Growth with New and Old Technologies

Bernardo SC Ribeiro*
Princeton University

This draft: February 15, 2025
First draft: November, 2023

Abstract

This paper proposes a semi-endogenous growth theory that incorporates technology vintages and the endogenous evolution of multiple technological paradigms through innovation. It provides a characterization of both balanced growth equilibrium and transitional dynamics in an environment where new technologies continuously emerge. From a positive perspective, the model rationalizes two distinct empirical patterns. Using two centuries of US patent data, I first document that the age profile of patents has a pronounced hump shape: most contemporary patents build upon technologies that are between 50 and 100 years old. Second, this age profile has remained stable throughout the past century. From a normative standpoint, the theory underscores a misallocation of research effort induced by the tendency among profit-maximizing firms to overinvest in further developing mature technologies. This yields a suboptimally slow development of emerging technologies. According to a calibrated version of the model, correcting such misallocation could raise long-run consumption by 30%.

Keywords: Growth; Emerging technologies; R&D misallocation

*Previously circulated as "Embracing the Future or Building on the Past? Growth with New and Old Technologies". I am very grateful to Fabrizio Zilibotti, Michael Peters, and Sam Kortum for their generosity with time and insights. I am also very grateful to Mohamad Adhami, Pedro Casavilca, Zhang Chen, Joel Flynn, Bernardo Guimaraes, Hugo Hopenhayn, Zhen Huo, Giammario Impullitti, Danial Lashkari, Claudio Michelacci, Ezra Oberfield, Ryungha Oh, Karthik Sastry, Todd Schoellman, Kjetil Storesletten, Weiliang Tan, Stephan Waizmann, Conor Walsh, Kai-Jie Wu, Emily Yankowitz, and numerous seminar participants. All mistakes are my own.

1 Introduction

From the telegraph to the telephone, from horse-drawn carriages to airplanes, capitalist economies are constantly reshaped by transitions between new and old technologies. This process is far from instantaneous. New technologies are not perfect at the start—they require ongoing innovation to resolve technical challenges and enhance their effectiveness. Meanwhile, older technologies continue to evolve, often for decades, through innovation and R&D efforts. This innovation race forges the transition between technologies, their rise or obsolesce, and the process of growth. The endogenous growth literature (Romer, 1990; Grossman and Helpman, 1991; Aghion and Howitt, 1992) established a model of *aggregate* technological change, but has overlooked the underlying technology transitions and heterogeneity.

This paper proposes an innovation-led growth model in which new technologies continually emerge, and the distribution of innovation efforts across technologies of varying ages arises as an equilibrium outcome. The theory endogenizes, in a growth model, not only well-known patterns of technology transitions (e.g., S-shaped adoption curves) but also a set of novel findings documenting innovation in old and new technologies over two centuries of patent data. Crucially, accounting for technology transitions within an innovation-led growth model has consequences for aggregate productivity and welfare. The theory reveals an inherent R&D misallocation, identifying a tendency of market economies to over-invest in older technologies relative to the ones closer to the frontier. These results provide a rationale for public policy to support investments in cutting-edge technologies, such as quantum computing or metabolic engineering.

The paper begins by documenting how the US economy has balanced innovation between new and old technologies over time. It does so by drawing from two centuries of patent data and using the US Patent Office's technological classification (USPC). These technology classes are not defined by the industries or sectors to which patents apply but rather by the scientific principles they incorporate. It is key that patents cite disproportionally more other patents within the same class than patents in different classes.[1] This means that a technology should be understood

[1] On average, over the two centuries spanning the early 19th to 21st centuries, patents within a given USPC technology class cite nearly seven times more other patents within the same class than

as a field of knowledge, expanded by cumulative inventions that take as an input the field's own knowledge more than any other. This will be the key defining property of a technology in this paper, including in the theoretical model, and is firmly grounded in the scientific and managerial literature.[2]

The analysis of innovation over time and across technologies reveals two new facts. First, the cross-sectional distribution of patents across technologies of different ages has a pronounced hump shape. For example, in 2000, most patents built on technologies that emerged 50-100 years earlier, such as 'Wave transmission lines and networks', which facilitated telephonic and telegraphic communications. Fewer patents built on newer technologies from the last decades of the 20th century, such as 'Software development,' or on technologies from the 19th century or earlier. Second, this hump-shaped distribution is stationary and remained stable throughout the entire 20th century. For instance, in 1900, nascent technologies like 'Wave transmission lines and networks' commanded a patent share comparable to what 'Software Development' would achieve in 2000, a century later.

Next, the paper introduces a growth model to address these facts and examine whether the allocation of innovation efforts across technologies of varying maturities is efficient. The first key building block of the model is the vintage structure: new and superior technologies (or vintages) emerge over time in the frontier. The arrival of a technology is represented by the exogenous emergence of a new set of knowledge, for example, due to a breakthrough innovation that opens a new field. Embodying this knowledge, new products are endogenously created over time, becoming tangible representations of the technology. The greater the knowledge that a technology emerges with, the more productive the products embodying it will be. This knowledge is referred to as the technology's *inherent productivity*, representing its fundamental principles and characteristics.

To illustrate this concept of newly emerging vintages, consider the example of magnetic tape recording technology. Fritz Pfleumer, a German engineer, invented the magnetic tape in 1928. He used the principle of magnetic encoding to store data as patterns of magnetic fields on the tape's surface. This approach was fundamen-

those in the most cited foreign class. Érdi et al. (2013) also provides evidence that the formation of citation clusters can predict the creation of new technological classes, while Liu and Ma (2023) offers additional evidence on the relative concentration of citations within technology classes.

[2]See Kuhn (1962)'s seminal work on scientific paradigms, and Dosi (1982). Redding (2002) reviews empirical evidence supporting a similar definition of technology.

tally different from earlier methods that used physical holes, notches, or mechanical positioning to encode data. Pfleumer's breakthrough was the creation of a set of knowledge—how to encode data using magnetic fields. Incorporating Pfleumer's key ideas, several magnetic tape models promptly emerged in the coming years, such as the Magnetophon, created by the AEG Company in 1935 to store audio, or the BASF company 'LH' series of tapes.

The theory assumes that each newly emerging technology has a higher inherent productivity than its predecessors. For example, magnetic tapes surpassed the efficiency of conventional punch card systems. Later, another revolution came in 1953 when IBM engineers introduced the concept of Random Access Storage with the first hard disk drive (HDD). Unlike the sequential data handling of magnetic tapes, this new paradigm of data storage allowed users to access data in any order, drastically boosting data retrieval efficiency.

The second key building block of the model is the cumulative nature of knowledge within a technology. Each newly created product not only incorporates the knowledge from the technology to which it is linked but also contains its own original ideas. These, in turn, contribute to the expansion of the technology stock of knowledge. This enables future inventors to leverage a larger knowledge base in order to create superior products (standing on the shoulders of giants externality). Over time, both the quantity and the average quality of varieties associated with a technology increase, representing its perfection process and underlying the gradual expansion in the expenditure share it commands.

The third key building block is directed technical change across technology vintages. Profit-maximizing researchers face a trade-off between directing innovation towards new technologies, whose intrinsic potential is higher, or further pushing the development of older technologies that are already very productive and for which the standing on shoulders of giants' effect is stronger. From this trade-off emerges a distribution of R&D effort across technologies of different ages, which is stationary in the economy's Balanced Growth Path.

Two key forces play a pivotal role in shaping this distribution. The first is the (exogenous) rate at which new technologies inherently become more productive. The second is the endogenous rate at which the perfection process of technologies occurs, which depends on the extent of an intertemporal knowledge spillover within each technology vintage. The first force shifts the distribution mode toward young vintages. The higher the intrinsic productivity newer vintages are endowed with, the more profitable opportunities they offer to researchers relative to older ones. On the other hand, within-technology knowledge spillovers yield an opposing effect. A stronger standing on the shoulders of giants' force leads to a faster accumulation of knowledge stock in aging technologies, making them attractive to researchers despite the advent of superior technologies at the frontier. The resulting equilibrium distribution is single-peaked, and the peak's location crucially depends on the balance between these two forces. In particular, if the first is sufficiently higher than the second, the distribution is monotonically declining, with most research concentrated on the youngest vintages. As the relative size of the second force increases, the distribution becomes hump-shaped.

The theory underscores a key inefficiency in research allocation under Laissez-faire conditions. The key externality present in the model is the standing on the shoulders of giants' effect. Private agents do not fully internalize knowledge spillovers their inventions have on future researchers. As a result—as it is well-known in the growth literature—the social value of innovation exceeds the private value, and this gap exists here within all technologies. However, in an economy where individual technologies eventually lose momentum and are gradually replaced by newer vintages, this gap is relatively higher in newer technologies in comparison to older ones. Within the latter, the knowledge spillovers missed by the market benefit only a smaller flow of future research, as the rate of innovation is declining. This asymmetry in knowledge spillovers implies that a social planner would allocate a higher share of total R&D resources to younger technologies in comparison to the market equilibrium allocation.

To evaluate the potential quantitative relevance of this research misallocation, the paper proceeds with a simple calibration exercise. Importantly, this analysis does not intend to estimate the best policy design when it concerns innovation across technologies and over time. For that, beyond knowledge spillovers (the main force in the model), the theory would need to account for several additional channels studied by the innovation literature, such as the impact on labor markets (Autor et al., 2024) and aggregate risk (Jovanovic and Ma, 2022). The goal here is to understand whether the novel misallocation discussed above is just a theoretical curiosity or if it can have sizable productivity impacts. Since this misallocation directly results from knowledge spillovers, whose empirical relevance is found to be large (Bloom et al., 2013), and which play a central role in most endogenous growth models, this question is of particular importance.

The model is calibrated to the above mentioned data on patent flows, as well as information on the age profile of patent valuations. Implementing the optimal research allocation initially causes a temporary growth slowdown, as resources shift toward nascent but initially less productive technologies. However, after 15 years, growth under the optimal plan surpasses that of the laissez-faire scenario and remains higher for several decades, ultimately increasing long-run consumption by 30%. The associated welfare gains amount to approximately 3% in consumption-equivalent terms, a result that holds across various parameter specifications.

Related Literature This paper builds on the tradition of vintage capital growth models, early developed by Johansen (1959), Solow (1960), and Arrow (1962). In these models, newly built capital goods automatically embodied the latest, more productive technology vintage. The productivity of a given technology—and thus of the capital goods embodying it—remained unchanged. Empirical and historical research, however, increasingly documented continued investment and improvement in old technologies, in parallel to the slow diffusion of newer ones (Griliches, 1957; Mansfield, 1961; Harley, 1973; Mokyr, 1990). This led to models featuring learning by doing (Jovanovic and Lach, 1989; Benhabib and Rustichini, 1991; Jovanovic and Nyarko, 1996; Atkeson and Kehoe, 2007) where technologies could continually become more efficient, discouraging immediate adoption of frontier ones. Alternatively, it also led to models where newly built (human) capital could embody older vintages and not necessarily the latest technology (see the seminal work from Chari and Hopenhayn 1991 and more recently Redding 2002, Comin and Hobijn 2010, and Jovanovic and Yatsenko 2012).

In this paper, instead, it is endogenous innovation that evolves technologies over time. In this sense, it bridges the traditional vintage capital literature with idea-based theories of growth building on Romer (1990), Grossman and Helpman (1991), Aghion and Howitt (1992), Jones (1995), and more recently Klette and Kortum (2004), Sampson (2015), Akcigit and Kerr (2018), and Peters (2020). This means that key externalities and drivers of economic growth emphasized by the latter, such as knowledge spillovers and the non-rivalry of ideas, which were also empirically shown to be crucial for growth (Bloom et al., 2013), can now be studied in the context of transitions between new and old technologies, which was not possible in the aforementioned vintage capital tradition. In turn, relative to idea-based models of growth, by introducing the vintage structure, this paper shows that knowledge spillovers are *asymmetric* depending on the life cycle stage of a technology in which an innovation happens.

This paper builds on and contributes to the directed technical change literature, as developed by Acemoglu (2002), Acemoglu and Zilibotti (2001), and more recently, Hopenhayn and Squintani (2021), Acemoglu et al. (2012), Acemoglu et al. (2016), and Donald (2024). In this literature, the set of research lines is typically fixed and, in most cases, limited to two technologies.[3] By contrast, the theory presented here features new technologies continuously emerging at the frontier, with an ever-expanding state space over time. As a result, this model characterizes innovation across endogenous cycles of technological rise and obsolescence, which was not the focus of the existing literature.

This paper relates to theories of General Purpose Technology (GPT) as conceptualized by Helpman and Trajtenberg (1994, 1996). In their framework, a technology serves as a platform for the creation of new varieties or applications over time. Similarly, here, a technology is modeled akin to a GPT in Helpman and Trajtenberg (1994). However, they focus on innovation happening only on the newest GPT, to which all research resources are allocated. Moreover, in Helpman and Trajtenberg (1994, 1996), there are no knowledge spillovers within a technology: current R&D does not impact the future cost of innovation on a technology. In contrast, this paper demonstrates that incorporating these spillovers is essential to replicate key empirical regularities, such as the hump-shaped distribution of innovation efforts across technologies and the S-shaped adoption curves.

In exploring the differentiation between innovation in emerging and mature technologies, this paper draws a parallel with Akcigit et al. (2020) framework, which considers basic and applied research. They introduce this distinction through three key channels. First, basic research enables spillovers that transcend the targeted industry, unlike applied research. Second, basic research may produce knowledge not immediately translatable into consumer products. Third, innovations originating from basic research enhance the efficiency of subsequent applied research. In

[3] An important exception is Hopenhayn and Squintani (2021), who consider a continuum of research lines.

this paper, however, the key distinction between alternative modes of innovation is solely the vintage structure of technology. It shows that there is an intrinsic misallocation of R&D even without assuming that the nature of spillovers differs when innovation occurs in old versus new technologies. As such, it adds a complementary perspective.

2 The Hump-Shape of US Innovation in the 20th Century

The life cycle of a technology, from emergence to obsolescence, is reflected in the pattern of innovation efforts directed toward it. To illustrate this, consider US patents containing keywords related to the steam engine technology in their text. Figure 1 shows the annual number of such patents from the 1830s, when US Patent Office (USPTO) records began, to the 1970s, when matched patents became negligible. During the 19th century, a period marked by the widespread adoption of steam engines (Crafts, 2004), there was a corresponding upward trend in the number of related patents. By the turn of the century, the advent of technologies such as the electric motor and the internal combustion engine led to its gradual obsolescence (Devine, 1983). Figure 1 captures this transition, showing a steady decline in the issuance of new steam-related patents as the 20th century progressed. However, albeit declining, such patent flows remained significant for a prolonged period, illustrating the steam engine's ongoing innovation even during obsolescence.

Are the patterns identified in specific technologies like the steam engine applicable on a broader scale? To address this question, I now turn to analyze nearly nine million US patents, classified into over 400 technological classes by the USPTO. This enables us to study how innovation was balanced between emerging and established technologies over nearly two centuries.

To identify the emergence date for each of these 400 technology classes, I adopt the methodology outlined by Griliches (1957). In this seminal paper, Griliches was primarily focused on technology diffusion—measuring the extent to which firms were adopting new technologies. He noted the characteristic S-shaped curve representing a technology's life cycle: an initial slow adoption phase, followed by a rapid uptake period, culminating in a plateau as the technology reaches peak

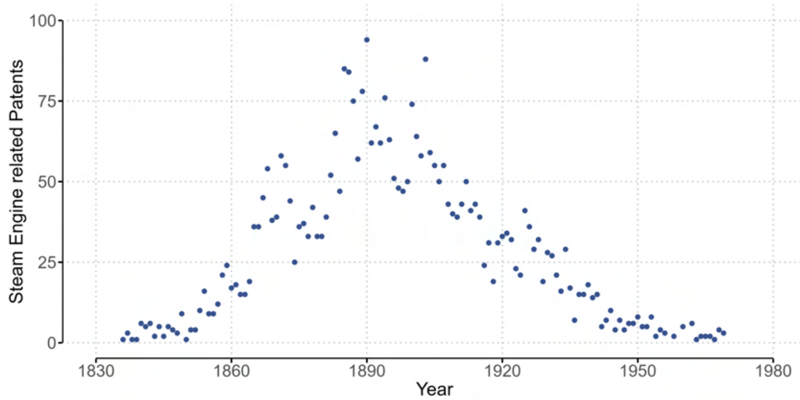

Figure 1: The Innovation Life Cycle—the Steam Engine case

Note: The Figure displays the annual count of patents identified as related to the Steam Engine technology based on the presence of the term 'Steam Engine' (and the absence of terms associated with other motor types) in the patent's text or title. The search criteria account for variations in keyword construction, such as plural forms and case sensitivity. I have narrowed the search to patents classified under the CPC class F04, 'Machines or Engines in general, Steam Engines,' as this classification comprehensively encompasses technologies related to steam engines.

adoption. Since S-shaped patterns align well with logistic functions, Griliches fitted a logistic trend tracing the adoption path of each technology. He defined the emergence date as the point at which this trend reached ten percent of its maximum value.

Here, I exploit the parallel between the life cycles of technology adoption and innovation, and employ Griliches' method to the patenting trajectory of each technology. This parallel is not only motivated by the case of the steam engine but is also corroborated by multiple case studies highlighting a consistent S-shaped pattern in patent timelines.[4] While a comprehensive exposition of this measurement procedure awaits in Section 4, here I present its main findings, which will shed light on the theory presented next in Section 3.

The left plot of Figure 2 displays the distribution of patents granted in the 2000 decade (2000-2009), categorized by the age of the associated technology classes. The horizontal axis denotes the age of these technologies, estimated using the Griliches (1957) methodology, while the vertical axis represents their corresponding share of the total patents issued in that decade. Each dot on the graph aggregates technology classes of the same estimated age, serving as a representation of a

[4]See Achilladelis et al. (1990), Achilladelis (1993), Andersen (1999) and Haupt et al. (2007).

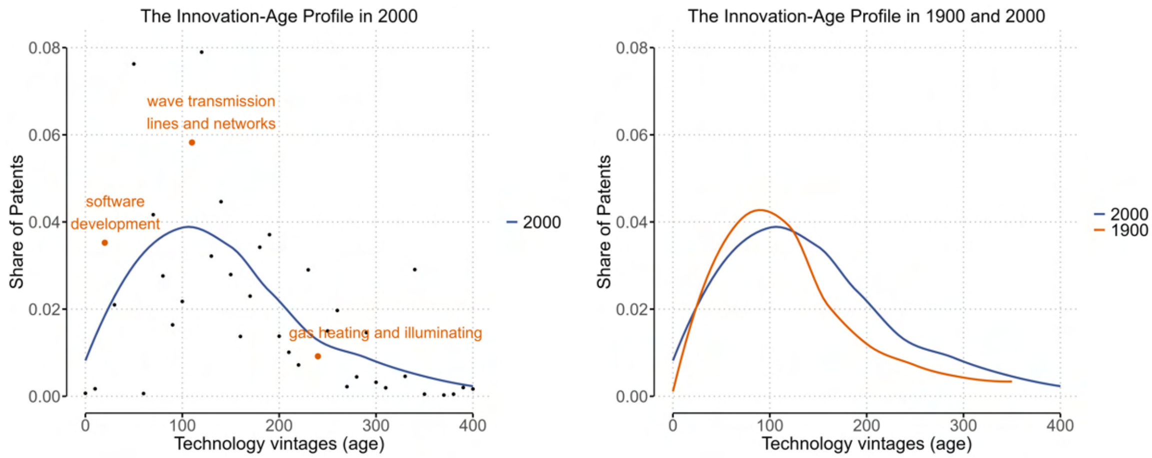

Figure 2: The Hump-Shaped Innovation Age Profile

Note: The left plot shows the distribution of patents issued between 2000 and 2009 across technology classes of different ages. The US Patent Office classifies each patent into one and only one principal technology class (USPC). Their age is estimated in accordance with Griliches (1957) methodology (see Section 4 for details). Each dot on the graph groups technologies with the same estimated age. The solid line represents a non-linear fit generated through local regression models from Cleveland et al. (1992). This line is also shown on the right panel (blue), together with the corresponding line for the 1900-1909 period (red).

particular technology vintage. For illustration, the Figure highlights three specific technology classes of different ages. This is the case of 'Software Development,' which emerged in the 1980s and falls into the category marked by the orange dot below its name, representing the 20-year-old technology vintage, along with other innovations that emerged in the same period. In contrast, 'Wave Transmission Lines and Networks' and 'Gas Heating and Illumination' represent older technologies, also highlighted in orange in the plot.

