Private Know-How, Creative Destruction, and Growth

Filip Mellgren
IIES, Stockholm University

17 November 2026

Latest version

Abstract

How important is product churn for aggregate growth? In this paper, I use Swedish manufacturing microdata to show that products added or dropped by a firm tend to be small in terms of sales. The benchmark framework of multiproduct innovating firms in Klette and Kortum (2004) abstracts from differences in product size, and higher product churn is linked to a higher aggregate growth rate. To bring the benchmark closer to the data, I introduce a parameter controlling the degree of private know-how, which introduces a competitive moat that shields incumbents from business stealing. In two economies calibrated to the same aggregate growth rate, own improvement accounts for about 60% of gross growth without private know-how and 83% with it. Creative destruction accounts for about 30% versus 2%, and new varieties for about 11% versus 15%.

1 Introduction

Technology growth is a major contributor to rising living standards. Leading theories point to creative destruction (Aghion and Howitt 1992; Grossman and Helpman 1991); the creation of new product varieties (Romer 1990); and the improvement of existing product varieties by incumbent firms. While much innovative activity may be specialized and product-specific, most empirical evidence relies on firm-level data. In this paper, I use firm-product panel data to motivate a new mechanism in an endogenous growth model: private know-how makes larger products harder to replace.

The Swedish manufacturing survey *Industrins Varuproduktion* records production value at the 8-digit product code level. I use these data to show three motivating facts. **Fact A:** sales per product, not product count, accounts for most variation in firm-level sales, and firm-level sales tend to be dominated by a single product category. **Facts B and C:** about 60% of newly added products and 69% of dropped products lie below the respective pooled median log sales after removing firm-year means from log sales. The highest sales firms thus tend to operate product portfolios consisting of one core product and a fringe of lower sales products that churn at a higher rate.

To rationalize these descriptive facts about manufacturing firms, I augment the canonical framework of Klette and Kortum (2004) with private know-how. In the model, private know-how grows with the size of a product and has the property that it is firm-product specific; a competitor is unable to learn from private know-how in order to improve the quality of incumbents' products. Setting private know-how to zero nests standard extensions to the Klette and Kortum (2004) framework.

In my model with private know-how, creative destruction contributes little to growth for two reasons. First, quality-improving inventions can fail to materialize. This stands in contrast to the benchmark model where all inventions are always successfully implemented. Second, while the calibrated private know-how model features a higher rate of creative destruction attempts, the quality-improving inventions that succeed tend to outcompete producers of small product categories with little importance for the macroeconomy.

The reason inventions fail to materialize in my model is that competitors don't have access to the private know-how of incumbent producers. A new invention thus needs to qualify by being good enough that the competitor can replace the incumbent without having access to the know-how of the incumbent. The hurdle a competitor needs to overcome rises with the total amount of quality embodied in a product, which also determines the sales. Consequently, some inventions are lost, and most inventions that materialize are applied to products that contribute less to aggregate growth for a given proportional quality improvement.

While private know-how results in lost inventions, the fact that it protects the incumbent producer strengthens incentives for the incumbent to invest in the quality of its own products. In my model, I therefore include own-innovation as a firm choice following Peters (2020). These incentives change the firm dynamics. In the benchmark model, firms grow large because they make an invention over an already high-quality product produced by a competitor. With private know-how, firms grow large because their products keep growing. In the data, adjusted log sales rise by 0.80 between ages one and ten among products observed continuously through age ten, controlling for cohort and industry-year effects and assuming no linear cohort trend.

The two economies also differ in which products dominate their largest firms. In the benchmark model, the sales of the largest firms tend to be dominated by recently added products that are large due to accumulated quality inventions by competing firms in the past. With private know-how, the largest firms instead operate portfolios dominated by a product the same firm has perfected over a long time period. Consistent with this mechanism, among Swedish multiproduct firms whose product ages can be identified, leading products tend to have been produced by the same firm for longer than the other products in their portfolios.

To make quantitative predictions, I estimate parameters for the benchmark model and the private know-how model using firm-level and aggregate moments. These moments include the firm turnover rate, entrant median size relative to mean firm size, the dispersion in size among entering firms, the aggregate growth rate, as well as the average product count and the share of firms with at least 20 employees. The two economies give similar predictions for these targeted firm-level moments. The composition of growth and the trajectory of accumulated quality inside firms are what differ between the two quantified models.

Quantitatively, I find that creative destruction accounts for 29.4% of growth before obsolescence without protection, versus 1.9% with protection. Own improvement accounts for 60.1% versus 82.7%, and new varieties for 10.5% versus 15.4%. Among the largest 1% of products, 11.7% versus 86.7% have never changed owner. Meanwhile, the mean age among the largest products is similar, 56.3 versus 57.5 years, but the current owner has held them for 12.3 versus 52.8 years.

The distinction also matters for innovation policy. I consider a permanent subsidy covering 10% of research costs for both expansion and own improvement. This policy raises growth and welfare in both calibrated economies, with consumption-equivalent welfare gains of 23.0% in the benchmark and 31.2% with private know-how. Subsidies to each activity separately clarify the channels: own-improvement support raises growth and welfare in both economies, while expansion support yields small gains in the benchmark but losses with protection, as gains from replacing small products fail to offset reduced own improvement. These experiments compare equal proportional cost reductions, with different fiscal costs.

2 Relation to literature

The empirical analysis builds on product-level studies of firm size and growth. A. 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. Selection across firms and within product portfolios also features in trade models (Melitz 2003; A. B. Bernard, Redding, and Schott 2011; Mayer, Melitz, and Ottaviano 2014). I connect the empirical trade literature's evidence on product selection within firms to a growth framework in which that selection shapes the contribution of creative destruction.

In the endogenous growth literature, previous studies that make use of product-level data include Argente, Lee, and Moreira (2024) and Berlingieri et al. (2026). 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. I document similarly concentrated product additions in Swedish manufacturing in the empirical appendix. 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. Relative to these two papers, I use evidence on the size of products added to and dropped from firms' portfolios to motivate a mechanism linking product size to the risk of being outcompeted.

My paper also relates to growth accounting that distinguishes creative destruction, new varieties and improvements to existing products (Garcia-Macia, Hsieh, and Klenow 2019; P. Klenow and Li 2021). Klenow and Li use employment at entering, exiting and continuing plants to infer these channels. In their benchmark, small plants fall below a viable relative-quality threshold as the economy advances. In my calibrated protected economy, small products face greater replacement risk even though their owners invest enough to increase their relative size.

Recent studies link imperfect knowledge diffusion to competition and the composition of innovation (Jo and Kim 2026; Shi 2025). I associate private know-how with relative product size and assume that a replacing producer does not inherit it. In the calibrated protected economy, large products develop mainly through improvements by their existing owners, while creative destruction concentrates on smaller products. I isolate and quantify how the association between product size and replacement gains reduces creative destruction's contribution to growth, even at a given frequency of replacement.

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 technological leads over competitors. This mechanism features in my

paper as protected products are less likely to be replaced, meaning that quality improvements on a firm's own products generate profits for longer. De Ridder (2024) shows how firm-specific advantages in deploying intangible capital can prevent higher-quality competitors from taking over products. Here, different products *within a firm* face different replacement risks. This has new implications for how firms reach the far end of the size distribution: not via firm-level process efficiency investments, but via product development efforts.

My paper also relates to research on how heterogeneity and research allocation shape growth and innovation policy (Lentz and Mortensen 2008; Acemoglu et al. 2018; Aghion et al. 2025). Relative to these papers, my focus on product heterogeneity provides an additional channel: in the calibrated protected economy, subsidies that help a firm increase its portfolio scope *lower* growth and welfare because the additional gains from capturing small products do not compensate for reductions in own product development.

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.

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* 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

[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.

[2]I extract the eight-digit CN8 code from the raw product code and harmonize categories over the sample

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.

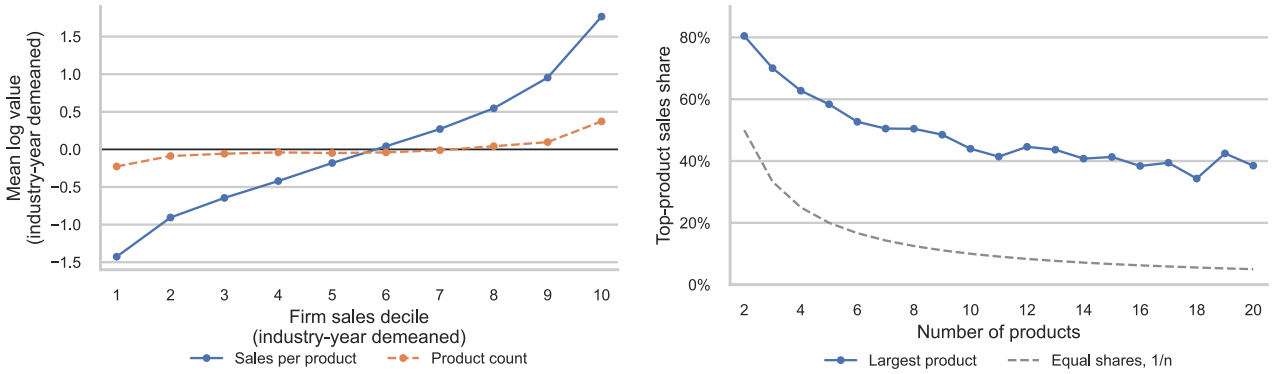
The baseline panel covers 9,518 firms, 89,721 firm-years and 254,409 firm-product-years over 1997–2021. Across firm-years, firms produce 2.84 products on average, and 45.9% of firm-years have a single product.

