Private Know-How, Creative Destruction, and Growth

Filip Mellgren
IIES, Stockholm University

1 November 2026

Latest version

Abstract

New-to-firm and exiting products tend to be small. I document this pattern in Swedish manufacturing data and study its implications for creative destruction's contribution to growth. I develop a growth model in which innovation creates private know-how that competitors cannot inherit. This know-how can protect large products from replacement, concentrating replacement among smaller products where the same proportional quality gain contributes less to growth. In the calibrated model, creative destruction contributes 0.58 percentage points less to annual growth than in an accounting benchmark that removes the association between productive scale and protection while preserving replacement frequency. A uniform increase in R&D subsidies raises growth mainly through improvements to existing products. Targeted subsidies also change research at firms receiving no additional support.

1 Introduction

New-to-firm and exiting products tend to be small. I study what this pattern implies for creative destruction's contribution to growth. In my model, innovation improves product quality and creates private know-how that competitors cannot inherit. This know-how can protect large products from replacement, concentrating turnover among smaller lines. In the calibrated model, this selection lowers creative destruction's contribution to annual growth by 0.58 percentage points relative to an accounting benchmark with the same replacement frequency.

Firm entry and exit counts miss changes within surviving firms' product portfolios. Using Swedish manufacturing data linking firms, products and administrative records, I find that differences in sales per product account for most variation in firm sales, and sales are concentrated within portfolios. After removing firm-year means from log sales, about 60% of new-to-firm products and 69% of exits lie below the respective pooled medians among eligible products. These findings complement the evidence on selective product dropping in Andrew B. Bernard, Redding, and Schott (2010).

A common replacement hazard, as in Klette and Kortum (2004), cannot by itself explain this concentration of exits among small products. Private know-how gives products unequal protection: a challenger must improve on the knowledge it can inherit by enough to overcome the incumbent's private know-how. When larger products are better protected, replacement is concentrated on smaller lines. A replacement product can then improve on its predecessor while remaining small relative to established products. Since the same proportional quality improvement contributes less to growth when it occurs in a smaller product, the size of replaced products matters alongside replacement frequency.

To quantify the growth implications, I extend the model to allow firms to improve existing products and create new varieties, with variable markups following Peters and Walsh (2026). Improvements also build private know-how, which can protect the returns to further investment. I calibrate the model to aggregate and firm-level moments, together with mean product scope and the sales position of new products within firms one year after introduction. For the accounting comparison, I pair productive scale and protection independently while preserving their distributions and the replacement frequency.

Protection also affects the response to R&D subsidies. I use the calibrated model to study Sweden's R&D payroll deduction, *forskningsavdrag*. Raising the deduction from 10% to 20% increases annual growth by 0.088 percentage points at a long-run fiscal cost of 0.48% of output. Improvements to existing products account for most of the growth gain, while replacement remains concentrated on small lines. Including the initial consumption loss and subsequent

transition, the welfare gain is equivalent to a permanent 4.66% increase in consumption. Entry adjustment is central: holding entrants per year at their pre-reform level turns the growth effect negative.

The effects of subsidies also extend beyond the firms receiving them. In experiments that subsidize one firm type at a time, growth gains per fiscal krona rise with the share of knowledge that remains private. Changes in research on existing products, including at other firms, account for much of this difference. These comparisons concern unobserved protection types and illustrate the model's mechanism rather than an implementable targeting rule.

2 Relation to literature

The empirical analysis builds on product-level studies of firm size and growth. Andrew B. Bernard, Redding, and Schott (2010) document selective dropping of smaller products, while Hottman, Redding, and Weinstein (2016) decompose differences in firm sales into product scope, appeal, costs and markups. Berlingieri et al. (2026) document bursts of product creation in French manufacturing and build a growth model in which innovations can add several products at once. Argente, Lee, and Moreira (2024) use retail product life cycles to discipline a growth model and document declining sales, including among long-lived products. In my manufacturing data, adjusted sales rise among product lines observed continuously through age ten. Selection across firms and within product portfolios also features in trade models (Melitz 2003; Andrew B. Bernard, Redding, and Schott 2011; Mayer, Melitz, and Ottaviano 2014). I study an innovation-based source of selection that links product size to replacement risk.

The paper also relates to growth accounting that distinguishes creative destruction, new varieties and improvements to existing products (Garcia-Macia, Hsieh, and Klenow 2019; Klenow and Li 2021). Klenow and Li use employment at entering, exiting and continuing plants to infer these channels. In their benchmark, small exiting plants point toward obsolescence partly because creative destruction is undirected; they also consider directed creative destruction. Building on Klette and Kortum (2004), I study how private know-how generates unequal replacement risk and use product-level size patterns to inform its growth implications.

This mechanism connects the paper to models in which limits to knowledge diffusion protect incumbent producers. Jo and Kim (2026) show how learning delays allow firms to protect products through improvements to existing technologies. I study a related mechanism in which part of each innovation remains specific to the firm and is not inherited by a replacing producer. Shi (2025) studies how general R&D generates spillovers and helps firms absorb outside knowledge. In Shi's model, entrants inherit the productivity of randomly replaced

incumbents. Here, challengers build only on transferable knowledge, so accumulated private know-how affects both replacement success and the net quality gain from a given successful innovation step.

Protection can also arise from legal rules or cost advantages. Acemoglu and Akcigit (2012) study formal patent protection and show how protection tied to technological leads shapes incentives to extend them. De Ridder (2024) studies firm-specific cost advantages from intangible inputs that can prevent higher-quality competitors from entering production. Here, quality-improving innovation itself creates private know-how, allowing protection to accumulate through continued product development.

The policy analysis connects differences in knowledge spillovers to the allocation of research across firms. Lentz and Mortensen (2008) link firm heterogeneity and selection to growth, while Acemoglu et al. (2018) study how innovation policy affects selection and the allocation of labor between production and research. Aghion et al. (2025) distinguish rents from quality improvements and private process efficiency when studying R&D misallocation. In my model, a given quality improvement can leave different amounts of transferable knowledge across firms. Atkeson and Burstein (2010) show that entry can offset changes in incumbent innovation following a reduction in trade costs, while Atkeson and Burstein (2019) study the aggregate and welfare effects of innovation subsidies. I study these equilibrium responses when firms differ in the share of innovation that remains private. The policy experiments trace how subsidies change entry and research on existing products, including at firms receiving no additional subsidy.

3 Swedish firm and product data

The main data set is *Industrins Varuproduktion* (IVP) for 1997–2021, merged with *Företagens Ekonomi* (FE). IVP is an annual survey of industrial production that records output by firm and product. It covers firms with at least 20 employees, or at least 10 employees in selected industries, as well as firms qualifying through turnover thresholds.[1] During 1999–2021, IVP firms account for 87.0–90.2% of value added among manufacturing firms in FE with positive value added, compared with about 80% in 1997–1998.

FE contains annual income statements and balance sheets for nonfinancial firms. Its unit of observation is the principal legal unit (*huvudenhet*) representing a statistical enterprise, which may comprise several legal units. I aggregate IVP legal units to this level using the *Underenhet*

[1]The turnover threshold was SEK 50 million from 2006, rising to SEK 75 million by 2017. For reference year 2019, the unweighted response rate at preliminary publication was approximately 88%. For response rate details, refer to SCB, *Kvalitetsdeklaration: Industrins varuproduktion 2019*, section 2.2.4.

register before merging with FE by firm and year.

IVP records annual production value in thousands of current-year SEK, using nine-character product codes based on the *Combined Nomenclature* (CN). Respondents are asked to report sales for all industrial goods and services covered by the survey, allocating these sales across product categories. The measure is based on the sales value of own-produced, delivered goods, including notional values for internal deliveries. I refer to this measured production value as sales.

Throughout the analysis, a product is a CN8 category harmonized across years.[2] Each category may contain several distinct economic products. A product is active at a firm when sales summed across its constituent raw codes are positive. Firm sales are the sum of sales across its active product categories.

I use the FE industry classification to restrict the baseline sample to manufacturing sectors 15-37 before 2007, and 10-33 from 2007. I keep firm-years with positive employment, positive value added, and a wage bill lower than value added, and product observations with positive sales as defined above. Observations imputed by the statistical agency are retained in the baseline sample.

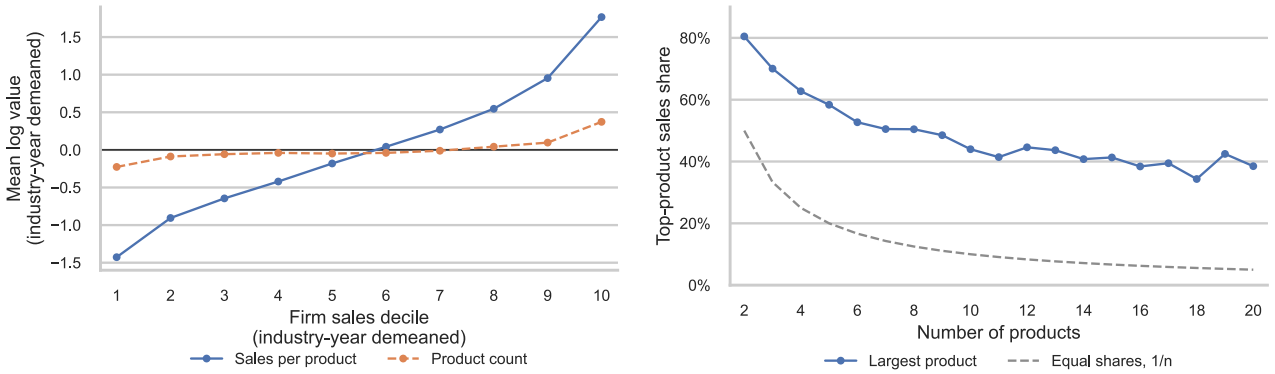
The baseline panel covers 9,518 firms, 89,721 firm-years and 254,409 firm-product-years over 1997–2021. Across firm-years, firms produce 2.84 products on average, and 45.9% of firm-years have a single product.

4 Facts on product size and turnover

Product turnover accompanies employment adjustment within continuing firms. The conditional comparisons in the empirical appendix show higher employment growth among firms that only add product categories and lower growth among firms that only drop categories, relative to firms with unchanged portfolios. This association motivates looking beyond firm entry and exit when studying changes in production. The following three facts describe how sales are distributed across products and which products firms add and drop.

[2] I extract the eight-digit CN8 code from the raw product code and harmonize categories over the sample period using the R package *harmonizer* (Baumgartner, Srhoj, and Walde 2023). Unmatched codes retain their original classification.

Fact A: Firm size and portfolio concentration. Large and small firms differ much more in sales per product than in the number of products they sell (Figure 1, panel i). Between the highest and lowest firm-sales deciles, the gap is 3.19 log units for sales per product and 0.60 for product count. Sales are also concentrated within firms. For example, among firms with five products, the largest product accounts for about 58% of sales, compared with 20% if sales were equally distributed (panel ii).



(i) Sales per product and product count

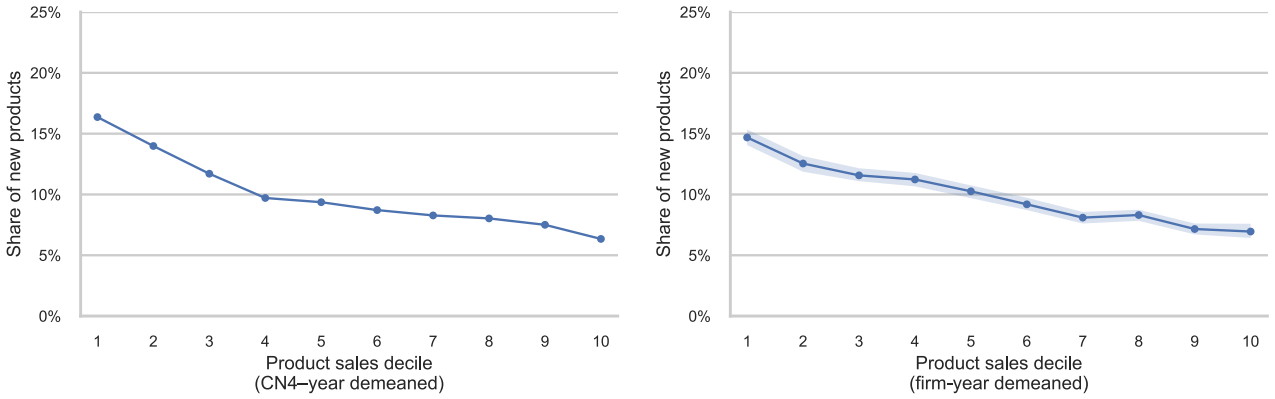
(ii) Top-product sales share

Figure 1: Firm size and portfolio concentration, 1997–2021. Panel i: industry-year-demeaned logs, averaged equally across firm-years within pooled firm-sales deciles. Panel ii: top-product shares averaged equally across industry-year cells for each exact product count; the dashed line is $1/N$.

Related patterns are documented by Andrew B. Bernard, Redding, and Schott (2010) and Hottman, Redding, and Weinstein (2016). The empirical appendix reports a complementary decomposition of firm-sales variation in the IVP data.

Fact B: New-to-firm products are small. New-to-firm products are concentrated in the lower sales deciles (Figure 2). About 60% fall below the pooled median of firm-year-adjusted log sales among eligible products (panel ii).

Because introduction-year sales may cover only part of a year, I measure sales in the year after first appearance ($t + 1$), among products still observed with positive sales at eligible firms. Panel i subtracts mean log sales for the same four-digit product group and year (CN4–year); panel ii subtracts each firm's mean log product sales in that year. In both panels, deciles are formed from pooled adjusted sales among all eligible products before selecting new products.

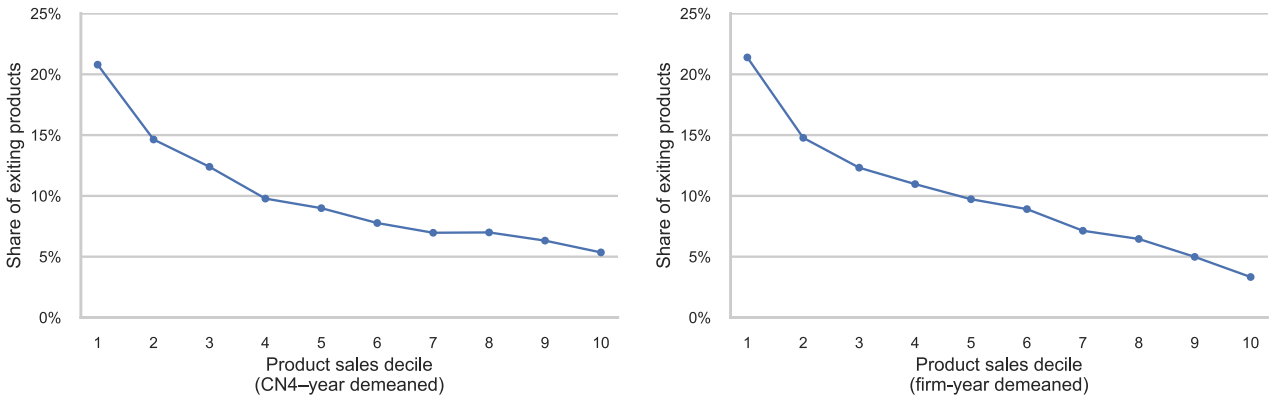


(i) Across products: CN4–year effects removed **(ii)** Within firms: firm-year effects removed

Figure 2: Distribution of new-to-firm products across sales deciles, 1997–2020 cohorts. If new products had the same adjusted-sales distribution as all eligible products, each decile would contain about 10%. New products exclude firm sample entry and reporting gaps; observations at $t+1$ require positive classified sales and an eligible firm (multiproduct in ii). Each panel sums to 100 percent and weights products equally. Shading in ii: pointwise 95 percent firm-cluster intervals; i: point estimates only.

