

Did R&D Misallocation Contribute to Slower Growth?*

Nils H. Lehr[†]

IMF

June 20, 2024

Abstract

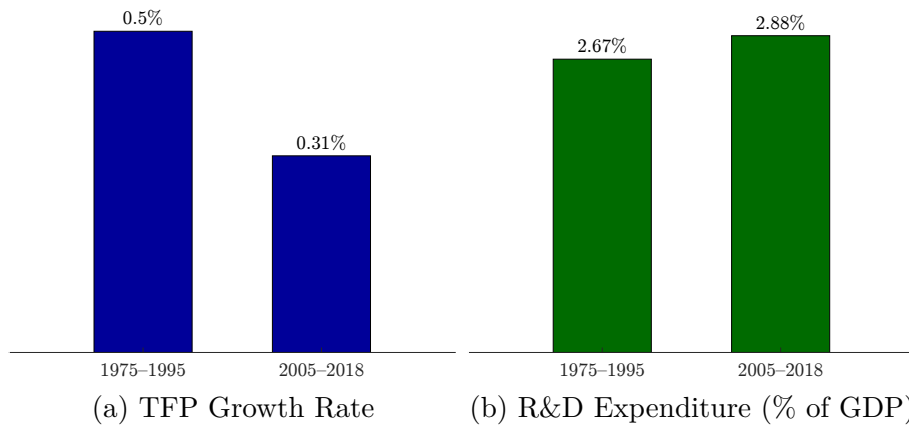
This paper provides evidence that rising misallocation in the R&D sector contributed to the recent slowdown in U.S. productivity growth. I develop a growth accounting framework allowing for misallocation of R&D resources across firms captured by wedges between their marginal cost and benefits of R&D. I show that R&D wedges can be measured from R&D returns and document large and persistent differences in R&D returns across US-listed firms. Combining data and model, I estimate that frictions reduced productivity growth by 18% over 1975–2014 and that rising misallocation in the R&D sector accounts for 25% of the growth slowdown.

*I especially thank Stephen J. Terry, Pascual Restrepo, Tarek A. Hassan, and David Lagakos for their continued guidance, feedback, and encouragement. I benefited from the advice and comments of many including 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: nilslehr@bu.edu

U.S. productivity growth has slowed down significantly in the last two decades. While productivity 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—has 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



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 productivity growth rooted in a simple decomposition. In endogenous growth models, aggregate 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 a function of 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” (Bloom et al., 2020). Instead, this paper focuses on the second channel

and provides evidence that declining allocative efficiency contributed to lower aggregate R&D productivity and, thus, economic growth. 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 25% of the 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 workhorse growth models (Romer, 1990; Aghion and Howitt, 1992). Firms hire R&D workers to maximize the private value created from innovation. I introduce frictions flexibly through exogenous *R&D wedges* in firms' first-order conditions. These distort firms' demand for R&D inputs such that marginal returns on R&D are controlled by R&D wedges rather than equalized across firms. 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 low impact-value factor create a lot of private value, while contributing little to productivity growth.

I show that the impact of private frictions, i.e., *R&D wedges*, on the equilibrium economic growth rate is captured by a sufficient statistic, which I refer to as *R&D Allocative Efficiency*. When private and growth incentives are aligned, the baseline case in the literature, variation in R&D wedges reduces allocative efficiency—an R&D sector equivalent of Hsieh and Klenow (2009). 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 multi-research lines per firms, however, the counterfactual holds constant their distribution. Second, frictions are costlier under free entry as they deter entry by reducing potential gains for entrants. Third, frictions are less costly when R&D inputs are imperfect substitutes across firms as there are fewer gains from input reallocation. Finally, the level of frictions has a direct effect on growth in the case of a positive aggregate R&D supply elasticity, while it does not in the case of fixed supply, in which only relative frictions matter.

Next, I estimate the model primitives from firms' patents and financial statements for a sample of US-listed firms for 1975–2014. I measure firms' R&D investment from their expenditure and the resulting private value created using patent valuations. In the model, the ratio of value created and 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 measures or citations. These data allow me to estimate *R&D Allocative Efficiency*, i.e., the impact of frictions on growth, in my sample.

