Growth through Innovation Bursts*

Giuseppe Berlingieri‖ ESSEC

Maarten De Ridder‡ LSE

Danial Lashkari§ FRBNY

Davide Rigo¶ LSE

21 March 2024

Please Click Here for the Latest Version.

Abstract

In theories of creative destruction, product innovation is a key driver of aggregate economic growth. In this paper, we confront the predictions of these models with the empirical patterns of product innovation, documented based on product-level data on the near-universe of French manufacturing firms. Two key patterns in our data stand in contrast to the workings of the conventional models. First, we find that product innovation is best described as a process of "bursts"–episodes where firms rapidly add a series of products to their portfolio. These bursts lead to substantial shifts in revenue and explain the majority of the variance in firm-level growth. Second, we find that the growth in product-level revenue declines over the course of a product's life cycle, indicating a diminishing effectiveness of process innovation. Including these features into a quantitative framework that nests the canonical models of creative destruction, we show that innovation bursts explain the concentration of production in a small number of large firms. Our model thus enables the joint study of the determinants of industry concentration and growth in a setting consistent with the empirical patterns of product dynamics. In a quantitative comparison against the conventional models, our theory attributes to creative destruction a substantially greater impact on aggregate productivity growth and on the concentration of production.

Keywords: Productivity, Endogenous Growth, Firms, Innovation

*We thank Xavier Gabaix and Michael Peters for useful discussion. This work is supported by a public grant overseen by the French National Research Agency (ANR) as part of the 'Investissements d'Avenir' program (reference: ANR-10-EQPX-1, CASD).

‡*Address:* ESSEC Business School, Department of Economics, 3 Avenue Bernard Hirsch, Cergy-Pontoise, France *E-mail*: giuseppe.berlingieri@essec.edu. *Web:* https://sites.google.com/site/giuseppeberlingieri/.

§*Address:* London School of Economics, Department of Economics, 32 Lincoln's Inn Fields, London, U.K. *E-mail*: m.c.de-ridder@lse.ac.uk. *Web:* http://www.maartenderidder.com/.

¶*Address:* Federal Reserve Bank of New York, Department of Research and Statistics, 33 Liberty St, New York, U.S.A. *E-mail*: danial.lashkari@ny.frb.org. *Web:* https://sites.google.com/view/danial-lashkari/home.

‖*Address:* London School of Economics, Department of Economic Geography, London, U.K. *E-mail*: d.rigo@lse.ac.uk. *Web:* https://daviderigo.weebly.com/.

1. Introduction

Product innovation lies at the heart of the modern theories of growth through creative destruction. In these theories, innovative firms invest in research and development (R&D) to acquire the ability to produce new products that are superior to those of their competitors. As their new products displace the old, innovating firms acquire market share and grow their output, revenue, and profits, while raising aggregate productivity and output in the process. Since the seminal contribution of Klette and Kortum (2004), the theories of creative destruction have been extended in many directions to account for key empirical patterns on the relationships between firm-level R&D, patenting, and growth. These empirical extensions have allowed researchers to study the consequences of the product innovation activity of firms for macro-level phenomena such as the aggregate productivity growth, business dynamism, competition, and product-market concentration (e.g. Lentz and Mortensen 2008, Akcigit and Kerr 2018, Acemoglu et al. 2018a, Cavenaile et al. 2020, Peters 2020, Aghion et al. 2023).

However, despite the wide applications of the theories of creative destruction, we know surprisingly little about how well they explain the dynamics of firm-level product entry and exit. Due to the unavailability of comprehensive, product-level data in innovation-intensive sectors, most researchers rely on the dynamics of firm-level revenue or employment, observed only as aggregates across all of of a firm's products, to empirically discipline their models.[1] In this paper, we fill this gap using comprehensive data from the French manufacturing sector, detailing the product composition of all major firms at a fairly disaggregated level. Using the data, we document novel facts on firm-level product innovation activity that shed light on the underlying mechanisms driving firm-level product innovation, but are difficult to explain under the standard models of creative destruction. We propose a simple theory based on Klette and Kortum (2004) that is compatible with these new facts. When quantitatively disciplined with our data, the new theory suggests a substantially larger role for creative destruction as the driver of aggregate growth compared to that implied by the standard framework.

To examine the dynamics of product creation and destruction, we rely on a comprehensive survey of French manufacturing firms containing information on the distribution of their sales across highly detailed product categories. Our first fact is that the distribution of the number of products produced by French manufacturing firms in highly concentrated among large firms and has a thick, Pareto-like right tail. Our second fact documents that the distribution of new products introduced in any given year also exhibits a substantial degree of concentration, which has a similar Pareto-like right tail. In other words, a small share of firms are responsible for a sizable share of all product innovation as they experience what we refer to as *innovation bursts*. We provide various robustness checks to ensure that such innovation bursts constitute episodes of meaningful product innovation, and are not the byproducts of the specifics of data construction or reporting. Together, these two facts point to the dynamics of product innovation as an important determinant of the high degree of production concentration typically observed in firm-level data.

[1] Two notable exceptions that empirically investigate the dynamics of firm-level products include Argente et al. (2024) and Argente et al. (2020), who focus on the nondurable, convenience consumer goods sector based on retail scanner data.

We next use our data to quantify the contribution of creative destruction and, more specifically, product innovation bursts to revenue growth at the firm and aggregate levels. We find a larger contribution for process innovation, defined as the growth or fall in revenue on each firm's continuing set of products, when compared with that of creative destruction, defined as the growth or fall in firm revenue due to the introduction of new products or loss of old products. Nevertheless, we also find substantial heterogeneity in the balance between the two contributions across firms. Importantly, we document that the balance steadily shifts toward higher contributions from creative destruction as we move from firms with low rates of growth toward firms in the either extreme ends of the distribution of revenue growth. In particular, the most extreme rates of growth belong to entering and exiting firms, whose revenue growth is by definition fully driven by creative destruction. More importantly, we also find that such firm-level entry and exit explains a substantial part of aggregate revenue growth.

Our last fact shows that, in addition to its direct contribution to aggregate growth highlighted above, creative destruction makes an additional indirect contribution by furnishing firms with new possibilities for process innovation. Using our data, We document a marked decline in the revenue growth of continuing products over the course of their life cycle. In other words, product-level revenue tends to grow the fastest in the first few years immediately after a new product is introduced to the market. Thus, not only does product innovation grow firm revenue in the first year of the introduction of a new product, it further enables the firm to sustain a faster rate of growth in the ensuing next few years through an initially higher rate of process innovation.

To rationalize these facts, we propose a new way of modeling endogenous growth through product and process innovation. Our first key novelty is that firm-level product innovation potentially arrive in bursts: episodes in which the firm's product portfolio expands rapidly. In standard models of creative destruction, firms that innovate expand their portfolio of products by a single product. In contrast, in our model of innovation bursts an innovating firm may acquire the technology to produce a large number of additional products, as innovations may have multiple applications. We characterize a stochastic process of product innovation bursts that is consistent with the high degree of concentration of new products observed among French manufacturing firms. The model naturally nests the canonical Klette and Kortum (2004) model for the special case in which an innovation only expands firms' product portfolio by a single product. In addition, process innovation in our model may exhibit diminishing returns in the sense that successive innovations improving the process of producing each product become increasingly incremental over the course of its life cycle. In other words, firms in our model find it increasingly harder to find ideas to improve the production process as a product ages.

Innovation bursts help explain the high degrees of production concentration typically observed in firm-level data, associated with the Pareto distribution of firm size. A well-known limitation of theories of creative destruction is that they have difficulty explaining the existence of large firms. The forces of creative destruction prevent the largest firms in the economy from significantly outgrowing their competitors. This is at odds with empirical evidence: numerous studies have found that the firm-size distribution also has a fat tail, with a Pareto index close to one (e.g. Axtell 2001, Gabaix 2009). We show that for sufficiently thick-tailed distribution of innovation bursts, the equilibrium firm-size distribution is Pareto. Thus, we propose the lumpy nature of innovative ideas as the main driver of the concentration of firm-level production, a reasoning that is fundamentally distinct from previous theories of firm dynamics that rationalize the fat-tailed distribution of firm size (e.g., Luttmer 2007, Luttmer 2011, Acemoglu et al. 2018b).[2]

Given the close links between our empirical analysis and the theoretical framework, we are able to calibrate the majority of the parameters of the model directly from the observed patterns in the data. The remainder are standard parameters that we calibrate based on common values used in the literature. For comparison, we also calibrate a conventional model of creative destruction featuring both product and process innovation without the presence of innovation bursts. We first confirm that, unlike the conventional model, ours can match the distribution of firm products (unconditionally or conditional on the lagged number of products), as well as the high degree of production concentration observed in the data. Our model also fits the evolution of average product revenue over the course of its life cycle. While remaining untargeted in our calibration, our model also provides a reasonable fit for the decomposition of revenue growth in the data to the contributions of product creation and destruction and that of continuing products.

Finally, we use the calibrated model to perform a quantitative investigation of the degree to which product and process innovation drive the dynamics of aggregate productivity growth. Following prior work (Garcia-Macia et al., 2019), we first confirm that relying on an additive decomposition, one may attribute a small contribution to aggregate growth from product innovation. However, through the lens of our model, this simple accounting exercise leaves out an important indirect contribution from product innovation to growth. As ideas for process innovation become harder to find over the life cycle of each product, the introduction of new products plays a crucial role in our model for facilitating sustained process innovation. In our calibration, the indirect contribution of product innovation to growth, mediated according to this channel through process innovation, is around twice larger than its direct contribution.

Using the calibrated models, we also perform a quantitative comparison of the effects of creative destruction on production concentration between our model and the conventional model. In both models, a rise in the rate of creative destruction increases the concentration of production among large firms. However, since product innovations arrive as bursts in our model, this link is tighter in our model than that in the conventional model. Measuring the concentration of production in terms of the revenue share of the top 10% of firms, we find that the response of this measure to a rise in the rate of creative destruction is around 30% stronger in our model than that in the conventional one. If we instead use as a proxy a measure of average firm size such as the average number of products per firm, the response is around 60% stronger in our model.

Taken together, our quantitative results suggest a far more important contribution from creative destruction to aggregate productivity growth and to the concentration of production than that implied by the conventional theories of creative destruction.

[2] Luttmer (2010) offers a comprehensive review of models of firm dynamics that lead to a Pareto-tailed distribution of firm size. For an alternative approach that leads to a Pareto distribution of firm size in a static assignment model, see Geerolf (2017).

Related Literature Our paper contributes to the large body of research in the tradition of the Schumpeterian growth models of creative destruction following Aghion and Howitt (1992) and Grossman and Helpman (1991) (see Aghion et al. 2014 for a review). Using this framework, Klette and Kortum (2004) build an analytically tractable model of innovating multi-product firms that it consistent with several empirical fact of firm dynamics. As already mentioned, many researchers have extended this model to apply it in different applications such as the study of firm heterogeneity (e.g., Lentz and Mortensen 2008), misallocation and innovation policy (e.g., Acemoglu et al. 2018b and Atkeson and Burstein 2019), competition and market power (e.g., Peters 2020), management and firm organization (e.g., Akcigit et al. 2021), business dynamism (e.g., Akcigit and Ates 2023), among others.

As already mentioned, we confront this class of models with data on the dynamics of product churn and turnover at the firm level. Prior attempts for testing the predictions of these models regarding the dynamics of product innovation have instead relied on indirect evidence drawn from patent data (e.g., Akcigit and Kerr 2018 and Argente et al. 2020).

Multiple recent studies have documented the fact that industry concentration has been globally rising over the past 30 years (e.g. Autor et al. 2020). Concurrently with the rise of large firms, there has been a gradual decline in the growth rate of total factor productivity, despite a sustained rise in innovative investments (Bloom et al. 2020). By constructing a theory of creative destruction that directly matches the firm-level empirical evidence on product innovation, our paper offers a laboratory for studying the interplay between concentration, innovation, and growth.

Outline The remainder of this paper proceeds as follows. Section 2 presents the details of data used in our paper and Section 3 discusses the empirical facts. Section 4 presents our model, Section 5 discusses our quantitative exercise, and Section 7 concludes the paper.

2. Data

We use detailed data on the composition of the product portfolio of firms in the French manufacturing sector. The data combines firm-level income statement and balance sheet data from tax records with a detailed survey of the product portfolio of firms, covering revenues and quantities sold at the level of detailed product categories.

Our source of information on firm product portfolios is the *Enquête Annuelle de Production* (EAP), which is a survey collected by the statistical office (INSEE). The EAP surveys the universe of French manufacturing firms with at least 20 employees or 5 million euros in revenue, comprising around 90% of the value of aggregate manufacturing output.[3] The data is available from 2009 to 2019. We start our analysis in 2010, when the survey methodology was finalized.[4] The survey contains revenue, quantity, and average unit values for each product category that a firm produces each year at the level of

[3]The survey additionally includes a random sample of small firms. As this sample is redrawn every year, it is unsuitable for our analysis of firm-level product dynamics, which requires panel data.

[4]The first year of the data, 2009, is excluded because the coverage of the survey appears incomplete as firms report substantially fewer numbers of products.

10-digit product codes. The high level of detail in these product codes, distinguishing among 4000 distinc products, enable us to investigate changes in each firm's product portfolio. The first six digits are the harmonized European activity codes, which are sufficiently narrow to identify applicable customs policies. The remaining digits contain a further sub-classification that is produced particularly for France by INSEE.[5] We exclude from our analysis the firms in EAP that do not belong to the manufacturing sector, such as those in mining, repairs, and installation industries.

We combine this product-level data with standard, firm-level balance sheet data from the *Fichier Approaché des Résultats d'Esane* (FARE). FARE provides detailed income statement and balance sheet data such as total sales, wage bill, and capital, for the universe of French firms. Our main dataset covers the set of firms that belong to the intersection of the FARE and EAP datasets. Additionally, we obtain total spending on research and development (R&D) from the research tax credit program and measure the wage bill of scientists that work at a firm using occupational codes of employees in the *Déclaration annuelle de Données Sociales*–the social security filings. All datasets are merged using their common firm identifier, the SIREN code.

Table 1 presents some summary statistics of the resulting dataset. The data contains 241,600 firm-year observations, covering 31,924 unique firms operating in 188 four-digit manufacturing industries. The number of products refers to the number of unique 10-digit product codes for which the firm reports positive revenue in a given year. A new product is a 10-digit product code for which the firm reported no revenue in the previous year and now reports positive revenue in the current year. For firms that appear in the data for the first time we set the number of new products to missing. This is because firms that appear for the first time in EAP may not be genuine entrants, as they may have only been included in the data because they have achieved the sufficient size needed for the inclusion in the data. A lost product, conversely, is a 10-digit product code for which the firm does not report positive revenue in the current year while it did report positive revenue in the previous period. For exiting firms, i.e., those exiting the comprehensive FARE dataset, we set the number of lost product in equal to the number of products that firm reported positive revenue for in its last period of presence in the data. We provide further details on variable definitions and cleaning in Online Appendix B.

3. Stylized Facts on Products and Product Innovation

This section establishes four stylized facts using the data from Section 2. First, we show that firms differ considerably in the number of products that they produce. This means that, besides differences in the amount of revenue that firms earn *per product*, concentration of product count is an important driver of the fact that economic activity concentrates in a limited number of firms. Second, we show that the number of products that firms produce is subject to significant episodes of product innovation and product destruction. We document that product innovation comes in bursts: times at which firms rapidly expand the number of products it produces. Third, we show that both product innovation and destruction are important drivers of overall firm growth, with product innovation steadily becoming

[5] For example, the 8-digit code 20421300 contains pedicure and manicure cremes, where the 10-digit codes 2042130010 and 2042130020 respectively identify cremes that should be applied to the skin and the nails.

Table 1: Summary Statistics

Variable	Mean	St. Dev.	Median	10th Pct.	90th Pct.	Observations
Sales	22291	283584	3319	584	29537	223887
Age	28	19	24	8	52	223887
Employees	67	307	20	4	120	223887
Revenue per Product	9607	179084	2001	316	15637	223887
Number of Products	1.93	2.66	1	1	4	223887
Lost Products	0.2	0.9	0	0	1	203159
New Products	0.14	0.89	0	0	0	193721
Continuing Products	1.81	2.32	1	1	4	193721

Notes: The table presents summary statistics for the merged EAP-FARE dataset between 2010 and 2019. Age and employment are defined in years and full-time-equivalent workers, respectively. Sales is measured in thousands of 2015 euros. Sales per new and lost product are the ratio of sales specific to new or lost products, respectively.

a more important driver of growth among firms that are growing or shrinking at higher rates. Finally, we show that there is a steady decline in the growth of revenue that firms earn on products that they continue to produce.

3.1. Concentration of Products

It is widely documented that production activity is highly concentrated among a small share of superstar firms, as the distribution of both employment and revenue across firms also resembles a power law (e.g. Axtell 2001, Luttmer 2007, Gabaix 2009, Garicano et al. 2016). Estimates yield that 60% to 90% of production is done by the largest 20% of firms. Among French firms in our sample, the largest 20% of firms are responsible for 84% of revenue and employ of 78% workers.[6]

In this section, we show that the distribution of the number of products that firms produce is highly concentrated: while the majority of firms only produce one or very few products, there is a thick tail of firms that produce numerous products. In other words, a minority of firms are responsible for the production of the majority of products sold in the market. We analyze the degree of concentration by plotting the natural logarithm of a firm's number of products against the natural logarithm of the ranking of the firm's product count in the sample. If the number of products of firms is well-described by a power law, as is the case for other measures of firm size such as employment or revenue, the relationship between log-product count and its log-rank should be linear.

Figure 1a plots the distribution for a firm's number of products, where a firm's product count is equal to the number of products for which it reports positive revenue. The figure shows that the relationship between rank and the number of products that firms produce is nearly perfectly linear. It therefore follows that the number of products that firms produce indeed resembles a power law and

[6] As it happens, this is closely in line with the heuristic 80/20% Pareto principle, stating that 80% of the outcome of interest typically belongs to 20% of agents. If we assume a Pareto distribution, the percentage y of the number of products that is produced by the largest x% of firms is given by $y = x^{(\theta-1)/\theta}$, where θ is the Pareto tail index. For the 80/20% rule, we find a Pareto tail index of around 1.16, fairly close to the tail index of unity implied by the Zipf law.

Figure 1. Decomposing Firm Concentration into Products and Revenue Per Product

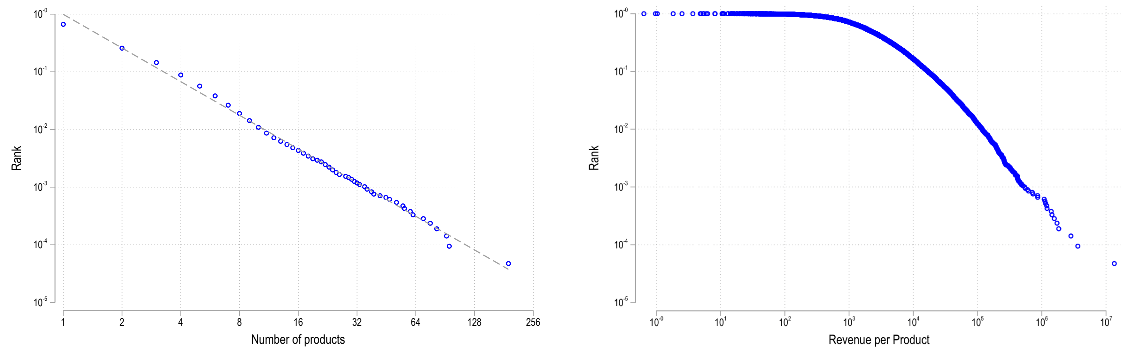

(a) Number of Products

(b) Revenue per Product

Notes: The figures plot the relationship between a firm's size (horizontal axes) and the firm's inverse rank (vertical axes). Size is measured through either revenue per product (in 2015 euros) or number of products. Inverse rank is measured as the ratio of firms' rank starting from the smallest firm, divided by the total number of observations in the data. Plots are based on the FARE-EAP sample for 2019. The distribution for 2010 to 2018 is plotted in Appendix Figure A3. Revenue per product is measured in thousands of euros.

can be approximated by a Pareto distribution in its right tail. The slope of the size-rank relationship implies that the largest 20% of firms produce 55% of all goods.[7]

The above fact implies that variation in the average revenue per product that firms produce is partly responsible for the concentration of economic activity. We plot the distribution of this average across firms in Figure 1b. The figure shows that revenue per product does not follow a power law everywhere, as the left tail of the distribution is highly dispersed. The number of products that firms produce, while not explaining the full degree of concentration, thus seems an important contributor to the overall dispersion in size.

