Did R&D Misallocation Contribute to Slower Growth?*

Nils H. Lehr†

IMF

May 18, 2026

Abstract

This paper provides evidence that rising misallocation in the R&D sector contributed to the recent slowdown in U.S. productivity growth. I develop a growth accounting framework allowing for misallocation of R&D resources across firms captured by wedges between their marginal costs and benefits of R&D. I show that R&D wedges can be measured from R&D returns and document large and persistent differences in R&D returns across U.S.-listed firms. Combining data and model, I estimate that frictions reduced productivity growth by 11% over 1975–2014 and that rising misallocation in the R&D sector accounts for 32% of the growth slowdown.

*The views expressed in this paper are those of the author and do not necessarily represent those of the IMF, its Executive Board, or IMF management. I benefited from the advice and comments of many including Stephen J. Terry, Pascual Restrepo, Tarek A. Hassan, David Lagakos, Stefania Garetto, Adam Guren, Maarten De Ridder, and the participants of the Macro Lunch Seminar, Green Line Macro Meeting, Econometric Society European Winter Meeting, and the Global Entrepreneurship and Innovation Research Conference.

†Email: nlehr@imf.org

U.S. total factor productivity (TFP) growth slowed down significantly in the last two decades. While TFP grew 0.5% per year in 1975–1995, its growth rate declined to 0.3% in 2005–2018. Surprisingly, investment in research & development (R&D)—commonly considered the driver of medium- and long-run productivity growth—remained stable over the same time horizon. R&D expenditure amounted to 2.7% of GDP for 1975–95, compared to 2.9% for 2005–18.

Figure 1: Declining Growth Despite Stable R&D Expenditure

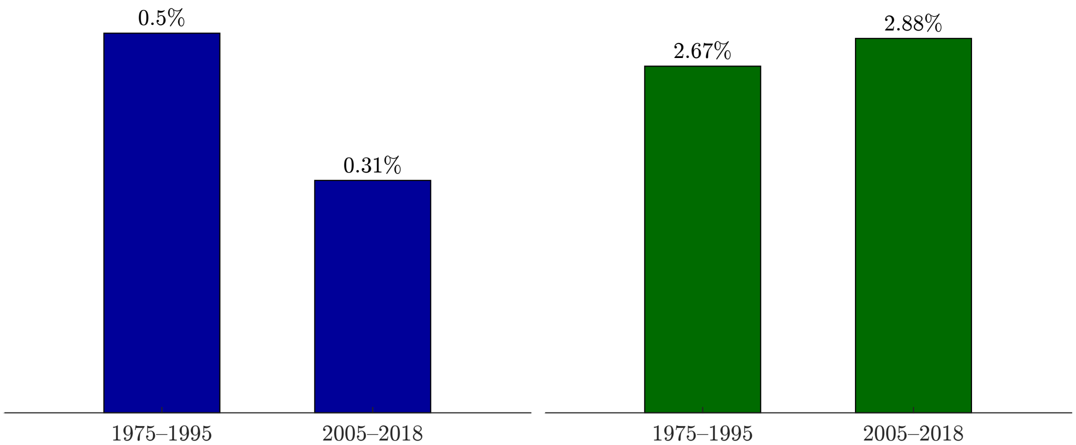

(a) TFP Growth Rate (b) R&D Expenditure (% of GDP)

Notes: Growth rate of Total Factor Productivity at Constant National Prices calculated using data from Penn World Table 10. R&D Expenditure as a share of Gross Domestic Product calculated using data from Bureau of Economic Analysis.

I provide evidence for an explanation of declining TFP growth at constant R&D investment rooted in a simple decomposition. In endogenous growth models, productivity growth is the product of two terms: aggregate R&D investment and aggregate R&D productivity—the rate at which these investments are translated into growth. Slower growth despite stable R&D investment can then only be rationalized by declining aggregate R&D productivity.

In turn, aggregate R&D productivity is shaped by two forces: the average R&D productivity of firms and the efficiency with which R&D resources are allocated across them, or *R&D Allocative Efficiency*. A growing literature highlights the first channel and finds that ideas are getting harder to find, inventors are becoming less productive, or ideas are less likely to feed into productivity (Bloom et al., 2020; Ekerdt and Wu, 2025; Fort et al., 2026). Instead, this paper focuses on the second channel and provides evidence for a significant decline in R&D Allocative Efficiency. While some firms invest too much in R&D relative to the inventions they produce, others do too little, and increasingly so. Quantitatively, this channel can account for 32% of the TFP growth slowdown. Thus, not only are ideas getting harder to find, but we are also increasingly looking for them in the wrong places.

I reach these conclusions based on a growth accounting framework nesting standard workhorse growth models (Romer, 1990; Aghion and Howitt, 1992). R&D labs hire inventors to maximize the private value created from innovation; however, their choices may be distorted through exogenous *R&D wedges*. These wedges prevent the equalization of marginal returns on R&D across firms, as predicted by frictionless and competitive R&D markets, and, thus, conveniently summarize the impact of a wide range of distortions such as financial frictions, adjustment costs, or market power on firms' R&D choices. Productivity growth occurs as a by-product of innovation; however, the private value created from an invention may not align perfectly with its contribution to productivity growth. I capture this divergence with an *impact-value factor* such that firms with a low impact-value factor create a lot of private value relative to their contribution to productivity growth.

I show that the impact of private frictions, i.e., *R&D wedges*, on TFP growth is captured by a sufficient statistic, which I label *R&D Allocative Efficiency*. When private and growth incentives are aligned, variation in R&D wedges reduces allocative efficiency. Intuitively, differences in marginal R&D returns due to R&D wedges imply that growth could be accelerated by redistributing R&D resources from low to high marginal R&D return firms. Heterogeneity in *impact-value factors* can amplify, dampen, or even overturn this result. If firms with low impact-value factors also have low R&D wedges, then misallocation is worsened as R&D wedges push firms that already invest too much in R&D from a growth-maximizing perspective to do even more. In contrast, the growth-maximizing R&D policy uses R&D wedges to offset differences in impact-value factors.

I consider several extensions. First, my results extend to a framework with multiple research lines per firm. Second, frictions reduce firm values and are thus costlier under free entry. Third, imperfect substitutability of R&D inputs across firms reduces the cost of frictions as there are fewer gains from input reallocation. Fourth, the level of frictions has a direct effect on growth under a positive aggregate R&D supply elasticity, while only relative frictions matter under a fixed aggregate supply. Lastly, knowledge externalities, which constitute a separate source of diverging incentives between firms and the planner, amplify the impact of declining R&D Allocative Efficiency when they are aligned with inventions' productivity impact.

Next, I estimate the model primitives from firms' patents and financial statements for a sample of U.S.-listed firms for 1975–2014. I measure firms' R&D investment using their expenditure and the resulting private value created using patent valuations. In the model, the ratio of value created to investment, which I refer to as *R&D return*, provides a direct measure of R&D wedges. I experiment with a range of proxies for impact-value factors based on profitability or patent quality measures.

At the micro-level, I find large and highly persistent differences in measured R&D returns across firms—in stark contrast to the prediction of the frictionless model, where firms equalize the marginal benefit to the marginal cost of R&D and, thereby, the R&D return as well. This finding closely parallels the misallocation literature in the production sector, which interprets dispersion in the return on capital as evidence of capital misallocation (David et al., 2016). The standard deviation of R&D returns is 45% larger than its counterpart for the return on capital, suggesting substantial R&D misallocation. Moreover, the strong persistence of R&D returns—with an implied annual autocorrelation coefficient around 0.9—is consistent with the idea that these differences reflect structural factors rather than statistical noise.

Nonetheless, R&D returns are surprisingly hard to explain using standard empirical proxies for frictions. They are uncorrelated with the return on capital, suggesting that R&D and capital investments are subject to different frictions, and they are not consistently correlated with measures of financial frictions. Instead, the strongest predictors are R&D employment—consistent with monopsony power that increases with firm size—and indicators of firm expansion, consistent with adjustment frictions. The relationship between R&D wedges and proxies for the impact-value factor is ambiguous, with most measures pointing towards a weakly positive correlation.

At the macro level, I estimate that TFP growth is significantly slower due to low R&D Allocative Efficiency, and increasingly so over time. For the full sample, dispersion in R&D wedges implies a 21% lower growth rate. For comparison, Hsieh and Klenow (2009) estimate that U.S. productivity would be 40% higher under the first-best factor allocation, while Berger et al. (2022) find that output would rise by 21% in the absence of monopsony in the production sector. Achieving the frictionless growth rate might not be feasible if R&D wedges reflect technological or information frictions that cannot be eliminated. I thus benchmark the aggregate estimates against the Life Sciences sector—for which I estimate stable R&D Allocative Efficiency throughout my sample—and obtain a preferred estimate of an 11% lower growth rate attributable to excess R&D misallocation.

Comparing 1975–90 to 2000–14, I find that declining R&D Allocative Efficiency can account for 13% slower growth in my preferred specification, which represents 32% of the $\frac{0.5\%-0.3\%}{0.5\%} \approx 40\%$ slowdown. While R&D Allocative Efficiency achieves benchmark levels in the early sample, it declines significantly during the 1990s without a subsequent recovery.

An important concern for these estimates is the presence of measurement error, which could explain why a large share of variation in R&D returns is not captured by direct measures of frictions. Measurement error mechanically reduces measured R&D Allocative Efficiency and, to the extent that it is time-varying, may also contribute to its observed trend. I thus conduct an extensive set of robustness exercises, including direct adjustments to the measurement of R&D inputs and output, statistical corrections such as the approach proposed in Bils et al. (2021), and alternative proxies for the impact value factor. Across these exercises, the main conclusions hold: even the most conservative estimates imply an 8% decline in R&D Allocative Efficiency, accounting for roughly 20% of the total decline in TFP growth. In addition, the allocation of R&D resources is worse—and deteriorating more rapidly—among smaller firms, which are relatively under-represented in my sample. At the industry level, R&D Allocative Efficiency predicts R&D expenditure and R&D returns in line with the model's predictions, providing direct evidence that it captures relevant aggregate dynamics. Taken together, these results indicate that slower TFP growth is partly driven by rising misallocation in the R&D sector.

Literature. This paper contributes to three strands of the literature. First, I contribute to the growing literature on the recent slowdown in economic growth by highlighting the importance of private frictions. Like Akcigit and Ates (2021), Olmstead-Rumsey (2025), and Ekerdt and Wu (2025), I argue that declining aggregate R&D productivity is a core driver; however, I attribute it to rising misallocation instead of declining micro-level R&D productivity or knowledge spillovers.[1] This perspective is closely related to de Ridder (2024), Aghion et al. (2023) and Ayerst (2023), who propose models in which rising misalignment between the private incentives for R&D and its growth impact leads to R&D misallocation and, thereby, a slowdown in economic growth. Instead, I focus on the contribution of private frictions, captured by R&D wedges, and follow a sufficient statistic approach allowing for a direct mapping between data and model, rather than relying on structural estimation.

Second, I provide a new framework to study the drivers of aggregate R&D productivity. The early endogenous growth literature identifies innovation as the main driver of economic growth and highlights the under- or over-provision of innovation due to externalities (Romer, 1990; Aghion and Howitt, 1992). Recent contributions study the distribution of R&D resources across firms, which might be inefficient with heterogeneity in spillovers or in firms' ability to benefit from inventions (de Ridder, 2024; Mezzanotti, 2021; Akcigit et al., 2022;

[1] Bloom et al. (2020) also argue that aggregate R&D productivity has declined; however, their focus is a long-run, steady decline in R&D productivity as "ideas are getting harder to find," in line with the predictions of semi-endogenous growth theory (Jones, 1995).

Manera, 2022; Aghion et al., 2025, 2023). My framework is closely connected to this literature but differs in several key dimensions. Most importantly, I allow for private frictions and estimate that they have a significant impact on economic growth. Furthermore, I show that these frictions interact with incentive misalignment between private value and growth, so that estimating the growth impact of either force requires a joint treatment.

Third, I contribute to the literature on factor misallocation by providing evidence on its pervasiveness in the R&D sector. Restuccia and Rogerson (2008) and Hsieh and Klenow (2009) first identified factor misallocation across firms—as captured by heterogeneity in factor returns—as an important determinant of aggregate productivity. The subsequent literature documents that dispersion in the return on capital is surprisingly difficult to attribute to individual mechanisms, which I also find for R&D returns, and that the persistence of frictions is striking in its own right (Banerjee and Moll, 2010; Asker et al., 2014; Midrigan and Xu, 2014; David et al., 2016). I complement this work by focusing on R&D investment rather than static production factors, thereby linking factor return heterogeneity to the productivity growth rate rather than its level. I further document that the sources of frictions differ for capital and R&D investment, as their returns are uncorrelated. I find limited evidence on the contribution of channels emphasized in the existing literature, such as government subsidies or financial frictions. In contrast, König et al. (2022) find that product-market frictions led to an inefficient R&D allocation in China, and Brown et al. (2009) argue that financial frictions are particularly severe for intangible investments such as R&D. My focus on large U.S. firms is a likely reason for the contrasting findings. In my data, the strongest predictor of R&D returns is inventor employment, which is consistent with increasing monopsony power exerted by large employers (Berger et al., 2022). Such a channel could explain rising frictions in light of increasing concentration (Autor et al., 2020).

1 Theory

This section introduces the theoretical framework for assessing the impact of frictions in the R&D sector on economic growth. Motivated by the subsequent empirical analysis, the model takes as given elements that will be measured directly in the data and is fully consistent with alternative growth theory traditions (Romer, 1990; Aghion and Howitt, 1992; Jones, 1995).[2]

[2] Online Appendix H develops an expanding variety growth model in line with the formulas developed in this section and confirms that a Schumpeterian microfoundation is equally compatible with the main results.

1.1 Model Setup

Time is infinite, discrete, and indexed by t. **Output** Y_t is a function of aggregate productivity A_t and production labor $L_{P,t}$, which is supplied inelastically:

$$Y_t = A_t\, L_{P,t}. \tag{1}$$

Aggregate productivity grows as the result of inventions made by a unit mass of **R&D labs**, indexed by $i \in [0,1]$. R&D labs have heterogeneous R&D productivity φ_{it} and hire R&D input ℓ_{it} to achieve mass z_{it} of innovations subject to the R&D production function

$$z_{it} = \varphi_{it}\, \ell_{it}^{\gamma} \quad \text{with} \quad \gamma \in (0,1). \tag{2}$$

R&D labs pay price W_t for hired inputs and assign value V_{it} to innovations. I take the value of innovations as given, since I measure firms' gains from innovation directly in the data.[3] In workhorse growth models, this value is linked to resulting profits and innovation opportunities, as discussed in Online Appendix G.

R&D labs' equilibrium input level ℓ_{it}^* is potentially distorted by R&D wedge Δ_{it}, s.t.

$$\left.\frac{\partial z_{it}}{\partial \ell_{it}}\right|_{\ell_{it}=\ell_{it}^*} V_{it} = (1+\Delta_{it})\, W_t. \tag{3}$$

The left-hand side is the marginal benefit of research input, while the right-hand side is the marginal cost adjusted for the R&D wedge. We recover the frictionless benchmark in which R&D labs equalize marginal benefit and cost with $\Delta_{it} = 0$. Otherwise, quantities are distorted with larger wedges resulting in lower R&D input levels.

R&D wedges can capture a range of distortions including financial frictions, adjustment costs, capacity constraints, market power in the R&D input market, and R&D subsidies, as discussed in Online Appendix G. For example, high R&D wedges arise when labs' R&D levels are constrained by financial frictions or adjustment costs. Conversely, R&D subsidies result in low R&D wedges as they reduce labs' marginal cost below the market price.

Market clearing requires that R&D labs' demand for R&D resources equals their

[3] In my counterfactuals, I assume this value to be exogenous up to a common factor, which could reflect any endogenous response of labs' discount rate to the aggregate productivity growth rate.

inelastic supply L_t:[4]

$$L_t = \int_0^1 \ell_{it}\, di. \tag{4}$$

Lastly, R&D labs' innovations determine the **growth rate** of aggregate productivity A_t. Innovations contribute at a rate of $V_{it}\, \zeta_{it}$, where impact-value factor ζ_{it} measures the growth contribution per dollar of private value created for a given firm. Productivity growth g_t is then the total private value created from innovation adjusted for the impact-value factor:

$$g_t \equiv \frac{A_{t+1} - A_t}{A_t} = A_t^{-\phi} \int_0^1 z_{it}\, V_{it}\, \zeta_{it}\, di. \tag{5}$$

The parameter $\phi \geq 0$ is the "fishing-out" effect that is necessary to achieve balanced growth in a semi-endogenous growth framework with R&D input growth (Jones, 1995).

Introducing the impact-value factor is necessary technically and conceptually. From a technical perspective, the TFP growth rate is unitless, while private value is measured in dollars. Thus, it is necessary to have an "exchange-rate" between the private value and growth impact. From a conceptual perspective, the term is important to capture potential misalignment between private value and productivity impact arising from differences across firms in their ability to profit from inventions of a given quality. For example, in Schumpeterian models with limit pricing, profits from an invention of step-size λ scale with $\frac{\lambda-1}{\lambda}$, while the productivity impact scales with $\lambda - 1$, such that step-size heterogeneity leads to a misalignment of profit and productivity impact (Aghion et al., 2014). Similarly, differences in firms' ability to protect their intellectual property, arising, e.g., due to their technology focus or the quality of patent lawyers they hire, will be reflected in the present discounted value of the invention independently of their growth impact (Mezzanotti, 2021). Importantly, the impact-value factor does not capture knowledge spillovers. I discuss these together with other extensions below and additional mechanisms in Online Appendix G.[5]

The definition of the growth rate encompasses a wide range of models in the literature. For example, in expanding variety models à la Romer (1990), the growth rate is the sum of new ideas, i.e., z_{it}, divided by the stock of new ideas, which we can recover by setting

[4] This assumption reflects the idea that R&D talent is scarce (Goolsbee, 2003; Wilson, 2009; Ekerdt and Wu, 2025). An alternative interpretation is that R&D policy already optimizes the size of the R&D sector, such that the allocation of resources within it remains the relevant margin of concern. The framework abstracts from any waste of R&D inputs linked to R&D wedges, e.g., due to "real" adjustment cost. Any adjustment costs are assumed to take the form of production labor or cash payments.

[5] de Ridder (2024) and Aghion et al. (2023) introduce models in which firms' profits from innovation differ due to heterogeneous markups, while König et al. (2022) consider product market wedges. Manera (2022) and Mezzanotti (2021) emphasize differences in intellectual property rights protection, while Akcigit and Ates (2021) and Olmstead-Rumsey (2025) investigate differences in knowledge diffusion over time.

$\zeta_{it} = 1/(V_{it}\ A_t)$. Similarly, in Schumpeterian growth models à la Aghion et al. (2014), the growth rate is the aggregate innovation rate of firms, i.e., z_{it}, adjusted for their productivity improvement $\lambda_{it} - 1$, which we can recover with $\zeta_{it} = (\lambda_{it} - 1)/V_{it}$. Note, however, that the framework does not fully capture the dynamics of models with endogenously evolving R&D productivity φ_{it} as in Klette and Kortum (2004), as discussed in the extensions below.

I use two simplified solution concepts in deriving the main results. The **Competitive Equilibrium** respects the equations detailed above, while the **Growth Maximizing Allocation** allocates R&D inputs to maximize growth.

Definition 1. *For given $\{Y_0, \{L_t, \{V_{it}, \varphi_{it}, \Delta_{it}, \zeta_{it}\}_{i\in[0,1]}\}_{t=0,\dots,\infty}\}$, a Competitive Equilibrium is a sequence $\{\{\ell_{it}\}_{i\in[0,1]}, W_t, g_t, Y_t\}_{t=0,\dots,\infty}$ satisfying (1)-(5).*

Definition 2. *For a given $\{Y_0, \{L_t, \{V_{it}, \varphi_{it}, \zeta_{it}\}_{i\in[0,1]}\}_{t=0,\dots,\infty}\}$, a Growth Maximizing Allocation is a sequence $\{\{\ell_{it}\}_{i\in[0,1]}, g_t, Y_t\}_{t=0,\dots,\infty}$ maximizing economic growth $\{g_t\}_{t=0,\dots,\infty}$ period-by-period and satisfying equations (1), (2), (4), and (5).*

1.2 Results

Proposition 1 establishes that the equilibrium economic growth rate can be decomposed into three terms. The first term reflects the economy's frontier growth rate at the growth-maximizing allocation. The second term, which I refer to as the *Incentive Alignment Index*, measures how closely private innovation value is aligned with the growth impact of R&D, as captured by the distribution of impact-value factors. Intuitively, when firms with high private innovation value also have high growth impact, private R&D incentives are closely aligned with the growth-maximizing allocation. The first two terms jointly characterize the growth rate in a frictionless, competitive equilibrium without R&D wedges. The third term, which I label *R&D Allocative Efficiency*, captures the impact of R&D wedges.[6]

Proposition 1. *Under equations (2)-(5), we can express the economic growth rate in a Competitive Equilibrium as the product of three terms:*

$$g_t = \underbrace{\frac{L_t^{\gamma}}{A_t^{\phi}} \left(\int_0^1 (\theta_{it}\,\zeta_{it})^{\frac{1}{1-\gamma}}\,di\right)^{1-\gamma}}_{=\ \textit{Frontier Growth Rate } g_t^F} \underbrace{\left(\int_0^1 \omega_{it}\,\tilde{\zeta}_{it}^{\frac{1}{1-\gamma}}\,di\right)^{\gamma-1}}_{\equiv\ \textit{Incentive Alignment } \Lambda_t} \underbrace{\frac{\int_0^1 \omega_{it}\,\tilde{\zeta}_{it}\,(1+\Delta_{it})^{-\frac{\gamma}{1-\gamma}}\,di}{\left(\int_0^1 \omega_{it}\,(1+\Delta_{it})^{-\frac{1}{1-\gamma}}\,di\right)^{\gamma}}}_{\equiv\ \textit{R\&D Allocative Efficiency } \Xi_t},$$

[6]Note that R&D Allocative Efficiency is independent of the average level of R&D wedges. Intuitively, because aggregate R&D supply is fixed, any excess or shortfall in aggregate R&D demand is absorbed by the R&D input price rather than affecting economic growth directly.

where $\tilde{\zeta}_{it} \equiv \zeta_{it} / \left(\int_0^1 \omega_{it} \, \zeta_{it} di \right)$ *and* $\omega_{it} \equiv \theta_{it}^{\frac{1}{1-\gamma}} / \left(\int_0^1 \theta_{it}^{\frac{1}{1-\gamma}} di \right)$ *are the normalized impact-value factor and an R&D productivity weight, with R&D productivity defined as* $\theta_{it} \equiv \varphi_{it} \, V_{it}$*.*

The economics captured by *R&D Allocative Efficiency* are best understood when considering the case of constant, or independently distributed, impact-value factors such that constrained firms do not systematically differ in their growth impact per unit of private value. Then, dispersion in R&D wedges strictly reduces economic growth as per Corollary 1. Intuitively, R&D wedges govern the marginal return on R&D. When R&D labs differ in their marginal returns, one can raise the aggregate return by moving resources from low to high marginal R&D return labs. Larger dispersion then suggests ever more unrealized opportunities for gainful reallocation and, thus, more misallocation—a growth rate analogue of the static misallocation measure developed in Hsieh and Klenow (2009).

