

Growth with New and Old Technologies

Bernardo S. C. Ribeiro*

EIEF

August 27, 2026

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 generate welfare gains of 7%.

Keywords: Growth; Emerging technologies; R&D misallocation

*bernardo.ribeiro@eief.it

I am deeply grateful to Samuel Kortum, Michael Peters, and Fabrizio Zilibotti for their invaluable guidance. I also thank Mohamad Adhami, Pedro Casavilca, Zhang Chen, Joel Flynn, Hugo Hopenhayn, Zhen Huo, Giammario Impullitti, Danial Lashkari, Ezra Oberfield, Ryungha Oh, Karthik Sastry, Todd Schoellman, Kjetil Storesletten, Weiliang Tan, Stephan Waizmann, Conor Walsh, Kai-Jie Wu, Emily Yankowitz and many seminar participants for helpful comments. Gabriele Licata and Liangnan Wang provided outstanding research assistance.

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 obsolescence, 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 disproportionately more other patents within the same class than patents in different classes.¹ This means that a technology should be understood

¹On 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 at 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.³ 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

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.

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

³Section 4 extends the baseline model to endogenize the arrival of technologies.

as patterns of magnetic fields on the tape's surface. This approach was fundamentally 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-the-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 (BGP).

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 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 quantitative exercise 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 to 20 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 61%. The associated welfare gains amount to approximately 7% in consumption-equivalent terms. A range of alternative calibrations delivers similarly significant long-run productivity and welfare gains.

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](#); [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 of [Chari and Hopenhayn 1991](#) and more recently [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 ([Acemoglu and Zilibotti, 2001](#); [Acemoglu et al., 2012, 2016](#); [Hopenhayn and Squintani, 2021](#); [Donald, 2024](#)).⁴ In this literature, the set of research lines is typically fixed and, in most cases, limited to two technologies.⁵ Complementing this work, the theory presented here features new technologies continuously emerging at the frontier, with an ever-expanding state space over time. This allows the paper to characterize innovation across endogenous cycles of technological rise and obsolescence and to study the misallocation of R&D across technologies of different ages.

Drawing on theories of General Purpose Technology (GPT) as conceptualized by [Helpman and Trajtenberg \(1994, 1996\)](#), this paper models a technology as a platform for the creation of new varieties or applications over time. However, [Helpman and Trajtenberg \(1994, 1996\)](#) focus on innovation happening only on the newest GPT, to which all research resources are allocated. Moreover, their framework features no knowledge spillovers within a technology. In contrast, this paper demonstrates that incorporating such 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.

This paper also relates to [Akcigit et al. \(2020\)](#), who distinguish basic from applied research. In their model these are intrinsically different types of innovation: basic research generates spillovers that transcend the target industry, raises the efficiency of subsequent applied research, and may not be immediately translated into consumption goods. Here, by contrast, innovations do not differ intrinsically across technologies, since the spillover parameters and the research production function are common to all. R&D misallocation nonetheless arises, driven by technologies' life-cycle position rather than by any difference in the nature of

⁴Most closely related, [Aghion et al. \(2026\)](#) study how financial frictions can slow technological transitions in a Schumpeterian model with a modern and an old vintage.

⁵An important exception is [Hopenhayn and Squintani \(2021\)](#), who consider a continuum of research lines.




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. Patent textual data are sourced from the Google Patents database ([Google, 2023](#)); see Section C.3 in the Supplementary Material.

innovation. The two perspectives are thus complementary.

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 applica-

ble on a broader scale? To address this question, I now turn to analyze nearly eight 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 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 apply 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.⁶ While a comprehensive exposition of this measurement procedure awaits in Section 5, here I present its main findings, which will shed light on the theory presented next in Section 3.

The left plot of Figure 2 reveals the first key empirical result of this paper: the cross-sectional distribution of patents across technologies of different ages has a pronounced hump shape. Each dot represents a vintage—the group of technology classes that share the same estimated age, computed following [Griliches \(1957\)](#). The horizontal axis shows this age, while the vertical axis shows the share of all patents issued during 2000–2009 that accrued to each vintage. A technology like ‘Software development,’ for instance, emerged in the late 20th century and falls in the orange dot on the left of the graph, together with a few other technologies that were among the newest in the early 2000s. These new technologies received significant innovation, but not as much as mid-aged ones that emerged 50 to 100 years earlier. The latter include examples like ‘Wave transmission lines and networks,’ which greatly facilitated telephonic and telegraphic communication, or ‘Drug, bio-affecting

⁶See [Achilladelis et al. \(1990\)](#), [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 2000–2009 across technology classes of different ages. The USPTO classifies each patent into one and only one principal technology class (USPC). Their age is estimated in accordance with [Griliches \(1957\)](#) methodology (see Section 5 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 1900–1909 (orange). Residual vintages outside [0,400] are omitted, accounting for less than 0.5% of patents in 2000–2009. The underlying data come from the USPTO Historical Patent Data Files ([Marco et al., 2015](#)); see Section 5.

and body treating compositions.’ Technologies linked to old classes, such as ‘Gas heating and illumination,’ are binned in the dots to the far right. They too receive significant innovation, in some cases comparable to very new technologies, but not as much as the mid-aged ones.

The right plot of Figure 2 contains a second key result: the hump shape of patenting is stable throughout the 20th century. Like the left plot, this graph shows the technology-age distribution of patents in the 2000s, but it adds the 1900–1909 decade for comparison. The hump is evident in both periods, and the two distributions are remarkably similar: in 1900, as in 2000, most patents were issued in technology classes about 50–100 years old. What changes is the set of classes occupying that peak. In 1900, many were tied to electricity, such as ‘Electrical generator or motor structure.’ The technologies that would dominate patenting in 2000, such as ‘Wave transmission lines and networks,’ were only just emerging around 1900: they already drew a positive but small share of patents.

Alternative measures of technology and age While the USPC classes satisfy the key defining property of technology used in this paper’s theory,⁷ they also have limitations. Technology class boundaries may reflect the patent office’s administrative goal of organizing inventions rather than a sharp technological distinction (USPTO, 2012). As a result, some classes span broad technical domains, with no clear economic interpretation.

Section B in the Supplementary Material revisits Figure 2 using alternative technology definitions and alternative procedures to proxy emergence dates.⁸ Across these exercises, the central qualitative patterns remain: the innovation-age profile is hump-shaped, most innovation is directed to technologies that are several decades old, and innovation in old technologies remains substantial. The exact location of the peak is more sensitive and varies with the granularity of the technology definition, ranging from roughly 20–40 years for narrower classifications (e.g., subclasses) to roughly 50–100 years for broader ones.

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 a 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 :

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

⁷See the Introduction for the definition and supporting evidence.

⁸These include, among others, more granular CPC subclasses and subgroups, breakthrough patents and their citation trees, the first patent granted in a (sub)class as an emergence proxy, and USPTO historical documents.

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: 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's 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

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

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 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:

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

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

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

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

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

Summarizing the extensive and intensive margins of each technology's develop-

ment, $\mu(t|\tau)$ and $Q(t|\tau)$ serve as our key state variables. 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.⁹ 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,

$$A_\tau = e^{\gamma\tau}, \quad (4)$$

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:

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

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

⁹A model extension in Section 4 introduces population growth and basic research to provide a microfoundation for the emergence of new technologies and their intrinsic potential A_τ .

technology τ increases 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

A competitive final-good sector places the following demand for intermediate varieties:

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

where P_t is the final good price and $p_t(\omega(\tau, q))$ is the price charged by the monopolist firm producing $\omega(\tau, q)$. Each monopolist solves

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

taking as given the demand schedule $x_t(\omega(\tau, q))$. To save notation, we assume hereafter $\psi = (1 - \beta)$. The solution to the monopolist’s pricing problem involves the firm charging $p_t(\omega(\tau, q)) = z(\omega(\tau, q)) P_t$, assembling $x(\omega(\tau, q)) = L$ units, and realizing profits $\pi_t(\omega(\tau, q)) = \beta z(\omega(\tau, q)) L P_t$. Varieties linked to more productive, highly perfected technologies yield higher profit flow.

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 firms' optimality conditions. $\square$

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 $Q_t(a)$; 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:

$$v_t(\omega(\tau, q)) = \int_t^\infty e^{(-\int_t^s r_v dv)} \pi_s(\omega(\tau, q)) ds. \quad (7)$$

Let $\bar{v}(t|\tau) = \mathbb{E}_q\{v_t(\omega(\tau, q))\}$ denote the ex-ante expected value of a successful innovation within technology τ . It is the crucial object for determining 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

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

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 $1/\epsilon$ 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 simpler microfoundation, assuming congestion forces in research as in Jones (1995) and Acemoglu et al. (2023). Specifically, the rate at which a scientist working on technology τ discovers a successful idea equals $R(t|\tau)^{-\epsilon}$, where ϵ 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 $R(t|\tau)^{1-\epsilon}$. Throughout, we assume that $\epsilon < 1$, ruling out zero or negative marginal productivity of total research labor.

In equilibrium, the allocation of researchers must meet a no-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 vintages.

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 $\bar{v}(t|\tau)$. 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

$$\bar{v}(t|\tau) = \beta e^{\gamma\tau} Q(t|\tau) \bar{\lambda} L \int_t^\infty \exp\left(-\int_t^s r_v dv\right) P_s ds. \quad (9)$$

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

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

where $\bar{Q}_t \equiv \left[\int_0^\infty (e^{-\gamma \tilde{a}} Q_t(\tilde{a}))^{\frac{1}{\epsilon}} d\tilde{a}\right]^\epsilon$ 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 $Q(t|\tau)$:

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

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 A.1 in the Supplementary Material 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).¹⁰ 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:

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

The representative household's dynamic problem closes the model, implying the familiar Euler equation $\dot{C}_t/C_t + \dot{P}_t/P_t = r_t - \rho$ and transversality conditions. Throughout, total consumption expenditure $P_t C_t$ is the numéraire, so that $r_t = \rho$ for all t .

Equilibrium A dynamic equilibrium can now be characterized.¹¹ The state variables are the distributions of quality and varieties $[Q_t(a), \mu_t(a)]_{a=0}^{\infty}$, which evolve under the profit-maximizing research allocation $[R_t(a)]_{a=0}^{\infty}$.

Definition 1. (Equilibrium) A trajectory for the cross-sectional distributions of quality, varieties, and research $[\Theta_t]_{t \geq t_0}$ (where $\Theta_t \equiv [Q_t(a), \mu_t(a), \bar{v}_t(a), R_t(a)]_{a=0}^{\infty}$) and for the aggregate variables $[Y_t, C_t, w_t]_{t \geq t_0}$ is an **equilibrium trajectory** if, given an initial condition $[Q_{t_0}(a), \mu_{t_0}(a)]_{a=0}^{\infty}$, the following conditions are met at every period

¹⁰For 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.

¹¹A list of all equilibrium equations and unknowns is provided in Section A.2 in the Supplementary Material.

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.

3.2.1. 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. At any given $t \geq t_0$, it satisfies:*

(i) If $\bar{\lambda} > 1$ and $\frac{1}{\bar{\lambda}-1} \neq \frac{1-\epsilon}{\epsilon}$,¹² the stationary technology-age distribution of quality is:

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

where c_1, c_2 are uniquely determined constants given by:

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

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

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

(ii) If $\bar{\lambda} > 1$, the stationary technology-age distribution of varieties is:

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

¹²In the knife-edge case $\frac{1}{\bar{\lambda}-1} = \frac{1-\epsilon}{\epsilon}$, the solution is the limit of (13), $Q(a) = Q_0 \exp\{c_5(1 - e^{-\frac{1-\epsilon}{\epsilon} \gamma a})\}$ with $c_5 \equiv \frac{\bar{\lambda}-1}{\frac{1-\epsilon}{\epsilon} \gamma} \frac{R^{1-\epsilon}}{\mu_0} (Q_0/\bar{Q})^{\frac{1-\epsilon}{\epsilon}}$.

If instead $\bar{\lambda} = 1$, then $\mu(a) = c_3 - c_4 e^{-\frac{1-\epsilon}{\epsilon} \gamma a} = \lim_{\bar{\lambda} \rightarrow 1} \mu(a)$, where:

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

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

Proof. See Section A.4 in the Supplementary Material. □

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.¹³ 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:

$$\frac{R'(a)}{R(a)} \propto \frac{Q'(a)}{Q(a)} - \gamma. \quad (15)$$

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 A.3 in the Supplementary Material 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 remains below 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

¹³See equation (11) and its related discussion.

equilibrium 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. Intuitively, because researchers are perpetually shifting from old to new technologies, growth must come from how much better the latter are relative to the former.

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 $(\bar{\lambda} - 1) / \gamma < \underline{b}$, then $R'(0) < 0$, whereas, if $(\bar{\lambda} - 1) / \gamma \geq \underline{b}$, then $R'(0) \geq 0$.

Corollary 2.1 shows that if $(\bar{\lambda} - 1) / \gamma$ 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 $(\bar{\lambda} - 1)$ 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, further reinforcing its slow growth. In this case, $R(a)$ is a decreasing function for




Figure 3: Stationary (Balanced Growth) equilibrium. *Note:* 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 blue, the case of a higher γ in orange, and the case of a higher $\bar{\lambda}$ in green.

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 decline 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 blue, alongside two alternative calibrations: one where the frontier growth rate γ is increased (orange curve) and another where the knowledge spillover parameter $\bar{\lambda}$ is higher (green curve). The calibration in all plots is such that $(\bar{\lambda} - 1) / \gamma$ is sufficiently high, implying that research shares initially increase with age.

Consider first the baseline illustrative calibration, represented by the blue 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 lose 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 orange 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 D.4 in the Supplementary Material for the case $\bar{\lambda} = 1$).

3.3 The Efficient Allocation of Research Across Vintages

This Section studies the problem of a welfare-maximizing planner. The economy exhibits three standard sources of inefficiency: (i) monopolistic competition in the production of varieties, (ii) a congestion externality in research, and (iii) knowl-

edge spillovers, which give rise to a classical standing-on-the-shoulders-of-giants externality. As we will see, the first two do not distort the direction of innovation, allowing us to focus on the role of knowledge spillovers in generating research misallocation. The planner solves:

$$\max_{\{[R(t|\tau), Q(t|\tau), \mu(t|\tau)]_{\tau \leq t}\}_{t \geq t_0}} \int_{t_0}^{\infty} e^{-\rho(t-t_0)} N \log c_t dt \quad (16)$$

s.t., for every $t \geq t_0$:

$$C_t = \tilde{\beta} L \int_{-\infty}^t e^{\gamma\tau} Q(t|\tau) \mu(t|\tau) d\tau, \quad (17)$$

$$\dot{Q}(t|\tau) = Q(t|\tau)(\bar{\lambda} - 1) \frac{R(t|\tau)^{1-\epsilon}}{\mu(t|\tau)} \quad \text{for every } \tau \leq t, \quad (18)$$

$$\dot{\mu}(t|\tau) = R(t|\tau)^{1-\epsilon} \quad \text{for every } \tau \leq t, \quad (19)$$

$$\int_{-\infty}^t R(t|\tau) d\tau = R, \quad (20)$$

where $N \equiv L + R$ is the household total population, $c_t = C_t/N$ is per capita consumption, $\mu(\tau|\tau) = \mu_0$ and $Q(\tau|\tau) = Q_0$ for all τ , and $[Q(t_0|\cdot), \mu(t_0|\cdot)]$ is given. Equation (17) gives the maximum feasible consumption level at time t for a given state of technology $[\mu(t|\tau), Q(t|\tau)]_{\tau \leq t}$. It coincides with the static equilibrium level of consumption in the decentralized economy except for the constant $\tilde{\beta} \equiv \beta(1 - \beta)^{-1/\beta}$, which corrects for the monopolistic-competition distortion.

