How do Firms Grow?
The Life Cycle of Products Matters*

David Argente [†] FRB Minneapolis

Munseob Lee[‡] UC San Diego

Sara Moreira[§] Northwestern University

February 2019

Abstract

We exploit detailed product- and firm-level data for the consumer goods industry and we document that: (1) the sales of individual products decline at a steady pace throughout most of their life cycle, (2) the evolution of product sales over the life cycle results mostly from changes in quantities sold and not prices, and (3) the decline in sales is steeper for products of firms and sectors with higher rates of product introduction. These patterns are robust across barcodes, brands, and very heterogeneous types of products, and contrast with the increasing size of firms over their life cycle. To understand the mechanism behind the decline in product sales and reconcile these seemingly opposing dynamics in the product and firm life cycles, we structurally estimate a model of oligopolistic competition between heterogeneous multi-product firms. Our estimates indicate that the decline in sales of products results from decreases in their quality relative to other options available to consumers, including new products of the same firm. We show that firms must introduce new products to grow, otherwise their portfolio becomes obsolete as their rivals introduce new products of their own. By introducing new products, however, firms accelerate the decline in sales of their own existing products.

JEL Classification Numbers: L1, L2, O3, O4.
Keywords: Product Life Cycle; Firm Dynamics; Multi-Product Firms.

*We are grateful to Fernando Alvarez, Anmol Bhandari, Paco Buera, Ariel Burstein, Meghan Busse, Peter J. Klenow, Jonathan Eaton, Marcela Eslava, Doireann Fitzgerald, John Haltiwanger, Hugo Hopenhayn, Thomas Hubbard, Erik Hurst, Benjamin F. Jones, Greg Kaplan, Marti Mestieri, Todd Schoellman, Robert Shimer, Joseph Vavra, Venky Venkateswaran and seminar participants at NYU, UCLA, Duke, UC San Diego, Northwestern, FRB Minneapolis, FRB NY, FRB Atlanta, GoetheU, UMannheim, WashU, UConn, LSE, Statistics Norway, EIEF, Bank of Spain, Barcelona GSE, NBER SI, SED, and LACEA. We would like to thank Olga Denislamova and Monica Toledo for excellent research assistance. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. Empirical results are calculated based on data from The Nielsen Company (US), LLC and marketing databases provided by the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researchers and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

[†]Email: david.argenteamaya@mpls.frb.org. Address: 90 Hennepin Ave, Minneapolis, MN 55401.

[‡]Email: munseoblee@ucsd.edu. Address: 9500 Gilman Drive #0519, La Jolla, CA 92093.

[§]Email: sara.moreira@kellogg.northwestern.edu. Address: 2211 Campus Drive, Evanston, IL 60208.

1 Introduction

Product introduction is a key vehicle through which innovation translates into economic growth. Most firms offer multiple varieties and frequently reshuffle their portfolios to add newer products. What is the role of product introduction in explaining firm growth? Existing theories of economic growth posit that new products could serve very different purposes: some represent the introduction of new varieties, others position the firm to compete against its rivals, and some others enhance or update the firms' own existing products. Each of these types of innovation has very different implications not only for innovation policy but also for firm growth. New products that impact the sales of older products produced by the firm may generate less incremental revenue than new products that impact the sales of products produced by competitors. This suggests that quantifying the margins of substitution between products over their life cycle is key for an accurate portrayal of the role that product introduction plays in explaining firm growth.

In this paper we provide large scale empirical evidence on the relationship between the life cycle of products and the introduction of new products, and how this relationship shapes firm growth. We build a detailed data set using scanner data from the consumer goods industry from 2006 to 2015 to measure the firms' and products' characteristics. When we use this data set to study the product life cycle, a striking pattern emerges – sales of products decline at a steady pace throughout most of their life cycle. We interpret this evidence through the lens of a parsimonious model of oligopolistic competition with heterogeneous multi-product firms, where products have heterogenous quality and cost. Our estimates indicate that, as existing products mature, they lose quality relative to other products available to consumers, which in turn results in losses of sales. Furthermore, we quantify that 55 percent of those losses accrue to products of competitors, while 40 percent are lost to other products of the same firm. We use the model to build counterfactuals that show that firms must introduce new products to grow, otherwise their products become obsolete as their rivals introduce new products of their own. By introducing new products, however, firms accelerate the decline in sales of their own existing products.

We begin our analysis by developing a simple statistical framework to measure how products of different ages contribute to the sales growth of a firm. We find that, on average, firms grow two percent a year, with new products accounting for a positive contribution of five percent and existing products accounting for a negative contribution of three percent. We find systematic evidence that existing products generate fewer sales every year, not only because they are more likely to be discontinued, but also because their average sales decline as they become older. Thus, while firms' sales grow throughout their life cycle, their individual product do not. This result is surprising because firms could potentially expand the markets and the customer base of their existing products without modifying their current product portfolio. To better illustrate these patterns, we plot the evolution of sales of one of the largest firms in our data set in Figure 1. The distinctive feature of this example is that the smooth and moderate growth in sales of this firm conceals massive product reallocation, as shown by the large share of revenue generated by new products and the large declines in the sales generated by older ones.

After documenting these empirical patterns, we then provide more systematic evidence on the evolution of product sales over their life cycles. The concept of the product life cycle has been relevant in many fields for decades, yet few studies describe it empirically and most have done so for specific products (e.g., Copeland and Shapiro (2016) for personal computers, Gowrisankaran and Rysman (2012) for the digital camcorder, and Abe, Ito, Munakata, Ohyama and Shinozaki (2016) for home electrical appliances and digital consumer electronics). Our data set allows us to measure products at their finest level of aggregation – barcodes consisting of a 12-digit number called the Universal Product Code (UPC). Although we present results that aggregate barcodes into broader groups of products (e.g. brands), our benchmark estimates use UPC codes since changes in any attribute of a good (e.g. forms, sizes, package, formula) result in the introduction of a new barcode. This feature allows us to study the life cycle of a product while holding all their physical attributes constant. We estimate the life cycle profile using econometric specifications that decompose longitudinal outcomes into age-period-cohort effects, while accounting for heterogeneity in the consumer goods industry. This methodology is well suited to measure age patterns while accounting for cohort specific effects and aggregate factors affecting sales.

Despite the substantial heterogeneity in the products we study, there is a common pattern that emerges when we study their life cycle – sales of products decline at a steady pace throughout most of their life cycle. This is persistent across products of all durations as even long lasting products have moderate increases in sales in their first year of activity and declines in sales thereafter. We also observe this pattern when we use broader definitions of products. Our findings challenge the standard view of the product's life cycle that assumes that product sales exhibit a bell-shaped evolution over time, where the market growth of a product lasts longer than its market decline (Levitt, 1965). Empirically, the framework has been applied to describe the evolution of aggregate sales of specific sectors. Our results show that by keeping the product attributes constant, most products virtually bypass the introduction and growth stages.

In order to uncover the potential mechanisms behind the decline in sales over the product life cycle, we study other outcomes such as quantity and prices, as well as the heterogeneity in the evolution of sales across different types of products. We find that while both quantities and prices decline, the evolution of sales is mostly explained by the evolution of quantities, which suggests a potentially important role for product-specific demand-side factors, such as changes in a product's quality relative to the qualities of other products available to consumers. We also find that the decline in sales is stronger for products from firms and sectors with higher rates of product introduction, which is consistent with our hypothesis that new products by the firm and its competitors affect the sales of existing products.

Guided by these empirical patterns, we develop a model of heterogeneous multi-product firms. The model shares several common features with Atkeson and Burstein (2008), Edmond, Midrigan and Xu (2015), and in particular with Hottman, Redding and Weinstein (2016). In the model, firms compete under conditions of oligopolistic competition, which introduces variable markups to our setup. The model features preferences for variety (horizontal differentiation), vertical differentiation, and a CES-nested structure to allow for greater substitution among products within a firm than among the products produced by different firms. This feature is important because it allows new products to potentially affect the sales of a firm's own products more than products by other firms. Following Hottman, Redding and Weinstein (2016), we structurally estimate the model using price and sales data for the products and firms.

We characterize products in terms of a variable, the quality-to-cost ratio, that captures the idiosyncratic time-varying relative quality and marginal cost. Our results show that the quality-to-cost of products changes systematically over their life cycle. As products mature in the market, they lose quality relative to other products available to consumers and become obsolete, which is consistent with diminishing product appeal and with consumer preferences for up-to-date varieties. We develop a framework that allows us to decompose the decline in product sales into that which is attributable to the product's quality-to-cost ratio and cannibalization as well as the competition and strategic interactions among firms. We quantify that about 55 percent of the dispersion in the sales of products is explained by changes in the product's quality-to-cost ratio, which result from losses of sales to competitors, and 40 percent is explained by changes in cannibalization effects, which result from losses to other products of the firm. The distinction between business stealing and cannibalization effects is key, because when we aggregate the sales of individual products to the level of the firm, we observe that cannibalization effects imply that an additional dollar of revenue accruing from a new product translates to less than a dollar of additional revenue for the firm.

Our framework also allows us to relate firms' total sales to the characteristics of their portfolio of products to study the implications of the product's life cycle for the firm's growth.

We decompose total sales into the effects of the number of products (scope), their average quality-to-cost ratio, their dispersion in quality-to-cost, the firms' markup, and measures of the size of the sector. We find that firm sales grow over the life cycle mostly because the average quality-to-cost ratio of their products increases, followed by increases in the number of product offerings and increases in the quality-to-cost dispersion. The contribution of markups and other components play a smaller role.

Our results indicate that firms grow by increasing the average quality-to-cost of their portfolio but the relative quality-to-cost of their individual products declines over time. These seemingly contradictory findings indicate that product introduction is a necessary condition for firm growth. New products not only increase the scope of firms but also increase their average quality-to-cost ratio which, in the absence of product introductions, declines over time as firms' existing products increasingly compete against newer and better versions of products. Thus, competition and, specifically, innovation by competing firms affects incentives to introduce new products and is an important force driving firms to keep improving their portfolio of products.

To illustrate and quantify the contribution of new products to firms' sales, we compute equilibrium outcomes in a counterfactual economy where firms are temporarily prevented from introducing new products and we compare its outcomes to those observed in the data. We find that, on average, the contribution of new products to firms' growth is approximately 50 percent smaller than their share of sales in the firms' portfolio.[1] Introducing new products increases firm sales by about 16 percent among very young firms and 2 percent among mature firms. This increase in sales is largely explained by improvements in the average quality-to-cost ratio of the firms rather than by increases in scope. As a result, adding new products is not a sufficient condition for growth since they must have a sufficiently high quality-to-cost ratio at entry to compensate for the reduction in the quality-to-cost ratio of existing products. Furthermore, the contribution of new products is smaller for older firms, not only because they represent a smaller fraction of the firms' total products, but also because they are less likely to increase their quality-to-cost ratio and are more likely to cannibalize older products.

Lastly, we show that the contribution of new products to firms' sales depends on how fast a product's quality-to-cost ratio depreciates relative to other products in the market. To show this, we conduct a counterfactual in which we estimate the contribution of new products holding constant the quality-to-cost ratio of the firms. We find that the effect of new products on sales in this case is significantly smaller; they are needed less when the quality-to-cost ratio

[1]For the firm in Figure 1, the revenue share of products introduced after 2006 was 80 percent in 2015, while their contribution to firm growth was 60 percent. The difference is given by the fact that older products would sell more in 2015 if no new products were introduced.

of the firm does not decrease over time. Therefore, by improving their portfolios with better and newer products, competing firms make their rivals' products obsolete and motivates them to constantly add new products to their portfolio. These product introductions, however, cannibalize the sales of existing products and contribute to decrease their sales even further.

The rest of the paper is organized as follows: Section 1.1 discusses our contribution to the literature. Section 2 presents the data and describes our procedure to identify the firm that supplies each product. In section 3, we develop a statistical model of growth as a function of the ages of the firms's products. In section 4, we show the evolution of revenue, prices, and quantities over the life cycles of products. We also show these life cycles have steeper decline for products in sectors with large product introduction. Section 5 presents a structural model and the counterfactuals used to understand the mechanisms behind our empirical findings. Section 6 concludes.

1.1 Related Literature

The product life cycle is a well known concept in marketing, management, and other business-related fields. It has been used to describe a bell-shaped evolution in sales; the process of a product's birth, evolution, maturity, and death. The shape of a product's life cycle, its determinants, and its implications for sales, prices, and profits have been discussed extensively in marketing and management. In economics, the literature on the product life cycle can be traced back to Vernon (1966) with important contributions by Brockhoff (1967), Polli and Cook (1969), and Greenstein and Wade (1998), among others. Hofer (1975) refers to a product's life cycle as "the most fundamental variable in determining an appropriate business strategy." An important distinction between our paper and most of these contributions is that we use the concept of product life cycle as a latent concept. We study the margins affecting sales and prices rather than treating these outcomes as simple time-dependent relationships. We show that the product life cycle profile depends on firms' decisions regarding the introduction of new products.

Due to data limitations, the empirical literature using variation on the product portfolio of firms has focused on studying a limited set of products and has used broader definitions of products (e.g., entire industries and product lines). For example, Bernard, Redding and Schott (2010) studies the extent of product switching within firms by using production classification codes (5-digit SIC codes). At this level of aggregation, the life cycle of a product line may resemble the life cycle of the firm. In the US, for example, 95% of firms have a single establishment and 75% of manufacturing firms have a single plant. Yet, single-plant firms are likely to produce several products. In our data, 74% of the firms are multi-product firms and the average firm supplies 12 products. The advantage of focusing on the life cycle profile of barcodes is that, by keeping their physical attributes constant over time, we rule out the possibility that their quality is changing. Thus, we provide direct evidence on how changes in the set of other products available to consumers induces changes in the demand for existing products. Moreover, our data cover very different types of goods, including non-durables and semi-durables, which allows us to compare their life cycle profiles.

Our paper is related to the growing literature studying firm dynamics and innovation (e.g. Klette and Kortum (2004); Akcigit and Kerr (2018)). Garcia-Macia, Hsieh and Klenow (2018), for instance, infer the sources of growth from the patterns of job creation and job destruction. Instead, we provide direct empirical evidence of the impact of new products and how differences in performance across newer and older products shape firm growth. Our framework allows us to capture the substitution of products occurring within firms and between firms that maps to theories of creative destruction, where innovating firms can improve existing products made by other firms (e.g., Grossman and Helpman (1991)) or improve their own products (e.g. Lucas and Moll (2014)). Atkeson and Burstein (2018) discuss the distinct policy implications of the different theories.

Our paper contributes to the growing body of work that emphasizes the firms' endogenous determination of the number of products. Examples of this work can be seen in the business cycle literature (Bilbiie, Ghironi and Melitz, 2012; Minniti and Turino, 2013) or in international trade (Arkolakis and Muendler, 2010; Mayer, Melitz and Ottaviano, 2014; Timoshenko, 2015). Most of this literature, however, has assumed that firms are homogeneous – there are no differences in the amount of products they produce or in the rate at which they introduce them to the market. Our empirical results emphasize the importance of allowing for these margins to understand the growth of firms. As a result, our framework complements the literature by studying the impact of new products in the context of multi-product firms with heterogeneous products.

The literature that studies the life cycle of the firm is vast both in international trade and macroeconomics. Recently the work by Hsieh and Klenow (2014) has highlighted the importance of this margin to explain productivity differences across countries. Atkeson and Kehoe (2005) show that firms' life cycle is driven by the accumulation of plant-specific organizational capital. In this interpretation, establishments grow with age as they invest in new technologies, develop new markets, and produce a wider array of higher quality products. Foster, Haltiwanger and Syverson (2016) show that even in commodity-like markets, the rising demand for a plant's products as it becomes older largely drives a firm's growth. Our paper complements this work by decomposing firm growth into the part attributable to new products and the contribution of existing ones to the firm. The importance of product entry on firm growth emphasized in our growth decomposition is also useful to reconcile the growing literature on product reallocation (Broda and Weinstein, 2010; Argente, Lee and Moreira, 2018) with the parallel literature on innovation (Klette and Kortum, 2004; Lentz and Mortensen, 2008; Akcigit and Kerr, 2018; Acemoglu, Akcigit, Alp, Bloom and Kerr, 2017).

