Industry Life Cycles in General Equilibrium*

Laurent Cavenaile[†]
University of Toronto

Ruben Gaetani[‡]
University of Toronto

Pau Roldan-Blanco[§]
Banco de España and CEMFI

Tom Schmitz[¶]
Queen Mary Univ. of London

February 1, 2024

PRELIMINARY AND INCOMPLETE

Abstract

Technological breakthroughs shape the life cycle of industries. Using patent data, we show that the ICT Revolution of the late 1980s generated new technological opportunities across non-ICT industries, leading to an initial phase of rapid growth in patenting and product innovation. Eventually, however, this gave way to a shakeout, characterised by a fall in the number of patenting firms and a greater prevalence of process innovation. To capture these patterns, we build a general equilibrium model of industry life cycles, featuring endogenous innovation choices and oligopolistic competition between an endogenously determined number of heterogeneous firms within each industry. We then use our model to study the welfare consequences of different R&D subsidy policies. Preliminary results show that subsidies have non-monotonic implications for welfare and accelerate life cycles.

Keywords: Industry Life Cycles, Product Innovation, Process Innovation, ICT, Technology.
JEL codes: E22, L13, L22, O31.

---

*We thank Sergio Inferrera for outstanding research assistance. This project received funding from the British Academy through Grant SRG2223\231245. The views expressed in this paper are our own and do not necessarily coincide with those of the Banco de España or the Eurosystem.

†Rotman School of Management, University of Toronto, Toronto, ON M5S 3E6, Canada. laurent.cavenaile@rotman.utoronto.ca.

‡Rotman School of Management, University of Toronto, Toronto, ON M5S 3E6, Canada. ruben.gaetani@rotman.utoronto.ca.

§Banco de España and CEMFI, Calle Alcalá 48, 28014, Madrid, Spain. pau.roldan@bde.es.

¶Queen Mary University of London and CEPR, Mile End Road, London E1 4NS, United Kingdom, t.schmitz@qmul.ac.uk.

1 Introduction

Technological breakthroughs give rise to innovation opportunities, enabling firms to create and introduce new products and production processes. An influential tradition in innovation economics argues that these events often follow a life cycle, with predictable time paths for entry, exit, innovation and market concentration. In the seminal paper of Klepper (1996), the author describes these patterns as follows: "When industries are new, there is a lot of entry, firms offer many different versions of the industry's product, the rate of product innovation is high, and market shares change rapidly. Despite continued market growth, subsequently entry slows, exit overtakes entry and there is a shakeout in the number of producers, the rate of product innovation and the diversity of competing versions of the product decline, increasing effort is devoted to improving the production process, and market shares stabilise."[1]

While these life cycles have been widely studied for individual industries (see Audretsch (1987); Jovanovic and MacDonald, 1994; Klepper and Simons, 2000), their aggregate importance and implications are not well understood. For instance, when life cycles are prevalent, aggregate business dynamism might be affected by the distribution of industries across life cycle stages. Moreover, life cycles have unclear implications for public policy: should R&D subsidies be concentrated on young or old industries? Should they target potential entrants trying to create new products, or incumbents aiming to improve process efficiency?

These questions are difficult to answer both because systematic empirical evidence on industry life cycles is scarce, and because there is no theoretical framework to study their quantitative implications in general equilibrium. Our paper aims to fill these gaps by (i) providing systematic empirical evidence on the life cycle behaviour of industries following the introduction of new innovation possibilities, and (ii) developing a rich general equilibrium framework to study industry life cycles and their implications.

We begin our analysis by studying how the arrival of new technological possibilities affects the dynamics of industry entry, exit, and innovation in the data. We leverage the ICT Revolution in the late 1980s as a major shock that generated heterogeneous opportunities across non-ICT industries to introduce new products and processes, while being largely exogenous to idiosyncratic factors that differentially affected industry dynamics. We use data on the universe of patents issued by the United States Patent and Trademark Office (USPTO) since 1976 to build an index of industries' ex-ante exposure to the ICT shock and to measure their response in terms of firms' entry and exit into patenting, and their innovation type. This approach allows us to overcome some of the limitations —such as the intrinsic endogeneity of innovation to the dynamics of entry and exit— which had confined existing empirical evidence on industry life cycles to studies of isolated episodes.

Our empirical findings show that, in response to the ICT shock, industry dynamics followed

[1] While Klepper focuses on new industries, our focus is slightly different, considering existing industries that are revolutionised by a technological breakthrough. However, the life cycle patterns triggered by this breakthrough are the same as the ones described in his quotation.

patterns that are strongly in line with the notion of industry life cycles. We document three key facts. First, following the shock, more exposed industries experienced a phase of rapid growth in the number of patenting firms. This phase was followed by a period in which exit outpaced entry, giving rise to a shakeout that reduced the number of active firms. Second, new inventions in more exposed industries tended to build upon more recent technologies, as inferred by patent citations. This suggests that a potential channel for the spike in entry is the fact that new entrants, being free from the constraints imposed on incumbents by their organisational capital, had an advantage in embracing the new technology. Third, the prevalence of process innovation increased steadily in more exposed industries, suggesting that as industries matured, firms shifted their focus from products to processes to minimise production costs and consolidate their market position.

To study the aggregate implications of these industry life cycles, we build a general equilibrium model with heterogeneous industries and firms. The model features a continuum of industries which are randomly hit by exogenous reset shocks. When an industry is hit by such a shock, it starts anew with a finite set of potential firms that can invest into product innovation in order to enter the market. Once they enter (with a heterogeneous initial productivity, reflecting the quality of their product's design), they can increase their productivity further through process innovation.[2] At every instant, firms with a design compete under oligopolistic competition, as in Atkeson and Burstein (2008). Crucially, each producing firm must pay a fixed cost. Therefore, the number of active firms in each industry is endogenous: some firms may decide to abstain from producing if their operating profit is insufficient to cover the fixed cost. When making production and innovation decisions, firms behave strategically: their profits depend not only on the firm's own design quality and productivity level, but also on those of all of the competitors in the industry. Therefore, there is a two-way interaction between innovation and market structure.

Our model is able to generate life cycles that resemble those observed in the data, with an inverted-U shape evolution of the number of producing firms over time for the average industry. Initially, the number of firms increases, as potential entrants successfully engage in process innovation. Eventually, however, the process innovations of some incumbents generate a shakeout, as the other lagging firms do not find it profitable any more to produce. This shakeout is persistent, as the low probability to (and low profits from) catching up with leaders discourages innovation by laggards and remaining potential entrants.

We use our framework to study the impact of different R&D subsidy policies on welfare and industry life cycles. In particular, we study subsidies to product innovation (which target potential entrants, and therefore younger industries) and process innovation (which target incumbents, and therefore older industries). Our preliminary results indicate that both of these subsidies can increase welfare, but that they affect industry life cycles differently. Product innovation subsidies accelerate the life cycle, leading to a greater initial increase in the number of producing firms and a

[2] Importantly, we assume that the possibilities for process innovation are eventually exhausted. Due to this assumption, ideas are getting harder to find as an industry ages. This is supported by empirical evidence (see e.g. Bloom *et al.*, 2020; Jones, 2022).

faster shakeout. Process innovation, on the other hand, also tends to speed up shakeouts, but lowers the number of producing firms in the average industry. Both subsidies decrease the aggregate markup and raise aggregate productivity and output. Moderate subsidies also increase consumer welfare, as their positive effect on output outweighs their resource cost.

Literature Review The idea that new technological innovations generate life cycle patterns goes back to the work of Schumpeter (1911, 1934). Schumpeter called for a wealth of industry-level studies showing how innovations entered the economy, and this call sparked a rich empirical and theoretical literature.[3] Many empirical studies have tried to uncover life cycle regularities. For instance, Gort and Klepper (1982) analysed the time path of 46 manufacturing products introduced between the 1890s and the 1940s, while Jovanovic and MacDonald (1994) and Klepper and Simons (2000) focused on the US tire industry.[4] Theoretical work has investigated the drivers of these patterns. For instance, Abernathy and Utterback (1978) claimed that shakeouts were driven by the eventual emergence of a "dominant design", forcing all firms which are unable to produce this design to exit. Jovanovic and MacDonald (1994) emphasise the heterogeneous success of firms in keeping up with (exogenous) technology shocks, while Klepper (1996) stresses the importance of incumbent process innovation.[5]

Our paper contributes makes both an empirical and a theoretical contribution to this literature. Empirically, we provide a broader perspective, focusing on the universe of (non-ICT) industries rather than a subsample of industries and products.[6] Theoretically, we introduce for the first time industry life cycles into a rich general equilibrium model. This allows us to go beyond the partial equilibrium models of the literature and to study important welfare questions, such as the implications of life cycles for R&D policy.

Besides the life cycle literature, our work is also related to a broader literature studying the interaction between innovation and market structure, and the resulting policy implications (see e.g. Acemoglu *et al.*, 2018; Cavenaile *et al.*, 2023; Akcigit and Ates, 2023). It also relates to a literature on technology waves (see e.g. Berkes *et al.*, 2023, Ma and Yang, 2023). We depart from this literature by introducing industry life cycles and studying their implications.

The rest of the paper is organised as follows: Section 2 presents our empirical results. Section 3 introduces our general equilibrium model. Section 4 proposes a parameterization of our model and discusses its properties. Section 5 reports the results of our policy experiments for both product and process innovation subsidies. Section 6 offers some concluding remarks.

[3]For a more thorough discussion of Schumpeter's ideas and research program, see McCraw (2010).

[4]Klepper (1997) studies the US car industry, and reviews studies for six other innovative industries: "typewriters, automobile tires, commercial aircraft for trunk carriers, televisions, television picture tubes and penicillin". Note that this literature has generally used the terms "industry life cycle" and "product life cycle" interchangeably.

[5]Jovanovic and Tse (2010) instead explain shakeouts through a model of vintage capital.

[6]Some existing studies do consider large samples, such as Dinlersoz and MacDonald (2009) for the US or Bos *et al.* (2013) for Europe. However, these studies do not rely on a clearly identified technology shock, and do not use patent data.

2 Empirical motivation

While recent history offers an abundance of episodes in which industries went through the life cycle patterns described by Klepper (1996), systematic empirical evidence documenting the properties of these life cycles is remarkably scarce. This is primarily due to the intrinsic endogeneity of technological possibilities to other industry-level shifters and the difficulty in defining credible control groups. These obstacles have confined empirical work on industry life cycles to descriptive studies of individual episodes, with a few notable exceptions (e.g., Gort and Klepper, 1982; Audretsch, 1987; Dinlersoz and MacDonald, 2009).

In this section, we propose an empirical approach to overcome those limitations and study systematically how industry dynamics respond to changes in technological possibilities. We leverage the arrival of the ICT Revolution in the late 1980s as a major shock that, while being largely exogenous to idiosyncratic industry shifters, differentially affected technological possibilities across non-ICT industries, to an extent that is predicted by industry characteristics *before* the onset of the shock. This allows us to compare the dynamics of sectors that were highly exposed to the ICT shock with those that were less exposed.

