Creative Destruction through Innovation Bursts

Giuseppe Berlingieri
ESSEC Business School

Maarten De Ridder
London School of Economics, CFM and CEPR

Danial Lashkari
Federal Reserve Bank of New York

Davide Rigo
London School of Economics

This version: 12 May 2026

Abstract

In theories of creative destruction, product innovation is a key driver of aggregate growth. This paper confronts the predictions of these theories about product dynamics with empirical patterns in product-level data on French manufacturing firms. We find that the process of product innovation frequently exhibits *bursts*—episodes in which firms rapidly add multiple products to their portfolio. As a result, 5% of firms are responsible for over 76% of product creation. Bursts lead to substantial shifts in revenue, and the process of product creation and destruction explains 88% of the variance in five-year firm growth. We introduce a model of firm product innovation that is compatible with the data while also nesting the canonical models of creative destruction. We show that innovation bursts alter the equilibrium composition of age, size, and innovation efficiency of firms, and explain the concentration of production among superstar 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.

*We thank three referees, Philippe Aghion, David Argente, Ben Moll, Nicolas Crouzet, Xavier Gabaix, Pete Klenow, Munseob Lee, Sara Moreira, Michael Peters, and Fabrizio Zilibotti, and numerous seminar participants for helpful feedback and discussions. 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).

1. Introduction

Product innovation lies at the heart of modern theories of firm growth through creative destruction (Aghion and Howitt, 1992; Grossman and Helpman, 1991; Aghion et al., 2014). According to these theories, innovative firms invest in research and development (R&D) to create new products superior to those of competitors, and to expand their market share, output, and profits. The canonical approach, pioneered by Klette and Kortum (2004), portrays product innovation as a gradual process, where firms incrementally add new products one at a time through independent innovation events.[1] This characterization implies steady firm growth, limiting the frequency of rapid expansions and the emergence of superstar firms.

In this paper, we use comprehensive product-level data from French manufacturing to show that such an account of product innovation may be at odds with empirical patterns. We show that firms experience what we term *innovation bursts*, rare but transformative episodes in which a firm rapidly expands its product portfolio, introducing many new products and gaining significant market share.

Such bursts of product creation may stem from innovations that spawn multiple commercial applications at once. Consider the cases of Laboratoire Science & Nature (LSN) and Picture Organic Clothing, two French manufacturers specializing in green and natural cosmetic products and sustainable outdoor apparel, respectively. For decades after its inception in the 1980s, LSN produced only a handful of products. Then, in the mid-2010s, it developed a way to use mineral-rich seawater to produce anti-aging skincare products such as deodorants, body lotions, and makeup removers. In a short span of time, LSN expanded its product range several-fold and substantially increased both employment and revenue.[2] A second example is Picture Organic Clothing, which shortly after its founding in 2008, developed a way to use sustainable inputs such as recycled polyester and bio-sourced polymers to produce a range of new outerwear, knitwear, helmets, and technical accessories. Picture subsequently registered 10 new product designs with the French intellectual property office in a single year, illustrating how a single idea around sustainable inputs enabled a rapid expansion of its product offering. Through the lens of canonical theories of creative destruction, the stories of LSN and Picture are exceedingly unlikely. Yet our data suggest that such bursts of product innovation are far more common than these theories predict.

What are the consequences of the burst-like nature of product innovation? We answer this question by constructing a model of growth and firm dynamics that matches the bursty distribution of product creation, and other facts on the contributions of entering/exiting vs.

[1] Since their seminal contribution, these theories have been extended in many directions to study the consequences of product innovation for many macro-level phenomena such as aggregate productivity growth, business dynamism, competition, and product-market concentration (e.g., Lentz and Mortensen 2008, Akcigit and Kerr 2018, Acemoglu et al. 2018, Cavenaile et al. 2020, Peters 2020, Aghion et al. 2023, Akcigit and Ates 2023).

[2] This information is in part based on an interview conducted in 2024 with Olivier Guilbaud, president of LSN.

continuing products to firm-level growth. Our model nests the canonical theories of creative destruction and allows us to examine the interplay between industry concentration and growth. We directly calibrate this model using the data on product dynamics, and show that the theory fits the data well. We then study two implications of innovation bursts for growth, showing that our model predicts a weaker relationship between firm size and innovation efficiency, yet a larger contribution of large firms to aggregate growth than the standard theories of creative destruction.

We begin the paper by using our product-level data from French manufacturing (introduced in Section 2) to present multiple facts on product creation and destruction (presented in Section 3). First, we verify some of the core assumptions of the standard theories of creative destruction on the invariance of the rates of product creation and destruction with respect to firm size and age. Next, we show that the distribution of the number of products is highly concentrated and has a thick, Pareto-like right tail. More importantly, we document that the distribution of newly created products also exhibits a substantial degree of concentration, with a similar Pareto-like right tail. In other words, we find bursts of product innovation, such as those exhibited by LSN and Picture, to be far from rare: in an average year, fewer than 5% of firms are responsible for over 76% of product creation. We provide extensive robustness checks to ensure that such bursts reflect true product creation, and are not the by-products of the data construction or reporting. Together, these facts point to innovation bursts as a key driver of the concentration of production among large firms.

We also use our data to measure the contribution of creative destruction to revenue growth at the firm and aggregate levels, defined as the share of growth explained by revenue changes from the introduction of new products or the loss of old ones. Unlike revenue growth on continuing products, creative destruction matters especially for firms experiencing rapid growth or decline, consistent with the role of innovation bursts in shaping firm dynamics and explaining observed variation in revenue growth across firms.

Finally, we examine the evolution of product-level revenue growth. The initial revenue of new products exhibits a persistent, firm-specific component. Beyond the first year, we find that growth in the revenue of a product exhibits a marked decline over the course of its life cycle. Thus, not only does product innovation grow firm revenue in the first year of the introduction of a new product, but it also enables the firm to sustain a faster rate of growth on this new product over the next few years than on a mature one.

In Section 4 we rationalize these facts in a model of endogenous growth in which firms invest in innovation to create new products or to improve the quality of their existing products. In the model, a firm's product innovation may bring about a technology with applications across multiple product lines, leading to a burst of new products. This burst-like stochastic process of product creation is consistent with the high concentration of new products in the data. It naturally nests the canonical Klette and Kortum (2004) model for the special case in which an innovation only expands firms' portfolio by a single product.

We model two additional features to help explain our facts on the evolution of product-level revenue. First, we allow for heterogeneity in innovation efficiency, to capture the observed persistent, firm-specific component of initial product revenue (Lentz and Mortensen, 2008; Acemoglu et al., 2018). Second, we allow investments in quality upgrading on existing products to exhibit diminishing returns; successive innovations to improve quality of a product using the same firm-level technology become increasingly incremental over its life cycle (Akcigit and Kerr, 2018; Acemoglu et al., 2022). As a product matures, firms find it increasingly *harder to find ideas* to improve the production process (Bloom et al., 2020).

Innovation bursts help explain the high concentration of production typically observed in firm-level data, associated with the Pareto distribution of firm size (Axtell, 2001; Luttmer, 2010). As is well known, canonical theories of firm dynamics through creative destruction can only match such high degrees of concentration if extended to allow for growth in the total number of products (Luttmer, 2011).[3] Here, we propose a distinct driver of the concentration of production, based on the lumpy nature of innovative ideas.

In Section 5, we quantify the model and assess its normative implications for innovation policy. Given the data's close links to the model, we calibrate most parameters directly from moments in our product-level data. For comparison, we similarly calibrate a conventional model of creative destruction without innovation bursts. We first confirm that, unlike the conventional model, ours can match the empirical distribution of product creation (unconditionally or conditional on the firm's product count), as well as the high degree of production concentration. While untargeted, the model also matches the average age of large firms, and the fraction of firms able to grow fast without an innovation burst.

Finally, we use the calibrated model to illustrate two implications of burst-like product creation. First, our model implies a weaker link between firm size and innovation efficiency than the conventional model. In the latter, high-efficiency firms are far more likely to grow by accumulating many products. In contrast, bursts can propel low-efficiency firms to suddenly grow large. Second, our model implies that large firms account for a greater share of growth than in the conventional model. This second finding stems from the fact that bursts generate a thick-tailed firm-size distribution, so that the greater concentration of production among larger firms more than offsets their lower share of high-efficiency types. Section 6 discusses how these findings can guide future work.

[3] In our data, the total number of products is fairly stable over time. Luttmer (2011) shows, even in an environment with sustained growth in the number of products, standard models predict that the largest firms in the data may take centuries to grow to their current scales. He proposes a theoretical solution in which young firms go through a phase of fast product creation before transitioning into their long-run behavior. For other examples of similar approaches to match empirical firm-size concentration, see, e.g., Cao et al. (2017); Acemoglu et al. (2022); Peters and Walsh (2021), and for an alternative approach based on a static assignment model, see Geerolf (2017).

Related Literature. We contribute to the extensive literature on the Schumpeterian growth models of creative destruction, confronting them with data on the dynamics of product churn and turnover. Most prior tests of these models' predictions regarding product innovation instead use indirect patent evidence (e.g., Akcigit and Kerr 2018).[4]

As two notable exceptions, Hottman et al. (2016) and Argente et al. (2024), using bar-code-level retail scanner data on non/semi-durable consumer products, show that product creation is a major contributor to firm growth. Hottman et al. (2016) study the drivers of cross-sectional firm size and growth and conclude that firms' product scope (and changes thereof) is the second largest determinant of both, after firm appeal. They also show that product churn accounts for the majority of firm growth in sectors with high product turnover. Argente et al. (2024) additionally document that revenue growth of existing products declines over time. They provide evidence that both creative destruction by competitors and cannibalization by the firm's own innovation drive these patterns, and present a model to study the interaction between firms' product creation and own-product innovation. Our paper is distinct in focusing on heterogeneity in the *distribution* of product creation and evidence for innovation bursts. Our data, furthermore, covers all of manufacturing (including durables) and defines products at a broader level than bar codes, ensuring clear differentiation among them. Theoretically, we attribute the life cycle of product revenue to the possibility that ideas to raise quality become harder to find as products mature.

Beyond theories of firm innovation and dynamics, Bernard et al. (2010) show that product dynamics, at the 5-digit industry level, are consistent with a multi-product extension of the Melitz (2003) model, while Broda and Weinstein (2010) use bar-code scanner data to study the implications of product dynamics for measuring aggregate prices. As customs records are widely available, exported products' dynamics have also been extensively studied (e.g., Fitzgerald et al., 2024; Albornoz et al., 2023). Others have explored implications of product dynamics over the business cycle (e.g., Dekle et al., 2015; Benguria et al., 2022).

Finally, multiple recent studies have documented the fact that industry concentration has been rising over the past few decades, in the US and globally (e.g. Autor et al. 2020; Kwon et al. 2024; Ma et al. 2024). Concurrently with the rise of large firms, there has been a gradual decline in the growth of total factor productivity, despite a sustained rise in R&D investments (Bloom et al. 2020). By constructing a creative destruction theory that directly matches the firm-level evidence on concentration and product innovation, our paper offers a laboratory for studying the interplay between concentration, innovation, and growth.[5]

[4] In a recent contribution, Argente et al. (2020) build an explicit empirical connection between firm-level patents and products and show their association is weaker among market leaders.

[5] Recent Schumpeterian models relating industry concentration to growth include Akcigit and Ates (2021, 2023), Olmstead-Rumsey (2022), Aghion et al. (2023), Cavenaile et al. (2023), Weiss (2023), De Ridder (2024).

2. Data

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

2.1. Sources

Our main source of information on firm product portfolios is the *Enquête Annuelle de Production* (EAP), a survey collected by the French statistical office (INSEE). The EAP covers the universe of manufacturing firms with at least 20 employees or 5 million euros in revenue, comprising around 90% of the aggregate gross output of manufacturing.[6] The data is available from 2009 to 2019. We start our analysis in 2010, when the survey methodology was finalized. We drop firms in EAP that do not belong to the manufacturing sector, such as those in mining, repairs, and installation industries.

The survey contains revenue, quantity, and average unit values for each product category that a firm produces each year at the level of 10-digit product codes following the PRODFRA classification, which is the official classification of French products published by INSEE. The high level of detail in these codes, distinguishing among 4500 distinct products, enables us to investigate changes in each firm's product portfolio. The first six digits are the harmonized European classification of products by activity codes (CPA), which are sufficiently narrow to identify customs policies. The remaining digits contain a further sub-classification that is produced particularly for France by INSEE and that nests the 8-digit European PRODCOM classification.[7] To ensure the consistency of product codes over time, in our baseline analysis, we use a modified 10c-digit product classification obtained through a concordance procedure that affects around 10% of products (see Online Appendix B.1 for further details). All main results are robust to using the original 10-digit classification, where we instead drop products with changes in codes, rather than attempt to concord them.

2.2. Summary Statistics

We combine this product data with the *Fichier Approché des Résultats d'Esane* (FARE). FARE provides firm-level income statement and balance sheet data such as total sales, wage bill, and capital, for the universe of French firms. Our baseline sample is the intersection of FARE and EAP. Additionally, we obtain the number of plants from social security filings (*Déclaration Annuelle de Données Sociales*, DADS) and determine firms' ownership status

[6]These firms are responsible for 86 to 88% of manufacturing value added. The survey also samples small firms but redraws that sample annually. The sample thus lacks the panel dimension needed to analyze product dynamics.

[7]For example, the 8-digit code 20421945 includes "pre-shave, shaving and after-shave preparations," while the 10-digit codes 2042194510 and 2042194520 distinguish between "lotions" and "foams and gels."

Table 1: Summary Statistics

Variable	Mean	St. Dev.	Median	10th Pct.	90th Pct.	95th Pct.	99th Pct.	Observations
Sales	22291	283587	3319	584	29533	62433	297729	223883
Age	28	19	24	8	52	58	113	223883
Employees	67	307	20	4	120	227	811	223883
Revenue per Product	9607	179085	2001	316	15635	30491	109024	223883
Number of Products	1.93	2.64	1	1	4	5	10	223883
New Products	.2	.81	0	0	1	1	3	203156
Lost Products	.13	.8	0	0	0	1	2	193718
Continuing Products	1.81	2.32	1	1	4	5	9	193718

Notes: Employment is measured in full-time equivalents. Sales are in thousands of 2015 euros. Age is defined as years since the firm was founded according to the firm census (FARE).

(whether independent or part of a group) using the *Liaisons Financières entre Sociétés* (LiFi) data. All data is merged using a common firm identifier, the SIREN code.

Table 1 provides the summary statistics. Focusing on the product-level information, we find that the median firm in a typical year has a single product and adds or loses no products every year. The picture is somewhat different for the average firm that produces 1.93 products, adds 0.13 new products, and loses 0.20.[8]

The table does not fully reflect the high degree of concentration of the distribution of product counts. Figure 1a plots the distribution and its degree of concentration by plotting the firm's product count against its ranking in the sample (logarithmic scale). The relationship between product count and its rank is approximately linear. Thus, the number of firm-level products follows an approximate power law with a Pareto-like right tail. The slope of the count-rank relationship implies that the largest 20% of firms produce 46% of all products.[9] This is in line with the widely documented fact that other measures of firm size (e.g., revenue and employment) also follow a power law and are highly concentrated among a small share of superstar firms (see, e.g. Axtell 2001, Luttmer 2007, Gabaix 2009, Garicano et al. 2016).[10] Figure 1b shows that the right tail of the distribution of revenue per product in our data also follows a power law.

[8]Because the EAP has a size threshold, firms may have been producing the same products before they first appear in the data. To avoid exaggerating product creation, we therefore set a new firm's number of new products to missing unless the firm simultaneously enters the EAP and the comprehensive FARE data. Similarly, we only include lost products of firms that exit the EAP if they also exit FARE.

[9]Table C3 reports Pareto tail indices for product count, employment, revenue, new products, and revenue per product.

[10]Among firms in our sample, the largest 20% of firms are responsible for 84% of revenue and employ 78% of workers. 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 equals $y = x^{(\theta-1)/\theta}$, where θ is the tail index. For the 80/20% rule, we find a Pareto tail index of around 1.16, fairly close to the tail index of 1 implied by Zipf's law.

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

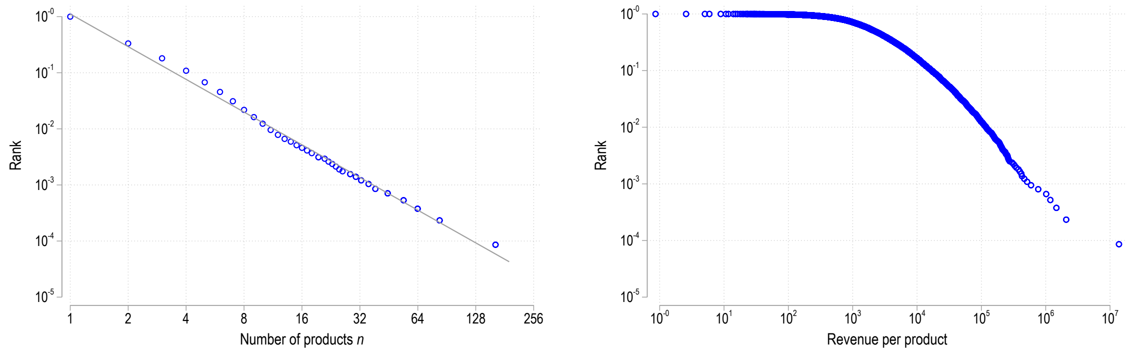

(a) Number of Products (b) Revenue per Product

Notes: The figures plot the relationship between a firm's size (horizontal axes) and the firm's rank (vertical axes). Size is either product count or revenue per product (in 1000s of 2015 euros). Rank is measured as the ratio of firms' rank starting from the largest firm, divided by the total number of observations in the data, so that rank equals the share of firms with higher products or revenues. Plots are based on the FARE-EAP for 2019. The distribution for 2010 to 2018 is plotted in the Companion Note.

3. Stylized Facts on Product Creation and Destruction

In this section, we use our data to document a number of facts on the dynamics of firm-level product portfolios and how they shape the dynamics of firm revenue growth.

3.1. How Do the Rates of Product Creation/Destruction Vary with Firm Size/Age?

We begin our investigation by testing one of the key predictions of the standard theories of creative destruction on the relationship between the rates of product creation, destruction, and firm size. Following Klette and Kortum (2004), these theories predict that the firm-level rates of product creation and destruction are independent of a firm's product count. In addition, they posit that product destruction occurs in an undirected and random fashion across all existing products. Together, these predictions yield a strong form of Gibrat's law: the growth rate of a firm's product portfolio is independent of its size.

Empirically, we define the product creation rate as the number of product codes for which a firm earns revenue at time t but not at $t-1$ divided by the number of product codes for which the firm earns revenue at $t-1$. The product destruction rate is the share of the firm's products at $t-1$ for which it does not report revenue at t.

Table 2 presents average product creation and destruction rates. We present unconditional averages and averages by number of products (top panel) and firms' age (bottom panel). We find that firms on average add 0.07 products to their portfolio per product that they initially produce, while they stop producing 0.104 products for every such product. The gap between incumbents' product creation and destruction rates means that the number of products that an incumbent firm produces on average shrinks over time. That is expected when new firms enter the economy, while the economy's total number of firm-products

Table 2: Product Creation and Destruction

By Size	All	Weighted	1	2	3	4-5	6-8	>8
Product creation rate	0.068	0.070	0.069	0.069	0.070	0.061	0.058	0.084
Product destruction rate	0.087	0.104	0.070	0.115	0.121	0.122	0.114	0.131

By Age Bins	All	0-5	5-10	10-15	15-20	20-25	25-50	> 50
Product creation rate	0.068	0.080	0.066	0.069	0.063	0.068	0.068	0.074
Product destruction rate	0.087	0.088	0.080	0.072	0.072	0.072	0.072	0.075

Notes: Product creation: number of products that a firm starts producing divided by its original number of products. Product destruction: number of products that a firm stops producing divided by its original number of products. Different columns report the results conditional on the firm's number of products in preceding period n_{t-1} (upper panel) or age bins (lower panel).

is approximately constant over time—which is the case in our data. The gap between the average rates of product creation and destruction, when weighted by product count, thus equals the contribution of entrants to product creation.

Turning to the results by initial product count, the table shows no particular pattern for product creation rates across firms with different initial sizes: the rate at which small and large firms expand their portfolio is similar, ranging from 0.058 to 0.084 per product initially produced. This means that the expected number of new products firms add to their portfolios scales linearly with the sizes of their portfolios. The rate of product destruction is similar for all firms with more than two products, but it is lower for firms that produce a single product. The unusually low rate among single-product firms is likely to stem from selection in the data: the EAP has a size threshold, implying that the single-product firms in the data earn higher revenue on their product compared to other firms.

Appendix Figure C1 plots average product creation rates net of four-digit industry fixed effects and for much wider initial sizes. That figure too shows no monotone or significant relationship between firms' initial size and either rate (as in Table 2). Table C1 additionally performs regressions to estimate the degree to which product creation scales in firm size, finding no consistent deviation from unity.

There is no systematic pattern in creation or destruction rates across age groups. The youngest firms (aged 0-5 years) have slightly higher churn rates than firms of nearby ages, but their average creation rates only marginally exceed those of firms aged 50+ years.

3.2. Innovation Bursts and the Concentration of Product Creation

We next investigate the distribution of product creation across firms. As mentioned in the introduction, all models building on the innovation process proposed by Klette and Kortum (2004) predict that, conditional on the firm's product count, the number of newly created products follows a distribution with a thin right tail. Thus, it is highly unlikely for firms to experience rapid product innovations in these theories. We focus on creation rather than destruction as product losses are bounded by firms' existing product count.

Figure 2. Distribution of Number of New Products

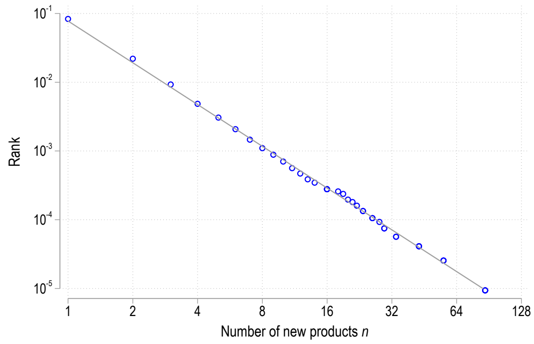

(a) All Firms

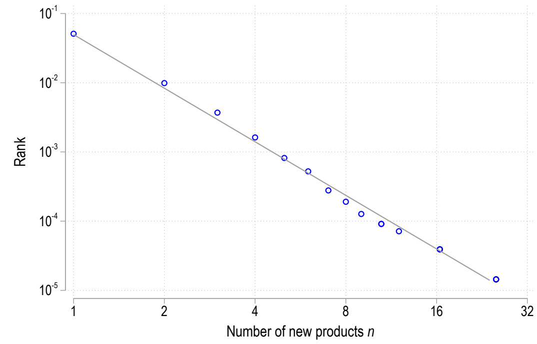

(b) Firms with $n_{t-1} = 1$

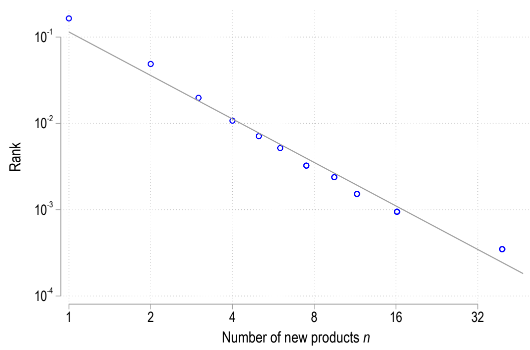

(c) Firms with $n_{t-1} = 5$

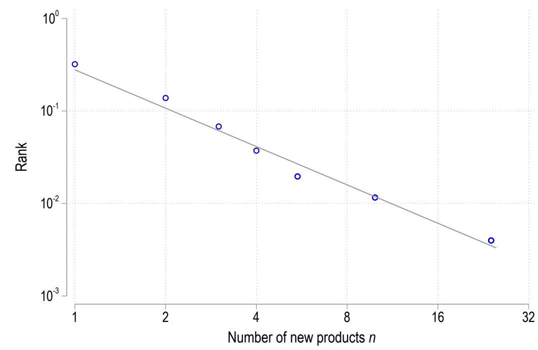

(d) Firms with $n_{t-1} = 10$

Notes: The figures plot the relationship between a firm's number of *new* products (horizontal axes) and the firm's rank (vertical axes), conditional on the firm's product count in the prior period (n_{t-1}). The (rescaled) rank is measured as the ratio of firms' rank starting from the largest firm, divided by the total number of observations in the data. Linear reference lines exclude the observation with the greatest number of added products; certain data points are combined to respect confidentiality rules.