Importantly, the left plot in Figure 2 shows the first key empirical result of this paper. There is a pronounced hump shape in the technology-age distribution of patenting at a point in time. In the 2000 decade, the majority of issued US patents built on technologies that emerged between 50 and 100 years earlier, such as 'Wave transmission lines and networks', which facilitated telephonic and telegraphic communications, and 'Drug, bio-affecting and body treating compositions'. There is also a significant volume of patents pertaining to older technologies that can be traced back much further in time. Even though older technology classes tend to accrue fewer patents as they age (evident from the declining right tail of the plot), this decrease is gradual. The overall volume of patents within this older technological space is comparable to that of patents associated with more recent technological advancements.

The right plot of Figure 2 contains a second key result: the hump shape of patenting is remarkably stable throughout the 20th century. Like the left plot, this graph illustrates the technology-age distribution of patents in the 2000s, but also includes data from the 1900-1909 decade for comparison. Not only is the hump shape evident in both periods, but strikingly the shape of the distribution is very similar: in 1900, as in 2000, most patents were issued in technological classes that were about 50-100 years old. For instance, in 1900, many technologies that emerged approximately one century before were tied to electricity, as 'Electrical generator or motor structure'. Technologies that would become the most important in 2000 were emerging around 1900 and already had some positive share in patenting by then but were still a minority in terms of innovation shares.

Further details and alternative measures of technology and age Section 4 discusses in detail the proxy for technology age based on Griliches (1957). In turn, Section A in the Appendix shows how the results above can be seen using alternative technology and age measures. This is the case, for example, of Section A.1, which leverages the citation network (not used at all above) and breakthrough patents identified by Kelly et al. (2021). It considers each of these breakthroughs (e.g., Samuel Morse's telegraphy wire patent and Google's PageRank algorithm patent) to represent the emergence of a new technology field. Finally, to sort all US patents into these different technologies, the analysis identifies, within the citation tree, the closest breakthrough each patent is building on. Table A.1.1 shows how, between the years 2000-2009, the share of newly issued patents building on breakthroughs from the past 10 years was equal to the share building on breakthroughs from 70 or more years in the past, with most patents building on mid-age technologies.

Section 3 now presents a theory where the distribution of innovation efforts across technologies of different maturities emerges as an equilibrium object. This allows us to study its shape, time evolution, and efficiency properties.

3 Theory

Consider an infinite-horizon economy in continuous time. The economy is populated by a representative household composed of production workers, with measure L, and researchers, with measure R. The household supplies production and research labor inelastically, has logarithmic instantaneous utility function, and discounts the future at rate ρ.

An evolving set Ω_t of intermediate varieties, together with production labor, constitute the inputs for the final good Y_t:

$$Y_t = \frac{L_t^\beta}{1-\beta}\int_{\Omega_t} z(\omega)x_t(\omega)^{1-\beta}d\omega. \tag{1}$$

Here, $x_t(\omega)$ represents the quantity, used in period t, of an intermediate variety $\omega \in \Omega_t$, while $z(\omega)$ denotes its time-invariant productivity. Each variety ω is supplied by the monopolist firm owning its blueprint. More efficient varieties are costlier to produce. Specifically, it takes $\psi z(\omega)$ units of the final good to produce each unit of ω. As we will see, this assumption guarantees that the firm profit is linear in $z(\omega)$ but has no additional implications.

Crucially, the efficiency $z(\omega)$ of a particular variety ω depends on its own quality $q \in \mathcal{Q}$ and on the technology $\tau \in \mathcal{T}$ to which it is linked. From now on, $\omega(\tau, q)$ will denote a variety ω linked to technology τ with quality q. Its productivity can be expressed as

$$z\big(\omega(\tau, q)\big) = A_\tau \times q,$$

where A_τ represents technology τ's intrinsic potential—incorporated by all varieties linked to it.

We can think of a variety linked to a technology τ as a product embedding the intrinsic properties and characteristics of τ. If τ refers, for instance, to the computer technology, examples of varieties linked to it would be the Dell XPS Laptop or the Apple's MacBook Air. They have their own quality and individual characteristics, but both embed the main principles and properties of the computer technology.

At each point in time t, there is a total measure $\mu(\Omega_t)$ of available variety blueprints. Since each variety is linked to one, and only one, technology τ, and has quality q, this measure is defined on the space $\mathcal{T} \times \mathcal{Q}$. Its density, denoted by

$f_t(\tau, q)$, satisfies:

$$\mu(\Omega_t) = \int_{\mathcal{T}} \int_{\mathcal{Q}} f(\tau, q) dq d\tau.$$

Consider one specific technology τ and let $\mu(t|\tau)$ denote the measure of all variety blueprints linked to it. Formally:

$$\mu(t|\tau) = \int_{\mathcal{Q}} f_t(\tau, q) dq. \tag{2}$$

Moreover, define by $Q(t|\tau)$ the average quality, at time t, of all blueprints linked to a technology τ:

$$Q(t|\tau) = \frac{\int_{\mathcal{Q}} q f_t(\tau, q) dq}{\mu(t|\tau)}. \tag{3}$$

Summarizing the extensive and intensive margins of each technology development, $\mu(t|\tau)$ and $Q(t|\tau)$ serve as the key state variables in this theory. As shown in Section 3.2, they are sufficient to compute the equilibrium, eliminating the need to track the entire distribution $f_t(\tau, q)$.

Technology vintages Each period, a new technology emerges exogenously. At calendar time t, technologies $\tau \in (-\infty, t]$ are known. Thus, the technology set $\mathcal{T}$ is the real line $\mathbb{R}$, and the density f_t's support satisfies $\text{supp}(f_t) \subseteq (-\infty, t] \times \mathcal{Q}$. Hereafter, the terms 'technologies' and 'vintages' will be used interchangeably, as τ also indexes the vintage of a technology by indicating the year of its emergence.

When technology τ emerges, the economy is exogenously endowed with the triple $\{A_\tau, \mu(\tau|\tau), Q(\tau|\tau)\}$, where $\mu(\tau|\tau) = \mu_0$ represents a small measure of breakthrough varieties—the first to embed τ—whose average quality is $Q(\tau|\tau) = Q_0$. Most importantly, the emergence of a new technology introduces fundamental new knowledge, with intrinsic potential measured by A_τ, enabling future research to endogenously create varieties incorporating it.

The intrinsic potential of the technology that just emerged, A_t, stands for the knowledge frontier. As such, I assume it increases over time at rate $\gamma > 0$. Formally,

$$A_\tau = e^{\gamma \tau}, \tag{4}$$

for every technology τ. A high γ implies that each successive generation of technology has significantly higher intrinsic potential than its predecessors.

Researchers conduct innovation to create new blueprints. Specifically, if a variety is invented at time t linked to a technology $\tau \in (-\infty, t]$, its quality level q is determined (i) by the quality of existing varieties within such technology, and (ii) by an original component λ drawn from a distribution $H(\cdot)$ at the time of the invention:

$$q = Q(t|\tau)\lambda, \quad \text{where } \lambda \sim H(\cdot) \text{ and } \bar{\lambda} \equiv \int \lambda dH(\lambda) \geq 1. \tag{5}$$

This implies that innovation at time t builds on the knowledge stock already accumulated within the targeted technology, $Q(t|\tau)$, characterizing a standing on the shoulders of giants externality. If $\bar{\lambda} > 1$, the average quality of varieties related to technology τ will increase over time with the arrival of new and, on average, better products. This represents the *perfection* process of a technology: τ may have a high intrinsic potential A_τ, but its first varieties may have been of poor quality. Only with R&D effort over time can it be perfected and streamlined. The larger $\bar{\lambda}$, the larger the advances made by each generation of varieties in a given technology. As such, $\bar{\lambda}$ represents the incentives to "build on the past."

The innovation possibilities frontier faced by researchers is discussed in detail in Section 3.2, alongside the model dynamics. Before that, Section 3.1 presents the static allocations, taking the set of varieties as given.

3.1 Equilibrium Production Allocation Given Technology

At a given time, optimal decisions by firms and households determine the allocation of consumption and production given the current state of technology, $[Q(t|\tau), \mu(t|\tau)]_{\tau \leq t}$. The final good Y_t serves as the numeraire in this economy. Its production occurs under perfect competition and places the following isoelastic demand for intermediate varieties:

$$x_t\big(\omega(\tau,q)\big) = L_t \left[\frac{p_t\big(\omega(\tau,q)\big)}{z\big(\omega(\tau,q)\big)} \right]^{-\frac{1}{\beta}} \quad \text{for all } \omega(\tau,q) \in \Omega_t.$$

Here, L_t is the optimal demand for labor, and $p_t\big(\omega(\tau,q)\big)$ is the price charged by the monopolist firm producing $\omega(\tau,q)$. It solves

$$\max_{p_t\left(\omega(\tau,q)\right)} x_t\big(\omega(\tau,q)\big)\Big[p_t\big(\omega(\tau,q)\big)-\psi z\big(\omega(\tau,q)\big)\Big],$$

taking as given the demand for $x_t\big(\omega(\tau,q)\big)$. To save notation, we assume hereafter $\psi=(1-\beta)$.

Since the household has a fixed measure L workers who supply labor inelastically, the market clearing condition $L_t = L$ pins down wages at every period. In turn, the solution to the profit-maximizing problem involves the firm charging $p_t\big(\omega(\tau,q)\big) = z\big(\omega(\tau,q)\big)$, assembling $x\big(\omega(\tau,q)\big) = L$ units, and realizing profits $\pi\big(\omega(\tau,q)\big) = \beta z\big(\omega(\tau,q)\big)L$.

Proposition 1 (Aggregate output)**.** *In equilibrium, aggregate output is given by:*

$$\begin{aligned} Y_t &= \frac{L}{1-\beta}\int_{-\infty}^{t} e^{\gamma\tau}Q(t|\tau)\mu(t|\tau)d\tau \\ &= \frac{A_t L}{1-\beta}\int_{0}^{\infty} e^{-\gamma a}Q_t(a)\mu_t(a)da, \end{aligned} \tag{6}$$

where $A_t = e^{\gamma t}$ is the knowledge frontier; $a = t-\tau$ denotes the age of technology τ at calendar time t; $Q_t(a) \equiv Q(t|t-a)$ and $\mu_t(a) \equiv \mu(t|t-a)$.

*Proof.*The result follows from combining (1)-(3) and the firm optimality conditions. □

Proposition 1 decomposes the contribution of different technologies to total output. The age-a technology (or $\tau = t-a$) equilibrium share in total GDP depends on three factors: how perfected it is, as represented by $Q(a)$; how diffused it is, as represented by $\mu_t(a)$; and how obsolete it is, as represented by $e^{-\gamma a}$.

Finally, let C_t denote the representative household's consumption and X_t the total quantity of final goods used in the production of intermediate varieties. Feasibility requires that $C_t = Y_t - X_t$. In equilibrium, substituting for X_t using the monopolist firm's decisions, it follows that $C_t = \beta(2-\beta)Y_t$.

3.2 Innovation and Equilibrium Dynamics

Researchers choose which technology to target in their projects at each period t. Successful research targeting technology τ generates a variety blueprint associated with τ, which is sold to firms through competitive auctions. The value of each blueprint, denoted by v_t, corresponds to the present discounted value of the profit stream it generates:

$$v_t\big(\omega(\tau,q)\big) = \int_t^\infty e^{\left(-\int_t^s r(v)\,dv\right)}\ \pi\big(\omega(\tau,q)\big)\,ds. \tag{7}$$

Ex-ante, however, a researcher targeting τ does not know the quality q of the invention. Let $\bar{v}(t|\tau) = \mathbb{E}_q\big\{v_t\big(\omega(\tau,q)\big)\big\}$ denote the ex-ante expected value of a successful innovation within technology τ. It is the crucial object to determine research incentives towards different technologies at t. Also, let $R(t|\tau)$ denote the density of researchers at time t targeting technology τ. The direction of innovation depends crucially on the relative supply of research effort, $R(t|\tau)/R(t|\tau')$, for every pair of technologies τ and τ'.

This paper proceeds with the equilibrium condition

$$\frac{R(t|\tau)}{R(t|\tau')} = \left(\frac{\bar{v}(t|\tau)}{\bar{v}(t|\tau')}\right)^{\frac{1}{\epsilon}} \quad \text{and } \epsilon > 0, \tag{8}$$

which, as discussed below, endogenously emerges under several alternative (and common) assumptions on the innovation possibilities frontier. Equation (8) simply states that the relative research supply, $R(t|\tau)/R(t|\tau')$, depends on the relative profitability of technologies, $\bar{v}(t|\tau)/\bar{v}(t|\tau')$. When a technology becomes more profitable relative to others, it attracts a proportionally greater number of researchers—with the elasticity of such response given by $1/\epsilon > 0$. As $1/\epsilon$ approaches zero, the distribution of research across technologies tends toward uniformity, with $R(t|\tau)/R(t|\tau')$ converging to 1. In contrast, when $1/\epsilon$ becomes large, relative research efforts become highly sensitive to differences in profitability, with the ratio $R(t|\tau)/R(t|\tau')$ diverging in the direction of the more valuable vintage.

The crucial restriction imposed in (8) is to rule out cases where all research is concentrated on a single technology or completely absent from another. In this model, therefore, the adjustment of R&D happens at the intensive (but not at the extensive) margin. If a technology becomes obsolete and less profitable compared to others, the share of researchers working on it shrinks and asymptotically converges to zero, but it remains at least infinitesimally positive at any given point in time t.

By satisfying this key assumption, several microfoundations are consistent with the relative supply Equation (8). These include models where researchers have heterogeneous abilities across different technologies or models featuring convex innovation costs (Klette and Kortum, 2004; Akcigit and Kerr, 2018). For concreteness, the paper adopts hereafter a specific and even simpler microfoundation. It assumes the existence of congestion forces in research as in Jones (1995) or Acemoglu et al. (2023). Specifically, the rate at which a scientist working on technology τ discoveries a successful is assumed to be $R(t|\tau)^{-\epsilon}$, where $\epsilon \geq 0$ represents the strength of congestion forces. These forces arise from the overlap or duplication of ideas—which reduces the likelihood that any researcher generates the next innovation herself. The aggregate flow of ideas within technology τ is, thus, equal to $R(t|\tau)^{1-\epsilon}$, and decreases with ϵ.

In equilibrium, the allocation of researchers must meet a non-arbitrage condition, ensuring that no individual researcher would benefit from shifting their focus to a different technology. This is guaranteed by Equation (8), which implies expected returns to research are equalized across fields.

Research allocation solution Despite being intuitive, the relative supply Equation (8) is not yet close to solving the model. The research distribution $R(t|\tau)$ has an unbounded support of technologies, each of them carrying its own forward-looking value functions $\bar{v}(t|\tau)$. Notably, however, after conditioning on technologies' current development state, $[Q(t|\tau)]_\tau$, and inherent productivity $[A_\tau]_\tau$, the model pins down the research distribution, at any time t, without forward-looking variables. To see that, observe how, by using the equilibrium value for profits and Equation (7), $\bar{v}(t|\tau)$ can be written as

$$\bar{v}(t|\tau) = \beta e^{\gamma\tau} Q(t|\tau)\bar{\lambda} L D_t, \quad \text{where } D_t \equiv \int_t^\infty \exp\left(-\int_t^s r(v)dv\right) ds. \tag{9}$$

The forward-looking component, after factoring out A_τ and $Q(t|\tau)$, is uniform across technologies and equal to D_t. Combining Equations (9) and (8) yields

$$R(t|\tau) = \left(\frac{e^{-\gamma(t-\tau)} Q(t|\tau)}{\bar{Q}_t} \right)^{\frac{1}{\epsilon}} R, \quad \text{or} \quad R_t(a) = \left(\frac{e^{-\gamma a} Q_t(a)}{\bar{Q}_t} \right)^{\frac{1}{\epsilon}} R, \tag{10}$$

where $\bar{Q}_t \equiv \left[\int_0^\infty \left(e^{-\gamma \tilde{a}} Q_t(\tilde{a}) \right)^{\frac{1}{\epsilon}} d\tilde{a} \right]^\epsilon$ is an age-discounted quality index for the economy.

Equation (10) highlights the interplay of two key factors influencing research allocation across technology vintages: obsolescence and accumulated knowledge (quality) stock. Older technologies, by being further from the frontier, exhibit lower inherent productivity and potential compared to their younger counterparts. This naturally biases research investment toward younger vintages, a tendency amplified by γ, which governs the frontier expansion rate. At the same time, older technologies have benefited from longer periods of development and refinement. Due to knowledge spillovers, this accumulated quality stock provides a strong incentive for researchers to innovate on them.

Laws of motion Given the innovation possibilities frontier and a research allocation $[R(t|\tau)]_\tau$, it is possible to write the laws of motion for $\mu(t|\tau)$ and $Q(t|\tau)$:

$$\frac{\dot{\mu}(t|\tau)}{\mu(t|\tau)} = \frac{R(t|\tau)^{1-\epsilon}}{\mu(t|\tau)}, \quad \frac{\dot{Q}(t|\tau)}{Q(t|\tau)} = (\bar{\lambda} - 1) \frac{R(t|\tau)^{1-\epsilon}}{\mu(t|\tau)}. \tag{11}$$

These equations hold for every τ, every $t \geq \tau$, and take $\mu(\tau|\tau) = \mu_0$ and $Q(\tau|\tau) = Q_0$ as given (see Section C.2 for the derivation). They reveal how the growth rates in $Q(t|\tau)$ and $\mu(t|\tau)$ are proportional to the innovation rate per variety within technology τ, $R(t|\tau)^{1-\epsilon}/\mu(t|\tau)$.

Crucially, this implies that the productivity of research decreases as $\mu(t|\tau)$ increases. To sustain a constant growth rate for both $Q(t|\tau)$ and $\mu(t|\tau)$, the number of researchers targeting technology τ, $R(t|\tau)$, must increase over time. This is not caused by the congestion force ϵ but actually stems from the assumption that innovations build not on the very best idea (i.e., on the maximum quality q among technology τ varieties), but on a moment of the distribution (we used the average,

$Q(t|\tau)$, although other percentiles would have the same effect).[5] When discussing the equilibrium, we will see how this property shapes the research allocation and what the consequences are if it is modified.

Finally, expressing the laws of motion for Q and μ as functions of age is beneficial because these age-specific schedules will remain constant along a Balanced Growth equilibrium, to be defined below. From (11), it can be shown that (see Section C.2):

$$\frac{\partial \mu_t(a)}{\partial t} + \frac{\partial \mu_t(a)}{\partial a} = R_t(a)^{1-\epsilon} \text{ and } \frac{\partial Q_t(a)}{\partial t} + \frac{\partial Q_t(a)}{\partial a} = Q_t(a)(\bar{\lambda} - 1)\frac{R_t(a)^{1-\epsilon}}{\mu_t(a)}. \quad (12)$$

To close the model, notice that the representative household dynamic problem implies the familiar Euler equation $\dot{C}_t = C_t(r_t - \rho)$ and transversality condition.

3.2.1. Equilibrium Definition

A dynamic equilibrium can now be characterized. In this economy, equilibrium objects are cross-sectional distributions that evolve over time. In particular, the state variables in a given period t are the distributions of quality and varieties $[Q_t(a), \mu_t(a)]_{a=0}^{\infty}$. They evolve over time as a result of the profit-maximizing distribution of research efforts $[R_t(a)]_{a=0}^{\infty}$.

Definition 1. (Equilibrium) A trajectory for the cross-sectional distributions of quality, varieties, and research $[\mathbf{\Theta}_t]_{t\geq t_0}$ (where $\mathbf{\Theta}_t \equiv [Q_t(a), \mu_t(a), \bar{v}_t(a), R_t(a)]_{a=0}^{\infty}$), for aggregate variables $[Y_t, C_t]_{t\geq t_0}$, and for prices $[r_t, w_t]_{t\geq t_0}$ is an **equilibrium trajectory** if, given an initial condition $[Q_{t0}(a), \mu_{t0}(a)]_{a=0}^{\infty}$, the following conditions are met at every period t: firms maximize profits; researchers maximize expected returns, the household Euler equation and transversality condition are satisfied, all markets clear; and the laws of motion in (12) are respected.

An equilibrium trajectory can be easily computed numerically, given its recursive nature. To see that, notice that the aggregate supply of innovation resources is constant and the relative supply of research across vintages at time t depends

[5] For the sake of intuition, consider the following example. Suppose that a scientist can produce one new variety per unit of time. She will impact the average quality significantly more when the pool of existing varieties is small (for example, when the technology has recently emerged) in comparison to when the set of available varieties is large, and the elasticity of the average quality to the marginal invention is small.

only on $[Q_t(a)]_{a\geq 0}$, as shown in (10). Therefore, after combining (10) and the laws of motion in (12), one is left with differential equations for $Q_t(a)$ and $\mu_t(a)$ that depend only on the current level of these distributions. Given an initial condition $[Q_{t_0}(a), \mu_{t_0}(a)]_{a=0}^{\infty}$, they can be solved forward in time, yielding a trajectory for the state variables. These, in turn, directly pin down all remaining equilibrium variables.