Following Nuño (2017) and Nuño and Moll (2018), the planner's problem is solved by setting up a Lagrangian and deriving the associated first-order conditions. The discussion below focuses on these conditions, while Section A.7 in the Supplementary Material provides the full derivations. Let $\psi^\mu(t|\tau)$ and $\psi^Q(t|\tau)$ denote the Lagrange multipliers associated with the laws of motion in equations (18) and (19), respectively. These multipliers have the usual interpretation: they give the social value, in utility units, of a marginal increase in the state variables $\mu(t|\tau)$ and $Q(t|\tau)$. It is useful to define:

$$\bar{v}^p(t|\tau) \equiv \psi^Q(t|\tau)(\bar{\lambda} - 1) \frac{Q(t|\tau)}{\mu(t|\tau)} + \psi^\mu(t|\tau). \quad (21)$$

We can interpret $\bar{v}^p(t|\tau)$ as the planner's social value of a new blueprint embodying technology τ at time t . By definition, new blueprints represent an expansion

in the measure of varieties, $\dot{\mu}(t|\tau)$. This horizontal expansion has social value $\psi^\mu(t|\tau)\dot{\mu}(t|\tau)$. Moreover, because new varieties are on average better than previously existing ones, they also increase the technology's quality:

$$\dot{Q}(t|\tau) = \dot{\mu}(t|\tau) \times \left[(\bar{\lambda} - 1) \frac{Q(t|\tau)}{\mu(t|\tau)} \right].$$

In turn, these additional $\dot{Q}(t|\tau)$ quality units have social value $\dot{Q}(t|\tau)\psi^Q(t|\tau)$. The total social value of $\dot{\mu}(t|\tau)$ new patents is therefore $\dot{\mu}(t|\tau)\bar{v}^p(t|\tau)$, where $\bar{v}^p(t|\tau)$ is given by (21) and represents the value of a unit flow of new blueprints.

The Lagrangian first order condition with respect to $R(t|\tau)$ gives:

$$\frac{R(t|\tau)}{R(t|\tau')} = \left(\frac{\bar{v}^p(t|\tau)}{\bar{v}^p(t|\tau')} \right)^{\frac{1}{\epsilon}} \quad \forall \tau, \tau' \leq t. \quad (22)$$

This is exactly the relative research supply equation (8) in the decentralized equilibrium, except that the social value of a patent $\bar{v}^p(t|\tau)$ replaces here the stock market value, $\bar{v}(t|\tau)$. The planner manages congestion externalities in the same way as the market: it equalizes, across all technologies, the value of the marginal product of a researcher—the probability that a researcher innovates, $R(t|\tau)^{-\epsilon}$, times the value of a successful innovation, $\bar{v}^p(t|\tau)$. If the social and market values of a patent coincide, the two allocations are identical. The congestion externality, per se, does not distort the distribution of innovative efforts across research lines, a result established in [Hopenhayn and Squintani \(2021\)](#).¹⁴

The key question is whether and how $\bar{v}^p(t|\tau)$ and $\bar{v}(t|\tau)$ differ. By taking the first order conditions with respect to $Q(t|\tau)$ and $\mu(t|\tau)$, we can solve for the Lagrangian multipliers $\psi^\mu(t|\tau)$ and $\psi^Q(t|\tau)$, and hence for the value $\bar{v}^p(t|\tau)$:

$$\bar{v}^p(t|\tau) = \int_t^\infty \underbrace{\frac{\tilde{\beta} L \bar{\lambda} A_\tau Q(t|\tau)}{C_s}}_{\equiv \pi_t^P(s|\tau)} e^{-\int_t^s \rho - (\bar{\lambda} - 1) g_\mu(v|\tau) dv} ds, \quad (23)$$

where, as in the rest of the paper, $g_x = \dot{x}/x$ denotes growth rates. For comparison,

¹⁴This result concerns the distribution of researchers across research lines, not the total number of researchers. In a model where households choose between research and production, the congestion externality could drive an inefficiently high level of entry into research—an extensive-margin distortion, as in [Jones \(1995\)](#).

recall the market value of a blueprint derived in equation (9):

$$\bar{v}(t|\tau) = \int_t^\infty \underbrace{\frac{\beta L A_\tau \bar{\lambda} Q(t|\tau)}{C_s}}_{\equiv \pi_t(s|\tau)} e^{-\int_t^s \rho dv} ds = \frac{\beta L A_\tau \bar{\lambda} Q(t|\tau)}{C_t} D_t,$$

with $D_t = \int_t^\infty \exp(-\int_t^s g_C(v) + \rho dv) ds$ summarizing the future discounting path.

To compare the two value functions, start with the flow payoffs $\pi_t^P(s|\tau)$ and $\pi_t(s|\tau)$. In the decentralized equilibrium, $\pi_t(s|\tau)$ corresponds to the expected profit flow a patent, created in time t , yields at time s . In the planner's allocation, $\pi_t^P(s|\tau)$ corresponds to such patent's direct impact on the household's marginal utility of consumption at time s .¹⁵ These flow terms are equal, except for $\tilde{\beta}$ and β which differ by the markup charged by monopolist firms. However, the allocation of research across technologies, $R(t|\tau)/R(t|\tau')$, depends on the relative valuation of a patent, and either β or $\tilde{\beta}$ cancels out in (22). The markup inefficiency, by being uniform across technologies, does not distort the allocation of researchers.

Next, consider how the flow payoff from a new patent, $\pi_t^P(s|\tau)$ or $\pi_t(s|\tau)$, is discounted back to the present. In the decentralized equilibrium, this flow is discounted at the constant rate ρ . In the planner's allocation the effective discount rate is $\rho - (\bar{\lambda} - 1)g_\mu(v|\tau) < \rho$. If researchers at time v invent new varieties linked to technology τ at rate $g_\mu(v|\tau)$, the planner knows that they build on the shoulders of past inventions by a factor $\bar{\lambda} - 1$. By internalizing these future knowledge spillovers, the planner's effective discount rate on an invention's future returns is smaller.

Most importantly, the planner's effective discount rate varies with the innovation intensity $g_\mu(v|\tau)$, being asymmetric across technologies. This is the difference driving the planner's research allocation away from the market allocation. To see

¹⁵Specifically,

$$\pi_t^P(s|\tau) = \left. \frac{\partial \log C_s}{\partial x_s(\omega(\tau, q))} \right|_{x_s^P(\omega(\tau, \bar{\lambda} Q(t|\tau))}$$

where, as defined in Section 3.1, $x_t(\omega(\tau, q))$ denotes the quantity produced of a variety ω linked to technology τ with quality q . The superscript P denotes its value in the planner allocation.

why, note that substituting (23) into (22) yields

$$\frac{R(t|\tau)^\epsilon}{R(t|\tau')^\epsilon} = \frac{\bar{v}^p(t|\tau)}{\bar{v}^p(t|\tau')} = \underbrace{\frac{A_\tau Q(t|\tau)}{A_{\tau'} Q(t|\tau')}}_{\text{decentralized equilibrium}} \frac{D^p(t|\tau)}{D^p(t|\tau')}, \quad (24)$$

where $D^p(t|\tau) \equiv \int_t^\infty \exp\left(-\int_t^s g_C(v) + \rho - (\bar{\lambda} - 1)g_\mu(v|\tau)dv\right) ds$. In the decentralized equilibrium, the discounting term D_t is identical across technologies, so the last term in equation (24) cancels out and only the direct effect given by the relative profit ratio remains. In contrast, in the planner allocation, $D^p(t|\tau)$ potentially varies across technologies as it depends on the entire technology-specific future path of spillovers $[(\bar{\lambda} - 1)g_\mu(t + s|\tau)]_{s \geq 0}$. The relative social value of these spillovers, summarized by the ratio $D^p(t|\tau)/D^p(t|\tau')$, is precisely the term in the planner's first-order condition absent in the market equilibrium.

It is immediate that, if there are no knowledge spillovers ($\bar{\lambda} = 1$) or if the planner does not value them ($\rho \rightarrow \infty$), then $D^p(t|\tau)/D^p(t|\tau') \rightarrow 1$ for all τ and τ' , and the decentralized equilibrium is efficient. When this is not the case,

$$\forall t, \tau, D^p(t|\tau) > \lim_{\tau' \rightarrow -\infty} D^p(t|\tau') = \int_t^\infty \exp\left(-\int_t^s g_C(v) + \rho dv\right) ds.$$

As one considers older technologies in the cross-section, $D^p(t|\tau)$ converges to its lower bound, approaching from above the functional form for D_t in the decentralized equilibrium. Intuitively, even though the planner discounts the future less than the market for all technologies, it discounts it relatively more for the very old ones, which have fewer future spillovers to be missed. For these, the planner's discount rate converges to the market's. For sufficiently old τ' , the ratio $D^p(t|\tau')/D^p(t|\tau)$ is strictly below one.

More generally, however, for an arbitrary pair of technologies τ and τ' , it is not necessarily true that the planner innovates relatively more on the younger of them. To see this, it is useful to rewrite $D^p(t|\tau)$ as:

$$D^p(t|\tau) \equiv \int_t^\infty \underbrace{\frac{Q(s|\tau)}{Q(t|\tau)}}_{\substack{\text{Spillovers} \\ \text{between } t \text{ and } s}} \underbrace{\frac{u'(C_s)}{u'(C_t)} e^{-\rho(s-t)}}_{\substack{\text{Intertemporal} \\ \text{discount factor}}} ds. \quad (25)$$

This representation shows how $D^p(t|\tau)$ can be interpreted as the present discounted value of future knowledge spillovers within technology τ . The term $Q(s|\tau)/Q(t|\tau)$ captures quality growth between t and s , i.e., the cumulative spillovers over that interval. This higher quality raises the consumption good one-for-one at date s , which is valued at $u'(C_s)$. Finally, dividing by $u'(C_t)$ and discounting at rate ρ converts this value to time t units.

As equation (25) makes clear, the inequality $D^p(t|\tau) \geq D^p(t|\tau')$ need not hold in general. A new technology may exhibit slow initial growth before accelerating later in its life cycle, generating small short-run spillovers. Even if it eventually produces large spillovers, an impatient planner will prefer an older technology with larger spillovers in the near term.

To sharpen the analytical characterization, Proposition 3 now examines the efficient research allocation along a balanced-growth trajectory.

Definition 3. (Efficient BGP trajectory) An efficient Balanced Growth Path (BGP) is a trajectory $[\mu(t|\cdot), Q(t|\cdot), \bar{v}^p(t|\cdot), R(t|\cdot), C_t]_t$ that satisfies the laws of motion and resource constraints in (17)–(20), the planner first order conditions (22)–(23) and is such that $Q(t|\tau)$ and $\mu(t|\tau)$ depend on t and τ only through $t - \tau$.

Proposition 3. *In an efficient BGP trajectory, consumption grows at rate γ , and the research allocation across technologies of different ages is time invariant. It satisfies, for every $a \geq 0$, $R'(a)/R(a) < \rho/\epsilon$ and $\lim_{a \rightarrow \infty} R(a) = 0$. Thus, if $\rho \rightarrow 0$, $R'(a) < 0 \forall a$.*

Proof. See Section A.7 in the Supplementary Material. □

Proposition 3 states that an infinitely patient planner chooses a strictly declining research distribution along a BGP. In the BGP, the life cycle of technologies is stationary, so a newer technology necessarily has a larger sum of remaining future spillovers than an older one. If the planner does not discount these spillovers—that is, if it is indifferent to how front-loaded or back-loaded they are—it allocates strictly more research to newer technologies. As the planner becomes impatient $\rho > 0$, this monotonicity does not necessarily hold and the research distribution might be hump-shaped, initially increasing in age before declining. The steepness of the initial increase is bounded by ρ/ϵ : the more impatient the planner, the less innovation young technologies may receive relative to those at the peak. As ρ increases, the hump-shaped profile may accentuate, until the distribution converges




Figure 4: Planner and decentralized equilibrium stationary research allocations. *Note:* In the left panel, the solid lines show efficient BGP research allocations (Definition 3) for different values of the discount rate ρ ; the dashed line shows the decentralized equilibrium BGP research allocation (Definition 2), which is independent of ρ . Holding ρ at 0.025, the right panel shows research allocations under the uncertainty extension of Section 4. Solid lines correspond to the case $\delta(a) = 0.5 \exp(-0.2a)$; dash-dotted lines correspond to the benchmark with no uncertainty ($\delta(a) = 0$). Green lines denote the planner allocation and pink lines the decentralized equilibrium. In both panels, parameters not varied across lines are held at their baseline values in Table 2.

to that of the decentralized equilibrium as $\rho \rightarrow \infty$. Figure 4 illustrates these results by plotting the planner and the decentralized equilibrium research distribution along a BGP under different ρ values.¹⁶

4 Alternative mechanisms and model extensions

The theory in Section 3 focused on knowledge spillovers, a central externality in endogenous growth models (Romer, 1990; Aghion and Howitt, 1992) with substantial empirical support (Bloom et al., 2013; Adhami, 2026). While the key allocative inefficiency discussed in Section 3.3 is portable to other settings featuring such spillovers, broader conclusions regarding underinvestment in new technologies may depend on additional forces not modeled above. This section discusses some of them along with a few model extensions.

Uncertainty Consider an extension of the baseline model where varieties are subject to a death shock with arrival rate $\delta(t - \tau)$ that is decreasing in technology

¹⁶The decentralized research distribution does not depend on ρ . This rate enters the value functions only through the common discount term D_t , which cancels out from the relative value between distinct technologies.

age, $\delta'(\cdot) < 0$, so that newer technologies carry greater risk.

In the decentralized equilibrium, the relative research supply schedule becomes:

$$\left(\frac{R(t|\tau)}{R(t|\tau')} \right)^\epsilon = \frac{e^{\gamma\tau} Q(t|\tau) D(t|\tau)}{e^{\gamma\tau'} Q(t|\tau') D(t|\tau')}, \quad (26)$$

with $D(t|\tau) = \int_t^\infty \exp(-\int_t^s \rho + g_C(v) + \delta(v - \tau) dv) ds$. As in the baseline model, the research allocation depends on each technology's profitability, $e^{\gamma\tau} Q(t|\tau)$, but now also on a risk-adjusted discount term $D(t|\tau)$. Because newer technologies face a riskier path ahead, their future profits are discounted more heavily, so $D(t|\tau)$ is lower and research is drawn away from them. The planner's discount term takes this same form, augmented by the spillover term $-g_Q(v|\tau)$ from Section 3.3, so it too penalizes riskier new technologies.