Before estimating the aggregate impact of R&D wedges, I investigate them at the micro-level. In absence of frictions, firms equalize the marginal benefit to the marginal cost of R&D and, thereby, the R&D return as well. In contrast, I find large and highly persistent differences in measured R&D returns and, by extension, R&D wedges. This finding is reminiscent of the literature on misallocation in the production sector, which argues that dispersion in the return on capital is a strong indicator for capital misallocation (David et al., 2016). I find that the standard deviation of R&D returns is 42% larger than its counterpart for the return on capital, suggesting significant R&D misallocation. Notably, R&D return dispersion is primarily among highly comparable firms. 78% of the variation remains when focusing on differences among firms within 6-digit industry×year cells only. Finally, the strong persistence of R&D returns—with an implied annual autocorrelation coefficient around 0.9—suggest structural factors rather statistical noise.

I perform numerous robustness exercises and find large measured R&D return dispersion throughout. First, I consider sales or employment growth as alternative measures of the private value of R&D output and find larger dispersion in the resulting R&D returns. Second, I estimate the prevalence of measurement error in complementary bootstrap and structural approaches, and find no evidence that it contributes significantly to R&D return dispersion. Lastly, I investigate additional mechanisms directly, including acquisitions, fixed costs, knowledge capital, and alternative assumptions around the patent valuations, and find no evidence that they significantly contribute to R&D return dispersion. The combined evidence thus suggests that dispersion in measured R&D wedges is not driven by measurement details, which leaves economic drivers as a candidate source.

R&D wedges measure frictions in the model, however, I find that they are surprisingly hard to predict with empirical proxies thereof. For example, they are uncorrelated with the return on capital, which suggests that these investments are subject to different frictions. Similarly, R&D returns are not consistently correlated with measures of financial frictions. The strongest predictors for R&D returns are R&D employment and measures of firm expansion such as rising R&D investment or TFP growth. The former is in line with a framework in which monopsony power over inventors increases with their employment. The latter suggests a potential role for adjustment frictions that lead to temporarily larger R&D returns during expansions due to a gradual adjustment process of R&D investment and vice versa. Lastly, I find an ambiguous relationship between R&D wedges and impact-value factors with different signed correlations across alternative measures of the latter.

At the aggregate level, I estimate that productivity growth is significantly slower due to low *R&D Allocative Efficiency* and increasingly so. For the full sample, I estimate an 18% lower growth rate due to dispersion in R&D wedges. For comparison, [Hsieh and Klenow \(2009\)](#) estimate that U.S. productivity would improve by 40% under the first-best factor allocation, while [Berger](#)

et al. (2022) estimate a 21% larger output in absence of monopsony in the production sector. Naturally, achieving the frictionless growth rate might not be feasible in practice if R&D wedges are the product of technological or information frictions that cannot, or should not, be adjusted for.

Comparing 1975–90 to 2000–14, I find that declining R&D Allocative Efficiency can account for 11% slower growth, which represents 25% of the $\frac{0.5\% - 0.3\%}{0.5\%} \approx 40\%$ overall slowdown. Thus, my estimates suggest that slower productivity growth, and lower R&D productivity, is partly due to 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. Similar to Akcigit and Ates (2021) and Olmstead-Rumsey (2022), I argue for declining aggregate R&D productivity as 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 similar to de Ridder (2023), Aghion et al. (2023) and Ayerst (2022), 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 firms' ability

¹Bloom 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).

to benefit from inventions (de Ridder, 2023; Mezzanotti, 2021; Aghion et al., 2023; Akcigit et al., 2022; Manera, 2022; Aghion et al., 2022). 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 and, thus, estimating the growth impact of either factor 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 factor return heterogeneity—as a significant force shaping aggregate productivity. The subsequent literature documented 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 surprising in and of itself (Banerjee and Moll, 2010; Asker et al., 2014; Midrigan and Xu, 2014; David et al., 2016). I complement the literature by focusing on R&D investments instead of static production factors such that factor return heterogeneity is linked to the productivity growth rate instead of its level. Furthermore, I show that the sources of frictions appear to be different for capital and R&D investment as their returns are uncorrelated. I also find limited evidence on the contribution of channels discussed prominently in the literature such as government subsidies or frictions. In contrast, König et al. (2022) find that product market frictions led to an inefficient R&D allocation in China. Similarly, financial frictions are often considered to be severe for intangible investments such as R&D (Brown et al., 2009). My focus on larger firms in the U.S. is a likely explanation for the contrasting findings. The strongest predictor of R&D returns is inventor employment, which could be explained by increasing monopsony power exerted by large employers (Berger et al., 2022). Such a channel could also explain rising frictions in light of rising overall concentration (Autor et al., 2020).

