

The Life Cycle of Products: Evidence and Implications*

David Argente[†]

Pennsylvania State University

Munseob Lee[‡]

University of California San Diego

Sara Moreira[§]

Northwestern University

September 2022

Abstract

We exploit detailed product- and firm-level data about the consumer goods industry to document that the sales of individual products decline at a steady pace throughout most of the product life cycle. This pattern is common across heterogeneous types of products and is mostly explained by declines in quantities sold rather than declining prices. We identify the margins that affect the product life cycle using a model-based decomposition of product sales. Our estimates indicate that the systematic decline in sales over time is mostly explained by declines in products appeal. The results are consistent with products quickly becoming obsolete as they face competition from newer products sold by competing firms (business stealing) and competition from newer products sold by the same firm (cannibalization). We build a tractable dynamic model of firm growth in which firms invest in creating new products that impact both its own existing products and the products of competitors. Our model aligns with empirical regularities in the product and firm life cycles, and our quantification shows that firms need to introduce new products to grow; otherwise their portfolios become obsolete as rivals introduce new products of their own. By introducing new products, however, firms accelerate the decline in sales of their own existing products.

*We are grateful to Fernando Alvarez, Anmol Bhandari, Paco Buera, Ariel Burstein, Meghan Busse, Jonathan Eaton, Marcela Eslava, Doireann Fitzgerald, John Haltiwanger, Hugo Hopenhayn, Chang-Tai Hsieh, Thomas Hubbard, Erik Hurst, Benjamin F. Jones, Greg Kaplan, Peter J. Klenow, Marti Mestieri, Yongseok Shin, Joseph Vavra, Venky Venkateswaran and participants at seminars and conferences for their feedback. We would like to thank Olga Denislamova and Xiaojie Liu for excellent research assistance. 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. The paper was previously circulated as “How do firms grow? The life cycle of products matters”.

[†]Email: dargente@psu.edu. Address: 403 Kern Building, University Park, PA 16801.

[‡]Email: munseobleee@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

The recent emergence of large-scale and granular product data sets has allowed economists to make significant advances in understanding the nature, extent, and cyclical properties of product churning.¹ Despite these advances in the literature, little is known about the dynamics behind the rise and fall of products and about the relationships between these product dynamics, firm growth, and competitive conditions in the market. Our paper fills these gaps in the literature by examining the life cycle of a large cross-section of products and by providing evidence about the role product performance plays in shaping both firm- and economic growth.

Our main empirical finding is that after a brief period of increasing sales that lasts approximately a year, most products see steadily declining sales throughout the remainder of their life cycles. This pattern holds across many different types of products and is driven chiefly by reductions in quantities sold, rather than being driven by reduced prices. We interpret this evidence through the lens of a structural decomposition of product sales that allows us to isolate the various margins that affect the evolution of product sales. The framework builds on a parsimonious model of firms that sell multiple products with heterogenous appeal and cost.² Our estimates indicate that as a product grows older, obsolescence is driven by the product’s falling appeal relative to other products. A product’s appeal wanes as competing firms introduce similar new products and as the firm improves upon its own products. We will refer to these processes as “business stealing” and “cannibalization”.

Building on these empirical findings, we create an endogenous growth model featuring innovation and creative destruction forces as in [Klette and Kortum \(2004\)](#). Firms invest in creating and introducing new products, and these introductions affect the sales of a firm’s own products as well as the sales of competitors’ products. Both in the model and in the data, sales of existing products decline steadily over time. On average, however, sales of new products compensate for this decline in full, accounting for the observed growth in overall sales of surviving firms. By introducing new products, a firm broadens its scope while preserving the average appeal of its product portfolio.

Our model highlights the pivotal trade-offs that multi-product firms face due to the tension between cannibalizing the sales of their own products versus competing against other firms. Firms must introduce new products if they want to grow because their product portfolios will otherwise become obsolete as rival firms introduce new products of their own. By introducing

¹See, for example, [Bernard, Redding and Schott \(2010\)](#) and [Broda and Weinstein \(2010\)](#).

²The notion of appeal captures the degree to which consumers prefer a specific product. The existing empirical literature refers to any shifter of demand conditional on price as “quality”. More recently, [Hottman, Redding and Weinstein \(2016\)](#) use the concept “appeal” to avoid taking a stand about whether the shift in demand arises from vertical quality differentiation or subjective differences in consumer taste.

new products, however, firms also accelerate the rate at which their existing products become obsolete. Our results show that competition can be characterized as a self-perpetuating innovation-obsolescence cycle whereby: (i) competitors introduce new products and erode the appeal of other products in the market; (ii) as the appeal of existing products decline, firms selling these products see increasing benefits in developing and introducing new products; and (iii) in introducing new products, firms accelerate the decline in sales and the eventual demise of their own existing products.

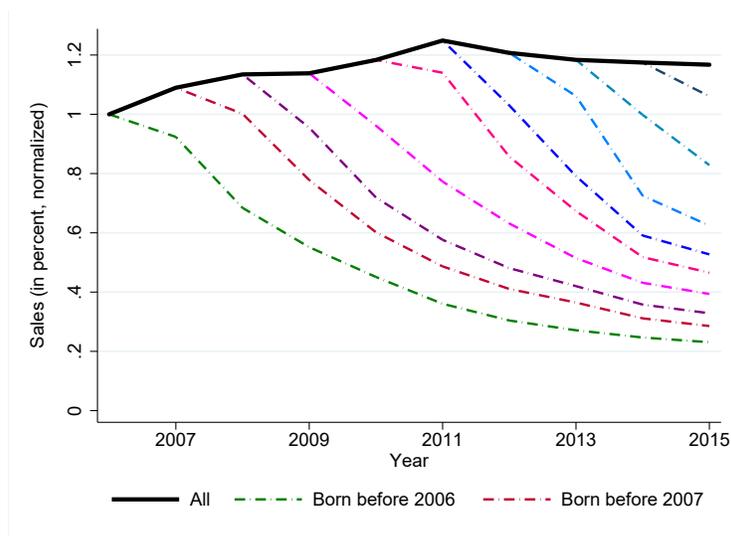
Our analysis is based on comprehensive retail scanner data from Nielsen’s Retail Measurement Services (RMS) that cover the consumer goods industry between 2006 and 2015. This data set covers a broad range of products and industries, including non-durable (e.g., cereals, drinks) and semi-durable consumer goods (e.g., razors, lamps). With these data, we can quantify the contributions that products of different vintages make to a firm’s total sales for more than 20,000 firms, many of which are less than five years old.³

The analysis starts with a simple accounting framework that quantifies the contributions that products with distinct ages make to the average growth in a firm’s sales. On average, firms grow 2% a year, conditional on surviving. New products account for a positive contribution of 12% for this growth and the sagging sales of existing products account for a negative contribution of 10%. Existing products generate fewer sales every year, partly due to discontinuations, but mostly because products’ average sales decline as they grow older. Thus, while each firm’s sales grow throughout the firm’s life cycles, sales of existing products do not. This result is striking, considering that new products take time to diffuse and that firms put effort into expanding the customer base for their existing products over time. To illustrate these patterns, we plot the sales of one of the largest firms in our data set in Figure 1. This firm’s smooth and moderate sales growth conceals massive product reallocation, which is evidenced by the large share of sales generated by new products and the large reductions in the sales generated by older ones.

After we document that the declining sales of existing products hinders firm growth, we employ a regression framework to uncover systematic evidence about how the sales of existing products evolve over the product life cycle. We employ econometric specifications that allow us to estimate how sales evolve as a product ages, while accounting for shocks that affect the sales of all products in a sector during a specific period and for systematic differences in sales across products introduced in different cohorts. Regardless of how long a product survives in

³In our empirical analysis, we use universal product codes (UPCs), or barcodes, as the baseline definition of a product. By using barcodes, we are able to precisely measure the characteristics of products and firms within narrowly-defined sectors and to capture any change in the physical attributes of a good (e.g., form, size, package, or formula). Because some barcode groups may be very close substitutes from the buyer’s perspective, we follow [Kaplan and Menzio \(2015\)](#) to aggregate barcodes into broader groups of products, which can be referred to as brands. Our findings are not qualitatively sensitive to using these alternative product definitions.

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



Notes: The figure shows the evolution of sales of different cohorts of products supplied by one of the largest firms in the data set for the period 2006–2015. The black line represents the evolution of total sales. Each of the dotted lines shows the evolution of sales for the products introduced up to each of the periods. The first green dotted line, for instance, represents the path of sales of products that existed in 2006 when the sales of products that entered the market subsequently are not added.

the market, a common pattern emerges: sales of products decline at a steady pace throughout most of their life cycle. Even within the set of long-lasting products, sales decline on average 30% per year after the first year.

We explore heterogeneity across products in terms of product novelty, success at entry (“superstar” products), and durability. We find that the pattern of declining sales holds for all these different types of products despite some differences in terms of the magnitude and speed of the decline.⁴ We further gauge the sensitivity of our results to the inclusion of firm-specific time-varying effects. We find that firm-specific factors explain an important fraction of the variation in sales across products. Yet, these factors do not affect the evolution of a product’s sales over its life cycle, which suggests that our results are driven not by firm-specific factors but rather by product-specific factors. Finally, we find that both prices and quantities decline over time and that most of the change in sales is driven by changes in quantities sold. Prices of existing products decline 2% per year on average, which is more than an order of magnitude smaller than the respective decline in quantities sold. This result suggests that demand shifters conditional on price, such as changes in a product’s appeal relative to other products, play an important role in the evolution of sales over time.

Next, we employ a model-based decomposition of product sales to quantify the underlying

⁴We also estimate the life cycle patterns of brands. Brands exhibit a less pronounced but still significant decline of sales (25% per year after the first year), which is consistent with firms introducing new barcodes within existing brands.

mechanisms that shape the life cycle of a product. We build upon the model of [Hottman, Redding and Weinstein \(2016\)](#) with product-specific demand and cost heterogeneity in a setting of multi-product firms with variable markups. The framework decomposes product sales into components attributable to variations in appeal and cost of the product relative to other products, markups, and degree of cannibalization. We find that changes in the appeal of a product relative to other products in the market (business stealing) and changes in the appeal of the product relative to other products sold by the same firm (cannibalization) are the most important determinants of the evolution of product sales. About three-fifths of the reduction in product sales are due to losses to products of other businesses, while the remaining two-fifths are due to cannibalization. Our estimation indicates that product costs are mostly constant over time and that firm-specific components that affect product sales (markups, scope, and dispersion) make only minor contributions to changes in product sales over time.

Having established that changes in product appeal are the main driver of changes in product sales, we further investigate whether these losses in product appeal relative to other products can be explained by the introduction of new products in the market. We examine cross-sectional variation across sectors and firms and find that the degree of decline in appeal is stronger for products from sectors and firms with higher rates of new-product entry.

We conclude the paper with a dynamic model of firm growth that includes endogenous product introduction decisions and in which the introduction of new products affects the sales of existing products. Firms invest in both external and internal innovations. The former create new products that improve those sold by competitors; the latter improve upon a firm's own products. When innovating internally, firms internalize the cannibalizing impact that new products have on their existing products, while in the case of external innovation, firms do not internalize the impact of business stealing on competitors' products. Importantly, the model's assumptions are flexible enough to allow a firm's investments in internal and external innovations to be complements or substitutes while allowing the model framework to remain analytically tractable.

We solve the dynamic model analytically and provide a precise theoretical characterization of the underlying determinants of firm and economic growth. The model predicts an innovation-obsolescence cycle in which firms innovate internally in response to obsolescence induced by external innovations introduced by competitors. This prediction has crucial implications for innovation policy and for the quantification of growth ([Atkeson and Burstein, 2019](#); [Garcia-Macia, Hsieh and Klenow, 2019](#)). The model's predictions suggest that negative shocks to the value of external innovation can hinder both internal and external innovations.

To the best of our knowledge, we are the first to use product- and firm-level data to quantify a dynamic model of firm growth with endogenous product introduction decisions.

We use product life cycle moments that capture the extent to which new products substitute for existing products, and differences in the degree of substitutability between products within firms and across firms. We find that surviving firms grow over their life cycles both because they increase the number of products in their portfolios and because of increases in the average appeal of the firm’s products. The introduction of new products plays an essential role in preserving firms’ average appeal which, in the absence of new products, would decline due to external innovations introduced by competitors.

Related Literature – The study product life cycles has been relevant in the fields of marketing and management for decades (e.g., [Levitt, 1965](#)). The relevant economics literature can be traced back to [Vernon \(1966\)](#). Yet, few studies have empirically examined patterns in the product life cycle, and those that do focus on very specific durable products, such as digital camcorders ([Gowrisankaran and Rysman, 2012](#)) and personal computers ([Copeland and Shapiro, 2016](#)). The broad coverage of our data set, which includes both non-durables and semi-durables, allows us to compare the life cycles of products across very different categories. Most importantly, we show that statistics about the product life cycle capture the intensity of the innovation-obsolescence cycle in a sector.

Our decomposition of the drivers of the product life cycle relates to recent literature that examines the determinants of heterogeneity among firms in terms of size and productivity (e.g., [Foster, Haltiwanger and Syverson, 2016](#); [Eslava and Haltiwanger, 2020](#)). We draw extensively from [Hottman, Redding and Weinstein \(2016\)](#), who find that differences in firm appeal explain most observed variance in firm size. Relative to their work, we focus more explicitly on the margins that affect product sales over the life cycle and on the connection between the margins that affect the product life cycle and the sources of heterogeneity among firms.

We also build on recent research about the pervasiveness of product churning within firms ([Bernard, Redding and Schott, 2010](#); [Broda and Weinstein, 2010](#); [Argente, Lee and Moreira, 2018](#)). More broadly, our work links product life-cycle dynamics to a growing literature about firm dynamics and innovation (e.g., [Klette and Kortum, 2004](#); [Perla, 2019](#)). For instance, [Garcia-Macia, Hsieh and Klenow \(2019\)](#) infer the sources of growth from patterns of job creation and job destruction. Instead, we provide direct empirical evidence of the impact of new products and that differences in the performance of newer and older products shape firm- and aggregate economic growth. Our paper is the first to calibrate an endogenous growth model with product-level data and offer direct evidence of the extent to which new products substitute existing products.

Our findings have implications for efforts to quantify the welfare effects of innovation policy. [Atkeson and Burstein \(2019\)](#) analyze the welfare effects of increasing investments in research in a model with own innovation and creative destruction. They find smaller welfare gains

from research investments when growth involves business stealing. In order to determine the welfare effects of innovation policy, it is therefore important to know the extent to which growth comes from external versus internal innovation. Our model shows that in the presence of an innovation-obsolescence cycle, internal and external innovation investments will have a strong complementarity in their responses to changes in innovation policies.

Lastly, our paper relates to the literature on the role that non-price strategies play in competition. [Nevo \(2001\)](#) and [Wollmann \(2018\)](#) show the importance of new product introduction as a hallmark of firm competition in the ready-to-eat cereal industry and truck manufacturing industry, respectively. Our paper is the first to explore the sales history of individual products across multiple sectors and show that competition is characterized by an innovation-obsolescence cycle whereby firms must introduce new products to compete and, significantly more so, when its competitors are more innovative themselves.