Fact C: Exiting products are small. Exiting products are likewise concentrated in the lower sales deciles (Figure 3). About 69% fall below the pooled median of firm-year-adjusted log sales among eligible products (panel ii). Both panels use the same sales adjustments as Fact B.

An exit is a product last observed at t, absent in the following year ($t+1$), with no later positive reappearance at that firm through 2021. I measure sales in the preceding year ($t-1$), because sales in the last observed year may already be reduced by discontinuation during that year.



(i) Across products: CN4–year effects removed **(ii)** Within firms: firm-year effects removed

Figure 3: Distribution of exiting products across sales deciles, 1999–2020 origins. Established products must have positive eligible observations at $t-2$, $t-1$ and t, with an eligible firm at $t+1$. Panel ii requires multiproduct portfolios at $t-1$ and t. Each panel sums to 100 percent and weights exits equally.

Exits are measured among firms that remain in IVP and meet the baseline restrictions at $t+1$. If losing a large product causes a firm to fall below the survey threshold or fail these restrictions, that loss will be excluded.

The model that follows examines how firm-specific know-how can link differences in product sales to the pattern of product turnover.

5 A model of product protection

The baseline model builds on the framework of Klette and Kortum (2004). The motivating facts document substantial sales differences within firms, with new-to-firm and exiting products concentrated at low sales positions. Whereas the original framework assigns equal sales to product lines, I use constant elasticity of substitution (CES) demand to allow quality differences to generate unequal sales. For tractability, producers charge constant markups, following Acemoglu et al. (2018).

With fully transferable quality and uniformly targeted innovation, products acquired through creative destruction improve on randomly sampled predecessors. Under CES demand and constant markups, their sales distribution at acquisition is therefore shifted toward larger products relative to the existing-product distribution, while replacement occurs independently of sales. To allow products to differ in how easily they can be replaced, I distinguish between *transferable knowledge*, which future innovators can build upon, and *firm-specific know-how*, which is product quality tied to the firm that developed it. Know-how can protect a product from replacement and prolong the inventor's profits while leaving less knowledge for subsequent innovators.

When larger products are better protected, replacement can concentrate among smaller targets, allowing new-to-firm products to remain small despite improving on their predecessors. Innovation steps are drawn from a common distribution that allows some attempts to overcome installed know-how and others to fail. The model abstracts from own innovation and new-variety creation by incumbents to isolate this mechanism. Following Peters and Walsh (2026), the quantitative extension relaxes these restrictions and the constant-markup assumption.

5.1 Environment

Time is continuous and a representative household consumes a final good with utility

$$\mathcal{W} = \int_0^\infty e^{-\rho t} \ln C_t \, dt.$$

The household owns all firms and inelastically supplies labor L, which can be used for entry, production, or research. In equilibrium,

$$L = L_t^E + L_t^P + L_t^X, \qquad C_t = Y_t.$$

5.1.1 Quality, production, and sales

A variety j denotes a persistent product line, and a product is a firm's version of that variety, with firm indexed by f. A product is produced at total quality $q_{fj}a_{fj}$. The first factor, $q_{fj} > 0$, denotes quality that stems from transferable knowledge that is observable by competing firms and provides a base for future innovations in the variety. The second factor, $a_{fj} \geq 1$, is a private quality multiplier stemming from firm-specific know-how. At $a_{fj} = 1$, all quality is transferable. Competing products within a variety are perfect substitutes in quality-adjusted units, while different varieties are imperfect substitutes.

The baseline model assumes the highest-quality producer is a monopolist supplier of a given variety and charges the constant CES markup $\bar{\mu} = \sigma/(\sigma - 1)$, with $\sigma > 1$. I maintain the pricing outcome of the two-stage participation and pricing game in Acemoglu et al. (2018, 3457–58). The quantitative extension introduces Bertrand limit pricing and heterogeneous markups.

Suppressing time and producer subscripts, let $\mathcal{A}$ denote the set of active varieties, of measure N, and ℓ_j^P the production labor employed by the sole producer of variety j. Aggregate output can then be expressed as

$$Y = \left[\int_{\mathcal{A}} (q_j a_j \ell_j^P)^{(\sigma-1)/\sigma}\, dj\right]^{\sigma/(\sigma-1)}.$$

Define the CES mean of total quality and normalized public quality as

$$Q = \left[\frac{1}{N}\int_{\mathcal{A}} (q_j a_j)^{\sigma-1}\, dj\right]^{1/(\sigma-1)}, \qquad \tilde{q}_j = \frac{q_j}{Q}.$$

Let S_j denote a product's sales. Its sales relative to the average product are

$$\tilde{S}_j \equiv \frac{S_j}{Y/N} = \left(\frac{q_j a_j}{Q}\right)^{\sigma-1} = (a_j \tilde{q}_j)^{\sigma-1}, \qquad \frac{1}{N}\int_{\mathcal{A}} \tilde{S}_j\, dj = 1.$$

Let $\ell = L^P/N$ be production labor per active variety. Output, the wage, and operating profits satisfy

$$Y = QN^{1/(\sigma-1)}L^P, \qquad w = \frac{QN^{1/(\sigma-1)}}{\bar{\mu}}, \qquad \pi_j = \frac{S_j}{\sigma}, \qquad \frac{\pi_j}{w} = \frac{\ell}{\sigma-1}\tilde{S}_j.$$

Operating profits in wage units are therefore proportional to a product's sales relative to the average product. $QN^{1/(\sigma-1)}$ measures the total factor productivity in the economy, and along a balanced growth path Q grows at a constant rate while N is constant.

5.1.2 Innovation steps and firm-specific know-how

In the baseline model, all firms share a common know-how share $p \in (0,1)$. Each research attempt draws a positive log innovation step $d \sim G$, independently of the innovating firm and its target. A constant share p of the log step becomes firm-specific know-how, while the remaining share $1-p$ becomes transferable knowledge.

Researching firms target active varieties uniformly. If an innovation arrives and the incumbent product has quality $q'a'$, the quality achievable by the new invention has

$$q^+ = e^{(1-p)d}q', \qquad a^+ = e^{pd}, \qquad q^+a^+ = e^d q'.$$

The new invention has a higher total quality than the incumbent product if its technological improvement exceeds the firm-specific know-how of the incumbent producer

$$q^+a^+ > q'a' \quad \Longleftrightarrow \quad e^d > a'.$$

Thus an inventor builds on the target's public quality and must improve the variety enough to exceed the installed firm-specific know-how. On success, the log total quality of the variety rises by $d - \log a'$, and the incumbent's know-how a' is lost. The innovation step d therefore measures an improvement over transferable quality.

For firm-specific know-how to protect a product from replacement, the innovation-step distribution must allow some research attempts to fall short of installed know-how. With a common deterministic step d, products would instead carry private log know-how $pd < d$, so every subsequent attempt would overcome it. Dispersion in step sizes allows protection to affect the probability of replacement. Failed attempts create neither a product nor a platform.

On success, relative sales become

$$\tilde{S}^+ = \left(\frac{e^d}{a'}\right)^{\sigma-1}\tilde{S}'.$$

5.1.3 Research platforms and research effort

Firms may operate several products and an accumulating research capacity, which is represented by a stock of research platforms. Let $\mathcal{A}_f$ be the set of active varieties supplied by firm f, $n_f = |\mathcal{A}_f|$ its active product count, and $K_f \geq n_f$ its stock of research platforms.

Each successful product introduction adds one platform. Losing a product to a competitor ends its operating profits but leaves its research platform with the displaced firm.[3] This reflects the idea that a firm retains knowledge from production even after it ceases to be the frontier producer of that variety. A firm can therefore continue researching without any active products.

All firms have the same research technology. As in Klette and Kortum (2004), a firm chooses a Poisson research-attempt arrival rate x_f per platform, giving a total attempt rate $K_f x_f$. Its research labor requirement is

$$\ell_f^X = \frac{K_f x_f^2}{\phi_X}, \qquad \phi_X > 0.$$

The parameter ϕ_X measures research efficiency. Convexity captures diminishing returns to research labor conditional on the number of platforms: raising the attempt rate requires increasing amounts of labor per additional attempt. The quadratic specification is chosen for tractability. A research platform's innovation opportunities do not depend on its associated product's sales, firm-specific know-how, or current operating status.

Varieties become obsolete at an exogenous rate δ, removing the active product and all research platforms associated with that variety. A new platform created by capture is assigned to the target variety; the platform generating the attempt remains associated with its original variety. A firm exits when its last platform disappears due to obsolescence.

5.1.4 Entry

Entry uses labor at a constant cost per entrant:

$$E_t = \phi_E L_t^E, \qquad \text{labor per entrant} = 1/\phi_E.$$

Because entry requires a fixed amount of labor, its cost in units of the final good rises with

[3] This departs from the research technologies tied to active product counts in Klette and Kortum (2004, 991–92). Retention removes an incentive to accept operating losses to preserve future research opportunities. If displacement also destroyed research capacity, pricing and replacement could depend on its continuation value as well as product quality.

wages and hence with productivity along the balanced growth path. Letting the cost of entry scale with productivity finds support in Klenow and Li (2025) and ties the total number of firms to overall employment rather than productivity.

After paying the entry cost, an entrant draws an innovation step $d \sim G$ and introduces a product in a new variety along with a new research platform. The new variety's initial quality equals the economy's per-product quality mean:

$$a^+ = e^{pd}, \qquad q^+ = Qe^{-pd}, \qquad q^+ a^+ = Q.$$

For entrants, the draw determines the split between transferable quality and firm-specific know-how, while total initial quality is fixed. An entrant's relative sales are therefore $\tilde{S}^+ = 1$; the quantitative extension introduces an entry scale parameter.

5.2 The firm's problem

A firm chooses research effort to maximize the discounted profits from its current products and future innovations. I characterize this choice on a balanced growth path, taking wages, aggregate growth and the distribution of research targets as given. Let $g = \dot{Q}/Q$ denote quality growth, which also equals output and wage growth along this path. All values below are expressed in wage units. Along the balanced growth path, the interest rate is $r = \rho + g$ and wages grow at rate g. Expressing future payoffs in wage units removes this common growth trend, so their effective discount rate is $r - g = \rho$.

5.2.1 Additive firm value

Operating profits add across products, research costs scale with the number of platforms, and every platform faces the same research opportunities. A firm's existing portfolio does not affect its expected payoff per research attempt, and there are no financing constraints or complementarities across its products. These assumptions allow firm value to be decomposed into the values of its products and research platforms.

Let $v(a, \tilde{S})$ be the value of an active product together with one research platform, and let B be the value of a platform without an active product. Firm value is

$$V_f = \underbrace{\sum_{j \in \mathcal{A}_f} v(a_j, \tilde{S}_j)}_{\text{active products and their platforms}} + \underbrace{(K_f - n_f)B}_{\text{platforms without active products}}.$$

The second term accounts for research capacity retained after product loss.

5.2.2 The value of a product

An active product and its platform have two sources of value: operating profits and future research opportunities. A product's relative sales and flow profit in wage units decline as it falls behind the rest of the economy. In addition, the flow profit faces the risk of creative destruction depending on its installed firm-specific know-how. Let λ_H be the incoming innovation-attempt arrival rate per market. Relative-sales drift and the product's capture hazard are then

$$\dot{\tilde{S}} = -(\sigma - 1)g\tilde{S}, \qquad h(a) = \lambda_H \Pr(e^d > a) = \lambda_H[1 - G(\log a)].$$

Products also face a constant risk of becoming obsolete. Each research platform can generate new products, including after its associated product has been lost to a competitor. Let J^X denote the expected value created by one research attempt:

$$J^X = \mathbb{E}\left[\mathbf{1}\{e^d > a'\}\, v(a^+, \tilde{S}^+)\right].$$

The expectation is over the uniformly sampled target and the independent innovation step, with the post-capture state $(a^+, \tilde{S}^+)$ defined above. Failed attempts contribute zero. The value of an active product and its platform then satisfies the Bellman equation

$$\begin{aligned}\rho v(a,\tilde{S}) = \quad & \underbrace{\frac{\ell\tilde{S}}{\sigma-1}}_{\text{current profit}} - \underbrace{(\sigma-1)g\tilde{S}\,\partial_{\tilde{S}}v(a,\tilde{S})}_{\text{decline in relative sales}} \\ & - \underbrace{h(a)[v(a,\tilde{S}) - B]}_{\text{capture}} - \underbrace{\delta v(a,\tilde{S})}_{\text{obsolescence}} \\ & + \max_{x\geq 0}\left\{ \underbrace{xJ^X}_{\text{innovation payoff}} - \underbrace{x^2/\phi_X}_{\text{research cost}} \right\}.\end{aligned}$$

Capture ends the operating profit stream but leaves the research platform, so the loss is $v - B$. Market obsolescence removes both the product and its platform, so the loss is the full value v. The last line describes the platform's research choice: a higher attempt rate increases the expected innovation payoff but uses more labor.

Operating profits are proportional to relative sales, and neither the capture hazard nor the platform's research opportunities depend on product scale. The finite admissible solution

therefore has the affine form

$$v(a, \tilde{S}) = B + U(a, \tilde{S}) = B + A(a)\tilde{S}.$$

Here U is operating value alone, and $A(a)$ is the discounted operating value per unit of current relative sales. More firm-specific know-how weakly lowers the capture hazard and weakly raises this operating value at a given sales level. The common platform value B is retained after capture.

5.2.3 Research effort

The research choice equates the marginal labor cost of another attempt to its expected value:

$$\frac{2x}{\phi_X} = J^X.$$

All firms face the same expected payoff per attempt and research technology, so they choose the same effort x per platform. Their total attempt rates, $K_f x$, differ with their accumulated research capacity.

5.3 Balanced growth path equilibrium

I now determine the stationary environment taken as given in the firm's problem. Along a balanced growth path (BGP), the mass of active varieties N and the labor allocations are constant, while quality, output, consumption and wages grow at the common rate $g = \dot{Q}/Q$. Firms choose research optimally, entry is free, and product and labor markets clear. I focus on equilibria with positive entry and production labor, finite private values and a finite stock of research platforms.