4 Facts on product size and turnover

Product turnover accompanies employment adjustment within continuing firms. The empirical appendix shows 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.

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).

period using the R package *harmonizer* (Baumgartner, Srhoj, and Walde 2023). Unmatched codes retain their original classification.



(i) Sales per product and product count

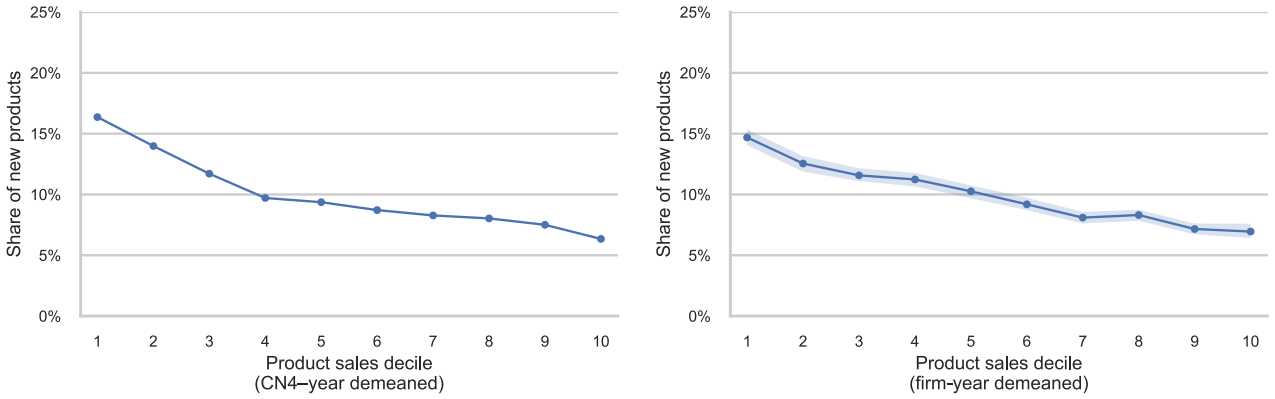
(ii) Top-product sales share

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 A. 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.

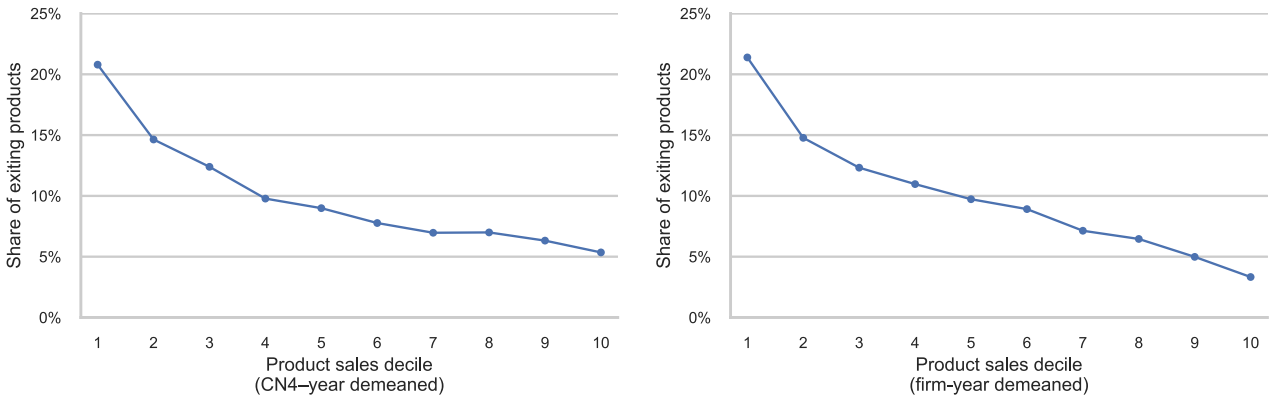


(i) Across products: CN4–year effects removed **(ii)** Within firms: firm-year effects removed

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 model builds on the framework of Klette and Kortum (2004). The motivating facts document substantial sales differences within firms and show that recently added and dropped products tend to be small. A first departure from Klette and Kortum (2004) is to relax the assumption that consumers have Cobb-Douglas preferences. Relaxing this assumption is crucial for a non-degenerate sales distribution. Instead, I introduce a competitive final goods aggregator with constant elasticity of substitution (CES) technology. 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. I let this protection increase with product size relative to the average product.

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. Discovery steps are drawn from a common distribution that allows some attempts to overcome installed know-how and others to fail. Following Peters and Walsh (2026), I also allow firms to improve their existing products and create new varieties, placing creative destruction alongside other sources of growth.

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,

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

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 firms 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 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 $\sigma > 1$. This assumption can be microfounded by a two-stage participation and pricing game described in Acemoglu et al. (2018).

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 Discovery steps and firm-specific know-how

In the model, all firms share a common size-protection rule indexed by $\theta \geq 0$:

$$a = a_\theta(\tilde{S}) \geq 1, \qquad a'_\theta(\tilde{S}) \geq 0, \qquad a_0(\tilde{S}) = 1.$$

Larger products therefore carry weakly more firm-specific know-how. The rule depends on quality-based productive capacity, measured by $\tilde{S} = (qa/Q)^{\sigma-1}$, rather than a discretionary choice of prices or reported sales. I assume the rule is differentiable and satisfies

$$(\sigma-1)\tilde{S}\frac{a'_\theta(\tilde{S})}{a_\theta(\tilde{S})} \leq 1,$$

so transferable quality $q = Q\tilde{S}^{1/(\sigma-1)}/a_\theta(\tilde{S})$ is weakly increasing in total quality at a given Q. Own innovation, introduced below, can raise total quality during an ownership spell. Protection rises when a product grows relative to the economy and falls when it falls behind. Thus q and a are derived from the same product-size state.

For the quantitative comparison, I use

$$a_\theta(\tilde{S}) = \left(1 + \tilde{S}^{1/(\sigma-1)}\right)^\theta, \qquad 0 \leq \theta \leq 1.$$

This rule satisfies the restriction that transferable quality be weakly increasing in total quality.

Each expansion attempt draws a positive log discovery step $d \sim G$, independently of the innovating firm and its target. The corresponding proportional discovery step is $e^d - 1$. The common distribution of log discovery steps satisfies $\mathbb{E}_G[e^{(\sigma-1)d}] < \infty$.

Attempts directed at existing varieties sample their targets uniformly. If an innovation arrives and the incumbent product has quality $q'a'$, the quality achievable by the new invention is

$$q^+ a^+ = e^d q'.$$

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 log discovery step d therefore measures an improvement over transferable quality.

Attempts that do not exceed installed know-how fail; ties leave the incumbent in place. Failed attempts create neither a product nor a platform.

On success, relative sales become

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

The new product's size determines its protection and transferable quality:

$$a^+ = a_\theta(\tilde{S}^+), \qquad q^+ = \frac{q' e^d}{a^+}.$$

5.1.3 Research platforms and research effort

Firms may operate several products and a stock of research platforms, with one platform per active product. Let $\mathcal{A}_f$ be the set of active varieties supplied by firm f, and $n_f = |\mathcal{A}_f|$ its active product count, which also equals its stock of research platforms.

Each successful product introduction adds one platform to the innovating firm. Losing a product to a competitor ends its operating profits and removes its research platform from the displaced firm. Successful capture therefore transfers a platform without changing the aggregate platform stock. The product and platform generating the attempt remain with the innovating firm.

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 $n_f x_f$. Its research labor requirement is

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

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 or firm-specific know-how.

Varieties become obsolete at an exogenous rate δ, removing the active product and its research platform. A firm exits when it loses its last active product, whether through capture or obsolescence.

5.1.4 Incumbent variety creation

A product new to a firm may outcompete another firm's product or open a new market. Following Peters and Walsh (2026), an incumbent research attempt targets an existing variety 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 is $\ell_f^X = n_f x_f^2/\phi_X$, where the firm chooses the Poisson arrival rate x_f per active product while taking α_X as given. An attempt directed at an existing variety follows the capture rule, $e^d > a_\theta(\tilde{S}')$. A new-variety attempt always introduces a product and one research platform. Entrants create new varieties only.

Both successful routes add a product and its research platform to the innovating firm. A capture also removes the target's product and platform from its previous owner, leaving the aggregate stock unchanged. The value acquired depends on the new product's size, and attempts at existing varieties may fail.