2 Data

2.1 Defining a Product

We use barcodes as our baseline definition of products. A barcode is a Universal Product Code (UPC) that consists of 12 digits and is uniquely assigned to each specific good available in stores. UPCs were created so retail outlets could determine prices and inventory accurately and to improve transactions along the supply chain ([Basker and Simcoe, 2021](#)). Barcodes offer a unique opportunity for economists to identify products at their 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., form, size, package, formula) result in a new barcode. By using barcodes, we ensure that we observe the exact same product at different points in time and that changes in performance do not result from changes in the attributes of the product. The most common alternative is to define goods by industry classification. Defining a product at that level can potentially aggregate very heterogeneous barcodes, which means that changes 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.

Second, barcodes are so widespread that our data is likely to cover all products in the consumer goods industry ([Basker and Simcoe, 2021](#)). 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. Further, because firms and products are included in the sample provided that a sale occurs, we observe a wide range of products and we explore several dimensions of heterogeneity.

Finally, by using barcodes as the baseline unit of analysis, we do not ex-ante distinguish major changes from minor changes in product characteristics. For example, 7.5 oz and 12 oz cans of Diet Coke are treated as two different products. In many settings these and other changes in packaging and size may not be a desirable definition of a product. In the context of this paper, measuring these types of changes in product characteristics matters to better understand the product’s life cycle. Minor changes in product characteristics such as packaging could differ across firms and types of products at different stages of their life cycle. Our analysis is flexible and implicitly accounts for this by estimating elasticities of substitution between barcodes produced by the same firm.⁵ Moreover, we also evaluate if the results of our analysis are qualitatively similar when we focus on barcodes with novel characteristics and when we define products using broader definitions as in [Kaplan and Menzio \(2015\)](#). Throughout the paper we use brands as an alternative product definition. We focus on brands because other work studying the consumer goods industry have used brands as their main unit of analysis either because advertising data is defined at the brand level or because firms’ internal organization aligns closely with their portfolio of brands and product lines ([Bronnenberg, Dhar and Dubé, 2009](#); [Bronnenberg and Dubé, 2017](#)).

2.2 Product Data

We rely primarily on the scanner data set from the RMS 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 $\text{UPC} \times \text{store} \times \text{week}$ level. Total sales in the RMS cover approximately \$220 billion per year, which is roughly 40% of the nationwide consumption in the consumer goods industry. This volume of sales 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 metropolitan statistical areas 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.⁶

⁵Barcodes that are perceived by consumers as indistinguishable will have very large elasticities of substitution, and generate outcomes isomorphic to treatment of products as the same in the first place.

⁶In 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

The data covers a wide range of products both in terms of type (e.g., from non-durables such as cereals to semi-durables like lamps) and in terms of sales share. The original data consist of more than one million distinct products that are identified by UPC and are organized into a hierarchical structure. Each UPC is classified into one of the 1,070 product modules that are organized into 104 product groups, that are then grouped into 10 major departments. For example, a 31 oz bag of Tide Pods (UPC 037000930389) is mapped to product module “Detergent-Packaged” in product group “Detergent” that belongs to the “Non-Food Grocery” department. Throughout the paper we refer to sectors as either product modules or product groups.⁷

Our baseline data set combines the sales of a product across all stores covered in the sample over a quarter. 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, we define quantity y_{ut} as the total quantity sold across all stores and weeks in the quarter, and price p_{ut} is the ratio of sales to quantity, which is equivalent to the quantity weighted average price. For some empirical analyses, we also use a data set of quarterly sales of a product in each store within a limited set of stores in the sample.

We identify the age and life cycle of a product by observing the timing of its initial transaction in the data set. Specifically, we define entry as the quarter in which the first sale of a product occurs and exit as the quarter following the last sale of the product. We cannot determine entry and exit for some products. We classify products that are already active in the first two quarters of the sample (2006q1 and 2006q2) as left-censored. This group of products includes some that were created just before 2006 and some others that were 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 we cannot measure how long they lasted in the market. To minimize concerns of potential mismeasurement of a product’s entry and exit, our baseline sample covers a balanced set of stores and excludes products without at least one transaction per quarter after entering as well as private label products and departments that are not representative.⁸

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 Nielsen Consumer Panel Data.

⁷By default we use product modules. In some exercises exploring heterogeneity in the product life cycle and in the estimation of the parameters of the model we use product groups because many product modules have insufficient observations.

⁸Our estimates of products’ entries and exits might be affected by the entries and exits of stores in the 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

Table 1: Summary Statistics of Products 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
Sales (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
95th percentile	342	122	482	107	833

Notes: The table presents the summary statistics for the products included in the baseline pooled sample for the period 2006q1-2015q4. Products that are already active in 2006q1 and 2006q2 are left-censored, and products with sales in 2015q3 and 2015q4 are right-censored. Products that enter and are discontinued in the period 2006q3-2015q2 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. Under duration, we report the average duration and the share of observations with duration below 4, 16 and above 28 quarters. The duration refers to the number of quarters for which we observe the products. The statistics for sales are computed by determining the average quarterly sales (in thousands of dollars), deflated by the Consumer Price Index for All Urban Consumers. The table presents the average and distribution statistics of this variable. Table A.I in the Online Appendix presents equivalent summary statistics of brands by censoring.

Our baseline sample includes approximately 650,000 products that are organized into 92 distinct groups and 904 modules. Products are very heterogeneous in their market shares and durations. Table 1 reports summary statistics for key product characteristics such as product sales, duration, and price. The table suggests that the distribution of sales by product is highly right-skewed. The average product generates 20 times more sales per quarter than the median product and the product in the 75th percentile of the distribution of sales generates 50 times more sales than the product in the 25th percentile of the distribution. Furthermore, we also show in Table A.III of the Online Appendix that this large dispersion in sales exists even within narrowly defined sectors.

Table 1 also presents summary statistics by type of censoring. We divide products into four categories: (i) complete, (ii) right-censored, (iii) left-censored, and (iv) both right- and left-censored. For example, our data set identifies a 12 count of 12oz cans of regular Coca-Cola as a “right- and left-censored” because the product already existed prior to the beginning of

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 qualitative role in our results.

our analysis and survived through our entire sample 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 in the data set when a product is uncensored or right-censored, but not simultaneously right- and left-censored. We are able to measure age and follow the life cycle of products in this group, which comprises approximately two-thirds of the sample. Among this group, we are able to measure both the age and duration of more than 50% for which we can identify exit. We cannot measure age for the remaining one-third of products that were already active in the first two quarters of our sample, but we can identify exit for 60% of the products in this group. When we measure the average quarterly sales of each product, we find that the total average quarterly sales of products for which we can determine age account for close to 60% of total average quarterly sales across all products in the sample. The summary statistics also show that products have short durations: the median product lasts between 12 and 16 quarters.

2.3 Firm Data

We study the implications of the life cycles of products for the growth of firms. We link firms and products with information obtained from GS1 US, which is the single official source of UPCs. With this link, we can conduct the analysis at the parent company level rather than at the level of the manufacturing firm. Because the GS1 US data contains all of the company prefixes generated in the US, we combine these prefixes with the UPC codes from the RMS. By linking firms to products, we are able to characterize the portfolio of every firm with products in our sample. Furthermore, we can identify the sales, price, and quantity of each product belonging to every firm and compute these variables at the firm level.⁹ We mostly focus on measures of firm size (number of products and total sales), and entry (frequency, number, and sales) and age of its products. 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 2 presents the characteristics of firms by type of censoring. Among the approximately 23,000 firms covered in the sample, we can measure the age of about 9,000, and the remaining 14,000 are firms that were born before 2006. As expected, firms that are not “left censored”

⁹To be able to interpret the aggregate sales of a firm at the retailer as the sales of a manufacturer firm, we make a few assumptions. First, we assume that the aggregation of sales across regions and retailers to the firm level averages out regional and retail-specific shocks in the cross-section. Second, we explore results at the quarter frequency and/or yearly frequency under the assumption that at this frequency variations in inventories at the retail level are less likely to affect overall sales. Third, we assume that the retailer’s markups are constant over time and over the life cycle of firms, consistent with recent evidence documented by [Anderson, Rebelo and Wong \(2018\)](#) and [Argente, Fitzgerald, Moreira and Priolo \(2021\)](#) respectively.

Table 2: Summary Statistics of Firms by Censoring

	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
Sales (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
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
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
95th percentile	4	2	2.3	3	6.6

Notes: The table presents summary statistics of firms included in the baseline pooled sample for the period 2006q1-2015q4. Firms that are already active in 2006q1 and 2006q2 are left-censored, and firms with sales in 2015q3 and 2015q4 are 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 firms, and number of sectors. Under duration, we report the average duration and the share of observations with duration below 4, 16 and above 28 quarters. The duration refers to the number of quarters for which we observe the firms. The statistics for sales are computed by determining the average quarterly sales (in thousands of dollars), deflated by the Consumer Price Index for All Urban Consumers. The table presents 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.

are smaller both in sales and in number of products and they are less diversified. Firm that are “right- and left-censored” have on average 27 products in their portfolios in four different product modules and two different product groups. Throughout the paper, we present evidence for both young firms and old firms, which we define as firms between one and four years of age and firms born before 2006 respectively.

3 Firms' Growth: New versus Existing Products

A well-established fact about firms' life cycle is that they start small and grow larger as they become older (e.g., Dunne, Roberts and Samuelson, 1989; Hsieh and Klenow, 2014). Consistent with this evidence, we also find that firms in our data set grow over time. The unique feature of our data set is that we have information to decompose each firm's sales into the sales of its individual products. Young firms necessarily have new products, but older firms usually sell a portfolio of both newer and older products. For example, Figure 1 illustrates the evolution of sales for one of the largest firms in the consumer goods industry. The solid black line represents the total sales while the colored lines represent the sales of each cohort of products of this firm. The black line indicates that the sales of this firm have grown smoothly, but the colored lines show that the sales of each cohort of products have declined rapidly over time. For example, products created through 2006 (green dashed line) account for about 90% of the sales in 2007 but less than 20% of the sales in 2015. This decline in sales of existing products is pervasive across all cohorts and is accompanied by a steady entry of new products. For this firm, the total sales generated by new products is larger than the decline in sales of older ones.

The patterns in Figure 1 motivate us to evaluate if the strong decline in the sales of older products is common across the firms in our data set. We decompose sales of firms in a sector into the sales of products by their respective age and apply this decomposition to all firms in our sample. After arranging the different components, and aggregating the sales of products into new and older products, we write sales growth (in percentage terms) for each firm \times sector i as:

$$\Delta S_{i,t} = \underbrace{n_{i,t}^{\text{new}} \times \bar{s}_{i,t}^{\text{new}}}_{\text{New Products}} + \underbrace{\Delta S_{i,t}^{\text{old,survive}} - \bar{S}_{i,t-1}^{\text{old,exit}}}_{\text{Product Life Cycle}}, \quad (1)$$

in which *New Products* is the share of sales from new products. We further decompose *New Products* into the product between the entry rate of new products, $n_{i,t}^{\text{new}}$, and the sales of the new products relative to the average sales of the products of the firm, $\bar{s}_{i,t}^{\text{new}}$. The *Product Life Cycle* term quantifies the contribution of older products to sales growth. We further decompose this component into the sales growth (in percentage) of existing products conditional on survival, $\Delta S_{i,t}^{\text{old,survive}}$, and the sales share of non-surviving products, $\bar{S}_{i,t-1}^{\text{old,exit}}$, representing the intensive margin of growth among surviving products and the extensive margin from exit, respectively.¹⁰

Table 3 (column 1) presents the weighted average contribution of each of these components

¹⁰A full derivation of equation (1) can be found in the Online Appendix B.1.

Table 3: Decomposition of Sales Growth over the Life Cycle of Firms

	(1) All	(2) Age 1	(3) Age 2-4	(4) Born before 2006
Growth Sales	0.016	0.328	0.172	0.014
Product Life Cycle Component	-0.102	-0.026	-0.035	-0.102
Growth of Surviving	-0.099	-0.019	-0.029	-0.099
Sales Share of Exit	-0.003	-0.006	-0.006	-0.003
New Products Component	0.118	0.354	0.207	0.116
Entry Rate	0.145	0.639	0.328	0.142
Entrants Relative Sales	0.718	0.587	0.727	0.718

Notes: The table presents the results from the decomposition of the annual growth of sales at the firm \times sector level, as defined in equation (1). For each firm \times sector \times year, we compute the contribution of entrants using data on the number and sales of products in their first full year of activity. The contribution of surviving and exit products is determined by the sales of products that have more than one full year of activity. The table presents the sales-weighted average across all firms, sectors and years. Because of censoring at entry and exit, the average is for the period 2007–2014. The first column groups the results for all firms in our sample. The second column shows the results for firms \times aggregates that are one year old or less (excluding firms at entry). Column 3 groups firms \times aggregates that are between two and four years of age. Column 4 shows the results of the sample of left-censored firms that are those that sold products before the beginning of our sample period. Sectors are defined according to Nielsen product groups. See Online Appendix B.2 for the results of the decomposition when using brands or an alternative definition of growth rates.

to sales growth, in which the weights of each firm are determined by their respective total share of sales in our sample. In our pooled sample, firms grow on average 2% a year during the sample period, conditional on survival. The contribution of new products to sales growth is about 12% a year. This positive impact of new products is the product of an average product entry rate of 15% and of average sales of new products that represent 72% of the average sales of existing products. The entry of new products more than accounts, on average, for the positive growth in firm sales. In line with our anecdotal example in Figure 1, we find that the growth rate of the sales of existing products is negative for the average firm in our sample. This pattern indicates that as products become older their sales decline. Furthermore, the negative contribution of the product’s life cycle to overall sales growth is mostly explained by the negative sales growth of surviving products rather than by the exit of products.

Columns (2) to (4) of Table 3 repeat the decomposition of the average annual growth rates across different groups of firms based on their respective age. The growth paths of 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 grow older. Firms in our sample grew almost 33% annually in their first year of activity but only approximately 17% from ages two to four. Among older firms created before 2006, the average rates of sales growth are even lower at approximately 1%. This decline in the average growth rates of sales as firms become older results both from a decline in

the product life cycle component and from a decline in the contribution of new products. The life cycle effect is negative and becomes even more negative among older firms as it changes from about 3-4% among young firms to about 10% among older firms. New products contribute positively to sales growth but their positive contribution also declines with age. The contribution of this component goes from approximately 35% among very young firms to approximately 12% among firms that existed before 2006. Most of this decline comes from older firms reducing their product entry rates since the average sales of new products remain approximately constant as firms become older.

Overall, the results from our statistical decomposition indicate that the patterns of Figure 1 are representative of the patterns of a broad sample of firms in our data. This accounting decomposition sheds light on the average negative contribution of existing products to the firm's growth but leaves open important questions. It does not ascertain if the decline in average sales of existing products is a systematic pattern associated with their age or is rather a result of composition and time effects that affect the sales of older products. It also does not offer empirical details about the important elements of the timing and rate of decline in sales of existing products. In the next section, we look deeper into these issues by using a regression framework that allows us to better understand the evolution of sales of existing products and their life cycle.