Let P denote the stationary, count-weighted distribution of active products over installed firm-specific know-how and relative sales, $(a, \tilde{S})$, with $\mathbb{E}_P[\tilde{S}] = 1$. Although the firm-specific know-how share p is common to all firms, products differ in their innovation steps and histories. The distribution P must reproduce itself under entry, successful replacement, obsolescence and the decline in relative sales during an ownership spell.

5.3.1 Research attempts and platform stocks

Each research platform generates attempts at the common rate x. Since targets are sampled uniformly, the probability that an attempt succeeds is

$$H = \Pr_{P,G}(e^d > a),$$

where the target state $(a, \tilde{S})$ is drawn from P and the innovation step d is drawn independently from G.

Let K denote the aggregate stock of platforms, so K/N is the number of platforms per active variety. Variety mass evolves according to $\dot{N} = E - \delta N$, so constant variety mass requires $E = \delta N$. Each entrant adds one platform, each successful incumbent attempt adds another, and every platform faces obsolescence at rate δ. The stationary platform balance is therefore

$$0 = E + HxK - \delta K.$$

Combining the variety and platform balances gives

$$\frac{K}{N} = \frac{\delta}{\delta - Hx}, \qquad \lambda_H = \frac{K}{N}x, \qquad Hx < \delta.$$

The restriction $Hx < \delta$ ensures a finite platform stock: obsolescence must offset platform creation by incumbent research. Whenever incumbent research succeeds at a positive rate, there are more research platforms than active varieties because displaced firms retain their platforms. Incoming research attempts per variety, λ_H, therefore reflect both effort per platform and accumulated research capacity. The average rate of successful capture per variety is $\lambda_H H$.

5.3.2 Entry and labor allocation

An entrant pays its labor cost before drawing its innovation step and starts with relative sales one. With positive entry, its expected value in wage units equals the labor cost of entry:

$$\mathbb{E}_{d\sim G}\left[v(e^{pd}, 1)\right] = \frac{1}{\phi_E}.$$

Production uses $N\ell$ workers, incumbent research uses Kx^2/ϕ_X, and entry uses E/ϕ_E. Using the platform ratio K/N derived above and $E = \delta N$, labor market clearing yields

$$N = \frac{L}{\ell + (K/N)x^2/\phi_X + \delta/\phi_E}.$$

The denominator is total labor required per active variety.

5.3.3 Aggregate growth

A successful attempt replaces a target's total quality qa with qe^d. Its CES-weighted quality, $(qa)^{\sigma-1}$, therefore rises by the proportion $(e^d/a)^{\sigma-1} - 1$. Aggregating these gains across uniformly sampled targets gives

$$(\sigma - 1)g = \lambda_H \, \mathbb{E}_{P,G}\left[\mathbf{1}\{e^d > a\}\left\{\left(\frac{e^d}{a}\right)^{\sigma-1} - 1\right\}\tilde{S}\right] \geq 0.$$

The indicator is one for successful creative destruction attempts. The weight $\tilde{S}$ reflects the target's relative CES-weighted quality: the same proportional quality improvement contributes more to aggregate growth when it occurs in a larger product.

Entry and obsolescence have offsetting direct contributions to quality growth on the BGP. Entrants start at the quality mean Q, with relative sales one, while obsolete varieties are sampled from P, whose mean relative sales is also one. Their equal flows therefore cancel in the growth accounting above. Entry remains necessary to replace obsolete varieties and research platforms, and its initial quality is tied to the endogenous quality mean.

These conditions determine the stationary environment jointly: the product distribution determines success probabilities and expected research payoffs; research determines incoming attempts; and growth changes relative sales and operating values.

5.4 Product size, creative destruction and growth

Firm-specific know-how shapes which products creative destruction replaces and how much these replacements contribute to aggregate growth. To see this, compare the economy with the benchmark $p = 0$, in which all quality is transferable, as in Klette and Kortum (2004).

Fact C: which products are replaced. With $p = 0$, every attempt succeeds, and all products face capture hazard λ_H and obsolescence hazard δ, regardless of sales. Instantaneous exits therefore sample active products without selection on sales. With $p > 0$, the capture hazard $h(a)$ falls with installed know-how, and know-how is installed by the same innovation step that determines a product's size. Holding the target's public quality q' and the quality mean Q fixed, a capture with step d yields

$$\tilde{S}^+ = \left(\frac{q'e^d}{Q}\right)^{\sigma-1}, \qquad a^+ = e^{pd}.$$

Larger steps therefore create larger products with weakly lower capture hazards, providing a

possible link between size and replacement. This link is conditional: sales also reflect inherited public quality, and relative sales decline during an ownership spell while protection does not.

Fact B: where new products land. Fact B concerns incumbents' launches, which in the model are successful creative destruction attempts: replacing a product with relative sales $\tilde{S}'$ and know-how a' yields $\tilde{S}^+ = (e^d/a')^{\sigma-1}\tilde{S}' > \tilde{S}'$. With $p = 0$, the factor $e^{(\sigma-1)d} > 1$ applies to a randomly drawn target. At the time of variety capture, the across-product sales distribution of new products is therefore shifted toward larger products relative to that of active products. With $p > 0$, replacement is tilted toward weakly protected products. If these targets are small, this selection can offset the size increase at replacement, allowing new products to remain small relative to existing products. Through the lens of the model, Facts B and C can thus reflect the same replacements, observed from the entering and exiting sides.

Where growth comes from. The growth equation weights each successful replacement by the target's relative sales. Let

$$\Gamma(a) = \mathbb{E}_{d\sim G}\left[\max\left\{\left(\frac{e^d}{a}\right)^{\sigma-1} - 1,\ 0\right\}\right]$$

denote the expected proportional gain in CES-weighted quality per attempt. Protection weakly lowers this gain by reducing both the success probability and the gain from a given successful step, since installed know-how is lost on capture.

Growth decomposition: product size and protection. Suppose $\mathbb{E}_G[e^{(\sigma-1)d}] < \infty$. In the model BGP,

$$\begin{aligned}(\sigma-1)g &= \lambda_H \mathbb{E}_P\left[\Gamma(a)\tilde{S}\right] \\ &= \lambda_H\Big\{\mathbb{E}_P[\Gamma(a)] + \mathrm{Cov}_P\big(\Gamma(a), \tilde{S}\big)\Big\}.\end{aligned}$$

Holding $\lambda_H > 0$, G and the marginal distributions of size and protection fixed, changing their association leaves the average capture rate $\mathbb{E}_P[h(a)]$ unchanged. A negative covariance lowers the growth contribution of replacement relative to independence of size and protection.

The decomposition follows from $\mathbb{E}_P[\tilde{S}] = 1$. Reassigning protection across product sizes preserves the average replacement rate and the count-weighted expected gain $\mathbb{E}_P[\Gamma(a)]$, but changes whether these gains occur in large or small products. Equally frequent replacements can therefore contribute differently to growth.

With $p = 0$, Γ is constant and the covariance vanishes. With $p > 0$, protection lowers the

average gain per attempt. When the covariance is negative, these gains are tilted toward smaller products and contribute less to aggregate growth. Fact A shows why the size of the products being improved matters. Facts B and C suggest that much of the observed turnover involves small products; the model examines whether protection can generate this pattern and what it implies for growth.

The quantitative extension assesses the size-protection association and the equilibrium responses of research, entry and the product distribution. It adds variety creation and own innovation to evaluate replacement's contribution alongside other sources of growth. Variable markups and demand variation also separate observed sales from the productive scale that carries the growth weights.

6 Extensions for quantitative analysis

I extend the analytical model with incumbent variety creation, own innovation and variable markups, following Peters and Walsh (2026). These additions place creative destruction alongside other sources of growth and allow research, entry and the product distribution to respond jointly. Heterogeneous launch qualities and market demand variation connect product histories to observed sales and turnover.

I first describe the extensions with a common know-how share p, then introduce permanent differences across firms. Parameter values, empirical targets and the construction of simulated moments are presented in the calibration section.

6.1 Incumbent variety creation

In the analytical model, incumbents expand only by replacing another firm's product. A product new to a firm may instead open a new market. SCB's innovation survey for 2016–2018, the Swedish component of the Community Innovation Survey (CIS), distinguishes products new to the firm's market from products new only to the firm. Among manufacturing respondents, these categories account for 8.0 and 11.9 percent of total sales revenue in 2018, respectively.[4] These categories motivate allowing both replacement and variety creation, although they do not identify the model channels: a product new to the market can displace an existing product, and both modeled routes add a product to the innovating firm.

Following Peters and Walsh (2026), an incumbent research attempt targets an existing variety

[4]Own calculations for manufacturing respondents with valid innovation shares and positive net sales in the linked FE accounts, including non-innovators. I weight the reported shares by these net sales. The shares refer to 2018 sales from products introduced during 2016–2018, including significantly improved products; survey design weights are unavailable.

with probability α_X and creates a new variety with probability $1 - \alpha_X$. This direction is assigned before the target and innovation draws are realized. The research technology remains $\ell_f^X = K_f x_f^2/\phi_X$, where the firm chooses the Poisson arrival rate x_f per platform while taking α_X as given. An attempt directed at an existing variety follows the baseline capture rule, $e^d > a'$. A new-variety attempt always introduces a product and one research platform. Entrants continue to create new varieties only, corresponding to $\alpha_E = 0$.

Both successful routes add a research platform of the same value. Their operating values depend on the product's initial state, while attempts at existing varieties may fail.

The mass of varieties and incoming capture attempts now satisfy

$$\dot{N} = E + (1 - \alpha_X) \int K_f x_f \, df - \delta N,$$
$$\lambda_H = \frac{\alpha_X}{N} \int K_f x_f \, df.$$

Along a BGP, entrant and incumbent variety creation together replace obsolete markets. The baseline restriction $E = \delta N$ therefore becomes $E + (1 - \alpha_X) \int K_f x_f \, df = \delta N$. A capture changes a variety's producer and adds a platform, but leaves the variety mass unchanged.

If H is the success probability conditional on targeting an existing variety, each incumbent attempt creates a platform with probability $\chi_X = (1 - \alpha_X) + \alpha_X H$. In the common-$p$ economy, the condition for a finite stationary platform stock becomes $x\chi_X < \delta$, accounting for reproduction through both routes.

6.2 Quality and protection at launch

The analytical baseline assigns every entrant total quality Q. This makes all entrants equally large at birth and ties their initial scale to that of the average existing product. In the data, entrants are typically small but differ substantially in size. Figure 4 shows their value added in the first full year after their first observation in the firm registry data FE: the mean is about one quarter of the all-age sector-year firm mean, while the median is about 9 percent. This complements the new-product landing evidence, which describes products introduced by firms already active in IVP. For the quantitative model, I allow new products to start above or below the baseline quality benchmark. Separate launch scales for entrants and incumbents also allow newly founded firms and diversifying firms to introduce products of different average scale.

A new variety, introduced by an entrant ($c = E$) or an incumbent ($c = X$), draws an innovation step from the common distribution $d \sim G$ and an independent quality draw ω. Total quality is

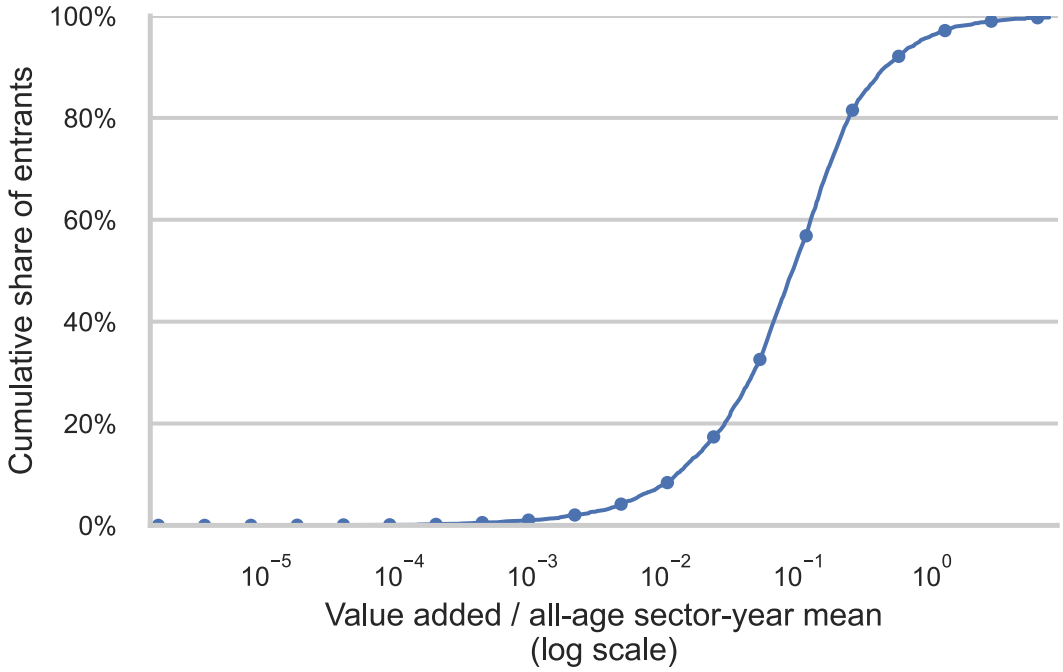


Figure 4: Entrant value added in the first full year, 2015–2019. Empirical cumulative distribution for 3,461 manufacturing entrants, with value added divided by the all-age mean in the same sector and year. Years receive equal weight and entrants equal weight within year. First appearance in FE defines age; age one avoids the entry year's partial exposure.

$$q^+a^+ = \bar{\omega}_c\, \omega e^d Q, \qquad c \in \{E, X\}.$$

The draw ω is lognormal,

$$\log \omega \sim \mathcal{N}(\mu_\omega, s_\omega^2), \qquad \mu_\omega = -\frac{\sigma - 1}{2} s_\omega^2, \qquad \mathbb{E}[\omega^{\sigma-1}] = 1.$$

The channel scale $\bar{\omega}_c$ sets average initial CES-weighted quality to m_c:[5]

$$\mathbb{E}\left[\left(\frac{q^+a^+}{Q}\right)^{\sigma-1}\right] = m_c.$$

Thus m_E and m_X govern average launch scale for entrants and incumbents, while s_ω governs dispersion conditional on the innovation step. The mapping from productive scale to sales also depends on markups. The quality draw is embodied in installed quality and is not redrawn during the product's life.

Differences in initial quality alone need not imply differences in protection from competitors.

[5]The normalization implies $\bar{\omega}_c = (m_c / \mathbb{E}_G[e^{(\sigma-1)d}])^{1/(\sigma-1)}$.