Figure 2 plots the distribution of a firm's number of new products. We plot the relationship between the firm's number of new products and its rank in that distribution, as we did for firm size in Figure 1. Figure 2(a) plots the distribution across all firms, showing that it is highly concentrated.[11] The number of new products follows a Pareto-like distribution, with the linear regression having a near perfect fit.

In Figures 2(b,c,d), we plot the distribution for the number of products, but *condition on the firm's initial product count.* Doing so is useful as we have already seen in Figure 1 that the cross-sectional distribution of firms' product count is highly concentrated. That, in combination with the fact that product creation rates in Table 2 were stable in size, could suffice to make the distribution of product creation resemble the distribution in Figure 2(a). Conditioning on size thus answers the question of whether the concentration of product creation is merely a reflection of the cross-sectional concentration of firms' product

[11] The figure is a scatter plot where each circle plots the rank of firms with a particular number n of new products against that number. For cases where fewer than 3 firms add n products, we adjust the scatters to comply with confidentiality requirements. Data appendix B.5 details the adjustment.

Figure 3. Average Revenue Per New Product

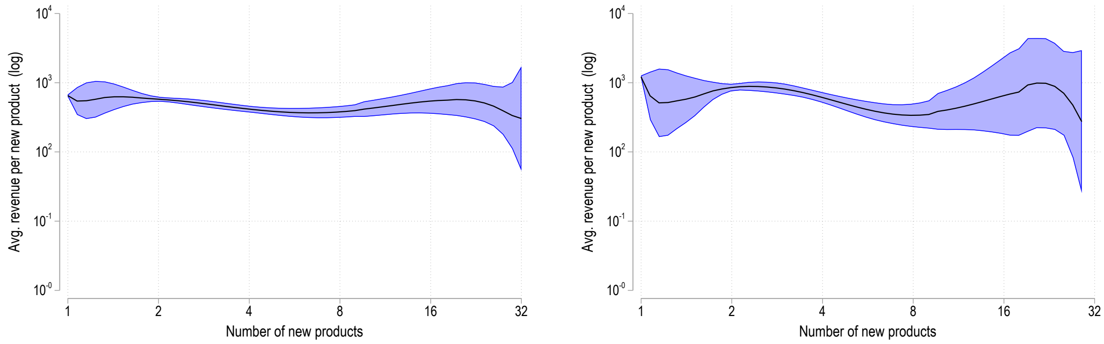

(a) All Firms (b) Firms with $n_{t-1} = 1$

Notes: Figures plot the logged 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. The line represents a kernel-weighted local polynomial smoothing of degree 5, together with its 95% confidence band.

counts.[12] Figures 2(b,c,d) show, to the contrary, that the high degree of concentration in product creation persists. Even among firms that initially produce 1 product, some add tens of new products. In other words, firms can experience sudden bursts of product creation that, as we will see, are exceedingly unlikely under benchmark theories of creative destruction. This pattern of product creation bursts is widespread across firms with any initial number of products in our data, including those with initially 5 or 10 products (Figures 2(c) and 2(d)).

Firms that experience an innovation burst also grow rapidly in sales. Figure 3 uses product-level revenue in the EAP data to plot average revenue earned *on new products* against the firm's number of new products. The left-hand figure does so for all firms; the right-hand figure conditions on size by restricting to single-product firms. The flat line implies that a firm creating, say, 10 products earns on average 10 times more revenue from new products than a firm creating one. Appendix Figure C2 shows the same pattern for firms starting with 5 and 10 products. If bursts simply reflected firms classifying similar products into many codes, or if the new products were close substitutes, one would expect revenue per new product to fall sharply with the number of new products. Instead, average revenue per product is similar across firms reporting different numbers of new products.

Robustness We next present various robustness checks to ensure that innovation bursts are indeed episodes in which firms expand their product portfolio. We show that innovation bursts are visible at different levels of aggregation, reflect organic growth rather than acquisitions, and do not cannibalize a firm's existing products. Moreover, we show that bursts are prevalent among firms of different ages and innovation histories, and across industries.

[12]The distribution in Figure 2(a) implies that 76% of creation comes from just 20% of firms (see Table C3). Note that this only accounts for firms that add at least one product to their portfolio, and 76.1% of firms report no product creation at all. Hence, 76% of product creation comes from just 5% of firms. For single-product firms concentration is even more extreme, with all product creation coming from just 7% of firms.

Figure 4. Distribution of Number of New Products by Level of Aggregation

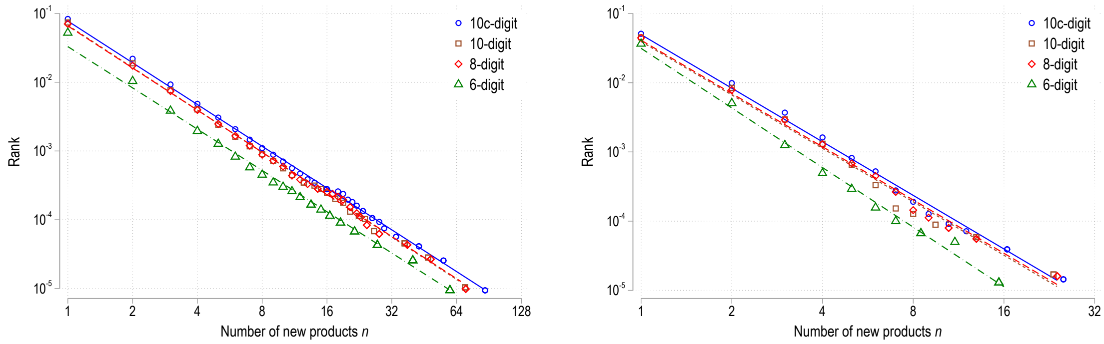

(a) All Firms (b) Firms with $n_{t-1} = 1$

Notes: The figures plot the relationship between a firm's number of new products (horizontal axes) and the firm's rank (vertical axes). The figure provides overlapping plots of the log rank against the log number of products added at the 10c-digit concorded product level (blue circles), 10-digit PRODFRA level (brown squares), 8-digit (red diamonds) and 6-digit level (green triangles).

Level of Aggregation Figure 4 shows that the distribution of product creation is similar if we define a "product" at different levels of aggregation. If innovation bursts are merely driven by the accuracy with which different firms classify and report their revenue across product codes, the distribution should become less concentrated if we define products at higher levels of aggregation. The figure shows that this is not the case: the distribution of the number of new products is similar at the 6-digit, 8-digit and the 10c-digit level. Figure C3 (Appendix C) presents similar results conditional on producing 5 and 10 products.

Mergers and Acquisitions We also verify that innovation bursts are driven by organic growth rather than mergers and acquisitions (M&A). If a firm's product portfolio changes due to the latter, those changes are unlikely to involve true product creation.

As M&A is difficult to track in the data, we use three different approaches to proxy for firm boundaries, two of which are plotted in Figure 5. As M&As are more prevalent among business groups, Figure 5(a) considers the set of single-product firms that do *not* belong to business groups; since M&As are likely to change the number of establishments a firm owns, Figure 5(b) considers the set of single-product firms that do *not* add any new plants between the two periods. In both cases, we find that the distribution remains fairly similar to our baseline, suggesting that innovation bursts are not likely to be driven by M&A. Appendix Figure C4 shows that this is also the case for firms of different sizes. Finally, we are also able to track whether firms experience a change in intangible capital. Intangible capital's chief component is goodwill, which increases when a firm acquires another for more than the target's book value. Appendix Figure C5 shows that, when excluding firms with changes in intangibles, the distribution of product creation remains almost unchanged.

Figure 5. Distribution of Number of New Products by Proxies of Fixed Firm Boundary ($n_{t-1} = 1$)

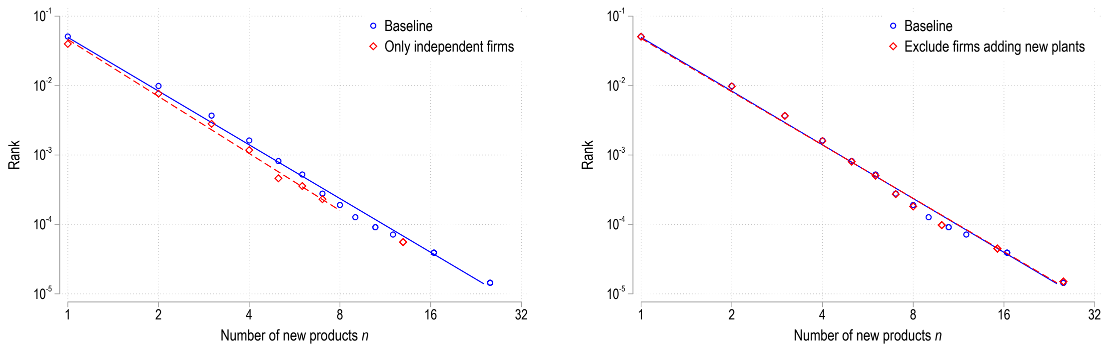

(a) Independent Firms Only (b) No Plant Change

Notes: The figures plot the relationship between a firm's number of new products (horizontal axes) and the firm's rank (vertical axes). The two panels plot the baseline Figure 2(b) against the distribution of new products of firms that do not belong to a business group (left panel) and of firms that do not add new plants over the period (right panel).

Existing Products We next show that innovation bursts have limited effects on firms' existing products. If bursts simply reflected firms reclassifying the product codes under which they report production, or if bursts predominantly cannibalized a firm's other products, then firms that add products would be more likely to cease producing their older goods or earn less revenue on them. Argente et al. (2024) find such cannibalization for nondurable, convenience consumer goods in scanner data. In our data, we only find modest effects of innovation bursts on existing products. Appendix Table C5 shows that the probability that a firm stops producing an existing product increases by only 8% if its product creation rate rises by 100%. If product creation merely reflected reassigning product codes, the odds would increase by 100%. On the intensive margin, Table C6 shows that product creation also has limited effects on the revenue of products that firms continue to produce.

Additional Robustness Online Appendix C shows that product creation is concentrated in bursts for firms of different ages, innovation histories, and sectors. Figure C6 shows that the slope of the log-rank, log-number of new products distribution of firms with below or above-median age is similar. Figure C7 finds the same for firms that innovated in the previous period and firms that did not. Moreover, we do not find any evidence for autocorrelation in bursts, i.e., firms adding many products in a previous episode of product creation are not more likely to be in the tail of product creation the next time they add products. Turning to sectoral differences, although our data are limited to manufacturing, we find that innovation bursts are prevalent across broad industries within the sector. In Figure 8 of the *Companion Note*, we identify the dominant broad-industry affiliation at different points in the product-creation distribution. Since all industries are represented throughout the distribution, innovation bursts are broad-based rather than concentrated in specific industries.

Online Appendix C presents several additional checks. Figure C8 restricts the sample to products that firms produce, design, and commercialize in-house, to rule out that bursts are driven by outsourcing or 'made-to-order' production. Figure C9 shows that bursts are also present among firms operating in consumer goods markets. Figure C10 shows that the distribution of product creation is unchanged when new products are defined as products a firm has *never* produced in the data, rather than as products not produced in the previous year. Finally, Figure C11 shows that firms are no more likely to stop producing products introduced as part of a burst than products introduced outside such bursts.

3.3. How Important Is Product Creation/Destruction for Firm Growth?

So far, we have focused on how firm product portfolios evolve through product creation and destruction. In this section, we use our data to study the contribution of such churn to firm dynamics, by decomposing firm-level revenue growth into the distinct contributions of product creation/destruction and the growth in the revenues of existing products.

We measure firm growth using the symmetric growth rate (Davis et al. 2006). 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 work on firm dynamics. We decompose the growth in revenue of firm i from time $t-h$ to t as

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

The first two terms capture growth in revenue due to product creation and destruction, the third term captures growth on products that firms continue to produce. Here, R_{it}^{N} denotes revenue on products produced at time t but not at $t-h$; R_{it-h}^{L} is revenue that the firm earned at $t-h$ on products that it stopped producing between $t-h$ and t; R_{it}^{-} is the revenue that the firm earns at time t on the products that it was already producing at $t-h$, and R_{it-h}^{+} is the revenue earned at $t-h$ on this same set of continuing products.

We quantify the contribution of product innovation and destruction in Table 3. In the top panel, we first present the average size of each term over a single and a five-year horizon. Observations are weighted by the denominator on the left-hand side of equation (1) to measure contributions to aggregate revenue growth. The first row shows that, even though each of the two creation/destruction terms have sizable effects, their net effect is small.[13]

This simple exercise may thus suggest from the table that the net effect of product creation and destruction, henceforth labeled "creative destruction," is small or even negative. However, this average result masks substantial heterogeneity across firms. The bottom panel of Table 3 measures how much of the *variation* in revenue growth is due to creative destruction. It does so using the Shapley value, which quantifies the marginal contribution of

[13]The contribution of product creation is more similar to product destruction over the longer horizon, which is likely because the sales of a new product does not reach its full scale in the first year. 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 (Bernard et al., 2010). This lowers the measured contribution of product innovation to revenue growth.

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	0.025	-0.026	0.005
5-year	0.027	0.094	-0.086	0.020
Shapley-Owen Contribution				
1-year	100.0	16.3	57.7	26.0
5-year	100.0	38.2	50.2	11.6

Notes: Decomposition of real revenue growth into revenue loss from product loss, gain from product gain, and changes in revenue on products the firm continues to produce. Observations in the upper panel are weighted by denominator $0.5(R_{it} + R_{it-h})$.

each component in decomposition (1) to the R^2 in regressions on all possible combinations of the components (see, e.g., Ozkan et al., 2023). Creative destruction explains 74% (88.4%) of the variation in revenue growth over 1-year (5-year) horizons. A variance decomposition as in Hottman et al. (2016) yields nearly identical results (see Table C2).[14]

In Figure 6, we show that this large contribution of creative destruction is due to its effect on the tails of firm revenue growth. Following the decomposition suggested first by Garcia-Macia et al. (2019), who apply this approach to account for the contribution of firms with different growth rates to overall job creation and job destruction, we split firms into 40 equally sized bins based on the firms' revenue growth. The top figure quantifies the importance of each bin in aggregate revenue growth. Moving from zero to the right side, the figure plots the ratio of the total change in revenue across the firms in each bin to the total change in revenue across all firms with positive growth. Thus, the height of the bars accounts for the contribution of firms in each bin to the overall growth of all growing firms. Moving from zero to the left, 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.

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. While most changes in overall revenue growth originate from the many firms that have modest growth rates, the tails of revenue growth still matter in the 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). As we move toward the tails of the distribution, the contribution of creative destruction gradually rises. In the extreme tails, the lion's share of the growth 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 creation is highly concentrated. A natural question is whether overall revenue growth is also highly concentrated and well

[14] Note that we exclude firms from the decomposition in their first year in the product-level EAP data if they have already appeared in the previous years of the comprehensive FARE dataset. Hence, firm-years with innovation bursts that push firms over the EAP size threshold are excluded. Thus, our results likely underestimate the importance of product creation for revenue growth. We also exclude lost products for firms that exit EAP but not FARE. This, too, may overstate the contribution of continuing products to revenue growth's variance.

Figure 6. Creative Destruction and Aggregate Revenue Growth

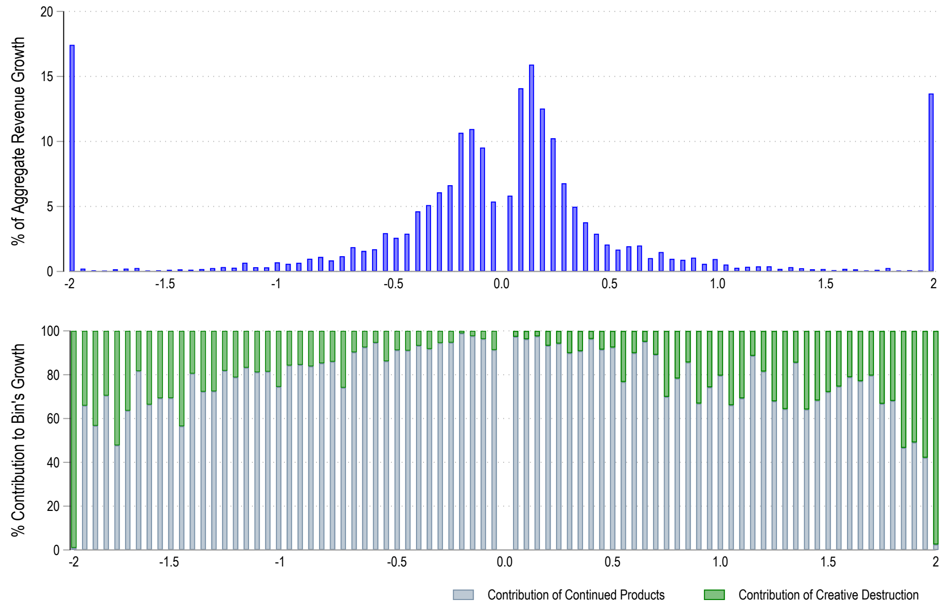

Notes: The horizontal axis measures firm growth through the symmetric growth rate, defined as the change in revenue between t and $t-1$ divided by average revenue in t and $t-1$. Growth rates are separated into 20 negative bins and 20 positive bins. The top figure presents the contribution of changes 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.

described by a power law. In the Companion Note, we confirm that the distributions of revenue and employment growth are much more concentrated than a log-normal distribution would predict.[15] Similar patterns have recently been documented using ORBIS (Jaimovich et al., 2023) and Compustat (Melcangi and Sarpietro, 2024) firm-level data.

3.4. Does Creative Destruction Matter for the Growth of Product-Level Revenue?

In Section 3.3 we saw that while creative destruction is key for the two tails of firm growth, the major share of the overall revenue growth comes from products that firms continue to produce. In the last part of our investigation, we study how creative destruction matters for the evolution of product-level revenue. We first show that the initial revenue of new products is driven by a firm-level factor, and then show that, after each instance of product creation, growth in product revenue subsequently falls over the product life cycle.

Revenues of Newly Created Products Are Auto-correlated at the Firm Level What determines the initial revenue that a new product brings to an innovating firm? Standard models of creative destruction (e.g., Lentz and Mortensen, 2008; Acemoglu et al., 2018) often

[15]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).

Figure 7. Firm-Level Persistence in Average Revenue from New Products

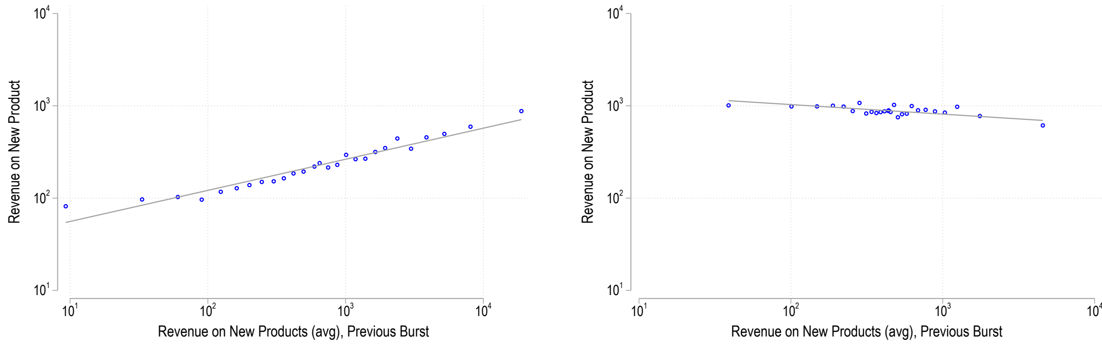

(a) Without Firm Effects

(b) With Firm Effects

Notes: The figure presents a binned scatter plot. Horizontal axes give average revenue per new product that an innovating firm earned the previous year that it added products to its portfolio. Vertical axes give revenue on a product that the respective firm is producing for the first time in that year. Both axes are on the log scale. We demean the data using 10-digit product fixed effects, and additionally by firm fixed effects on the right-hand figure. Appendix Table C4 gives the regression table for the linear fits.

assume ex-ante heterogeneity in innovative capacity of firms, which influences how much their new products upgrade the quality of the prior state-of-the-art techniques. Through the lens of these models, the initial revenue brought about by a new product reflects this degree of upgrading. Thus, we begin by studying whether the initial product-level revenue is autocorrelated for firms that experience multiple episodes of product creation.[16]

For firms that experience at least two episodes of product creation, Figure 7 shows the binscatter plot of the average revenue per new product in the current period against the average revenue per new product in their previous episode of product creation. We absorb 10-digit product fixed effects. The left-hand panel shows that on average, revenue on new products is higher for firms whose past instances of product creation had brought in relatively more revenue. The coefficient from a corresponding linear regression is 0.38 and is highly significant. Moreover, the positive relationship disappears when we additionally control for firm fixed effects in the right-hand panel, suggesting that the driver of the persistence is indeed a firm-level factor. This is in line with the existence of ex ante firm-level heterogeneity in the quality improvements embodied in product creation.

Product-Level Revenue Growth Falls over the Product's Life Cycle Lastly, we show that the growth of revenue of a continuing product gradually falls over the course of its life cycle. Figure 8 plots the relationship between revenue growth and product tenure—the time since a firm began producing a good. The first year of production is excluded to avoid mismeasurement of initial revenue growth coming from partial-year effects. Because firms might stop producing goods with low revenue growth sooner, we first demean revenue growth (ΔR_{ijt}) by the average revenue growth over the spell for which firm i produces

[16] Since the rate of product creation is fairly small, it is difficult to use the data to uncover systematic firm-level differences in the rates of product creation given the limited time span of our sample.

