pi_2'(P_1^*)+\tfrac{\beta}{P_1^*}V_2'(1)\tfrac{\Delta\sigma}{\sigma_\varepsilon^2}\left(\xi_1\mathbb{E}_\varepsilon\left[x_{\max}+\varepsilon|\varepsilon\in E_3\right]-\mathbb{E}_\varepsilon\left[x_{\max}+\varepsilon|\varepsilon\in E_1\right]\right)\geq$$

$$
\begin{aligned}
(1-\eta)(1-\lambda_0)\Pi_2'(P_1^*) + \tfrac{\beta}{P_1^*}V_2'(1)\Bigg(\xi_1\mathbb{E}_\varepsilon\left[\frac{\partial b_1(\lambda_0,\log(P),\varepsilon)}{\partial \log(p)}\Bigg|_{P=P_1^*}\ \Bigg|\ \varepsilon\in E_3\right]& \\
-\mathbb{E}_\varepsilon\left[\frac{\partial b_1(\lambda_0,\log(P),\varepsilon)}{\partial \log(p)}\Bigg|_{P=P_1^*}\ \Bigg|\ \varepsilon\in E_1\right]\Bigg) &> \\
(1-\eta)(1-\lambda_0)\Pi_2'(P_1^*) + \frac{\beta}{P_1^*}\mathbb{E}_\varepsilon\left[h(\varepsilon,\lambda_0)\frac{\partial b_1(\lambda_0,\log(P),\varepsilon)}{\partial \log(P)}\Bigg|_{P=P_1^*}\ \Bigg|\varepsilon\in E_1\cup E_3\right]&
\end{aligned}
$$

where the weak inequality follows from lemma 2 and the last strict inequality from the fact that $V_2'(1) > V_2'(\lambda')$ for any $\lambda' < 1$. This means that we are done whenever we can show:

$$
(1-\eta)(1-\lambda_0)\Pi_2'(P_1^*) + \tfrac{\beta}{P_1^*}V_2'(1)\tfrac{\Delta\sigma}{\sigma_\varepsilon^2}\left(\xi_1 x_{\max} + \xi_1\mathbb{E}_\varepsilon\left[\varepsilon|\varepsilon\geq\varepsilon(\lambda_0)\right] - \mathbb{E}_\varepsilon\left[\varepsilon|\varepsilon\leq -x_{\max}\right]\right) < 0
$$

Using the previous finding on expectations of truncated standard normal random variables, the latter inequality is satisfied whenever the following condition holds:

$$
(1-\eta)(1-\lambda_0)\Pi_2'(P_1^*) + \tfrac{\beta}{P_1^*}V_2'(1)\tfrac{\Delta\sigma}{\sigma_\varepsilon^2}\left(\xi_1 x_{\max} + \varphi(0)\frac{1+\xi_1}{\Phi\left(\frac{-x_{\max}}{\sigma_\varepsilon}\right)}\right) < 0 \qquad (\mathbf{B2})
$$

Whenever we define $\tilde{x}_{\min}$ and $\tilde{x}_{\max}$ as:

$$
\begin{aligned}
\tilde{x}_{\min} &\equiv \log(P_2^*)\Delta\sigma - \Delta\mu - \varphi(0)\frac{1+\xi_2}{1-\Phi\left(-\frac{\log(P_2^*)\Delta\sigma-\Delta\mu}{\sigma_\varepsilon}\right)} < 0, \\
\tilde{x}_{\max} &\equiv \xi_1\big(\log(P_1^*)\Delta\sigma - \Delta\mu\big) + \varphi(0)\frac{1+\xi_1}{\Phi\left(-\frac{\log(P_1^*)\Delta\sigma-\Delta\mu}{\sigma_\varepsilon}\right)} > 0.
\end{aligned}
$$

then, it is clear that **B1** and **B2** are satisfied whenever $V_2'(1)$ is bounded from above. More precisely, we get:

$$
V_2'(1) < \overline{\mathcal{V}} \equiv \frac{\sigma_\varepsilon^2}{\beta\Delta\sigma}\min\left\{\frac{\eta\lambda_0\Pi_1'(P_2^*)}{-\tilde{x}_{\min}}, \frac{(1-\eta)(1-\lambda_0)\left(-\Pi_2'(P_1^*)\right)}{\tilde{x}_{\max}}\right\}. \qquad (\mathbf{B})
$$

Thus, we have shown $\mathbf{B} \implies (\mathbf{B1}$ and $\mathbf{B2}) \implies (\mathbf{A1}$ and $\mathbf{A2})$. However, we concluded in the beginning of the proposition that $P^*(\lambda_0) \in \text{int}(\mathcal{P})$ whenever **A1** and **A2** hold. This is exactly what we wanted to show. □

A.5 Numerical Algorithm

PARAMETERS. β, δ, σ_1, σ_2, λ_0, m, s, σ_ε, ρ and σ_ζ.

NUMÉRAIRE. $W = 1$.

ALGORITHM: PSEUDO-CODE.

We assume consumer taste shocks to satisfy $\varepsilon_k \sim N(m, s^2)$. As a result, integrals over consumer taste shocks are approximated using Gaussian quadrature methods. With some abuse of notation, let the weights and nodes be denoted by $\{\omega_j^{GH}\}_{j=1}^M$ and $\{\zeta_j^{GH}\}_{j=1}^M$ respectively.[2] The quadrature weights and nodes are chosen "optimally". The nodes $\{\zeta_j^{GH}\}_{j=1}^M$ are the roots of the Hermite polynomial $H_M(\zeta)$ that is defined as $H_M(\zeta) = \oint \frac{M!}{2\pi i} \exp\left(-t^2 + 2t\zeta\right) t^{-(M+1)} dt$ and the weights are equal to:

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

We approximate the continuous AR(1) process for idiosyncratic productivity with a finite state Markov process by following the Tauchen (1986) procedure. The number of Markov states is denoted by N_T. Let the Markov transition density be denoted by $\mathcal{M}(z_i, z_j)$ for $i, j \in \{1, 2, \ldots, N_T\} \times \{1, 2, \ldots, N_T\}$.

I (initialization). Set $\overline{P}_1^0$, $\overline{P}_2^0$, χ^0 and convergence criteria $\Delta_\varepsilon, \Delta_\varphi > 0$. Let the counter k equal zero.

II (out). Given k, set $\overline{P}_1^k$, $\overline{P}_2^k$ and χ^k.

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

$$H^\eta = \frac{1}{\chi^k} \frac{1}{H + \Pi^k}$$

Then, set $S^k = H + \Pi^k$. Define the aggregate state as $\omega^k = (\overline{P}_1^k, \overline{P}_2^k, S^k)$.

2. 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.

IV. Solve the firm's problem by obtaining $V(\lambda, z, p_{-1})$:

$$
\begin{aligned}
V(\lambda, z, p_{-1}) &= \max\left\{V^A(\lambda, z), V^N(\lambda, z, p_{-1})\right\} \text{ where} \\
V^A(\lambda, z) &= \max_{p\geq 0}\,(p - \tfrac{1}{z})\left[\lambda\eta\tfrac{p^{-\sigma_1}}{\overline{P}_1^{1-\sigma_1}} + (1-\lambda)(1-\eta)\tfrac{p^{-\sigma_2}}{\overline{P}_2^{1-\sigma_2}}\right]\tfrac{1}{\overline{P}_1^{\eta}\overline{P}_2^{1-\eta}\chi} - \tfrac{\psi}{\overline{P}_1^{\eta}\overline{P}_2^{1-\eta}} \\
&\quad + \beta\lambda\sum_{i=1}^{N_T}\sum_{j=1}^{M}\mathcal{M}(z_i, z)\tfrac{1}{\sqrt{\pi}}\omega_j^{GH}V\left(b_1(\lambda, \log(\tfrac{p}{1+\tilde{\pi}}), \sqrt{2}\sigma_\alpha\zeta_j^{GH} + \mu_\alpha\right), z_i, \tfrac{p}{1+\tilde{\pi}})
\end{aligned}
$$

$$
\begin{aligned}
&\quad + \beta(1-\lambda)\sum_{i=1}^{N_T}\sum_{j=1}^{M}\mathcal{M}(z_i, z)\tfrac{1}{\sqrt{\pi}}\omega_j^{GH}V\left(b_2(\lambda, \log(\tfrac{p}{1+\tilde{\pi}}), \sqrt{2}\sigma_\alpha\zeta_j^{GH} + \mu_\alpha\right), z_i, \tfrac{p}{1+\tilde{\pi}}) \\
V^N(\lambda, z, p_{-1}) &= (p_{-1} - \tfrac{1}{z})\left[\lambda\eta\tfrac{p_{-1}^{-\sigma_1}}{\overline{P}_1^{1-\sigma_1}} + (1-\lambda)(1-\eta)\tfrac{p_{-1}^{-\sigma_2}}{\overline{P}_2^{1-\sigma_2}}\right]\tfrac{1}{\overline{P}_1^{\eta}\overline{P}_2^{1-\eta}\chi} \\
&\quad + \beta\lambda\sum_{i=1}^{N_T}\sum_{j=1}^{M}\mathcal{M}(z_i, z)\tfrac{1}{\sqrt{\pi}}\omega_j^{GH}V\left(b_1(\lambda, \log(\tfrac{p_{-1}}{1+\tilde{\pi}}), \sqrt{2}\sigma_\alpha\zeta_j^{GH} + \mu_\alpha\right), z_i, \tfrac{p_{-1}}{1+\tilde{\pi}}) \\
&\quad + \beta(1-\lambda)\sum_{i=1}^{N_T}\sum_{j=1}^{M}\mathcal{M}(z_i, z)\tfrac{1}{\sqrt{\pi}}\omega_j^{GH}V\left(b_2(\lambda, \log(\tfrac{p_{-1}}{1+\tilde{\pi}}), \sqrt{2}\sigma_\alpha\zeta_j^{GH} + \mu_\alpha\right), z_i, \tfrac{p_{-1}}{1+\tilde{\pi}}) \\
&\text{with } b_i(\lambda, \log(p), \epsilon) = \left[1 + \frac{1-\lambda}{\lambda}\frac{F'(\mu_i - \mu_2 + (\sigma_2 - \sigma_i)\log(p) + \epsilon)}{F'(\mu_i - \mu_1 + (\sigma_1 - \sigma_i)\log(p) + \epsilon)}\right]^{-1} \\
&\text{and } \mu_i = (\sigma_i - 1)\log(\overline{P}_i) + \log(\eta_i)
\end{aligned}
$$

IV. Store the optimal pricing policy function $P^*(\lambda, z)$ for every $(\lambda, z) \in [0,1] \times Z$.

V. Simulation. Simulate a panel of $N = 50000$ firms who use the policy function $P^*(\lambda, z)$.

Simulation initialization. The initial distribution $\varphi_{i,0}(\lambda, z)$ for $i = 1, 2$ is degenerate at (λ_0, z_0). For each firm $n \in \{1, 2, \ldots, N\}$, we assign it to be a firm of type $\sigma_n = \sigma_1$ with probability λ_0 and set time counter t to zero.

V.a. Given a firm's belief $\lambda_{n,t}$, let firm n set price $P^*(\lambda_{n,t}, z_n)$. We generate log sales by drawing log demand shocks $\varepsilon_{n,t} \sim N(m, s^2)$ through:

$$q_{n,t} = -\sigma_n P^*(\lambda_{n,t}, z_n) + \mu_i + s^k + \varepsilon_{n,t}$$

where $\mu_i = (\sigma_i - 1)\log(\overline{P}_i^k) + \log(\zeta_i)$. Update firm n's posterior to:

$$\lambda_{n,t+1} = B(\lambda_{n,t}, P^*(\lambda_{n,t}, z_n), q_{n,t}, S^k)$$

We apply exogenous death shocks δ to each firm. If a firm exits, then we replace it with a new firm that we assign as type σ_1 firm with probability λ_0. Its prior becomes λ_0.

V.b. We calculate $\varphi_{i,t+1}(\lambda, z)$ for each $i = 1, 2$ and stop the simulation when the distribution of beliefs settles in both measures of active firms or when the number of simulation periods exceed some upper bound $T > 1$, i.e. $\sup\limits_{\lambda\in(0,1),z\in Z} ||\varphi_{i,t+1}(\lambda, z) - \varphi_{i,t}(\lambda, z)|| < \Delta_\varphi$ for $i \in \{1, 2\}$ and/or $t = T$. Otherwise, we set $t := t + 1$ and repeat step **V.a**.

VII. Calculate $\overline{P}_i^{temp}$ with the *simulated* density $\tilde{\Phi}_i(\lambda)$:

$$\overline{P}_i^{temp} = \left(\sum_\lambda P^*(\lambda, z)^{1-\sigma_i} \tilde{\Phi}_i(\lambda, z)\right)^{\frac{1}{1-\sigma_i}}$$

where $\tilde{\Phi}_i(\lambda, z)$ is the empirical cross-sectional probability distribution function of beliefs and idiosyncratic productivity. Also, calculate the total amount of labor in the economy as:

$$H^{temp} = S^k \sum_{i=1}^{2} \frac{\eta_i}{z} \left[\frac{\sum_{\lambda,z} P^*(\lambda, z)^{-\sigma_i} \tilde{\Phi}_i(\lambda, z)}{\sum_{\lambda,z} P^*(\lambda, z)^{1-\sigma_i} \tilde{\Phi}_i(\lambda, z)}\right]$$

If $\sup_i |\overline{P}_i^{temp} - \overline{P}_i^k| < \Delta_\varepsilon$ **and** $|H - H^{temp}| < \Delta_\varepsilon$, then stop; otherwise, we set $\overline{P}_1^{k+1} = \overline{P}_i^{temp}$ and $\overline{P}_2^{k+1} = \overline{P}_2^{temp}$. Let $\chi^{k+1} > \chi^k$ if and only if $H - H^{temp} < 0$. We update the counter to $k := k + 1$ and repeat step **II**.

A.6 Robustness exercises and extensions to framework

A.6.1 Bayesian learning with a continuum of types

Our baseline framework in section 1.3 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 3.2. 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 $\mathcal{Q}_{t-1} = \{q_1, q_2, \ldots, q_{t-1})$ 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:

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

$$\Sigma_{t+1} = \Sigma_t - \frac{\Sigma_t X_t X_t' \Sigma_t}{X_t' \Sigma_t X_t + \sigma_\varepsilon^2} \tag{K2}$$

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 $B : \Theta \times \mathcal{P} \times \mathbb{R}_+ \to \Theta$. 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:

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

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:

$$\begin{aligned} p^{\text{my}}(\theta) &= \arg\max_{p \in \mathcal{P}} \Pi(p; \theta) \\ &= \frac{a + sc}{2s} \end{aligned}$$

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.[3] 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:

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

By backward induction, we obtain:

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

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.