Both the market and the planner therefore shift research toward older technologies, as shown in the right panel of Figure 4, where the solid (uncertainty) lines lie to the right of the dash-dotted (no-uncertainty) lines for both.¹⁷ However, the planner–market gap in favor of new technologies does not narrow, and may even widen. Knowledge spillovers remain the only source of misallocation, so the planner continues to direct relatively more research toward new technologies. Moreover, since $\delta'(\cdot) < 0$, risk weighs more heavily on the near term than on the long run of a technology's life cycle. Private agents have a higher effective discount rate than the planner, as established in Section 3.3, and are therefore more sensitive to changes in short-run expected profits. Near-term risk thus deters the market from new technologies disproportionately more than it deters the planner, widening the gap.

This conclusion rests on firms fully internalizing the risk carried by new technologies, as the death shock $\delta(\cdot)$ assumes. If some risk goes uninternalized, as with the existential risks of transformative technologies discussed by Jones (2024) and Acemoglu and Lensman (2024), a channel opposite to knowledge spillovers arises: the uninternalized risk pushes the planner to innovate less in new technologies than the market (Acemoglu and Lensman, 2024), leaving the net effect on the planner–market gap ambiguous.¹⁸

¹⁷Section A.9 in the Supplementary Material presents all equilibrium equations characterizing this extension.

¹⁸Experimentation externalities are another non-internalized aspect of uncertainty, but operate

Switching costs Switching innovation across vintages is costly when firms accumulate vintage-specific know-how. [Aghion et al. \(2026\)](#) models an economy where there is an inverse relationship between the number of product lines a firm leads within a technology and the cost of directing innovation to it. This generates path dependence. Established firms, with prior experience in the old vintage, find further innovation there cheaper, which can slow the transition to the new vintage. Relatedly, [Jovanovic and Nyarko \(1996\)](#) shows that learning-by-doing in an inferior technology can delay switching and even produce a no-growth trap in which producers never adopt the new one. For a sufficiently patient planner, such lock-in effects can reinforce the underinvestment in new vintages. On the other hand, within-firm knowledge accumulation partially internalizes the future-cost benefits of today's innovation, a channel that could offset this underinvestment, especially among young firms.

Cross-technology knowledge spillovers The baseline model assumes that knowledge spillovers operate exclusively within each technology's boundaries. Section [A.5](#) in the Supplementary Material extends the model by allowing for cross-technology spillovers, provided they are weaker than their within-technology counterparts.

Population Growth and Basic Research Population growth and basic research can serve as microfoundations for the emergence of better technologies over time. Suppose an exogenous share α of the research population works in basic research, say in universities or government labs; the remaining share $1 - \alpha$ develops existing technologies, as in the baseline model of Section [3](#). I endogenize α along a BGP below.

If $R_{B,t}$ researchers work in basic research, they discover new fundamental ideas at the flow rate $\mu(t|t) \equiv R_{B,t}^{1-\epsilon}$. These ideas comprise the breakthroughs launching a new vintage at time t . The intrinsic potential of the new vintage, A_t , exceeds that of the previous vintage, $A_{t-\Delta}$, by the aggregate mass of breakthroughs arriving over

in the same direction as knowledge spillovers. When agents can free-ride on others' experiments with risky new technologies, waiting to learn whether they work, equilibrium experimentation, and hence innovation in new technologies, may be inefficiently low ([Bolton and Harris, 1999](#); [Keller et al., 2005](#)).

the interval $[t - \Delta, t]$:

$$A_t - A_{t-\Delta} = \int_{t-\Delta}^t \mu(s|s) ds = \int_{t-\Delta}^t R_{B,s}^{1-\epsilon} ds.$$

Intuitively, as $\Delta \rightarrow 0$, the productivity gain of the newest vintage equals precisely the number of breakthroughs that originated it: $A_{t+\Delta} - A_t = \Delta R_{B,t}^{1-\epsilon}$. Integration yields $A_t = \int_{-\infty}^t R_{B,s}^{1-\epsilon} ds$. Assuming that the populations of production workers and researchers grow at a constant rate n , the knowledge frontier satisfies

$$A_t \propto e^{\gamma t}, \quad \text{where } \gamma = n(1 - \epsilon).$$

The long-run growth rate is proportional to the population growth rate as in the semi-endogenous growth literature (Jones, 1995).

The remaining parts of the model do not change qualitatively. The share of non-basic researchers working on each vintage, $R_t(a)/(1 - \alpha)R_t$, still follows the expression derived in (10), and the laws of motion for $\mu_t(a)$ and $Q_t(a)$ retain the same form as in (11). The model features a balanced growth path in which the quality life cycle $Q(a)$ is stationary, as is the normalized number of varieties $\tilde{\mu}(a) = \mu_t(a)/\mu_t(0)$. Section A.6 in the Supplementary Material presents the complete solution.

Endogenizing α . The basic-research share is pinned down along a BGP by an arbitrage condition between basic and applied research. From equation (9), the expected return to an applied researcher, $\tilde{v}_t \equiv R(t|\tau)^{-\epsilon} \bar{v}(t|\tau)$, is common across existing vintages. The average basic researcher earns $\bar{v}_t^B = R_{B,t}^{-\epsilon} \beta A_t Q_0 L_t D_t$, reflecting that a breakthrough blueprint has potential A_t and initial quality Q_0 . Setting $\tilde{v}_t = \bar{v}_t^B$ yields

$$\frac{\alpha}{1 - \alpha} = \left(\frac{Q_0}{\bar{\lambda} \bar{Q}} \right)^{\frac{1}{\epsilon}}, \quad (27)$$

where $\bar{Q}$ is the stationary value of $\bar{Q}_t$ along the BGP.¹⁹

¹⁹One may ask whether a planner choosing α along a BGP would direct all researchers to basic research. The following argument suggests not. Since γ is pinned down by n , α affects only the level of the BGP: a higher α raises the breakthrough mass $(\alpha R_t)^{1-\epsilon}$ with which each vintage is born, but leaves fewer researchers to improve vintages over their life cycle. At $\alpha = 1$ every vintage is frozen at birth, and since $R^{1-\epsilon}$ satisfies an Inada condition at $R = 0$, the first applied researcher on any vintage has infinite marginal value while the marginal value of basic research is finite.

5 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. The baseline dataset comprises approximately eight 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 noted, 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 USPTO classifies each patent into one and only one principal technology class. The analysis here focuses exclusively on that classification.

5.1 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 concerns 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,




Figure 5: Patenting share and proxied emergence date. *Note:* The dots in each panel 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 panel by the maximum value achieved. The solid line represents the fitted logistic trend (see this Section text and Supplementary Material’s Section C.2). The dashed lines represent the year in which the trend achieved 10% of the ceiling value. In all panels, the underlying data come from the USPTO Historical Patent Data Files (Marco et al., 2015); see this Section text.

which are the focus here. For instance, notice in Panel (a) of Figure 5 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 5. More generally, such a trend is corroborated by multiple case studies highlighting a consistent S-shaped pattern in patent timelines (Achilladelis et al., 1990; 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 5 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 unity 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 dashed 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 5 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 inno-

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	0.09
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: Examples of USPC technology classes with estimated dates of origin and patenting shares in two decades. Dates of origin are estimated via logistic regression following [Griliches \(1957\)](#) (see Supplementary Material, Section C.2).

vation 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 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 5, and ‘Pipe joints,’ in Figure D.3 in the Supplementary Material. The highest innovation share achieved by ‘Tuners’ was in 1950, and it is only half the value for the earliest observation of ‘Pipe joints’ (Figure D.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 6 shows a time series with the number of technologies emerging in each decade. Three waves 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




Figure 6: Technologies and patents across and within time. *Note:* The left plot shows the number of technology classes (USPC classification system) that emerged in each decade. The emergence date is estimated with a logistic regression, following [Griliches \(1957\)](#) (see this Section text and Supplementary Material’s Section C.2). 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. Residual vintages outside [0,400] are omitted, accounting for less than 0.5% of patents in 2000–2009. The underlying data come from the USPTO Historical Patent Data Files ([Marco et al., 2015](#)).

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 extended into 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 1968.²⁰

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 B.4 in the Supplementary Material). 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

²⁰For 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.

stable for more than a century. This is shown in the right plot in Figure 6, 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 D in the Supplementary Material redoes Figures 1–2 controlling for patent quality.

6 Quantitative Analysis

This Section explores the potential quantitative relevance of the research misallocation identified in the theory. The model in Section 3 has nine parameters ($\epsilon, \bar{\lambda}, \gamma, \beta, \rho, Q_0, \mu_0, L, R$), whose baseline calibration is discussed below. The welfare analysis that follows, however, considers a range of parameter values rather than relying on the baseline alone.

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.

Assume that the observed number of patents is an indicator, albeit noisy, of the total number of new ideas generated in a specific time period:

$$\text{Patents}(t|\tau) = \kappa_\tau \times \dot{\mu}(t|\tau) \times u(t|\tau). \quad (28)$$

Here, $\text{Patents}(t|\tau)$ denotes the number of patents issued in a given period within the technological class τ , $\dot{\mu}(t|\tau)$ is the number of new ideas in the model, κ_τ is the technology-specific patent propensity, and $u(t|\tau)$ is a scalar disturbance term.

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

Table 2: Baseline Calibration

Parameter	Description	Value	Parameter	Description	Value
γ	Frontier growth rate	2%	β	Labor share	0.66
$\bar{\lambda}$	Knowledge spillover	2.56	ρ	Discount rate	0.025
ϵ	Research elasticity	0.65	L, R	Population	1, 1
μ_0	Breakthrough varieties	0.55	Q_0	Initial quality	1

equation:

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

where δ_τ and δ_t represent technology and time fixed effects,²¹ while the regressor $\bar{v}(t|\tau)$ 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 reduced form correlation between patenting activity and valuation, after filtering out technology and time fixed effects. It calibrates $\epsilon = 0.65$ to match the coefficient on $\bar{v}(t|\tau)$ estimated from regression (29). Section C.1 in the Supplementary Material presents the complete results. The quantitative exercises below use a range of calibrations for ϵ .

Calibrating γ and $\bar{\lambda}$ As established in Section 3.2.1, the peak of the cross-sectional distribution of research depends crucially on the parameter ratio $(\bar{\lambda} - 1)/\gamma$, while γ alone pins down the long-run per capita growth rate. Setting $\gamma = 2\%$, and given the baseline values of all remaining parameters (discussed below), the calibration chooses $(\bar{\lambda} - 1)/\gamma$ —and hence $\bar{\lambda}$ —so that the model matches the peak age of 106 years observed in the cross-sectional distribution for the 2000s, depicted in Figure 2. The resulting value is $\bar{\lambda} = 2.56$.

Other parameters and Model Fit The labor force L and the supply of researchers R are both normalized to one. To calibrate the measure of breakthrough varieties μ_0 , consider the basic research extension described in Section 4. There, the R researchers working on existing vintages represent only a share $1 - \alpha$ of the total research population, with the remaining $R_B = R\alpha/(1 - \alpha)$ allocated to basic research to

²¹ $\delta_\tau = \log \kappa_\tau - \frac{1-\epsilon}{\epsilon} \log(\beta\bar{\lambda})$, and $\delta_t = -\frac{1-\epsilon}{\epsilon} \gamma t - \frac{1-\epsilon}{\epsilon} \log(D_t \bar{Q}_t)$.




Figure 7: The Hump-Shaped Innovation-Age Profile—Data (2000s decade) and model. *Note:* The dots represent US data from the 2000 decade (see Sections 2 and 5), trimmed at the top 1% to remove outliers; the solid line is a non-linear fit to these data. The dashed line represents the model’s competitive equilibrium stationary patenting distribution $\hat{\mu}(a)$ under the baseline calibration (Table 2).

launch new vintages, so that $\mu_0 = (R_B)^{1-\epsilon}$. The share of total US R&D spent on basic research has averaged approximately 15% since the 1950s (NSF, 2026), which yields $\mu_0 = [R\alpha / (1 - \alpha)]^{1-\epsilon} \approx 0.55$.

Turning to Q_0 , Proposition 2 establishes that the normalized schedules $\hat{Q}(a) \equiv Q(a)/Q_0$ and $\hat{\mu}(a) \equiv \mu(a)/\mu_0$ are independent of Q_0 : like L , it scales the level of model aggregates without affecting the research allocation or technologies’ life cycle. The calibration thus sets $Q_0 = 1$. Following commonly used values in the literature, the discount rate ρ is set to 2.5%, and β , representing the labor share in total output, to 0.66.

Figure 7 compares the empirical distribution of patents in the 2000 decade (orange line) with the stationary age profile of innovation produced by the calibrated model (blue dashed line, representing the competitive equilibrium). The model is successful in fitting the data, as demonstrated by the proximity of the curves. The peak age is targeted through the calibration of $\bar{\lambda}$ and the model also matches other untargted features of the distribution, such as the share of patents accruing to frontier technologies (age zero).




Figure 8: Transition Dynamics and Consumption Gains from the Efficient Allocation. *Note:* The left panel plots the consumption growth rate along the efficient transition path starting from the decentralized BGP at $t = 0$. Each line represents results computed under a different ρ value, with all other parameters held at the baseline calibration (Table 2). The right panel plots the ratio of consumption along the efficient BGP to consumption along the decentralized BGP, for alternative values of $\bar{\lambda}$, ρ , and ϵ . Each column presents an alternative value of the parameter indicated on the horizontal axis, holding the remaining parameters fixed at their baseline calibration.

6.1 The Costs of Research Misallocation

We now consider an efficient trajectory implemented by a benevolent planner who takes the decentralized equilibrium BGP as the starting point at $t = 0$. To isolate the gains from reallocating research across technologies to internalize knowledge spillovers, we assume the planner cannot correct the static markup distortion.²² Section A.8 in the Supplementary Material shows how the planner's choices can be decentralized under a budget-neutral policy that subsidizes research across technologies and taxes firms' profits.

Each line in Figure 8's left plot traces the path of consumption growth along the efficient trajectory for a different value of ρ , with all remaining parameters fixed at the baseline calibration (Table 2). At $t = 0$, the economy is at the decentralized BGP and the research distribution is as depicted in Figure 7. The planner immediately shifts research toward new technologies, whose innovations yield larger future knowledge spillovers. In the short run, however, diverting research toward these not-yet-perfected technologies slows growth.

How large an initial growth slowdown the planner tolerates to reap future

²²In the planner optimality conditions derived in Section 3.3, this amounts to replacing $\tilde{\beta}$ with its decentralized counterpart $\beta(2 - \beta)/(1 - \beta)$.

Table 3: Welfare Gains at Different Calibrations

ρ	Welfare Gain	$\bar{\lambda}$	Welfare Gain	ϵ	Welfare Gain
0.01	24.04%	2.7	9.35%	0.61	10.27%
0.025	7.33%	2.0	2.78%	0.70	5.26%
0.05	1.55%	1.5	0.67%	0.80	2.65%

Note: This Table reports the consumption-equivalent welfare gains yielded by the efficient transition path departing from the decentralized BGP at $t = 0$. 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 right parts of the Table repeat this procedure, but varying individually $\bar{\lambda}$ and ϵ , respectively.

knowledge spillovers depends crucially on the discount rate. When $\rho = 0.01$, for instance, the growth rate falls to less than 0.018, a 10% reduction relative to the long-run baseline of 0.02. It only recovers after approximately 20 years, after which it rises significantly above 0.02 for several decades. A less patient planner ($\rho = 0.05$) reallocates fewer researchers to new technologies: the growth rate decreases modestly to 0.019 and rebounds within 12 years. After that, however, the growth effects experienced by the less patient planner are also more modest—but still significant and persistent.