Figure 8. Life Cycle of Revenue Growth

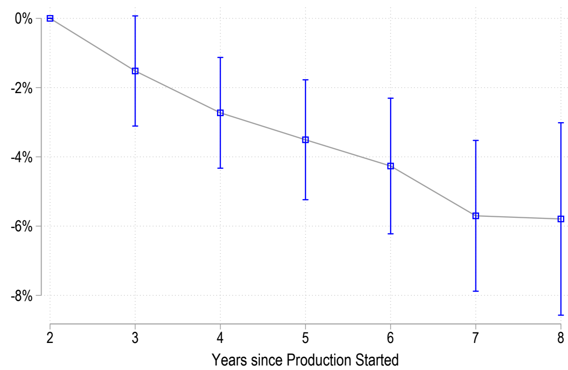

Notes: Avg. product revenue growth relative to growth at tenure 2 and after absorbing firm-product fixed effects, vs years since firm started producing the product. 90%-bounds are based on clustered std. errors at the firm-product level. To avoid mis-measuring tenure as a result of concordance, the figure uses the original PRODFRA codes. Figure 9 of the *Companion Note* shows similar results using concorded product codes.

good j, denoted $\overline{\Delta R}_{ij}$. This removes permanent differences in average growth rates across firm–product pairs and focuses on within-product changes over time. The figure then plots, for each tenure τ, the average demeaned revenue growth relative to the second year of production: $\mathbb{E}\left[\Delta R_{ijt} - \overline{\Delta R}_{ij} \,\middle|\, \text{tenure}_{ijt} = \tau\right] - \mathbb{E}\left[\Delta R_{ijt} - \overline{\Delta R}_{ij} \,\middle|\, \text{tenure}_{ijt} = 2\right]$.

Thus, the figure shows how a firm-product's revenue growth deviates from its spell average as tenure increases. This can be interpreted as the life cycle profile of revenue growth under the assumption that the exit of products is not driven by transitory growth shocks.[17] The figure shows a significant decline in revenue growth over a product's life cycle, averaging a decrease of about a percentage point per year.

There are several potential drivers of this pattern of revenue growth over the life cycle. First, firms may increase the appeal of their products by improving quality or by cutting the production costs of the products that they sell. The results in Figure 8 suggest that, if this is the driver of the life cycle of revenue growth, firm-level opportunities to sustain growth through the upgrading of own products decline as the product matures. The decline in revenue growth in a product's tenure is present across firms of different initial size, age, and industry (see Figures 10, 11, and 12 in the *Companion Note*, respectively).

But other forces may equally be at play. The gradual introduction of new products by firms could, for example, cannibalize their existing products. As mentioned in the introduction, Argente et al. (2024) document a similar pattern to Figure 8 at the bar-code level using scanner data from the US. They find that roughly two-fifths of the life-cycle pattern is driven by cannibalization. In our informal test of the effect of new products on firms' existing

[17] Appendix Figure C12 plots the life-cycle of growth in unit values. In canonical theories of firm dynamics and growth, and the model in Section 4, changes in revenue growth driven by variation in a product's quality are isomorphic to those driven by variation in its relative marginal cost: both have identical implications for revenue, quality-adjusted prices, and for aggregate productivity. Prices in the data, however, are not quality adjusted, so that the two mechanisms have distinct effects on unit values. Falling marginal costs raise revenue by shifting supply outward and reducing prices, whereas quality raises revenue by shifting demand outward and increasing prices.

product revenue (Appendix Table C6), we find little evidence of cannibalization. This may partly reflect the fact that we analyze 10-digit product codes rather than bar codes. Revenue growth may also decline over the life cycle in the presence of customer acquisition frictions: growth is initially high as new customers begin consuming a product, but diminishes once the product reaches its full customer base. Evidence on the role of customer acquisition frictions can be found in, e.g., Eaton et al. (2011), Einav et al. (2021), and Argente et al. (2025). Ignaszak and Sedlácek (2026) provide a comprehensive review.

4. Model

This section describes our parsimonious model of creative destruction through innovation bursts that rationalizes the facts documented in Section 3. We start by describing the environment and then characterizing the model's balanced growth equilibrium.

4.1. Environment

The main equations that characterize the economic environment are given in Table 4.

4.1.1. Households, Firms, and the Aggregate Economy

An infinitely-lived representative household has log utility over consumption C_t, discounting at a rate ρ over continuous time t. It supplies a single unit of labor inelastically. Consumption is a CES aggregate of a unit-measure continuum of differentiated goods with elasticity of substitution $\epsilon > 1$, as in equation (2). It nests the Cobb-Douglas aggregator in Klette and Kortum (2004) for the case where ϵ approaches one.

We denote by q_{ijt} the level of quality with which firm i produces good j at time t and I_{jt} is the set of firms that own a technology to produce good j at that time. Firm i is defined by the combination of its innovation efficiency and the set of products for which it produces at the leading quality level. Regardless of quality, all firms produce quantity y_{ijt} of product j with the same unit labor productivity, i.e., $y_{ijt} = l_{ijt}$ where l_{ijt} denotes production labor. Firms that can produce each j compete à la Bertrand, but subject to an infinitesimally small cost for operating in that market. Given the identical marginal costs across firms, only the firm with the highest-quality technology will enter and produce j in equilibrium, charging the monopoly markup $\epsilon/(\epsilon - 1)$, as in Acemoglu et al. (2018).

Letting the aggregate price index be the numeraire, the demand for product j from the firm i with the highest quality is $y_{ijt} = C_t\, p_{ijt}^{-\epsilon}\, q_{ijt}^{\epsilon-1}$ where p_{ijt} denotes the product price. As we saw this price is a constant markup $\epsilon/(\epsilon - 1)$ over marginal cost, which equals the wage

Table 4: Main Equations of the Economic Environment

Object	Equation	
Aggregator of aggregate consumption (C_t)	$C_t^{\frac{\epsilon-1}{\epsilon}} = \int_0^1 \left(\sum_{i \in I_{jt}} q_{ijt} y_{ijt}\right)^{\frac{\epsilon-1}{\epsilon}} dj$	(2)
Product-level demand (y_{ijt}) by quality (q_{ijt})	$y_{ijt} = Y_t \left(\frac{\epsilon}{\epsilon-1} w_t\right)^{-\epsilon} q_{ijt}^{\epsilon-1}$	(3)
Aggregate productivity (Q_t)	$Q_t^{\epsilon-1} = \int_0^1 q_{jt}^{\epsilon-1} dj$	(4)
R&D costs by bursts arrival rate ($x_{b,it}$)	$z_{b,it} = \eta_b \, x_{b,it}^{\psi} \, n_{it}^{-(\psi-1)\sigma}$	(5)
Innovation bursts size (n_{it}^c) distribution	$P(n_{it}^c = n) = \frac{n^{-\theta}}{\zeta(\theta)}$, where $\zeta(\theta) = \sum_{i=1}^{\infty} \frac{1}{i^\theta}$	(6)
Distribution of quality improvements ($\lambda_{b,ijt}$)	$\Pr(\lambda_{b,ijt} \leq \lambda \mid h) = H_{b,h}(\lambda)$	(7)
Distribution of firm types (h) over entrants	$\Pr(h_i = h) = G(h)$	(8)

w_t. Thus, product demand is an isoelastic function of quality given by equation (3), where we have also used the fact that consumption C_t in equilibrium equals output Y_t.[18]

Aggregate productivity Q_t is the ratio of output to production labor. It is given by equation (4), where q_{jt} is the highest quality for product j at time t, so that $q_{jt} = \max_{i \in I_{jt}} q_{ijt}$. Wages are a constant markdown over productivity, $w_t = (\epsilon - 1)/\epsilon \cdot Q_t$. Substituting this result in equation (3) allows us to write the per-product profit of firm i as a function of product j's quality relative to aggregate productivity $\widehat{q}_{ijt} \equiv q_{ijt}/Q_t$ following

$$\pi_t(\widehat{q}_{ijt}) = \tfrac{1}{\epsilon} \, \widehat{q}_{ijt}^{\,\epsilon-1} \, Y_t.$$

As we will see below, incumbent firms and potential entrants also hire workers to conduct research and development (R&D). Defining total production labor L_t as all labor used to produce goods, $L_t = \int_0^1 \int l_{ijt} \, di \, dj$, labor market equilibrium requires

$$L_t = 1 - L_t^e - L_t^{rd},$$

where L_t^{rd} and L_t^e denote labor involved in R&D by incumbents and entrants, respectively.

4.1.2. Innovation Bursts

Firms expand their product portfolio through innovation bursts that occur according to a stochastic Poisson process with endogenous arrival rates. Each burst generates an "idea" that enables the production of one or more products the firm has not previously offered, each drawn randomly from the economy's continuum of product lines. Each of these new

[18] Products in our model differ only in quality, not in firms' production efficiency. This is because quality and productivity are isomorphic: improvements in either lower the quality-adjusted price at which products are sold and have identical effects on the aggregate productivity index in equation (4). See Aghion et al. (2025) for a setting where this isomorphism does not apply, as quality and cost-enhancing innovations have different spillovers.

products surpasses the quality of the prior incumbent in the corresponding product line. Large innovation bursts are breakthroughs with applications across a broader range of technologies, enabling firms to advance the quality of multiple product lines using a single idea.

Formally, we model each innovation burst as a set with n^c_{it} elements, where each element corresponds to a randomly chosen, newly created product $j \in [0,1]$ with a corresponding quality improvement $\lambda_{b,ijt} > 1$ over the product line's incumbents' quality q_{-ijt}, so that the quality q_{ijt} at which the innovator produces its new products is $q_{ijt} = q_{-ijt}\,\lambda_{b,ijt}$.

The number n^c_{it} of new products in the burst is a random variable that follows a Zeta distribution, with the probability mass function given by equation (6), in which $\zeta(\theta)$ is the zeta function. The Zeta distribution is the discrete counterpart of the Pareto distribution. Tail parameter $\theta > 1$ determines the thickness of the tail of innovation bursts. As θ approaches one, innovation bursts follow Zipf's law. As θ increases, bursts become less dispersed. The average number of product innovations in a burst equals the ratio of $\zeta(\theta-1)$ and $\zeta(\theta)$, which is finite if θ exceeds two. In the limiting case where θ approaches infinity, every burst has a single product, and the model reduces to the standard process of Klette and Kortum (2004).

To achieve a Poisson arrival rate $x_{b,it} \geq 0$ of innovation bursts, firm i must hire R&D researchers $z_{b,it}$. The relationship between the arrival rate and R&D is given by equation (5), where n_{it} denotes the firm's current product count. The number of researchers required is convex in the burst rate ($\psi > 1$) and declines in the number of goods that the firm produces ($\sigma > 0$). The latter implies that current access to more leading technologies makes firms more productive innovators. Following Klette and Kortum (2004), this reflects the idea that a firm's product count also proxies for its knowledge or organizational capital, making its R&D more productive and ensuring that innovation rates scale with product count. As in Akcigit and Kerr (2018), we allow the degree of scaling to be governed by σ.[19]

4.1.3. Firm Heterogeneity

In addition to the heterogeneity in the number and qualities of their products, firms are also heterogeneous in terms of the efficiency with which they use R&D researchers to generate innovation bursts. We model this heterogeneity through variations in types h of the distribution $H_{b,h}(\lambda)$ from which firm i draws the quality improvement $\lambda_{b,ijt}$ embodied in its new product j (equation 7, following Lentz and Mortensen 2008).

Firms with a higher type h on average achieve greater quality improvements in their bursts, in the sense that the average moment $\overline{\lambda}_{b,h}$ defined as

$$\overline{\lambda}_{b,h} \equiv \mathbb{E}_{b,h}[\lambda^{\epsilon-1}]^{1/(\epsilon-1)},$$

[19]The standard case in which product creation rates scale linearly with product count requires $\sigma = 1$, which is a natural benchmark. When $\sigma < 1$, the marginal option value of R&D from adding a product declines with firm size. This generates the unintuitive implication that large firms have an incentive to split up as they grow.

is increasing in h. As subsequent innovators improve the quality of the product in proportion to its current quality, these firms also generate a greater positive externality to other firms. We assume that the firm's innovation efficiency is a fixed firm characteristic, although it is straightforward to allow types to evolve stochastically (as in, e.g., Acemoglu et al. 2018).

4.1.4. Firm Innovation on Own Products

In addition to investing in innovation bursts, firms also engage in targeted R&D to upgrade the quality of the products currently in their portfolios. Quality raises demand, so higher quality raises revenue on the upgraded product.

Successful own-product innovation raises the quality by which good j is produced by a factor $\lambda_{o,ijt} \equiv 1 + \lambda\,\beta^{s_{ijt}}$. Here λ is a random component drawn from a distribution $H_o(\lambda)$, s_{ijt} denotes the firm's number of past own-product innovations on good j. Similar to how diminishing returns over product life cycles are modeled by Acemoglu et al. (2022), own-quality improvements diminish as innovations accumulate when $0 < \beta < 1$. The average moment of own-product innovations,

$$\overline{\lambda}_o\,(s) \equiv \mathbb{E}_{H_o}[\left(1 + \lambda\,\beta^s\right)^{\epsilon-1}]^{1/(\epsilon-1)},$$

then decreases in the number of steps that existing producers have implemented.

To achieve an own-product innovation rate of $x_{o,ijt}$, firm i must hire $z_{o,ijt}$ researchers along

$$z_{o,ijt} = \eta_o\, x^{\psi}_{o,ijt}\, c_o(\widehat{q}_{ijt}, s_{ijt}),$$

where $\eta_o > 0$ shifts the cost of the quality upgrading innovation on own products. R&D on own-product innovation rises in the relative quality of the product and falls in the number of previous process innovations. For ease of analysis, and because own-product innovation is not a novelty of the paper, we choose a functional form for $c_o(\widehat{q}_{ijt}, s_{ijt})$ so that the equilibrium own-product innovation rate does not depend on $\widehat{q}_{ijt}$ and s_{ijt}. Hence firms choose equal innovation rates across their product portfolio. We also assume, for ease of exposition, that firms face a fixed operating cost for producing a good that exactly equals the option value of investing in own-product innovation.

Besides quality changes arising from innovation, we assume product quality may also change over time due to independently and identically distributed consumer taste shocks. Shocks are needed in order to match the variance of revenue growth on continuing products. Importantly, we assume such shocks apply to all potential producers of a good, so that they affect the producer's revenue and profits, while not altering the identity of the producer.

4.1.5. Entry and Exit

There is a unit measure of potential entrants that hire researchers to create a new product. To achieve an innovation arrival rate, x_{et}, the potential entrant must hire $z_{et} = \eta_e x^{\psi}_{et}$ researchers, as in Lentz and Mortensen (2008) and Acemoglu et al. (2018). Entrants are ex-ante identical: they know the probability $G(h)$ that they will have innovation efficiency type h but only learn about their type after creating their new product (eq. 8).

Firms stop producing a good when a different firm comes up with a higher-quality version of that product. This happens at the rate of creative destruction, x_{dt}, which is endogenously determined by the rates of entry and of innovation bursts by incumbents. As usual, firms exit when innovation by other firms causes them to cease producing their sole good.[20]

4.2. Balanced Growth Path Equilibrium

We next characterize the model's *balanced-growth* equilibrium path, along which aggregate productivity Q_t grows at a constant rate g. As labor supply is constant, the growth in Q_t is the sole source of output growth in the balanced growth equilibrium. Household optimization yields the standard Euler equation, $r - g = \rho$, where r is the interest rate.

The remainder of this section describes firms' value functions and optimal innovation policies along the balanced growth path, and derives the resulting distribution of firms by product count and innovation efficiency. Finally, we derive an expression for the balanced growth rate g, to understand how the burstiness of innovation affects aggregate growth.[21]

4.2.1. Firm Value, Innovation Policies, and Creative Destruction

The following proposition describes the firm value function along a balanced growth path equilibrium for a firm that chooses the value-maximizing rate of innovation bursts:

Proposition 1. The value of a firm with innovation efficiency h that produces a portfolio J with relative qualities $\widehat{q}$, where $n \equiv \mid \widehat{q} \mid$ is the product count, grows at rate g along the balanced growth path and equals

$$V_{h,t}(\widehat{q}) = w_t \left(v^* \sum_{j \in J} \widehat{q}_j^{\epsilon - 1} + n \, O_h(n) \right), \tag{9}$$

[20] In practice, some product creation may reflect the introduction of new varieties rather than the immediate destruction of existing products. In our data, however, the number of firm-products does not rise over time. Hence, newly created varieties would have to be offset by the destruction of old varieties. That could emerge, for example, if production comes with a fixed cost, so that there is a minimum level of quality below which a variety is no longer profitable to produce (see, e.g., Acemoglu et al. 2018).

[21] We can also characterize firms' value functions, their optimal investments in innovation bursts and own-product innovation, and the implied evolution of the mass of firms across product counts and types along the transitional dynamics of the equilibrium paths leading to the balanced growth paths presented in this section. For details, see Appendices A.1.1, A.1.2 and A.1.3.

where $v^* \equiv \frac{L/(\epsilon-1)}{\rho+(\epsilon-1)g+x_d}$ captures the present value of expected per-product profits while $O_h(n)$ captures the per-product option value of product creation for a firm of type h with n products. This option value is the solution to the recursion

$$\begin{aligned}(\rho + n x_d)\, O_h\,(n) &= (n-1)x_d\, O_h(n-1) \\ &\quad + \Psi\left(\frac{\zeta(\theta-1)}{\zeta(\theta)}\overline{\lambda}_{b,h}^{\epsilon-1}\, v^* + \sum_{k=1}^{\infty}\frac{k^{-\theta}}{\zeta(\theta)}(n+k)\, O_h(n+k) - n\, O_h(n)\right)^{\frac{\psi}{\psi-1}} n^{\sigma-1},\end{aligned} \tag{10}$$

where $\Psi \equiv (\psi-1)\psi^{-1}\left(\psi\eta_b\right)^{-\frac{1}{\psi-1}}$. The optimal rate of product creation $x_{b,h}(n)$ for a firm of type h with n products, the optimal rate x_o of own-product quality upgrading, and the entry rate x_e that satisfies the free-entry condition are respectively given by

$$x_{b,h}(n) = \left(\frac{1}{\psi\eta_b}\left(\frac{\zeta(\theta-1)}{\zeta(\theta)}\overline{\lambda}_{b,h}^{\epsilon-1}\, v^* + \sum_{k=1}^{\infty}\frac{k^{-\theta}}{\zeta(\theta)}(n+k)\, O_h\,(n+k) - n\, O_h(n)\right)\right)^{\frac{1}{\psi-1}} n^{\sigma}, \tag{11}$$

$$x_o = \left(\frac{v^*}{\psi\eta_o}\right)^{\frac{1}{\psi-1}}, \tag{12}$$

$$x_e = \left(\frac{1}{\psi\eta_e}\sum_h G\,(h)\left(\overline{\lambda}_{b,h}^{\epsilon-1}\, v^* + O_h\,(1)\right)\right)^{1/(\psi-1)}. \tag{13}$$

Proof*. See Appendix A.2.1.*

Equation (11) shows that the innovation-burst arrival rate rises with the average size of innovation bursts (i.e., falls in θ), falls with the innovation cost η_b, and rises with the present value of expected profits v^*, expected product improvement $\overline{\lambda}_{b,h}$, and product count n. Equation (12) shows that own-product innovation falls with the innovation cost η_o and rises with the present value of expected profits v^*. Free-entry condition (13) averages the value of successful single-product bursts over the innovation efficiency types.

The balanced-growth rate of creative destruction is given by

$$x_d = x_e + \frac{\zeta(\theta-1)}{\zeta(\theta)}\sum_{h,n} M_h(n)\, x_{b,h}(n), \tag{14}$$

where $M_h(n)$ is the stationary measure of firms with efficiency h and product count n. Creative destruction due to incumbents, the second term on the right, equals their rate of innovation bursts times the expected number of new products $\zeta(\theta-1)/\zeta(\theta)$ per burst.

To provide a more intuitive interpretation of the option value of R&D and the first order condition for the burst arrival rate, let us focus on the case of $\sigma = 1$. Under this parametrization, the per-product option value of product creation, $O_h(n)$, becomes independent of n

as in Luttmer (2011). Inserting the constant option value into (11), in turn, yields that the optimal burst rate per product, $x_{b,h}(n)/n \equiv \overline{x}_{b,h}$, is constant across firm sizes within types h:

$$(\rho + x_d)\, O_h = \Psi \left(\frac{\zeta(\theta-1)}{\zeta(\theta)} \left(\overline{\lambda}_{b,h}^{\epsilon-1} v^* + O_h \right) \right)^{\frac{\psi}{\psi-1}} \; ; \; \overline{x}_{b,h} = \left(\frac{1}{\psi \eta_b} \frac{\zeta(\theta-1)}{\zeta(\theta)} \left(\overline{\lambda}_{b,h}^{\epsilon-1} v^* + O_h \right) \right)^{\frac{1}{\psi-1}} . \tag{15}$$

The first order condition yields that firms equate the marginal cost of the burst arrival rate to bursts' marginal effect on firm value, which is the sum of the present value of expected additional profits and the additional R&D option value that the added products deliver.

As a per-product burst rate that is constant in product count aligns with the evidence that creation rates scale in size (Table 2), we focus on the $\sigma = 1$ case in the rest of this section.

4.2.2. Equilibrium Distributions of Firm Size and Innovation Efficiency

We next discuss how innovation bursts shape the firm-size and type distribution along the balanced growth path. This yields our first key theoretical result: in an economy where product innovations arise in bursts ($\theta < \infty$), the tail of the stationary firm-size distribution follows a power law. By contrast, we confirm that the firm-size distribution decays exponentially in an economy without innovation bursts, as noted in Luttmer (2010).

To find the stationary composition of firms, we use the fact that $M_h(n)$ evolves along the following Kolmogorov forward equation along the balanced growth path:

$$\dot{M}_h(n) = \overline{x}_{b,h} \sum_{j=1}^{n-1} j \, \frac{(n-j)^{-\theta}}{\zeta(\theta)} \, M_h(j) \; + \; (n+1)\, x_d \, M_h(n+1) - \; n \left(\overline{x}_{b,h} + x_d \right) M_h(n) \; + \; x_e \, G(h) \, \mathbb{I}\{n=1\}. \tag{16}$$

The first term on the right-hand side of the equality captures the flow of firms into product count n from smaller product counts due to innovation bursts, the second term captures the inflow from firms that lose one of their $n+1$ products due to creative destruction, the third term captures the outflow of size n firms due to product creation or creative destruction, and the last term captures the inflow of entrants into product count $n = 1$.

The stationary composition of firms is found by setting $\dot{M}_h(n) = 0$ for all sizes and innovation efficiencies. For single product firms, the resulting measure of firms by type is $M_h(1) = (x_e/x_d)G(h)$. For multi-product firms, the measure is given by the sequence

$$M_h(n+1) = \frac{1}{n+1} \left(n M_h(n) \left(\overline{x}_{b,h} + x_d \right) - \sum_{j=1}^{n-1} M_h(j) \frac{(n-j)^{-\theta}}{\zeta(\theta)} j \cdot \overline{x}_{b,h} \right)$$

The sequence yields a firm-size distribution of which the right tail is characterized by a power law, as it inherits the tail of the Zeta distribution of the innovation bursts:

Proposition 2 (Zeta tail of the firm-size distribution)**.** Suppose $\sigma = 1$ and $\theta > 2$. Then, for each type h, the stationary measure of type-h firms with n products satisfies

$$M_h(n) \sim \frac{\overline{x}_{b,h}\, m_h^2}{x_e\, G(h)\,(\theta-1)\,\zeta(\theta)}\, n^{-\theta} \qquad \text{as } n \to \infty, \tag{17}$$

where m_h is the mass of products of type-h firms, $m_h = \sum_{n=1}^{\infty} n, M_h(n)$, and satisfies

$$m_h = \frac{x_e\, G(h)}{x_d - \frac{\zeta(\theta-1)}{\zeta(\theta)}\, \overline{x}_{b,h}}. \tag{18}$$

The aggregate measure of firm-size satisfies

$$M(n) \equiv \sum_h M_h(n) \sim \left(\sum_h \frac{\overline{x}_{b,h}\, m_h^2}{x_e\, G(h)\,(\theta-1)\,\zeta(\theta)}\right) n^{-\theta} \qquad \text{as } n \to \infty. \tag{19}$$

Proof. *See Appendix A.2.2.*

Equations (17) and (19) show that the stationary distributions of firm size (whether unconditional or conditional on type) also have an asymptotic Zeta tail with the same power-law exponent θ as that of the bursts (6). The intuition is simple: a single large burst of size k can propel a new entrant to size n in one step, producing the same power-law decay in the firm-size distribution.[22] If every burst adds only one product ($\theta \to \infty$, as in Klette and Kortum 2004), the firm-size distribution decays exponentially rather than following a power law. Multi-product bursts are thus essential for generating the Pareto-like concentration of production observed in the data (e.g. Figure 1).