Our model builds on recent papers that have examined the determinants of heterogeneity among firms and have characterized the contribution of several margins (e.g., Melitz (2003); Manova and Zhang (2012); Feenstra and Romalis (2014); Fitzgerald, Haller and Yedid-Levi (2017); Eslava and Haltiwanger (2017)). In particular, we draw extensively from Hottman, Redding and Weinstein (2016) who find that differences in firm quality explain most of the observed variance in firm size. Our analysis differs from Hottman, Redding and Weinstein (2016) in that we use our model to study the margins affecting the sales of products over their life cycles and we make an explicit connection between the margins that affect the life cycle of products and the sources of firm heterogeneity. This connection allow us to draw two important conclusions. First, existing products lose quality relative to other products available to consumers over their life cycles, which, in turn, has a negative effect on the firms' quality. And, second, new products are key to explain heterogeneity in firm growth due to their contribution to increase both the scope of the firms and their average quality.

2 Data

2.1 Defining a Product

We use barcodes as our baseline definition of products. A barcode is a 12-digit Universal Product Code (UPC) consisting of 12 numerical digits that is uniquely assigned to each specific good available in stores. UPCs were created to allow retail outlets to determine prices and inventory accurately and improve transactions along the supply chain distribution (Basker and Simcoe, 2017). This offers a unique opportunity for economists to identify products at its finest level of disaggregation.

Defining products as barcodes has some important advantages. First, barcodes are by design unique to every product: changes in any attribute of a good (e.g. forms, sizes, package, formula) result in a new barcode.[2] The most common alternative is to define goods by industry classifications that can potentially aggregate very heterogeneous barcodes. Changes

[2]Firms have strong incentives not to reuse barcodes. Assigning more than one product to a single barcode can interfere with a store's inventory system and pricing policy; it is rare that a meaningful quality change occurs without resulting in a UPC change. Nonetheless, a possible concern 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.

in industry-level outcomes can result from changes in the composition of quality within those industries. In fact, our data show that large firms typically sell hundreds of different products within narrowly defined categories. By using barcodes, we ensure that we observe the exact same product at different points in time and thus changes in performance cannot result from changes in the attributes of the product.

Second, barcodes are so widespread that our data is likely to cover all products in the consumer goods industry (Basker and Simcoe, 2017). Producers have a strong incentive to purchase barcodes for all products that have more than a trivial amount of sales because the codes are inexpensive, and they allow sellers to access stores with scanners as well as internet sales. Further, because firms and products are included in the sample provided that a sale occurs in which we observe a wide range of products and we can explore several dimensions of heterogeneity.

We also provide results for an alternative definition of products, where we aggregate all barcodes into brands. The average brand in our data has nine different barcodes. Barcodes within a brand vary because of differences in attributes; thus, comparing a brand at different points in time will naturally reflect changes in its quality that result from the entry and exit of barcodes. Brands, however, are a good comparison benchmark and are useful for understanding some underlying mechanisms discussed throughout the paper. Brands are commonly used in research featuring the consumer goods industry (Bronnenberg, Dhar and Dubé, 2009; Bronnenberg and Dubé, 2017) both because of advertising data and the firm internal organization of the firm.

2.2 Product Data

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 data is generated by point-of-sale systems in retail stores. Each individual store reports weekly sales and the quantities of every barcode that had any sales volume during that week. We use data for the period from 2006 to 2015.

The main advantage of this data set is its size and coverage. Overall, the RMS consists of more than 100 billion unique observations at the UPC $\times$ store $\times$ week level that cover approximately \$2 trillion in sales. This volume represents about 53% of all sales in grocery stores, 55% in drug stores, 32% in mass merchandisers, 2% in convenience stores, and 1% in liquor stores. A key distinctive feature of this database is that the collection points include more than 40,000 distinct stores from around 90 retail chains in 371 MSAs and 2,500 counties. As a result, the data provide good coverage of the universe of products and of the full portfolio of firms in this sector.[3]

The data covers a wide range of products both in terms of type (e.g. from non-durables such as milk to semi-durables like printers) and in terms of revenue share. The original data consist of more than one million distinct products identified by UPC, organized into a hierarchical structure. Each UPC is classified into one of the 1070 product modules, that are organized into 104 product groups, that are then grouped into 10 major departments.[4] For example, a 31-ounce bag of Tide Pods (UPC 037000930389) is mapped to product module "Detergent-Packaged" in product group "Detergent", which belongs to the "Non-Food Grocery" department. Throughout the paper we refer to product groups as sectors.

Our baseline data set combines all sales at the national and quarterly level, although we also conduct some exercises at the store quarterly level when studying staggered entry. For each product u in quarter t, we define sales Y_{ut} as the total sales across all stores and weeks in the quarter. Likewise, quantity y_{ut} is defined as the total quantities sold across all stores and weeks in the quarter. Price P_{ut} is defined by the ratio of revenue to quantity, which is equivalent to the quantity weighted average price.[5]

We identify the state of the life cycle of a product through information on its age. Scanner data sets do not directly measure the age of a product. We infer the age by observing the timing of its initial transaction in the data set. More specifically, we define entry as the first quarter of sales of a product and exit as the quarter after we last observe it being sold. We cannot determine entry and exit for some products. For products that are already active in the first two quarters of the sample (2006q1 and 2006q2), we classify them as left censored. These products can include those created just before 2006 or very established products. Likewise, we classify products that have transactions in the last two quarters of the sample (2015q3 and 2015q4) as right censored. For those, we cannot determine exit and thus cannot measure how long they last in the market. Moreover, in order to minimize concerns of potential measurement error in the calculation of a product's entry and exit, our baseline sample includes the balanced set of stores and products with at least one transaction per quarter after entering, while excluding private label products and departments that are not representative.[6]

[3]In comparison to other scanner data sets collected at the store level, the RMS covers a much wider range of products and stores. 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. Argente, Lee and Moreira (2018) presents a full comparison of the different scanner data sets, including IRI Symphony and the Kilts-Nielsen Homescan.

[4]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).

[5]We use the weight and the volume of the product to compute unit values.

[6]Our estimates of products' entries and exits might be affected by the entries and exits of stores in the

In our baseline sample we have around 650,000 products that are organized into 92 distinct groups and in a total of 904 modules. There is substantial heterogeneity in the market share and duration across production. Table 1 provides summary statistics for products in terms of their duration, and revenue. Products in the 75th percentile of revenue generate 50 times more revenue than a product in the 25th percentile. Some of that heterogeneity in sales can result from using data from very different types of products. Table 2 shows that the dispersion in sales occurs even in narrowly defined sectors. In a given quarter, the average interquartile range across products in the same sector is 4.6 log points, and the 95-5 ratio is 10.9 log points.

Table 1 also gives the summary statistics by the type of censoring. We divide products into four categories: (i) complete, (ii) right-censored, (iii) left-censored, and (iv) both left- and right-censored. For example, our data set identifies a 12 count of 12oz cans of regular Coca-Cola as a "both left- and right-censored" because the product already existed in the beginning of the period under analysis and survived the entire period. By contrast, a 12 oz bottle of Coca-Cola BlāK (a coffee-flavored soft drink) is "left-censored" because it was available in the beginning of our sample period but was discontinued during the years covered by our data. We observe product entry within the data when the product is uncensored or right censored. To investigate a product's life cycle for certain periods after the entry, we mainly use those products. For two-thirds of the products we can measure age. Among those, for more than 50% we can also identify exit, and thus we are able to measure both age and duration. The remaining third of the products were already active in the first periods of our data set, and thus we cannot measure age. Among those, we can identify exit for 60%.[7] If we account for the average sales that the different products generate, then the products for which we can determine age account for close to 60% of sales. The summary statistics also show that products have short durations. The median product lasts between 12 and 16 quarters.

sample. Therefore, we consider only a balanced sample of stores during our sample period. 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 little sales. 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. And, finally, we exclude Alcohol and General Merchandise because these are the departments for which the coverage in our data is smaller and less likely to be representative. These exclusions do not play any role in our results.

[7] On average there are almost 250,000 products active per quarter. Every quarter, uncensored products represent 21.9% of the total, and right censored represent 26.2%.

2.3 Firm Data

We study the implications of the life cycle of products for the growth of firms. In order to do that, we link firms and products with information obtained from GS1 US, the single official source of UPCs. This link allows us to perform the analysis at the parent company level as opposed to at the level of the manufacturing firm. 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.[8] By linking firms to products, we can characterize their portfolio. With this data set, we can identify the revenue, price, and quantity of each product in a firm's portfolio and aggregate them to the level of the firm. We mostly focus on measures of size (number of products and total revenue), product introduction (frequency, number, and revenue), and the product's age. We also use this data set to identify the entry and exit of firms. The product-firm baseline data set allows us to study how size and product introduction change over a firm's life cycle.

Table 3 describes the characteristics of firms by censoring. Among the 23,000 firms in the sample period, we can measure age for about 9,000, and the remaining 14,000 are firms that were born before 2006. As expected, young firms are smaller both in revenue and number of products and are less diversified. Older firms have on average 27 products in their portfolios, from 4 different product modules, and 2 different product groups. Young firms have on average 2 to 3 products in a single product group. Given the richness of the data, throughout the paper, we present evidence for young firms (age 1 to 4 years old) and older ones, which are those firms born before 2006.

3 Firms' Growth: New versus Older Products

Existing research has established that firms tend to start small and they grow as they age (e.g. Dunne, Roberts and Samuelson (1989), Hsieh and Klenow (2014)). Figure 2 shows the estimated total sales over the life cycle of the firms in our baseline data set. Conditional on surviving, young firms more than double their size in their first four years of activity. Older firms are more than two times larger than new firms. The growth paths of the firms in our sample are similar to those of the representative firm in the US economy, that is, firms grow fast in their initial years of activity, but their growth rates subsequently decline as they become older. The unique feature of our data set, is that we have information to decompose each firm's sales into the sales of its products. Young firms necessarily have new products, but older firms may have new and older products. For example, Figure 1 plots the evolution

[8] Argente, Lee and Moreira (2018) provide more details on these data and their advantages compared to alternative sources.

of sales for one of the largest firms in the consumer goods industry. The black solid line represents total sales. This firm has smooth and moderate growth in sales, but the revenue of each cohort of products is declining at a fast pace. For example, products created up to 2006 (green dashed line) account for about 90% of the revenue in 2007 but less than 20% of the sales in 2015. The decline in sales is pervasive across all cohorts of products and is accompanied by a continuum of new products. For this firm, the total revenue generated by new products is larger than the decline in revenue of older ones.

Motivated by the patterns in Figure 1, we next evaluate if the strong decline in the revenue of older products is common across the firms in our data set. We develop a decomposition of the firm's sales into the sales of products by their ages and apply it to all our firms. After arranging the different components, we can write sales growth (in percentage) for each firm i as

$$\Delta_{i,t} \quad = \quad n^E_{i,t} \times \tilde{r}^E_{i,t} \quad + \quad \sum_{u=0}^{t-1} \omega_{u,i,t-1} G_{u,i,t} \quad , \tag{1}$$

where $n^E_{i,t}$ is the entry rate, $\tilde{r}^E_{i,t}$ is the average revenue of the new products relative to the average revenue of the products of the firm, $G_{u,i,t}$ is the growth rate of the revenue of products of age u conditional on survival, and $\omega_{u,i,t-1}$ is the weight of each cohort.[9] We refer to the first part as the *entrants component* and to the second part as the *product life cycle component*. The entrants component is the product of the entry rate and the revenue of the new products relative to the average revenue of the products of the firm, which is a measure of the quantity and the quality of each new product. The product life cycle component quantifies the contribution of older products to sales growth.

Table 4 (column 1) presents the average contribution of each of these components after weighting each firm by their total revenue share in our data. Firms grow on average 2% a year among the pooled sample of firms during our sample period. This positive revenue is entirely explained by the introduction of new products whose contribution to revenue growth is about 5% a year. Like in our example above, the growth rate of existing products is negative for the average firm in our sample, which indicates that as products age their revenue declines on average. At the same time, firms are adding products to their portfolio at a rate of 14% a year that generate 34% of the average revenue of existing products.

Columns (2) to (4) of Table 4 show the average annual growth rates across the firms' ages.[10] Consistent with Figure 2, firms grow fast in their initial years of activity but their growth rates subsequently decline as they become older. Firms in our sample grew more than 83% annually in their first year of activity but only approximately 5% from ages 2 to

[9] A full derivation of equation 1 can be found in the Appendix C.1.

[10] See Appendix C.2 for robustness checks: (i) using brands and (ii) excluding top and bottom 5% of firms.

4. Among firms created before 2006, this value is even smaller, about 1%. This decline in growth rates results from both a decline in the product life cycle effect and the contribution of new products.[11] With the exception of firms in their first year of activity, the life cycle effect is negative and becomes particularly negative among older firms. The entrants effect is positive and declines with age at a slower pace, from about 24% among very young firms to 5% among firms that existed before 2006. Most of this decline comes from the reduction in product entry rates as firms get older since the entrants' relative quality remains fairly constant over the firm's life cycle.[12]

The results from the statistical decomposition indicate that the patterns in Figure 1 are very representative of the patterns of most of the older firms in our data. The results also indicate that the decline in the total revenue of each cohort results from both the extensive margin, due to the exit of products, and from the intensive margin, due to the decline in revenue among surviving products. The only exception occurs with very young products (e.g., those that are particularly prominent in young firms' portfolios). This decline in sales, which is conditional on survival, is surprising as one would expect that firms' growth could also result from the increase in revenue among the surviving products. This is certainly the case if we were presenting the decomposition of aggregate revenue across firms, instead of products of different cohorts. In that case, the size of different cohorts of firms increases due to the growth of surviving firms. In the next Section we develop an econometric approach to study patterns in product sales with their age, and reconcile them with the results in this decomposition.

The statistical decomposition quantifies the revenue contribution of new and existing products. There is, however, important channels between these two components, as the competitors actions and the past performance of firms' existing products may affect the introduction of new products and, in turn, the introduction of new products may affect the future performance of those existing products. In Section 5 we show that their contributions are not independent and we quantify the impact of new products when we properly account for these channels.

[11] We analyze the relation between firm's age and the revenues generated by its new products in more detail in Appendix C.3. We show that: i) total revenues accruing from new products increase in the first year of activity of the firm and remain mostly constant over time, and ii) the share of revenue of new products in the total revenue of the firm declines with age.

[12] In Appendix C.3 we show both that the total number of new products introduced every quarter declines approximately 20% in the first 16 quarters of activity and that the ratio between the average revenue of new products and the average revenue of older products initially declines during the first two years but remains constant thereafter.

4 The Life Cycle of Products

4.1 Measurement

We use our data and information about the products' ages to study common patterns that characterize their life cycles by estimating how the products' performance evolves as a function of their age. To properly isolate the effect of age we must account for the fact that we observe products in different quarters and want to capture the effect of age irrespective of the particular period at which a product is observed. Likewise, we want to control for the fact that for otherwise comparable products, they might behave differently depending on the timing of their introduction. In order to address such issues, we estimate age effects by implementing age-period-cohort models. These specifications allow us to estimate the evolution of revenue, quantities, and prices since the time of the product's introduction, while accounting for cohort-specific differences in outcomes and any calendar effects specific to the moment when we measure the outcomes (e.g., business cycles affecting all products). This methodology is commonly used in the literature on individuals' life cycle consumption and income dynamics and was recently used for firms in Moreira (2017).

In the baseline specification, we estimate the outcome of interest (Y) of product u observed at time t as a function of age (a), interacted period and product category (mt), and cohort (c) fixed effects:[13]

$$\ln Y_{j,t} = \alpha + \sum_{a=2}^{A} \beta_a D_a + \lambda_{mt} + \theta_c + u_{j,t} \tag{2}$$

We control for heterogeneity by allowing the time fixed effect to be specific to the product modules.