By growing above the 0.02 long-run rate for several decades, the planner converges to a BGP trajectory with a higher consumption level. The right plot of Figure 8 shows how much higher consumption along the efficient BGP is relative to the decentralized equilibrium. Each bar varies one parameter at a time (as indicated on the x-axis), holding all others at the baseline calibration. Under the baseline $\rho = 0.025$, long-run consumption is 61% higher than in the decentralized BGP. For the more patient planner $\rho = 0.01$, this gain rises to almost 70%; even for the less patient planner, it remains close to 50%. We can also see how long-run consumption gains increase in $\bar{\lambda}$ and decrease in ϵ .²³ Intuitively, the more knowledge spillovers there are for the planner to take advantage of (higher $\bar{\lambda}$), and the more flexibly the planner can redirect research toward them (lower ϵ), the greater the cumulative consumption gains.

Finally, Table 3 reports the welfare gains the planner achieves. The left panel presents these gains for different values of ρ , keeping all remaining parameters at the baseline calibration (Table 2). Under the baseline $\rho = 0.025$, the welfare gain

²³Table D.1 in the Supplementary Material presents results for other model parameters.

equals 7.33%. This means the representative household is indifferent between (i) the consumption path in Figure 8's left plot and (ii) an alternative trajectory in which consumption grows at rate 0.02 for all $t \geq 0$ but at a level 7.33% higher than in the decentralized BGP. With a lower discount, such as $\rho = 0.01$, the welfare gain rises to 24%; with a more conservative choice, $\rho = 0.05$, it equals 1.55%, still significant. The central and right panels of Table 3 vary $\bar{\lambda}$ and ϵ , respectively. The same patterns observed for long-run consumption in the right plot of Figure 8 carry over to welfare: gains increase in $\bar{\lambda}$ and decrease in ϵ .

7 Conclusion

Research agencies, such as the National Science Foundation or the European Research Council, allocate a significant share of their budget to support research in new technological paradigms. However, the rationale for this *selective* approach remains understudied in the growth literature. While the importance of basic research has long been recognized (Nelson, 1959; Akcigit et al., 2020), it remains unclear why market forces, even within narrow categories of applied, market-based research, may misallocate R&D across technologies of different maturities. In standard innovation-led growth models, knowledge accumulation and the resulting knowledge externalities are uniform, occurring within a single representative technology. In contrast, this paper develops a model that features the endogenous evolution of several technologies, which gradually replace each other. Knowledge externalities are no longer uniform, but rather asymmetric, depending on the life cycle stage of the technology in which they arise. It is significant that even when the total supply of researchers and the long-run growth rate are fixed, a benevolent planner can achieve substantial welfare gains by reallocating research to new technologies.

References

Daron Acemoglu and Todd Lensman. Regulating transformative technologies. *American Economic Review: Insights*, 6(3):359–76, September 2024.

- 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, Albert Schwarzkopf, and Martin Cines. The dynamics of technological innovation: the case of the chemical industry. *Research Policy*, 19(1):1–34, 1990.
- Mohamad Adhami. Quantifying knowledge spillovers using firm and product dynamics. Working Paper, Stanford University, 2026. URL https://adhamimohamad.github.io/papers/Adhami_JMP.pdf.
- Philippe Aghion and Peter Howitt. A model of growth through creative destruction. *Econometrica*, 60(2):323–351, 1992.
- Philippe Aghion, Antonin Bergeaud, Maarten De Ridder, and John Van Reenen. Lost in transition: Financial barriers to new technologies. Working paper, March 2026.
- 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.
- Enrico Berkes. Comprehensive universe of U.S. patents (CUSP): Data and facts. Working Paper, University of Maryland, 2016.
- Enrico Berkes, Kristina Manysheva, and Martí Mestieri. Global innovation spillovers and productivity: Evidence from 100 years of world patent data. (2022-15), May 2022a. Working paper. Version dated June 14, 2024.
- Enrico Berkes, Kristina Manysheva, and Martí Mestieri. Fields of knowledge [dataset], ‘Fields_of_knowledge.csv’. GitHub, 2022b. Accessed November 10, 2025. <https://github.com/manysheva/MyWebsite>.
- 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.
- Patrick Bolton and Christopher Harris. Strategic experimentation. *Econometrica*, 67(2):349–374, 1999. doi: <https://doi.org/10.1111/1468-0262.00022>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/1468-0262.00022>.
- 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.
- Cooperative Patent Classification Committee. Cpc to ipc concordance [dataset]. United States Patent and Trademark Office and European Patent Office, 2025. Accessed November 3, 2025. <https://www.cooperativepatentclassification.org/cpcConcordances>.
- 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.
- Google. Google patents public data [dataset], table ‘patents-public-data.patents.publications’. Google BigQuery public datasets, 2023. Accessed August 27, 2023. https://console.cloud.google.com/marketplace/product/google_patents_public_datasets/google-patents-public-data.
- 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.
- Charles I. Jones. The ai dilemma: Growth versus existential risk. *American Economic Review: Insights*, 6(4):575–90, December 2024. doi: 10.1257/aeri.20230570. URL <https://www.aeaweb.org/articles?id=10.1257/aeri.20230570>.
- 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.
- Godfrey Keller, Sven Rady, and Martin Cripps. Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68, 2005. doi: <https://doi.org/10.1111/j.1468-0262.2005.00564.x>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-0262.2005.00564.x>.
- Bryan Kelly, Dimitris Papanikolaou, Amit Seru, and Matt Taddy. Replication data for: Measuring technological innovation over the long run [dataset]. American

- Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor], 2021a. Accessed August 27, 2023. <https://doi.org/10.3886/E119043V1>.
- Bryan Kelly, Dimitris Papanikolaou, Amit Seru, and Matt Taddy. Measuring technological innovation over the long run. *American Economic Review: Insights*, 3(3): 303–20, 2021b.
- 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.
- Leonid Kogan, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. Technological innovation, resource allocation, and growth: Extended data [dataset]. GitHub, 2022. Patents granted through 2022. Accessed February 21, 2023. <https://github.com/KPSS2017/Technological-Innovation-Resource-Allocation-and-Growth-Extended-Data>.
- 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*, 29607, 2023.
- Robert E. Lucas and Benjamin Moll. Knowledge growth and the allocation of time. *Journal of Political Economy*, 122(1):1–51, 2014.
- David G. Luenberger. *Optimization by Vector Space Methods*. John Wiley & Sons, New York, 1969. ISBN 047155359X.
- 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.

- 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.
- NSF. National patterns of R&D resources: 2023–24 data update. Technical Report NSF 26-313, National Science Foundation, Alexandria, VA, 2026. URL <https://nces.nsf.gov/data-collections/national-patterns/2023-2024>.
- Galo Nuño. Optimal social policies in mean field games. *Applied Mathematics & Optimization*, 76(1):29–57, 2017.
- Galo Nuño and Benjamin Moll. Social optima in economies with heterogeneous agents. *Review of Economic Dynamics*, 28:150–180, April 2018.
- PatentsView. Cooperative patent classification of us patents [dataset], table ‘g_cpc_current’. United States Patent and Trademark Office, Open Data Portal, Bulk Data Directory, 2023. Vintage covering patents granted through mid-2023. Accessed mid-2023. <https://data.uspto.gov/bulkdata>.
- PatentsView. Us patent citations [dataset], table ‘g_us_patent_citation’. United States Patent and Trademark Office, Open Data Portal, 2025. Release of November 21, 2025. Accessed April 1, 2026. <https://data.uspto.gov/bulkdata>.
- 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.
- 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.

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

USPTO. Classification of patents, 1838 [dataset]. Reproduced at page 189 of *Development and Use of Patent Classification Systems*. Washington, DC: For sale by the Superintendent of Documents, U.S. Government Printing Office, 1966, 1838. Digitized volume: <https://hdl.handle.net/2027/uiug.30112101591045>.

USPTO. Classification index of subjects of invention [dataset], editions of 1872 and 1895. Washington, DC: U.S. Government Printing Office, 1872–1895. Digitized volumes: 1872, <https://hdl.handle.net/2027/loc.ark:/13960/t3pv7p12s>; 1895, <https://hdl.handle.net/2027/loc.ark:/13960/t29893g8p>. Consulted through HathiTrust, February to July 2023.

USPTO. Classification bulletin [dataset], volumes 29–33 and 44–78. Washington, DC: U.S. Government Printing Office, 1912–1940. Covering 1912 to 1914 and 1920 to 1940. HathiTrust catalogue record 004052778, <https://catalog.hathitrust.org/Record/004052778>; handles for the individual volumes are listed in the replication package. Consulted February to July 2023.

USPTO. Manual of classification of subjects of invention [dataset], editions of 1912, 1916, 1920, 1923 and 1947. Washington, DC: U.S. Government Printing Office, 1912–1947. Consulted February to July 2023. Handles for the five editions are listed in the replication package; the 1912 edition was consulted through Google Books and the rest through HathiTrust.

USPTO. *The classification of patents*. U.S. Government Printing Office, 1915.

USPTO. Classes in the u.s. patent classification system: Dates established [dataset]. United States Patent and Trademark Office, 2009. The page states “This information is complete through 2009.” Accessed October 2022. <https://www.uspto.gov/patents/classes-us-patent-classification-system-dates-established>.

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

USPTO. Historical patent data files [dataset], `historical_masterfile.dta`. USPTO Open Data Portal, 2015. Accessed August 18, 2026. <https://developer.uspto.gov/product/historical-masterfile>.

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 Model derivations, proofs, and extensions

A.1 Derivations of the laws of motion for $\mu(t|\tau)$, $Q(t|\tau)$, $\mu_t(a)$, $Q_t(a)$

Fix a technology τ and define $F_t^\tau(q) \equiv F_t(\tau, q)$, where $F_t(\tau, q)$ is the cdf of the density f_t (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 $\Delta \rightarrow 0$, gives:

$$\mu'(t|\tau) = R(t|\tau)^{1-\epsilon}. \quad (\text{A.30})$$

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 $\Delta \rightarrow 0$, and using (A.30), gives:

$$Q'(t|\tau) = \frac{Q(t|\tau)}{\mu(t|\tau)} (\bar{\lambda} - 1) R(t|\tau)^{1-\epsilon} \quad (\text{A.31})$$

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:

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

A.2 Summary of Equilibrium Conditions

This section collects the firms' and households' decision problems from the model of Section 3. It then maps each equilibrium unknown to the equation that determines it.

Intermediate varieties differ by their (τ, q) type. Varieties linked to the same technology and with the same quality level are consumed at the same rate and yield the same profit flow. To save notation, as standard, we can then index varieties by (τ, q) and change the space of integration from Ω_t to $\mathcal{T} \times \mathcal{Q}$ using the distribution $f_t(\tau, q)$.²⁴ The model decision problems can be summarized as:

Final Good Producer Production function:

$$Y_t = \frac{L_t^\beta}{1-\beta} \int_{\mathcal{T} \times \mathcal{Q}} e^{\gamma\tau} q x_t(\tau, q)^{1-\beta} f_t(\tau, q) dq d\tau. \quad (\text{A.32})$$

²⁴Recall that $f_t(\tau, q)$ is not a probability distribution as it does not integrate to one, but rather to the total mass of varieties.

Problem and FOCs:

$$\max_{L_t, \{x_t(\tau, q)\}} P_t Y_t - \int_{\mathcal{T} \times \mathcal{Q}} p_t(\tau, q) x_t(\tau, q) f_t(\tau, q) dq d\tau - w_t L_t.$$

$$\frac{x_t(\tau, q)}{L_t} = \left[\frac{p_t(\tau, q) / P_t}{e^{\gamma \tau} q} \right]^{-\frac{1}{\beta}}, \quad (\text{A.33})$$

$$\frac{w_t L_t}{\beta} = P_t \frac{L_t^\beta}{1 - \beta} \int_{\mathcal{T} \times \mathcal{Q}} e^{\gamma \tau} q x_t(\tau, q)^{1-\beta} f_t(\tau, q) dq d\tau. \quad (\text{A.34})$$

Monopolist Firms Problem and FOC:

$$\pi_t(\tau, q) = \max_{p_t(\tau, q)} x_t[p_t(\tau, q)] [p_t(\tau, q) - P_t \psi e^{\gamma \tau} q].$$

$$\frac{dx_t[p_t(\tau, q)]}{dp_t(\tau, q)} [p_t(\tau, q) - P_t \psi e^{\gamma \tau} q] + x_t[p_t(\tau, q)] = 0. \quad (\text{A.35})$$

Amount of final good used in the production of varieties:

$$X_t = \int \psi e^{\gamma \tau} q x_t(\tau, q) f_t(\tau, q) dq d\tau. \quad (\text{A.36})$$

Equal return on all assets: no-arbitrage conditions in financial markets (HJB equations):

$$r_t = \frac{\pi_t(\tau, q) + \dot{v}_t(\tau, q)}{v_t(\tau, q)}. \quad (\text{A.37})$$

This leads to:

$$v_t(\tau, q) = \int_t^\infty e^{-\int_t^s r_v dv} \pi_s(\tau, q) ds.$$

Researchers Problem (Innovation direction):

$$\max_{\tau} R(t|\tau)^{-\epsilon} \mathbb{E}_q \{v_t(\tau, q)\}.$$

Optimality (no-arbitrage across vintages):

$$R(t|\tau)^{-\epsilon} \mathbb{E}_q \{v_t(\tau, q)\} = R(t|\tau')^{-\epsilon} \mathbb{E}_q \{v_t(\tau', q)\} \quad \forall \tau, \tau'. \quad (\text{A.38})$$

Household Problem and budget constraint:

$$\begin{aligned} & \max_{C_t, V_t} \int_{t_0}^{\infty} e^{-\rho t} \log C_t dt \\ & \text{subject to} \\ & \dot{V}_t = r_t V_t + w_t L + \Pi_t^R - P_t C_t. \end{aligned} \quad (\text{A.39})$$

Here, V_t is total household financial wealth, and Π_t^R is the total value of new patents discovered at time t , both endogenously through researchers and exogenously, through the breakthrough patents that start a new vintage:

$$\Pi_t^R = \int_{-\infty}^t R(t|\tau)^{1-\epsilon} \mathbb{E}_q \{v_t(\tau, q)\} d\tau + \mu_0 v_t(t, Q_0).$$

Optimality conditions:

$$\frac{\dot{C}_t}{C_t} + \frac{\dot{P}_t}{P_t} = r_t - \rho \quad \text{and} \quad \lim_{t \rightarrow \infty} \left[\exp \left(- \int_{t_0}^t r(s) ds \right) V_t \right] = 0. \quad (\text{A.40})$$

Innovation possibilities frontier Laws of motion:

$$\dot{Q}(t|\tau) = Q(t|\tau)(\bar{\lambda} - 1) \frac{R(t|\tau)^{1-\epsilon}}{\mu(t|\tau)}, \quad (\text{A.41})$$

$$\dot{\mu}(t|\tau) = R(t|\tau)^{1-\epsilon}. \quad (\text{A.42})$$

Market Clearing and Resource Constraints

$$L_t = L, \quad (\text{labor}) \quad (\text{A.43})$$

$$\int_{-\infty}^t R(t|\tau) d\tau = R, \quad (\text{researchers}) \quad (\text{A.44})$$

$$Y_t = C_t + X_t, \quad (\text{goods}) \quad (\text{A.45})$$

$$V_t = \int_{\mathcal{T} \times \mathcal{Q}} v_t(\tau, q) f_t(\tau, q) dq d\tau. \quad (\text{financial assets}) \quad (\text{A.46})$$

Equilibrium An **equilibrium trajectory** is a collection of time paths

$$\begin{aligned} &\{Y_t, C_t, X_t, L_t, V_t, r_t, w_t\}_{t \geq t_0}, \\ &\{[x_t(\tau, q), p_t(\tau, q), v_t(\tau, q)]_{\tau \leq t, q \geq 0}\}_{t \geq t_0}, \\ &\{[R(t|\tau), Q(t|\tau), \mu(t|\tau)]_{\tau \leq t}\}_{t \geq t_0} \end{aligned}$$

such that, for every $t \geq t_0$:

1. $[x_t(\tau, q)]_{\tau \leq t, q \geq 0}$ and L_t solve the final good producer problem, satisfying (A.33)–(A.34). In turn, the production function (A.32) pins down Y_t and (A.36) pins down X_t .
2. $[p_t(\tau, q)]_{\tau \leq t, q \geq 0}$ solves the monopolist firms problem, satisfying (A.35);
3. $[v_t(\tau, q)]_{\tau \leq t, q \geq 0}$ satisfy the no-arbitrage/HJB (A.37).
4. $[R(t|\tau)]_{\tau \leq t}$ satisfies the research optimality condition (A.38) for all $\tau, \tau' \leq t$ and the research market-clearing condition (A.44).
5. C_t and V_t solve the household problem, satisfying (A.39) and (A.40).
6. $[Q(t|\tau), \mu(t|\tau)]_{\tau \leq t}$ satisfy (A.41)–(A.42).
7. w_t clears the labor market (A.43), and r_t clears the asset market (A.46).