The mass of varieties and incoming capture attempts satisfy

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

5.1.5 Entry and launch quality

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

After paying the entry cost, an entrant introduces a product in a new variety along with a new research platform.

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 8.5 percent. This complements the new-product landing evidence, which describes products introduced by firms already active in IVP. I allow new products to start above or below the quality mean, with separate launch scales for entrants and incumbents.

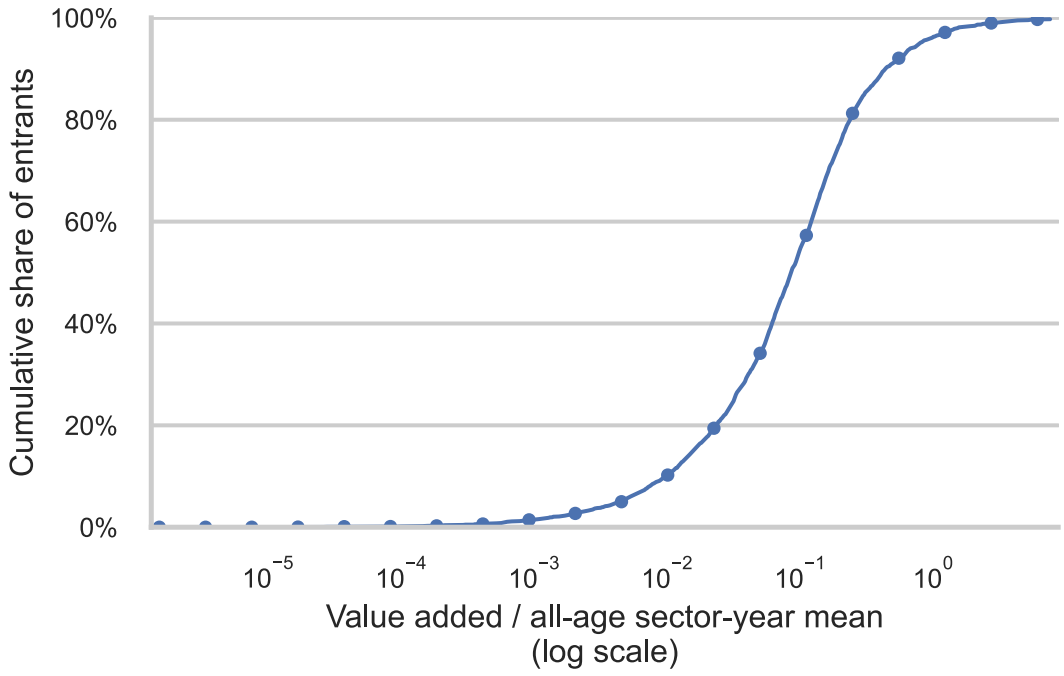


Figure 4: Entrant value added in the first full year, 2015–2019. Empirical cumulative distribution for 3,726 manufacturing entrants with positive employment and value added. Value added is divided by the all-age mean in the same sector and year. Years receive equal weight and entrants equal weight within year. First appearance in FE defines age; age one avoids the entry year's partial exposure.

For quantitative analysis, the proportional discovery step $e^d - 1$ follows a lognormal distribution with a fixed upper truncation. Both channels share the induced distribution G of log discovery steps.[3]

[3]The upper truncation is at the 99.99th percentile of the underlying lognormal distribution, with probabilities renormalized. The distributional parameters and cutoff are assigned and common across the two economies.

A new variety, introduced by an entrant ($c = E$) or an incumbent ($c = X$), draws a log discovery step $d \sim G$ and an independent quality draw ω. Total quality is

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

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:[4]

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

Thus m_E and m_X govern average launch scale for entrants and incumbents, while s_ω governs dispersion conditional on the log discovery step. Under constant markups, normalized CES-weighted quality equals relative sales. The quality draw is embodied in installed quality and is not redrawn during the product's life.

A new product's relative sales and protection are

$$\tilde{S}_c^+ = (\bar{\omega}_c \omega e^d)^{\sigma-1}, \qquad a_c^+ = a_\theta(\tilde{S}_c^+), \qquad q_c^+ = \frac{\bar{\omega}_c \omega e^d Q}{a_c^+}.$$

Thus a larger launch draw produces a larger and weakly better protected product. The quality draw is embodied in the product, while protection subsequently adjusts with its relative size. Successful captures instead inherit their target's transferable quality and follow the capture law.

5.1.6 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 P. 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.[5]

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

[5] 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

Atkeson and Burstein (2019) study own improvement alongside product acquisition and entry. Here, size-dependent protection makes both the incentive to improve a product and its replacement risk depend on its relative size.

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 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.

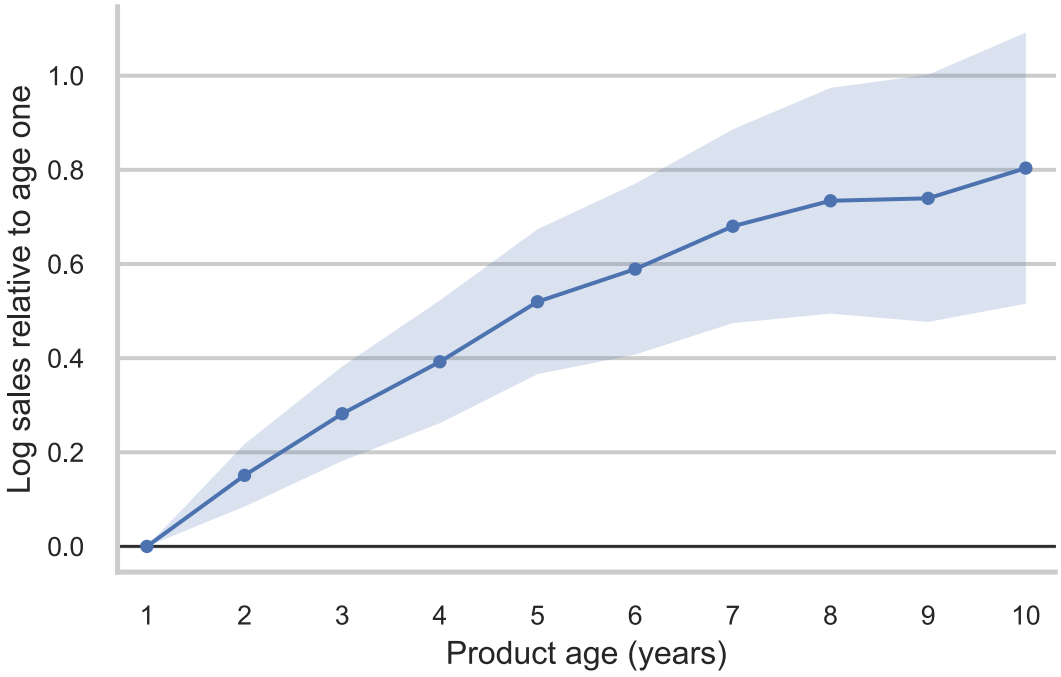


Figure 5: Sales over the life cycle of continuing products. The line shows estimated log-sales differences from age one for the same 2,507 firm-product lines at every age. The regression includes industry-year and cohort effects, with the linear cohort trend set to zero and no initial-sales control. 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. The resulting improvement raises total quality at rate I_j:

$$\frac{d\log(q_j a_j)}{dt} = I_j, \qquad \dot{\tilde{S}}_j = (\sigma - 1)\left(I_j - \frac{\dot{Q}}{Q}\right)\tilde{S}_j.$$

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.

Own innovation increases the quality of the current product without creating a new variety or research platform. Protection continues to satisfy $a_j = a_\theta(\tilde{S}_j)$: a product gains protection as it grows relative to the economy and loses protection as it falls behind. Positive own innovation therefore need not increase protection if $I_j < \dot{Q}/Q$.

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

$$\ell_j^I = \tilde{S}_j \frac{I_j^2}{\phi_I}, \qquad \phi_I > 0.$$

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. Together with separate costs for own innovation and expansion and no shared firm research constraint, this preserves additive firm value. Product value is generally nonlinear in relative size because protection also depends on size. I characterize the own-innovation policy below.

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, expansion costs scale with the number of active products, and own-innovation costs are separable. 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 active products and their associated research platforms.