2.1 Data and Empirical Setting

Our main source of data is the universe of patents granted by the USPTO between 1976 and 2019 and provided by PatentsView.[7] For each patent grant, we collect information on the identity of the assignee, the set of listed Cooperative Patent Classification (CPC) technology classes, the year of filing, and the patents cited as prior art. PatentsView provides a disambiguation of unique assignees (i.e., the firm or organization to which the patent is granted) in the form of an assignee identifier. We attach to each assignee the 6-digit industry NAICS code corresponding to the modal sector in its own distribution, using the crosswalk developed by Lybbert and Zolas (2014) to map CPC classes into industry codes. We label CPC classes as ICT-related if they appear in the list of ICT technologies in Braguinsky *et al.* (2023) (Table A2).[8] Finally, we measure the intensity of process innovation using the dataset developed by Bena and Simintzi (2022), which provides the count of process-related claims in each granted patent.

To minimize measurement error, we construct industry-level observations at a 5-year frequency between 1985 and 2019 (e.g., 1985-89, 1990-94, etc...) and restrict the sample to industries with at least 10 patents in each of the 5-year windows. Our approach relies on the exogeneity of the ICT shock with respect to other idiosyncratic determinants of industry-level dynamics throughout the sample period. For this reason, we restrict the sample to industries that do *not* appear in the list of ICT sectors in Braguinsky *et al.* (2023) (Table A1). This delivers a balanced panel of 167 6-digit NAICS industries.

[7]The data can be downloaded from `https://patentsview.org/download/data-download-tables`.

[8]We assign patents with multiple CPC classes to the set of ICT inventions in proportion to their share of ICT-related classes.

Figure 1: Share of ICT patents in non-ICT industries

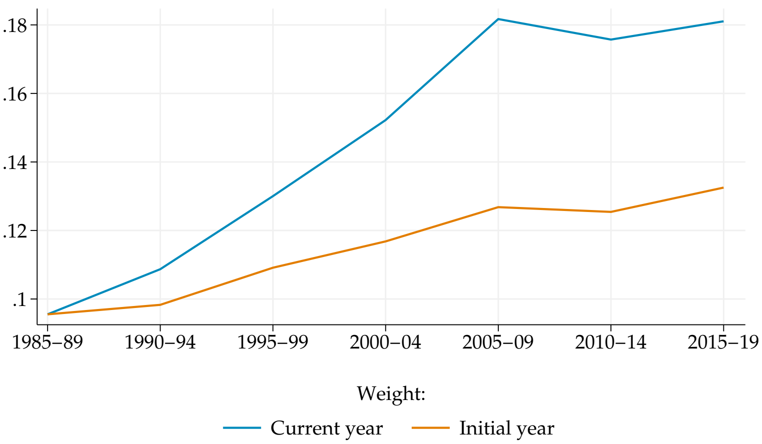

Notes: The sample includes only industries that do not appear in the list of ICT sectors in Braguinsky *et al.* (2023) (Table A1). The blue line weights industries by their number of patents in the current period, while the orange line weights industries by their number of patents in the first period (1985-89).

The blue line in Figure 1 displays the share of patents in non-ICT industries belonging to ICT-related technology classes over the time period covered by our sample. This share increased from 9.6% in 1985-89 to 18.1% in 2015-19. The orange line in the same figure shows the corresponding share when each industry is weighted in proportion to their number of patents in the first period. The orange line explains roughly half of the overall increase, suggesting that the expansion of ICT patenting was due both to an increase in the prevalence of ICT patenting for the average sector as well as to a shift in patenting towards more ICT-intensive sectors.

The data reveal large heterogeneity across industries in the extent to which the ICT Revolution was reflected in innovation output. The top-left and top-right panels of Figure 2 show the distribution across sectors (weighted by the total number of patents in the first period) of the growth rate of ICT-related patenting between the initial period (1985-89) and the intermediate (2000-04, left) and final (2015-19, right) periods of the sample. Both distributions display sizeable variation, with standard deviations as large as 0.92 and 0.71 of their respective means.

This heterogeneity is visible not only in the variation across industries in the *realized* uptake of ICT-related ideas, but also in the *ex-ante* exposure of those industries to the ICT revolution. The bottom panel of Figure 2 shows the distribution across industries of the share of patents in ICT-related classes in the first period of the sample (1985-89), which displays a standard deviation (0.127) equal to 1.33 times its mean (0.096). Among the industries with more than 1,000 patents in 1985-89, the ones with the highest degree of exposure are NAICS codes 515120 ("Television Broadcasting"), 334613 ("Blank Magnetic and Optical Recording Media Manufacturing"), and 339112 ("Surgical and Medical Instrument Manufacturing").[9]

The initial share of patents in ICT-related classes is a simple measure of an industry's ex-ante

[9] Note that, as explained above, ICT industries are not included in the sample.

Figure 2: The ICT shock across non-ICT industries

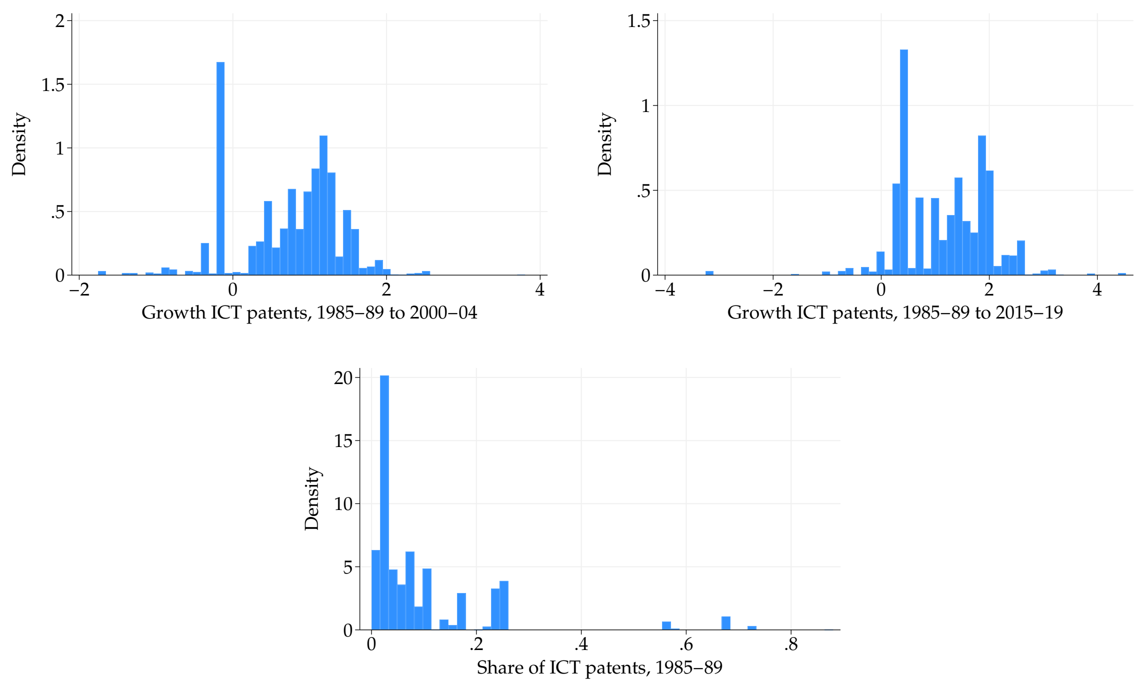

Notes: Top panels: Growth in the number of ICT patents between 1985-89 and 2000-04 (left) and 2015-19 (right). Bottom panel: share of ICT patents in 1985-89. All histograms are weighted by the number of patents in 1985-89. The sample includes only industries that do not appear in the list of ICT sectors in Braguinsky *et al.* (2023) (Table A1).

exposure to the arrival of the ICT revolution and it turns out to be a strong predictor of the degree of industry innovativeness in the decades that follow. Figure 3 shows weighted scatter plots of the share of 1985-89 patents in ICT-related classes (horizontal axis) and the growth rate of total patenting between 1985-89 and the intermediate (2000-04, left panel) and final (2015-19, right panel) periods of the sample. Both scatter plots show a strong correlation, with the R^2 of the underlying regressions equal to 0.14 and 0.10, respectively.

2.2 Results

In this section, we explore how the differential arrival of technological opportunities across sectors, as proxied by their ex-ante exposure to the ICT revolution, affected the patterns of entry, exit, and innovation. To this end, we estimate local projection models of the following form:

$$y_{i,t} = \alpha_t + \beta_t \,\text{Share ICT}_{i,1985-89} + \gamma_t \, X_{i,1985-89} + \epsilon_{i,t}, \qquad (1)$$

where $y_{i,t}$ denotes any of the outcomes observed for industry i in the 5-year window t. The main explanatory variable, Share $\text{ICT}_{i,1985-89}$, is defined as the share of ICT patenting for industry i in the initial period of the sample (1985-89). The vector $X_{i,1985-89}$ includes controls for the log of the initial number of patents, the log of the initial number of assignees, and the initial value of the outcome variable, $y_{i,1985-89}$.

Figure 3: The ICT shock and industry-level patent growth

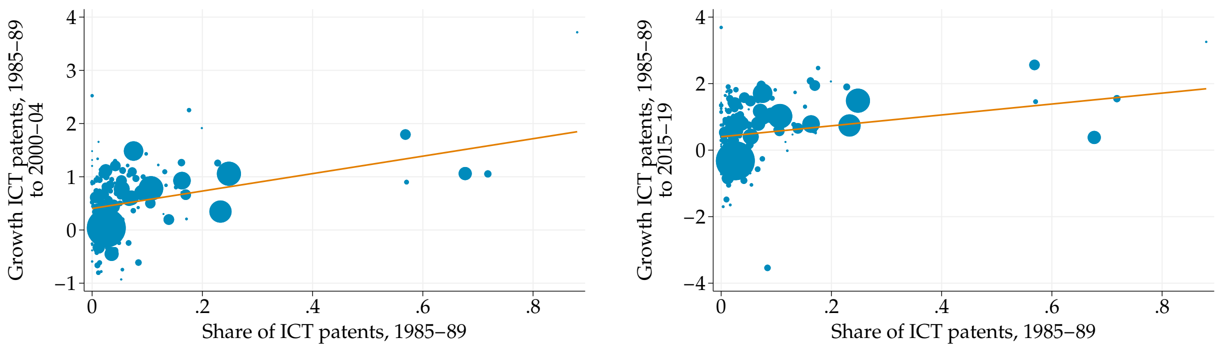

Notes: Share of ICT patents in 1985-89 (horizontal axis) and growth in total patenting (vertical axis) between 1985-89 and 2000-04 (left panel, $N = 167$, $R^2 = 0.14$) and 2015-19 (right panel, $N = 167$, $R^2 = 0.10$), weighted by the number of patents in 1985-89. The sample includes only industries that do not appear in the list of ICT sectors in Braguinsky *et al.* (2023) (Table A1).