If a higher quality draw enters entirely as transferable knowledge, it makes the product larger without making it harder to replace. To allow part of this quality to be firm-specific, I let the same draw also influence initial know-how:

$$\log a^+ = p \max\{0,\, d + \theta_\omega(\log \omega - \mu_\omega)\},$$
$$q^+ = \frac{\bar{\omega}_c\, \omega e^d Q}{a^+}.$$

Here $\theta_\omega \geq 0$ governs the strength of the link between quality draws and protection. At $\theta_\omega = 0$, initial know-how depends only on the innovation step, $a^+ = e^{pd}$. The benchmark assigns $\theta_\omega = 1$, so p splits the centered log quality draw in the same way as the step. A sufficiently low draw leaves the product with $a^+ = 1$.[6]

Conditional on p and d, a higher quality draw therefore produces a larger and weakly better protected product at launch. It affects expected launch values and hence entry and incumbent research, as well as the joint distribution of product scale and exposure to capture.

These launch draws apply only to new varieties. A successful capture instead builds on its target's public quality, installing $a^+ = e^{pd}$ and $q^+ = e^{(1-p)d} q'$ as in the analytical model.

6.3 Own innovation

In addition to adding products, existing firms may invest in improvements to their current products, as in Luttmer (2007) and Atkeson and Burstein (2010). Garcia-Macia, Hsieh, and Klenow (2019) and Klenow and Li (2021) attribute most aggregate growth to incumbent own innovation. In Swedish manufacturing, SCB's 2013 business R&D survey reports that improvements to existing products account for 43.2 percent of product-directed intramural R&D expenditure.[7]

This motivates a separate research margin for improving existing products alongside expansion of the portfolio. It also allows firm-specific know-how to accumulate over time as a firm deepens its knowledge about its own products.

Figure 5 shows the sales profile of products observed continuously through age ten. The estimates follow the same 2,507 firm-product lines at every age, accounting for industry-year

[6] The numerical approximation caps initial log know-how at a boundary above every supported innovation step, preserving total quality and assigning public quality as the residual. Capped launches remain immune to capture, but the cap can bring forward markup compression by failed challenges and affect private values.

[7] Statistics Sweden, *Forskning och utveckling inom företagssektorn 2013*, Table 4, pp. 36–38. Improving existing products accounts for SEK 22,688 million of SEK 52,521 million spent on new or improved products. Products include goods and services; process R&D and general knowledge are excluded. The survey mainly covers firms with at least ten employees, and some purpose allocations are imputed.

and cohort effects under the normalization described in the empirical appendix. The estimated age profile rises by about 0.80 log points between ages one and ten. Growth within continuing product lines motivates allowing firms to improve their existing products. Sales also reflect demand and markups, so the profile does not separately measure quality improvement.

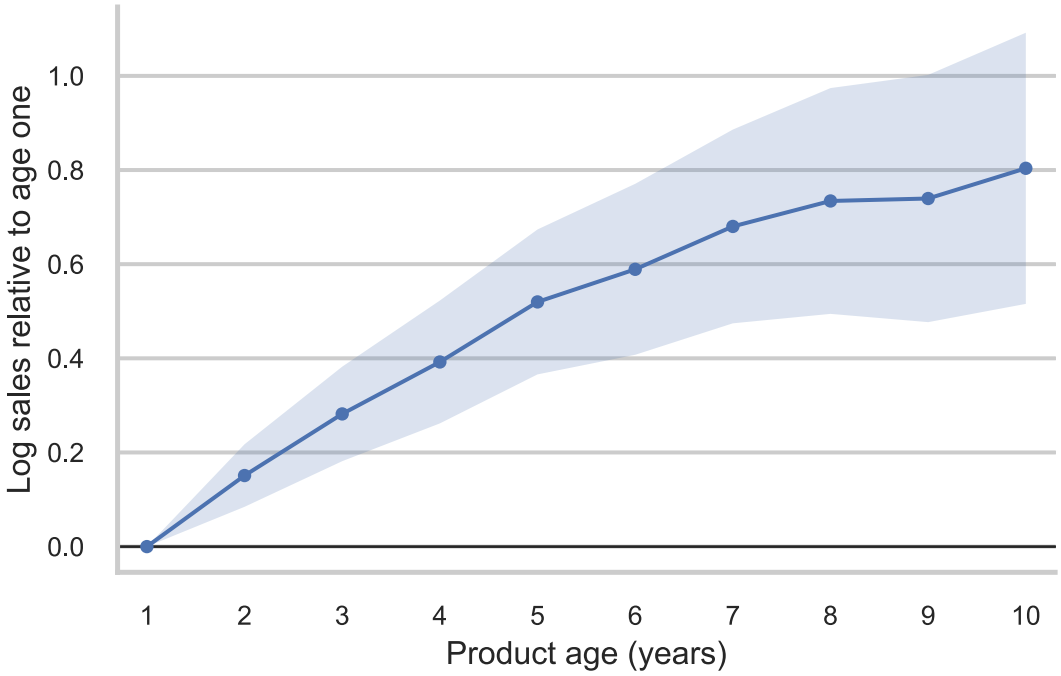


Figure 5: Sales over the life cycle of continuing products. The line shows estimated log-sales differences from age one for the same 2,507 firm-product lines at every age. The regression includes industry-year and cohort effects, with the linear cohort trend set to zero and no initial-sales control. The partial launch year is omitted. Shading shows pointwise 95 percent confidence intervals clustered by harmonized CN8 category.

A firm chooses a nonnegative own-innovation rate I_j for each active product. As with an innovation step, the resulting improvement is divided between transferable knowledge and firm-specific know-how:

$$\frac{\dot{q}_j}{q_j} = (1-p)I_j, \qquad \frac{\dot{a}_j}{a_j} = pI_j, \qquad \frac{d\log(q_j a_j)}{dt} = I_j.$$

Own innovation increases the quality of the current product without creating a new variety or research platform. It raises productive efficiency and can lower the product's future capture hazard by increasing installed firm-specific know-how.

Following the CES research-cost specification in Peters (2020, equation (16)), own innovation requires labor

$$\ell_j^I = (a_j \tilde{q}_j)^{\sigma-1} \frac{I_j^2}{\phi_I}, \qquad \phi_I > 0.$$

The quadratic cost parallels the expansion technology. A product with twice the normalized CES-weighted quality requires twice as much research labor to achieve the same growth rate. Because operating returns scale in the same way, this assumption preserves an operating value linear in product scale and an own-innovation policy that depends on protection and competitive pressure rather than scale. Together with separate costs for own innovation and expansion and no shared firm research constraint, it preserves additive firm value. I characterize the policy after specifying competition below. The demand extension then scales both operating returns and improvement costs with market demand.

6.4 Bertrand competition and variable markups

Constant markups make product sales proportional to CES-weighted quality. They also abstract from the competition that a displaced or unsuccessful producer may still exert. I therefore replace the baseline pricing assumption with Bertrand competition within varieties, following Peters (2020). This allows research to affect profits through both productive quality and market power.

Figure 6 provides an empirical motivation for allowing markups to evolve. It shows firm-age profiles of revenue productivity (TFPR) and value added relative to labor and capital costs. Between ages one and twenty, the estimated differences correspond to increases of 9.1 and 6.7 percent, respectively. The cost ratio divides value added by the wage bill plus an imputed capital cost; TFPR uses a Cobb–Douglas index of capital and the wage bill. These revenue-based proxies admit a markup interpretation under the assumptions described in the empirical appendix.

Let $a_j^C \tilde{q}_j^C$ denote the normalized total quality of the best external competing product within a variety. Inferior products remain available as suppliers with no overhead cost, but earn zero operating profits and cannot improve their products. Transferable knowledge enters subsequent research attempts; it does not automatically upgrade these old products.

The leading producer supplies the product and charges the lower of the CES monopoly markup and the quality gap to its best competitor. Define the capped log gap

$$\Delta_j = \min\left\{\log \bar{\mu}, \log\left(\frac{a_j \tilde{q}_j}{a_j^C \tilde{q}_j^C}\right)\right\},$$
$$\mu_j = e^{\Delta_j}, \qquad \bar{\mu} = \frac{\sigma}{\sigma - 1}.$$

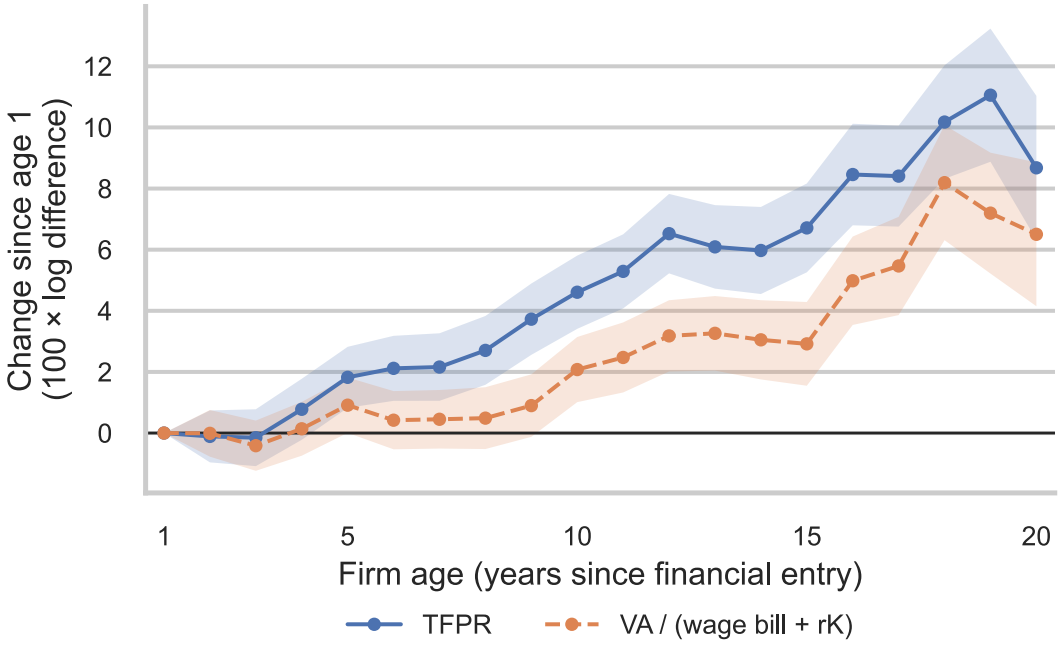


Figure 6: Revenue productivity and the value-added cost ratio over firm age. Lines show estimated log differences from age one, multiplied by 100, from regressions with firm and year effects. The normalization sets the linear cohort trend in firm effects to zero. Shading shows pointwise 95 percent firm-cluster confidence intervals for each age-one contrast. The cost ratio permits zero measured capital; TFPR requires positive capital.

A product without a competitor in the same variety charges the monopolistic markup $\bar{\mu}$. Between research attempts, own innovation increases Δ_j at rate I_j until the monopoly cap is reached. A successful challenger, $d > \log a'$, becomes the producer and faces the displaced leader as its best competitor. Its initial gap is

$$\Delta^+ = \min\{\log \bar{\mu},\, d - \log a'\}.$$

An unsuccessful challenger leaves ownership, installed quality and platform stocks unchanged. It can nevertheless become a better competing supplier and reduce the incumbent's gap to

$$\Delta^+ = \min\{\Delta',\, \log a' - d\}, \qquad d \le \log a'.$$

Thus failed research can compress markups even when firm-specific know-how prevents replacement. Ties retain the incumbent as producer and imply a markup of one. Retention of research platforms continues to separate the loss of operating profits from the loss of future expansion opportunities.

Competition at launch. New varieties also need an initial competitive gap. Allowing this gap

to fall below the monopoly cap accommodates positive margins at launch followed by modest markup growth. The firm-age profiles motivate allowing this margin but do not directly measure initial product markups. Entrant and incumbent launches start at a common log gap

$$\Delta^+ = \Delta_0 \in (0, \log \bar{\mu}], \qquad \mu^+ = e^{\Delta_0}.$$

This reduced-form initial condition sets pricing pressure separately from launch quality and protection. Subsequent gaps follow the same improvement and challenge laws as other products. Setting $\Delta_0 = \log \bar{\mu}$ recovers launch pricing without an initially binding competitor.

Sales, profits and aggregate output. Variable markups change the mapping from quality to sales and profits. Let P now denote the count-weighted stationary product distribution over $(a, \tilde{q}, \Delta)$, with $\mathbb{E}_P[(a\tilde{q})^{\sigma-1}] = 1$, and define

$$H_{\sigma-1} = \mathbb{E}_P[(a\tilde{q})^{\sigma-1}\mu^{1-\sigma}],$$
$$H_\sigma = \mathbb{E}_P[(a\tilde{q})^{\sigma-1}\mu^{-\sigma}].$$

Relative sales become

$$\tilde{S}_j = \frac{(a_j\tilde{q}_j)^{\sigma-1}\mu_j^{1-\sigma}}{H_{\sigma-1}}.$$

Operating profits are $\pi_j = (1 - \mu_j^{-1})S_j$, so the profit share now varies across products. For a given productive quality, a higher markup reduces sales. Aggregate output and wages satisfy

$$Y = MQN^{1/(\sigma-1)}L^P, \qquad w = \Lambda\frac{Y}{L^P},$$
$$M = \frac{H_{\sigma-1}^{\sigma/(\sigma-1)}}{H_\sigma}, \qquad \Lambda = \frac{H_\sigma}{H_{\sigma-1}}.$$

Here $M \leq 1$ measures the production loss from markup dispersion relative to a common-markup allocation at the same qualities and production labor. The production labor share is Λ, so $1 - \Lambda$ is the operating-profit share before research costs. With a common markup $\bar{\mu}$, these expressions reduce to $M = 1$, $\Lambda = 1/\bar{\mu}$ and the baseline sales equation. With variable markups, productive scale and observed relative sales must be distinguished.

6.5 Firm value and research choices

The firm continues to choose research to maximize discounted profits. As in the analytical model, values are measured in wage units and discounted at ρ. Retained platforms and the scaling of operating returns and improvement costs preserve the decomposition

$$V_f = K_f B + \sum_{j \in \mathcal{A}_f} (a_j \tilde{q}_j)^{\sigma-1} A(a_j, \Delta_j).$$

Here B is the value of a research platform and $A(a, \Delta)$ is operating value per unit of normalized CES-weighted quality, net of future own-innovation costs. The operating value ends at capture or obsolescence; the platform survives capture. The coefficient A incorporates future changes in protection and markups, including compression from failed challenges.

Own innovation. A unit increase in I raises the growth rate of productive scale by $\sigma - 1$, the growth rate of know-how by p, and the drift of the log competitive gap by one until its cap. Where A is differentiable, the optimal rate satisfies

$$I(a, \Delta) = \frac{\phi_I}{2} \left[(\sigma - 1)A + pa\, \partial_a A + \mathbf{1}\{\Delta < \log \bar{\mu}\} \partial_\Delta A \right]_+ ,$$

where $[z]_+ = \max\{z, 0\}$ and A and its derivatives are evaluated at (a, Δ). The three terms are the marginal values of increasing productive scale, building protection, and widening the competitive gap. The firm equates their sum to marginal research cost $2I/\phi_I$ when $I > 0$. The common scale factor cancels, leaving a policy that depends on protection and the gap, rather than product size.