Corollary 1. *Suppose* $\zeta_{it} \perp (1+\Delta_{it})^{-\frac{\gamma}{1-\gamma}}$*, then R&D Allocative Efficiency is characterized by the joint distribution of R&D wedges and productivity:*

$$\Xi_t = \frac{\int_0^1 \omega_{it} \, (1+\Delta_{it})^{-\frac{\gamma}{1-\gamma}} \, di}{\left(\int_0^1 \omega_{it} \, (1+\Delta_{it})^{-\frac{1}{1-\gamma}} \, di \right)^{\gamma}}.$$

Up to a 2^{nd}*-order approximation, it is strictly decreasing in the dispersion of R&D wedges and achieves a maximum of* 1 *if they are equalized.*

Relaxing the assumption of orthogonal impact-value factors qualifies this result, as established in Proposition 2. Depending on their correlation with R&D wedges, impact-value factors can either amplify or dampen the impact of R&D wedge dispersion. Under positive correlation, i.e., if particularly constrained firms also achieve a high growth impact per private value created, impact-value factors amplify the misallocation caused by R&D wedges. Conversely, they dampen this effect in the case of (weak) negative correlation, i.e., if constrained firms have lower growth impact per private value created.

Proposition 2. *Let the* ω_{it}*–weighted covariance of log R&D wedges and impact-value factors,* $\sigma_{\Delta,\zeta}$ *be greater than minus half the* ω_{it}*–weighted variance of log R&D wedges,* σ_{Δ}^2 *and define* $\eta = \sqrt{1 + 2 \, \sigma_{\Delta,\zeta} / \sigma_{\Delta}^2}$*. Then, up to a* 2^{nd}*-order approximation,*

$$\Xi_t = \frac{\int_0^1 \omega_{it} \, (1+\Delta_{it})^{-\frac{\gamma}{1-\gamma} \, \eta} \, di}{\left(\int_0^1 \omega_{it} \, (1+\Delta_{it})^{-\frac{1}{1-\gamma} \, \eta} di \right)^{\gamma}}.$$

Two examples illustrate the underlying economics. Suppose young firms' R&D labs are financially constrained and less skilled at capturing rents from their inventions—e.g., because of fewer legal resources to defend their patents—such that they have large R&D wedges and impact-value factors. Then, R&D wedges are more costly for growth since such firms would have underinvested in R&D even absent wedges and are pushed still further from the optimum. Alternatively, suppose large firms' R&D labs have market power over inventors and excel at capturing rents from inventions, such that labs with large R&D wedges (from monopsony power) tend to have low impact-value factors. Then, market power pushes R&D workers away from low impact-value-factor firms and, thus, is less costly for growth.

R&D wedges are closely linked to optimal growth policy, as shown in Proposition 3. The planner could maximize growth by neutralizing variation in the product of impact-value factors and R&D wedges with firm-specific R&D subsidies. The resulting allocation achieves the frontier growth rate by setting $\Lambda_t \; \Xi_t = 1$, so that any $\Lambda_t \; \Xi_t < 1$ implies a growth rate below the growth possibility frontier. This result clarifies the role of Λ_t in the decomposition: it captures the frictionless alignment between private innovation value and growth impact, while the paper's empirical focus is on the losses from R&D wedges, summarized by Ξ_t.

Proposition 3. *Let g_t^* be the growth rate at the growth-maximizing allocation. Then $g_t^* = g_t^F$. This growth rate can be achieved by setting the R&D subsidy component of Δ_{it} to equalize $\zeta_{it}\,(1 + \Delta_{it})$ across firms.*

1.3 Extensions

I consider several extensions in Online Appendix D. First, allowing for **entry** can amplify the cost of frictions. R&D wedges reduce labs' profits by raising the equilibrium input price W_t and, thus, depress entry. Having fewer operating labs implies more inventors per lab and, due to decreasing returns to scale, lower productivity.

Second, the formulae extend to frameworks with **multi-research-line firms** as in Klette and Kortum (2004); however, the counterfactual holds constant the distribution of research lines across firms. The estimated growth impacts are conservative if R&D labs that expand absent frictions also tend to be more productive.

Third, recent work argues that amenities or **specialization** make firms imperfect substitutes for workers, which limits gains from reallocation (Card et al., 2018; Berger et al., 2022). The formulae are preserved in this case with specialization reflected in a lower effective scale elasticity γ, such that R&D wedges tend to be less costly.

Fourth, the model assumes fixed aggregate R&D inputs, implying that the level of R&D

wedges does not affect growth. The formulae extend to the case with positive **input supply elasticity**; however, there is a supply adjustment term depending on the average R&D wedge. Reallocation gains are unchanged if this average remains constant. The formulae also extend to **multiple R&D production factors** as long as their supply is equally inelastic and frictions are common at the firm-level or apply to only one input factor. In the latter case, the relevant elasticity for the marginal impact of misallocation becomes the scale elasticity of the misallocated factor.

Lastly, **knowledge externalities** appear distinctly from the impact-value factor in the model and provide a separate rationale for divergence between the planner allocation and the competitive equilibrium. As I show in Online Appendix D, the formulae derived above constitute a lower bound for the growth impact of frictions when knowledge externalities are fully aligned with the productivity impact, as commonly assumed in the literature.

2 Data and Measurement

2.1 Data

I assume that an R&D lab is equivalent to a firm and focus my empirical analysis on research-active firms listed on U.S. stock exchanges. The latter choice is primarily motivated by the availability of sufficient data to measure the model primitives and, thus, estimate R&D Allocative Efficiency. I discuss selection and measurement concerns in Section 4.

I obtain annual firm-level R&D expenditure from WRDS Compustat, which collects the data from mandatory filings. This data includes firms' industry classification and accounting data such as profits, sales, and employment. I adjust profits for SG&A expenditure to capture flow profits linked to the firm's production activity and its inventions.

I use patents granted by the U.S. Patent and Trademark Office (USPTO) to measure R&D outputs. Patents are arguably the most direct measure of R&D output available to researchers. A patent captures an invention deemed new and useful by the patent office, and grants the owner exclusive rights to its use, giving firms strong incentives to patent. Nonetheless, patents may present an incomplete picture as not all inventions are patented (Cohen et al., 2000; Mezzanotti and Simcoe, 2025). To address this concern, I focus on firms that tend to patent and consider measures independent of the patent system for robustness.

I use patent valuation estimates from Kogan et al. (2017) to measure the private value created from innovation. Their methodology uses the firm's stock returns around the patent grant announcement to estimate the patent's value such that larger returns translate into higher valuations. Patent valuations capture the private value of an invention, which is directly linked to firms' incentives to innovate. In contrast, other patent-based measures of innovation, such as raw counts or citations, capture the quantity of innovation, but not its value to the firm. As discussed above, divergence between the two concepts is an important object of interest when estimating the aggregate impact of private frictions.

I consider two measures of patent quality as proxies for their growth impact: the text-based measure from Kelly et al. (2021) and forward citations. The former measures the degree of patent novelty by comparing textual similarity with prior and future patents. For the latter, I construct forward citations using the USPTO's citations files as citations received within the first 5 years since the patent grant date and normalize them by their average value within an application year to ensure that they are comparable across years.

I aggregate patent-based measures to the firm-year level using the mapping in Kogan et al. (2017). Patents are recorded in their application year to reflect the timing of innovation.

I restrict the sample to 1975–2018 and drop firms with consistently low R&D expenditure (less than 2.5m 2012 USD per year), low patenting (less than 2.5 patents per year) or less than 5 years in sample. The start year reflects that USPTO data is available for patents granted after 1976, while the end year is chosen such that grant decisions are likely final for relevant patent applications. The final sample covers more than 80% of R&D expenditure in Compustat and patent valuations in Kogan et al. (2017), and 40% of the R&D recorded in BEA accounts. See Appendix B for details.

2.2 Measurement

R&D Allocative Efficiency depends on four parameters: $\{\gamma, \{\omega_{it}, \Delta_{it}, \zeta_{it}\}\}$. As is common in the literature, I set $\gamma = 1/2$ (Acemoglu et al., 2018; Akcigit and Kerr, 2018).[7]

R&D wedges can be measured up to a factor from the average R&D return, i.e., the ratio of value created from R&D divided by its cost:

$$\frac{z_{it}\, V_{it}}{W_t\, \ell_{it}} = \frac{1}{\gamma}\,(1 + \Delta_{it}).$$

In the model, firms with high R&D returns are more constrained, implying larger wedges. Key for this interpretation are common, log-linear production and cost functions, which yield proportional marginal and average returns, and are standard in the literature (Gancia and

[7] Terry (2023) estimates a γ of 0.4, while the estimates in Dechezleprêtre et al. (2023) and Lichter et al. (2025) imply a value of 2/3 or larger; however, the authors attribute the large estimate to financial frictions. Online Appendix E.4 discusses the standard estimation approach and reports estimates from my sample in line with the chosen value of 1/2. A larger γ increases the cost of misallocation because firms' demand for researchers becomes more sensitive to R&D wedges: the elasticity of R&D employment with respect to the wedge is $-\frac{1}{1-\gamma}$. Online Appendix Table C.3 reports robustness for alternative values of γ.

Zilibotti, 2005; Aghion et al., 2014).

I measure R&D wedges over 5-year windows with a 1-year lag between R&D expenditure and patent valuations:

$$\widehat{1+\Delta_{it}} = \gamma\, \frac{\sum_{s=0}^{4} \texttt{Patent Valuations}_{it+s}}{\sum_{s=0}^{4} \texttt{R\&D Expenditure}_{it-1+s}}.$$

I further refine this measure along several dimensions to address potential measurement concerns. First, under the assumption that technology is similar within industry × year cells, I residualize measured R&D wedges with respect to these cells. This addresses the concern that differences in the scale elasticity γ across firms could be confounded with differences in R&D wedges.[8] Second, I restrict the sample to observations with at least 50 patents over the 5-year window. This leverages the law of large numbers to close the gap between realized and expected R&D returns, which is the relevant object as firms equalize expected marginal benefits to marginal costs.[9] This approach does not account for firm-level shocks that yield common variation in realized patent valuations, which I investigate separately. Third, following Bloom et al. (2020), I use non-negative changes in profits or sales as alternative measures of innovation output and measure R&D wedges as

$$\widehat{1+\Delta_{it}} = \gamma\, \frac{\sum_{s=0}^{4} \max\{\, X_{it+s} - X_{it+s-1}, 0\}}{\sum_{s=0}^{4} \texttt{R\&D Expenditure}_{it-1+s}} \text{ with } X \in \{\texttt{Profits, Sales}\}.$$

These measures bypass the patent system entirely and thus address the concern that differential patenting rates across firms could generate variation in measured R&D returns unrelated to frictions (Cohen et al., 2000).

R&D productivity can be measured from firms' first-order conditions as

$$\theta_{it} = (1+\Delta_{it}) \times (W_t\, \ell_{it})^{1-\gamma}\; W_t^{\gamma}.$$

Note that the formula for R&D Allocative Efficiency is scale independent in θ_{it} (and Δ_{it}) such that we can drop the common wage intercept. I thus measure R&D productivity as

[8] In Online Appendix E.4 I further confirm that estimated elasticities are constant across the R&D return distribution, suggesting that differences therein are not a main driver of R&D returns.

[9] The model in Section 1 assumes no uncertainty around the value or quantity of inventions. In general, the appropriate value is the expected discounted value created from innovation, which is proportional to the expected value under homogeneous discount rates and a common gap between investment and realization.

$$\hat{\theta}_{it} = \widehat{(1 + \Delta_{it})} \left(\sum_{s=0}^{4} \texttt{R\&D Expenditure}_{it-1+s} \right)^{1-\gamma}.$$

I consider three approaches to measuring the **impact-value factor**. In the first approach, I follow the workhorse growth models and assume a constant factor across firms. In the second approach, I measure them from markups, as suggested by a limit-pricing setup and in line with the discussion above and in Online Appendix G. I obtain markups either directly from Loecker et al. (2020) ($\hat{\mu}_{it}$) or, alternatively, via the value implied by firms' profit rates. The impact-value factor is then given by

$$\hat{\zeta}_{it} = \hat{\mu}_{it} \quad \text{or} \quad \hat{\zeta}_{it} = \frac{\sum_{s=0}^{4} \texttt{Sales}_{it}}{\sum_{s=0}^{4} \texttt{Sales}_{it} - \sum_{s=0}^{4} \texttt{Profit}_{it}}. \tag{6}$$

In the third approach, I measure the impact-value factor as the ratio of combined patent quality to valuations, thereby implicitly equalizing patent-quality with the growth impact:

$$\hat{\zeta}_{it} = \frac{\sum_{s=0}^{4} \texttt{Patent Quality}_{it+s}}{\sum_{s=0}^{4} \texttt{Patent Valuations}_{it+s}}. \tag{7}$$

I residualize impact-value factors with respect to industry-year fixed effects to account for differences across either dimension that are unrelated to true quality. I confirm n Online Appendix C.2 that this measure of the impact-value factor indeed mitigates the relationship between R&D expenditure and profit rates as required by the theory.

3 R&D Returns, Growth, and Welfare

In this section, I report descriptive facts around measured R&D wedges, including their correlation with measures of frictions, before presenting my baseline estimates for aggregate R&D Allocative Efficiency. I use the terms R&D wedges and R&D returns interchangeably here owing to their measurement equivalence in my context. I discuss robustness in Section 4.

3.1 Facts on R&D Returns

In the frictionless benchmark economy, R&D wedges are equalized across firms and, thus, measured R&D returns should be as well. Instead, I find large dispersion in measured R&D returns as highlighted by their histogram in Figure 2 (a). A firm at the 75th percentile of the distribution has close to twice the median return in levels with a similar gap between the median and 25th percentile. The standard deviation of log R&D returns is 0.9 and has been rising over time, as documented in Table 1.

Large dispersion in investment returns echoes the literature on capital investment. Hsieh and Klenow (2009) document large differences in the return on capital across U.S. firms, while the frictionless investment model predicts none. Importantly, this literature argues that the empirical return on capital dispersion implies large losses in aggregate production relative to a return-equalizing allocation. Comparing dispersion in both returns in Table 1, I find that R&D return dispersion is $0.93/0.64 - 1 \approx 45\%$ larger in my sample.

Figure 2: R&D Returns are Highly Dispersed and Persistent

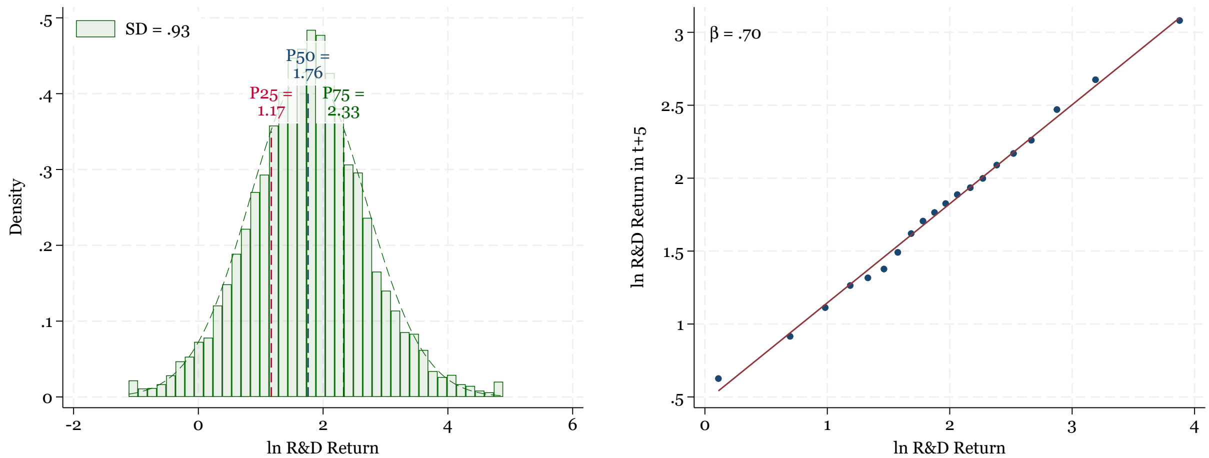

(a) Histogram of R&D Returns

(b) Autocorrelation of R&D Returns

Notes: Left panel reports histogram of R&D returns and density function of a normal distribution with the same mean and variance. Right panel reports a binscatter plot of R&D returns against their 5-year lag. See text for details.

Figure 2 (b) further confirms that R&D returns are highly autocorrelated over time and, thus, appear to capture persistent forces rather than idiosyncratic shocks as implied, e.g., by expectation-realization gaps or classical measurement error. The autocorrelation coefficient of 0.7 at the 5-year level suggests an annual autocorrelation above 0.9. I confirm that this implied persistence level remains at longer time horizons in Online Appendix E.1.

I next investigate economic drivers of R&D return dispersion in Table 2.[10] First, Panel A examines general investment frictions. If investment is distorted across multiple margins, measures of capital investment frictions should correlate with R&D returns; however, results are mixed: R&D returns are uncorrelated with the return on capital, a summary measure for investment frictions (David et al., 2016), but highly correlated with Tobin's Q, another established measure (Peters and Taylor, 2017). Financial frictions also yield mixed results. High-liquidity firms have lower returns, consistent with being less constrained, yet firms

[10]See Online Appendix G for the associated theoretical foundations.

Table 1: Return Dispersion Across Time

Period	R&D Return	Return on Capital	
	SD	SD	Δ%
1975 – 2014	0.93	0.64	45%
1975 – 1990	0.74	0.46	62%
2000 – 2014	0.98	0.73	33%

Note: Column headers SD report standard deviations of return measure. Columns headers Δ% indicate percent difference of Return on R&D dispersion with respect to return in consideration. Returns are measured in logs.

with large dividend payments—presumably unconstrained—have high R&D returns. These inconclusive findings are surprising given evidence that intangible investments such as R&D are particularly susceptible to financial frictions (Brown et al., 2009).

Second, R&D return dispersion could reflect firm-specific risk-premia driven by heterogeneous exposure to aggregate risk (David et al., 2022). As reported in Panel B, I find no evidence that stock market βs, a measure of systematic risk, explain R&D returns. However, firms with volatile patent valuations tend to have higher R&D returns. Such "risk-premia" could arise if firms' decision makers cannot fully diversify innovation risk.

Third, I find that R&D subsidies do not explain R&D return variation in Panel C. Investment subsidies distort returns by reducing the true investment costs relative to reported costs such that firms with large subsidies earn low reported returns (Hsieh and Klenow, 2009). Using data on state-level R&D tax credits from Lucking et al. (2019), I find that the induced variation in R&D user costs only weakly correlates with R&D returns. Similarly, public patent co-ownership does not explain a significant share of R&D return dispersion.

Fourth, Panel D investigates the link between R&D returns and firm dynamics. Factor return dispersion arises naturally in models with adjustment costs (Asker et al., 2014). With such frictions, a positive shock to R&D productivity leads to a temporary rise in R&D returns as R&D outputs, which reflect R&D inputs and productivity, adjust faster than inputs, leading to a positive correlation between R&D growth and returns. In line with this prediction, long-term growth in R&D—from $t-1$ to $t+6$—accounts for approximately 10% of R&D return dispersion. Long-term TFP growth shows a similar correlation, rationalized by the same mechanism. Finally, I find that prior TFP growth and stock returns predict R&D returns, suggesting that high returns are associated with firm expansion.

Lastly, a growing literature finds that monopsony power plays an important role in shaping the allocation of workers (Card et al., 2018). This literature also finds that high-

Table 2: Correlations with R&D Returns

Variable	Estimate	Std. err.	R^2	Observations
A. Frictions				
Return on Capital	0.038	(0.066)	0.1%	11,932
Tobin's Q	0.200***	(0.029)	6.2%	10,538
Liquidity	-0.041*	(0.021)	0.2%	10,654
Dividend rate	35.609***	(7.196)	1.5%	11,572
B. Risk				
CAPM β	-0.009	(0.062)	0.0%	6,799
Valuation risk	0.616***	(0.098)	3.6%	10,993
C. Taxation				
R&D user cost $1-\tau$	-0.644	(0.619)	0.1%	11,241
Public patent involvement	1.499	(1.067)	0.2%	11,933
D. Dynamics				
Long-term R&D growth	0.422***	(0.054)	10.5%	6,535
Long-term TFP growth	0.463***	(0.050)	4.1%	5,409
Prior excess stock return	0.263***	(0.031)	1.1%	10,127
Prior TFP growth	0.338***	(0.043)	1.3%	7,285
E. Inventors				
Inventors	0.194***	(0.027)	5.5%	11,880
Firm dominance	0.314***	(0.033)	12.3%	11,877
Inventor specialization	0.241**	(0.099)	0.4%	11,876
F. Impact-Value Factor				
Scientific impact-based	0.054	(0.040)	0.2%	11,838
5-year citations-based	0.183***	(0.049)	1.7%	11,859
Lifetime citations-based	0.201***	(0.049)	2.0%	11,859
Profit-rate	0.021***	(0.007)	0.7%	11,928

Note: Each row in Panel A–E reports the regression coefficient of a separate regression with dependent variable log R&D returns. In panel F, log R&D returns are the independent variable. All regressions control for NAICS3× Year fixed effects and standard errors are clustered at the NAICS6 level. See text and Appendix for details.

skilled workers, a group likely including inventors, are more affected by monopsony power and that larger firms tend to have more of it (Seegmiller, 2025; Berger et al., 2022; Yeh et al., 2022). Monopsony power over inventors could drive R&D return dispersion as firms with greater market power restrict hiring more aggressively to keep wages low and, as a result, create more value per unit of cost (Lehr, 2025). In line with this idea, Panel E shows that firms hiring more inventors have larger R&D returns.[11] In addition, firms dominating their

[11]This finding would also be in line with a size-dependent R&D scale elasticity s.t. scaling up is costlier

inventor labor market and hiring more specialized inventors tend to have larger R&D returns, in line with firm-specific human capital or limited outside options as sources of monopsony power (Acemoglu, 1997; Schubert et al., 2025).