4.2.3. Aggregate Growth

Turning to the growth rate g of aggregate productivity, we can show that it satisfies

$$g = \frac{\zeta(\theta-1)}{\zeta(\theta)} \sum_h \left(\frac{\overline{\lambda}_{b,h}^{\epsilon-1} - 1}{\epsilon-1}\right) \left(\sum_n M_h(n)\, x_{b,h}(n)\right) + \sum_h \left(G(h) \frac{\overline{\lambda}_{b,h}^{\epsilon-1} - 1}{\epsilon-1}\right) x_e + \left(\frac{\mathbb{E}_s\left[\overline{\lambda}_o(s)^{\epsilon-1}\right] - 1}{\epsilon-1}\right) x_o,$$

where each of the three terms on the right-hand side accounts for one of the three underlying sources of growth: innovation bursts, innovation by entrants, and own-product innovation by incumbents. In each case, the contribution of that source is given by the product of the rate of innovation and the innovation's expected quality improvement.

[22] The proposition characterizes the tail of the stationary *measure* $M_h(n)$. The probability that a randomly chosen type-h firm has n products is $\widehat{M}_h(n) = M_h(n)/\overline{M}_h$, where $\overline{M}_h = \sum_n M_h(n)$ is the type-h firm mass. Since $\overline{M}_h$ is a positive constant, the probability has the same tail exponent θ.

4.2.4. Innovation Bursts and the Properties of the Balanced Growth Path

Which elements of the balanced growth path change when product creation is more or less burst-like? As Proposition 2 shows, one key aspect of the equilibrium, i.e., the *distribution* of firm-level product counts, strongly depends on parameter θ. A lower θ means that large bursts are more likely, so some firms grow very large in a single step. As a result, innovation bursts lead to a thick right tail of the firm-size distribution and allow the model to fit the observed concentration of production. It follows that innovation bursts alter any analysis for which the firm-size distribution is relevant.

For the benchmark case where $\sigma = 1$, the model completely separates the effect of innovation bursts on the firm-size distribution from aggregate variables such as aggregate growth or the aggregate composition of firms in terms of their innovation efficiency. For instance, we can show that in this case product creation affects the balanced rate of aggregate growth only through its *effective* rate–the arrival rate of bursts multiplied by their average size, $\frac{\zeta(\theta-1)}{\zeta(\theta)}\,\overline{x}_{b,h}$. As such, without size-dependence, a process in which ideas are rare but broad in scope (burst-like) grows along the balanced growth path similarly to one in which ideas are frequent but narrow (no bursts).[23] The same equivalence extends to certain other macroeconomic outcomes such as the rate of creative destruction, entry, and the fraction of products produced by each type. In fact, we can further show that, for any value of θ, we can adjust the R&D cost parameter η_b to keep the effective rate of product creation fixed and also preserve the values of all these aggregate quantities (see Appendix A.2.4 for the proof). Thus, this benchmark disentangles the firm-size distribution from long-run aggregate growth: innovation bursts reshape the former without affecting the aggregate implications of an otherwise standard neo-Schumpeterian model.

By the same token, the burst-like nature of product creation directly affects aggregate outcomes in any departure from this benchmark under which the firm-size distribution matters in the aggregate. Even when $\sigma = 1$ such cases can arise, for example if the economy is characterized by size-dependent wedges, or in applications where size is used as a proxy for firm characteristics that are typically unobservable to outsiders, such as a firm's innovation-efficiency type. Think of size-dependent R&D policy, for example, where subsidies depend on product count. When innovation-efficiency types are unobserved, such policies can be used to target high-type firms indirectly, which is desirable because R&D by those firms generates larger positive externalities. When θ is high, a firm with many products is likely to have reached that size through a sustained history of efficient R&D investments. A subsidy that targets larger firms is therefore more cost effective than a uniform policy in such an economy. When θ is low, by contrast, a firm may become large because of a single, unusually broad innovation burst. Innovation bursts, therefore, make firm size less informative about innovation efficiency. Section 5.5 elaborates on this compositional effect.

[23] As innovation rates are constant in size, the average type-specific arrival rate of bursts in equation (20) is given by $\sum_n M_h(n)\,x_{b,h}(n) = \frac{\zeta(\theta-1)}{\zeta(\theta)}\,\overline{x}_{b,h}\,m_h$ (see eq. 15 and 18). Appendix A.2.3 provides a formal derivation.

Table 5: Summary of Parameter Values

Parameter	Description	Value	Method
θ	Tail parameter of the zeta distribution	3.06	Grid search
σ	Degree of returns to product counts in product innovation	1.00	Direct
η_e	Product creation R&D cost scalar for entrants	10.5	Direct
η_b	Product creation R&D cost scalar for incumbents	11.0	Direct
ψ	Convexity of R&D costs in innovation rate	2.00	External
ρ	Discount rate	.020	External
β	Decline rate of follow-up process innovation size	.783	Direct
x_o	Poisson rate of process innovations	1.00	Direct
λ	Process innovation step size	.061	Direct
ϵ	Elasticity of substitution	4.00	External
Ξ	Standard deviation of idiosyncratic quality shocks	.354	Grid search
G	Entrant share of type L and H	[0.90,0.10]	Grid search
$\overline{\lambda}_{b,h}$	Avg. quality improv. size of product innovation L and H	[1.014,1.098]	Grid search

In settings where innovation rates do not scale proportionally with firm size, the effect of the burstiness of product creation on aggregate outcomes such as growth is immediate. In an economy with $\sigma < 1$, for example, larger firms innovate less than proportionally to their size, so a more concentrated firm-size distribution reduces growth.[24] Thus, if bursts allow firms to attain substantial size, they affect concentration and, through it, growth.

5. Quantification

We next quantify the model and show that it replicates the empirical facts documented in Section 3. We also use the quantified model to illustrate how the properties of an economy with growth through creative destruction depend on the burstiness of product creation.

5.1. Calibration

We infer most innovation-related parameters directly from the data on product dynamics. This contrasts with the common approach to calibrating models of creative destruction, which typically relies only on indirect inference from firm-level information. Table 5 lists all parameters, and their calibration method. Table 6 summarizes the targeted moments.

Product Innovation The key novel parameter of our model is the tail of the zeta distribution, θ, which determines the thickness of the tail of the distribution of product creation. We calibrate θ to match the relationship between the number of newly created products among single-product firms and their corresponding rank in the distribution, as in Figure 2b. At $\theta = 3.06$, the model matches the slope of -2.36 in the figure.[25]

[24] This is the case, e.g., in the models of Akcigit and Kerr (2018), Aghion et al. (2023), and De Ridder (2024).

[25] As firms may experience multiple bursts within a year, θ does not equal the slope in Figure 2. Instead we perform a grid search to find the θ that delivers the log-rank, log-new products relationship from a regression à la Gabaix and Ibragimov (2011), choosing the value of θ for which the regression estimate (-2.36) is matched.

Table 6: Targeted Moments and Related Parameters

Parameter	Moment	Target	Model
θ	Gabaix-Ibragimov reg. of new products (log) on rank (log) if N_{t-1}=1	-2.36	-2.35
σ	Scaling of number of innovation bursts in firm size	1.00	1.00
η_b	Arrival rate of new products (per existing product)	.070	.070
η_e	Implied contribution of entrants to creative destruction	.037	.037
x_o, λ	Revenue growth on continuing products	.005	.005
β	Change in revenue growth in product tenure	-.020	-.020
Ξ	Shapley-Owen contrib. of own products to revenue growth (%)	26.0	26.0
$G(h)$	Ratio of the variance of log revenue on new products over squared avg.	.890	.890
$\overline{\lambda}_{b,L}/\overline{\lambda}_{b,H}$	Persistence of revenue per new product within firms	.316	.316
$\overline{\lambda}_b$	Growth rate of manufacturing total factor productivity	.021	.021

Turning to the frequency of innovation bursts, we have two key parameters: (1) the cost scalar η_b which governs the average arrival rate of bursts, and (2) σ, which controls how product creation scales with size. We set σ to 1 to match the lack of a clear relationship between firm product count and product creation rates in Table 2. We calibrate η_b to match the size-weighted average of incumbent product creation rates of 0.070, while the average number of products per burst, given θ, is 1.33. The model delivers a burst arrival rate of 0.053, the ratio of these two values, when η_b equals 11.0. We set the curvature parameter ψ to 2, to deliver the standard R&D cost elasticity of unity (see, e.g., Bloom et al. 2002).

We do not observe product creation by entrants directly, as not all entrants immediately appear in the EAP. Instead, we infer the rate of entry from the fact that the share of all products in the economy single-product firms produce is pinned down by the entry rate: $M(1) = x_e/(x_e + \tilde{x}_b)$, where $\tilde{x}_b$ is the 0.070 average rate of incumbent product creation. $M(1)$ equals 34.5% in our data, implying a 0.037 entry rate, which is matched at $\eta_e = 10.5$.

Firm types To capture ex-ante heterogeneity in innovation efficiency, we assume that there are two firm types, $h \in L, H$. Both draw quality improvements from Pareto distributions $H_{b,h}(\lambda)$ with minimum value 1 but different tail parameters, with the respective average quality improvements $\overline{\lambda}_{b,L}$, $\overline{\lambda}_{b,H}$, where $\overline{\lambda}_{b,L}$ is smaller than $\overline{\lambda}_{b,H}$.

Three moments help calibrate the three corresponding parameters, $\overline{\lambda}_{b,L}$, $\overline{\lambda}_{b,H}$, and the share of low-type entrants G_L. The first is the aggregate productivity growth rate of 2.1% in the data, which pins down the average quality improvement from product creation.[26]

The other two moments jointly pin down the two remaining degrees of freedom, that is, the gap between $\overline{\lambda}_{b,L}$ and $\overline{\lambda}_{b,H}$, and the share of low-type entrants G_L. The first moment is the variance of the log of new-product revenue, normalized by squared average new-product log revenue to reduce sensitivity to miscalibration of the substitution elasticity ϵ. This ratio equals 0.89. The second is a regression coefficient analogous to the slope in Figure 7, designed to capture the extent to which some firms persistently earn higher revenue on newly introduced products than others. We measure this by regressing the log of revenue of

[26] Consistent with the model, we use average French manufacturing TFP growth in labor-augmenting form (i.e. raw TFP growth divided by the capital share) in EU-KLEMS over 2010–2019, following Aghion et al. (2025).

a new product j' on that of another new product j observed for the same firm, after taking 10-digit product code fixed effects. The regression coefficient equals 0.316.

Intuitively, greater heterogeneity in innovation efficiency raises both the dispersion of new-product revenue and its within-firm persistence, as high-efficiency types on average generate larger improvements than other firms. Its magnitude depends on the gap between $\lambda_{b,L}$ and $\lambda_{b,H}$ (estimated at 1.014 and 1.098, respectively), and on the equilibrium type composition, governed by G_L (89.5%). See Appendix A.3 for the details of the calibration.

Own-Product Innovation The quality improvements due to own-product innovation $\overline{\lambda}_o$ and their rate of arrival x_o play an interchangeable role in the model. We set the Poisson arrival rate x_o to 1, such that firms on average improve the quality of their existing products once per year. We then calibrate $\overline{\lambda}_o$ to 0.061 to match the average rate 0.5% of growth in real revenue on firms' continuing products, as found in the upper panel of Table 3.[27] The degree of decreasing returns to follow-up innovation, β, is calibrated to 0.783, to match expected annual decline in the growth of revenue with tenure in Figure 8.

Other Parameters The idiosyncratic product-level quality shocks cause changes in sales unrelated to innovation. Greater volatility in these shocks raises the contribution of continuing products to the overall variance of firm revenue growth. We assume a mean-zero, log-normal distribution for quality shocks with standard deviation Ξ, set at 0.35 to match the 26% Shapley value of continuing products in revenue growth (Table 3). The discount rate ρ is set to 0.02. Finally, we set the elasticity of substitution ϵ to 4 to deliver a markup of 1.33, in line with evidence for French manufacturing (De Ridder et al. 2026).

Parametrization of the Benchmark Model We compare the performance of the model with innovation bursts to a benchmark model where product creation is a Poisson process, which is the assumption underlying standard models such as Klette and Kortum (2004), Akcigit and Kerr (2018), Acemoglu et al. (2018) and Garcia-Macia et al. (2019). To implement this, we set θ to a large number and adjust the innovation costs such that firms still add an average of 0.070 products per initial product. All other parameters are unchanged.

5.2. Firm-Size Distribution in the Quantified Model

We begin the assessment of the quantified model by examining the predictions of Proposition 2 on the distribution of product counts. The proposition shows that innovation bursts are vital for explaining the observed concentration of production among large firms. Figure 9 confirms this by comparing the distribution of product count under the model with bursts

[27] In principle, population growth can also drive growth in revenue on continuing products. We omit this from the model because working-age population grew at -0.01% per year on average over our sample.

Figure 9. Firm Size Distribution

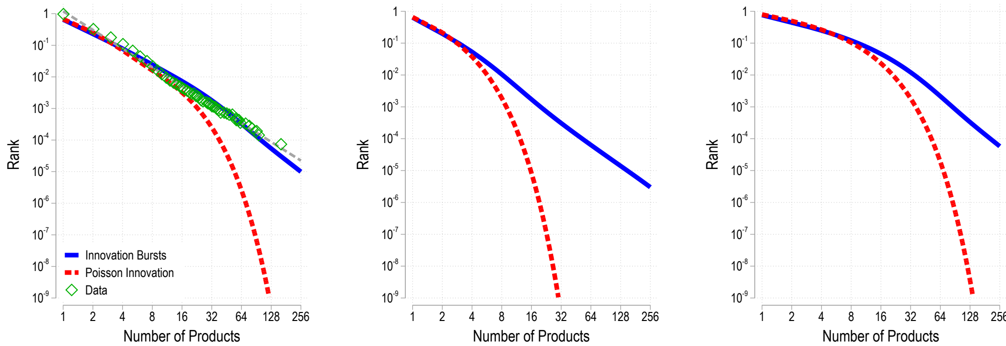

(a) Full distribution (b) Firms with low $\overline{\lambda}_{b,h}$ (c) Firms with high $\overline{\lambda}_{b,h}$

Notes: The figure plots the firm-size distribution in the log-log space. The vertical axis measures firms' rank divided by the number of observations. The horizontal axis measures the number of products that firms produce. Data is from the EAP.

($\theta = 3.06$, the solid blue lines) to the case where the same rate of product creation arises under a Poisson process ($\theta \to \infty$, dashed red lines), and the data (green diamonds).

The shape of the firm-size distribution closely resembles that in the data. Consistent with Proposition 2, the model only replicates the Pareto-like tail of product count when innovation comes in bursts. Instead, in the benchmark model of one product per innovation, product counts follow a thin-tailed *logarithmic* distribution—a familiar shortcoming of standard models of creative destruction (see, e.g., Akcigit and Kerr 2018, Luttmer 2010).

Bursts thus offer an explanation for the emergence of large firms through product innovations. Luttmer (2011), alternatively, shows that creative-destruction models can generate Pareto-tailed firm-size distributions when the total number of products grows over time (see also Cao et al. 2017 and Peters and Walsh 2021). As noted above, however, we do not observe such growth in our data. More directly relevant is Luttmer's observation that large firms also emerge if firms experience occasional episodes of rapid growth—a mechanism sometimes referred to as the "Luttmer rocket" (see also Jones and Kim 2018). Our evidence suggests that innovation bursts may be the empirical counterpart of such episodes.

Figures 9b and c separately plot the size distributions of the two innovation types. Firms with lower average quality improvements choose lower innovation rates and therefore tend to remain smaller than firms whose innovations generate larger quality improvements. The tails indices of the two distributions nevertheless coincide (as established in Proposition 2), because a single large innovation burst can propel even a low-innovation-rate firm into the tail of the firm-size distribution. With Poisson innovation, by contrast, firms become large only from innovating at persistently above-average rates.

We should note that the original Klette and Kortum (2004) model assumes that firms with higher average quality improvements also face higher innovation costs, scaled so that all firms innovate at the same expected frequency and thus have identical size distributions.

Figure 10. Firm Size Distribution with Klette and Kortum (2004)'s Firm Types

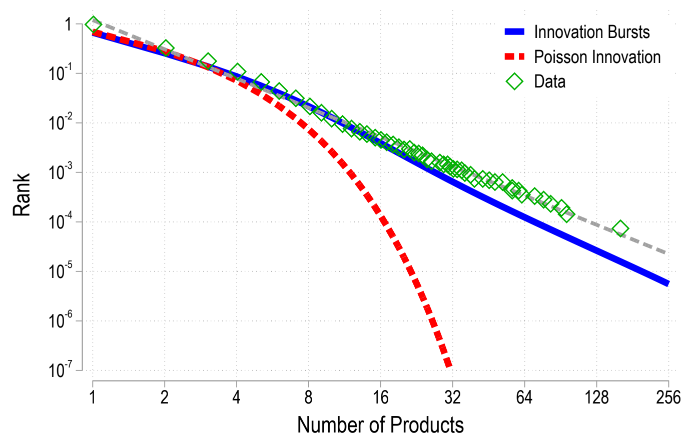

Notes: Firm-size distribution for a recalibrated model where firm types have equal average innovation rates. Vertical axis: firms' rank divided by the number of observations (log). Horizontal axis measures the number of products that firms produce (log).

Figure 10 presents the associated firm-size distribution, both in the model with innovation bursts and the benchmark model with Poisson innovation. Under innovation bursts, the firm-size distribution remains reasonably well matched, while firms are exceedingly unlikely to have a significant number of products under Poisson innovation.

5.3. Model Performance on Other Untargeted Moments

We next assess the model's ability to match a series of additional moments, including those presented in Section 3. Unlike the firm-size results in Section 5.2, generating the moments needed for this comparison requires us to simulate the model. We draw 100 Monte Carlo simulations of 22,000 firms over 10 years, matching the structure of the micro data.

Product Creation Figure 11 plots the distribution of product creation in the data and the model. The empirical distribution is plotted with green diamonds, while simulations from the model with innovation bursts are plotted with blue circles. To understand the importance of the fact that product innovations come in bursts, the figures also plot the distribution of product creation for the benchmark model where product creation is a Poisson process.

The figure shows that the model accurately matches the high concentration of product creation in the data. Pooling firms of all sizes in the left-hand figure, the linear relationship between the level and the rank of the number of new products fits almost perfectly with an R^2 of 0.99. Conditioning on size, the model with innovation bursts again closely matches the empirical distributions. The benchmark model in which the instances of product creation do not come in bursts, in contrast, fits poorly. Single-product firms in the right-hand figure

Figure 11. Distribution of Number of New Products: Data and Model

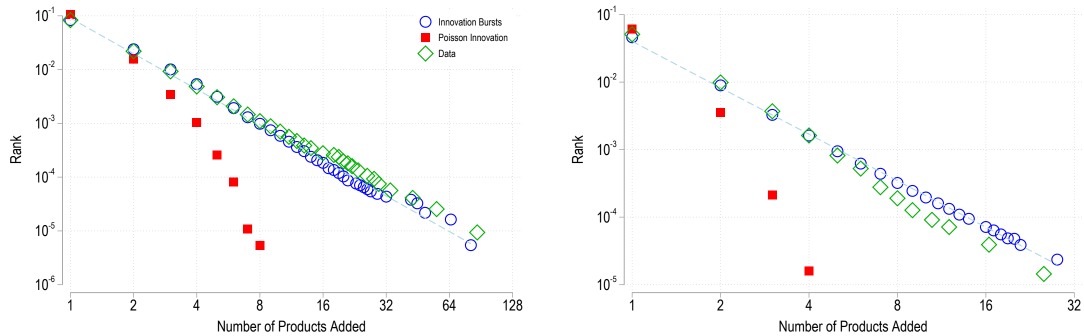

(a) All Firms (b) Firms with $n_{i,t-1} = 1$

Notes: Figures plot the distribution of the number of products firms add to their portfolio. Green diamonds are data. Blue circles are from a model with innovation bursts. Red squares are based on a model with Poisson product innovation. See Online Appendix Figure C13 for figures that include the distribution for firms that initially produce 5 or 10 products.

see negligible levels of product creation each year in the benchmark model, while such firms can add over 20 products in a single year in the data.

Tail Parameters Table 7 details the model's performance on untargeted moments such as the tail parameters of the three components of observable firm heterogeneity: revenue, revenue per product, and product count. Tail indices are obtained using the estimator of Gabaix and Ibragimov (2011) above the 90th percentile of the distribution. Quantitatively, the model closely matches the tail index of revenue. Qualitatively, the model matches revenue concentration as driven by revenue per product and product count: the tail parameter of product count is the highest, followed by similar tails for revenue per product and revenue, although the model somewhat overestimates concentration in revenue per product and underestimates concentration in product count. Proposition 2 implies that the tail index of product count as $n \to \infty$ should equal $\theta - 1$, that is 2.06.

Table 7: Untargeted: Tails of Revenue, Revenue per Product, and Product Count

	Revenue		Revenue per Product		Number of Products	
	Data	Model	Data	Model	Data	Model
Tail Parameter	1.12	1.15	1.24	1.14	1.97	2.25
	(.006)	(.011)	(.010)	(.010)	(.044)	(.111)
R^2	0.99	0.98	0.99	0.99	0.99	0.95

Notes: The tail parameter is estimated on observations above the 90th percentile of the variable's distribution. Standard errors are in parentheses. The estimates are based on the OLS regression approach proposed by Gabaix and Ibragimov (2011) and involves calculating the rank of each observation for the columns' variable, and then running a regression of the log of rank minus 0.5 on the log of the variable. Model estimates, standard errors, and R^2s are averages across 100 Monte Carlo simulations.

Table 8: Untargeted: Odds that Firm adds Majority of 10-Year New Products in Single Year

	Any 2010 Product Count		Product Count of 1 in 2010	
Expansion over 2010-2019:	Data	Model	Data	Model
Added >5 times original product count	0.48	0.65 (0.03)	0.49	0.64 (0.04)
Added >10 times original product count	0.67	0.78 (0.05)	0.69	0.77 (0.06)
Added >20 times original product count	0.75	0.87 (0.12)	0.82	0.87 (0.12)

Notes: The table plots the probability that firms add more than half of their new products in a single year. Sample: firms continuously in the EAP data between 2010 and 2019. Rows differ in the number of products that firms have added. Model estimates: avg. across 100 Monte Carlo simulations. Standard devs. of the average across Monte Carlo draws are given in parentheses.