1 Theory

This section introduces the theoretical framework for assessing the impact of friction in the R&D sector on economic growth. The framework nests alternative growth theory traditions (Romer, 1990; Aghion and Howitt, 1992).

1.1 Model Setup

Time is infinite, discrete, and indexed by t .

Production. Output Y_t is a function of aggregate productivity A_t and production labor $L_{P,t}$, which is supplied inelastically:

$$Y_t = A_t \cdot L_{P,t}, \quad (1)$$

Firms. A unit mass of innovative firms $i \in [0, 1]$ with R&D productivity φ_{it} hires R&D input ℓ_{it} at price W_t to achieve mass z_{it} of innovations:²

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

Firms assign value V_{it} to innovations, which I take as given. In workhorse growth models, this value is linked to resulting profits and innovation opportunities.³ Firms are subject to R&D wedge Δ_{it} such that their equilibrium R&D input choice ℓ_{it}^* satisfies

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

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. If $\Delta_{it} = 0$, we recover the frictionless benchmark in which firms equalize marginal benefit and cost. Otherwise, firms' choices are distorted relative to the benchmark with larger wedges resulting in lower demand for R&D inputs.

In theory, there are a range of distortions captured by R&D wedges in-

² z_{it} can also be interpreted as the arrival rate of inventions in a continuous time setup.

³I provide examples of V_{it} for alternative microfoundations in Appendix E.

cluding financial frictions, adjustment costs or capacity constraints, market power in the R&D input market, and R&D subsidies. For example, high R&D wedges arise when firms face constraints on their choice of R&D inputs due to financial frictions or adjustment costs. Similarly, low R&D wedges can capture R&D subsidies, which reduce firms’ marginal cost below the market price. I discuss mechanisms further in Appendix E and take R&D wedges as given.

Factor markets. Aggregate R&D input L_t is supplied inelastically, capturing the idea that research talent is scarce (Goolsbee, 2003; Wilson, 2009):

$$L_t = \int_0^1 \ell_{it} \cdot di. \quad (4)$$

An alternative interpretation is that R&D policy already fixes the size of the R&D sector at its optimal level, such that the allocation of resources within it remains the relevant margin of concern.⁴

Growth. Productivity grows through innovation and its growth rate is the aggregate of the mass of inventions times their growth impact. The latter is linked to an inventions’ private value via the impact-value factor ζ_{it} . Firms with a large impact-value factor contribute more productivity growth per dollar of private value created. The growth rate is given by

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

where $\phi \geq 0$ is “phishing-out” effect that is necessary to achieve balanced growth in a semi-endogenous growth framework (Jones, 1995).

The impact-value factor plays a prominent role in the growth literature as it determines the degree to which firms’ incentives are aligned with a growth-maximizing planner. The early endogenous growth literature emphasizes that firms might not be able to appropriate the full value generated from their innovation to society such that social value exceeds private value (Romer,

⁴This assumption abstracts from any waste of R&D inputs linked to R&D wedges, e.g., due to “real” adjustment cost. Any adjustment costs captured here are assumed to take the form of either production labor or cash payments.

1990). On the other hand, the Neo-Schumpeterian literature argues that the business stealing effect acts as a counterbalancing force as innovators do not internalize the economic damage imposed on firms made obsolete by innovation (Aghion and Howitt, 1992). In recent contributions, heterogeneity in impact-value factors arises due to differences in firms' ability to earn larger profits from inventions, protect their intellectual property, and withstand challengers (Akcigit and Ates, 2021; Mezzanotti, 2021; König et al., 2022; Manera, 2022; Olmstead-Rumsey, 2022; Aghion et al., 2023; de Ridder, 2023). I discuss these mechanisms in Online Appendix E and take the factors as given here.