3.2. Product Innovation and Product Destruction

Having established that firms differ considerably in the number of products that they produce, we now assess how the composition of a firm's product portfolio changes over time. We first assess the rate at which firms create new products, and the rate at which they lose them. We define the rate of product innovation as the ratio of the number of products that firms report positive revenue on at time t but not at time $t-1$, divided by the number of products that the firm reported positive revenue on at time $t-1$. We define the rate of product destruction as the share of the firm's products at time $t-1$ that the firm does not report revenues on at time t.

Table 2 shows the product innovation and destruction rates for firms, depending on their number of products in the previous period. The first row gives the average of the rates, while the bottom row gives their bootstrapped standard errors. When firms of all sizes are included, we find that firms on average add 0.066 products to their portfolio for every product that they initially produce, while they stop producing 0.082 products for every such product. The wedge between incumbents' product innova-

[7]We can find the corresponding tail index from the coefficient of a least-squares regression of the log-rank on log-firm size relationship, yielding an index of 1.58.

Table 2: Overview of Product Innovation and Destruction

Product Innovation Rate							Product Destruction Rate						
All	1	2	3	4-5	6-8	8+	All	1	2	3	4-5	6-8	8+
.066	.066	.067	.068	.058	.056	.083	.082	.065	.111	.117	.117	.110	.129
(.001)	(.002)	(.001)	(.002)	(.002)	(.003)	(.004)	(.001)	(.001)	(.001)	(.002)	(.002)	(.003)	(.004)

Notes: Product addition rate is the number of new products that a firm starts producing divided by its original number of products. Product loss rate is the number of products that a firm stops producing over the horizon divided by its original number of products. Different columns report the results conditional on the firm's number of products in the preceding period N_{t-1}. The top row lists the product innovation and creation rates while bootstrapped standard errors from 1000 replications are given in parentheses.

tion and destruction rates means that the number of products that a typical incumbent firm produces on average shrinks over time. Such a wedge is expected if new firms enter the economy while the total number of firm-products in the economy is approximately constant over time–which is the case in our data. Columns differ in the initial product count of firms that are included, grouping the larger firms to maintain a sufficient number of observations. For the product innovation rates, the table shows no particular pattern across firms with different initial sizes: the rate at which small and large firms expand their portfolio is similar, ranking from 0.056 to 0.083 per products initially produced. This means that the expected number of new products firms add to their portfolios scales with the sizes of their portfolios. The rate of product destruction is similar for firms with at least two products, though it is almost 50% lower for firms that produce a single product. This means that single product firms have higher average growth rates than larger firms, which is consistent with previous evidence that the smallest firms grow faster than larger firms in terms of employment (e.g. Akcigit and Kerr 2018).

The average rates of product innovation reported in Table 2 mask substantial dispersion in the number of products that firms add to their product portfolio. The product destruction rate is naturally

Figure 2. Distribution of Number of New Products

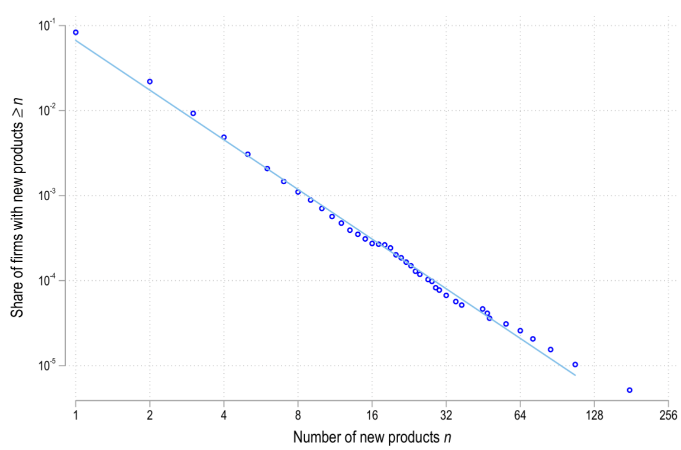

Notes: The figures plot the relationship between a firm's size (horizontal axes) and the firm's inverse rank (vertical axes). Size is measured through either revenue (in 2015 euros) or employment (in full-time equivalence). Inverse rank is measured as the ratio of firms' rank starting from the smallest firm, divided by the total number of observations in the data. Plots are based on the FARE-EAP sample for 2019. To find the fraction of firms x that is responsible for y% of product innovations we calculate $x = y^{\alpha/(\alpha-1)} \times 0.34$, where α is the Pareto tail parameter and where the multiplication adjusts for the fact that only 34% of firms report product innovation. We find an α of -1.94.

bounded to one, as firms can only lose as many goods as they produce. The product innovation rate has no upper bound, however, and firms may expand well beyond their initial size between periods. Figure 2, which plots the distribution of a firm's number of new products, shows that this is the case in our data. The vertical axis gives the log of rank while the horizontal axis gives the log of the number of new products, as we did in Figure 1a. The figure shows that product innovation is highly concentrated. The distribution in Figure 2 implies that 75% of product innovation comes from just 20% of firms.

The high degree of concentration in product innovation persists when we separate the sample by firms' initial size in Figure 3. This matters because firms of different sizes have constant innovation *rates*, which means that concentration of overall product innovation can simply arise from the fact that firms differ in their initial sizes. The figure shows that even among firms that initially produce only a single product, there are product innovators that add tens of products to their portfolio. Note that both Figure 2 and 3 only plot the distribution of product innovation for firms that add at least one product to their portfolio. A large share of firms (66%) report no product innovation at all. For single-product firms concentration is even more extreme, with all product innovation coming from just 7% of firms. At the same time, it is common for firms to add a large number of products to the market when they do engage in product innovation. We henceforth labels these episodes "innovation bursts."

Robustness We perform various robustness checks to assure that innovation bursts are indeed episodes in which firms expand their product portfolio, and not an artifact of reporting. We first show that the revenue that firms report on new products does not depend on the number of new products that firms report. If innovation bursts were an artifact of firms reporting similar products separately, one would expect average revenue per new product to fall sharply in the number of new products. This would also be the case if the new products were close substitutes. To the contrary, Appendix Figure A4 shows that average revenue per product is constant in the number of new products that firms introduce.

As a further robustness test, in Appendix Figure A5, we show that the distribution of product innovation is similar across different level of aggregation. The idea is that if innovation bursts are driven by firms differing in the care with which they report revenue separately for each 10-digit product, the distribution should become less concentrated at higher levels of aggregation. The figure shows that this is not the case: the distribution of the number of products that firms add is similar at the 6-digit, 8-digit and the 10-digit level. We thus conclude that innovation bursts are episodes of rapid firm expansion, rather than an artifact of the way that some firms report changes to their product portfolio.

Finally, we verify that innovation bursts are not driven by mergers and acquisitions. While firms may expand their portfolio rapidly by purchasing other firms, such acquisitions are unlikely to involve product innovations. While it is difficult to track mergers and acquisitions in the data, we are able to track whether firms have seen a change in their intangible capital. The chief component in intangible capital is goodwill, which rises when firms acquire a competitor. Appendix Figure A6 shows that the concentration of product innovation is not driven by mergers and acquisitions: when excluding firms with for changes in intangibles, the figure shows that the distribution is almost unchanged.

Figure 3. Distribution of Number of New Products

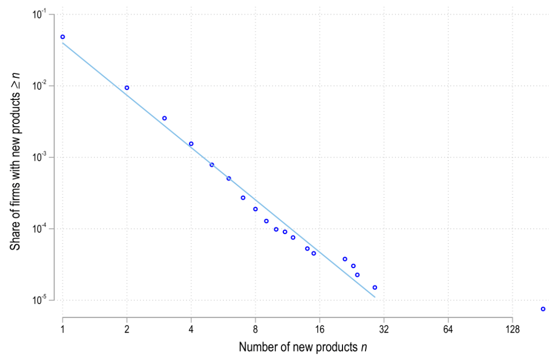

(a) Distribution of New Products ($N_{t-1} = 1$)

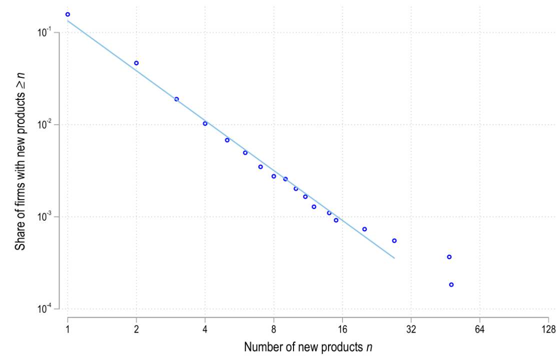

(b) Distribution of New Products ($N_{t-1} = 5$)

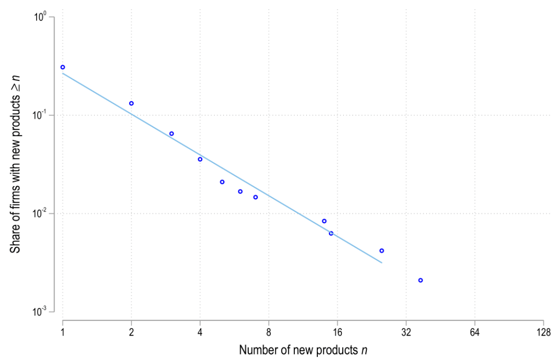

(c) Distribution of New Products ($N_{t-1} = 10$)

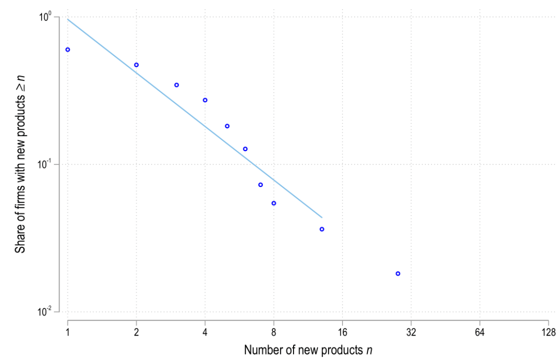

(d) Distribution of New Products ($N_{t-1} = 20$)

Notes: The figures plot the relationship between a firm's size (horizontal axes) and the firm's inverse rank (vertical axes). Size is measured through either revenue (in 2015 euros) or employment (in full-time equivalence). Inverse rank is measured as the ratio of firms' rank starting from the smallest firm, divided by the total number of observations in the data. Plots are based on the FARE-EAP sample for 2019. Linear reference lines exclude the observation with the greatest number of added products.

Conditioning on Size, Age or Sector The product innovation and destruction patterns are present across firms of different sizes, ages and industries. We have thus far have shown that rates of product innovation and destruction, as well as the high concentration of product innovation, are similar for firms with various initial product counts. In Appendix Figure A7, we show that product innovation is similarly concentrated with separating the sample in three age bins, although young firms are on average smaller and thus on average add fewer products to their portfolio. While all results are for manufacturing because of the nature of our data, we find similar rates (and concentration) of product innovation and destruction rates across our sample's three primary Broad Economic Classification industries – capital good production, industrial supplies production and consumer good production.

3.3. Firm Growth through Product Innovation and Destruction

How important is product innovation as a driver of firm growth? In addition to introducing new products, firms can further engage in innovation by improving the process of producing their existing prod-

Table 3: Contribution of Product Innovation and Destruction to Revenue Growth

	Overall Growth	Product Innovation	Product Destruction	Continuing Products
Mean value				
1-year	0.004	0.022	-0.024	0.006
5-year	0.027	0.076	-0.075	0.026
Shapley-Owen Contribution				
1-year	100.0	16.0	55.2	28.8
5-year	100.0	37.3	49.8	12.9

Notes: The table decomposes total revenue growth into revenue loss from product loss, revenue gain from product gain, and changes in revenue on products that the firm is continuing to produce. Bootstrapped standard errors from 1000 replications are given in parentheses. Revenue is deflated by the GDP deflator. Observations in the upper panel are weighted by the denominator $0.5(R_{it} + R_{it-h})$.

ucts. In fact, Garcia-Macia et al. (2019) have recently argued, using evidence drawn based on the distribution of the overall employment growth of firms, that such process innovation plays a substantially larger role in driving aggregate productivity growth compared to that of product innovation. In this section, we rely on our data to decompose the contributions of product and process innovation to the dynamics of firm growth.

We measure firm growth using the symmetric growth rate introduced by Davis et al. (2006), which divides the change in firm size between period t and $t-h$ by the average of size across both periods. In contrast to changes in log size, this growth rate can be calculated for continuing firms as well as firms that enter or exit, which is why it has become a standard measure in recent work on firm and employment dynamics. Using our data, we can decompose the revenue growth of each firm into changes in revenue that come from the entry and exit of products from the firm's product portfolio, and from changes in the revenue that firms earn on their existing product portfolio. We use the following decomposition for the change in the revenue of R_{it} of firm i at time t:

$$\frac{R_{it} - R_{it-h}}{\frac{1}{2}(R_{it} + R_{it-h})} = \frac{R_{it}^{N}}{\frac{1}{2}(R_{it} + R_{it-h})} - \frac{R_{it-h}^{L}}{\frac{1}{2}(R_{it} + R_{it-h})} + \frac{R_{it}^{-} - R_{it-h}^{+}}{\frac{1}{2}(R_{it} + R_{it-h})}, \tag{1}$$

over some horizon of length h. In the decomposition, the first two terms capture growth in revenue due to product innovation and destruction, while the final term captures revenue growth on products that firms continue to produce. The term R_{it}^{N} denotes revenue on products that the firm produces at time t but did not produce at time $t-h$, the term R_{it-h}^{L} is the revenue that the firm earned at time *t-h* on products that it stopped producing between time $t-h$ and time t, the term R_{it}^{-} is the revenue that the firm earns at time t on the products that it was already producing at time $t-h$, and the term R_{it-h}^{+} is therevenue the firm earns at time $t-h$ on this same set of continuing products.

We quantify the contribution of product innovation and destruction using the decomposition in Equation (1) in three ways. In the top panel of Table 3, we first present the average size of each term, both for a single and a five year horizon. Observations are weighted by the denominator on the left hand side of Equation (1) in order to measure contributions to aggregate revenue growth. The first row shows that the effect of product innovation and product destruction on overall growth are an order of magnitude greater than the effect of revenue growth on continuing products. Through the lens

Figure 4. Creative Destruction and Aggregate Revenue Growth

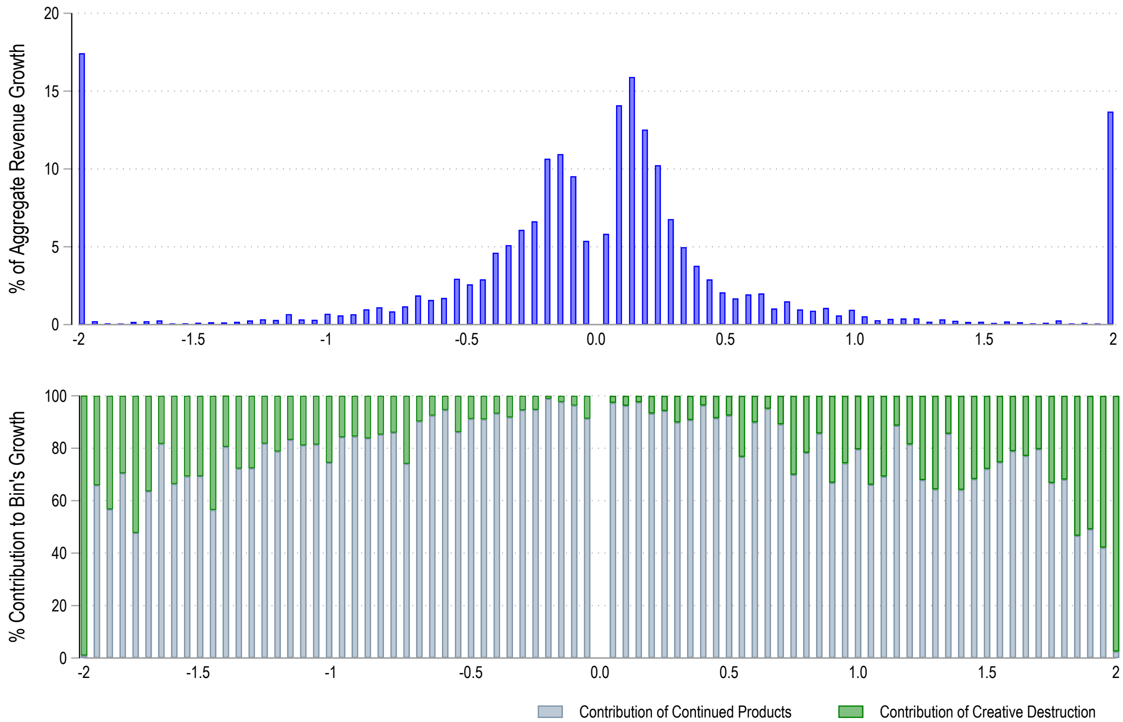

Notes: The horizontal axis measures firm growth through the symmetric growth rate, defined as the change in revenue between t and $t-1$ divided by average revenue in t and $t-1$. Growth rates are separated into 20 negative bins and 20 positive bins. The top figure presents the contribution of changes in revenue across firms in a particular growth bin as a percentage of total revenue creation (the sum of increases in revenue across growing firms) for positive bins or as a percentage of total revenue destruction (the sum of decreases in revenue across shrinking firms) for negative bins. The bottom panel decomposes a bin's overall revenue change into changes coming from continuing products and the net of product innovation and destruction – creative destruction.

of this simple exercise, it may thus appear from the table that the net effect of product innovation and product destruction–henceforth, labeled creative destruction–is small or even negative.[8] As we will show below, this average result masks a substantial degree of heterogeneity across firms in the contribution of creative destruction to their growth.

However, the picture changes when we consider a decomposition of revenue growth at the firm level. The bottom panel of Table 3 measures the contribution of product innovation and destruction to overall revenue growth of firms through the Shapley value (see, e.g., Ozkan et al., 2023). The Shapley value assigns importance to each of the components in revenue growth in explaining firms' overall revenue growth. To do so, it quantifies each component's average marginal contribution to the R^2 in regressions on all possible combinations of the components. The table shows that creative destruction explains over 70% of revenue growth over 1-year horizons, and more 87% over 5-year horizons. The contribution of product innovation is more similar to product destruction over the longer horizon,

[8]Note that we exclude firms from the decomposition in their first year in the product-level EAP data if they appeared in earlier years of the census of firms. As a result, innovation bursts that push firms over the size threshold to feature in the EAP are therefore excluded from the data. This quantification on the importance of product innovation is thus likely to be a lower bound.

which is likely because product sales do not reach their full scale in the first year in the firm's product portfolio.[9]

The large contribution of creative destruction is largely due to its effect on the tails of firm revenue growth. We illustrate the contribution of creative destruction to the tails of growth and the importance of those tails for aggregate revenue growth in Figure 4. We split the sample of firms into 40 equally-sized bin based on the firms' revenue growth. The top figure quantifies the importance of the tails of revenue growth. On the right side of the figure, for firms with positive growth, the figure plots the ratio of the total change in revenue across the firms in each bin by the total change in revenue across all firms with positive growth. Thus, the height of the bars account for the contribution of firms in each bin to the overall growth of all growing firms. On the left side, we similarly report the ratio of the total fall in the revenue of all firms within each bin to the overall fall in the revenues of all shrinking firms.[10] The figure shows that 18% of revenue destruction comes from firms with revenue growth between -1.95 and -2, while 14% of revenue creation comes from firms with revenue growth between 1.95 and 2. This means that, while most changes in overall revenue growth originates from the many firms that have modest growth rates, the tails of revenue growth matter in aggregate. In the bottom figure we decompose a bin's total change in revenue into changes in revenue for continuing products (grey) and the net change from creative destruction (green). The figure shows that the larger the change in revenue growth, the greater the share of overall growth that comes from creative destruction.