Age-Size Relationship Luttmer (2011) emphasizes that episodes of rapid growth are necessary for firm-dynamics models to match the average age of large firms. Validating the model's performance on the age-size relationship is complicated by the fact that, for firms founded before 1950, birth years recorded in the French administrative data (FARE) contain substantial measurement error. We therefore exclude firms aged 70 or older when computing these moments, and impose the same restriction in the model. We find that the average age of large firms in our quantified model is comparable with the data. The average age of firms with more than 5, 10, or 20 products is 34.5, 35.8, and 38.4 years in the data, respectively, compared with 32.0, 35.3, and 37.9 years in the simulations.

Contribution of Bursts to Rapid Growth We next show that the model replicates the contribution of innovation bursts to episodes of high firm growth. We ask what percentage of firms, among those that grew rapidly, i.e., adding 5, 10, or 20 times their original number of products over the 10 years in our data, create the *majority* of these new products in a single year. Table 8 confirms that the model replicates the average importance of a single burst for high-growth firms (*vis à vis* frequent incremental product creation). This offers a further confirmation that innovation bursts are a key element of product dynamism.

Tenure Profile Figure 12 presents the life cycle of revenue growth in the model and in the data. Recall that we have targeted the decline in revenue growth in the data between the first and the second years, while the remaining data points are untargeted. The figure shows that the model's decline in innovation efficiency over the product life cycle delivers a path of revenue growth beyond the second year that is in agreement with the data.

5.4. Empirical Validity of R&D Predictions

Section 4 also yields predictions for firms' R&D spending, which are tested in Online Appendix D. Many of these overlap with the stylized facts that originally motivated the Klette and Kortum (2004) framework. We focus on predictions that distinguish our model or require product-level (rather than firm-level) data. For instance, Online Appendix D shows that R&D spending scales linearly in firm product count, and also provides evidence supporting our modeling assumption that R&D spending does not determine the burst size.

Figure 12. Life Cycle of Revenue Growth

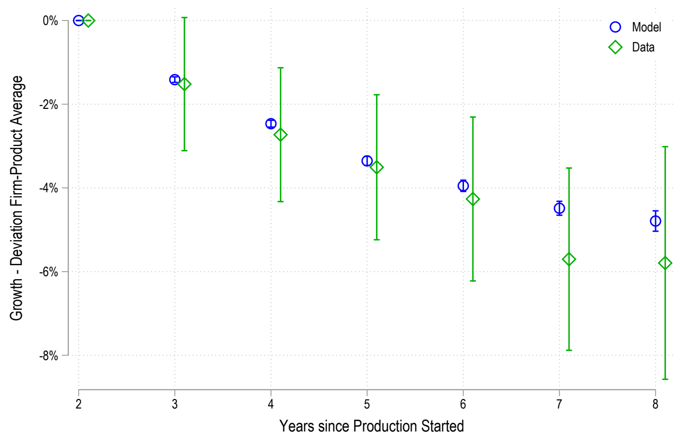

Notes: The vertical axis plots average growth of product revenue among all firm-products in the sample, where the avg. growth rate of a firm-product's revenue is subtracted from growth at each horizon. The horizontal axis plots the number of years that have passed since the firm first started producing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

5.5. Applications

Finally, we illustrate the implications of innovation bursts in our model in two examples.

Size–Innovation Efficiency Relationship Figure 13 plots the share of high-innovation-efficiency firms at each product count for different values of θ, with innovation costs recalibrated to keep identical rates of aggregate product creation across all values of θ, as explained in Appendix A.2.4.

Without innovation bursts (as in Klette and Kortum 2004, red dashed line), large firms are almost exclusively high-innovation-efficiency types: the only way to grow large is through the creation of many successive products and high-efficiency firms create products at a higher rate. This pattern changes with innovation bursts. A single large burst allows even a low-innovation-efficiency firm to grow large, weakening the link between innovation efficiency and size.

For extreme values of firm size, the figure shows that the type composition converges to the same limit for all finite θ, as predicted by Proposition 2. The asymptotic mass for each type only depends on its entry share and effective rate of product creation. The intuition is that firms appear this far in the tail only when they draw an enormous innovation burst, and it is the effective product creation rate that governs the probability of such events.

However, the histogram shows that this tail behavior occurs far beyond the range of product counts observed in our data. Within the empirically relevant range, stronger innovation bursts make firm size a less informative signal of innovation efficiency. This matters for innovation policies that target firms based on size. In a setting with innovation types similar to ours, Lentz and Mortensen (2016) show that reallocating innovation resources toward high-type firms is optimal because their innovations generate larger positive spillovers.

Figure 13. Effect of Innovation Bursts on Size–Innovation Efficiency Relationship

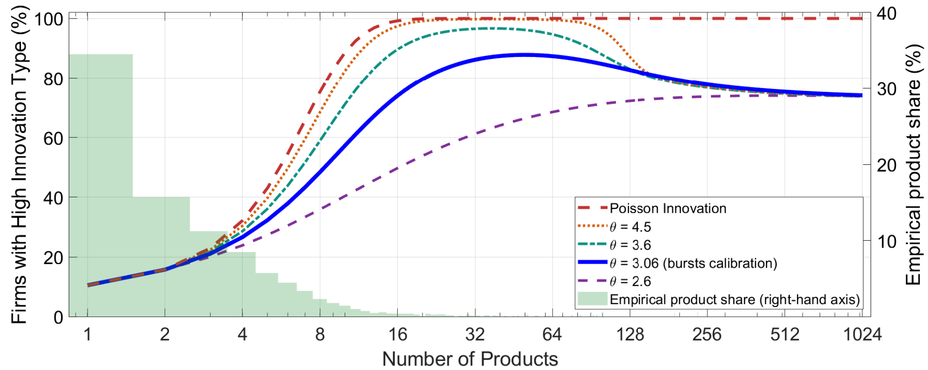

Notes: Lines plot the probability that a firm is of the high innovation efficiency type by product count n at various parametrizations of θ. The background histogram illustrates the empirical distribution of firm product count n.

Since types are likely unobservable, firm size would be a useful proxy under the benchmark Poisson model of product creation. Innovation bursts weaken the information content of this proxy, cautioning against such size-dependent policies.[28]

Contribution of Large Firms to Growth Figure 14 presents a second application of our framework, by quantifying the contribution of large firms to aggregate productivity growth. To do so, it plots what share of overall productivity growth from product creation originates from firms with at least n products, for various calibrations of θ (again, holding overall product creation constant).

The figure shows that in a model with Poisson innovation, such as Klette and Kortum (2004), large firms play a negligible role in aggregate growth. Firms with at least 16 products, for example, are responsible for 13% of growth. Under our calibration for innovation bursts, they contribute 29% of growth. The difference is striking given the fact that larger firms are less likely to have higher innovation efficiencies under innovation bursts, and thus contribute less to growth for every successful innovation. The driver is the shift in the stationary distribution of firm size: there are simply too few firms in the far right tail to deliver a sizable contribution to growth. Although large firms have higher product-creation rates under Poisson innovation than under bursts, Poisson innovation leaves too few firms in the far right tail for them to make a sizable contribution.

[28] Revenue is an even noisier proxy than product count: in the quantified model, the correlation between a firm's high-innovation-efficiency type and log revenue is 0.18, compared with 0.29 for product count.

Figure 14. Contribution of Large Firms to Aggregate Productivity Growth

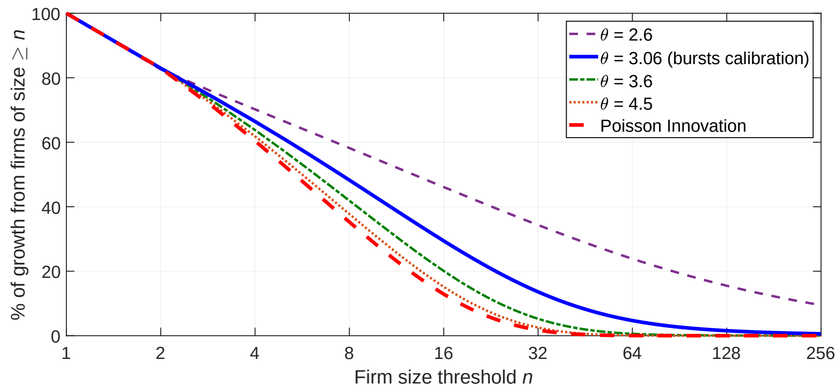

Notes: Lines plot the share of aggregate productivity growth from product creation accounted for by firms with at least n products, holding aggregate product creation fixed across values of θ.

6. Conclusion

This paper contributes to the growing literature on firm-level innovation and its aggregate consequences. Using data from the French manufacturing sector, we document a novel fact about product innovation: the distribution of firm product creation exhibits a thick, Pareto-like tail. This empirical pattern is difficult to reconcile with standard theories of creative destruction.

We propose a model of product creation through innovation bursts to address this gap. Innovations can arrive in clusters, enabling firms to rapidly expand their product portfolios. The model offers a new explanation for the observed Pareto distribution of firm size, which the standard theories of creative destruction often struggle to match.

These findings open several avenues for future research, including studies on whether product creation occurs in bursts outside of manufacturing or in countries other than France. Further research could explore the normative implications of innovation bursts, or their effect on transitional dynamics in economies with firm dynamics and creative destruction.

References

ACEMOGLU, D., U. AKCIGIT, H. ALP, N. BLOOM, AND W. KERR (2018): "Innovation, reallocation, and growth," *American Economic Review*, 108 (11), 3450–3491.

ACEMOGLU, D., U. AKCIGIT, AND M. A. CELIK (2022): "Radical and incremental innovation: The roles of firms, managers, and innovators," *American Economic Journal: Macroeconomics*, 14 (3), 199–249.

AGHION, P., U. AKCIGIT, AND P. HOWITT (2014): "What do we learn from Schumpeterian growth theory?" in *Handbook of economic growth*, Elsevier, vol. 2, 515–563.

AGHION, P., A. BERGEAUD, T. BOPPART, P. KLENOW, AND H. LI (2025): "Good rents versus bad rents: R&D misallocation and growth," Tech. rep., NBER Working Paper 34190.

AGHION, P., A. BERGEAUD, T. BOPPART, P. J. KLENOW, AND H. LI (2023): "A Theory of Falling Growth and Rising Rents," *Review of Economic Studies*.

AGHION, P. AND P. HOWITT (1992): "A Model of Growth through Creative Destruction," *Econometrica*, 60 (2), 323–351.

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

——— (2023): "What happened to US business dynamism?" *Journal of Political Economy*, 131 (8), 2059–2124.

AKCIGIT, U. AND W. R. KERR (2018): "Growth through heterogeneous innovations," *Journal of Political Economy*, 126 (4), 1374–1443.

ALBORNOZ, F., H. F. C. PARDO, G. CORCOS, AND E. ORNELAS (2023): "Sequentially exporting products across countries," *Journal of International Economics*, 142, 103735.

ARGENTE, D., S. BASLANDZE, D. HANLEY, AND S. MOREIRA (2020): "Patents to products: Product innovation and firm dynamics," .

ARGENTE, D., D. FITZGERALD, S. MOREIRA, AND A. PRIOLO (2025): "How do entrants build market share? The role of demand frictions," *American Economic Review: Insights*, 7 (2), 160–176.

ARGENTE, D., M. LEE, AND S. MOREIRA (2024): "The life cycle of products: Evidence and implications," *Journal of Political Economy*, 132 (2), 000–000.

AUTOR, D., D. DORN, L. F. KATZ, C. PATTERSON, AND J. VAN REENEN (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.

BEHRENS, K. AND J. MARTIN (2015): "Concording large datasets over time: The C3 method." *Unpublished paper.*

BENGURIA, F., H. MATSUMOTO, AND F. SAFFIE (2022): "Productivity and trade dynamics in sudden stops," *Journal of International Economics*, 139, 103631.

BERGEAUD, A., A. GUILLOUZOUIC, E. HENRY, AND C. MALGOUYRES (2025): "From Public Labs to Private Firms: Magnitude and Channels of Local R&D Spillovers*," *The Quarterly Journal of Economics*.

BERNARD, A. B., S. J. REDDING, AND P. K. SCHOTT (2010): "Multiple-product firms and product switching," *American economic review*, 100 (1), 70–97.

BLOOM, N., R. GRIFFITH, AND J. VAN REENEN (2002): "Do R&D tax credits work? Evidence from a panel of countries 1979–1997," *Journal of Public Economics*, 85 (1), 1–31.

BLOOM, N., C. I. JONES, J. VAN REENEN, AND M. WEBB (2020): "Are ideas getting harder to find?" *American Economic Review*, 110 (4), 1104–44.

BRODA, C. AND D. E. WEINSTEIN (2010): "Product creation and destruction: Evidence and price implications," *American Economic Review*, 100 (3), 691–723.

CAO, D., H. R. HYATT, T. MUKOYAMA, AND E. SAGER (2017): "Firm growth through new establishments," *SSRN 3361451*.

CAVENAILE, L., M. A. CELIK, J. PERLA, AND P. ROLDAN-BLANCO (2023): "A theory of dynamic product awareness and targeted advertising," Tech. rep., Working Paper.

CAVENAILE, L., M. A. CELIK, AND X. TIAN (2020): "Are Markups Too High? Competition, Strategic Innovation, and Industry Dynamics," *Working Paper*.

DAVIS, S. J., J. HALTIWANGER, R. JARMIN, J. MIRANDA, C. FOOTE, AND E. NAGYPAL (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.

DE RIDDER, MAARTEN, BASILE GRASSI, AND GIOVANNI MORZENTI (2026): "The hitchhiker's guide to markup estimation: assessing estimates from financial data," *Econometrica*, 94 (1), 137–168.

DEKLE, R., A. KAWAKAMI, N. KIYOTAKI, AND T. MIYAGAWA (2015): "Product Dynamics and Aggregate Shocks: Evidence from Japanese product and firm level data," *USC-INET Research Paper*, (16-01).

EATON, J., S. KORTUM, AND F. KRAMARZ (2011): "An Anatomy of International Trade: Evidence from French Firms," *Econometrica*, 79 (5), 1453–1498.

EINAV, L., P. J. KLENOW, J. D. LEVIN, AND R. MURCIANO-GOROFF (2021): "Customers and retail growth," Tech. rep., National Bureau of Economic Research.

FITZGERALD, D., S. HALLER, AND Y. YEDID-LEVI (2024): "How exporters grow," *Review of Economic Studies*, 91 (4), 2276–2306.

FLAJOLET, P. AND R. SEDGEWICK (2009): *Analytic Combinatorics*, Cambridge: Cambridge University Press.

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

GABAIX, X. AND R. IBRAGIMOV (2011): “Rank- 1/2: a simple way to improve the OLS estimation of tail exponents,” *Journal of Business & Economic Statistics*, 29 (1), 24–39.

GARCIA-MACIA, D., C.-T. HSIEH, AND P. J. KLENOW (2019): “How destructive is innovation?” *Econometrica*, 87 (5), 1507–1541.

GARICANO, L., C. LELARGE, AND J. VAN REENEN (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 E. HELPMAN (1991): “Quality ladders in the theory of growth,” *The review of economic studies*, 58 (1), 43–61.

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

HOTTMAN, C. J., S. J. REDDING, AND D. E. WEINSTEIN (2016): “Quantifying the sources of firm heterogeneity,” *The Quarterly Journal of Economics*, 131 (3), 1291–1364.

IGNASZAK, M. AND P. SEDLÁCEK (2026): “Customer acquisition, business dynamism and aggregate growth,” Tech. rep., Working paper.

JAIMOVICH, N., S. J. TERRY, AND N. VINCENT (2023): “The empirical distribution of firm dynamics and its macro implications,” Tech. rep., National Bureau of Economic Research.

JONES, CHARLES I AND JIHEE KIM (2018): “A Schumpeterian model of top income inequality,” *Journal of political Economy*, 126 (5), 1785–1826.

KLETTE, T. J. AND S. KORTUM (2004): “Innovating firms and aggregate innovation,” *Journal of political economy*, 112 (5), 986–1018.

KWON, S. Y., Y. MA, AND K. ZIMMERMANN (2024): “100 years of rising corporate concentration,” *American Economic Review*, 114 (7), 2111–2140.

LENTZ, R. AND D. T. MORTENSEN (2008): “An empirical model of growth through product innovation,” *Econometrica*, 76 (6), 1317–1373.

——— (2016): “Optimal growth through product innovation,” *Review of Economic Dynamics*, 19, 4–19.

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

——— (2010): "Models of growth and firm heterogeneity," *Annual Review of Economics*, 2 (1), 547–576.

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

MA, Y., M. ZHANG, AND K. ZIMMERMANN (2024): "Business Concentration around the World: 1900–2020," .

MELCANGI, D. AND S. SARPIETRO (2024): "Nonlinear firm dynamics," Tech. rep., Staff Report.

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

NIST (2024): "Polylogarithms," In *Digital Library of Mathematical Functions*, Chapter 25, Section 25.12. https://dlmf.nist.gov/25.12.

OLMSTEAD-RUMSEY, J. (2022): "Market Concentration and the Productivity Slowdown," .

OZKAN, S., J. HUBMER, S. SALGADO, AND E. HALVORSEN (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.

PETERS, M. AND C. WALSH (2021): "Population Growth and Firm Dynamics," *Working Paper*.

SCHWEITZER, C. (2019): "The R&D Survey: Measuring the R&D Effort of Enterprises, Beyond the Research Tax Credit," Tech. Rep. Insee Références, December 3, 2019, INSEE, accessed: 2025-08-29.

SILVA, J. S. AND S. TENREYRO (2006): "The log of gravity," *The Review of Economics and statistics*, 641–658.

WEISS, J. (2023): "Market Concentration, Growth, and Acquisitions," .

WOOD, D. C. (1992): "The Computation of Polylogarithms," Tech. Rep. 15-92, University of Kent Computing Laboratory.

“Creative Destruction through Innovation Bursts” Online Appendix

Appendix A. Theoretical Appendix

A.1. Additional Theoretical Results

A.1.1. Firm Value and Innovation Decisions

Let $V_{h,t}(\widehat{q}_i)$ denote the value function for a firm with innovation efficiency h and with a product portfolio $\widehat{q}_i$. The portfolio is a matrix with n_i rows, where each row has information on a product for which the firm has the highest quality product: the relative quality $\widehat{q}_{ij}$, which determines the profit, and the number of prior own-product innovations s_{ij}, which determines the efficacy of its own-product innovation activity. Along the balanced growth path, aggregate quality Q grows at rate g and firm value satisfies:

$$r_t V_{h,t}(\widehat{q}_i) - \dot{V}_{h,t}(\widehat{q}_i) = \tag{20}$$
$$\max_{x_{b,i},\{x_{o,ij}\}} \left\{ \begin{array}{c} \frac{Y_t}{\epsilon} \sum_{j\in J_i} \widehat{q}_{ij}^{\epsilon-1} - g_t \sum_{j\in J_i} \widehat{q}_{ij} \frac{\partial V_{h,t}(\widehat{q}_i)}{\partial \widehat{q}_{ij}} \\ + \sum_{j\in J_i} x_{d,t} \left[V_{h,t}\left(\widehat{q}_i \setminus \left\{\widehat{q}_{ij},\, s_{ij}\right\}\right) - V_{h,t}(\widehat{q}_i)\right] \\ + \sum_{j\in J_i} x_{o,ij}\, \mathbb{E}\left[V_{h,t}\left(\widehat{q}_i \setminus \{\widehat{q}_{ij}, s_{ij}\} \cup_+ \left\{\lambda_{o,j}(s_{ij})\,\widehat{q}_{ij},\, s_{ij}+1\right\}\right) - V_{h,t}(\widehat{q}_i)\right] \\ + x_{b,i} \sum_{k=1}^{\infty} \frac{k^{-\theta}}{\zeta(\theta)} \mathbb{E}_h \left[V_{h,t}\left(\widehat{q}_i \cup_+ \{\lambda_{b,i\ell}\,\widehat{q}_{i\ell}, 0\}_{\ell=1}^{k}\right) - V_{h,t}(\widehat{q}_i)\right] \\ - w_t\, \eta_b\, x_{b,i}^{\psi}\, n_i^{-\sigma(\psi-1)} - w_t\, \eta_o \sum_{j\in J_i} x_{o,ij}^{\psi}\, c_o(\widehat{q}_{ij}, s_{ij}) - F(\widehat{q}_i) \end{array} \right\},$$

where r_t and g_t stand for the interest rate and the rate of productivity growth and where $\dot{V}_{h,t}$ denotes the time derivative of firm value $V_{h,t}$, all at time t.

The first line on the right-hand side contains the sum of the flow of profits, $\frac{Y_t}{\epsilon}\widehat{q}_{ij}^{\epsilon-1}$, and the decline in profits over time from the gradual increase in average quality. The second line contains the expected change in value if the firm stops producing *j* because of creative destruction. The term $V_{h,t}\left(\widehat{q}_i \setminus \left\{\widehat{q}_{ij}, s_{ij}\right\}\right)$ denotes the value of producing the portfolio of products $\widehat{q}_i$ except product *j* with relative quality $\widehat{q}_{ij}$, on which the firm has implemented s_{ij} own-product improvements. The third line contains change in value from own-product innovation. The fourth row contains the expected increase in value due to innovation bursts. This is equal to the arrival rate of innovation bursts, $x_{b,i}$, multiplied by the expected increase in the firm's value if it acquires an innovation burst. The term $V_{h,t}\left(\widehat{q}_i \cup_+ \{\lambda_{b,i\ell}\,\widehat{q}_{i\ell}, 0\}_{\ell=1}^{k}\right)$ denotes the rise in value if the burst contains k new products, with the ℓ-th product having an initial relative quality $\widehat{q}_{i\ell}$. This value is weighted by the probability density function of *k*, which is $k^{-\theta}/\zeta(\theta)$. The first and second terms in the last row contain the firm's total R&D expenditure. The final term $F(\cdot)$ is a fixed cost that firms must pay to operate, which we assume exactly equals the option value of own-product innovation, to simplify exposition.

A.1.2. Evolution of the Measures of Firm Size

Let $x_{b,ht}(n)$ denote the optimal innovation burst rate $x_{b,i}$ for firms with n products and with innovation efficiency h at time t, solving the problem in Equation (21), and $M_{ht}(n)$ stand for the measure of such firms. Then, the change in the measure of single-product firms over time is given by

$$\dot{M}_{ht}(1) = G(h)\, x_{e,t} + 2M_{ht}(2)\, x_{d,t} - M_{ht}(1)\, \left(x_{b,ht}(1) + x_{d,t}\right), \tag{21}$$

where $x_{e,t}$ and $x_{d,t}$ denote the rates of entry and creative destruction at time t. The first term on the right-hand side captures new h-type firms from entry, the second term captures the inflow from firms of type h that used to produce two products but that have lost one due to creative destruction, and the final term captures outflow either through creatively destroyed firms exiting or through expansion by firms that draw an innovation burst.

For multi-product firms, the equation is

$$\dot{M}_{ht}(n) = \sum_{k=1}^{n-1} M_{ht}(k)\, \frac{(n-k)^{-\theta}}{\zeta(\theta)}\, x_{b,ht}(k) + (n+1)\, M_{ht}(n+1)\, x_{d,t} - M_{ht}(n)\, \left(x_{b,ht}(n) + n\, x_{d,t}\right).$$

The first term on the right-hand side shows how bursts alter the composition of firms. In a world without innovation bursts, i.e. when product creation is a Poisson process ($\theta \to \infty$), only firms that produce $n-1$ products can become producers of n products. Because of larger bursts, however, some firms are able to jump from being small producers to being large firms, and the likelihood that this happens increases as θ declines.