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

Figure A.18: Numerical Example of the Two-Period Model: Continuum of Types

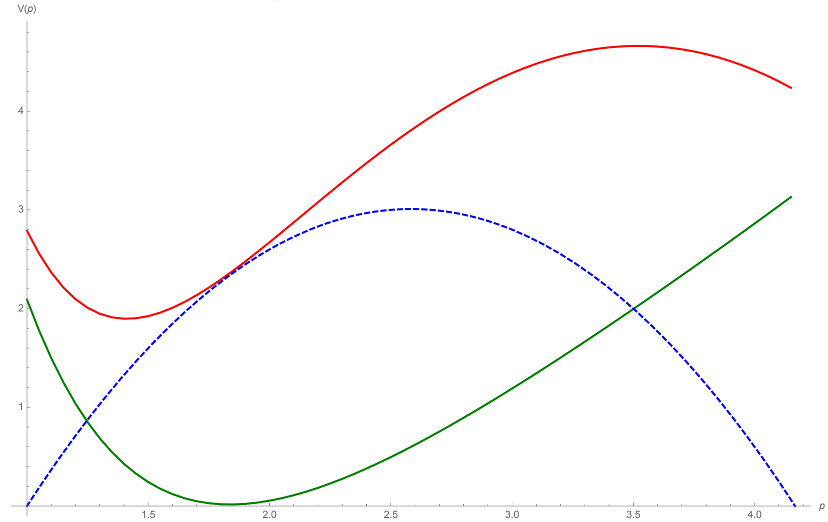

Note: The figure shows the static profits (blue), the continuation value (green), and the total payoff (red) of the two-period model with a continuum of types. The y-axis represents the total payoff and the x-axis the price.

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 1.3.2. 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 3.7 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 $(\overline{\theta}', \overline{p})$ must satisfy a set of three properties that we outline below:

$$\overline{\theta} = B(\overline{\theta}, \overline{p}, \alpha - \sigma\overline{p} + \varepsilon) \tag{B1}$$

$$\Pi(\overline{p}, \overline{\theta}) = \max_{p \in \mathcal{P}} \Pi(p; \overline{\theta}) \tag{B2}$$

$$\mathbb{E}(\alpha|\overline{\theta} - \mathbb{E}(\sigma|\overline{\theta})\overline{p} = \alpha - \sigma\overline{p} \tag{B3}$$

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 $\overline{p}$. If 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 $\overline{\theta}$ as described in equation B2. Kiefer and Nyarko (1989) refer to this policy as one-period optimization. Further, if prices are forever held at $\overline{p}$, then a firm will at least infer the true amount of sales associated at the price $\overline{p}$. Equation B3 is also known as the mean prediction property.

The solution $(\overline{\theta}', \overline{p})$ 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. $\frac{\overline{v}_{as}^2}{\overline{v}_a \overline{v}_s} = 1$.

Uncertainty. $\overline{v}_a, \overline{v}_s > 0$.

Limit actions. $\overline{p} = -\frac{\overline{v}_{as}}{\overline{v}_s} = -\frac{\overline{v}_a}{\overline{v}_{as}}$.

As a result, there is a set of incorrect, confounding beliefs under the continuum of types case. Recall from section 1.3 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:

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

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

A.6.2 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. In this section, we show that the assumption of a constant exit rate does not significantly bias the results generated by the baseline framework.

Therefore, we first show in the IRI Symphony 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.

Figure A.19: Exit Probability by Age

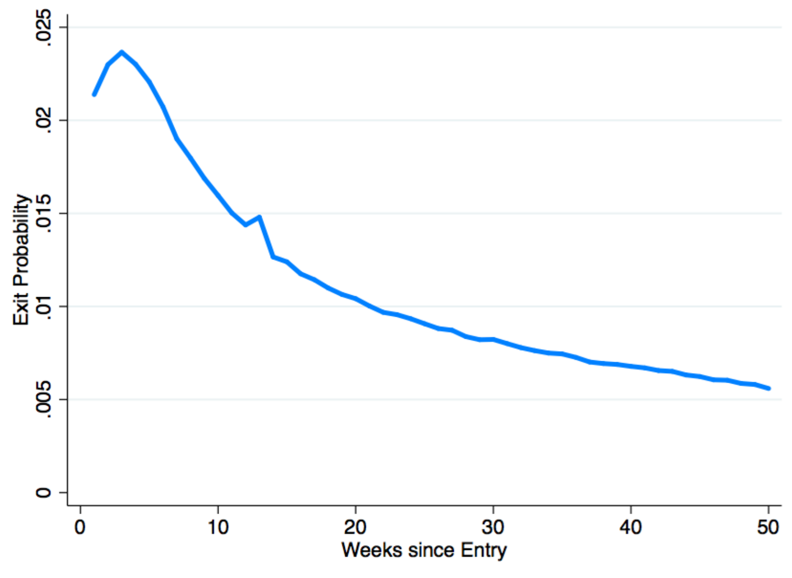

Note: The graph plots the average probability of exit of a UPC-store pair as a function of its age. We first compute the probability of exit for each category in the Symphony IRI data set. We then aggregate across categories using equal weights. The y-axis denotes the average probability exit in a given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

An alternative way of showing this fact is by estimating the product hazard function. There are multiple ways of doing this estimation, but we adopt a parsimonious one. Under this method, 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 $T \sim \mathcal{W}(\lambda, p)$:

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

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 8326.5. Then, we can form a system of two equations in the pair of unknowns (λ, p). Its solution is given by $\widehat{\lambda} = 48.13$ and $\widehat{p} = 0.8922 < 1$. Note that our calibrated value for p is not too far away from unity. This distance means that even though hazard rates decline, the 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:

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

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 1.3 by assuming that product-level exit rates depend on age as follows:

$$\delta(a) = \delta_0 \exp\left(-\delta_1 a\right)$$

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.

A.6.3 Age profile for product-level sales

In section 3.2, we showed that entering products play a substantial role in the aggregate economy. Approximately 45 percent of products in the US market entered in the last five years and they account for about 30 percent of total expenditures. However, entering products do not immediately reach these high levels of sales because, for example, the need to build customer bases. Paciello, Pozzi, and Trachter (2014) 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. In the following sections, we present two extensions to alleviate these concerns.

Demand shocks with age-dependent trend

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 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:

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

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 dataset and occurs within the first three months after the product enters the market. As a result, we use the following functional form:

$$\omega^i(k,a) = \iota a^{\upsilon}$$

where we calibrate the initial level ι and age-dependent slope υ 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. As a result, under this specification, the effects on real ouput remain quantitatively similar to those obtained in our benchmark model.

Customer base

In this section, we extend the canonical framework of Golosov and Lucas (2007) by adding a customer base. We show that such a model can rationalize the fact that product-level sales are dependent on age but this fact is not consistent with the documented life cycle patterns on product pricing.

Therefore, we model the customer base by incorporating external habits from the consumer side as is done in Gilchrist, Schoenle, Sim, and Zakrajšek (2017). Under this setup, the aggregate consumption good C_t consists of a continuum of monopolistic 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. 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). Thus, 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 $\{b_{it}\}_i$ 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:

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

Identical to the baseline framework, we assume that each firm produces with a labor 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 profits are equal to:

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

Furthermore, we assume that firms are faced with a nominal rigidity in the form of a menu cost (denoted in units of labor). A firm's dynamic programming problem is summarized by the following Bellman equation:

$$V(b_{-1}, z, p_{-1}) = \max\left\{V^A(b_{-1}, z), V^N(b_{-1}, z, p_{-1})\right\}$$

where the value of adjusting is given by:

$$\begin{aligned} V^A(b_{-1}, z) = \max_p & \left(\frac{p}{P} - \frac{W}{zP}\right) \left(\frac{p}{P}\right)^{-\sigma} b_{-1}^{\eta(1-\sigma)} C - \psi \tfrac{W}{P} \\ & + \beta \int_{z'} V\left(b, z', \frac{p}{1+\tilde{\pi}}\right) dG(z', z) \\ & \text{s.t. } b = B(b_{-1}, p) = (1-\delta^C) b_{-1} + \delta^C \left[\left(\frac{p}{P}\right)^{-\sigma} b_{-1}^{\eta(1-\sigma)} C)\right] \end{aligned}$$

which can be rewritten as:

$$\begin{aligned} V^A(b_{-1}, z) = \max_p & \left(\frac{p}{S} - \frac{\omega}{z}\right) \left(\frac{p}{S}\right)^{-\sigma} b_{-1}^{\eta(1-\sigma)} \left(\frac{P}{S}\right)^{\sigma-2} - \frac{\psi\omega S}{P} \\ & + \beta \int_{z'} V\left(B\left(b_{-1}, \frac{p}{1+\tilde{\pi}}\right), z', \frac{p}{1+\tilde{\pi}}\right) dG(z', z) \end{aligned}$$

The two crucial parameters that govern the customer base are η and δ^C. To verify whether a standard price-setting model with customer base incentives are consistent with our stylized facts, we simulate the model by choosing the parameters η and δ^C externally. Foster, Haltiwanger, and Syverson (2016) structurally estimate these parameters and find values of 0.92 and 0.188 for $\eta(1-\sigma)$ and δ^C respectively. However, this depreciation rate is based on an annual basis. Our framework's unit of time is at the weekly level. Therefore, we set δ^C to satisfy:

$$(1-\delta^C)^{52} = 1 - 0.188$$

which gives $\delta^C \simeq 0.003997$. The following figure calculates the fraction of positive and negative price changes by product age in the simulated data.

Figure A.20: Fraction of Positive and Negative Price Changes

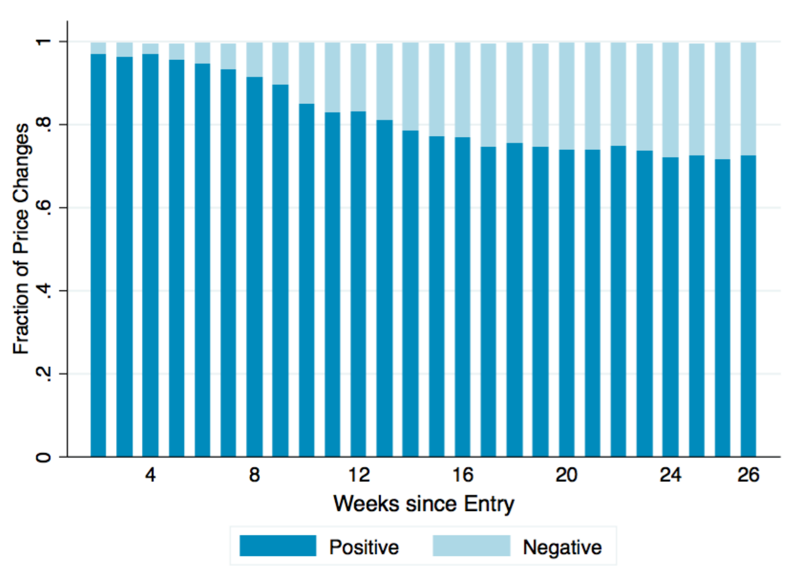

Note: The graph plots the share of regular price increases and regular price decreases conditional on adjustment. The data is generated by simulating the Golosov-Lucas model with a costumer base under the calibration described in section A.6.3. It considers the first six months after the entry of a product in the model.

The figure shows that the majority of the price changes for young products is positive, which the data contradicts. In the IRI data, we observe that roughly 60 percent of all price changes are positive, and this fraction is largely stable over the product's life cycle.

In a model with customer retention, firms have "invest" and "harvest" motives. By construction in a customer base model, current sales are dependent on the level of previous sales. This dependence means that recently entered products have relatively low prices that induce a high volume of sales and, hence, a large customer base. Once this customer base is built up to a sufficiently high level, a firm exercises its market power by increasing its price and generates profits. The simulated data indicates that customer bases are built up extremely fast. Firms set low prices to attract customers and then immediately exercise their market power afterward in a gradual fashion. This is also reflected in the frequency of price adjustments by product age. In the early stage of a product's life cycle, firms' incentives for harvesting are extremely large. Thus, they are willing to pay the menu cost to increase their prices. This incentive then slows down over time, which is consistent with the IRI data. However, we also observe in the data that the frequency of negative price adjustments weakly decreases with product age. This is contradicted by the customer base model because the frequency of negative price changes actually weakly increases.

Figure A.21: Frequency of Price Increases and Decreases at Entry

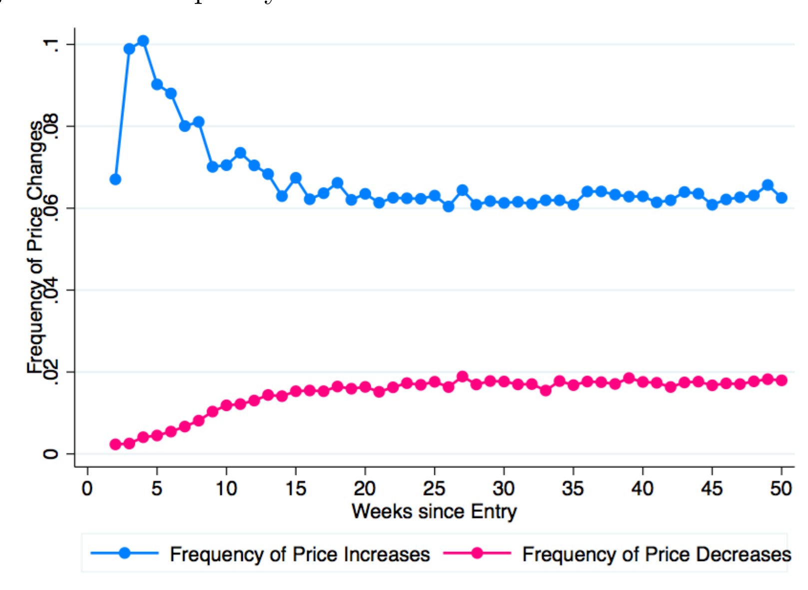

Note: The graph plots the average weekly frequency of price adjustments of products entering the market. The data is generated by simulating the Golosov-Lucas model with a costumer base under the calibration described in section A.6.3. The y-axis denotes the probability that the product adjusts prices in a given week and the x-axis denotes the number of weeks the product has been observed in the data after it entered the market. The blue line indicates the frequency of positive price adjustments and the red line the frequency of negative price adjustments.

We also show in the data that the absolute size of price changes declines with the product's age. However, we see from the figure above that a customer base does not generate such a pattern. In fact, it seems to show that the absolute size of price changes is independent of a product's age.