Concentration of Revenue Growth We have so far shown that creative destruction plays an important role in the tails of firm growth and that product innovation is highly concentrated. A natural question is whether overall revenue growth is therefore highly concentrated, and well described by a power law, as well. In Online Appendix C.4 we confirm that the distributions of revenue and employment growth are much more concentrated than, say, a log-normal distribution would predict. The log-normal distribution is a natural starting point as, motivated by Gibrat (1931), firm dynamics models with random shocks to productivity typically assume that firm growth is log normal (see, e.g., Hopenhayn 2014).[11]

Conditioning on Size, Age or Sector The patterns we describe above are similar for firms of different sizes, ages, and primary sectors. In Online Appendix C.4, we show that the contribution of creative destruction to overall revenue growth is similar for firms of different ages and sizes, although the contribution is (by definition) exhaustive for entrants. The appendix also shows that revenue growth is fat tailed for firms of different ages and sizes.

[9] For example, if firms start producing new goods uniformly over the year, revenue for new products will on average only capture half a year of sales. This lowers the measured contribution of product innovation to revenue growth.

[10] This approach is similar to how the contribution of firms with different growth rates to overall job creation and job destruction is calculated in Garcia-Macia et al. (2019).

[11] Our results echo recent findings in the literature on income dynamics which shows that individual income growth is also better described by a fat-tailed distributions such as the Pareto distribution (see, e.g., Guvenen et al. (2021). Gabaix et al. (2016) discuss the importance of the income growth process to explain stationary distributions of income and wealth.

3.4. Revenue Growth over a Product's Life Cycle

While creative destruction is key for the two tails of firm growth, the bulk of overall revenue growth appears to come from products that firms continue to produce. This may imply that process innovation, enabling firms to produce their existing products at higher quality or lower costs, could be the primary driver of economic growth. We next show, however, that the growth of revenue of a continuing product gradually falls over the course of its life cycle. Thus, by introducing new products that have higher potential for process innovation, creative destruction makes an indirect contribution to growth both at the firm and at the aggregate levels.

In Figure 5, we plot the relationship between average revenue growth and product tenure–the time that has passed since a firm started producing the good. Before calculating average revenue growth by tenure, we first subtract a firm-product's average revenue growth over the entire spell of its life cycle obverved in our data. This is to avoid a selection bias arising from the fact that revenue growth might affect how long a firm produces a product. Hence, the figure plots how a firm-product's revenue growth deviates from its life cycle average as tenure increases. Each point in the figure gives the average of the de-meaned revenue growth for a given years of tenure, while 90% confidence intervals based on bootstrapped standard errors.

The figure shows that there is a significant decline in the growth rate of revenue over the life cycle of a product. When firms have been producing a good for one year, revenue growth on their product is on average 2 percentage points higher than the spell's average. For comparison, average revenue growth across all continuing firm-products is -2.1% in the data. In subsequent years, revenue growth declines by an average of 0.5 percentage points per year. By year 8, which is the maximum horizon over which we can study products' life cycle with our 9-year dataset, revenue growth is on average 2 percentage points below the spell average.

The decline of revenue growth over the life cycle of a product means that creative destruction and revenue growth on continuing products interact. For firms to sustain growing real revenue growth on their products, they must improve the quality of the products that they sell or cut the costs at which they produce, in order to sell at lower prices. When firms are just starting to produce a new product, their revenue growth is rapid and exceeds the spells average. It thus appears that opportunities for product innovation are large at the start of a product's life cycle within of a firm, but growth falls over the product's life cycle – eventually turning significantly negative. Opportunities to sustain growth through process innovation thus appears to decline. This means that creative destruction offers two benefits: besides directly contributing to revenue growth of the firm that engages in product innovation, it also enables a sequence of process innovations that deliver further revenue growth.

Conditioning on Size, Age or Sector The decline in revenue growth in a product's tenure is present across firms of different initial sizes, ages, and industries. In Appendix Figure A8 and A9, we respectively show that revenue growth declines in product tenure for firms across firms of different sizes (as measured through their product count) and firms of different ages. In Appendix Figure A10, we further show that the decline in revenue with tenure is visible across products across our three main Broad

Figure 5. Life Cycle of Revenue Growth

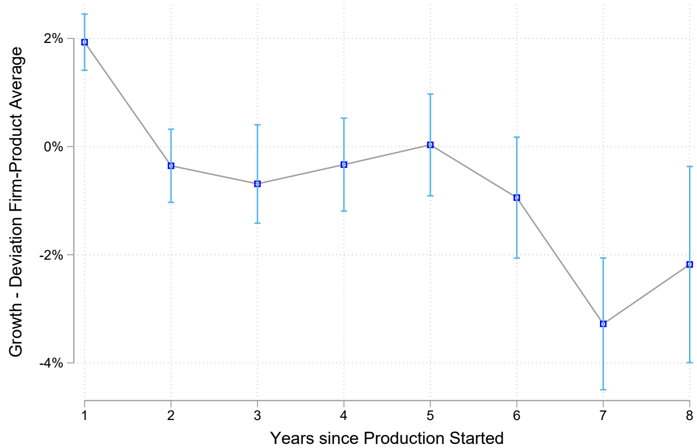

Notes: The vertical axis plots average growth of product revenue among all firm-products in the sample, where the average growth rate of a firm-product's revenue is subtracted from growth at each horizon. The horizontal axis plots the number of years that have passed since the firm first started producing the product. Confidence bounds are at the 90% level and based on bootstrapped standard errors.

Economic Classification industries, though more strongly in consumer good production than in capital good production.

3.5. Discussion

This section has established three new facts about products and product innovation. First, we show that the concentration in economic activity in a small fraction of firms comes from a combination of concentration in the number of products that firms produce, and high variation in the amount of revenue that firms earn per product. We then show that the product portfolio of firms changes frequently,. We furthermore show that creative destruction is an important contributor to overall revenue growth, through two channels: 1) large bursts of product innovation and creative destruction cause fat tails in the distribution of revenue growth, and 2) innovation bursts are a necessary ingredient for sustained process innovation, as firms are best at improving their continuing products early in their life cycle.

4. The Baseline Model

In this section we describe our parsimonious model of endogenous growth through innovation bursts. We start by describing the static and dynamic optimization of the firm, and then characterize the model's stationary balanced growth path equilibrium.

4.1. Preferences, Market Structure and Intermediate Good Production

An infinitely-lived representative household has log utility over instantaneous consumption C_t, which is discounted at a rate ρ over time. Time is continuous and indexed by t, which is omitted when convenient. The household owns a single unit of labor which it supplies inelastically. Consumption is an aggregate of a continuum of differentiated goods with a constant elasticity of substitution:

$$C_t = \left[\int_0^1 \left(\sum_{i \in I_j} q_{ij} y_{ijt} \right)^{\frac{\epsilon-1}{\epsilon}} dj \right]^{\frac{\epsilon}{\epsilon-1}} \tag{2}$$

where $\epsilon > 1$ is the elasticity of substitution, goods are indexed by j, producers are indexed by i, and I_j denotes the set of firms that owns the technology that enables them to produce good j at a particular level of quality q_{ij}. Output of firm i on product j is denoted y_{ijt}.

The outputs of the potential producers of good j are perfectly substitutable. Firms compete à la Bertrand, which means that the firm with the technology to produce the good at the highest level of quality is the good's sole producer. This firm produces j with the linear production technology

$$y_{ij} = l_{ij},$$

where l_{ijt} denotes the quantity of labor that firm i dedicates to the production of good j. It operates on a competitive labor market and takes the wage w as given. The monopolist producer of good j sets a price p_{ij} that maximizes profits $p_{ij} y_{ij} - w l_{ij}$ subject to demand

$$y_{ij} = p_{ij}^{-\epsilon} Y q_{ij}^{\epsilon-1}.$$

Following Acemoglu et al. (2012) and Acemoglu et al. (2018a), we assume that any firm that would like to offer good j on the market faces an infinitely small price posting costs. As non-producing are unwilling to incur this cost, only the firm with the highest level of quality to produce good j posits a price. Hence, the price for good j is the familiar markup $\epsilon/(\epsilon-1)$ over marginal costs w. Profits are therefore

$$\pi(q_{ij}) = \left(\frac{1}{\epsilon-1}\right)\left(\frac{\epsilon-1}{\epsilon}\right)^{\epsilon} Y \hat{q}_{ij}^{\epsilon-1} \tag{3}$$

where $\hat{q}_{ij} \equiv q_{ij}/w$ is a measure of the relative quality of good j produced by firm i.

4.2. Product Innovation and Innovation Bursts

While each good is produced by a single firm, firms are able to produce all of the goods for which they own the highest-quality patents.

Incumbent product innovation Incumbent firms expand the portfolio of patents through bursts of innovation. An innovation burst is a collection of product innovations, each of which contains the technology that enables a firm to produce a product that it did not yet produce. The innovating firm can then produce each of the products in its innovation burst at a higher level of quality than the respective product's incumbent. Formally, an innovation burst can be represented as a set with n_{it}^c elements, where each element contains a quality improvements $\lambda_{i\tilde{j}}$ and where each $\tilde{j} \in [0,1]$ denotes a randomly chosen product that i acquires a technology to produce. The quality $q_{i\tilde{j}}$ at which the innovator produces its new goods is given by

$$q_{i\tilde{j}} = q_{-i\tilde{j}} + \lambda_{i\tilde{j}}\overline{q},$$

where $q_{-i\tilde{j}}$ is the incumbent's quality while $\overline{q}$ denotes the average quality of the economy's products.[12]

The number of products for which a quality improvement is contained in the firm's innovation burst at time t is n^c_{it}. We assume that n^c_{it} is a random variable that follows a Zeta distribution, such that probability mass function of n^c_{it} is given by

$$P(n^c_{it} = n) = \frac{n^{-\theta}}{\zeta(\theta)}, \quad \text{where} \quad \zeta(\theta) = \sum_{i=1}^{\infty} \frac{1}{i^{\theta}},$$

which is the discrete counterpart of the probability density function of a Pareto distribution. The tail parameter $\theta > 1$ determines the thickness of the tail of innovation bursts. As θ approaches one, innovation bursts have a Zipf's distribution. As θ increases, innovation bursts become less dispersed. The average number of product innovations in a burst is given by the ratio of $\zeta(\theta-1)$ and $\zeta(\theta)$, which is finite as long as θ exceeds two. In the limiting case where θ becomes large, firms only add a single product to their product portfolio when they acquire an innovation burst.

Firms choose the Poisson flow rate $x_i \geq 0$ at which innovation bursts arrive. In order to achieve x_i, firms employ rd^x researchers along

$$rd^x(x_i, n_i) = \eta x_i^{\psi} n_i^{-\sigma(\psi-1)}$$

where $\psi > 1$ and $0 \leq \sigma \leq 1$. The parameter σ governs the degree to which the probability that a firm's optimal innovation rate scales with its size. For $\sigma = 1$, firms choose an arrival rate of innovation bursts that is proportional to the number of products that they currently produce. For $\sigma = 0$ a firm's optimal arrival rate of innovation bursts does not depend on its current size, which means that the expected percentage growth rate in firm size rapidly declines with size.

Our setup nests the canonical Klette and Kortum (2004) model of endogenous growth through product innovation. In that model, firms choose the Poisson rate with which they acquire the technology to produce a single additional good. As a consequence, the number of additional products that a firm produces over a discrete time interval has a Poisson distribution. For large values of the innovation burst parameter θ, firms in our model also acquire only a single production technology when they innovate. The model would, in that case, also predict that the number of additional products has a Poisson distribution. In addition, firms' innovation capacity in Klette and Kortum scales perfectly with size, $\sigma = 1$ in our model, while products are aggregated along a Cobb-Douglas aggregator, i.e. $\epsilon = 1$.

[12] By assuming that quality improvements are multiplied by the average quality rather than a good's quality we follow Acemoglu et al. (2018a) and assure that the distribution of qualities in equilibrium is non-degenerate. This is important because quality determines a product's demand and thus revenue under a CES aggregator with $\epsilon > 1$.

Entrant product innovation There is a continuum of potential entrants that employ researchers in order to acquire the technology to produce one of the incumbents' products at a higher level of quality. Their employment of researchers rd^e is related to the rate of entry s along

$$rd^e = \eta^e s^{\psi_s}$$

where $\eta^e > 0$ and $\psi \geq 1$. Potential entrants that draw an innovation improve the quality of an incumbent's product in the same way that innovation by incumbents raises a good's quality.

4.3. Process Innovation

Incumbent firms also improve the quality of their existing products, which we henceforth label process innovation. Firms choose the rate at which they improve the quality of each of their products, and when a product's process innovation is successful its quality rises along

$$q_{ij}(t+\Delta t) = q_{ij}(t)(1+\lambda\gamma^{s_{ij}}).$$

where $\lambda > 0$ and where s_{ij} counts the previous process innovation that firm i has achieved on good j. As $0 < \gamma \leq 1$, innovation steps on process innovation may become smaller over time, which represents the idea that process innovations become increasingly incremental as a product within a firm matures.

To improve the quality of existing product j at a Poisson rate $z_{ij} > 0$, the firm employs

$$rd^z(z_{ij}, s_{ij}) = \eta_z z_{ij}^{\psi} \left(\left(\widehat{q}_{ij}(1+\gamma^{s_{ij}}) \right)^{\epsilon-1} - \widehat{q}_{ij}^{\varepsilon-1} \right)$$

researchers, where $\eta_z > 0$. The R&D expense on process innovation rises in the quality of the product and falls in the number of previous process innovations. The former implies that it is more expensive to improve a more advanced product. As we show in the next section, these assumptions combined assure that the returns to process innovation do not depend on $\widehat{q}_{ij}$ and s_{ij}, such that firms choose equal process innovation rates across their product portfolio.

4.4. Growth and Creative Destruction

Firms stop producing a good if a different incumbent or an entrant comes up with a higher-quality version of that product. The rate at which this happens to a product is the rate of creative destruction. The rate of creative destruction is endogenous, as it is determined by the rate of entry and the effort to acquire innovation bursts by incumbents:

$$\tau = e + \sum_{n=1}^{\infty} \left(\frac{M_n x_n}{\zeta(\theta)} \sum_{k=1}^{\infty} k^{-\theta} \right)$$

The first term reflects creative destruction due to entry while the second term reflects creative destruction due to innovation by incumbents. The outer summation of the incumbent innovation term is over the number of products (size) of each incumbent firm. Within the summation, total creative destruction by firms of size n is the product of the mass of firms of that size (M_n), the flow rate at which these firms draw innovation bursts (x_n), multiplied by the number of products that they expect to acquire a technology for as part of the innovation burst–which is captured by the final summation. Using the mean of the Zeta distribution, the rate of creative destruction can be written as

$$\tau = e + \frac{\zeta(\theta-1)}{\zeta(\theta)} \sum_{n=1}^{\infty} M_n x_n \tag{4}$$

Firms exit if innovation by incumbents or entrants causes them to cease producing their sole good. The rate at which this happens is the creative destruction rate, which depends on both the rate at which innovation bursts arrive and the size of innovation bursts, as explained in the next section.

4.5. Equilibrium

We now characterize the full balanced growth path equilibrium where productivity, consumption, output and wages grow at a rate g.

Value Functions and First Order Conditions/ Equilibrium Definition Firms employ the number of researchers that maximizes the firm's value along the value function

$$rV_t(J_i) - \dot{V}_t(J_i) = \max_{x_i, z_{ij}} \left\{ \begin{array}{c} \sum_{j\in J_i} \pi_t(q_{ij}) \\ +\sum_{j\in J_i} \tau \left[V_t\left(J_i \setminus \{q_{ij}, s_{ij}\}\right) - V_t(J_i)\right] \\ +\sum_{j\in J_i} z_{ij} \left[V_t\left(J_i \setminus \{q_{ij}, s_{ij}\} \cup_+ \{q_{ij}(1+\gamma^{s_{ij}}), s_{ij}+1\}\right) - V_t(J_i)\right] \\ +x_i \sum_{h=1}^{\infty} \frac{h^{-\theta}}{\zeta(\theta)} \mathbb{E}\left[V_t(J_i \cup_+ \{q_{ik}, k=1,..,h\}) - V_t(J_i)\right] \\ -w_t \eta(x_i)^{\psi} n_i^{-\sigma(\psi-1)} - w_t \eta_z \sum_{j\in J_i} (z_{ii})^{\psi} \left(\left(\widehat{q}_{ij}(1+\gamma^{s_{ij}})\right)^{\epsilon-1} - \widehat{q}_{ij}^{\varepsilon-1}\right) - F(n_i) \end{array} \right\} \tag{5}$$

where r is the interest rate and where $\dot{V}_t$ denotes the change in firm value V_t with time. The first line on the right-hand side contains the sum of the flow of profits, $\pi_t(q_{ij})$. The second line contains the expected change in firm value if the firm stops producing good j because of creative destruction. The term $V_t\left(J_i \setminus \{q_{ij}, s_{ij}\}\right)$ denotes the value of producing the portfolio of goods J_i except good j that has quality q_{ij} and on which the firm has implemented s_{ij} process innovations. The third line contains the change in value that arises from process innovation, which is again summed across all of the products that the firm produces. The fourth row contains the expected increase in firm value from successful innovation. This is equal to the Poisson rate at which the firm chooses to obtain innovation bursts, $_i$, multiplied by the expected increase in the firm's value if it acquires an innovation burst. $V_t(J_i \cup_+ \{q_{ik}, k=1,..,h\})$ denotes the increase in value if the innovation burst contains the technology to add h products, with the k-th products being of quality q_{ik}. This is weighted by the probability density function of h, which is $h^{-\theta}/\zeta(\theta)$. The first and second term in the bottom row contain the firm's total expenditure on researchers, which are paid a wage w. The final term $F(n_i)$ is a fixed costs that firms must pay in order to operate, which we assume equals the option value of research and development.[13]

The following proposition characterizes the value function as the solution to the problem in Equation (5).

Proposition 1. *The value function of a firm with intangible cost ϕ_i that produces a portfolio of goods J_i with cardinality n_i grows at rate g along the balanced growth path and is given by*

$$V(J_i) = \frac{\sum_{j\in J_i} \hat{\pi}\hat{q}_{ij}^{\epsilon-1}}{r-g+\tau},$$

where $\hat{\pi} \equiv \left(\frac{1}{\epsilon-1}\right)\left(\frac{\epsilon-1}{\epsilon}\right)^{\epsilon} Y$, which grows at rate g. The optimal rate of innovation reads as

$$x_{n_i} = \left(\left[\frac{\hat{\pi}\mathbb{E}\left[\hat{q}_{ij}^{\epsilon-1}\right]}{r-g+\tau}\right]\frac{1}{\eta\psi w}\frac{\zeta(\theta-1)}{\zeta(\theta)}\right)^{\frac{1}{\psi-1}} n_i^{\sigma}. \tag{6}$$

The optimal process innovation rate reads as

$$z = \left(\left[\frac{\hat{\pi}}{r-g+\tau}\right]\frac{1}{\eta_z\psi w}\right)^{\frac{1}{\psi-1}}. \tag{7}$$

The optimal startup rate is given by

$$s = \left(\left[\frac{\hat{\pi}\mathbb{E}\left[\hat{q}_{ej}^{\epsilon-1}\right]}{r-g+\tau}\right]\frac{1}{\eta^{e}\psi w}\right)^{\frac{1}{\psi-1}}. \tag{8}$$

Proof: Appendix A.

Quality Shocks To mimic the empirical variability of product-level revenue we introduce idisyncratic quality shocks. These shocks are log-normally distributed with zero mean and standard deviation Ξ.

Household Optimization The household maximizes life-time utility with respect to consumption and saving. Given log utility over consumption with discount rate ρ, this yields the Euler equation:

$$\frac{\dot{C}}{C} = r-\rho. \tag{9}$$

Along the balanced growth path, consumption grows at the rate g, such that $r-g=\rho$.