A.1.3. Balanced Growth Path Definition

Definition 1. *The economy is in a balanced growth path equilibrium if for every t the variables $\{r, x_d, x_e, x_o, L, L^e, L^i, g\}$ and functions $\{x_{b,h}, M_h\}$ are constant, $\{Y, C, Q, w\}$ grow at a constant rate g that satisfies (20), interest rates follow from $r = g + \rho$, Q is given by Equation (4), Y is given by $Y = QL = \frac{\epsilon}{\epsilon-1} wL$, innovation rates $x_{b,h}$ and x_o satisfy (11) and (12), the entry rate satisfies the free entry condition (13) which maximizes the profits of entrants, and together they satisfy Equation (20). The stationary distributions of product types M_h satisfy Equation (16), the rate of creative destruction x_d satisfies (14), and both goods and labor markets are in equilibrium so that $Y = C$ and $L = 1 - L^i - L^e$, with $L^e = \eta_e x_e^{\psi}$ and the labor hired in incumbent R&D sectors given by*

$$L^i = \eta_b \sum_{h,n} M_h(n)\, x_{b,h}(n)^{\psi}\, n^{-\sigma(\psi-1)} + \eta_o\, x_o^{\psi}\, \mathbb{E}\left[c_o(\widehat{q}, s)\right], \tag{22}$$

A.2. Proofs and Derivations

A.2.1. Proof of Proposition 1:

The dynamic optimization problem is given by the HJB equation (21). Guess that the solution takes the form $V_{h,t}(\widehat{q}) = w_t \left[\sum_{j\in J_i} v_h(\widehat{q}_{ij}, s_{ij}) + nO_h(n)\right]$. Substituting for this expression in the HJB equation, we find

$$(r-g)\left[\sum_{j\in J_i} v_h(\widehat{q}_{ij}, s_{ij}) + n\,O_h(n)\right] = \max_{x_{b,i},\{x_{o,ij}\}}\left\{\begin{array}{c} \frac{L}{\epsilon-1}\sum_{j\in J_i}\widehat{q}_{ij}^{\epsilon-1} - g\sum_{j\in J_i}\widehat{q}_{ij}\frac{\partial v_h(\widehat{q}_{ij},s_{ij})}{\partial \widehat{q}_{ij}} \\ -\sum_{j\in J_i} x_d\left[v_h(\widehat{q}_{ij}, s_{ij}) + nO_h(n) - (n-1)O_h(n-1)\right] \\ +\sum_{j\in J_i} x_{o,ij}\,\mathbb{E}\left[v_h\left((1+\lambda\beta^{s_{ij}})\widehat{q}_{ij}, s_{ij}+1\right) - v_h(\widehat{q}_{ij}, s_{ij})\right] \\ +x_{b,i}\left(\sum_{k=1}^{\infty}\frac{k^{-\theta}}{\zeta(\theta)}\left(k\,\mathbb{E}_h\left[v_h\left(\lambda_{b,i}\widehat{q},0\right)\right] + (n+k)O_h(n+k)\right) - nO_h(n)\right) \\ -\eta_b\,x_{b,i}^{\psi}\,n_i^{-\sigma(\psi-1)} - \eta_o\sum_{j\in J_i} x_{o,ij}^{\psi}\,c_o(\widehat{q}_{ij}, s_{ij}) - F(\widehat{q}_i)\end{array}\right\}, \tag{23}$$

Now substitute $v_h\left(\widehat{q}, s\right) \equiv \widetilde{v}\,(s)\,\widehat{q}^{\epsilon-1}$ and let $c_o\left(\widehat{q}, s\right) \equiv \widetilde{c}_o\,(s)\,\widehat{q}^{\epsilon-1}$ and $F(\widehat{q}_i) \equiv \sum_{j\in J_i}\widetilde{F}\,(s)\,\widehat{q}^{\epsilon-1}$ in the above equation and collect the terms involving $\widetilde{v}\,(\cdot)$ to find

$$(r-g)\,\widetilde{v}\,(s) = \max_{x_o}\left\{\begin{array}{c}\frac{L}{\epsilon-1} - g\,(\epsilon-1)\,\widetilde{v}\,(s) - x_d\widetilde{v}\,(s) \\ x_o\,\left(\mathbb{E}_{H_o}\left[(1+\lambda_o\beta^s)^{\epsilon-1}\right]\widetilde{v}\,(s+1) - \widetilde{v}\,(s)\right) - \eta_o x_o^{\psi}\widetilde{c}_o\,(s) - \widetilde{F}\,(s)\end{array}\right\}. \tag{24}$$

Collecting the terms involving $O_h\,(\cdot)$ then yields

$$n(r-g+n\,x_d)\,O_h(n) = \max_{x_b}\left\{\begin{array}{c} x_b\,\sum_{k=1}^{\infty}\frac{k^{-\theta}}{\zeta(\theta)}\,\left((n+k)\,O_h(n+k) - n\,O_h(n) + k\,\widetilde{v}(0)\,\mathbb{E}_h\left[\lambda^{\epsilon-1}\right]\right) \\ -\eta_b\,x_b^{\psi}\,n^{-\sigma(\psi-1)} + n(n-1)\,x_d\,O_h(n-1)\end{array}\right\}. \tag{25}$$

The first order condition for x_o yields

$$x_o\,(s) = \left(\frac{1}{\psi\eta_o}\frac{\mathbb{E}_{H_o}\left[(1+\lambda\beta^s)^{\epsilon-1}\right]\widetilde{v}(s+1) - \widetilde{v}(s)}{\widetilde{c}_o(s)}\right)^{\frac{1}{\psi-1}}, \tag{26}$$

while the first order condition for x_b gives

$$x_{b,h}(n) = \left(\frac{1}{\psi\eta_b}\left(\frac{\zeta(\theta-1)}{\zeta(\theta)}\overline{\lambda}_{b,h}^{\epsilon-1}\widetilde{v}\,(0) + \sum_{k=1}^{\infty}\frac{k^{-\theta}}{\zeta(\theta)}(n+k)\,O_h(n+k) - n\,O_h(n)\right)\right)^{\frac{1}{\psi-1}} n^{\sigma}. \tag{27}$$

Noting $\rho = r-g$, from Equation (24), we find

$$\left(\rho + (\epsilon-1)\,g + x_d\right)\widetilde{v}(s) = \frac{L}{\epsilon-1} + \frac{\psi-1}{\psi}\left(\psi\eta_o\widetilde{c}_o(s)\right)^{-\frac{1}{\psi-1}}\left(\mathbb{E}_{H_o}\left[(1+\lambda\beta^s)^{\epsilon-1}\right]\widetilde{v}\,(s+1) - \widetilde{v}\,(s)\right)^{\frac{\psi}{\psi-1}} - \widetilde{F}(s).$$

Next, we make the following simplifying assumptions

$$\widetilde{c}_o(s) \equiv \mathbb{E}_{H_o}\left[(1+\lambda\beta^s)^{\epsilon-1}\right]-1,$$
$$\widetilde{F}(s) \equiv \tfrac{\psi-1}{\psi}\left(\psi\eta_o\right)^{-\frac{1}{\psi-1}}\widetilde{c}_o(s)\left(\tfrac{L/(\epsilon-1)}{\rho+(\epsilon-1)g+x_d}\right)^{\frac{\psi}{\psi-1}}.$$

Given these assumptions, Equations (24) and (26) now together imply

$$\widetilde{v}(s) = v^* \equiv \tfrac{L/(\epsilon-1)}{\rho+(\epsilon-1)g+x_d},$$

and Eqs. (27) and (12). Substituting these results in Eq. (25), we find Eq. (10).

A.2.2. Proof of Proposition 2:

We specialize to the benchmark $\sigma = 1$. When $\sigma = 1$, the factor $n^{\sigma-1}$ in the option value recursion (10) equals unity, and the per-product option value $O_h(n)$ is in fact constant in n. To see this, conjecture $O_h(n) = \overline{O}_h$ for all n. The burst-related sum inside the bracket of (10) then simplifies: $\sum_{k=1}^{\infty}\frac{k^{-\theta}}{\zeta(\theta)}(n+k)\,\overline{O}_h - n\,\overline{O}_h = \frac{\zeta(\theta-1)}{\zeta(\theta)}\,\overline{O}_h$, because the terms proportional to n cancel. The bracket reduces to $\frac{\zeta(\theta-1)}{\zeta(\theta)}\big(\overline{\lambda}_{b,h}^{\epsilon-1}\,v^*+\overline{O}_h\big)$, which is independent of n, and the recursion collapses to the fixed-point equation $(\rho+x_d)\,\overline{O}_h = \frac{(\psi-1)/\psi}{(\psi\eta_b)^{1/(\psi-1)}}\big[\frac{\zeta(\theta-1)}{\zeta(\theta)}(\overline{\lambda}_{b,h}^{\epsilon-1}\,v^* + \overline{O}_h)\big]^{\psi/(\psi-1)}$, confirming the conjecture. Substituting $O_h(n) = \overline{O}_h$ into (11) gives $x_{b,h}(n) \propto n^{\sigma} = n$, so the per-product burst rate $\overline{x}_{b,h} \equiv x_{b,h}(n)/n$ is constant in n for each type h. Substituting $x_{b,h}(n) = \overline{x}_{b,h}\,n$ into equation (16) yields the type-h balance equation with constant per-product rate $\overline{x}_{b,h}$.

Define the type-h generating function

$$\mathbb{H}_h(z) = \sum_{n=1}^{\infty} M_h(n)\,z^n, \qquad \mathbb{H}_h'(1) = m_h. \tag{28}$$

Multiplying the resulting balance equation by z^n and summing over $n \geq 1$ gives

$$D_h(z)\,\mathbb{H}_h'(z) = \frac{x_e\,G(h)}{x_d}\,(1-z), \text{ where} \tag{29}$$

$$D_h(z) = 1-\left(1+\frac{\overline{x}_{b,h}}{x_d}\right)z + \frac{\overline{x}_{b,h}}{x_d\,\zeta(\theta)}\,z\,\mathrm{Li}_\theta(z), \tag{30}$$

and $\mathrm{Li}_\theta(z) = \sum_{k=1}^{\infty} k^{-\theta}z^k$ is the polylogarithm of order θ.

When $|\mathcal{H}| = 1$, write $x_1 = x$, $x_e\,G(1) = x_e$, $m_1 = 1$ (product normalization), $M_1(n) = M(n)$, $\mathbb{H}_1 = \mathbb{H}$, and $D_1 = D$. The creative destruction identity (14) becomes $x_d = \zeta(\theta-1)\,x/\zeta(\theta)+x_e$, and the ODE reduces to $D(z)\,\mathbb{H}'(z) = (x_e/x_d)\,(1-z)$.

We first prove the result in the single-type case, where the normalization $\mathbb{H}'(1) = 1$ simplifies the algebra, and then explain how it extends to general type h. In the single-type case, (17) reduces to

$$M(n) \sim \frac{x}{x_e\,(\theta-1)\,\zeta(\theta)}\,n^{-\theta}. \tag{31}$$

The proof extracts the large-n behavior of $M(n)$ from the generating function $\mathbb{H}(z) = \sum M(n)\,z^n$. Because $M(n) \geq 0$ and $\sum n\,M(n) = 1$, the generating function converges absolutely for all $|z| \leq 1$ and can be decomposed near $z = 1$ as

$$\mathbb{H}(z) = \underbrace{\overline{M} - (1-z) + a_2(1-z)^2 + \cdots}_{\text{analytic (Taylor series in }(1-z))} + \mathbb{H}^{(\text{sing})}(z), \tag{32}$$

where $\overline{M} = \sum M(n)$ is the total firm mass, the linear coefficient equals 1 because $\mathbb{H}'(1) = \sum n\,M(n) = 1$ (product normalization), and $\mathbb{H}^{(\text{sing})}$ collects every term that is *not* an integer power of $(1-z)$. Any function analytic at $z = 1$ has Taylor coefficients that decay exponentially, so for the power-law asymptotics of $M(n)$ only the singular part matters. Writing $[z^n]\,f(z)$ for the coefficient of z^n in the power series $f(z)$:

$$M(n) = [z^n]\,\mathbb{H}^{(\text{sing})}(z) + \underbrace{[z^n]\,(\text{analytic part})}_{\text{exponentially small}}.$$

The classical transfer theorem of Flajolet and Sedgewick (2009) then converts the form of $\mathbb{H}^{(\text{sing})}$ into the asymptotics of $M(n)$.

The strategy is therefore: (i) expand $\mathrm{Li}_\theta(z)$ near $z = 1$ to find its singular part; (ii) propagate this through $D(z)$ and the ODE to obtain $\mathbb{H}^{(\text{sing})}(z)$; (iii) verify that $z = 1$ is the only singularity of $\mathbb{H}$ on the closed unit disk; (iv) apply the transfer theorem to extract $M(n) \sim C\,n^{-\theta}$. For non-integer and integer orders we have (NIST, 2024, Eq. 25.12.12) and Wood (1992, Eq. 9.5) for $|\log z| < 2\pi$:

$$\mathrm{Li}_\theta(z) = \sum_{\substack{m=0\\ m\neq\theta-1}}^{\infty} \frac{\zeta(\theta-m)}{m!}\,(\log z)^m + \begin{cases} \Gamma(1-\theta)\,\big(-\log z\big)^{\theta-1}, & \theta \notin \mathbb{Z}_+,\\[2ex] \dfrac{(\log z)^{s-1}}{(s-1)!}\,\big(H_{s-1} - \log(-\log z)\big), & \theta = s \in \mathbb{Z}_+, \end{cases} \tag{33}$$

where $H_{s-1} = \sum_{j=1}^{s-1} j^{-1}$. As $z \to 1^-$, $-\log z = (1-z) + O((1-z)^2)$, so the $m = 0$ and $m = 1$ terms give the common analytic contributions $\zeta(\theta)$ and $-\zeta(\theta-1)(1-z)$. The leading non-analytic term is:

$$\mathrm{Li}_\theta(z) = \zeta(\theta) - \zeta(\theta-1)\,(1-z) + O\big((1-z)^2\big) + \begin{cases} \Gamma(1-\theta)\,(1-z)^{\theta-1}, & \theta \notin \mathbb{Z}_+,\\[2ex] \kappa_s\,(1-z)^{s-1}\log(1-z), & \theta = s \in \mathbb{Z}_+, \end{cases} \tag{34}$$

with $\kappa_s = \frac{(-1)^s}{(s-1)!}$. In both cases the leading singular order is $\theta - 1$.

Substituting (34) into (30) with $\overline{x}_{b,h} = x$, the constant term vanishes because $D(1) = 1 - (1 + x/x_d) + (x/x_d)\ \mathrm{Li}_\theta(1)/\zeta(\theta) = 0$. Using the creative destruction identity (14) for the linear coefficient:

$$D(z) = \frac{x_e}{x_d}\,(1-z) + O\big((1-z)^2\big) + \begin{cases} \frac{x\,\Gamma(1-\theta)}{x_d\,\zeta(\theta)}\,(1-z)^{\theta-1}, & \theta \notin \mathbb{Z}_+, \\ \frac{x\,\kappa_s}{x_d\,\zeta(s)}\,(1-z)^{s-1}\log(1-z), & \theta = s \in \mathbb{Z}_+. \end{cases} \tag{35}$$

Dividing the ODE by (35) gives

$$\mathbb{H}'(z) = 1 + O(1-z) + \begin{cases} -\frac{x\,\Gamma(1-\theta)}{x_e\,\zeta(\theta)}\,(1-z)^{\theta-2}, & \theta \notin \mathbb{Z}_+, \\ -\frac{x\,\kappa_s}{x_e\,\zeta(s)}\,(1-z)^{s-2}\log(1-z), & \theta = s \in \mathbb{Z}_+. \end{cases}$$

Integrating from z to 1 separates $\mathbb{H}(z)$ into analytic terms and a singular part:

$$\mathbb{H}^{(\mathrm{sing})}(z) = \begin{cases} \dfrac{x\,\Gamma(1-\theta)}{x_e\,(\theta-1)\,\zeta(\theta)}\,(1-z)^{\theta-1}, & \theta \notin \mathbb{Z}_+, \\ \dfrac{x\,\kappa_s}{(s-1)\,x_e\,\zeta(s)}\,(1-z)^{s-1}\log(1-z), & \theta = s \in \mathbb{Z}_+. \end{cases} \tag{36}$$

Next, we verify that $z = 1$ is the only singularity on the closed unit disk. Because $M(n) \geq 0$, the series $\mathbb{H}(z) = \sum M(n)\,z^n$ is bounded by $\sum M(n) = \overline{M} < \infty$ at every point with $|z| \leq 1$. If $D(z_0) = 0$ at any $z_0 \neq 1$ with $|z_0| \leq 1$, the ODE would force $\mathbb{H}'(z_0) = \infty$ (since $(x_e/x_d)(1-z_0) \neq 0$), contradicting the finiteness of $\mathbb{H}$. Therefore $D(z) \neq 0$ on $\{|z| \leq 1\} \setminus \{1\}$, and $\mathbb{H}$ is analytic on the entire closed unit disk except at $z = 1$.

The transfer theorem of Flajolet and Sedgewick (2009, Chapter VI) converts the singular part into coefficient asymptotics:

$$[z^n](1-z)^\alpha \sim \frac{n^{-\alpha-1}}{\Gamma(-\alpha)}\ \ (\alpha \notin \mathbb{Z}_{\geq 0}), \qquad [z^n](1-z)^m\log(1-z) \sim (-1)^{m+1}\,m!\,n^{-m-1}\ \ (m \in \mathbb{Z}_{\geq 0}). \tag{37}$$

Applying these to (36):

$$M(n) \sim \begin{cases} \frac{x\,\Gamma(1-\theta)}{x_e\,(\theta-1)\,\zeta(\theta)} \cdot \frac{n^{-\theta}}{\Gamma(1-\theta)} = \frac{x}{x_e\,(\theta-1)\,\zeta(\theta)}\,n^{-\theta}, & \theta \notin \mathbb{Z}_+, \\ \frac{x\,\kappa_s}{(s-1)\,x_e\,\zeta(s)} \cdot (-1)^s\,(s-1)!\,n^{-s} = \frac{x}{x_e\,(s-1)\,\zeta(s)}\,n^{-s}, & \theta = s \in \mathbb{Z}_+, \end{cases} \tag{38}$$

where $\Gamma(1-\theta)$ cancels in the first case, and $\kappa_s = (-1)^s/(s-1)!$ cancels against $(-1)^s(s-1)!$ in the second. Both cases recover (31).

For a general type h, the polylogarithm expansion (34) is unchanged (it depends only on θ), and $D_h(z)$ has the same functional form as $D(z)$ with $\overline{x}_{b,h}$ replacing x. The only difference is that $\mathbb{H}_h'(1) = m_h$ rather than 1. This modifies the linear coefficient of D_h: the ODE constraint $D_h'(1)\,\mathbb{H}_h'(1) = -x_e\,G(h)/x_d$ gives $D_h'(1) = -x_e\,G(h)/(m_h\,x_d)$. Equivalently, if $D_h(z) = a_h(1-z) + \cdots$ near $z = 1$, then $a_h\,m_h = x_e\,G(h)/x_d$, so $a_h = x_e\,G(h)/(m_h\,x_d)$ and

$$D_h(z) = \frac{x_e\,G(h)}{m_h\,x_d}\,(1-z) + O\big((1-z)^2\big) + (\text{singular terms as in (35)}).$$

Dividing (29) by $D_h(z)$ and integrating as above produces a singular part with the prefactor $\overline{x}_{b,h} m_h^2/(x_e\,G(h))$ in place of x/x_e: one factor of m_h comes from $\mathbb{H}_h'(1) = m_h$, and the second from the denominator $x_e\,G(h)/(m_h\,x_d)$ when dividing by D_h. The regularity argument and transfer theorem apply identically, yielding (17). Summing over h gives (19). □

In the Klette–Kortum limit ($\theta \to \infty$), bursts add exactly one product ($\mathrm{Li}_\theta(z)/\zeta(\theta) \to z$), and the singularity of the generating function moves from the branch point at $z = 1$ to a simple pole at $z = x_d/x > 1$, yielding exponential rather than power-law decay.

A.2.3. Invariance of Aggregate Balanced Growth to Varying θ

Here, we show that under the market equilibrium of the benchmark case where $\sigma = 1$, the rate of balanced growth does *not* depend on the thickness of the tail of innovation bursts, θ. To see this, note that in this case where $x_{b,h}(n) = n\overline{x}_{b,h}$, we can write the growth arising from product creation by incumbents (g_b) and by entrants (g_e) as the first two terms on the right-hand side of equation (20), as

$$g_b + g_e = \tfrac{\zeta(\theta-1)}{\zeta(\theta)} \sum_h \left(m_h \frac{\overline{\lambda}_{b,h}^{\epsilon-1} - 1}{\epsilon - 1}\,\overline{x}_{b,h} \right) + \sum_h \left(G(h) \frac{\overline{\lambda}_{b,h}^{\epsilon-1} - 1}{\epsilon - 1} \right) x_e.$$

Inserting the stationary fraction of products produced by type-h firms given in (18) enables us to write the growth from creative destruction in terms of the rate of creative destruction x_d, the entry rate x_e, the entry type distribution $G(h)$, and the type-specific burst rates $\overline{x}_{b,h}$:

$$g_b + g_e = x_e \left[\sum_h \left(G(h) \frac{\overline{\lambda}_{b,h}^{\epsilon-1} - 1}{\epsilon - 1} \right) \left(1 + \frac{\frac{\zeta(\theta-1)}{\zeta(\theta)}\overline{x}_{b,h}}{x_d - \frac{\zeta(\theta-1)}{\zeta(\theta)}\,\overline{x}_{b,h}} \right) \right]. \tag{39}$$

Note that we also have $x_e = \left(\sum_h \frac{G(h)}{x_d - \frac{\zeta(\theta-1)}{\zeta(\theta)}\,\overline{x}_{b,h}} \right)^{-1}$.

A.2.4. Re-calibrating Product-Creation Costs with Varying θ

In the $\sigma = 1$ benchmark, define the effective per-product rate of product creation by $\widetilde{x}_{b,h} \equiv \frac{\zeta(\theta-1)}{\zeta(\theta)}\,\overline{x}_{b,h}$. Using the second equation in (15), the burst first-order condition can be rewritten in terms of $\widetilde{x}_{b,h}$, so the benchmark system in Equations (14), (15), (18), and (39) depends on θ and $\overline{x}_{b,h}$ only through $\widetilde{x}_{b,h}$ and the combination $\eta_b\,[\zeta(\theta)/\zeta(\theta-1)]^\psi$. The free-entry and labor-market conditions inherit the same dependence because they only involve v^*, O_h, and $\widetilde{x}_{b,h}$.

It follows that if θ is changed from θ to θ' while adjusting η_b according to

$$\eta_b(\theta') \;=\; \eta_b(\theta)\left(\frac{\zeta(\theta'-1)/\zeta(\theta')}{\zeta(\theta-1)/\zeta(\theta)}\right)^{\psi}, \tag{40}$$

the quantity $\eta_b\,[\zeta(\theta)/\zeta(\theta-1)]^{\psi}$ remains unchanged, so $\widetilde{x}_{b,h}$, O_h, v^*, m_h, x_e, x_d, and g all take the same benchmark equilibrium values.