Figure A.22: Absolute Value of Price Changes at Entry

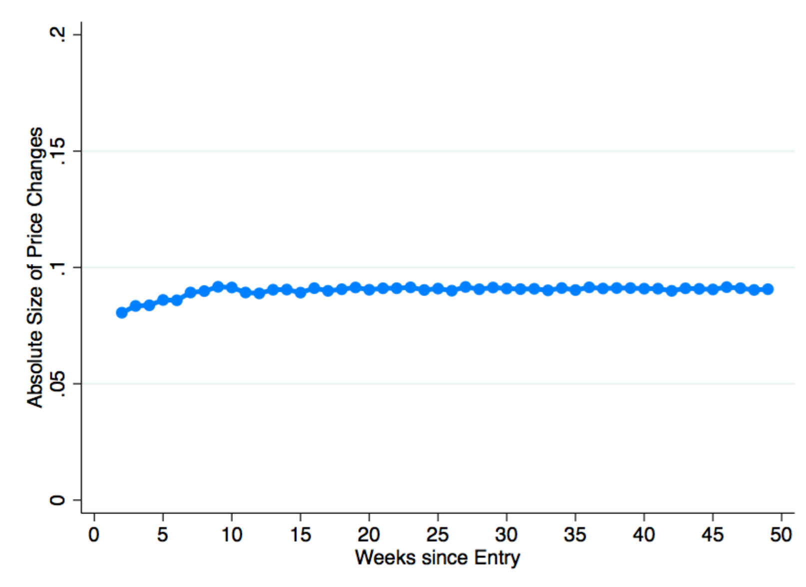

Note: The graph plots the average absolute size of price adjustments of entering products. The data is generated by simulating the Golosov-Lucas model with a costumer base under the calibration described in section A.6.3. The y-axis is the absolute value of the log price changes in that week, and the x-axis denotes the number of weeks after the product entered.

A.6.4 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 show in section 3.2, it does not capture whether the magnitude of the previously mentioned 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 (2015), 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. Aggregate productivity Z_t follows a symmetric, two state Markov chain. Thus, we have $Z_t \in \{Z_L, Z_H\}$.

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 characterized by:

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

As before, we assume that aggregate spending grows deterministically at the rate $\tilde{\pi}$, i.e. $S_t = \exp(\tilde{\pi} \cdot t)$. Recall that a price-adjusting incumbent firm has a two-dimensional idiosyncratic state. We denote this state by $\mathbf{v}_t^i(k) = (\lambda_t^i(k), z_t^i(k))$. Then, a firm's ex-ante expected profits can be written as a function of the idiosyncratic state $\mathbf{v}_t^i(k)$ and the aggregate state $\xi_t \equiv (S_t, Z_t)$:

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

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

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

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

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

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 $\xi_t = (S_t, Z_t)$. 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 labor. 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:

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

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:

$$\begin{aligned} V_t^E(z_t^*; \xi_t) &= W_t \cdot c_E \\ &= \omega S_t \cdot c_E \end{aligned}$$

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

To analyze the model in general equilibrium, we need to consider an environment in which consumers and firms engage in optimal behavior while the markets for goods and labor clear. The optimization behavior 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 labor market.

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

$$L_t^{d,p} = N_t \cdot \int_{\mathbf{v},p} \varphi_t(\mathbf{v}, p_{-1}) d\mu_t(\mathbf{v}, p_{-1})$$

where N_t denotes the actual mass of potential entrants in period t. Furthermore, labor is used for the costs of entry. Thus, total labor 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 labor supply. Then, the market for labor clears when the labor supply equals labor demand.

Lastly, we describe how to calibrate 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 $Z_L = 0.9974$.[4] 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 $\vartheta = 1 - \frac{1}{140} \simeq 0.99286$.

4. 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.

APPENDIX B

INNOVATION AND PRODUCT REALLOCATION IN THE GREAT RECESSION

B.1 Firm Data Set

Using the product level data set and the information on the firm, we construct a data set with the following information for each firm per quarter:

- type of firm Γ_{it} is defined as 1 if it is the first time we see the firm (or all products of the firm are entrant products), -1 if the last period we saw the firms was in t (or all products of the firm are exiting products), and 0 otherwise.
- total number of entrant products $N_{it} = \sum_j N_{jit}$
- number of exiting products $X_{it} = \sum_j X_{jit}$
- number of continuing products $C_{it} = \sum_j C_{jit}$
- number of all products
  - (i) T_{it} computed as number of observations of products of firm i in t, defined as defined as $N_{it} + C_{it}$
  - (ii) $T2_{it}$ computed as number of observations of products of firm i in $t-1$ plus the changes in the number of products as defined in $\Delta T2_{it}$, that is, $T2_{it-1} + \Delta T2_{it}$ with $T_{i06Q1} = T2_{i06Q1}$
- change in the number of products
  - (i) ΔT_{it} is computed as number of products in t minus number of products in $t-1$
  - (ii) $\Delta T2_{it}$ is computed as number of entrants in t minus number of exits in t, that is, change in products is defined as $N_{it} - X_{it}$
- entry rate $n_{it} = N_{it}/T_{it}$ and $n2_{it} = N_{it}/T2_{it}$
- exit rate $x_{it} = X_{it}/T_{it-1}$ and $n2_{it} = N_{it}/T2_{it-1}$
- reallocation within firm $r_{it} = n_{it} + x_{it}$ and $r_{it} = n2_{it} + x2_{it}$

- total revenue of all products rev_{it}
- total revenue of entrant a products $revN_{it}$
- total revenue of entrant b products $revNa_{it}$
- total revenue of entrant c products $revNb_{it}$
- total revenue of entrant d products $revNc_{it}$
- total revenue of entrant d products $revNd_{it}$
- total revenue of entrant e products $revNd_{it}$
- total revenue of exit products $revX_{it}$
- dummy variable to indicate if it is a multi-product firm (i.e. if number of products >1) $multi_{it}$
- number of different modules $modules_{it}$
- dummy variable to indicate if it is a multi-module firm (i.e. if number of module >1) $multimodules_{it}$
- number of different groups $groups_{it}$
- dummy variable to indicate if it is a multi-group firm (i.e. if number of group >1) $multigroups_{it}$
- number of different departments $departments_{it}$
- dummy variable to indicate if it is a multi-department firm (i.e. if number of depart. >1) $multidepartments_{it}$
- main module (defined as the module that generates the most revenue) $mainmodule_{it}$
- share of revenue in main module $revmainmodule_{it}$
- main group (defined as the group that generates the most revenue) $maingroup_{it}$
- share of revenue in main group $revmaingroup_{it}$
- main department (defined as the department that generates the most revenue) $maindepartment_{it}$
- share of revenue in main department $revmaindepartment_{it}$

- simple average quality of all products (using the variable *quality0* to *quality7* in the original data set) $q0_{it}$ $q1_{it}$ $q2_{it}$ $q3_{it}$ $q4_{it}$ $q5_{it}$ $q6_{it}$ $q7_{it}$

- simple average quality of all exit products $q0_{it}$ $q1_{it}$ $q2_{it}$ $q3_{it}$ $q4_{it}$ $q5_{it}$ $q6_{it}$ $q7_{it}$

B.2 Reallocation Measures

B.2.1 Aggregate Entry, Exit and Reallocation

We define entry, exit and reallocation as follows

Aggregate entry rate

$$n_t = \frac{\sum_i N_{it}}{\sum_i T_{it}} \tag{B.1}$$

Aggregate exit rate

$$x_t = \frac{\sum_i X_{it}}{\sum_i T_{it-1}} \tag{B.2}$$

Aggregate reallocation rate

$$r_t = n_t + x_t \tag{B.3}$$

Aggregate entry rate by firms' net growth of product

(Note that $n_t = n_t^{expanding} + n_t^{contracting} + n_t^{stable}$)

$$n_t^{expanding} = \frac{\sum_i N_{it} 1_{\{i \in \Delta T_{it} > 0\}}}{\sum_i T_{it}} \tag{B.4}$$

$$n_t^{contracting} = \frac{\sum_i N_{it} 1_{\{i \in \Delta T_{it} < 0\}}}{\sum_i T_{it}} \tag{B.5}$$

$$n_t^{stable} = \frac{\sum_i N_{it} 1_{\{i \in \Delta T_{it} = 0\}}}{\sum_i T_{it}} \tag{B.6}$$

We can define the revenue-weighted version as follows

$$\tilde{n}_t^{expanding} = \frac{\sum_i Nrev_{it} 1_{\{i \in \Delta T_{it} > 0\}}}{\sum_i Rev_{it}} \tag{B.7}$$

$$\tilde{n}_t^{contracting} = \frac{\sum_i Nrev_{it} 1_{\{i \in \Delta T_{it} < 0\}}}{\sum_i Rev_{it}} \tag{B.8}$$

$$\tilde{n}_t^{stable} = \frac{\sum_i Nrev_{it} 1_{\{i \in \Delta T_{it} = 0\}}}{\sum_i Rev_{it}} \tag{B.9}$$

Aggregate exit rate by firms' net growth of product

(Note that $x_t = x_t^{expanding} + x_t^{contracting} + x_t^{stable}$)

$$x_t^{expanding} = \frac{\sum_i X_{it} 1_{\{i \in \Delta T_{it} > 0\}}}{\sum_i T_{it-1}} \tag{B.10}$$

$$x_t^{contracting} = \frac{\sum_i X_{it} 1_{\{i \in \Delta T_{it} < 0\}}}{\sum_i T_{it-1}} \tag{B.11}$$

$$x_t^{stable} = \frac{\sum_i X_{it} 1_{\{i \in \Delta T_{it} = 0\}}}{\sum_i T_{it-1}} \tag{B.12}$$

We can define the revenue-weighted version as follows

$$\tilde{x}_t^{expanding} = \frac{\sum_i Xrev_{it} 1_{\{i \in \Delta T_{it} > 0\}}}{\sum_i Rev_{it}} \tag{B.13}$$

$$\tilde{x}_t^{contracting} = \frac{\sum_i Xrev_{it} 1_{\{i \in \Delta T_{it} < 0\}}}{\sum_i Rev_{it}} \tag{B.14}$$

$$\tilde{x}_t^{stable} = \frac{\sum_i Xrev_{it} 1_{\{i \in \Delta T_{it} = 0\}}}{\sum_i Rev_{it}} \tag{B.15}$$

Using the definitions above we can define the reallocation counterparts: $r_t^{expanding}$, $r_t^{contracting}$, r_t^{stable}, $\tilde{r}_t^{expanding}$, $\tilde{r}_t^{contracting}$, $\tilde{r}_t^{stable}$

Aggregate entry rate by firms' net growth of product and their type

(Note that $n_t = n_t^c + n_t^d + x_t^e = n_t^{expanding,entrant} + n_t^{expanding,incumbent} + n_t^{contracting,incumbent} + n_t^{stable,incumbent}$)

$$n_t^{expanding,entrant} = \frac{\sum_i N_{it} 1_{\{i \in \Delta T_{it} > 0 \bigcap i \in \Delta \Gamma_{it} > 0\}}}{\sum_i T_{it}} \tag{B.16}$$

$$n_t^{expanding,incumbent} = \frac{\sum_i N_{it} 1_{\{i \in \Delta T_{it} > 0 \bigcap i \in \Delta \Gamma_{it} = 0\}}}{\sum_i T_{it}} \tag{B.17}$$

$$n_t^{contracting,incumbent} = \frac{\sum_i N_{it} 1_{\{i \in \Delta T_{it} < 0 \bigcap i \in \Delta \Gamma_{it} = 0\}}}{\sum_i T_{it}} \tag{B.18}$$

$$n_t^{stable,incumbent} = \frac{\sum_i N_{it} 1_{\{i \in \Delta T_{it}=0 \bigcap i \in \Delta\Gamma_{it}=0\}}}{\sum_i T_{it}} \tag{B.19}$$

Aggregate exit rate by firms' net growth of product and their type

(Note that $x_t = x_t^c + x_t^d + x_t^e = x_t^{expanding,incumbent} + x_t^{contracting,incumbent} + x_t^{contracting,exit} + x_t^{stable,incumbent}$)

$$x_t^{expanding,incumbent} = \frac{\sum_i N_{it} 1_{\{i \in \Delta T_{it}>0 \bigcap i \in \Delta\Gamma_{it}=0\}}}{\sum_i T_{it-1}} \tag{B.20}$$

$$x_t^{contracting,incumbent} = \frac{\sum_i N_{it} 1_{\{i \in \Delta T_{it}<0 \bigcap i \in \Delta\Gamma_{it}=0\}}}{\sum_i T_{it-1}} \tag{B.21}$$

$$x_t^{contracting,exit} = \frac{\sum_i N_{it} 1_{\{i \in \Delta T_{it}<0 \bigcap i \in \Delta\Gamma_{it}<0\}}}{\sum_i T_{it-1}} \tag{B.22}$$

$$x_t^{stable,incumbent} = \frac{\sum_i N_{it} 1_{\{i \in \Delta T_{it}=0 \bigcap i \in \Delta\Gamma_{it}=0\}}}{\sum_i T_{it-1}} \tag{B.23}$$

As with the previous definitions of entry and exit, we can define the revenue weighted counterparts. Likewise, using the definitions above we can define the reallocation: $r_t^{expanding,incumbent}$, $r_t^{expanding,entrant}$, $r_t^{contracting,incumbent}$, $r_t^{contracting,exit}$, $r_t^{stable,incumbent}$.

Aggregate entry rate by firms' revenue level

$$n_t^{Q(rev)=k} = \frac{\sum_i N_{it} 1_{\{i \in Q(rev)=k\}}}{\sum_i T_{it}}, k = 1, 2, 3, 4 \tag{B.24}$$

Aggregate exit rate by firms' revenue level

$$x_t^{Q(rev)=k} = \frac{\sum_i N_{it} 1_{\{i \in Q(rev)_{it}=k\}}}{\sum_i T_{it}}, k = 1, 2, 3, 4 \tag{B.25}$$

Using the definitions above we can define the reallocation: $r_t^{Q(rev)=k} = n_t^{Q(rev)=k} + x_t^{Q(rev)=k}$.

Aggregate entry rate by firms' multi-product status

$$n_t^{status=k} = \frac{\sum_i N_{it} 1_{\{i \in status_{it}=k\}}}{\sum_i T_{it}}, k = 1, 2, 3, 4, 5 \tag{B.26}$$

where K=1 if a single-product firm, K=2 if multi-product but single-module firm, K=3 if multi-module but single-group firm, K=4 if multi-group but single-department firm, and K=5 if multi-department firms.

Aggregate exit rate by firms' multi-product status

$$x_t^{status=k} = \frac{\sum_i X_{it} 1_{\{i \in status_{it}=k\}}}{\sum_i T_{it}}, k = 1, 2, 3, 4, 5 \tag{B.27}$$

where K=1 if a single-product firm, K=2 if multi-product but single-module firm, K=3 if multi-module but single-group firm, K=4 if multi-group but single-department firm, and K=5 if multi-department firms.

Using the definitions above we can define the reallocation: $r_t^{status=k} = n_t^{status=k} + x_t^{status=k}$.