Let $v(\tilde{S})$ be the value of an active product together with its research platform. Protection is determined by relative sales, so it is not an additional state. Firm value is

$$V_f = \sum_{j \in \mathcal{A}_f} v(\tilde{S}_j).$$

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. Both sources of value face the risk of creative destruction depending on the product's protection. Let λ_H be the incoming innovation-attempt arrival rate per market. Relative-sales drift and the product's capture hazard are then

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

Here $a = a_\theta(\tilde{S})$ changes with relative sales. Products also face a constant risk of becoming obsolete. Each research platform can generate new products while its associated product remains active. Let J^X denote the expected value created by one research attempt:

$$\begin{aligned} J^X = (1 - \alpha_X)\mathbb{E}_{d,\omega}[v(\tilde{S}_X^+)] \\ + \alpha_X \mathbb{E}_{P,G}\left[\mathbf{1}\{e^d > a_\theta(\tilde{S}')\}\, v(\tilde{S}^+)\right]. \end{aligned}$$

The first term averages over incumbent launch draws. The second averages over the uniformly sampled target's relative sales, drawn from the stationary product distribution P, and the independent log discovery step, with the post-capture state $\tilde{S}^+$ defined above. Failed attempts contribute zero. The value of an active product and its platform then satisfies the Bellman equation

$$\begin{aligned} \rho v(\tilde{S}) = \underbrace{\frac{\ell \tilde{S}}{\sigma - 1}}_{\text{current profit}} - \underbrace{h(a_\theta(\tilde{S}))v(\tilde{S})}_{\text{capture}} - \underbrace{\delta v(\tilde{S})}_{\text{obsolescence}} \\ + \max_{I \geq 0} \left\{ \underbrace{(\sigma - 1)(I - g)\tilde{S}\, v'(\tilde{S})}_{\text{change in relative sales}} - \underbrace{\tilde{S} I^2 / \phi_I}_{\text{own-innovation cost}} \right\} \\ + \max_{x \geq 0} \left\{ \underbrace{x J^X}_{\text{expansion payoff}} - \underbrace{x^2 / \phi_X}_{\text{expansion cost}} \right\}. \end{aligned} \tag{1}$$

Capture removes the product and its research platform from the displaced firm, so the loss is

the full value v. Market obsolescence also removes both. The derivative $v'(\tilde{S})$ includes the change in protection as relative sales change. Own innovation raises relative-sales growth at a labor cost. The last line describes the platform's research choice: a higher attempt rate increases the expected innovation payoff but uses more labor. Although operating profits are proportional to relative sales, the size-dependent capture hazard generally makes product value nonlinear in relative sales.

5.2.3 Research effort

Own innovation. A unit increase in I raises the growth rate of relative sales by $\sigma - 1$. Where v is differentiable, the optimal rate satisfies

$$I(\tilde{S}) = \frac{\phi_I(\sigma - 1)}{2}[v'(\tilde{S})]_+,$$

where $[z]_+ = \max\{z, 0\}$. The derivative includes the value of larger operating profits and the change in protection of both profits and future research opportunities. The firm equates this marginal value to marginal research cost when $I > 0$. Although the common size factor cancels from the first-order condition, v' generally varies with size, so the own-innovation policy does too.

These returns create scope for complementarity between protection and own innovation: protection prolongs the return to improving a product, while faster improvement raises protection relative to the alternative. They do not imply that $I(\tilde{S})$ increases monotonically with size. Additional protection may have little value for a product already insulated from capture. The policy and product value must therefore be determined jointly.

Expansion. 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 $x = \phi_X J^X/2$ per platform. Their total attempt rates, $n_f x$, differ with their active product counts, while own-innovation effort depends on the sizes of their products.

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 mass of active products.

Let P denote the stationary, count-weighted distribution of active products over relative sales $\tilde{S}$, with $\mathbb{E}_P[\tilde{S}] = 1$. Protection is given by $a_\theta(\tilde{S})$. The distribution P must reproduce itself under the specified entrant and incumbent launches, successful replacement, obsolescence and own-improvement laws.

5.3.1 Research attempts and platform stocks

Each research platform generates attempts at the common rate x. Conditional on targeting an existing variety, the probability that an attempt succeeds is

$$H = \Pr_{P,G}(e^d > a_\theta(\tilde{S})),$$

where the target's relative sales $\tilde{S}$ are drawn from P and the log discovery step d is drawn independently from G.

There is one platform per active product, so N also denotes the aggregate stock of platforms. Entrant and incumbent launches add products and platforms, while obsolescence removes both. Successful captures transfer them between owners. Constant variety mass therefore requires

$$E + (1 - \alpha_X)Nx = \delta N.$$

Positive entry requires $E/N = \delta - (1 - \alpha_X)x > 0$. There are $\alpha_X Nx$ capture attempts and N active varieties, so incoming attempts per variety satisfy

$$\lambda_H = \alpha_X x.$$

The average rate of successful capture per variety is $\alpha_X x H$.

5.3.2 Entry and labor allocation

An entrant pays its labor cost before drawing launch quality. With positive entry, its expected value in wage units equals the labor cost of entry:

$$\mathbb{E}_{d,\omega}[v(\tilde{S}_E^+)] = \frac{1}{\phi_E}.$$

Entry and expansion need not have the same expected payoff: entrants create only new varieties, while incumbents may capture existing ones. Production, expansion, entry and own innovation employ the available labor:

$$L = N\ell + \frac{Nx^2}{\phi_X} + \frac{E}{\phi_E} + N\mathbb{E}_P\left[\tilde{S}\frac{I(\tilde{S})^2}{\phi_I}\right].$$

Using the stationary variety balance gives

$$N = \frac{L}{\ell + x^2/\phi_X + [\delta - (1-\alpha_X)x]/\phi_E + \mathbb{E}_P[\tilde{S}I(\tilde{S})^2/\phi_I]}.$$

The denominator is total labor required per active variety.

5.3.3 Aggregate growth

A successful capture replaces a target's total quality qa with qe^d. Its CES-weighted quality, $(qa)^{\sigma-1}$, therefore rises by the proportion $(e^d/a)^{\sigma-1} - 1$. 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 per attempt at an existing variety. 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. Adding the contributions of own improvement and new varieties, and subtracting losses from obsolescence, gives

$$\begin{aligned}(\sigma-1)g = {} & \underbrace{(\sigma-1)\mathbb{E}_P[\tilde{S}I(\tilde{S})]}_{\text{own improvement}} \\ & + \underbrace{\lambda_H \mathbb{E}_P[\tilde{S}\Gamma(a_\theta(\tilde{S}))]}_{\text{creative destruction}} \\ & + \underbrace{\frac{E}{N}m_E}_{\text{entrant launches}} + \underbrace{(1-\alpha_X)xm_X}_{\text{incumbent launches}} - \underbrace{\delta}_{\text{obsolescence}}.\end{aligned} \tag{2}$$

The capture term counts the increase in CES-weighted quality over the target's existing product, while the new owner acquires the whole product. 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.

I report the composition of growth as shares of growth before obsolescence: the sum of the own-improvement, creative-destruction and new-variety contributions, before subtracting $\delta/(\sigma-1)$. These three shares sum to one; subtracting obsolescence gives net growth g.

With a common launch mean $m = m_E = m_X$, launches and obsolescence together contribute $\delta(m-1)/(\sigma-1)$ to g. Own improvement and creative destruction are weighted by relative product size.

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.

I solve the product-value equation and stationary size distribution jointly with research choice, free entry and the mean-size normalization. The zero-protection economy supplies an analytical check on the numerical solution; the protected economy generally requires a nonlinear value function. Simulated firm histories then translate the product-level solution into annual observations.

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 $\theta = 0$, in which all quality is transferable.

Fact C: which products are replaced. With $\theta = 0$, every attempt at an existing variety 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 size-based protection, the capture hazard $h(a_\theta(\tilde{S}))$ is weakly decreasing in relative sales. Holding the target's public quality q' and the quality mean Q fixed, a successful capture with step d yields

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

Larger steps therefore create larger products with weakly lower capture hazards. More

generally, the protection rule directly links size and replacement risk. During an ownership spell, protection adjusts with relative size, so capture risk can rise as a product falls behind the economy or fall as it pulls ahead.

Fact B: where new products land. Fact B concerns new-to-firm products added by existing firms, through either new-variety launches or creative destruction. For the capture component, 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 $\theta = 0$, 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 captured additions is therefore shifted toward larger products relative to that of active products. With size-based protection, replacement is tilted toward smaller, weakly protected products. This selection can offset the size increase at replacement, allowing captured additions to remain small relative to existing products. New-variety launches can be small in either economy. Through the lens of the model, the capture components of Facts B and C 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. Protection weakly lowers $\Gamma(a)$, the expected proportional gain per capture attempt, 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, let g_{CD} denote creative destruction's contribution to aggregate growth. With $a = a_\theta(\tilde{S})$,

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

Since a_θ is nondecreasing and Γ is nonincreasing, the covariance is nonpositive. Holding $\lambda_H > 0$, G and the marginal distributions of size and protection fixed, a negative covariance lowers the growth contribution of replacement relative to an accounting reference in which size and protection are independent. The average capture rate $\mathbb{E}_P[h(a)]$ is the same in that reference.

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. This reassignment is an accounting comparison: it need not satisfy the rule $a = a_\theta(\tilde{S})$ or describe another stationary equilibrium.

With $\theta = 0$, Γ is constant and the covariance vanishes. With size-based protection, expected proportional gains are weakly smaller for larger products. When the covariance is negative, these gains are tilted toward smaller products and contribute less to aggregate growth than in the independence reference. This does not sign the equilibrium growth response to protection, because research effort and the product distribution also change. 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.