Overall, the explanatory power of the mechanisms considered here is low, echoing similar findings for the return on capital (David et al., 2016). This finding makes the interpretation of measured R&D wedges difficult, since we do not fully know their source. One possibility difficult to rule out is measurement error. Consequently, I discuss extensive robustness checks and direct adjustment methods to mitigate the impact of such variation in Section 4.

Finally, I find inconclusive evidence on the link between R&D wedges and the impact-value factor in Panel F of Table 2. Across all measures of the impact-value factor I find positive correlations, however, not always significant. Magnitudes are small for the profit-rate- and patent-quality-based measures and moderately larger for the citation-based measures.[12] However, these correlations change over time and can be negative at times.

3.2 Combining Data and Model

I next turn to estimating the macroeconomic impact of R&D wedges. Per Proposition 2, I **measure** R&D Allocative Efficiency as

$$\hat{\Xi}_t = \frac{\sum_{i=1}^{N_t} \hat{\omega}_{it} \left(\widehat{1+\Delta_{it}}\right)^{-\frac{\gamma}{1-\gamma}\hat{\eta}_t}}{\left(\sum_{i=1}^{N_t} \hat{\omega}_{it} \left(\widehat{1+\Delta_{it}}\right)^{-\frac{\gamma}{1-\gamma}\hat{\eta}_t}\right)^{\gamma}} \quad \text{with} \quad \hat{\omega}_{it} = \frac{\hat{\theta}_{it}^{\frac{1}{1-\gamma}}}{\sum_{i=1}^{N_t} \hat{\theta}_{it}^{\frac{1}{1-\gamma}}}.$$

This approach assumes a representative sample for the U.S. R&D sector. Thus, the estimates are biased towards a milder impact of R&D wedges if large, established firms tend to be less impacted by frictions. I confirm in Section 4.2 that R&D Allocative Efficiency is indeed lower among smaller firms in my sample.

For my baseline estimates I assume that R&D wedges and impact-value factors are independent and set $\hat{\eta}_t = 1$. Alternatively, I follow Proposition 2 and estimate its value as $\hat{\eta}_t = \sqrt{1 + 2\,\hat{\beta}_t}$, where $\hat{\beta}_t$ is estimated over a centered rolling 10-year window by regressing profit-based R&D returns on the patent-quality-based impact-value factor as in Table 3.

I collapse annual estimates using geometric averages to analyze longer-run developments. I consider the full sample from 1975 to 2014 as well as the early and late periods, 1975–90 and 2000–14.

for large companies. However, I do not find a systematic link between R&D returns and the R&D scale elasticity, as discussed in Online Appendix E.4.

[12] I replace patent valuations with changes in profits in regressions with patent-quality-based measures to avoid mechanical correlation arising from using patent valuations in R&D returns and impact-value factors.

Finally, I calculate bootstrapped standard errors for the estimates. For each year, I sample observations with replacement until I reach the true sample size and calculate the annual aggregates. I repeat this exercise for 1000 bootstrap samples and report the standard deviation of the resulting estimates together with non-parametric 95% confidence intervals.

Proposition 2 allows us to estimate the short-run impact of R&D wedges by comparing the growth rate under the measured impact $\hat{\Xi}_t$ to its hypothetical value under $\Xi_t = 1$. This **counterfactual** assumes that offsetting R&D wedges is feasible in principle. Even if not, the estimates are still informative about whether changes in the economic growth rate arise from R&D wedges or the frictionless growth rate.

I estimate the long-run impact on economic growth and welfare for two scenarios. In the endogenous growth scenario, I set $L_t = L$ and $\phi = 0$, such that setting $\Xi_t = 1$ achieves the frictionless growth rate g^C, which I calibrate as $g^C = 1.5\% \; \Xi^{-1}$ to match the long-run U.S. growth rate. Misallocation reduces the long-run growth rate in this scenario.

In the second scenario—the semi-endogenous growth case—I assume that the frictionless growth rate and population dynamics take the form

$$g_t^C = A_t^{-\phi} \, L_t^{\gamma} \, g^C \qquad \text{with} \quad \phi > 0 \quad \text{and} \quad L_{t+1} = (1+n) \, L_t.$$

Parameter $\phi > 0$ determines the degree to which "ideas are getting harder to find" over time, which is key to achieving constant long-run productivity growth with a growing population (Jones, 1995). The long-run growth rate in this economy is pinned down by $g = (1+n)^{\gamma/\phi} - 1$; however, the short-run growth rate and the long-run productivity level respond to changes in the environment.[13] I assume that the economy is on its balanced growth path prior to the policy change and trace subsequent changes in productivity and consumption. I set population growth to $n = 1\%$ and calibrate ϕ to achieve a long-run growth rate of 1.5%.

3.3 R&D Misallocation and Economic Growth

Consider the case of unrelated R&D wedges and impact-value factors first. The black line in Figure 3 plots the annual estimates of R&D Allocative Efficiency Ξ_t, while Panel A of Table 3 reports long-run values and their associated consumption-equivalent welfare cost.

The estimated R&D Allocative Efficiency suggests that frictions significantly reduce economic growth. I estimate an average growth impact of -21% for the full sample, which suggests a growth rate of 1.9% in the absence of R&D wedges given realized annual productivity growth of 1.5%. Unsurprisingly, such a stark slowdown has large welfare consequences.

[13] Stable growth requires constant $A_t^{\phi} \, L_t^{\gamma}$, such that we can solve for A_t conditional on L_t.

Table 3: The Impact of R&D Wedges on Economic Growth and Welfare

Time Horizon	Growth Impact $\Xi - 1$		Welfare Cost	
	Estimate	Std. Err.	Endogenous	Semi-End.
A. Baseline				
1975–2014	-20.9%	(0.37%)	12.0%	11.4%
1975–1990	-14.3%	(0.45%)	7.4%	7.2%
2000–2014	-23.8%	(0.64%)	14.4%	13.6%
Δ Change	-11.2%		5.5%	5.4%
B. Adjusted				
1975–2014	-21.0%	(0.38%)	12.1%	11.5%
1975–1990	-13.4%	(0.42%)	6.9%	6.7%
2000–2014	-26.1%	(0.70%)	16.4%	15.4%
Δ Change	-14.7%		7.7%	7.4%
C. Preferred				
1975–2014	-10.8%	(0.41%)	5.4%	5.2%
1975–1990	-2.5%	(0.47%)	1.1%	1.1%
2000–2014	-15.0%	(0.76%)	7.8%	7.6%
Δ Change	-12.7%		6.5%	6.3%

Notes: Table reports estimates for impact of R&D wedges across samples and their welfare impact in consumption equivalent terms. Standard errors and confidence intervals are calculated using a bootstrapping procedure. Adjusted values account for the covariance of R&D wedges and impact-value factors, using the patent-quality measure for the latter. Preferred values further adjust for measurement error and are benchmarked against estimates in the life sciences industry. See text and Appendix for details.

The model suggests that welfare would be 11–12% higher in the absence of R&D wedges. For comparison, Berger et al. (2022) estimate that monopsony in the production sector reduces U.S. output by 21% and welfare by 8%, while Hsieh and Klenow (2009) estimate 30–40% larger U.S. output in the absence of production factor misallocation.

The annual estimates further suggest that declining R&D Allocative Efficiency contributed to the productivity growth slowdown. The estimates declined throughout the sample period with a sharp downturn and partial recovery during the 1990s, potentially reflecting the Dot-Com boom. Comparing the estimates for 1975–90 and 2000–14, I find that R&D Allocative Efficiency declined from -14% to -24%. This decline implies an $\frac{24\%-14\%}{1-14\%} \approx 11\%$ slower short-run growth rate, which accounts for $\frac{11\%}{40\%} \approx 25\%$ of the growth slowdown.[14]

[14] Annual TFP growth declined from 0.5% in 1976–1995 to 0.3% in 2005–2018, a $\frac{0.5\%-0.3\%}{0.5\%} \approx 40\%$ reduction.

Figure 3: The Impact of R&D Wedges $\Xi_t - 1$ for 1975–2014

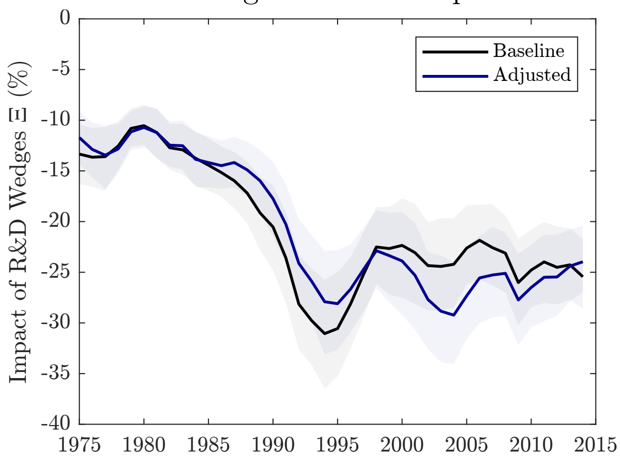

(a) Estimates for Economy

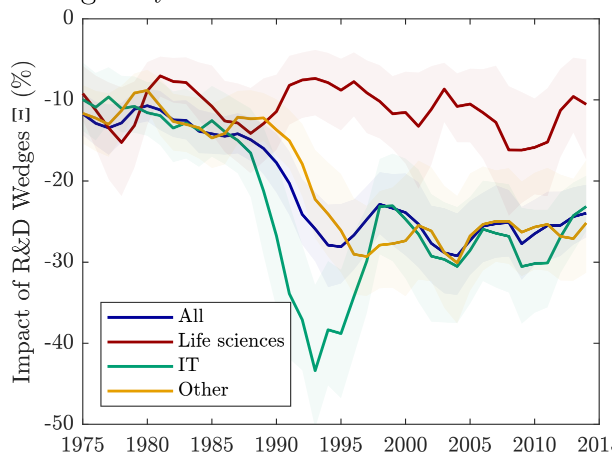

(b) Adjusted Estimates by Sector

Notes: Panel (a) reports estimates for $\Xi - 1$ without and with adjustment for the impact-value factor. The adjustment factor is estimated using 5-year citations over sales growth to measure the impact-value factor. Panel (b) reports estimates by sector. Sectors are defined following Mezzanotti and Simcoe (2025). 95% confidence intervals shaded and estimated via bootstrapping. See text and Appendix for details.

Adjusting for the impact-value factor has a negligible impact on the average estimated growth impact, but implies a marginally stronger decline in R&D Allocative Efficiency. As reported in Panel B, the long-run estimated impact is -21% and the evolution of R&D Allocative Efficiency implies 15% deceleration in the productivity growth rate.

Panel (b) of Figure 3 highlights that the decline was not universal across industries.[15] In particular, average R&D Allocative Efficiency was approximately stable in Life Sciences, while it declined in the IT sector and other manufacturing industries. In line with a general rise in misallocation of R&D during the dot-com boom, the decline in R&D Allocative Efficiency was particularly pronounced for the IT sector during this period; however, the sector also experienced a strong subsequent recovery. I exploit the consistency of the estimates for the Life Sciences sector in Panel C of Table 3, where I benchmark the estimates for the full economy against this sector, in addition to a minor adjustment for measurement error discussed in the next section. The resulting estimates measure R&D misallocation relative to a potentially feasible benchmark: the R&D allocation in Life Sciences. The estimates suggest an overall impact of -11% and an implied decline in the growth rate of -12.7% as misallocation levels were initially close to the benchmark, but subsequently diverged. This preferred estimate suggests an average welfare loss of 5% from R&D misallocation and at-

[15] The industry classification follows Mezzanotti and Simcoe (2025), but combines "other" and "other manufacturing" to achieve a sufficient sample size.

tributes $\frac{12.7\%}{40\%} \approx 32\%$ of the growth slowdown to declining R&D Allocative Efficiency.

Declining R&D Allocative Efficiency is also potentially important for estimating or calibrating the degree to which "ideas are getting harder to find." Bloom et al. (2020) use data for 1930–2010 to estimate an average decline of research productivity, defined as TFP growth divided by the effective number of researchers, of 5.1% per year. Together with an average TFP growth rate of around 1.65%, they conclude that the effective degree of fishing-out, $\frac{\phi}{\gamma}$, is given by $5.1\%/1.65\% \approx 3.1$.[16] Using their data and focusing on the decades starting in 1970 yields a similar value of 3.3; however, this estimate does not take into account that declining R&D productivity is partly driven by rising misallocation rather than fishing out alone. Adjusting their estimates yields alternative values of 2.4 or 2.7 depending on whether we also adjust the average TFP growth rate. Thus, taking into account rising R&D misallocation is not only important for our understanding of the recent decline in productivity growth, but also for estimating the degree to which "ideas are getting harder to find."[17]

4 Discussion

The baseline estimates for R&D Allocative Efficiency suggest that it contributed significantly to the U.S. productivity slowdown. In this section, I confirm that this conclusion is robust to a large number of robustness checks, investigate the role of firm heterogeneity, and discuss the implications of also taking into account the evolution of the Incentive Alignment Index.

4.1 Measurement concerns

As discussed in Proposition 2, estimated R&D Allocative Efficiency is inversely related to variation in R&D returns. The latter may be subject to measurement error, which could overstate the degree of R&D misallocation. This concern is particularly pertinent as R&D returns are not easily explained by direct measures of frictions. While my main focus is changes over time, rather than levels, this concern remains valid to the extent that measurement error also changed over time. I thus consider four types of adjustments in Table 4—to R&D inputs, to R&D outputs, econometric approaches to measurement error, and alternative impact-value factors. The main estimates are robust, with estimates for the change in allocative efficiency falling in $[-21.8\%, -8.2\%]$, compared to the preferred estimate of -12.7%. For comparison, the first row reports the estimates from Panel B in Table 3. Technical details on the adjustments are provided in Appendix C.1.

[16] See Table 7 in their paper and the accompanying text. Their model assumes $\gamma = 1$; however, it is straightforward to show that the equivalent measure in my model is the ratio of ϕ and γ.

[17] Similarly, Ekerdt and Wu (2025) find that accounting for the decline in researcher quality significantly affects the implied degree of fishing out.

Panel A reports adjustments to **R&D inputs**, which I measure using R&D expenditure at baseline. First, I confirm that results are robust to alternative input measures such as **SG&A expenditures**, which include complementary investments such as marketing, or capitalized innovation expenditure, as represented by the **knowledge capital** measure from Ewens et al. (2024). The former is appropriate if the value of R&D outputs also reflects complementary investments, while the latter is in line with models where R&D outputs depend on past and current R&D expenditure (Peters and Taylor, 2017). Second, the first-order condition for R&D investment (2.2) implies that the ratio of R&D output to variable R&D costs should be equalized in the absence of frictions; however, in practice, reported R&D expenditure may also entail **fixed costs**. To address this potential measurement concern, I propose to estimate R&D fixed costs by industry×year by comparing the measured R&D returns for very large firms, whose returns should be unaffected by comparatively small fixed costs, with those of very small firms, whose measured returns would be artificially depressed by the presence of fixed costs. I then adjust total R&D costs for the fixed cost component to estimate variable costs and re-estimate the aggregate measures. Finally, it is well known that R&D-intensive firms frequently acquire smaller innovative firms (Phillips and Zhdanov, 2013; Fons-Rosen et al., 2024). Such **acquisitions** could distort measured returns if firms patented acquired inventions, such that they would be reflected in the value of innovation, but not the R&D expenditure. I show that the degree of induced measurement error is a function of the acquisition intensity, i.e., the ratio of acquisitions to R&D expenditure, and propose a model-consistent methodology to estimate the average share of acquisition costs linked to R&D returns. In the last step, I adjust R&D expenditure for this share and recalculate the aggregate measures. The results in Panel A suggest that these adjustments to measured R&D inputs are immaterial with adjusted average estimates ranging from -23% to -21% and changes ranging from -17% to -11%, compared to -21% and -15% in the baseline.

Panel B investigates robustness around the measurement of **R&D outputs**. I split this investigation into adjustments to patent valuations and alternative measures of R&D output. Kogan et al. (2017) measure **patent valuations** by filtering stock market returns around the patent grant announcement. Their estimates incorporate several assumptions whose violation may lead to measurement error.[18] While uncorrelated measurement error may

[18] In their framework, stock market returns follow a normal distribution, while patent valuations are drawn from a conditional normal with support over non-negative values. Investors know the value of the patent, but it is only granted with a certain probability, which is assumed constant across patents. It then follows that the grant announcement increases the value of the firm by the fraction of the patent value that exceeds its expected value, i.e., in excess of the product of the grant probability times the patent value.

wash out when averaged across patents—consistent with my findings in Online Appendix F—systematic error may not. I thus investigate a range of potential adjustments.

First, the literature has long recognized that **outlier ideas** significantly contribute to the return on innovation (Akcigit and Kerr, 2018; Berlingieri et al., 2026). Due to their infrequent nature, these may contribute to return dispersion as firms equalize expected values, but the econometrician only observes realizations. I confirm that outliers do not materially impact my estimates by winsorizing patent valuations at the 95%-level on a year-by-year basis. Next, the estimation procedure applies strictly positive values to patents by assumption, even if they may not have contributed to the firm value (Jaffe and Lerner, 2007). I confirm that patents with **low valuations**, which could be disproportionately due to measurement error and functional form assumptions, do not drive my results by excluding patents valued at less than \$500k. Another simplifying assumption behind the valuation estimates is that all patents have the same ex-ante **grant probability**. I confirm that my results are robust to either allowing the grant rate to vary by patent technology class or by allowing for a relationship between the patents' true value and their grant probability, e.g., because higher quality patents are more likely to be granted. For the latter, I assume an elasticity of grant probability to value of 0.05 here and report robustness for alternative values in Appendix Table C.1. As Panel B.1 confirms, these adjustments do not materially affect the estimates, which range from -22 to -20% for the average and -16 to -15% for the change over time.

While the results appear robust to adjustments to the patent valuation measure, it is possible that the measure is conceptually incomplete as it has long been recognized that not all inventions are patented (Cohen et al., 2000). R&D returns might then reflect whether firms patent their inventions rather than their quantity or value. While any cross-industry or time variation is accounted for by industry-year fixed effects, there may be intra-industry differences. Following Bloom et al. (2020), I thus propose to measure **R&D outputs alternatively** as non-negative changes in profits or sales. Both measures tend to increase the measured impact of R&D misallocation, while confirming a significant decline in R&D Allocative Efficiency over time as reported in Panel B.2.

In Panel C, I explore three methods to directly estimate the degree of **measurement error** and adjust the estimates accordingly. Such methods are particularly useful when measurement error is ubiquitous across measures, e.g., due to the expectation-realization gap arising from measuring realized returns, while firms equalize expected returns. First, I estimate potential measurement error arising from the expectation-realization gap using bootstrapping in Online Appendix F.1. The results suggest minimal measurement error:

because firms hold many patents, the law of large numbers implies that realized portfolio values closely approximate expected values, limiting the scope for expectation-realization mismatch. Second, I model R&D returns as an AR(1) process with i.i.d. measurement error, e.g., due to the expectation-realization gap, and derive a GMM estimator.[19] The appropriate adjustment is then to dampen R&D return dispersion according to the degree of estimated measurement error. As detailed in Appendix F.2, the estimator suggests very little measurement error when applied to the full sample, but period-specific estimates suggest some measurement error. The adjusted estimates suggest slightly lower R&D misallocation and dampen the fall from -14.7% to -13.3%. This adjustment is implemented in Panel C of Table 3. Third, I implement the adjustment proposed in Bils et al. (2021), who assume additive measurement error. The adjustment reduces the estimated average misallocation and its increase over time, from -21 to -18% and from -15 to -8%, but continues to suggest a quantitatively important role.[20]

Finally, I investigate alternative measures of the **impact-value factor** in Panel D. Using citations instead of Kelly et al. (2021)'s impact measure significantly increases the estimates, as does using the ratio of TFP changes to profit changes. On the other hand, using profit rates or markups increases the average level of misallocation, but decreases the measured change. Using the markup measure from Loecker et al. (2020) to proxy for the impact-value factor implies a decline in R&D Allocative Efficiency by 8.2%, which would attribute $8.2/40 \approx 21\%$ of the decline in productivity growth to rising R&D misallocation.

In summary, the adjustments considered do not materially alter the conclusion that R&D misallocation contributed significantly to declining TFP growth. The most conservative estimates suggest that misallocation reduced growth by around 18% relative to a frictionless, and likely unachievable, benchmark and that rising R&D misallocation accounted for at least 20% of declining U.S. TFP growth. Appendix Table C.3 confirms that these conclusions hold under alternative fixed effect specifications for returns, when requiring a larger number of patents, and when extending the time window over which R&D expenditure and valuations are accumulated. The table also provides estimates for alternative values of the scale elasticity γ and confirms that larger values imply more misallocation and vice versa.

[19] A telltale sign of low measurement error for the AR(1) model is the high persistence of R&D returns over time. I confirm in Appendix E.1 that persistence is stable across time and that the implied AR(1) coefficient remains stable even at long time horizons. The Appendix also discusses the expectation-realization gap.

[20] An important caveat to the methodology is that it can confound size-dependent frictions with additive measurement error and, thus, may be excessively aggressive in its correction of R&D returns. See Online Appendix F.3 for additional details.