In total, there are 13 unknowns and 14 equations. Notice that:

- Every unknown corresponds to one equation, except for the research density $R(t|\tau)$, which requires two: (A.38) and (A.44). Equation (A.38) is a no-arbitrage condition that determines only *relative* research levels $R(t|\tau)/R(t|\tau')$ across technologies τ and τ' . To pin down the levels themselves, we need the aggregate researcher resource constraint (A.44).
- The economy's resource constraint (A.45) is the only equation not used to pin down the equilibrium. The definition of the aggregate intermediate input (A.36) yields $X_t = (1 - \beta)^2 Y_t$, so the resource constraint holds if and only if $C_t = Y_t - X_t = \beta(2 - \beta)Y_t$. That this is indeed the case follows from Walras'

Law: given equilibrium in all other markets, the goods market clears automatically. To verify, start from the household budget constraint. From (A.37) and the profit flow implied by the monopolist's problem, $v_t(\tau, q) = \beta A_\tau q L D_t$, where D_t was defined in the main text. Imposing the financial-market clearing condition (A.46) yields $V_t = \beta L \int_{-\infty}^t A_\tau Q(t|\tau) \mu(t|\tau) d\tau D_t$. Next, rewrite Π_t^R as

$$\Pi_t^R = \int_{-\infty}^t R(t|\tau)^{1-\epsilon} \beta L \bar{\lambda} A_\tau Q(t|\tau) d\tau D_t + \mu_0 \beta L A_t Q_0 D_t.$$

Substituting for V_t and Π_t^R in the budget constraint and simplifying yields $C_t = \beta(2 - \beta)Y_t$, confirming that the resource constraint (A.45) holds exactly.

A.3 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 on 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:

$$\partial_t \mu_t(a) = -\partial_a \mu_t(a) + R_t(a)^{1-\epsilon} - \delta \mu_t(a) \tag{A.47}$$

$$\partial_t Q_t(a) = -\partial_a Q_t(a) + Q_t(a)(\bar{\lambda} - 1) \frac{R_t(a)^{1-\epsilon}}{\mu_t(a)}. \tag{A.48}$$

Throughout this Section, the parameters are assumed to satisfy $\bar{\lambda} < 1/(1 - \epsilon)$, so that the exponent in (13) is positive.²⁵

Proposition 4. *Suppose $\bar{\lambda} < 1/(1 - \epsilon)$. A Balanced Growth Path exists if and only if*

$$\delta(\bar{\lambda} - 1) < \gamma. \quad (\text{A.49})$$

It is unique and characterized by:

1. *Aggregate variables grow at the rate γ .*
2. *Innovation shares converge to zero as technologies age:*

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

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

$$Q(a) = \left\{ c_1 - c_2 e^{-\left(\frac{1-\epsilon}{\epsilon} \gamma - \delta\right)a} \right\}^{\frac{1}{\frac{1}{\bar{\lambda}-1} - \frac{1-\epsilon}{\epsilon}}} \quad (\text{A.50})$$

where c_1, c_2 are

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

$$c_1 \equiv Q_0^{\frac{1}{\bar{\lambda}-1} - \frac{1-\epsilon}{\epsilon}} + c_2$$

*and $\bar{Q}$ is the stationary level of $\bar{Q}_t$, which is unique.*²⁶

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

$$\frac{\mu(a)}{\mu_0} = e^{-\delta a} \left(\frac{Q(a)}{Q_0} \right)^{\frac{1}{\bar{\lambda}-1}}. \quad (\text{A.51})$$

²⁵Outside this region, the characterization becomes more involved when we extend the model with the death shock $\delta > 0$. The baseline calibration in Table II satisfies $\bar{\lambda} < 1/(1 - \epsilon)$.

²⁶If $\delta = \frac{1-\epsilon}{\epsilon} \gamma$, which (A.49) allows, (A.50) is replaced by its limit, $Q(a) = \{Q_0^{\frac{1}{\bar{\lambda}-1} - \frac{1-\epsilon}{\epsilon}} + c_6 a\}^{1/\left(\frac{1}{\bar{\lambda}-1} - \frac{1-\epsilon}{\epsilon}\right)}$ with $c_6 \equiv \left(\frac{1}{\bar{\lambda}-1} - \frac{1-\epsilon}{\epsilon}\right) \frac{Q_0^{\frac{1}{\bar{\lambda}-1}} R^{1-\epsilon}}{\bar{Q}^{\frac{1-\epsilon}{\epsilon}} \mu_0} (\bar{\lambda} - 1)$.

Proof. Equations (A.50)–(A.51) are derived in Section A.4 for a general $\delta \geq 0$. Write $\nu \equiv \frac{1}{\bar{\lambda}-1} - \frac{1-\epsilon}{\epsilon} > 0$. Since $\nu > 0$, the term inside the braces in (A.50) is positive for every a and every $\bar{Q} > 0$, so $Q(a)$ is finite at every age. Its long-run behavior depends on δ . If $\delta < \frac{1-\epsilon}{\epsilon}\gamma$, $Q(a)$ converges to $c_1^{1/\nu}$. If $\delta = \frac{1-\epsilon}{\epsilon}\gamma$, it grows polynomially. If $\delta > \frac{1-\epsilon}{\epsilon}\gamma$, it grows exponentially at the asymptotic rate

$$g \equiv \frac{\delta - \frac{1-\epsilon}{\epsilon}\gamma}{\nu},$$

and $g < \gamma$ if and only if (A.49) holds. Three conditions must be met for a BGP, and each is equivalent to $g < \gamma$. First, from (10), $R(a) \propto (e^{-\gamma a} Q(a))^{1/\epsilon}$ converges to zero if and only if $Q(a)$ grows slower than γ . Second, for the same reason, the integral defining $\Xi(\bar{Q})$ in Section A.4 converges if and only if $g < \gamma$. Third, by Proposition 1 and (A.51), the output integrand is proportional to $e^{-(\gamma+\delta)a} Q(a)^{1+\frac{1}{\bar{\lambda}-1}}$, which decays if and only if $(1 + \frac{1}{\bar{\lambda}-1})g < \gamma + \delta$. Using $1 + \frac{1-\epsilon}{\epsilon} = 1/\epsilon$, this inequality is again equivalent to (A.49). Hence, if (A.49) fails, no BGP exists. If it holds, $\Xi(\bar{Q})$ is finite for every $\bar{Q} > 0$ and the argument in Section A.4 applies: Ξ is continuous and strictly decreasing, tends to infinity as $\bar{Q} \rightarrow 0$ (since $c_2 \rightarrow \infty$), and to $Q_0(\epsilon/\gamma)^\epsilon$ as $\bar{Q} \rightarrow \infty$. The stationary $\bar{Q}$ therefore exists and is unique, which pins down the BGP. $\square$

Condition (A.49) compares two growth rates. As we know, γ represents how better new technologies are, while $\delta(\bar{\lambda} - 1)$ is how fast a technology can steadily grow in the long run.²⁷ If the latter reaches or exceeds the former, very old technologies never lose salience: their quality keeps up with the frontier, the share of research they attract stops declining, and the research distribution cannot integrate to the finite supply R . Under the maintained assumption $\bar{\lambda} < 1/(1 - \epsilon)$, this is also exactly the condition under which output would be unbounded at every finite t , so no separate case arises. A steady-state in which old technologies remain permanently salient is also inconsistent with the empirical evidence that the patent share of older technologies declines over time.

²⁷If a constant number of researchers work on technology τ , its quality $Q(t|\tau)$ growth rate converges in the long run to $\delta(\bar{\lambda} - 1)$.

A.4 Proof of 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 Section A.3).

Using the BGP definition (Definition 2), plug (A.47) into (A.48) and differentiate with respect to time to obtain:

$$\mu'(t|\tau) = \mu(t|\tau) \left[\frac{1}{\epsilon} \frac{Q'(t|\tau)}{Q(t|\tau)} - \frac{Q''(t|\tau)}{Q'(t|\tau)} - \frac{1-\epsilon}{\epsilon} \gamma \right] \quad (\text{A.52})$$

Use (A.47) and (A.48) to substitute for $\mu'(t|\tau)/\mu(t|\tau)$ in (A.52) and obtain:

$$\begin{aligned} \left[\frac{1-\epsilon}{\epsilon} \gamma - \delta \right] + \left[\frac{1}{\bar{\lambda}-1} - \frac{1}{\epsilon} \right] \frac{Q'(t|\tau)}{Q(t|\tau)} + \frac{Q''(t|\tau)}{Q'(t|\tau)} &= 0 \\ e^{(\frac{1-\epsilon}{\epsilon} \gamma - \delta)(t-\tau)} Q(t|\tau)^{\frac{1}{\bar{\lambda}-1} - \frac{1}{\epsilon}} Q'(t|\tau) &= c_\tau \end{aligned} \quad (\text{A.53})$$

where the second line follows from integrating the first and $c_\tau \equiv Q(\tau|\tau)^{\frac{1}{\bar{\lambda}-1} - \frac{1}{\epsilon}} Q'(\tau|\tau)$.

Equation (A.53) is a separable first-order differential equation in time t . We can hence apply the well-known solution for separable equations and obtain:

$$Q(t|\tau) = \left[\frac{\frac{1}{\bar{\lambda}-1} - \frac{1-\epsilon}{\epsilon}}{\frac{1-\epsilon}{\epsilon} \gamma - \delta} \frac{Q(\tau|\tau)^{\frac{1}{\bar{\lambda}-1}}}{\bar{Q}^{\frac{1-\epsilon}{\epsilon}}} \frac{R^{1-\epsilon}}{\mu(\tau|\tau)} (\bar{\lambda}-1) \left(1 - e^{-(\frac{1-\epsilon}{\epsilon} \gamma - \delta)(t-\tau)} \right) + Q(\tau|\tau)^{\frac{1}{\bar{\lambda}-1} - \frac{1-\epsilon}{\epsilon}} \right]^{\frac{1}{\frac{1}{\bar{\lambda}-1} - \frac{1-\epsilon}{\epsilon}}} \quad (\text{A.54})$$

where we have used (A.48) 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 , (A.54) gives the BGP solution for $Q(t|\tau)$ presented in equation (13).

Uniqueness of $\bar{Q}$

An equilibrium value for $\bar{Q}$ must pin down $Q(a)$ in equation (13) such that

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

exactly equals $\bar{Q}$. We thus look for a fixed point in the space $\Xi(\bar{Q}) \times \bar{Q}$. The derivations below take $\delta = 0$; the case $\delta > 0$ is treated in Section A.3.

The domain of Ξ must first be established. Define

$$\underline{Q} \equiv \begin{cases} Q_0 \left[\frac{R^{1-\epsilon}}{\gamma \mu_0} \left(\bar{\lambda} - 1 - \frac{\epsilon}{1-\epsilon} \right) \right]^{\frac{\epsilon}{1-\epsilon}} & \text{if } \bar{\lambda} > \frac{1}{1-\epsilon}, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.55})$$

If $\bar{\lambda} \leq 1/(1-\epsilon)$, $Q(a)$ is finite for every a and every $\bar{Q} > 0$: the term inside the braces in (13) is positive when the inequality is strict, and in the knife-edge case the exponential solution in footnote 12 in the main paper is bounded by $Q_0 e^{c_5}$. If $\bar{\lambda} > 1/(1-\epsilon)$, then $c_2 < 0$ and the term inside the braces decreases from $Q_0^{\frac{1}{\bar{\lambda}-1} - \frac{1-\epsilon}{\epsilon}}$ toward c_1 . Since the exponent is negative, $Q(a)$ is bounded if and only if $c_1 > 0$, which, from the definition of c_1 , is equivalent to $\bar{Q} > \underline{Q}$. If $c_1 = 0$, $Q(a)$ is finite at every age but grows exponentially at the rate $\frac{1-\epsilon}{\frac{1-\epsilon}{\epsilon} - \frac{1}{\bar{\lambda}-1}} \gamma > \gamma$, so $\Xi(\bar{Q}) = \infty$. If $c_1 < 0$, $Q(a)$ diverges at a finite age and again $\Xi(\bar{Q}) = \infty$. Hence, in all cases, no fixed point exists on $(0, \underline{Q}]$, and it suffices to study Ξ on $(\underline{Q}, \infty)$, where it is finite and continuous.

From (A.54), $\partial Q(a)/\partial \bar{Q} < 0$ for every $a > 0$, which implies that (i) $\Xi'(\bar{Q}) < 0$ on $(\underline{Q}, \infty)$.²⁸ Moreover, when $\bar{Q} \rightarrow \infty$, then $Q(a) \rightarrow Q_0 \forall a$. Hence, (ii) $\lim_{\bar{Q} \rightarrow \infty} \Xi(\bar{Q}) = Q_0(\epsilon/\gamma)^\epsilon$. On the other hand, (iii) $\lim_{\bar{Q} \downarrow \underline{Q}} \Xi(\bar{Q}) = \infty$. If $\underline{Q} = 0$, this follows because $c_2 \rightarrow \infty$ as $\bar{Q} \rightarrow 0$, so $Q(a) \rightarrow \infty$ for every $a > 0$. If $\underline{Q} > 0$, it follows because $c_1 \rightarrow 0$ as $\bar{Q} \downarrow \underline{Q}$, in which case $Q(a)$ converges pointwise to $|c_2|^{-\left(\frac{1-\epsilon}{\epsilon} - \frac{1}{\bar{\lambda}-1}\right)^{-1}} \exp \left\{ \frac{1-\epsilon}{\frac{1-\epsilon}{\epsilon} - \frac{1}{\bar{\lambda}-1}} \gamma a \right\}$, which grows at a rate strictly above γ , so that the integral defining Ξ diverges by monotone convergence. Conditions (i)–(iii), together with the continuity of $\Xi(\bar{Q})$ on $(\underline{Q}, \infty)$, guarantee a unique fixed point in the space $\Xi(\bar{Q}) \times \bar{Q}$. This fixed point is greater than $\max\{\underline{Q}, Q_0(\epsilon/\gamma)^\epsilon\}$. In particular, $\bar{Q} > \underline{Q}$ implies $c_1 > 0$, so $Q(a)$ in (13) is finite and strictly increasing for every a , as claimed in the Proposition.