We are interested in the coefficients β_a that capture the average evolution of the aging process of the product relative to the level of the outcome in the first full quarter of activity. Further, the evolution of the outcomes of products over the life cycle is affected by selection. The main sample includes all barcodes from their full first quarter in the market until the quarter before they exit. Table 1 presents the statistics on the duration of a product in the market. There are substantial differences in the duration of barcodes, which means that the estimated effects are conditional on survival if we use all active observations regardless of their duration. Because the products that are discontinued early are different from those that are discontinued later, the unconditional estimated effects will be different from the

[13]Because there is an exact linear relation between the three effects, we normalize the cohort effect as suggested in Deaton (1997). The normalization averages the cohort effects to zero over the sample period and orthogonalizes the cohort trends such that all growth is attributed to age and time effects. In the appendix we check the robustness of this normalization by considering alternative specifications.

conditional estimated effects. In order to ensure that the estimation results are not sensitive to the selection bias that result from the inclusion of short-lived products, we repeat the empirical analysis on products that survive past the median age for products. We also check the robustness to alternative criteria. Most of our results are computed with a balanced sample of products that survived at least 16 quarters. In Appendix D.1, we present the results with semi-balanced and unbalanced samples.

We consider an alternative specification that explicitly accounts for selection. To obtain information on the nature of selection, while also isolating true dynamics, we examine how the initial outcomes forecast survival and also document the dynamics conditioning on the ex-post duration. In the alternative specification, we estimate the outcome of interest of product j observed at time t as follows:

$$\ln Y_{j,t} = \alpha + \sum_{d=2}^{D} \sum_{a=1}^{d} \gamma_{ad} D_{ad} + S_{j,t} + \lambda_{mt} + \theta_c + u_{j,t} \tag{3}$$

where λ_{mt} are the interacted product category and time fixed effects, θ_c are Deaton's normalized cohort effects, $\gamma_{ad} D_{ad}$ are dummies for age interacted with duration, and $S_{j,t}$ is a dummy for observations censored observations.[14] Exponentiated, linear combinations of the estimates γ_{ad} allow us to characterize both variations at the initial level of the outcome variable, and the evolution with age over the different durations. This approach has some similarities with the approach of Fitzgerald, Haller and Yedid-Levi (2017) who work with exporter dynamics and Altonji and Shakotko (1987) who deals with selection in estimating the effect of job tenure on wages.

4.2 Stylized Facts on the Product Life Cycle

Fact 1: Sales of products decline throughout most of their life cycle

We estimate equation (2) by using the level of quarterly revenue as the dependent variable (in logs) for products that were active for at least 16 quarters. Table 5 reports the estimated age fixed effects. The results presented in column 1 show that the coefficients for the age fixed effects are positive and significant for the first periods and negative and significant for later periods. We plot the estimated coefficients in Figure 3. The revenue of products mostly declines with age except for the first four to five quarters. By the end of the fourth year of activity, revenue declines by more than 50%.

[14] We include a indicator variable for both right and left censored to help identify the fixed effects.

Brands as products Our results are robust to using brand as a broader definition of a product. Figure 4 shows that the life cycle of revenue is similar for brands, as provided by the Nielsen data and under a broader definition we have constructed after combining similar brands (e.g. Tide Pods and Tide Heavy Duty combined into Tide). Although the increase in sales during the first quarters of the life cycle is larger, revenue declines throughout most of the life cycle of a brand.

Sales decline for both short and long lasting products We also estimate equation (3) for products with different durations. Figure 5 shows that for short-lived products, revenue declines throughout their entire life cycle until exit and the negative growth rates are larger. Moreover, sales at exit is several orders of magnitude lower than the level observed at entry. These results mean that products of a short duration, have larger decreases in revenue, even several periods prior to exit. The figure also shows that the level of revenue at entry, which in the figure is normalized relative to products that lasted only one period, is a good predictor of how long a products stays in the market.

Left-censored products Figure 6 shows that revenue also declines for products that lasted the entire period covered in our data (i.e., products that entered the market before 2006 and exit after 2015). Given that we cannot determine the age of these products, we plot the overall path of revenue after controlling for module and quarter fixed effects. This is relevant because these products are very likely to be successful products given their duration in the market.

Nondurables and semi-durables Revenue declines over most of the life cycles of both nondurable goods and semi-durables. The research on marketing and industrial organization has documented this fact for specific durable goods such as semi-conductors (Byrne, Oliner and Sichel, 2018), home appliances (Abe, Ito, Munakata, Ohyama and Shinozaki, 2016) personal computers (Copeland and Shapiro, 2016), and digital camcorders (Gowrisankaran and Rysman, 2012). These studies have argued that a combination of process innovation along with the introduction of more up-to-date products drive down the revenue of existing products in the durable goods markets. However, our data cover a broader set of products that vary in their durability. We approximate the durability of each category, using the Nielsen Consumer Panel Data, by counting the average number of shopping trips made by households in a given year to purchase products in each category.[15] We call categories with few trips per year durable categories. Examples of durable categories are sun exposure trackers (1.00),

[15] A full description of these data can be found in Appendix Table B.I.

bathroom scales (1.03), and printers (1.03), where the average number of shopping trips per year is in parenthesis. Examples of nondurable categories are refrigerated milk (23.61), cigarettes (19.19), and fresh bread (18.76). Figure 7 shows revenue over the life cycle by durability. Although revenue declines faster for durable categories, we also find a large drop in revenue for nondurable goods that indicates the patterns that we identify are robust to a broader set of goods than those found in the literature.

Alternative specifications We consider some alternative approaches to evaluate the robustness of our estimates to different assumptions on the data generating process. These specifications are also useful in shedding light on the potential reasons behind the decline in the revenue of products over time. Figure 8 shows that the life cycle of sales is similar after controlling for UPC effects, firm×time effects, or firm×module×time effects. The graph shows that our findings are not sensitive to our baseline specification. Furthermore, several potential mechanisms, such as competition or changes in the firms' quality, can be ruled out as the main force behind the decline in revenue given that our findings do not substantially change after controlling for them.

Other robustness checks In Appendix D we show that the decline in revenue over the product life cycle is also present when we split products along dimensions such as the degree of novelty they bring to the market that we measure by the number of new attributes a product has at the time of its introduction, or by its ranking at entry that we approximate by whether they are in the top decile of the revenue distribution of new products.

Fact 2: Both prices and quantities decline, but the decline in prices is quantitatively small relative to the decline in quantities

An advantage of our data is that we observe prices and quantities. Hence, we are able to determine the contribution of both to the decline in revenue over the product's life cycle. As described in Section 5, the variation in sales can be attributed to cost or to demand factors. Both have different implications for product sales; the marginal cost affects revenue through prices, while the firm's demand affects revenue conditional on prices. Figure 9 shows the estimated life cycle profile of the price of products (those that last at least 16 quarters). The age fixed effects show that prices decline 2% a year on average. The decline in price happens at a fairly constant pace and, by the end of the fourth year of activity, the price is almost 8% lower than the price at entry. Because our empirical specification controls for aggregate effects (e.g., inflation) specific to particular modules of products, the decline happens on top of the average fluctuations in prices. Contrary to the evolution of revenue, product prices decline at a constant pace later in the life cycle. The figure also shows the percent decline in quantities sold after the product was introduced. Similar to revenue, quantities decline after the first year of activity. Our results show that 16 quarters after introduction, quantities had declined by more than 50% relative to the initial quantities sold when the products were launched. When comparing the magnitudes of the decline in quantities and in prices, we conclude that the decline in sales, comes mostly from the decline in quantities.[16]

Fact 3: The introduction of new varieties steepens the decline in sales

What factors are behind the decline in sales over the product's life cycle? Two possibilities suggested in the literature are increases in competition within a category and cannibalization by new products. The products in categories with high entry rates could experience faster declines in sales if consumers substitute them with newer varieties. Similarly, multi-product firms can introduce new varieties to the market that not only affect the sales of other firms but also the sales of their own products. In this subsection we provide reduced form evidence that both of these channels affect the sales of the products over their life cycle. We use this evidence to inform our modeling choices presented in section 5.

We begin by splitting the sectors in our data into three groups according to their average quarterly entry rate (weighted by revenue). Panel (a) in figure 10 shows substantial differences in the life cycle of products sold in sectors with low entry rates; the revenue of products starts declining almost a year after that of products sold in high-entry sectors, and declines at a slower pace.[17] This indicates that, as existing products mature, they lose sales because of the increasing business stealing from new products. Therefore, the frequent adoption of higher quality varieties introduced by other firms increases the rate at which products become obsolete.

Similarly, a firm's introduction of new product varieties affects not only the sales of competitors but also those of its own existing varieties. To approximate these cannibalization effects, we compute for each firm the revenue-weighted share of new products in each sector-quarter. We then split firms into three groups according to the level of this variable and estimate equation 2 that allows for category fixed effects. This procedure allows us to compare, within the same sector, the life cycle of products supplied by firms that are very active in introducing new products against those that are not. Panel (b) shows the results. The lines show the estimated effect of age for each of the three groups and shows that the sales of products that belong to firms that are very active in introducing new products start declining

[16] Appendix D presents the same robustness checks presented in the previous subsection.

[17] Prices also decline at a slower pace in low-entry rate sectors. This is consistent with the literature that studies the life cycle of prices for durable goods.

substantially sooner. This suggests that when firms introduce new products, they affect the sales of their own existing products.

5 Structural Approach

In this section, we study the mechanisms behind the decline of sales over the product's life cycle by building a model of oligopolistic competition among a finite number of heterogeneous multi-product firms. The model allows for changes in quality and costs over the life cycle of products and explicitly models the effects of future product introduction on sales of existing products.

We start this section with a description of our economy, followed by a characterization of the static equilibrium. Our equilibrium yields conditions similar to those developed by Hottman, Redding and Weinstein (2016). We use these conditions to estimate the parameters governing the per period quality and cost and employ the estimated parameters to isolate different margins affecting the product life cycle. We conclude by studying counterfactuals that allow us to better understand the role that the product life cycle plays in shaping firm growth.

5.1 Model

5.1.1 Demand for products

The economy consists of households with nested preferences for products. Preferences are given by $E_0 \sum_{t=0}^{\infty} \beta^t u_t$, where β is the discount factor, and u_t is the per period utility function. The utility is defined as $\log\left(c_t^{\mu}(1-l_t)^{1-\mu}\right)$ where c_t is the aggregate consumption, and l_t is the total working hours. The aggregate consumption is a Cobb-Douglas function over the consumption index of a continuum of sectors $j \in [0,1]$ that is defined as

$$\ln(c_t) = \int_0^1 \varphi_{jt} \ln(c_{jt}) dj \tag{4}$$

where φ_{jt} is the time-varying quality of sector j in period t. In each sector there is a finite number of active firms indexed by $i \in \Theta_{jt}$, and the consumption of each sector is given by the following CES specification over the consumption index of each firm

$$c_{jt} = \left[\sum_{i \in \Theta jt} c_{ijt}^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}} \tag{5}$$

where η is the elasticity of substitution across firms (assumed to be greater than one), and c_{ijt} is the real consumption index of each firm. In this economy, firms can have a single product or multiple ones. Each firm has a finite set of active products indexed by $u \in \Omega_{ijt}$, and the real consumption of each firm is given by:

$$c_{ijt} = \left[\sum_{u \in \Omega ijt} (\gamma_{uijt} c_{uijt})^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \tag{6}$$

where γ_{uijt} is the time-varying quality of product u in period t, c_{uijt} is the quantity of product u in period t, and σ (assumed to be finite) is the elasticity of substitution across products within the same firm.

In this nested specification the upper level relates to the consumption of a large number of sectors using a Cobb-Douglas. This means that the firm problem is effectively separable across sectors and that multi-sector firms do not have incentives to decide prices in one sector to influence the behavior in another sector. This assumption implies that a firm is assumed small relative to the aggregate economy even if it is large in a sector. In line with this assumption, we refer to a firm-sector combination simply as a firm. Using a Cobb-Douglas functional form, the consumers' first order conditions (given the price index of each sector and wages) are such that the share of expenditure for each sector at time s_{jt} is given by φ_{jt}. Thus, the implicit demand of each sector is $c_{jt} = \varphi_{jt}(\frac{P_{jt}}{P_t})^{-1} c_t$ and the price-index of the aggregate consumption is given by:

$$\ln(P_t) = \int_0^1 \varphi_{jt} \ln(P_{jt}) dj. \tag{7}$$

The second level relates to the consumption of each sector to the consumption index of a finite number of firms. A structure with a finite number of firms allows for strategic interactions among firms that produce in the same sector. This feature allows firms to be large relative to product groups, which in turn means that they can internalize their effects on the consumption and price index of a sector. The CES structure generates a theoretical price index for the final consumption of each sector that is given by:

$$P_{jt} = \left[\sum_{i \in \Theta_{jt}} (P_{ijt})^{1-\eta} \right]^{\frac{1}{1-\eta}}. \tag{8}$$

Each firm faces the demand function $C_{ijt} = (\frac{P_{ijt}}{P_{jt}})^{-\eta} c_{jt}$ and as a result, the market share of each firm is given by:

$$S_{ijt} \equiv \frac{P_{ijt}c_{ijt}}{P_{jt}c_{jt}} \equiv \frac{C_{ijt}}{C_{jt}} = \frac{(P_{ijt})^{1-\eta}}{\sum_{k\in\Theta_{jt}}(P_{kjt})^{1-\eta}} \tag{9}$$

The third level of the nested specification relates the consumption index of the firm to products. This stage is critical to the model because we want to allow the elasticity of substitution for products supplied within the same firm (σ) to be different than that for products supplied by competing firms (η). The CES structure embedded in this stage implies that the price index of each firm is given by:

$$P_{ijt} = \left[\sum_{u\in\Omega_{ijt}} \left(\frac{P_{uijt}}{\gamma_{uijt}}\right)^{1-\sigma}\right]^{\frac{1}{1-\sigma}}. \tag{10}$$

Each product faces the demand function $c_{uijt} = (\gamma_{uijt})^{\sigma-1}(\frac{P_{uijt}}{P_{ijt}})^{-\sigma}c_{ijt}$, and thus the share of sales of product u within firm i is given by:

$$S_{uijt} \equiv \frac{P_{uijt}c_{uijt}}{P_{ijt}c_{ijt}} \equiv \frac{C_{uijt}}{C_{ijt}} = \frac{\left(\frac{P_{uijt}}{\gamma_{uijt}}\right)^{1-\sigma}}{\sum_{k\in\Omega ijt}\left(\frac{P_{kijt}}{\gamma_{kijt}}\right)^{1-\sigma}}. \tag{11}$$

Another important component of our specification is that products have heterogeneous effects on the firm that depends on their quality; equation 11 shows that if prices are held constant, a product with higher quality will have a larger market share. Using the consumer first order conditions, we can write the demand for product u as:

$$c_{uijt} = (\gamma_{uijt})^{\sigma-1}(P_{uijt})^{-\sigma}(P_{ijt})^{\sigma-\eta}(P_{jt})^{\eta-1}C_{jt}. \tag{12}$$

5.1.2 Supply of Products

In our economy, firms have technologies that are separable across products and that use labor as the only input. The production function of each product u is given by:

$$y_{uijt} = (a_{uijt})^{\frac{-1}{1+\delta}}(L_{uijt})^{\frac{1}{1+\delta}}$$

where $(a_{uijt})^{\frac{-1}{1+\delta}}$ is the time-varying product-specific productivity, $\frac{1}{1+\delta}$ is the economies of a scale parameter, and $\delta > 0$ is the decreasing returns to scale. The marginal cost of product u is given by:

$$z_{uijt} = a_{uijt}(1+\delta)(y_{uijt})^{\delta} \tag{13}$$

In addition to the production costs, firms pay an operational cost in units of labor to produce each of its products h_{ijt}. We allow for scope economies, by introducing an operational cost in units of labor, f_{jt}, to produce in sector j.