B.3 Revenue-based Measures

Using the RMS data set we can define the revenue of a UPC as the total revenue across all stores and weeks of the quarter. The revenue of a entering (exiting) product is defined as the revenue in the first(last) quarter of sales in any store. We define the revenue-weighted versions of product reallocation as follows:

$$n_t^* = \frac{\sum_i RevN_{it}}{\sum_i Rev_{it}} \tag{B.28}$$

$$x_t^* = \frac{\sum_i RevX_{it}}{\sum_i Rev_{it-1}} \tag{B.29}$$

$$r_t^* = n_t^* + x_t^* \tag{B.30}$$

where $RevN_{it}$, $RevX_{it}$, and Rev_{it} are the revenues of entering products, exiting products, and total products, respectively. The entry rate is defined as the revenue of entering products in period t as a share of the total revenue in period t. The exit rate is defined as the revenue of exiting products in period t (i.e., the last time we observe a transaction was in $t-1$) as a share of the total revenue of products in period $t-1$. Likewise, we define

$$\bar{n}_t^* = \frac{1}{\gamma_t}\sum_{i=1} n_{it}^* \tag{B.31}$$

$$\bar{x}_t^* = \frac{1}{\gamma_{t-1}}\sum_{i=1} x_{it}^* \tag{B.32}$$

$$\bar{r}_t^* = \bar{n}_t^* + \bar{x}_t^* \tag{B.33}$$

where $n_{it}^* = \frac{RevN_{it}}{Rev_{it}}$, $x_{it} = \frac{RevX_{it}}{Rev_{it-1}}$, and γ_t is the number of firms active in t.

Tables B.15 and B.16 present the aggregate and within reallocation measures as defined above. When compared to the measures based on a count of products, we observe that the average level of quarterly reallocation is lower: 0.9% compared to 7.9% for the aggregated reallocation, and 7.7% compared to 10.7% for the within reallocation. This is not surprising given the limitations documented above. The measures might be underestimating the importance of entries and exits within the portfolios of firms. In terms of the evolution of these revenue weighted measures over the period from 2007 to 2013, we observe an upward trend in the aggregate measures. Moreover, the deviations from trend show that both entry and exit rates are pro-cyclical, although the magnitude is quantitatively less important. When we compare the cyclical patterns of the average revenue-weighted reallocation rates we find strong pro-cyclicality, which is consistent with the results for the count-based measures.

Because of the limitations in evaluating the importance of entry and exit products, we construct alternative measures of revenue. For each firm, we estimate the average revenue of each product (excluding the entry and exit quarters) and define the revenue of entries, exits, and total products as follows:

$$RevN_{it}^{**} = \sum_j \frac{\sum_{k=t+1}^{dur_j-1} rev_{ijk}}{dur_j - 2} \times 1_{\{j \text{ entry } t, j \text{ in } i\}} \tag{B.34}$$

$$RevX_{it}^{**} = \sum_j \frac{\sum_{k=t+1}^{dur_j-1} rev_{ijk}}{dur_j - 2} \times 1_{\{j \text{ exit } t, j \text{ in } i\}} \tag{B.35}$$

$$Rev_{it}^{**} = \sum_j \frac{\sum_{k=t+1}^{dur_j-1} rev_{ijk}}{dur_j - 2} \times 1_{\{j \text{ in } i\}} \tag{B.36}$$

where Rev_{itk} is the revenue of product j of firm i in period k, and dur_j represents the duration of product j. The summation is from the quarter after entry until the quarter prior to exit. The main limitation of these measures is that the duration is censored in the last quarter of the data set. This censoring problem means that we observe several quarters of revenue for products that enter early in the period from 2007 to 2013, while we have just a few periods for products that were created closer to the end of the period. This problem affects the results substantially because of the interesting dynamics of revenue over the life cycle of products. Using these measures we can define reallocation as above. The results are presented in Tables B.17 and B.18. The level of aggregate reallocation is now more than three times bigger than before and the within measures are similar.

Our interpretation of the results is that the revenue-weighted rates of aggregate product creation and destruction are large. Over a typical 12-month interval, the revenue of entering and exiting products in the US consumer goods sector can be as large as 13% in the period from 2007 to 2013. The firm-specific revenue-weighted rates of product reallocation display substantial heterogeneity. On average, the share of firms' revenue from entering (exiting) products in a year was 18% (15%) in the period from 2007 to 2013. This share indicates that firms display substantial dynamics of product creation and destruction both in the intensive and extensive margins. When we evaluate how these evolve over the business cycle, we conclude that both extensive and intensive margins are highly pro-cyclical. A significant fraction of this pro-cyclicality is explained by a substantial decline in the number and revenue of entering products during the Great Recession and a subsequent increase in the post-recession period.

B.4 Construction of Variables: Compustat

We use the Compustat fundamental annual data from 2006 to 2014. Our sample consists of all firms that have positive data on sales that do not belong to the finance or real estate sectors (SIC classification between 6010 and 6800). To compute total factor productivity (TFP) we follow the methodology of İmrohoroğlu and Tüzel (2014). TFP is estimated using the semi-parametric method by Olley and Pakes (1996). We supplement the Compustat data with output and investment deflators from the Bureau of Economic Analysis and wage data from the Social Security Administration.

We compute R&D as the ratio of research and development expenses (Compustat item XRD) to net sales (Compustat item SALE). We winsorize the variable at the 1% level and use the one-year lagged value of in our regressions. Advertising expense (Compustat item XAD) to net sales is constructed following the same specification. Selling, general, and administrative spending (SG&A) to net sales is measured using Compustat item XSGA minus research and development expenses (item XRD) and in process research and development expenses (item RDIP). When XRD exceeds XSGA but is less than the costs of goods sold (Compustat item COGS), we measure SG&A as XSGA with no further adjustments. We make no adjustments because Compustat adds R&D and SG&A even if firms report SG&A and R&D separately.[1] We winsorize the variable at the 1% level and use the one-year lagged value of in our regressions. *Size* is the natural logarithm of total sales. Given that total staff expenses (Compustat item XLR) is mostly missing in our data, we define labor share as the ratio of the number of employees (Compustat item EMP) times the national wage (obtained from the Social Security Administration) over net sales. *Std Sale* is defined as the volatility of annual growth of sales on a quarterly basis winsorized at the 1% level.

We follow Gorodnichenko and Weber (2016) to construct the rest of our control variables. Specifically, the price to cost margin (PCM) is the ratio of net sales minus costs of goods sold to net sales. Rating (Rat) is the S&P domestic long-term issuer credit rating (Compustat item SPLTICRM) where the highest rating category, AAA, is assigned a value of 4.33, which decreases by 1/8 with every rating notch. We use the one-year lagged value of the mean ratings within the year. In order to control for financial constraints, we include the Kaplan-Zingales index (KZ, Kaplan and Zingales (1997)) that is defined as follows:

$$KZ_{i,t} = 1.002\frac{CF_{i,t}}{AT_{i,t1}} - 39.368\frac{Div_{i,t}}{AT_{i,t-1}} - 1.315\frac{C_{i,t}}{AT_{i,t1}} + 3.139Lev_{i,t} + 0.283Q_{i,t} \tag{B.37}$$

1. More details can be found in Peters and Taylor (2017).

where cash flow (CF) is the sum of income before extraordinary items (Compustat item IB) and depreciation and amortization, dividends (Div) are measured as common and preferred dividends (Compustat items DVC and DVP), C is cash and short-term investments (Compustat item CHE), leverage (Lev) is the ratio of long-term debt and debt in current liabilities (Compustat items DLTT and DLC) to stockholders equity (Compustat item SEQ), and Q is the market value of equity from CRSP as of fiscal year end minus the book value of equity and deferred taxes (Compustat items CEQ and TXDB) to total assets. The first three variables are normalized by lagged total assets. We winsorize all variables at the 1% level before calculating the index and use the one-year lagged values of the index in our regressions.

B.5 Placebo Test for Regressions

We find that an increase in R&D expenditure is associated with an increase in the reallocation rate in the next year. However, there could be a confounding factor such as time-varying firm-level shocks that affect both variables at the same time. In order to address the concern, we test a placebo specification as follows:

$$r_{f,t} = \alpha + \beta \mathrm{R\&D}_{f,t+1} + \Gamma X_{f,t} + \mu_f + \lambda_t + \epsilon_{f,t} \tag{B.38}$$

where $r_{f,t}$ represents the reallocation rate of firm f in year t. R&D represents the ratio of research and development expenses to total sales for firm f at time $t+1$. Different from the main specification, we use future R&D instead of past R&D as a placebo test.

Table B.29 reports the placebo test's results. Under all specifications, future R&D is negatively associated with the current reallocation rates. An increase in R&D expenditure leads to an increase in product reallocation in the next year, but not vice versa.

In the same way, we also conduct a placebo test for future TFP. The specification is as follows:

$$TFP_{f,t-1} = \alpha + \beta\, r_{f,t} + \Gamma X_{f,t} + \mu_f + \lambda_t + \epsilon_{f,t} \tag{B.39}$$

where as before f is the firm, and t is the year. We replace future TFP in the main specification with past TFP. Table B.30 shows that past TFP is not associated with any current reallocation activities.

B.6 Intangible Capital Investments: Alternative Measures

In this section, we evaluate the role of other intangible capital investments on product reallocation. Although we have shown that investments that lead to product improvements, such as research and development expenses, increase product reallocation; it remains to explore whether other forms of intangible capital investments have similar effects. In particular, our goal is to separate the role of demand factors (e.g., advertising, marketing expenses) that play in shaping the dynamics of product entry and exit. This is because firms also innovate by investing in product design, marketing, and customer support especially in the consumer goods sector (Corrado and Hulten, 2010). In fact, the average firm in the sample spends slightly more money on advertisement (4.8% of total sales) than on R&D (3.2% of total sales).

In order to test whether marketing efforts are associated with product reallocation, we take advantage of two useful measures available in Compustat: advertising expenses; and selling, general, and administrative spending (SG&A). Advertising expenses capture firms' desire to expand sales. Nonetheless, its main drawback is that it does not measure other expenditures associated with selling expenses such as marketing expenses. On the other hand, SG&A not only includes advertising expenses but also marketing expenses and the salaries and commissions of sales personnel.[2] Gourio and Rudanko (2014), for instance, use this measure to capture the importance of concerns on the customer base. For this reason, our preferred specification uses SG&A as dependent variable but we show that the results that follow are robust to using advertising expenses.

We follow the same specification as in equation 2.13 and use SG&A to measure the direct impact of the marketing effort on product reallocation. Table B.31 shows that, with the inclusion of firm level controls (column(3)-(4)), SG&A has a positive but not significant impact on reallocation. We obtain similar findings if we use advertisement expenses instead (see Table B.32). For robustness, and in order to separate the effect of marketing expenses and product improvements on product reallocation, we include both R&D and SG&A as dependent variables. Table B.33 shows that both variables are positive, which indicates that both product improvement concerns and demand factors are linked to product reallocation. The table also shows that, although slightly less precise, the coefficient for R&D remains significant.

2. See Peters and Taylor (2017) for more details on the interpretation of this variable and appendix B.4 for details on the construction of this measure using Compustat.

B.7 Tables and Figures

Figure B.1: Example of a Company Prefix

Note: This figure shows examples of a 6- and a 9-digit firm prefix. The source is the GS1-US website (http://www.gs1-us.info/company-prefix).

Figure B.2: Aggregate Entry, Exit, and Reallocation Rates

Figure B.3: Within-firm Entry, Exit, and Reallocation Rates

Figure B.4: Aggregated and Within Reallocation

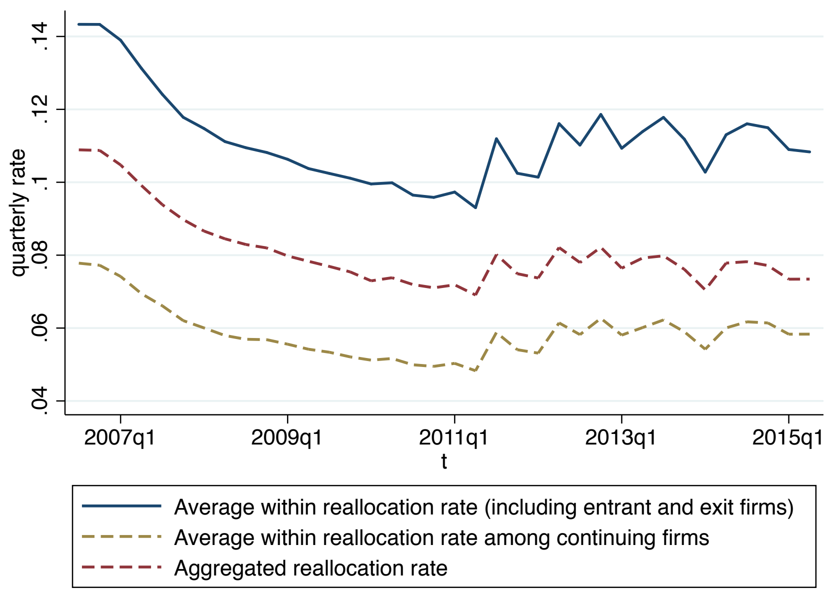

Table B.1: Comparison between Three Barcode-level Data-sets

The table compares three data sets: the Nielsen Retail Measurement Services (RMS), Nielsen Homescan (HMS), and the IRI Symphony. The three data sets report barcode transactions. The RMS and the IRI data sets are collected at the store level. The HMS is collected directly from households. The 31 categories in the IRI represent roughly 143 modules in the Nielsen data sets.

	RMS	HMS	IRI
Time period	2006-2015	2004-2015	2001-2012
Coverage	1,071 modules, 114 groups	1,077 modules, 115 groups	31 categories
Observational units	Store	Households	Store
# stores/households	35,510 stores	61,557 households	2,945 stores
# states	49	49	43
# counties	2,550	2,699	503
# products in 2006	724,211	392,455	50,434
Frequency	Weekly, average	Daily, transaction	Weekly, average
Tag temporary sales	none	deal flag household	sales flag IRI

Table B.2: Top 20 Firms in Total Revenue

The table reports the top 20 firms in total revenue for 2006 to 2014. The total revenue is in billion USD and market share is in percent. Top 3, 5, 10, and 20 firms represent 13.95, 18.35, 26.86, and 37.93 percent of total sales respectively.

Ranking	Name of Firm	Total Revenue	Market Share
1	Procter & Gamble Company	68.94	6.35
2	Kraft Heinz Foods Company	53.50	4.93
3	General Mills, Inc.	29.14	2.68
4	Nestle USA Inc.	24.35	2.24
5	Coca-Cola USA Operations	23.39	2.15
6	Frito-Lay Company	21.47	1.98
7	Pepsi-Cola North America Inc.	20.65	1.90
8	Philip Morris USA/Tobacco Products	20.52	1.89
9	Kimberly-Clark Household Products Div.	15.78	1.45
10	The Kellogg Company	13.99	1.29
11	Unilever Home and Personal Care USA	13.37	1.23
12	The J.M. Smucker Company	13.08	1.20
13	Conagra Brands, Inc.	12.80	1.18
14	The Hershey Company	12.58	1.16
15	Campbell Soup Company	12.57	1.16
16	Nabisco Biscuit Company	12.06	1.11
17	Produce Marketing Association	11.86	1.09
18	Bimbo Bakeries USA, Inc.	11.77	1.08
19	Dr. Pepper/Seven Up, Inc.	10.54	0.97
20	Nestle Purina PetCare Company	9.65	0.89

Table B.3: Summary Statistics - Publicly Traded Firms vs Matched Sample

The table reports summary statistics for all publicly traded firms and those matched with the Nielsen RMS from 2006 to 2014. The construction of these variables is detailed in Appendix B.4. R&D along with the rest of the financial variables are reported relative to net sales. We report the natural logarithm of TFP. Total sales are reported in millions. All statistics are weighted by total sales except the number of firms, employees, and sales.