3.2.2. Stationary (or Balanced Growth) Equilibrium Characterization

A stationary (or balanced growth) equilibrium is characterized by cross-sectional technology-age distributions that are time-invariant. Formally:

Definition 2. (Stationary Equilibrium) A stationary equilibrium is an equilibrium trajectory that satisfies, for every $t \geq t_0$, $\partial_t Q_t(a) = \partial_t \mu_t(a) = 0$.

Proposition 2. *There exists a unique stationary (or balanced growth) equilibrium. It is such that:*

(i)If $\bar{\lambda} > 1$, the stationary technology-age distribution of quality at any given $t \geq t_0$ is:

$$Q(a) = \left\{ c_1 - c_2 e^{-\frac{1-\epsilon}{\epsilon}\gamma a} \right\}^{\frac{1}{\frac{1}{\bar{\lambda}-1} - \frac{1-\epsilon}{\epsilon}}} \tag{13}$$

where c_1, c_2 are uniquely determined constants given by:

$$c_2 \equiv \frac{Q_0^{\frac{1}{\bar{\lambda}-1}}(\bar{\lambda}-1)}{\mu_0 \frac{1-\epsilon}{\epsilon}\gamma \bar{Q}^{\frac{1-\epsilon}{\epsilon}}} \left(\frac{1}{\bar{\lambda}-1} - \frac{1-\epsilon}{\epsilon} \right), \quad c_1 \equiv Q_0^{\frac{1}{\bar{\lambda}-1} - \frac{1-\epsilon}{\epsilon}} + c_2,$$

and $\bar{Q}$ is the unique stationary level of $\bar{Q}_t$. If $\bar{\lambda} = 1$, then $Q(a) = Q_0$. Hence, for every $\bar{\lambda} \geq 1$,

$$\lim_{a\to\infty} Q(a) < \infty, \quad \textit{and} \ \lim_{a\to\infty} R(a) = 0.$$

(ii)If $\bar{\lambda} > 1$, the stationary technology-age distribution of varieties at any given $t \geq t_0$ is:

$$\frac{\mu(a)}{\mu_0} = \left(\frac{Q(a)}{Q_0} \right)^{\frac{1}{\bar{\lambda}-1}}. \tag{14}$$

Conversely, if $\bar{\lambda} = 1$, *then* $\mu(a) = c_3 - c_4 e^{-\frac{1-\epsilon}{\epsilon}\gamma a} = \lim_{\bar{\lambda}\to 1}\mu(a)$, *where:*

$$c_4 \equiv \left(\frac{\gamma}{\epsilon}\right)^{1-\epsilon}\frac{1}{\gamma^{\frac{1-\epsilon}{\epsilon}}}, \text{ and } c_3 \equiv \mu_0 + c_4.$$

(iii)Output Y_t *and consumption* C_t *grow at rate* γ.

*Proof.*See the Appendix. □

The first part of Proposition 2 characterizes the stationary quality-age profile $\{Q(a)\}_a$. The average quality $Q(a)$ increases as a technology ages but eventually converges to a finite upper bound. This occurs for two key reasons. First, as established above, even if the number of researchers working on a technology remains constant, its quality growth rate declines over time due to diminishing returns.[6] Second, while the growth in $Q(a)$ slows down, frontier technologies continue to improve at a rate γ. Hence, the profitability of an aging technology steadily declines relative to the newest ones. In response, research efforts gradually shift away from the old vintage, reinforcing the slowdown in $Q(a)$. This can be seen from manipulating (10) to write:

$$\frac{R'(a)}{R(a)} = \frac{Q'(a)}{Q(a)} - \gamma. \tag{15}$$

In the long run, this feedback mechanism drives the share of researchers working on older technologies to zero.

Crucially, as seen in Equation (15), the obsolescence of old technologies hinges on the assumption that their long-run growth rate remains bounded below the frontier growth rate γ. While diminishing returns, which drive a technology's growth rate to zero over time, provide a sufficient condition for obsolescence, they are *not* necessary. Section B.1 in the Appendix extends the model by allowing a constant number of researchers to sustain a steady growth rate in $Q(a)$ indefinitely. This represents the potential long-run growth rate a technology can achieve. However, if this potential growth rate does not sufficiently exceed γ, the economy still converges to an equilibrium in which old technologies are eventually abandoned. Conversely, if a technology's potential growth rate is sufficiently higher than γ, an equilibrium

[6]See Equation (11) and its related discussion.

with multiple coexisting technologies cannot emerge. The profitability of old technologies would never become dominated by younger vintages. The latter would grow boundlessly, always receiving more innovation than young technologies.

The final part of Proposition 2 contains an additional important result. In the BGP, the economy grows at rate γ, the rate at which new technologies are inherently superior to older ones. Since they arrive exogenously, this is a model of semi-endogenous growth (Jones, 1995). Intuitively, because researchers are perpetually shifting from old to new technologies, growth has to come from how better the latter are relative to the former. Old technologies do not steadily grow in the very run.

So far, we have examined why the share of research devoted to a technology declines to zero as it ages. But does it initially rise when the technology is young? From (15), we know that this will happen if $Q(a)$ grows faster than γ when the technology is young. Under what conditions does this happen in equilibrium?

Corollary 2.1 (Innovation in the frontier)**.** For any vector of parameters, there exists $\underline{b} > 0$ (where $\underline{b}$ is a function of the parameter vector) such that, if $\left(\bar{\lambda} - 1\right)/\gamma < \bar{b}$, then $R'(0) < 0$, whereas, if $\left(\bar{\lambda} - 1\right)/\gamma \geq \bar{b}$, then $R'(0) \geq 0$.

Corollary 2.1 shows that if $\left(\bar{\lambda} - 1\right)/\gamma$ is sufficiently high, the share of researchers targeting a technology increases in the first periods after its emergence. As we know, in the long run, decreasing returns drive the growth rate in Q to zero. However, $\bar{\lambda} - 1$ crucially controls how fast it can grow in the short and medium run. For instance, if $\bar{\lambda} - 1 = 0$, newly created varieties have, on average, the same quality as existing ones, and Q's growth rate is null from the beginning. On the other hand, larger values for $\bar{\lambda}$ imply that the average quality will significantly increase as the next better innovation is introduced. Therefore, saying that $\left(\bar{\lambda} - 1\right)$ is sufficiently higher than γ effectively means that the initial growth rate in Q is higher than the frontier. As a young technology ages, it becomes increasingly more profitable relative to the age zero technology. As a result, the share of research directed to the former increases relative to the latter. Over time, however, as the technology keeps getting older, decreasing returns drive down the growth in Q, which falls below γ, triggering the process of obsolescence discussed above.

If $(\bar{\lambda} - 1)/\gamma$ is small, productivity growth within a technology already starts at a lower level relative to the frontier. As a result, it loses research share from the outset,

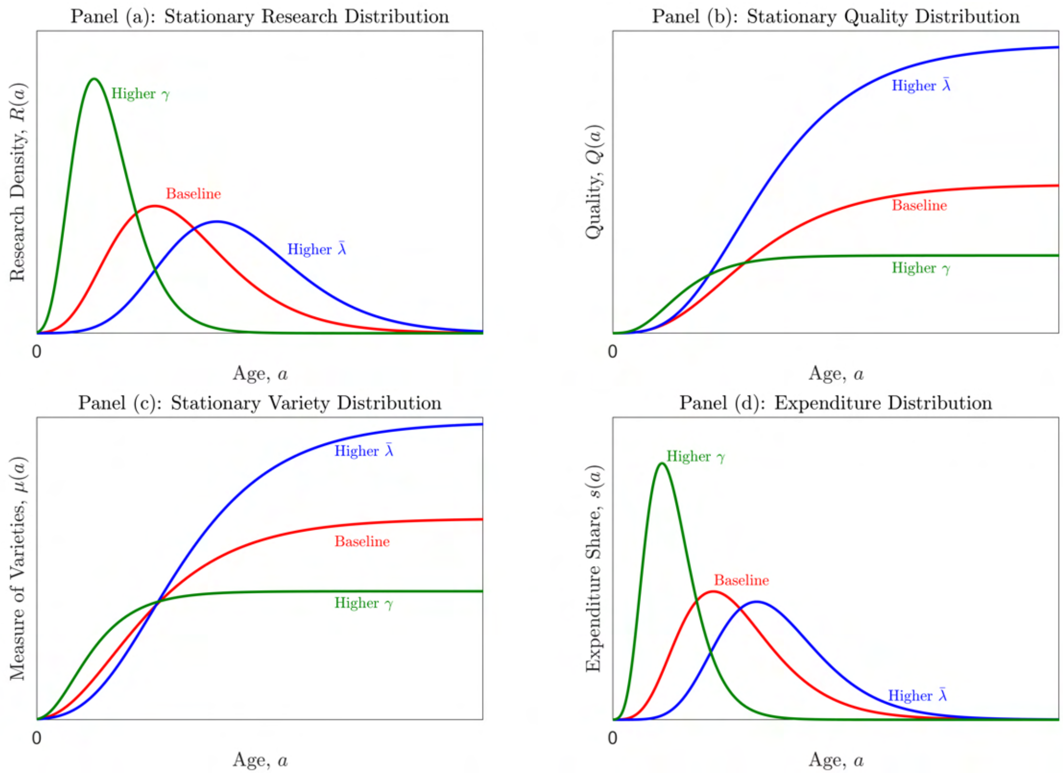

Figure 3: Stationary (Balanced Growth) equilibrium

Notes: The Figure displays the stationary distributions of research (Panel a) and expenditure shares (Panel d), and the age profiles for quality (Panel b) and varieties (Panel c). Each plot shows an illustrative calibration in red, the case of a higher γ in green, and the case of a higher $\bar{\lambda}$ in blue.

further reinforcing its slow growth. In this case, $R(a)$ is a decreasing function for all a, and no hump-shaped pattern emerges. For instance, from Proposition 2 and Equation (10), if $\bar{\lambda} = 1$, it follows directly that $R(a) \propto \exp(-a\gamma/\epsilon)$, implying an exponential decay.

Therefore, to match the hump-shaped research distribution documented in Section 2, it is key that research productivity—the growth rate in a technology quality produced by a constant number of researchers—declines over time. It should start sufficiently above the frontier growth rate γ, and declines below it. Falling research productivity is consistent with the evidence presented by Bloom et al. (2020) for several industries and goods. In this model, it resulted from the within technology decreasing returns to quality accumulation.

Illustrative calibration Exploiting Proposition 2's analytical characterization, Figure 3 illustrates the stationary distributions of research (Panel a) and expenditure shares (Panel d), along with the age profiles for quality (Panel b) and varieties (Panel c). Each plot shows an illustrative calibration in red, alongside two alternative calibrations: one where the frontier growth rate γ is increased (green curve) and another where the knowledge spillover parameter $\bar{\lambda}$ is higher (blue curve). The calibration in all plots is such that $\left(\bar{\lambda}-1\right)/\gamma$ is sufficiently high, implying that research shares initially increase with age.

Consider first the illustrative calibration, represented by the red curves. Suppose we track a fixed technology τ over time. In the stationary equilibrium shown in the plots, when τ is very young, the share of researchers targeting it is small (Panel a), as older technologies have already become very productive and offer higher knowledge spillovers. Consequently, while τ's average quality $Q(a)$ increases, it does so at a modest growth rate (Panel b). The same holds for the expenditure share it commands (Panel d).

Nonetheless, because technology τ average quality initially grows at least as fast as the frontier γ (Corollary 2.1), its research share continuously rises as it ages. This process accelerates as previously dominant older vintages face increasing diminishing returns and loses researchers. Technology τ then gains momentum, as can be seen in Panel (a) when its research share rises fast to the peak. Its quality growth rate accelerates (Panel b), and its expenditure share rises accordingly (Panel d). However, as the cohort continues to age, its growth rate declines due to diminishing returns, eventually falling below that of some younger vintages, triggering the process of obsolescence.

Finally, consider the blue and green curves in Figure 3. As discussed above, a higher γ (or lower $\bar{\lambda}$) makes frontier technologies relatively more attractive, leading researchers to abandon older technologies earlier. This shifts the mode of the research distribution to the left, favoring younger technologies. If γ continues to increase or $\bar{\lambda}$ to decrease, the distribution eventually becomes monotonically declining (e.g., see Figure E.5 in the Appendix for the case $\bar{\lambda}=1$).

3.3 The Efficient Allocation of Research Across Vintages

This Section characterizes the efficient allocation of resources. The economy exhibits three standard sources of inefficiency. The first is a static inefficiency arising from monopolistic competition in the production of varieties. The second is the congestion, or duplication, externality in the research sector, where multiple researchers may independently pursue the same idea. Third, dynamic externalities arise as individual researchers do not consider how their discoveries affect the development of quality and variety within a technology, thereby influencing the value of future innovations.

Importantly, given the assumption that research is carried out by a separate group of workers whose supply is inelastic, the first two sources do not distort the allocation of researchers across technology vintages. A uniform change in the profitability of the monopolistic producers across technologies neither changes the total investment in innovation nor affects how it is distributed. Similarly, while the duplication externality might typically lead to over-investment in innovation, as noted by Jones (1995), this is not the case here due to the fixed supply of researchers. Moreover, because expected returns are equalized across technologies in equilibrium, as guaranteed by Equation (8), there are no distortions in the allocative margin—a result formally proven by Hopenhayn and Squintani (2021).

The efficient allocation is characterized by the following optimization program:

$$\underset{\left[[R_t(a),\mu_t(a),Q_t(a)]_{a\geq 0},C_t\right]_{t\geq t_0}}{\text{Max}} \int_{t_0}^{\infty} e^{-\rho(t-t_0)} \ln C_t \, dt \quad \text{s.t.} \tag{16}$$

$$C_t = Y_t - X_t = \tilde{\alpha} \int_0^{\infty} e^{-\gamma a} Q_t(a)\mu_t(a) da \tag{17}$$

$$\frac{\partial \mu_t(a)}{\partial t} = -\frac{\partial \mu_t(a)}{\partial a} + R_t(a)^{1-\epsilon} \tag{18}$$

$$\frac{\partial Q_t(a)}{\partial t} = -\frac{\partial Q_t(a)}{\partial a} + Q_t(a)(\bar{\lambda}-1)\frac{R_t(a)^{1-\epsilon}}{\mu_t(a)} \tag{19}$$

$$R = \int_0^{\infty} R_t(a) da, Q_t(0) = Q_0, \mu_t(0) = \mu_0 \tag{20}$$

In this program, the objective function (16) represents the household's intertemporal utility. The resource constraints are given by (17) for the final good and (20) for researchers. Finally, (18) and (19) delineate the innovation possibilities frontier. The planner takes into account how the research allocation across vintages $[R(a)]_a$ impacts the evolution of $[\mu(a), Q(a)]_a$ through the laws of motion (18)-(19), effectively internalizing the dynamic externalities.

To gain insight into the market inefficiencies and how they can be corrected, it is worth focusing on Balanced Growth Path (BGP) trajectories, in which consumption and output grow at constant rates, and the age distributions are time-invariant. Formally, I will consider the set with all feasible BGP trajectories (which contains, for instance, the decentralized equilibrium BGP previously derived) and look for the one that maximizes the household intertemporal utility. As such, this analysis has the spirit of the well-known Solow Golden Rule, whose goal is to maximize consumption in the BGP.

Definition 3. (BGP feasible trajectories) A trajectory for the cross-sectional distributions of quality and varieties, $[\mathbf{\Theta}_t]_{t\geq t_0}$ (where $\mathbf{\Theta}_t \equiv [Q_t(a), \mu_t(a)]_{a=0}^{\infty}$), for the cross-sectional distribution of R&D, $[\mathbf{R}_t]_{t\geq t_0}$ (where $\mathbf{R}_t \equiv [R_t(a)]_{a=0}^{\infty}$), and for aggregate consumption $[C_t]_{t\geq t_0}$ is Balanced Growth feasible if, $\forall t \geq t_0, \forall a \geq 0$:

1. The distributions are time-invariant: $\partial_t Q_t(a) = \partial_t \mu_t(a) = \partial_t R_t(a) = 0$.
2. The innovation possibilities frontier is respected: given the trajectory for $R_t(a)$, the trajectories for $Q_t(a)$ and $\mu_t(a)$ satisfy the laws of motion in (12).
3. All resource constraints are respected.

Corollary 3.1. In every feasible Balanced Growth trajectory, aggregate consumption grows at the constant rate γ. Namely, for every $t \geq t_0$, $C_t = A_t\hat{C}$, where $A_t = e^{\gamma t}$, and $\hat{C}$ is a (trajectory specific) positive constant, which denotes the detrended level of consumption. The household intertemporal utility is hence $U = \int_{t_0}^{\infty} e^{-\rho(t-t_0)} \ln C_t dt = \frac{1}{\rho} \ln \hat{C} + B$, where B is a constant that depends on γ and ρ.

Corollary 3.1 shows that the Balanced Growth feasible trajectory that maximizes the intertemporal utility of the household must be the one that delivers the maximum *level* of consumption, $\hat{C}$. This problem can be written as:

$$\underset{[R(a)]_{a\geq 0}}{\text{Max}} \ \kappa \int_0^\infty e^{-\gamma a} Q(a) R(a)^{1-\epsilon} da, \quad \text{s.t.} \tag{21}$$

$$Q'(a) = Q(a)(\bar{\lambda} - 1)\frac{R(a)^{1-\epsilon}}{\mu(a)}, \tag{22}$$

$$\mu'(a) = R(a)^{1-\epsilon}, \tag{23}$$

$$\text{and } R = \int_0^\infty R(a) da, Q(0) = Q_0, \text{ and } \mu(0) = \mu_0, \tag{24}$$

where (21) represents the level of consumption $\hat{C}$ being maximized. It was obtained from the resource constraint (17) and transformed through integration by parts.[7] In turn, (22) and (23) represent the laws of motion for $Q_t(a)$ and $\mu_t(a)$ after imposing the condition $\partial_t Q_t(a) = \partial_t \mu_t(a) = \partial_t R_t(a) = 0$.

The objective function (21) shows that the direct impact of allocating $R(a)$ researchers to the age a technology is given by $e^{-\gamma a} Q(a) R(a)^{1-\epsilon}$. This expression arises because these $R(a)$ researchers produce a flow $R(a)^{1-\epsilon}$ of new varieties—which have average quality proportional to $Q(a)$, and intrinsic productivity (detrended by the long run aggregate growth rate) equal to $e^{-\gamma a}$.

In turn, the law of motion in (22) captures the impact of innovation on the average quality evolution. The more innovations are made in a technology of age a, the higher its average quality $Q(a')$ will be at older ages $a' > a$. This raises the direct impact of researchers who will work on the technology at such older ages, $e^{-\gamma a'} Q(a') R(a')^{1-\epsilon}$. Such mechanism represents a classic standing on the shoulder of giants' externality, and I will hereafter refer to it as such, or, for brevity, as the *Shoulder's externality ($\mathcal{E}_S$)*. It is captured by constraint (22) in the above problem.

One of the reasons this way of writing the problem is helpful is that $\mu(a)$ does not appear directly in the objective function, only in the constraints. This makes it clear how innovation, although having a direct positive effect on aggregate consumption (21), also has a cost. It moves the technology further into the region of diminishing returns. Specifically, innovation ($R(a)$) increases $\mu(a)$, as seen in Equation (23), which in turn will make proportional improvements on the average quality more difficult: in (22), the higher $\mu(a)$, the smaller $Q'(a)/Q(a)$ is in the face

[7] The expression in (21) equals the consumption level $\hat{C}$ up to a positive and finite additive constant (which does not distort optimality conditions).

of a flow of research $R(a)$. This is the second externality not taken into account in the decentralized equilibrium, which I will refer to as the *Dilution externality* ($\mathcal{E}_D$). It is captured by constraint (23) in the problem above.

Proposition 3. *Within the set of feasible BGP trajectories (see Definition 3), consider the problem of choosing the one that maximizes the household intertemporal utility. In a solution, the allocation of innovation must satisfy, for every $a \geq 0$:*

$$R(a)^{\epsilon} \propto \underbrace{e^{-\gamma a}Q(a)}_{\substack{\text{decentralized} \\ \text{equilibrium}}}\kappa + \underbrace{\psi_Q(a)(\bar{\lambda}-1)\frac{Q(a)}{\mu(a)}}_{\mathcal{E}_S(a)} + \underbrace{\psi_\mu(a)}_{\mathcal{E}_D(a)} \tag{25}$$

where $\psi_Q(a)$ and $\psi_\mu(a)$ are the costate functions associated with constraints (22) and (23), respectively.