Table 4: R&D Wedges, Economic Growth and Welfare — Robustness

Specification	Growth Impact $\Xi - 1$				Welfare Cost of Δ	
	1975–2014	1975–90	2000–14	Δ	End.	Semi-End.
Baseline	-21.0%	-13.4%	-26.1%	-14.7%	7.7%	7.4%
A. R&D Input						
S,G& A	-22.7%	-16.3%	-25.5%	-11.0%	5.4%	5.3%
Knowledge capital	-23.1%	-14.5%	-28.9%	-16.9%	9.1%	8.7%
Fixed costs	-22.0%	-14.1%	-27.1%	-15.1%	7.9%	7.7%
Acquisitions	-20.7%	-13.0%	-25.7%	-14.5%	7.6%	7.3%
B. R&D Output						
B.1. Patent Valuations						
Winsoring at 5%	-20.2%	-12.9%	-26.4%	-15.5%	8.2%	7.9%
Exclude if $< 500k$	-21.3%	-13.1%	-27.3%	-16.3%	8.7%	8.4%
CPC-Class grant rate	-21.1%	-12.8%	-27.0%	-16.3%	8.7%	8.4%
Value-Grant Elasticity $\eta = 0.05$	-21.8%	-14.0%	-27.0%	-15.1%	7.9%	7.7%
B.2. Alternative Measures						
Δ Profits	-23.4%	-19.9%	-29.5%	-12.0%	6.0%	5.9%
Δ Sales	-29.5%	-22.1%	-37.4%	-19.6%	11.0%	10.5%
C. Econometric Approaches						
GMM	-20.4%	-13.3%	-24.8%	-13.3%	6.8%	6.6%
Additive error	-17.9%	-13.7%	-20.8%	-8.3%	3.9%	3.9%
D. Private vs Social Return						
Short-run citations	-25.6%	-14.8%	-32.9%	-21.2%	12.2%	11.6%
Long-run citations	-25.9%	-14.8%	-33.3%	-21.8%	12.7%	12.0%
Profit-rates	-21.3%	-14.8%	-23.5%	-10.3%	5.0%	4.9%
Markups	-21.3%	-15.4%	-22.3%	-8.2%	3.9%	3.8%
ΔTFP/ΔProfit	-21.5%	-15.2%	-28.3%	-15.5%	8.2%	7.9%

Notes: Table reports estimates for impact of R&D wedges across samples together with their implications for welfare. Changes in welfare are in consumption equivalent terms. See text and Appendix for details.

4.2 Firm Heterogeneity

I next investigate the role of firm heterogeneity in R&D misallocation in Table 5. This also serves as a robustness check, since my sample is skewed towards larger firms. First, I confirm that the entry and exit of firms does not materially impact the main estimates. The implied level and trajectory of R&D misallocation are not materially different when focusing on continuing firms only. See Appendix C.3 for implementation details. Second, Panel B confirms that misallocation appears to be significantly worse in level and trajectory among firms with below-median R&D expenditure at a given point in time. This finding is in line with a world in which smaller firms face more frictions, which is commonly assumed (Brown et al., 2009), and suggests that a representative sample of firms, covering all firms in the U.S. economy, might feature lower estimated R&D Allocative Efficiency as well as a faster decline therein. Thus, the main estimates reported may be conservative regarding the impact of frictions in the R&D sector. See Online Appendix C.5 for a discussion of this insight. Finally, Panel C investigates methodologies to weight observations. Compared to a baseline of using implied R&D productivity, applying no weight worsens estimates in level and trajectory as does using changes in sales to calculate R&D productivity.

Table 5: R&D Wedges, Economic Growth and Welfare — Firm-Heterogeneity

Specification	Growth Impact $\Xi - 1$				Welfare Cost of Δ	
	1975–2014	1975–90	2000–14	Δ	End.	Semi-End.
Baseline	-21.0%	-13.4%	-26.1%	-14.7%	7.7%	7.4%
A. Entry & Exit						
Continuing firms	-20.4%	-13.0%	-26.1%	-15.1%	7.9%	7.6%
B. Firm Size in R&D						
High R&D	-19.3%	-11.7%	-24.1%	-14.1%	7.3%	7.1%
Low R&D	-30.3%	-18.9%	-38.1%	-23.7%	14.2%	13.4%
C. Weighting						
Unweighted	-24.7%	-17.3%	-30.0%	-15.4%	8.1%	7.9%
Δ Sales	-27.1%	-18.9%	-31.4%	-15.3%	8.1%	7.8%

Notes: Table reports estimates for impact of R&D wedges across samples together with their implications for welfare. Changes in welfare are in consumption equivalent terms. See text and Appendix for details.

Online Appendix E.3 further investigates whether changes in R&D Allocative Efficiency have predictive power at the industry level. I first show that the model predicts that R&D Allocative Efficiency is negatively correlated with industry-level R&D expenditure and R&D returns, which captures the intuition that high allocative efficiency implies less waste. I then test this prediction focusing on 10-year changes within industries and confirm a strong negative correlation of both aggregates with R&D Allocative Efficiency. Such a relationship does not arise mechanically as R&D Allocative Efficiency is scale independent, i.e., homogeneous of degree 0, in R&D expenditure and patent valuations. The evidence thus suggests that the model has predictive power at the aggregate level, at least for U.S. industries.

4.3 An Alternative Approach

Lastly, the main formulae developed in Section 3.3 also feature the Incentive Alignment Index, which is tightly linked to the impact-value factor. While my discussion so far has abstracted from this index, I provide an additional set of results including its evolution in Online Appendix C.6. Incorporating the evolution of the Incentive Alignment Index implies significantly larger estimates for R&D misallocation and a stronger decline in combined Allocative Efficiency. However, these conclusions are tightly linked to rising dispersion in the relevant proxies and, thus, crucially depend on whether such a rise is partially accounted for by measurement error. Importantly, the results above can be thought of as the impact of rising R&D return dispersion alone and provide a comprehensive assessment of this channel.

5 Conclusion

This paper presents evidence that frictions, and their impact on the allocation of R&D resources, contributed to the recent decline in U.S. productivity growth. I reach this conclusion based on a growth accounting framework capturing frictions flexibly through a wedge between the private marginal costs and benefits of R&D. In the model, the impact of frictions is captured through a summary statistic, R&D Allocative Efficiency.

I measure the model fundamentals for a sample of U.S.-listed firms over the 1975–2014 period. R&D wedges can be measured from R&D returns, i.e., the ratio of value created from R&D to its costs, which I measure as the ratio of patent valuations to R&D expenditure. I document large and persistent differences in R&D returns contrary to the frictionless prediction of return equalization; the model links such dispersion to frictions. Regression analysis suggests adjustment costs, financial frictions, and monopsony power over inventors as potential drivers of R&D return dispersion; however, most variation remains unexplained.

Next, I estimate the aggregate impact of R&D wedges by combining model and data. My estimates suggest that frictions reduce U.S. economic growth significantly and increasingly so. I estimate for the full sample that economic growth was 11% slower due to frictions, benchmarked against the Life Sciences sector, implying a welfare cost of 5% in consumption-equivalent terms. Furthermore, I find that rising frictions can account for a 13% lower growth rate for 2000–14 compared to 1975–90, which accounts for 32% of the observed productivity slowdown. A large set of robustness checks confirms the main conclusions and suggests a lower bound of -8% for the growth impact of rising R&D misallocation, amounting to 20% of the observed TFP slowdown.

References

Acemoglu, Daron, "Training and Innovation in an Imperfect Labour Market," *Review of Economic Studies*, 1997, *64*, 445–464.

__ , **Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William R. Kerr**, "Innovation, Reallocation, and Growth," *American Economic Review*, 2018, *108*, 3450–3491.

Aghion, Philippe and Peter Howitt, "A Model of Growth Through Creative Destruction," *Econometrica*, 1992, *60*, 323–351.

__ , **Antonin Bergeaud, Timo Boppart, Peter J. Klenow, and Huiyu Li**, "A Theory of Falling Growth and Rising Rents," *Review of Economic Studies*, 10 2023, *90*, 2675–2702.

__ , __ , __ , **Peter J Klenow, and Huiyu Li**, "Good Rents versus Bad Rents: R&D Misallocation and Growth," 2025.

__ , **Ufuk Akcigit, and Peter Howitt**, *What Do We Learn From Schumpeterian Growth Theory?*, Vol. 2, Elsevier B.V.,

Akcigit, Ufuk and Sina T. Ates, "Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory," *American Economic Journal: Macroeconomics*, 2021, *13*, 257–298.

__ **and William R. Kerr**, "Growth Through Heterogeneous Innovations," *Journal of Political Economy*, 2018, *126*, 1374–1443.

__ , **Douglas Hanley, and Stefanie Stantcheva**, "Optimal Taxation and R&D Policies," *Econometrica*, 2022, *90*, 645–684.

Asker, John, Allan Collard-Wexler, and Jan De Loecker, "Dynamic Inputs and Resource (Mis)Allocation," *Journal of Political Economy*, 2014, *122*, 1013–1063.

Autor, David, David Dorn, Christina Patterson, and John Van Reenen, "The Fall of the Labor Share and the Rise of Superstar Firms," *The Quarterly Journal of Economics*, 2020, *135*, 645–709.

Ayerst, Stephen, "Innovator Heterogeneity, R&D Misallocation and the Productivity Growth Slowdown," 2023.

Banerjee, Abhijit V and Benjamin Moll, "Why Does Misallocation Persist?," *American Economic Journal: Macroeconomics*, 2010, *2*, 189–206.

Berger, David W., Kyle Herkenhoff, and Simon Mongey, "Labor Market Power," *American Economic Review*, 2022, *112*, 1147–1193.

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

Bils, Mark, Peter J. Klenow, and Cian Ruane, "Misallocation or mismeasurement?," *Journal of Monetary Economics*, 2021, *124*, 39–56.

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb, "Are Ideas Getting Harder to Find?," *American Economic Review*, 2020, *110*, 1104–1144.

_ , Mark Schankerman, and John Van Reenen, "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 2013, *81*, 1347–1393.

Brown, James R., Steven M. Fazzari, and Bruce C. Petersen, "Financing Innovation and Growth: Cash Flow, External Equity, and the 1990s R&D Boom," *Journal of Finance*, 2 2009, *64*, 151–185.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline, "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 2018, *36*, S13–S70.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber, "Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities," *NBER Macroeconomics Annual*, 2012, *27*, 1–56.

Cohen, Wesley M, Richard R Nelson, and John P Walsh, "Protecting their intellectual assets: Appropriability conditions and why U.S. manufacturing firms patent (or not)," *NBER Working Paper 7552*, 2000.

David, Joel M., Hugo A. Hopenhayn, and Venky Venkateswaran, "Information, Misallocation, and Aggregate Productivity," *The Quarterly Journal of Economics*, 5 2016, *131*, 943–1005.

_ , Lukas Schmid, and David Zeke, "Risk-adjusted capital allocation and misallocation," *Journal of Financial Economics*, 2022, *145*, 684–705.

de Ridder, Maarten, "Market Power and Innovation in the Intangible Economy," *American Economic Review*, 1 2024, *114*, 199–251.

Dechezleprêtre, Antoine, Elias Einiö, Ralf Martin, Kieu Trang Nguyen, and John Van Reenen, "Do Tax Incentives Increase Firm Innovation? An RD Design for R&D, Patents, and Spillovers," *American Economic Journal: Economic Policy*, 2023, *15*, 486–521.

Ekerdt, Lorenz K F and Kai-Jie Wu, "Self-Selection and the Diminishing Returns of Research," 2025.

Ewens, Michael, Ryan Peters, and Sean Wang, "Measuring Intangible Capital with Market Prices," *Management Science*, 4 2024, *71*, 407–427.

Fons-Rosen, Christian, Pau Roldan-Blanco, and Tom Schmitz, "The Effects of Startup Acquisitions on Innovation and Economic Growth," 2024.

Fort, Teresa, Nathan Goldschlag, Jack Liang, Peter Schott, and Nikolas Zolas, "Growth is Getting Harder to Find, Not Ideas," Technical Report, National Bureau of Economic Research 5 2026.

Gancia, Gino and Fabrizio Zilibotti, "Horizontal Innovation in the Theory of Growth and Development," *Handbook of Economic Growth*, 2005, *1*, 111–170.

Goolsbee, Austan, "Investment subsidies and wages in capital goods industries: To the workers go the spoils?," *National Tax Journal*, 2003, *56*, 153–165.

Hsieh, Chang-Tai and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *The Quarterly Journal of Economics*, 2009, *124*, 1403–1448.

Jaffe, Adam B. and Josh Lerner, *Innovation and Its Discontents: How Our Broken Patent System is Endangering Innovation and Progress, and What to Do About It*, Princeton University Press, 2007.

Jones, Charles I., "Time Series Tests of Endogenous Growth Models," *The Quarterly Journal of Economics*, 1995, *110*, 495–525.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy, "Measuring Technological Innovation over the Long Run," *American Economic Review: Insights*, 9 2021, *3*, 303–20.

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

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, "Technological Innovation, Resource Allocation, and Growth," *The Quarterly Journal of Economics*, 5 2017, *132*, 665–712.

König, Michael, Kjetil Storesletten, Zheng Song, and Fabrizio Zilibotti, "From Imitation to Innovation: Where Is All That Chinese R&D Going?," *Econometrica*, 2022, *90*, 1615–1654.

Lehr, Nils H, "Does Monopsony Matter for Innovation?," 2025.

Lichter, Andreas, Max Löffler, Ingo E. Isphording, Thu Van Nguyen, Felix Poege, and Sebastian Siegloch, "Profit Taxation, R&D Spending, and Innovation," *American Economic Journal: Economic Policy*, 2025, *17*, 432–463.

Loecker, Jan De, Jan Eeckhout, and Gabriel Unger, "The Rise of Market Power and the Macroeconomic Implications," *The Quarterly Journal of Economics*, 2020, *135*, 561–644.

Lucking, Brian, Nicholas Bloom, and John Van Reenen, "Have R&D Spillovers Declined in the 21st Century?," *Fiscal Studies*, 12 2019, *40*, 561–590.

Manera, Andrea, "Competing for Inventors: Market Concentration and the Misallocation of Innovative Talent," 2022.

Mezzanotti, Filippo, "Roadblock to Innovation: The Role of Patent Litigation in Corporate R&D," *Management Science*, 2021.

_ **and Timothy Simcoe**, "Innovation, Patenting and Appropriability: Survey Evidence from a Nationally Representative Sample of U.S. Firms," *NBER Working Paper No. 31428*, 2025.

Midrigan, Virgiliu and Daniel Yi Xu, "Finance and misallocation: Evidence from plant-level data," *American Economic Review*, 2014, *104*, 422–458.

Olmstead-Rumsey, Jane, "Market Concentration and the Productivity Slowdown," 2025.

Peters, Ryan H. and Lucian A. Taylor, "Intangible capital and the investment-q relation," *Journal of Financial Economics*, 2017, *123*, 251–272.

Phillips, Gordon M and Alexei Zhdanov, "R&D and the Incentives from Merger and Acquisition Activity," *The Review of Financial Studies*, 2013, *26*, 34–78.

Restuccia, Diego and Richard Rogerson, "Policy distortions and aggregate productivity with heterogeneous establishments," *Review of Economic Dynamics*, 10 2008, *11*, 707–720.

Romer, Paul M, "Endogenous Technological Change," *Journal of Political Economy*, 1990, *98.*

Schubert, Gregor, Anna Stansbury, Bledi Taska, and Matthew Dey, "Employer Concentration and Wages: Evidence from Large Firms' Hiring Shocks," 2025.

Seegmiller, Bryan, "Valuing Labor Market Power : The Role of Productivity Advantages," 2025.

Terry, Stephen J., "The Macro Impact of Short-Termism," *Econometrica*, 2023, *91*, 1881–1912.

_ **, Toni M. Whited, and Anastasia A. Zakolyukina**, "Information Versus Investment," *Review of Financial Studies*, 2022, *36*, 1148–1191.

Wilson, Daniel J., "Beggar Thy Neighbor? The In-State, Out-of-State, and Aggregate Effects of R&D Tax Credits," *The Review of Economics and Statistics*, 2009, *91*, 431–436.

Yeh, Chen, Claudia Macaluso, and Brad Hershbein, "Monopsony in the U.S. Labor Market," *American Economic Review*, 2022, *112*, 2099–2138.

Appendix

A Proofs

Proof of Proposition 1. The proof of the proposition is entirely algebraic. Firstly, defining $\theta_{it} = \varphi_{it}\, V_{it}$ we can solve for firms' demand for R&D inputs as

$$\ell_{it} = \left(\frac{\theta_{it}\,\gamma}{(1+\Delta_{it})\,W_t}\right)^{\frac{1}{1-\gamma}}.$$

Plugging into the R&D resource constraint, we can solve for the R&D input price:

$$\frac{W_t}{\gamma} = L_t^{-(1-\gamma)}\,\left(\int_0^1 (\theta_{it}/(1+\Delta_{it}))^{\frac{1}{1-\gamma}}\,di\right)^{1-\gamma}.$$

Next, using the firm's first-order condition, we can express the economic growth rate as

$$g_t = A_t^{-\phi}\int_0^1 \zeta_{it}\,\ell_{it}\,\frac{W_t}{\gamma}\,(1+\Delta_{it})\,di.$$

Plugging in the definition of the wage and firms' R&D labor demand, we have

$$g_t = A_t^{-\phi}\,L_t^{\gamma}\,\frac{\int_0^1 \zeta_{it}\,\theta_{it}^{\frac{1}{1-\gamma}}\,(1+\Delta_{it})^{-\frac{\gamma}{1-\gamma}}\,di}{\left(\int_0^1 \theta_{it}^{\frac{1}{1-\gamma}}\,(1+\Delta_{it})^{-\frac{1}{1-\gamma}}\,di\right)^{\gamma}}.$$

Some rearrangement yields the formulae in the proposition. □

Proof of Corollary 1. The formula follows immediately since the terms in the numerator and denominator are expected values with normalized R&D productivity ω_{it} acting as a probability weight. Furthermore, and by Jensen's inequality, $\Xi_t \leq 1$ with equality in absence of dispersion in R&D wedges. The final statement follows immediately from the second order approximation provided in Lemma 1. □

Proof of Proposition 3. The planner problem is given by

$$\max \quad g_t = A_t^{-\phi}\int_0^1 \zeta_{it}\,z_{it}\,V_{it}\,di$$
$$\text{s.t.} \quad L_t = \int_0^1 \ell_{it}\,di \quad \text{and} \quad z_{it} = \varphi_{it}\,\ell_{it}^{\gamma}$$

The first-order conditions give rise to R&D input demand

$$\ell_{it} = \left(\frac{\zeta_{it}\, \theta_{it}\, \gamma}{\lambda_t^W} \right)^{\frac{1}{1-\gamma}}, \tag{A.1}$$

where λ_t^W is the shadow wage.

One can confirm immediately, that the implied allocation coincides with the competitive equilibrium iff $\zeta_{it}\,(1+\Delta_{it})$ is a constant. All proportional level differences are absorbed into the shadow wage and, thus, do not affect the allocation across firms. Thus, the planner can implement the growth maximizing allocation by setting $1+\Delta_{it} = 1/\zeta_{it}$. □

Lemma 1. *The second-order approximation of Ξ_t around $\zeta_{it} = \zeta$ and $\Delta_{it} = \Delta$ is given by*

$$\Xi_t \approx \exp\left(-\frac{1}{2}\, \frac{\gamma}{1-\gamma} \left(\sigma_\Delta^2 + 2\, \sigma_{\Delta,\zeta} \right) \right), \tag{A.2}$$

where σ_Δ^2 is the weighted variance of log R&D wedges and $\sigma_{\Delta,\zeta}$ is the ω_{it}-weighted covariance of log R&D and Impact-Value factors. If all variables are jointly log-normal, the approximation is precise and observation weights are unnecessary for calculating the variance and covariance.

Proof. The result follows immediately from the 2nd order approximation of Ξ_t around a no-dispersion point. □

Proof of Proposition 2. The proof for proposition follows by noting that the second-order approximation of Ξ_t in Lemma 1 can be expressed as

$$\Xi_t \approx \exp\left(-\frac{1}{2} \frac{\gamma}{1-\gamma} \sigma_\Delta^2\, \tilde{\beta} \right) \quad \text{with} \quad \tilde{\beta} = 1 + 2\, \frac{\sigma_{\Delta,\zeta}}{\sigma_\Delta^2}.$$

In turn, it is straightforward to show that a second order approximation of the formula in Proposition 2 yields the same expression. □

B Data Appendix

B.1 Data Definitions

Inventor employment. Let $\mathcal{P}_{it\to t+4}$ be the set of successful patent applications for firm i between t and $t+4$ and $\mathcal{I}_{it\to t+4}$ be the set of associated inventors. I denote the number of patents assigned to firm i and listing inventor j at time t as P_{ijt} and the total number of patents listing j as inventor as P_{jt}. My measure of inventors is then given by

$$\texttt{Inventors}_{it\to t+4} = \sum_{j\in\mathcal{I}_{it\to t+4}} \frac{\sum_{s=0}^{4} P_{ijt+s}}{\sum_{s=0}^{4} P_{jt+s}}. \tag{B.1}$$

Return on Capital. Following David et al. (2016), I measure the return on capital as the ratio of sales to beginning-of-period capital stock. Analogously to the R&D return, I construct the measure at the 5-year level:

$$\text{Return on Capital}_{it} \equiv \frac{\sum_{s=0}^{4} \texttt{Sales}_{it+s}}{\sum_{s=0}^{4} \texttt{Capital}_{it+s}}. \tag{B.2}$$

Tobin's Q. I define the (physical) investment Q as the ratio of firm valuation, defined as stock price times outstanding shares plus debt, net of cash holdings, ($\texttt{prcc\_f} \times \texttt{csho} + \texttt{dltt} + \texttt{dlc} - \texttt{act}$), to physical capital ($\texttt{ppeqgt}$).

Liquidity. I define liquidity as cash holdings divided by assets $\texttt{ch}/\texttt{at}$.

Dividend rate. I define the dividend rate as dividends over assets $\texttt{dvt}/\texttt{at}$.

Profit rate. I calculate the profit rate as earnings before extraordinary items ($\texttt{ib}$) divided by revenue ($\texttt{revt}$), both of which I aggregate at the 5-year horizon.

Public patent involvement. Patents are considered connected to public actors if assigned to a government entity, research lab, or university, or if they include a government interest statement. Public involvement is the share of patent valuations connected to public actors for the 5-year window.

Firm dominance is constructed in two steps. First, for each of a firm's new patents within a 5-year window, I calculate the share of inventors working for the firm among those that worked on patents in exactly the same technology classification. For the latter, I use the complete CPC classification of the patent, which contains more than 600 technology classes that are non-exclusive at the patent level. Patents of the same technology class are thus those that have exactly the same classifications as the patent under consideration. Second, I aggregate to the firm-level by taking a simple average over the firm's new patents. Note that the resulting measure is between 0 and 1 by construction with 1 implying maximal dominance and 0 implying no dominance.

Inventor specialization is constructed in two steps. First, inventor specialization is defined as the average cosine similarity between patent classifications in their portfolio of new patents within a 5-year window. For each patent, I create an indicator vector over the set of patent classifications with individual categories weighted by their inverse frequency. I then calculate the average cosine similarity across all patents in the portfolio. This measure is between 0 and 1, with 0 implying completely different patents and 1 implying that all patents have the same technology classification. Second, I aggregate to the firm-level by taking a patent-weighted average across the firm's inventors.