Firms are engaged in imperfect competition. Each firm $i \in \Theta_{jt}$ has a set of potential products B_{uijt}. The firm must decide what to produce ($\Omega_{uijt} \in B_{uijt}$) and at what price conditional on the exogenously given quality γ_{uijt} and marginal cost shifter a_{uijt}. In our baseline specification, firms choose prices under Bertrand competition. Each firm plays a static game of price competition. Specifically, each firm chooses the set of products and their prices to maximize profits.

$$\max_{\{I_{kijt}, P_{kijt}\} \in B_{ijt}} \quad \Pi_{ijt} = \sum_{k \in B_{ijt}} \Big[\big[P_{kijt} - (1+\delta) z_{kijt} \big] y_{uijt} - h_{ijt} \Big] I_{kijt} - f_{jt} \tag{14}$$

where I_{kijt} is an indicator that equals one if the firm decides to supply the product k. The number of products produced in period t is defined as $N^U_{ijt} = \sum_{k \in B_{ijt}} I_{kijt}$, and we refer to M_{jt} as the number of all products in sector j in period t. Although the firm's problem is separable across sectors, the product's problem is not separable within firm. As a result, the decisions made for one product affect the other products of the firm, and thus the firm needs to internalize that.

From the first order conditions with respect to the price of product u we obtain:

$$\sigma\big(1 - \frac{1}{\mu_{uijt}}\big) = \sum_{k \in B_{ijt}} \Big[I_{kijt} S_{kijt} \big(1 - \frac{1}{\mu_{kijt}}\big) \Big] \Big[(\eta - 1) S_{ijt} + (\sigma - \eta) \Big]$$

where the markup for each product is given by $\mu_{uijt} = \frac{p_{uijt}}{z_{uijt}}$. This condition makes it clear that the markups are the same across the different products of the firm, $\mu_{uijt} = \mu_{ijt}$. The intuition for this result is that the firm internalizes that it is the monopoly supplier of its real output within the sector. Solving the equation, we find that:

$$\mu_{ijt} = \frac{\eta - (\eta - 1) S_{ijt}}{\eta - (\eta - 1) S_{ijt} - 1} \tag{15}$$

and the equilibrium price is given by

$$P_{uijt} = \mu_{ijt} (1+\delta) a_{uijt} \left(y_{uijt} \right)^{\delta} \tag{16}$$

This equation represents the supply function of product u. This is the same pricing rule as in Hottman, Redding and Weinstein (2016), and it is in line with Atkeson and Burstein (2008) and Edmond, Midrigan and Xu (2015). Each product has a variable elasticity of demand that decreases with the expenditure share of the firm. For the firms with a small market share, the markup is close to η, which is the assumption under monopolistic competition.

5.1.3 Equilibrium

Given the entry/exit decisions of firms and conditional on quality and costs, our economy is in equilibrium under the following conditions:

1. Firms play a static game of price competition and maximize profits. The equilibrium conditions are such that the set of active products $\{I_{kijt}\} \in B_{ijt}$ solves the problem (14) and the supply of products u is given by equation (16), $u \in \Omega_{ijt} \in \Theta_{jt}$, and equation (15).

2. Consumers maximize utility by making labor and consumption decisions. The optimum conditions imply sector shares are given by φ_{jt}, demand for product u is determined by equation 12, $u \in \Omega_{ijt} \in \Theta_{jt}$, and aggregate, sectoral, and the firms' price indexes are determined by equations (7), (8), and (10), respectively.

3. Product and labor markets clear. The clearing condition in the product market is $c_{uijt} = y_{uijt}$, $u \in \Omega_{ijt} \in \Theta_{jt}$. In the labor market, the supply for labor equals the demand for labor given by $\int_0^1 \Big(\sum_{\Theta_{jt}} \big(f_{jt} + \sum_{\Omega_{ijt}} L_{uijt} + h_{ijt}\big)\Big) dj$.

The equilibrium conditions indicate that by holding the set of firms and products constant, the only decision the firm makes is to set the prices of its products, taking as given the prices of other firms. All firms behave similarly; so, in equilibrium, the prices of other firms must correspond to their optimal prices. Thus, computing the equilibrium is equivalent to finding the fixed point.

The economy becomes substantially more complex when firms also decide to introduce new products. In this case, firms must not only decide the prices of products but also which products to produce. This scenario could lead to multiple equilibria. We follow Atkeson and Burstein (2008) and focus on the equilibria in which firms sequentially decide which products to introduce, starting with the highest profitability product to the lowest. Using the optimality conditions of the firm's problem we can write the profit of each product as

$$\pi_{uijt} = \frac{\mu_{ijt}(1+\delta) - 1}{\mu_{ijt}(1+\delta)} Y_{uijt} - h_{ijt} \tag{17}$$

where $Y_{uijt} \equiv y_{uijt} P_{uijt}$ is the revenue of product u. Furthermore, the only component that is product-specific is revenue, which implies that the highest revenue products are the most profitable. Moreover, using equations (12) and (16), the ratio of any two products offered by the same firm can be defined as:

$$\frac{Y_{kijt}}{Y_{uijt}} = \left(\frac{\gamma_{kijt}}{\gamma_{uijt}}\right)^{\frac{(\sigma-1)(1+\delta)}{1+\sigma\delta}} \left(\frac{a_{kijt}}{a_{uijt}}\right)^{\frac{-(\sigma-1)}{1+\sigma\delta}}. \tag{18}$$

Thus, our firms rank products according to their profit impact based only on exogenous parameters.[18] This feature implies that firms sequentially introduce products until they reach the product whose additional contribution to the total profit of the firm is lower than its fixed cost of introduction. In this case, in equilibrium, no firm has an incentive to change its set of products.

5.2 Parameter Estimation

Our structural estimation procedure closely follows that of Hottman, Redding and Weinstein (2016). The model assumes a continuum of product groups, where each firm is small relative to the aggregate expenditure in that product group category. We approximate the model to the data by mapping sectors to 83 product group categories and we allow for heterogeneity in preferences and technology parameters across sectors by estimating sector-specific elasticities. The estimation has two steps. In the first step, we use product shares S_{uijt}, product prices P_{uijt}, and firm shares S_{ijt} to estimate the model parameters for each sector σ_j, η_j, and δ_j. In the second step, we recover γ_{uijt} and a_{uijt}.

5.2.1 Sector Elasticities

We use product data (product shares and prices) to recover the elasticity of substitution of products within firms σ_j and the technology parameter δ_j. Our estimation relies on the assumption that demand and supply shocks at the product level are uncorrelated once we control for firm-time specific effects, that is, we assume that the double-differenced product quality $\Delta_t^k \gamma_{uijt}$ and cost shifters $\Delta_t^k a_{uijt}$ are orthogonal. More specifically, we follow Broda and Weinstein (2006), which developed a generalization of Feenstra (1994), and use the double differenced demand and supply equations 12 and 16:

$$\Delta_t^k \ln S_{uijt} = (1 - \sigma_j)\Delta_t^k \ln P_{uijt} + \Delta_t^k \gamma_{uijt}$$
$$\Delta_t^k \ln P_{uijt} = \frac{\delta_j}{1 + \delta_j} \Delta_t^k \ln S_{uijt} + \Delta_t^k a_{uijt}$$

[18] In Atkeson and Burstein (2008), this maps to productivity differences.

Thus, we define a set of moment conditions by assuming that the perceived quality shocks are orthogonal to shocks to the production technology of specific goods.[19]

$$E_T\left[\Delta_t^k \gamma_{uijt} \Delta_t^k a_{uijt}\right] = 0$$

Next, we follow a similar procedure to the one presented earlier and use market shares to estimate the elasticities of substitution across firms η_j. The difference is that in this case, we assume that demand and supply shocks at the firm level are uncorrelated because we control for sector-time specific effects.[20]

Table 6 presents the summary of the distribution of the estimated parameters across sectors. The median elasticity of substitution across products within a firm is 13.8, and the average is 22.9. The elasticity of substitution of products varies substantially across sectors, the interquartile range is 24.6. These numbers indicate that on average the products supplied by the same firm are imperfect substitutes but highly substitutable. The median elasticity of substitution across firms is 7.1, the average is 10.1, and the interquartile range is 7.3, all of which are substantially lower than the corresponding statistics for the elasticity of substitution across products. In terms of magnitudes, our estimates for the elasticities are larger than the elasticities typically estimated with trade data, but within the levels of the elasticities that use similar data. For example, Broda and Weinstein (2010) reports a median elasticity of substitution within brands of 11.5 and a median elasticity across brands of 7.5. Hottman, Redding and Weinstein (2016) reports a median elasticity of substitution within a firm of 6.9 and 3.9 across firms.[21]

Our results are consistent with our model assumption that the elasticity of substitution across products within a firm is larger than the elasticity of substitution between firms. For most sectors we can statistically reject the hypothesis that these elasticities are equal. But, we also observe substantial heterogeneity across sectors in the difference between the two parameters. In 42 out of 83 sectors, the difference does not exceed 1, whereas it is greater than 5 in 33 out of 83 sectors. The difference in these elasticities is critical to establish the differential impact of product introduction by the firm or by its competitors in explaining

[19] We use the entire sample period for the estimation. For each firm we use as reference the products that are the firm's best sellers. For a small number of groups, the GMM approach produces imaginary estimates and estimates with the wrong sign, so we follow Broda and Weinstein (2006) and use a grid search over the space of the elasticities. The derivation of the moment conditions and details on the estimation can be found in Appendix H.2.

[20] In Appendix J.3, we also implement a robustness exercise in which we follow the estimation procedure in Fally and Faber (2017) and use data at the region level to estimate the elasticities.

[21] The differences seem to result from the fact that we use a different data set with wider coverage. Each firm in our data has on average 23.5% more products in their portfolios compared to the Nielsen Homescan.

the sales patterns of products over their life cycle.

To better understand the heterogeneity in the products and the elasticities of substitution across sectors, we correlate the estimates with information about the entry rates and durability of products for each sector (as defined in Section 4). Both elasticities as well as the difference between them are statistically larger in sectors with higher entry rates of products and firms, and in sectors whose products are relatively more durable.

5.2.2 Product Demand and Cost Shifters

In the second step of our procedure, we recover γ_{uijt} and a_{uijt}. Given the estimated parameters $\{\hat{\sigma}_j, \hat{\eta}_j, \hat{\delta}_j\}$ and observed data $\{S_{uijt}, P_{uijt}, S_{uijt}, N_{ijt}, M_{jt}\}$, the quality and cost shifters are determined as the structural residuals that ensure that the model equilibrium replicates the observed data. In the case of the cost/productivity shifter a_{uijt}, we use equation 16 and obtain:

$$a_{uijt} = {P_{uijt}}^{1+\delta_j} \left(S_{uijt}\right)^{-\delta_j} \left(\frac{\eta_j - (\eta_j - 1)S_{ijt}}{\eta_j - (\eta_j - 1)S_{ijt} - 1}\right)^{-1} (1+\delta_j)^{-1} \tag{19}$$

The identification of product quality requires an extra normalization because the utility function is homogenous of degree 1 in product quality. Given the CES structure of our model, we normalize by setting the geometric mean of product quality within each sector and time period equal to one:

$$\gamma_{jt}^{Norm} = \Bigg(\prod_{u\in\Omega_{ijt}\in\Theta_{jt}} \gamma_{uijt}\Bigg)^{\frac{1}{M_{jt}}} = 1.$$

The key advantage of this normalization is that we can compare product quality within the firm and across firms within a sector and time period. It also allows us to compute a measure of the firm's overall quality that we define as the geometric average of the quality of the products in its portfolio. Our normalization procedure implies that changes in γ_{uijt} over time result from changes in the quality of a product or firm relative to changes in the geometric mean of the qualities of all products within their respective sector and time period. Using this normalization, we obtain the product specific quality as:[22]

$$\gamma_{uijt} = \left(\frac{S_{uijt} \times S_{ijt}}{\prod_{u,j}(S_{uijt} \times S_{ijt})^{\frac{1}{M_{jt}}}}\right)^{\frac{1}{\sigma_j - 1}} \left(\frac{S_{ijt}}{\prod_{u,j}(S_{ijt})^{\frac{1}{M_{jt}}}}\right)^{\frac{\sigma_j - \eta_j}{(1-\eta_j)(1-\sigma_j)}} \left(\frac{P_{uijt}}{\prod_{u,j}(P_{uijt})^{\frac{1}{M_{jt}}}}\right). \tag{20}$$

Table 6 presents the summary statistics for the estimates. The quality shifter, γ_{uijt}, is

[22] See Appendix H.1 for the derivation of this result.

approximately log normal because of the normalizations that we implemented. The statistics for the cost shifter parameter, a_{uijt}, indicate that its mean value is below zero and that its distribution is very dispersed. The mapping between data and estimated parameters does not impose any functional form assumption on the distributions of the structural residuals of the product demand and supply, and thus allows us to study their properties. We compare the cross-sectional correlation between the estimated quality shifters γ_{uijt}, the estimated cost shifters a_{uijt}, and the observed sales and prices of products for our baseline sample and we find that the correlation between the demand and cost shifters is 0.1, meaning that it is more costly for firms to produce higher quality products. Moreover, the correlation between between sales and quality is 0.35, and sales and cost is -0.24.

5.3 Margins Affecting the Product Life Cycle

Next, we quantify the sources that shape the products' sales over their life cycle. We use the demand function in equation (12), the firm price index, and the equilibrium pricing rule to decompose sales of products into six terms:

$$\begin{aligned} \ln Y_{uijt} = &(\eta-1)\ln\frac{\gamma_{uijt}}{z_{uijt}} - (\eta-1)\ln\mu_{ijt} \\ &-\frac{\sigma-\eta}{\sigma-1}\ln N_{ijt} - \frac{\sigma-\eta}{\sigma-1}\ln\left[\frac{1}{N_{ijt}}\sum_{k\in\Omega_{jit}}\left(\frac{\frac{\gamma_{kijt}}{z_{kijt}}}{\frac{\gamma_{uijt}}{z_{uijt}}}\right)^{\sigma-1}\right] \\ &+(\eta-1)\ln P_{jt} + \ln Y_{jt} \end{aligned} \tag{21}$$

We begin our analysis by considering the case of a product that is produced by a single-product firm. In this case, sales are represented by:

$$\begin{aligned} \ln Y^{SP}_{uijt} = &(\eta-1)\ln\frac{\gamma_{uijt}}{z_{uijt}} - (\eta-1)\ln\mu_{ijt} \\ &+(\eta-1)\ln P_{jt} + \ln Y_{jt} \end{aligned}$$

The first term captures the effect of a product's *quality-to-cost ratio*, $\frac{\gamma_{uijt}}{z_{uijt}}$, on its sales. Holding everything else constant, a 1% increase in this ratio increases sales by $(\eta-1)$%. A higher elasticity of substitution between firms, η, makes this effect stronger because sales react more to an increase in the product's quality-to-cost ratio if consumers can easily substitute products sold by other firms. Throughout our analysis, the quality-to-cost ratio of products assumes a key role. It allows us to characterize and compare products, while summarizing into a single variable the two sources of exogenous heterogeneity across products (the demand and cost shifters). The second component captures the effect of an increase in a product's *markup*, μ_{ijt}, on its sales. An increase in the markup has an impact that is symmetric to that of the product's quality-to-cost ratio. The third term captures the sector's price index, P_{jt}, and summarizes the actions of competing firms. A 1% increase in the sector's price index increases the sales of the firm by $(\eta - 1)$%. The last term captures the size of the sector. Product sales increase one-to-one with the aggregate expenditures in that sector (due to homogeneity of degree 1), which is a one-to-one function of the sector's exogenous shares and the aggregate expenditure of the economy. This exercise shows that in the case of single-product firms, the critical parameter is the elasticity of substitution across firms.