If $\bar{\lambda} = 1$, there are no externalities, $\mathcal{E}_S(a) = \mathcal{E}_D(a) = 0$ for every $a \geq 0$, and the decentralized equilibrium allocation of research is efficient in the BGP.

If $\bar{\lambda} > 1$, the two externalities are active and satisfy:

(i)(Building on the shoulder of giants) For every $a \geq 0$:

$$\mathcal{E}_S(a) > 0 \text{ and } \mathcal{E}_S'(a) < 0 \tag{26}$$

The building on the shoulder of giants' externality leads a benevolent planner to innovate relatively more on young technologies compared to the decentralized equilibrium.

(ii)(Dilution effect) For every $a \geq 0$:

$$\mathcal{E}_D(a) < 0 \text{ and } \mathcal{E}_D'(a) > 0 \tag{27}$$

The second externality (the dilution effect) leads a benevolent planner to innovate relatively less on young technologies compared to the decentralized equilibrium.

*Proof.*See the Appendix. □

Proposition 3 contains three important results concerning the efficient allocation of research in the BGP. First, it shows in (25) how the two externalities highlighted above, Shoulders and Dilution, are the forces deviating the decentralized equilibrium from the efficient allocation. Were they not active, the equilibrium allocation of research would be efficient. This occurs if $\bar{\lambda} = 1$, and hence $Q(t|\tau)$ is constant in time, indicating the absence of any perfection process within each technology. In this case, the benefits of research are limited to the horizontal expansion in the mass of varieties, with its direct effects on aggregate consumption fully internalized in private profits. If $\bar{\lambda} > 1$, however, knowledge accumulates within a technology in the form of its perfection process. By taking into account the externalities involved in this process (Shoulders and Dilution), a benevolent planner can lead to a greater accumulation of quality within each vintage, and hence to a higher level of detrended consumption.

It is, therefore, crucial to understand how these two externalities shape the optimal allocation of research in the case of $\bar{\lambda} > 1$. The third important result of Proposition 3, contained in (26), does so for the first. Proposition 3 shows that although the building on the shoulder of giants' externality results in a benevolent planner valuing innovation more than the market within every technology ($\mathcal{E}_S(a) > 0$ for every a), it values it *relatively* more than the market for younger technologies ($\mathcal{E}'_S(a) < 0$ for every a). In other words, the distribution of research efforts is biased towards old technologies in the decentralized economy—as far as the building on the shoulder of giants' effect is concerned.

Finally, the fourth important result from Proposition 3 relates to the second externality, the Dilution effect $\mathcal{E}_D$. In contrast to the first one, it leads the planner to value innovation less than individual researchers, and it values it relatively less for young technologies. This is intuitive: early inventions will make a technology average quality less elastic to the contribution of older inventions. By internalizing this cost, a benevolent planner would be less restrictive on innovating in older technologies than younger ones. For example, the cost of a bad idea outcome is much smaller when the technology is old, and many inventions have already been made: the new one will have only a marginal influence on the average quality.

The net effect of externalities $\mathcal{E}_S(a)$ and $\mathcal{E}_D(a)$ for a given technology a depends on the parameters. It can be shown, for example, that when $\bar{\lambda}$ is sufficiently high, the building on the shoulders of giants' externality will dominate, in the sense that $\mathcal{E}_S(a) + \mathcal{E}_D(a) > 0$. Regardless of the net signal of such effects, however, an important additional result regarding the efficient allocation can be derived, which is the subject of Proposition 4.

Proposition 4. *Within the set of feasible BGP trajectories, consider the problem of choosing* *the one that maximizes the household intertemporal utility. In a solution, the allocation of innovation must satisfy for every* $a \geq 0$*:*

$$R'(a) < 0 \text{ and } \lim_{a \to \infty} R(a) = 0$$

The optimal allocation of research is not hump-shaped. It monotonically falls with age.

*Proof.*See the Appendix. □

Proposition 4 states that the planner distribution of research is decreasing across vintages ($R'(a) < 0$): older technologies receive less weight and are abandoned asymptotically ($\lim_{a \to \infty} R(a) = 0$). This is in contrast to the decentralized equilibrium hump-shaped allocation of scientists. When a technology emerges with little accumulated knowledge (low average quality), there is value for individual scientists in waiting until others develop the foundations of the new field, so that they can jump in later, when returns in older fields are small and the knowledge accumulated in the new one is higher. This waiting value is not present for a planner, who prefers to develop the new field as soon as possible.

Discussion It is well known in the growth literature that the building on the shoulder of giants' externality leads to underinvestment in innovation, as inventors' profits do not incorporate positive spillovers on future research in a market equilibrium. This brings no surprise to the result that $\mathcal{E}_S(a) > 0$ for every $a \geq 0$, meaning that the planner values more innovation in every technology. The novel mechanism uncovered by this theory, however, is that in an economy where individual technologies eventually lose momentum and are gradually replaced by newer vintages, the non-internalized spillovers of research are higher for young than for older technology, $\mathcal{E}'_S(a) < 0$. In contrast to the literature on creative destruction following Grossman and Helpman (1991) and Aghion and Howitt (1992), which considers the accumulation of knowledge to be uniform, I assume there are different blocks of knowledge (technologies) that gradually and perpetually replace each other. The size of knowledge spillovers is no longer homogenous but rather asymmetric depending on the life cycle stage of the technology. Section B.2 in the Appendix relaxes the assumption that knowledge spillovers happen exclusively within a technology. It presents a model extension where cross-technology knowledge spillovers are allowed but to a smaller degree than the within-technology ones. The normative results discussed here still hold.

Depending on the theory one has in hand, other externality channels may be present in addition to this one, making the final results regarding underinvestment in innovation ambiguous. This is the case of the business stealing effect in the creative destruction literature (Aghion and Howitt, 1992) and will be the case of decreasing returns within vintage in this paper. Nonetheless, the widespread presence of the standing on giants' shoulders externality on growth models, either alone or combined with others, makes its study important.

4 Measurement: Data and Technology Vintages

This Section describes the patent data used throughout the paper, including its technology classes and their proxied age presented in Section 2. Section A in the Appendix develops alternative technology definitions and age measures. Specifically, Section A.1 uses the citation network, while Section A.2 leverages historical USPTO documents to identify when new technologies are added to the code.

4.1 Baseline Data: patents and technologies

The baseline dataset comprises approximately nine million patents, representing virtually the universe of utility patents issued by the United States Patent Office (USPTO) from 1836 to 2010. It is sourced from the USPTO Historical Patent Data Files (Marco et al., 2015). Unless otherwise noticed, patents are binned by decade based on their issue date.

A key aspect of the theory is the technological vintage structure. To implement this concept, the paper relies on the USPTO's technological classification system (USPC). Each of its 401 utility classes represents a distinct technology in the analysis that follows. The classification assigned to each patent, from the oldest to the newest, is consistent with the latest version of the USPC code. Whenever a class is created or abolished, the USPTO retroactively reclassifies all previously issued patents to ensure consistency with the updated classification system. Moreover, the USPC system classifies each patent into one and only one principal technology class. The analysis here focuses exclusively on that classification.

4.2 Measuring the emergence of new technologies

In the theory, a vintage or technology is defined by its emergence date. Hence, a procedure is needed to determine the age of technologies at a given point in time. As first discussed in Section 2, I follow the seminal approach of Griliches (1957).

Technologies diffuse in the innovation space in a manner qualitatively similar to their diffusion in product and consumption markets. The latter refers to the gradual adoption of a technology by firms and consumers, as studied by Griliches and much of the diffusion literature. The former pertains to the increasing share of innovations related to a given technology. Over time, researchers and firms increasingly adopt it as a platform for further development. Starting from low levels, patenting activity within a technology grows, and eventually stabilizes (in relative terms), or even decreases. This S-shaped diffusion pattern naturally lends itself to a logistic fit, as explored by Griliches (1957).

This intuition is supported by the example of individual technologies, such as the steam engine case in Figure 1, or by the technology classes of the USPC system, which are the focus here. For instance, notice in Panel (a) of Figure 4 how the share of patents for 'Amplifiers' was small and only slowly increased by the turn of the 19th and 20th centuries. It then gained momentum and significantly increased before reaching a peak in 1960, a trend also observed in the other panels of Figure 4. More generally, such a trend is corroborated by multiple case studies highlighting a consistent S-shaped pattern in patent timelines (Achilladelis et al., 1990; Achilladelis, 1993; Andersen, 1999; Haupt et al., 2007).

Following Griliches, I run a logistic regression for each technology class so as to fit its diffusion trend in the innovation space. Specifically, for each class, I define its diffusion period as the period that ends when its share reaches the maximum. I then fit a logistic curve to it and use the 10 percent of such maximum value as a crossing point to define the origin date proxy. Similarly to the notion of availability in Griliches, the focus is not on the first discovery associated with a technology, but on identifying when it becomes a significant platform for subsequent innovations.

Figure 4 further clarifies this procedure by providing examples of the estimation. To begin, first notice that the dots in Panels (a)-(d) depict the share of patents accrued by each technology over time—and have been normalized by the maximum value achieved in each case. For instance, the share of patents associated with the

‘Amplifiers’ technology, among all patents issued in a given decade, reached a maximum in the 1960s, being normalized to unit by then. This maximum value also delimits the sample used for the fit: only a technology’s expansion (diffusion) period is used, which I define as the entire period preceding the maximum share point. The resulting logistic trend for the diffusion of each technology is represented by the solid line. Finally, the dotted lines indicate the point (and the year) in which the estimated trend reached 10% of its maximum value. For the ‘Amplifiers’ technology, this happened in 1911.

Second, observe in Figure 4 the possibility of extrapolating the date of origin into the pre-sample period, which becomes clear when comparing Panels (a) and (b). In the ‘Amplifiers’ case, the first observations of its patent share are relatively far from the peak value, slowly increasing at first and then accelerating to approach the peak. This is the typical start of the diffusion pattern documented extensively in the S-shape literature. As a result, the logistic fit interprets the first observations as, in fact, the start of the technology diffusion process. In the ‘Coating’ technology case, however, when the records of patents from the US patent system started, at around 1830, the patenting share in Coating was already relatively stable and close to its peak. The logistic fit interprets this as an indication that the technology had emerged further in the past, having already surpassed the early stages of its diffusion period among inventors in previous decades. As a result, the logistic fit extrapolates the 10 percent mark to the year 1779.

From a total of 401 technology classes, the origin date was estimated by the above method for 343 of them. This is because, for some classes, the maximum achieved patenting share is such that no more than one observation remains for the logistic estimation (observations for patents are binned into 10-year periods). Additionally, to remove outlier estimates projecting the origin of technologies far back in the past or ahead in the future, the results were winsorized at the 2.5 and 97.5 percentiles, leaving me with 325 classes. They represented 93% of all patents issued in the 2000 decade and 76% in 1900.

Crucially, this empirical approach does not impose whether the share of innovation in a given technology, which is, for example, 1 year old, will be higher or lower than the share in a 50-year-old one. Similarly, it does not impose that two technologies will have the same patent share when each of them has a given age, say 10 years old. This is because, when establishing the 10 percent cutoff, the date

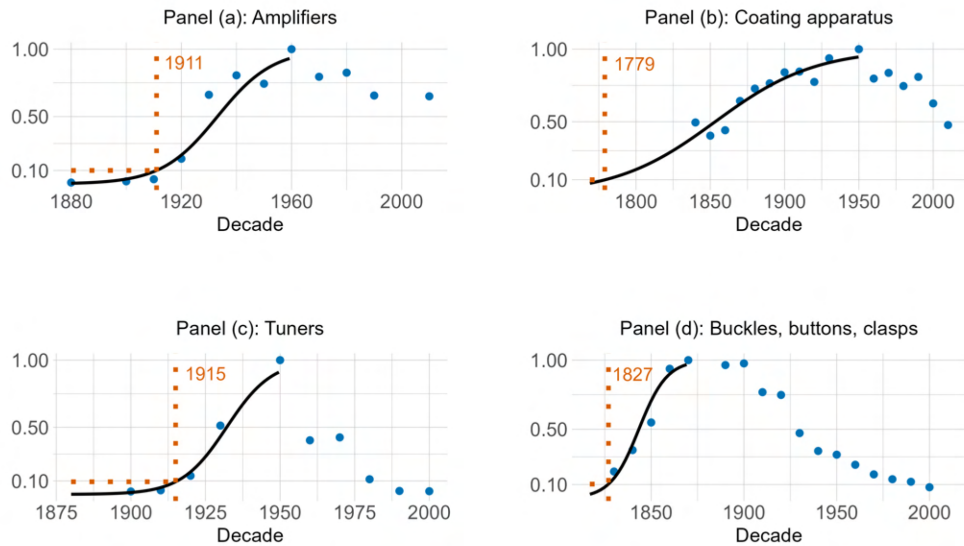

Figure 4: Patenting share and proxied emergence date

Note: The dots in each Figure represent the share of patents accrued by the respective technology class among all utility patents issued in a given decade. For each technology, observations below the 5th and above the 95th percentile were removed to avoid the influence of outliers. Shares are normalized in each Figure by the maximum value achieved. The solid line represents the fitted logistic trend (see this Section text and Appendix Section F). The dashed lines represent the year in which the trend achieved 10% of the ceiling value.

of origin estimation takes as a reference the maximum share achieved by each technology individually, without direct relation to others. Take as an example the cases of technologies 'Tuners', represented in Panel (c) of Figure 4, and 'Pipe Joints', in Figure E.3 in the Appendix. The highest innovation share achieved by 'Tuners' was in 1950, and it is less than five times the value for the earliest observation of Pipe Joints (for easy comparison, Figure E.3 presents the patent share evolution of both technologies with the actual, i.e., non-normalized values). The fact that a technology is young or old does not imply mechanically that its innovation share is higher or lower than others, at a fixed moment in time or at different points of their life cycle.

Results The second column in Table 1 presents the estimated emergence date for a set of key technologies. More generally, the left plot in Figure 5 shows a time series with the number of technologies emerging in each decade. Three waves

Table 1: The emergence of key technologies

Technology class	Origin	Patenting share (%)	
		1950s	2000s
Gas: heating and illumination	1766	0.14	0.02
Aeronautics	1871	0.55	0.22
Railway draft appliances	1826	0.10	0.00
Electric lamp	1878	0.06	0.08
Television	1929	0.36	0.81
Semiconductor device manufacturing	1964	0.05	3.07
Data Processing: artificial intelligence	1968	0.00	0.15
Software development, and installation	1985	0.00	0.34
Information security	1987	0.00	0.25

Note: This Table presents examples of technology classes (USPC classification system) with their corresponding estimated date of origin and patenting shares in two different decades. The date of origin is estimated with a logistic regression, following the approach in Griliches (1957) (see this Section text and Appendix Section F).

can be distinguished, representing periods with many technologies being created. The first wave happened early in the 19th century, coinciding with the first industrial revolution in the United States. In my empirical analysis, it represented the emergence of several technologies serving as a base for general manufacturing activities, such as 'Endless belt power transmission system', and 'Turning'. It also marked the emergence of the first classes associated with railways and electricity. The second wave can be distinguished late in the 19th century, with a peak during the 1880s, and expanded the period until the 1910s. It is dominated by the emergence of several classes related to electricity and the chemical industry, such as 'Battery or capacitor charging or discharging' and 'Synthetic resins or natural rubbers,' coinciding with the period of the second industrial revolution. The third wave happened in the second half of the 20th century, with its peak in the 1960s. Within it, several technologies related to computing and telecommunications were estimated to emerge, as is the case of 'Data Processing: artificial intelligence' in 1968.[8]

Initially, patents related to a new technology are often misclassified, typically assigned to residual categories such as "Miscellaneous", as they do not fit into the existing classification framework (see the event study in Section A.2 in the

[8]For comparison, the origin of AI is often attributed to the seminal conference in 1956 in which the term was first used: the Dartmouth Summer Research Project on Artificial Intelligence.

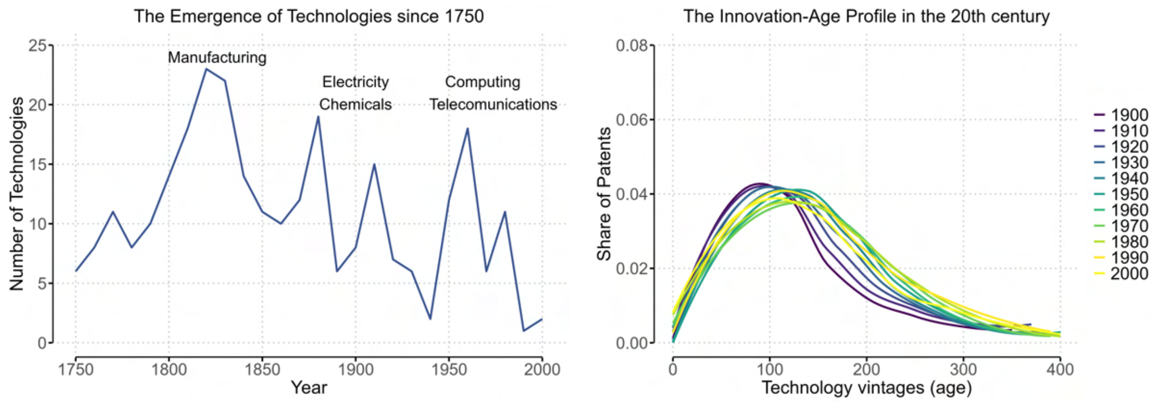

Figure 5: Technologies and patents across and within time

Note: The left plot shows the number of technology classes (UPSC classification system) that emerged in each decade. The emergence date is estimated with a logistic regression, following Griliches (1957) (see this Section text and Appendix Section F). The right plot shows, for different decades, the distribution of patents across USPC technology classes sorted by their estimated age. Each line plots a non-linear fit of the share of patents across the different technology-age groups.

Appendix). Therefore the number of technologies that emerged in recent decades is downward biased. However, the quantitative analysis in this paper mainly focuses on cross-sectional distributions, which have been qualitatively stable for more than a century. This is shown in the right plot in Figure 5, which depicts, for different decades, the share of patents accrued by technologies of different ages. As discussed in Section 2, for each point in time, the distribution displays a pronounced hump shape. Here, however, we can observe this fact holds not only in the 1900 and 2000 decades, as shown in Figure 2, but all across the 20th century. The hump shape and the location of the distribution present a distinct stability. Finally, Section E in the Appendix redo the analysis of Figures 1-2 controlling for patent quality.

5 Quantitative Analysis

This Section explores the potential quantitative relevance of the research misallocation identified in the theory. This is done using a range of calibrations.

5.1 Baseline calibration

The model presented in Section 3 contains a total of six parameters $(\epsilon, \bar{\lambda}, \gamma, \beta, \rho, Q_0, \mu_0)$ whose baseline calibration is sequentially discussed below.

Calibrating ϵ The parameter ϵ governs how elastic research efforts are to a technology's profitability relative to others. Low values of ϵ indicate that even small increases in a technology's relative profitability lead to significant redirection of research toward it. As a result, its calibration can be informed by the observed comovement between research intensity and patent valuation across different technologies. For that, it is important to clearly establish a mapping between the data and the model's objects.

The paper assumes the observed number of patents is an indicator, albeit noisy, of the total number of new ideas generated in a specific time period. Formally,

$$Patents(t|\tau) = \kappa_\tau \times \dot{\mu}(t|\tau) \times u(t|\tau), \tag{28}$$

where $Patents(t|\tau)$ denotes the number of patents issued in a given period within the technological class τ, while $\dot{\mu}(t|\tau)$ is the number of new ideas in the model. In turn, κ_τ is a constant unique to the technological class, indicative of its patenting propensity, and $u(t|\tau)$ is a scalar disturbance term. This specification allows for variation in patenting rates across technologies, an essential flexibility when analyzing innovation across diverse fields.

Using the model, substitute $\bar{v}(t|\tau)$ for $\dot{\mu}(t|\tau)$ and rewrite (28) as the regression equation

$$\log Patents(t|\tau) = \delta_\tau + \delta_t + \frac{1-\epsilon}{\epsilon} \log \bar{v}(t|\tau) + \log u(t|\tau). \tag{29}$$

Here δ_τ and δ_t represent technology and time-fixed effects, grouping parameters and aggregate variables from the model.[9] The number of new patents, representing the left-hand side variable, is directly observed. On the right-hand side, $\bar{v}(t|\tau)$ is constructed using Kogan et al. (2017) estimates for patents' market value.