Equilibrium. I use two simplified equilibrium definitions in deriving the main results. The Competitive Equilibrium respects the equation detailed above, while the Planner Equilibrium 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 Planner Equilibrium 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 two terms jointly characterize the growth rate in a frictionless, competitive equilibrium without R&D wedges. The third term, which I refer to as *R&D Allocative Efficiency*, captures the impact of R&D wedges. *R&D Allocative Efficiency* depends on the distribution of R&D wedges, but not their average level. Intuitively, any aggregate excess or insufficient demand for R&D resources is balanced by the R&D input price due to the fixed aggregate supply thereof and, thus, does not have a direct impact on economic growth.

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

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

where $\tilde{\zeta}_{it} \equiv \zeta_{it} / \left(\int_0^1 \omega_{it} \cdot \zeta_{it} di \right)$ 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, respectively, with R&D productivity defined as $\theta_{it} \equiv \varphi_{it} \cdot V_{it}$.

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 create more or less growth impact per private value. Then, dispersion in R&D wedges strictly reduces economic growth as shown in Corollary 1. Intuitively, R&D wedges control the marginal return on R&D. If there are firms with different marginal returns, one could raise the aggregate return by moving resources from low to high marginal R&D return firms. Larger dispersion then suggests ever more unrealized opportunities for gainful reallocation and, thus, more misallocation—a growth equivalent to [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:*

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

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

The analysis is more nuanced once we relax the assumption of orthogonal impact-value factors as established in Proposition 2. Under some conditions on

their correlation with R&D wedges, we can understand them as either amplifying or dampening their impact. Under positive correlation, i.e., if particularly constrained firms also achieve a high growth impact per private value created, then impact-value factors amplify the misallocation created from R&D wedges. Conversely, they dampen their 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 $\eta = \sqrt{1 + 2 \cdot \sigma_{\Delta,\zeta} / \sigma_{\Delta}^2}$. Then, up to a 2nd-order approximation,*

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

A couple of examples highlight the underlying economics. Suppose young firms are financially constrained and less skilled in 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 invested insufficiently in R&D even without wedges, but are now pushed even further away from the optimum. Alternatively, suppose large firms have market power over inventors but are also better at capturing rents from inventions, such that firms with large R&D wedges (from monopsony power) tend to have lower impact-value factors. Then, market power pushes the allocation of R&D workers away from low impact-value factor firms and, thus, is less costly in terms of growth.

R&D wedges are closely linked to optimal growth policy as shown in Proposition 3. To achieve maximal growth, the planner uses R&D subsidies to perfectly offset any variation in the product of impact-value factors and R&D wedges. This allocation achieves the frontier growth rate by setting $\Lambda_t \cdot \Xi_t = 1$, which suggests that any $\Lambda_t \cdot \Xi_t < 1$ yields a growth rate within the growth possibility frontier. Intuitively, firms maximize value, while the planner wants

to maximize growth. As long as both are not perfectly aligned, the allocation of R&D inputs is inefficient from the perspective of growth maximization.

Proposition 3. *Let g_t^* be the growth frontier achieved in the Planner Growth Equilibrium, then $g_t^* = g_t^F$. Furthermore, this allocation can be achieved by setting the R&D subsidy component of Δ_{it} to equalize $\zeta_{it} \cdot (1 + \Delta_{it})$ across firms.*

1.3 Extensions

I consider several extensions in Appendix D. First, I show that allowing for **entry**, rather than holding the mass of firms constant, can amplify the cost of private frictions. R&D wedges reduce the expected profits of innovative firms due to rising input costs, leading to lower entry. Fewer firms implies more inventors per firm and, due to decreasing returns to scale, lower productivity.

Second, I show that the formulae developed above 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 firms that expand in absence of frictions also tend to be more productive.

Third, the recent literature on labor market power emphasizes that firms might be imperfect substitutes for workers due to amenities or **specialization**, which may limit the gains from reallocation (Card et al., 2018; Berger et al., 2022). I show that the formulae developed above are preserved in this case with a lower implied scale elasticity γ reflecting specialization. Resultingly, R&D wedges tend to be less costly due to lower gains from reallocation.

Finally, the model assumes fixed aggregate R&D inputs, implying that the level of R&D wedges does not affect growth. I show that the formulae derived above extend to the case with positive **input supply elasticity**, however, there is a supply adjustment term depending on the average R&D wedge. The gains from reallocation thus coincide as long as 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.