These returns create scope for complementarity between protection and own innovation: protection prolongs the return to improving a product, while improvement builds further protection. They do not imply that $I(a, \Delta)$ increases monotonically with installed know-how. Additional protection may have little value for a product already insulated from capture, and the markup return disappears at the monopoly cap. The policy and operating value must therefore be determined jointly.

Expansion and entry. Write $v(a, \tilde{q}, \Delta) = B + (a\tilde{q})^{\sigma-1} A(a, \Delta)$ for the value of a product together with its new platform. The payoff per incumbent attempt, J^X, averages this value over incumbent launches with weight $1 - \alpha_X$ and attempts at existing varieties with weight α_X; failed attempts yield zero to the challenger. The research condition and free-entry condition are

$$\frac{2x}{\phi_X} = J^X, \qquad \mathbb{E}_{d,\omega}\left[v(a_E^+, \tilde{q}_E^+, \Delta_0)\right] = \frac{1}{\phi_E}.$$

The second condition averages over the entrant launch draws specified above. Entry and expansion need not have the same expected payoff: entrants create only new varieties and have their own launch scale. With common p, all platforms face the same opportunities and choose the same x. Platform value satisfies $(\rho + \delta)B = xJ^X - x^2/\phi_X$, with no loss term for capture because the platform is retained.

6.6 Permanent firm heterogeneity

The quantitative benchmark lets each entrant draw a permanent know-how share $p_f \sim F_p$ after paying the entry cost. All products developed by the firm use this share to divide improvements into transferable knowledge and private know-how. At a given target and innovation step, firms generate the same total candidate quality but differ in the protection they install and the knowledge left to future innovators.

Research effort, platform value and own-innovation policies can therefore vary by type. Free entry equates the entry cost to expected value before the type draw, rather than imposing zero profit separately on every realized type. The distribution of types among producing firms, products and research platforms is endogenous and need not equal the distribution at entry. The equilibrium conditions below are written for common p; the quantitative benchmark applies the firm problem by type and aggregates across types.

6.7 Equilibrium and aggregate growth

For the common-p economy described above, a balanced growth path combines optimal research and free entry with stationary product and platform distributions and market clearing. The product law P reproduces itself under own improvement, challenges, launches and obsolescence. Let $K = \int K_f \, df$ be the aggregate platform stock. The variety and platform balances are

$$E + (1 - \alpha_X)Kx = \delta N, \qquad E + \chi_X Kx = \delta K,$$

where $\chi_X = (1 - \alpha_X) + \alpha_X H$ and $H = \mathbb{E}_P[1 - G(\log a)]$. Production, expansion, entry and own innovation employ the available labor:

$$L = L^P + \frac{Kx^2}{\phi_X} + \frac{E}{\phi_E} + N\mathbb{E}_P\left[(a\tilde{q})^{\sigma-1}\frac{I(a,\Delta)^2}{\phi_I}\right].$$

Aggregate quality growth follows by adding the contributions of improvement, replacement and new varieties, and subtracting losses from obsolescence. Let $\Gamma(a) = \mathbb{E}_G[\max\{(e^d/a)^{\sigma-1} - 1, 0\}]$ denote the expected proportional gain in CES-weighted quality from an attempt at an existing variety, as in the analytical decomposition. Since N is constant on the BGP,

$$\begin{aligned}(\sigma-1)g = &\underbrace{(\sigma-1)\mathbb{E}_P[(a\tilde{q})^{\sigma-1}I(a,\Delta)]}_{\text{own improvement}} \\ &+ \underbrace{\lambda_H \mathbb{E}_P[(a\tilde{q})^{\sigma-1}\Gamma(a)]}_{\text{creative destruction}} \\ &+ \underbrace{\frac{E}{N}m_E}_{\text{entrant launches}} + \underbrace{\frac{(1-\alpha_X)Kx}{N}m_X}_{\text{incumbent launches}} - \underbrace{\delta}_{\text{obsolescence}}.\end{aligned} \tag{1}$$

Both launch channels make positive contributions to the quality stock. Obsolescence occurs independently of innovation and removes products with average normalized CES-weighted quality one. Launches replace their number on the BGP, but their combined quality contribution can fall short of the quality lost to obsolescence. When this happens, launches and obsolescence make a negative net contribution to quality growth, even though each launch adds quality.

Own improvement and creative destruction are weighted by productive scale. Writing normalized productive scale as $\nu = (a\tilde{q})^{\sigma-1}$, with $\mathbb{E}_P[\nu] = 1$, the creative-destruction term retains the analytical covariance decomposition:

$$\lambda_H \mathbb{E}_P[\nu\Gamma(a)] = \lambda_H\Big\{\mathbb{E}_P[\Gamma(a)] + \mathrm{Cov}_P\big(\nu, \Gamma(a)\big)\Big\}.$$

A negative covariance means that expected replacement gains are concentrated on smaller products. It lowers creative destruction's contribution to growth relative to independence of scale and protection, holding the attempt rate, innovation-step distribution and both marginal distributions fixed. The relevant gain $\Gamma(a)$ reflects both the probability of capture and the quality gain conditional on success. Own innovation and variety creation add further growth contributions without removing this allocation mechanism. Variable markups separate productive scale from sales and also affect incentives for research and entry.

Since M, N and L^P are stationary on the BGP, output and wages grow at the same rate g as Q. Values, research policies, entry, the product distribution and growth are determined jointly by these conditions.

6.8 Market demand and annual sales

Sales within continuing product lines also vary for reasons unrelated to improvements in quality. I capture this variation with a market-level demand process.

Each variety carries a demand shifter $\theta_j > 0$, common to all its potential suppliers. It starts at one when the variety is created and subsequently follows

$$d \log \theta_j = -\frac{\sigma_\theta^2}{2}\, dt + \sigma_\theta\, dW_j, \qquad \theta_j = 1 \text{ at variety birth,}$$

where Brownian motions are independent across varieties and of innovation histories. The shifter is a mean-one martingale. It stays with the variety when a new producer takes over and disappears at obsolescence. Unlike the quality draw ω, which is embodied in quality, θ_j changes current demand without adding public knowledge or private know-how.

The shifter multiplies the variety's CES demand, so sales and production labor scale with θ_j. I also scale own-innovation cost with demand-weighted product size:

$$\ell_j^I = \theta_j (a_j \tilde{q}_j)^{\sigma-1} \frac{I_j^2}{\phi_I}.$$

This scaling keeps operating value linear in demand-weighted quality, $\theta_j (a_j \tilde{q}_j)^{\sigma-1} A(a_j, \Delta_j)$, and the own-innovation policy independent of θ_j. The mean-one demand process then leaves expected research payoffs and aggregate resource requirements unchanged. Capture and pricing depend on quality comparisons within a variety, where the common demand shifter cancels. Under these assumptions, demand risk changes the distribution of sales and employment while leaving research policies, aggregate growth and labor allocation unchanged at given structural parameters. Its volatility σ_θ is disciplined by sales-growth dispersion in the calibration.

Finally, I construct annual product sales by integrating the demand-adjusted sales flow over the portions of each year in which a firm supplies the variety, and sum across products to obtain firm sales. This includes partial-year production and changes of producer within the year. Firm scope, sales ranks and growth are constructed from these annual observations using the empirical definitions described in the calibration section.

7 Model estimation

I estimate the model by matching aggregate, firm-level and firm-product moments for Swedish manufacturing. The benchmark allows permanent differences in firms' know-how shares, with the type distribution assigned before estimation. Six parameters governing innovation, entry and product scale are chosen jointly, conditional on the remaining parameters. Model observations use annual timing and sample restrictions corresponding to the data.

7.1 Parameters and estimation strategy

The reference economy is Swedish manufacturing around 2019. A model firm corresponds to the FE reporting unit described in the data section, and an active product to a harmonized CN8 category with positive sales at that firm, following Berlingieri et al. (2026). A category may contain several economic varieties; an observed addition or exit does not by itself identify innovation or creative destruction. Product moments use annual sales and simulated firms with at least 20 employees, approximating IVP coverage. Firm-level moments use the broader FE sample.

Table 1 lists assigned parameters. The elasticity, research curvature and discount rate follow Peters and Walsh (2026); the remaining entries are maintained assumptions. Entrants create new varieties only, consistent with their small observed size. The ex-ante distribution of protection types is assigned.

Table 1: Parameters and distributions assigned before estimation.

Parameter	Interpretation	Value
σ	Elasticity of substitution	4
ρ	Annual discount rate	.012
ζ	Research-cost curvature	2
α_E	Entrant share of attempts at existing markets	0
$e^d - 1$	Discovery increment	Lognormal
$\log \omega$	Log quality draw	Normal
θ_ω	Link between quality draw and protection	1
p at entry	Permanent protection heterogeneity	Beta(2,6)

Notes: The discovery increment has mean .25 and coefficient of variation 1. The quality draw is normalized so that $\mathbb{E}[\omega^{\sigma-1}] = 1$. Protection uses eight equal-mass conditional-mean nodes.

Four parameters are held at values disciplined in earlier fits (Table 2). Those earlier moments therefore do not provide independent validation of the model.

Table 2: Other parameters held fixed during estimation.

Parameter	Value	Basis
Dispersion of quality draws, s_ω	.4588	Earlier fit to the dispersion of entrant sales
Capture-attempt share, α_X	.8004	Earlier fit to CIS2020 new-to-firm sales share
Initial log competitive gap, Δ_0	.1111	Earlier fit to the age profile of the value-added cost ratio
Demand volatility, σ_θ	.1080	Earlier fit to growth-variance increments

Notes: Entrant dispersion uses log first-full-year value added relative to the industry-year mean. The age-profile target is the age-20 minus age-1 log value-added cost ratio, with firm and year effects and a zero cohort-trend normalization. The empirical denominator is the wage bill plus imputed capital costs; the model counterpart excludes research labor. The empirical appendix describes the proxy and its limitations.

I combine stationary-equilibrium calculations with simulated annual firm histories. Growth, firm turnover, relative entrant size and mean employment are computed from the equilibrium; product counts and the size distribution of new products are measured in the simulated histories. The parameters are adjusted jointly to match these six moments.[8] Independent simulations and a finer grid assess numerical error, not parameter uncertainty.

7.2 Target moments

The moments jointly discipline the model; the parameter pairings in Table 3 describe their main economic roles rather than separate identification.

Productivity growth. The mean of twenty EU KLEMS manufacturing TFP changes ending in 2000–2019, each divided by the corresponding average labor compensation share, gives labor-augmenting growth of 3.21 log percentage points per year. This informs own-innovation efficiency jointly with the other sources of growth. The empirical appendix gives the construction.

Firm turnover. I target 3.64 percent, averaging annual entry and net exit over 2015–2019. Entry averaged 3.37 percent and net exit 3.92 percent; their average accommodates the stationary model's equal flows. The model counterpart is the annual founding share among producing

[8]Conditional on incumbent launch scale, an inner step matches growth, turnover, entrant size and an approximation to mean product count. The common research-efficiency scale matches employment. An outer step updates this approximation and incumbent launch scale using simulated product counts and new-product size distributions.

firms, which depends on obsolescence, research and survival rather than obsolescence alone.[9]

Relative entrant size. Firms first observed in FE in 2018 have mean value added in 2019 equal to .207 of the all-age mean in their two-digit industry. This informs entrant launch scale. Model revenue corresponds to value added because the model has no intermediate inputs; its counterpart averages over founding dates and survival during the first full calendar year.[10]

Mean employment. The 2019 FE sample contains 21,488 manufacturing firms with positive employment and net sales, employing 511,666 workers: a mean of 23.8. Model employment per producing firm includes production and incumbent research, excluding prospective entrants and platforms without active products. This moment fixes the common research-efficiency scale and affects selection into the product sample.

Mean product count. The 3,283 firms in the 2019 IVP sample produce 2.95 positive-sales categories on average. Simulated scope counts products with positive sales during the year, including products that disappear before year-end. It informs expansion efficiency jointly with product acquisition and survival.

Size of new products. For additions by continuing firms in 1997–2020, I subtract the owner's mean log product sales in the following year ($t+1$) from each new product's log sales in that year. Deciles are constructed by pooling adjusted sales across all eligible products, firms and years, as in Fact B. Among positive classified followups in multiproduct firms, the third-decile share minus the tenth-decile share is 4.62 percentage points. This contrast summarizes the tilt toward smaller introductions without using the bottom two deciles. It informs incumbent launch scale jointly with capture, growth and survival; the model applies the same annual timing and normalization.

7.3 Parameter estimates and targeted fit

Table 3 reports the jointly estimated parameters. Changing the common scale ϕ_E changes all three research efficiencies while preserving the two efficiency ratios.

[9]The broad FE sample requires positive employment and net sales. Entry is first appearance in FE; net exit is inferred from entries and the change in the firm stock. These measures also reflect returns, reclassification and eligibility changes, so first appearance proxies founding.

[10]Eligible firms have positive employment and value added exceeding the wage bill; industry-year normalization cells require at least ten firms. Founding dates are uniform within the entry year. The model integrates cohort revenue and producing exposure over ages zero to two using triangular age weights, rather than evaluating a single age-one snapshot.

Table 3: Jointly estimated parameters and informative moments.

Parameter	Value	Informative target moment
Own-innovation efficiency, ϕ_I/ϕ_E	.007761	Annual productivity growth
Obsolescence rate, δ	.03924	Annual firm turnover
Entrant launch moment, m_E	.2443	Relative entrant value added
Common research-efficiency scale, ϕ_E	.07661	Mean producing-firm employment
Expansion efficiency, ϕ_X/ϕ_E	.09582	Mean products per firm
Incumbent launch moment, m_X	.6844	New-product size within firms at $t+1$: D3 minus D10

Notes: Moment pairings summarize joint information. Efficiency ratios and the common scale refer to the research technologies specified above.

7.3.1 Targeted moments

The first four moments in Table 4 are computed from the stationary equilibrium. Product counts and new-product size use simulated annual histories; the independent sample illustrates their simulation variation.

Table 4: Targeted moments and model fit.

Target moment	Data	Model	Independent simulation
Annual productivity growth, log percentage points	3.21	3.21	Same
Annual firm turnover, percent	3.64	3.64	Same
Relative entrant value added	.207	.207	Same
Mean producing-firm employment	23.8	23.8	Same
Mean products per firm	2.95	$3.03 \pm .07$	$2.85 \pm .08$
New-product size within firms at $t+1$: D3 minus D10, percentage points	4.62	$4.89 \pm .31$	$4.66 \pm .48$

Notes: The primary and independent simulations contain 40,000 and 20,000 firm families. Reported $\pm$ values are one Monte Carlo standard error from 100 family-cluster bootstrap draws. They exclude empirical and parameter uncertainty. All six moments enter estimation, so their fit is not an independent test of the model.