Next, we consider the general case of a multi-product firm. Equation (21) includes two extra terms that capture the effect of product substitution within a firm. The first captures the effect of an increase in the portfolio of products of the firm, N_{ijt}, on sales that we call *size cannibalization*. A 1% increase in the number of products affects the sales of an existing product by $-\frac{\sigma-\eta}{\sigma-1}$%. The effect is negative when products are more substitutable within firms than across firms ($\sigma > \eta$). The second component captures the effect of changes in the quality-to-cost ratios of products in the firm's portfolio. We refer to this effect as *quality-to-cost cannibalization*. This source of product-specific cannibalization varies with changes in the allocation of products (captured by changes in the dispersion of the quality-to-cost ratios of the products of the firm) and with changes in the product's own quality-to-cost ratio.[23] For example, this component captures the negative effect on a product's sales that accrues from the firm's introduction of new products with higher quality-to-cost ratios. The term also captures the effect of a change in a product's quality-to-cost ratio on its own sales resulting from consumers adjusting their consumption in favor of other products of the firm. Overall, a decrease in the relative quality-to-cost ratio of a product while keeping the quality of the other products of the firm constant has two effects: a $(\eta - 1)$% decrease because consumers substitute the product with products from other firms, and $(\sigma - \eta)$% decrease in sales because consumers will substitute the products with other products of the firm. One way to interpret the first effect is to consider the case in which the increase in the quality of product u is followed by a proportional increase in the quality of all other products of the firm. This effect then is captured by the component quality-to-cost ratio, while second effect is captured by the component quality-to-cost cannibalization.

These six components capture the most direct margins through which products differ in sales. In the context of our model, some of these margins, such as the product quality and the cost shifters, are independent of the decisions of firms. Firms endogenously determine other

[23]The numerator $\left(\frac{1}{N_{ijt}} \sum_{k \in \Omega_{jit}} \left(\frac{\gamma_{kijt}}{z_{kijt}}\right)^{\sigma-1}\right)^{\frac{1}{\sigma-1}}$ is the power mean of the quality-to-cost ratio, which is a measure of entropy that captures the dispersion in products' quality-to-cost ratios.

margins, such as the scale and markups. Moreover, among the six components, the quality-to-cost ratio and the quality-to-cost cannibalization effects are product-specific, the markup and size cannibalization are common across all products of the firm, and the price and sector size effects are common across all firms within a sector. Table I.I in the appendix shows the cross-sectional correlation between sales of products and each of these components. A product with higher sales has higher quality-to-cost ratio, higher markup, higher size cannibalization, and lower quality-to-cost cannibalization. Moreover, the quality-to-cost effect and the quality-to-cost cannibalization are only weakly correlated, which suggests that they are driven by different sources of variation.

In Section 4, we find that the sales of a product decline throughout most of its life cycle. We document such effects by estimating age fixed effects after controlling for cohort- and sector-time fixed effects. Here, we build on that decomposition to examine whether the negative estimated age effects arise from changes in the effects of quality-to-cost, markups, size cannibalization, or quality-to-cost cannibalization. The other two components in the decomposition of product sales, sector size and price indices, cannot explain the declining estimated age fixed effects because they are absorbed by the cohort- and sector-time fixed effects. We identify and quantify the role of these four components in shaping the evolution of revenue by replicating the econometric specification used for sales and estimate equation 2 for each of them.

Figure 11 plots the estimated age fixed effects of each component over the life cycle of the product. The decline in sales over the life cycle has two main drivers: changes in the quality-to-cost ratio and changes in the quality-to-cost cannibalization. In the first year of activity, revenue increases 0.35 log points relative to entry, while the effects of the quality-to-cost of the product and quality-to-cost cannibalization account for 0.26 and 0.09 log points of such increase respectively. Similarly, in the following three years, the declines of 0.73 and 0.34 log points in these two components account for most of the decline in total revenue (1.35 log points).

Our estimates further indicate that the changes in the quality-to-cost ratio over a product's life cycle are mostly driven by changes in its relative quality. In the first year of activity, we estimate that both the increases in quality (3%) and declines in the marginal cost (2%) contribute positively to an increase in the quality-to-cost ratio (5%). After the first year, we estimate that the strong decline in the quality-to-cost ratio is largely determined by the steady decline in the estimated quality (on average 5% per year), which more than offsets the countervailing effect associated with the 2% decline per year in the cost component. This decline shows that as the product matures in the market, it loses its quality relative to other products available to consumers. Thus, it becomes obsolete and loses sales to other products sold in the market. Our interpretation is that this pattern is consistent with diminishing product appeal and consumer preferences toward up-to-date varieties.

The effect of the quality-to-cost cannibalization on product sales also plays an important role in explaining the decline in sales. This component increases during the first year a product is on the market, which means that the product is being relatively less cannibalized than it was at entry. As the product's quality-to-cost ratio declines relative to the firm's average, the product loses sales to other products of the firm. This effect results both from the decline in the product's quality and from the firm's product turnover, the latter allows the firm to increase the average quality-to-cost ratio of its product portfolio.[24]

The remaining components, in particular the markup effect, do not affect the life-cycle pattern of sales significantly. There are no substantial systematic changes in the markup of firms as the product matures in the market. The analysis also shows that the number of product offerings by the firm plays a small role in explaining the products life cycle. The size of the firm, as measured by its total number of products, increases over the products life cycle but only by a small amount so that it plays only a minor role in determining the patterns in the products life cycle.

Consistent with the patterns in Figure 11, we estimate that most of the dispersion in sales post-entry relative to entry is explained by the quality-to-cost ratio and the quality-to-cost cannibalization. We use equation 21 and a procedure similar to that in Eaton, Kortum and Kramarz (2004) and Eslava and Haltiwanger (2017)to quantify the contribution of each component to the dispersion of products' sales over their life cycle (Appendix J). The variance decomposition indicates that roughly 55% of the overall change in sales post-entry can be attributed to changes in the quality-to-cost ratio of products, and 40% is attributed to changes in the quality-to-cost cannibalization. The rest of the components play a reduced role.

5.4 Implications for Firms' Growth

5.4.1 Firm Sales Decomposition

We now examine how the total sales of a firm are affected by the characteristics of its product portfolio. We use the demand structure of the model and the equilibrium pricing rule to show that the total sales of a firm can be decomposed into six terms defined by the following equation:

[24] According to this decomposition, the effect of changes in the underlying quality-to-cost cannibalization should be larger in sectors with a greater difference between the degree of substitution within the firm and between firms. We test this assumption by exploring heterogeneity in the life cycle of products across sectors. Consistent with our model, we find that the sales of products decline at a substantially faster pace in sectors with higher differences in the elasticity of substitution within the firm relative to the elasticity of substitution between firms. Appendix I.2 presents the details of the robustness exercise.

$$\ln Y_{ijt} = (\eta-1)\ln\frac{\tilde{\gamma}_{ijt}}{\tilde{z}_{ijt}} - (\eta-1)\ln\mu_{ijt} + \left(\frac{\eta-1}{\sigma-1}\right)\ln N_{ijt} + \left(\frac{\eta-1}{\sigma-1}\right)\ln\left[\frac{1}{N_{ijt}}\sum_{k\in\Omega_{jit}}\left(\frac{\frac{\gamma_{kijt}}{z_{uijt}}}{\frac{\tilde{\gamma_{ijt}}}{\tilde{z_{ijt}}}}\right)^{\sigma-1}\right] + (\eta-1)\ln P_{jt} + \ln Y_{jt} \tag{22}$$

where $\tilde{\gamma}_{ijt}$ represents the geometric mean of the quality of the firms' products, and $\tilde{z}_{ijt}$ is the geometric mean of their marginal costs.[25]

The first component captures the effect on sales from the *average quality-to-cost* of the products of the firm. A 1% increase in the average quality-to-cost of the products of the firm increases sales by $(\eta-1)$%. This component allows us to characterize and compare firms in terms of the exogenous characteristics of the products in their portfolio. While the product's quality-to-cost ratio are determined by exogenous sources of variation (their quality and cost shifters), the quality-to-cost of a firm is mostly endogenous, since firms decide which products to keep in their portfolio. The second component captures the role played by the *markup* of the firm. As is standard in the literature, because the elasticity of substitution between firms is lower than one, an increase in the markup reduces sales.

The next two terms represent effects that are specific to multi-product firms: a product scope measure and a quality-to-cost dispersion measure. The product *scope* measures the effect on the total sales of the number of products supplied by the firm. The percentage gain in sales that accrues to a firm that adds 1% new products will be less than 1% because those new products cannibalize the sales of existing products of the firm. The *quality-to-cost dispersion* is a function of the ratio of the power mean with exponent $\sigma-1$ to the geometric mean of the quality-to-cost ratio. This is a measure of entropy that captures the dispersion in the quality-to-cost ratio of a firm's products, relative to their average quality-to-cost ratio. Firms sell more in equilibrium if the dispersion in the quality-to-cost ratio across its products increases. For example, consider two firms that have the same number of products and same average (log) quality and (log) cost, but differ in that one of them supplies all products with the same quality and cost while the other has dispersed characteristics. The second firm will be larger because it is able to supply its production bundle more cheaply as it is able to shift resources to products with higher quality-to-cost ratios. The last two components capture sector-specific effects. The sector price index summarizes the prices of competing varieties supplied by other firms. A 1% increase in the sector price index increases the sales of the firm by $(\eta-1)$%. Finally, the firm's sales rise one-to-one with the aggregate expenditure of

[25] $\tilde{\gamma}_{ijt} = \left(\prod_{k\in\Omega_{jit}}\gamma_{kijt}\right)^{\frac{1}{N_{ijt}}}$ and $\tilde{z}_{ijt} = \left(\prod_{k\in\Omega_{jit}} z_{kijt}\right)^{\frac{1}{N_{ijt}}}$.

the sector.

The sales decomposition has a similar structure to the sales decomposition for products. However, some important differences exist in the interpretation of the terms that represent the effect of the number of products and the dispersion of the quality-to-cost ratio. In the case of product sales, these terms reflect cannibalization effects: an increase in the number of product offerings or an increase in the dispersion of quality-to-cost decrease product sales. Conversely, when we consider the firm, an increase in the number of product varieties and in the dispersion of the products' quality-to-cost has positive effects on a firm's sales. This reflects the fact that, despite the role of cannibalization, firms gain from increasing the number of varieties in their portfolio so long as products within the firm are not perfect substitutes.

Section 3 documents the pace of the firms' growth over their life cycle by estimating age fixed effects after controlling for cohort- and sector-specific effects. We now study how the components of sales evolve with the age of the firm. Similar to the product-level analysis, we do so by estimating equation 2 for each of the components of sales. Figure 12 plots the results of the estimation. The increase in sales over time results from increases in the average quality-to-cost, scope, and quality-to-cost dispersion of the firm's portfolio. Our results indicate that the component of the firm's total sales that increases the most with age is the average quality-to-cost component. This result is seemingly at odds with the finding that products experience large declines in their quality-to-cost ratios over their life cycles. Because the quality-to-cost ratio changes endogenously with the portfolio of products, firms over their life cycle must make reallocation decisions that raise their average quality-to-cost and offset the decline in sales that result from the declining quality-to-cost ratios of its existing products.

Our results also show that firms continuously increase their number of product offerings. As described earlier, the effect of the number of product offerings on sales depends on the relation between the elasticity of substitution within a firm and between firms. On average, firms increase the number of products by approximately 40% in the first four years of activity, which accounts for an increase of roughly 20% in sales. As firms increase the number of products they are offering, they also increase their portfolio's quality-to-cost dispersion. This component also increases with age, and its magnitude is only slightly smaller than that of the scope of the firm. The markup has a relatively small and insignificant negative effect on the firm's sales over time; we only find an effect for this component when we restrict the analysis to the set of very large firms.

We also implement a variance decomposition and find results that are consistent with the patterns in Figure 12 (Appendix J). The variance decomposition indicates that the changes in average quality-to-cost of firms explain the largest share of the differences in growth across firms (35%). We also find that changes in the number of product offerings over the life cycle explains a large proportion of the differences in growth across firms over their life cycle (20%), followed by the quality-to-cost dispersion (account 5-7%) and markups (2-3%).

5.4.2 Role of Product Introduction

In the context of the model, firms set prices and make product reallocation decisions by selecting which products to produce every period. Such decisions mean that the firm's average quality-to-cost ratio, markup, scope, and quality-to-cost dispersion change endogenously.

Taking the exogenous quality and cost shifter of products as given, firms decide what products to produce and, thus, determine the role of these components in the evolution of total sales. A key result from the decomposition earlier is that firms grow by increasing the average quality-to-cost of their portfolio. In this section, we use that framework to show that firms increase their quality-to-cost by adding products to the portfolio with higher quality-to-cost ratios.

We quantify the effect of product entry on firms' sales by comparing our economy with an alternative economy where all firms are temporarily prevented from introducing new products. Specifically, every year and for all firms, we compare their sales in the actual equilibrium, in which the firms produce the set of products Ω_{ijt}, with sales as the counterfactual equilibrium in which firms produce Ω^*_{ijt}. The difference between the two set of products represents the new products each firm introduced in t, that is, $\Omega_{ijt} - \Omega^*_{ijt}$. We solve for the counterfactual equilibrium taking as given the product quality and cost shifters and using an iterative procedure to find a fixed point according to the description in Section 5.1.3. We hold constant the aggregate expenditure abstracting from general equilibrium effects from firm entry and exit, as well as from redistribution of profits and changes in labor input prices.[26] The counterfactual equilibrium prices and sales of each product allow us to compute the new sales level of firms and the corresponding price index.

Product introduction has a positive effect on the firm's sales. As documented in Section 3, new products account for an average of 5% of sales. This is an upper bound of their true effect. Our model allows us to quantify that almost two-thirds of those sales are obtained by cannibalizing the firms' own products, while the remaining sales are obtained by capturing sales from other firms. This finding means that capturing the effect of new products on the performance of existing products is crucial to correctly evaluate their role in the firm's growth.

Product introduction has strong redistributive effects across firms. Figure 13 plots the

[26] Section K.1 of Appendix describes the details of the algorithm.

difference in total sales in the two economies. The black line plots the total effect on sales and shows that introducing new products increases sales by about 16% among very young firms and 2% among mature firms. The bars in the figure show the effect of new products on each of the components that affect sales. New products increase growth by increasing the firm's average quality-to-cost, scope, and quality-to-cost dispersion and do not have a significant effect on the markup of their products.[27] The effect of new products on the quality-to-cost ratio is large in magnitude and fairly constant throughout the life cycle. In contrast, the effect on scope is larger among young firms. This effect happens because while the total number of products increases over the entire life cycle, the rate of new product introduction declines with age. Lastly, new products also increase the dispersion of the firms' quality-to-cost ratio. The dispersion effect also declines with the age of the firm.

The most striking result is that throughout the firm's life cycle, new products increase its size mostly by increasing the average quality-to-cost of its portfolio of products. This component captures the role of product reallocation, as firms introduce products with higher quality-to-cost than that of older vintages. Our interpretation is that product innovation allows firms to remain appealing to consumers.

5.4.3 Relationship Between the Life Cycle of Products and Firm Growth

After having quantified the positive role of new products to growth, in this subsection we show that the evolution of older products have an entirely different effect on the firm's sales. This is because the quality-to-cost ratio of older products declines over their life cycle affecting, in turn, the average quality-to-cost ratio of the firm. As a result, firms must necessarily add newer and better products to their portfolio in order to grow.

Most of the firms have products of different ages in their portfolios. Our results above indicate that after one year post-entry, products' quality-to-cost ratio starts declining, which plays an important role in explaining why the sales of products decline over their life cycle. This depreciation in the quality-to-cost ratio is likely affecting the firm's growth. We study the effect of this depreciation by developing a counterfactual exercise in which the firms' quality and cost do not change. Specifically, every year t and for all firms, solve this new equilibrium by holding constant the aggregate expenditure, the labor input prices, and the portfolio of each firm Ω_{ijt}, while imposing that the products' quality-to-cost ratios do not change over time (i.e., the demand shifter is given by $\gamma^{**}_{uijt} = \gamma_{uij,t-1}$, and the cost shifter is $a^{**}_{uijt} = a_{uij,t-1}$).

[27] The difference between the black line and the height of the bars comes from changes in the sector price index. Our model indicates that new products generate a decline of 2.5% in the aggregate price index. In Appendix K.3 we compute the welfare implications of product introduction.