Firm-Size Distribution The mass of single product firms expands through the entry of new firms, which happens at a rate e, and because the mass of firms that produce two products declines through

[13]We borrow this ad-hoc restriction from Akcigit and Kerr (2018). It assures that the value function is linear in the number of goods that firms produce, so that the model admits analytical first-order conditions for research and development. **?** shows that the economic effects of this assumption are offset by the calibration of σ, as both have a similar effect of the degree to which optimal research and development scales with the number of products that firms produce.

creative destruction. Simultaneously, the mass of single-product firms declines because these firms also expand (through innovation bursts at rate x_1) or exit (through creative destruction at rate τ_1). The flow equation of the single-product firms is given by

$$\dot{M}_1 = e + M_2\tau_2 - M_1(x_1+\tau_1) \tag{10}$$

where, given that $\tau_n = n\tau$, we find $\tau_1 = \tau$. For larger firms, the flow equation looks similar. The main difference is that the mass of firms of size n expands through the arrival of innovation bursts at firms of any size smaller than n:

$$\dot{M}_n = M_{n+1}\tau_{n+1} + \left(\sum_{k=1}^{n-1} M_k \cdot x_j \cdot \frac{[n-k]^{-\theta_k}}{\zeta(\theta_k)}\right) - M_n\,(x_n+\tau_n) \quad \forall n > 1. \tag{11}$$

The first two terms in Equation (11) capture the two sources through which the mass of size n firms expands. First, firms of size $n+1$ lose a product through creative destruction at Poisson rate τ_{n+1}, which is the first term in (11). The second term captures the increase in the mass of size n firms through innovation bursts by firms that initially produce $k \in \{1,..,n-1\}$ goods that increase the size of those firms' product portfolio by $n-k$ goods. The final term captures the two sources through which the mass of size n declines: creative destruction, at rate τ_n, and innovation bursts, at rate x_n.

Along the balanced growth path the firm-size distribution is stationary, such that $\dot{M}_n = 0$ for any size n. The overall mass of firms is also constant, which means that exit through creative destruction is exactly offset by entry such that $e = \tau_1 M_1$. This pins down the mass of single-product firms:

$$M_1 = \frac{e}{\tau_1}.$$

Inserting the mass of single-product firms into Equation (10) gives the mass of two-product firms:

$$M_2 = M_1\left(\frac{x_1+\tau_1}{\tau_2}\right)$$

Setting the left-hand side of (10) to zero gives a sequence for the mass of any size bin $n > 2$ as a function of the mass of smaller-sized firms:

$$M_{n+1} = \frac{1}{\tau_{n+1}}\left(M_n(x_n+\tau_n) - \left(\sum_{j=1}^{n-1} M_j \cdot x_j \cdot \frac{[n-j]^{-\theta_j}}{\zeta(\theta_j)}\right)\right). \tag{12}$$

Aggregate Variables and Aggregate Productivity Growth As the representative household supplies a single unit of labor, equilibrium on the labor market requires that

$$1 = L^p + L^{rd} + L^e,$$

where L^p, L^{rd}, and L^e respectively denote labor used to produce intermediate goods, to conduct research for incumbents, and to conduct research for entrants. Inserting the production function, demand function and optimal prices, gives a straightforward expression for total production labor:

$$L^p = \frac{Y}{w}\left(\frac{\epsilon}{\epsilon-1}\right)^{-\epsilon}\left(\int_0^1 \int \mathbb{I}_{j\in J_i}\hat{q}_{ij}^{\epsilon-1} di dj\right)$$

where $\mathbb{I}_{j\in J_i}$ is the indicator function that equals one when firm i produces good j. Total employment of researchers by incumbents and entrants is given by

$$L^{rd} = \sum_{n=1}^{\infty}\left(M_n \eta x_n^{\psi} n^{-\frac{\sigma}{\psi-1}}\right) \quad \text{and} \quad L^e = \eta^e s^{\psi}.$$

The equilibrium wage is found by inserting the production function, demand function and optimal prices into the aggregator gives the equilibrium wage:

$$w = \left(\frac{\epsilon-1}{\epsilon}\right)\left(\int_0^1 \int \mathbb{I}_{j\in J_i} q_{ij}^{\epsilon-1} di dj\right)^{\frac{1}{\epsilon-1}},$$

As usual, the equilibrium wage is a constant mark-down over the CES productivity index Q:

$$Q = \left(\int_0^1 \int \mathbb{I}_{j\in J_i} q_{ij}^{\epsilon-1} di dj\right)^{\frac{1}{\epsilon-1}}. \tag{13}$$

This aggregate index grows at a rate g, where the growth rate is determined by the innovation decisions of incumbents and entrants. Aggregate productivity growth $g \equiv \partial \ln Q/\partial t$ is therefore given by:

$$g = \frac{1}{\epsilon-1}\frac{1}{Q}\int_0^1 \int \mathbb{I}_{j\in J_i}\left(\frac{\partial q_{ij}^{\epsilon-1}}{\partial t}\right) di dj = \tau \cdot \mathbb{E}[\lambda_{ij}]. \tag{14}$$

Equilibrium Definition Combining all of the conditions state above, we present the definition of the equilibrium in our model in Appendix A.

5. Quantification

We next detail how we quantify the model. The aim of the quantification is twofold. First, we show that when the model can be disciplined with moments that are directly obtained from our empirical analysis, and that the model is able to qualitatively and quantitatively replicate the empirical facts on creative destruction and process innovation from Section 3. Second, we use the resulting calibrated model to understand the importance of creative destruction and process innovation for overall productivity growth, and to demonstrate how matching the empirical facts documented in this paper reshapes our understanding of the sources of growth.

5.1. Calibration

We infer as many parameters as possible directly from the data on firm-and-product-level dynamics. This contrasts with the common approach to calibrating the models of creative destruction, which often relies only on firm-level information. Our data enables us to directly infer parameters from empirical rates of creative destruction, product innovation, and process innovation. The inferred parameters are summarized in Table 4.

Product Innovation The model has eight innovation parameters. The first parameter that we calibrate is the tail parameter of zeta distribution, θ. This parameter determines the average size of innovation bursts in the model, as well as the thickness of the tail of the innovation burst distribution. We calibrate θ to match the relationship between rank and the number of product innovations for single-product firms in Figure 3a. Because firms may experience multiple innovation bursts over the course of a year, the distribution of the number of product innovations differs from the distribution of the size of innovation bursts. Hence, we cannot infer θ directly from the slope of the rank-size relationship in Figure 3. At 3.05, θ matches the negative slope of 2.48 for the relationship between the number and rank of product innovations for single-product firms in the data.

Turning to the frequency of innovation bursts, there are two parameters: the frequency of innovation bursts relative to the number of products that firms currently produce (x), and the degree to which the frequency of bursts scales with size (σ). We set the latter to 1, to match the fact that there is no clear relationship between the number of products that firms initially produce and the ratio of their new products over their existing products in Table 2.

We next use data on the rate of product innovation to calibrate x. As firms on average achieve 0.066 product innovations for each of their existing products (Table 2), and because the average size of bursts $\zeta(\theta-1)/\zeta(\theta)$ is approximately 1.35, we set x to 0.049. Note that by directly calibrating the rate of product innovation to match the data, we do not have to separately estimate the unobserved parameters that drive optimal product innovation (the cost shifter η, curvature ψ and discount rate ρ).

As for the frequency of innovation bursts, we infer entry rate e from the data. As it is difficult for us to measure entry directly due to the exclusive sampling of large firms in EAP, we rely on the fact that the number of products in the model (and largely in the data) is constant over time. Firms lose an average of 0.082 products for every product that they initially produce, while they gain 0.066 products 2). As $\tau = x + e$, this means that entry e contributes 0.012 to the creative destruction rate.

For the average innovation step size of product innovation, η, we choose the value that delivers an overall productivity growth rate of 1.6%, which is the average French productivity growth in the Penn World Tables. Together with the process innovation parameters, this is achieved by setting η to 0.031.

Process Innovation The process innovation step size λ and the process innovation frequency z play a largely interchangeable role in the model. We set the Poisson rate of process innovations z to 1, such that firms develop improved versions of their existing projects on average once per year. This has the convenient implication of aligning the average number of improvements that firms have made in

Table 4: Summary of Parameter Values

Parameter	Description	Value
θ	Tail parameter of the zeta distribution	3.05
σ	Rate of innovation bursts scaling with firm size	1.00
x	Product innovation rate	.049
γ	Decline rate of follow-up process innovation size	.870
e	Entry rate adding to creative destruction rates	.016
λ	Process innovation step size	.041
z	Poisson rate of process innovations	1.00
η	Average innovation step size of product innovation	.031
Ξ	Standard deviation of idiosyncratic quality shocks	.271
ϵ	Elasticity of substitution	5.00

any given year equal to the tenure of their product, which gives the rate at which follow-up process innovation declines (γ) a straight-forward interpretation. We then calibrate λ to 0.041 to match the average growth of 0.06% of revenue on firms' continuing products, as found in the upper panel of Table 3. As the average tenure of a product is 13.2 years, the expected step size of process innovation is 0.0135. This means that the average productivity gain from successful process innovation is smaller than the average gain from product innovation, in line with the patent-level evidence in Akcigit and Kerr (2018).

The final parameter for process innovation to calibrate is the parameter controlling follow-up process innovation, γ, which determines the rate at which quality growth declines in a product's tenure. The decline in the growth of revenue between the first and second year that a firm produces a good is given by

$$\frac{\partial^2 \ln r_{ijt}}{\partial t \partial s} = \lambda \gamma^s (\epsilon - 1) \ln \gamma$$

where s is the product tenure. We calibrated parameter γ to 0.83, which delivers the 0.5 percentage-point rate of decline in revenue growth over the product life-cycle observed in Figure 5.

Productivity Shocks The idiosyncratic product-level productivity shocks create changes in sales that are unrelated to process and product innovation. Greater volatility of productivity shocks raises the contribution of continuing products' revenue growth to variation in overall firm revenue growth. We assume a log-normal distribution for productivity shocks with a log-average of zero and a standard deviation of Ξ. We set Ξ to 0.271 in order to match the 28.8% Shapley value of continuing products in overall revenue growth in Table 3.

Other Parameters Finally, the elasticity of substitution ϵ determines the degree to which dispersion in quality translates to dispersion in revenue per product and in revenue growth. It also controls the markup, which is frequently estimated in empirical work. We set ϵ to 5 for a markup of 1.25.

Table 5: Performance on Targeted Moments

Parameter	Moment	Target	Model
Product Innovation			
θ	Regression of number of new products (log) on rank (log) if N_{t-1} =1	-2.43	-2.43
σ	Scaling of number of innovation bursts in firm size	1.00	1.00
x	Arrival rate of innovation bursts (per existing product)	.048	.048
e	Entry rate	.012	.012
η	Growth rate of total factor productivity	.016	.016
Process Innovation			
z, λ	Revenue growth on continuing products	.006	.006
γ	Change in revenue growth in product tenure	-.005	-.005
Productivity shocks			
Ξ	Shapley-Owen contribution of continuing products to revenue growth	28.8%	28.8%
Other parameters			
ϵ	Markup	1.25	1.25

5.2. Model Performance

Before analyzing the sources of aggregate growth and the effect of innovation policy, we assess the model's ability to match the targeted moments from the previous section, as well as its ability to match untargeted moments, such as the stylized facts of Section 3. The results below come from a simulation of the quantified economy for 10,000 firms over 40 years.[14] The model's performance on targeted moments is summarized in Table 5, which shows that all moments are matched with precision.

We next assess the degree to which the model is able to replicate the empirical facts on products and product innovation from Section 3, which were largely untargeted in the calibration. In Figure 6 we compare the distribution of product innovations in the data and the model. Left-hand figures are from the data, while right-hand figures are based on simulations of the model. To understand the role of innovation bursts in the model, right-hand figures contain both patterns of product innovations with innovation bursts (blue-circled scatters) and patterns of product innovations if all episodes of product innovation added 1 product to a firm's portfolio (red-squared scatters). This is a natural benchmark, as innovation bursts are a main deviation from previous models of firm dynamics and creative destruction, such as Klette and Kortum (2004), Akcigit and Kerr (2018) and Acemoglu et al. (2018b).[15]

Figure 6 shows the model's ability to accurately match the high concentration of product innovation in the data. Pooling firms of all sizes in the topmost panelt, the linear relationship between the level and the rank of the number of product innovations fits almost perfectly with an R^2 of 0.99. Moreover, the slope of the relationship is -2.12, which is close to the untargeted slope of -1.94 in the data,

[14]We simulate the economy for 40 years rather than the 10 years of the data that we have for the EAP survey. This is because the fat-tailed distribution of product innovation implies that in small samples, convergence of sample moments to true moments may be slow.

[15]To implement the model without bursts, we set ζ to a large number and adjust the innovation rate x such that the model still matches the overall rate of product innovation in the data, as well as the overall growth rate of productivity

Figure 6. Distribution of Number of New Products: Data and Model

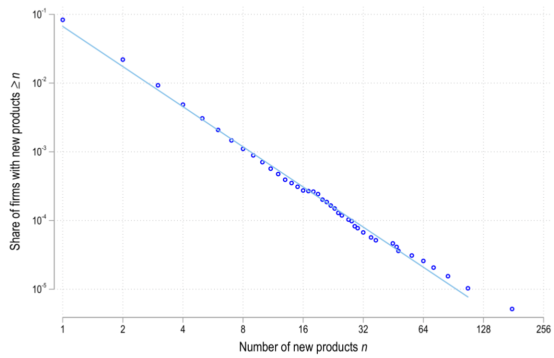

(a) Data: New Products (Any N_{t-1})

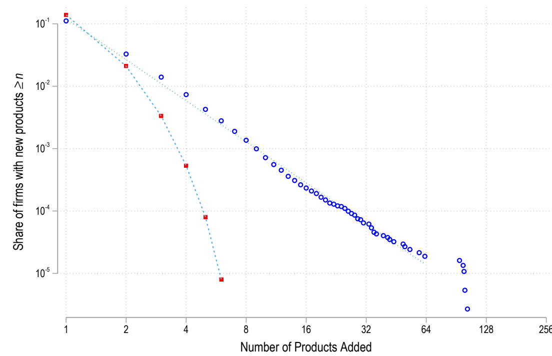

(b) Model: New Products (Any N_{t-1})

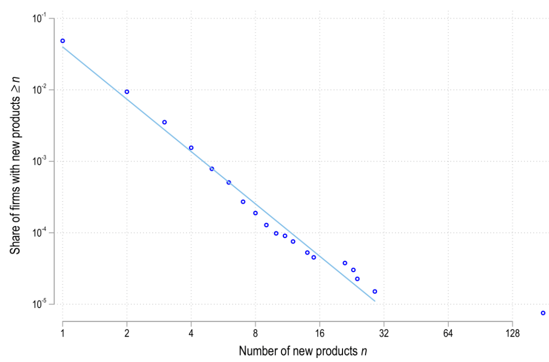

(c) Data: New Products ($N_{t-1} = 1$)

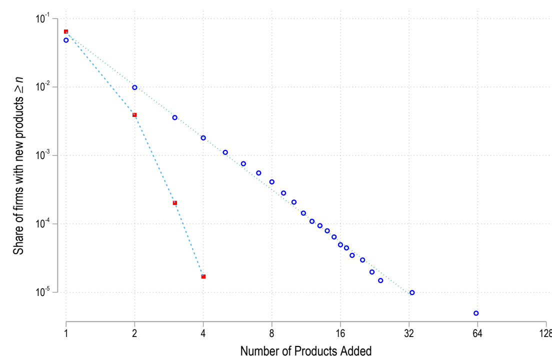

(d) Model: New Products ($N_{t-1} = 1$)

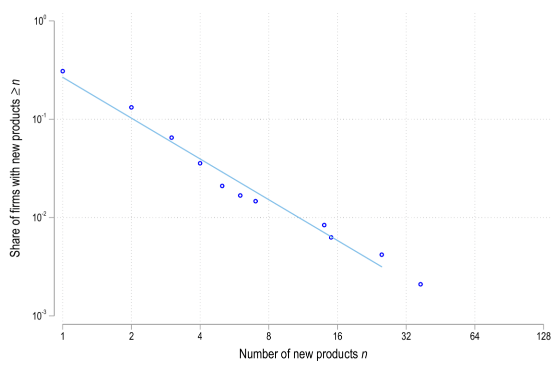

(e) Data: New Products ($N_{t-1} = 10$)

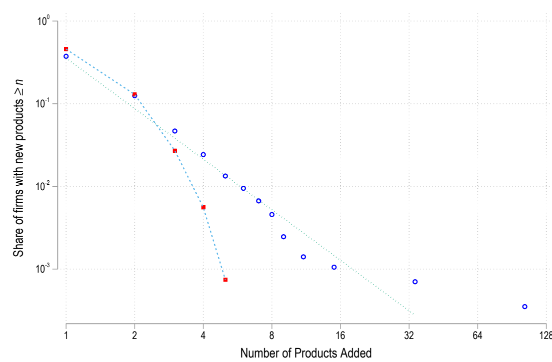

(f) Distribution of New Products ($N_{t-1} = 10$)

Notes: Left-hand figures are based on the EAP data. Blue-circled scatters on the right-hand side are from a model with innovation bursts. Red-squared scatters are based on a model with Poisson-distributed product innovation. Figures plot log-rank, log-size relationships.

suggesting that product innovation has a similar degree of concentration in the model and the data. When conditioning on size in the middle and lower panels, the model with innovation bursts again closely matches the empirical distributions. If product innovations do not come in bursts, in contrast, the model fits the empirical pattern of product innovations poorly. In this case, the resultaing distri-

Figure 7. Untargeted Moments: Product Concentration and Revenue Per Product

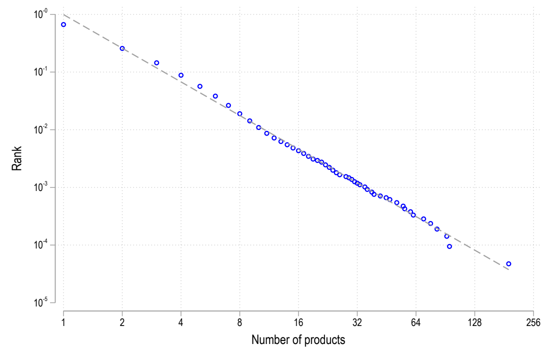

(a) Data: Number of Products

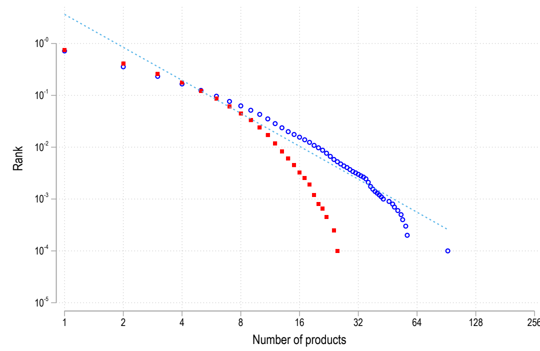

(b) Model: Number of Products

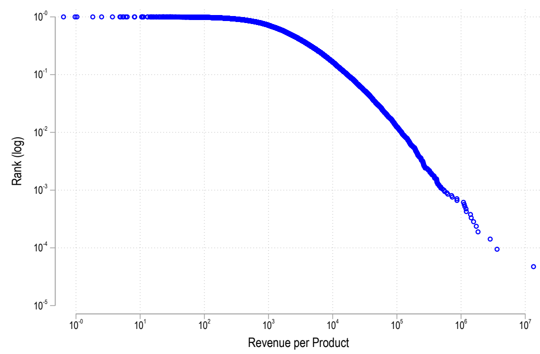

(c) Data: Revenue per Product

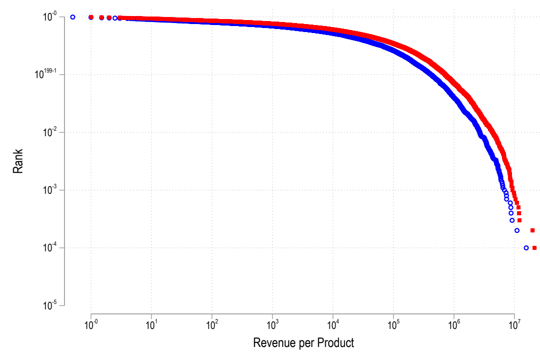

(d) Model: Revenue per Product

Notes: Left-hand figures are based on the EAP data. Blue-circled scatters on the right-hand side are from a model with innovation bursts. Red-squared scatters are based on a model with Poisson-distributed product innovation. Figures plot log-rank, log-size relationships.

bution of product innovation is Poisson, which leads to substantially lower degrees of concentration under the conventional model.