2 Data and Measurement

2.1 Data

I focus my empirical analysis on research-active firms listed on US stock exchanges. This 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 concerns together with the main results.

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 sales and employment.

I use patents granted by the US Patent and Trademark Office (USPTO) to measure R&D outputs. Patents are arguably the most direct measure of R&D output available to researchers. They capture an invention that the patent office deemed new and useful, and grant the owners 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). I propose to address this concern by focusing on firms that tend to patent and by investigating robustness using measures independent of the patent system.

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 its 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 forward-citations as a measure of the growth impact of R&D as I discuss below. I construct forward-citations, i.e., citations received, by the patent within the first 5 years since its grant date using the USPTO's

citations files and normalize them by their average value within an application year. These adjustments ensure that citations remain comparable across years.

I aggregate forward-citations and valuations to the firm-year-level using the patent-to-firm 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 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 more details.

2.2 Measurement

R&D Allocative Efficiency depends on four parameters: $\{\gamma, \{\omega_{it}, \Delta_{it}, \zeta_{it}\}\}$. The consensus in the literature is to set $\gamma = 1/2$, which implies an R&D unit-cost elasticity of -1 ([Acemoglu et al., 2018](#); [Akcigit and Kerr, 2018](#)).⁵

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:

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

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 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:

⁵[Terry \(2023\)](#) estimates a γ of 0.4, while [Terry et al. \(2023\)](#) estimate 0.8 at the county-level. The estimates in [Dechezleprêtre et al. \(2023\)](#) imply a value for γ of 2/3 or larger, however the authors attribute the large estimate to the presence of financial frictions.

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

Three measurement concerns arise immediately. First, R&D return dispersion may arise due to differences in the scale elasticity of the innovation production function γ across firms. Without adjustment, one would confound those for differences in R&D wedges. I, thus, residualize measured R&D wedges with respect to industry \times year cells under the assumption that technology is similar within them. Secondly, measurement of R&D wedges requires ex-ante expected R&D returns as firms equalize expected marginal benefits to marginal costs.⁶ Realized R&D returns might differ from their expected value due to the stochastic nature of the innovation process. Thus, I restrict the sample to observations with at least 50 patents over the 5-year window to leverage the law of large numbers in closing the gap between averages and expectations. This approach does not account for firm-level shocks that yield common variation in realized patent valuations, which I investigate separately. Finally, not all inventions are patented and patenting rates may differ across firms (Cohen et al., 2000). Such differences could lead to variation in measured R&D returns due to patenting choices rather than R&D wedges. Following Bloom et al. (2020), I use non-negative changes in sales or employment as alternative measures of innovation output following the idea that inventions lead firms to expand. My alternative measure of R&D wedges is then

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

I discuss additional measurement concerns in Section 4.3.

⁶The 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.

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

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

Note that this formula is scale independent such that we can drop the common wage intercept. I thus measure R&D productivity as

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

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 use measure impact-value factors from markups guided by their link in a limit-pricing setup. In these models, firms' innovation quality is only partially reflected in markups such that they are not able to fully capture the additional social value created.⁷ I obtain markups either 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

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

In the third approach, I propose to use forward-citations as a direct measure of the growth-impact of inventions (up to a constant factor) such that the impact-value factor is given by the ratio of patent citations to valuations:⁸

$$\hat{\zeta}_{it} = \frac{\sum_{s=0}^4 \text{Patent Citations}_{it+s}}{\sum_{s=0}^4 \text{Patent Valuations}_{it-1+s}}. \quad (7)$$

One potential concern with this measure is heterogeneity in citation conven-

⁷Let $\lambda_i > 1$ be the quality improvement over a potential competitor in a model with limit pricing. Then, profits are proportional to markup $1 - \lambda_i^{-1}$, however, the growth impact is proportional to $\lambda_i - 1$. The impact-value factor is the ratio of both, which is proportional to λ_i itself. Thus, we can use either markups or profit rates to back-out the implied λ_i .

⁸See [Ayerst \(2022\)](#) and [Akcigit and Kerr \(2018\)](#) for similar interpretations of citations.

tions across industries or time that affect the relative frequency of citations even if growth impacts are comparable. I thus residualize impact-value factors with respect to industry-year fixed effects. Another concern when relating impact-value factor to the R&D return is that they might be related by construction due to the use of patent valuations. I, thus, replace them with changes in sales when relating the impact-value factor to R&D wedges.