Figure 14 summarizes the effect of changes in the quality and cost shifters of existing products on the firms' sales and on its components by splitting the effect by the firms' age. On average, firms are affected through their average quality-to-cost and dispersion components. There are also effects on the markup component but the magnitudes are very small. The size of the effect on the average quality-to-cost ratio of firms is not surprising given the nature of the exercise. What is relevant is the fact that the impact is quite heterogeneous depending on the age of the firm. A one-year-old firm supplies products that are less than a year old and that are still experiencing increases in their quality-to-cost ratio (Figure 11). If these ratios are held fixed before they reach their peak, as in this counterfactual, then firms on average will experience losses in sales. At the other extreme, older firms have products already experiencing a large depreciation in their quality-to-cost ratios. When the quality-to-cost ratios of these products are held constant, older firms on average experience gains in total sales.

The results of this counterfactual show that older products have a negative effect on firms' growth because the average quality-to-cost ratio declines as a result of decreases in the quality-to-cost of products over their life cycle. As a result, firms on average do not rely on their older products to generate positive growth, and introducing new products is a necessary condition for firms to grow. In fact, because firms face this negative effect from older products, they are likely to have incentives to introduce new products.

We evaluate whether this is the case by comparing the effect of new products summarized in the previous subsection with the effect of new products in the alternative scenario where products' quality and cost do not change over time. We solve for the this new equilibrium holding constant the aggregate expenditure and labor input prices, while setting the set of products equal to Ω^*_{ijt} and their quality-to-cost ratios fixed over time (i.e. $\{\gamma^{**}_{uijt}, a^{**}_{uijt}\}$).

Figure 15 summarizes the estimated effect of new products. As expected, only with the exception of young firms, the overall effect of new products is lower in this economy. This is because new products have a smaller role in generating increases in the average quality-to-cost ratio of firms. On the other hand, the effect of introducing a new product on the scope, the quality-to-cost dispersion, and the markup of the firm is roughly the same with or without depreciation of quality-to-cost ratios. Since in this counterfactual exercise older products do not depreciate, it is harder for new products to be substantially better than them. As a result, the fact that a product's quality depreciates increases the benefits of introducing new products, and this partially explains why firms rely on the reallocation of products to fuel growth.

Our model emphasizes the relation between the margins that affect the product's life cycle and the firm's growth. On one side, the decline in the quality-to-cost ratio of existing products affects the decision for and overall effect of new products on the firm's growth. On the other side, the introduction of new products affects the existing products by cannibalizing them and making their overall sales decline even further.

6 Conclusion

We study the life cycles of products and firms jointly. We find that for a wide range of products, sales decline at a fast pace throughout most of their life cycle. The decline in sales is mostly driven by declines in the quantities sold, as opposed to prices, and is more dramatic for larger firms. We develop and structurally estimate a model of heterogeneous multi-product firms and use it to recover the contribution to product sales of quality, cost, markup, the number of products a firm supplies, and the effects of cannibalization. Our results point to the decline in the perceived quality of a product and to cannibalization effects as being the principal reasons for the decline in product sales.

We show that firms counteract these life cycle effects by introducing new products. Product introduction is a necessary condition for their growth. New products not only increase the scope of firms but also help maintain the firms' quality-to-cost ratios. In order to grow, firms must launch products with a sufficiently high quality-to-cost ratio to compensate for the depreciation of older ones. This condition is harder to meet for larger firms not only because they are less active at launching new products, but also because their new varieties are less likely to increase their quality-to-cost ratio and more likely to cannibalize older products. Our counterfactual exercises indicate that the introduction of new products is more beneficial for younger firms and for firms facing faster depreciation in their products' quality. As a result, product innovation is the main force contributing to young firms growing at a faster pace than older firms.

References

Abe, Nobuhiro, Yojiro Ito, Ko Munakata, Shinsuke Ohyama, and Kimiaki Shinozaki, "Pricing Patterns over Product Life-Cycle and Quality Growth at Product Turnover: Empirical Evidence from Japan," *Bank of Japan Working Paper Series 16-E-5*, 2016.

Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William R Kerr, "Innovation, Reallocation and Growth," *NBER Working Paper No. 18993*, 2017.

Akcigit, Ufuk and William R Kerr, "Growth through heterogeneous innovations," *Journal of Political Economy*, 2018, *126* (4), 1374–1443.

Altonji, Joseph G and Robert A Shakotko, "Do wages rise with job seniority?," *Review of Economic Studies*, 1987, *54* (3), 437–459.

Argente, David, Munseob Lee, and Sara Moreira, "Innovation and product reallocation in the great recession," *Journal of Monetary Economics*, 2018, *93*, 1–20.

Arkolakis, Costas and Marc-Andreas Muendler, "The extensive margin of exporting products: A firm-level analysis," *NBER Working Paper No. 16641*, 2010.

Atkeson, Andrew and Ariel Burstein, "Pricing-to-market, trade costs, and international relative prices," *American Economic Review*, 2008, *98* (5), 1998–2031.

_ and _ , "Aggregate Implications of Innovation Policy," *NBER Working Paper No. 17493*, 2018.

_ and Patrick J Kehoe, "Modeling and measuring organization capital," *Journal of Political Economy*, 2005, *113* (5), 1026–1053.

Basker, Emek and Timothy Simcoe, "Upstream, Downstream: Diffusion and Impacts of the Universal Product Code," *NBER Working Paper No. 24040*, 2017.

Bernard, Andrew B, Stephen J Redding, and Peter K Schott, "Multiple-Product Firms and Product Switching," *American Economic Review*, 2010, *100* (1), 70–97.

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

Brockhoff, Klaus, "A test for the product life cycle," *Econometrica*, 1967, pp. 472–484.

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

— and — , "Product Creation and Destruction: Evidence and Price Implications," *American Economic Review*, 2010, *100* (3), 691–723.

Bronnenberg, Bart J and Jean-Pierre Dubé, "The formation of consumer brand preferences," *Annual Review of Economics*, 2017, *9*, 353–382.

— , Sanjay K Dhar, and Jean-Pierre H Dubé, "Brand history, geography, and the persistence of brand shares," *Journal of Political Economy*, 2009, *117* (1), 87–115.

Byrne, David M, Stephen D Oliner, and Daniel E Sichel, "How fast are semiconductor prices falling?," *Review of Income and Wealth*, 2018, *64* (3), 679–702.

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

Copeland, Adam and Adam Hale Shapiro, "Price setting and rapid technology adoption: The case of the PC industry," *Review of Economics and Statistics*, 2016, *98* (3), 601–616.

Deaton, Angus, *The analysis of household surveys: a microeconometric approach to development policy*, World Bank Publications, 1997.

Dunne, Timothy, Mark J Roberts, and Larry Samuelson, "The growth and failure of US manufacturing plants," *The Quarterly Journal of Economics*, 1989, *104* (4), 671–698.

Eaton, Jonathan, Samuel Kortum, and Francis Kramarz, "Dissecting trade: Firms, industries, and export destinations," *American Economic Review*, 2004, *94* (2), 150–154.

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu, "Competition, markups, and the gains from international trade," *American Economic Review*, 2015, *105* (10), 3183–3221.

Eslava, Marcela and John Haltiwanger, "The Life-cycle Growth of Plants in Colombia: Fundamentals vs. Distortions," 2017.

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

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

— and John Romalis, "International prices and endogenous quality," *The Quarterly Journal of Economics*, 2014, *129* (2), 477–527.

Fitzgerald, Doireann, Stefanie Haller, and Yaniv Yedid-Levi, "How exporters grow," *NBER Working Paper No. 21935*, 2017.

Foster, Lucia, John Haltiwanger, and Chad Syverson, "The slow growth of new plants: Learning about demand?," *Economica*, 2016, *83* (329), 91–129.

Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J Klenow, "How Destructive is Innovation?," *NBER Working Paper No. 22953*, 2018.

Gowrisankaran, Gautam and Marc Rysman, "Dynamics of consumer demand for new durable goods," *Journal of Political Economy*, 2012, *120* (6), 1173–1219.

Greenstein, Shane M and James B Wade, "The product life cycle in the commercial mainframe computer market, 1968-1982," *The RAND Journal of Economics*, 1998, pp. 772–789.

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

Hofer, Charles W, "Toward a contingency theory of business strategy," *Academy of Management journal*, 1975, *18* (4), 784–810.

Hottman, Colin J, Stephen J Redding, and David E Weinstein, "Quantifying the sources of firm heterogeneity," *Quarterly Journal of Economics*, 2016, *131* (3), 1291–1364.

Hsieh, Chang-Tai and Peter J Klenow, "The life cycle of plants in India and Mexico," *The Quarterly Journal of Economics*, 2014, *129* (3), 1035–1084.

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

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

Levitt, Theodore, *Exploit the product life cycle*, Vol. 43, Graduate School of Business Administration, Harvard University, 1965.

Lucas, Robert E and Benjamin Moll, "Knowledge growth and the allocation of time," *Journal of Political Economy*, 2014, *122* (1), 1–51.

Manova, Kalina and Zhiwei Zhang, "Export prices across firms and destinations," *The Quarterly Journal of Economics*, 2012, *127* (1), 379–436.

Mayer, Thierry, Marc J. Melitz, and Gianmarco I.P. Ottaviano, "Market size, competition, and the product mix of exporters," *American Economic Review*, 2014, *104* (2), 495–536.

Melitz, Marc J, "The impact of trade on intra-industry reallocations and aggregate industry productivity," *Econometrica*, 2003, *71* (6), 1695–1725.

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

Moreira, Sara, "Firm Dynamics, Persistent Effects of Entry Conditions, and Business Cycles," 2017.

Polli, Rolando and Victor Cook, "Validity of the product life cycle," *The Journal of Business*, 1969, *42* (4), 385–400.

Timoshenko, Olga A, "Product switching in a model of learning," *Journal of International Economics*, 2015, *95* (2), 233–249.

Vernon, Raymon, "International Investment and International Trade in the Product Cycle," *The Quarterly Journal of Economics*, 1966, *80* (2), 190–207.

Table 1: Summary Statistics of Products by Censoring

The table reports summary statistics for the products included in the baseline pooled sample for the period 2006q1-2015q4. For each product, we determine if it has sales in 2006q1 and in 2015q4 to determine if is left- and/or right- censored. Products that enter and are discontinued in the period under analysis are classified as "Complete", products for which we can determine entry but not exit are classified as "Right", products for which we do not observe entry but we observe exit are classified as "Left", and products for which both entry and exit cannot be determined are both right and left-censored ("Right&Left"). For each of these categories, we report the total number of observations, statistics on duration, and statistics on sales. The duration refers to the number of quarters for which we observe the products. Only for products products that enter and are discontinued in the period under analysis ("Complete") it can also be interpreted as the length of life the products. The statistics for the revenue are computed by determining the average quarterly sales (in thousands of dollars), deflated by the Consumer Price Index for All Urban Consumers. The table reports the average and distribution statistics of this variable. Table B.II in the Appendix reports summary statistics of brands by censoring.

	All	By Censoring Type			
		Complete	Right	Left	Right&Left
Total # of products	655,205	225,583	214,554	128,424	86,644
Duration (quarters)					
average	15	7.4	13	13	40
less than 4	33	52	29	31	0
less than 16	68	90	71	70	0
above 28	19	1.3	11	11	100
Revenue (quarterly, $1,000)					
mean	79	27	105	25	180
25th percentile	.5	.2	1	.1	2.2
median	3.8	1.9	7.7	1	13
75th percentile	29	13	54	7.7	89
90th percentile	147	56	233	42	407
95th percentile	342	122	482	107	833
Relative Price					
mean	-.02	-.082	.11	-.14	-.02
25th percentile	-.46	-.56	-.35	-.54	-.39
median	-.0031	-.043	.094	-.079	-.028
75th percentile	.45	.44	.56	.32	.35
90th percentile	.95	.97	1.1	.77	.78
95th percentile	1.3	1.3	1.4	1.1	1.1

Table 2: Dispersion Across Sectors of Within-Sector (log) Sales Distribution

The table summarizes the within-sector moments of log sales across 92 product sectors. We identify product sectors according to their Nielsen classification of product group. Sales are computed by determining the average quarterly sales (in thousands of dollars), deflated by the Consumer Price Index for All Urban Consumers. For each sector, we compute the (log) sales moments across products. We use products included in the baseline pooled sample for the period 2006q1-2015q4. We provide results for all observations, and for sets of observations according to their age (new products, products with 16 quarters, and products with 28 quarters). Columns summarize the information across sectors in terms of the weighted mean (weighted by total number of products within the sector), mean, standard deviation, and inter-quartile range.

	Within-Sector Moments	Weighted Mean	Mean	Std. Dev.	IQ Range
All					
	Median	1.18	1.22	1.00	1.18
	IQ range	4.64	4.65	0.82	1.24
	90-10 percentile range	8.79	8.67	1.13	1.61
	95-5 percentile range	10.87	10.67	1.01	1.55
At entry					
	Median	1.71	1.40	1.35	1.61
	IQ range	4.44	4.18	0.87	1.05
	90-10 percentile range	8.11	7.65	1.20	1.25
	95-5 percentile range	10.02	9.49	1.31	1.25
At age 16					
	Median	1.62	1.44	1.10	1.59
	IQ range	5.01	4.72	1.13	1.58
	90-10 percentile range	8.83	8.44	1.29	1.93
	95-5 percentile range	10.74	10.23	1.24	1.55
At age 28					
	Median	1.84	1.48	1.38	1.88
	IQ range	4.93	4.60	1.23	1.36
	90-10 percentile range	8.74	8.13	1.58	1.86
	95-5 percentile range	10.62	9.99	1.50	1.80

Table 3: Summary Statistics of Firms by Censoring

The table reports summary statistics of firms included in the baseline pooled sample for the period 2006q1-2015q4. For each firm, we determine if it has sales in 2006q1 and in 2015q4 to determine if is left- and/or right- censored. Firms that enter and exit in the period under analysis are classified as "Complete", firms for which we can determine entry but not exit are classified as "Right", firms for which we do not observe entry but we observe exit are classified as "Left", and firms for which both entry and exit cannot be determined are both right and left-censored ("Right&Left"). For each of these categories, we report the total number of observations and statistics on duration, sales, number of products, and number of product categories. The duration refers to the number of quarters for which we observe the firms. Only for firms that enter and exit in the period under analysis ("Complete") it can also be interpreted as duration of the firm. The statistics for the revenue are computed by determining the average quarterly sales (in thousands of dollars), deflated by the Consumer Price Index for All Urban Consumers. The table reports the average and distribution statistics of the variables total sales, number of products and number of sectors. Sectors refers to the number of different product groups as classified by Nielsen.

	All	By Censoring Type			
		Complete	Right	Left	Right&Left
Total # of firms	22,938	4,425	4,726	6,107	7,680
Duration (quarters)					
average	23	11	17	16	40
less than 4	16	35	18	20	0
less than 16	44	81	58	61	0
above 28	43	3.9	18	18	100
Revenue (quarterly, $1,000)					
mean	1,183	8.4	24	111	3,425
25th percentile	.6	.1	.1	1.3	8.9
median	6	.5	1.1	6.8	57
75th percentile	52	3.3	7.7	36	366
90th percentile	359	15	36	162	2,218
95th percentile	1,177	32	87	350	7,387
Products (quarterly)					
mean	12	2.1	3.2	5.3	27
25th percentile	1	1	1	1.3	2.7
median	2.8	1	1.8	3	6.7
75th percentile	6.6	2.5	3.5	5.5	18
90th percentile	18	4.2	6.7	10	52
95th percentile	37	5.8	10	16	98
Sectors (quarterly)					
mean	1.7	1.1	1.3	1.4	2.4
25th percentile	1	1	1	1	1
median	1	1	1	1	1.4
75th percentile	1.7	1	1.1	1.5	2.5
90th percentile	2.8	1.5	2	2.1	4.5
95th percentile	4	2	2.3	3	6.6

Table 4: Decomposition of Sales Growth over the Life Cycle of Firms

The table reports results from the decomposition of the growth of sales at the firm level. The results are obtained after applying equation 1 to each firm × group and aggregating the results using revenue weights. The first column groups the results for all firms in our sample. The second column shows the results for firms that are one year old or less. Column 3 groups firms that are between two and four years of age. Column 4 shows the result of the sample of left-censored firms, those that sold products before the beginning of our sample period (2006q1).