	Compustat	Sample
Num. Firms	11594	479
Employees	7.70	30.77
Sales	2444	10515
R&D	0.03	0.03
TFP	0.0003	0.05
Price Cost Margin	0.70	0.67
Share of Labor	0.13	0.13
Std. Sales	0.15	0.13
KZ	1.98	1.49
Cash	0.11	0.11
Cash Flow	0.10	0.12
Dividends	0.02	0.03
Leverage	0.98	0.81
Market Value Equity	0.78	1.14
Q Ratio	0.43	0.78
S&P	2.61	3.17

Table B.4: Summary Statistics - Nielsen RMS vs Matched Sample

The table reports summary statistics for all firms in the Nielsen RMS and those matched with the sample of publicly traded firms in Compustat from 2006-2014. Average modules (groups or departments) indicates the average number of modules (groups or departments) in which the firm sells products. Revenue main module (groups or departments) indicates the share of the firm's total revenue that comes from the module (groups or departments) with the highest share of its revenue. Revenue is reported in millions. The construction of the rest of the variables is detailed in the main text. All statistics are weighted by revenue except the number of firms, number of products and revenue.

	RMS	Sample
Num. Firms	25587	372
Num. Products	15.58	225.2
Revenue	4.75	148.40
Reallocation	0.26	0.27
Entry Rate	0.24	0.24
Exit Rate	0.13	0.13
Avg. Departments	2.90	3.76
Avg. Groups	13.73	22.28
Avg. Modules	45.31	78.91
Rev. Main Group	0.66	0.50
Rev. Main Module	0.48	0.33
Rev. Main Dept.	0.87	0.78

Table B.5: Summary Statistics of Firm-level Revenue and Product Quality

The table shows summary statistics of firm-level revenue (in logs) and product quality (benchmark and percentile). To be consistent with the regressions, we keep a balanced sample of firms for the sample period.

		num. of obs.	mean	std. dev.	10th percentile	median	90th percentile
$\text{Revenue}_{f,t}$	main specification	250,561	10.74	2.99	10.71	6.97	14.54
	with full sample	345,782	10.63	3.10	6.69	10.64	14.57
$Q_{f,t}^{\text{benchmark}}$	main specification	250,561	0.04	0.66	-0.64	0.02	0.78
	with full sample	345,782	0.10	0.67	-0.59	0.08	0.86
$Q_{f,t}^{\text{percentile}}$	main specification	250,688	52.25	23.06	20.61	52.59	83.27
	with full sample	346,048	55.26	21.70	25.36	56.30	83.66

Table B.6: Characteristics by Type of Products

The table reports descriptive statistics on entering, continuing and exiting products with regards to their revenue share within the firm and to their entering quality. Quality is measured using our benchmark quality definition as described in the main text.

	All 2007-2013	(1) 2007	(2) 2010	(3) 2013
Entering products				
Within-firm revenue share	4.61%	4.75%	4.10%	5.55%
Benchmark quality	0.0030	0.0041	0.0010	0.0047
Continuing products				
Within-firm revenue share	91.96%	91.17%	92.94%	91.22%
Benchmark quality	-0.0079	-0.0043	-0.0105	-0.0055
Exiting products				
Within-firm revenue share	3.79%	4.54%	3.28%	3.65%
Benchmark quality	-0.0100	-0.0140	-0.0087	-0.0060

Table B.7: Summary Statistics of Products (with Full Sample)

The table reports summary statistics of products included in the full sample. Products are defined as a unique barcode/UPC with firm information. The variables are defined at the quarter level and grouped as averages for the period 2007Q1–2013Q4, and for the subperiods 2007Q1–2007Q4, 2010Q1–2010Q4, and 2013Q1–2013Q4. Number of products refers to the average number of products per quarter. "Entrants" refers to the average share of products that are identified for the first time in the data set in each quarter. "Exits" refers to the average share of products that are identified for the last time in the data set in each quarter. The diversification statistics report the average number of products per module, group, and department. The revenue of a product is the total sales (in dollars) across all stores and weeks of the quarter, deflated by the Consumer Price Index for All Urban Consumers. The revenue table presents the share of products with revenue in each revenue interval and is computed using all surviving products.

		2007-2013	2007	2010	2013
Number of products		564440	570289	563651	556295
By status (share of products)					
	Entrants	0.038	0.047	0.032	0.040
	Exits	0.038	0.037	0.036	0.042
By diversification (average)					
	Per module	522	544	512	521
	Per group	5037	5194	4970	4999
	Per department	52724	54301	51962	52713
By Revenue (share of products)					
	[0,500[	0.698	0.702	0.699	0.694
	[500,5000[	0.191	0.189	0.190	0.194
	[5000,50000[	0.097	0.095	0.096	0.098
	[50000,500000[	0.014	0.014	0.014	0.014

Table B.8: Summary Statistics of Products (using Brands as Product)

The table reports summary statistics of products included in the specification using brands as product. Products are defined as a unique brand with firm information. Our baseline sample excludes private label products, considers products with at least one transaction per quarter after entering, and excludes products in the Alcohol and General Merchandise departments. The variables are defined at the quarter level and grouped as averages for the period 2007Q1–2013Q4, and for the subperiods 2007Q1–2007Q4, 2010Q1–2010Q4, and 2013Q1–2013Q4. Number of products refers to the average number of products per quarter. "Entrants" refers to the average share of products that are identified for the first time in the data set in each quarter. "Exits" refers to the average share of products that are identified for the last time in the data set in each quarter. The diversification statistics report the average number of products per module, group, and department. The revenue of a product is the total sales (in dollars) across all stores and weeks of the quarter, deflated by the Consumer Price Index for All Urban Consumers. The revenue table presents the share of products with revenue in each revenue interval and is computed using all surviving products.

		2007-2013	2007	2010	2013
Number of products		62645	58946	60927	70810
By status (share of products)					
	Entrants	0.027	0.029	0.023	0.032
	Exits	0.020	0.025	0.017	0.018
By diversification (average)					
	Per module	6	6	6	7
	Per group	18	17	18	20
	Per department	140	130	136	159
By Revenue (share of products)					
	[0,500[	0.517	0.517	0.510	0.540
	[500,5000[	0.246	0.243	0.248	0.239
	[5000,50000[	0.165	0.162	0.168	0.158
	[50000,500000[	0.072	0.077	0.074	0.063

Table B.9: Summary Statistics of Firms in the Data (with Full Sample)

The variables are defined at the quarter level and then grouped as averages. The diversification statistics report the share of firms with a single product, multi-products in a single group, multi-products by multiple modules within the same group, multi-products in multiple groups within the same department, and products in multiple departments. We report the avg. and median number of products for multi-product firms. The revenue of a firm is the total sales deflated by the CPI for All Urban Consumers.

			2007-2013	2007	2010	2013
Number of firms (average)			20126	20262	19693	20571
By status (share of firms)	Entrants		0.019	0.026	0.015	0.022
	Exits		0.018	0.016	0.016	0.020
	Continuers Expanding		0.129	0.159	0.114	0.135
	Continuers Contracting		0.172	0.150	0.170	0.176
By diversification	Single Product	Share firms	0.262	0.259	0.268	0.259
	Multi-Product & Single Module	Share firms	0.279	0.283	0.277	0.278
		Mean number products	10.4	9.7	11.0	10.2
		Median number products	4	4	4	4
	Multi-Module & Single Group	Share firms	0.129	0.122	0.128	0.135
		Mean number products	19.8	19.6	20.9	19.2
		Median number products	8	8	9	8
	Multi-Group & Single Department	Share firms	0.158	0.161	0.158	0.155
		Mean number products	32.4	31.9	33.8	31.1
		Median number products	12	12	12	12
	Multi-Department	Share firms	0.172	0.174	0.169	0.174
		Mean number products	99.1	100.7	101.8	95.0
		Median number products	22	21	22	22
By Revenue (share continuing firms)	$[0,10^4[$		0.533	0.531	0.534	0.534
	$[10^4,10^5[$		0.227	0.228	0.226	0.229
	$[10^5,10^6[$		0.157	0.158	0.156	0.155
	$>=10^6$		0.083	0.084	0.083	0.082

Table B.10: Summary Statistics of Firms in the Data (using Brands as Product)

Products are defined as a unique brand. Firms are identified as defined by the GS1 company. The variables are defined at the quarter level and then grouped as averages. The number of firms refers to the average number of firms per quarter. The diversification statistics report the share of firms with a single product, multi-products in a single group, multi-products by multiple modules within the same group, multi-products in multiple groups within the same department, and products in multiple departments. We report the average and median number of products for multi-product firms. The revenue of a firm is the total sales (in dollars) across all products, deflated by the Consumer Price Index for All Urban Consumers.

			2007-2013	2007	2010	2013
Number firms (average)			14268	13965	13671	15142
By status (share firms)	Entrants		0.021	0.024	0.016	0.025
	Exits		0.017	0.021	0.014	0.016
	Continuers Expanding		0.052	0.050	0.046	0.061
	Continuers Contracting		0.041	0.049	0.038	0.038
By diversification	Single Product	Share firms	0.478	0.498	0.479	0.464
	Multi-Product					
	Single Module	Share firms	0.015	0.015	0.015	0.016
		Mean number products	2.1	2.1	2.1	2.2
		Median number products	2	2	2	2
	Multi-Module & Single Group	Share firms	0.019	0.018	0.020	0.020
		Mean number products	2.3	2.3	2.3	2.3
		Median number products	2	2	2	2
	Multi-Group & Single Department	Share firms	0.057	0.054	0.056	0.059
		Mean number products	6.1	5.6	6.3	6.3
		Median number products	3	3	3	3
	Multi-Department	Share firms	0.431	0.415	0.430	0.441
		Mean number products	8.3	8.1	8.3	8.5
		Median number products	4	4	4	4
By Revenue (share continuing firms)	$[0,10^4[$		0.435	0.439	0.425	0.447
	$[10^4,10^5[$		0.265	0.260	0.268	0.263
	$[10^5,10^6[$		0.192	0.193	0.197	0.187
	$>=10^6$		0.107	0.109	0.111	0.103

Table B.11: Aggregate Entry, Exit, and Reallocation Rates (with Full Sample)

The table reports the aggregate entry, exit, and reallocation rates, as defined in Section 2.3.1, for the full sample. Products are defined as a unique barcode/UPC with firm information. The entry, exit, and reallocation rates are computed at the quarter level, and then grouped as averages for the periods 2007Q1–2013Q4, and for the subperiods 2007Q1–2007Q4, 2010Q1–2010Q4, and 2013Q1–2013Q4.

	All 2007-2013	(1) 2007	(2) 2010	(3) 2013	Change (2)/(1)-1	(3)/(2)-1
Reallocation	0.076	0.084	0.068	0.081	-19%	19%
Entry	0.038	0.047	0.032	0.040	-31%	23%
Exit	0.038	0.038	0.036	0.042	-5%	16%

Table B.12: Aggregate Entry, Exit, and Reallocation Rates (using Brands as Product)

The table reports the aggregate entry, exit, and reallocation rates, as defined in Section 2.3.1, for the specification using brands as product. Products are defined as a unique brand with firm information. The sample excludes private label products, considers products with at least one transaction per quarter after entering, and excludes products in the Alcohol and General Merchandise departments. The entry, exit, and reallocation rates are computed at the quarter level, and then grouped as averages for the periods 2007Q1–2013Q4, and for the subperiods 2007Q1–2007Q4, 2010Q1–2010Q4, and 2013Q1–2013Q4.

	All 2007-2013	(1) 2007	(2) 2010	(3) 2013	Change (2)/(1)-1	(3)/(2)-1
Reallocation	0.047	0.054	0.041	0.051	-24%	25%
Entry	0.027	0.029	0.023	0.033	-19%	40%
Exit	0.020	0.025	0.018	0.019	-30%	6%

Table B.13: Average Within Entry, Exit, and Reallocation Rates (with Full Sample)

The table reports the average entry, exit, and reallocation firm specific rates, as defined in Section 2.3.1, for the full sample. Products are defined as a unique barcode/UPC with firm information. The entry, exit, and reallocation rates are computed for all firms (including entering and exiting firms) at the quarter level, and then grouped as averages for the period 2007Q1–2013Q4, and for the subperiods 2007Q1–2007Q4, 2010Q1–2010Q4, and 2013Q1–2013Q4.

	All 2007-2013	(1) 2007	(2) 2010	(3) 2013	Change (2)/(1)-1	(3)-(2)-1
Reallocation	0.100	0.111	0.088	0.110	-21%	25%
Entry	0.050	0.062	0.041	0.057	-35%	39%
Exit	0.052	0.052	0.049	0.056	-5%	14%

Table B.14: Average Within Entry, Exit, and Reallocation Rates (using Brands as Product)

The table reports the average entry, exit, and reallocation firm specific rates, as defined in Section 2.3.1, for the specification using brands as product. Products are defined as a unique brand with firm information. The sample excludes private label products, considers products with at least one transaction per quarter after entering, and excludes products in the Alcohol and General Merchandise departments. The entry, exit, and reallocation rates are computed for all firms (including entering and exiting firms) at the quarter level, and then grouped as averages for the period 2007Q1–2013Q4, and for the subperiods 2007Q1–2007Q4, 2010Q1–2010Q4, and 2013Q1–2013Q4.

	All 2007-2013	(1) 2007	(2) 2010	(3) 2013	Change (2)/(1)-1	(3)-(2)-1
Reallocation	0.062	0.072	0.052	0.067	-27%	29%
Entry	0.034	0.037	0.028	0.042	-24%	50%
Exit	0.029	0.037	0.025	0.027	-32%	7%

Table B.15: Aggregate Entry, Exit, and Reallocation Rates (with Revenue Weight)

The table reports the aggregate entry, exit, and reallocation rates, as defined in Appendix B.3, with revenue weights. Products are defined as a unique barcode/UPC with firm information. The sample excludes private label products, considers products with at least one transaction per quarter after entering, and excludes products in the Alcohol and General Merchandise departments. The entry, exit, and reallocation rates are computed at the quarter level, and then grouped as averages for the periods 2007Q1–2013Q4, and for the subperiods 2007Q1–2007Q4, 2010Q1–2010Q4, and 2013Q1–2013Q4.

	All 2007-2013	(1) 2007	(2) 2010	(3) 2013	Change (2)/(1)-1	(3)/(2)-1
Reallocation	0.0092	0.0087	0.0085	0.0096	-2%	13%
Entry	0.0088	0.0083	0.0082	0.0091	-1%	11%
Exit	0.0004	0.0004	0.0003	0.0005	-21%	82%

Table B.16: Average Within Entry, Exit, and Reallocation Rates (with Revenue Weight)

The table reports the average entry, exit, and reallocation firm specific rates, as defined in Appendix B.3, with revenue weights. Products are defined as a unique barcode/UPC with firm information. The sample excludes private label products, considers products with at least one transaction per quarter after entering, and excludes products in the Alcohol and General Merchandise departments. The entry, exit, and reallocation rates are computed at the quarter level, and then grouped as averages for the periods 2007Q1–2013Q4, and for the subperiods 2007Q1–2007Q4, 2010Q1–2010Q4, and 2013Q1–2013Q4.