B.2 Sample Coverage

Table B.1: Sample Coverage

	Full sample			At least 50 patents		
	1975–2014	1975–90	2000–14	1975–2014	1975–90	2000–14
A. Patents						
Patents (KPSS)	63.1	67.4	59.8	58.8	62.1	56.9
Valuations (KPSS)	82.5	85.4	78.1	80.3	82.8	76.3
Patents (USPTO)	41.3	40.0	44.7	38.5	36.9	42.5
Citations (USPTO)	40.5	38.5	43.4	37.1	34.6	40.9
B. Compustat Firms with R&D						
Firms	36.0	35.6	37.5	13.5	13.0	15.3
R&D expenditure	90.3	91.8	89.2	76.0	78.1	77.2
Revenue	90.2	91.7	88.5	71.4	72.9	72.1
Capital Investment	94.7	94.5	96.0	76.3	76.9	81.0
Capital Stock	94.1	94.3	95.4	76.0	77.6	79.2
C. All Compustat Firms						
Firms	15.1	13.4	17.8	5.7	4.9	7.2
R&D expenditure	90.3	91.8	89.2	76.0	78.1	77.2
Revenue	48.0	54.7	42.2	38.1	43.7	34.4
Capital Investment	46.1	54.3	38.3	37.2	44.7	32.3
Capital Stock	43.5	51.3	35.9	35.3	42.6	29.8
D. BEA Statistics						
R&D expenditure	45.0	36.3	52.3	37.8	30.7	45.3

Note: Sample coverage in percent for final sample and sample restricted to observations with at least 50 patents for different periods compared to different population samples. Statistics are calculated annually and averaged over periods. Compustat sample excludes firms in the financial and utilities sectors and firms without a U.S. location, as proxied for by a missing state in Compustat.

C Empirical Appendix

C.1 Measurement Robustness

This section provides details for the adjustments made to R&D returns in Section 4.

Acquisitions may lead to measurement error in R&D returns due to not counting all R&D costs associated with granted patents. Suppose that the firm acquires some inventions that are subsequently patented and added to total value created $z_{it}\ V_{it}$; however, the costs are recorded as acquisition costs Aqc._{it} instead of R&D expenditure $\text{R\&D}_{it}$. Assuming that the firm is otherwise unconstrained, the measured R&D return then becomes

$$\frac{V_{it}\ z_{it}}{\text{R\&D}_{it}} = \frac{1}{\gamma}\ \left(1 + \frac{\text{Aqc.}_{it}}{\text{R\&D}_{it}}\right),$$

which may yield measured R&D return dispersion, even when true R&D returns are equalized, to the extent that acquisition intensities differ across firms.

I propose the following approach to investigating the importance of acquisitions for R&D return dispersion. First, I assume that firms use a fixed fraction s of total reported acquisition expenditures on innovative products such that $\text{Aqc.}_{it} = s\ \text{Total aqc.}_{it}$. Total acquisition expenditures are reported in Compustat. Second, assuming that the acquisition intensity is relatively small, we can estimate s as the semi-elasticity of R&D returns with respect to the total acquisition intensity using OLS. I find $s \in \{6.3\%, 8.6\%\}$ depending on the fixed effects. Finally, we can construct adjusted R&D returns as

$$\frac{V_{it}\ z_{it}}{\text{R\&D}_{it} + \hat{s}\ \text{Total aqc.}_{it}} = \frac{1}{\gamma}.$$

Fixed costs of R&D. Suppose firms face R&D fixed costs $f_i\ W_t$. Then, total R&D expenditure is $(f_i + \ell_{it})\ W_t$ and the frictionless R&D return is

$$\frac{V_{it}\ z_{it}}{(f_i + \ell_{it})\ W_t} = \frac{1}{\gamma}\ \frac{\ell_{it}}{f_i + \ell_{it}}. \tag{C.1}$$

Consequently, as long as firms face some fixed costs, their average R&D return will be increasing in their quantity of R&D conducted ℓ_{it}, which generates R&D return dispersion unrelated to frictions. Note, however, that the average R&D return for very large firms, i.e., $\ell_{it} >> f_i$, is still approximately constant.

I propose a simple approach to investigate the importance of fixed costs. First, I assume that fixed costs are identical within a NAICS3×5-year cell. Second, let $\bar{\Delta}$ be the average R&D return for a firm in the top 75th percentile and $\underline{\Delta}$ be the average R&D return for a firm in the 25th percentile. I can then estimate the industry-specific $\hat{\gamma}$ as the inverse of the average R&D return for firms in or above the 75th percentile of R&D expenditure. Finally, let $\underline{TC}$ be the average total R&D expenditure of a firm in the 25th percentile of the R&D

cost distribution. I can then estimate fixed costs and adjusted R&D returns as

$$\hat{f}_i\, W_t = \underline{TC}_i \left(1 - \frac{\underline{\Delta}_i}{\bar{\bar{\Delta}}_i}\right) \quad \text{and} \quad \frac{V_{it}\, z_{it}}{TC_{it} - \hat{f}_i\, W_t} = \frac{1}{\gamma}.$$

The measure will estimate larger fixed costs if firms with high R&D expenditure also tend to have much larger R&D returns and vice versa. The corrected R&D returns should no longer exhibit the size-related component of dispersion driven by fixed costs.

Patent valuations. Kogan et al. (2017) measure patent valuations based on the idea that a patent grant should increase the value of the firm by the unexpected part of the patent value. Let M be the valuation of the firm, V be the value of the patent, and π the ex-ante probability that the patent is granted. Then the change in firm valuation ΔM at the grant date should equal

$$\Delta M = (1 - \pi)\, V \quad \text{or, equivalently,} \quad V = \frac{\Delta M}{1 - \pi}. \tag{C.2}$$

They infer ΔM from stock market returns and assume that the probability that a patent is granted is constant across all patents. This assumption is restrictive for at least two reasons. First, patents of different patent classes may have different probabilities of being granted. For example, grant rates within 3.5 years range from below 30% to above 80% across CPC subclasses. Second, patent grant decisions are assumed to be independent of the value of the patent. Such an assumption would not hold, for example, if higher-quality patents are more valuable but also more likely to be granted.

I test whether either mechanism contributes to R&D return dispersion as follows. First, I use data on patent applications and grant decisions from the USPTO for 1991–2014 to calculate the grant probability at the CPC-subclass level and construct an adjusted valuation that takes into account differences in grant probabilities. For a given CPC subclass, I calculate the grant probability as the share of patents granted within 3.5 years of filing. Second, I allow the probability of patent rejection to depend on patent value and assume the form $1 - \pi_p = \pi_0\, V_p^{-\eta}$, where η measures the degree to which higher-value patents are also more likely to be granted. The adjusted patent valuation is

$$\tilde{V}_p = \left(V\, \frac{1 - \bar{\pi}}{\pi_0}\right)^{\frac{1}{1-\eta}}. \tag{C.3}$$

I calibrate π_0 at the annual level to keep the average patent valuation constant and experiment with alternative values for η. Ideally, η would be estimated directly; however, we only observe valuations for granted patents.

Table C.1: R&D Wedges, Economic Growth and Welfare — Additional Specification Robustness

Specification	Growth Impact $\Xi - 1$				Welfare Cost of Δ	
	1975–2014	1975–90	2000–14	Δ	End.	Semi-End.
Baseline	-21.0%	-13.4%	-26.1%	-14.7%	7.7%	7.4%
A. Acquisitions						
$s = 8.1\%$	-20.7%	-13.0%	-25.7%	-14.5%	7.6%	7.3%
$s = 6.5\%$	-20.7%	-13.1%	-25.7%	-14.5%	7.6%	7.3%
$s = 100\%$	-23.9%	-16.0%	-28.6%	-15.0%	7.9%	7.6%
B. Grant-Value Elasticity						
$\eta = 0.05$	-21.8%	-14.0%	-27.0%	-15.1%	7.9%	7.7%
$\eta = 0.25$	-27.1%	-17.7%	-33.2%	-18.8%	10.5%	10.0%
$\eta = 0.50$	-40.4%	-28.1%	-49.8%	-30.1%	20.3%	18.8%
$\eta = 1.50$	-84.4%	-81.0%	-87.9%	-36.4%	27.8%	25.1%

Notes: Table reports estimates for impact of R&D wedges across samples together with their implications for welfare. Changes in welfare are in consumption equivalent terms. See text and Appendix for details.

C.2 Profits and the Impact-Value Factor

The impact-value factor captures the wedge between the private incentives to conduct R&D and the productivity impact of that R&D. In standard expanding-variety growth models, heterogeneity in this wedge can arise only when differences in profitability reflect variation in market power rather than in demand shifters. In such models, heterogeneity in the quality of inventions—effectively demand shifters—induces proportional variation in both the private value and the productivity impact of inventions. Consequently, the impact-value factor is constant across firms and drops out of the expression for the Impact of R&D Wedges.

In contrast, in Schumpeterian growth models, heterogeneity in the step size generates a divergence between private value and productivity impact. The profit function is concave in the step size—because step size affects markups—while the productivity impact is linear. This mechanism parallels the insight in Aghion et al. (2025), who show that heterogeneity in markups across R&D projects produces misallocation, whereas pure differences in R&D productivity do not. Standard theory therefore implies that the impact-value factor should adjust for differences in profitability arising from heterogeneous profit rates.

I evaluate this implication empirically in Table C.2. Column (1) shows that R&D expenditure is positively correlated with profitability, while Column (2) shows that the impact-value factor is negatively correlated with it — consistent with firms whose large wedge captures a smaller share of the social value of their inventions. Columns (3)–(4) show that controlling for the impact-value factor substantially attenuates the R&D–profitability correlation, with the same pattern arising when measuring productivity impact via citations rather than the text-based measure.

Taken together, these results indicate that the empirical proxies for the impact-value factor behave as predicted by theory.

Table C.2: Profitability, R&D, and the Impact Value Wedge

	(1)	(2)	(3)	(4)
		Profit Rates		
R&D Expenditure	0.188***	0.087**	0.084**	0.133***
	(0.045)	(0.036)	(0.032)	(0.034)
Impact-Value Wedge		-0.276***	-0.279***	-0.149***
		(0.045)	(0.041)	(0.033)
Interaction			-0.111***	-0.111***
			(0.017)	(0.018)
Impact Measure	Scientific	Scientific	Scientific	Citations
Observations	9,082	9,082	9,082	9,082

Note: Profit rates are profits excluding R&D expenditure divided by sales. All variables in logs. All regressions control for NAICS3× Year fixed effects. Standard errors clustered at the NAICS6 level. See text for details.

Standard errors in parentheses. Significance levels: * 10% , ** 5%, *** 1%.

C.3 Constructing Estimates for Continuing Firms

I examine whether entry and exit contributed to the evolution of R&D Allocative Efficiency by constructing the measure using only continuing firms. I begin with the baseline estimate in 1975 and compute annual changes in the Impact of R&D wedges using only firms that are active in both adjacent years. These annual changes are then accumulated forward.

Formally, let $\hat{\Xi}_t$ denote the baseline estimate of the Impact of R&D wedges in year t. Let $\hat{\Xi}_t^{t,t-1}$ denote the estimate computed using only firms active in both t and $t-1$, and let $\hat{\Xi}_{t-1}^{t,t-1}$ be the corresponding estimate for $t-1$, computed on the same set of continuing firms.

The time series for the Impact of R&D wedges for continuing firms, $\hat{\Xi}_t^C$, is constructed as

$$\hat{\Xi}_t^C = \begin{cases} \hat{\Xi}_t, & t = 1975, \\ \hat{\Xi}_{t-1}^C \left(\frac{\hat{\Xi}_t^{t,t-1}}{\hat{\Xi}_{t-1}^{t,t-1}} \right), & t = 1976, \ldots, 2014. \end{cases} \tag{C.4}$$

This procedure isolates the contribution of continuing firms and identifies the portion of the trend in R&D Allocative Efficiency attributable solely to changes within incumbents.

C.4 Robustness for Aggregate Measures

Table C.3 reports estimates of R&D efficiency for alternative specifications.

Table C.3: R&D Wedges, Economic Growth and Welfare — Specification Robustness

Specification	Growth Impact $\Xi - 1$				Welfare Cost of Δ	
	1975–2014	1975–90	2000–14	Δ	End.	Semi-End.
Baseline	-21.0%	-13.4%	-26.1%	-14.7%	7.7%	7.4%
A. Fixed Effects						
Year	-25.6%	-18.3%	-30.5%	-15.0%	7.9%	7.6%
NAICS6 × Year	-16.3%	-9.3%	-21.5%	-13.4%	6.9%	6.7%
Technology Class × Year	-19.1%	-10.8%	-24.9%	-15.8%	8.4%	8.1%
B. Minimum Patents						
100 Patents	-20.3%	-12.0%	-25.9%	-15.8%	8.3%	8.1%
200 Patents	-19.4%	-10.7%	-25.8%	-16.8%	9.1%	8.7%
C. Time Horizon						
10–Year	-21.6%	-18.6%	-24.8%	-7.6%	3.6%	3.5%
20–Year	-21.4%	-21.6%	-24.2%	-3.3%	1.5%	1.5%
D. Scale Elasticity						
$\gamma = 0.6$	-28.8%	-19.1%	-35.3%	-20.0%	11.3%	10.8%
$\gamma = 0.4$	-15.1%	-9.2%	-19.2%	-10.9%	5.4%	5.3%

Notes: Table reports estimates for impact of R&D wedges across samples together with their implications for welfare. Changes in welfare are in consumption equivalent terms. See text and Appendix for details.

C.5 Unobserved Firms and Size Heterogeneity

The main results rely on the assumption that the sample is representative of firms conducting R&D. However, firms covered in Compustat—and especially those patenting frequently—tend to be substantially larger than the average firm in the economy, including many smaller firms that also conduct R&D. In this section, I develop a decomposition of R&D Allocative Efficiency for subsamples to investigate the implications of unrepresentative sampling.

Corollary 2. *Let there be a total mass M_t of firms whereof M_t^U are unobserved and M_t^O are observed. Then, one can decompose overall R&D Allocative Efficiency as:*

$$\Xi_t = \Omega_t^U \, \Xi_t^U + \Omega_t^O \, \Xi_t^O, \tag{C.5}$$

where $\omega_{it}^X = \theta_{it}^{\frac{1}{1-\gamma}} / (\int_{M_t^X} \theta_{it}^{\frac{1}{1-\gamma}} \, di)$ *for* $X \in \{U, O\}$ *and*

$$\Omega_t^X = \frac{L_t^X}{L_t} \left(\frac{\int_{M_t^X} \omega_{it}^X \, (1 + \Delta_{it})^{-\frac{1}{1-\gamma}} \, di}{\int_{M_t} \omega_{it} \, (1 + \Delta_{it})^{-\frac{1}{1-\gamma}} \, di} \right)^{-(1-\gamma)} .$$

Proof. The proof follows algebraically. □

Two insights follow from this Corollary. First, Ξ_t^O is an unbiased estimator for Ξ if distributions are comparable across unobserved and observed firms. Second, the bias in the evolution of estimated Ξ_t depends on whether unobserved firms are subject to similar trends or not. The estimates are conservative if the evolution of R&D Allocative Efficiency is worse for unobserved firms and vice versa. Assuming that unobserved firms resemble smaller firms in the Compustat sample and taking into account the estimates in Table 5, the Corollary would thus suggest that the main estimates are conservative and that R&D Allocative Efficiency may have declined even faster in a representative sample of firms.

C.6 An Alternative Decomposition

This section expands the main-text analysis to the Incentive Alignment Index Λ_t. Figure C.1 reports the evolution of the individual indices and their product, while Table C.4 reports average levels and changes. I consider several alternatives for the impact-value factor. Accounting for the Incentive Alignement Index unambiguously increases the implied costs of R&D misallocation and exacerbates its estimated rise, regardless of the measure of the impact-value factor considered. The patent-quality measure from Kelly et al. (2021) yields the most conservative estimates for the decline, 24% compared to 15% when only considering R&D Allocative Efficiency, and an average loss of 44%, compared to 21% in the baseline.

An important caveat to this analysis is that measurement error in the impact-value factor directly affects the Incentive Alignment Index, since it is a function of its variance, while it does not impact R&D Allocative Efficiency, since it only depends on its covariance with R&D returns. To the degree that measurement error in the impact-value factor is independent of the measured R&D return, it would not influence the covariance, but it would still significantly increase the variance of the impact-value factor itself.

Table C.4: The Impact of R&D Wedges and Incentive Alignment on Economic Growth and Welfare

Time Horizon	Growth Impact $\Xi - 1$		Incentive Alignment $\Lambda - 1$		Combined Effect $\Lambda\,\Xi - 1$	
	Estimate	Std. Err.	Estimate	Std. Err.	Estimate	Std. Err.
A. Scientific Impact						
1975–2014	-21.0%	(0.38%)	-29.3%	(0.62%)	-44.1%	(0.61%)
1975–1990	-13.4%	(0.42%)	-21.7%	(0.77%)	-32.2%	(0.74%)
2000–2014	-26.1%	(0.70%)	-30.1%	(1.24%)	-48.4%	(1.13%)
Δ Change	-14.7%		-10.7%		-23.8%	
B. Citations						
1975–2014	-25.6%	(0.44%)	-39.5%	(0.80%)	-54.9%	(0.66%)
1975–1990	-14.8%	(0.46%)	-23.2%	(0.81%)	-34.6%	(0.78%)
2000–2014	-32.9%	(0.82%)	-50.6%	(1.53%)	-66.8%	(1.09%)
Δ Change	-21.2%		-35.6%		-49.3%	
C. Long-run Citations						
1975–2014	-25.9%	(0.45%)	-40.5%	(0.79%)	-55.9%	(0.65%)
1975–1990	-14.8%	(0.46%)	-24.5%	(0.94%)	-35.7%	(0.89%)
2000–2014	-33.3%	(0.82%)	-50.6%	(1.44%)	-67.1%	(1.03%)
Δ Change	-21.8%		-34.6%		-48.8%	
D. Profit-rate						
1975–2014	-21.3%	(0.38%)	-60.6%	(0.42%)	-69.0%	(0.34%)
1975–1990	-14.8%	(0.46%)	-51.1%	(0.78%)	-58.3%	(0.64%)
2000–2014	-23.5%	(0.63%)	-62.8%	(0.61%)	-71.5%	(0.49%)
Δ Change	-10.3%		-23.8%		-31.7%	

Notes: Table reports estimates for impact of R&D wedges across samples and their welfare implications expressed in consumption equivalent terms. Standard errors and confidence intervals are calculated using a bootstrapping procedure. Adjusted values account for the covariance of R&D wedges and impact-value factors. Normalized values are further benchmarked against estimates in the life sciences industry. See text and Appendix for details.

Figure C.1: The Impact of Frictions and the Incentive Alignment Index

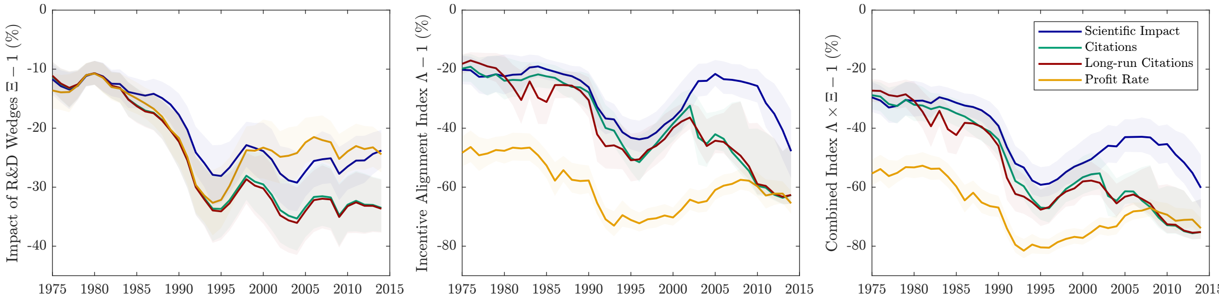

(a) Impact of Frictions (b) Incentive Alignment (c) Combined Impact

Notes: Figure reports annual estimates of R&D Allocative Efficiency (a), the Incentive Alignment Index (b), and their product (c) transformed into the implied impact on growth. See text for details.

Online Appendix

Not for publication

D Model Extensions

Specialization of R&D inputs. Workers might not be perfectly substitutable across firms and vice versa (Card et al., 2018). Such forces can be incorporated in the model by augmenting the R&D resource constraint to

$$L_t = \left(\int_0^1 \ell_{it}^{1+\xi} \, di \right)^{\frac{1}{1+\xi}}, \tag{D.1}$$

where $\xi > 0$ captures increasing marginal costs of R&D inputs to a given firm. As a result, firms' wages are potentially heterogeneous and take the form $W_{it} = W_t \, \ell_{it}^{\xi}$, where W_t is a common factor clearing the labor market. Firms' first-order conditions are given by

$$\frac{\partial z_{it}}{\partial \ell_{it}}\Big|_{\ell_{it}=\ell_{it}^*} V_{it} = (1+\Delta_{it}) \, W_t \, \ell_{it}^{\xi}. \tag{D.2}$$

Proposition 4 highlights that the main results carry over to this alternative setup; however, the effective scale elasticity is lower. Consequently, frictions tend to be less costly for larger ξ as reallocation of resources becomes less beneficial in a world with specialized inputs.