Besides matching the empirical pattern of product innovation, innovation bursts are also vital to match the concentration of firm size observed in the data. Figure 7 plots the distribution of the number of products that firms produce and their average average revenue per product, where, as before, the x- and y-axes show revenue and rank, respectively, in logarithmic scale. The figure shows that innovation bursts are vital for explaining the observed concentration of production across large firms. When product innovation comes in bursts, the distribution of the number of products in the model is similar to the distribution of the model. Instead, if product innovation is a Poisson process, the number of large firms is negligible. This is a familiar shortcoming of standard models of creative destruction (see, e.g., Akcigit and Kerr 2018).

Table 6 summarizes the model's ability to explain the fraction of firm growth that is respectively explained by product innovation, product destruction and revenue growth on continuing products. The model matches the fact that within the contribution of creative destruction, product innovation explains a smaller part of the variation in revenue growth than product destruction. This is due to the

Table 6: Untargeted Moments: Shapley-Owen Decomposition in the Model vs the Data

	Overall Growth	Product Innovation	Product Destruction	Continuing Products
1-year				
Data	100.0	16.0	55.2	28.8 (target)
Model	100.0	29.8	41.5	28.8 (target)
5-year				
Data	100.0	37.3	49.8	12.9
Model	100.0	35.5	58.6	5.89

Notes: The table decomposes total revenue growth into revenue loss from product loss, revenue gain from product gain, and changes in revenue on products that the firm is continuing to produce.

fact that firms do not earn revenue for the full year in which they start producing a good. The wedge between the contribution of product innovation and product destruction is greater in the data than in the model, which is likely due to the fact that we underestimate entry in the EAP data.[16] Over 5-year horizons, the model matches the (untargeted) contribution of product innovation and product innovation closely.

As a final untargeted moment, Figure 9 presents the relationship between the magnitude of firms' revenue changes. The model is able to replicate the broad empirical pattern that the bulk of overall revenue destruction and creation occurs at firms with modest changes in sales, and that continuing products are responsible for the bulk of that growth. The model is also able to match that the tails of growth, in particular firm entry and exit, make up a significant part of overall revenue creation and destruction. The model also qualitatively matches the fact that the larger a firm's change in revenue, the greater the contribution of creative destruction.

Figure 8. Life Cycle of Revenue Growth

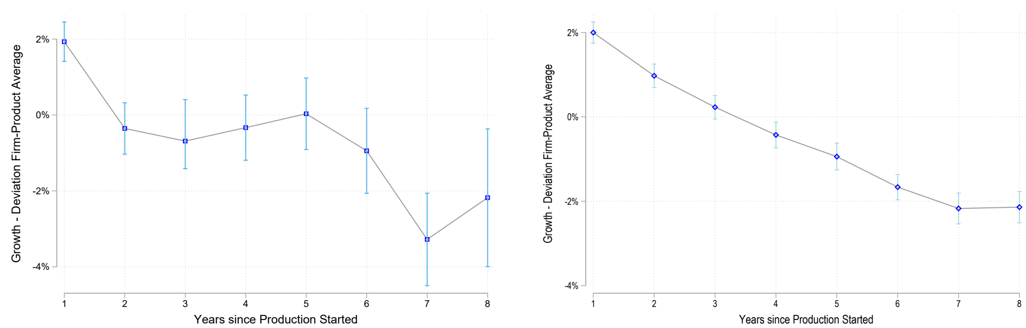

(a) Data (b) Model

Notes: The vertical axis plots average growth of product revenue among all firm-products in the sample, where the average growth rate of a firm-product's revenue is subtracted from growth at each horizon. The horizontal axis plots the number of years that have passed since the firm first started producing the product. Confidence bounds are at the 90% level and based on bootstrapped standard errors.

[16] Firms only enter the EAP data if they earn at least 5 million euros of revenue or employ at least 20 firms. We only consider a product a new product if we observe the firm prior to the introduction to a product, which means that we under-count entry of new products when firms add a product to their portfolio in the same year that they first appear in EAP.

Figure 9. Untargeted Moments: Creative Destruction and Aggregate Revenue Growth

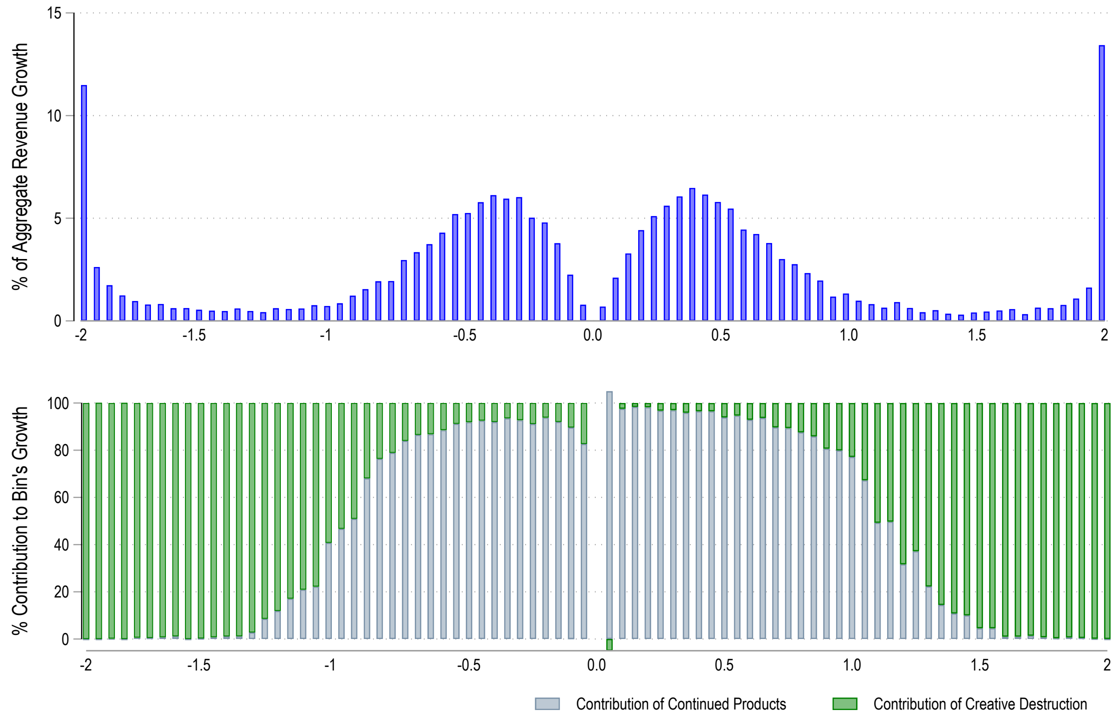

Notes: The horizontal axis measures firm growth through the symmetric growth rate, defined as the change in revenue between t and $t-1$ divided by average revenue in t and $t-1$. Growth rates are separated into 20 negative bins and 20 positive bins. The top figure presents the contribution of changes in revenue across firms in a particular growth bin as a percentage of total revenue creation (the sum of increases in revenue across growing firms) for positive bins or as a percentage of total revenue destruction (the sum of decreases in revenue across shrinking firms) for negative bins. The bottom panel decomposes a bin's overall revenue change into changes coming from continuing products and the net of product innovation and destruction – creative destruction.

6. Application: Creative Destruction as the Driver of Growth

We now apply the framework to understand the impact of creative destruction on aggregate productivity growth and market concentration. The calibrated model features various new interactions between creative destruction, process innovation, concentration, and aggregate growth. Informed by the stylized facts, product innovation is highly concentrated–which in turn means that an economy with frequent innovation bursts features greater concentration of production among large firms. Similarly, we find that the efficacy of process innovation declines in a product's tenure, which means that frequent innovation bursts can increase the effectiveness of process innovation by rejuvenating the product mix of firms. In this section we quantify these interactions and show how they reshape our thinking about the macroeconomic effects of creative destruction.

6.1. Additive Decomposition

A first way of quantifying the macroeconomic importance of creative destruction on growth is to use the model to decompose aggregate productivity growth into the contribution of process innovation and creative destruction. In particular, the model implies that, in the long run, aggregate growth is given by

$$g = \tau \cdot \eta + z \cdot \lambda \cdot \mathbb{E}[\gamma^s]. \quad (15)$$

In Equation (15), we may account for the direct contribution of creative destruction to productivity growth by dividing the term involving creative destruction, the product of the creative destruction rate and the product innovation step size, $\tau\eta$, by the overall rate of productivity growth, g. Although this ratio abstracts from potential linkages between creative destruction and the rate and effectiveness of process innovation, it offers a starting point for quantifying the contribution of creative destruction to growth, if we counterfactually keep the contributions of process innovation constant. This exercise is similar to that performed by Garcia-Macia et al. (2019) for quantifying the contribution of creative destruction.

The top row in Table 7 presents the results of this approach to quantifying the direct contribution to growth of creative destruction in our model. The first column presents the contribution in our calibrated model, which is 15.3%. This is much smaller than that in models without process innovation (e.g. Klette and Kortum 2004), where the contribution of creative destruction to growth is 100%. In a similarly calibrated model in which the process innovation step size remains constant in the product's tenure (e.g., Akcigit and Kerr 2018, Garcia-Macia et al. 2019), the direct contribution of creative destruction would also be 15.3%. For comparison, in a model disciplined by the U.S. establishment data on employment flows rather than product churn, Garcia-Macia et al. (2019) find that around 25% of growth originates from creative destruction. Thus, using the richer data on firm-level product dynamics, we confirm that the simple decomposition above leads to the conclusion that the majority of growth in the long run stems from process innovation, once we allow for it to play a role in our models.

6.2. Accounting for Endogeneity in Process Innovation

Our empirical exercise and quantitative framework show that process and product innovation are intertwined: innovation bursts add a wide set of products to the firms' product portfolios, enabling new waves of reinvigorated process innovation that further increase growth. At the same time, creative destruction raises the firms' discount rates, which negatively affects their incentive to engage in process innovation.

The clearest way to see the interaction between creative destruction and process innovation is to consider the impacts on the aggregate growth in the extreme scenario in which creative destruction ceases altogether, such that τ is zero. In models with constant product innovation, the fall in growth is given by the direct contribution of creative destruction, which is 15.3% under our calibration. When process innovation declines with product tenure, however, growth eventually ceases entirely when $\tau = 0$ (Table 7, row 2). This is because process innovation becomes gradually ineffective over time as products become older without the introduction of new products. In our model, the step size of process innovation would gradually decline as prior instances of process innovation accumulate over time. Thus, the direct effect of creative destruction on growth underestimates its true relevance.

Table 7: Effect of Creative Destruction on Productivity Growth

	Full Model	No Process Innovation (e.g. Klette and Kortum)	Constant Step Process Inn. (e.g. Akcigit and Kerr)
Simple decomposition			
Percent of growth due to creative destruction	15.3%	100%	15.3%
Growth in absence of creative destruction			
Growth decline when $\tau = 0$ in % of baseline	100%	100%	15.3%
Effect of creative destruction on growth			
Direct effect	3.1pp	3.1pp	3.1pp
Incl. endo. effect on of process innovation step	13.2pp	3.1pp	3.1pp
Incl. endo. effect on of process innovation rate	4.9pp	3.1pp	-5.2pp
Effect of process innovation on growth			
Direct effect	1.3pp	N.A.	1.3pp

Notes: Contributions of growth are obtained from the quantified model in Section 5. For all models, the step size of product innovations η is 3.1%. For process innovation in the full model, the process innovation step size λ is 4.1% while γ is 0.88, giving rise to an average process innovation step of 1.3%. For the model with constant process innovation steps, we set λ to 1.3% and γ to 1.

An alternative way to express the interaction between process and product innovation by examining the effect of a marginal increase in creative destruction on growth, taking into account both its direct effects as well as the indirect effects mediated by the impact on process innovation. The latter can, in turn, be separated into two parts. First, an increase in creative destruction raises the expected step size of process innovation and thus its effect on growth. Second, a higher pace of creative destruction raises the rate at which the benefits to process innovation are discounted, as firms expect to lose their product sooner when they are replaced by another firm through creative destruction. These three components are terms in the the following total derivative of growth with respect to τ:

$$\frac{dg}{d\tau} = \underbrace{\eta}_{\text{direct effect}} + \underbrace{\lambda z \left(\frac{\partial \mathbb{E}[\gamma^s]}{\partial \tau} \right)}_{\text{step-size effect} >0} + \underbrace{\lambda \mathbb{E}[\gamma^s] \left(\frac{\partial z}{\partial \tau} \right)}_{\text{discount effect} <0} . \tag{16}$$

The third panel of Table 7 presents the decomposition of the effects of creative destruction on aggregate long-run growth implied by Equation (16). The direct effect is simply given by the step-size of quality improvements from creative destruction, η, which we estimate to be 3.1 percentage points. The steps-size effect, the second term in the decomposition, equals 10.1 percentage points. That means that the indirect effect of innovation bursts on growth, by igniting a sequence of highly productive process innovation, is more than three times larger than their direct effect. Finally, we observe that the negative effect of creative destruction on the rate of process innovation can also be sizable. From the first-order condition of process innovation, this response is

$$\frac{\partial z}{\partial \tau} = -\frac{1}{\psi - 1} \left(\frac{z}{\rho + \tau} \right).$$

Hence, to quantify the discount effect, we need to calibrate ψ and ρ. We set the former to 2.6 and the latter to 0.02, which are standard values in the literature. The negative effect is sizable: the decline

Figure 10. Interaction between Creative Destruction, Process Innovation, and Growth

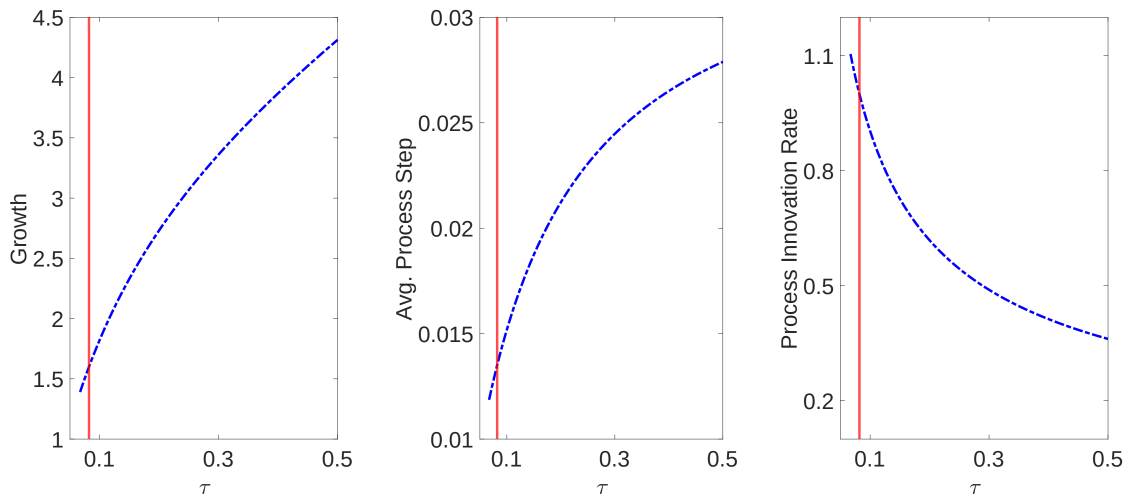

(a) Growth (b) Process Step Size (c) Process Rate

Notes: The figure shows the changes in (a) the aggregate rate of productivity growth, (b) the average step size of process innovation, and (c) the rate of process innovation for different rates of creative destruction in the calibrated model.

in process innovation reduces growth by 8.3 percentage points. The total derivative of growth with respect to creative destruction, from the innovation bursts of either incumbents or entrants, is 4.9 percentage points. In Figure 10, we illustrate the relationship between creative destruction, process innovation's step size, and the incentive to engage in process innovation for a wide range of changes in the rate of creative destruction, τ, keeping fixed all other model parameters.

The exercise above shows that an assessment of the effect of creative destruction on growth may be highly sensitive to the details of the model considered. In our model, which is compatible with the documented facts on the dynamics of product-level revenue growth over the life cycle, creative destruction ignites a series of highly productive process innovations. As such, the indirect effect of a rise in creative destruction, mediated through the impact on process innovation, may be far larger than its direct impact on aggregate productivity growth.

6.3. Creative Destruction and Industry Concentration

Finally, in this section we examine the interaction between creative destruction and industry concentration. As explained in Section 5, innovation bursts are a natural explanation for why production concentrates in a small set of large firms. We next show that creative destruction has large effects on the degree of industry concentration, which means that introducing innovation bursts into the model is key when studying the joint determinants of productivity growth and industry concentration, which is the topic of a growing body of research (e.g. Olmstead-Rumsey 2022, Akcigit and Ates 2021, Akcigit and Ates 2023, Aghion et al. 2023, De Ridder 2024).

An increase in the rate at which firms obtain innovation bursts, x, raises the rate of creative destruction, τ, and reshapes the distribution of production across firms. Figure 11 shows this distribution

Figure 11. Firm-Size Distribution: Interaction between Creative Destruction and Concentration

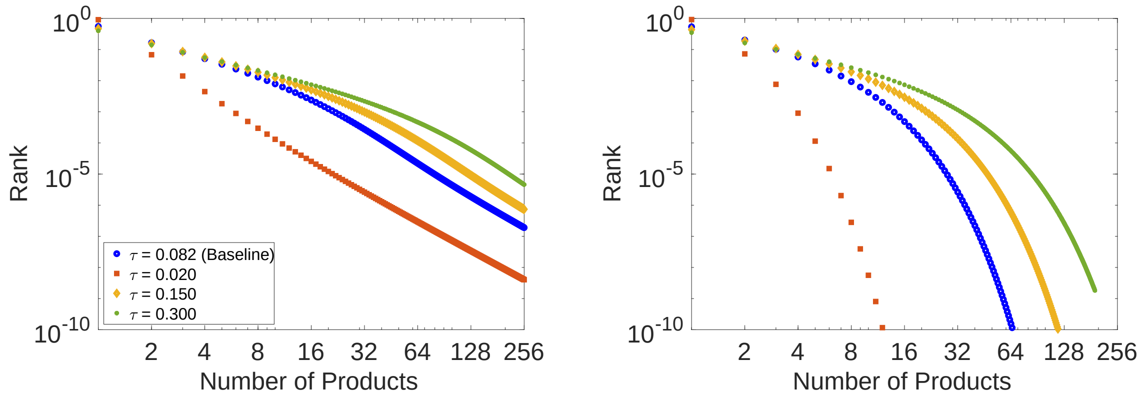

(a) Model with Innovation Bursts

(b) Model without Innovation Bursts

Notes: The figures plot the relationship between a firm's size (horizontal axes) and the firm's inverse rank (vertical axes).

for the number of firm products, as a log-rank-vs.-log-size relationship, as we change the rate of creative destruction τ in the model while keeping fixed all other parameters. As we have already seen, this relationship is close to linear in the data–a fact that the model with innovation bursts (the left-hand figure) is able to replicate while that without innovation bursts (the right-hand figure) cannot. First, the figure shows that this result holds across many different rates of creative destruction. Moreover, we observe that higher rates of creative destruction are associated with the emergence of larger firms in both models, but particularly so in the presence of innovation bursts.

Figure 12 presents the resulting variations in two different proxies of concentration of the distributions of production as we change the rate of creative destruction, τ, in the two models. Under both models, this relationship is increasing. However, the response of concentration to a rise in the rate of creative destruction is more pronounced in the model featuring innovation bursts. As we can see in Figure 12(a), the revenue share of the top 10% of firms increases in the rate of creative destruction, and this rise is stronger under the full model featuring innovation bursts. At the calibrated value of τ, the impact of a marginal rise in the rate of creative destruction on concentration are 2.36 and 1.83 in our model and the conventional model, respectively. As an alternative proxy, Figure 12(b) shows the average number of products of firms, which similarly rises in the rate of creative destruction, and does so much more strongly in the presence of innovation bursts. Here, the impact of a marginal rise in the rate of creative destruction on average firm size are 24.3 and 15.3 in our model and the conventional model, respectively.