4 Life Cycle of Products: Evidence and Causes

We start by using information about the age of each product to empirically study common patterns that characterize the life cycles of products. We document that the sales of products decline at steady pace throughout most of the product life cycle, and the pattern is common across heterogeneous types of products. After describing the evidence, we investigate the margins that affect the product's life cycle. We build on [Hottman, Redding and Weinstein \(2016\)](#) and develop a product sales decomposition by structurally estimating a model of oligopolistic competition between heterogeneous multi-product firms. Guided the results of the decomposition, we show that the decline is associated with the introduction of new products.

4.1 Descriptive Evidence

Measurement – We characterize the life cycles of products by estimating the evolution of the performance of product sales as a function of age. To properly isolate the effect of age, we account for the fact that we observe products in different quarters and that we want to capture the effect of age irrespective of the particular period in which we observe a product.

Likewise, we want to control for the fact that otherwise comparable products might behave differently depending on the timing of their entry. In order to address such issues, we estimate age effects by implementing age-period-cohort models. These models allow us to estimate the evolution of sales, quantities, and prices since the entry of the product, while accounting for cohort-specific differences in outcomes and any calendar effects specific to the period in which we measure the outcomes (e.g., business cycles that affect all products).¹¹

In the baseline specification, we estimate the outcome of interest (Y) of product u observed at time t as a function of age (a), period \times sector (jt), and cohort (c) fixed effects:

$$\log Y_{u,t} = \alpha + \sum_{a=2}^A \beta_a D_a + \lambda_{jt} + \theta_c + u_{u,t} \quad (2)$$

We are interested in the series of coefficients, β_a , that capture the average aging process of the products relative to the level of the outcome in the first full quarter of activity. Because there is an exact linear relation between the three effects, we normalize the cohort effect as suggested in [Deaton \(1997\)](#).¹² Also, because the products in our sample belong to heterogeneous categories, we allow the time fixed effect to be specific to each sector.

The evolution of the products' outcomes over their life cycle is affected by selection. Our sample of products contains all barcodes from their full first quarter of activity until the quarter before they exit. The statistics on the duration of a product that we report in [Table 1](#) show substantial differences in the duration of barcodes. This large heterogeneity in product duration means that our estimated effects are conditional on survival if we use all active observations regardless of their duration. Because products that are discontinued earlier are different from those that are discontinued later, the unconditional estimated effects will be different from the conditional estimated effects. In order to ensure that our estimation results are not driven by selection biases that result from the inclusion of short-lived products, we conduct the baseline empirical analysis on products that survive at least 16 quarters, which is just above the median survival age. In [Online Appendix C.1](#), we present the results with alternative specifications and samples. We also consider an alternative sample and specification that explicitly accounts for selection to better understand its nature. We use a sample with both short- and long-lived products, and we allow for the age fixed effects to be distinct depending on the duration of the product. In this alternative specification, we estimate the

¹¹These models are commonly used in the literature on individuals' life cycle consumption and income dynamics and was also used for firms in [Moreira \(2017\)](#). [Schulhofer-Wohl \(2018\)](#) provides a general discussion of these models in the context of structural estimation.

¹²The normalization averages the cohort effects to zero over the sample period and orthogonalizes the cohort trends such that the linear component of growth is attributed to age and time effects. As is common in the literature, we check the robustness of this normalization by considering alternative specifications and found that the estimated age effects are qualitatively robust to this normalization ([Online Appendix C.1](#)).

outcome of interest of product u observed at time t as follows:

$$\log Y_{u,t} = \alpha + \sum_{d=2}^D \sum_{a=1}^d \gamma_{ad} D_{ad} + S_{u,t} + \lambda_{jt} + \theta_c + u_{u,t} \quad (3)$$

in which λ_{jt} are the interacted sector and time fixed effects, θ_c are Deaton’s normalized cohort effects, D_{ad} are dummies for age interacted with duration, and $S_{u,t}$ is a dummy for censored observations. The specification not only isolates the life cycle dynamics conditional on the ex-post duration of a product but also allows us to examine if the initial outcomes forecast the product’s own survival.¹³

Average sales patterns – We estimate equation (2) using the quarterly sales (in logs) of products that were active for at least 16 quarters as our main dependent variable. Table 4 presents the estimated age fixed effects. Column (1) shows the results of the baseline specification in the absence of age fixed effects. When we compare the results presented in column (2) with those of column (1) we see that the age fixed effects explain some of the variation in the quarterly sales. More importantly, in column (2), we find that the coefficients of the series of age fixed effects are positive and statistically significant in the early stages of a product’s life cycle and negative and statistically significant later on. We plot the series of estimated coefficients in Figure 2. Product sales mostly decline with age except during the first four to five quarters of the life of a product. Between the first and fourth year of activity, product sales decline on average 30% per year. When we run the same specification for each sector (defined as product groups) separately, we find that 86 out of 92 sectors show a decline in sales between the first and fourth year of activity. Our results indicate that the phase of growth of a typical product is, in fact, very short.

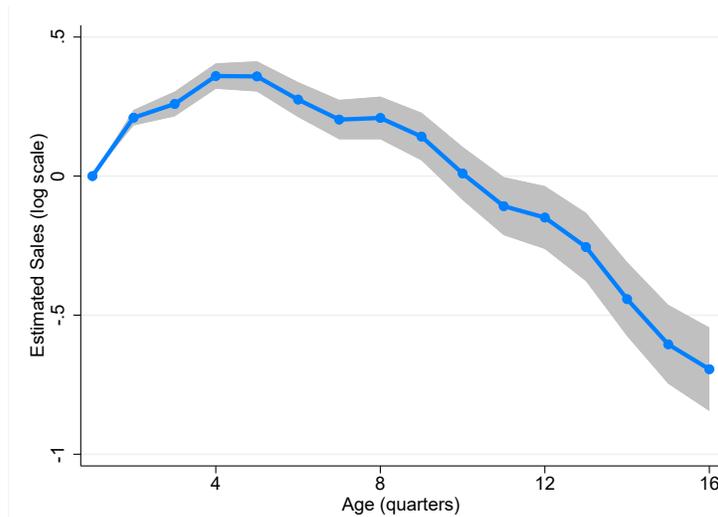
By contrast with the conventional view that product sales follow a bell-shaped evolution (Levitt, 1965), we find a steady decline in sales throughout the greater part of the life cycle of products whose market longevity exceeds 16 quarters.¹⁴

Sales for short and long lasting products – To better understand the nature of selection in our data set, we study the life cycle patterns conditional on the product’s ex-post duration

¹³This approach has some similarities with the approach of Fitzgerald, Haller and Yedid-Levi (2016) who work with exporter dynamics and Altonji and Shakotko (1987) who deal with selection in estimating the effect of job tenure on wages.

¹⁴We show the robustness of this result in Online Appendix C. In Online Appendix C.1, we show our results are similar when we use other data sets (i.e. Nielsen Homescan Measurement System) or alternative data samples (e.g. unbalanced sample). In Online Appendix C.2, we also show that M&A events do not affect the average patterns of product and firms dynamics. Lastly, in Online Appendix C.3, we show that the life cycle patterns are qualitatively similar after exploring the role of generic brands.

Figure 2: Sales over the Product Life Cycle



Notes: The figure shows the estimated age fixed effects of revenue over the life cycle of products identified by their barcodes and computed using equation (2). The estimation includes cohort and time effects that are specific to product modules and cohort effects. We keep a balanced sample with 16 quarters or above durations. The gray area indicates the 95% confidence interval. Standard errors are clustered at product category level.

by estimating equation (3) for products with durations of between 2 and 28 quarters. Figure 3 shows that sales of short-lived products decline throughout their life cycles and that the negative growth rates of these products accelerate as they approach exit. Moreover, short-lived products also generate fewer sales at entry, which indicates that sales during the first few quarters of activity are an important determinant of the expected duration of products. Long-lived products see mild increases in sales in the first quarters of activity and declines in sales thereafter; this reduction in sales occurs at a slower pace than that of short-lived products. We also find that sales at exit are significantly lower than sales in the first quarter of activity across all durations that we consider in the analysis and most products experience declines in sales for several periods prior to exit.¹⁵

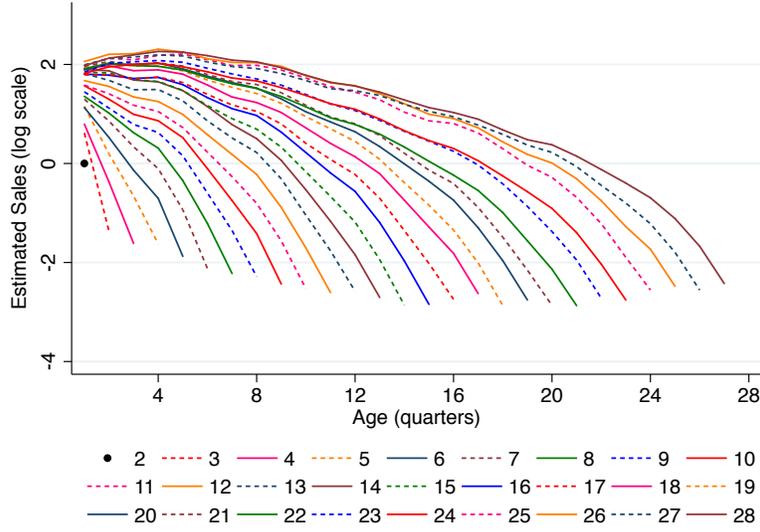
¹⁵We also examine very long lasting products. Figure C.11 of Online Appendix shows the path of sales for products that lasted the entire period covered in our data (i.e., products that entered the market before 2006 and exited after 2015). Given their longevity, we might conjecture that these are the most successful products. Because we cannot determine their age, we plot the evolution of these products' sales after controlling for sector and quarter fixed effects, and we show that sales consistently decline for these products as well.

Table 4: The Life Cycle of Products: Sales, Price, and Quantity

	(1) Log(Sales)	(2) Log(Sales)	(3) Log(Sales)	(4) Log(Price)	(5) Log(Quantity)
1[Age = 2]		0.210*** (0.0141)	0.263*** (0.0270)	-0.00988*** (0.00152)	0.220*** (0.0135)
1[Age = 3]		0.259*** (0.0226)	0.334*** (0.0412)	-0.0120*** (0.00201)	0.271*** (0.0224)
1[Age = 4]		0.360*** (0.0232)	0.455*** (0.0445)	-0.0168*** (0.00295)	0.376*** (0.0227)
1[Age = 5]		0.358*** (0.0276)	0.459*** (0.0525)	-0.0252*** (0.00378)	0.384*** (0.0272)
1[Age = 6]		0.275*** (0.0319)	0.375*** (0.0587)	-0.0311*** (0.00461)	0.306*** (0.0312)
1[Age = 7]		0.203*** (0.0361)	0.303*** (0.0636)	-0.0333*** (0.00551)	0.237*** (0.0352)
1[Age = 8]		0.209*** (0.0390)	0.321*** (0.0673)	-0.0380*** (0.00648)	0.247*** (0.0376)
1[Age = 9]		0.142*** (0.0436)	0.263*** (0.0732)	-0.0448*** (0.00730)	0.187*** (0.0416)
1[Age = 10]		0.00952 (0.0487)	0.128 (0.0808)	-0.0483*** (0.00800)	0.0579 (0.0464)
1[Age = 11]		-0.108** (0.0531)	0.0123 (0.0862)	-0.0505*** (0.00878)	-0.0575 (0.0508)
1[Age = 12]		-0.149*** (0.0575)	-0.0198 (0.0925)	-0.0544*** (0.00968)	-0.0946* (0.0546)
1[Age = 13]		-0.255*** (0.0626)	-0.107 (0.100)	-0.0623*** (0.0107)	-0.193*** (0.0590)
1[Age = 14]		-0.442*** (0.0677)	-0.288*** (0.108)	-0.0676*** (0.0115)	-0.375*** (0.0639)
1[Age = 15]		-0.605*** (0.0722)	-0.426*** (0.113)	-0.0681*** (0.0122)	-0.537*** (0.0683)
1[Age = 16]		-0.694*** (0.0765)	-0.456*** (0.119)	-0.0753*** (0.0130)	-0.619*** (0.0721)
Constant	8.847*** (0.00340)	8.864*** (0.0411)	8.802*** (0.0690)	-0.551*** (0.00701)	9.415*** (0.0390)
Observations	1,290,208	1,290,208	1,228,544	1,290,208	1,290,208
R-squared	0.192	0.200	0.695	0.789	0.391
Cohort FE	Yes	Yes	Yes	Yes	Yes
Sector&Time FE	Yes	Yes	No	Yes	Yes
Firm&Sector&Time FE	No	No	Yes	No	No
Sample	Balanced	Balanced	Balanced	Balanced	Balanced
Products	UPC	UPC	UPC	UPC	UPC

Notes: The table presents the coefficients for the age fixed effects of OLS regressions. The dependent variable in columns 1–3 is sales (in logs), in column 4 is price (in logs), and in column 5 it is quantity. Age is the number of quarters since we first observe sales for a product ($\mathbb{1}[Age = i]$ represents an indicator variable that equals one if the product is i quarters of age). Other controls include cohort variables (using Deaton’s normalization) and sector-quarter fixed effects except column 3 with firm \times sector \times quarter fixed effects. Sector refers to Nielsen’s module. The sample used in this table comprises all products in the baseline balanced sample that were born between 2006q3 and 2012q2 and their outcomes for 16 quarters. The ***, **, and * represent statistical significance at 1%, 5%, and 10% levels, respectively. Standard errors are clustered at product category level.

Figure 3: Sales over the Product Life Cycle: by Duration



Note: The figure shows the life cycles for products that lasted between 2 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. The estimation includes time effects that are specific to product modules and cohort effects.

Brands as products – A potential concern of defining products as barcodes is that their life cycle is short because many products are superseded by very similar products that represent minor changes to their predecessor but nevertheless originate a new barcode. We address this concern by defining products as brands instead of barcodes using the brand information provided by Nielsen. The average brand in our data contains nine different barcodes. Panel (a) of Figure 4 shows that the evolution of brand sales over time is very similar to that of barcodes in that sales also decline throughout most of their life cycles. The most noticeable difference is a slightly larger increase in sales during their first year and a smaller decline thereafter (an average of 25% per year). Barcodes within a brand have different attributes; thus, comparing a brand at different points in time could reflect changes in its composition that result from the entry and exit of barcodes within the brand. The smaller decline in sales throughout the life cycle of a brand is consistent with firms renewing the brand with new barcodes that do not fully compensate for the declining sales in the existing barcodes of the brand. This pattern of declining brand sales also indicates that the decline in sales of existing barcodes is unlikely to be fully explained by the entry of new barcodes that cannibalize their sales.