Proposition 4. *Under equations (2), (D.1), (D.2), and (5), we can express the economic growth rate in a Competitive Growth Equilibrium as the product of three terms:*

$$g_t = \underbrace{\frac{L_t^{\gamma}}{A_t^{\phi}} \left(\int_0^1 (\theta_{it} \, \zeta_{it})^{\frac{1}{1-\tilde{\gamma}}} \, di \right)^{1-\tilde{\gamma}}}_{=\text{ Frontier Growth Rate } g_t^F} \underbrace{\left(\int_0^1 \omega_{it} \, \tilde{\zeta}_{it}^{\frac{1}{1-\tilde{\gamma}}} \, di \right)^{\tilde{\gamma}-1}}_{\equiv \text{ Incentive Alignment } \Lambda_t} \underbrace{\frac{\int_0^1 \omega_{it} \, \tilde{\zeta}_{it} \, (1+\Delta_{it})^{-\frac{\tilde{\gamma}}{1-\tilde{\gamma}}} di}{\left(\int_0^1 \omega_{it} \, (1+\Delta_{it})^{-\frac{1}{1-\tilde{\gamma}}} di \right)^{\tilde{\gamma}}}}_{\equiv \text{ R\&D Efficiency } \Xi_t}, \tag{D.3}$$

where $\tilde{\zeta}_{it} = \zeta_{it} / \left(\int_0^1 \omega_{it} \, \zeta_{it} di \right)$ *and* $\omega_{it} = \theta_{it}^{\frac{1}{1-\tilde{\gamma}}} / \left(\int_0^1 \theta_{it}^{\frac{1}{1-\tilde{\gamma}}} di \right)$ *are the normalized impact-value factor and an R&D productivity weight, respectively, and* $\tilde{\gamma} \equiv \frac{\gamma}{1+\xi}$ *is the adjusted scale elasticity.*

Proof. R&D input demand is given by

$$\ell_{it} = \left(\frac{\theta_{it} \, \gamma}{(1+\Delta_{it}) \, W_t} \right)^{\frac{1}{1-\gamma+\xi}}.$$

We can then solve for the growth rate using the R&D input demand and supply constraint:

$$g_t = \frac{L_t^\gamma}{A^\phi} \frac{\int_0^1 \zeta_{it}\, \theta_{it}^{\frac{1+\xi}{1-\gamma+\xi}}\, (1+\Delta_{it})^{-\frac{\gamma}{1-\gamma+\xi}}\, di}{\left(\int_0^1 \theta_{it}^{\frac{1+\xi}{1-\gamma+\xi}}\, (1+\Delta_{it})^{-\frac{1+\xi}{1-\gamma+\xi}}\, di\right)^{\frac{\gamma}{1+\xi}}}.$$

Defining $\tilde{\gamma} = \frac{\gamma}{1+\xi}$ yields the formulae in the proposition. □

Note that the unit cost elasticity of R&D identifies $\tilde{\gamma}$ in this setup. Thus, this extension does not necessarily change the quantitative implications as long as γ is calibrated to match the unit cost elasticity in the baseline model.

Multiple R&D lines. Consider an alternative version of the model with multiple R&D lines per firm. I index a firm by $i \in \mathcal{I}$ and an R&D line by $j \in \mathcal{J}_i$. The production function is given by

$$z_{ij} = \varphi_{ij}\, \ell_{ij}^\gamma. \tag{D.4}$$

Firms' first-order conditions for R&D inputs at the R&D line level are

$$\gamma \ell_{ij}^{1-\gamma}\, \theta_{ij} = (1+\Delta_{ij})W. \tag{D.5}$$

We can solve for the R&D wage as

$$\frac{W}{\gamma} = L^{-(1-\gamma)} \left(\int_{\mathcal{I}} \Big(\sum_{j\in\mathcal{J}_\rangle} (\theta_{ij}/(1+\Delta_{ij}))^{\frac{1}{1-\gamma}} \Big) di \right)^{1-\gamma}. \tag{D.6}$$

The economic growth rate is then

$$g = \int_0^1 \left(\sum_{j\in\mathcal{J}_i} \zeta_{ij}\ z_{ij}\ V_{ij} \right)\ di = \frac{L^\gamma}{A_t^\phi} \frac{\int_0^1 \left(\sum_{j\in\mathcal{J}_i} \zeta_{ij}\, \theta_{ij}^{\frac{1}{1-\gamma}}\, (1+\Delta_{ij})^{-\frac{\gamma}{1-\gamma}} \right)\ di}{\left(\int_0^1 \left(\sum_{j\in\mathcal{J}_i} \theta_{ij}^{\frac{1}{1-\gamma}}\, (1+\Delta_{ij})^{-\frac{1}{1-\gamma}} \right)\ di \right)^{\gamma}}. \tag{D.7}$$

Next, consider the inputs at the firm level, measured as

$$\begin{aligned} 1+\Delta_i &= \frac{\sum_{j\in\mathcal{J}_i} \theta_{ij}\ \ell_{ij}^\gamma}{W\ \sum_{j\in\mathcal{J}_i} \ell_{ij}} = \sum_{j\in\mathcal{J}_i} \frac{\ell_{ij}}{\ell_i}\ (1+\Delta_{ij}) \\ \zeta_i &= \frac{\sum_{j\in\mathcal{J}_i} z_{ij}\ V_{ij}^P}{\sum_{j\in\mathcal{J}_i} z_{ij}\ V_{ij}} = \sum_{j\in\mathcal{J}_i} \frac{\theta_{ij}\ \ell_{ij}^\gamma}{\sum_{j\in\mathcal{J}_i} \theta_{ij}\ \ell_{ij}^\gamma}\ \zeta_{ij} \\ \theta_i &= (1+\Delta_i)\ \tilde{W}^\gamma\ (\tilde{W}\ \ell_i)^{1-\gamma} \end{aligned} \tag{D.8}$$

Some algebra confirms the familiar growth rate formula

$$g = \frac{L^{\gamma}}{A^{\phi}} \frac{\int_0^1 \zeta_i \, \theta_i^{\frac{1}{1-\gamma}} \, (1+\Delta_i)^{-\frac{\gamma}{1-\gamma}} \, di}{\left(\int_0^1 \theta_i^{\frac{1}{1-\gamma}} \, (1+\Delta_i)^{-\frac{1}{1-\gamma}} \, di\right)^{\gamma}}. \tag{D.9}$$

Thus, the growth rate abstracting from the product-line level heterogeneity recovers the growth rate under full heterogeneity under the proposed measurement approach.

Abundant resources. Suppose aggregate supply of L_t responds to productivity-adjusted wage W_t such that

$$L_t = \bar{L}_t \, \left(\frac{W_t}{Y_t}\right)^{\frac{\xi}{1-\gamma}}, \tag{D.10}$$

where $\bar{L}_t$ is given exogenously and $\xi/(1-\gamma)$ is the aggregate supply elasticity. Also, let L_t^* be the supply in the absence of frictions, i.e., when the R&D wage is at its frictionless level.

Proposition 5. *Under equations (2)-(5) and (D.10), we can express the economic growth rate in a Competitive Growth Equilibrium using the sample decomposition as in Proposition 1 with two adjustments. First, the frontier growth rate g_t^F reflects the frictionless R&D input supply,*

$$g_t^F = \frac{L_t^{*\gamma}}{A_t^{\phi}} \, \left(\int_0^1 (\theta_{it} \, \zeta_{it})^{\frac{1}{1-\gamma}} \, di\right)^{1-\gamma}, \tag{D.11}$$

and, second, R&D efficiency also reflects the potential effect on labor supply

$$\Xi_t = \frac{\int_0^1 \omega_{it} \, \tilde{\zeta}_{it} \, (1+\Delta_{it})^{-\frac{\gamma}{1-\gamma}} di}{\left(\int_0^1 \omega_{it} \, (1+\Delta_{it})^{-\frac{1}{1-\gamma}} di\right)^{\gamma}} \, \left(\int_0^1 \omega_{it} \, (1+\Delta_{it})^{-\frac{1}{1-\gamma}} \, di\right)^{\frac{\xi \, \gamma}{1+\xi}}. \tag{D.12}$$

Note that the supply elasticity only appears in the second term, which depends on the productivity-weighted average level of frictions. Any change in frictions or policy that keeps constant this average thus has the same effect on growth as in the case of $\xi = 0$.

Proof. The proof follows from the same steps as in the derivation of the baseline results. □

Note, however, that the adjusted formulas tend to be less sensitive to variation in R&D returns. Intuitively, with flexible labor supply, excess demand for R&D workers tends to lead to more aggregate R&D employment instead of crowding-out demand from other firms.

Proposition 6. *Suppose that R&D returns, impact-value factors, and R&D productivity are jointly log-normally distributed and that R&D returns and impact-value factors are either* *positively or uncorrelated. Then, R&D efficiency is declining in the dispersion of log-R&D wedges as long as the supply of R&D inputs is sufficiently inflexible:* $\frac{1}{\gamma} > \frac{\xi}{1-\gamma}$*. Furthermore, holding constant the average level of R&D wedges, the Impact of R&D wedges is declining in the dispersion of R&D wedges as long as* $\gamma \geq \frac{\xi}{1+2\xi}$*.*

Proof. Solving for Ξ_t under log-normal distribution and setting $\mu_\Delta = 0$, we have

$$\ln \Xi_t = -\frac{1}{2} \frac{\gamma}{(1-\gamma)^2} \left(\gamma - \frac{1}{1+\xi}\right) \sigma_\Delta^2.$$

It is straightforward to show that this term is decreasing in σ_Δ^2 if and only if $\frac{1}{\gamma} > \frac{\xi}{1-\gamma}$. Alternatively, setting $\mu_\Delta = -\frac{1}{2}\sigma_\Delta^2$ to maintain the average level of $1 + \Delta_{it}$, we have

$$\ln \Xi_t = -\frac{1}{2} \left(\frac{\gamma}{1-\gamma} - \frac{\xi}{1+\xi}\right) \sigma_\Delta^2,$$

which is declining in σ_Δ^2 as long as the condition in the proposition holds. □

Importantly, aggregate estimates suggest that $\frac{\xi}{1-\gamma}$ is around 0.5 and, thus, satisfies the more stringent constraint given that $\gamma = 0.5$ (Chetty et al., 2012).

Free Entry. Suppose that the mass M_t of innovative firms is potentially responsive to changes in the economic environment and let M_t^* be the mass of firms in the absence of frictions. The equilibrium wage satisfies

$$\frac{W_t}{\gamma} = \left(\frac{M_t}{L_t}\right)^{1-\gamma} \left(\int_0^1 \theta_{it}^{\frac{1}{1-\gamma}} (1+\Delta_{it})^{-\frac{1}{1-\gamma}} di\right)^{1-\gamma}. \tag{D.13}$$

I assume that all firm types are permanent and that frictions Δ_{it} show up directly in the firm's cost function. The current period profits of an innovative firm are given by

$$\begin{aligned} \pi_{it} &\equiv \max \{\theta_{it}\, \ell_{it}^\gamma - W_t\, \ell_{it}\, (1+\Delta_{it})\} \\ &= (1-\gamma)\, \theta_{it}^{\frac{1}{1-\gamma}} \left((W_t/\gamma)\, (1+\Delta_{it})\right)^{-\frac{\gamma}{1-\gamma}}. \end{aligned}$$

Assuming a constant discount factor and permanent types implies that current and expected, discounted value are proportional by factor $R/(R-1)$, where R is the discount rate. The expected value of an R&D firm is then given by

$$\begin{aligned}\mathcal{V}_t &= \mathbb{E}_t\left[\frac{R}{R-1}\,\pi_{it}\right]\\ &= \frac{R\,(1-\gamma)}{R-1}\,\left(\frac{L_t}{M_t}\right)^{\gamma}\,\left(\int_0^1 \theta_{it}^{\frac{1}{1-\gamma}}\,di\right)^{1-\gamma}\,\frac{\int_0^1 \omega_{it}\,(1+\Delta_{it})^{-\frac{\gamma}{1-\gamma}}\,di}{\left(\int_0^1 \omega_{it}\,(1+\Delta_{it})^{-\frac{1}{1-\gamma}}\,di\right)^{\gamma}}\end{aligned}$$

Assuming that entrants draw values from a random firm in the existing distribution, entrants receive expected value $\mathcal{V}_t$ and in turn need to pay entry cost. I consider two alternatives. In the first case, entry costs are in units of the output and given by $\phi_t^E\,\frac{R\,(1-\gamma)}{R-1}\,M_t^{\frac{\gamma}{\varphi}}$. The free entry condition is

$$\mathcal{V}_t = \phi_t^E\,\frac{R\,(1-\gamma)}{R-1}\,M_t^{\frac{\gamma}{\varphi}}$$

Using the formula for value of entry, we can then solve for equilibrium entry:

$$\frac{M_t}{M_t^*} = \left(\frac{\int_0^1 \omega_{it}(1+\Delta_{it})^{-\frac{\gamma}{1-\gamma}}di}{\left(\int_0^1 \omega_{it}(1+\Delta_{it})^{-\frac{1}{1-\gamma}}di\right)^{\gamma}}\right)^{\frac{1}{\gamma}\frac{\varphi}{1+\varphi}} \quad \text{s.t. } M_t^* = \left(\frac{L_t}{\phi_t^{E\,\frac{1}{\gamma}}}\left(\int_0^1 \theta_{it}^{\frac{1}{1-\gamma}}di\right)^{\frac{1-\gamma}{\gamma}}\right)^{\frac{\varphi}{1+\varphi}}. \tag{D.14}$$

Note that $\varphi \to 0$ recovers the baseline model with $M_t = 1$, while $\varphi \to \infty$ yields a standard free entry condition. In general, larger values of φ make the mass of firms more responsive to the economic environment.

In the second case, I assume that entry costs are linked to the R&D wage and given by $\phi_t^E\,(1-\gamma)\,M_t^{\frac{1}{\varphi}}\,\frac{W_t}{\gamma}$. The free entry condition is

$$\mathcal{V}_t = \phi_t^E\,\frac{R\,(1-\gamma)}{R-1}\,M_t^{\frac{1}{\varphi}}\,\frac{W_t}{\gamma}$$

Using the formula for value of entry, we can then solve for equilibrium entry:

$$\begin{aligned}\frac{M_t}{M_t^*} &= \left(\frac{\int_0^1 \omega_{it}(1+\Delta_{it})^{-\frac{\gamma}{1-\gamma}}di}{\left(\int_0^1 \omega_{it}(1+\Delta_{it})^{-\frac{1}{1-\gamma}}di\right)^{\gamma}}\,\left(\int_0^1 \omega_{it}(1+\Delta_{it})^{-\frac{1}{1-\gamma}}di\right)^{\gamma-1}\right)^{\frac{\varphi}{1+\varphi}}\\ \text{s.t. } M_t^* &= \left(\frac{L_t}{\phi_t^E}\right)^{\frac{\varphi}{1+\varphi}}.\end{aligned} \tag{D.15}$$

Proposition 7. *Under equations (2)-(5) and (D.14) or (D.15), we can express the economic growth rate in a Competitive Growth Equilibrium using the sample decomposition as in Proposition 1 with two adjustments. First, the frontier growth rate g_t^F reflects frictionless* *entry,*

$$g_t^F = \frac{L_t^\gamma}{A_t^\phi} M_t^{*1-\gamma} \left(\int_0^1 (\theta_{it}\ \zeta_{it})^{\frac{1}{1-\gamma}}\ di\right)^{1-\gamma}, \tag{D.16}$$

and, second, R&D efficiency also reflects potential effects on entry

$$\Xi_t = \frac{\int_0^1 \omega_{it}\ \tilde{\zeta}_{it}\ (1+\Delta_{it})^{-\frac{\gamma}{1-\gamma}} di}{\left(\int_0^1 \omega_{it}\ (1+\Delta_{it})^{-\frac{1}{1-\gamma}} di\right)^\gamma} \frac{M_t}{M_t^*}, \tag{D.17}$$

where M_t/M_t^* is given by the respective formulas.

Proof. The proof follows the same steps as before apart from taking the number of firms as a variable and using the entry condition. □

Private frictions now have an additional detrimental effect on growth through the number of firms. Notably, the entry-effect does not depend on the impact-value factor, which is irrelevant to firms' decision to enter or exit the economy. It is straightforward to show that the impact of frictions is more negative in the economy with free entry holding constant the average level of R&D returns.

Knowledge externalities: Knowledge spillovers are central to modern growth theory and constitute a divergence between the private and public incentives to conduct R&D. However, as discussed in the main text, they are different from the impact-value factor. In this section, I briefly discuss their impact in theory and propose a methodology to filter out knowledge externalities, if any, from the empirical measures of the impact-value factor.

The model presented in the main text implicitly hides knowledge externalities as an intercept of R&D productivity, as is standard in the literature. However, we can consider a more explicit formulation potentially allowing for heterogeneous knowledge spillovers. Consider the following extension of the formula for economic growth presented in the main text:

$$g_t = \theta_t \int_0^1 z_{it}\ V_{it}\ \zeta_{it}\ di \quad \text{with} \quad \theta_t = \left(\int_0^1 z_{it}\ V_{it}\ \zeta_{it}\ \theta_{it}\ di\right)^\phi. \tag{D.18}$$

The term θ_t captures knowledge externalities, which are assumed instantaneous to ease the exposition, while the term θ_{it} captures the degree to which they differ from the productivity impact. The optimization problem from the perspective of the firm remains identical, since it ignores knowledge spillovers. From the perspective of the growth-maximizing allocation, the first-order conditions now state

$$\frac{\partial z_{it}}{\partial \ell_{it}} \, \theta_t \, V_{it} \, \zeta_{it} \, (1 + \tilde{\phi}_t \, \theta_{it}) = \tilde{W}_t \quad \text{where} \quad \tilde{\phi}_t = \phi \, \frac{\int_0^1 z_{it} \, V_{it} \, \zeta_{it} \, di}{\int_0^1 z_{it} \, V_{it} \, \zeta_{it} \, \theta_{it} \, di}$$

A couple of insights follow. First, note that if θ_{it} is constant across firms, i.e., when productivity impact and knowledge spillovers perfectly align, then the additional term is a constant and, thus, does not influence the optimal allocation of R&D workers across firms. This case is the standard in the literature. Second, even if this is the case, the impact of achieving a first-best allocation on growth would be amplified, since it would increase the value of θ_t. Third, if the productivity and knowledge spillovers diverge, then the first-best allocation may differ from the one established in the main text, since it would shift resources towards firms with high relative knowledge spillovers.

Beyond the theoretical implication, the exercise also raises the question of whether the empirical proxies for the impact value wedge are also capturing heterogeneity in knowledge spillovers. If so, they would potentially bias estimates for the R&D Allocative Efficiency, to the degree that they are not purely orthogonal to R&D wedges. I investigate this concern empirically in Online Appendix E.5.

E Additional Empirical Results

E.1 The Realization-Expectation Gap

The frictionless firm R&D investment model equalizes the expected R&D return across all firms. Variation in realized returns could still arise due to the stochastic nature of innovation:

$$\underbrace{\ln\left(\frac{z_{it}\,V_{it}}{\ell_{it}\,W_t}\right)}_{\text{Realized R\&D Return}} = \underbrace{\ln\left(\frac{\mathbb{E}[z_{it}\,V_{it}]}{\ell_{it}\,W_t}\right)}_{\text{Expected R\&D Return}} + \underbrace{\ln\left(\frac{z_{it}\,V_{it}}{\mathbb{E}[z_{it}\,V_{it}]}\right)}_{\text{Realization-Expectation Gap}}, \tag{E.1}$$

where expectations reflect the firm's information set at the time of the investment decision.

When firms optimize and have rational expectations, the realization-expectation gap should be an i.i.d. random variable and thus should not be predictable using information within the firm's information set at the time of the investment decision. Furthermore, this property should extend to realized R&D returns when expected returns are equalized.

This prediction is not borne out by the data as past R&D returns are a consistent predictor of future R&D returns. I confirm this by estimating a simple autoregressive model:

$$\ln \text{R\&D Return}_{it} = \alpha_{j(i)t} + \beta_h \ln \text{R\&D Return}_{it-h} + \epsilon_{it}, \tag{E.2}$$

where $\alpha_{j(i)t}$ are industry-year fixed effects. Following the argument above, we expect $\beta_h = 0$ as long as $\ln \text{R\&D Return}_{it-h}$ was in the information set of the firm making its investment decision and expected R&D returns are equalized.

I test this hypothesis for several values of horizon h. I choose 5 years as the minimum horizon to ensure that the information contained in both returns is non-overlapping; however, longer horizons might be more reliable since they provide a clearer delineation between the lagged and current returns.

The estimates in Figure E.2 suggest that R&D returns are highly predictable, even at longer horizons, using lagged values as predictors. The autocorrelation coefficient at the 5-year horizon is 0.7 and declines to 0.3 for the 15-year horizon as one would expect, e.g., in a standard AR(1) model. I thus transform the coefficients by taking their h-th root to get an implied annual autocorrelation coefficient. The estimated values are consistently above 0.9 confirming that R&D returns are indeed highly persistent and, thus, predictable. We can therefore reject the hypothesis that variation in R&D returns is primarily driven by the stochastic nature of realization vs expectation.

Figure E.2: R&D Returns are Highly Persistent

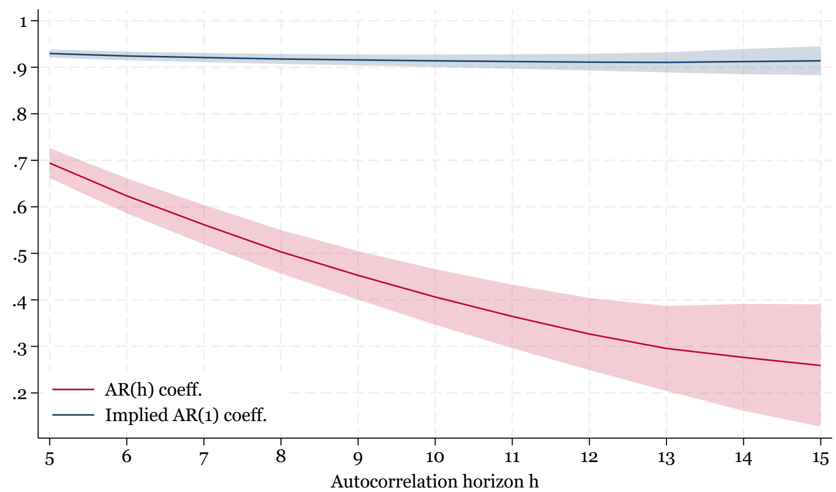

Notes: This figure plots estimated autocorrelation coefficients together with their implied annualized values. 95% confidence intervals are shaded. All regressions control for industry-year fixed effects and standard errors are clustered at the NAICS3 level. Standard errors for the implied coefficients are calculated using the Delta method.

E.2 R&D Wedges and Impact-Value Factors Over Time

Figure E.3 investigates the evolution of the relationship between impact-value factors and R&D wedges over time. The baseline estimate, based on the text-based patent quality measure, declines until 1990 before increasing until 2005 and subsequently declining again. The implied adjustment factor would dampen R&D return variation during the 1990s and slightly amplify it around 2005. In my quantitative exercise, I limit the factor to take values weakly less than 1. Citation-based measures suggest a much stronger and more monotonic increase with significant implied amplification post 2000. In contrast, the profit-rate based measure suggests essentially no adjustment, while the markup-based measure suggests some dampening around 2005.

E.3 Model Prediction and Industry Trends

Beyond the prediction linking aggregate R&D misallocation to productivity growth, the model also makes predictions about R&D performance at the industry level that can be tested empirically. Below, I confirm the model's predictions for the link between R&D Allocative Efficiency and industry-level R&D expenditure or R&D returns.