6 Calibration

I estimate both the benchmark economy where $\theta = 0$ and the model with private know-how for which I let $\theta = 1$, to match the same aggregate and firm moments in the Swedish manufacturing data. The models thus give similar predictions for the targeted firm-level moments and the aggregate growth rate. What can differ between the two economies is then the composition of aggregate growth and the product dynamics.

6.1 Annual sales and employment

I construct annual product sales by integrating the 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. A product captured during a year contributes sales to its previous owner before capture and to its new owner afterwards. Firm scope counts products with positive annual sales, including products that disappear before year-end.

Employment includes production and both research activities. At a given date, an active firm's labor use is

$$\ell_f = \sum_{j \in \mathcal{A}_f} \left[\ell \tilde{S}_j + \tilde{S}_j \frac{I(\tilde{S}_j)^2}{\phi_I} \right] + \frac{n_f x^2}{\phi_X}.$$

I integrate these flows over the year and divide by the employment observation scale c, discussed below, to obtain observed employment. Firm scope, sales ranks and growth are constructed from the annual observations using the empirical definitions. Product moments apply the 20-employee owner threshold and the corresponding timing and eligibility restrictions. The balanced product lifecycle uses the same continuing ownership spells across ages; selection into survival remains part of that comparison.

6.2 Parameters and empirical counterparts

To begin, I assign common parameters standard values. I set exogenously the elasticity of substitution between varieties $\sigma = 4$, as well as letting the research cost be quadratic as is common in the literature. For the discount rate I use $\rho = .012$ as in Peters and Walsh (2026). I normalize entry efficiency to $\phi_E = 1$. I assume that entrants create new varieties only and set $\alpha_X = .9$, meaning most expansion research is attempting creative destruction. Both economies impose $m_E = m_X$ and use the launch and protection laws specified above. The proportional discovery step $e^d - 1$ follows the same upper-truncated lognormal distribution in both economies, with mean .25 and coefficient of variation one after truncation.

The remaining parameters are discussed in turn.

6.2.1 Expansion efficiency, ϕ_X

The expansion efficiency governs the amount of labor required for a given invention arrival rate. Since successful expansion adds products to a firm, I target an average number of products moment. I observe product count among responding firms in the IVP survey. To approximate this sample, I match the simulated average product count moment among model firms with at least 20 employees.

Following Berlingieri et al. (2026), I let one distinct product code in the data map to one product in the model. The mean among 3,283 manufacturing firms observed in IVP in 2019 gives a product count target at **2.949 products per firm**.

6.2.2 Own-innovation efficiency, ϕ_I

Similar to the expansion efficiency, the own-innovation efficiency ϕ_I governs the amount of labor required for a given amount of own-innovation. As own-innovation complements expansion research in generating aggregate growth, a data moment sensitive to the own-innovation efficiency is the aggregate growth rate. I fit total growth through the stationary growth equation, without simulation.

To measure an aggregate growth rate in the data, I average twenty Swedish manufacturing TFP changes ending in 2000–2019, each converted to labor-augmenting units using its adjacent-year average labor compensation share. The empirical appendix shows the annual series and construction. The series is noisy, motivating a longer panel. The aggregate growth target is **3.21 log percentage points per year**.

6.2.3 Market obsolescence, δ

Market obsolescence, δ, makes products obsolete at a constant hazard, regardless of a product's state. Single-product firms in the model will thus periodically be forced to exit due to obsolescence, making the firm exit rate an informative target. Capture also contributes to last-product loss. As the model assumes the economy is in a steady state, the exit rate coincides with the entry rate.

In the broad manufacturing sample of the administrative FE data, I infer net exits from new entrants minus the change in the number of active firms. Averaging entry and net exit rates over 2015–2019 gives a turnover target of **3.64 percent per year**.

If T is a firm's lifetime, the maintained model measure is $F = \int_0^1 \Pr(T > u)\, du / \mathbb{E}[T]$: first-year-end survivors relative to continuous producing-firm exposure, averaging over founding dates. It can be computed from firm survival without simulation.

6.2.4 Common launch scale, $m = m_E = m_X$

For the launch scale of new varieties, I begin by assuming that this is the same for entering firms as it is for incumbent firms who invent new varieties, i.e. $m = m_E = m_X$. Since entrants enter with one product, this assumption allows me to measure a related moment in the data by measuring the size of entrants.

In the data I use value added relative to the same two-digit sector-year's all-age mean, pooling 2015–2019 with equal year weights. The target is a first-full-year entrant median equal to **8.47 percent of mean firm size**.

6.2.5 Launch dispersion, s_ω

The same entrant sample has a **standard deviation of log relative value added of 1.607**. The independent launch-quality draw ω adds dispersion beyond the log discovery step d. In the launch distribution, $\mathrm{Var}(\log \tilde{S}^+) = (\sigma - 1)^2\{\mathrm{Var}(d) + s_\omega^2\}$, so s_ω directly governs dispersion at birth conditional on the distribution of log discovery steps d. Observed entrants have already experienced development, selection and possible acquisitions, so I fit their first-full-year dispersion through simulation rather than equating it to launch dispersion. With common firm primitives and no demand shocks, this parameter absorbs dispersion that the model assigns to launch opportunities.

6.2.6 Employment observation scale, c

In broad FE manufacturing data, **15.25 percent of positive-employment, positive-sales firm-years have at least 20 employees**, pooling 2015–2019. Conditional on the simulated annual payroll distribution, choosing c is a direct threshold-matching calculation, up to discrete sample shares. It is repeated in each history, so the threshold share is fitted mechanically rather than independently validated. The inferred scale and the resulting selected scope remain subject to simulation variation.

6.3 Computation and precision

I adjust ϕ_X, m and s_ω jointly using simulated scope and entrant moments. At every evaluation, an inner numerical solve adjusts δ and ϕ_I to match turnover and aggregate growth while re-solving equilibrium. Employment scale is then chosen from the simulated payroll distribution. Deterministic calculations have no Monte Carlo error conditional on parameters, but retain numerical approximation error; fitted parameters also inherit uncertainty from the simulated targets.

Calibration uses one common random seed with 16,000 firm families across parameter evaluations. The protected point was accepted manually after further updates moved the simulated scope around the target. I assess the accepted parameters on two new independent histories of 16,000 families each, recalibrating c in each. These histories measure simulation variation, not empirical or parameter uncertainty.

6.4 Parameters and targeted fit

Table 1 reports the fitted parameter values. The large difference in expansion efficiency compensates for inventions lost due to protection in the $\theta = 1$ economy. In addition, product obsolescence is higher in the $\theta = 1$ economy, partially offsetting the size-dependent firm-product separation rate.

Table 1: Calibrated parameters, with $\phi_E = 1$. Employment scale is shown for the calibration history and is recalculated in the confirmation histories.

Parameter	No protection	Protection
Expansion efficiency, ϕ_X	.02691	2.592
Obsolescence, δ	.02678	.04272
Own-innovation efficiency, ϕ_I	.01856	.01511
Common launch scale, m	.4820	.5012
Launch dispersion, s_ω	.4964	.5075
Employment scale, c	.02677	.06967

6.4.1 Targeted moments

Table 2: Targeted fit on the two independent confirmation histories. Growth is deterministic; reported turnover is its simulated counterpart. The employment share is refitted in each history.

Target moment	Data	No protection	Protection
Mean product count	2.949	2.901	2.849
Annual firm turnover (%)	3.64	3.65	3.66
Entrant median / mean firm size	.0847	.0859	.0856
Entrant log-size SD	1.607	1.613	1.620
Annual growth (log pp)	3.21	3.21	3.21
Employment ≥ 20 share (%)	15.25	15.25	15.25