	All 2007-2013	(1) 2007	(2) 2010	(3) 2013	Change (2)/(1)-1	(3)/(2)-1
Reallocation	0.077	0.088	0.067	0.082	-24%	22%
Entry	0.043	0.047	0.037	0.051	-21%	39%
Exit	0.036	0.044	0.032	0.033	-27%	3%

Table B.17: Aggregate Entry, Exit, and Reallocation Rates (with Average Revenue Weight)

The table reports the aggregate entry, exit, and reallocation rates, as defined in Appendix B.3, with average revenue weight. Products are defined as a unique barcode/UPC with firm information. The sample excludes private label products, considers products with at least one transaction per quarter after entering, and excludes products in the Alcohol and General Merchandise departments. The entry, exit, and reallocation rates are computed at the quarter level, and then grouped as averages for the periods 2007Q1–2013Q4, and for the subperiods 2007Q1–2007Q4, 2010Q1–2010Q4, and 2013Q1–2013Q4.

	All 2007-2013	(1) 2007	(2) 2010	(3) 2013	Change (2)/(1)-1	(3)/(2)-1
Reallocation	0.030	0.028	0.029	0.032	2%	12%
Entry	0.024	0.025	0.022	0.025	-14%	13%
Exit	0.006	0.003	0.007	0.007	145%	9%

Table B.18: Average Within Entry, Exit, and Reallocation Rates (with Average Revenue Weight)

The table reports the average entry, exit, and reallocation firm specific rates, as defined in Appendix B.3, with average revenue weights. Products are defined as a unique barcode/UPC with firm information. The sample excludes private label products, considers products with at least one transaction per quarter after entering, and excludes products in the Alcohol and General Merchandise departments. The entry, exit, and reallocation rates are computed at the quarter level, and then grouped as averages for the periods 2007Q1–2013Q4, and for the subperiods 2007Q1–2007Q4, 2010Q1–2010Q4, and 2013Q1–2013Q4.

	All 2007-2013	(1) 2007	(2) 2010	(3) 2013	Change (2)/(1)-1	(3)/(2)-1
Reallocation	0.069	0.070	0.063	0.074	-11%	19%
Entry	0.037	0.038	0.033	0.045	-13%	37%
Exit	0.032	0.033	0.030	0.030	-8%	0%

Table B.19: Decomposition (with Full Sample)

The table reports decomposition exercises as defined in Section 2.4, in the full sample. The sample excludes private label products, considers products with at least one transaction per quarter after entering, and excludes products in the Alcohol and General Merchandise departments. We decompose the first differences of the aggregate entry and exit rates. The decomposed series are seasonally adjust and then summed over the periods 2007Q1–2009Q4, and 2010Q1–2012Q4. The decomposition of the reallocation rate is computed by adding the subcomponents of the entry and exit decompositions.

	Within (+)	Between (+)	Cross (+)	Entry (+)	Exit (-)	Change
OP - Non dynamic						
Entry Rate						
07Q1 - 09Q4	-4.5	0.0	-	-	-	-3.1
10Q1 - 12Q4	1.8	-1.4	-	-	-	1.0
Exit Rate						
07Q1 - 09Q4	0.0	0.9	-	-	-	0.3
10Q1 - 12Q4	0.3	0.3	-	-	-	0.4
Reallocation Rate						
07Q1 - 09Q4	-4.5	1.0	-	-	-	-2.8
10Q1 - 12Q4	2.1	-1.0	-	-	-	1.4
OP - Dynamic						
Entry Rate						
07Q1 - 09Q4	-23.2	18.2	-	2.0	0.1	-3.1
10Q1 - 12Q4	-16.2	15.3	-	2.0	0.0	1.0
Exit Rate						
07Q1 - 09Q4	16.6	-15.2	-	0.0	1.0	0.3
10Q1 - 12Q4	18.4	-16.9	-	0.0	1.1	0.4
Reallocation Rate						
07Q1 - 09Q4	-6.7	3.0	-	2.0	1.0	-2.8
10Q1 - 12Q4	2.3	-1.7	-	2.0	1.1	1.4
GR						
Entry Rate						
07Q1 - 09Q4	-8.3	3.3		2.0	0.1	-3.1
10Q1 - 12Q4	-4.4	3.5		2.0	0.0	1.0
Exit Rate						
07Q1 - 09Q4	3.8	-2.5		0.0	1.0	0.3
10Q1 - 12Q4	3.9	-2.4		0.0	1.1	0.4
Reallocation Rate						
07Q1 - 09Q4	-4.5	0.8	-	2.0	1.0	-2.8
10Q1 - 12Q4	-0.5	1.1	-	2.0	1.1	1.4
FHK						
Entry Rate						
07Q1 - 09Q4	-10.3	1.4	3.9	2.0	0.1	-3.1
10Q1 - 12Q4	-6.4	1.5	4.0	2.0	0.0	1.0
Exit Rate						
07Q1 - 09Q4	2.4	-3.9	2.8	0.0	1.0	0.3
10Q1 - 12Q4	2.4	-3.9	3.0	0.0	1.1	0.4
Reallocation Rate						
07Q1 - 09Q4	-7.9	-2.5	6.7	2.0	1.0	-2.8
10Q1 - 12Q4	-4.0	-2.4	7.0	2.0	1.1	1.4

Table B.20: Decomposition (using Brands as Product)

The table reports decomposition exercises as defined in Section 2.4, in the specification using brands as product. The sample excludes private label products, considers products with at least one transaction per quarter after entering, and excludes products in the Alcohol and General Merchandise departments. We decompose the first differences of the aggregate entry and exit rates. The decomposed series are seasonally adjust and then summed over the periods 2007Q1–2009Q4, and 2010Q1–2012Q4. The decomposition of the reallocation rate is computed by adding the subcomponents of the entry and exit decompositions.

			Within (+)	Between (+)	Cross (+)	Entry (+)	Exit (-)	Change
OP - Non dynamic								
	Entry Rate							
		07Q1 - 09Q4	-1.8	0.6	-	-	-	-1.2
		10Q1 - 12Q4	1.7	-0.6	-	-	-	1.2
	Exit Rate							
		07Q1 - 09Q4	-1.4	0.6	-	-	-	-0.9
		10Q1 - 12Q4	-0.1	0.0	-	-	-	-0.1
	Reallocation Rate							
		07Q1 - 09Q4	-3.3	1.1	-	-	-	-2.1
		10Q1 - 12Q4	1.6	-0.6	-	-	-	1.1
OP - Dynamic								
	Entry Rate							
		07Q1 - 09Q4	-20.0	13.5	-	6.1	0.6	-1.2
		10Q1 - 12Q4	-18.2	13.6	-	6.1	0.3	1.2
	Exit Rate							
		07Q1 - 09Q4	16.8	-13.0	-	0.6	5.3	-0.9
		10Q1 - 12Q4	14.6	-11.1	-	0.3	3.9	-0.1
	Reallocation Rate							
		07Q1 - 09Q4	-3.2	0.5	-	6.1	5.3	-2.1
		10Q1 - 12Q4	-3.6	2.5	-	6.1	3.9	1.1
GR								
	Entry Rate							
		07Q1 - 09Q4	-9.1	2.6		6.0	0.6	-1.2
		10Q1 - 12Q4	-8.1	3.5		6.0	0.3	1.2
	Exit Rate							
		07Q1 - 09Q4	6.6	-2.8		0.6	5.3	-0.9
		10Q1 - 12Q4	5.6	-2.1		0.2	3.9	-0.1
	Reallocation Rate							
		07Q1 - 09Q4	-2.5	-0.2	-	6.0	5.3	-2.1
		10Q1 - 12Q4	-2.4	1.4	-	6.0	3.9	1.1
FHK								
	Entry Rate							
		07Q1 - 09Q4	-11.4	0.4	4.5	6.0	0.6	-1.2
		10Q1 - 12Q4	-10.6	0.9	5.1	6.1	0.3	1.2
	Exit Rate							
		07Q1 - 09Q4	4.5	-4.9	4.2	0.6	5.3	-0.9
		10Q1 - 12Q4	4.1	-3.6	3.1	0.2	3.9	-0.1
	Reallocation Rate							
		07Q1 - 09Q4	-6.9	-4.5	8.7	6.0	5.3	-2.1
		10Q1 - 12Q4	-6.5	-2.7	8.2	6.1	3.9	1.1

Table B.21: Reallocation Activities and R&D Expenses (with Full Sample)

The table reports the coefficients of OLS regressions with revenue weights. The dependent variable reallocation at $t+1$, defined as the product entry rate plus the product exit rate, at the firm level as defined in the main text. The main independent variable is the ratio of R&D expenses to total sales at t. The construction of the rest of the control variables is described in Appendix B.4. Other controls include firm fixed-effects and year fixed-effects. Revenue is winsorized at the 1% level. The sample used include the set of publicly traded firms that are found in the Nielsen RMS. Standard errors are presented in parentheses. ***, **, and * represent statistical significance at 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)
Dep. var.: $r_{f,t+1}$				
R&D	0.571**	0.571**	0.770***	0.771***
	(0.241)	(0.242)	(0.270)	(0.270)
Size	-0.001	-0.001	-0.000	-0.002
	(0.018)	(0.019)	(0.021)	(0.021)
Price Cost Margin		-0.004	-0.008	-0.004
		(0.110)	(0.117)	(0.118)
Std. Sale			-0.022	-0.022
			(0.067)	(0.067)
Kaplan-Zingales				0.000
				(0.001)
Observations	937	937	869	863
R-squared	0.639	0.639	0.646	0.648
Year Effects	Yes	Yes	Yes	Yes
Firm Effects	Yes	Yes	Yes	Yes

Table B.22: Reallocation Activities and R&D Expenses (using Brands as Product)

The table reports the coefficients of OLS regressions with revenue weights. The dependent variable reallocation at $t+1$, defined as the product entry rate plus the product exit rate, at the firm level as defined in the main text. The main independent variable is the ratio of R&D expenses to total sales at t. The construction of the rest of the control variables is described in Appendix B.4. Other controls include firm fixed-effects and year fixed-effects. Revenue is winsorized at the 1% level. The sample used include the set of publicly traded firms that are found in the Nielsen RMS. Standard errors are presented in parentheses. ***, **, and * represent statistical significance at 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)
Dep. var.: $r_{f,t+1}$				
R&D	0.377	0.378	0.835***	0.832***
	(0.278)	(0.279)	(0.309)	(0.310)
Size	0.056**	0.056**	0.062**	0.061**
	(0.022)	(0.022)	(0.025)	(0.025)
Price Cost Margin		-0.020	0.021	0.008
		(0.129)	(0.136)	(0.138)
Std. Sale			0.018	0.020
			(0.077)	(0.078)
Kaplan-Zingales				0.001
				(0.001)
Observations	896	896	830	824
R-squared	0.529	0.529	0.540	0.540
Year Effects	Yes	Yes	Yes	Yes
Firm Effects	Yes	Yes	Yes	Yes

Table B.23: Reallocation Activities and Revenue Growth (with Full Sample)

The table reports the coefficients of OLS regressions with revenue weights. The dependent variable is the revenue growth in the next quarter. The reallocation rate at t of firm f, $r_{f,t}$, is defined as the product entry rate plus the product exit rate at the firm level as defined in the main text. Revenue is winsorized at the 1% level. Standard errors are presented in parentheses. ***, **, and * represent statistical significance at 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dep. var.: Revenue$_{f,t+1}$							
$r_{f,t}$	0.1595***						
	(0.015)						
$n_{f,t}$		0.8499***					
		(0.019)					
$n_{f,t}$ (in module)			0.9850***				
			(0.022)				
$n_{f,t}$ (beyond module)				0.7128***			
				(0.040)			
$x_{f,t}$					-1.1852***		
					(0.023)		
$x_{f,t}$ (in module)						-0.2954***	
						(0.029)	
$x_{f,t}$ (beyond module)							-2.9092***
							(0.039)
Revenue$_{f,t}$	0.6998***	0.6968***	0.6960***	0.6996***	0.6890***	0.6899***	0.6955***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	335,064	335,211	335,211	335,211	335,087	335,087	335,087
R-squared	0.946	0.947	0.947	0.946	0.946	0.946	0.947
Year Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table B.24: Reallocation Activities and Revenue Growth (using Brands as Product)

The table reports the coefficients of OLS regressions with revenue weights. The dependent variable is the revenue growth in the next quarter. The reallocation rate at t of firm f, $r_{f,t}$, is defined as the product entry rate plus the product exit rate at the firm level as defined in the main text. Revenue is winsorized at the 1% level. Standard errors are presented in parentheses. ***, **, and * represent statistical significance at 1%, 5%, and 10% levels, respectively.

Dep. var.: Revenue$_{f,t+1}$	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$r_{f,t}$	0.1883***						
	(0.014)						
$n_{f,t}$		0.3968***					
		(0.018)					
$n_{f,t}$ (in module)			0.4776***				
			(0.039)				
$n_{f,t}$ (beyond module)				0.3709***			
				(0.021)			
$x_{f,t}$					-0.3542***		
					(0.022)		
$x_{f,t}$ (in module)						0.1779***	
						(0.066)	
$x_{f,t}$ (beyond module)							-0.4189***
							(0.023)
Revenue$_{f,t}$	0.7650***	0.7639***	0.7640***	0.7639***	0.7537***	0.7538***	0.7538***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	274,390	274,529	274,529	274,529	274,476	274,476	274,476
R-squared	0.964	0.964	0.964	0.964	0.963	0.963	0.963
Year Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table B.25: Reallocation Activities and Benchmark Quality Improvement (with Full Sample)

The table reports the coefficients of OLS regressions with revenue weights. The dependent variable is the benchmark quality improvement in the next quarter. The reallocation rate at t of firm f, $r_{f,t}$, is defined as the product entry rate plus the product exit rate at the firm level as defined in the main text. Revenue is winsorized at the 1% level. Standard errors are presented in parentheses. ***, **, and *, represent statistical significance at 1%, 5%, and 10% levels, respectively.

Dep. var.: $Q^{\text{benchmark}}_{f,t+1}$	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$r_{f,t}$	0.0417***						
	(0.004)						
$n_{f,t}$		0.0694***					
		(0.005)					
$n_{f,t}$ (in module)			0.0627***				
			(0.006)				
$n_{f,t}$ (beyond module)				0.0922***			
				(0.011)			
$x_{f,t}$					0.0061		
					(0.006)		
$x_{f,t}$ (in module)						-0.0063	
						(0.008)	
$x_{f,t}$ (beyond module)							0.0329***
							(0.011)
Revenue$_{f,t}$	0.0070***	0.0067***	0.0066***	0.0069***	0.0071***	0.0070***	0.0071***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	334,802	334,948	334,948	334,948	334,825	334,825	334,825
R-squared	0.915	0.915	0.915	0.915	0.915	0.915	0.915
Year Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table B.26: Reallocation Activities and Benchmark Quality Improvement (using Brands as Product)

The table reports the coefficients of OLS regressions with revenue weights. The dependent variable is the benchmark quality improvement in the next quarter. The reallocation rate at t of firm f, $r_{f,t}$, is defined as the product entry rate plus the product exit rate at the firm level as defined in the main text. Revenue is winsorized at the 1% level. Standard errors are presented in parentheses. ***, **, and *, represent statistical significance at 1%, 5%, and 10% levels, respectively.