	(1) All	(2) Age 1	(3) Age 2-4	(4) Born before 2006
Growth Sales	0.02	0.83	0.05	0.01
Product Life Cycle Effect	-0.03	0.59	-0.02	-0.03
Entrants Total Effect	0.05	0.24	0.06	0.05
- Entrant Rate	0.14	0.91	0.28	0.13
- Entrants Relative Sales	0.34	0.17	0.32	0.35

Table 5: The Life Cycle of Products: Revenue, Price, and Quantity

The table reports the coefficients for the age fixed-effects of OLS regressions. The dependent variable in columns 1–2 is revenue (in logs), in columns 3–4 is price (in logs), and in column 5–6 it is quantity. For each dependent variable we estimated the regression with and without the age fixed-effect. Age is the number of quarters since we first observe sales for a product (i.age represents an indicator variable that takes the value of one if the product is i quarters of age). Other controls include cohort variables (using Deaton's normalization) and module-quarter fixed-effects. The variables are described in Section 2. The sample used in this table comprises all products in the baseline balanced sample that were born between 2006q2 and 2012q3 and their outcomes for 16 quarters. Standard errors are presented in parentheses and are clustered at state-level. The ***, **, and *, represent statistical significance at 1%, 5%, and 10% levels, respectively.

	(1) logrevenue	(2) logrevenue	(3) logprice	(4) logprice	(5) logquantity	(6) logquantity
age 2		0.210***		-0.00988**		0.220***
		(0.0138)		(0.00428)		(0.0144)
age 3		0.259***		-0.0120***		0.271***
		(0.0139)		(0.00431)		(0.0145)
age 4		0.360***		-0.0168***		0.376***
		(0.0140)		(0.00434)		(0.0146)
age 5		0.358***		-0.0252***		0.384***
		(0.0140)		(0.00434)		(0.0146)
age 6		0.275***		-0.0311***		0.306***
		(0.0141)		(0.00438)		(0.0147)
age 7		0.203***		-0.0333***		0.237***
		(0.0142)		(0.00440)		(0.0147)
age 8		0.209***		-0.0380***		0.247***
		(0.0142)		(0.00442)		(0.0148)
age 9		0.142***		-0.0448***		0.187***
		(0.0143)		(0.00443)		(0.0149)
age 10		0.00952		-0.0483***		0.0579***
		(0.0144)		(0.00446)		(0.0149)
age 11		-0.108***		-0.0505***		-0.0575***
		(0.0144)		(0.00448)		(0.0150)
age 12		-0.149***		-0.0544***		-0.0946***
		(0.0145)		(0.00451)		(0.0151)
age 13		-0.255***		-0.0623***		-0.193***
		(0.0146)		(0.00453)		(0.0152)
age 14		-0.442***		-0.0676***		-0.375***
		(0.0147)		(0.00458)		(0.0154)
age 15		-0.605***		-0.0681***		-0.537***
		(0.0149)		(0.00463)		(0.0155)
age 16		-0.694***		-0.0753***		-0.619***
		(0.0152)		(0.00471)		(0.0158)
Constant	8.847***	8.864***	-0.591***	-0.551***	9.438***	9.415***
	(0.00237)	(0.0103)	(0.000731)	(0.00321)	(0.00246)	(0.0108)
Observations	1,290,208	1,290,208	1,290,208	1,290,208	1,290,208	1,290,208
R-squared	0.192	0.200	0.789	0.789	0.386	0.391
Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Module&Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Sample	Balanced	Balanced	Balanced	Balanced	Balanced	Balanced
Products	UPC	UPC	UPC	UPC	UPC	UPC

Table 6: Statistics on Estimated Parameters

This table reports the statistics for the estimated elasticities and product quality and cost from 83 sectors. We use GMM to estimate the elasticities. The quality and cost are backed out of the given elasticities and observed prices/quantities. Details on the estimation can be found in Appendix H.

	Mean	Std.Dev.	P25	P50	P75
Estimated Elasticities					
σ_j	22.9	26.2	6.5	13.8	31.1
δ_j	0.8	1.4	0.1	0.2	1.0
η_j	10.1	7.9	5.0	7.1	12.3
$\sigma_j - \eta_j$	12.8	24.6	0.1	0.9	17.0
Estimated Product Variables					
$\ln \gamma_{uijt}$	0.0	1.1	-0.6	0.0	0.6
$\ln a_{uijt}$	-9.2	11.8	-11.4	-4.3	-2.7

Figure 1: Evolution of Sales by Cohort of Products: An Example

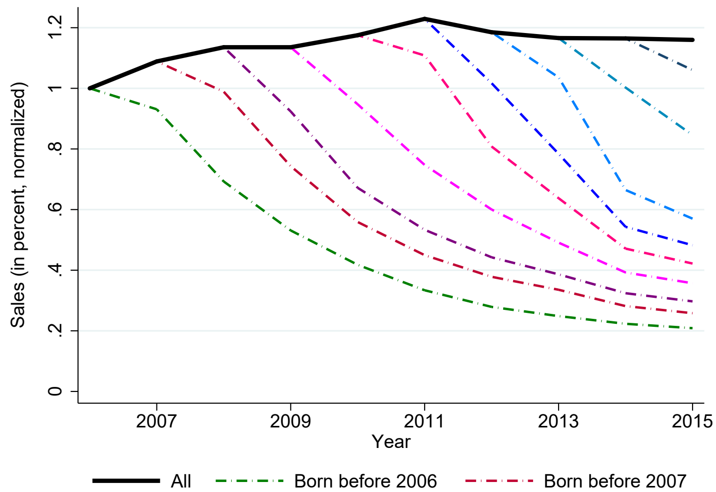

Note: The figure show the evolution of sales of different cohorts of products supplied by one of the largest firms for the period 2006-2015. A black line shows the total sales of the firm over this period. Each of the dotted lines shows the evolution of sales for the products that existed up to each of the periods. The first green dotted line, for instance, represents the path of sales of products that existed in 2006 where the sales of products that enter the market subsequently are not added.

Figure 2: Sales over the Firm Life Cycle

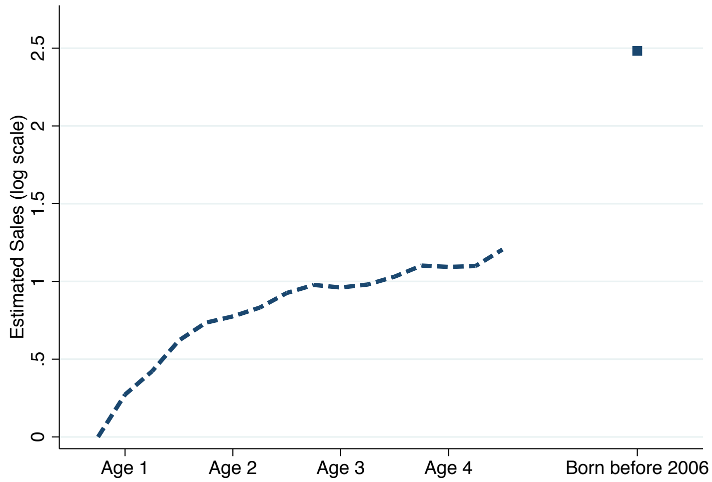

This figure plots the evolution of sales over the firms life cycle. We estimate sales with OLS equation 2 for the baseline balanced sample of firm with at least four years (16 quarters) of duration. We add a fixed-effect to include firms for which we cannot determine age (left-censored firms in our sample), since they introduce a large fraction of the products. The observations are defined as a firm in a sector (defined as product group) in a quarter. The plot presents the estimated age fixed-effects by normalizing the level to zero at entry.

Figure 3: Sales over the Product Life Cycle

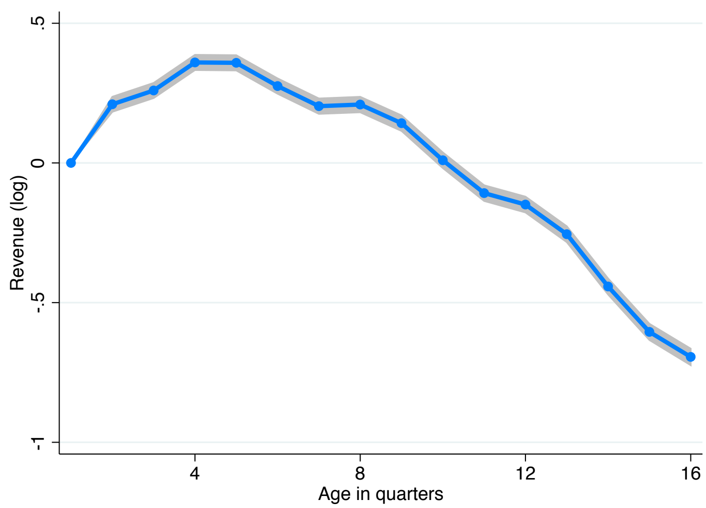

Note: The figure shows the estimated age fixed effects ($\hat{\beta}_a$) of revenue over the life cycle of products identified by their UPCs and computed using equation 2. We keep a balanced sample with 16 quarters or above durations.

Figure 4: Revenue over the Product Life Cycle: Brands

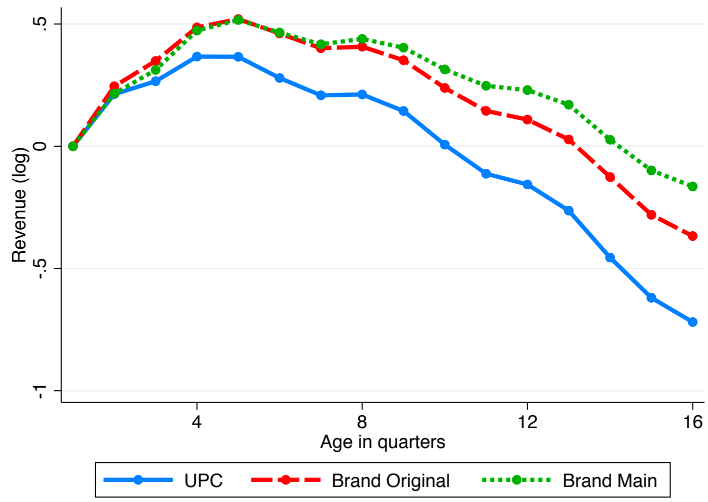

Note: The figure shows the estimated age fixed effects ($\hat{\beta}_a$) of revenue over the life cycle of products defined as brands computed using equation 2. We keep a balanced sample with 16 quarters or above durations. Brand (original) is given by the Nielsen data set, brand (main) combines brands with similar names within firms.

Figure 5: Revenue over the Product Life Cycle: by Duration

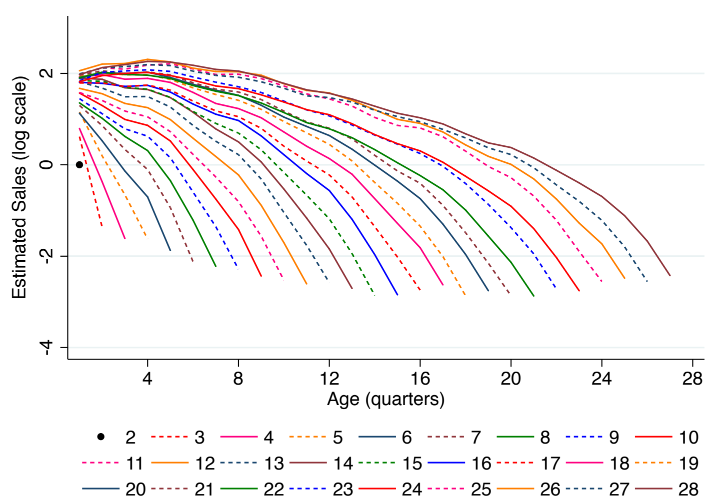

Note: The figure shows the life cycles for products that lasted between 1 and 28 quarters in the market. Every line is estimated using equation 3 and is plotted taking as reference the level of sales of products that lasted only one quarter in the market.

Figure 6: Revenue over the Product Life Cycle (Alternative Samples)

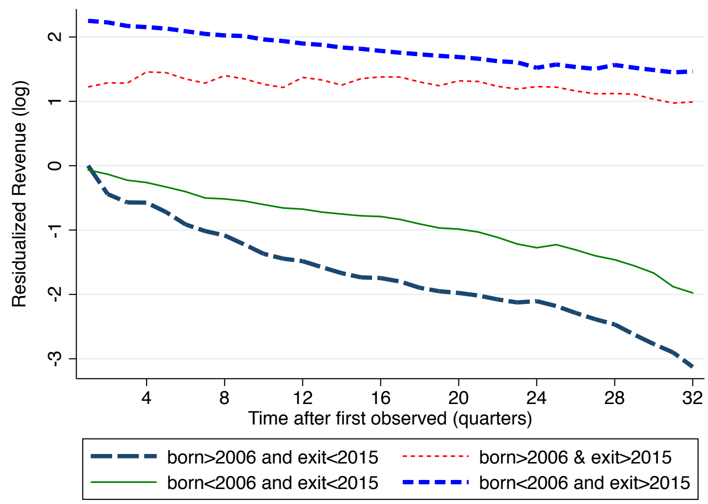

Note: The figure shows the path of revenue for different samples of products after controlling for module and time effects. The graph shows the estimates for four different samples of products: only left censored, only right censored, not censored, and left and right censored.

Figure 7: Revenue over the Product Life Cycle: Durability

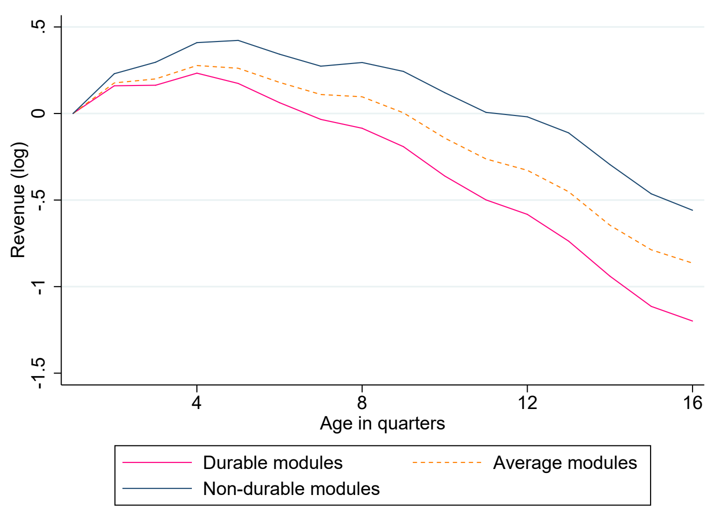

Note: The figure shows the estimated age fixed effects ($\hat{\beta}_a$) of revenue over the life cycle of products that is computed using equation 2 by durability. We keep a balanced sample of at least 16 quarters of duration.

Figure 8: Revenue over the Product Life Cycle: Robustness

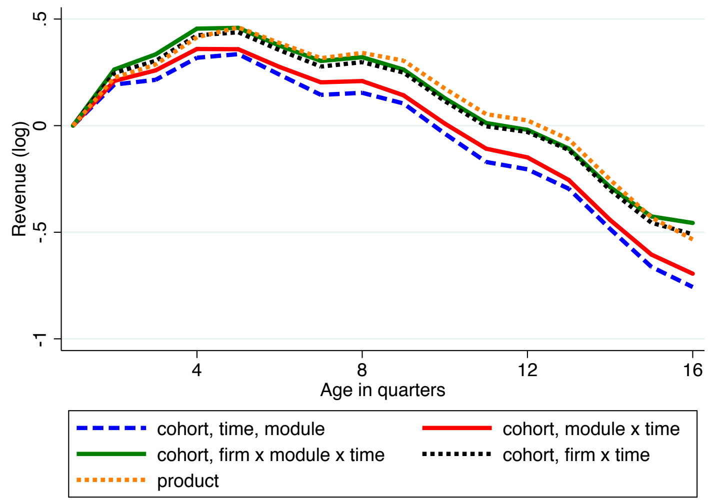

Note: The figure shows the estimated age fixed effects ($\hat{\beta}_a$) of revenue over the life cycle of products using several alternative specifications to that used in equation 2. We keep a balanced sample with 16 quarters or above durations.