²⁸When $\bar{\lambda} > 1/(1-\epsilon)$, a higher $\bar{Q}$ lowers $|c_2|$, raises the term inside the braces, and, given the negative exponent, lowers $Q(a)$.

A.5 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 as $\bar{Q}_t$ 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, $\bar{Q}_t \equiv \left[\int_0^\infty (e^{-\gamma \tilde{a}} Q_t(\tilde{a}))^{\frac{1}{\epsilon}} d\tilde{a} \right]^\epsilon$ 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-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 to find within a technology, still draws researchers away from old technologies in the very long run.

A.6 Population growth and basic research

This section extends the baseline model of Section 3 with two additional assumptions:

- **Population growth.** The production-worker and researcher populations grow at the constant rate $n \geq 0$:

$$L_t = L_0 e^{nt} \quad \text{and} \quad R_t = R_0 e^{nt}.$$

- **Basic research.** An exogenous share $\alpha \in (0, 1)$ of researchers engages in basic research, generating a flow $\mu(\tau|\tau) = (\alpha R_\tau)^{1-\epsilon}$ of breakthroughs that launch the newest technology vintage at each date τ . The remaining share $1 - \alpha$ pursues applied research on existing technologies. The vintage born at τ has intrinsic productivity $A_\tau = e^{\gamma\tau}$, where $\gamma = n(1 - \epsilon)$.

All remaining assumptions are as in the baseline model.

Static allocation Population growth leaves the static equilibrium allocation of Section 3.1 unchanged. That allocation depends only on the current population

and the current state of technology, $\{L_t, [\mu(t|\tau), Q(t|\tau)]_{\tau \leq t}\}$, not on how L_t evolves over time or on the research assumptions that generated $[\mu(t|\tau), Q(t|\tau)]_{\tau \leq t}$.

Innovation and Dynamics Population growth affects variety values only through the profit flow $\pi_s(\omega(\tau, q)) = P_s \beta A_\tau q L_s$, since now L_s grows at the rate n . A researcher innovating in technology τ expects to discover a variety of quality $q = \bar{\lambda} Q(t|\tau)$, so the expected patent value is

$$\bar{v}(t|\tau) = \beta A_\tau \bar{\lambda} Q(t|\tau) L_t \int_t^\infty e^{-\int_t^s (r_v - n) dv} P_s ds,$$

where the effective discount rate shifts from r_v to $r_v - n$ because $L_s = L_t e^{n(s-t)}$. Substituting into the relative research-supply condition (8) shows that applied research shares take the same form as in the baseline:

$$\frac{R(t|\tau)}{R_{A,t}} = \left(\frac{e^{-\gamma(t-\tau)} Q(t|\tau)}{\bar{Q}_t} \right)^{\frac{1}{\epsilon}}, \quad (\text{A.56})$$

where $R_{A,t} \equiv (1 - \alpha) R_t$ and

$$\bar{Q}_t \equiv \left[\int_{-\infty}^t \left(e^{-\gamma(t-\tilde{\tau})} Q(t|\tilde{\tau}) \right)^{\frac{1}{\epsilon}} d\tilde{\tau} \right]^\epsilon = \left[\int_0^\infty \left(e^{-\gamma \tilde{a}} Q_t(\tilde{a}) \right)^{\frac{1}{\epsilon}} d\tilde{a} \right]^\epsilon.$$

The laws of motion for the state variables are unchanged:

$$\mu'(t|\tau) = R(t|\tau)^{1-\epsilon}, \quad (\text{A.57})$$

$$Q'(t|\tau) = \frac{Q(t|\tau)}{\mu(t|\tau)} (\bar{\lambda} - 1) R(t|\tau)^{1-\epsilon}. \quad (\text{A.58})$$

Household The representative household maximizes the intertemporal utility

$$U = \int_{t_0}^\infty e^{-\rho(t-t_0)} N_t \log(c_t) dt,$$

where $N_t = L_t + R_t$ is the total household population and $c_t = C_t/N_t$ is per capita consumption ($\rho > n$ ensures effective discounting). The Euler equation and

transversality condition are:

$$\dot{P}_t/P_t + \dot{C}_t/C_t = r_t + n - \rho, \quad \lim_{T \rightarrow \infty} V_T \exp \left(- \int_{t_0}^T r_v dv \right) = 0,$$

where V_t is total household wealth (in levels, not per capita). Taking total consumption expenditure as numeraire, $P_t C_t = 1$, implies $\dot{P}_t/P_t + \dot{C}_t/C_t = 0$, so the Euler equation yields

$$r_t = \rho - n.$$

Equilibrium and stationary equilibrium Equilibrium is defined as in Definition 1 (see also Section A.2). A stationary trajectory is an equilibrium trajectory such that $\dot{Q}_t(a) = 0$ for all a —equivalently, an equilibrium trajectory in which $Q(t|\tau)$ depends on t and τ only through their difference $t - \tau$.

The stationary trajectory admits an analytical solution via the same steps as those taken in the derivation of as Proposition 2. Combining (A.56), (A.57), and (A.58) yields:

$$\frac{Q'(t|\tau)}{Q(t|\tau)} \frac{1}{\bar{\lambda} - 1} = \frac{1}{\epsilon} \frac{Q'(t|\tau)}{Q(t|\tau)} - \frac{Q''(t|\tau)}{Q'(t|\tau)} + \gamma \left(1 - \frac{1 - \epsilon}{\epsilon} \right) \quad (\text{A.59})$$

This is a second-order differential equation in $Q(t|\tau)$; solving it following Section A.4 gives the stationary quality-age schedule:

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

where $\tilde{\epsilon} \equiv (1 - \epsilon)(1 - 2\epsilon)/\epsilon$ and

$$c_2 \equiv \left(\frac{1}{\bar{\lambda} - 1} - \frac{1 - \epsilon}{\epsilon} \right) Q(\tau|\tau)^{\frac{1}{\bar{\lambda}-1}} \left(\frac{1 - \alpha}{\alpha} \right)^{1-\epsilon} \bar{Q}^{\frac{1-\epsilon}{\epsilon}} (\bar{\lambda} - 1) \frac{1}{n\tilde{\epsilon}}$$

and $c_1 = c_2 + Q(\tau|\tau)^{\frac{1}{\bar{\lambda}-1} - \frac{1-\epsilon}{\epsilon}}$. Given $Q(a)$, combining the laws of motion (A.57)–(A.58) yields:

$$\frac{\mu_t(a)}{\mu_t(0)} = \left(\frac{Q(a)}{Q(0)} \right)^{\frac{1}{\bar{\lambda}-1}}.$$

The normalized variety-age profile $\tilde{\mu}(a)$ is thus also time invariant. Finally, since

$$Y_t = \frac{A_t L_t}{1 - \beta} \int_0^\infty e^{-\gamma a} Q_t(a) \mu_t(a) da,$$

it follows that output per capita is proportional to $A_t \mu_t(0)$ along the stationary path. The per capita GDP growth rate equals 2γ .

A.7 Derivations and Proofs of Section 3.3 – the planner problem

In the planner problem of Section 3.3, the control and state variables are $R(t|\cdot)$, $\mu(t|\cdot)$, and $Q(t|\cdot)$. These are functions defined over the continuum of technology vintages. The problem is therefore infinite-dimensional: at each point in time, the planner chooses an entire function $R(t|\cdot)$ rather than a finite-dimensional vector of controls. As a consequence, standard notions of marginal variation used in finite-dimensional optimal control problems—such as perturbing a scalar state or control at time t —are no longer adequate. Optimality must instead be characterized through variations of functions, a broader notion of differentiation known as the Gâteaux differential (see, e.g., [Luenberger, 1969](#); [Nuño, 2017](#); [Nuño and Moll, 2018](#); [Lucas and Moll, 2014](#)).

Definition A.7.1. (*Gâteaux differential, from [Luenberger \(1969\)](#)*) Let F be a real-valued functional defined on a vector space $\mathcal{X}$. The Gâteaux differential of F at $x \in \mathcal{X}$ in the direction $h \in \mathcal{X}$ is

$$\delta F(x; h) \equiv \lim_{\varepsilon \rightarrow 0} \frac{F(x + \varepsilon h) - F(x)}{\varepsilon},$$

whenever the limit exists.

Remark on notation One could equivalently perform the derivations that follow using age notation, $a \in [0, \infty)$, such that $R_t(a)$, $\mu_t(a)$, and $Q_t(a)$ lie in the space of continuous functions on $[0, \infty)$ for each t . The derivations below use instead the τ notation predominant in the main text.

Solving the planner problem The solution follows [Nuño \(2017\)](#), [Nuño and Moll \(2018\)](#), and [Lucas and Moll \(2014\)](#), who extend standard optimal control techniques

used in economics to infinite dimensions. Starting from the optimization program in equations (16)–(20) the first step is to set up the Lagrangian:

$$\begin{aligned}
H[Q, \mu, R, \psi^Q, \psi^\mu, \psi^R] &= \int_0^\infty e^{-\rho t} \log \left(\beta(1-\beta)^{-\frac{1}{\beta}} L \int_{-\infty}^t e^{\gamma \tau} Q(t|\tau) \mu(t|\tau) d\tau \right) dt \\
&+ \int_0^\infty \int_{-\infty}^t e^{-\rho t} \psi^Q(t|\tau) \left(\frac{(\bar{\lambda}-1)Q(t|\tau)R(t|\tau)^{1-\epsilon}}{\mu(t|\tau)} - \frac{\partial Q(t|\tau)}{\partial t} \right) d\tau dt \\
&+ \int_0^\infty \int_{-\infty}^t e^{-\rho t} \psi^\mu(t|\tau) \left(R(t|\tau)^{1-\epsilon} - \frac{\partial \mu(t|\tau)}{\partial t} \right) d\tau dt \\
&+ \int_0^\infty e^{-\rho t} \psi^R(t) \left(R - \int_{-\infty}^t R(t|\tau) d\tau \right) dt \\
&\text{given } [Q(0|\tau), \mu(0|\tau)]_{\tau \leq 0}, \quad \text{and } Q(\tau|\tau) = Q_0, \mu(\tau|\tau) = \mu_0 \text{ for every } \tau.
\end{aligned}$$

A necessary condition for the household welfare to achieve a maximum at Q, μ, R is that there exist functions ψ^Q and ψ^μ such that $\delta H[Q, \mu, R, \psi^Q, \psi^\mu, \psi^R] = 0$. We now derive these first-order conditions. These conditions are necessary but not sufficient. As in [Nuño and Moll \(2018\)](#), I therefore do not establish existence or uniqueness; when multiple trajectories satisfy the first-order conditions, the efficient allocation is the one attaining the highest social welfare.

First order condition with respect to Q Using Definition [A.7.1](#), the Gâteaux derivative of H with respect to Q in the direction of an arbitrary function h is:

$$\begin{aligned}
&\int_0^\infty \int_{-\infty}^t e^{-\rho t} \frac{\tilde{\beta} L}{C_t} e^{\gamma \tau} h(t|\tau) \mu(t|\tau) d\tau dt \\
&+ \int_0^\infty \int_{-\infty}^t e^{-\rho t} \psi^Q(t|\tau) (\bar{\lambda}-1) h(t|\tau) \frac{R(t|\tau)^{1-\epsilon}}{\mu(t|\tau)} d\tau dt \\
&- \int_0^\infty \int_{-\infty}^t e^{-\rho t} \psi^Q(t|\tau) \frac{\partial h(t|\tau)}{\partial t} d\tau dt.
\end{aligned}$$

The last term can be modified by changing the order of integration and perform-

ing integration by parts:

$$\begin{aligned}
& \int_0^\infty \int_{-\infty}^t e^{-\rho t} \psi^Q(t|\tau) \frac{\partial h(t|\tau)}{\partial t} d\tau dt = \int_{-\infty}^\infty \int_{\max(0,\tau)}^\infty e^{-\rho t} \psi^Q(t|\tau) \frac{\partial h(t|\tau)}{\partial t} dt d\tau \\
& = \int_{-\infty}^\infty \left[e^{-\rho t} \psi^Q(t|\tau) h(t|\tau) \right]_{\max(0,\tau)}^\infty d\tau \\
& - \int_{-\infty}^\infty \int_{\max(0,\tau)}^\infty e^{-\rho t} h(t|\tau) \left[\frac{\partial \psi^Q(t|\tau)}{\partial t} - \rho \psi^Q(t|\tau) \right] dt d\tau.
\end{aligned}$$

Combining all terms, the first-order condition with respect to Q becomes

$$\begin{aligned}
& \int_0^\infty \int_{-\infty}^t e^{-\rho t} h(t|\tau) \left[\frac{\tilde{\beta} L e^{\gamma \tau}}{C_t} \mu(t|\tau) + \psi^Q(t|\tau) (\bar{\lambda} - 1) \frac{R(t|\tau)^{1-\epsilon}}{\mu(t|\tau)} \right. \\
& \quad \left. + \frac{\partial \psi^Q(t|\tau)}{\partial t} - \rho \psi^Q(t|\tau) \right] d\tau dt - \lim_{T \rightarrow \infty} \int_{-\infty}^\infty e^{-\rho T} \psi^Q(T|\tau) h(T|\tau) d\tau.
\end{aligned}$$

where I used the fact that $h(\max(0, \tau)|\tau) = 0$ since the initial conditions at $t = 0$ are given, as well as the vintage birth condition $Q(\tau|\tau)$ (perturbations have to respect the given data).

For the Lagrangian to be stationary, $\delta H = 0$, the differential above must vanish for all directions h . This yields:

$$[\rho - g_Q(t|\tau)] \psi^Q(t|\tau) = \frac{\tilde{\beta} L e^{\gamma \tau} \mu(t|\tau)}{C_t} + \partial_t \psi^Q(t|\tau), \quad (\text{A.60})$$

$$\lim_{T \rightarrow \infty} e^{-\rho T} \psi^Q(T|\tau) = 0. \quad (\text{A.61})$$

First-order condition with respect to μ Applying analogous steps with respect to μ yields the necessary conditions:

$$\rho \psi^\mu(t|\tau) = \frac{\tilde{\beta} L e^{\gamma \tau} Q(t|\tau)}{C_t} - \psi^Q(t|\tau) (\bar{\lambda} - 1) Q(t|\tau) \frac{R(t|\tau)^{1-\epsilon}}{\mu(t|\tau)^2} + \partial_t \psi^\mu(t|\tau), \quad (\text{A.62})$$

$$\lim_{T \rightarrow \infty} e^{-\rho T} \psi^\mu(T|\tau) = 0. \quad (\text{A.63})$$

First order condition with respect to R

$$\begin{aligned} R(t|\tau)^\epsilon &= (1 - \epsilon) \frac{\psi^Q(t|\tau) \frac{(\bar{\lambda}-1)Q(t|\tau)}{\mu(t|\tau)} + \psi^\mu(t|\tau)}{\psi^R(t)} \\ &= (1 - \epsilon) \frac{\bar{v}^p(t|\tau)}{\psi^R(t)}, \end{aligned} \quad (\text{A.64})$$

where the second line uses the definition of $\bar{v}^p(t|\tau)$ in equation (21).