We conclude that a model that is compatible with the observed dynamics of firm-level production in the data, as captured by the presence of innovation bursts, implies a tight link between creative destruction and the concentration of production. In this model, stimulating product innovation raises the likelihood of the emergence of large, superstar firms much more strongly than in a model with a conventional account of product innovation.

Figure 12. Interaction between Creative Destruction and Concentration

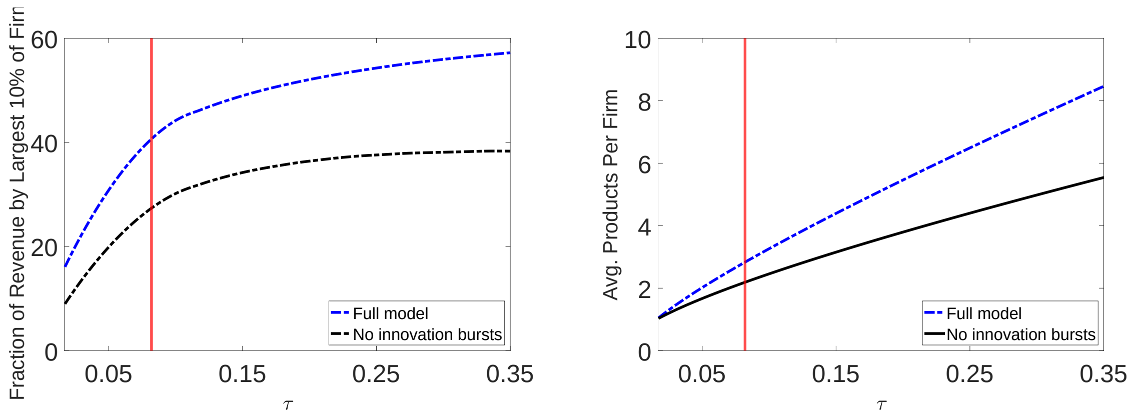

(a) Top 10% Market Share

(b) Products per Firm

Notes: The vertical axis of the left-hand figure plots the fraction of all products that is produced by the largest 10% of firms. The vertical axis of the right-hand figure plots the average number of products that firms produce.The horizontal axis measures the rate of creative destruction, τ.

7. Conclusion

This paper contributes to the ongoing research on the dynamics of firm-level innovation activity and its impact on aggregate-level phenomena such as productivity growth, business dynamism, and the firm-size distribution. By using comprehensive data from the French manufacturing sector, we have documented a novel fact on firm-level product innovation activity, revealing that the distribution of new products introduced by firms in each year has a thick, Pareto-like tail. This finding is inconsistent with the predictions of standard theories of creative destruction, which have difficulty fitting this empirical pattern.

To address this gap, we propose a new way of modeling endogenous growth through product innovation, which incorporates the concept of innovation bursts. Our model suggests that innovations can come in bursts, with the potential for immediate applications that enable firms to expand their portfolio of products rapidly. This notion of innovation bursts is consistent with the novel fact we have documented on product innovation among French manufacturing firms.

Our proposed model offers a new explanation for the Pareto distribution of firm size commonly observed in the data, which is a well-known limitation of theories of creative destruction. By showing that the equilibrium firm-size distribution can be Pareto for a sufficiently thick-tailed distribution of innovation bursts, we propose the lumpy nature of innovative ideas as the main driver of the concentration of firm-level production.

Overall, this paper contributes to the literature on firm-level innovation activity, highlighting the importance of incorporating the concept of innovation bursts into models of endogenous growth through product innovation. By shedding new light on the interplay between concentration, innovation, and growth, this paper offers a promising avenue for future research in this area.

References

Acemoglu, D., Akcigit, U., Alp, H., Bloom, N., and Kerr, W. (2018a). Innovation, reallocation, and growth. *American Economic Review*, 108(11):3450–91.

Acemoglu, D., Akcigit, U., Alp, H., Bloom, N., and Kerr, W. (2018b). Innovation, reallocation, and growth. *American Economic Review*, 108(11):3450–3491.

Acemoglu, D., Gancia, G., and Zilibotti, F. (2012). Competing engines of growth: Innovation and standardization. *Journal of Economic Theory*, 147(2):570–601.

Aghion, P., Akcigit, U., and Howitt, P. (2014). What do we learn from schumpeterian growth theory? In *Handbook of economic growth*, volume 2, pages 515–563. Elsevier.

Aghion, P., Bergeaud, A., Boppart, T., Klenow, P. J., and Li, H. (2023). A theory of falling growth and rising rents. *Review of Economic Studies*.

Aghion, P. and Howitt, P. (1992). A model of growth through creative destruction. *Econometrica*, 60(2):323–351.

Akcigit, U., Alp, H., and Peters, M. (2021). Lack of selection and limits to delegation: firm dynamics in developing countries. *American Economic Review*, 111(1):231–275.

Akcigit, U. and Ates, S. T. (2021). Ten facts on declining business dynamism and lessons from endogenous growth theory. *American Economic Journal: Macroeconomics*, 13(1):257–98.

Akcigit, U. and Ates, S. T. (2023). What happened to us business dynamism? *Journal of Political Economy*, 131(8):2059–2124.

Akcigit, U. and Kerr, W. R. (2018). Growth through heterogeneous innovations. *Journal of Political Economy*, 126(4):1374–1443.

Argente, D., Baslandze, S., Hanley, D., and Moreira, S. (2020). Patents to products: Product innovation and firm dynamics.

Argente, D., Lee, M., and Moreira, S. (2024). The life cycle of products: Evidence and implications. *Journal of Political Economy*, 132(2):000–000.

Atkeson, A. and Burstein, A. (2019). Aggregate implications of innovation policy. *Journal of Political Economy*, 127(6):2625–2683.

Autor, D., Dorn, D., Katz, L. F., Patterson, C., and Van Reenen, J. (2020). The fall of the labor share and the rise of superstar firms. *The Quarterly Journal of Economics*, 135(2):645–709.

Axtell, R. L. (2001). Zipf distribution of us firm sizes. *science*, 293(5536):1818–1820.

Bloom, N., Jones, C. I., Van Reenen, J., and Webb, M. (2020). Are ideas getting harder to find? *American Economic Review*, 110(4):1104–44.

Cavenaile, L., Celik, M. A., and Tian, X. (2020). Are markups too high? competition, strategic innovation, and industry dynamics. *Working Paper*.

Davis, S. J., Haltiwanger, J., Jarmin, R., Miranda, J., Foote, C., and Nagypal, E. (2006). Volatility and dispersion in business growth rates. *NBER macroeconomics annual*, 21:107–179.

De Ridder, M. (2024). Market power and innovation in the intangible economy. *American Economic Review*, 114(1):199–251.

Gabaix, X. (2009). Power laws in economics and finance. *Annu. Rev. Econ.*, 1(1):255–294.

Gabaix, X., Lasry, J.-M., Lions, P.-L., and Moll, B. (2016). The dynamics of inequality. *Econometrica*, 84(6):2071–2111.

Garcia-Macia, D., Hsieh, C.-T., and Klenow, P. J. (2019). How destructive is innovation? *Econometrica*, 87(5):1507–1541.

Garicano, L., Lelarge, C., and Van Reenen, J. (2016). Firm size distortions and the productivity distribution: Evidence from france. *American Economic Review*, 106(11):3439–3479.

Geerolf, F. (2017). A theory of pareto distributions. *UCLA manuscript*.

Gibrat, R. (1931). Les inégalits économiques. *Sirey*.

Grossman, G. M. and Helpman, E. (1991). Quality ladders in the theory of growth. *The review of economic studies*, 58(1):43–61.

Guvenen, F., Karahan, F., Ozkan, S., and Song, J. (2021). What do data on millions of us workers reveal about lifecycle earnings dynamics? *Econometrica*, 89(5):2303–2339.

Hopenhayn, H. A. (2014). Firms, misallocation, and aggregate productivity: A review. *Annu. Rev. Econ.*, 6(1):735–770.

Klette, T. J. and Kortum, S. (2004). Innovating firms and aggregate innovation. *Journal of political economy*, 112(5):986–1018.

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

Luttmer, E. G. (2007). Selection, growth, and the size distribution of firms. *The Quarterly Journal of Economics*, 122(3):1103–1144.

Luttmer, E. G. (2010). Models of growth and firm heterogeneity. *Annu. Rev. Econ.*, 2(1):547–576.

Luttmer, E. G. (2011). On the mechanics of firm growth. *The Review of Economic Studies*, 78(3):1042–1068.

Olmstead-Rumsey, J. (2022). Market concentration and the productivity slowdown.

Ozkan, S., Hubmer, J., Salgado, S., and Halvorsen, E. (2023). Why are the wealthiest so wealthy? a longitudinal empirical investigation.

Peters, M. (2020). Heterogeneous markups, growth, and endogenous misallocation. *Econometrica*, 88(5):2037–2073.

'Growth Through Innovation Bursts'
Appendix - For Online Publication Only

Appendix A. Theoretical Appendix

The flow equation for the mass M_n of firms with different numbers of product n is given by

$$\dot{M}_n = \left(\sum_{j=1}^{n-1} M_j \cdot x_j \cdot \frac{[n-j]^{-\theta_j}}{\zeta(\theta_j)}\right) + \left(\sum_{k=n+1}^{\infty} M_k \cdot \tau_k \cdot \frac{[k-n]^{-s_k}}{\zeta(s_k)}\right) - M_n (x_n + \tau_n).$$

In the steady state, the mass of firms in a size bin is constant. This implies the following expresion characterizing the steady state distribution:

$$M_n = \frac{1}{x_n + \tau_n}\left[\left(\sum_{j=1}^{n-1} M_j \cdot x_j \cdot \frac{[n-j]^{-\theta_j}}{\zeta(\theta_j)}\right) + \left(\sum_{k=n+1}^{\infty} M_k \cdot \tau_k \cdot \frac{[k-n]^{-s_k}}{\zeta(_k)}\right)\right].$$

Equilibrium Definition The following provides the definition of the equilibrium in our model.

Definition 1. *The economy is in a balanced growth path equilibrium if for every t the variables* $\{r, e, L^p, g\}$ *and functions* $\{x_{n_i}, M_{n_i}, \tau\}$ *are constant,* $\{Y, C, Qw,\}$ *grow at a constant rate g that satisfies (14), interest rates follow from (9), Q is given by (13), Y is* $Y = QL^p$*, innovation rates* x_n *satisfy (7), the entry rate e satisfies (8), firm distribution* M_n *is constant and satisfies (12), the rate of creative destruction* τ *satisfies (4), and both goods and labor markets are in equilibrium so that* $Y = C$ *and* $L^p = 1 - L^s - L^{rd} - L^e$.

Appendix B. Data Appendix

In this section, we describe the data sources used in our analysis and the procedure we use to merge them and clean the resulting dataset of outliers.

B.1. EAP

The data are based on an annual survey of firms' production activities, called the *Enquête Annuelle de Production* (EAP), which is administered by the Institut National de la Statistique et des Études Économiques (INSEE). In accordance with the EU regulation, the survey must encompass at least 90 per cent of the annual production of each 4-digit industry. The data contain comprehensive information on sales and the volume of goods. The volume is recorded in units of measurement (number of items, kilograms, litres) that are product-specific, while the value is recorded in current euros. As for the level of product aggregation, the PRODFRA classification contains roughly 4300 10-digit product codes (see below). The survey covers the entire manufacturing sector (NACE rev. 2 section C), except for the agri-food industries (section 10, 11 and 12) and the manufacture of wood (16). While also including the extractive industry; electricity, gas, steam and air-conditioning supply; water supply, sewerage, waste management and remediation; we exclude these specific industries.

Table A1: Examples of products in the PRODFRA classification

1812125000	Advertising and similar printed matter (excluding commercial catalogues)
1812199010	Administrative or commercial printed matter, at or continuous, customised or not, and directories
2511235040	Industrial boiler products: not including tanks, boilers, nuclear equipment
3102100010	Wooden kitchen furniture: by mounted elements, including custom
310912502B	Dining and living room furniture other than tables: buffets, credenzas and livings, bookcases, cabinets by element.

B.1.1. Sampling Framework

The survey contains an exhaustive sample of firms with at least 20 employees or revenues higher than 5 million Euros. The sample size varies over time but it is usually around 25'000 units. To ensure a good level of coverage the survey must cover at least 90 per cent of the total production value of each 4-digit industry (NACE rev. 2). If this threshold is not reached more enterprises are surveyed. Additionally, the survey contains a random sample of firms with less than 20 employees. Its size varies year by year and it is usually around 8/9000 units. Because of sample attrition, we drop this second set of smaller firms from our analysis.

B.1.2. Product Classification

As for the level of product aggregation, the PRODFRA classification contains roughly 4300 10-digit product codes. Table A1 presents examples of products in the PRODFRA classification. The first eight positions of PRODFRA represent the PRODCOM classification, where the last two positions are used to refine the nomenclature. A few special products, which account for around 5 per cent of the total number of product categories, are identified by a letter (H, N, S, Y) instead of a number in the 9th position of PRODFRA. These product categories are dropped when working with the PRODCOM classification. The PRODCOM classification can be directly linked to the classification of industrial activities in the EU (NACE Rev. 2), as the first four digits of the PC code identify a 4-digit NACE industry.

B.1.3. Product Concordance

A common feature of PRODFRA is to change over time (from 3 to 5% of product categories change each year). The use of these product codes in longitudinal studies requires harmonizing the product classification system over time. To do so, we use the algorithm developed by **?**, called "connected components concordance", or *C3* for short. *C3* uses the graph theory to identify stable and comparable groups of products over time while minimizing the size of each group. The identified groups of products are then assigned to a single, time-consistent, code. The vast majority of products (almost 90%) are not affected by this concordance procedure, and a marginal fraction of the new product groups include more than three PRODFRA10 codes.

B.2. FARE

The *Fichier approché des résultats d'Esane* (FARE) contains a coherent set of statistics on the universe of French private companies. It combines administrative data obtained from annual profit declarations made by companies to the tax authorities and from annual social data which provide information on employees and data obtained from a sample of companies surveyed by a specific questionnaire to produce structural business statistics (ESA).

Appendix C. Additional Figures

Figure A1. Size Distribution - Revenue and Employment

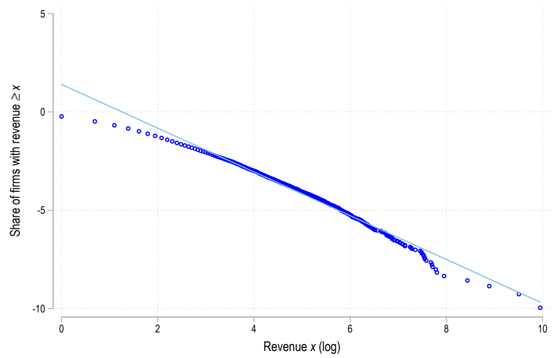

(a) Distribution: Revenue

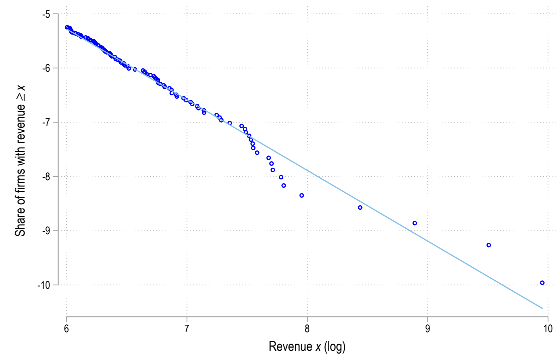

(b) Distribution: Revenue - Tail

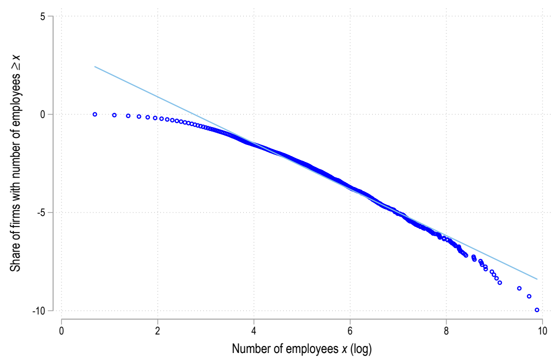

(c) Distribution: Employment

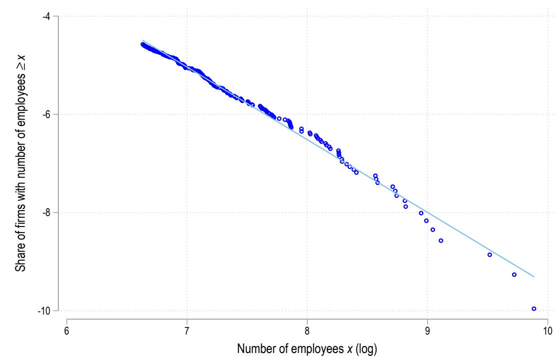

(d) Distribution: Employment - Tail

Notes: The figures plot the relationship between a firm's size (horizontal axes) and the firm's inverse rank (vertical axes). Size is measured through either revenue (in 2015 euros) or employment (in full-time equivalence). Inverse rank is measured as the ratio of firms' rank starting from the smallest firm, divided by the total number of observations in the data. Plots are based on the FARE-EAP sample for 2019. Slopes by figure: (a) Slope: -1.1. Standard error: .01. R^2: 0.98; (b) Slope: -1.30. Standard error: .02. R^2: 0.99; (c) Slope: -1.2. Standard error: .01. R^2: 0.97; (d) Slope: -1.5. Standard error: .01. R^2: 0.99.

C.1. Number of products

Figure A2. Robustness for Figure 1

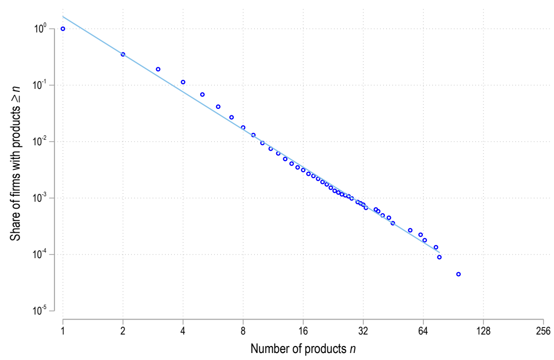

(a) Distribution: Products 2010

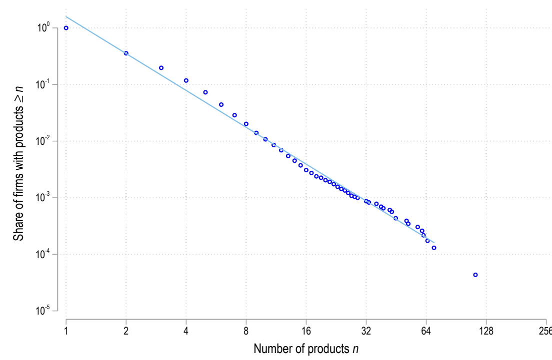

(b) Distribution: Products 2012

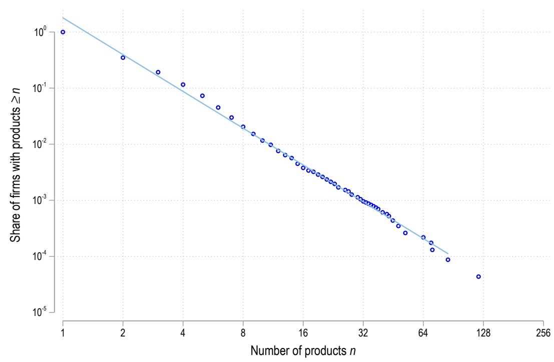

(c) Distribution: Products 2014

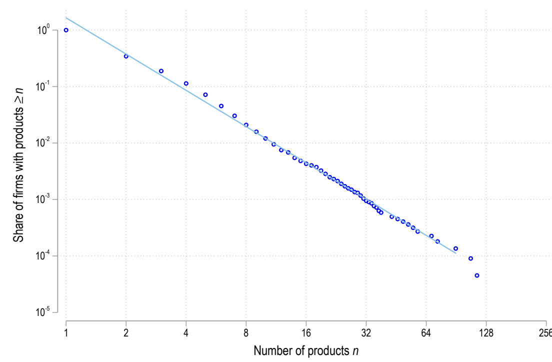

(d) Distribution: Products 2016

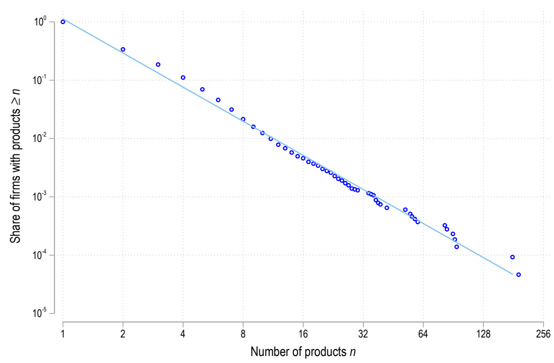

(e) Distribution: Products 2018

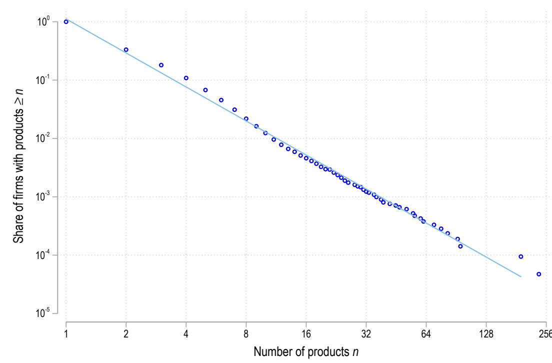

(f) Distribution: Products 2019

Notes: The figures plot the relationship between a firm's size (horizontal axes) and the firm's inverse rank (vertical axes). Left-hand figures plot the entire distribution, while right-hand figures plot the tail of the distributions by narrowing in on the top 25% of firms. All figures use 2019 data. Size is measured through either revenue (in 2015 euros) or employment (in full-time equivalence). Inverse rank is measured as the ratio of firms' rank starting from the smallest firm, divided by the total number of observations in the data. Plots are based on the FARE-EAP sample for 2019. ADD OLS SLOPES.