The model predicts a positive correlation between R&D allocative efficiency and produc-

Figure E.3: Evolution of the Relationship between R&D Wedges and Impact-Value Factors

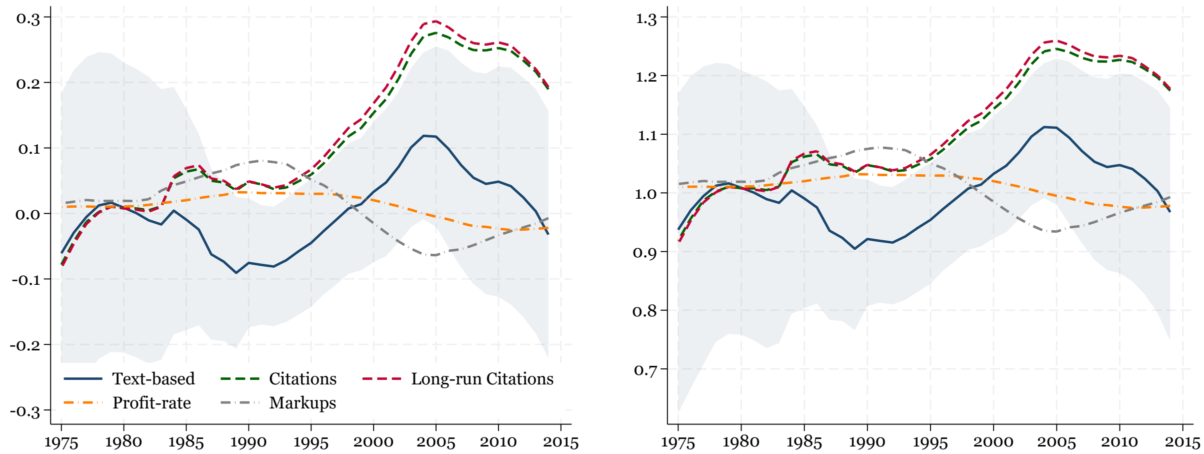

(a) Regression Coefficients

(b) Implied Adjustment Factor

Notes: Left panel reports regression coefficient of impact-value factors on R&D returns as in Table 2 estimated over rolling 10-year windows. Standard errors are clustered at the industry level and regressions control for industry-year fixed effects. The right panel transforms the regression coefficients into implied adjustment factors and calculates their standard error using the Delta method. Shaded area reports 95% confidence interval for the main estimate.

tivity growth, all else equal. As per the model formula,

$$\ln g_t = \ln g_t^F + \ln \Gamma_t + \ln \Xi_t.$$

There are two challenges with investigating this relationship at the industry level. First, it is not clear at what time horizon R&D flows into productivity growth in practice. R&D is typically associated with technology development, which is different from deployment and, thus, productivity growth. Second, the formula need not hold at the industry level, e.g., if R&D workers can freely move across industries. In that case, misallocation is not reflected in industry R&D wages and, resultingly, the relationship between productivity growth and misallocation could even go in the opposite direction. In particular, suppose that $(1 + \Delta_{it})$ and θ_{it} are log-normally distributed in cross-section and $\Delta_{it} \perp \zeta_{it}$. Then, we have

$$\ln \Xi_{it} = -\frac{1}{2} \frac{\gamma}{1-\gamma} \sigma_{i,\Delta}^2.$$

Furthermore, we can solve for the industry growth rate as

$$\ln g_{it} = -\frac{\gamma}{1-\gamma} \ln W_t + \frac{1}{2} \left[\left(\frac{\gamma}{1-\gamma} \right)^2 \sigma_{i,\Delta}^2 + \left(\frac{1}{1-\gamma} \right)^2 \sigma_{i,\theta}^2 - 2 \frac{\gamma}{(1-\gamma)^2} \sigma_{i,\Delta\theta} \right]$$

It follows that there could be even a strong negative correlation between both measures unless $\sigma_{i,\Delta\theta}$ and $\sigma^2_{i,\Delta}$ are strongly correlated across industries. Key to this conclusion is that $\ln W_t$ does not co-vary with $\sigma^2_{i,\Delta}$.

Thus, I instead focus on total R&D expenditure and returns. In particular, one can show that the total R&D return for an industry i is given by

$$\ln\left(\frac{\int_{j\in\mathcal{I}_i} z_{jt}\, V_{jt}dj}{\int_{j\in\mathcal{I}_i} \ell_{jt}\, W_t dj}\right) = -\frac{1}{2}\left(\frac{1-\gamma^2}{(1-\gamma)^2}\,\sigma^2_{i,\Delta} - 2\,\frac{1+\gamma}{1-\gamma}\,\sigma_{i,\Delta\theta}\right).$$

Furthermore, total R&D expenditure is given by

$$\ln\left(\int_{j\in\mathcal{I}_i} \ell_{jt}\, W_t dj\right) \approx -\frac{\gamma}{1-\gamma}\ln W_t + \frac{1}{2}\frac{1}{1-\gamma}\left(\sigma^2_{i,\Delta} + \sigma^2_{i,\theta} - 2\,\sigma_{i,\Delta\theta}\right)$$

Thus, both measures should be negatively correlated with R&D Efficiency at the industry level if the covariance of $\sigma_{i,\Delta\theta}$ and $\sigma_{i,\Delta}$ is sufficiently small.

I investigate this relationship for 10-year differences at the industry level using OLS. Differences ensure that I focus on changes within industries rather than permanent heterogeneity. As reported in Table E.1, the predictions are borne out in the data. Industries with lower allocative efficiency spend more on R&D and have weakly higher aggregate R&D returns. Intuitively, the documented relationship reflects the waste from R&D misallocation.

Table E.1: R&D Allocative Efficiency and Performance at the Industry Level

	(1)	(2)	(3)	(4)
	Δ R&D Expenditure			
Δ R&D Allocative Efficiency	-0.894***	-0.786***	-0.764***	-0.630***
	(0.180)	(0.156)	(0.174)	(0.158)
	Δ R&D Return			
Δ R&D Allocative Efficiency	-0.926***	-1.007***	-0.102	-0.112
	(0.297)	(0.312)	(0.232)	(0.245)
Industry FEs		✓		✓
Year FEs			✓	✓
Observations	900	900	900	900

Note: All variables in 10-year differences over 5-year aggregates. An observation is an industry-year. Robust standard errors in parentheses.

Standard errors in parentheses. Significance levels: * 10% , ** 5%, *** 1%.

E.4 R&D Returns and the Scale Elasticity

The R&D return is the product of the inverse R&D scale elasticity and the R&D wedge:

$$\frac{z_{it}\, V_{it}}{\ell_{it}\, W_t} = \frac{1}{\gamma_{it}}\,(1+\Delta_{it}). \tag{E.3}$$

To identify the R&D wedge, I assume that the R&D scale elasticity is common within industry-year cells. Under this assumption, we can recover the R&D wedge by removing industry-year fixed effects from the observed R&D return. This procedure also removes any systematic variation in R&D returns across industry-year cells and therefore yields a conservative measure of the overall dispersion in R&D wedges, provided the identification assumption is valid.

A natural concern is that the scale elasticity may vary within industry-year cells, which would generate measurement error in the recovered R&D wedges. As discussed in the main text, this type of measurement error would bias the estimated R&D Allocative Efficiency downwards and therefore bias upward my estimates of the role of R&D misallocation.

To investigate this concern, I estimate the scale elasticity directly and examine whether it varies systematically with the R&D return. In brief, I find no evidence that the R&D scale elasticity differs systematically across the R&D return distribution. This finding suggests that heterogeneity in the scale elasticity is unlikely to be a major contributor to R&D return dispersion or measured misallocation.

The dominant approach in the literature to estimating the R&D scale elasticity exploits variation in R&D tax credits. Let τ_{it} be the firm-specific R&D tax credit. In the absence of frictions, equilibrium gross R&D expenditure satisfies

$$\ln(\ell^*_{it}\, W_t) = \frac{1}{1-\gamma_{it}}\ln(\gamma_{it}\,\theta_{it}) - \frac{\gamma_{it}}{1-\gamma_{it}}\ln W_t - \frac{1}{1-\gamma_{it}}\,\ln(1-\tau_{it}) - \frac{1}{1-\gamma_{it}}\,\ln(1-\Delta_{it}).$$

If $\gamma_{it}=\gamma$ is homogeneous across firms, then γ can be estimated by regressing log R&D expenditure on $\ln(1-\tau_{it})$. The literature follows this approach using state-level variation in R&D tax credits (Lucking et al., 2019). I follow the same approach and merge the effective R&D tax credits from Lucking et al. (2019) to firms via their headquarters state.

Heterogeneous scale elasticities introduce additional complications. The elasticity cannot be identified at the firm level, and it is not clear how to group firms ex-ante into cells with homogeneous elasticities other than through industry classifications—where the fixed-effects approach is already more conservative.

Because the central concern is whether the scale elasticity varies systematically with the R&D return, I estimate the elasticity by quantiles of the R&D return distribution. Specifically, I estimate

$$\ln \text{R\&D Expenditure}_{it} = \alpha_i + \gamma_t + \sum_{i=1,\ldots,10} \beta_j \left\{\text{R\&D Return}_{it} \in \mathcal{D}_{jt}\right\} \ln\left(1 - \tau_{s(i)t}\right) + \epsilon_{it}, \quad \text{(E.4)}$$

including firm and year fixed effects. The set $\mathcal{D}_{jt}$ denotes the j-th decile of the (tax-adjusted) R&D return distribution in year t. Firms are linked to R&D tax credits through their headquarters state, and industry affiliation is absorbed in the firm fixed effects. The implied scale elasticities follow from $\frac{\hat{\beta}_j + 1}{\hat{\beta}_j}$ with standard errors calculated via the Delta-method. The estimates are unbiased so long as the true elasticity does not vary within each decile.

Panel A of Figure E.4 shows that the sensitivity of R&D expenditure to R&D tax credits does not vary systematically across deciles of the R&D return distribution. Panel B shows that the corresponding implied scale elasticities cluster tightly around 0.5. Together, these findings indicate that heterogeneity in the R&D scale elasticity is unlikely to be a primary driver of R&D return dispersion.

Figure E.4: Estimates for β and γ across the R&D Return Distribution

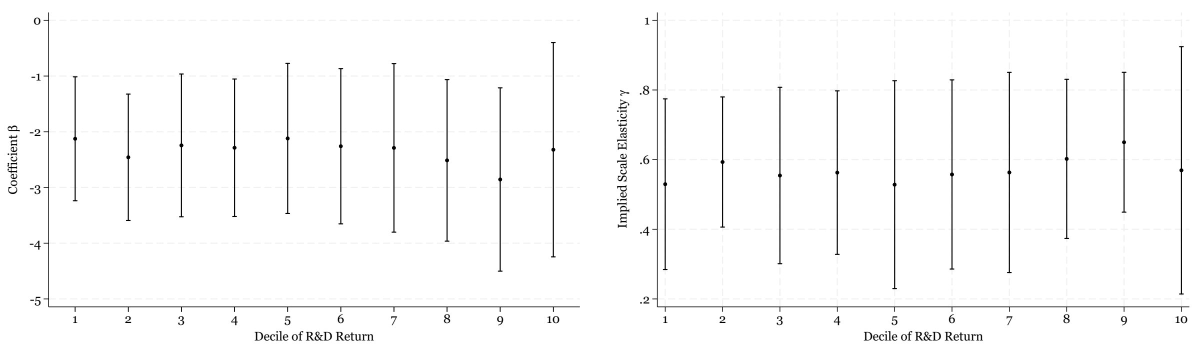

A. Estimates of β_j B. Implied Estimates of γ_j

Notes: Panel A reports coefficient estimates from regressing log R&D expenditure on the implied log user costs of R&D for a sample of Compustat firms. Regression controls for firm, state, industry, and year fixed effects. Standard errors are clustered at the NAICS6 level. The sample restricts to firms with at least 10 patents over the subsequent 5-year window and within-year deciles of the R&D return distribution are taken over 1-year R&D returns.

E.5 Spillovers

Knowledge externalities can create a separate wedge between private and planner incentives as discussed in Online Appendix D. Furthermore, they might also be reflected in the empirical measures of the impact-value factor. To adjust for such confounders, I propose a direct proxy for the knowledge spillover term and adjust the measured impact-value term accordingly:

$$\tilde{\zeta}_{it} = \frac{\hat{\zeta}_{it}}{1 + \tilde{\phi}\,\hat{\theta}_{it}}.$$

For the empirical proxy of knowledge spillovers, I follow the approach pioneered in Bloom et al. (2013), who proxy for knowledge spillovers received by the firm using R&D activity of firms working on similar technologies as proxied for by their technology class distribution in patent data. Here, a measure of technology spillovers is necessary, rather than of the spillovers received. I thus adjust their measure to aggregate over knowledge contributions given rather than received.

Let V_i be a firm's technology vector, i.e., the distribution of citations received across patent classes, normalized to $|V_i| = 1$ and $\omega_{ij} = V_i\, V_j$ be the inner vector product. Then, I define the empirical spillover measure as

$$\hat{\theta}_{it} = \frac{\sum_{j \in N_t \setminus i} \omega_{ij}}{\frac{1}{N_t} \sum_{k \in N_t} \sum_{j \in N_t \setminus i} \omega_{kj}}, \tag{E.5}$$

where N_t is the set of currently active firms for a given year.[21] I implement the methodology using CPC group codes, which include more than 200 classification codes. When a patent has multiple codes, I distribute its weight equally across classes. Following Bloom et al. (2013), I weight patents by citations received.

Table E.2 confirms that adjusting for spillovers does not materially alter my estimates. For alternative specifications of the impact-value factor, I experiment with different relative importance of the spillover measure and find that, if anything, larger importance increases the implied Growth Impact of Wedges.

[21] Note that Bloom et al. (2013) weight firms by their R&D stock. In my case, θ is defined as a measure per unit of R&D and, thus, I refrain from such weighting.

Table E.2: R&D Wedges, Economic Growth and Welfare — Filtering Out Spillovers

Specification	Growth Impact $\Xi - 1$				Welfare Cost of Δ	
	1975–2014	1975–90	2000–14	Δ	End.	Semi-End.
Baseline	-25.6%	-14.8%	-32.9%	-21.2%	12.2%	11.6%
A. Citations						
No spillovers	-25.6%	-14.8%	-32.9%	-21.2%	12.2%	11.6%
Some spillovers ($\omega = 1$)	-25.4%	-14.8%	-32.7%	-21.0%	12.0%	11.5%
Strong spillovers ($\omega = 100$)	-26.4%	-15.9%	-33.8%	-21.3%	12.3%	11.7%
B. Life-time Citations						
No spillovers	-26.1%	-15.6%	-33.4%	-21.2%	12.2%	11.6%
Some spillovers ($\omega = 1$)	-25.9%	-15.5%	-33.3%	-21.0%	12.1%	11.5%
Strong spillovers ($\omega = 100$)	-26.9%	-16.7%	-34.4%	-21.3%	12.3%	11.7%
C. Scientific Impact						
No spillovers	-21.0%	-13.4%	-26.1%	-14.7%	7.7%	7.4%
Some spillovers ($\omega = 1$)	-20.9%	-13.5%	-26.1%	-14.5%	7.6%	7.3%
Strong spillovers ($\omega = 100$)	-22.0%	-14.7%	-27.4%	-15.0%	7.8%	7.6%

Notes: Table reports estimates for impact of R&D wedges adjusting for knowledge spillovers together with their implications for welfare. Changes in welfare are in consumption equivalent terms. See text and Appendix for details.

F Measurement Error

This section considers adjustments for two sources of measurement error in R&D returns: Uncertainty across R&D projects within a firm and firm-level uncertainty in R&D outcomes. The former arises when firms conduct R&D projects whose ex-post value is uncertain, e.g., because some inventions turn out more valuable than others. The latter arises when there are firm-level shocks to the value of R&D outputs after investments are made, e.g., general taste shocks for the firm's products. I propose a bootstrapping procedure to address the former and a structural GMM approach to address the latter. Finally, I also consider the adjustment procedure proposed in Bils et al. (2021).

F.1 Bootstrapping

Suppose the value of individual research projects, as captured by patents, is ex-ante uncertain. Ex-post variation in valuations then might give rise to dispersion in measured R&D returns even with equalized ex-ante expectations. I propose a simple bootstrapping procedure to estimate the variability in R&D returns induced by this variation.

I establish the realized portfolio of patent valuations for each firm $\times$ 5-year interval in which the firm has at least 50 patents. For each of 1000 bootstrap samples, I then implement the following procedure:

1. For each firm and 5-year window in which the firm has at least 50 patents:
   (a) From the realized portfolio for the firm-period, draw with replacement an alternative portfolio with the same number of patents.
   (b) Calculate the return gap as the log of the ratio of valuations in the alternative portfolio to the valuation of the true portfolio.
2. Calculate the within-period standard deviation of return gaps for the simulated data.

One way to interpret this approach is that the realized patent portfolio is a good approximation for the true uncertainty faced by the firm around its innovation outcomes. The procedure ignores all variation coming from shifts in the level of expected patent valuation and instead considers the dispersion conditional on the average value only. As a result, the procedure will overstate the associated measurement error if firms are aware that certain projects are low or high expected value within their research portfolio. On the other hand, the procedure ignores all uncertainty around the number of realized patents.

Table F.1 reports the estimates. I find an average standard deviation of the return gap of around 0.06, which suggests that uncertainty around patent valuation might have contributed $(0.06/0.93)^2 \approx 0.4\%$ to the variance of R&D returns. Uncertainty across patents thus does not appear to contribute much to that dispersion. Note that this is not necessarily surprising, since averages should converge to the true mean with a sufficiently large number of independent observations by the law of large numbers.

Table F.1: Bootstrapping Estimates for Measurement Error

Measure	Period		
	1975-2014	1975-1990	2000-2014
Standard deviation	0.060	0.056	0.060
	[0.052,0.069]	[0.048,0.066]	[0.053,0.070]
Adjustment factor	0.998	0.997	0.998
	[0.997,0.998]	[0.996,0.998]	[0.997,0.998]

Notes: Table reports bootstrapping estimates for noise in R&D returns. See text for details.

F.2 GMM Approach

The bootstrapping approach can address variation across projects; however, it cannot adjust for correlated shocks to the firms' patent valuations or citations, which could arise, e.g., due to the expectation-realization gap, correlated errors in patent valuation estimation, or misreporting of R&D expenditure.[22] I propose to investigate the importance of such variation using a structural decomposition of the variation in R&D returns.

Consider a stationary, AR(1) process $\{y_{it}\}$:

$$y_{it} = (1-\rho)\mu_i + \rho y_{it-1} + \varepsilon_{it} \text{ with } \varepsilon_{it} \overset{iid}{\sim} N(0, \sigma_\varepsilon^2) \text{ and } \mu_i \sim N(0, \sigma_\mu^2). \tag{F.1}$$

The econometrician observes the process with i.i.d. normal measurement error:

$$\tilde{y}_{it} \equiv y_{it} + \nu_{it} \quad \nu_{it} \overset{iid}{\sim} N(0, \sigma_\nu^2). \tag{F.2}$$

[22] R&D expenditure is expensed in US GAAP accounting, giving firms an incentive to fully report it to reduce their tax liability. Terry et al. (2022) argue that managers might still misreport to hit short-run earnings targets or smooth earnings.

Lemma 2. *Define* $\Delta\tilde{y}_{it} \equiv \tilde{y}_{it} - \tilde{y}_{it-1}$, *then under* $\rho \in (0,1)$, *we have*

$$\begin{aligned} m_1 &\equiv Cov(\tilde{y}_{i,t}, \Delta\tilde{y}_{it}) = \frac{1}{1+\rho}\sigma_\varepsilon^2 + \sigma_\nu^2 \\ m_2 &\equiv Cov(\tilde{y}_{i,t}, \Delta\tilde{y}_{it-1}) = \frac{\rho}{1+\rho}\sigma_\varepsilon^2 \\ m_3 &\equiv Cov(\tilde{y}_{i,t}, \Delta\tilde{y}_{it-2}) = \frac{\rho^2}{1+\rho}\sigma_\varepsilon^2 \\ m_4 &\equiv Cov(\tilde{y}_{i,t}, \tilde{y}_{it-1}) = \sigma_\mu^2 + \frac{\rho}{1-\rho^2}\sigma_\varepsilon^2. \end{aligned}$$

Proof. The results follow immediately from the assumptions. □

Proposition 8. *If* $\rho \in (0,1)$, *we can solve for* $\{\rho, \sigma_\mu, \sigma_\varepsilon, \sigma_\nu\}$ *using the population auto-covariance structure of* $\tilde{y}_{it}$ *and* $\Delta\tilde{y}_{it} \equiv y_{it} - y_{it-1}$*:*

$$\beta \equiv \begin{bmatrix} \rho \\ \sigma_\varepsilon^2 \\ \sigma_\mu^2 \\ \sigma_\nu^2 \end{bmatrix} = \begin{bmatrix} \frac{m_3}{m_2} \\ \frac{(m_2)^2}{m_3} + m_2 \\ m_4 - \frac{(m_2)^2}{m_2 - m_3} \\ m_1 - \frac{(m_2)^2}{m_3} \end{bmatrix}$$

Let Ω *be the covariance matrix of* m *and* $\hat{m}$ *denote the sample moments. Then,*

$$\hat{\beta} \sim N(\beta, \Sigma) \quad \textit{and a feasible estimator is} \quad \hat{\Sigma} = \left(\frac{\partial\hat{\beta}}{\partial m}\right)' \hat{\Omega} \left(\frac{\partial\hat{\beta}}{\partial m}\right),$$

where $\partial\beta/\partial m$ *is evaluated at* $\hat{m}$ *and given by*

$$\frac{\partial\beta}{\partial m} = \begin{bmatrix} 0 & 0 & 0 & 1 \\ -\frac{m_3}{(m_2)^2} & 2\frac{m_2}{m_3} + 1 & m_2\left(\frac{m_2 - 2m_3}{(m_2 - m_3)^2}\right) & -2\frac{m_2}{m_3} \\ \frac{1}{m_2} & -\left(\frac{m_2}{m_3}\right)^2 & -\left(\frac{m_2}{m_2 - m_3}\right)^2 & -\left(\frac{m_2}{m_3}\right)^2 \\ 0 & 0 & 1 & 0 \end{bmatrix}.$$

Proof. The first part follows by rearranging the moments expressions. The second part follows from the Law of Large Numbers for the moment vector and the Delta method. □

I report estimates for two measures of R&D returns in Table F.2. I find no contribution of transitory shocks to the overall variation for my baseline measure of R&D returns in column 1; however, the estimates are imprecise. Using sales growth to measure R&D output yields a contribution of 22%.