7 Results

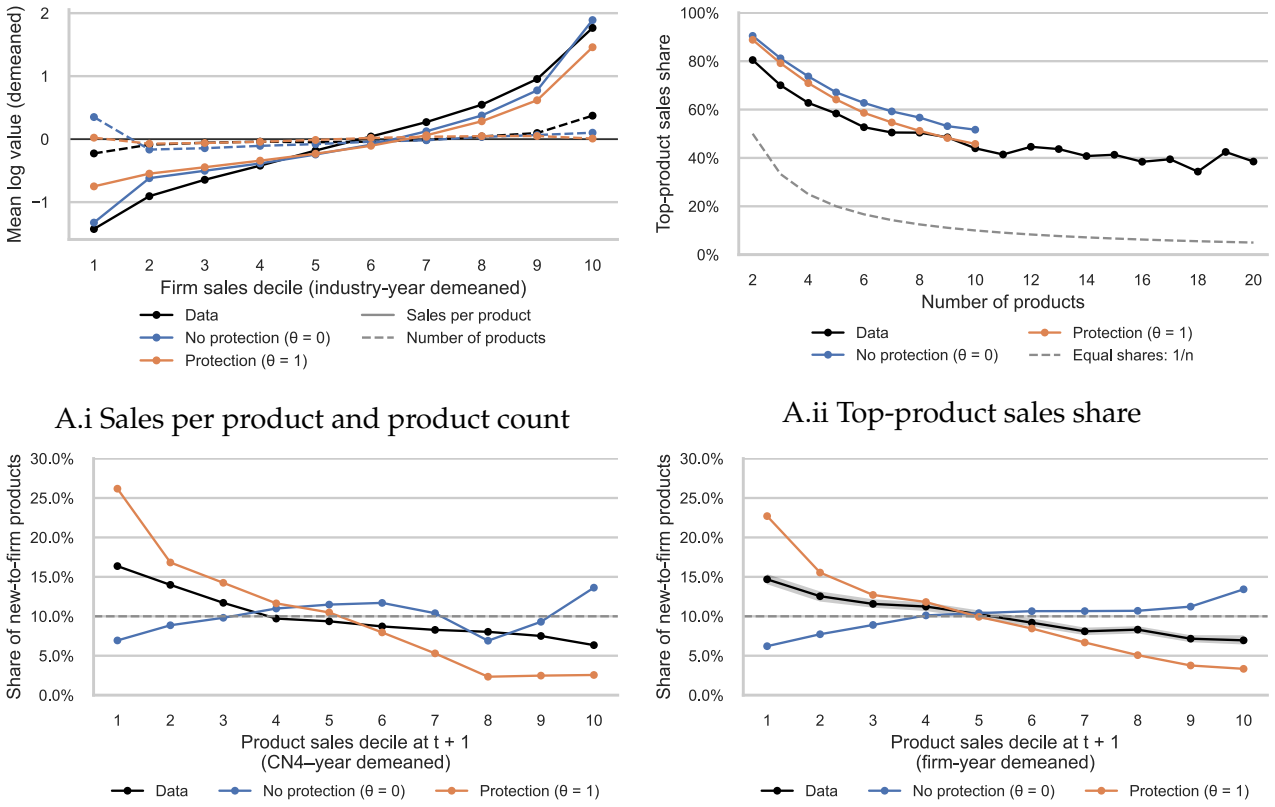
I compare the two calibrated economies to learn how private know-how changes product dynamics and the sources of growth. Both economies match the same aggregate growth rate and target the same firm-level moments.

7.1 Product size, additions and exits

Private know-how moves product additions towards smaller products (Figure 6). Without protection, newly added products tend to be large. With protection, competitors instead replace smaller products, so additions move towards the lower sales deciles, although more strongly than in the data. Both economies produce concentrated portfolios, but neither reproduces the increase in product count with firm size.

Product exits require more care because losing a large product can also remove its owner from the observed sample. This selection makes the benchmark appear to protect large products. I therefore emphasize exit rates among products whose owners' other sales imply at least 20 employees, using lagged employment and sales shares. This restriction reduces selection caused by losing the focal product.

Exit rates are nearly flat without protection and decline with size under protection (Figure 7). The decline remains weaker than in the data, and constant obsolescence keeps exits among the largest products too frequent.



A.i Sales per product and product count

A.ii Top-product sales share

B.i Additions across products

B.ii Additions within firms

Figure 6: Product size and additions in the data and the two economies, using the annual observation rules. Black: data; blue: no protection; orange: protection. In A.i, solid lines show sales per product and dashed lines product count. B shows the distribution of additions across sales deciles. Shading in B.ii: pointwise 95% intervals.

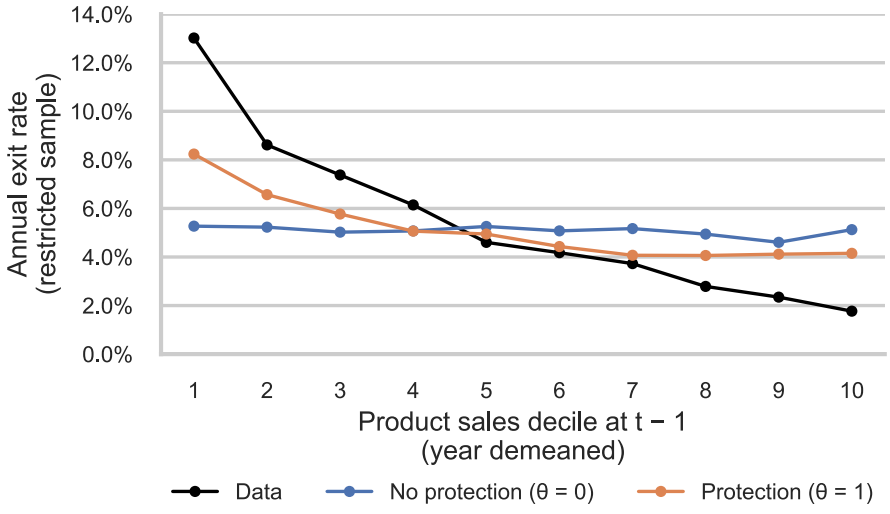


Figure 7: Annual product exit rates in the employment-at-least-20 sample, by pooled year-demeaned sales decile at $t-1$. Lagged owner employment times one minus the focal product's lagged sales share must be at least 20. Colors follow Figure 6.

7.2 Product development over the life cycle

Protection also changes what happens to products that survive. By strengthening the incentive to improve existing products, it allows them to grow under the same owner. Between ages one and ten, adjusted log sales rise by .20 with protection, compared with a decline of .24 without it (Figure 8). The direction agrees with the data, where the increase is .80, but the model explains only part of the life-cycle growth.

I estimate the data profile following Argente, Lee, and Moreira (2024), using products observed continuously under the same owner from launch through age ten. I regress log sales on age, birth-cohort and owner-industry-year effects, omitting the potentially partial launch year and assuming no linear cohort trend. The same products contribute at every age, so the profile describes growth among continuing products. It does not isolate technological improvement from demand or other sources of sales growth. The empirical appendix gives the specification and sample construction.

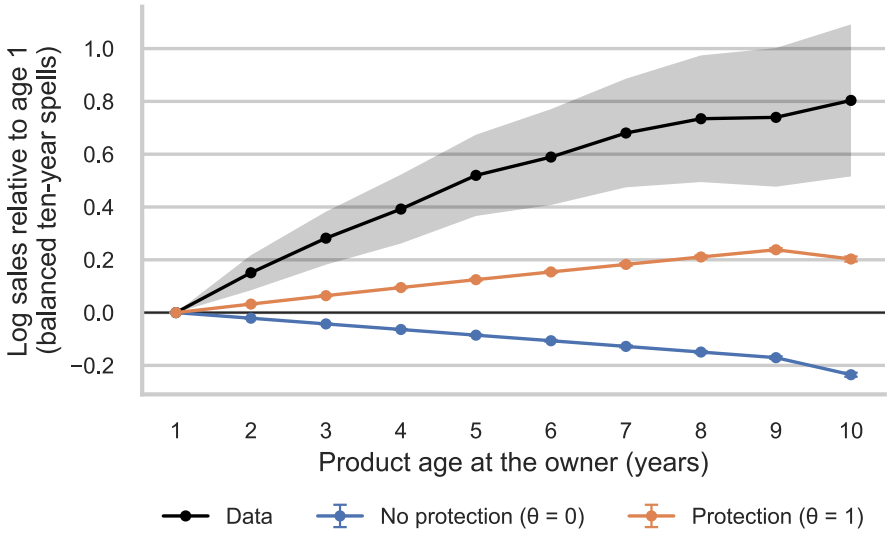


Figure 8: Product sales over balanced ten-year life cycles, relative to age one, controlling for cohort and industry-year effects with no linear cohort trend. Bands and bars show pointwise 95% intervals. Colors follow Figure 6.

7.3 Sources of aggregate growth

The two economies have the same net growth rate, but for different reasons (Table 3). With private know-how, own improvement accounts for 82.7% of growth before obsolescence, compared with 60.1% without it. Creative destruction contributes 1.9% versus 29.4%, and new varieties 15.4% versus 10.5%. In both economies, new varieties start small and add less quality than is lost to obsolescence.

Table 3: Sources of annual aggregate growth, in log percentage points. Parentheses report shares of growth before obsolescence, the sum of own improvement, creative destruction and new-variety creation. New varieties include entrant and incumbent launches.

Source	No protection	Protection
Own improvement	2.468 (60.1%)	3.833 (82.7%)
Creative destruction	1.205 (29.4%)	.087 (1.9%)
New varieties	.430 (10.5%)	.714 (15.4%)
Growth before obsolescence	**4.103 (100.0%)**	**4.634 (100.0%)**
Obsolescence	-.893	-1.424
Net growth	**3.210**	**3.210**

The analytical decomposition explains why the size of the products being replaced matters. Protection makes the expected proportional gain from a capture attempt smaller at large products. Relative to an accounting reference with size and protection independent, the covariance term contributes -.42 log percentage points to annual growth in the protected economy and is zero in the benchmark. Creative destruction contributes less when its gains fall on small products, even for a given frequency of replacement.