The benchmark also provides a check on moments used in the earlier fits. In the primary simulation, the standard deviation of log relative entrant sales is 1.54, against 1.53 in the data. The age-20 minus age-1 contrast in the log value-added cost ratio is 6.52 log percentage points, against 6.50; the independent simulation gives 6.10. The new-to-firm sales share is 4.40

percent, against 6.09, and the primary simulation understates growth-variance increments at horizons of two to five years. These comparisons assess how well the earlier fits carry over to the benchmark; they are not independent validation.

7.4 Additional model implications

Matching selected moments leaves differences elsewhere in the distributions. Single-product firms account for 47.8 percent of the data but 31.7 percent of the primary simulation (Figure 7), despite the close fit to mean scope.

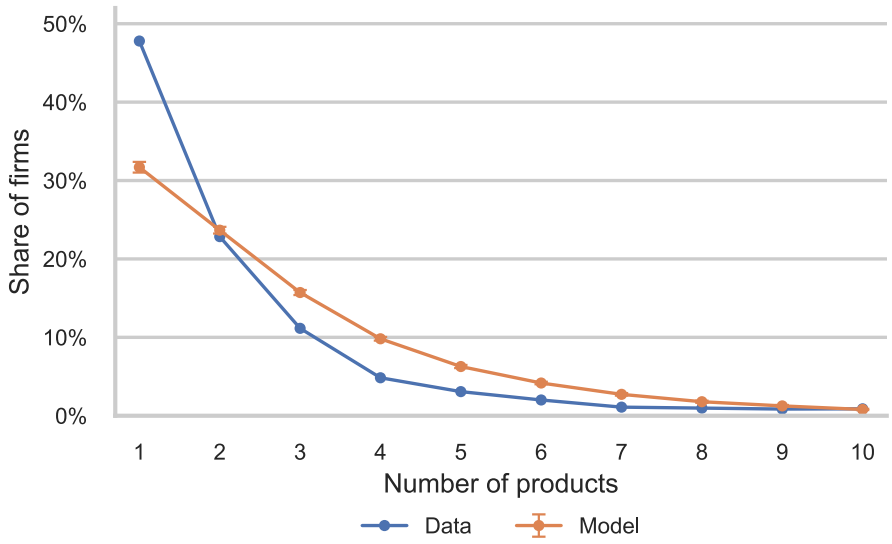
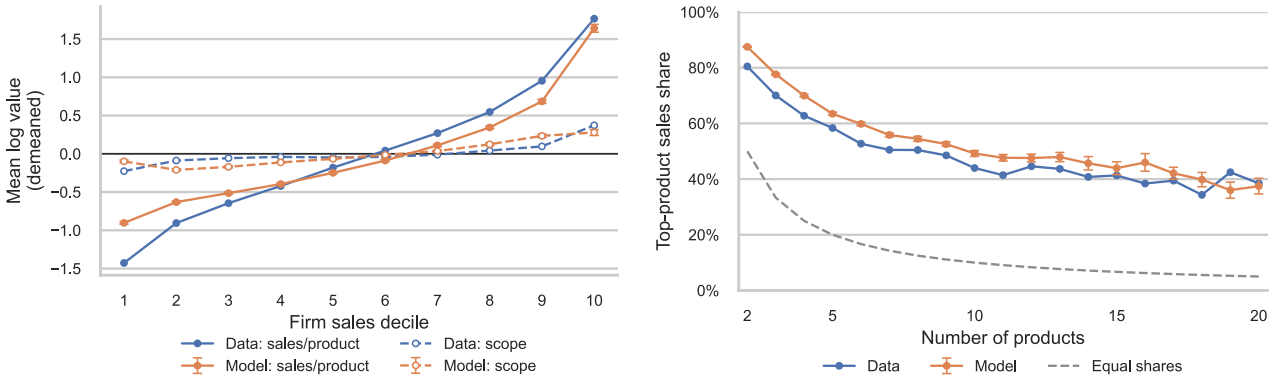


Figure 7: Distribution of products per firm in the 2019 IVP sample and the model. Shares use the full population, including firms with more than ten products. Model bars show ± one Monte Carlo standard error. Only the mean product count is targeted.

The following figures compare the model with the motivating facts, using the same empirical definitions and corresponding annual model observations. Within-firm panels use the same firm-year normalization in both samples. Across-product panels remove CN4–year effects in the data but use pooled model ranks, providing an approximate correspondence in the one-sector model.[11]

Firm size and portfolios. Sales per product account for more size variation than scope, but both profiles are flatter in the model (Figure 8). The top-to-bottom firm-decile gaps are 2.54 and .38 log points, compared with 3.19 and .60 in the data. Among two-product firms, the largest product accounts for 87.5 percent of model sales and 80.5 percent in the data.

[11]Model profiles use the primary simulation. For Fact A, the model weights eligible firm-years equally; empirical weighting follows the original figures. Model error bars use the same family-cluster bootstrap as Table 4 and reflect simulation error only. Each draw recomputes the centering mean for centered profiles and both event counts and their denominator for event shares.

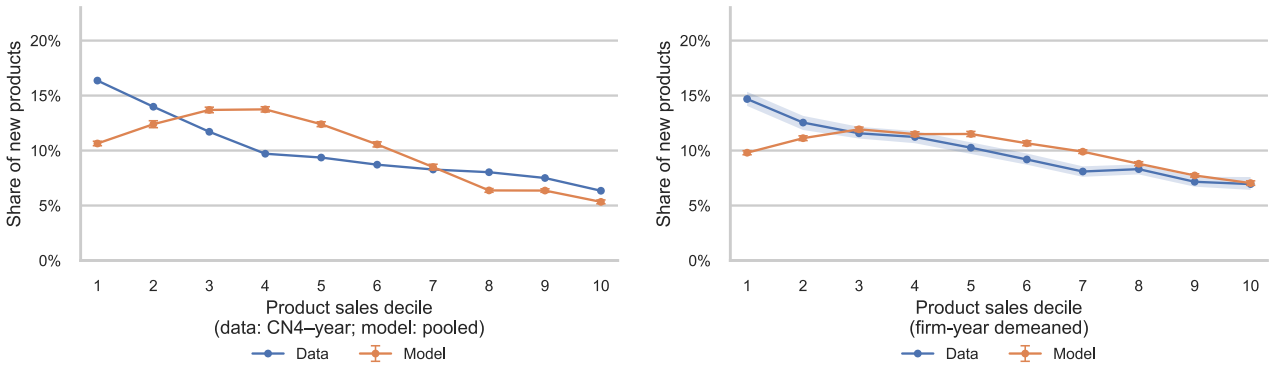


(i) Sales per product and scope

(ii) Top-product sales share

Figure 8: Model comparison with Fact A, 1997–2021 data. Panel i shows demeaned log sales per product (solid) and product count (dashed) across firm-sales deciles. Panel ii shows mean top-product shares; the dashed curve is $1/N$. Model bars show $\pm$ one Monte Carlo standard error. Empirical point estimates have no intervals.

New products. The model understates introductions in the smallest within-firm sales decile: 9.8 percent versus 14.7 percent (Figure 9). Matching the third-minus-tenth-decile contrast leaves this part of the distribution unfitted. Both profiles condition on followup observation.

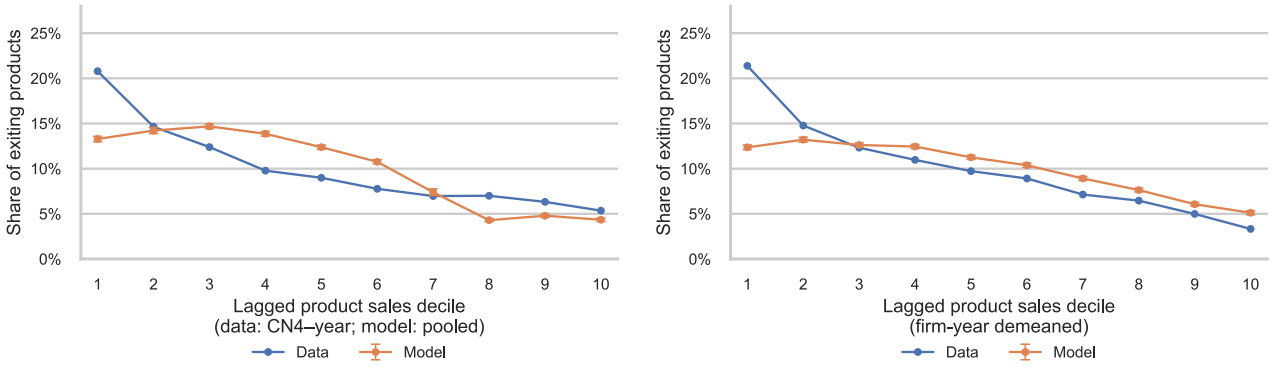


(i) Across products

(ii) Within firms

Figure 9: Model comparison with Fact B, 1997–2020 introduction cohorts. Sales are measured at $t+1$, conditional on positive classified sales at an eligible firm (multiproduct in ii). Each curve weights products equally and sums to 100 percent. Blue shading in ii shows pointwise 95 percent firm-cluster empirical intervals; model bars show $\pm$ one Monte Carlo standard error. Only the third-minus-tenth-decile contrast in ii is targeted.

Exiting products. The smallest within-firm decile accounts for 12.4 percent of model exits and 21.4 percent in the data (Figure 10). These are sales positions among exits, not exit probabilities conditional on size. The comparison conditions on continuing eligible owners and excludes losses accompanying firm disappearance. Neither exit profile is targeted.



(i) Across products **(ii)** Within firms

Figure 10: Model comparison with Fact C, 1999–2020 exit origins among continuing eligible owners. Sales ranks are measured in the year before the last observed year. Each curve weights exits equally and sums to 100 percent. Empirical point estimates have no intervals; model bars show ± one Monte Carlo standard error. Neither exit distribution is targeted.

7.5 Product selection and growth

For a given proportional quality gain, replacing a smaller product contributes less to aggregate growth. I compare the calibrated allocation with an accounting benchmark in which productive scale and protection are independent. The benchmark preserves both distributions, the research attempt rate and the innovation-step distribution, so replacement frequency and average proportional gains per successful replacement remain unchanged. Only the association between product scale and replacement gains changes.

Productive scale measures a product's weight in aggregate quality: $\nu = (a\tilde{q})^{\sigma-1}$, normalized so that $\mathbb{E}_P[\nu] = 1$. Observed sales also depend on markups and demand. Let $\Gamma(a)$ denote the expected proportional gain in CES-weighted quality per attempt, as defined in the quantitative growth decomposition. Both calculations use the count-weighted distribution P of all active products across firms, including the endogenous distribution of firm types. Creative destruction's contributions under the calibrated allocation and independent pairing are

$$g_{\text{CD}} = \frac{\lambda_H}{\sigma - 1}\mathbb{E}_P[\nu\Gamma(a)],$$
$$g_{\text{CD}}^{\text{ind}} = \frac{\lambda_H}{\sigma - 1}\mathbb{E}_P[\Gamma(a)].$$

Table 5: Selection into replacement and growth. Deterministic accounting at the calibrated parameters; replacement rates are event flows.

Quantity	Calibrated allocation	Independent pairing
Replacements per 100 product-years	2.319	2.319
Mean displaced-product scale (% of product mean)	28.2	100.0
Creative destruction's growth contribution (pp/year)	0.216	0.791

Replacement occurs at the same rate in both allocations, but displaced products in the calibrated model have, on average, only 28.2% of the mean productive scale (Table 5). Creative destruction contributes 0.216 percentage points to annual growth, compared with 0.791 under independent pairing, a difference of 0.58 percentage points. The comparison measures the accounting contribution of selection. Because the reassigned distribution need not satisfy equilibrium conditions, the gap is not a predicted growth gain from removing protection or changing policy.

8 Policy experiment

I use the calibrated model to study how an increase in Sweden's R&D payroll deduction affects research allocation, growth and welfare. I then compare alternative uses of the same long-run subsidy budget. Evidence that R&D tax incentives raise research expenditure motivates the exercise (Bloom, Griffith, and Van Reenen 2002).

Institutional setting. Sweden's *forskningsavdrag* reduces employer contributions for qualifying commercial R&D. In 2019, the deduction was 10 percent, capped at SEK 230,000 per corporate group per month. From January 2024, the rate is 20 percent and the monthly cap SEK 3 million, subject to the pension-contribution floor. I compare these two endpoints. The policy specifies eligibility as requiring systematic and qualified R&D work with a commercial purpose.[12]

Mapping to the model. For simplicity, I omit the group cap to preserve additive firm value and research policies independent of firm scale. The resulting proportional subsidy continues to lower marginal research costs for firms that would have otherwise reached the cap.

I construct a proxy for research covered by the subsidy using SCB data. In 2019, full-time equivalent internal R&D personnel represented 6.05 percent of manufacturing employment. Allocating this total using the 2013 expenditure shares by research purpose gives implied employment shares of 2.98 percent for expansion and 2.27 percent for own improvement.[13] I map both new-to-firm and new-to-market research to expansion. The combined R&D employment share is 5.25 percent; I exclude process research and general knowledge from the model mapping. This calculation assumes that the 2019 R&D mix resembles that of 2013, and that expenditure shares approximate labor shares across research purposes.

Dividing these shares by their model counterparts, 2.91 and 15.87 percent, gives fixed scaling weights $\kappa_X = 1.0255$ and $\kappa_I = 0.1430$. I assume that entry research receives no subsidy.

[12] Government Proposition 2019/20:68, Section 4.4; Skatteverket, *Forskningsavdrag*, checked 28 October 2026. Standard employer contributions of 31.42 percent and the pension-contribution floor of 10.21 percent permit at most 21.21 percentage points of relief. All three alternative designs exceed this limit; the rates near 38 percent also exceed total contributions.

[13] SCB's internal R&D personnel database reports 30,978 FTE in 2019; the administrative manufacturing sample has 511,666 employees. The survey mainly covers firms with at least ten employees, while the denominator includes smaller employers. Purpose shares use *Forskning och utveckling inom företagssektorn 2013*, Table 4, pp. 36–38, summed over manufacturing industries SNI 10–33. Of SEK 60,550 million in total R&D, SEK 11,962 million concerns products new to the firm, 17,871 million products new to the market, and 22,688 million improvements. Multiplying 30,978/511,666 by $(11{,}962 + 17{,}871)/60{,}550$ gives the expansion share; using 22,688/60,550 gives the own-improvement share. Product-directed research accounts for 86.74 percent of expenditure; the remaining implied employment share of 0.80 percentage points covers process research and general knowledge, excluded here by assumption.