Figure A3. Robustness for Figure 1

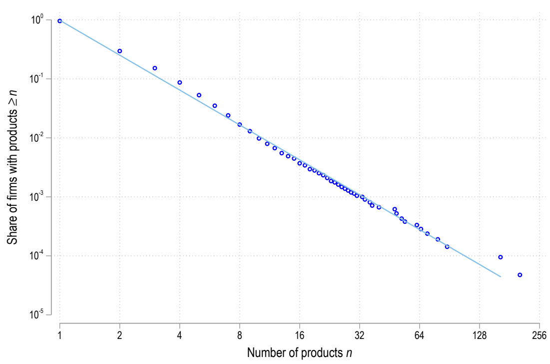

(a) Distribution: Products 8 digit

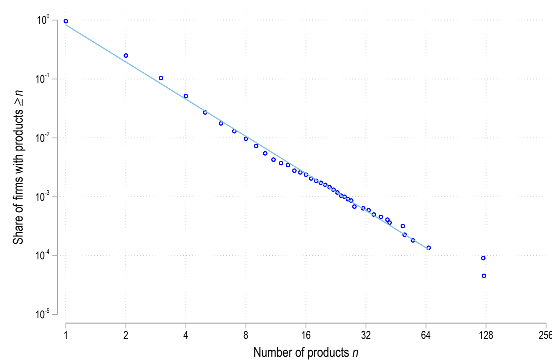

(b) Distribution: Products 6 digit

Notes: The figures plot the relationship between a firm's size (horizontal axes) and the firm's inverse rank (vertical axes). Size is measured through either revenue (in 2015 euros) or employment (in full-time equivalence). Inverse rank is measured as the ratio of firms' rank starting from the smallest firm, divided by the total number of observations in the data. Plots are based on the FARE-EAP sample for 2019. ADD OLS SLOPES.

C.2. Product Innovation Bursts Robustness

Figure A4. Revenue Per Product and Number of New Products

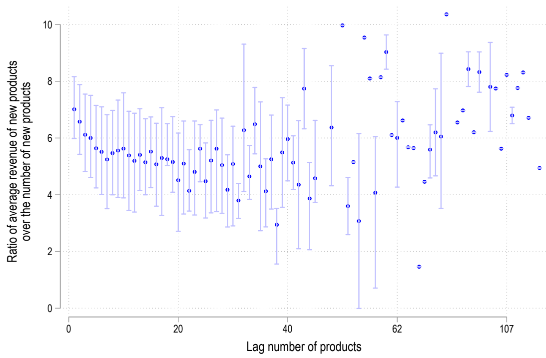

(a) Average Revenue of New products (All Sizes)

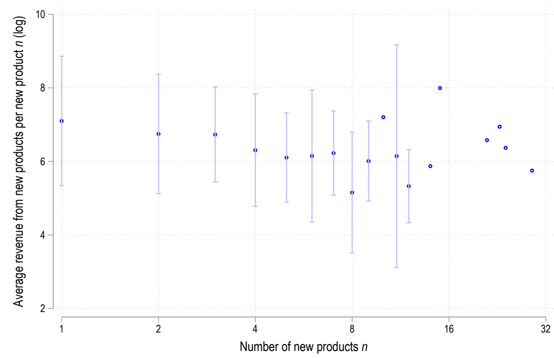

(b) Average Revenue of New products ($N_{t-1} = 1$)

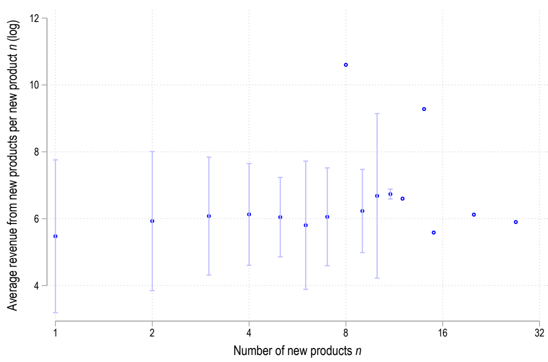

(c) Average Revenue of New products ($N_{t-1} = 5$)

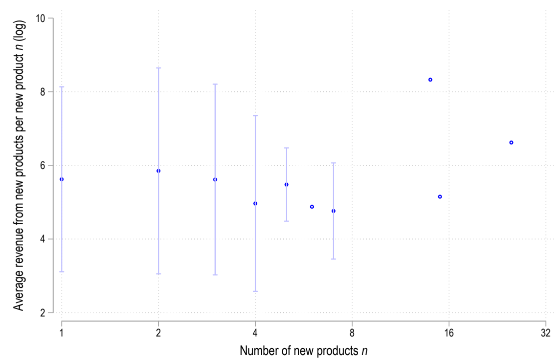

(d) Average Revenue of New products ($N_{t-1} = 10$)

Notes: Figures plot the ratio of total revenue earned on new products divided by the total number of new products on the vertical axis, against the number of new products on the horizontal axis. Dots present the average ratio in the sample, while confidence bounds are plus/minus one standard deviation.

Figure A5. Distribution of Number of New Products by Level of Aggregation

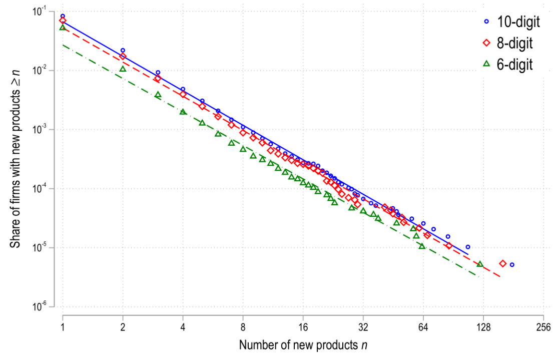

(a) Distribution of New Products (All Sizes)

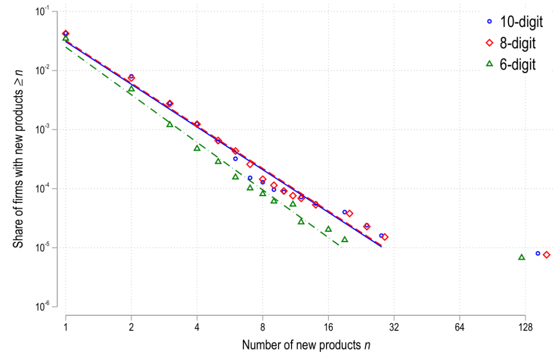

(b) Distribution of New Products ($N_{t-1} = 1$)

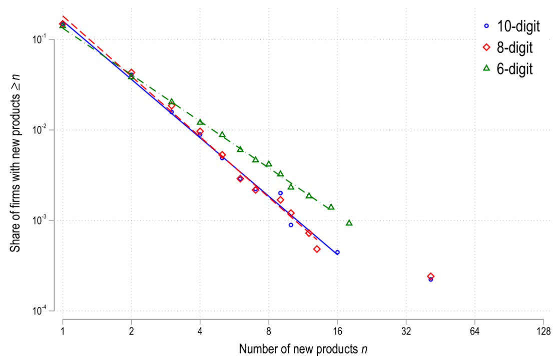

(c) Distribution of New Products ($N_{t-1} = 5$)

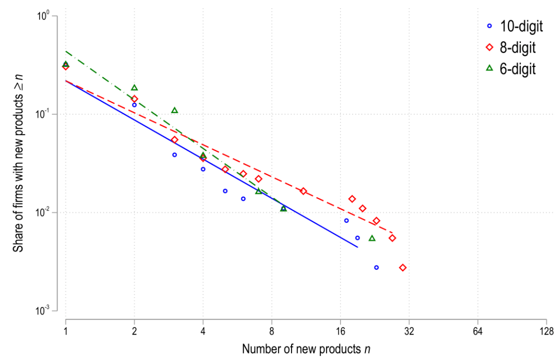

(d) Distribution of New Products ($N_{t-1} = 10$)

Notes: The figures plot the relationship between a firm's size (horizontal axes) and the firm's inverse rank (vertical axes). The figure provides overlapping plots of the log rank against the log number of products added at the 10 digit level (blue circles), 8-digit (red triangles) and 6 digit level (green diamonds). The idea is that if innovation bursts are driven by firms differing in the care with which they report revenue separately for each 10-digit product, the distribution should become less thick tailed at higher levels of aggregation. Size is measured through either revenue (in 2015 euros) or employment (in full-time equivalence). Inverse rank is measured as the ratio of firms' rank starting from the smallest firm, divided by the total number of observations in the data. Plots are based on the FARE-EAP sample for 2019.

Figure A6. Distribution of Number of New Products - No Mergers

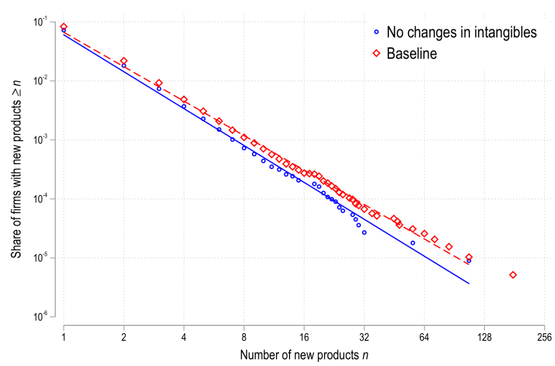

Notes: The figures plot the relationship between a firm's number of new products and its rank. Inverse rank is measured as the ratio of firms' rank starting from the firm with the fewest new products, divided by the total number of observations in the data. Plots are based on the FARE-EAP sample for 2019.

Figure A7. Product Innovation by Age

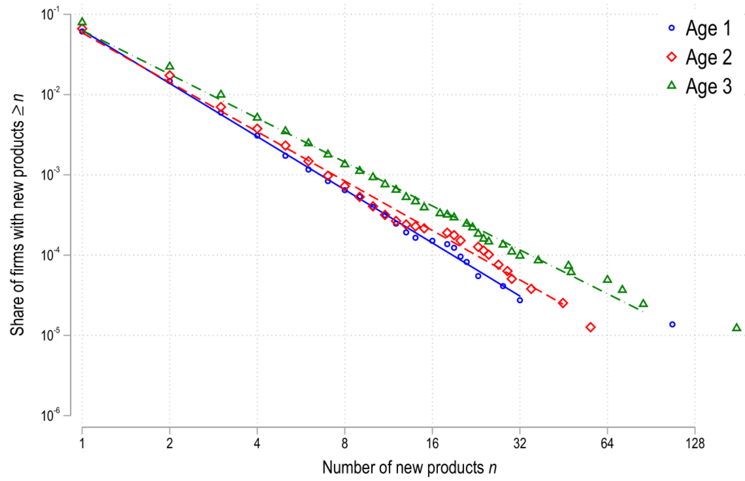

Notes: The figures plot the relationship between a firm's number of new products and its rank. Inverse rank is measured as the ratio of firms' rank starting from the firm with the fewest new products, divided by the total number of observations in the data. Plots are based on the FARE-EAP sample for 2019. Firms are split into three equally-sized age bins, such that scatters for Age 1 belong to the youngest third of firms.

C.3. Product-Level Revenue Growth by Tenure for Type of Firm

Figure A8. Life Cycle of Revenue Growth by Product Count

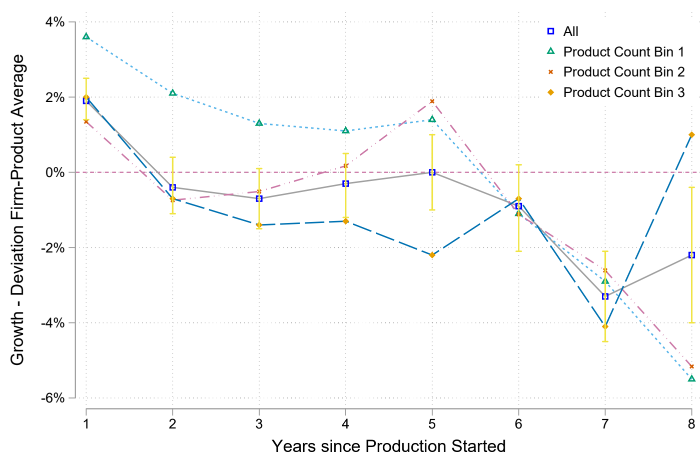

Notes: The vertical axis plots average growth of product revenue among all firm-products in the sample, where the average growth rate of a firm-product's revenue is subtracted from growth at each horizon. The horizontal axis plots the number of years that have passed since the firm first started producing the product. Confidence bounds are at the 90% level and based on bootstrapped standard errors.

Figure A9. Life Cycle of Revenue Growth by Age

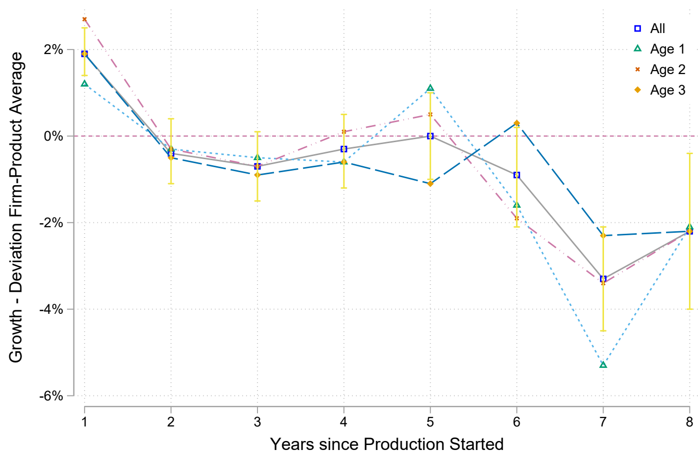

Notes: The vertical axis plots average growth of product revenue among all firm-products in the sample, where the average growth rate of a firm-product's revenue is subtracted from growth at each horizon. The horizontal axis plots the number of years that have passed since the firm first started producing the product. Confidence bounds are at the 90% level and based on bootstrapped standard errors.

Figure A10. Life Cycle of Revenue Growth by Sector

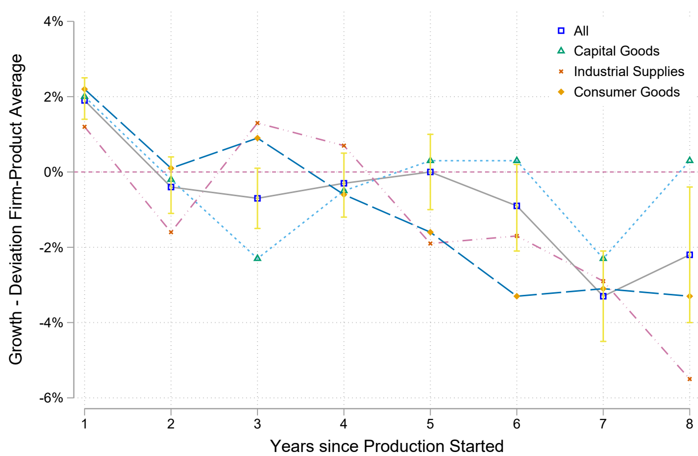

Notes: The vertical axis plots average growth of product revenue among all firm-products in the sample, where the average growth rate of a firm-product's revenue is subtracted from growth at each horizon. The horizontal axis plots the number of years that have passed since the firm first started producing the product. Confidence bounds are at the 90% level and based on bootstrapped standard errors.

Table A2: Concentration of Firm Growth

	1 Year				5 Year			
	-3 σ	-2 σ	+2 σ	+3 σ	-3 σ	-2 σ	+2 σ	+3 σ
Product Count								
Data	0.98	1.08	1.09	1.00	0.90	0.93	0.99	0.96
Normal Dist.	0.13	2.28	2.28	0.13	0.13	2.28	2.28	0.13
Ratio	7.3	0.5	0.5	7.4	6.7	0.4	0.4	7.1
Revenue								
Data	0.47	0.97	0.84	0.42	0.13	0.34	0.55	0.34
Normal Dist.	0.13	2.28	2.28	0.13	0.13	2.28	2.28	0.13
Ratio	3.5	0.4	0.4	3.1	1.0	0.2	0.2	2.5
Employment								
Data	0.77	1.51	0.95	0.38	0.27	0.75	0.65	0.19
Normal Dist.	0.13	2.28	2.28	0.13	0.13	2.28	2.28	0.13
Ratio	5.7	0.7	0.4	2.8	2.0	0.3	0.3	1.4

Notes: The table presents the percentage of observations in the tail of the data and the percentage of observations in the tail under the normal distribution. Colomn headers indicate tail in terms of standard deviations from the mean.

C.4. Distribution of firm growth

We now turn to the first new fact of the paper: the distribution of firm *growth* is fat tailed. As with size, we measure firm growth through changes in the number of 10-digit products that it sells, changes in revenue and changes in employment. To measure firm growth, we measure either year-on-year or five-year changes in the natural logarithm of each of these variables: $\ln y_{it} - \ln y_{it-h}$, where $h = 1,5$. By focusing on log changes we are able to study the entire distribution of firm growth. The main alternative measure of firm growth is the symmetric growth proposed by Davis et al. (2006), which is bounded by $[-2,2]$ regardless of the underlying distribution of firm size. This makes it impractical for studying the tail of the firm-growth distribution. We return to symmetric growth rates in Section **??**.

To show that the distribution of growth is fat tailed, we compare the distribution of log-change in firm size to the log-normal distribution. The log-normal distribution is a natural starting point as, motivated by Gibrat (1931), firm dynamics models with random shocks to productivity typically assume that such shocks (and thus firm growth) are log normal (see, e.g., Hopenhayn 2014, Danial: add references).[17] In models of firm dynamics and creative destruction in the spirit of Klette and Kortum (2004) firm growth arises as the result of a Poisson process, which yields a distribution of firm growth that is even more thin tailed (see, e.g., Luttmer 2010).