Integrating $R(t|\tau)$ over τ in equation (A.64) and imposing the resource constraint for researchers eliminates $\psi^R(t)$:

$$R(t|\tau) = \frac{\bar{v}^p(t|\tau)^{\frac{1}{\epsilon}}}{\int_{-\infty}^t \bar{v}^p(t|\tau')^{\frac{1}{\epsilon}} d\tau'} R. \quad (\text{A.65})$$

Notice that the relative research supply equation of the planner problem, equation (22), follows directly from (A.65).

Combining (A.60), (A.62), and the definition of $\bar{v}^p(t|\tau)$ in equation (21), one obtains the planner HJB:

$$\rho \bar{v}^p(t|\tau) = \frac{\tilde{\beta} \bar{\lambda} L e^{\gamma \tau} Q(t|\tau)}{C_t} + \partial_t \bar{v}^p(t|\tau). \quad (\text{A.66})$$

Proposition 3 proof Notice that, on a BGP, the planner HJB (A.66) in age notation reads:

$$\bar{v}^{p'}(a) = \rho \bar{v}^p(a) - \pi^p(a), \quad \pi^p(a) \equiv \frac{\tilde{\beta} \bar{\lambda} L Q(a) e^{-\gamma a}}{C_0} > 0,$$

and the planner FOC gives $R(a) = c \bar{v}^p(a)^{1/\epsilon}$ for some $c > 0$. Multiplying the HJB by $e^{-\rho a}$ yields

$$\frac{d}{da} [e^{-\rho a} \bar{v}^p(a)] = -e^{-\rho a} \pi^p(a) \leq 0,$$

so $e^{-\rho a} \bar{v}^p(a)$ is non-increasing in a .

Strict negativity at every finite a gives $\bar{v}^{p'}(a)/\bar{v}^p(a) < \rho$, and hence

$$R'(a)/R(a) = (1/\epsilon) \bar{v}^{p'}(a)/\bar{v}^p(a) < \rho/\epsilon.$$

Since $\bar{v}^p \geq 0$, the limit $\tilde{L} \equiv \lim_{a \rightarrow \infty} e^{-\rho a} \bar{v}^p(a) \geq 0$ exists, and by monotonicity

$\bar{v}^p(a) \geq (\tilde{L}/2) e^{\rho a}$ for every a . If $\tilde{L} > 0$, then $R(a) \geq c (\tilde{L}/2)^{1/\epsilon} e^{\rho a/\epsilon}$ for every a , contradicting $\int_0^\infty R(a) da = R < \infty$. Hence $\tilde{L} = 0$, the transversality condition $\lim_{a \rightarrow \infty} e^{-\rho a} \bar{v}^p(a) = 0$ holds, and the integral form

$$\bar{v}^p(a) = \int_a^\infty e^{-\rho(s-a)} \pi^p(s) ds \leq \int_a^\infty \pi^p(s) ds$$

applies. The BGP requirement $\int_0^\infty e^{-\gamma a} Q(a) \mu(a) da < \infty$, combined with $\mu(a) \geq \mu_0$, gives $\int_0^\infty \pi^p(s) ds < \infty$. The tail of a convergent integral vanishes, so $\bar{v}^p(a) \rightarrow 0$ and therefore $R(a) \rightarrow 0$.

Planner value function Finally, equation (23) follows from solving the planner's HJB equation (A.66) forward. The transversality condition for the value function follows from equations (A.61)–(A.63) and necessarily holds for solutions in the space where the ratio $Q(a)/\mu(a)$ stays bounded.

A.8 Decentralizing the efficient allocation

The planner allocation can be implemented with a budget-neutral tax-and-subsidy policy. Consider a technology-time-specific subsidy rate $S(t|\tau)$. An inventor who sells a technology τ patent to firms for market value V earns an additional payment $S(t|\tau)V$ from the government. Under this subsidy, the relative research supply equation becomes:

$$\left(\frac{R(t|\tau)}{R(t|\tau')} \right)^\epsilon = \frac{e^{\gamma\tau} Q(t|\tau)}{e^{\gamma\tau'} Q(t|\tau')} \frac{1 + S(t|\tau)}{1 + S(t|\tau')}. \quad (\text{A.67})$$

Equation (A.67) generalizes the laissez-faire research allocation: a technology-specific subsidy tilts the direction of innovation by raising the effective return to research in the subsidized vintage. If the subsidy rate is identical across technologies, $S(t|\tau) = S(t|\tau')$, the ratio cancels and the allocation coincides with the decentralized equilibrium.

To finance the subsidy, the government levies a time-varying tax ζ_t on firms' profits, uniform across technologies. Importantly, this tax is neutral. Because it applies as a proportional levy on monopolistic profits, it does not distort firms' pricing decisions—monopolist first-order conditions remain unchanged. Total profits

of a firm owning patent $\omega(\tau, q)$ equal $\beta z(\omega(\tau, q))LP_t$, so aggregate government revenues are

$$\zeta_t \int_{\mathcal{T} \times \mathcal{Q}} \beta A_\tau q LP_t f_t(\tau, q) dq d\tau = \zeta_t \beta (1 - \beta) \frac{Y_t}{C_t},$$

where the second equality uses the fact that total expenditure serves as the numéraire, $P_t C_t = 1$.

The government's total outlays under the subsidy policy are

$$\int_{-\infty}^t S(t|\tau) \bar{v}(t|\tau) R(t|\tau)^{1-\epsilon} d\tau,$$

where $\bar{v}(t|\tau)$ denotes the average value of new patents linked to technology τ at time t . In the presence of the profit tax, this value becomes

$$\bar{v}(t|\tau) = \frac{\beta L A_\tau \bar{\lambda} Q(t|\tau)}{C_t} D^\zeta(t),$$

with $D^\zeta(t) \equiv \int_t^\infty \exp(-\int_t^s g_{C_v} + \rho dv) (1 - \zeta_s) ds$. The term $D^\zeta(t)$ captures how the profit tax reduces the present discounted value of future earnings: it plays the same role as D_t in the laissez-faire equilibrium, except that the flow payoff at each future date s is scaled down by the after-tax share $(1 - \zeta_s)$.

Equating revenues and expenditures yields the budget-balance condition:

$$\frac{\zeta_t}{D^\zeta(t)} = \bar{\lambda} \frac{\int_{-\infty}^t e^{\gamma\tau} Q(t|\tau) \mu(t|\tau) g_\mu(t|\tau) S(t|\tau) d\tau}{\int_{-\infty}^t e^{\gamma\tau} Q(t|\tau) \mu(t|\tau) d\tau}. \quad (\text{A.68})$$

This condition pins down ζ_t as a function of the subsidy schedule $S(t|\tau)$ and the state of the economy.

With these instruments in place, the efficient allocation can be decentralized. Given an initial condition $[Q(t_0|\tau), \mu(t_0|\tau)]_{\tau \leq t_0}$, let $\mathcal{P}_{t_0} \equiv \{[R^p(t|\tau), Q^p(t|\tau), \mu^p(t|\tau)]_{\tau \leq t}\}_{t \geq t_0}$ denote the efficient trajectory chosen by the planner. A trajectory of taxes and subsidies $\{[S(t|\tau)]_{\tau \leq t}, \zeta_t\}_{t \geq t_0}$ implements the efficient allocation under competitive markets if, given $\mathcal{P}_{t_0}$, the subsidy $S(t|\tau)$ satisfies equation (A.67) for every $t \geq t_0$ and τ , and the profit tax ζ_t satisfies equation (A.68) for every $t \geq t_0$.

A.9 Uncertainty extension

All blueprints are subject to an idiosyncratic death shock. The Poisson arrival rate of this shock is denoted by $\delta(t|\tau)$ and is time-technology dependent. For instance, we can think that $\delta(t|\tau)$ is decreasing in $t - \tau$, so the older the technology, the less risky it is.

The static allocation derived in Section 3.1 remains unchanged. Turning to the model dynamics, consider the expected value of a successful innovation directed to technology τ at time t :

$$\bar{v}(t|\tau) = \beta A_\tau \bar{\lambda} Q(t|\tau) L \int_t^\infty \exp\left(-\int_t^s \rho + \delta(v|\tau) dv\right) P_s ds$$

This is the same expected value as equation (9) in the baseline model, except that the discount rate of future profits includes now the death shock.

The equilibrium relative research supply satisfies:

$$\left(\frac{R(t|\tau)}{R(t|\tau')}\right)^\epsilon = \frac{e^{\gamma\tau} Q(t|\tau) D(t|\tau)}{e^{\gamma\tau'} Q(t|\tau') D(t|\tau')} \quad (\text{A.69})$$

where $D(t|\tau) = \int_t^\infty \exp\left(-\int_t^s \rho + g_C(v) + \delta(v|\tau) dv\right) ds$. Differently from the baseline model, the forward looking component in the value function is now technology-specific and does not cancel out. Research allocation thus satisfies:

$$R(t|\tau) = \frac{\left(e^{-\gamma(t-\tau)} Q(t|\tau) D(t|\tau)\right)^{\frac{1}{\epsilon}}}{\int_{-\infty}^t \left(e^{-\gamma(t-\tilde{\tau})} Q(t|\tilde{\tau}) D(t|\tilde{\tau})\right)^{\frac{1}{\epsilon}} d\tilde{\tau}}$$

The laws of motion for $\mu(t|\tau)$ and $Q(t|\tau)$ become:

$$\dot{Q}(t|\tau) = Q(t|\tau)(\bar{\lambda} - 1) \frac{R(t|\tau)^{1-\epsilon}}{\mu(t|\tau)} \quad \text{for every } \tau \leq t \quad (\text{A.70})$$

$$\dot{\mu}(t|\tau) = R(t|\tau)^{1-\epsilon} - \delta(t|\tau)\mu(t|\tau) \quad \text{for every } \tau \leq t \quad (\text{A.71})$$

A dynamic equilibrium can again be defined as in Definition 1, and a stationary (or balanced growth) equilibrium path is once more characterized as a path where the endogenous variables depend on (t, τ) only through $t - \tau$. However, the life

cycle of technologies in this balanced growth cannot be solved in closed form as in Proposition 2. This is because the research allocation $R(t|\tau)$ cannot be written as a function of time t state variables only, but now also includes the discounted future path of risk in $D(t|\tau)$. The analytical solution tractability hinged on such property.

Planner problem The planner problem outlined in equations (16)–(20) does not change, but now includes the laws of motion (A.70)–(A.71). In the planner solution, relative research supply takes the same form as in (A.69) but with $D^p(t|\tau)$ substituting $D(t|\tau)$, where:

$$D^p(t|\tau) = \int_t^\infty \exp\left(-\int_t^s \rho + \delta(v|\tau) + g_C(v) - g_Q(v|\tau) dv\right) ds.$$

Here, as in Section 3.3, the distinction between planner and market optimality conditions in research allocation lies in the knowledge spillover term $-g_Q(v|\tau)$ in the discount rate. While both the planner and the market may discount more the future for new technologies than for old technologies due to the risk $\delta(t|\tau)$, this is not driving research misallocation.

B Alternative Measures of Technology and Age

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

Kelly et al. (2021b) use textual analysis to construct measures of forward and backward textual similarity with respect to subsequent and prior inventions, respectively. They define the ratio of these measures as a patent importance index. The more novel (dissimilar from prior art) and influential (similar to subsequent inventions) a patent, the more important it is. Following their analysis, I refer to patents in the top 5% of this index as breakthroughs and use them as the empirical counterparts to the breakthrough inventions $\mu(\tau|\tau)$ that inaugurate new technology fields in the model of Section 3.²⁹ These include, for example, Lawrence Page’s PageRank algorithm

²⁹Using the top 1% yields qualitatively similar patterns, with the age distribution shifted further toward older technologies. As in Kelly et al. (2021b), I residualize the importance index on issue-year fixed effects before computing the percentile thresholds.

patent (No. 6,285,999) and John Bardeen’s transistor patent (No. 2,524,035).

For each non-breakthrough patent, I locate the closest breakthrough it builds on via a breadth-first search on the citation network. To illustrate, suppose patent A cites $\{B, C\}$, B cites $\{D, E\}$, and C cites $\{F\}$. This means that A ’s first-degree citations are $\{B, C\}$ and its second-degree citations are $\{D, E, F\}$. I assign A to whichever of these is a breakthrough at the lowest citation order. If B and F are breakthroughs, A is assigned to B ; if both B and C are breakthroughs, I split A equally across them. If none of $B, \dots, F$ is a breakthrough, I expand the search to third-degree citations and so on, until finding a breakthrough. The age of the technology on which A builds is the difference between A ’s filing year and its assigned breakthrough’s filing year.

The citation data combines two sources. The USPTO PatentsView database provides citations for all patents granted from 1976 onward ([PatentsView, 2025](#)). To extend coverage to earlier decades, I supplement it with citations from the Comprehensive Universe of U.S. Patents (CUSP) constructed by [Berkes \(2016\)](#), who extract citations from patent texts and front pages going back to the early 19th century. Overall, the breadth-first search described above assigns to a breakthrough approximately 88 percent of all non-breakthrough patents in [Kelly et al. \(2021a\)](#) filed between 1950 and 2009. These comprise the sample of patents analyzed below.

Figure [B.1.1](#) plots the distribution of non-breakthrough patents across breakthrough technologies of different ages, by decade, from the 1950s—soon after citations begin to appear systematically in patent documents—through the 2000s.³⁰ The distributions are hump-shaped, with most patents building on technologies that emerged a few decades earlier. In the 1990s, the distribution shifts notably towards young technologies, driven by the emergence of several influential breakthroughs related to the ICT revolution.

Finally, Table [B.1.1](#) shows how the shares of innovation in very recent breakthroughs are similar to (or even smaller than) the share building on breakthroughs 50 or more years old. This reinforces the importance of old technologies for innovation and the hump-shape pattern.

³⁰Before 1947, there was no formal requirement for patents to disclose references to prior art. After 1947, a mandatory section dedicated to citations was introduced, and the number of citations per patent increased considerably over time since then.




Figure B.1.1: The Hump-Shaped Innovation–age profile (Alternative Technology Definition). *Note:* Each panel plots the share of non-breakthrough patents (within a given decade) building on breakthrough technologies of different ages. Breakthrough patents are defined as those in the top 5% of the residualized forward similarity score measure from [Kelly et al. \(2021b\)](#) and retrieved from [Kelly et al. \(2021a\)](#). Non-breakthrough patents are assigned to their closest breakthrough using a breadth-first search on the citation network. Age is defined as the difference between the patent’s filing year and the breakthrough’s filing year. Citation data combine the USPTO PatentsView database (1976 onward) ([PatentsView, 2025](#)) with historical citations from [Berkes \(2016\)](#).