The calibrated costs already incorporate the 10 percent deduction. With employer contributions of 31.42 percent, raising the headline rate from $\tau_0 = 0.10$ to $\tau_1 = 0.20$ reduces private research costs by

$$\xi_c = \frac{\kappa_c(\tau_1 - \tau_0)}{1.3142 - \kappa_c \tau_0}, \qquad c \in \{X, I\}.$$

The reductions are 8.46 percent for expansion and 1.10 percent for own improvement. The model assigns substantially more labor to own improvement than the empirical mapping implies, so the common headline increase translates into a smaller reduction in model costs for this channel. Resource costs and non-policy primitives remain fixed, and lump-sum taxes finance the incremental subsidy. The calibration absorbs the existing deduction into costs, approximating baseline private and resource costs as equal.

Alternative policies. I direct the additional subsidy toward expansion alone, own improvement alone, or research above the benchmark choice for each firm type and product state (Table 6). Each design matches the uniform reform's long-run subsidy expenditure relative to output, approximately 0.48 percent. All retain the baseline deduction. The alternatives are hypothetical subsidies that can exceed the statutory limits on payroll deductions.

The threshold rule uses the benchmark research choice per platform and product line, preserving additivity. I hold these state-specific rules fixed after the reform. Implementing them would require detailed information about firm types, product states and their benchmark research choices. Restricting additional support to research above these thresholds allows a higher headline rate at the same long-run fiscal cost: 38.9 percent in both channels. Research below the thresholds retains the baseline 10 percent rate.

Table 6: Policy comparisons with free entry and a common long-run subsidy cost of approximately 0.48 percent of output. Welfare is the consumption equivalent of the full transition; dashes denote uncomputed transitions. In the last row, the higher rate applies above state-specific benchmark research; the baseline 10 percent rate applies below.

Policy	Headline rates, X / I (%)	Growth change (pp/year)	Welfare gain (%)
Uniform reform	20.0 / 20.0	+0.088	+4.66
Expansion only	24.0 / 10.0	+0.082	—
Own improvement only	10.0 / 38.0	+0.104	—
Above state-specific benchmark research	38.9 / 38.9	+0.675	—

At the same long-run subsidy cost, support for own improvement raises growth slightly more than the uniform reform, while support for expansion raises it slightly less. The effects extend across research channels: the expansion-only subsidy raises own-improvement growth by 0.034 percentage points. Subsidizing research above state-specific benchmark choices produces a much larger growth increase.

Uniform reform. Annual growth rises by 0.088 percentage points. Entrants per year fall 8.9 percent, the number of products falls 2.8 percent, and profit per product rises 2.3 percent. Holding entrants per year fixed instead gives a growth decline of 0.089 points.[14] Allowing entry to adjust therefore changes the sign of the reform's long-run growth effect. The uniform-reform bars in Figure 12 show this sign reversal after normalizing by fiscal cost.

Most of the growth gain comes from improvements to existing products, while replacement remains concentrated on small products. Own improvement contributes 0.0575 percentage points to the increase, compared with 0.0084 points from creative destruction. Successful replacement rises from 2.32 to 2.43 events per 100 product-years. The mean productive scale of products displaced in successful replacements is about 28 percent of the economy-wide product mean in each equilibrium, using the measure defined in the accounting exercise.

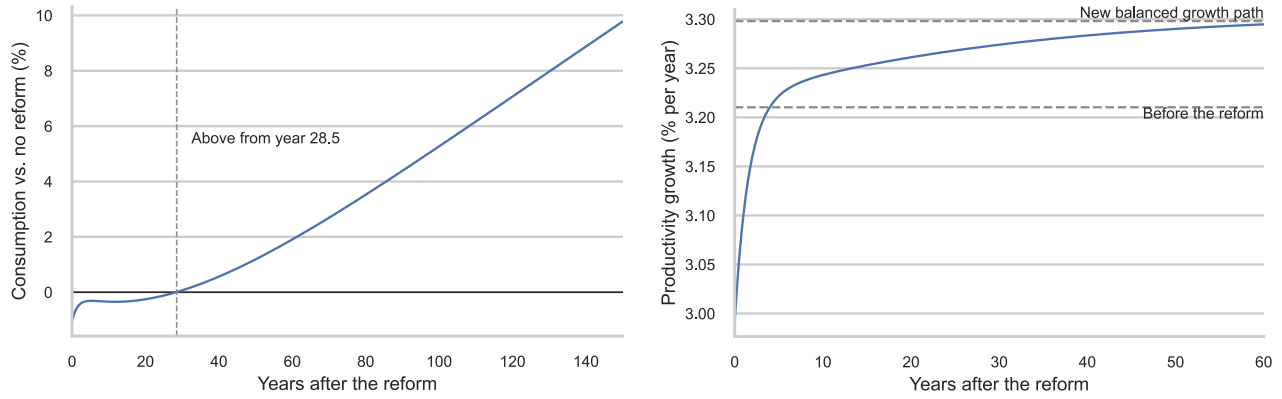
Transition and welfare. The long-run growth gain comes with an initial consumption loss. Following an unanticipated, permanent uniform reform, consumption falls about 1 percent as labor moves into research. Productivity growth also falls on impact and exceeds its pre-reform rate from year 4, but consumption overtakes its no-reform path only after about 29 years (Figure 11). Including this transition, the welfare gain is equivalent to a permanent 4.66 percent increase in consumption.

With log utility and $\rho = 0.012$, the consumption-equivalent welfare gain is

$$\lambda = \exp\left\{\rho \int_0^\infty e^{-\rho t} \log\left(\frac{C_t^{\text{reform}}}{C_t^{\text{no reform}}}\right) dt\right\} - 1.$$

This measure is the permanent proportional increase in no-reform consumption that gives the household the same discounted utility as the reform path, including its initial consumption loss.

[14] This diagnostic experiment replaces the free-entry condition with a fixed annual entrant flow at its benchmark value. Entry value can therefore differ from entry cost. The comparison shows how results depend on the assumption about entry adjustment.



(i) Consumption relative to no reform

(ii) Productivity growth

Figure 11: Transition following an unanticipated, permanent increase in the deduction from 10 to 20 percent, with free entry and lump-sum financing. Panel i plots $100 \log(C_t / C_t^{\text{no reform}})$; the vertical line marks consumption overtaking its no-reform path after about 29 years. Panel ii shows annual productivity growth; dashed lines mark the old and new balanced-growth rates. Deterministic model paths; no uncertainty bands.

Firm heterogeneity. The growth effect of a subsidy depends strongly on which firms receive it. I raise the headline deduction by 10 percentage points for one firm type at a time, where a higher p means that more of each innovation remains private. Targeting would require observing this latent share. These experiments keep the rate increment fixed and normalize growth effects by each design's fiscal cost, rather than matching budgets across types. Growth effects per fiscal krona rise with p under both free and fixed entry: subsidizing the lowest type reduces growth, while subsidizing the highest raises it (Figure 12). Improvements to existing products account for much of this difference. Their contribution falls for the lowest type and rises for the highest, where it more than offsets lower growth from new varieties and replacement.

The response extends beyond firms receiving additional support. Under free entry, subsidizing the lowest-p type reduces average own-improvement effort among other types, while subsidizing the highest-p type raises it.[15] These responses reflect changes in profits, entry and product composition, so they do not isolate the effect of displacement risk. The aggregate response therefore depends on changes in research both at firms receiving additional support and at other firms.

[15]Mean effort is weighted by product capacity within types, and its changes are weighted by each type's pre-reform capacity. These accounting terms include changes in product-state composition within types, rather than holding product states fixed. The pattern relates to the offsetting innovation responses under free entry in Atkeson and Burstein (2010).

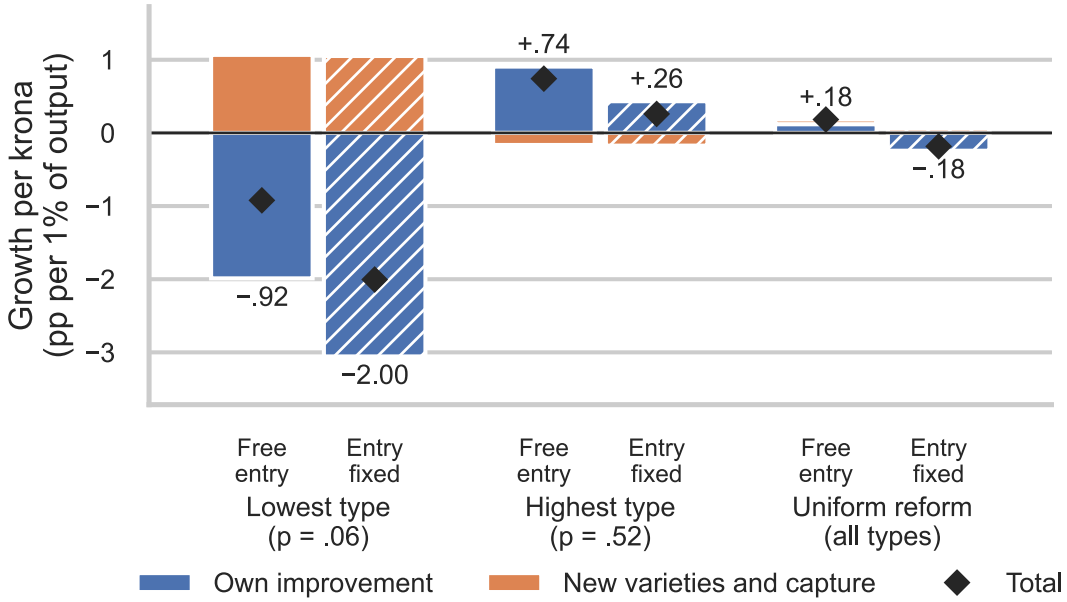


Figure 12: Growth effects per fiscal cost, in percentage points of annual growth per 1 percent of output spent. Blue bars show own improvement; orange bars show new varieties and capture. Diamonds show totals, and hatching denotes fixed entrants per year. Each design uses a 10 percentage-point rate increase and its own fiscal cost. Deterministic balanced-growth comparisons; no uncertainty intervals.

References

Acemoglu, Daron, and Ufuk Akcigit. 2012. "Intellectual Property Rights Policy, Competition and Innovation." *Journal of the European Economic Association* 10 (1): 1–42. https://doi.org/10.1111/j.1542-4774.2011.01053.x.

Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr. 2018. "Innovation, Reallocation, and Growth." *American Economic Review* 108 (11): 3450–91. https://doi.org/10.1257/aer.20130470.

Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter J Klenow, and Huiyu Li. 2025. "Good Rents Versus Bad Rents: R&D Misallocation and Growth."

Argente, David, Munseob Lee, and Sara Moreira. 2024. "The Life Cycle of Products: Evidence and Implications." *Journal of Political Economy* 132 (2): 337–90. https://doi.org/10.1086/726704.

Atkeson, Andrew, and Ariel Burstein. 2019. "Aggregate Implications of Innovation Policy." *Journal of Political Economy* 127 (6): 2625–83.

Atkeson, Andrew, and Ariel Tomás Burstein. 2010. "Innovation, Firm Dynamics, and International Trade." *Journal of Political Economy* 118 (3): 433–84. https://doi.org/10.1086/653690.

Baumgartner, Christoph, Stjepan Srhoj, and Janette Walde. 2023. "Harmonization of Product Classifications: A Consistent Time Series of Economic Trade Activities." *Jahrbücher Für Nationalökonomie Und Statistik* 243 (6): 643–62. https://doi.org/10.1515/jbnst-2022-0034.

Berlingieri, Giuseppe, Maarten De Ridder, Daniel Lashkari, and Davide Rigo. 2026. "Creative Destruction Through Innovation Bursts."

Bernard, Andrew B., Stephen J Redding, and Peter K Schott. 2011. "Multiproduct Firms and Trade Liberalization." *The Quarterly Journal of Economics* 126 (3): 1271–1318. https://doi.org/10.1093/qje/qjr021.

Bernard, Andrew B, Stephen J Redding, and Peter K Schott. 2010. "Multiple-Product Firms and Product Switching." *American Economic Review* 100 (1): 70–97. https://doi.org/10.1257/aer.100.1.70.

Bloom, Nick, Rachel Griffith, and John 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. https://doi.org/10.1016/S0047-2727(01)00086-X.

De Ridder, Maarten. 2024. "Market Power and Innovation in the Intangible Economy." *American Economic Review* 114 (1): 199–251. https://doi.org/10.1257/aer.20201079.

Gabaix, Xavier, and Rustam Ibragimov. 2008. "Rank-1/2: A Simple Way To Improve the OLS Estimation of Tail Exponents."

Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J. Klenow. 2019. "How Destructive Is Innovation?" *Econometrica* 87 (5): 1507–41. https://doi.org/10.3982/ECTA14930.

Hottman, Colin J., Stephen J. Redding, and David E. Weinstein. 2016. "Quantifying the Sources of Firm Heterogeneity *." *The Quarterly Journal of Economics* 131 (3): 1291–1364. https://doi.org/10.1093/qje/qjw012.

Jo, Karam, and Seula Kim. 2026. "Heterogeneous Innovations and Growth Under Imperfect Technology Spillovers." *Journal of Monetary Economics* 158 (March): 103899. https://doi.org/10.1016/j.jmoneco.2026.103899.

Klenow, Peter J., and Huiyu Li. 2021. "Innovative Growth Accounting." *NBER Macroeconomics Annual* 35 (May): 245–95. https://doi.org/10.1086/712325.

———. 2025. "Entry Costs Rise with Growth." *Journal of Political Economy Macroeconomics* 3 (1): 43–74. https://doi.org/10.1086/733979.

Klette, Tor Jakob, and Samuel Kortum. 2004. "Innovating Firms and Aggregate Innovation." *Journal of Political Economy* 112 (5): 986–1018.

Lentz, Rasmus, and Dale T. Mortensen. 2008. "An Empirical Model of Growth Through Product Innovation." *Econometrica* 76 (6): 1317–73. https://doi.org/10.3982/ECTA5997.

Luttmer, Erzo G. J. 2007. "Selection, Growth, and the Size Distribution of Firms." *The Quarterly Journal of Economics* 122 (3): 1103–44. https://www.jstor.org/stable/25098869.

Mayer, Thierry, Marc J. Melitz, and Gianmarco I. P. Ottaviano. 2014. "Market Size, Competition, and the Product Mix of Exporters." *American Economic Review* 104 (2): 495–536. https://doi.org/10.1257/aer.104.2.495.

Melitz, Marc J. 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71 (6): 1695–725. https://doi.org/10.1111/1468-0262.00467.

Peters, Michael. 2020. "Heterogeneous Markups, Growth, and Endogenous Misallocation." *Econometrica* 88 (5): 2037–73. https://doi.org/10.3982/ECTA15565.

Peters, Michael, and Conor Walsh. 2026. "Population Growth and Firm Dynamics." *Journal of Political Economy Macroeconomics* 4 (2): 256–96. https://doi.org/10.1086/740426.

Shi, Chenchuan. 2025. "Knowledge Generality, Competition and Growth."