Novel products – Another potential approach to ensure that our results are not driven by products that represent minor improvements over those already available in the market is to use information about the characteristics of each product to identify those that have novel features. We use a similar approach to that of [Argente and Yeh \(2017\)](#) to construct a *novelty*

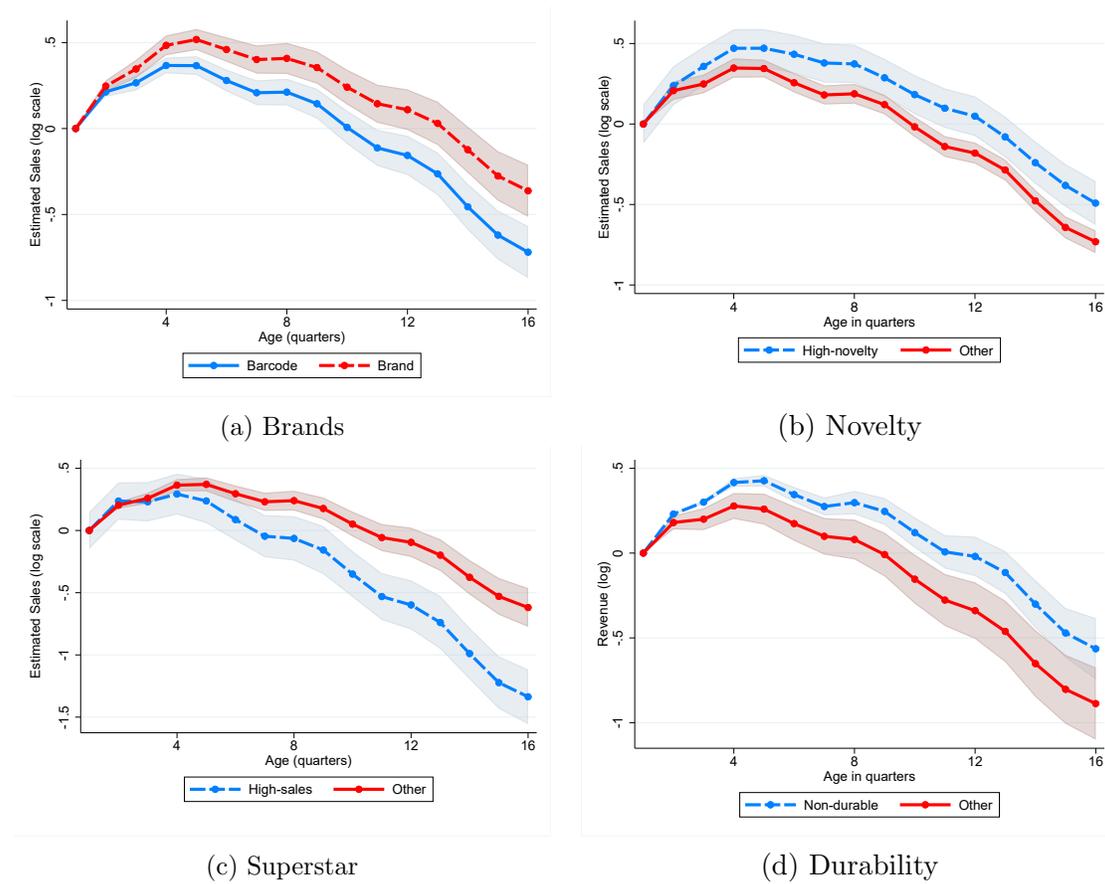
index based on the detailed information about the characteristics of each UPC provided in the Nielsen RMS data set. The index counts the number of characteristics of a product that are new and unique relative to those of all other products ever sold within the same sector. Specifically, we define a product u in sector j as a vector of characteristics $V_{uj} = [v_{u1}, v_{u2}, \dots, v_{uA_j}]$ in which A_j denotes the number of attributes we observe in sector j . For example, we observe $A_{\text{soft drinks}} = 8$ attributes for each barcode in the sector “soft drinks - carbonated”: brand, flavor, firm, size, type (sparkling soda or natural soda), container (e.g., can or bottle), formula, generic (i.e. private label). Let Ω_{jt} contain the set of product characteristics for each product ever sold in j at time t . The *novelty index* of a product u in j that is launched at time t is defined as: $I_{ut}^j = \frac{1}{A_j} \sum_{a=1}^{A_j} \mathbf{1}[v_{ua} \notin \Omega_{jt}]$. For instance, if a new product within the soft drinks category enters with a flavor and size that has never been sold in any store before, its novelty index is $2/8$.¹⁶ We classify products as “High-novelty” if their novelty index is in the top quartile of the distribution of the index and classify products as “Other” if their novelty index is below the top quartile. Panel (b) in Figure 4 presents the estimated age fixed effects for these two groups of products. The figure shows that the product life cycle patterns of both groups is similar. Products that bring new attributes to the market experience a slightly larger increase in sales during the first year but the sales of both groups continuously drop thereafter. The slightly more favorable life cycle of novel products may reflect the fact that these products are less substitutable than those already available in the market.

Superstar products – We further explore heterogeneity across products by studying the life cycle of the most successful new products.¹⁷ For each sector within a cohort (measured by the year of entry), we classify products as superstars if they are “High-sales”, i.e. if their sales in the first year of activity is in the top decile of the distribution of sales of their respective sector and cohort. By contrast, we classify products as “Other” if their sales is below the top decile. Superstar products generate approximately four times the sales of “Other” products at entry, approximately nine times the sales of “Other” products over their entire lifecycle, and they also last an additional year in the market relative to “Other” products. Panel (c) in Figure 4 shows the estimated age-fixed effects for “High-sales” and “Other”. The plot shows that even for very successful products, the growth phase of a product’s life cycle is on average short, as sales peak at around one year after entry. Perhaps more surprising, we find a steeper decline in sales among superstar products. After the first year in the market, sales of superstar products

¹⁶On average, we observe 7.2 product attributes in each product module. The minimum attributes we observe for each sector is 5 and the maximum is 12. Online Appendix C.4.5 has more details on the construction of the index, summary statistics, and alternative indexes we constructed for robustness.

¹⁷The recent literature has highlighted the relevance of superstar firms, which are the largest firms in the economy in terms of sales and employment, and has contrasted their behavior relative to the rest of the firms in the economy (e.g., Autor, Dorn, Katz, Patterson and Reenen, 2020).

Figure 4: Sales over the Product Life Cycle: Brands and Heterogeneity



Notes: The figure shows the estimated age fixed effects of sales over the life cycles of products using equation (2). Panel (a) uses two distinct definitions of products (barcodes and brands) and in Panels (b)-(d) each line is estimated by dividing products according to the characteristics novelty, superstar, and durability. In Panel (a), the blue line depicts the estimates for barcodes (equivalent Figure 2). The red line depicts the estimates for products defined as brands \times product modules. Panel (b) presents the estimated age fixed effects for sales for “High-novelty” and “Other” products. We classify products as “High-novelty” if their novelty index is in the top quartile of the distribution of the index. “Other” refers to products with novelty index below the top quartile. Panel (c) presents the estimated age fixed effects for sales for “High-sales” and “Other” products. To classify products in these groups, we measure the sales of new products in the first year of activity (summing the first four full quarters of sales). Within a sector and within a cohort of products (measured as year of entry), we classify products as “High-sales” if their sales in the first year of activity is in the top decile of the sales distribution. “Other” refers to products with sales below the top decile. Panel (d) presents the age fixed effects for sales for products divided according to their durability. We classify products as “Non-durable” if their durability is in the top quartile of the distribution. “Other” refers to products with durability below the top quartile. Durability is determined at the level of the product module. We approximate the durability of each sector by using the Nielsen Household Consumer Panel Data, and count the average number of shopping trips made by households. In all regressions, the estimation includes time effects that are specific to product modules and cohort effects. We keep a balanced sample with 16 quarters or above durations. Standard errors are clustered at product category level.

decline almost twice as much than non-superstar products. Superstar products may have a steeper decline in sales because these products may be more affected by fiercer competition from new products introduced in the market.

Nondurables and semi-durables – The research on marketing and industrial organiza-

tion has documented that sales decline over most of the life cycle of products for specific durable goods such as 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 entry of more up-to-date products drive down the sales of existing products in the durable goods markets. With our data, we study a broader set of products that vary in their durability. We approximate the durability at the sector level by using the Nielsen Consumer Panel Data to count the average number of shopping trips made by households in a given year to purchase products in that sector. We call sectors with few trips per year durable categories. Examples of durable categories are sun exposure trackers (1.00), bathroom scales (1.03), and printers (1.03), whose 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). Panel (d) in Figure 4 shows sales over the life cycle by durability. For both non-durable and more semi-durable products the growth phase of a product’s life cycle is on average short, and sales decline throughout most of the life cycle. Although sales decline faster for durable categories, we also find a large drop in sales for non-durable goods that indicates the patterns that we identify are common to a broader set of goods than those previously considered in the literature.

4.2 Margins Affecting the Product Life Cycle

In this section, we investigate the margins that affect the product’s life cycle. We start by using reduced form evidence to evaluate the role of firm and potential demand and supply forces at play, followed by a structural approach that decomposed sales into different margins.

Role of the firm – We consider alternative specifications that are useful in shedding light about potential reasons behind the decline in the sales of products over time associated with firm characteristics. These are also useful approaches to evaluate the robustness of our estimates of the product life cycle to different assumptions about the data generating process. In particular, we estimate the life cycle pattern of products after conditioning on firm-specific time-varying factors. Column (3) of Table 4 presents the estimated age fixed effects, when we include firm \times sector \times time fixed effects. The R-squared of this specification is substantially higher than that of the specification that only includes sector \times time fixed effects (column 2). This increase suggests that an important fraction of the variation in sales across products can be explained by firm specific factors. Yet, in spite of this increase in fraction of the variation in sales explained by this specification, the estimated age fixed effects follow a similar pattern in that they increase during the first year of activity and subsequently decline at a

very fast pace. Our results show that firm-specific factors, such as changes in competition or changes in firms' quality, likely play a small role in determining the trajectory of sales over the product's life cycle and points toward product-specific mechanisms as the potential drivers.

Prices and quantities – Next, we use prices and quantities to examine their separate contribution to the decline in sales. Columns (4) and (5) of Table 4 present the estimated age fixed effects for price and quantity. The estimation for prices as dependent variable shows that prices decline 2% a year on average. By the end of the fourth year of activity, prices are almost 8% lower than prices at entry (column (4) of Table 4). Thus, unlike the evolution of sales, prices decline slowly and steadily over the product's life cycle. Because our empirical specification conditions on aggregate effects (e.g., inflation) specific to particular sectors of products, this estimated decline already accounts for average fluctuations in prices. Column (5) of Table 4 show that the quantities sold begin to decline following the first year of activity. Our results show that between the first and fourth year of activity, quantities decline on average 28% per year. 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.¹⁸

Decomposition of product sales using structural approach – We use the framework developed by [Hottman, Redding and Weinstein \(2016\)](#) based on a model of oligopolistic competition between heterogeneous multi-product firms with rich product heterogeneity – in terms of demand and cost shifters. The model uses a simple setup that allows to decompose product sales into components.

Under the [Hottman, Redding and Weinstein \(2016\)](#) framework, sales of product u from firm i at time t can be expressed as the sum of the following margins (Online Appendix shows derivation):

$$\log Y_{uit} = \underbrace{(\eta - 1) \log \frac{\gamma_{uit}}{z_{uit}}}_{\text{appeal-to-cost}} \underbrace{- (\eta - 1) \log \mu_{it}}_{\text{markup}} \underbrace{+ (\eta - 1) \log p_t}_{\text{sector price}} \underbrace{+ \log Y_t}_{\text{sector size}} \\ \underbrace{- (\sigma - \eta) \log \frac{\tilde{\gamma}_{it}}{\tilde{z}_{it}}}_{\text{appeal-to-cost cannibalization}} \underbrace{- \frac{\sigma - \eta}{\sigma - 1} \log N_{it}}_{\text{scope cannibalization}} \underbrace{- \frac{\sigma - \eta}{\sigma - 1} \log \left(\frac{1}{N_{it}} \sum_{k=1}^{N_{it}} \left(\frac{\gamma_{ki}}{\tilde{z}_{it}} \right)^{\sigma-1} \right)}_{\text{dispersion cannibalization}} \quad (4)$$

in which $\tilde{\gamma}_{uit}$ and \tilde{z}_{uit} are the demand and marginal cost shifter of the product, respectively. The variables $\tilde{\gamma}_{it}$ and \tilde{z}_{it} represents the geometric mean of the appeal of all the firm's products and geometric mean of all their marginal costs, respectively. The variable N_{it} is the number of

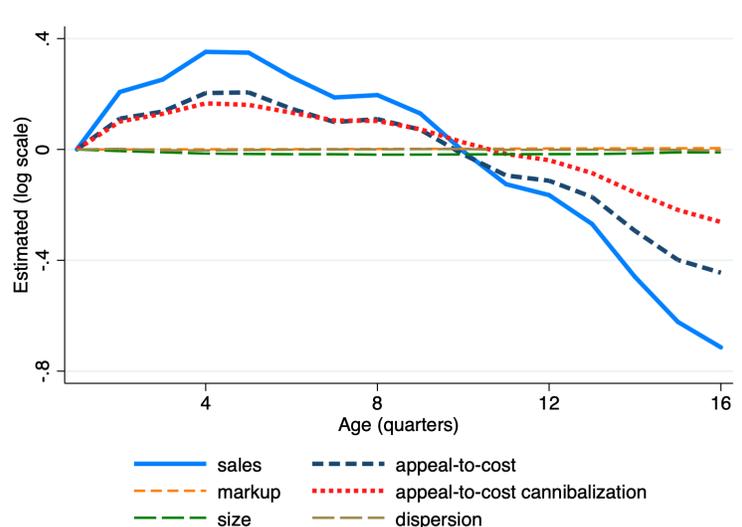
¹⁸Online Appendix C presents the estimated age fixed effects for price and quantities for the different types of products discussed above.

products of the firm and μ_{it} is the markup of the firm. The first term captures the importance of product-specific characteristics on sales. In this stylized framework, the profitability of a specific product is fully characterized by the ratio of appeal to marginal cost. We refer to this component as appeal-to-cost, and later we further decompose it into the contributions of appeal and marginal costs. Holding everything else constant, a 1% increase in the appeal (or reduction in marginal cost) increases sales by $(\eta - 1)\%$. An increase in the elasticity of substitution between firms, η , makes this effect stronger because sales react more to an increase in product appeal (or decline in marginal cost) when consumers easily substitute across firms. The second component captures the effect of the markup on sales. The third component refers to the effect of the sector's price index 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 fourth term captures the effect of the sector's size. Product sales increase one-to-one with the aggregate expenditures in a sector (due to homogeneity of degree 1). The fifth component (appeal-to-cost cannibalization) captures product-specific cannibalization that results from changes in the ratio of the appeal-to-cost of the product relative to the average appeal-to-cost of the firm. The effects of the appeal (cost) cannibalization are negative (positive) when products are more substitutable within firms than across firms ($\sigma > \eta$). This component varies with changes in the appeal and marginal cost of other products of the firm and with changes in a product's own appeal and marginal cost. The sixth component (scope cannibalization) captures the cannibalization due to changes in product scope. This component measures the effect that the number of products sold by a firm has on sales of each of their products. The last component, dispersion cannibalization, is a function of the ratio of the power mean with exponent $\sigma - 1$ to the geometric mean of the appeal to marginal cost ratio. This ratio is a measure of entropy that captures the dispersion in the appeal to marginal cost ratio of a firm's products relative to their average appeal-to-cost ratio.¹⁹ All cannibalization effects are negative when products are more substitutable within firms than across firms ($\sigma > \eta$).