Table F.2: GMM Parameter Estimates for AR(1) with Noise

Parameter	Valuations/ R&D	ΔSales/ R&D
ρ	0.580***	0.764***
	(0.080)	(0.148)
σ^2_ϵ	0.691***	0.423***
	(0.095)	(0.096)
σ^2_μ	0.239**	−0.002
	(0.106)	(0.575)
σ^2_ν	−0.003	0.504***
	(0.103)	(0.097)
Observations	7,996	7,627
Adjustment factor	1.001	0.817

Notes: Table reports parameters estimates for AR(1) with Noise in logs using a General Methods of Moments approach. See text for details.

F.3 Bils et al. (2021) Adjustment

The previous two adjustments primarily deal with log-additive measurement error. Bils et al. (2021) instead propose a methodology for additive measurement error. Applying their approach to this context suggests the following procedure to account for additive measurement error:

1. Define deciles k of the R&D returns distribution and estimate:

$$\Delta \ln \text{R\&D Output}_{it} = \alpha_{k(i)} + \sum_{k=1,10} \beta_k \, \Delta \ln \text{R\&D Expenditure}_{it} \, \mathbb{I}\{i \in k\} + \epsilon_{it}, \quad \text{(F.3)}$$

   where $\alpha_{k(i)}$ is a decile fixed effect and $\mathbb{I}\{i \in k\}$ an indicator for whether firm i belongs to decile k.

2. Create adjusted R&D returns as

$$\ln \widehat{\text{R\&D Return}}_{it} = \ln \text{R\&D Return}_{it} + \ln \hat{\beta}_k(i) + \sigma_\beta \, \epsilon_{it}, \quad \text{(F.4)}$$

   where $\epsilon_{it} \sim N(0,1)$ and

$$\sigma^2_\beta = -\sigma\left(\ln \text{R\&D Return}_{it}, \ln \hat{\beta}_{k(i)}\right) - \sigma^2\left(\ln \beta_{k(i)}\right). \quad \text{(F.5)}$$

Implementation for full sample. I implement this approach for the full sample and report estimates in Table F.3. As in Bils et al. (2021), I find that the estimates decline with the decile while being broadly centered around 1. Following step 2, I find that the standard deviation of adjusted R&D returns is 0.80 compared to the unadjusted value of 0.93.

Table F.3: Estimates from Bils et al. (2021) Specification

Decile	1	2	3	4	5	6	7	8	9	10
Estimate	1.367	1.174	1.214	1.026	1.001	0.916	0.934	0.802	0.728	0.532
Std. Err.	(0.12)	(0.16)	(0.09)	(0.09)	(0.10)	(0.10)	(0.11)	(0.10)	(0.08)	(0.07)

Note: Coefficients estimates for full sample for specification (F.3). Regression controls for NAICS3× Year fixed effects and standard errors are clustered at the NAICS6 level.

Implementation for annual estimates. I implement their methodology for the measurement of alternative R&D wedges to construct alternative estimates for R&D Allocative efficiency. I estimate coefficients over time using 15-year windows centered on the main window when possible to ensure a reasonable sample size. For early and late observations I use the first and last available 15-year window, respectively.

G Mechanisms Driving R&D Wedges

In this section, I highlight mechanisms captured in R&D returns and impact-value factors. I rely on a two-period growth model for simplicity.

G.1 Baseline Model

Setup. The final good producer creates consumption good Y_t by combining inputs y_{jt} from a unit mass of product lines according to:

$$\ln Y_t = \int_0^1 \ln y_{jt} \, dj.$$

Each input is supplied by a single monopolist with constant marginal costs ψ/A_{jt}. The monopolist is free to choose any price p_{jt}; however, there is a competitive fringe of firms with constant unit costs $\lambda_{jt}\ (\psi/A_{jt})$ that limit the monopolists' price setting power. Consequently, the monopolist sets limit price equal to the marginal costs of the competitive fringe and earns profits

$$\pi_{jt} = Y_t\ (1 - \lambda_{jt}^{-1}).$$

There is a unit mass of innovative firms at time 0, which may hire inventors ℓ_i at wage W to produce an invention at time 1 with probability z_i:

$$z_i = \varphi_i\ \ell_i^{\gamma}.$$

An invention improves technology in a random product line by λ_i such that $A_{j1} = \lambda_i\ A_{j0}$ in a product line with a successful invention. The competitive fringe then absorbs the knowledge of the previous monopolist, such that its unit cost gap to the monopolist is λ_i as well. As a result, the innovation yields profits π_i in period 1, which firms discount at rate R. The value of innovation to the firm is thus given by $V_i = \pi_i/R$ and its optimization problem

$$\max_{\ell_i} \{V_i\ z_i - W\ \ell_i\}$$

There is a fixed number of research workers, whose labor market clearing condition determines the R&D wage in equilibrium:

$$L = \int_0^1 \ell_i \, di.$$

Finally, I define the productivity index A_t such that $\ln A_t = \int_0^1 \ln A_{jt} \, dj$. Consequently, its growth rate is given by

$$g = \ln(A_1/A_0) \approx \int_0^1 (\lambda_i - 1)\; z_i \; di,$$

where the approximation relies on $\ln \lambda_i \approx \lambda_i - 1$.

The planner maximizes economic growth subject to the same technological constraints as firms:

$$g^* = \max \int_0^1 z_i \; (\lambda_i - 1) \; di \quad \text{s.t.} \quad L = \int_0^1 \ell_i \; di.$$

R&D returns and impact-value factors. It is straightforward to show that in this setup R&D returns are equalized across firms:

$$\frac{V_i \; z_i}{W \; \ell_i} = \frac{1}{\gamma} \qquad \text{and} \quad \ell_i = \left(\frac{V_i \; \varphi_i}{(W/\gamma)} \right)^{\frac{1}{1-\gamma}}.$$

Furthermore, one can show that this allocation is also the solution to

$$g = \max \int_0^1 z_i \; V_i \; di \quad \text{s.t.} \quad L = \int_0^1 \ell_i \; di.$$

Defining $\zeta_i \equiv (\lambda_i - 1)/V_i$, we can thus rearrange the planner problem as

$$g^* = \max \int_0^1 z_i \; V_i \; \zeta_i \; di \quad \text{s.t.} \quad L = \int_0^1 \ell_i \; di. \tag{G.1}$$

From the formulation of V_i it then follows immediately that planner and private allocation coincide iff ζ_i is a constant across firms.

G.2 Mechanisms for R&D Return Dispersion

R&D Subsidies or Taxes. Suppose firms face R&D subsidies τ_i on their gross R&D expenditure. The firm problem is then given by

$$\max_{\ell_i} \left\{ V_i \; z_i - (1 - \tau_i) \; W \; \ell_i \right\}.$$

Consequently, firms' R&D returns directly reflect differences in subsidy rates:

$$\frac{V_i \; z_i}{W \; \ell_i} = \frac{1}{\gamma} \; (1 - \tau_i) \qquad \text{and} \quad \ell_i = \left(\frac{V_i \; \varphi_i}{(W/\gamma) \; (1 - \tau_i)} \right)^{\frac{1}{1-\gamma}}.$$

Capacity constraints. Suppose firms face exogenous capacity constraint $\ell_i \leq \bar{\ell}_i$. The firm problem is then given by

$$\max_{\ell_i} \left\{ V_i \; z_i - W \; \ell_i \quad \text{s.t.} \quad \ell_i \leq \bar{\ell}_i \right\}.$$

Consequently, firms' R&D returns directly reflect the tightness of the capacity constraint $\tilde{\lambda}_i$:

$$\frac{V_i \; z_i}{W \; \ell_i} = \frac{1}{\gamma} \; (1 + \tilde{\lambda}_i) \qquad \text{and} \quad \ell_i = \left(\frac{V_i \; \varphi_i}{(W/\gamma) \; (1 + \tilde{\lambda}_i)} \right)^{\frac{1}{1-\gamma}}.$$

Discount Rates. Suppose firms have heterogeneous discount rates R_i reflecting, e.g., risk or financial constraints, which are not observed in the data. Let $V_i = \pi_i / R$ with $R = \mathbb{E}[R_i]$, then the firm problem is given by

$$\max_{\ell_i} \left\{ (R/R_i) \; V_i \; z_i - W \; \ell_i \right\}.$$

Consequently, firms' measured R&D returns directly reflect these differences:

$$\frac{V_i \; z_i}{W \; \ell_i} = \frac{1}{\gamma} \; \frac{R_i}{R} \qquad \text{and} \quad \ell_i = \left(\frac{V_i \; (R/R_i) \; \varphi_i}{(W/\gamma)} \right)^{\frac{1}{1-\gamma}}.$$

Adjustment costs. Suppose firms face exogenous adjustment costs $\phi \; W \; (\ell_i - \bar{\ell}_i)^2$. The firm problem is then given by

$$\max_{\ell_i} \left\{ V_i \; z_i - W \; \ell_i - \phi \, W \; (\ell_i - \bar{\ell}_i)^2 \right\}.$$

Consequently, firms' R&D returns directly reflect the adjustment costs:

$$\frac{V_i \; z_i}{W \; \ell_i} = \frac{1}{\gamma} \; (1 + 2 \; \phi \; (\ell_i - \bar{\ell}_i)) \quad \text{and} \quad \frac{V_i \; z_i}{W \; \ell_i + \phi \, W \; (\ell_i - \bar{\ell}_i)^2} = \frac{1}{\gamma} \; \frac{1 + 2 \; \phi \; (\ell_i - \bar{\ell}_i)}{1 + \phi \frac{(\ell_i - \bar{\ell}_i)^2}{\ell_i}}$$

Firms with high R&D relative to their reference point have higher returns.

Monopsony Power. Suppose R&D labor is specialized across fields. R&D labor is perfectly mobile across firms within a field, but not across fields, such that the labor market clearing condition is given by

$$L = \int_0^1 \ell_i \; \left(\frac{\frac{1}{N_i} \sum_{i \in \mathcal{N}_i} \ell_j}{L} \right)^{\xi} \; di, \tag{G.2}$$

where N_i is the number of firms in a given field.

As a result, wages may differ across fields and are generally increasing in the average demand for R&D input within a given field:

$$W_i = W \left(\frac{\frac{1}{N_i} \sum_{j \in \mathcal{N}_i} \ell_j}{L} \right)^{\xi} \tag{G.3}$$

Firms internalize the impact of labor demand on wages and, consequently, their first-order conditions under symmetry ($\ell_j = \ell_i$ for $j \in \mathcal{N}_i$) are given by

$$\gamma \, \theta \, \ell_i^{\gamma-1} = \left(1 + \frac{1}{N_i} \, \xi\right) \, W_i \tag{G.4}$$

R&D return is given by $(1/\gamma) \, \left(1 + \frac{1}{N_i} \, \xi\right)$ with $\Delta_i = \frac{1}{N_i} \, \xi$. Variation in R&D returns is thus directly linked to the degree of competition in the firm-specific labor market. Firms with more competition for R&D workers have lower R&D returns and vice versa.

G.3 Mechanisms for Dispersion in Impact-Value Factors

Patent Protection. Suppose that the competitive fringe learns with probability $1 - P_i$ about the new technology of a monopolist such that the monopolist is only able to profit from the innovation with probability P_i. In this case, the private value of the invention is $V_i = P_i \; \pi_i/R$, while the public value remains $\lambda_i - 1$. As a result, variation in P_i induces variation in ζ_i.

Exogenous Markup Differences. Suppose that firms differ in their unit cost parameter ψ_i due to, e.g., technological differences or complementarities across product lines. The profit of an invention is then given by $\pi_i = Y_1 \; (1 - (\psi/\psi_i) \; \lambda^{-1})$. As a result, variation in ψ_i across firms yields variation in the private value a firm creates from innovation without changing the growth impact $\lambda_i - 1$, which induces variation in the impact-value factor ζ_i.

Endogenous Markup Differences. Suppose that firms differ in their step-size λ_i, then $\zeta_i \propto \lambda_i$ such that variation in step-sizes yields variation in the impact-value factor. Intuitively, the growth gains of λ_i are linear, while the profit gains are concave, such that firms with high-quality innovation under-invest in R&D.

Frictions in the Product Market. It is straightforward to see that any frictions in the product market that affect π_i without changing the growth impact of an invention naturally yield variation in ζ_i as well. Firms with artificially low profits under-provide innovation.

H An Endogenous Growth Model with Misallocation

This section confirms that the formulas in the main text arise in a fully specified endogenous growth model. I present a fully specified expanding variety growth model à la Romer (1990) and discuss alternative assumptions to derive equivalent results in a Schumpeterian setup.

H.1 Expanding Variety Model

Time is discrete and infinite, and indexed by t. There is a representative household with CRRA utility in consumption C_t that supplies a mass L_t of R&D workers and a mass L_P of production workers inelastically to the innovation and production sectors. The household maximizes discounted utility and faces a standard budget constraint:

$$\max \sum_{s=0}^{\infty} \beta^s \, \frac{C_{t+s}^{1-\sigma} - 1}{1-\sigma}$$
$$B_{t+1} + C_t = R_t \, B_t + W_{Pt} \, L_P + W_t \, L_t + \Pi_t, \tag{H.1}$$

where W_{Pt} and W_t are wages, B_t is the riskless bond with interest rate R_t that is in zero net-supply, and Π_t are profits. First-order conditions yield the standard Euler equation:

$$\left(\frac{C_{t+1}}{C_t}\right)^{\sigma} = \beta \, R_{t+1}. \tag{H.2}$$

The competitive **final goods sector** combines production labor L_P and intermediate goods x_{it} from a mass A_t different varieties with quality z_{it} in the final good according to

$$Y_t = L_P^{1-\alpha} \int_{A_t} z_i^{1-\alpha} \, x_{it}^{\alpha} \, di. \tag{H.3}$$

Production labor is paid wage W_{Pt}, while intermediate goods are priced at P_{it}. The final goods sector treats these prices as given and maximizes profits. First-order conditions imply

$$W_{Pt} = (1-\alpha) \int_{A_t} z_i^{1-\alpha} \left(\frac{x_{it}}{L_P}\right)^{\alpha} di \quad \text{and} \quad P_{it} = \alpha \left(\frac{L_P \, z_i}{x_{it}}\right)^{1-\alpha}. \tag{H.4}$$

Each **intermediate good** variety is controlled by a monopolist that transforms ψ units of the final goods into one unit of the intermediate good. The monopolist maximizes profits taking into account the demand curve from the final goods sector. The first-order conditions imply an equilibrium price $P_{it} = \frac{\psi}{\alpha}$ such that quantities and profits are given by

$$x_{it} = z_{it}\, L_P \left(\frac{\psi}{\alpha^2}\right)^{-\frac{1}{1-\alpha}} \quad \text{and} \quad \pi_i = (1-\alpha)\,\alpha \left(\frac{\psi}{\alpha^2}\right)^{-\frac{\alpha}{1-\alpha}} L_P\, z_i. \tag{H.5}$$

Thus, **aggregate production and consumption**—production net of costs—are

$$Y_t = L_P\, Z_t \left(\frac{\psi}{\alpha^2}\right)^{-\frac{\alpha}{1-\alpha}} \text{ and } C_t = (1-\alpha^2)\, L_P\, Z_t \left(\frac{\psi}{\alpha^2}\right)^{-\frac{\alpha}{1-\alpha}} \text{ with } Z_t = \int_{A_t} z_{it}\, di \tag{H.6}$$

The growth rate of productivity in this economy is thus given by the growth rate of Z_t. Let a_t be the set of novel varieties in each period and $z_t = \int_{a_t} z_{it}\, di$, then

$$g_t = \frac{A_t - A_{t-1}}{A_{t-1}} = \frac{z_t}{Z_{t-1}}. \tag{H.7}$$

The net-present value of profits for each variety Π_{it} is given by

$$\Pi_{it} = z_i\, \pi\, \mathcal{R}_t \quad \text{where } \mathcal{R}_t = \sum_{s=0}^{\infty} R^{-1}_{t,t+s}, \tag{H.8}$$

where $R^{-1}_{t,t+s}$ is the discount rate between t and $t+s$, s.t., $R_{t,t} = 1$ and $R_{t,t+s} = \prod_{k=1}^{s} R_{t+k}$, where R_t is the discount rate between $t-1$ and t. The parameter π is π_i evaluated at $z = 1$.

The **innovation sector** is populated by a unit mass of R&D labs. Labs have heterogeneous R&D productivity and can hire researchers ℓ_{kt} to produce a mass of new inventions M_{kt+1} in the subsequent period according to

$$M_{kt+1} = Z_t\, \varphi_{kt}\, \ell_{kt}^{\gamma}. \tag{H.9}$$

The labs know the quality z_{kt+1} of inventions in advance. Upon receiving them, they sell them to the intermediate goods sector for a fraction $\tilde{\zeta}^{-1}_{kt+1}$ of the net present value of the resulting profits such that they receive discounted value $V_{kt+1} = R^{-1}_{t+1}\, \tilde{\zeta}^{-1}_{kt+1}\, \Pi_{kt+1}$.

Lastly, each lab faces an R&D wedge $1 + \Delta_{kt}$ such that its optimization problem is

$$\max M_{kt+1}\, V_{kt+1} - (1 + \Delta_{kt})\, W_t\, \ell_{kt}. \tag{H.10}$$

The resulting first-order conditions are thus given by

$$\gamma\, Z_t\, \varphi_{kt}\, \ell_{kt}^{\gamma-1} = (1 + \Delta_{kt})\, W_t. \tag{H.11}$$

Labor market clearing requires

$$L_t = \int_0^1 \ell_{kt} \, dk \tag{H.12}$$

Consider the growth rate of the economy. It follows that

$$g_t = \int_0^1 \frac{M_{kt}}{Z_{t+1}} \, z_{kt} \, dk = \int_0^1 M_{kt} \, V_{kt} \, \zeta_{kt} \, dk \quad \text{where} \quad \zeta_{kt} = \frac{\tilde{\zeta}_{kt}}{\mathcal{R}_t \, \pi \, Z_t}, \tag{H.13}$$

which is the formula assumed in the main text. The impact-value factor ζ_{kt} captures the degree to which R&D labs fail to appropriate value created in the production sector. Differences therein could be due to firms' ability to protect their intellectual property or fringe competition limiting markups. Differences in the discount factor $\mathcal{R}_t$ would also yield variation in ζ_{kt}. Finally, note that the model allows for heterogeneity in z_{kt} and φ_{kt}. From the perspective of the firm and the growth-maximizing allocation, only their product matters.

H.2 A Schumpeterian Foundation

The same insights extend to a Schumpeterian framework with limit pricing à la Aghion and Howitt (1992). In the Schumpeterian model, the **final production function** is given by a Cobb-Douglas aggregator over a unit mass of intermediate goods:

$$Y_t = \exp\left(\int_0^1 \omega_{it} \, \ln Y_{it} \, di\right) \quad \text{s.t.} \quad \int_0^1 \omega_{it} \, di = 1 \tag{H.14}$$

Consequently, demand is given by

$$\omega_{it} \, Y_t = P_{it} \, Y_{it} \tag{H.15}$$

In each product line, there is an incumbent hiring production workers L_{Pit} to produce the intermediate good with a linear production function and productivity A_{it}:

$$Y_{it} = A_{it} \, L_{it} \tag{H.16}$$

The **incumbent** faces a competitive fringe with productivity A_{it}/λ_{it} and pays labor W_{Pt}. Per limit pricing, the incumbent prices at the fringe's marginal costs such that

$$\pi_{it} = \left(\frac{\lambda_{it} - 1}{\lambda_{it}}\right) \omega_{it} \, Y_t \tag{H.17}$$

In turn, the aggregate production function simplifies to

$$Y_t = \underbrace{\exp\left(\int_0^1 \omega_{it} \;\ln A_{it}\; di\right)}_{\equiv A_t} \; L_t \; \underbrace{\frac{\exp(\int_0^1 \omega_{it} \ln(\omega_{it}/\lambda_{it})\; di)}{\int_0^1 \frac{\omega_{it}}{\lambda_{it}}\; di}}_{\equiv \Lambda_t} \tag{H.18}$$

There is a unit mass of R&D labs in the **innovation sector**. In each period, the lab can hire researchers ℓ_{kt} to innovate with probability z_{kt+1} in the subsequent period:

$$z_{kt+1} = \varphi_{kt}\; \ell_{kt}^{\gamma}. \tag{H.19}$$

Innovations improve upon the productivity of an incumbent by factor λ_{kt+1}, which the firm knows upon making the innovation decision, and replaces the incumbent upon successful innovation. I denote the present discounted value of incumbency by V_{kt+1}:

$$V_{kt} = \pi_{kt}\; \mathcal{R}_{kt} \text{ with } \mathcal{R}_{kt} = \left(1 + \sum_{s=1}^{\infty} \prod_{m=1}^{s} \left(\frac{1 - \bar{z}_{kt+m}}{R_{kt+m}}\right)\right), \tag{H.20}$$

where $\bar{z}_{kt+m}$ is the probability of being replaced by another inventor. This probability could be potentially heterogeneous across product lines; however, it has to hold that

$$\int_0^1 z_{kt+1}\; dk = \int_0^1 \bar{z}_{kt+1}\; dk \tag{H.21}$$

Innovative firms then receive the expected value of invention and face R&D wage W_t. Furthermore, they are subject to wedge Δ_{kt} as before, such that their optimization problem is given by

$$\max\{z_{kt+1}\; V_{kt+1} - (1 + \Delta_{kt})\; W_t\; \ell_{kt}\} \tag{H.22}$$

Finally, approximating TFP growth through log-difference and noting that $\ln \lambda_{kt} \approx (\lambda_{kt} - 1)$, we have

$$g_t = \ln\left(\frac{A_t}{A_{t-1}}\right) = \int_0^1 z_{kt}\; \omega_{kt}\; (\lambda_{kt} - 1)\; dk = \int_0^1 z_{kt}\; V_{kt}\; \zeta_{kt}\; dk \quad \text{with} \quad \zeta_{kt} = \frac{\lambda_{kt}}{\mathcal{R}_{kt}\; Y_t} \tag{H.23}$$

The equation clarifies that in a limit pricing setup, expected differences in quality naturally appear in the impact-value factor. Furthermore, as in the previous setup, differences in the discount factor across firms would also appear. Otherwise, the model gives rise to the same formulae as developed in the main text.