A Supplementary evidence on the data

The 1997–2021 IVP–FE baseline contains 89,721 firm-years. For positive sales r_{fjt} of concorded CN8 product categories (CN8+), define firm sales $S_{ft} = \sum_j r_{fjt}$, product count N_{ft}, and sales per product $A_{ft} = S_{ft}/N_{ft}$.

A.1 Firm-size and product-scope tails

Define $z^X_{ft} = \log X_{ft} - \overline{\log X}_{g(f,t),t}$ for $X \in \{S, A, N\}$, where the bar averages baseline firm-years in owner industry g and year t. In each variable's own pooled upper 10%, estimate

$$\log(i/n) = c_X - \alpha_X z^X_{(i)} + u_i,$$

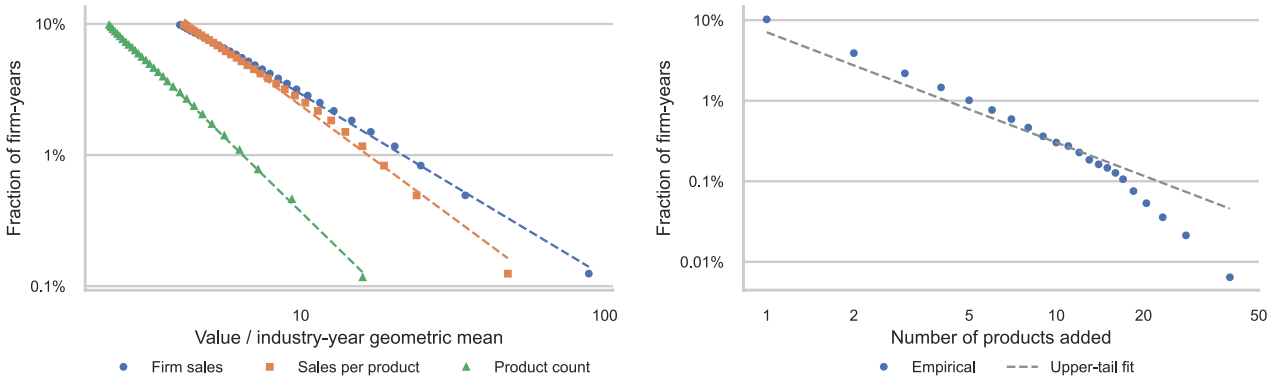
by OLS, with slope $-\alpha_X$, intercept c_X, descending rank i and $n = 89{,}721$. Cutoff ties remain included and retain ordinary ranks; industry-year means are fixed before selection. Figure 13 (i) shows complementary cumulative distributions (CCDFs) and descriptive fits.

A.2 Bursts of product additions

Following Berlingieri et al. (2026) (Table C3), let B_{ft} count categories positive at t but not at $t-1$, including returning product codes. For 76,330 equally weighted firm-year transitions in 1998–2021, I estimate by OLS

$$\log[(i - 1/2)/n_B] = c_B - \alpha_B \log B_{(i)} + v_i,$$

where i ranks additions from largest to smallest and $n_B = 76{,}330$ includes zeros. The subtraction of one-half follows Gabaix and Ibragimov (2008). Ties at one addition expand the upper 10% to 10.23%. Averaging across years, the top 5% of firms ranked by the number of product code additions account for 80.0% of additions.



(i) Firm-size and product-scope tails

(ii) Annual product additions

Figure 13: Tail distributions on log-log axes. Markers aggregate ranks; dashed lines fit underlying observations. Panel i: upper-10% slopes are -1.39 for firm sales, -1.71 for sales per product and -2.26 for product count. Panel ii: slope -1.37; zeros remain in the probability denominator.

A.3 Product turnover and employment

I relate portfolio changes to employment growth, $\Delta \log L_{ft} = \log L_{f,t+1} - \log L_{ft}$, using 76,330 equally weighted adjacent-year transitions at 7,982 firms with origins in 1997–2020. Both endpoints must satisfy the IVP–FE baseline and have resolved product sales. An addition is a CN8+ category with positive sales at $t+1$ but not t; a drop reverses these conditions. Additions include returns and drops include temporary disappearances, unlike the identified births and terminal disappearances in Facts B–C.

Table 7 reports group means and conditional differences from unchanged portfolios. For the latter, I estimate by OLS

$$\begin{aligned}\Delta \log L_{ft} = \lambda_{g(f,t),t} + \beta_{+} I^{+}_{ft} + \beta_{-} I^{-}_{ft} + \beta_{\pm} I^{\pm}_{ft} \\ + \gamma_L \log L_{ft} + \gamma_N \log N_{ft} + \varepsilon_{ft}.\end{aligned}$$

The mutually exclusive indicators I^{+}_{ft}, I^{-}_{ft} and $I^{\pm}_{ft}$ denote additions only, drops only and both, respectively; unchanged portfolios are omitted. Controls are origin-year employment L_{ft}, product count N_{ft} and industry-year effects $\lambda_{g(f,t),t}$. The table reports 100β with pointwise 95 percent confidence intervals clustered by firm. Added-sales shares divide added categories' sales by total product sales at $t+1$; dropped-sales shares use sales at t. These ratios are averaged equally across transitions within each group.

Table 7: Product turnover and employment among continuing eligible firms. Growth and conditional differences are in log units multiplied by 100; sales shares are percentages.

Product changes	Transitions	Mean growth	Conditional difference [95% CI]	Added sales (%)	Dropped sales (%)
Unchanged	64,206	0.74	Reference	0.0	0.0
Additions only	3,269	4.80	4.46 [3.59, 5.32]	18.0	0.0
Drops only	4,318	−2.06	−2.37 [−3.15, −1.59]	0.0	12.8
Add and drop	4,537	2.04	0.79 [0.03, 1.55]	48.5	47.5

These associations concern net changes in firm employment, not observed worker movements between products or causal effects of turnover. Requiring continued eligibility can exclude contractions that take firms below the survey threshold or other baseline restrictions.

A.4 Product sales over the life cycle

The product-sales profile in the quantitative model follows a balanced sample of continuing firm-product lines. I estimate an annual counterpart to Argente, Lee, and Moreira (2024), equation (2):

$$\log r_{fjt} = \lambda_{g(f,t),t} + \gamma_{c(f,j)} + \sum_{a=2}^{10} \beta_a \mathbf{1}\{t - c(f,j) = a\} + \varepsilon_{fjt}.$$

Here r_{fjt} is sales of firm f in harmonized CN8 category j, $c(f,j)$ is the product's birth year, a is calendar age, and λ denotes owner-industry-year effects. The coefficient β_a measures the log-sales difference from age one; the potentially partial launch year is omitted. Because age equals calendar year minus birth year, I identify the profile by normalizing cohort effects to have zero mean and assuming no linear cohort trend: $\sum_c \gamma_c = \sum_c (c - \bar{c})\gamma_c = 0$, with equal weights across observed cohorts.

Identified births in 1997-2011 require prior-year firm reporting and positive, eligible launch sales. I retain product lines observed continuously with positive sales and an eligible owner from launch through age ten, using histories through 2021. A product absence, nonpositive or missing sale, reporting gap or accounting failure ends the spell; later returns do not fill gaps. Product lines affected by an industry-year cell containing a single observation are removed in full, repeating until no such cell remains. The final sample contains 2,507 product lines at 1,254 firms, each contributing all ten estimated ages. OLS gives equal weight to the resulting 25,070 product-years, and confidence intervals cluster by harmonized CN8 category.

The estimated age-ten contrast is 0.804 log points, with a pointwise 95 percent confidence

interval of [0.516, 1.092]. This is an adjusted age profile under the stated cohort normalization, rather than raw cumulative sales growth. Since the same products contribute at every age, the rising profile cannot be explained by changing sample membership. It describes products that remain observed and eligible through age ten, rather than all product launches. Demand, markups and changes in the goods sold within a category can also affect sales, so the profile does not isolate technological improvement. Annual industrial categories and owner-industry controls also differ from the quarterly retail barcodes and product-market controls in Argente et al.

A.5 Simplified HRW decomposition

Following Hottman, Redding, and Weinstein (2016) (HRW, equations 18–22), define geometric-mean sales $G_{ft} = \exp[N_{ft}^{-1} \sum_j \log r_{fjt}]$ and dispersion $D_{ft} = \log(A_{ft}/G_{ft})$. The identity $\log S = \log N + \log G + D$ implies

$$\widetilde{x}_{k,ft} = a_k + b_k \widetilde{\log S}_{ft} + e_{k,ft}, \qquad \sum_k \widehat{b}_k = 1,$$

where $x_k \in \{\log N, \log G, D\}$ and tildes remove baseline industry-year means. Equal-weight OLS slopes equal $\widehat{\mathrm{Cov}}(\widetilde{x}_k, \widetilde{\log S}) / \widehat{\mathrm{Var}}(\widetilde{\log S})$, allocating covariance equally. In HRW, assuming negligible market shares and equal substitution elasticities within and across firms, gives unit scope weight and no cannibalization. Sales alone cannot separate product appeal from process efficiency.

Table 8: HRW accounting shares (point estimates). The last two rows split sales per product. Tail regressions use 8,973 sales-selected firm-years and baseline means; shares describe within-tail variation.

Component	Baseline (%)	Sales top 10% (%)
Product count, $\log N$	14.7	32.2
Sales per product, $\log A$	85.3	67.8
Of which: geometric mean, $\log G$	71.7	44.4
Of which: dispersion, D	13.6	23.5

A.6 Revenue productivity and the value-added cost ratio

The firm-age profiles in the quantitative model use two revenue-based proxies from FE. Let W_{ft} be the wage bill and $\mathcal{K}_{ft}$ machinery and equipment book value. Define

$$R_{ft} = \frac{VA_{ft}}{W_{ft} + r\mathcal{K}_{ft}}, \qquad r = .15,$$
$$\text{TFPR}_{ft} = \frac{VA_{ft}}{\mathcal{K}_{ft}^{\beta_{st}} W_{ft}^{1-\beta_{st}}}, \qquad \beta_{st} = \frac{\sum_{f\in(s,t)} r\mathcal{K}_{ft}}{\sum_{f\in(s,t)} (W_{ft} + r\mathcal{K}_{ft})}.$$

The capital-cost share β_{st} uses all-age eligible firms with nonnegative capital in industry s and year t, before life-cycle restrictions. Capital is lagged when the preceding year is observed and otherwise contemporaneous; a missing observed lag stays missing. R&D payroll remains in W. The cost ratio permits zero capital, whereas TFPR requires positive capital, giving different samples.

Estimation. I regress each log proxy on firm effects, year effects and age indicators for ages zero through twenty. Firms require uncensored FE entry and at least two eligible observations. TFPR uses 157,835 firm-years from 22,316 firms; the cost ratio uses 181,976 firm-years from 24,928 firms. Each sample spans 24 calendar years and 23 entry cohorts and differs from the IVP–FE product sample used elsewhere in this appendix.

Since age equals calendar year minus entry cohort, identifying its linear trend requires a normalization: equally weighted firm effects have zero linear trend across entry cohorts. The figure reports $100(\gamma_a - \gamma_1)$, where γ_a is the age coefficient and age one the first full financial year. Pointwise 95 percent intervals use the firm-clustered covariance of each contrast, including covariance with the age-one coefficient. The profiles depend on the effects, sample selection and normalization and do not identify causal effects of aging.

Markup interpretation. Suppose production is constant-returns Cobb–Douglas with capital elasticity β, cost minimization, flexible inputs and competitive factor prices. Let $C = W + r\mathcal{K}$ denote variable cost. When value added measures output revenue and C includes all variable costs, $VA/C = \mu$. Cost minimization gives $r\mathcal{K} = \beta C$ and $W = (1-\beta)C$, so

$$\text{TFPR} = \frac{VA}{\mathcal{K}^{\beta} W^{1-\beta}} = \mu \frac{r^{\beta}}{\beta^{\beta}(1-\beta)^{1-\beta}}.$$

At fixed β and r, TFPR moves proportionately with the markup; physical productivity cancels. An efficiency gain at an unchanged markup need not raise TFPR. Industry-year changes in β, input adjustment, overhead and R&D labor, capital measurement and multiproduct composition can affect this benchmark mapping. Neither proxy directly measures permanent firm productivity.

A.7 Manufacturing productivity growth

The aggregate growth moment in the estimation section uses the EU KLEMS national manufacturing aggregate for Sweden (NACE Rev. 2 section C). This population differs from the IVP firm-product sample. I convert annual TFP growth to labor-augmenting units using the average labor compensation share in the origin and destination years. Writing g_t^{TFP} for the annual log change and $s_{L,t}$ for labor compensation divided by nominal value added,

$$\bar{s}_{L,t} = \frac{s_{L,t-1} + s_{L,t}}{2}, \qquad g_{Q,t} = \frac{g_t^{\mathrm{TFP}}}{\bar{s}_{L,t}}.$$

The twenty annual changes ending in 2000–2019 receive equal weight. Their mean is .0321024, or 3.2102 log percentage points per year. Figure 14 shows the annual observations and this mean, which is the model's growth target. The estimation matches the mean rather than annual fluctuations. The TFP residual can reflect utilization, reallocation and measurement as well as innovation, so this conversion does not identify the contribution of product innovation alone.

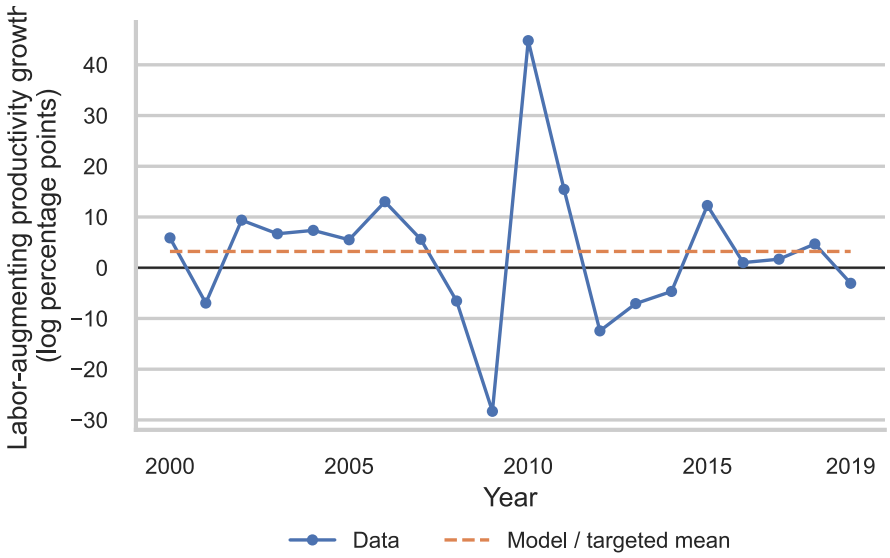


Figure 14: Annual manufacturing productivity growth. EU KLEMS TFP growth is converted to labor-augmenting units using adjacent-year average labor compensation shares. The dashed line is the equally weighted 2000–2019 mean, used as the model's growth target. No sampling intervals are shown.