First, we show that more exposed industries experience a sharp increase in net entry after the shock, followed by a shakeout that reduces the number of active firms. Second, we show that, following the shock, new inventions in more exposed industries build more intensely upon more recent technologies, suggesting a potential channel that prompts an increase in entry after the shock. Third, we show that the prevalence of process innovation gradually increases over time in more exposed industries relative to less exposed ones. Taken together, the evidence we uncover suggests that new technological possibilities give rise to rich dynamics resembling those proposed by existing theories of industry life cycles.

Fact 1: More exposed industries experience an increase in net entry, followed by a shakeout

We first explore how the arrival of the ICT Revolution affected the dynamics of entry and exit across industries.

The top-left panel of Figure 4 shows the estimates of β_t in Equation (1), when the industry-level entry rate—defined as the ratio of the number of *new* assignees in the current period over the total number of assignees active in the previous period—is used as the dependent variable.[10] Entry rates in 1990-94 increase sharply in more exposed industries. The size of the estimated coefficient implies that a one-standard deviation increase in exposure to the ICT shock leads to a 1.13 percentage points (0.083 standard deviations) increase in the entry rate. The effect reaches a peak in the following period (1995-99), with a one-standard deviation increase in exposure leading to a 3.83 percentage points (0.248 standard deviations) increase in the entry rate. The effect becomes weaker in later periods, until the point estimates turn negative from 2010-14 onwards.

[10] Since all regressions control for the initial value of the outcome variable, the coefficient in 1985-89 is equal to zero by construction.

Figure 4: The effect of the ICT shock on entry and exit rates

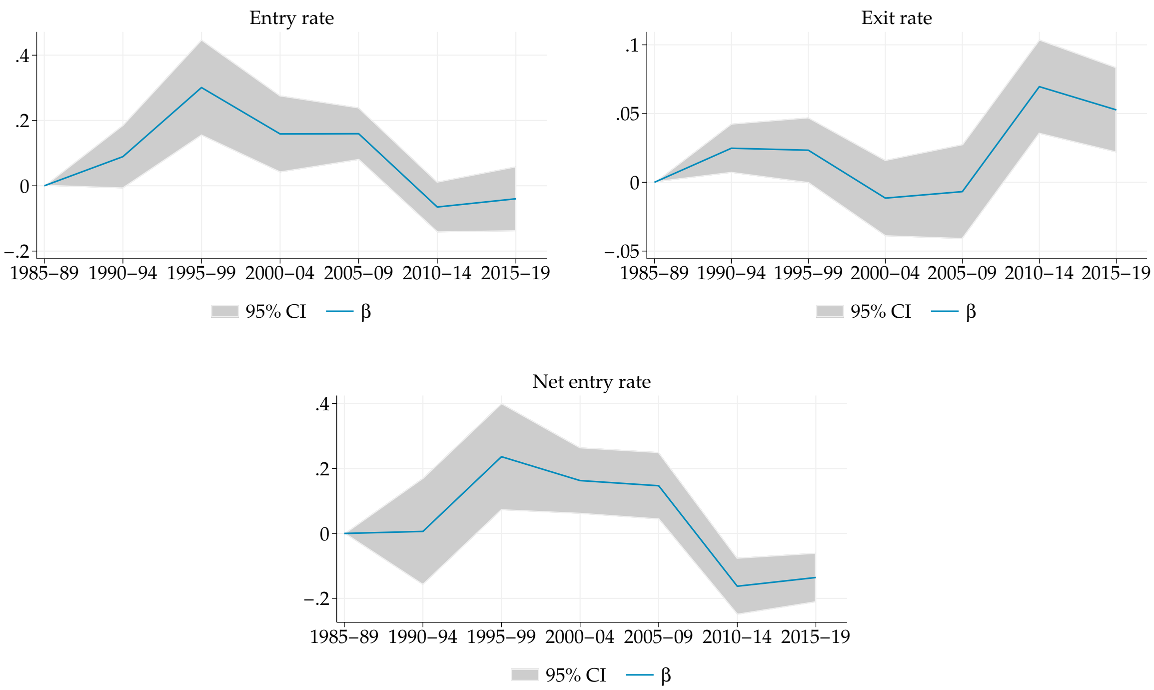

Notes: Estimates of β_t in Equation (1). The outcome is the entry rate (top-left panel), the exit rate (top-right panel), and the net entry rate (bottom panel). Confidence intervals are computed by clustering standard errors at the 2-digit NAICS level. All regressions are weighted by the number of patents in 1985-89. The sample includes only industries that do not appear in the list of ICT sectors in Braguinsky *et al.* (2023) (Table A1). $N = 167$.

The top-right panel of Figure 4 shows the corresponding estimates when the industry-level exit rate—defined as the ratio of the number of assignees in the previous period that are no longer active over the total number of assignees active in the previous period—is used as the dependent variable. After a small rise in 1990-94 (0.32 percentage points for each standard deviation in the exposure measure), there is no significant difference in exit rates between more and less exposed industries until later in the sample. The effect becomes visible in 2010-14, where a one standard deviation increase in exposure leads to a 0.89 percentage points (0.12 standard deviations) increase in the exit rate.

Finally, the bottom panel of Figure 4 shows the impact of the ICT shock on the net entry rate, defined as the difference between the entry and the exit rate. In the early periods following the shock, the effect on entry is significantly larger than the one on exit, leading to a positive effect on the net entry rate. At its peak, in 1995-99, a one standard deviation increase in exposure leads to a 3.01 percentage points (0.193 standard deviations) increase in the net entry rate. In later periods, more exposed industries experience shakeouts in which exit overtakes entry and the effect on net entry becomes negative. In 2010-14, a one standard deviation increase in exposure leads to a 2.07 percentage points (0.133 standard deviations) drop in the net entry rate.

What explains these life-cycle dynamics following a major shock to technological opportunities in an industry? In the remainder of this section, we provide suggestive evidence that changes in the nature of firm's innovation play a significant role in shaping these dynamics.

Fact 2: Innovation in more exposed industries builds upon newer technologies

A large literature in managerial and organizational economics has argued that new entrants are often in a better position to take advantage of new technological opportunities compared to incumbents (e.g., Henderson and Clark, 1990; Christensen, 2013). Proposed explanations emphasize the constraints that incumbents' organizational capital and innovation capabilities, which are optimized for the old technology, impose on their ability to embrace new opportunities. We now show that, consistently with this intuition, new inventions in more exposed industries tend to build upon more recent technologies. This observation can explain the sharp positive response of the entry rate after the onset of the shock.

Figure 5: The effect of the ICT shock on the age of patent citations

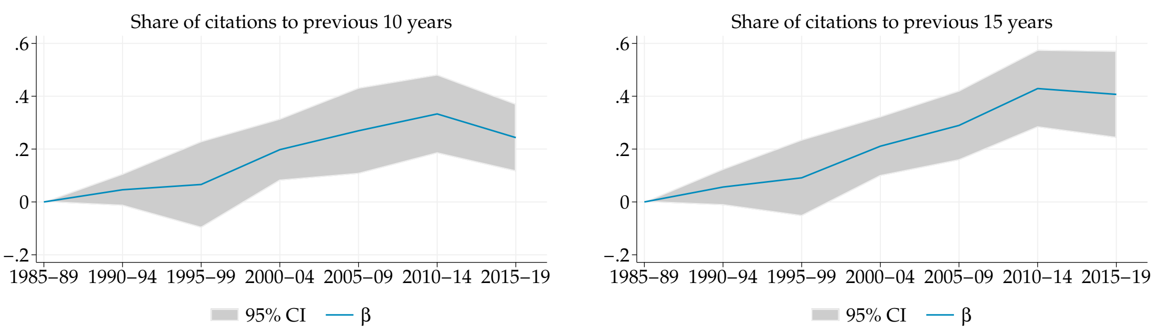

Notes: Estimates of β_t in Equation (1). The outcome is the share of citations to patents filed no earlier than 10 years (left panel) or 15 years (right panel) before the filing year of the citing patent. Confidence intervals are computed by clustering standard errors at the 2-digit NAICS level. All regressions are weighted by the number of patents in 1985-89. The sample includes only industries that do not appear in the list of ICT sectors in Braguinsky *et al.* (2023) (Table A1). $N = 167$.

Figure 5 shows the estimates of β_t in Equation (1), when the outcome variable is constructed as the share of backward citations to grants filed in the 10 (left panel) and 15 (right panel) years before the filing year of the citing patent. Positive estimates of β_t imply that, in more exposed industries, new inventions tend to build upon more recent ideas. The results show a positive response in the share of citations to more recent patents, which becomes larger in magnitude over time. At its peak (in 2010-14), a one-standard deviation increase in exposure is associated to a 4.24 percentage point (0.556 standard deviations) increase in the share of citations to patents that are less than ten years old, and a 5.47 percentage point (0.566 standard deviations) increase in the share of citations to patents that are less than fifteen years old.

Fact 3: Over time, more exposed industries shift from product to process innovation.

As a final step, we show that the prevalence of process innovation increases gradually but steadily over time in more exposed industries relative to less exposed ones. This finding suggests that as the new technology matures, firms have an opportunity to consolidate their position via process innovation, making entry less attractive and driving out of the market firms that are unsuccessful at bringing down their cost of production. These two forces will play a major role in generating rich industry dynamics in our model.