Informed by this relation, the baseline calibration chooses ϵ to account for the

[9] $\delta_\tau = \log \kappa_\tau + \log \eta - \frac{1-\epsilon}{\epsilon} \log(\beta \bar{\lambda})$, and $\delta_t = -\frac{1-\epsilon}{\epsilon} \gamma t - \frac{1-\epsilon}{\epsilon} \log(D_t \bar{Q}_t)$.

reduced form correlation between patenting activity and valuation, after filtering out time and technology-fixed effects. It calibrates $\epsilon = 0.65$ to match the coefficient on $\bar{v}(t|\tau)$ estimated from regression (29). Section D.1 in the Appendix presents the complete results. Most importantly, later in this Section, a range of ϵ values will be tested.

Calibrating the frontier growth rate parameter γ The parameter γ represents the rate at which new technologies are intrinsically better than previous vintages. Along a balanced growth, it equals the output growth rate (see Proposition 2). The baseline calibration, thus, sets γ to 2%, consistently with the approximate average GDP per capita growth in the US economy over the past century.

Calibrating the knowledge spillover parameter $\bar{\lambda}$ The parameter $\bar{\lambda}$ critically determines how the average quality of varieties within a technology evolves. If $\bar{\lambda} = 1$, newly added varieties are of the same quality as existing ones, which means the average quality $Q(a)$ remains flat as the technology ages. Conversely, when $\bar{\lambda}$ is significantly larger than 1, new varieties are substantially better, and the average quality increases over time. In principle, therefore, $\bar{\lambda}$ could be calibrated to match the age profile $\{Q(a)\}_{a=0}^{\infty}$. Unfortunately, the latter is not directly observable.

The theory, however, establishes a direct relationship between a technology's average quality and its market profitability. From Equation (9), fixing a given time t, the expected value of innovating on an age a technology—relative to the frontier vintage—is given by:

$$\frac{\bar{v}_t(a)}{\bar{v}_t(0)} = e^{-\gamma a} \left(\frac{Q_t(a)}{Q_t(0)} \right). \tag{30}$$

Using information on the stock market value of new patents (Kogan et al., 2017) to back up the left-hand side in (30), and having set γ to 2%, we can construct an empirical measure for the average quality (relative to $Q_t(0)$) on the right-hand side.

Given the baseline calibration for the remaining parameters, $\bar{\lambda}$ is chosen such that the model's stationary quality age schedule $\{Q(a)/Q(0)\}$ matches its empirical counterpart recovered from Equation (30). Section D.2 in the Appendix provides the complete details. The resulting value for $\bar{\lambda}$ is 2.34.[10]

[10]While this may look large, remember $\bar{\lambda}$ denotes the step size over the average quality—and not over the best idea. Most importantly, later in this Section, a range of values will be tested.

Table 2: Baseline Calibration

Parameter	Description	Value	Parameter	Description	Value
γ	Frontier growth rate	2%	β	Labor share	0.66
$\bar{\lambda}$	Knowledge spillover	2.34	ρ	Discount rate	2.5%
ϵ	Research elasticity	0.65	$Q_0,\ \mu_0$	Initial conditions	1, 1

Other parameters and Model Fit Following commonly used values in the literature, the discount rate ρ is set to 2.5%, although results are reported for a range of values. The parameter β, representing the labor share in total output, is set to 0.66. The labor force L, the supply of researchers R, and the initial conditions Q_0 and μ_0 are set to one. These latter parameters do not influence the allocation of research across technologies, as they affect all technologies symmetrically. Table 2 summarizes the baseline calibration.

Figure 6 compares the empirical distribution of patents in the 2000 decade (solid green line) with the stationary age profile of innovation produced by the calibrated model (dashed blue line, representing the competitive equilibrium). The model is successful in fitting the data, as demonstrated by the proximity of the curves. Notably, the peak of the distribution generated by the model occurs around technologies aged approximately 50-100 years, mirroring the empirical peak, despite this moment not being targeted during calibration.

5.2 The Costs of Research Misallocation

To quantify the costs of research misallocation between new and old technologies, this Section uses the calibrated model to compare two balanced growth path (BGP) trajectories. The first, the Laissez-faire BGP, represents the decentralized equilibrium described in Proposition 2. The second, the efficient BGP, reflects the socially optimal allocation determined by a welfare-maximizing planner, as described in Proposition 3.

Figure 6 shows how the efficient allocation of research differs from the Laissez-faire equilibrium. The dot-dashed purple line shows the share of patents accrued to technologies of different ages in the efficient BGP. As predicted by the theory, contrary to the hump-shaped distribution observed in laissez-faire conditions (dashed blue line) and in the data (solid green line), the socially optimal number of patents

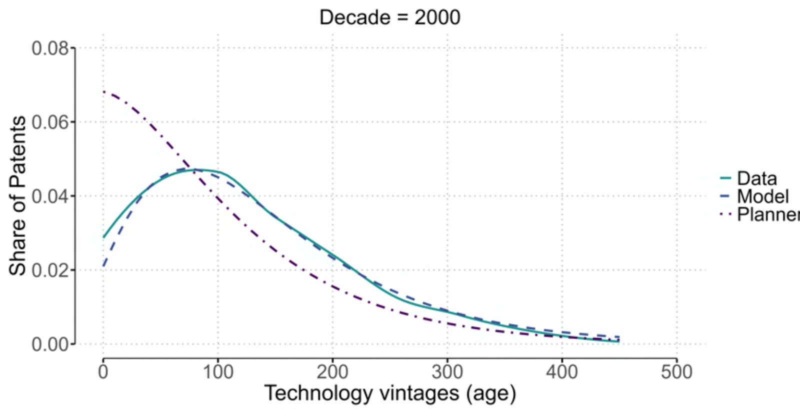

Figure 6: The Hump-Shaped Innovation-Age Profile—Competitive Equilibrium, Planner, and Data

Note: The three lines in the plot represent the share of patents allocated to technologies of different ages. The solid (green) line shows a non-linear fit based on US data from the 2000 decade (see Sections 2 and 4), with observations trimmed at the top 1% to remove outliers. The dashed (blue) line represents the model's competitive equilibrium stationary patenting distribution ($\dot{\mu}(a)$) under the baseline calibration (see Table 2). The dot-dashed (purple) line represents the optimal stationary patent allocation derived from the planner's problem (see Sections 3.3 and 5.2).

monotonically declines as technologies age. Moreover, the share of patents accruing to the frontier technology (age zero) is larger in the planner allocation than in the competitive equilibrium.

Correcting the research misallocation has significant long-run impacts on consumption. The left-hand side of Table 3 reports the ratio of consumption levels in the two BGP trajectories. Each row presents the results under a different calibration for $\bar{\lambda}$, while the remaining parameters are held at their baseline values in Table 2. Notice, in particular, that when $\bar{\lambda}$ also equals its baseline value, 2.34, consumption along the efficient BGP is 30% higher at every point in time, in comparison to what it would be in the Laissez-faire BGP. As the different rows show, this gain monotonically increases with $\bar{\lambda}$ and can reach up to 60% when $\bar{\lambda}$ is slightly raised to 2.7. Intuitively, this happens because the planner takes advantage of the higher knowledge spillovers offered by new technologies, shifting research in their direction. The larger the spillovers (i.e., the larger $\bar{\lambda}$), the more gains there are to exploit with this reallocation.

Finally, the right-hand side of Table 3 examines the gains in consumption under alternative calibrations for ϵ while keeping all other parameters, including $\bar{\lambda}$, at

Table 3: Comparing Consumption in Different BGP trajectories

$\bar{\lambda}$	$\frac{\hat{C}_{\text{efficient}}}{\hat{C}_{\text{Laissez-faire}}}$	ϵ	$\frac{\hat{C}_{\text{efficient}}}{\hat{C}_{\text{Laissez-faire}}}$
2.70	1.60	0.60	1.34
2.34	1.30	0.65	1.30
2.00	1.13	0.70	1.24
1.50	1.03	0.80	1.13

Note: This Table presents the ratio of detrended consumption ($\hat{C} = e^{-\gamma t} C_t$) between the two BGP trajectories (Laissez-faire and BGP). The Table's left side studies changes in $\bar{\lambda}$, holding fixed the remaining parameters (including ϵ) in the baseline calibration (see Table 2). Each line presents an alternative $\bar{\lambda}$ value used and the resulting gain in detrended consumption. The right part of the Table repeats this procedure, but varying ϵ, while keeping $\bar{\lambda}$ constant at the baseline.

their baseline values. The larger the congestion force ϵ, the smaller the impact brought by the efficient allocation. As we know, to obtain consumption gains by exploiting spillovers, the planner needs to reallocate research across technologies. If congestion forces are very high, this reallocation becomes prohibitively costly, significantly reducing the potential gains. However, even with high levels of ϵ such as 0.80, the increase in consumption is still significant and greater than 10%.

5.2.1. Transition dynamics

Although the efficient allocation of research significantly increases consumption in the BGP, it may take long for such gains to be realized. In the short run, the growth rate may even decrease as we shift resources to young technologies that are not yet very productive.

This Section studies a policy that gradually steers the economy from the Laissez-faire BGP toward the efficient one. Specifically, consider the following simple rule:

$$R_t^{PM}(a) = (1-\alpha)R_t^{LF}(a) + \alpha R^*(a). \tag{31}$$

Here, $R_t^{PM}(a)$ denotes the allocation chosen by the policymaker, $R^*(a)$ represents the long-run efficient allocation, and $R_t^{LF}(a)$ is the allocation that would arise under free markets, given the economy's state $[Q_t(a), \mu_t(a)]_a$ at time t. In words, at every point in time t, the policy maker tilts the research allocation that would prevail under free markets, bringing it closer to the long-run efficient allocation. The degree

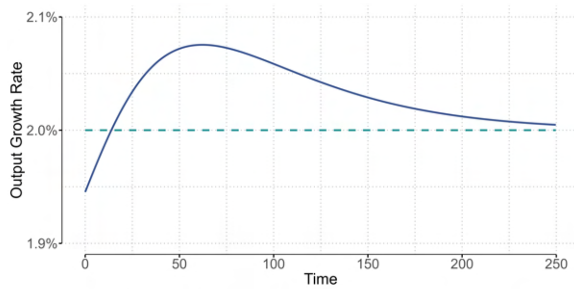

Figure 7: Transition dynamics

Note: This Figure shows (solid line) the trajectory of the consumption growth rate as the economy moves from the Laissez-faire BGP to the efficient BGP—under the policy rule defined in Equation (31). The model calibration follows the baseline values presented in Table 2, and α is set to 0.6.

Table 4: Welfare Gains at Different Calibrations

ρ	Welfare Gain	$\bar{\lambda}$	Welfare Gain	ϵ	Welfare Gain
0.01	10%	2.7	5%	0.60	4%
0.025	3%	2.0	2%	0.70	3%
0.05	1%	1.5	1%	0.80	1%

Note: This Table reports the consumption-equivalent welfare gains resulting from the policy described in Section 5.2.1. The Table's left side studies changes in ρ, holding fixed the remaining parameters (including $\bar{\lambda}$ and ϵ) in their baseline calibration values (see Table 2). Each line presents an alternative ρ value used, and the resulting welfare gain. The central and left parts of the Table repeat this procedure, but varying individually $\bar{\lambda}$ and ϵ, respectively.

of this adjustment is given by $\alpha \in [0,1]$. If $\alpha = 1$, the policy maker radically sets, at time zero and once and for all, the efficient distribution $R^*(a)$. If $\alpha = 0$, the policymaker decides not to act, and the economy is always on its decentralized equilibrium path.

Assume our calibrated economy is at the Laissez-faire BGP for $t \leq 0$. Thus, the actual research allocation at $t = 0$ equals $R_0^{LF}(a)$, which is the stationary hump-shaped distribution depicted in Figure 6. Starting at $t = 0$, a policymaker implements the rule in (31), choosing a constant α to maximize household welfare. Figure 7 presents the resulting transition dynamics. The solid blue line traces the trajectory for the growth rate of output when α equals 0.6, the welfare-maximizing choice. The dashed line represents the exogenous long-run growth rate.

When the policymaker starts to shift research toward new technologies, we observe at first a temporary slowdown: the growth rate immediately jumps down and remains below 2% for 15 years. Instead of creating varieties linked to technologies that are already very productive, researchers start to create more varieties linked to new technologies, which have not been perfected, and have lower productivity and market share. This impacts growth negatively at first. However, after 15 years, growth under the counterfactual scenario becomes higher than under Laissez-faire for the next several decades.

Importantly, correcting the research misallocation improves welfare. Table 4 presents the welfare gains in consumption equivalent units associated with this policy. As shown in the first column, the size of these gains depends on the discount rate, which reflects the household's willingness to accept an initial growth slowdown in exchange for a permanent consumption level increase. Under the baseline calibration, $\rho = 0.025$, welfare gain equals 3%. This means the representative household is indifferent between (i) the transition dynamics in Figure 7 and (ii) an alternative trajectory in which consumption would remain growing at 2%, but at a level 3% higher than in the Laissez-faire BGP. With a lower discount, such as $\rho = 0.01$, the welfare gain becomes as large as 10%, and with a more conservative choice, $\rho = 0.05$, it equals 1%, but still positive. Finally, the central and right-hand side parts of Table 4 study how welfare gains vary when we change $\bar{\lambda}$ and ϵ, respectively. The same patterns observed in Table 3, when we studied the BGP only, apply here as well for the transition dynamics. Namely, welfare impacts increase on $\bar{\lambda}$ and decrease on ϵ.

6 Conclusion

It is common practice for national and international research agencies, such as the National Science Foundation or the European Research Council, to sponsor dedicated funding programs for the advancement of cutting-edge research in new technological paradigms. Despite such policies, the existing economic literature on growth has yet to reveal a reasoning behind this *selective* approach—or the potential biases within market forces that might lead to the disproportionate allocation of research efforts towards established, mature technologies. Knowledge externalities are a standard justification for public funding of research and have underlaid seminal contributions on the importance of basic, as opposed to applied, R&D (Nelson, 1959; Akcigit et al., 2020). However, in standard endogenous growth models (Romer, 1990; Grossman and Helpman, 1991; Aghion and Howitt, 1992), knowledge externalities are assumed to be homogeneous, thereby muting the issue of determining the relative desirability of different types of innovation from a societal perspective. This paper addresses this limitation by introducing an innovation-driven growth model with vintage technologies. It shows how the equilibrium distribution of R&D efforts across technologies of different ages emerges and why it may fail to be efficient.

In focusing on knowledge spillovers, however, the model abstracted from important considerations that should motivate future work. Key among them is understanding the inherent risks associated with emerging technologies, vital for policy assessment.

References

Daron Acemoglu. Directed technical change. *The Review of Economic Studies*, 69(4): 781–809, 2002.

Daron Acemoglu. *Introduction to Modern Economic Growth*. Princeton University Press, New Jersey, 2009.

Daron Acemoglu and Fabrizio Zilibotti. Productivity differences. *The Quarterly Journal of Economics*, 116(2):563–606, 2001.

Daron Acemoglu, Philippe Aghion, Leonardo Bursztyn, and David Hemous. The environment and directed technical change. *American Economic Review*, 102(1): 131–66, 2012.

Daron Acemoglu, Ufuk Akcigit, Douglas Hanley, and William Kerr. Transition to clean technology. *Journal of Political Economy*, 124(1):52–104, 2016.

Daron Acemoglu, Philippe Aghion, Lint Barrage, and David Hemous. Climate change, directed innovation, and energy transition: The long-run consequences of the shale gas revolution. *NBER Working Paper*, w31657, 2023.

Basil Achilladelis. The dynamics of technological innovation: the sector of antibacterial medicines. *Research Policy*, 22(4):279–308, 1993.

Basil Achilladelis, Albert Schwarzkopf, and Martin Cines. The dynamics of technological innovation: the case of the chemical industry. *Research Policy*, 19(1):1–34, 1990.

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

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

Ufuk Akcigit, Douglas Hanley, and Nicolas Serrano-Velarde. Back to Basics: basic Research Spillovers, Innovation Policy, and Growth. *The Review of Economic Studies*, 88(1):1–43, 2020.

Birgitte Andersen. The hunt for s-shaped growth paths in technological innovation: a patent study. *Journal of Evolutionary Economics*, 9(4):487 – 526, 1999.

Kenneth J. Arrow. The economic implications of learning by doing. *The Review of Economic Studies*, 29(3):155–173, 1962.

Andrew Atkeson and Patrick Kehoe. Modeling the transition to a new economy: lessons from two technological revolutions. *The American Economic Review*, 97(1): 64–88, 2007.

David Autor, Caroline Chin, Anna Salomons, and Bryan Seegmiller. New frontiers: the origins and content of new work, 1940–2018. *The Quarterly Journal of Economics*, 139(3):1399–1465, 2024.

Jess Benhabib and Aldo Rustichini. Vintage capital, investment, and growth. *Journal of Economic Theory*, 55(2):323–339, 1991.

Nicholas Bloom, Mark Schankerman, and John Van Reenen. Identifying technology spillovers and product market rivalry. *Econometrica*, 81(4):1347–1393, 2013.

Nicholas Bloom, Charles I. Jones, John Van Reenen, and Michael Webb. Are ideas getting harder to find? *The American Economic Review*, 110(4):1104–1144, 2020.

V. V. Chari and Hugo Hopenhayn. Vintage human capital, growth, and the diffusion of new technology. *Journal of Political Economy*, 99(6):1142–1165, 1991.

William Cleveland, Eric Grosse, and William Shyu. *Local Regression Models*, chapter 8. Routledge, 1992.

Diego Comin and Bart Hobijn. An exploration of technology diffusion. *The American Economic Review*, 100(5):2031–2059, 2010.

Nicholas Crafts. Steam as a general purpose technology: a growth accounting perspective. *The Economic Journal*, 114(495):338–351, 2004.

Warren D. Devine. From shafts to wires: Historical perspective on electrification. *The Journal of Economic History*, 43(2):347–372, 1983.

Eric Donald. Spillovers and the direction of innovation: an application to the clean energy transition. *Working Paper*, 2024.

Giovanni Dosi. Technological paradigms and technological trajectories: a suggested interpretation of the determinants and directions of technical change. *Research Policy*, 11(3):147–162, 1982.

Zvi Griliches. Hybrid corn: an exploration in the economics of technological change. *The American Economic Review*, 25(4):501–522, 1957.

Gene Grossman and Elhanan Helpman. *Innovation and Growth in the Global Economy*. MIT Press, Cambridge, Mass., 1991.

C. K. Harley. On the persistence of old techniques: the case of north american wooden shipbuilding. *The Journal of Economic History*, 33(2):372–398, 1973.

Reinhard Haupt, Martin Kloyer, and Marcus Lange. Patent indicators for the technology life cycle development. *Research Policy*, 36(3):387–398, 2007.

Elhanan Helpman and Manuel Trajtenberg. A time to sow and a time to reap: growth based on general purpose technologies. *NBER working paper series*, 4854, 1994.

Elhanan Helpman and Manuel Trajtenberg. Diffusion of general purpose technologies. *NBER working paper series*, 5773, 1996.

Hugo Hopenhayn and Francesco Squintani. On the direction of innovation. *Journal of Political Economy*, 129(7):1991–2022, 2021.

Leif Johansen. Substitution versus fixed production coefficients in the theory of economic growth: a synthesis. *Econometrica*, 27(2):157–176, 1959.

Charles I. Jones. R & d-based models of economic growth. *Journal of Political Economy*, 103(4):759–784, 1995.

Boyan Jovanovic and Saul Lach. Entry, exit, and diffusion with learning by doing. *The American Economic Review*, 79(4):690–699, 1989.

Boyan Jovanovic and Sai Ma. Uncertainty and growth disasters. *Review of Economic Dynamics*, 44:33–64, 2022.

Boyan Jovanovic and Yaw Nyarko. Learning by doing and the choice of technology. *Econometrica*, 64(6):1299–1310, 1996.

Boyan Jovanovic and Yuri Yatsenko. Investment in vintage capital. *Journal of Economic Theory*, 147(2):551–569, 2012.

Bryan Kelly, Dimitris Papanikolaou, Amit Seru, and Matt Taddy. Measuring technological innovation over the long run. *American Economic Review: Insights*, 3(3): 303–20, 2021.

Tor Jakob Klette and Samuel Kortum. Innovating firms and aggregate innovation. *Journal of Political Economy*, 112(5):986–1018, 2004.

Leonid Kogan, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. Technological Innovation, Resource Allocation, and Growth*. *The Quarterly Journal of Economics*, 132(2):665–712, 2017.

Thomas S. Kuhn. *The Structure of Scientific Revolutions*. University of Chicago Press, Chicago, 1962.

Francois Lafond and Daniel Kim. Long-run dynamics of the u.s. patent classification system. *Journal of Evolutionary Economics*, 29:631–664, 2019.

Ernest Liu and Song Ma. Innovation networks and r&d allocation. *NBER working paper series*, 29607, 2023.