7.4 The products behind large firms

Large products have different histories in the two economies (Table 4). Without protection, they grow large through repeated replacement by competing firms. With protection, they tend to remain with one owner and grow through own improvement. Among the largest 1% of products, 86.7% have never changed owner, compared with 11.7% without protection. Their mean ages are similar, but current-owner tenure rises from 12.3 to 52.8 years.

Table 4: Product and firm tails in the model's stationary cross-section, before applying annual survey rules. Products and firms are ranked separately by sales in each economy.

Tail statistic	No protection	Protection
Largest 1% of products		
Never captured (%)	11.7	86.7
First entered the tail through own growth (%)	0.0	65.3
Mean product-line age (years)	56.3	57.5
Mean current-owner tenure (years)	12.3	52.8
Largest 1% of firms		
Mean product count	3.18	2.87
Mean top-product sales share (%)	94.5	93.6
Top product acquired through capture (%)	92.0	10.8

At the firm level, both economies produce large firms dominated by one product. What differs is where that product came from. In the benchmark, 92.0% of the largest firms acquired their leading product from a competitor. With protection, this share falls to 10.8%; most instead sell a leading product they launched themselves. Similar firm-level concentration can thus reflect very different processes of innovation and growth.

The data also link leading products to longer ownership tenure. Among multiproduct firms whose product ages can all be identified, the leading product's observed ownership age averages 9.1 years in 2016–2019, compared with 7.3 years for the other products. I cannot date products already present when observation begins in 1997, or those first seen when their owner enters the sample. Requiring a dated history for every product therefore leaves a small, selected sample.[6] The model supplies complete product histories under its assumptions, allowing the analysis to extend beyond the observation window.

7.5 Innovation policy

The different sources of growth matter for innovation policy. I consider an unexpected, permanent subsidy covering 10% of research costs, directed at expansion, own improvement, or both. Entry research remains unsubsidized, entry adjusts endogenously, and lump-sum taxes finance the policy. Each economy starts from its calibrated balanced growth path. The subsidy gives the same proportional incentive in both economies, but fiscal costs differ. The policy appendix describes the transition equilibrium and welfare calculation.

Supporting own improvement raises growth and welfare in both economies (Table 5). Expansion support has a much smaller benefit without protection and lowers welfare with protection. Subsidizing expansion reduces the contribution of own improvement to growth in both calibrated economies. What differs is the growth gained from captures. With protection, more expansion research mostly produces failed attempts on protected products, and the successful ones win small products. Creative destruction then makes up for only about a sixth of the lost own improvement, compared with more than all of it without protection.

[6]The comparison gives equal weight to 850 firm-years from 291 firms, covering 13.5% of eligible firm-years with 2–10 products. Ownership age is time since the first identified report at the current owner. First appearances after reporting gaps also have unknown ages. The sample is not restricted to the largest firms.

Table 5: Effects of a permanent 10% research-cost subsidy. Growth changes compare balanced growth paths, in percentage points per year; the rows below each split the change by source. New varieties and obsolescence contribute no long-run change: total launches per product equal the exogenous exit rate on a balanced growth path, and the subsidy only shifts launches from entrants to incumbents. Welfare includes the full transition and is expressed as a permanent percentage increase in no-reform consumption. Fiscal cost is discounted subsidy expenditure relative to discounted output, using the household's discount factor.

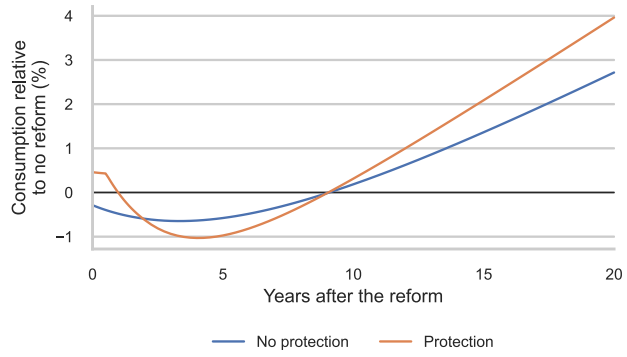
Policy and outcome	No protection	Protection
Both activities		
Long-run growth change (pp/year)	+.282	+.365
of which: own improvement	+.176	+.354
of which: creative destruction	+.106	+.011
Welfare gain (%)	+22.97	+31.18
Fiscal cost (% of output)	1.81	2.07
Expansion only		
Long-run growth change (pp/year)	+.018	-.053
of which: own improvement	-.088	-.064
of which: creative destruction	+.106	+.011
Welfare gain (%)	+.25	-4.26
Fiscal cost (% of output)	1.02	.32
Own improvement only		
Long-run growth change (pp/year)	+.274	+.425
of which: own improvement	+.274	+.425
of which: creative destruction	.000	.000
Welfare gain (%)	+23.46	+37.71
Fiscal cost (% of output)	.87	1.78

With log utility, the consumption-equivalent welfare gain is

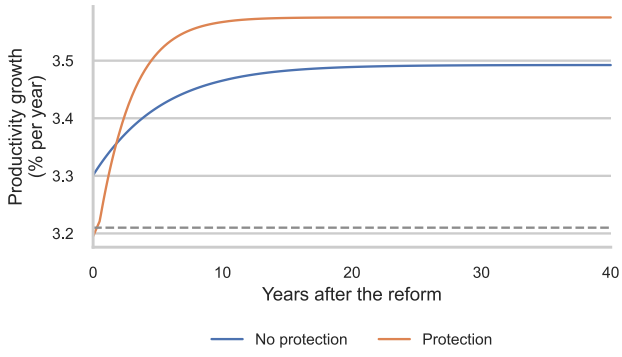
$$\omega = \exp\left\{\rho \int_0^\infty e^{-\rho t} \log\left(\frac{C_t^{\text{reform}}}{C_t^{\text{no reform}}}\right) dt\right\} - 1.$$

The large gains reflect permanent growth increases discounted at $\rho = .012$. They are conditional on the calibrated research technologies and this discount rate.

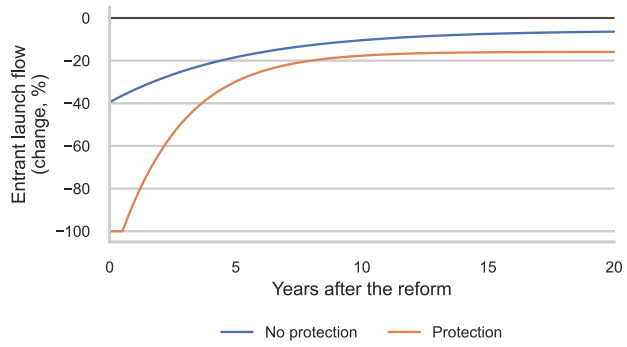
The joint subsidy illustrates the adjustment (Figure 9). Without protection, consumption falls .29% on impact. With protection, entry stops for about .6 years, freeing enough labor for consumption initially to rise .46%. Consumption subsequently falls below its no-reform path as entry resumes and product mass declines. In both economies it remains above the no-reform path after about nine years. The long-run growth gains outweigh these temporary losses, giving welfare gains of 23.0% and 31.2%.



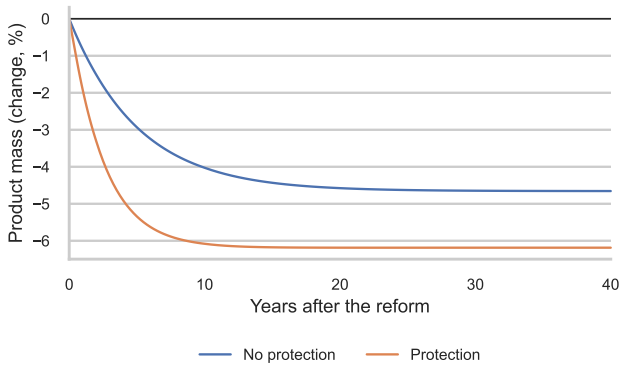
(i) Consumption relative to no reform



(ii) Productivity growth



(iii) Entrant product launches



(iv) Product mass

Figure 9: Transition after the joint subsidy. Blue: no protection; orange: protection. Consumption is $100(C_t/C_t^{\text{no reform}} - 1)$; entry and product mass are percentage changes from their pre-reform levels. The dashed growth line marks the common pre-reform rate. Panels use different time horizons to show the adjustment. Deterministic model paths.

8 Conclusion

In Swedish manufacturing, firms' sales are concentrated in a few products, while products added to or dropped from their portfolios tend to be small. These facts motivate a model in which private know-how rises with product size. This makes successful products harder for competitors to replace and strengthens incumbents' incentives to improve them. The growth contribution of creative destruction then depends on the size of the products it reaches as well as the frequency of replacement.