The firm-size distribution $M_h(n)$, however, is *not* invariant under this recalibration, because the transition kernel $P_\theta(k) = k^{-\theta}/\zeta(\theta)$ in the balance equation (16) changes shape with θ. The recalibration therefore isolates the distributional role of burstiness: it changes how a fixed aggregate amount of product creation is allocated across firms of different sizes while leaving the aggregate rate of growth unchanged.

A.2.5. Asymptotic Distribution of Innovation Efficiency and Size

Proposition 2 implies that, for any finite value of θ, the asymptotic shares of each type h in the extreme upper tail converge to a limit that, under the recalibration described above, is itself invariant to θ.[29] What burstiness changes is how quickly the distribution of firm types transitions from its entry composition to its tail composition.

A.3. Calibration Moments

This appendix derives the relationship between the type-specific parameters of the model and the moments we use to calibrate them. There are three parameters to be calibrated: the average quality improvement of successful innovations by the low innovative type ($\overline{\lambda}_{b,L}$), the average improvement by the high type ($\overline{\lambda}_{b,H}$), and the fraction of entrants that is of the low type ($G(L)$) – where the latter governs the share of products that type L produces, that is m_L, in equilibrium. The moments that we use to discipline these parameters are i) a regression measuring the within-firm covariance of revenue on new products; ii) the variance of revenue on new products; iii) the growth rate of manufacturing productivity.

To see how the first moment is informative about the parameters, note that it follows from (3) that the revenue that firm i earns on product j at time t, r_{ijt}, can be written as

$$\ln r_{ijt} = (\epsilon-1)\ln q_{ijt} + c,$$

where c is a constant. For a product that a firm first produces at time t, i.e. after successful innovation, we furthermore have that $\ln q_{ijt} = \ln q_{-ijt} + \ln \lambda_{ijt}$, where q_{-ijt} is the

[29] From Proposition 2, the tail constant for type h is $C_h = \overline{x}_{b,h}\, m_h^2/[x_e\, G(h)\,(\theta-1)\,\zeta(\theta)]$. Under the recalibration that keeps $\widetilde{x}_{b,h} \equiv [\zeta(\theta-1)/\zeta(\theta)]\,\overline{x}_{b,h}$ fixed, we can write $\overline{x}_{b,h} = \widetilde{x}_{b,h}\,\zeta(\theta)/\zeta(\theta-1)$ and $m_h = x_e\, G(h)/(x_d - \widetilde{x}_{b,h})$. Substituting gives $C_h = [x_e/((\theta-1)\zeta(\theta-1))]\,\widetilde{x}_{b,h} G(h)/(x_d-\widetilde{x}_{b,h})^2$, so the ratios $C_h/\sum_\ell C_\ell$ are independent of θ because the only remaining θ-dependence is common across types.

quality at which the previous producer of good j produced the good, and where λ_{ijt} is the size of the quality improvement that firm i accomplished for good j.

The correlation of revenue on two new products, j and j', is proportional to the correlation of quality improvements λ_{ijt} and $\lambda_{ij't}$. That is because each innovation involves the draw of a random product, so that the initial qualities of the products are orthogonal:

$$\begin{aligned}\mathbb{C}\text{ov}\left(\ln r_{ijt}, \ln r_{ij't}\right) &= (\epsilon-1)^2\, \mathbb{C}\text{ov}\left(\ln \lambda_{ijt}, \ln \lambda_{ij't}\right) \\ &= (\epsilon-1)^2 \Big(\mathbb{E}_h \left[\mathbb{C}\text{ov}\left(\ln \lambda_{ijt}, \ln \lambda_{ij't} \mid h\right)\right] + \mathbb{C}\text{ov}\left(\mathbb{E}[\ln \lambda_{ijt} \mid h], \mathbb{E}[\ln \lambda_{ij't} \mid h]\right) \Big),\end{aligned}$$

where the second line follows from the law of total covariance. Since step size draws are independent across products within a type, the first term on the second line is zero. The second term captures covariance *across* types and therefore reflects heterogeneity in firms' average step sizes. Using the fact that the share of products belonging to the low-innovation-efficiency type is m_L, we can express the covariance as a function of the model parameters:

$$\begin{aligned}\mathbb{C}\text{ov}\left(\ln r_{ijt}, \ln r_{ij't}\right) &= (\epsilon-1)^2\, \mathbb{C}\text{ov}\left(\mathbb{E}[\ln \lambda_{ijt} \mid h], \mathbb{E}[\ln \lambda_{ij't} \mid h]\right) \\ &= (\epsilon-1)^2\, \mathbb{V}\text{ar}\left(\mathbb{E}[\ln \lambda_{ijt} \mid h]\right). \\ &= (\epsilon-1)^2 \left(m_L(1-m_L) \left(\frac{\overline{\lambda}_{b,L}-1}{\overline{\lambda}_{b,L}} - \frac{\overline{\lambda}_{b,H}-1}{\overline{\lambda}_{b,H}} \right)^2 \right). \end{aligned} \tag{41}$$

Dividing the covariance by the variance of log revenue on a new product yields the slope coefficient from a regression of log revenue on one new product on log revenue on another new product within the same firm. The $(\epsilon-1)^2$ terms cancel in the ratio, so that the regression coefficient does not depend on the calibration of the elasticity of substitution–whereas using the covariance would. Since the data contain revenues rather than step sizes, we include product fixed effects in the regression, and assume that these absorb variation in the initial quality level q_{-ijt} of each product (which is orthogonal to the step size in the model but would otherwise contribute to revenue variation in the data). Let $\hat{r}_{ijt}$ denote log revenue after absorbing product fixed effects. The regression coefficient $\beta = \mathbb{C}\text{ov}\left(\ln \hat{r}_{ijt}, \ln \hat{r}_{ij't}\right) / \mathbb{V}\text{ar}\left(\ln \hat{r}_{ijt}\right)$ is then solely a function of the three moments:

$$\beta = \frac{m_L(1-m_L)\left(\frac{\overline{\lambda}_{b,L}-1}{\overline{\lambda}_{b,L}} - \frac{\overline{\lambda}_{b,H}-1}{\overline{\lambda}_{b,H}}\right)^2}{m_L(1-m_L)\left(\frac{\overline{\lambda}_{b,L}-1}{\overline{\lambda}_{b,L}} - \frac{\overline{\lambda}_{b,H}-1}{\overline{\lambda}_{b,H}}\right)^2 + m_L\left(\frac{\overline{\lambda}_{b,L}-1}{\overline{\lambda}_{b,L}}\right)^2 + (1-m_L)\left(\frac{\overline{\lambda}_{b,H}-1}{\overline{\lambda}_{b,H}}\right)^2}. \tag{42}$$

The second moment we target is the relative variance of log revenue on new products, defined as the ratio of the variance to the squared mean. This moment again has the attractive property that it is robust to the calibration of the elasticity of substitution. To see this, note that we have $\mathbb{V}\text{ar}(\ln \hat{r}_{ijt}) = (\epsilon - 1)^2 \mathbb{V}\text{ar}(\ln \lambda_{ijt})$ and $\mathbb{E}[\ln \hat{r}_{ijt}] = (\epsilon - 1)\mathbb{E}[\ln \lambda_{ijt}]$, so

$$\frac{\mathbb{V}\text{ar}\left(\ln \hat{r}_{ijt}\right)}{\mathbb{E}\left(\ln \hat{r}_{ijt}\right)^2} = \frac{(\epsilon - 1)^2 \mathbb{V}\text{ar}\left(\ln \lambda_{ijt}\right)}{(\epsilon - 1)^2 \mathbb{E}\left(\ln \lambda_{ijt}\right)^2}$$

where the $(\epsilon - 1)^2$ terms cancel. The model counterpart is:

$$\frac{\mathbb{V}\text{ar}\left(\ln \hat{r}_{ijt}\right)}{\mathbb{E}\left[\ln \hat{r}_{ijt}\right]^2} = \frac{m_L(1 - m_L)\left(\ln \overline{\lambda}_{b,L} - \ln \overline{\lambda}_{b,H}\right)^2}{\left(m_L \ln \overline{\lambda}_{b,L} + (1 - m_L) \ln \overline{\lambda}_{b,H}\right)^2}. \tag{43}$$

This moment captures the dispersion in average innovation outcomes across firm types, expressed in scale-free terms as a squared coefficient of variation in log-space.

Finally, the third moment is the growth rate of manufacturing productivity. Since aggregate productivity growth, in addition to own-product innovation, is driven by the average size of quality improvements, this moment naturally pins down the mean step size $m_L \overline{\lambda}_{b,L} + (1 - m_L)\overline{\lambda}_{b,H}$.

Appendix B. Data Appendix

B.1. EAP

The data are based on an annual survey of firms' production activities, *Enquête Annuelle de Production* (EAP), administered by the Institut National de la Statistique et des Études Économiques (INSEE). In accordance with EU regulation, the survey encompasses at least 90 per cent of 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. The survey provides information at the 10-digit product level, classified according to the PRODFRA system—the official French product classification published by INSEE, which includes approximately 4,500 product codes (see Section B.1.2). The survey covers the entire manufacturing sector (NACE section C), except for agri-food (sections 10, 11, and 12) and the manufacturing of wood (16). The survey also samples the extractive industry; electricity, gas, steam, and air-conditioning supply; and water supply, sewerage, waste management, and remediation, which we exclude.

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 firms. To ensure a good level of coverage the survey must cover at least 90 per cent of the

Table B1: Examples of products in the PRODFRA classification

1812125000	Advertising and similar printed matter (excluding commercial catalogues)
1812199010	Administrative or commercial printed matter, flat or continuous, customised or not, and directories
2042194510	Lotions for pre-shaving, shaving, or after-shaving
2042194520	Pre-shaving, shaving, and after-shaving foams and gels
2511235040	Industrial boiler products: not including tanks, boilers, nuclear equipment
3102100010	Wooden kitchen furniture: by mounted elements, including custom

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 firms. Because this sample does not allow us to observe product creation or loss, we drop this second set of smaller firms from our analysis.

B.1.2. Product Classification

As for the level of product aggregation, the survey classifies products following the PRODFRA system—the official French product classification published by INSEE, which includes roughly 4,500 10-digit product codes. The first eight positions of PRODFRA represent the European PRODCOM classification, where the last two positions are used to refine the nomenclature.[30] Table B1 presents examples of products in the PRODFRA classification. A few special products, which account for around 5% 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 first four digits of the PC code identify a 4-digit NACE industry.

B.1.3. Product Concordance

A feature of PRODFRA is that it changes over time (from 3 to 5% of the 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 Behrens and Martin (2015), called "connected components concordance," or *C3* for short. *C3* uses 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.

[30] For additional information see `https://www.insee.fr/en/metadonnees/definition/c1097`.

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).

B.3. DADS

We obtain the number of R&D employees and the number of establishments from *Déclaration Annuelle de Données Sociales* (DADS), which is a matched employer-employee dataset that covers the whole population of private sector workers in France. We utilize DADS Poste, which offers data at the individual job spell level. Each worker in the dataset is associated with an establishment identifier, and if the same employee works in two different establishments during the same year, only the main job is included in our analysis.

B.4. LiFi

We determine a firm's ownership status (i.e. independent or part of a group) using the *Liaisons Financières entre Sociétés* (LiFi) data. This dataset collects information on the financial links between enterprises incorporated in France as well as their foreign owners and affiliates. LiFi serves as the French directory of corporate groups and is based on the most comprehensive knowledge of capital ownership links between companies (or financial connections). It is constructed using multiple sources: data from the *Banque de France* collected as part of corporate credit ratings, tax data (DGFIP), commercial data (ORBIS), and information available in activity reports published by corporate groups. We define independent firms as those not listed in LiFi, either as head of the group or as affiliates.

B.5. Confidentiality

When plotting the distributions of variables, such as in the scatter plots in Figure 2, we are bound by confidentiality requirements that prohibit the disclosure of results relying on fewer than three observations. In Figure 2, as well as its robustness checks, this means that a small number of data points where the number of firms adding n products is fewer than three cannot be directly reported. To comply with the confidentiality requirement, we bin two (or three) observations in those cases. The number of new products that these firms add in the plot is the geometric average of the original number of products that they added. This adjustment leaves the plot qualitatively unaltered, and the linear fit in the figure is estimated before the adjustment.

Appendix C. Additional Tables and Figures

Table C1: Relationship between Product Creation Rates and Firm Size

	Unweighted Regressions				Weighted Regressions			
Sample:	All	$N > 1$	All	$N > 1$	All	$N > 1$	All	$N > 1$
Scaling estimates	1.083	1.157	0.858	0.959	1.189	1.157	0.996	1.014
	(.029)	(.038)	(.043)	(.038)	(.061)	(.038)	(.048)	(.036)
Fixed Effects								
Industry-Year	No	No	Yes	Yes	No	No	Yes	Yes

Notes: The table presents estimates from a PPML regression (Silva and Tenreyro 2006). Columns with industry-year fixed effects control for fixed-effects at the level of four-digit NACE industries, based on the primary industry code that is assigned to the firm in FARE. Standard errors in parentheses are clustered by firm and year. Weighted regressions weigh observations by a firm's product count.

Table C2: Contribution of Product Innovation and Destruction to Revenue Growth

	Overall Growth	Product Innovation	Product Destruction	Continuing Products
Hottman et al. (2016) Decomposition				
1-year	100	15.3	58.6	26.1
		(0.10)	(0.10)	(0.10)
5-year	100	38.2	50.3	11.5
		(0.10)	(0.10)	(0.10)

Notes: The table decomposes revenue growth into revenue loss from product loss, revenue gain from product gain, and changes in revenue on products the firm continues to produce. As an alternative to the Shapley-Owen values that we calculate in the main text, here we use a variance decomposition as in Hottman et al. (2016). The decomposition is based on OLS regressions of each component on overall firm revenue growth. Standard errors are given in parentheses.

Table C3: Tail Parameters for Key Variables

Gabaix and Ibragimov (2011) *Pareto Tail Estimator*	Revenue	Employment	Product Count	New Products	Revenue per Product
Tail parameter	1.12	1.37	1.97	1.85	1.24
	(.002)	(.004)	(.025)	(.100)	(.002)
R^2	0.99	0.99	0.99	0.96	0.99

Notes: The table presents the Pareto tail parameters for key variables. The tail parameter estimated on observations in the 90th percentile of the variable's distribution, although both absolute and relative tail parameters are similar when widening the sample up to the median. Standard errors are in parentheses. The estimates are based on the OLS regression approach proposed by Gabaix and Ibragimov (2011) and involves calculating the rank of each observation for the columns' variable, and then running a regression of the log of rank minus 0.5 on the log of the variable. Gabaix and Ibragimov (2011) show that small sample performance of this tail estimator is superior to a simple log rank, log size regression.

Table C4: Persistence in Revenue per Product of New products

	(1)	(2)
Persistence of product-level revenue on new products	0.38	-0.13
	(0.03)	(0.05)
Product (10-digit) fixed effects	Yes	Yes
Firm fixed effects	No	Yes
R-squared	0.58	0.75

Notes: The table plots the linear regression coefficient from a regression where the dependent variable is the revenue that a firm earns on a newly added product, while the explanatory variable is the average revenue that the firm earned on its new product in its previous episode of product creation. Additional controls include age, number of new products and 4-digit industry × year fixed effects. Standard errors in parentheses. Both regressions include 10-digit fixed effects.

Figure C1. Average Product Creation Rate by Number of Products with and without Industry Effects

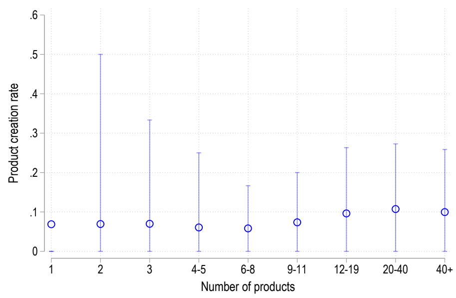

(a) Raw Data

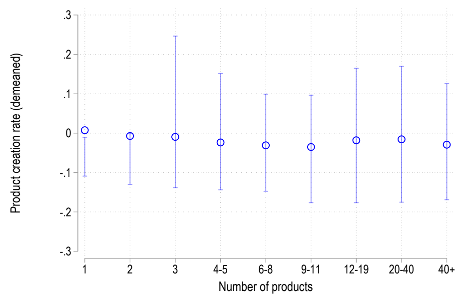

(b) Demeaned

Notes: The figures plot average product creation rates for firms of different initial sizes. The left-hand figure plots raw data, the right-hand figure demeans product creation rates by year and four-digit NACE industries using firms' primary industry. Confidence bounds give the 10th and 90th percentile of product creation rates for each size bin.

Figure C2. Average Revenue Per New Product: Different Sizes

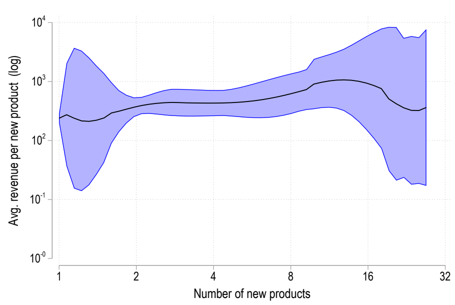

(a) Firms with $N_{t-1} = 5$

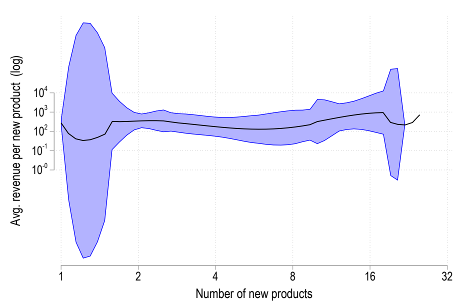

(b) Firms with $N_{t-1} = 10$

Notes: The 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. The line represents a kernel-weighted local polynomial smoothing of degree 5, together with its 95% confidence band.

Figure C3. Distribution of Number of New Products by Level of Aggregation: Different Sizes

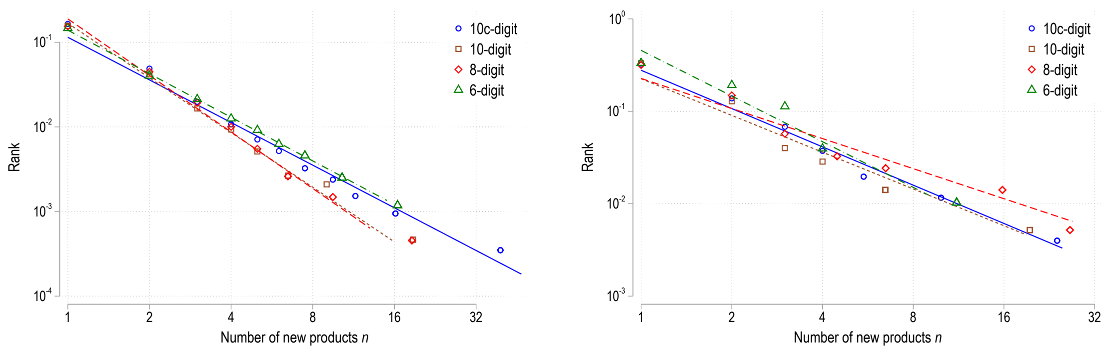

(a) Firms with $n_{it-1} = 5$ (b) Firms with $n_{it-1} = 10$

Notes: The figures plot the relationship between a firm's number of new products (horizontal axes) and the firm's rank (vertical axes). The figure provides overlapping plots of the log rank against the log number of products added at the 10c-digit concorded product level (blue circles), 10-digit PRODFRA level (brown squares), 8-digit (red diamonds) and 6-digit level (green triangles).

Figure C4. Distribution of Number of New Products by Proxies of Fixed Firm Boundary (All Firms)

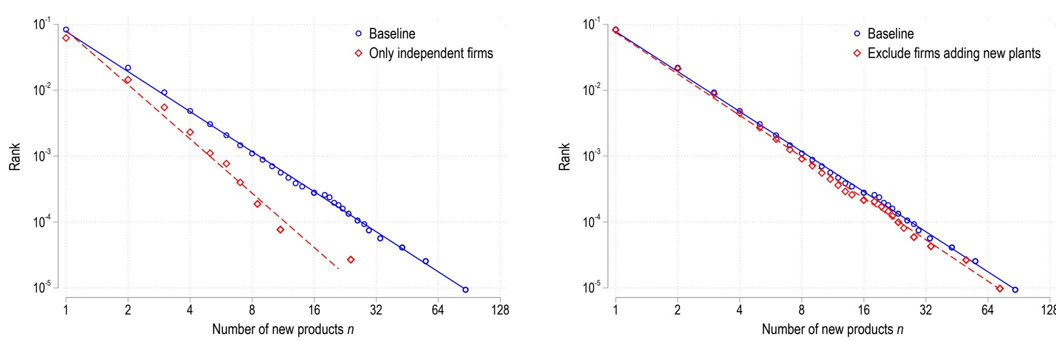

(a) Independent Firms Only (b) No Plant Change

Notes: The two panels plot the baseline Figure 2(a) against the distribution of new products of firms that do not belong to a business group (left panel) and of firms that do not add new plants over the period (right panel). Note that independent firms are, on average, smaller than firms that are part of a business group.

Figure C5. Distribution of Number of New Products with Fixed Intangible Capital

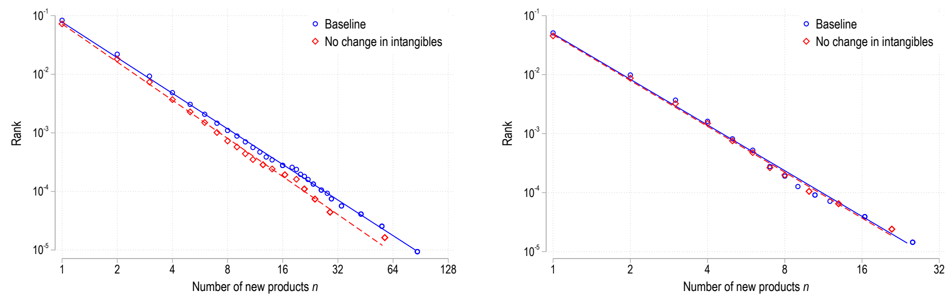

(a) All Firms (b) Firms with $N_{t-1} = 1$

Notes: The two panels plot the baseline Figure 2(a) (left panel) and Figure 2(b) (right panel) against the distribution of new products of firms that do not change intangible capital from one period to the next.

Figure C6. Distribution of Number of New Products by Age

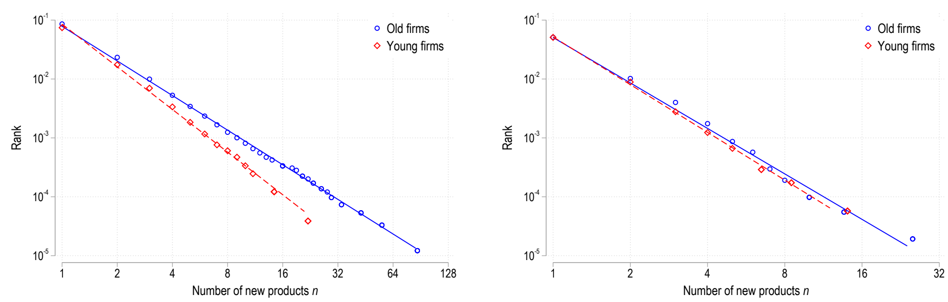

(a) All Firms (b) Firms with $N_{t-1} = 1$

Notes: The two panels replicate Figure 2(a) (left panel) and Figure 2(b) (right panel) by splitting firms into two groups by age (young if age ≤ 10 years). Figures for firms with 5 or 10 initial products can be found in the Companion Note.