Figure 6: The effect of the ICT shock on process innovation

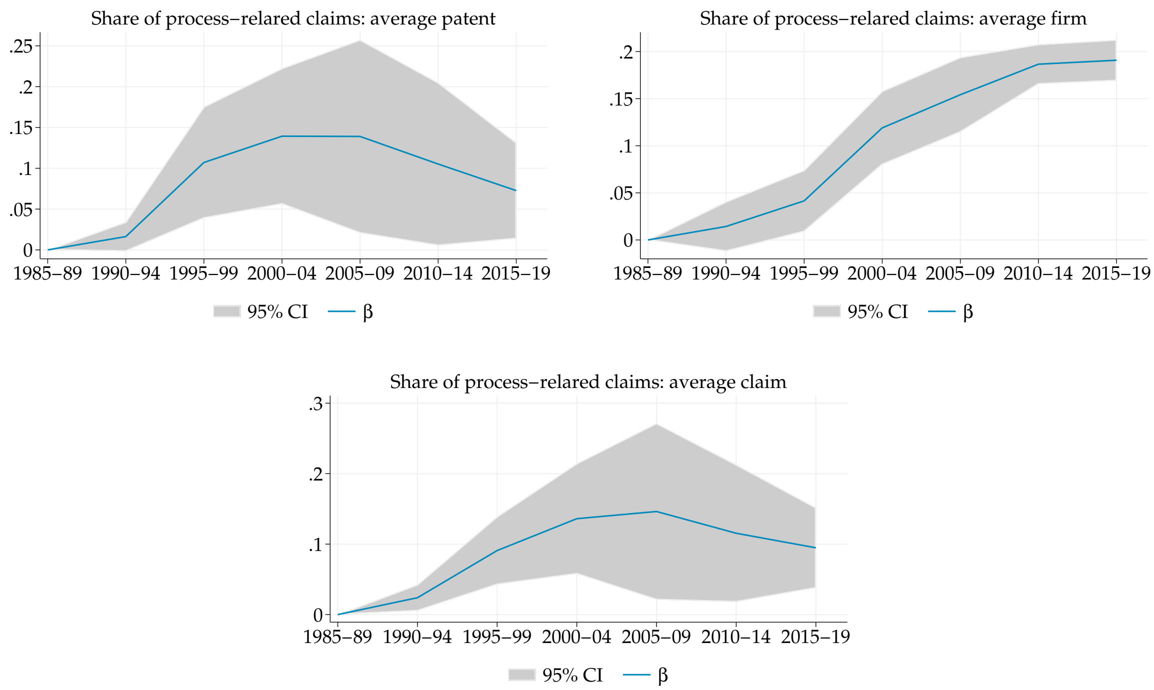

Notes: Estimates of β_t in Equation (1). The outcome is the share of patent claims that are identified as process-related by Bena and Simintzi (2022), aggregate at the industry-time level as the average patent (top-left), the average firm (top-right), and the average claim (bottom) in the industry. Confidence intervals are computed by clustering standard errors at the 2-digit NAICS level. All regressions are weighted by the number of patents in 1985-89. The sample includes only industries that do not appear in the list of ICT sectors in Braguinsky *et al.* (2023) (Table A1). $N = 167$.

The three panels in Figure 6 display the estimates of β_t in Equation (1), when the outcome variable is defined as the prevalence of process innovation in an industry's patenting output. To measure process innovation, we leverage the dataset developed by Bena and Simintzi (2022), who use textual analysis to construct counts of the number of process-related claims in all USPTO patents since 1976. The three panels of Figure 6 correspond to different ways of aggregating this patent-level measure of process innovation to the industry level, as the fraction of process-related claims for the average patent (top-left panel), the average firm (top-right panel), and the industry overall (bottom panel).

The results across these different measures convey a common message. The prevalence of process innovation in more exposed industries increases steadily over the sample period. In 2010-14, a one-standard deviation increase in exposure is associated to a 2.37 percentage point (0.149 standard deviations) increase in the share of process-related claims for the average firm in the industry. As industries mature and uncertainty on the product and market conditions attenuates, firms leverage process innovation to lower the cost of production and consolidate their position in the industry.

3 Model

In this section, we build a general equilibrium model of industry life cycles that can help us rationalize the patterns described in Section 2. Our model builds on the oligopolistic competition framework of Atkeson and Burstein (2008), adding endogenous entry, exit, product and process innovation and allowing for rich firm and industry dynamics. It combines features of the dominant design model of Abernathy and Utterback (1978) and of the life cycle theory of Klepper (1996) into a general-equilibrium macroeconomic framework. This allows us to study the determinants of the life cycle and how the distribution of industries over different life cycle states matters for aggregate outcomes.

3.1 Environment

Households Time is continuous, infinite, and indexed by $t \in \mathbb{R}^+$. An infinitely-lived representative household maximises lifetime utility:

$$\int_0^{+\infty} e^{-\rho t} \ln(C_t) \mathrm{d}t, \quad \text{subject to } \dot{\mathcal{A}}_t = r_t \mathcal{A}_t + w_t L - C_t, \tag{2}$$

where $\rho > 0$ is the discount rate, C_t is consumption of a single final good (the numeraire), $\mathcal{A}_t$ stands for the household's asset holdings, and w_t is a wage earned from inelastically supplied labour services. Aggregate labour supply is fixed at L.

The final good is produced by a representative firm using as inputs the goods produced by a continuum of industries. These inputs are combined in a constant-elasticity-of-substitution (CES) technology:

$$Y_t = \left(\int_0^1 \left(Y_{i,t}\right)^{\frac{\eta-1}{\eta}} \mathrm{d}i \right)^{\frac{\eta}{\eta-1}} \tag{3}$$

where $\eta > 1$ is the elasticity of substitution between industries. Each industry i has a finite number $\overline{N} \in \mathbb{N}$ of firms (a parameter). Each of these firms can produce one variety $n \in \{1, \ldots, \overline{N}\}$ of a differentiated good. The industry good is a CES aggregator of varieties:

$$Y_{i,t} = \left(\sum_{n=1}^{\overline{N}} \left(\omega_{in,t}\right)^{\frac{1}{\varepsilon}} \left(y_{in,t}\right)^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}, \tag{4}$$

where $\varepsilon > \eta$ is the elasticity of substitution between varieties within an industry. In this expression, $\omega_{in,t}$ captures the quality of the *design* of firm n in industry i at time t. The actual number of operating firms within the industry is endogenous, as will be described later.

Firms can produce with a technology that is linear in labour:

$$y_{in,t} = q_{in,t} \ell_{in,t}, \tag{5}$$

where $q_{in,t} > 0$ is the labour productivity of firm n in industry i at time t. This productivity can be built over time through innovation, in a process that we describe later. Finally, all firms face a flow operating cost equal to ϕY_t, where $\phi > 0$ is a parameter.

Within-period competition At each instant t, the $\overline{N}$ firms in each industry engage in static Bertrand competition. That is, they choose a price that maximises their instantaneous profit, taking the prices of all other firms as given. To characterise the equilibrium, in what follows it will be useful to define the quality-adjusted productivity of a firm as

$$a_{in,t} \equiv (\omega_{in,t})^{\frac{1}{\varepsilon-1}} q_{in,t}.$$

As there is a fixed cost of production, some firms may decide not to produce in equilibrium. Moreover, the fixed cost can create multiple equilibria in the static Bertrand game. To select one of them, we assume (as in Atkeson and Burstein, 2008) that firms decide sequentially whether to operate or not, in the ascending order of their quality-adjusted productivity $a_{in,t}$.

Industry life cycles and R&D Each industry is reset at a Poisson arrival rate $\delta > 0$, which we treat as a parameter. When this shock occurs, the industry (and all its firms) disappear, and a new industry appears in its place. When industry i starts off anew at some time t_0, we assume that no firm has yet found a design to produce. Thus, for any firm $n \in \{1, \ldots, \overline{N}\}$, we have $\omega_{in,t_0} = 0$. Firms discover their designs through endogenous product R&D. Once they have found their design, they further increase their labour productivity through process R&D.

Consider first a firm that has not yet found its design (i.e., with $\omega_{in,t} = 0$ at some time $t \geq t_0$). By investing $\zeta x^{\psi} Y_t$ units of the final good into product R&D, this firm generates a Poisson arrival rate $x > 0$ for finding a design, where $\zeta > 0$ and $\psi > 1$ are R&D cost parameters. When the firm finds a design, it draws a design quality ω from some distribution over a finite set $\Omega \equiv \{\omega_1, \ldots, \omega_{k_{\max}}\}$ of possible design qualities. The probability of drawing a given design $\omega_k \in \Omega$ is an exogenous parameter, denoted by $\tau_k \geq 0$, with $\sum_k \tau_k = 1$. The firm is assumed to keep this design forever. After the draw, the firm can start to produce output with a labour productivity equal to $\underline{q} > 0$.

Next, consider a firm that already has a design. These firms can expand their productivity on a finite quality ladder by doing process innovation. We assume that by investing $\chi z^{\psi} Y_t$ into process R&D, this firm generates a Poisson arrival rate $z > 0$ for a process innovation, where $\chi > 0$ is a parameter. A successful process innovation increases the firm's labour productivity by a factor of $\lambda > 1$. We assume that a firm can do at most $(j_{\max} - 1)$ process innovations, so that there are $j_{\max}$ distinct productivity states, $q_{in,t} \in \{q_0, q_1, \ldots, q_{j_{\max}-1}\}$, where $q_j \equiv \lambda^j \underline{q}$.

3.2 Equilibrium

To solve the model, we compute a Nash equilibrium in policies: every firm $n \in \{1, \ldots, \overline{N}\}$ chooses its optimal R&D policy taking the R&D policies of all other firms in its industry as given.[11] Specifically, we focus on a steady-state equilibrium, in which aggregate output and wages are constant (i.e. $Y_t = Y$ and $w_t = w$). We will first characterise the model's static solution for a given level of quality and labour productivity for every firm in the economy, and then turn to the dynamic decisions which pin down the equilibrium distributions of quality and labour productivity. Finally, we will derive the stationary distribution of industries, and use it to compute economic aggregates.

3.2.1 Static Industry Equilibrium

The demand for the output of each industry is given by $Y_{i,t} = P_{i,t}^{-\eta} Y$, where $P_{i,t}$ is the price of the final output of industry i at time t. Within each industry, in turn, demand for the good produced by firm n holds

$$y_{in,t} = \omega_{in,t} \left(\frac{p_{in,t}}{P_{i,t}} \right)^{-\varepsilon} Y_{i,t}. \tag{6}$$

Conditional on producing, firm $n \in \{1, \ldots, \overline{N}\}$ optimally chooses a markup

$$m_{in,t} = \frac{\varepsilon - (\varepsilon - \eta) \frac{p_{in,t} y_{in,t}}{P_{i,t} Y_{i,t}}}{\varepsilon - 1 - (\varepsilon - \eta) \frac{p_{in,t} y_{in,t}}{P_{i,t} Y_{i,t}}}, \tag{7}$$

where $\frac{p_{in,t} y_{in,t}}{P_{i,t} Y_{i,t}}$ is the market share of firm n in industry i at time t, computed as the sales of the firm relative to the total sales of the industry, since $P_{i,t} Y_{i,t} = \sum_{i=1}^{\overline{N}} p_{in,t} y_{in,t}$. Formula (7) says that the markup is higher when the firm is relatively larger in its industry, with a sensitivity that depends on the relative elasticity of substitution between firms within industries relative to the elasticity of substitution across industries. In turn, the market share holds

$$\frac{p_{in,t} y_{in,t}}{P_{i,t} Y_{i,t}} = \left(\frac{m_{in,t} w}{a_{in,t} P_{i,t}} \right)^{1-\varepsilon}. \tag{8}$$

The industry price index is given by

$$P_{i,t} = M_{i,t} \frac{w}{A_{i,t}}, \quad \text{where } A_{i,t} \equiv \left(\sum_{n=1}^{\overline{N}} (a_{in,t})^{\varepsilon - 1} \mathbb{1}_{in,t} \right)^{\frac{1}{\varepsilon - 1}}. \tag{9}$$

In this expression, $A_{i,t}$ is the quality-adjusted productivity of the industry, and $\mathbb{1}_{in,t}$ is a indicator variable equal to 1 if firm n of industry i produces at time t, and to zero otherwise. Equation (9)