Dep. var.: $Q^{\text{benchmark}}_{f,t+1}$	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$r_{f,t}$	0.0157**						
	(0.007)						
$n_{f,t}$		0.0387***					
		(0.008)					
$n_{f,t}$ (in module)			0.0400**				
			(0.018)				
$n_{f,t}$ (beyond module)				0.0463***			
				(0.010)			
$x_{f,t}$					-0.0276***		
					(0.010)		
$x_{f,t}$ (in module)						-0.0020	
						(0.031)	
$x_{f,t}$ (beyond module)							-0.0308***
							(0.011)
Revenue$_{f,t}$	0.0035***	0.0034***	0.0034***	0.0034***	0.0035***	0.0035***	0.0035***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	274,310	274,449	274,449	274,449	274,396	274,396	274,396
R-squared	0.807	0.807	0.807	0.807	0.807	0.807	0.807
Year Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table B.27: Reallocation Activities and Firm-Level Productivity (with Full Sample)

The table reports the coefficients of OLS regressions with revenue weights. The dependent variable is the natural logarithm of the total factor productivity at the firm-level at $t+1$. The reallocation at t is defined as the product entry rate plus the product exit rate at the firm level as defined in the main text. The construction of the control variables is described in Appendix B.4. Other controls include firm fixed-effects and year fixed-effects. Revenue is winsorized at the 1% level. The sample used include the set of publicly traded firms that are found in the Nielsen RMS. Standard errors are presented in parentheses. ***, **, and * represent statistical significance at 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$\text{TFP}_{f,t+1}$										
$r_{f,t}$	0.531***	0.529***	0.513***	0.506***						
	(0.110)	(0.110)	(0.113)	(0.114)						
$n_{f,t}$					0.105					
					(0.120)					
$n_{f,t}$ (in module)						0.016				
						(0.122)				
$n_{f,t}$ (beyond module)							2.480***			
							(0.604)			
$x_{f,t}$								0.858***		
								(0.173)		
$x_{f,t}$ (in module)									0.894***	
									(0.180)	
$x_{f,t}$ (beyond module)										0.768
										(0.793)
Size	0.170***	0.154***	0.194***	0.190***	0.225***	0.225***	0.226***	0.189***	0.187***	0.227***
	(0.027)	(0.028)	(0.029)	(0.029)	(0.027)	(0.027)	(0.027)	(0.029)	(0.029)	(0.027)
Price Cost Margin		-0.323*	-0.306	-0.247	-0.700***	-0.693***	-0.686***	-0.174	-0.176	-0.684***
		(0.170)	(0.196)	(0.198)	(0.188)	(0.188)	(0.187)	(0.198)	(0.198)	(0.188)
Std. Sale			-0.229**	-0.217**	-0.257**	-0.262***	-0.265***	-0.252**	-0.253**	-0.262***
			(0.102)	(0.102)	(0.100)	(0.100)	(0.099)	(0.102)	(0.102)	(0.100)
Kaplan-Zingales				0.000	0.000	0.000	-0.000	0.000	0.000	0.000
				(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	1,177	1,177	1,099	1,094	1,226	1,226	1,226	1,094	1,094	1,226
R-squared	0.845	0.845	0.850	0.851	0.834	0.833	0.836	0.852	0.852	0.834
Year Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table B.28: Reallocation Activities and Firm-Level Productivity (using Brands as Product)

The table reports the coefficients of OLS regressions with revenue weights. The dependent variable is the natural logarithm of total factor productivity at the firm-level at $t+1$. Reallocation at t is defined as the product entry rate plus the product exit rate at the firm level as defined in the main text. The construction of the control variables is described in Appendix B.4. Other controls include firm fixed-effects and year fixed-effects. Revenue is winsorized at the 1% level. The sample used include the set of publicly traded firms that are found in the Nielsen RMS. Standard errors are presented in parentheses. ***, **, and * represent statistical significance at 1%, 5%, and 10% levels, respectively.

$\text{TFP}_{f,t+1}$	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$r_{f,t}$	0.344***	0.343***	0.358***	0.353***						
	(0.072)	(0.072)	(0.073)	(0.073)						
$n_{f,t}$					0.257***					
					(0.089)					
$n_{f,t}$ (in module)						0.120				
						(0.132)				
$n_{f,t}$ (beyond module)							0.506***			
							(0.138)			
$x_{f,t}$								0.418***		
								(0.124)		
$x_{f,t}$ (in module)									0.650***	
									(0.179)	
$x_{f,t}$ (beyond module)										0.113
										(0.163)
Size	0.165***	0.149***	0.192***	0.189***	0.225***	0.223***	0.231***	0.185***	0.182***	0.224***
	(0.028)	(0.029)	(0.030)	(0.030)	(0.028)	(0.028)	(0.028)	(0.030)	(0.030)	(0.028)
Price Cost Margin		-0.329*	-0.347*	-0.277	-0.721***	-0.715***	-0.668***	-0.255	-0.219	-0.699***
		(0.175)	(0.200)	(0.202)	(0.192)	(0.194)	(0.191)	(0.204)	(0.204)	(0.193)
Std. Sale			-0.264**	-0.251**	-0.270***	-0.279***	-0.261**	-0.261**	-0.277***	-0.274***
			(0.104)	(0.104)	(0.102)	(0.102)	(0.102)	(0.105)	(0.105)	(0.102)
Kaplan-Zingales				0.000	-0.000	0.000	0.000	0.000	0.000	0.000
				(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	1,121	1,121	1,051	1,046	1,172	1,172	1,172	1,046	1,046	1,172
R-squared	0.845	0.845	0.851	0.852	0.835	0.834	0.836	0.850	0.850	0.834
Year Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table B.29: Placebo Test for Reallocation Activities and R&D Expenses (using future R&D)

The table reports the coefficients of placebo regressions with revenue weights. The dependent variable reallocation at t, defined as the product entry rate plus the product exit rate, at the firm level as defined in the main text. The main independent variable is the ratio of R&D expenses to total sales at $t+1$ (in the future). The construction of the rest of the control variables is described in Appendix B.4. Other controls include firm fixed-effects and year fixed-effects. Revenue is winsorized at the 1% level. The sample used include the set of publicly traded firms that are found in the Nielsen RMS. Standard errors are presented in parentheses. ***, **, and * represent statistical significance at 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)
Dep. var.: $r_{f,t}$				
Future R&D	-0.475	-0.649	-0.256	-0.260
	(0.444)	(0.450)	(0.526)	(0.528)
Size	0.053	0.052	0.058	0.059
	(0.047)	(0.047)	(0.050)	(0.051)
Price Cost Margin		0.554**	0.665**	0.664**
		(0.274)	(0.295)	(0.298)
Std. Sale			0.011	0.013
			(0.077)	(0.077)
Kaplan-Zingales			-0.236	-0.227
			(0.151)	(0.152)
Observations	522	522	474	470
R-squared	0.576	0.580	0.587	0.589
Year Effects	Yes	Yes	Yes	Yes
Firm Effects	Yes	Yes	Yes	Yes

Table B.30: Placebo Test for Reallocation Activities and Firm-Level Productivity (using Past Productivity)

The table reports the coefficients of placebo regressions with revenue weights. The dependent variable is the natural logarithm of total factor productivity (TFP) at the firm-level at $t-1$. As a placebo test, we use past TFP instead of future TFP as a dependent variable. Reallocation at t is defined as the product entry rate plus the product exit rate at the firm level as defined in the main text. The construction of the control variables is described in Appendix B.4. Other controls include firm fixed-effects and year fixed-effects. Revenue is winsorized at the 1% level. The sample used include the set of publicly traded firms that are found in the Nielsen RMS. Standard errors are presented in parentheses. ***, **, and * represent statistical significance at 1%, 5%, and 10% levels, respectively.

$\text{TFP}_{f,t-1}$	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$r_{f,t}$	0.005	0.019	0.010	0.007						
	(0.045)	(0.039)	(0.039)	(0.040)						
$n_{f,t}$					0.058					
					(0.059)					
$n_{f,t}$ (in module)						0.072				
						(0.068)				
$n_{f,t}$ (beyond module)							-0.000			
							(0.105)			
$x_{f,t}$								-0.048		
								(0.065)		
$x_{f,t}$ (in module)									-0.055	
									(0.086)	
$x_{f,t}$ (beyond module)										-0.029
										(0.086)
Size	0.351***	0.334***	0.342***	0.342***	0.344***	0.343***	0.342***	0.343***	0.342***	0.343***
	(0.034)	(0.030)	(0.030)	(0.030)	(0.030)	(0.030)	(0.030)	(0.030)	(0.030)	(0.030)
Price Cost Margin		-2.844***	-2.930***	-2.937***	-2.944***	-2.941***	-2.936***	-2.931***	-2.935***	-2.933***
		(0.200)	(0.205)	(0.207)	(0.207)	(0.207)	(0.207)	(0.207)	(0.207)	(0.207)
Std. Sale			-0.347***	-0.353***	-0.349***	-0.348***	-0.354***	-0.353***	-0.356***	-0.352***
			(0.107)	(0.108)	(0.108)	(0.108)	(0.108)	(0.108)	(0.108)	(0.108)
Kaplan-Zingales				0.000	0.000	0.000	0.000	0.000	0.000	0.000
				(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	734	734	683	679	679	679	679	679	679	679
R-squared	0.875	0.906	0.909	0.910	0.910	0.910	0.910	0.910	0.910	0.910
Year Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table B.31: Reallocation Activities and SG&A Spending

The table reports the coefficients of OLS regressions with revenue weights. The dependent variable reallocation at t, defined as the product entry rate plus the product exit rate, at the subsidiary level as defined in the main text. The main independent variable is the ratio of SG&A spending to total sales at $t-1$. The construction of the rest of the control variables is described in Appendix B.4. Other controls include firm fixed-effects and year fixed-effects. Revenue is winsorized at the 1% level. The sample used include the set of publicly traded firms that are found in the Nielsen RMS. Standard errors are presented in parentheses. ***, **, and * represent statistical significance at 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)
Dep. var.: $r_{f,t+1}$				
SG&A Spending	0.062	0.123	0.239	0.236
	(0.167)	(0.170)	(0.185)	(0.185)
Size	0.016	0.027	0.033	0.034
	(0.033)	(0.033)	(0.036)	(0.036)
Price Cost Margin		0.400*	0.502**	0.495**
		(0.214)	(0.229)	(0.230)
Std. Sale			-0.218*	-0.210
			(0.127)	(0.128)
Kaplan-Zingales				0.000
				(0.001)
Observations	661	661	599	595
R-squared	0.560	0.563	0.573	0.576
Year Effects	Yes	Yes	Yes	Yes
Firm Effects	Yes	Yes	Yes	Yes

Table B.32: Reallocation Activities and Advertisement Expenses

The table reports the coefficients of OLS regressions with revenue weights. The dependent variable reallocation at t, defined as the product entry rate plus the product exit rate, at the subsidiary level as defined in the main text. The main independent variable is the ratio of advertisement expenses to total sales at $t-1$. The construction of the rest of the control variables is described in Appendix B.4. Other controls include firm fixed-effects and year fixed-effects. Revenue is winsorized at the 1% level. The sample used include the set of publicly traded firms that are found in the Nielsen RMS. Standard errors are presented in parentheses. ***, **, and * represent statistical significance at 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)
Dep. var.: $r_{f,t+1}$				
Advertisement	0.007	0.124	0.321	0.322
	(0.460)	(0.464)	(0.492)	(0.492)
Size	0.023	0.025	0.019	0.020
	(0.031)	(0.031)	(0.034)	(0.034)
Price Cost Margin		0.359*	0.454**	0.459**
		(0.209)	(0.225)	(0.226)
Std. Sale			-0.260**	-0.256**
			(0.121)	(0.121)
Kaplan-Zingales				-0.000
				(0.001)
Observations	774	774	695	690
R-squared	0.554	0.556	0.557	0.560
Year Effects	Yes	Yes	Yes	Yes
Firm Effects	Yes	Yes	Yes	Yes

Table B.33: Reallocation Activities, R&D, and SG&A Spending

The table reports the coefficients of OLS regressions with revenue weights. The dependent variable reallocation at t, defined as the product entry rate plus the product exit rate, at the subsidiary level as defined in the main text. The main independent variables are the ratios of R&D and SG&A spending to total sales at $t-1$. The construction of the rest of the control variables is described in Appendix B.4. Other controls include firm fixed-effects and year fixed-effects. Revenue is winsorized at the 1% level. The sample used include the set of publicly traded firms that are found in the Nielsen RMS. Standard errors are presented in parentheses. ***, **, and * represent statistical significance at 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)
Dep. var.: $r_{f,t+1}$				
R&D	0.747*	0.636	0.808*	0.811*
	(0.402)	(0.407)	(0.439)	(0.439)
SG&A Spending	-0.085	-0.011	0.068	0.064
	(0.185)	(0.190)	(0.207)	(0.207)
Size	0.003	0.014	0.018	0.018
	(0.033)	(0.034)	(0.036)	(0.036)
Price Cost Margin		0.341	0.425*	0.418*
		(0.217)	(0.233)	(0.234)
Std. Sale			-0.176	-0.169
			(0.129)	(0.129)
Kaplan-Zingales				0.000
				(0.001)
Observations	661	661	599	595
R-squared	0.563	0.565	0.576	0.579
Year Effects	Yes	Yes	Yes	Yes
Firm Effects	Yes	Yes	Yes	Yes

APPENDIX C

COST OF LIVING INEQUALITY DURING THE GREAT RECESSION

C.1 Tables and Figures

Table C.1: Distribution of Expenditures over Departments

Code	Description	Product Groups	No. of UPCs	Expenditure
0	Health and Beauty Aids	baby care, cosmetics, cough & cold remedies, deodorant, hair care, oral hygiene, pain remedies, skin care, shaving	17.4%	10.2%
1	Dry Grocery	baby food, baking mixes, bottled water, candy, carbonated beverages, cereal, coffee, condiments, crackers, pet food, prepared foods, snacks, soup, canned vegetables	29.4%	39.3%
2	Frozen Foods	ice cream, frozen pizza, frozen vegetables	5.0%	8.8%
3	Dairy	cheese, eggs, yogurt	3.2%	9.0%
4	Deli		1.7%	2.0%
5	Packaged Meat		1.3%	3.2%
6	Fresh Produce		1.0%	2.6%
7	Non-Food Grocery	detergent, diapers, fresheners/deodorizers, household cleaners, laundry supplies, pet care	11.2%	12.9%
8	Alcohol	beer, wine, liquor, coolers	2.9%	4.0%
9	General Merchandise	batteries/flashlights, candles, computer/electronic, cookware, film/cameras, insecticides, lawn & garden, motor vehicle, office supplies	26.9%	8.0%

Note: The Nielsen Consumer data are organized into departments, product groups, product modules, and UPC codes. The departments, product groups, and product modules are all Nielsen defined codes, while the UPC codes are defined by manufacturers. The table reports the share of UPCs and expenditures in each department in the Nielsen Homescan.