Quantile-Quantile plots We plot the distribution of firm growth in Figure A11. Each sub-figure contains a Quantile-Quantile (QQ) plot comparing the distribution of a particular measure of firm growth to a log-normal distribution with the same mean and standard deviation. To produce these figures, we

[17] Our results echo recent findings in the literature on income dynamics which shows that individual income growth is also better described by a fat-tailed distributions such as the Pareto distribution (see, e.g., Guvenen et al. (2021). Gabaix et al. (2016) discuss the importance of the income growth process to explain stationary distributions of income and wealth.

Figure A11. Distribution of Growth Rates

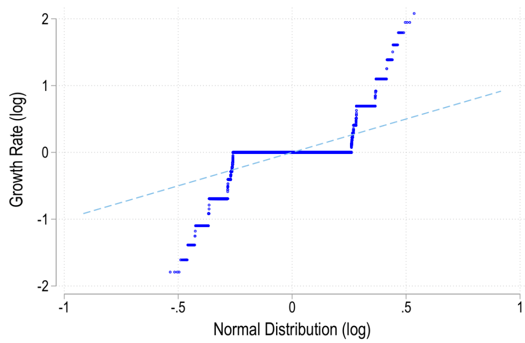

(a) Product Count Growth

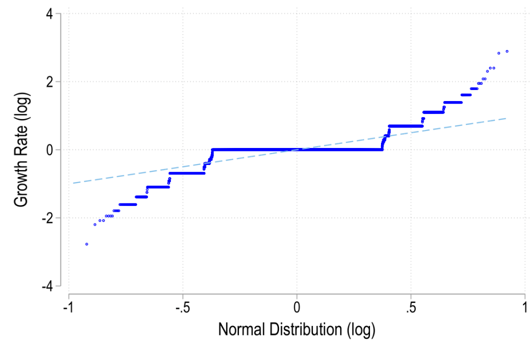

(b) Product Count Growth (5 year)

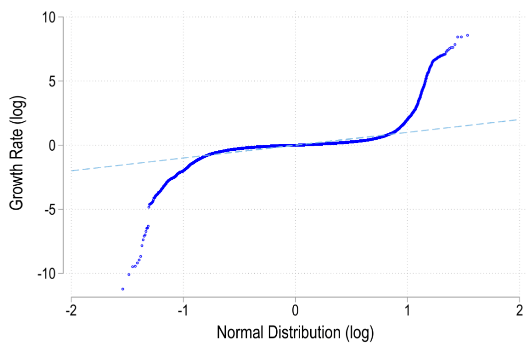

(c) Revenue Growth

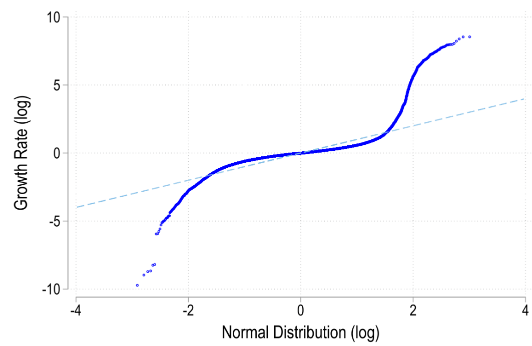

(d) Revenue Growth (5 year)

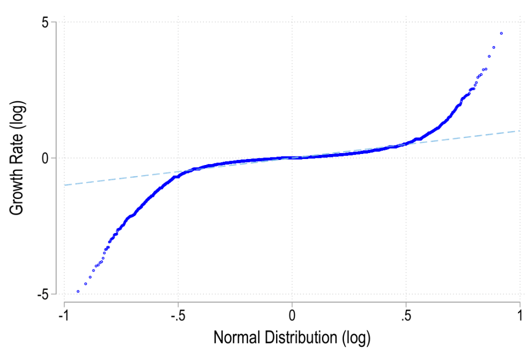

(e) Employment Growth

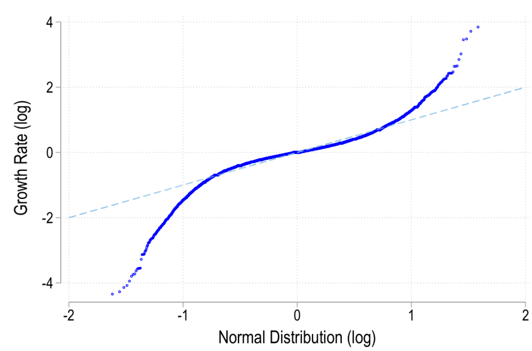

(f) Employment Growth (5 year)

Notes: The figures plot Quantile-Quantile plots comparing the empirical distribution of log change in product count, revenue and employment against the log-normal distribution. Dashed light-blue lines present the reference log-normal distribution.

sort firm-year observations by growth rate and calculate their rank. The plot than compares a particular quantile's growth rate on the vertical axis to what that s quantile's growth rate would be under the log-normal distribution. If growth rates are log-normally distributed, the scatters should be positioned on the diagonal reference lines. If the growth rates are fat tailed, scatters of the left tail are below the reference line, while scatters on the right tail are above the reference line.

Figure A11 shows that firm growth is fat tailed in the data. The upper, middle, and lower figures respectively plot the distribution product count, revenue, and employment. For each measure, both the 1-year growth rate (left-hand) and the 5-year growth rate show that large positive and negative growth rates occur more often than a log-normal distribution predicts. The deviation is particularly clear for the largest changes in size. For employment, for example, the deviation becomes graphically clear around ± 0.5 log change, which translate to -45% and +65%.

Tail mass Table A2 quantifies how much the tails of firm growth deviates from the log-normal distribution. It compares the percentage of observations to the left (right) of minus (plus) 2 and 3 standard deviations from the mean for each variable and horizon. For product count, the data shows a close alignment with the normal distribution at the +2 standard deviations and +3 standard deviations levels in the 1-year period, with ratios of 0.5 and 7.4 respectively, indicating that extreme changes in a firm's product count are relatively common. Revenue data again displays a marked deviation from the normal distribution, with changes in annual revenue growth outside the ±3 standard deviations occurring more than three times as often. Results for employment growth are similar.

Log-rank versus log-size Having established that extreme changes in firm size are much more common than under the log-normal distribution, we next explore what distribution fits firm growth better. As firm growth can be both positive and negative, the standard Pareto or Zeta distribution from Section 2 are poor fits. Instead, we explore whether a two-sided version of these distribution can explain both the left and right-hand tails of firm growth. To see if this is the case, we separately calculate the rank of firm growth firms among the firms with negative growth ($y_{it} < y_{it-h}$) and positive growth ($y_{it} > y_{it-h}$) and plot the log of the inverse rank against the log of firm growth (y_{it}/y_{it-h}). This relationship is linear at either side of zero if the firm growth distribution is indeed two-sided Pareto.

Figure A12 plots the results. The subfigures have the same order as in Figure A11. The horizontal axis range is fixed at -4 to 4 in all plots. The figure shows that the two-sided Pareto distribution fits the empirical firm growth well. For product count, the linear reference lines are a near perfect fit for positive growth in both the 1-year and 5-year growth plots, although negative product growth seems somewhat less thick-tailed than the Pareto distribution at the 5-year horizon. For revenue growth and employment growth the tails appear to be *even fatter* than the two-sided Pareto distribution. This is particularly visible for the right tail of the distribution: as the scatter plot is mostly above linear reference lines, extreme growth in employment and revenue is more common in the data than in a Pareto distribution.

Conditioning on age or size We next assess how the tails of the firm growth distribution change with age and size. Figure A13 presents the QQ plots for all measures of 1-year growth, dividing the sample into three equally large age groups on the left-hand figures. Blue-round scatters plot the distributions for the youngest third of firms, while red-square and green-diamond scatters respectively belong to the middle and oldest age group. (describe result: tails decline with size, stronger result than by age).

Figure A12. Two-Sided Pareto Distribution of Firm Growth

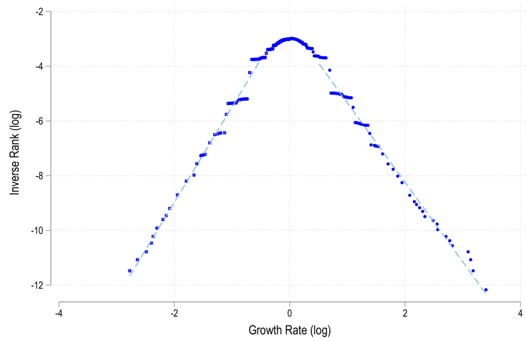

(a) Product Count Growth

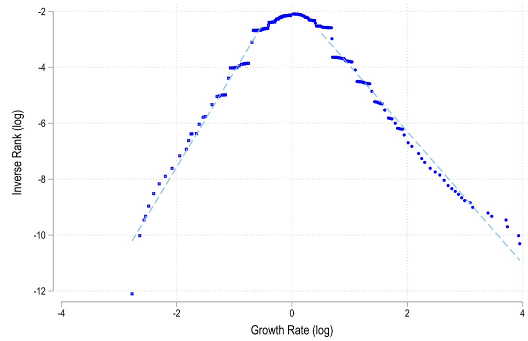

(b) Product Count Growth (5 year)

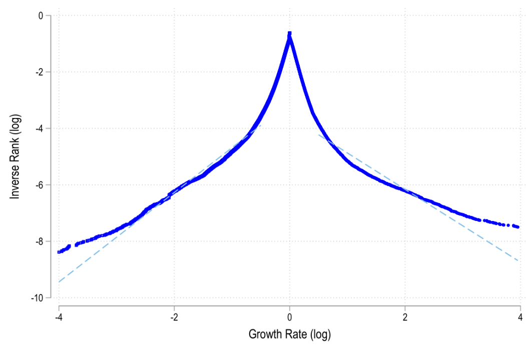

(c) Revenue Growth

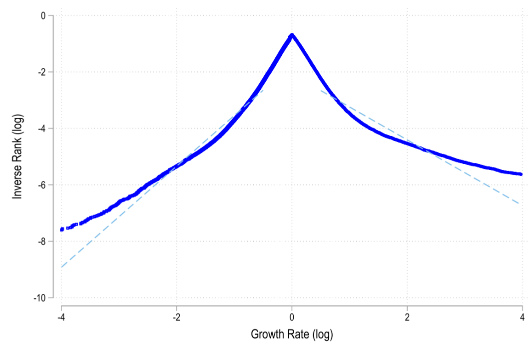

(d) Revenue Growth (5 year)

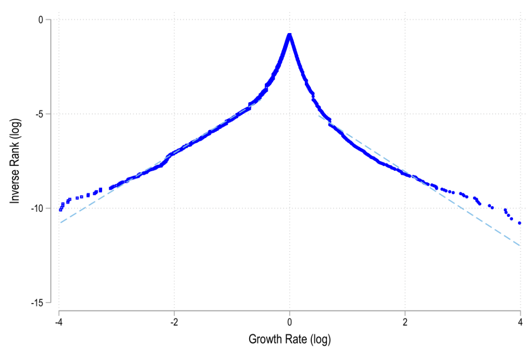

(e) Employment Growth

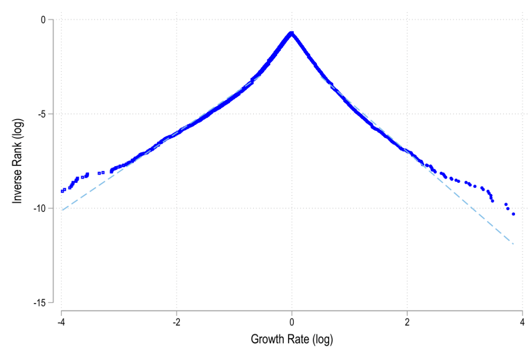

(f) Employment Growth (5 year)

Notes: The figures plot the relationship between the log of a firm's inverse rank for growth and the log of growth. Growth is defined as y_{it}/y_{it-h}. Rank is calculated separately for firms with negative growth and firms with positive growth. Blue-dashed lines are illustrative linear reference lines to asses whether

Figure A13. Distribution of Growth Rates by Age and Initial Size

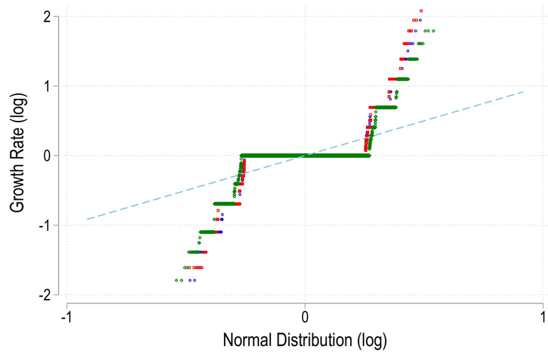

(a) Product Count Growth - by Age

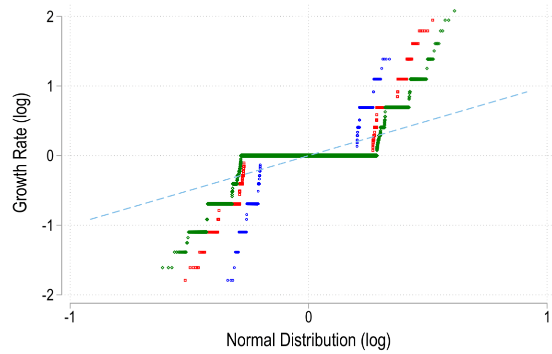

(b) Product Count Growth - by Size

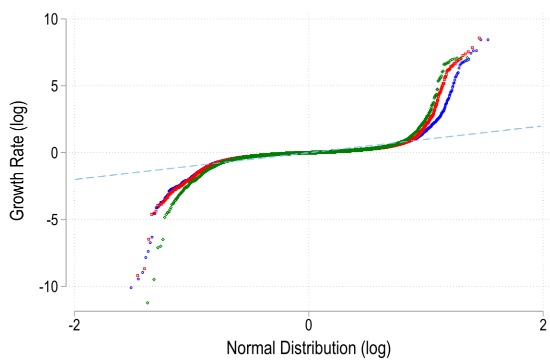

(c) Revenue Growth - by Age

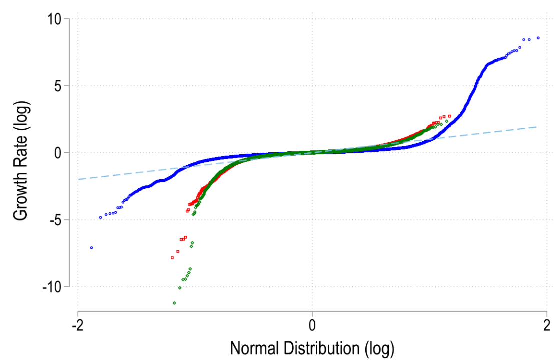

(d) Revenue Growth - by Size

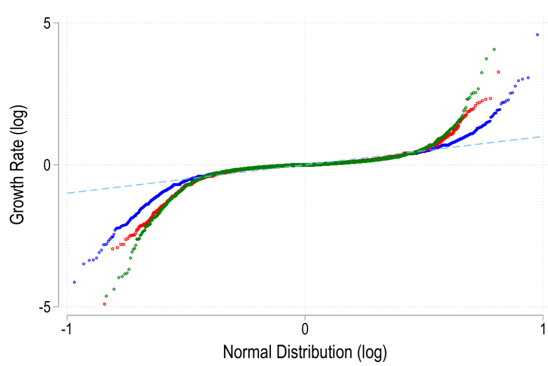

(e) Employment Growth - by Age

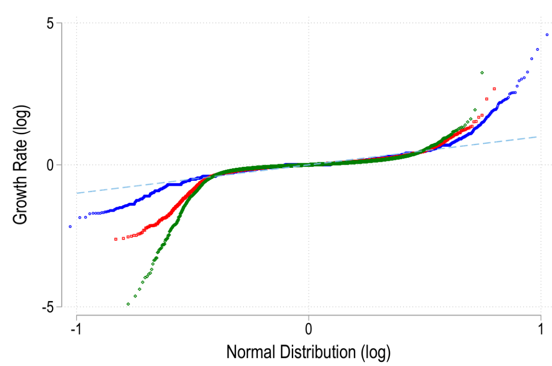

(f) Employment Growth - by Size

Notes: The figures plot Quantile-Quantile plots comparing the empirical distribution of log change in product count, revenue and employment against the log-normal distribution. Dashed light-blue lines present the reference log-normal distribution.

Figure A14. Symmetric Shapley-Owen Decomposition of Revenue Growth by Age and Size

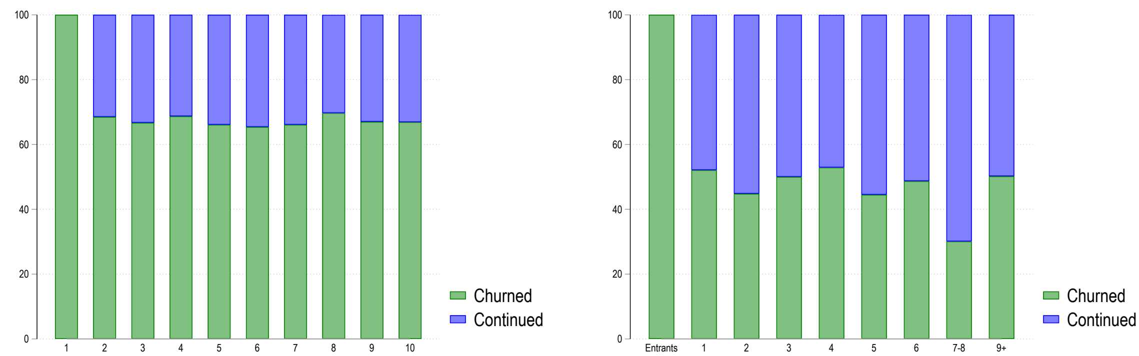

(a) Age

(b) Product Count

Notes: The figures decompose total revenue growth into revenue loss from product loss, revenue gain from product gain, and changes in revenue on products that the firm is continuing to produce.

Figure A14 presents the symmetric Shapley-Owen decomposition separately for firms of different age deciles (left-hand) and initial sizes (right-hand), grouping entry and exit into churn. Size is measured through the number of products that firms produced in the previous year.[18] The figure shows that the contribution of churn is similar for firms of different ages and sizes. The only outlier is entrants, which inherently see their entire growth attributed to churn as they have no continuing products.

[18]Results are similar when measuring size through revenue and employment.

Table A3: Overview of Churn

	Product Addition Rate					Product Loss Rate				
Horizon (years):	1	2	3	4	5	1	2	3	4	5
All firms	0.066	0.057	0.050	0.047	0.044	0.082	0.077	0.073	0.071	0.069
	(0.513)	(0.328)	(0.217)	(0.196)	(0.174)	(0.250)	(0.328)	(0.374)	(0.406)	(0.429)
All firms	0.055	0.049	0.046	0.045	0.044	0.075	0.072	0.069	0.070	0.069
Balanced	(0.317)	(0.259)	(0.229)	(0.203)	(0.174)	(0.240)	(0.318)	(0.367)	(0.404)	(0.429)
$N_{t-h} = 1$	0.055	0.051	0.048	0.048	0.048	0.059	0.059	0.058	0.060	0.060
$N_{t-h} = 2$	0.057	0.048	0.043	0.042	0.041	0.102	0.094	0.089	0.088	0.087
$N_{t-h} = 3$	0.060	0.049	0.041	0.038	0.037	0.110	0.100	0.094	0.092	0.091
$N_{t-h} = 4$	0.052	0.041	0.036	0.032	0.033	0.106	0.095	0.090	0.088	0.086
$N_{t-h} = 5$	0.049	0.044	0.038	0.034	0.031	0.116	0.102	0.095	0.091	0.090
$N_{t-h} = 6$	0.042	0.035	0.031	0.029	0.027	0.098	0.087	0.080	0.080	0.080
$N_{t-h} = 7, 8$	0.052	0.041	0.033	0.034	0.036	0.104	0.094	0.089	0.088	0.086
$N_{t-h} > 8$	0.073	0.055	0.048	0.045	0.043	0.115	0.102	0.096	0.096	0.094

Notes: Product addition rate is the number of new products that a firm starts producing divided by its original number of products. This number is divided by the length of the horizon to obtain an annual rate. Product loss rate is the number of products that a firm stops producing over the horizon divided by its original number of products. This rate is annualized by taking $\tau = 1 - (1 - \overline{\tau})^{1/h}$, where $\overline{\tau}$ is the raw fraction of products lost over horizon h.