[11] Note that there is no interaction between the static price choices and the dynamic R&D choices, as price choices do not affect any state variable of the firm. Therefore, the equilibrium that we consider here (with static price decisions) is Markov-perfect. This is not to say, of course, that there might not be other equilibria (e.g. with dynamic collusion), but for simplicity we do not focus on these type of solutions.

shows that the industry-level price is a markup over the industry-level marginal cost. Moreover, the industry-level markup holds

$$M_{i,t} \equiv \left(\sum_{n=1}^{\overline{N}} \left(m_{in,t}\right)^{1-\varepsilon} \left(\frac{a_{in,t}}{A_{i,t}}\right)^{\varepsilon-1} \mathbb{1}_{in,t} \right)^{\frac{1}{1-\varepsilon}}. \tag{10}$$

That is, the industry-level markup is a weighted average of the markups of producing firms. Putting the previous two equations together, we can write market shares as

$$\frac{p_{in,t} y_{in,t}}{P_{i,t} Y_{i,t}} = \left(\frac{a_{in,t} / A_{i,t}}{m_{in,t} / M_{i,t}} \right)^{\varepsilon-1}. \tag{11}$$

With this, profits hold

$$\pi_{in,t} = \left(\left(m_{in,t} - 1\right) m_{in,t}^{-\varepsilon} \left(a_{in,t}\right)^{\varepsilon-1} \left(\frac{M_{i,t}}{A_{i,t}}\right)^{\varepsilon-\eta} w^{1-\eta} - \phi \right) Y. \tag{12}$$

Given a wage w and the quality-adjusted productivity of all firms, $\{a_{in,t}\}_{n=1}^{\overline{N}}$, the above equations allow us to compute the static equilibrium for an industry (that is, for the production decision of all firms, as well as the market shares and markups of all producing firms).

To compute this solution, we first rank firms by their quality-adjusted productivity. For a given guess for the set of producing firms, replacing equation (7) into equation (11) gives a system of $\overline{N}$ non-linear equations that can be solved numerically for the market shares of all firms. From this, we can deduce markups and profits. The set of producing firms is pinned down by the condition that (i) the least productive producing firm must cover its fixed cost ϕ, and (ii) either the most productive non-producing firm would make a loss if it were to enter, or all firms produce.

The quality-adjusted productivity of firms is determined endogenously, through firms' R&D decisions. In the next section, we characterise these decisions.

3.2.2 Dynamic Equilibrium Conditions

Firms with a design We first consider firms which already have a design. Throughout, we will label points on the design quality grid by k and points on the labour productivity grid by j. In total, there are $k_{\max} \times j_{\max}$ potential states in which a firm with a design might find itself.

The value function of a firm with a design not only depends on the two endogenous state variables k and j, but also on a set of state variables (exogenous from the point of view of the firm) summarising the state of its competitors, which we denote by $(N_0^C, \mathbf{N}^C)$. Here, N_0^C stands for the number of competitor firms within the industry which have not discovered a design yet. In turn, $\mathbf{N}^C$ is a $k_{\max}$-by-$j_{\max}$ matrix in which the row k', column j' element, denoted $N_{k'j'}^C$, corresponds to the number of competitor firms with design quality $\omega_{k'}$ and labour productivity $q_{j'}$. For reference, henceforth we let $\mathbf{I}_{kj}$ denote a $k_{\max}$-by-$j_{\max}$ indicator matrix in which the line k, column j element is 1 and all other elements are zero.

As profits and all cost functions scale linearly in aggregate output, the value function takes the form $V(k, j, N_0^C, \mathbf{N}^C)Y$.[12] Moreover, firms which have not yet discovered a design are irrelevant for static profits (as they are unable to produce), so the profit function π does not depend on N_0^C. Then, the Hamilton-Jacobi-Bellman (HJB) equation of the firm with a design is

$$\begin{aligned}\rho V\left(k, j, N_0^C, \mathbf{N}^C\right) = \max_z \Bigg\{ & \pi(k, j, \mathbf{N}^C) - \chi z^\psi + z\left[V\left(k, \min\left(j+1, j_{\max}\right), N_0^C, \mathbf{N}^C\right) - V\left(k, j, N_0^C, \mathbf{N}^C\right)\right] \\ & + N_0^C\, \widetilde{x}\left(k, j, N_0^C, \mathbf{N}^C\right)\left[\sum_{k'=1}^{k_{\max}} \tau_{k'} V\left(k, j, N_0^C - 1, \mathbf{N}^C + \mathbf{I}_{k'1}\right) - V\left(k, j, N_0^C, \mathbf{N}^C\right)\right] \\ & + \sum_{k'=1}^{k_{\max}} \sum_{j'=1}^{j_{\max}} N_{k'j'}^C \widetilde{z}_{k'j'}\left(k, j, N_0^C, \mathbf{N}^C\right)\left[V\left(k, j, N_0^C, \mathbf{N}^C - \mathbf{I}_{k'j'} + \mathbf{I}_{k', \min(j'+1, j_{\max})}\right) - V\left(k, j, N_0^C, \mathbf{N}^C\right)\right]\Bigg\}.\end{aligned} \tag{13}$$

In this expression, the first line is the profit flow from sales net of costs from process R&D, and the change in the value of the firm when it performs a process innovation, which allows the firm to advance in the productivity space. The second line is the change in the value of the firm that results from a successful product innovation by any of the N_0^C competitor firms in the industry that do not yet have a design. On this line, $\widetilde{x}\left(k, j, N_0^C, \mathbf{N}^C\right)$ stands for the process innovation rate chosen by firms without a design if the incumbent is in state $\left(k, j, N_0^C, \mathbf{N}^C\right)$.[13] In equilibrium, this rate will be directly related to the policy function of firms without a design, which we derive below. Finally, the third line of equation (13) is the change in the value of the firm due to the process innovation of the competitors that have a design, where $\widetilde{z}_{k'j'}\left(k, j, N_0^C, \mathbf{N}^C\right)$ is the process innovation rate chosen by competitors of type (k', j') when the incumbent firm is in state $\left(k, j, N_0^C, \mathbf{N}^C\right)$.

In equilibrium, it must be that

$$\widetilde{z}_{k'j'}\left(k, j, N_0^C, \mathbf{N}^C\right) = z\left(k', j', N_0^C, \mathbf{N}^C - \mathbf{I}_{k'j'} + \mathbf{I}_{kj}\right),$$

so that the process innovation policy rate of a firm's competitor in a given competitor state (k', j') is consistent with the process innovation policy that the firm itself makes when it finds itself in that state (k', j'). The first-order condition in equation (13) is:

$$z(k, j, N_0^C, \mathbf{N}^C) = \left(\frac{V\left(k, \min\left(j+1, j_{\max}\right), N_0^C, \mathbf{N}^C\right) - V\left(k, j, N_0^C, \mathbf{N}^C\right)}{\chi\psi}\right)^{\frac{1}{\psi-1}}. \tag{14}$$

Firms without a design Next, consider the problem of a firm that has not yet found a design. This firm is spending resources on product R&D in order to find a design. The state variables of

[12]Note that wages w are an aggregate state variable for the firm as well, as they affect flow profits (equation (12)). To alleviate notation, however, we do not write profit, value, and policy functions explicitly as functions of w.

[13]Here and everywhere, we use tildes on rates when we refer to the policy of a competitor (which is taken as given by the firm in question).

such a firm are N_0^C and $\mathbf{N}^C$, just as for an incumbent firm with a design, as the firm must take into account the competitive landscape of the industry when making entry decisions. The HJB equation of this firm is

$$\rho V^E\left(N_0^C,\mathbf{N}^C\right) = \max_x \Bigg\{ -\zeta x^\psi + x\left[\sum_{k=1}^{k_{\max}} \tau_k V\left(k,1,N_0^C,\mathbf{N}^C\right) - V^E\left(N_0^C,\mathbf{N}^C\right)\right]$$
$$+ N_0^C \widetilde{x}^E\left(N_0^C,\mathbf{N}^C\right)\left[\sum_{k'=1}^{k_{\max}} \tau_{k'} V^E\left(N_0^C-1,\mathbf{N}^C+\mathbf{I}_{k'1}\right) - V^E\left(N_0^C,\mathbf{N}^C\right)\right]$$
$$+ \sum_{k'=1}^{k_{\max}}\sum_{j'=1}^{j_{\max}} N_{k'j'}^C \widetilde{z}_{k'j'}^E\left(N_0^C,\mathbf{N}^C\right)\left[V^E\left(N_0^C,\mathbf{N}^C-\mathbf{I}_{k'j'}+\mathbf{I}_{k',\min(j'+1,j_{\max})}\right) - V^E\left(N_0^C,\mathbf{N}^C\right)\right]\Bigg\}. \tag{15}$$

The first term on the right-hand side of the first line shows the process R&D costs and the change in value due to finding a design. This occurs at rate x, after which the firm finds design ω_k with probability τ_k and starts with the baseline productivity level $\underline{q}$ (corresponding to state $j = 1$). The second line includes the change in firm value that results from a product innovation by any of the N_0^C competitor firms without a design. Finally, the last line is the change in value resulting from a successful process innovation by firms with a design. Here, $\widetilde{x}^E\left(N_0^C,\mathbf{N}^C\right)$ and $\widetilde{z}_{k'j'}^E\left(N_0^C,\mathbf{N}^C\right)$ are the competitor innovation rates faced by firms without a design in state $\left(N_0^C,\mathbf{N}^C\right)$. In equilibrium, these innovation rates must be consistent with the decisions of other firms, so that

$$\widetilde{z}_{k'j'}^E\left(N_0^C,\mathbf{N}^C\right) = z\left(k',j',N_0^C+1,\mathbf{N}^C-\mathbf{I}_{k'j'}\right),$$
$$\widetilde{x}^E\left(N_0^C,\mathbf{N}^C\right) = x\left(N_0^C,\mathbf{N}^C\right).$$

Moreover, we can now note that

$$\widetilde{x}\left(k,j,N_0^C,\mathbf{N}^C\right) = x\left(N_0^C-1,\mathbf{N}^C+\mathbf{I}_{kj}\right).$$

That is, the product innovation efforts of a firm's competitor also coincide with what the firm itself would choose if it were to find itself in its competitor's state.

Using (15), the first-order condition for firms without a design implies

$$x(N_0^C,\mathbf{N}^C) = \left(\frac{\sum_{k=1}^{k_{\max}} \tau_k V\left(k,1,N_0^C,\mathbf{N}^C\right) - V^E\left(N_0^C,\mathbf{N}^C\right)}{\chi\psi}\right)^{\frac{1}{\psi-1}}. \tag{16}$$

To compute the industry equilibrium, we proceed as follows. Given a wage guess w, first we compute market shares and markups for all possible states in which an industry can be in (i.e., for all combinations of quality-adjusted productivity and for all possible sets of producing firms).