These seven components capture the most direct margins through which products differ in sales. The first four components govern the reallocation of products across firms and, therefore, capture business stealing effects. The last three components capture sources of variation specific to the multi-product firm case and reflect the economic margins that induce substitutability within the set of products of a firm. Unlike all other components in the decomposition, the appeal-to-cost and appeal-to-cost cannibalization are product specific. The distinction between them is, however, very important. Suppose that a product experiences a

¹⁹For example, consider two firms that have the same number of products and same average (log) appeal and (log) cost, but differ in that one of them supplies all products with the same appeal and cost while the other has dispersed appeals and costs across its products. The latter will be larger because it is able to supply its production bundle more cheaply because it will shift resources to products with higher appeal-to-cost ratios.

Figure 5: Evolution of the Components of Product Sales



Notes: This figure shows the components of the product decomposition over the product’s life cycles. The solid line plots the estimated age fixed effects 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 2). Likewise, the dashed lines reflect 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 appeal-to-cost, the markup, the size cannibalization, and the appeal-to-cost cannibalization) is weighted by combinations of the elasticities of substitution σ_j and/or η_j as derived in equation (4). The estimated age fixed effects of all four components add up to the estimated age effects for sales.

decrease (increase) of 1% in its appeal (marginal cost), while the appeal (marginal cost) of all other products of the firm and its competitors are kept constant. In this case, the sales of the product will experience both a decline in sales of $(\eta - 1)\%$ because consumers substitute toward relatively more attractive products of other firms and a decline in sales of $(\sigma - \eta)\%$ because consumers will also substitute towards relative more attractive products within the firm. We refer to the former effect as business stealing and to the latter effect as cannibalization. By contrast, if the increase in appeal of a product is matched by a proportional increase in the appeal of all other products of the firm, there is no appeal-to-cost cannibalization, and only appeal-to-cost business stealing effect. These two components, therefore, capture distinct and independent events affecting product and firm growth.

To understand how each of these components evolves over the life cycles of products, we repeat the estimation of the specification of equation (2) by using each component of the decomposition as a dependent variable.²⁰ The decomposition relies on parameters that are unobserved but can be recovered from our rich product data under some assumptions. Online Appendix E.2 provides details on the structural estimation of the elasticities and the

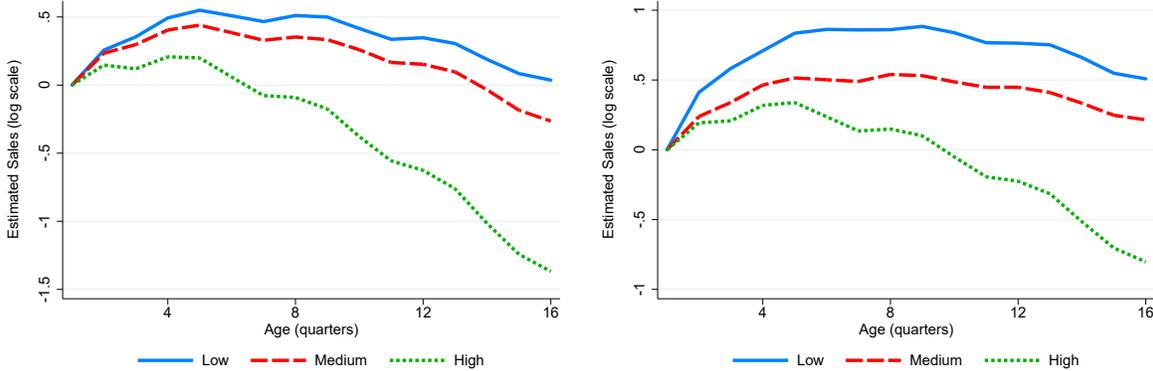
²⁰The sector’s price and size in the decomposition of product sales cannot explain the decline in the estimated age fixed effects because they are absorbed by the sector-time effects.

individual appeal and cost components. Figure 5 illustrates the series of estimated age fixed effects for each component of the decomposition. The evolution in the product’s sales over its life cycle is largely explained by changes in two components: changes in the estimated appeal-to-cost and changes in the estimated appeal-to-cost cannibalization. In the first year of activity, sales increase almost 0.4 log points relative to entry, with the appeal-to-cost and appeal-to-cost cannibalization components contributing similarly for this increase. In the following three years, sales decline by about 1.1 log points. The appeal-to-cost components accounts for a decline of 0.65 log points, whereas the remainder is explained by declines in the appeal-to-cost cannibalization component. Changes in firm-specific components (markup, scope and dispersion) make only minor contributions to changes in product sales as products become older. Despite evidence of substantial cross-sectional differences in the levels of these components, this result indicates that changes in these components within firms over time are neither systematic nor large enough to meaningfully affect the evolution of sales of individual products over their life cycles. This pattern is in line with our empirical findings that the evolution of sales over a product’s life cycle is not significantly affected by the inclusion of firm-specific covariates (column 3 of Table 4).

Product Introduction and the Product Life Cycle – Our previous results have shown that the appeal of a product is the most important factor determining how sales of a product evolve over its life cycle. Our decomposition, however, takes changes in product appeal as given and is silent on the sources of such changes. We now examine the conjecture that when a firm or its competitors introduce new products, the existing products become relatively less attractive, which results in a decline in their appeal.

We exploiting variation in product entry rates across sectors to show that products in sectors with high product entry rates experience faster declines in sales. We partition the sectors into three groups according to their average quarterly entry rate (weighted by sales). Panel (a) of Figure 6 shows substantial differences in the life cycle of products sold in sectors with low entry rates. The plot shows that the decline in sales of products in this group lags that of products sold in sectors with high entry rates by almost a year and that the pace of their decline in sales is much slower. These patterns are consistent with the conjecture that maturing products lose sales as they face new rival products that increasingly steal business. The frequent adoption of varieties by other firms accelerates the rate at which existing products become obsolete. While consistent with our conjecture, this correlation between decline in sales of existing products and the entry of new products, could be explained by confounding factors associated with heterogeneity in the nature of the products across sectors (e.g. some sectors are more subject to changes in fashion than others). Therefore, we see these patterns as simply indicative that the intensity of adoption of new varieties at a sectoral level is linked

Figure 6: Heterogeneity by Product Introduction Rates



(a) Sector's Product Introduction

(b) Firm's Product Introduction

Notes: Panel (a) shows the product life cycle after splitting the sectors in the Nielsen data by their average quarterly product introduction rates (weighted by sales). Panel (b) shows the product's life cycle after splitting firms by the average share of new products in a given group and quarter. In each of the figures, the estimated age fixed effects of sales are computed by using equation (2). The estimation includes time effects that are specific to product modules and cohort effects. We keep a balanced sample with 16 quarters or above durations. Product groups with higher entry rates have on average higher sales. Therefore, the average life cycle is closer to the product's life cycle of high entry category in this figure.

to the evolution of the appeal of existing products which, in turn, is the most important factor determining a product's life cycle.

We also estimate that, consistent with appeal cannibalization, the decline in sales of existing products is even more pronounced when their respective firms introduce more products within their own sectors (irrespective of their competitors' introduction rates). We compute the sales-weighted share of new products in each sector for each firm and we partition firms into three groups according to the level of this share. With this procedure, we compare, within the same sector, the life cycles of products supplied by firms that are very active in introducing new products against those that are not. Panel (b) of Figure 6 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 introducing new products start declining sooner. This steeper decline indicates that when firms introduce more new products, those new products have on average higher appeal and they cannibalize the sales of their own existing products.²¹

²¹A potential weakness of the evidence presented above is that confounding factors, such as firm-specific demand shifts, could affect both the life cycles of existing products and the decisions of firms to introduce new products.

5 Dynamic Model

Thus far, we have been focused on documenting that sales of products decline throughout most of the life cycle, and on understanding the margins that drive this decline. We have shown that the introduction of new products makes existing products obsolete by reducing their appeal and sales. In this section, we show that these findings are indissociable from the margins that affect a firm’s innovation decisions and aggregate growth. We build a dynamic model of endogenous firm growth that features the forces of innovation and creative destruction. The model allows us to study and quantify how firms’ anticipation of obsolescence effects on existing products affect decisions about innovation. In the model, firms can invest both in new products that improve on their existing products (*internal innovation*) and in new products that improve on competitors’ products (*external innovation*). Both internal and external innovation decisions are endogenous and interdependent.

Our model has two crucial differences relative to other models that include both forces of creative destruction and internal innovation (e.g., [Akcigit and Kerr, 2018](#); [Peters and Walsh, 2021](#)). First, our model features partial substitutability between products such that new products do not fully substitute for existing products. We introduce this feature because our empirical results suggest that when a firm or a competitor introduces new product, the now-obsolete existing product does not immediately exit but rather gradually loses appeal over time. Second, the decision to introduce new products depends on the performance of existing products, which means that while innovation induces obsolescence, obsolescence also affects innovation decisions. By contrast, other models typically do not take the rich interconnection between internal innovation and creative destruction forces into account.

5.1 Static Equilibrium

The model economy is in continuous time and is populated with representative households with preferences

$$U = \int_0^{\infty} e^{-\rho t} \ln c_t dt$$

where $\rho > 0$ is the discount factor and c_t is the CES consumption aggregate of appeal-weighted continuous products, given by

$$c_t = \left(\int_0^1 (q_{ut} c_{ut})^{\frac{\sigma-1}{\sigma}} du \right)^{\frac{\sigma}{\sigma-1}}$$

where c_{ut} represents quantity and q_{ut} represents the appeal of product u . $\sigma > 1$ is the elasticity of substitution between products. We follow the convention in the endogenous growth literature of having a mass of identical individuals L_t who each inelastically supply one unit

of labor. We assume that the each product is produced according to a linear technology given by $y_{ut} = l_{ut}$, where labor is the only factor of production and l_{ut} represents the number of employees used in the production of u . Each product u is produced by the firm that introduced it to the market.

At any point in time, there are two different sets of firms: (i) a set of incumbent firms of measure K_t that own at least one product and (ii) a set of potential entrants of measure one that currently do not sell any products. Incumbent firms can sell multiple products simultaneously. We denote the set of the appeals of products produced by firm i as $[Q_{it}] \equiv \{q_{i1t}, \dots, q_{int}\}$, with Ω_{it} representing the set of active products and n_{it} is the cardinality of this set. Across all firms, the quality index of the economy is given by

$$Q_t \equiv \left(\int_0^1 (q_{ut})^{\sigma-1} du \right)^{\frac{1}{\sigma-1}}$$

Firms face a common wage in a competitive labor market W_t .²² Given this monopolistic competitive market environment, solving the consumer and firm problem leads to the following

$$y_{ut} = q_{ut}^{\sigma-1} \left(\frac{p_{ut}}{P_t} \right)^{-\sigma} Y_t$$

$$p_{ut} = \frac{\sigma}{\sigma-1} W_t$$

The market share of each product and profits are determined respectively by

$$s(q_{ut}) = \left(\frac{q_{ut}}{Q_t} \right)^{\sigma-1}$$

$$\Pi(q_{ut}) = q_{ut}^{\sigma-1} \pi_t$$

where π_t is an aggregate profit shifter given by $\pi_t \equiv \sigma^{-1} Q_t^{2-\sigma} L_t$.

5.2 Product Introduction

Incumbent firms can improve upon existing products within their portfolios or improve upon competitors' products. We adopt a stochastic formulation whereby firms choose the flow rates of creating new products that improve upon existing products.

Firms improve on their own products by choosing a Poisson rate x_{ut}^I that determines the rate of arrival of new products. Conditional on the arrival of a new product, the firm gains a new product with appeal $q_{u'}(t+) = q_{ut} + \lambda^I q_{ut}$ (where $\lambda^I > 0$ is given) and retains an old existing product whose sales are then cannibalized by the new product. We model

²²Without loss of generality, we normalize the wage rate to be $W_t = Q_t$.

cannibalization by replacing the existing product u by a new product with a lower level of appeal determined by $q_u(t+) = q_{ut} - \lambda^C q_{ut}$ (with $\lambda^C > 0$). We refer to λ^C as the cannibalization step size.

Opportunities to replace existing products of competitors arrive at rate x_t^E . External innovation is undirected in the sense that any resulting innovation is realized in any product with equal probability.²³ Conditional on the introduction of a new product whose sales affect an existing competitor's product, firms gain a new product with appeal $q_u'(t+) = q_{ut} + \lambda^E \bar{q}_t$ (where $\lambda^E > 0$ is given, and \bar{q}_t is the average appeal across the economy) and make the competitor's product partially obsolete. Creative destruction is partial, as new products do not completely replace existing products right away. We model business stealing by replacing the competitor's product u with a new product of lower appeal that is given by $q_u(t+) = q_{ut} - \lambda^S q_{ut}$, with $\lambda^S > 0$. We refer to λ^S as the business-stealing step size. Product introductions impact existing products at rate τ_t , the total rate of external innovation. Existing products can also exit the market if they receive an exogenous exit shock with probability $\psi_t > 0$.

Internal and external innovation are both costly activities. In their internal innovation activities, firms spend resources that depend on the quality of a new product. The cost of internal innovation is determined by

$$c^I(x_{ut}^I, q_{ut}, Q_t) = \xi^I (x_{ut}^I)^{\frac{1}{1-\alpha}} q_{ut}^{\sigma-1} \pi_t + F_t x_{ut}^I (Q_t^{\sigma-1} - f q_{ut}^{\sigma-1})$$

This cost function has two components. The first component increases with the likelihood of internal product improvement, x_{ut}^I , and with the profit of the existing product, $q_{ut}^{\sigma-1} \pi_t$. The exogenous cost shifter $\xi^I > 0$ governs the importance of this term and $\alpha < 1$ is the elasticity of the cost to the Poisson rate of the arrival of new product improvements. Importantly, this component of cost increases in the appeal of the existing product, q_{ut} , in order to reflect the idea that generating and implementing improvements upon better products will be more costly (from an R&D standpoint) than improving upon less successful products (e.g., [Bloom, Jones, Van Reenen and Webb, 2020](#)).

The second component of the internal innovation cost is linear with the likelihood of product improvement, x_{ut}^I , and also depends on the difference between the product's quality, q_{ut} , and aggregate quality, Q_t . F_t governs the importance of this term and is determined in equilibrium while $f \geq 0$ is a parameter that determines the quality threshold below which this component is positive. When $f > 0$, this cost component declines as the appeal of the product increases. This term implies that improving on a product whose appeal is far below

²³This assumption has two main implications. First, business stealing is just as likely for low and high quality products, and thus external innovation impacts the average quality product. Second, a firm does not innovate over their own products through external innovations.

the aggregate quality in the market is more costly than improving a product whose appeal is closer to the aggregate quality of the products in the market. The intuition is that the overall cost of investing in improving on products that are under-performing relative to other products in the market is higher than the cost of investing in improvements to products that have proven to have consumer appeal. For instance, this term could reflect the idea that retailers and distributors will be reluctant to market a new product that replaces a product with very low appeal in their customer bases and that firms will have to spend additional resources to effectively promote and launch these new products.