I compare two economies calibrated to the same aggregate growth rate and firm-level targets. With private know-how, creative destruction accounts for 1.9% of growth before obsolescence, compared with 29.4% without it. Own improvement accounts for 82.7% versus 60.1%, and new varieties for 15.4% versus 10.5%. Large products also more often develop under the same owner. Similar firm-level outcomes can therefore conceal different histories of product development and different sources of aggregate growth.

This distinction matters for innovation policy. In the calibrated experiments, a subsidy covering 10% of the cost of improving existing products raises growth and welfare in both economies. Subsidizing expansion research yields small gains without protection but lowers growth and welfare with private know-how. Subsidizing expansion reduces the contribution of own improvement to growth in both calibrated economies. With private know-how, the additional growth from replacement is too small to offset this loss.

The strength of the link between product size and private know-how remains to be estimated. Linking product histories to evidence on knowledge transfer could help distinguish this mechanism from other sources of persistent product success. Allowing firms to share or license know-how would also make it possible to study how policies that facilitate knowledge diffusion affect both incentives for product development and opportunities for subsequent innovation.

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 Klenow, and Huiyu Li. 2025. "Good Rents Versus Bad Rents: R&D Misallocation and Growth." https://doi.org/10.3386/w34190.

Aghion, Philippe, and Peter Howitt. 1992. "A Model of Growth Through Creative Destruction." *Econometrica* 60 (2): 323–51. https://doi.org/10.2307/2951599.

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, Danial Lashkari, and Davide Rigo. 2026. "Creative Destruction Through Innovation Bursts."

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

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. 2011. "Rank - 1 / 2: A Simple Way to Improve the OLS Estimation of Tail Exponents." *Journal of Business & Economic Statistics* 29 (1): 24–39. https://doi.org/10.1198/jbes.2009.06157.

Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter Klenow. 2019. "How Destructive Is Innova-

tion?" *Econometrica* 87 (5): 1507–41. https://doi.org/10.3982/ECTA14930.

Grossman, Gene M., and Elhanan Helpman. 1991. "Quality Ladders in the Theory of Growth." *The Review of Economic Studies* 58 (1): 43–61. https://doi.org/10.2307/2298044.

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

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

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.

Romer, Paul M. 1990. "Endogenous Technological Change." *Journal of Political Economy* 98 (5, Part 2): S71–102. https://doi.org/10.1086/261725.

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

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

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

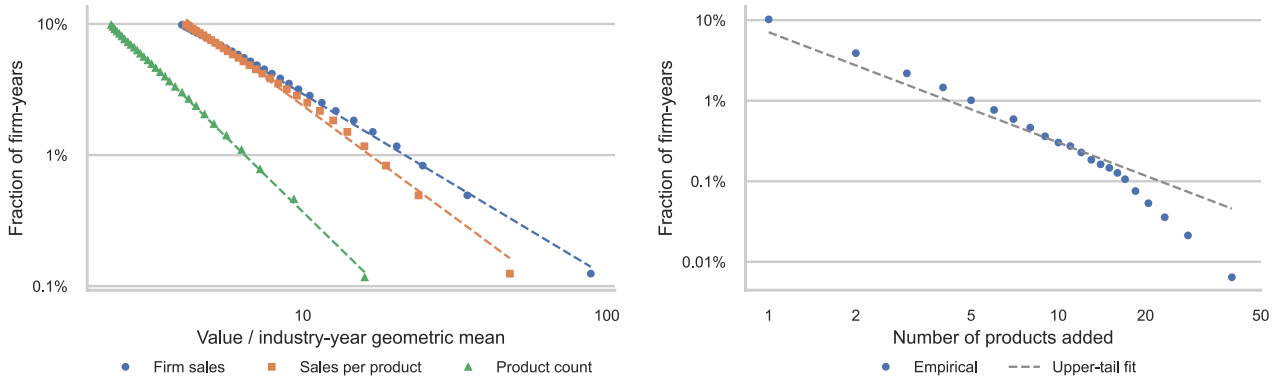
by OLS, with slope $-\alpha_X$, intercept c_X, descending rank i and $n = 89{,}721$. Cutoff ties remain included and retain ordinary ranks; industry-year means are fixed before selection. Figure 10 (i) shows complementary cumulative distributions (CCDFs) and descriptive fits.

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 firm-years with no additions. The subtraction of one-half follows Gabaix and Ibragimov (2011). 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 10: 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; probabilities use all firm-year transitions, including those with no additions.

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 6 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 6: 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 model section 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 7: 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 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 11 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.

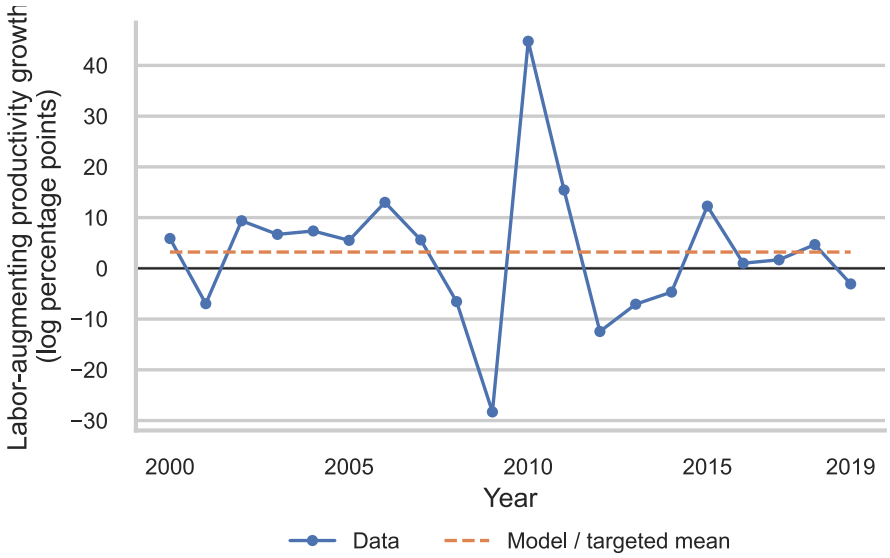


Figure 11: Annual manufacturing productivity growth. EU KLEMS TFP growth is converted to labor-augmenting units using adjacent-year average labor compensation shares. The dashed line is the equally weighted 2000–2019 mean, used as the model's growth target. No sampling intervals are shown.

B Policy transitions

Each experiment starts from the economy's calibrated BGP and introduces an unexpected, permanent subsidy at $t = 0$. Calibrated primitives remain fixed. Let ξ_X and ξ_I denote the fractions of expansion and own-improvement costs covered by the subsidy: each is either zero or .10. In the private research problem, replace ϕ_c by $\phi_c/(1-\xi_c)$ for $c \in \{X, I\}$. Physical labor requirements retain the original ϕ_c, so labor clearing uses the unsubsidized research technologies. Entry is unsubsidized, and lump-sum taxes finance subsidy payments. These payments are transfers rather than an additional resource cost.

Product values v_t, the count-weighted product distribution P_t, and aggregate allocations now

vary over time. The Bellman equation has the same returns and transition events, evaluated at current policies and P_t, but its left-hand side becomes $\rho_t v_t - \partial_t v_t$. Wage normalization gives

$$\rho_t = r_t - \frac{\dot{w}_t}{w_t} = \rho + \frac{d\log(C_t/w_t)}{dt} = \rho + \frac{d\log(N_t\ell_t)}{dt}.$$

Thus the effective discount rate equals ρ only when production labor $N_t\ell_t$ is constant. Relative sales drift at $(\sigma-1)(I_t - g_{Q,t})\tilde{S}$, where $g_{Q,t} = \dot{Q}_t/Q_t$ keeps $\mathbb{E}_{P_t}[\tilde{S}] = 1$.

The distribution evolves forward under own improvement, captures, launches and obsolescence. Product mass need not remain constant:

$$\dot{N}_t = E_t + (1-\alpha_X)N_t x_t - \delta N_t.$$

Aggregate productivity growth is therefore $g_{Q,t} + \dot{N}_t/[(\sigma-1)N_t]$. Labor clearing determines the feasible entry flow. Free entry allows periods without entry:

$$E_t \geq 0, \qquad \frac{1}{\phi_E} - \mathbb{E}_{d,\omega}[v_t(\tilde{S}_E^+)] \geq 0,$$
$$E_t\left(\frac{1}{\phi_E} - \mathbb{E}_{d,\omega}[v_t(\tilde{S}_E^+)]\right) = 0.$$

The initial distribution, product mass and quality are inherited from the pre-reform BGP; research and production allocations can adjust immediately. I solve values backward from the post-reform BGP and the distribution forward from the initial state, iterating until research choices, entry and labor clearing are consistent. Consumption follows $C_t = Q_t N_t^{1/(\sigma-1)} N_t\ell_t$. Welfare compares this path with continued growth on the original BGP. Beyond the computational horizon, I continue consumption at the new BGP growth rate, retaining the accumulated level difference. A longer horizon and finer time and size grids leave the reported welfare effects essentially unchanged.