From this, we can recover the profit functions of all firms as a function of the wage. Then, we can evaluate payoffs at every point of the quality-productivity grid, and for every possible configuration of competitor states, and solve for the value function using value function iteration.

3.2.3 Distribution of Industries and Market Clearing

To close the model, we need to compute market-clearing wages. Denote by $(N_0, \mathbf{N})$ the state of an industry, where N_0 is the number of firms without a design and $\mathbf{N}$ is a matrix where the row-k, column-j element corresponds to the number of firms with design quality ω_k and labour productivity $\underline{q}\lambda^{j-1}$. To simplify notation, it is sufficient to enumerate industry states by $s \in \{1, \ldots, \overline{S}\}$. Further, we denote the invariant equilibrium distribution of industries over these states by $h(s)$. This distribution can be inferred from innovation policies, as described in Appendix A.1.

To compute the wage, we can use the definition of the aggregate price index to get

$$w = \frac{A}{M}, \tag{17}$$

where A is aggregate quality-adjusted productivity and M is a markup wedge, holding

$$A \equiv \left(\sum_{s=1}^{\overline{S}} \left(A(s)\right)^{\eta-1} h(s) \right)^{\frac{1}{\eta-1}}, \quad \text{and} \quad M \equiv \left(\sum_{s=1}^{\overline{S}} \left(M(s)\right)^{1-\eta} \left(\frac{A(s)}{A}\right)^{\eta-1} h(s) \right)^{\frac{1}{1-\eta}}.$$

where $A(s)$ and $M(s)$ were defined in equations (9) and (10), respectively. Next, we can compute aggregate output using labour market clearing. For this, we first compute labour demand for each industry. For each producing firm $n \in \{1, \ldots, \overline{N}\}$ in an industry in state $s \in \{1, \ldots, \overline{S}\}$, we have

$$\ell_n(s) = \left(a_n(s)\right)^{\varepsilon-1} \left(m_n(s)\right)^{-\varepsilon} \left(\frac{A(s)}{M(s)}\right)^{\varepsilon} Y(s). \tag{18}$$

Summing up across all producing firms in the industry, we get that

$$Y(s) = \left(\frac{\mathcal{W}(s)}{M(s)}\right)^{\varepsilon} A(s) L(s), \quad \text{where } \mathcal{W}(s) \equiv \left(\sum_{n=1}^{\overline{N}} \left(m_n(s)\right)^{-\varepsilon} \left(\frac{a_n(s)}{A(s)}\right)^{\varepsilon-1} \mathbb{1}_n(s) \right)^{-\frac{1}{\varepsilon}}. \tag{19}$$

In this equation, $L(s)$ is the total labour employed in an industry in state s, $\mathcal{W}(s)$ is an industry-level measure of relative productivity, and that $\mathbb{1}_n(s)$ is a indicator variable equal to 1 if firm n produces, and to zero otherwise. Aggregate labour market clearing then implies

$$Y = \mathcal{W} A L, \tag{20}$$

where

$$\mathcal{W} \equiv \left(\sum_{s=1}^{\overline{S}} h(s) \left(\frac{A(s)}{A}\right)^{\eta-1} \left(\frac{\mathcal{W}(s)}{M(s)}\right)^{-\varepsilon} \left(\frac{M(s)}{M}\right)^{-\eta} \right)^{-1} \tag{21}$$

is an aggregate efficiency wedge. Finally, aggregate output, pinned down by equation (20), is used for consumption, R&D spending and the fixed cost of production, so feasibility requires that:

$$\frac{C}{Y} + \sum_{s=1}^{\overline{S}} h(s) \Big[\phi N(s) + X(s) + Z(s) \Big] = 1 \tag{22}$$

where $N(s)$, $X(s)$ and $Z(s)$ stand for the number of producing firms, the product R&D spending share and the process R&D spending share, respectively, in industry state s. These variables are defined as follows

$$N(s) \equiv \sum_{n=1}^{\overline{N}} \mathbb{1}_n(s), \qquad X(s) \equiv \zeta \sum_{n=1}^{\overline{N}} \big(x_n(s)\big)^{\psi}, \qquad Z(s) \equiv \chi \sum_{n=1}^{\overline{N}} \big(z_n(s)\big)^{\psi}.$$

Finally, note that in this economy there exist static efficiency losses from input misallocation coming from the fact that innovation induces concentration within industries, allowing firms to make positive profits. In particular, not all income is absorbed by labour, and the labour income share is given by $\frac{wL}{Y} = \big(\mathcal{W}M\big)^{-1}$, a measure of the static efficiency losses in equilibrium.

4 Equilibrium and model features

4.1 Parameterization

To illustrate some of the predictions of our model, from now on we rely on a parameterization shown in Table 1.

Table 1: Illustrative parameterization

Parameter	Value	Meaning
η	2	Elasticity of substitution across industries
ε	7.54	Elasticity of substitution within industries
$\overline{N}$	6	Maximum number of firms per industry
$k_{\max}$	2	Number of designs
$j_{\max}$	4	Maximum number of process innovations
δ	0.02	Industry reset shock frequency
ρ	0.04	Discount factor
ϕ	0.01	Fixed cost of production
χ	26.16	Product innovation cost scale
ζ	0.67	Process innovation cost scale
ψ	2	Innovation cost curvature
λ	1.3	Process innovation step size
ω_L	0.99	Low design quality
ω_H	1	High design quality
τ_L	0.95	Low design probability
$\underline{q}$	1	Starting labour productivity

Notes: This table lists the parameter values used in the remainder of the paper.

Note that these parameter values are indicative, and not chosen to match any data moments. Therefore, the following discussion focuses on the qualitative rather than the quantitative implications of the model.

In the next sections, we first discuss the production and innovation decisions of firms, before turning to industry-level and aggregate outcomes.

4.2 Production and innovation decisions

Production For a given distribution of quality-adjusted productivity at the industry-level, each firm needs to decide whether to produce or not. Then, conditional on producing, firms need to decide on their markup.

Figure 7: Market shares of a firm in state $(1,1)$

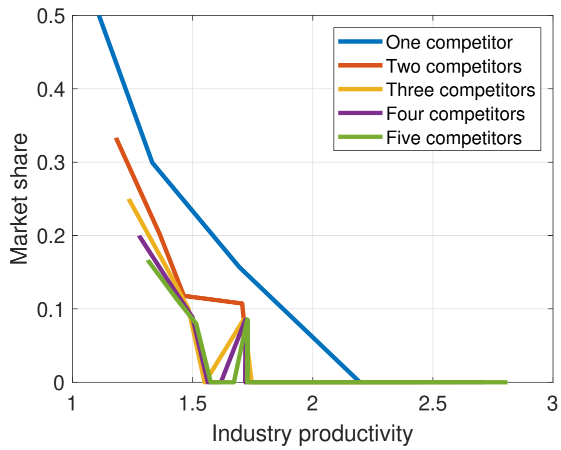

Notes: Market shares of a firm in the lowest productivity, lowest design state, as a function of the aggregate productivity of the industry (horizontal axis) and the number of competitor firms with a design (different lines).

Figure 7 illustrates production decisions by considering the case of a firm in state $(k, j) = (1,1)$ (i.e., with a bad design and the lowest level of labour productivity). The figure plots the market share of this firm as a function of the number of competitors with a design, and the aggregate productivity of the industry (capturing the strength of competitors). This shows that market shares are generally decreasing both in the number of competitors and in their aggregate productivity.[14] When there are too many competitors and/or competitors are too productive, the firm stops producing –its market share becomes zero– as its operating profits no longer cover the fixed cost. This endogenous exit will be the source of industry-level shakeouts.

Innovation In order to gain market share (and thereby increase profits), firms engage in innovation. Figure 8 summarises these innovation decisions. The left panel considers the situation of

[14]However, the non-monotonicity in Figure 7 shows that the number of competitors and their aggregate productivity are not sufficient statistics for the state of the industry. Instead, the distribution matters as well. For instance, the firm at $(1,1)$ produces when facing two other competitors at $(1,1)$ and one at $(1,3)$, but does not produce with three competitors at $(1,2)$, even though the aggregate productivity of its competitors is lower in the latter scenario.

a potential entrant investing in product innovation in order to enter the industry, while the right panel considers the situation of an incumbent in state $(1,1)$, investing in process innovation to increase its productivity.

Figure 8: Product and process innovation decisions

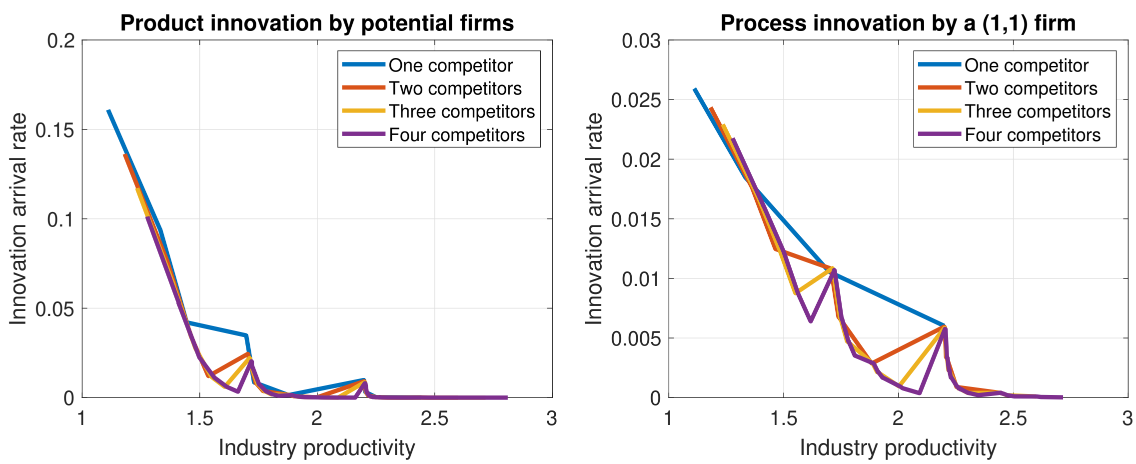

Notes: Product innovation of a potential entrant (left panel) and process innovation of a firm in the lowest productivity, lowest design state (right panel), as a function of the aggregate productivity of the industry (horizontal axis) and the number of competitor firms with a design (different lines).

Both panels show that innovation intensities are generally decreasing in the number of competitors and in their aggregate productivity. This is due to a discouragement effect: if there are already many productive competitors, entry and product innovation are less attractive, because they can at best (after a series of successful innovation) lead to splitting the market with the best competitors. On the other hand, if there are few productive competitors, it is attractive to try to move first, become the most productive producer and enjoy the ensuing high market shares and profits.

However, Figure 8 also shows that this relationship is not fully monotonic: in some parts of the state space, firms might actually innovate more when facing more and/or more productive competitors. This is due to two effects. First, firms are engaged in an innovation race. Thus, if there are already a number of competitors, it becomes more likely that one of them will innovate and pull ahead, increasing in turn the incentives to escape competition by innovating. Second, as the elasticity of substitution between industries exceeds one, productive competitors increase the overall sales of the industry, providing further incentives for innovation.