The cost of external innovation is assumed to be

$$c^E(x_t^E, Q_t) = \xi^E (x_t^E)^{\frac{1}{1-\alpha}} Q_t^{\sigma-1} \pi_t$$

where ξ^E is an exogenous cost shifter and governs the importance of this cost. This cost is increasing in both the likelihood of external product improvement, x_t^E , the aggregate quality in the economy, Q_t , and the aggregate profit shifter, π_t .

Firm's decisions are forward-looking, and the value function of incumbent firm i is

$$\begin{aligned} rV_t([Q_{it}]) - \dot{V}_t([Q_{it}]) &= \sum_{u \in \Omega_{it}} \Pi(q_{ut}) \\ &+ \max_{x_{ut}^I} \sum_{u \in \Omega_{it}} \left[x_{ut}^I \left(V_t([Q_{it}] \setminus \{q_{ut}\} \cup \{q_{ut} - \lambda^C q_{ut}\} \cup \{q_{ut} + \lambda^I q_{ut}\}) - V_t([Q_{it}]) \right) \right. \\ &\quad \left. - \xi^I (x_{ut}^I)^{\frac{1}{1-\alpha}} q_{ut}^{\sigma-1} \pi_t - F_t x_{ut}^I (Q_t^{\sigma-1} - f q_{ut}^{\sigma-1}) \right] \\ &+ \max_{x_t^E} \sum_{u \in \Omega_{it}} \left[x_t^E \left(\mathbb{E}_{u'} V_t([Q_{it}] \cup \{q_{u'} + \lambda^E \bar{q}\}) - V_t([Q_{it}]) \right) - \xi^E (x_t^E)^{\frac{1}{1-\alpha}} Q_t^{\sigma-1} \pi_t \right] \\ &+ \sum_{u \in \Omega_{it}} \tau_t \left(V_t([Q_{it}] \setminus \{q_{ut}\} \cup \{q_{ut} - \lambda^S q_{ut}\}) - V_t([Q_{it}]) \right) \\ &+ \sum_{u \in \Omega_{it}} \psi_t \left(V_t([Q_{it}] \setminus \{q_{ut}\}) - V_t([Q_{it}]) \right) \end{aligned} \quad (5)$$

The value of a firm (net of capital gain in case the value function increases over time) consists of multiple additively separable parts. First, the value of the firm is increased by the current flow profits, which is simply the sum of profits across all products. The second part is the change in firm value after internal innovation and the corresponding research costs. The term $V_t([Q_{it}] \setminus \{q_{ut}\} \cup \{q_{ut} - \lambda^C q_{ut}\} \cup \{q_{ut} + \lambda^I q_{ut}\})$ represent the firm value of two products: the old product whose appeal was reduced by size λ^C and the new product that improves on product u by size λ^I . Firms innovate on each existing product separately. The third part shows the expected change in firm value following a successful external innovation, $\mathbb{E}_{u'} V_t([Q_{it}] \cup \{q_{u'} + \lambda^E \bar{q}\})$, which is the net research cost. Note that the expectation is about the level of appeal of

the product that was improved upon and that the external innovation is scaled to the number of products of the firm n_{it} . The fourth part is the reduction to the value of the firm when a competitor has created a more appealing version of one of the firm's products. The rate of creative destruction from competitors is determined endogenously but is taken as given by incumbent firms. The last part of the value function corresponds to the value of the firm if any of its n_{it} products suffers a random exit shock.

The economy also has potential entrant firms of measure one. Potential entrant firms have the same opportunities as incumbents. Entrants do not currently sell any products that they can improve upon but they do engage in costly external innovation. Innovation by entrants is also undirected in the sense that any resulting innovation is realized in any existing product with equal probability, and can only choose a Poisson rate x_t^N that determines the rate of arrival of a new product that improves upon a random existing product. Conditional on successfully creating a product that replaces that of an incumbent, the quality of the new product to the firm is deterministically defined as $q_u(t+) = q_{ut} + \lambda^N \bar{q}$.

The cost of external innovation by entrants is similar to the cost of external innovation for an incumbent firm

$$c^N(x_t^N, Q_t) = \xi^N (x_t^N)^{\frac{1}{1-\alpha}} Q_t^{\sigma-1} \pi_t$$

where ξ^N is an exogenous cost shifter that governs the importance of this cost.

Potential entrant's decisions are forward-looking, and the value function for entrants V_t^N can be expressed as

$$r_t V_t^N - \dot{V}_t^N = \max_{x_t^N} \left[x_t^N \left(\mathbb{E}_{u'} V_t(q_{u'} + \lambda^N \bar{q}) - V_t^N \right) - \xi^N (x_t^N)^{\frac{1}{1-\alpha}} Q_t^{\sigma-1} \pi_t \right] \quad (6)$$

where $\mathbb{E}_{u'} V_t(q_{u'} + \lambda^N \bar{q})$ is the expected value of a new product that improves upon a random incumbent's existing product.

5.3 Dynamic Equilibrium

We now characterize the Markov-perfect equilibria of the economy that make strategies a function solely of payoff-relevant states. We focus on the steady state in which aggregate variables grow at a constant rate. We start by solving for the optimal innovation rates and by identifying a simple closed-form solution that follows from functional-form specifications of the internal and external innovation costs and some additional assumptions.

Proposition 1 *Consider the value function $V_t([Q_{it}])$ given in (5). $V_t([Q_{it}])$ is given by $V_t([Q_{it}]) = \sum_{u=1}^{n_{it}} \Gamma_t(q_{ut})$. Under the assumption that the cost-shifter of internal innovation cost $F_t = B_t$,*

we can write $\Gamma_t(q_{ut})$ as

$$\Gamma_t(q_{ut}) = A_t q_{ut}^{\sigma-1} + B_t Q_t^{\sigma-1}$$

with

$$A_t = \frac{\pi_t + \Lambda_t^I}{r - g_A + \tau[1 - (1 - \lambda^S)^{\sigma-1}] + \psi} \quad (7)$$

$$\Lambda_t^I = \xi^I \frac{\alpha}{1 - \alpha} (x_t^I)^{\frac{1}{1-\alpha}} \pi_t \quad (8)$$

$$x_t^I = \left[\frac{1 - \alpha}{\xi^I} \left(\frac{A_t}{\pi_t} [(1 - \lambda^C)^{\sigma-1} + (1 + \lambda^I)^{\sigma-1} - 1] + \frac{B_t}{\pi_t} f \right) \right]^{\frac{1-\alpha}{\alpha}} \quad (9)$$

and

$$B_t = \frac{\Lambda_t^E}{r - g_B + \psi} \quad (10)$$

$$\Lambda_t^E = \xi^E \frac{\alpha}{1 - \alpha} (x_t^E)^{\frac{1}{1-\alpha}} \pi_t \quad (11)$$

$$x_t^E = \left[\frac{1 - \alpha}{\xi^E} \left(\frac{A_t}{\pi_t} (1 + \lambda^E)^{\sigma-1} + \frac{B_t}{\pi_t} \right) \right]^{\frac{1-\alpha}{\alpha}} \quad (12)$$

Proof: See Online Appendix F.1.

Proposition (1) includes three important results. First, the incumbent's value function (5) is additively separable with respect to products. Second, the value is itself the sum of two components: (i) the present discounted value of the flow profits and the option value of internal innovation; (ii) the present discounted option associated with external innovation. Third, the optimal internal and external innovation rates do not depend on the levels of q_{ut} and under the condition that A_t , B_t and π_t grow at the same pace, they do not change with Q_t .

The condition that the second component of the cost of internal innovation satisfies $F_t = B_t$ is necessary to guarantee perfect scaling of the value function, and renders the firm problem tractable (Klette and Kortum, 2004). Intuitively, this condition relates the opportunity cost of investing in a firms' own products to the option value associated with external innovation.²⁴ Changes in this opportunity cost are, however, heterogeneous by product if $f > 0$; otherwise, opportunity costs are homogeneous. When $f > 0$, products that have relatively low appeal have a higher opportunity cost than products that are performing well in the market. Equation (9) shows that when $f \neq 0$, the rate of internal innovation depends on both the present discounted option associated with internal innovation (weighted by $(1 - \lambda^C)^{\sigma-1} + (1 + \lambda^I)^{\sigma-1} - 1$,

²⁴In other words, $F = B$ and $f > 0$ firms internalize the change in their innovation capacity when they add a new product to their portfolio; by owning a product line, the firm has a franchise value of extending into new product lines through external innovations. Thus, allowing for interdependence between internal and external innovation.

which represents the new appeal growth) and with external innovation (weighted by f). This feature is attractive as it captures the interdependence of external and internal innovations. For example, when the cannibalization step λ^C is substantially larger than the internal innovation step size λ^I and f is high, the rate of internal innovation can be quite large. This result confirms the intuition that a firm may want to continue to improve upon its own products to reduce the future opportunity cost of improving the products further. If firms do not improve their own products, the cost of future improvement will be even higher as the products will lag further behind the aggregate quality Q that is growing with external innovation (as captured by B). Another way to see this interdependence is by noticing that internal innovation can respond directly to the costs of external innovation. As the cost of external innovation decreases, external innovation rates increase and internal innovation can also increase.

We solve the entrants problem to determine the arrival rate of external innovation by potential entering firms x_t^N . The value function for entrants can be expressed as $V_t^N = C_t Q_t^{\sigma-1}$, where

$$C_t = \frac{\Lambda_t^N}{r - g_C} \quad (13)$$

$$\Lambda_t^N = \xi^N \frac{\alpha}{1 - \alpha} (x_t^N)^{\frac{1}{1-\alpha}} \pi_t \quad (14)$$

$$x_t^N = \left[\frac{1 - \alpha}{\xi^N} \left(\frac{A_t}{\pi_t} (1 + \lambda^N)^{\sigma-1} + \frac{B_t}{\pi_t} - \frac{C_t}{\pi_t} \right) \right]^{\frac{1-\alpha}{\alpha}} \quad (15)$$

The value function of entrants is similar to the present discounted option associated with external innovation from incumbents, with potential differences arising from differences in the innovation steps and innovation costs, and the discount factor. This reflects that the forces that determine innovation by entrants are similar to those that determine external innovation by incumbents.

After solving for the value functions of the incumbent and entrant, we characterize the equilibrium. Along the balanced growth path (BGP), the equilibrium growth rate is constant.

Proposition 2 *On the balanced growth path (BGP), the following conditions hold:*

(i) A_t and B_t grow at the same pace as π_t

$$g_A = g_B = g_C = g_\pi = (2 - \sigma)g_Q \quad (16)$$

(ii) *The number of products available to consumers is constant, and thus*

$$\psi = x^I + \tau \quad , \quad \tau = x^E + x^N \quad (17)$$

(iii) The aggregate growth rate is

$$\begin{aligned}
g_Q &= x^I \left[\frac{(1 + \lambda^I)^{\sigma-1} + (1 - \lambda^C)^{\sigma-1} - 2}{\sigma - 1} \right] \\
&+ x^E \left[\frac{(1 + \lambda^E)^{\sigma-1} + (1 - \lambda^S)^{\sigma-1} - 2}{\sigma - 1} \right] \\
&+ x^N \left[\frac{(1 + \lambda^N)^{\sigma-1} + (1 - \lambda^S)^{\sigma-1} - 2}{\sigma - 1} \right]
\end{aligned} \tag{18}$$

Proof: See Online Appendix F.1.

Proposition (2) highlights the key determinants of growth. Increases in the step size of internal (λ^I) and external innovation (λ^E , λ^N) will increase aggregate growth both directly and indirectly, as they increase internal and external innovation rates, respectively. Increases in the cannibalization step size λ^C decrease growth both directly and indirectly, as incentives toward creating new products that build on existing products, as in Arrow's replacement effect, are reduced. Increases in the business-stealing step size λ^S have a direct negative impact on growth rates. Importantly, the difference in the impacts of λ^C and λ^S is related to the fact that the innovating firm internalizes the impact of an internal innovation, but does not internalize the impact on its competitor's products when a new product is introduced.

Our model has a closed-form solution. By combining equations (7), (8), (9), (10), (11), (12), (13), (14), (15), replacing terms using conditions (16) and (17), and using equation (18), we can solve for $\{A, B, C, \Lambda^I, \Lambda^E, \Lambda^N, x^I, x^E, x^N, g_Q\}$. Then, we can express the value function for a product of any appeal level.

Next, we solve for the firm-size distribution. Our economy has a continuum of products of measure one. Suppose the measure of firms of size n (the number of products sold) is $\mu(n)$. Then, by definition, $\sum_{n=1}^{\infty} \mu(n)n = 1$. Define $\sum_{n=1}^{\infty} \mu(n) = K$. We determine the inflow and outflow of $\mu(n)$ for all n , which results in an explicit solution of the invariant distribution of product numbers.

Proposition 3 *The invariant distribution is given by*

$$\mu(n) = \frac{1}{n} \beta (1 - \beta)^{n-1} \quad , \quad \text{where} \quad \beta = \frac{x^N}{x^I + x^E + x^N} \tag{19}$$

and the measure of firms is given by

$$K = \frac{\beta}{1 - \beta} \ln \left(\frac{1}{\beta} \right)$$

Proof: See Online Appendix F.1.

The firm-size distribution indicates that as β increases, meaning that external innovation by entrants is relatively more intense, $\mu(n)$ decreases faster as n increases. This captures that when there is a lot of entry, fewer large firms will thrive and the economy is more competitive. Note that internal innovation forces also affect the size distribution because internal innovation creates one additional product and the firm increases its size by one. In such an economy, if we observe a mass of large firms, we should expect that the firms are innovating internally at an intense rate.

5.4 Quantitative Analysis

We estimate our model using the product-level data described in Sections 3 and 4 that have direct counterparts in the model. Section 5.4.1 describes our calibration procedure and describes the main results. Section 5.4.2 compare our quantified model against untargeted features of the data, and provides a characterization of the economy.

5.4.1 Calibration

Our model has 12 structural parameters $\{r, \alpha, \sigma, \lambda^I, \lambda^E, \lambda^N, \lambda^C, \lambda^S, \xi^I, \xi^E, \xi^N, f\}$. We identify these parameters in three ways. First, we fix two parameters (r, α) using standard values from the literature. We set the interest rate equal to 2 percent and the elasticity of the arrival of a product improvement (internal or external) to innovation costs to 0.5, which implies quadratic curvature. This is the standard elasticity used in the literature (see for example in Acemoglu, Akcigit, Alp, Bloom and Kerr, 2018; Akcigit and Kerr, 2018). We also fix the external innovation step size from entrants equal to the external innovation step size from incumbents, i.e. $\lambda^E = \lambda^N$.²⁵ Second, we use our estimated elasticities of substitution for the consumer-goods sector to determine σ . In particular, we use the value of 6 as a baseline which corresponds to the average elasticity pooling the elasticities both within firms and across firms in all sectors (Online Appendix E.2). Finally, for the remaining eight parameters, we calibrate the model to the product-level moments.

Our calibration procedure allows us to use the product-level data described in Sections 3 and 4 that have direct counterparts in the model. We choose eight moments that have closed-forms expressions in the theory. Table 5 summarizes the moments.