Table B.1.1: The Hump-Shaped Innovation–age profile (Alternative Technology Definition)

		Share of Non-Breakthrough Patents Building on Breakthroughs of Age:			
		[0,10)	[0,20)	50+	70+
Decade	1950	16%	48%	1%	1%
	1960	11%	31%	5%	1%
	1970	9%	24%	18%	2%
	1980	14%	27%	29%	5%
	1990	27%	45%	27%	10%
	2000	27%	62%	18%	9%

Note: Breakthrough patents are defined as those in the top 5% of the residualized forward similarity score measure from [Kelly et al. \(2021b\)](#) and retrieved from [Kelly et al. \(2021a\)](#). All remaining patents are assigned to breakthroughs using a breadth-first search on the citation network (see the text of this section for details). Age is defined as the difference between the patent’s filing year and the breakthrough’s filing year. Citation data combine the USPTO PatentsView database ([PatentsView, 2025](#)) with historical citations from [Berkes \(2016\)](#).

B.2 CPC technological classification, granular measures of technology, and alternative emergence dates

This section revisits the paper’s main empirical facts using the Cooperative Patent Classification (CPC) system to define technologies. CPC data are obtained from [PatentsView \(2023\)](#).

The CPC system is hierarchical. Classes correspond to broad technological areas, subclasses represent narrower fields within classes, and the most disaggregated units are groups and subgroups. For example, the code B61C11/04 decomposes into class B61 (railways), subclass B61C (locomotives and motor railcars), group 11 (steam locomotives), and subgroup /04 (steam locomotives with steam accumulators). This section treats CPC *subgroups*—the most granular level of the classification—as individual technologies.

First Patent Granted as the Proxy for Technology Emergence First, consider the emergence of a technology subgroup to be the grant year of the first patent classified into it. Because all patents since the early 19th century have been classified under the current CPC scheme, the emergence date of a technology subgroup can predate

the classification's official creation.³¹ The vintage (age) of a technology subgroup in a given decade is the difference between that decade and the decade of its first patent.

The left panel of Figure B.2.1 shows the share of patents granted in the 2000–2009 decade accounted for by technology subgroups of different vintages. Subgroups that emerged in the 1990s or 2000s account for a smaller share of patenting activity than those whose first patents date to the 1980s. Beyond this peak, patent shares decline gradually and eventually stabilize at lower levels for older vintages. A similar hump-shaped pattern is observed in other recent decades (right panel). Over the course of the twentieth century, the peak of patenting activity consistently occurs for vintages between 20 and 40 years old (see Table B.2.1). While still decades old, these peak ages are smaller than those in the main analysis of Figure 2.




Figure B.2.1: The Innovation–age profile using CPC Subgroups (as technologies) and first patent granted (as their emergence date). *Note:* The left panel shows the distribution of patents issued between 2000 and 2009 across CPC subgroups of different ages. Each dot aggregates subgroups with the same estimated age. Age is defined as the difference between the 2000–2009 decade and the birth decade of the subgroup, where birth corresponds to the grant decade of the first patent classified into that subgroup. The vertical axis reports the share of all patents issued in the 2000s accounted for by each technology vintage. The right panel compares the innovation–age profiles across three decades–1990, 2000, and 2010. Each curve represents a local polynomial smooth fitted to the underlying data. The underlying data for both panels come from [PatentsView \(2023\)](#).

Using logistic estimation to date technology emergence. Consider now applying the logistic regression methodology described in Section 5 to CPC subgroups.

³¹The CPC classification scheme was introduced in the 2010s.

Table B.2.1: Peak Technology Vintage by Decade (First-Patent Emergence Proxy)

Decade	Peak vintage (age)	Share of patents
1900	40	0.231
1950	30	0.134
1990	30	0.101
2000	20	0.109
2010	40	0.123

Note: This table reports the technology vintage with the highest patent share in each decade, using CPC subgroups as the unit of technology. Vintage is defined as the difference between the decade of patent grant and the birth decade of the CPC subgroup, where birth is the grant decade of the first patent classified into that subgroup. The underlying data come from [PatentsView \(2023\)](#).

Doing so yields emergence date estimates for approximately 55,000 subgroups, covering about 60% of all patents granted during the twentieth century.³² For each subgroup, I compute its vintage in each decade and the corresponding share of patenting activity. The left panel of Figure B.2.2 plots the distribution of patent shares across technology vintages for patents granted in the 2000s and the 1900s. The main empirical patterns emphasized in the paper—a pronounced hump-shaped innovation–age profile, substantial patenting activity in relatively old technology subgroups, and the stability of the distribution over the twentieth century—are all present. The peak of the distribution continues to occur at vintages several decades old (around 40 years), but it appears earlier than in the main analysis, which relies on more aggregated measures of technology.

³²These figures correspond to the final sample after winsorizing the estimated emergence dates at the 2.5% tails to remove outliers, following the procedure in Section 5.




Figure B.2.2: The Innovation–age profile using CPC Subgroups. *Note:* The blue (orange) line in the left panel shows the share of all the patents in the 2000s (1900s) decade accrued by technologies of different ages. Technologies are CPC subgroups; age is estimated with the logistic methodology described in Section 5. Residual vintages outside $[0, 500]$ are omitted. The right panel presents the coefficients $\{\beta_{age}\}$ estimated from regression (B.1) for positive ages. The underlying data for both panels come from PatentsView (2023).

Within-subclass dynamics. I now study the evolution of a technology’s patenting share within its own CPC subclass. Returning to the example of steam locomotives with steam accumulators (B61C11/04), instead of measuring this technology’s share relative to all patents, I focus on its share within subclass B61C (locomotives and motor railcars).

Consider the following regression specification:

$$\text{share_within_subclass}_{a,sc,t} = \Gamma_t + \Gamma_{sc} + \sum_{age>0} \beta_{age} 1\{a = \text{age}\} + e_{a,sc,t}, \quad (\text{B.1})$$

where $\text{share_within_subclass}_{a,sc,t}$ denotes the patenting share of vintage- a technologies within subclass sc in decade t ; Γ_t and Γ_{sc} are decade and subclass fixed effects; and the coefficients β_{age} capture differences in patenting intensity across technology vintages relative to newly emerged technologies ($age = 0$).

The right panel of Figure B.2.2 plots the estimated β_{age} coefficients. Patenting intensity within narrowly defined technology subclasses exhibits a clear hump-shaped pattern as a function of age.

B.3 Fields of Knowledge Classification from [Berkes et al. \(2022a\)](#)

This section revisits the main empirical facts from Figure 2 of the paper using the Fields of Knowledge classification constructed by [Berkes et al. \(2022a\)](#). The authors cluster International Patent Classification (IPC) technology classes into 164 broader fields based on the propensity of the same inventor to innovate across multiple classes. Fields of knowledge therefore consist of groups of IPC classes that tend to be jointly patented by the same inventors. Examples include ‘Vehicle Brake Control Systems or Parts Thereof’ and ‘Chemical Treatment of Skins.’

To link patents to fields of knowledge, I use the official CPC–IPC crosswalk ([Cooperative Patent Classification Committee, 2025](#)). To proxy for the emergence date of each field, I apply the same logistic estimation procedure described in Section 5 of the paper.

Figure B.3.1 presents the results. The left panel shows the distribution of patents across knowledge fields of different proxied ages for the 2000–2009 decade. The distribution exhibits a clear hump shape, with a peak around fields that are nearly a century old, closely mirroring the pattern documented in Figure 2 of the paper. The right panel compares this distribution to the corresponding one for the 1900–1909 decade and shows that both the shape and the location of the peak are remarkably stable over the course of the twentieth century.




Figure B.3.1: The Innovation–age profile using fields of knowledge from [Berkes et al. \(2022a\)](#). *Note:* The left panel shows the distribution of patents issued between 2000 and 2009 across fields of knowledge of different proxied ages. The fields of knowledge classification is constructed by [Berkes et al. \(2022a\)](#) and available at [Berkes et al. \(2022b\)](#). The age of each field is computed using the logistic estimation procedure described in Section 5. The solid line corresponds to a local polynomial smooth fitted to the underlying data. The right panel reproduces the same distribution from the left panel and overlays the corresponding smoothed distribution for patents issued between 1900 and 1909. The patent data underlying both panels come from [PatentsView \(2023\)](#).

B.4 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 ([USPTO, 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 Commerce, 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 (USPTO, 2009). However, many current classes were introduced as relabels of previously existing ones, involving only minor organizational changes. Take, for instance, the class 297 ‘Chairs and seats,’ established only in 1961, an implausibly late 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 No. 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 (USPTO, 1838); (ii) the editions of 1872 and 1895 of the Classification Index of Subjects of Invention (USPTO, 1872–1895); (iii) the editions of 1912, 1916, 1920, 1923 and 1947 of the Manual of Classification of Subjects of Invention (USPTO, 1912–1947);³³ and (iv) the Classification Bulletin editions in the period 1912–1940 (not all volumes) (USPTO, 1912–1940). Sources (i)–(iii) provide a snapshot of the complete classification code at a given point in time, and (iv) provides a flow of new class creation.

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

³³Other editions of the Manual of Classification of Subjects of Invention and the Classification Index of Subjects of Invention, to the best of my knowledge, are not available in US libraries.

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 be the same as 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 a 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.³⁴ 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.³⁵ 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:

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

³⁴Recall the USPC was discontinued in 2013, from when the latest version is.

³⁵Notice 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.




Figure B.4.1: Technology Classes: Textually Identified Emergence Dates. *Note:* The left panel 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 are described in the text of this Section. The right panel 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. In the right panel, the underlying patent data come from the USPTO Historical Patent Data Files ([Marco et al., 2015](#); [USPTO, 2015](#)).

Results The left panel in Figure B.4.1 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 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 B.4.1 shows the share of patents attributed to these old technologies. Consistent with the discussions in Sections 2 and 5, 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:

$$PatentShare_{t,\tau} = \sum_a \beta_a 1_{\{t-\tau=a\}} + \Gamma_\tau + e_{t,\tau} \quad (B.2)$$

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 on a technology's patenting share. While there are no time fixed effects due to collinearity, the fact that we are using patent shares as the dependent variable already normalizes the observations by the aggregate trend in patents.

Figure B.4.2 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.




Figure B.4.2: Technology class and innovation. *Note:* The vertical axis displays the estimated coefficients for age dummies in regression (B.2), along with their 95% confidence intervals. Standard errors are computed using a heteroscedasticity-consistent covariance matrix. The sample comprises the technologies whose class creation is dated to 1900 or later. The documents, data, and procedure to identify technologies and their age are described in the text of this Section. The data on patents come from the USPTO Historical Patent Data Files (Marco et al., 2015; USPTO, 2015). Observations are trimmed at the bottom and top 1% of the patenting share to remove outliers.




Figure B.4.3: 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. Each dot groups technologies of the same estimated age. The documents, data, and procedure to identify technologies and their age are described in the text of this Section. The sample comprises the technologies whose class creation is dated to 1900 or later. The data on patents come from the USPTO Historical Patent Data Files (Marco et al., 2015; USPTO, 2015). Observations are trimmed at the bottom and top 1% of the patenting share to remove outliers.

Finally, still focusing on 20th-century technologies, Figure B.4.3 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 B.4.1). 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.

C Measurement and Calibration

C.1 Additional details: ϵ calibration

Table C.1.1 presents the results from the estimation of regression (29). In Column (1), the estimated coefficient on $\log \bar{v}(t|\tau)$ is statistically significant and equal to 0.381. In words, on average, a one percent 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 percent 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.

C.2 Logistic trend estimation

The next paragraphs provide additional details on the logistic trend estimation in Section 5.

Table C.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 (0.190) $\epsilon = 0.724$	0.690 (0.170) $\epsilon = 0.591$
Year FE	✓	✓
Tech FE	✓	✓
N	205	187
R^2	0.024	0.101
F -Stat	4.040	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. The age of technologies is computed by the logistic trend fit described in Section 5 and comprises only technologies whose R-squared in the logistic regression is greater than 0.5. Outliers with emergence dates beyond 2023 were excluded. Column (1) uses patent valuations (the explanatory variable) from period t , while Column (2) uses their one-period lag. Patent valuation data are from Kogan et al. (2017) and retrieved from Kogan et al. (2022). The data on patents' technology class come from the USPTO Historical Patent Data Files (Marco et al., 2015; USPTO, 2015).

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

$$pat_share(t) = \frac{K}{1 + e^{-(\alpha + \beta t)}} \quad (C.1)$$

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 right-hand side 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 (C.1), one can solve for t^* .

C.3 Additional details: Figure 1

Figure 1 plots the annual count of U.S. steam engine patents. The underlying data come from Google Patents, accessed through the BigQuery public patents dataset (Google, 2023). I identify steam engine patents by keyword search. A patent qualifies if its title or description contains the term ‘steam engine.’ I then exclude patents that also reference competing technologies, namely combustion engines, internal combustion engines, and wind. The search is case insensitive and allows for variations in spacing and plural forms, so that ‘steam engine,’ ‘steam-engine,’ ‘steamengine,’ and ‘steam engines’ are all captured. I further restrict the sample to patents whose primary classification falls under CPC class F01, ‘Machines or Engines in general, Steam Engines,’ as this classification comprehensively encompasses technologies related to steam engines. I keep only U.S. patents published on or before 1970, and I count each publication number once to avoid double counting. The figure reports the resulting counts by publication year.

D Additional Figures and Tables

Table D.1: Robustness: Counterfactual Policy under Alternative Parameter Values

γ	$\frac{C_{\text{eff}}^{BGP}}{C_{\text{dec}}^{BGP}}$	Welfare Gain	μ_0	$\frac{C_{\text{eff}}^{BGP}}{C_{\text{dec}}^{BGP}}$	Welfare Gain
0.015	1.63	4.65%	0.25	1.89	10.77%
<i>0.02</i>	<i>1.61</i>	<i>7.33%</i>	<i>0.55</i>	<i>1.61</i>	<i>7.33%</i>
0.025	1.59	9.74%	0.85	1.47	5.65%

Note: This Table reports the long-run consumption ratio between the efficient and decentralized BGP trajectories ($C_{\text{eff}}^{BGP}/C_{\text{dec}}^{BGP}$) and the consumption-equivalent welfare gain from the efficient transition path starting at the decentralized BGP at $t = 0$. The Table’s left panel varies γ and the right panel varies μ_0 , in each case holding all remaining parameters at their baseline calibration values (Table II). Rows in italics denote the baseline calibration.




Figure D.1: The Innovation Life Cycle of the Steam Engine (Quality Adjusted Patents). *Note:* This Figure displays the weighted annual count of patents containing keywords related to the Steam Engine technology. Patent textual data are sourced from the Google Patents database (Google, 2023); see Section C.3. Each patent is weighted by its respective importance index (*kpst_5*) from Kelly et al. (2021b) and retrieved from Kelly et al. (2021a).




Figure D.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. (2021b) and sourced from Kelly et al. (2021a). The graphs show the distribution of quality-weighted patents across technologies of different ages. See Section 5 for more details. The data on patents' technology class come from the USPTO Historical Patent Data Files (Marco et al., 2015; USPTO, 2015).




Figure D.3: Patenting share and proxied emergence date. *Note:* The dots in each panel 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. The solid line represents the fitted logistic trend whose estimation is described in Section 5. The dashed lines represent the year in which the trend achieved 10% of the ceiling value. The underlying patent data are from the USPTO Historical Patent Data Files (Marco et al., 2015; USPTO, 2015).




Figure D.4: The cross-sectional stationary distributions with $\bar{\lambda} = 1$. *Note:* The figure presents the illustrative calibration discussed in Section 3.