Edwin Mansfield. Technical change and the rate of imitation. *Econometrica*, 29(4): 741–766, 1961.

Alan Marco, Michael Carley, Steven Jackson, and Amanda Myers. The uspto historical patent data files: two centuries of innovation. *USPTO Economic Working Paper*, (2015-1), 2015.

Phillippe Michel. On the transversality condition in infinite horizon optimal problems. *Econometrica*, 50(4):975–985, 1982.

Joel Mokyr. *The Lever of the riches*. Oxford University Press, New York, 1990.

Richard R. Nelson. The simple economics of basic scientific research. *Journal of Political Economy*, 67(3):297–306, 1959.

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

Stephen Redding. Path dependence, endogenous innovation, and growth. *International Economic Review*, 43(4):1215–1248, November 2002.

Paul M. Romer. Endogenous technological change. *Journal of Political Economy*, 98 (5):S71–S102, 1990.

Thomas Sampson. Dynamic selection: an idea flows theory of entry, trade, and growth. *The Quarterly Journal of Economics*, 131(1):315–380, 2015.

Atle Seierstad and Knut Sydsæter. *Optimal Control Theory with Economic Applications*. Elsevier Science, 1987.

Robert Solow. Investment and technical progress. In Kenneth Arrow, Samuel Karlin, and Patrick Suppes, editors, *Mathematical Methods in the Social Sciences*. Stanford University Press, Stanford, California, 1960.

United States Patent and Trademark Office. *The classification of patents*. U.S. Government Printing Office, 1915.

US Patent and Trademark Office and US Department of Comerce. *Development and use of patent classification systems*. U.S. Government Printing Office, 1966.

USPTO. Overview of the u.s. patent classification system (uspc). 2012.

Péter Érdi, Kinga Makovi, Zoltán Somogyvári, Katherine Strandburg, Jan Tobochnik, Péter Volf, and László Zalányi. Prediction of emerging technologies based on analysis of the us patent citation network. *Scientometrics*, 95:225–242, 2013.

A Alternative Measures of Technology and Age

A.1 Using Breakthrough Patents to Define Technologies and their emergence date

Using textual analysis, Kelly et al. (2021) identify a small set of breakthrough patents that are very novel and influential (e.g., Samuel Morse's telegraphy wire patent; Google's PageRank algorithm patent). Here, I consider each of these breakthrough patents to represent the emergence of a new technology field. The emergence date of a technology, thus, is the year its associated breakthrough patent was issued. Then, for each of the non-breakthrough patents (representing the vast majority), I use the citation network and find its closest breakthrough. To make this procedure clear, consider a simple example. Suppose a given patent A cites patents $\{B, C\}$, while B cites $\{D, E\}$ and C cites $\{F\}$. Hence, X's first degree citations are $\{B, C\}$ and X's second degree citations are $\{D, E, F\}$, with the same procedure applying for higher orders. The exercise here involves finding the breakthrough patent cited by x with the lowest citation order. For example, if B and F are breakthrough patents, I assign X to the breakthrough technology B. In other words, I consider patent X as an innovation building on the technology field born with B. If both $\{B, C\}$ were breakthrough patents, I would consider X as $1/2$ a patent building on technology B and $1/2$ building on technology C.

After assigning patents to technologies (breakthroughs), it is possible to study how patents issued at a given period are distributed across breakthrough technologies of different ages. Figure A.1.1 plots the resulting distributions for different decades, from the 1960s, soon after the USPTO starts to record citations, until the 2000-2009 decade. The results show a hump shape pattern, with most patents building on technologies that emerged few decades in the past. In the 1990s, the distribution significantly shifts towards young technologies. This is driven by the emergence of several influential breakthroughs in that decade related to the ICT

Table A.1.1: The Hump-Shaped Innovation age Profile (Alternative Technology Definition)

		Share of Patents Building on Breakthroughs of Age:			
		[0,10)	[0,20)	50+	70+
Decade	1960	6%	31%	15%	6%
	1970	5%	22%	29%	8%
	1980	11%	26%	32%	11%
	1990	18%	43%	27%	14%
	2000	9%	54%	19%	10%

Note: Breakthrough patents are identified by Kelly et al. (2021). All remaining patents are assigned to breakthroughs using the citation network (see the text of this session for details).

revolution. Table A.1.1 shows how the shares of innovation in very recent breakthroughs are similar (or even smaller) than the share building on breakthroughs 70 or more years old. This reinforces the importance of old technologies for innovation and the hump-shape pattern.

A.2 New technologies and the US classification code evolution

This Section leverages several historical publications from the USPTO to construct a new dataset documenting the evolution of the patent classification code in the United States (USPC). It proxies the emergence of a technology class by the year it was included in the USPTO classification code.

The classification code over time In 1830, the United States Congress issued one of the earliest known patent classification schemes. This initial system comprised 16 groups, ranging from 'Wheel carriages' and 'Lever and screw power' to 'Factory Machine' (Lafond and Kim, 2019). As the world evolved over the subsequent two centuries, so too did the classification system, "periodically amended to accommodate new technologies" (USPTO, 2012). By 1836, it had grown to encompass 22 classes, including, for example, 'Steam and Gas Engines.' In 1868, a new edition expanded the number of classes to 36. However, the most transformative amendment arrived in 1872 when a revised classification system was adopted (United States

Patent and Trademark Office, 1915). The number of classes increased from 36 to 145 to include, for the first time, examples such as 'Railways' and 'Photography'. Importantly, the 1872 reclassification "created the framework upon which the modern US system was built" (US Patent and Trademark Office and US Department of Comerce, 1966). This modernization process culminated in 1898 when the 'Classification Division' was established within the patent office, further professionalizing the process (Lafond and Kim, 2019).

Data The USPTO provides official creation dates for USPC technology classes.[11] However, many current classes were introduced as relabels of previously existing ones, involving only minor organizational changes. Take, for instance, the class 297 'Chair and seats,' established only in 1961, an implausibly later date for its emergence in the innovation space. It turns out that, by 1895, one can verify the existence of a class with number 155 named 'Chairs'—which was later abolished by the Classification Order n. 319 in 1961 for the creation of the current 297 class.

Therefore, rather than relying solely on official class creation dates, I supplement them with several historical classification documents. This is the case of the Classification Bulletin, a periodical publication by the USPTO providing systematic updates and revisions to the classification code, including descriptions of newly created classes and subclasses. Specifically, I gather the following datasets: (i) the US patent classification in 1838;[12] (ii) the editions of 1872 and 1895 of the Classification Index of Subjects of Invention (henceforth CI); (iii) the editions of 1912, 1916, 1920 and 1923, and 1947 of the Manual of Classification of Subjects of Invention (henceforth MC); (iv) the 1980 and 1987 editions of the Index to the US Patent System (henceforth IPS);[13] (v) the USPTO Official Gazette in the years 1878-1912 (not all volumes); and (vi) the Classifications Bulleting editions in the period 1912-1945 (not all volumes). Sources (i)-(iv) provide a snapshot of the complete classification code at a given point in time, and (v)-(vi) provide a flow of new class creation.

Using these documents, I will search for the first time a technology class was

[11] Available at USPTO directory.

[12] Available in US Patent and Trademark Office and US Department of Comerce (1966).

[13] Other editions for the IPS document series, although readily available, do not contain the list of existent classes. Other editions of the MC and CI documents, to the best of my knowledge, are not available in US libraries.

mentioned and use it as a proxy for its emergence. However, searching for the exact names can be misleading due to the relabeling of classes with essentially the same content (e.g., 'Chair and seats' and 'Chairs'). Searching for keywords (e.g., 'Chairs') also poses its own challenges. For example, one may not want to attribute the emergence of the technology class 'Incremental printing of symbolic information' to the be the same of an old class like 'Printing', already present in the 1872 classification code. Significant differences between the names of a current class and an older one may indicate that they represent distinct knowledge fields. To operationalize such distinction, I apply cosine similarity, a widely used text proximity measure in text analysis. Cosine similarity quantifies the similarity between two text strings, returning a value between 0 (completely dissimilar) and 1 (identical). I set an arbitrary threshold of 0.5 to determine whether an older class represents the same technology as a current one. The following enumerated list outlines the complete procedure.

Methodology As in the main text, the technologies considered here correspond to all the utility classes in the latest version of the USPC classification code.[14] Let $\vec{X}$ denote the vector containing all these technology names. To determine, for each $x \in \vec{X}$, a proxied emergence date, τ_x, we follow the steps:

1. Let τ_x^{uspto} be the official class creation date available on the USPTO webpage.

2. From the historical documents, extract a list $\vec{Y}$ with the technology classes mentioned.[15] For each of these classes $y \in \vec{Y}$, record the year of the document when they were first mentioned in the dataset. Denote this year by τ_y^{hist}.

3. Compute the cosine similarity $\rho(x, y)$ between every x and y.

4. Finally, to determine τ_x, use the USPTO official date as a default, but update with the historical documents' evidence:

$$\tau_x = \min \left\{ \tau_x^{uspto}; \left\{ \tau_y^{hist},\ \rho(x,y) \geq 0.5 \right\} \right\}$$

[14] It can be found on: https://www.uspto.gov/web/patents/classification/selectnumwithtitle.htm. Recall the USPC was discontinued in 2013, from when the latest version is.

[15] Notice that we are considering only technology classes and not all historical documents' text. Exploiting the textual richness of these documents is an interesting avenue for future research.

Results The left panel in Figure A.2.1 shows the number of new technology classes identified each year. The time series displays a clear truncation pattern, with many classes dated to the 19th century or early 20th century (1901-1920). By construction, since the first classification system was established in the 1830s, all technology classes are dated after that point. Additionally, the classification system remained virtually static until 1872, creating a period during which no new classes were introduced. In the 1870s, the USPTO increased the number of classes increased from 36 to 145, which is reflected in the spike identified in the plot. Moreover, the official establishment of a Classification Division in 1898—and the heightened activity in the following decade—significantly increased the likelihood that new technology classes are found during that period, regardless of their true, or approximate, emergence date.

In the face of this truncation, it is more informative to treat technologies already present in the earliest classification schemes as a group rather than focusing on their specific introduction dates. For this purpose, we define as old technologies those dated to the 19th century (by the procedure described above). These include, for example, the previously mentioned case of 'Chairs and seats.' The right panel of Figure A.2.1 shows the share of patents attributed to these old technologies. Consistent with the discussions in Sections 2 and 4, innovation in these technologies declines over time but remains substantial for an extended period.

I now turn to the set of technology classes whose emergence date falls within the 20th century. With the classification division already established and the baseline code largely in place, these technologies were introduced more gradually, making their identified emergence dates more likely to reflect the true periods when they gained attention. The choice of 1900 as the threshold is somewhat arbitrary. In fact, significant classification activity would still occur during the first decade of the 20th century following the division's creation. While setting a later threshold would be equally valid and would not alter the results, 1900 is used here for simplicity. In the analysis that follows, I restrict attention to patents attributed to these newer technologies, defined as those with emergence dates in the 20th century or later. Patent shares refer to this subset.

To investigate how the path of innovation correlates with technology age, taking as a reference the class creation date, I run the following event study regression:

$$PatentShare_{t,\tau} = \sum_a \beta_a 1_{\{t-\tau=a\}} + \Gamma_\tau + e_{t,\tau} \tag{A.32}$$

Here, Γ_τ's are technology fixed effects, and the β_a's are our coefficients of interest. Controlling for technology-specific factors, they represent the mean effect of age in a technology patenting share. While there are no time-fixed effects due to colinearity, the fact that we are using patent shares as the dependent variable already normalizes the observations by the aggregate trend in patents.

Figure A.2.2 presents the estimated age coefficients $\{\hat{\beta}_a\}$ along with their 95% confidence intervals. The Figure exhibits a pronounced hump-shaped pattern. On average, patenting activity within a technology begins to rise a few decades before the USPTO creates a new technology class. Remember the USPTO retroactively classifies patents, ensuring that all historical patents are consistently categorized according to the most up-to-date classification system. This retroactive classification enables observation of patent activity within a technology prior to its formal class creation.

The age effect peaks around age zero, indicating that the creation of a technology class typically coincides with the period when the technology is gaining the most attention and popularity. As the technology continues to age, its share of patenting activity gradually declines (see Section 4 for a more detailed discussion of Figure A.2.2).

Finally, still focusing on 20th-century technologies, Figure A.2.3 shows the cross-sectional patent distribution across technologies of different ages in the 2000 decade. The distribution shows two distinct peaks: at ages 20–30 and 90. These correspond to technologies whose class creation occurred in the 1910s and 1970–1980s, respectively. The larger patent shares associated with these vintages can be attributed to the higher number of new technology classes introduced during these periods (see Figure A.2.1). The latter peak reflects the emergence of many classes related to ICT technologies, while the earlier peak, which includes a wide array of technologies, likely results from the truncation effects discussed earlier. Without this truncation, the distribution would likely exhibit a more distinct hump-shaped pattern in research allocation.

B Model extensions

B.1 Exogenous death shock faced by varieties

For any given technology τ, the growth rates for $Q(t|\tau)$ and $\mu(t|\tau)$ are proportional to the innovation rate per variety, $R(t|\tau)^{1-\epsilon}/\mu(t|\tau)$ (see Equation 11). Thus, if the number of researchers working at technology τ is kept constant, the growth rates fall as $\mu(t|\tau)$ increases. The most straightforward way to offset this effect is to assume intermediate varieties may be hit by a death shock with Poisson intensity δ.

At a given point in time, the death process removes varieties whose average quality is $Q(t|\tau)$, while innovation adds new varieties whose quality is on average $Q(t|\tau)\bar{\lambda}$. As such, the case where $\delta > 0$ creates a selection process that allows the average quality to grow in the long run. If $\delta = 0$, we are back to the baseline model in Section 3, and long-run growth within a technology necessarily peters out.

Solving the model with $\delta \geq 0$

Allowing a general $\delta \geq 0$ does not change the static equilibrium allocation presented in Section 3.1. Nor does it change the research allocation given by equation (10). The equilibrium is characterized as in Definition 1, with the laws of motion for the state variables being now:

$$\partial_t \mu_t(a) = -\partial_a \mu_t(a) + R_t(a)^{1-\epsilon} - \delta \mu_t(a) \tag{B.1}$$

$$\partial_t Q_t(a) = -\partial_a Q_t(a) + Q_t(a)(\bar{\lambda} - 1)\frac{R_t(a)^{1-\epsilon}}{\mu_t(a)}. \tag{B.2}$$

Proposition 5. *If and only if*

$$\frac{\delta}{\gamma} \leq \max\left\{\frac{1-\epsilon}{\epsilon}; \frac{1}{\epsilon\bar{\lambda}}\right\} \tag{B.3}$$

there exists a Balanced Growth Path for this economy. It is unique and characterized by:

1.*Aggregate variables grow at the rate γ.*

2.Innovation shares converge to zero as technologies age:

$$\lim_{a \to \infty} R(a) = 0$$

3.The stationary technology-age distribution of quality at any given t is:

$$Q(a) = \left\{ c_1 - c_2 e^{-\left(\frac{1-\epsilon}{\epsilon}\gamma - \delta\right)a} \right\}^{\frac{1}{\frac{1}{\bar{\lambda}-1} - \frac{1-\epsilon}{\epsilon}}} \tag{B.4}$$

where c_1, c_2 are

$$c_2 \equiv \left(\frac{\frac{1}{\bar{\lambda}-1} - \frac{1-\epsilon}{\epsilon}}{\frac{1-\epsilon}{\epsilon}\gamma - \delta} \right) \frac{Q_0^{\frac{1}{\bar{\lambda}-1}}}{\mu_0} (\bar{\lambda} - 1) \bar{Q}^{-\frac{1-\epsilon}{\epsilon}}$$
$$c_1 \equiv Q_0^{\frac{1}{\bar{\lambda}-1} - \frac{1-\epsilon}{\epsilon}} + c_2$$

and $\bar{Q}$ is the stationary level of $\bar{Q}_t$, which is unique.

4.The stationary technology-age distribution of varieties at any given t is:

$$\frac{\mu(a)}{\mu_0} = e^{-\delta a} \left(\frac{Q(a)}{Q_0} \right)^{\frac{1}{\bar{\lambda}-1}}. \tag{B.5}$$

*Proof.*See Section C.1. There, we considered a general $\delta \geq 0$ and derived (B.4)-(B.5). In turn, Assumption (B.3) is needed to guarantee output is bounded for a finite t, i.e., $Y_t < \infty$ for every t. □

If δ/γ is not sufficiently small, either (i) $R(a)$ is non-decreasing for large a, and hence $\int R(a)da \to \infty$; (ii) or it decreases but not fast enough to guarantee that output stays bounded for every finite t. As we know, γ represents how better new technologies are, while δ controls how fast a technology can steadily grow in the long-run.[16] If the latter force is sufficiently larger than the former, such that (B.3) is violated, very old technologies never lose salience. They keep attracting a significant share of research, in many cases superior to the share attracted by younger technologies, which is inconsistent with a bounded equilibrium. This

[16] If a constant number of researchers work on technology τ, its quality $Q(t|\tau)$ growth rate converges in the long run to $\delta(\bar{\lambda} - 1)$.

is also inconsistent with the empirical evidence that the patent share of older technologies declines over time.

B.2 Cross-technology spillovers

This Section relaxes the assumption that knowledge spillovers happen exclusively within a technology. In particular, it considers the case where innovators working on technology τ learn not only from τ's own accumulated quality $Q(t|\tau)$, but also from the aggregate quality index $\bar{Q}_t$. The specific weighting scheme in computing the average $\bar{Q}_t$ can be changed as long $\bar{Q}_t$ is well defined.

The economy environment described in Section 3 remains unchanged, except for equation (5), which is replaced by the following assumption:

Assumption 1. *The quality q of a variety invented at t within a technology $\tau \in (-\infty, t]$ is*

$$q = \lambda Q(t|\tau)\bar{Q}_t,$$

where, as before, $\bar{Q}_t \equiv \left[\int_0^\infty \left(e^{-\gamma \tilde{a}} Q_t(\tilde{a})\right)^{\frac{1}{\epsilon}} d\tilde{a}\right]^\epsilon$ is an age-discounted aggregate quality index;

Here, λ represents the innovator's own talent, $Q(t|\tau)$ represents the within technology spillovers, and $\bar{Q}_t$ introduces cross-technology spillovers, indicating that innovators can always benefit from the average level of quality development.

The static allocation, given a state for technology $[Q(t|\tau), \mu(t|\tau)]_\tau$, is still characterized by Proposition 1. Regarding the dynamic allocation, equation (9), representing the expected value of an innovation directed to technology τ, now becomes:

$$\bar{v}(t|\tau) = \beta e^{\gamma\tau} \bar{\lambda} Q(t|\tau) L \bar{Q}_t D_t.$$

Most importantly, the allocation of research across technologies is preserved:

$$R(t|\tau) = \left(\frac{e^{-\gamma(t-\tau)} Q(t|\tau)}{\bar{Q}_t}\right)^{\frac{1}{\epsilon}} R$$

While the law of motion for the measure of varieties $\mu(t|\tau)$ remains unchanged, the same is not true for the average quality $Q(t|\tau)$ case. Equation (11) now becomes:

$$\frac{\dot{Q}(t|\tau)}{Q(t|\tau)} = (\bar{\lambda}\bar{Q}_t - 1)\frac{R(t|\tau)^{1-\epsilon}}{\mu(t|\tau)}$$

In words, the effective step size in quality accumulation is no longer $\bar{\lambda} - 1$, but rather $\bar{\lambda}\bar{Q}_t - 1$.

A stationary equilibrium exists as before and is characterized by cross-sectional technology-age distributions of varieties and quality that are time-invariant. This means that $\bar{Q}_t$ becomes a constant $\bar{Q}$, and the new step size $\bar{\lambda}\bar{Q} - 1$ is also time time-invariant. The stationary equilibrium preserves the main features of the benchmark model in Section 3 but with an effective higher $\bar{\lambda}$. The peak in the research-age distribution shifts right, with technologies being developed for longer before losing steam. In other words, cross spillovers retard the take-off of new technologies. What is crucial is that cross-knowledge spillovers are not enough to preclude a technology from becoming obsolete in the very long run. They may retard it, but the fact that new technologies are increasingly better, combined with ideas getting harder within technology, still draw researchers away from old technologies in the very long run.

Importantly, the normative analysis is also not changed. In the long run BGP, the infinitely patient planner still faces the same problem, but with an augmented $\bar{\lambda}$. The results, therefore, extend naturally.

C Proofs and Additional Derivations

C.1 Step by Step proof for Proposition 2

(In the derivations below, the case with $\delta = 0$ represents the baseline model from Section 3, while $\delta > 0$ corresponds to the extension discussed in B.1).