Figure 9: Revenue, Prices, and Quantities over the Product the Life Cycle

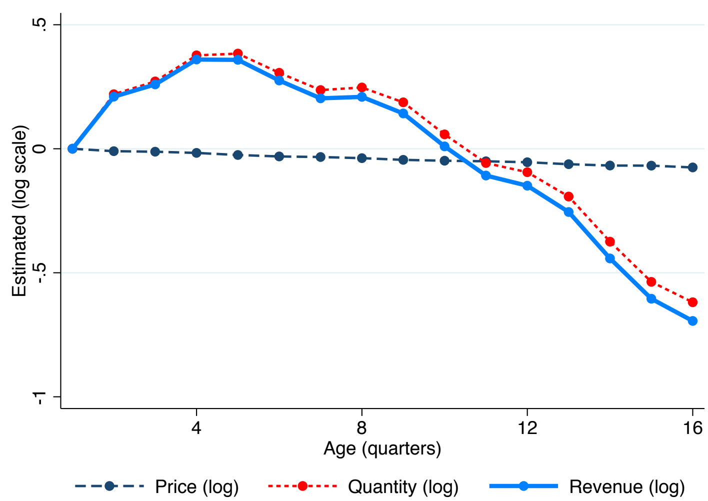

Note: These figures show revenue, price, and quantity of products over their life cycle. Each line is estimated using equation 2 where we keep a balanced sample with 16 quarters or above durations.

Figure 10: Introduction of New Products and the Product Life Cycle

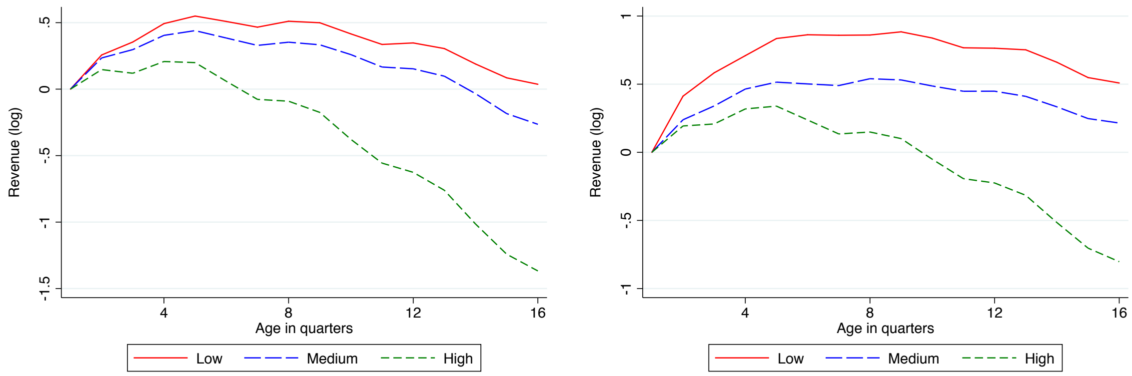

(a) Entry Rate by Category

(b) Firm's Share of Entry

Note: Panel (a) shows the product life cycle after splitting the groups in the Nielsen data by their average quarterly entry rate (weighted by revenue). Each line is estimated using equation 2 for each of the different sets of groups. Panel (b) shows the product life cycle after splitting firms by the average share of new products in a given group and quarter. Each line is estimated using equation 2 and includes sector fixed effects.

Figure 11: Evolution of the Components of Product Sales

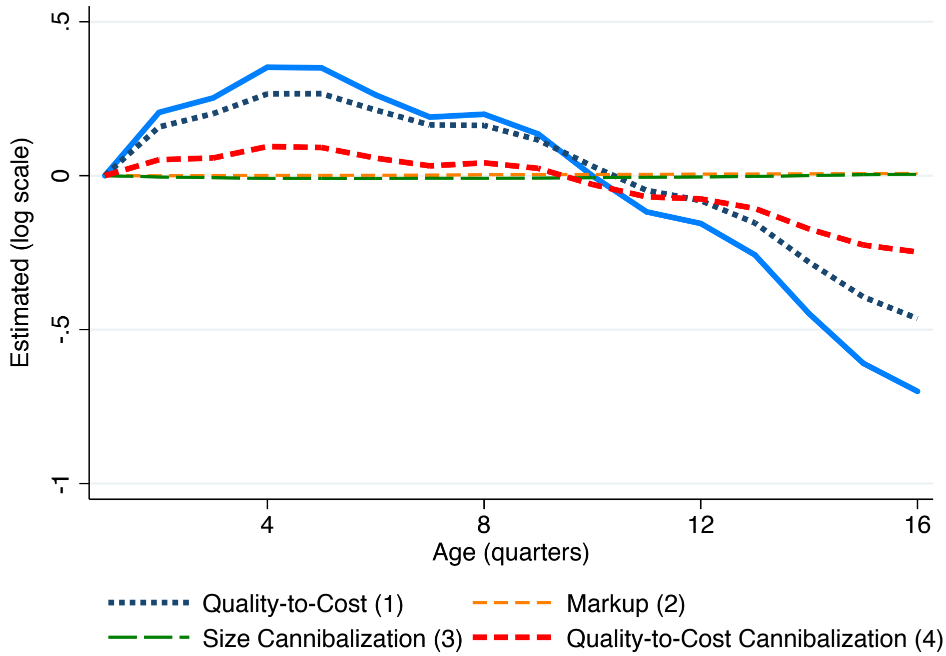

Note: This figure plots the components of the product decomposition over the products life cycle. The solid line plots the estimated age fixed effects ($\hat{\beta}_a$) of sales over the life cycle of products that are computed using equation 2 for the baseline balanced sample of products with at least 16 quarters of duration used in the model (same as in Figure 3). Likewise, the dashed lines plot the fixed effects for each of the components of product sales one-at-a-time as described in the paper. For all variables, the level of the variable is normalized to zero at entry (when products have age equal to one quarter old), and thus a negative fixed effect reflects that the value of the variable is estimated to be below the level at entry. Each of the components (the quality-to-cost, the markup, the size cannibalization, and the quality-to-cost cannibalization) is weighted by combinations of the elasticities of substitution σ_j and/or η_j as derived in equation 21. The estimated age fixed effects of all four components add up to the estimated age effects for sales.

Figure 12: Evolution of the Components of Firm Sales

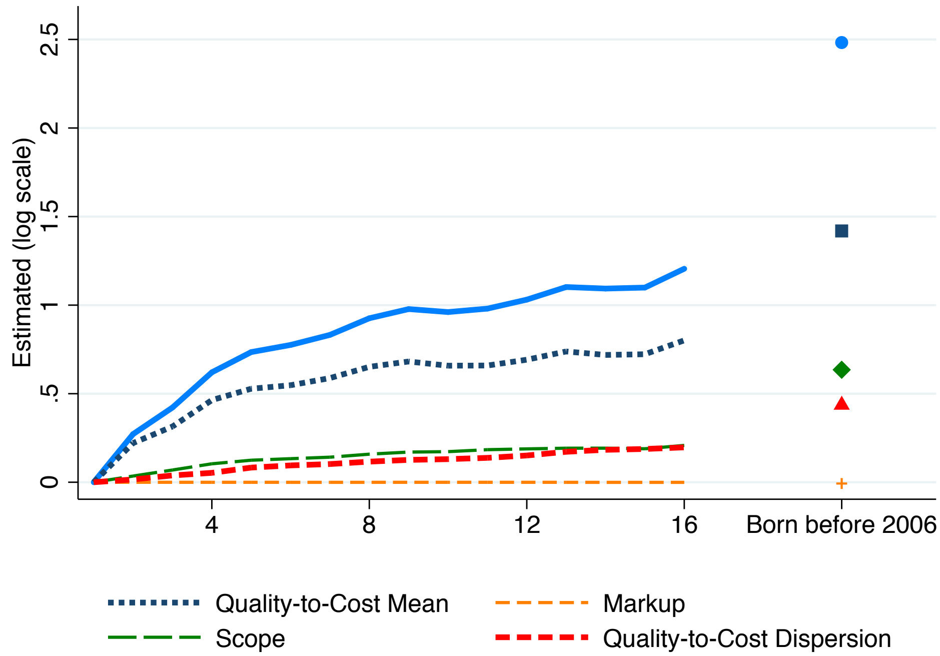

Note: This figure plots the components of the firm decomposition over the product life cycle. The solid blue line plots the estimated age fixed effects ($\hat{\beta}_a$) of sales over the life cycle of products that are computed using equation 2 for the baseline balanced sample of firm with at least four years (16 quarters) of duration. We then add a fixed effect to include firms for which we cannot determine age (left-censored firms in our sample), since they introduce a large fraction of the products. Likewise, the dashed lines plot the fixed effects for the each of the components of firm sales one-at-a-time. For all variables, the level of the variable is normalized to zero at entry (when firms have age equal to one quarter old). Each of the components (the quality-to-cost mean, the markup, the size cannibalization, and the quality-to-cost cannibalization) is weighted by combinations of the elasticities of substitution σ_j and/or η_j as derived in equation 22. The estimated age fixed effects of all four components add up to the estimated age effects for sales.

Figure 13: Difference in Sales and its Components: Observed and Counterfactual Economy without New Products

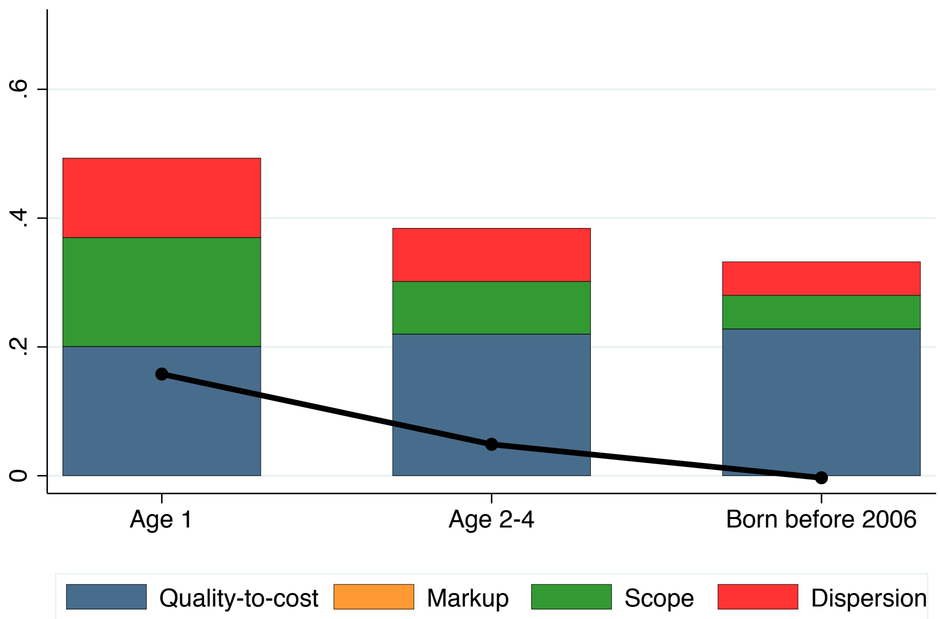

Note: The figure plots the evolution of the estimated average percentage differences in total sales and their components in the observed and counterfactual economy without new products. Age refers to the number of years since entry. Born before 2006 refers to firms that are left-censored in our data. The black dotted line represents total sales, and bar graphs represent four components of firm growth: (i) quality-to-cost, (ii) markup, (iii) scope, and (iv) dispersion.

Figure 14: Difference in Sales and its Components: Observed and Counterfactual Economy without Quality-to-cost Depreciation

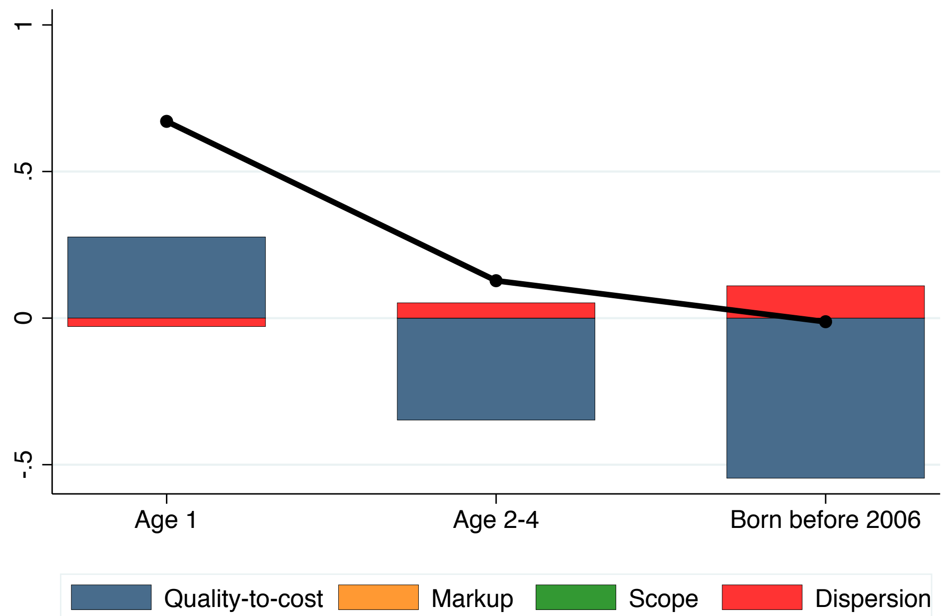

Note: The figure plots the evolution of the estimated average percentage differences in total sales and its components in the observed and counterfactual economy without quality-cost-depreciation as firms become older. Age refers to the number of years since entry. Born before 2006 refers to firms that are left-censored in our data. The black dotted line represents total sales, and the bar graphs represent four components of firm growth: (i) quality-to-cost, (ii) markup, (iii) scope, and (iv) dispersion.

Figure 15: Difference in Sales and its Components: Observed and Counterfactual Economy without New Products or Quality-to-cost Depreciation

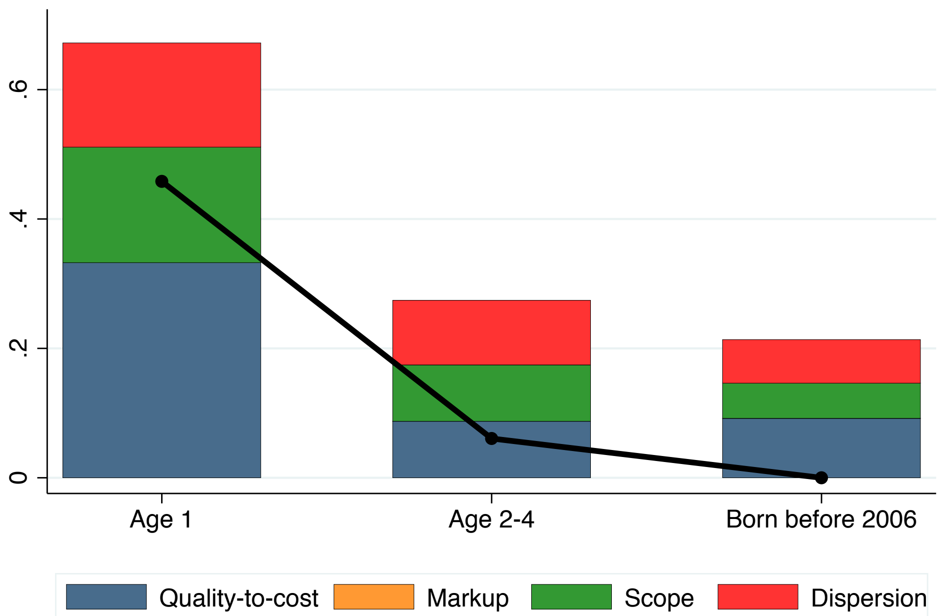

Note: The figure plots the evolution of the estimated average percentage differences in total sales and their components in the observed and counterfactual economy without new products and quality-cost-depreciation as firms become older. Age refers to the number of years since entry. Born before 2006 refers to firms that are left-censored in our data. The black dotted line represents total sales, and the bar graphs represent four components of firm growth: (i) quality-to-cost, (ii) markup, (iii) scope, and (iv) dispersion.