Figure C7. Distribution of Number of New Products by Innovation History

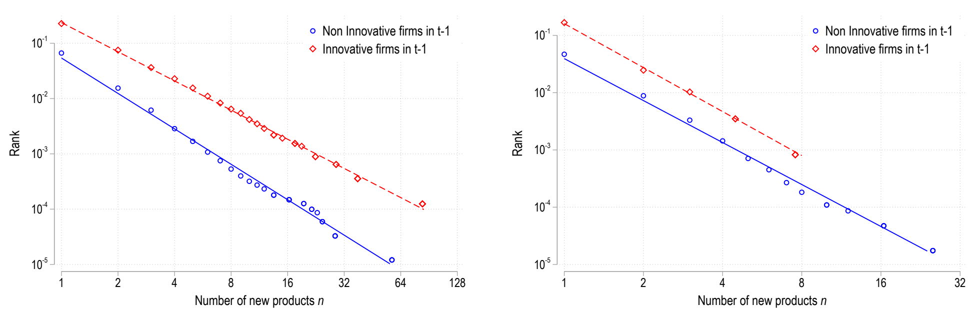

(a) All Firms (b) Firms with $N_{t-1} = 1$

Notes: The two panels replicate Figure 2(a) (left panel) and Figure 2(b) (right panel) by splitting firms into two groups by innovative status (innovative firms if they introduced new products at $t - 1$). Figures for firms with 5 or 10 initial products can be found in the Companion Note.

Figure C8. Distribution of Number of New Products by Production Mode

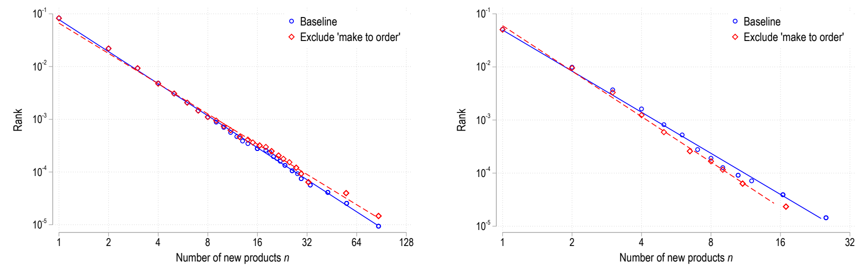

(a) All Firms (b) Firms with $N_{t-1} = 1$

Notes: The two panels replicate Figure 2(a) (left panel) and Figure 2(b) (right panel) using the original data (baseline) and data where new products are only counted if they are produced according to "Production Mode 3" in the EAP survey. Production mode 3 requires that the reporting firm has commercial and technological responsibility for the product. That is the case for the majority of the products in EAP. Alternative production modes involve the firm either outsourcing the product, or producing the product on behalf of a different firm to order, with the order firm either carrying commercial responsibility, design responsibility, or both. Figures for firms with 5 or 10 initial products can be found in the Companion Note.

Figure C9. Distribution of Number of New Products - Consumer Goods

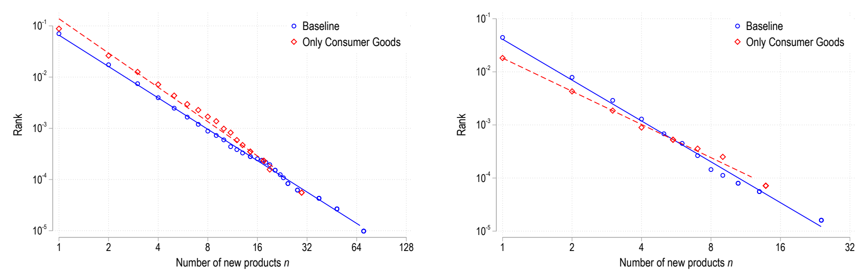

(a) All Firms (b) Firms with $N_{t-1} = 1$

Notes: The panels replicate Figure 2(a) (left panel) and Figure 2(b) (right panel) for the full sample, and exclusively for products in the Consumer Goods sector (BEC6). Figures for firms with 5 or 10 initial products are found in the Companion Note.

Figure C10. Distribution of Number of New Products - Excluding Product Re-entry

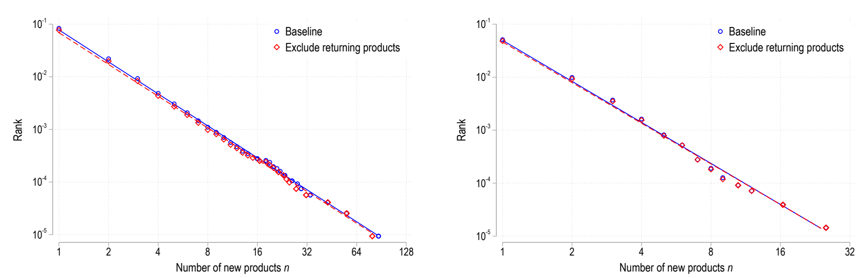

(a) All Firms (b) Firms with $N_{t-1} = 1$

Notes: The two panels replicate Figure 2(a) (left panel) and Figure 2(b) (right panel) both for the original sample, in which a product is new if it was not produced the previous year, and for the sample of new products that firms have never produced before. Figures for firms with 5 or 10 initial products can be found in the Companion Note.

Figure C11. Probability that Production Ends by Size of Innovation Burst at Product Creation

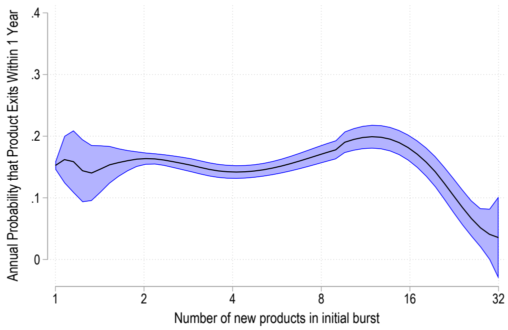

(a) All Firms, First Year

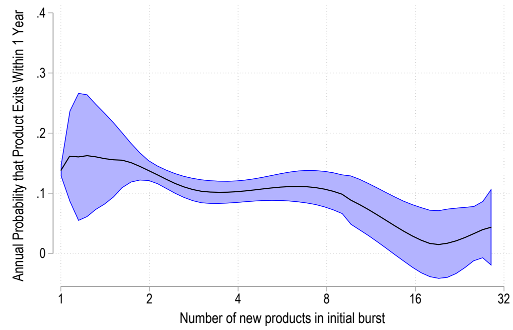

(b) Firms with $N_{t-1} = 1$, First Year

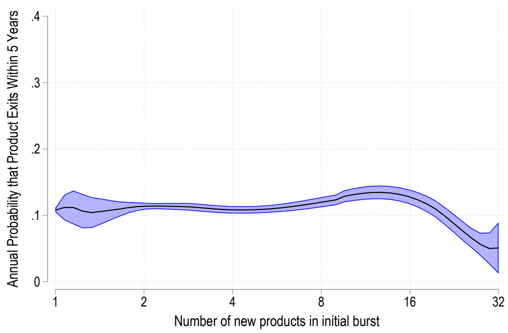

(c) All Firms, First 5 Years

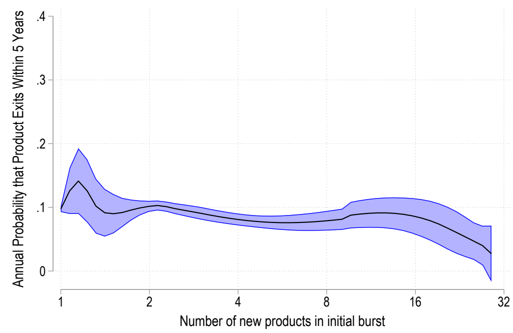

(d) Firms with $N_{t-1} = 1$, First 5 Years

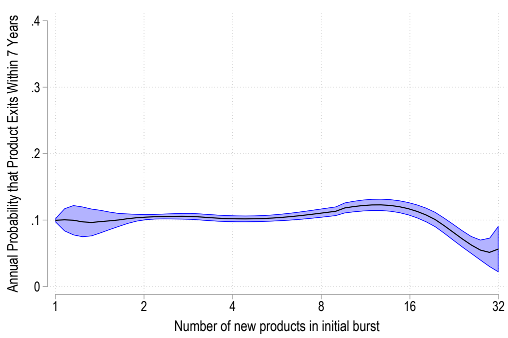

(e) All Firms, First 7 Years

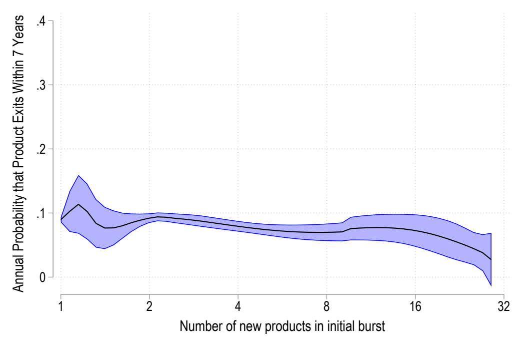

(f) Firms with $N_{t-1} = 1$, First 7 Years

Notes: The figure plots a local polynomial regression on the relationship between the number of products that were added at the time that a firm started producing good *j*, and the probability that the firm stops producing good *j*. Left-hand figures plot the regressions for firms of any initial size, right-hand figures only include firms that produced 1 product before they started producing good *j*. Rows of figures differ in the number of years of product tenure that are included in the plot. Confidence intervals are at the 95% level.

Table C5: Cannibalization: Effect of New Products on Probability of Product Exit

	(1)	(2)	(3)	(4)	(5)	(6)
β	0.082	0.082	0.075	0.101	0.135	0.074
	(0.029)	(0.029)	(0.027)	(0.036)	(0.014)	(0.000)
R-squared	0.004	0.015	0.029	0.109	0.339	0.721
Controls						
Product count FE	No	Yes	Yes	Yes	Yes	No
Product FE	No	No	Yes	Yes	Yes	No
Firm FE	No	No	No	Yes	Yes	Yes
Additional controls	No	No	No	No	Yes	Yes
Firm-product FE	No	No	No	No	No	Yes

Notes: The table presents regression coefficients for the linear regression along $\mathbb{I}^{exit}_{ijt} = \beta \left(\frac{n^c_{it}}{N_{it-1}} \right) + \gamma' X_{ijt} + \nu_{ijt}$, where $\mathbb{I}^{exit}_{ijt-1}$ is the indicator function that equals 1 for a product that has exited; i is the firm index; t denotes time; N_{it} is the number of products that the firm produces; n^c_{it} is the number of products the firm adds. A coefficient $\beta = 1$ implies full cannibalization, $\beta = 0$ implies no cannibalization. 361,000 observations.

Table C6: Cannibalization: Effect of New Products on Continuing Products

	(1)	(2)	(3)	(4)	(5)	(6)
β	-0.093	-0.095	-0.095	-0.083	-0.085	-0.083
	(0.007)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
R-squared	0.001	0.007	0.020	0.086	0.088	0.171
Controls						
Product count FE	No	Yes	Yes	Yes	Yes	No
Product FE	No	No	Yes	Yes	Yes	No
Firm FE	No	No	No	Yes	Yes	Yes
Additional controls	No	No	No	No	Yes	Yes
Firm-product FE	No	No	No	No	No	Yes

Notes: The table presents regression coefficients for the linear regression along

$$\frac{R^-_{ijt} - R^+_{ijt-1}}{0.5(R^-_{ijt} + R^+_{ijt-1})} = \beta \left(\frac{R^N_{it} / n^-_{it-1}}{0.5(R^+_{ijt} + R^-_{ijt-1})} \right) + \gamma' X_{ijt} + \nu_{ijt},$$

where R^N_{it} denotes total revenue from new products and R^-_{ijt} denotes revenue for kept product j while R^+_{ijt-1} denotes the revenue that the firm earned on that product at $t-1$; i is the firm index; t denotes time; n^-_{it} denotes the number of products that firm i produces and keeps producing. A coefficient $\beta = 1$ implies full cannibalization, $\beta = 0$ implies no cannibalization.

Figure C12. Life Cycle of Price Growth

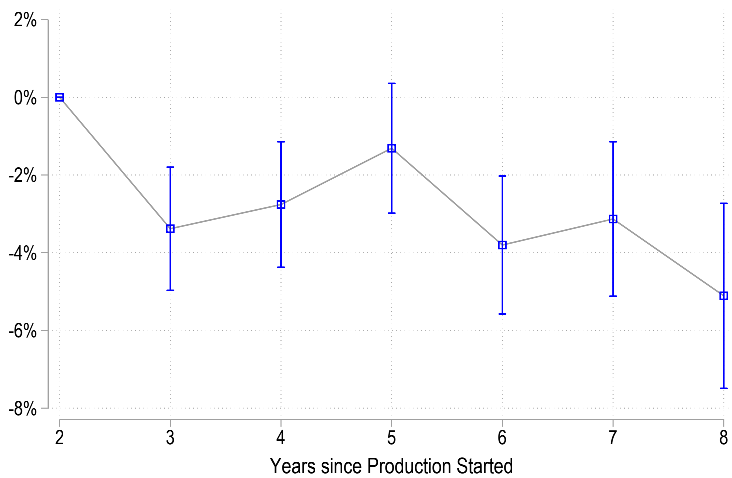

Notes: Vertical axis: avg. product unit value growth relative to growth at tenure 2 and after absorbing firm-product fixed effects. Horizontal axis: years since firm started producing the product. Confidence bounds (90% level) are based on clustered std. errors.

Figure C13. Distribution of Number of New Products: Data and Model

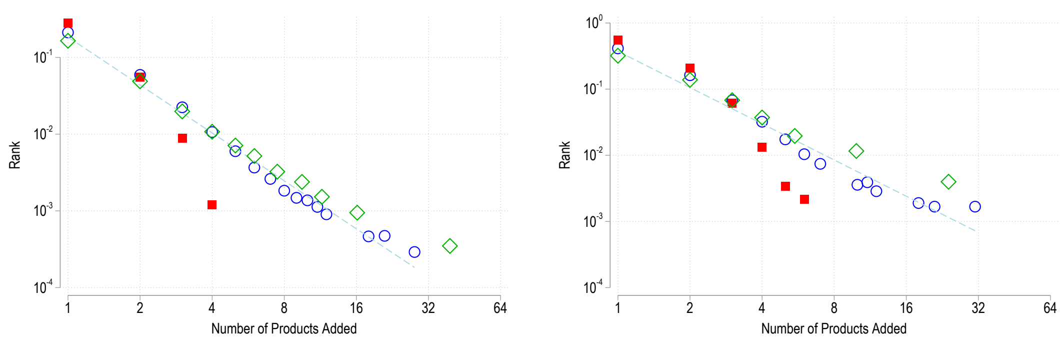

(a) Firms with $n_{i,t-1} = 5$

(b) Firms with $n_{i,t-1} = 10$

Notes: Figures plot the distribution of the number of products firms add to their portfolio. Green diamonds are data. Blue circles are from a model with innovation bursts. Red squares are based on a model with Poisson product innovation.

Appendix D. Predictions for Research and Development Spending

The model in Section 4 yields a set of predictions for firms' research and development (R&D) spending. Many of these overlap with established stylized facts in the literature on firm dynamics and growth, which motivated the Klette and Kortum (2004) framework. Below, we examine those predictions that are distinctive to our model and/or can only be assessed with product-level data (and not with the firm-level data available to prior work).

We obtain data on firm-level R&D spending from two sources. As our primary measure, we use data on total salaries paid to researchers employed by the firm from the administrative social security filings (DADS, see Appendix B). We identify researchers using occupational codes following the classification of R&D labor used by Bergeaud et al. (2025) and define salaries as all gross salary expenditures associated with the corresponding employees. As a robustness check, we also use data on firms' eligible R&D expenses from France's R&D subsidy scheme (Crédit Impot Recherche, CIR).[31]

The main prediction of our model on R&D is that a firm's choice of R&D spending scales in the number of products it produces. Previous work on the relationship between R&D intensity and size typically uses sales or employment. Our data enables us to instead examine this prediction with data on product count, which is a more direct approach to testing this prediction. In the quantified model where $\sigma = 1$, for all firms belonging to the same type in terms of innovative efficiency, R&D spending scales linearly in product count, and the ratio of R&D spending to a firm's product count should not depend on the latter.[32]

Table D1 presents an assessment of the degree to which R&D spending scales in the firm's number of products. It estimates the scaling parameter $\tilde{\sigma}$ in the constant-elasticity model $\mathbb{E}(\text{R\&D}_{it} \mid N_{it}) = N_{it}^{\tilde{\sigma}} \epsilon_{it}$ (where ϵ_{it} is the residual) using a PPML estimator as proposed by Silva and Tenreyro (2006).[33]

The table reports different estimates of the scaling parameter depending on different samples and weighting schemes. The top row is based on the DADS employment-based R&D data, and the bottom row based on the R&D subsidy data. The estimates, which range from 0.81 to 1.12, are broadly consistent with the linear scaling assumption of the model ($\tilde{\sigma} = 1$). With the exception of two unweighted DADS regressions that include single-product firms, none of the estimates differ from unity at conventional significance levels. This particular result is likely due to the positive selection of innovative single-product firms in both datasets. The patterns are quantitatively similar across the DADS-based regressions and those using the R&D subsidy data, so our subsequent results focus on the DADS data.

[31] An advantage of using the CIR data is that non-labor expenses on research are included. A downside of using the CIR data is that analysis by INSEE shows that there is under-reporting of overall R&D, causing firms to not appear in the survey when they do have positive R&D spending in the DADS data (Schweitzer 2019).

[32] Note that differences in firms' R&D efficiency can create a positive correlation between product count and R&D, because productive innovators invest more and therefore grow larger. However, as we explain with Figure 13, innovation bursts weaken the relationship between R&D and size, as a single burst can make a firm large.

[33] Silva and Tenreyro (2006) note that this is a consistent estimator of such models also in the presence of heteroskedasticity and frequent zeros, while log-linear OLS estimation may be severely biased.

Table D1: Relationship between R&D Spending and Firm Product Count

	Unweighted Regressions				Weighted Regressions			
Sample:	All	$N > 1$	All	$N > 1$	All	$N > 1$	All	$N > 1$
DADS Data	0.828	1.098	0.811	0.987	1.071	1.120	0.998	0.981
	(.077)	(.084)	(.070)	(.070)	(.080)	(.086)	(.069)	(.081)
R&D Subsidy Data	0.880	0.995	0.863	1.028	1.091	1.115	1.019	0.964
	(.102)	(.111)	(.0946)	(.0809)	(.135)	(.148)	(.120)	(.137)
Industry-Year Fixed Effects	No	No	Yes	Yes	No	No	Yes	Yes

Notes: The table presents estimates from a PPML regression (Silva and Tenreyro 2006). Columns with industry-year fixed effects control for fixed-effects at the level of four-digit NACE industries, based on the primary industry code that is assigned to the firm in FARE. Standard errors in parentheses are clustered by firm and year.

Figure D1. Relationship between R&D Spending per Product and Firm Size

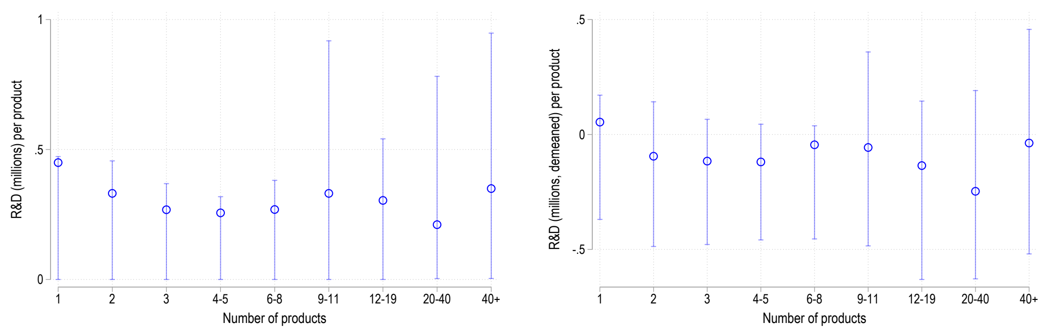

(a) Original Data

(b) Demeaned by NACE

Notes: Scatters plot averages. Confidence bounds plot 90th and 10th percentiles. DADS data. CIR data results are similar.

Given that R&D spending in the model scales linearly with product count, unweighted and weighted regressions should yield similar results. This is largely borne out in the data: estimates of $\tilde{\sigma}$ are closer to unity when weighting by product count, but the differences are typically within a standard error, especially once industry fixed effects are included.

As an alternative illustration of how R&D spending scales with product count, Figure D1 plots R&D per product by size bin. Each figure contains the average R&D per product (scatters) as well as the 90th and 10th percentile (confidence bounds). The left-hand figure plots raw data while the right-hand figure demeans R&D per product by 4-digit NACE industry codes and years. Both figures show that there is no monotone relationship between R&D per product and size, consistent with the linear scaling that we find in Table D1.

Turning to our framework's novel predictions, a distinctive feature is that the number of products firms add to their portfolios has a thick right tail. In our model, although firms can increase R&D to raise the arrival rate of innovation bursts, the number of products added when a burst occurs is not determined by firm choices. To test this assumption, we examine the relationship between the firm's R&D spending *per product* and the number of products it creates. The expected number of new products depends both on the arrival rate

Figure D2. Relationship between R&D Spending per Product and Number of New Products

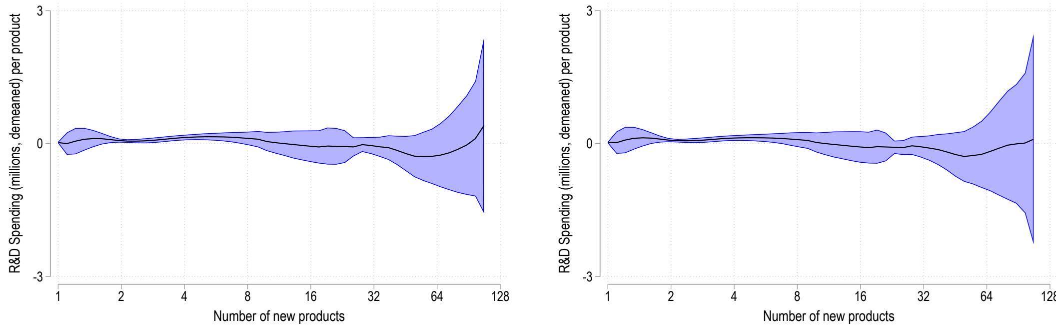

(a) Firm and Year Fixed Effects (b) Firm and Industry-Year Fixed Effects

Notes: Industry fixed effects are at the four-digit NACE level. The figure plots local polynomial regressions. DADS data.

of bursts and on the size of the bursts firms draw. According to the first-order condition (11), the arrival rate is pinned down by firm and time fixed effects. Therefore, residualized R&D should be orthogonal to the number of products firms add.

Figure D2 presents the results of local polynomial regressions of R&D per product on the number of products that firms create. The left-hand figure demeans R&D by firm fixed effects and time fixed effects.[34] In the right-hand figure we demean by firm and industry-time fixed effects, to control for potential time-varying heterogeneity in innovation costs parameters. The flat line in the figure indicates that R&D per product indeed is a poor predictor of the number of products that firms add, consistent with the model.

[34]The plot is qualitatively unchanged when removing firm fixed effects.