Using the BGP definition 2, plug (B.1) into (B.2) and differentiate with respect to time to obtain:

$$\mu'(t|\tau) = \mu(t|\tau)\left[\frac{1}{\epsilon}\frac{Q'(t|\tau)}{Q(t|\tau)} - \frac{Q''(t|\tau)}{Q'(t|\tau)} - \frac{1-\epsilon}{\epsilon}\gamma\right] \tag{C.1}$$

Use (B.1) and (B.2) to substitute for $\mu'(t|\tau)/\mu(t|\tau)$ in (C.1) and obtain:

$$\begin{aligned} &\left[\frac{1-\epsilon}{\epsilon}\gamma - \delta\right] + \left[\frac{1}{\lambda-1} - \frac{1}{\epsilon}\right]\frac{Q'(t|\tau)}{Q(t|\tau)} + \frac{Q''(t|\tau)}{Q'(t|\tau)} = 0 \\ &e^{\left(\frac{1-\epsilon}{\epsilon}\gamma-\delta\right)(t-\tau)} Q(t|\tau)^{\frac{1}{\lambda-1}-\frac{1}{\epsilon}} Q'(t|\tau) = c_\tau \end{aligned} \tag{C.2}$$

where the second line follows from integrating the first and defining

$$c_\tau \equiv Q(\tau|\tau)^{\frac{1}{\lambda-1}-\frac{1}{\epsilon}} Q'(\tau|\tau).$$

Equation (C.2) is a separable first-order differentiation equation in time t. We can hence apply the well-known solution for separable equations and obtain:

$$Q(t|\tau) = \left[\frac{\frac{1}{\lambda-1} - \frac{1-\epsilon}{\epsilon}}{\frac{1-\epsilon}{\epsilon}\gamma - \delta}\frac{Q(\tau|\tau)^{\frac{1}{\bar{\lambda}-1}}}{\mu(\tau|\tau)}(\bar{\lambda}-1)\bar{Q}^{-\frac{1-\epsilon}{\epsilon}}\left(1 - e^{-\left(\frac{1-\epsilon}{\epsilon}\gamma-\delta\right)(t-\tau)}\right) + Q(\tau|\tau)^{\frac{1}{\lambda-1}-\frac{1-\epsilon}{\epsilon}}\right]^{\frac{1}{\frac{1}{\lambda-1}-\frac{1-\epsilon}{\epsilon}}} \tag{C.3}$$

where we have used (B.2) to substitute for c_τ. With the change of variables $a = t - \tau$, and assuming that $Q(\tau|\tau) = Q_0$ and $\mu(\tau|\tau) = \mu_0$ for every τ, (C.3) gives the BGP solution for $Q(t|\tau)$ presented in (13).

Uniqueness of $\bar{Q}$

To show the uniqueness of $\bar{Q}$, assume here $\delta = 0$, as in the baseline model. The case with $\delta > 0$ is considered in Section B.1. An equilibrium value for $\bar{Q}$ must pin down $Q(a)$ in (13) such that

$$\Xi(\bar{Q}) \equiv \left[\int_0^\infty \left(e^{-\gamma a} Q(a)\right)^{\frac{1}{\epsilon}} da\right]^\epsilon$$

exactly equals $\bar{Q}$. We thus look for a fixed point in the space $\Xi(\bar{Q}) \times \bar{Q}$.

From (13), $\partial Q(a)/\partial \bar{Q} < 0$, which implies that (i) $\Xi'(\bar{Q}) < 0$. Moreover, when $\bar{Q} \to \infty$, then $Q(a) \to Q_0 \ \forall a$. Hence, (ii) $\lim_{\bar{Q}\to\infty} \Xi(\bar{Q}) = Q_0(\epsilon/\gamma)^\epsilon$. On the other hand, when $\bar{Q} = Q_0(\epsilon/\gamma)^\epsilon$, equation (13) shows that $Q(a) > Q_0 \ \forall a > 0$. Hence, (iii) $\Xi\big(Q_0(\epsilon/\gamma)^\epsilon\big) > Q_0(\epsilon/\gamma)^\epsilon$. Conditions (i)-(iii), together with the continuity of $\Xi(\bar{Q})$, guarantee a unique fix point in the space $\Xi(a) \times \bar{Q}$. This fixed point is greater than $Q_0(\epsilon/\gamma)^\epsilon$.

C.2 Derivations of the laws motion for $\mu(t|\tau)$, $Q(t|\tau)$, $\mu_t(a)$, $Q_t(a)$

Fix a technology τ and define $F_t^\tau(q) \equiv F(\tau, q)$, where $F(\tau, q)$ is the cdf of the density f (recall that it is not a probability density, but rather integrates to the total mass of varieties). Using the innovation possibilities frontier of the economy described in Section 3:

$$F_{t+\Delta}^\tau(q) = F_t^\tau(q) + \Delta R(t|\tau)^{1-\epsilon} H\left(\frac{q}{Q(t|\tau)}\right),$$

where $H(\cdot)$ is the distribution from which λ is drawn at the time of an innovation. From (2), one can write:

$$\begin{aligned}\mu(t+\Delta|\tau) &= \int dF_{t+\Delta}^\tau(q) = \int dF_t^\tau(q) + \Delta R(t|\tau)^{1-\epsilon} \int dH\left(\frac{q}{Q(t|\tau)}\right) \\ &= \mu(t|\tau) + \Delta R(t|\tau)^{1-\epsilon},\end{aligned}$$

where the last line used the fact that $H(\cdot)$ is a probability density and integrates to 1. Subtracting $\mu(t|\tau)$ from both sides, dividing by Δ and taking the limit $\Delta \to 0$, gives:

$$\mu'(t|\tau) = R(t|\tau)^{1-\epsilon}. \tag{C.4}$$

From (3), one can write:

$$\begin{aligned}Q(t+\Delta|\tau) &= \frac{\int q dF_{t+\Delta}^\tau(q)}{\mu(t+\Delta|\tau)} = \frac{\int q dF_t^\tau(q) + \Delta R(t|\tau)^{1-\epsilon} \int q dH\left(\frac{q}{Q(t|\tau)}\right)}{\mu(t+\Delta|\tau)} \\ &= Q(t|\tau)\frac{\mu(t|\tau)}{\mu(t+\Delta|\tau)} + \frac{\Delta R(t|\tau)^{1-\epsilon} Q(t|\tau)\bar{\lambda}}{\mu(t+\Delta|\tau)}\end{aligned}$$

Subtracting $Q(t|\tau)$ from both sides, taking the limit $\Delta \to 0$, and using (C.4), gives:

$$Q'(t|\tau) = \frac{Q(t|\tau)}{\mu(t|\tau)}(\bar{\lambda} - 1)R(t|\tau)^{1-\epsilon} \tag{C.5}$$

It remains to derive the laws of motion for $Q_t(a)$ and $\mu_t(a)$, which consider that time also affects the age of technologies. For brevity, as the cases of $Q_t(a)$ and $\mu_t(a)$ involve exactly the same steps, I will show below only the former. Recall that, by definition, $Q_t(a) = Q(t|t-a)$. Therefore:

$$\frac{\partial Q_t(a)}{\partial t} = Q_1(t|t-a) + Q_2(t|t-a) \quad \text{and} \quad \frac{\partial Q_t(a)}{\partial a} = -Q_2(t|t-a)$$

Hence:

$$\begin{aligned}\frac{\partial Q_t(a)}{\partial t} &= -\frac{\partial Q_t(a)}{\partial a} + Q_1(t|t-a) \\ &= -\frac{\partial Q_t(a)}{\partial a} + \frac{Q_t(a)}{\mu_t(a)}(\bar{\lambda}-1)R_t(a)^{1-\epsilon}.\end{aligned}$$

C.3 Proofs of Proposition 3 and 4

From Corollary 2, maximizing the household present discounted utility, as of time t_0, is equivalent to maximizing $\hat{C}$. Using the resource constraint, this can be written as:

$$\underset{[R(a)]_{a>0}}{\text{Max}} \; \hat{C} = \kappa \int_0^\infty e^{-\gamma a} Q(a)\mu(a)da \quad \text{s.t.} \tag{C.6}$$

$$Q'(a) = Q(a)(\bar{\lambda}-1)\frac{R(a)^{1-\epsilon}}{\mu(a)} \tag{C.7}$$

$$\mu'(a) = R(a)^{1-\epsilon} \tag{C.8}$$

$$1 = \int_0^\infty R(a)da \tag{C.9}$$

where $Q(0) = Q_0$ and $\mu(0) = \mu_0$. The presence of the constraint (C.9) can be treated by introducing a new state variable and solving the modified maximization problem, as common in the solution to isoperimetric problems. Define $\Gamma(a) \equiv -\int_0^a R(a')da'$. Then:

$$\lim_{a\to\infty} \Gamma(a) = -1 \quad \text{and } \Gamma'(a) = -R(a) \tag{C.10}$$

The problem at hand consists of maximizing the objective function below (which follows from integrating by parts (C.6)):

$$\kappa \int_0^\infty e^{-\gamma a} Q(a) R(a)^{1-\epsilon} da \tag{C.11}$$

A collection of functions $[R(a), Q(a), \mu(a), \Gamma(a)]$ is admissible in this problem if the integral in (C.11) converges and the following conditions are satisfied:

$$Q'(a) = Q(a)(\bar{\lambda} - 1)\frac{R(a)^{1-\epsilon}}{\mu(a)} \quad \text{(C.12)}$$

$$\mu'(a) = R(a)^{1-\epsilon} \quad \text{(C.13)}$$

$$\Gamma'(a) = -R(a) \quad \text{(C.14)}$$

$$Q(0) = Q_0;\ \mu(0) = \mu_0;\ \Gamma(0) = 0 \quad \text{(Initial conditions)}$$

$$\lim_{a\to\infty} \Gamma(a) = -1;\ \lim_{a\to\infty} Q(a) \geq Q_0 \text{ and } \lim_{a\to\infty} \mu(a) \geq \mu_0 \quad \text{(Terminal conditions)}$$

The terminal conditions above use the fact that the limits for $Q(a)$ and $\mu(a)$ exist and are finite within the set of admissible functions. This is not ex-ante trivial (notice how it rules out divergence to ∞ or oscillating trajectories) and is proved in the following Lemma.

Lemma 1. *(Convergence of μ and Q) $\lim_{a\to\infty} \mu(a)$ and $\lim_{a\to\infty} Q(a)$ exist and are finite.*

Proof. From (C.13):

$$\lim_{a\to\infty} \mu(a) = \mu_0 e^{\int_0^\infty g_\mu(s)ds} \quad \text{(C.15)}$$

where $g_\mu(s) = \frac{R(s)^{1-\epsilon}}{\mu(s)}$. Let $f(s) = \frac{R(s)}{\mu_0}$ and notice that $\int_0^\infty f(s)ds = 1/\mu_0$. Since $\int f(s)ds$ converges and $0 \leq g_\mu(s) \leq f(s)$ for every s, it follows that $\int_0^\infty g_\mu(s)ds$ also converges, and hence, from (C.15), one can conclude that $\lim_{a\to\infty} \mu(a)$ exists and is finite.

Finally, notice that $g_Q(a) = g_\mu(a)(\bar{\lambda} - 1)$. Then, $\int_0^\infty g_Q(s)ds$ converges, $\lim_{a\to\infty} Q(a)$ exists and is finite. □

The problem of finding an admissible collection of functions $[R(a), Q(a), \mu(a), \Gamma(a)]$ that maximizes (C.11) lends itself to the Pontryagin's maximum principle (see Seierstad and Sydsæter (1987)). Define the Hamiltonian of the problem as:

$$H = e^{-\gamma a} Q(a)R(a)^{1-\epsilon}\kappa + \pi_Q(a)Q(a)(\bar{\lambda} - 1)\frac{R(a)^{1-\epsilon}}{\mu(a)} + \pi_\mu(a)R(a)^{1-\epsilon} - \pi_\Gamma(a)R(a) \quad \text{(C.16)}$$

where $\pi_Q(a)$ and $\pi_\mu(a)$ are the costate functions, representing the shadow values of relaxing constraints (C.12) and (C.13), respectively.

An admissible solution for $[R(a), Q(a), \mu(a), \Gamma(a)]$ satisfies the necessary conditions:

$$R(a)^{\epsilon}\pi_{\Gamma} = (1-\epsilon)\left(e^{-\gamma a}Q(a)\kappa + \pi_Q(a)\frac{Q(a)}{\mu(a)}(\bar{\lambda}-1) + \pi_{\mu}(a)\right) \tag{C.17}$$

$$\frac{d\pi_Q(a)}{da} = -\frac{\beta L\bar{\lambda}}{\delta+\gamma}\frac{\kappa}{1-\epsilon}e^{-\gamma a}R(a)^{1-\epsilon} - \pi_Q(a)(\bar{\lambda}-1)\frac{R(a)^{1-\epsilon}}{\mu(a)} \tag{C.18}$$

$$\frac{d\pi_{\mu}(a)}{da} = \pi_Q(a)\frac{Q(a)}{\mu(a)}(\bar{\lambda}-1)\frac{R(a)^{1-\epsilon}}{\mu(a)} \tag{C.19}$$

Notice that $\pi_{\Gamma}(a) = \pi_{\Gamma}$ is a constant since $\partial_{\Gamma}H = 0$. The terminal condition for $\Gamma(a)$, corresponding to the scientist's resource constraint (C.14), allows us to substitute for π_{Γ} when necessary. Moreover, the problem meets the conditions from Michel (1982) and Seierstad and Sydsæter (1987) for an additional necessary (transversality) condition to be established:

$$\lim_{a\to\infty} H(a) = 0 \tag{C.20}$$

This leads to the following Lemma - which, beyond providing additional information on the optimal solution, will be useful later.

Lemma 2. *(Innovation in old technologies dies in the long-run) In a solution, it must be that:*

$$\lim_{a\to\infty} R(a) = 0$$

*Proof.*The result follows from combining (C.20) and (C.17):

$$\lim_{a\to\infty} H(a) = 0$$

$$\lim_{a\to\infty} R(a)^{1-\epsilon}\left(e^{-\gamma a}Q(a)\kappa + \pi_Q(a)\frac{Q(a)}{\mu(a)}(\bar{\lambda}-1) + \pi_{\mu}(a)\right) - \phi_{\Gamma}R(a) = 0$$

$$\lim_{a\to\infty} \frac{\epsilon}{1-\epsilon}R(a) = 0$$

□

At this point, I can prove:

Proof.(Decentralized equilibrium is efficient if $\bar{\lambda} = 1$) If $\bar{\lambda} = 1$, the constraint (C.12) in the maximazation problem discussed above disappears, and $Q(a) = Q(0)$ for every

a. Without the term $\pi_Q(a)Q(a)(\bar{\lambda}-1)\frac{R(a)^{1-\epsilon}}{\mu(a)}$ in the Hamiltonian (C.16), it turns out that the necessary condition (C.19) becomes now:

$$\frac{d\pi_\mu(a)}{da} = 0$$

which implies $\pi_\mu(a) = K$, where K is a constant. Then, the transversality condition (C.20), together with Lemmas 1 and 2 imposes $K = 0$.

The first order condition (C.17) can thus be rewritten as:

$$R(a)^\epsilon = \frac{e^{-\gamma a}Q(a)}{\bar{Q}}$$

where $\bar{Q} = \int_0^\infty e^{-\gamma a'}Q(a')da'$. For any given age profile $[Q(a), \mu(a)]$, this corresponds to the decentralized equilibrium allocation of innovation for $R(a)$. As $Q(0) = Q_0$ and $\mu(0) = \mu_0$ in the decentralized and in the optimal path, the whole age profile $[Q(a), \mu(a), R(a)]$ will coincide in both solutions. It is also possible to explicitly solve for $\bar{Q}$ in such a solution given that $Q(a) = Q_0$ for every a. In fact, $R(a) = \gamma e^{-\gamma a}$. □

I now move forward to finish proving all the statements contained in the remaining propositions.

Lemma 3. *The limit* $\lim_{a\to\infty} \pi_Q(a)$ *exists, is finite and is non-negative.*

*Proof.*The coestate variable $\pi_Q(a)$ has a continuous path and satisfies the necessary condition (C.18), a linear first-order differential equation. As such, it can be expressed as:

$$\pi_Q(a)Q(a) = \left(K - \int_0^a \kappa e^{-\gamma s}Q(s)R(s)^{1-\epsilon}ds\right) \tag{C.21}$$

where K is the constant of integration. Hence:

$$\lim_{a\to\infty} \pi_Q(a) = \frac{1}{Q^*}\left(K - \int_0^\infty \kappa e^{-\gamma s}Q(s)R(s)^{1-\epsilon}ds\right) \leq \infty \tag{C.22}$$

where $Q^* \equiv \lim_{a\to\infty} Q(a)$, which exists and is finite by Lemma 1 - and the inequality follows from the fact that the integral above converges for all admissible functions:

notice how it equals the objective function (C.11) of the problem, which in turn equals the consumption level (up to a constant of difference) and converges.

To prove the last claim, let the optimized value of the problem with initial conditions $Q(a)$ and $\mu(a)$, objective function $\int_a^\infty e^{-\gamma a'} Q(a')R(a')^{1-\epsilon}da'$, and the constraints (C.12)-(C.14) be denoted by:

$$V(a, Q(a), \mu(a)) = \kappa \int_a^\infty e^{-\gamma a'} \hat{Q}(a')\hat{R}(a')^{1-\epsilon}da'$$

where $\hat{x}(\cdot)$ denotes the optimal choice for the function $x(\cdot)$ among all feasible alternatives. The exact same steps used in the proof of Theorem 7.9 in Acemoglu (2009) and omitted here allow one to write:

$$\pi_Q(a) = \frac{\partial V(Q(a), \mu(a))}{\partial Q}$$

In words, $\pi_Q(a)$ measures the value of a marginal increase in the state $Q(a)$, of a relaxation in the constraint (C.12), the shadow value of $Q(a)$ (Acemoglu, 2009). I claim that:

$$\frac{V(a, Q(a), \mu(a))}{\partial Q} \geq 0 \quad \forall a \geq 0$$

Let $[\hat{R}(a')]_{a'}$ be the optimal choice of research allocation when the initial conditions are Q and μ. For an arbitrary $\epsilon > 0$, consider the problem with initial conditions $Q + \epsilon$ and μ. As the supply of scientists does not change with a marginal increase in the initial condition, $[\hat{R}(a')]_{a'}$ is still a feasible allocation. Observe that, by choosing it, since the initial condition for μ is unchanged, (C.12)-(C.13) show that the resulting sequence $[Q(a')]_a'$ under the new scenario must satisfy $Q(a') \geq \hat{Q}(a')$ for every a. It follows that:

$$V(a, Q + \epsilon, \mu) \geq \kappa \int_a^\infty e^{-\gamma a'} \hat{Q}(a')\hat{R}(a')^{1-\epsilon}da' = V(a, Q, \mu)$$

Rearranging terms and taking $\epsilon \to 0$, gives $\partial_Q V(a, Q, \mu) \geq 0$ for every a, as claimed. In turn, this implies that $\pi_Q(a) \geq 0$ for every $a \geq 0$, and hence it must be that $\lim_{a\to\infty} \pi_Q(a) \geq 0$.

Finally, notice that □

Lemma 4. *The limit* $\lim_{a\to\infty} \pi_\mu(a)$ *exists, is finite and is non-positive.*

*Proof.*Lemmas 2 together with the first order condition (C.17) imply that:

$$\lim_{a\to\infty} \pi_Q(a)(\bar{\lambda}-1)\frac{Q(a)}{\mu(a)} + \pi_\mu(a) = 0$$

In turn, Lemmas 1 and 3 imply that:

$$\lim_{a\to\infty} \pi_Q(a)(\bar{\lambda}-1)\frac{Q(a)}{\mu(a)} = G$$

where G is a constant satisfying $0 \leq G < \infty$.

It follows that, necessarily:

$$\lim_{a\to\infty} \pi_\mu(a) = -G \leq 0 \tag{C.23}$$

□

From (C.21) and (C.22):

$$\begin{aligned}
&\lim_{a\to\infty} \pi_Q(a) = \frac{1}{Q^*}\left(K - \int_0^\infty \kappa e^{-\gamma s} Q(s) R(s)^{1-\epsilon} ds\right) \leq \infty \\
&K = Q^* \pi_Q^* + \int_0^\infty \kappa e^{-\gamma s} Q(s) R(s)^{1-\epsilon} ds \\
&\pi_Q(a) Q(a) = \left(K - \int_0^a \kappa e^{-\gamma s} Q(s) R(s)^{1-\epsilon} ds\right) \\
&\pi_Q(a) Q(a) = e^{-\gamma a} \int_a^\infty \kappa e^{-\gamma(s-a)} Q(s) R(s)^{1-\epsilon} ds + cte
\end{aligned}$$

where $cte \geq 0$ is a non-negative constant.