This discussion illustrates that firms in our model evolve in a rich competitive environment, where industry structure both shapes and is shaped by their production and innovation decisions. In the next sections, we analyse how this firm-level choices translate into industry dynamics over the life cycle.

4.3 The life cycle of the average industry

In our model, each industry has a finite number of firms, and innovation is stochastic. Therefore, we cannot invoke the law of large numbers at the industry-level: depending on the realisation

Figure 9: The life cycle of the average industry

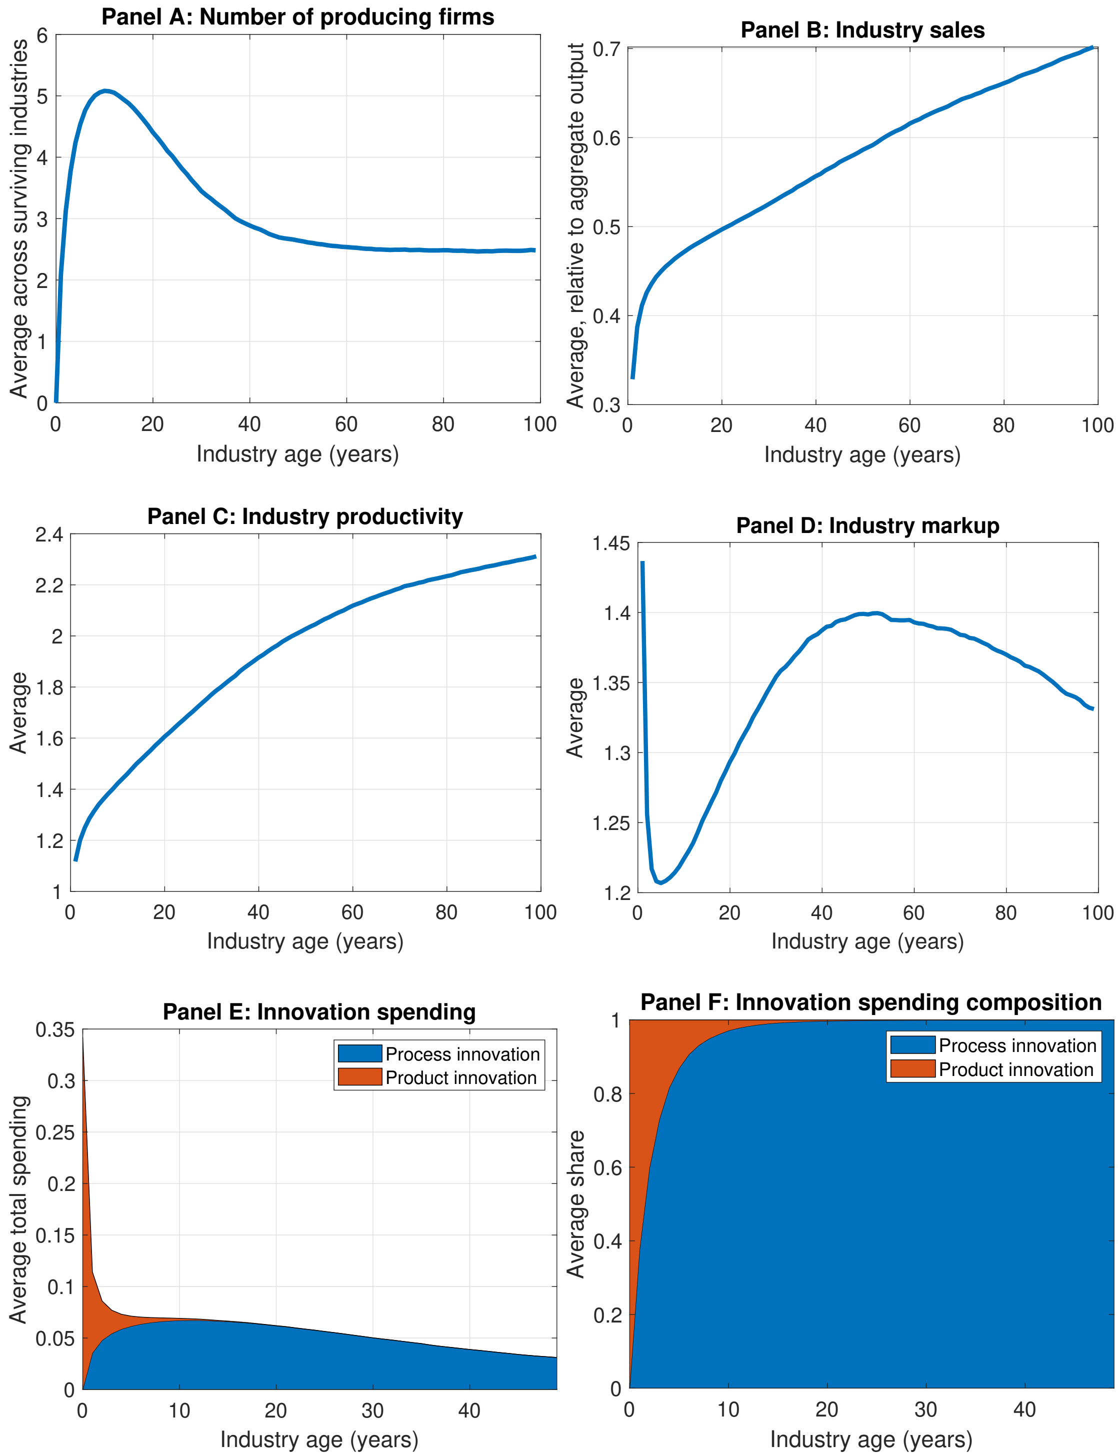

Notes: Selected equilibrium industry aggregates as a function of industry age. Each plot shows statistics for the average surviving industry obtain from the numerical simulation of a large cohort of industries over time.

of the innovation processes, a given industry might evolve in many different ways over time. However, as our model has an infinite number of industries, outcomes for the average industry are not stochastic. Therefore, to understand industry life cycles, we will focus on how average industry-level statistics change with age for a large sample of industries.[15]

Figure 9 plots some key statistics. Panel A plots the average number of producing firms over time, showing that it follows the standard life cycle pattern. Initially, firm numbers increase, as firms find a design and start production. Over time, as more firms enter and incumbent firms improve their productivity (leading to an increase in the industry-level productivity A_i, as shown in Panel C), competition increases. Thus, the firms with the lowest levels of quality-adjusted productivity choose to cease production, as their operating profit no covers the fixed cost. This leads to a shakeout, i.e., a decline in the number of operating firms. As in Klepper (1996), the driving force of this shakeout turns out to be the process innovation of incumbents.

This life cycle is associated with systematic changes in industry structure, showing up in markups and innovation spending. Panel D shows the evolution of average industry-level markups. Average markups start at a high level as the first firm entering the industry is a monopoly. As more firms enter, the industry-level markup falls. However, this fall is short-lived: as the first incumbents start to improve their labour productivity, they pull ahead of their competitors, increasing markups and pushing some competitors out of the market. When the number of firms has stabilised, finally, markups start falling again, as the lagging firms slowly catch to initial innovator, decreasing market share heterogeneity and average markups.

Finally, Panels E and F display the composition of innovation spending. Overall, innovation decreases with industry age, as ideas are becoming harder to find. Moreover, as all potential entrants eventually enter, the composition of innovation shifts from product to process innovation.

5 R&D subsidy policies

In this section, we study implications of industry life cycles for R&D subsidy policies. To do so, we assume assume that the government reimburses a fraction ξ_{product} of all product innovation costs and a fraction ξ_{process} of all process innovation costs to firms. These policies are financed by a lump-sum tax on the representative agent.[16]

Even though we consider flat-rate subsidies, product and process innovation subsidies differently affect industries at different stages of their life cycle. As young industries rely more on product innovation (recall Panels E and F of Figure 9), a subsidy to product innovation will affect these industries more. Therefore, subsidies to product innovation disproportionately subsidise potential entrants, whereas process innovation subsidies mainly affect incumbent firms.

[15]To obtain these statistics, we numerically simulate the trajectory of a large cohort of industries over time. All statistics refer to the sample of surviving industries.

[16]Our specification also allows for R&D taxes, when ξ_{product} or ξ_{process} are negative parameters. In this case, the proceeds of the tax are rebated lump-sum to the representative consumer.

5.1 Product innovation subsidies

Figure 10 displays our results for subsidies (or taxes, if $\xi_{\text{product}} < 0$) to product innovation. The left panel shows the evolution of the number of active firms for the average industry along its life cycle, for different levels of the subsidy. Larger subsidies to product innovation are naturally associated with an increase in product innovation. As industries rely relatively more on this type of innovation spending in their earlier stages of rapid growth, higher subsidy rates translate into a faster entry of new producing firms on average. As a result, as the product innovation subsidy increases, the number of active producers in the average young industry grows more rapidly, and the shakeout starts earlier on.

Figure 10: The impact of product innovation subsidies

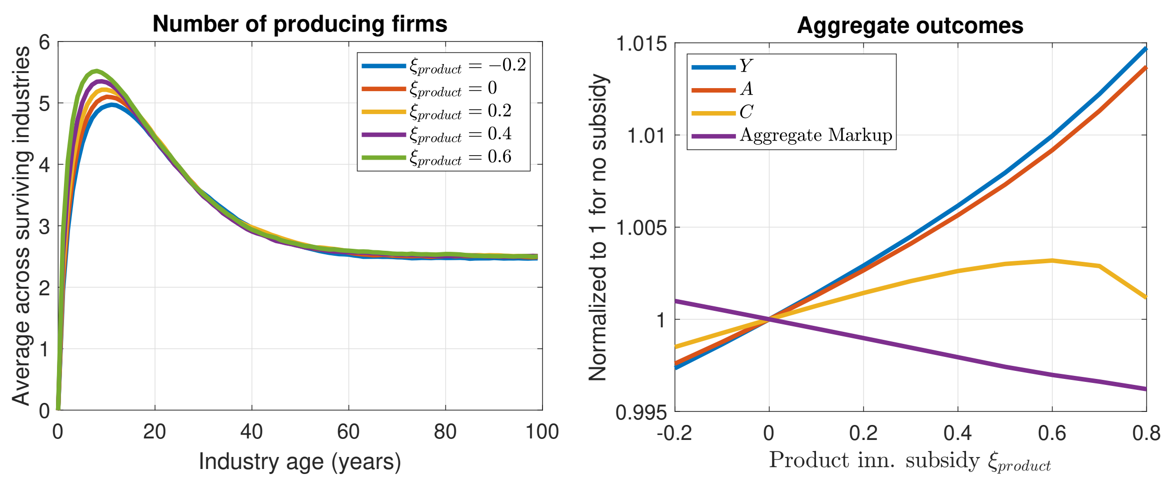

Notes: Number of producing firms as a function of industry age (left), and aggregate outcomes (right), for different levels of the product innovation subsidy.