Table C.2: Summary Statistics - Shopping Behavior by Income Group

	Less than 25k	25k to 50k	50k to 100k	Over 100k
Exp. to Income	0.26	0.10	0.06	0.04
Sales	0.25	0.25	0.26	0.26
Coupons	0.06	0.06	0.07	0.07
Generics	0.19	0.17	0.15	0.14
Num. of Trips per Store	4.37	4.11	3.79	3.58
Num. of Stores Visited	8.67	9.28	9.63	9.76
Num. of Chains	8.28	8.84	9.15	9.27
Num. of UPCs	156.44	182.01	193.75	185.93
Num. of Categories	90.98	103.40	109.07	105.21
Brands per Category	3.44	3.61	3.69	3.73

Note: This table reports the summary statistics of the shopping behavior of different income groups. The expenditure to income ratio is the average of the total expenditures of the households in a given income group divided by their total annual incomes. Sales is the share of expenditures of goods bought on sale. Coupons is the share of expenditures of goods bought with coupons. Generics is the share of expenditures on generic products. Number of trips per store is the average number of times a household visited a specific store. Number of stores (chains) is the average number of different stores (chains) at which the households report purchases in a given quarter. Number of UPCs (categories) is the average number of products (categories) purchased by the household every quarter. Brands per category indicate the average number of brands purchased by households of different income groups in a product category. All of the variables, except the expenditure to income ratio are at the quarterly level. The source is the Nielsen Homescan from 2004 to 2010.

Table C.3: Expenditure Shares of Top 15 Firms

Rank	Firm	Total Expenditure	Expenditure by Income Group			
			Less than 25k	25k to 50k	50k to 100k	over 100k
1	Procter & Gamble	9.16	7.85	8.61	9.53	10.53
2	Wal-Mart	8.57	10.97	10.14	7.84	4.92
3	Nestle	8.36	8.58	8.36	8.31	8.29
4	Kraft Foods	5.64	5.48	5.59	5.72	5.68
5	General Mills	5.14	4.65	4.98	5.33	5.39
6	Kroger	4.95	5.31	5.19	4.84	4.41
7	Conagra Foods	3.09	3.58	3.23	2.93	2.74
8	Pepsico	2.97	3.32	3.14	2.91	2.44
9	Coca-Cola	2.74	2.69	2.64	2.78	2.91
10	Campbell Soup	2.36	2.43	2.40	2.34	2.27
11	Kellogg	2.13	1.82	1.97	2.27	2.40
12	Costco	2.09	0.87	1.46	2.38	3.79
13	Unilever Group	2.09	1.89	2.04	2.16	2.17
14	Supervalu	2.07	2.20	2.06	2.04	2.04
15	Heinz	1.69	1.61	1.66	1.74	1.67
	Sum	63.05	63.25	63.49	63.13	61.66

Note: This table shows the expenditure shares by firms in the Homescan data. We link manufacturers and firms by merging the barcode provided by Nielsen and the information provided by GS1 US. The table shows the top 15 firms in terms of expenditure shares.

Figure C.1: Household Income over Head Age

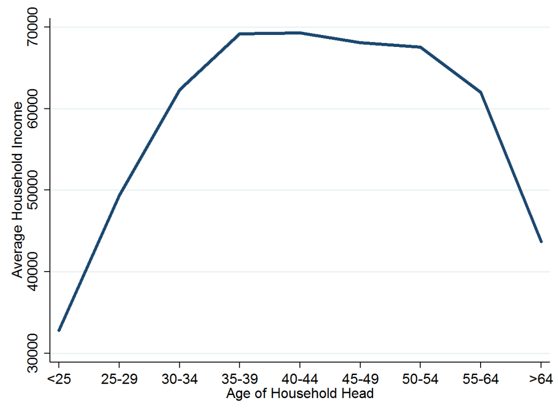

Note: This graph depicts annual household income over the age of household head. The y-axis depicts average nominal annual household income, and the x-axis depicts household head's age. Source: Nielsen Homescan 2004-2010

Figure C.2: Laspeyres Index vs CPI-U Food at home

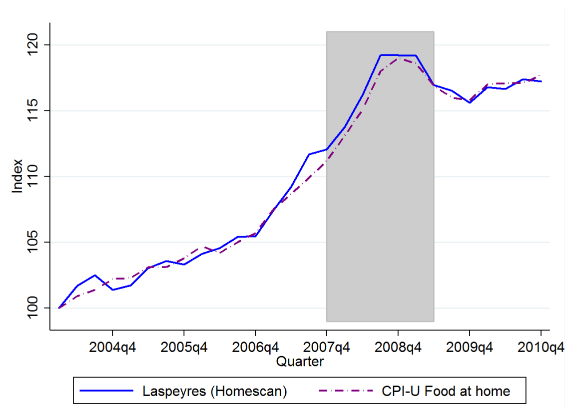

Note: This graph plots a Laspeyres index from the first quarter of 2004 to the last quarter of 2010. The graph also plots the CPI for all urban consumers: Food at home. From the Nielsen Homescan data, we pick product groups that match the food at home category in the CPI-U construction (dry grocery, frozen foods, dairy, deli, packaged meat, and fresh produce.). The shaded areas indicate periods designated as recessions by the National Bureau of Economic Research. The source is the Nielsen Homescan.

Figure C.3: Difference between the Price Index of the Highest vs. the Lowest Quartile after Controlling for their Shopping Behavior Cumulatively

Panel A: Benchmark

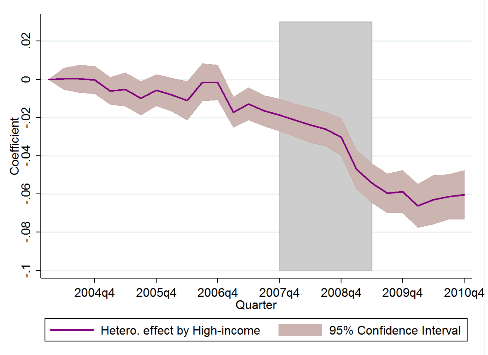

Panel B: Add Shopping Intensity

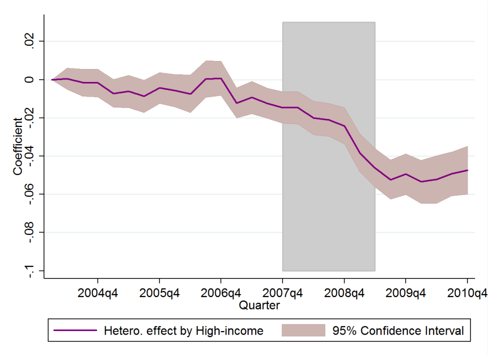

Panel C: Add Store Control

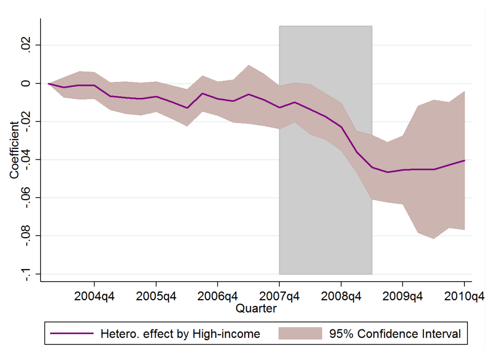

Panel D: Add UPC Control

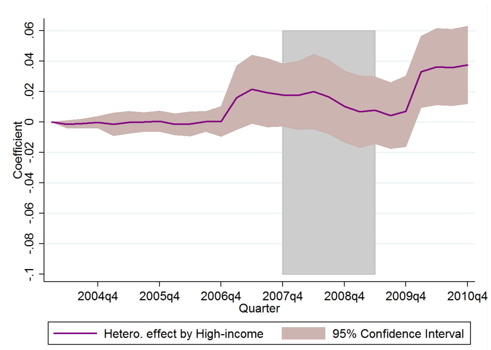

Note: This graph plots the coefficients of the interaction term in equation (6). The y-axis indicates the mean percent deviations of the prices paid by the highest quartile with respect to the lowest quartile of the income distribution. Panel A is the benchmark specification in equation (6). Panel B reports the results after controlling for shopping intensity: the fraction of goods bought on sale, the fraction of goods bought with coupons, the share of expenditures on generics, the number of trips in a given quarter, and the mean number of units purchased in a given quarter and product category. Panel C depicts the results after controlling for the shopping intensity and the store fixed effects. Panel D reports the results obtained after introducing the UPC fixed effects in addition to the shopping intensity and the store fixed effect controls. The standard errors are clustered at the product category level. The pink area indicates the 95% confidence interval. The shaded areas indicate the periods designated as recessions by the National Bureau of Economic Research. The source is the Nielsen Homescan. The number of observations is 103,940,344 and the R-squareds are 0.81 (Panel A), 0.82 (Panel B), 0.86 (Panel C), and 0.98 (Panel D) respectively.

Figure C.4: Example of a Firm Prefix

Note: This figure shows examples of a 6-digit and a 9-digit firm prefix. The source is the GS1-US website (http://www.gs1-us.info/company-prefix).

C.2 Price Discrimination - Robustness

In Section 3.6.2 we provide evidence that the price of high-priced products does not decrease more than that of low-priced products when economic conditions worsen. Here we expand the analysis and show that our results are robust to different definitions of high-priced goods and to alternative methods of calculating the posted price's inflation.

We previously defined high-priced products as those whose price per volume is on average 25% above the benchmark price (the average price per volume at the store-week level within a category). We now use less restrictive cutoffs. Table C.4 reports our results under a 15% cutoff. The results are similar when we use a 20% cutoff. The table shows that our conclusions on the differential effect of economics on high-priced and low-priced products are essentially unaltered under different cutoffs.

We also explore the sensitivity of our results to alternative methods. Gagnon, Lopez-Salido, and Sockin (2015) argue that, relative to the method used by Coibion, Gorodnichenko, and Hong (2015b), the sensitivity of the posted price's inflation to economic conditions increases if one imputes missing price observations, excludes clearance sales, and/or varies the degree of truncation of the price changes. This sensitivity could in principle affect our results if these choice of methods affect the price observations of the high-priced products differently relative to the prices of low-priced products. In what follows we show that our conclusions are robust to considering these alternatives.

In order to compare our results with those of Coibion, Gorodnichenko, and Hong (2015b) and Gagnon, Lopez-Salido, and Sockin (2015), we reproduce our analysis by using the IRI Symphony data that both these studies use. The data set is extremely similar to the Nielsen RMS (Retail Measurement Services) scanner data set. It includes information about the number of units sold and total revenues at the week-store-barcode level.[1] The two main differences between the two data sets are that the IRI data cover a larger sample of years, 2001 to 2011, but have substantially smaller geographic and product category coverage. The IRI Symphony data cover only 31 product categories over 50 Metropolitan Statistical Areas (MSA) whereas the RMS considers more than 1,000 product categories in over 200 MSAs.[2]

Because the method used in Section 3.6.2 to construct the posted price's inflation follows Coibion, Gorodnichenko, and Hong (2015b), here we describe in detail the adjustments made by Gagnon, Lopez-Salido, and Sockin (2015). First, instead of censoring price changes they discard the whole price history for a given UPC in a given store if any of the price changes for this good are greater than 666% in absolute value (annualized). Second, if the price of an item is missing, they impute the last price observation by using the inflation rate in a given city and product category. And, third, they remove clearance sales by dropping the last quarter of observations from every price's history. We present our results using both methods in Table C.5.[3] Consistent with Gagnon, Lopez-Salido, and Sockin (2015), we find that their adjustments boost the measured sensitivity of the posted price's inflation to local slack for both high-priced and low-priced products relative to the method developed by Coibion, Gorodnichenko, and Hong (2015b). Nonetheless, under both methods, the prices of high-priced products do not decrease more relative to low-priced products when local economic conditions deteriorate. That is, the coefficient of the interaction term remains positive. Together, the different definitions we construct and the various alternative methods we explore strongly support our conclusions that the difference in the cost of living of high- and low-income consumers do not come from the retailer's price discrimination in favor of consumers of high-priced products.

1. The data set contains around 2.4 billion transactions that represents roughly 15% of the household spending in the Consumer Expenditure Survey (Stroebel and Vavra (2016)). All estimates and analysis in this paper are based on data provided by SymphonyIRI Group, Inc. The conclusions are the responsibility of the authors and not SymphonyIRI Group, Inc. The data set is discussed in more detail in Bronnenberg, Kruger, and Mela (2008).

2. The product categories in the IRI Symphony data are 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 comprises 25-30% of the sales in consumer packaged goods in grocery stores.

3. For consistency, Table C.5 shows our results for the same sample period of 2004 to 2010. Our results are unchanged if we consider all years available in the IRI symphony data.

Table C.4: Posted Inflation and Local Demand Shocks by Product Type - 15% Cutoff

	Equal (1)	Market (2)	Common (3)	Equal (4)	Market (5)	Common (6)
UR	-0.152***	-0.224***	-0.219***	-0.044***	-0.066***	-0.070***
	(0.014)	(0.019)	(0.018)	(0.007)	(0.009)	(0.008)
UR×High Priced	0.067***	0.134***	0.142***	0.067***	0.133***	0.141***
	(0.006)	(0.008)	(0.008)	(0.006)	(0.008)	(0.008)
Market ×Category×Type FE				✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓

Notes: We use market×category×quality-type fixed effects and month fixed effects. The number of observations is 10,909,370. Driscoll and Kraay (1998) standard errors are in parenthesis. The ***, **, and * denote significance at 0.01, 0.05, and 0.10 levels respectively.

Table C.5: Posted Inflation and Local Demand Shocks by Product Type

	Coibion, Gorodnichenko, and Hong (2015b)			Gagnon, Lopez-Salido, and Sockin (2015)		
	Equal	Market	Common	Equal	Market	Common
	(1)	(2)	(3)	(4)	(5)	(6)
UR	-0.112**	-0.155**	-0.149**	-0.345***	-0.448***	-0.504***
	(0.043)	(0.066)	(0.069)	(0.064)	(0.068)	(0.088)
UR×High Priced	0.029	0.073	0.080	0.205*	0.271***	0.291***
	(0.030)	(0.050)	(0.057)	(0.105)	(0.088)	(0.103)
Market ×Category×Type FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓

Notes: We use market×category×quality-type fixed effects and month fixed effects. The number of observations is 239,465. Driscoll and Kraay (1998) standard errors are in parenthesis. The ***, **, and * denote significance at 0.01, 0.05, and 0.10 levels respectively.