Hence, using the definition $\mathcal{E}_1(a) = \pi_Q(a)\frac{Q(a)}{\mu(a)}(\bar{\lambda}-1)$, it is straightforward to verify that $\mathcal{E}_1(a) \geq 0$ and $\mathcal{E}_1'(a) \leq 0$.

Finally, notice from (C.19) and from Lemma 4 that

$$\pi_\mu(a) = -\int_a^\infty \mathcal{E}_1(s) g_\mu(s) ds - cte2$$

where $cte2 \geq 0$ is a non-negative constant. It is straightforward to see that $\mathcal{E}_2(a) \leq 0$ and $\mathcal{E}_2'(a) \geq 0$ (the last condition can also be directly verified from (C.19) with the results derived for $\mathcal{E}_1(a)$.

One last final point concerning Proposition 4 is now proved: namely, the fact that the optimal allocation of research features $R'(a) < 0$. Differentiating (C.17) with respect to a and using (C.18)-(C.19), the following expression can be derived:

$$R(a)^{\epsilon} \propto -\frac{e^{-\gamma a} Q(a)}{\epsilon g_R(a)}$$

where $g_R(a) \equiv \dot{R}(a)/R(a)$. Hence, as $0 \leq R(a) < \infty$ for every $a > 0$, we must have $g_R(a) < 0$.

D Calibration

D.1 Additional details: ϵ calibration

Table D.1.1 presents the results from regression (29) estimation. In Column (1), the estimated coefficient on $\log \bar{v}(t|\tau)$ is statically significant and equal to 0.381. In words, on average, a one percentage point increase in technology τ new patents' market valuation (in addition to the overall stock market trend) is associated with an extra increase of 0.381 percentage points in patents related to τ (in addition to the aggregate increase in patenting). This implies a value for ϵ equal to 0.72. Column (2) reports results when the explanatory variable is lagged by one period (decade), examining the relationship between $\log Patents(t|\tau)$ and $\bar{v}(t-1|\tau)$. While the model assumes an instantaneous reallocation of research efforts and immediate outcomes in terms of patenting, real-world adjustments in research direction and the completion of projects typically require time. Consistent with this lagged adjustment, Column (2) finds a stronger comovement between patenting intensity and past market valuation. The corresponding implied value for ϵ is 0.59. The baseline calibration adopts an intermediate value, $\epsilon = 0.65$, while all quantitative exercises report results across a range of lower and higher values for robustness.

D.2 Additional details: $\bar{\lambda}$ calibration

First, notice that, when using equation (30), we have fixed t to compute the quality age profile at a given point in time. We can, however, exploit the longitudinal variation to obtain more precise estimates. Namely, use (30) in different points in

Table D.1.1: Estimates for the Research Supply Elasticity ϵ

	New Patents: $\log \text{Patents}(t\|\tau)$	
	(1)	(2)
Avg Patent Value: $\log \bar{v}(t\|\tau)$	0.381	0.690
	(0.190)	(0.170)
	$\epsilon = 0.724$	$\epsilon = 0.591$
Year FE	✓	✓
Tech FE	✓	✓
N	205	187
R^2	0.024	0.101
F-Stat	4.028	16.44

Note: This table presents the fixed effects (within) estimates of regression (29). Time periods correspond to decades and technology vintages represent technology classes grouped by proxied age (see Section 4). The data used corresponds to the more conservative set of technologies whose logistic trend fit (used to establish emergence dates) achieved an R square value greater than 0.5, excluding outliers with emergence dates beyond 2023. Column (1) uses patent valuations (the explanatory variable) from period t, while Column (2) uses their one-period lag.

time to compute $Q_t(a)/Q_t(0)$ and regress it on time and age-fixed effects:

$$\log\left(Q_t(a)/Q_t(0)\right) = \Gamma_t + \Gamma_a + e_{t,a}. \tag{D.1}$$

We are interested in the age-fixed effect estimates $\{\Gamma_a\}_{a\geq 0}$. They represent the average, over time, of the schedule $\{Q(a)/Q(0)\}$, after filtering out time effects and idiosyncratic noise. Figure E.4 in Section E plots the estimated results. Reassuringly, and without it being imposed by the analysis, the age profile $\{\Gamma_a\}_{a\geq 0}$ is upward sloping, as is $Q(a)$ in the theory when $\bar{\lambda} > 1$. This pattern holds for different definitions of technology and measures of age, as shown in the Figure's two different panels.

Intuitively, combining equations (30) and (D.1), an upward-sloping profile for $\{\Gamma_a\}_a$ indicates a positive stock market premium for older technologies relative to younger vintages (after controlling for the inherent productivity advantage of the latter). In the model, this can be rationalized if $\bar{\lambda}$ is sufficiently high. By representing the building on the shoulder of giants effect, $\bar{\lambda}$ introduces a premium for building on older technologies, which is reflected on the stock market of patents, after controlling for γ.

The final step, therefore, is to specify the exact choice of $\bar{\lambda}$ given the age-fixed effects $\{\Gamma_a\}_a$. While different alternatives are possible, the baseline calibration chooses its value to minimize the sum of squared deviations between the model's steady quality age profile and its average empirical counterpart measured above:

$$\min_{\bar{\lambda}} \sum_a \left[\Gamma_a - \left(\frac{Q(a)}{Q(0)} \right) \right]^2 ,$$

where all remaining parameters are kept at their baseline calibration levels, as discussed in this Section. The resulting value for $\bar{\lambda}$, reported in Table 2, is of 2.34. All quantitative exercises performed later will use a range of $\bar{\lambda}$ values above and below its baseline calibration.

E Additional Figures and Tables

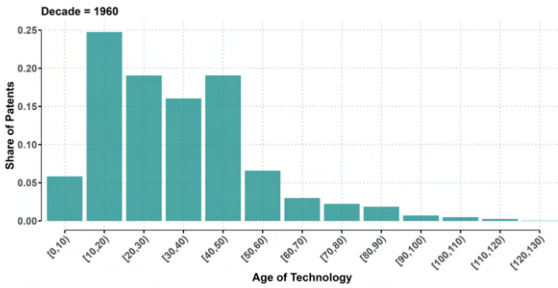

(a)

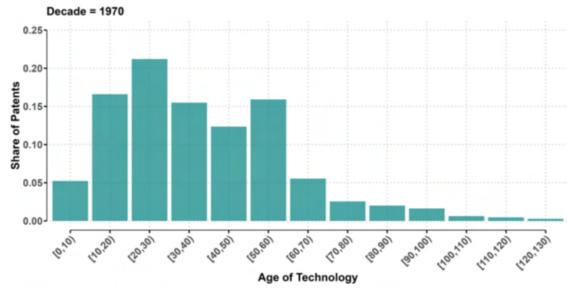

(b)

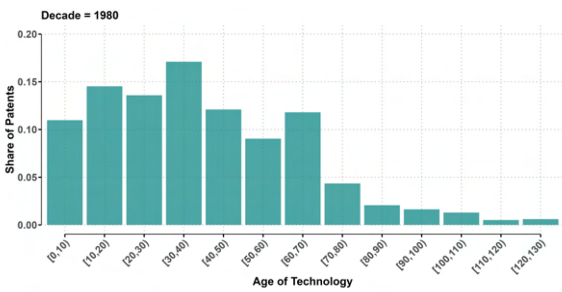

(c)

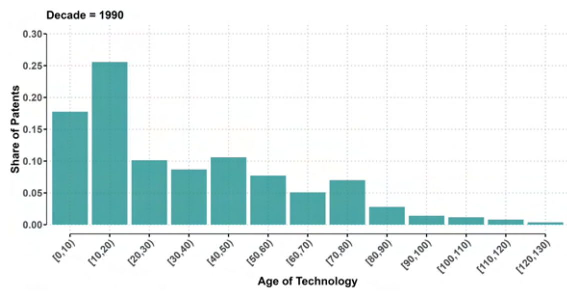

(d)

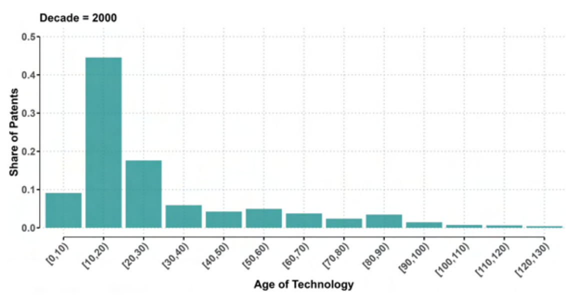

(e)

Figure A.1.1: The Hump-Shaped Innovation age Profile (Alternative Technology Definition)

Note: This Figure plots the share of patents (within a given decade) building on technologies of different ages. Technologies are defined by the breakthrough patents identified by Kelly et al. (2021). To determine the specific technology a given patent builds on, I use the citation network (see the text of this session for details).

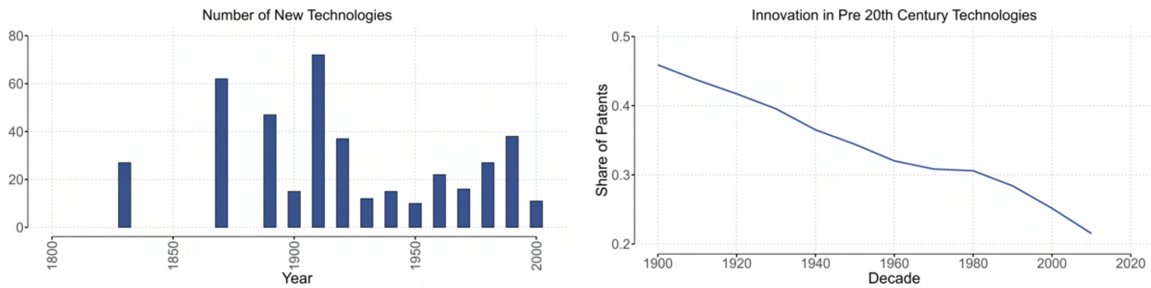

Figure A.2.1: Technology Classes: Textually Identified Emergence Dates

Note: The left plot shows the number of new technologies identified in USPTO classification documents every year. Technologies correspond to the technological classes present in the latest version of the USPC classification code. The documents, data, and procedure to identify their first appearance is described in the text of this Section. The right plot shows the evolution, starting in 1900, of the share of patents attributed to old technologies. Old technologies are defined as those present in the USPTO classification scheme by 1900.

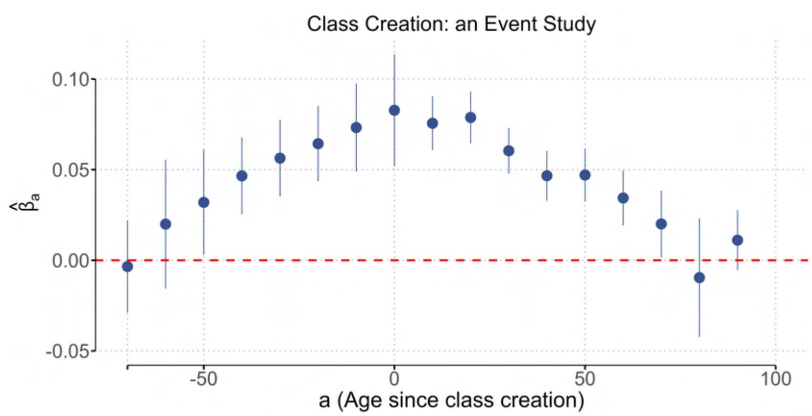

Figure A.2.2: Technology class and innovation

Note: The vertical axis displays the estimated coefficients for age dummies in regression (A.32), along with their 95% confidence intervals. The data includes 20th-century technologies and starts in 1900. Standard errors are computed using a heteroscedasticity-consistent covariance matrix. Observations are trimmed at the bottom and top 1% to remove outliers.

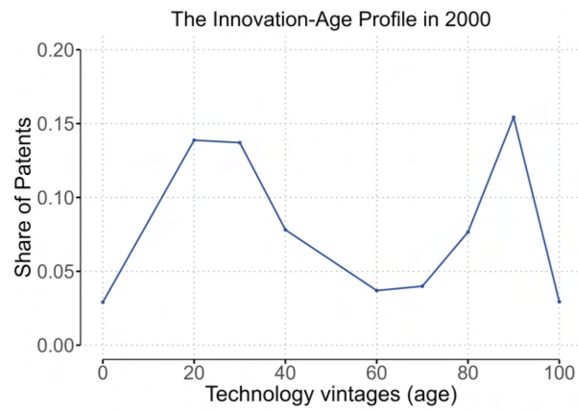

Figure A.2.3: Cross-sectional patent distribution across technology classes

Note: The Figure presents the share of patents in the 2000 decade accrued by technologies of different ages—where age is measured with respect to the technology class creation date (see the text of this Section). Each dot bins technologies with the same estimated age. The sample of technologies comprehends the ones whose class creation was dated in the 20th century. Observations are trimmed at the bottom and top 1% to remove outliers.

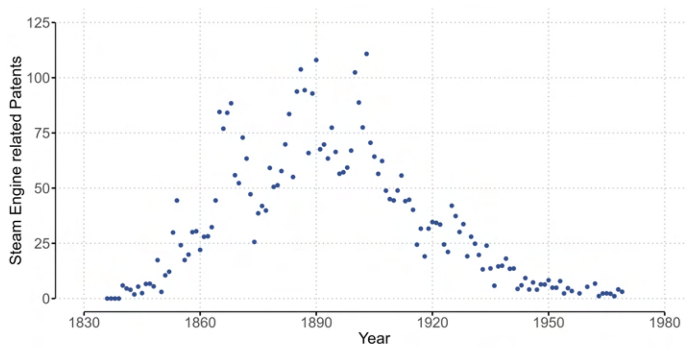

Figure E.1: The Innovation Life Cycle of the Steam Engine (Quality Adjusted Patents)

Note: This Figure displays the adjusted annual count of patents containing keywords related to the Steam Engine technology. This adjustment is made by weighting each patent with its respective importance index (*kpst_5*) from Kelly et al. (2021).

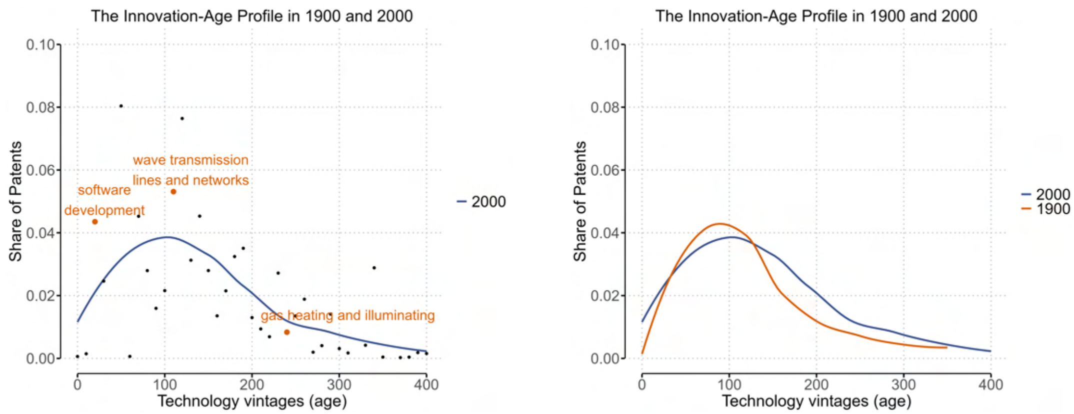

Figure E.2: The Hump-Shaped Innovation Age Profile (Quality Adjusted Patents)

Note: This Figure replicates Figure 2 in the main text, but weighting patents by their quality/importance index from Kelly et al. (2021). The graphs show the distribution of quality-weighted patents across technologies of different ages. See Section 4 for more details

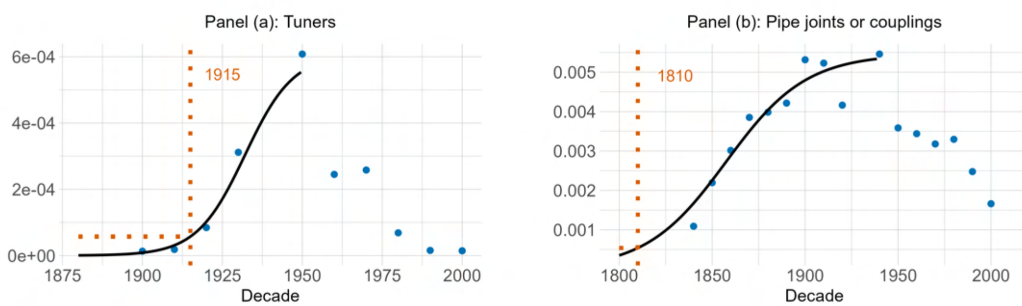

Figure E.3: Patenting share and proxied emergence date

Note: The dots in each Figure represent the share of patents accrued by the respective technology class among all utility patents issued in a given decade. For each technology, observations below the 5th and above the 95th percent data were removed to avoid the influence of outliers. The solid line represents the fitted logistic trend whose estimation is described in Section 4. The dashed lines represent the year in which the trend achieved 10% of the ceiling value.

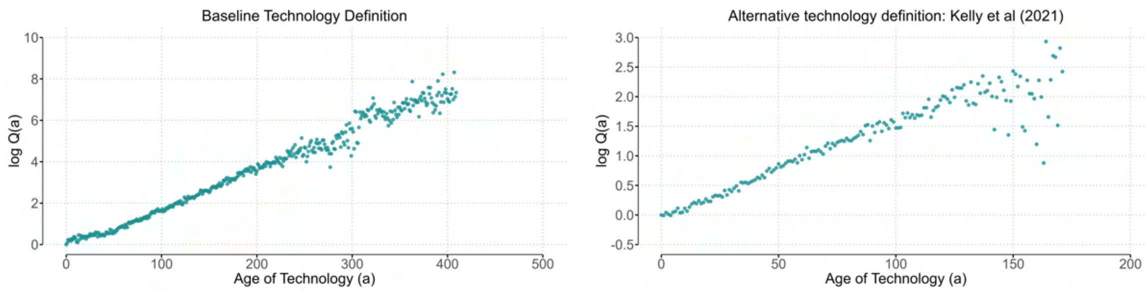

Figure E.4: Inferred $Q(a)$ from Patent Valuation Data

Note: This Figure presents the age-fixed effects obtained when estimating regression (D.1). To run this regression, we need to first sort patents into technologies of different ages. The left plot shows results when we use the baseline technology and age definitions (see Section 4. The right plot shows the results when we use the technology and age measures presented in Section A.1).

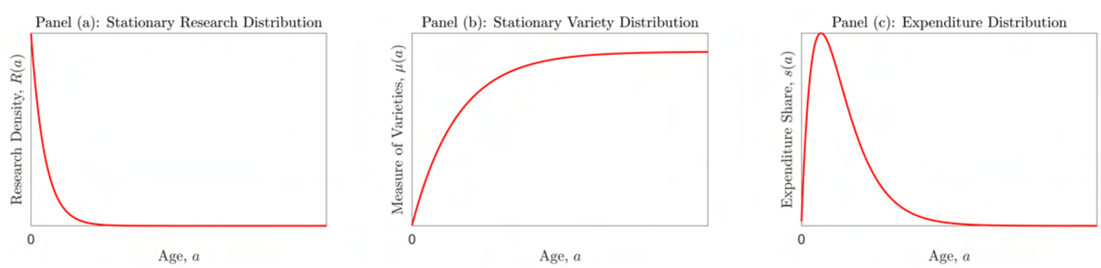

Figure E.5: The cross-sectional stationary distributions with $\bar{\lambda} = 1$

Note: The figure presents the illustrative calibration discussed in Section 3.

F Logistic trend estimation

This Section provides additional details on the logistic trend estimation in Section 4.

Use a logistic function to model a technology's patent share evolution:

$$pat\_share(t) = \frac{K}{1 + e^{-(\alpha+\beta t)}} \tag{F.1}$$

Here, K is the ceiling value (maximum share achieved by a technology), β is the growth rate coefficient, and α is a location parameter. The logistic equation can be rewritten as

$$\log\left(\frac{pat\_share(t)}{K - pat\_share(t)}\right) = \alpha + \beta t.$$

The left-hand side variables are observed in the data, while the RHS contains a constant and a time trend. We can thus use an OLS regression to find the values of $\hat{\alpha}$ and $\hat{\beta}$ that best fit the data. This is done individually for each technology.

The proxied emergence date of a technology corresponds to the time t^* the logistic trend reaches 10% of its peak. Plugging $pat\_share(t^*) = 0.1K$ and the estimated values for $\hat{\alpha}$ and $\hat{\beta}$ into (F.1), one can solve for t^*.