One crucial advantage of our calibration procedure is that we can use moments that allow us to distinguish new from existing products. We start by measuring the share of new products by incumbent and entrant firms (see Table 3 for details). Incumbents firms exhibit average product introduction rates of 14.5%, and new products created by new firms corresponds to

²⁵We relax this assumption in Online Appendix F.4

Table 5: Moments

Moment	Equation	Data	Model
Incumbents innovation rate	$x^I + x^E$ (equations 9 and 12)	0.1450	0.1450
Share of products by entrants	x^N (equation 15)	0.0108	0.0101
Product life cycle	$x^I(\sigma - 1)\log(1 - \lambda^C)$ $+ (x^E + x^N)(\sigma - 1)\log(1 - \lambda^S)$	-0.1506	-0.1492
Product life cycle decomposition	$(x^E + x^N)\log(1 - \lambda^S) / x^I\log(1 - \lambda^C)$	1.4858	1.4858
Incumbent sales growth	$x^I\left((1 + \lambda^I)^{\sigma-1} + (1 - \lambda^C)^{\sigma-1} - 2\right)$ $+ (x^E + x^N)\left((1 + \lambda^E)^{\sigma-1} + (1 - \lambda^S)^{\sigma-1} - 2\right)$	0.0160	0.0160
Entrants sales share	$x^N(1 + \lambda^E)^{\sigma-1}$	0.0114	0.0114
Ratio innovation costs to sales	$\xi^I\sigma^{-1}(x^I)^{\frac{1}{1-\alpha}} + Bx^I\sigma^{-1}\pi^{-1}(1 - f)$ $+ \xi^E\sigma^{-1}(x^E)^{\frac{1}{1-\alpha}} + \xi^N\sigma^{-1}(x^N)^{\frac{1}{1-\alpha}}$	0.0988	0.0976
Ratio costs of entrants to sales	$\xi^N\sigma^{-1}(x^N)^{\frac{1}{1-\alpha}}$	0.0115	0.0140

Notes: Table presents the moments used in the calibration algorithm described in detail in Online Appendix 5.4.1. The moments are estimated using data across all sectors, weighting sectors by their sales share. Sectors are defined according to Nielsen product groups. The moments are computed using data at the firm \times sector \times year level, for the period 2007–2014, as in Table 3. Entrants are defined from the first observation of sales of that firm \times sector in a particular sector.

about 1.1% of total products. These moments are central equilibrium objects that depend on multiple parameters.

Next, we use information on the performance of existing products, crucial to identify the obsolescence step sizes $\{\lambda^C, \lambda^S\}$. In particular, we match the expected average decline in log sales in the first year after a product reaches its maximum sales. In the model, the expected decline can be described as a closed-form solution of the cannibalization and business stealing step sizes and the external and internal innovation rates. In order to help identify the contribution of internal and external innovation for the decline, we use our product sales decomposition estimates to determine the relative contributions of business stealing and cannibalization, after controlling for time fixed-effects.

For the innovation step sizes $\{\lambda^I, \lambda^E\}$, we use information about the annual average growth rate of a firm’s total sales and the sales contribution of new products by entrants (Table 3). In the model, these moments have closed form solutions and are informative of the step size parameters for internal λ^I and external λ^E innovation, conditional on the obsolescence rates and equilibrium internal and external innovation rates.

Finally, we use data on product introduction costs to discipline the internal and external

cost shifters $\{\xi^I, \xi^E, \xi^N, f\}$ (together with the product innovation rates). Product introduction costs depend on R&D and marketing and advertising expenses, such as investments in the promotion of new products such as lump-sum costs for shelf space.²⁶ We obtain data from Compustat for the firms in the Nielsen dataset (as described in [Argente, Lee and Moreira \(2018\)](#)). We compute innovation expenses using data on R&D and general and administrative expenses. Overall, we find that our estimates of product-introduction costs of all firms represent about almost 10% of total sales, and about 1% of total sales are incurred by entrant firms.²⁷

In this exactly identified system (8 moments and 8 internally calibrated parameters), we put equal weights across eight target moments and find the parameters that minimize a sum of the percent deviations of simulated moments from target moments:

$$\min \sum_{i=1}^8 \frac{|\text{Model}(i) - \text{Data}(i)|}{|\text{Data}(i)|} \quad (20)$$

where $\text{Model}(i)$ is a simulated i -th moment and $\text{Data}(i)$ is a target value of i -th moment.

Table 5 reports the empirical and data moments and Table 6 reports the estimated estimated parameters. Overall, the model closely matches the targeted moments. The estimated parameters conform well with economic intuitions from theory, where the introduction of new products has sizable impacts on reducing the appeal of older existing products.

Our economy exhibits a rate of internal innovation of 12.4%, followed by a rate of external innovation by incumbents of 2.1% and a rate of external innovation by entrants of 1%. Our estimates indicate that the innovation steps from internal innovation net of the cannibalization obsolescence step is 0.023 and it is about the same as the innovation step from external innovation 0.024. Note that while firms internalize impacts on old products with internal innovation, they do not do the same when doing external innovation. External innovation has low probability but, conditional on impacting an existing product from a competitor, has a great effect on sales. Internal innovation is relatively more common, and cannibalizes upon old products relatively less. In line with the more radical impacts of external innovation, the incumbent’s external innovation cost shifter is about three times larger (and the entrant’s cost shifter is five times) than the internal-innovation cost shifter, which partially explains why external innovation is less common.

²⁶[Argente, Fitzgerald, Moreira and Priolo \(2021\)](#) study advertising spending in this sector and find that significant marketing expenses are involved with product introductions. Likewise, [Granja and Moreira \(2021\)](#) shows that product introduction often requires physical capital investments, where expansions into new product lines are associated with greater investments in building plant capacity.

²⁷We provide more details in the Online Appendix 5.4.1.

Table 6: Model Parameters

Parameter	Identification	Value
Baseline		
r Interest rate	External calibration	0.02
α Elasticity of new product to innov. costs	External calibration	0.5
σ Elasticity of substitution	Estimated from the data	6
Innovation Steps		
λ^I Internal	Internal calibration	0.114
λ^E External Incumbents	Internal calibration	0.024
λ^N External Entrants	External calibration	0.024
Obsolescence Steps		
λ^C Cannibalization	Internal calibration	0.092
λ^S Business Stealing	Internal calibration	0.438
Innovation Costs		
ξ^I Internal	Internal calibration	161.327
ξ^E External Incumbents	Internal calibration	459.667
ξ^N External Entrants	Internal calibration	817.894
f Opportunity Cost of Investment	Internal calibration	17.622

Notes: Table presents all calibrated parameters.

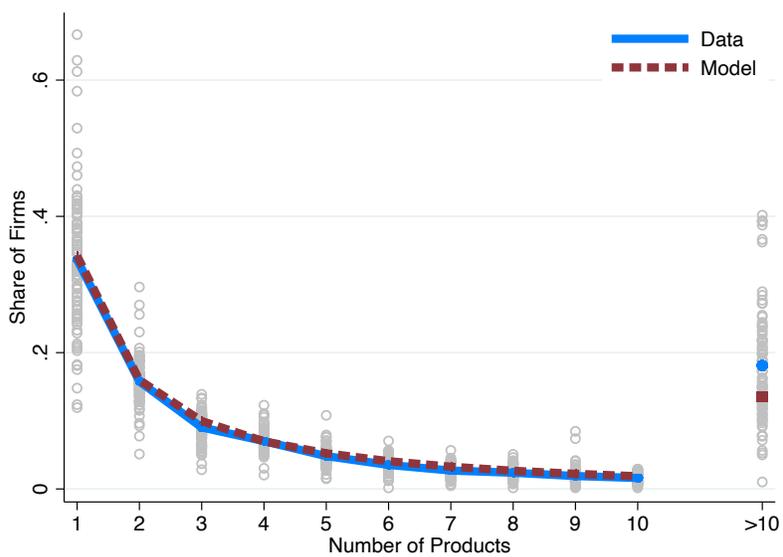
5.4.2 Untargeted moments and additional results

We next compare our quantified model against untargeted features of the data. We start by studying the properties of the distribution of firms and then various other moments using simulated data. While all the moments used in the calibration have a closed form solution, we produce several other non-targeted moments from a simulation exercise. We use these additional simulation-derived moments to evaluate the performance of the model. We use simulations covering 40 quarters, as our original data do, and we evaluate both product- and firm-level statistics.

Firm size distribution – Our model yields an analytical solution for the firm distribution. Proposition 3 shows that the distribution is dictated by the relative likelihood of external innovation by entrants $\beta = \frac{x^N}{x^I + x^E + x^N}$. Our calibration implies a β equal to 0.065, which indicates that the share of external innovation by entrants is small relative to incumbents’ innovation. Figure 7 shows the empirical and model distributions. The blue line shows the average share of firms by the number of products across all sectors, where each dot represents a same statistic for a particular sector in the data. We compare the data to the model and find that the model approximates well the empirical distribution of small firms, and account partially for a long tail of very large firms. In the data, about 18% of firms have more than 10 products, compared with 14% predicted by the model.

Through the lenses of our model, the fact that the data exhibits a low β has two important implications. First, firms seem to innovate more internally than externally. Nevertheless,

Figure 7: Share of Firms by the Number of Products



Notes: The hallow dots indicates the share of firms with that number of products in a particular sector, and the color-full dots plot the average and the predicted under θ equal to 0.05, 0.08, and 0.10.

business stealing makes a large contribution to declining sales of existing products, which indicates that while the likelihood of having a product impacted by a competitor or entrant is small, conditional on the event, the new competing product has a very negative impact on the sales of the existing product (i.e. λ^S should be large relative to λ^C , and indeed our results exactly that). The second implication of a lower β is the indication of the high levels of concentration in the firm-size distribution, as the top of that distribution is dominated by large firms that maintain their position by continuing to improve on their own existing products.²⁸

Product-level moments – We estimate the product life cycle with our simulated data. Our model estimation is able to account for more than half of the average yearly decline in product sales over the product life cycle (Online Appendix Figure F.1). It is not surprising that our model generates declining sales conditional on survival since our calibration directly targets the decline in sales after the first year a product is on the market. It is nevertheless reassuring that even in the absence of staggered entry-exit observations across markets and selection forces (at entry and exit), the model is able to generate a sizable decline in product sales as products grow obsolete. Online Appendix F.2 presents additional product-level statistics.

²⁸This finding is consistent with evidence from other industries. For instance, [Wollmann \(2018\)](#) shows that accounting for the incumbent automakers’ competitive decisions about whether to introduce new models is critically important when evaluating the effects of mergers in the auto industry.

Overall, the results of the model are consistent with these non-targeted moments.

Firm-level moments – We explore two main dimensions of the firm-level moments: heterogeneity across firm size and across the firm life cycle. In our model, firm size is determined by a combination of the number of products and the distribution of product quality.²⁹ In particular, we can express log sales as:

$$\log s_{it} = \underbrace{\log N_{it}}_{\text{Scope}} + \underbrace{\log \left(\frac{1}{N_{it}} \sum (q_{ut})^{\sigma-1} \right)}_{\text{Average sales per product (Appeal)}} + \underbrace{(2 - \sigma) \log Q_t + \log L_t}_{\text{Time FE}} \quad (21)$$

Firms manage their average appeal through decisions to reallocate resources among the products they sell.³⁰ Panel (a) of Figure 8 shows that both in our model and in the data, product appeal explains most differences in average size across firms. We use nonparametric regressions that compare variables across the firm size distribution. We include indicator variables by firm size decile, with the smallest firm size category serving as the reference group. We also include time-fixed effects. Our results show very large differences in product appeal across the firm-size deciles. While the magnitudes differ, both the model and the data point to large differences in firm size and to a prominent role for innovation in driving changes in product appeal.

Panel (b) estimates the roles of scope and appeal over the firm life cycle. We study the evolution of the components in equation (21) over the firm life cycle by estimating age-fixed effects after controlling for cohort- and time-fixed effects. The results indicate consistent firm growth over the life cycle in both model and data, though the pace of growth is faster in the data, especially for the very young firms. Firm scope makes a positive non-negligible contribution to firm growth, though product appeal explains most of the increase in firm sales in the latter part of the life cycle.³¹

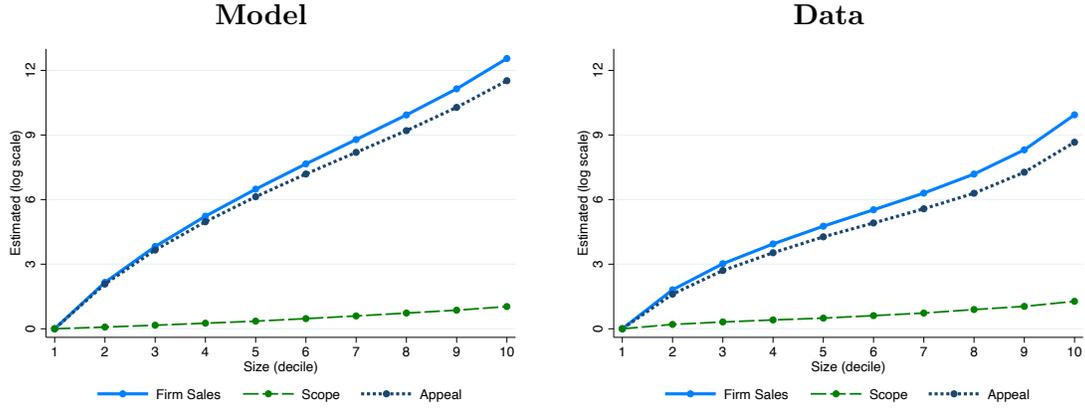
This model shows that the decline in sales over a product’s life cycle due to business stealing and cannibalization has several important implications for firm growth. Our estimation generates the steady decline of individual product sales simultaneous with the steady increase

²⁹We explore some additional non-targeted moments in Online Appendix F.2. For example, we consider the association between size and scope, on average. In the model and in the data, the average relationships between scope and firm size are remarkably similar.

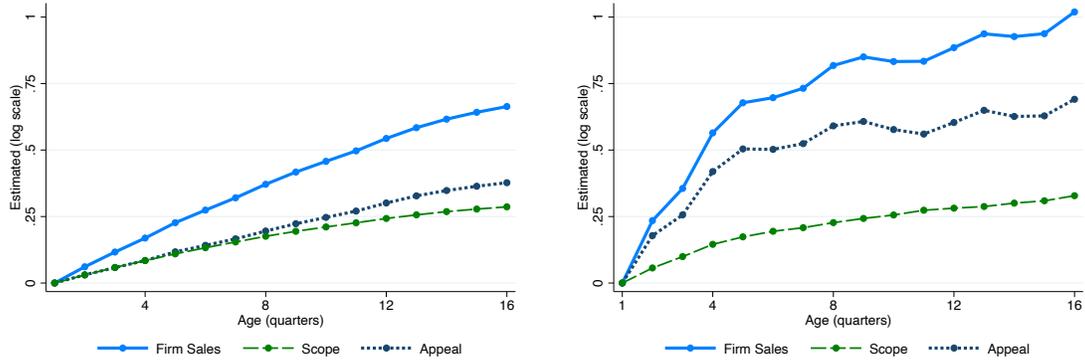
³⁰The contribution from product quality comes from both the geometric mean of the appeal of the firms’ products and the dispersion of product quality: $(\sigma - 1) \log \bar{q}_{it} + \log \left(\frac{1}{N_{it}} \sum \left(\frac{q_{ut}}{\bar{q}_{it}} \right)^{\sigma-1} \right)$.