The right panel of Figure 10 displays how economic aggregates differ for different levels of the product innovation subsidy. When product innovation subsidies are higher, there are on average more active firms per industry at all industry ages. As a result, higher product innovation subsidies reduce the average level of market concentration at the industry level, leading to lower aggregate markups. While the reduction in market concentration reduces the associated wedge on aggregate TFP, most of the increase in output comes from the increase in aggregate quality-adjusted productivity (A). Consumption (and, therefore, welfare) follows an inverted-U relationship with product innovation subsidies. While higher subsidies increase output, their financing also reduces consumption. At low levels of the subsidy, the marginal gain from raising subsidies in terms of increased competition and higher productivity more than offset the consumption costs of the subsidy. However, at high levels of subsidies, this relationship is reversed, implying the existence of an optimal product innovation subsidy rate.

5.2 Process innovation subsidies

Figure 11 shows the results for process innovation subsidies. The results are significantly different from those of product innovation subsidies. Higher subsidies are now associated with fewer firms per industry over most of their life cycle, though the shakeout also starts earlier on when subsidy rates are higher. Process innovation subsidies result in a faster increase in productivity for incumbent firms. As a result, potential entrants now face stronger competition from much more technologically-advanced incumbents, which reduces their incentives to invest in product innovation. At the same time, incumbents successfully improving their productivity reduce the market share of unsuccessful incumbents as well, which prompts some of them to exit, generating the shakeout earlier.

Figure 11: The impact of process innovation subsidies

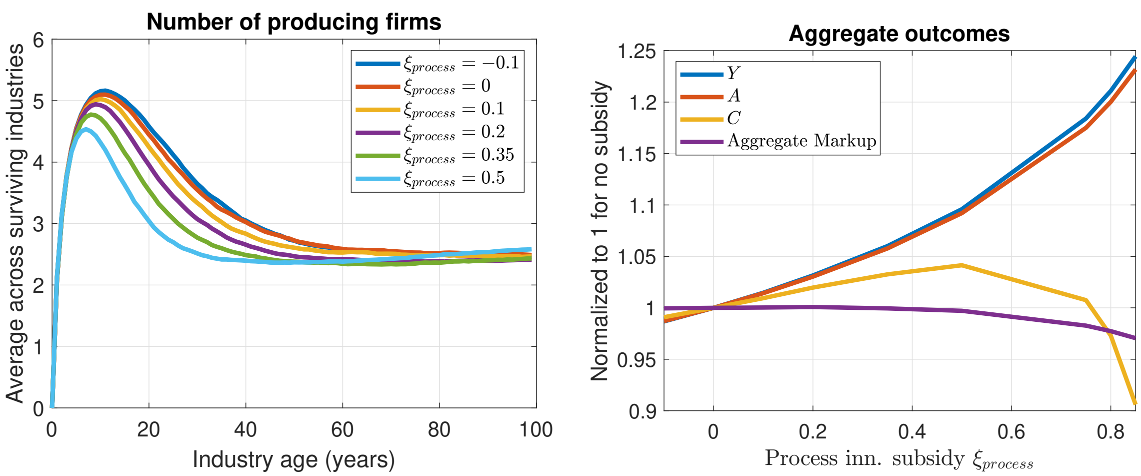

Notes: Number of producing firms as a function of industry age (left), and aggregate outcomes (right), for different levels of the process innovation subsidy.

Interestingly, despite leading to an overall reduction in the average number of firms per industry, process innovation subsidies nevertheless decrease aggregate markups. While a decrease in the average number of firms per industry would reduce competition, subsidies simultaneously decrease productivity dispersion within the industry. This generates less market power for each individual firms, thereby reducing aggregate markups. Output increases with process innovation subsides as they directly increase average firm productivity and decrease the wedge associated with larger aggregate markups. As was the case for product innovation subsidies, the marginal cost of financing the subsidy eventually more than offsets the marginal gains in terms of increased output, leading to an inverted-U relationship between the subsidy and welfare.

6 Conclusion

In this paper, we study the life cycle patterns triggered by technological breakthroughs. Empirically, we show that the industries that were most exposed to the ICT Revolution experienced an increase in net entry rates of patenting firms, followed by a shakeout that reduced the number of active patenting firms. Moreover, we find that new inventions in those industries tended to build upon more recent technologies, and that maturing industries substituted product for process innovation.

To understand the macroeconomic implications of these patterns, we build a general equilibrium model with industry life cycles, oligopolistic competition and endogenous innovation choices. Under our current parameterization, the average industry undergoes an initial phase of rapid positive net entry of firms, in which productivity improves at a fast pace and markups decline due to increased competition. Over time, as more firms enter and some of the existing firms improve their productivity through via process innovation, the firms with the lowest productivity levels can no longer sustain positive profits after paying for fixed costs, and they choose to stop production. This triggers a shakeout, until the industry eventually consolidates and the number of active firms stabilises, with markups falling again as lagging firms slowly catch up to the initial innovators.

We use our parameterized model to study the impact of R&D subsidies. We find that these subsidies generally speed up industry life cycles, and that they can increase consumer welfare.

References

ABERNATHY, W. J. and UTTERBACK, J. M. (1978). Patterns of Industrial Innovation. *Technology Review*, **80**, 41–47.

ACEMOGLU, D., AKCIGIT, U., ALP, H., BLOOM, N. and KERR, W. (2018). Innovation, reallocation, and growth. *American Economic Review*, **108** (11), 3450–91.

AKCIGIT, U. and ATES, S. T. (2023). What Happened to US Business Dynamism? *Journal of Political Economy*, **131** (8), 2059–2124.

ATKESON, A. and BURSTEIN, A. (2008). Pricing-to-Market, Trade Costs, and International Relative Prices. *American Economic Review*, **98** (5), 1998–2031.

AUDRETSCH, D. B. (1987). An Empirical Test of the Industry Life Cycle. *Weltwirtschaftliches Archiv*, **123** (2), 297–308.

BENA, J. and SIMINTZI, E. (2022). Machines Could not Compete with Chinese Labor: Evidence from US Firms' Innovation. *Working Paper*.

BERKES, E., GAETANI, R. and MESTIERI, M. (2023). Technological Waves and Local Growth. *Working Paper*.

BLOOM, N., JONES, C. I., VAN REENEN, J. and WEBB, M. (2020). Are Ideas Getting Harder to Find? *American Economic Review*, **110** (4), 1104–44.

BOS, J. W., ECONOMIDOU, C. and SANDERS, M. W. (2013). Innovation over the Industry Life-Cycle: Evidence from EU Manufacturing. *Journal of Economic Behavior & Organization*, **86**, 78–91.

BRAGUINSKY, S., CHOI, J., DING, Y., JO, K. and KIM, S. (2023). *Mega Firms and Recent Trends in the US Innovation: Empirical Evidence from the US Patent Data*. Tech. rep., National Bureau of Economic Research.

CAVENAILE, L., ROLDAN-BLANCO, P. and SCHMITZ, T. (2023). International Trade and Innovation Dynamics with Endogenous Markups. *The Economic Journal*, **133** (651), 971–1004.

CHRISTENSEN, C. M. (2013). *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Harvard Business Review Press.

DINLERSOZ, E. M. and MACDONALD, G. (2009). The Industry Life-Cycle of the Size Distribution of Firms. *Review of Economic Dynamics*, **12** (4), 648–667.

GORT, M. and KLEPPER, S. (1982). Time Paths in the Diffusion of Product Innovations. *The Economic Journal*, **92** (367), 630–653.

HENDERSON, R. M. and CLARK, K. B. (1990). Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms. *Administrative Science Quarterly*, pp. 9–30.

JONES, C. I. (2022). The Past and Future of Economic Growth: A Semi-Endogenous Perspective. *Annual Review of Economics*, **14** (1), 125–152.

JOVANOVIC, B. and MACDONALD, G. M. (1994). The Life Cycle of a Competitive Industry. *Journal of Political Economy*, **102** (2), 322–347.

— and TSE, C.-Y. (2010). Entry and Exit Echoes. *Review of Economic Dynamics*, **13** (3), 514–536.

KLEPPER, S. (1996). Entry, Exit, Growth, and Innovation over the Product Life Cycle. *American Economic Review*, **86** (3), 562–83.

— (1997). Industry Life Cycles. *Industrial and Corporate Change*, **6** (1), 145–181.

— and SIMONS, K. L. (2000). The Making of an Oligopoly: Firm Survival and Technological Change in the Evolution of the U.S. Tire Industry. *Journal of Political Economy*, **108** (4), 728–760.

LYBBERT, T. J. and ZOLAS, N. J. (2014). Getting Patents and Economic Data to Speak to Each Other: An 'Algorithmic Links with Probabilities' Approach for Joint Analyses of Patenting and Economic Activity. *Research Policy*, **43** (3), 530–542.

MA, Y. and YANG, S. (2023). Technology Driven Market Concentration through Idea Allocation. *Working Paper*.

MCCRAW, T. K. (2010). *Prophet of Innovation: Joseph Schumpeter and Creative Destruction*. Harvard University Press, Cambridge, Massachusetts.

SCHUMPETER, J. A. (1911). *Theorie der wirtschaftlichen Entwicklung*. Verlag von Duncker und Humblot, Leipzig.

— (1934). *The Theory of Economic Development*. Harvard Economics Studies, Cambridge, Massachusetts.

Industry Life Cycles in General Equilibrium

by Laurent Cavenaile, Ruben Gaetani, Pau Roldan-Blanco and Tom Schmitz

Appendix Materials

A Model Appendix

A.1 The invariant distribution of industries

We define the state of an industry as a $(k_{\max} \cdot j_{\max}) + 1$ vector, in which the first element contains the number of firms without a design, and the subsequent elements contain the number of firms with a design for each possible firm state (k, j). The total number of states holds

$$S = \binom{k_{\max} \cdot j_{\max} + \overline{N}}{k_{\max} \cdot j_{\max}}. \tag{A.1}$$

Firm innovation rates define the transition probabilities from each state s to any other state s'.[17] In an equilibrium with an invariant distribution, inflows into each industry state s must equal outflows from that state. These conditions define a linear system of equations that can easily be solved numerically.

A.2 Industry life cycles

To solve our model, we only need to know the invariant distribution of industries over states $\{1, 2, ..., S\}$. However, in order to characterise industry life cycles, we need to know the joint distribution of industries over states $\{1, 2, ..., S\}$ and age.

We obtain this joint distribution by numerically simulating a cohort of 10'000 industries over time, using the equilibrium transition rates implied by the model. Once we have the joint distribution of industries over industry states and age, we can easily compute all necessary statistics as a function of industry age (as industry and indeed firm outcomes are fully pinned down by the industry state s).

[17] Moreover, the reset shock also triggers transitions from any state s to the state in which all $\overline{N}$ firms do not have a design.