³¹In Online Appendix, we use the firm size decomposition suggested by [Hottman, Redding and Weinstein \(2016\)](#) and find similar results. Similarly, [Eslava and Haltiwanger \(2020\)](#) use data from manufacturing plants in Colombia and find evidence that appeal explains the bulk of firms’ sales growth. We also use a variance decomposition of firm sales over the life cycle. These results are presented in Online Appendix E.5.

Figure 8: Firm Size Distribution: Data - Model Comparison



(a) Firm Size Distribution (sales deciles)



(b) Firm Life Cycle (age)

Notes: Panel (a) shows the estimated differences in average sales/scope/appeal across size deciles. We use a nonparametric regression that includes indicator variables for each firm size decile, with the smallest firm size category serving as the reference group. We also control for time and cohort fixed effects. Panel (b) shows the estimated age fixed effects of sales/scope/appeal over the life cycle of firms using equation (2). The estimation includes cohort and time effects. We keep a balanced sample with 16 quarters or above durations. In both data and simulated data, we use the same definition of age, survival and censoring.

in a firm’s overall sales figures over the life cycle. These seemingly contradictory trends co-exist because, both in the data and in the model, sales from new products compensate for the declining sales of existing products.³² New products impact a firm’s total sales by increasing the firm’s scope and increasing the firm’s appeal component (equation 21). While scope also makes an important contribution, both in the model and in the data, a firm’s growth is strongly associated with the increase in the average appeal of its products (Panel (b) of Figure 8). This means that the appeal of new products is sufficiently high to compensate for the decline in appeal of existing products as innovation proceeds.

³²Our statistical decomposition in Section 3 shows a very negative contribution from the sales of existing products and a very large positive contribution coming from the introduction of new products. Online Appendix Table F.I shows that a decomposition using simulated data is similar to the empirical decomposition.

5.5 The innovation-obsolescence cycle

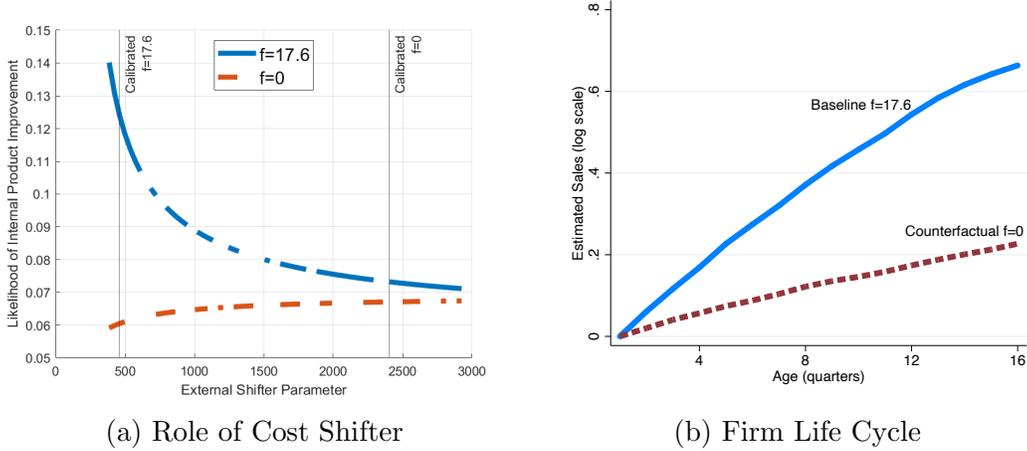
The simultaneous decline in the appeal of existing products and increase in firm appeal suggest that the data are well-characterized by an economy in which: (i) competitors introduce new products that erode the appeal of other products in the market; (ii) as the appeal of existing products declines, firms selling these products see increasing benefits to introducing new improved products; and (iii) in introducing new products, firms accelerate the decline in sales and eventual demise of their existing products. We call this mechanism the innovation-obsolescence cycle. Our model framework allows for this cycle to characterize the economy but does not impose the dynamic with its formulation. Our quantifications of the innovation and obsolescence steps shows that (i) and (iii) characterize the economy well, and allow us to conclude that innovation induces obsolescence strongly. It is less clear how strongly obsolescence induces innovation, however, if indeed the benefits from introducing new products increase as business stealing on the part of competitors increases (ii).

We now study the relationship between rates of internal and external innovation. The model captures rich interdependence between internal innovation and external innovation, where internal innovation and external innovation may be complements (i.e. internal innovation increases in response to an increase the rate of external innovation) or substitutes (i.e. internal innovation increases in response to a decline in the implicit cost of external innovation). Most endogenous growth models that feature both types of innovation predict that as the rate of external innovation increases (say driven by a decline in cost of external innovation), the expected life-span of products will decrease, which would reduce a firm's incentives to innovate internally (Akcigit and Kerr, 2018). While this force is at play in our model (captured by A), we also allow for a countervailing force, as the internal innovation rate is also dependent on the present discount value associated with external innovation (captured by B , as specified in proposition 1).³³ The importance of each of these forces depend on the estimated parameters, particularly on the cost shifter f . If we impose this cost shifter to be equal to zero, we would exclude this countervailing force. The intuition here is that a positive f makes the marginal cost of internal innovation smaller for products with high appeal, and incentivizes firms to improve the quality of their products. Otherwise internal innovation would become more costly over time. A reduction in the marginal cost of innovation induces an increase in the optimal rate of internal innovation (equation (9) of Proposition 1).

Our quantification indicates that the data is better described by internal and external innovation being strongly complementary, by which we mean that firms respond with more internal innovation when there is an increase in obsolescence from business stealing. This

³³Note that $x^I = \left[\frac{1-\alpha}{\xi^I} \left(\frac{A}{\pi} [(1-\lambda^C)^{\sigma-1} + (1+\lambda^I)^{\sigma-1} - 1] + \frac{B}{\pi} f \right) \right]^{\frac{1-\alpha}{\alpha}}$ and the standard force is captured by changes in A , and the countervailing force is captured by changes in B .

Figure 9: Impact of Cost Shifter f on the Firm Life Cycle



Notes: Panel (a) shows the optimal internal innovation rate for the simulated baseline and counterfactual f , while varying the cost shifter of external innovation. Panel (b) shows the estimated age-fixed effects of sales over the life cycle of firms using equation (2) for a simulated baseline and counterfactual f . The estimation includes cohort and time effects. We maintain a balanced sample with durations of 16 quarters or more. In both empirical and simulated data, we use the same definitions of age, survival, and censoring.

point can be clarified if we study the sensitivity of internal innovation to a change in the cost shifter of external innovation (ξ^E) while keeping all other parameters constant. Figure 9 (Panel a) shows a negative association between internal innovation and the cost shifter for external innovation in our baseline model, where this cost shifter is estimated to be $f = 17.6$.

We consider a counterfactual economy where we impose $f = 0$, and we re-calibrate the other parameters to match the moments in the data.³⁴ Figure 9 (Panel a) shows that in this counterfactual economy, the relationship between internal innovation and the cost shifter of external innovation is positive.

To further understand the implications of the complementarity between internal and external innovation, we compare the product life cycle and the firm life cycle in the baseline and counterfactual economies. Not surprisingly, we estimate a very similar decline in sales in the baseline and counterfactual economies because both were calibrated to match the data moments of the product life cycle. Figure 9 (Panel b) shows that the two economies differ substantially in their expected evolutions of sales over the firm life cycle. In the counterfactual economy, competitors' business stealing reduces incentives toward internal innovation. Indeed, in this economy, the overall level of innovation is smaller, as firms are less willing to cannibalize their own existing products.

The model clarifies how the declining product life cycle coexists with growth over the firm

³⁴Online Appendix F.3 provides the estimated parameters. The fit to the data is substantially worse. Specifically, it is very hard to match the incumbent innovation rate, the firm size distribution, and the ratio of innovation costs to sales.

life cycle. Firms optimally adjust by investing in new products to offset losses in the appeal of existing products. In fact, our results indicate that firms on average do not rely on older products to generate positive growth, instead introducing newer and better products (those with higher appeal) as a necessary condition for growth. New products broaden a firm’s scope, and, more importantly, they preserve the firm’s average appeal, which would decline over time in the absence of new-product introductions.

Our model establishes a new mechanism that ties together the empirical results that we documented above. It indicates that a firm’s incentives to renew its appeal through the introduction of new products are directly tied to the rate of decline of the appeal of existing products. Obsolescence induces innovation, in short. We show, in Section 4.2 and model quantification, that these incentives are driven by the introduction of new products on the part of competitors. Innovation also induces obsolescence. The decline in the appeal of existing products is, therefore, an important force behind the introduction of new products: firms are more likely to introduce new products when competitors are also introducing more products of their own, even in the presence of cannibalization forces.

6 Conclusion

We study the product life cycle as described by sales, quantities sold, and prices. We find that sales decline at a fast pace throughout the product life cycle, for a wide range of products. The decline in sales is mostly driven by declines in quantities sold, as opposed to prices, and cannot be explained by firm-specific factors. These are novel findings that are relevant to theories of product and firm dynamics.

We structurally estimate a model of heterogeneous multi-product firms and use it to recover the contributions to product sales from appeal, cost, markup, and the effects of cannibalization. Our results point to the decline in appeal relative to other products in the market (business stealing) and relative to other products of the same firm (cannibalization) as the most important determinants of the evolution of product sales. Guided by the results of this decomposition, we explore the underlying causes for the decline in appeal. We show that the decline in product appeal over the life cycle is closely related to the introduction of new products both internally and externally.

The decline in a product’s appeal and sales over the life cycle has several important implications for firm growth. When we examine the firm and product life cycles jointly in a dynamic model with endogenous innovation, we find that firms counteract the effects of a product’s life cycle by introducing new products. Firms must introduce new products to compete, significantly more so when facing innovative competitors. Otherwise, a firm’s portfolio

becomes obsolete as competitors introduce new products of their own. However, by introducing new products, firms accelerate the decline in sales of their existing products, which partially explains why a product's sales decline throughout most of the life cycle.

Our findings are relevant to the vibrant debate concerning the evolution of competition and market power in the economy. They emphasize the critical role that non-price strategies play in shaping the modern competitive environment. In the context of our framework, a firm can respond to a competitor by introducing new products. When business stealing is relatively prevalent, firms will find it more profitable to respond by introducing a new product than by reducing the prices of existing products. Our headline finding that the sales of individual products decline throughout most of the life cycle suggests that an arms race to introduce the most appealing and consumer-enticing products is a hallmark of firm competition across a wide variety of sectors.

The present work can be extended in several directions. First, with price and quantity data at the product level covering other sectors, it would be possible to quantify the pervasiveness of the patterns documented above across the broader economy. Second, one could explore the persistent welfare consequences of the innovation-obsolescence cycle. For instance, when competitors suffer shocks that slow product introduction, non-affected firms may respond by introducing less products of their own, which would further reduce the rate of innovation in the economy. We leave these extensions for future research.

References

- Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr, “Innovation, reallocation, and growth,” *American Economic Review*, 2018, *108* (11), 3450–91.
- 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.
- Anderson, Eric, Sergio Rebelo, and Arlene Wong, “Markups across space and time,” *NBER Working Paper No. 24434*, 2018.
- Argente, David and Chen Yeh, “Product Life Cycle, Learning, and Nominal Shocks,” 2017.
- , Doireann Fitzgerald, Sara Moreira, and Anthony Priolo, “How Do Firms Build Market Share?,” 2021.
- , Munseob Lee, and Sara Moreira, “Innovation and product reallocation in the great recession,” *Journal of Monetary Economics*, 2018, *93*, 1–20.
- Atkeson, Andrew and Ariel Burstein, “Aggregate implications of innovation policy,” *Journal of Political Economy*, 2019, *127* (6), 2625–2683.
- Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen, “The Fall of the Labor Share and the Rise of Superstar Firms,” *Quarterly Journal of Economics*, 2020, *135* (2).
- Basker, Emek and Timothy Simcoe, “Upstream, Downstream: Diffusion and Impacts of the Universal Product Code,” *Journal of Political Economy*, 2021, *129* (4), 1252–1286.
- Bernard, Andrew B, Stephen J Redding, and Peter K Schott, “Multiple-Product Firms and Product Switching,” *American Economic Review*, 2010, *100* (1), 70–97.
- Bloom, Nicholas, Charles I Jones, John Van Reenen, and Michael Webb, “Are ideas getting harder to find?,” *American Economic Review*, 2020, *110* (4), 1104–44.
- Broda, Christian and David E Weinstein, “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.
- 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 microeconomic 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.
- Eslava, Marcela and John Haltiwanger, “The Life-cycle Growth of Plants: The Role of Productivity, Demand and Wedges,” *NBER Working Paper No. 27184*, 2020.
- Fitzgerald, Doireann, Stefanie Haller, and Yaniv Yedid-Levi, “How exporters grow,” *NBER Working Paper No. 21935*, 2016.
- 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?,” *Econometrica*, 2019, *87* (5), 1507–1541.
- Gowrisankaran, Gautam and Marc Rysman, “Dynamics of consumer demand for new durable goods,” *Journal of Political Economy*, 2012, *120* (6), 1173–1219.
- Granja, Joao and Sara Moreira, “Product Innovation and Credit Market Disruptions,” *Available at SSRN 3477726*, 2021.
- 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.
- Kaplan, Greg and Guido Menzio, “The morphology of price dispersion,” *International Economic Review*, 2015, *56* (4), 1165–1206.
- Klette, Tor Jakob and Samuel Kortum, “Innovating Firms and Aggregate Innovation,” *Journal of Political Economy*, 2004, *112* (5).

- Levitt, Theodore, *Exploit the product life cycle*, Vol. 43, Graduate School of Business Administration, Harvard University, 1965.
- Moreira, Sara, “Firm Dynamics, Persistent Effects of Entry Conditions, and Business Cycles,” 2017.
- Nevo, Aviv, “Measuring market power in the ready-to-eat cereal industry,” *Econometrica*, 2001, *69* (2), 307–342.
- Perla, Jesse, “A Model of Product Awareness and Industry Life Cycles,” Working Paper 2019.
- Peters, Michael and Conor Walsh, “Population growth and firm dynamics,” Technical Report, National Bureau of Economic Research 2021.
- Schulhofer-Wohl, Sam, “The age-time-cohort problem and the identification of structural parameters in life-cycle models,” *Quantitative Economics*, 2018, *9* (2), 643–658.
- Vernon, Raymon, “International Investment and International Trade in the Product Cycle,” *The Quarterly Journal of Economics*, 1966, *80* (2), 190–207.
- Wollmann, Thomas G, “Trucks without bailouts: Equilibrium product characteristics for commercial vehicles,” *American Economic Review*, 2018, *108* (6), 1364–1406.