3 Exploring R&D Returns

Before estimating aggregate measures, I investigate the behavior of R&D wedges in the data. I use the terms R&D wedges and R&D returns interchangeably here given their measurement equivalence in my context.

3.1 Basic Facts


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 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.

R&D return dispersion has increased throughout the sample. Comparing the early and late sample in Table 1.B, I find that R&D return dispersion has risen by 32%. Rising dispersion is a broad phenomenon with about 64% of industries having more dispersed R&D returns in the later sample.

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.A, I find that R&D return dispersion is $0.93/0.64 - 1 \approx 46\%$ larger in my sample.

Large dispersion in R&D returns is not surprising if we expect significant

Figure 2: R&D Returns are Highly Dispersed



Notes: Histogram of log R&D returns and density function of a normal distribution with same mean and variance. See text for details.

Table 1: Return Dispersion Across Comparison Groups

Within Cell	R&D Return	Return on Capital	
	SD	SD	$\Delta\%$
<i>A. Across Industries</i>			
—	1.09	0.77	42%
Year	1.05	0.74	43%
NAICS3 \times Year	0.93	0.64	46%
NAICS6 \times Year	0.85	0.58	45%
<i>B. Across Time</i>			
1975 – 2014	0.93	0.64	45%
1975 – 1990	0.74	0.46	62%
2000 – 2014	0.98	0.73	33%
<i>C. Across Measures</i>			
Patent valuations	0.93	0.64	46%
Δ Revenue	1.11	0.64	73%
Δ Employment	1.35	0.64	111%

Note: Return measures residualized with respect to fixed effects indicated in first column. Column headers SD report standard deviations of return measure. Columns headers $\Delta\%$ indicate percent difference of Return on R&D dispersion with respect to return in consideration. Returns are measured in logs.

measurement error therein. At least three concerns may motivate that belief. First, we might be concerned about the use of patents to measure innovation as it has been long 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. One important dimension for this consideration is industry differences in patenting conventions that could contribute to R&D return dispersion. For example, patents are considered quite important in life sciences, but less so in manufacturing (Mezzanotti and Simcoe, 2023). Empirically, I find that the contribution of such cross-industry differences is small relative to overall R&D return dispersion. Variation across firms in the same 6-digit industry and year accounts for $0.85/1.09 \approx 78\%$ of the overall R&D return dispersion as reported in Panel A of Table 1. It is also not the case that the importance of patenting is a strong predictor for within-industry dispersion. For example, R&D returns are similarly dispersed within life sciences and manufacturing as reported in Appendix Table C.2, even though both industries differ significantly in the degree to which they consider patents essential to their intellectual property protection strategy. Furthermore, R&D return dispersion is a robust finding across alternative measures of R&D output that do not rely on patents. For example, the dispersion is $1.11/0.93 - 1 \approx 19\%$ larger when using revenue growth instead of patent valuations to measure R&D output as reported in Panel C.⁹

A second potential concern relates to the use of patent valuations. Kogan et al. (2017) estimate these by using a non-linear transformation of stock-market returns around the patent announcement window. This procedure likely entails some measurement error as it is impossible to disentangle other events impacting the firm concurrently from the patent value. To the degree that these confounding events are quantitatively important and independent of each other over time, we might thus expect that R&D returns are partly driven by classical measurement error and, thus, uncorrelated over time. In-

⁹Naturally, these measures have issues of their own, however, I find that they are highly correlated with my preferred measure as reported in Panel B of Table 2.

stead, I find that R&D returns are highly persistent, as shown in Table 2, Panel A, with an implied annual autocorrelation coefficient of $0.697^{1/5} \approx 0.93$. Importantly, all measures of R&D returns are highly persistent. A structural variance decomposition of R&D returns leveraging this insight confirms that transitory shocks, such as classical measurement error, contributes almost none of the overall variation in R&D returns (Appendix F). Another concern might be that patent valuations are strictly positive by construction, while some patents might be actually worthless (Jaffe and Lerner, 2007). I show in Appendix C.3 that excluding low value patents from R&D returns does not reduce their dispersion. Lastly, Kogan et al. (2017) assume that all patents have an equal likelihood of being granted ex-ante to back out the value of innovation from stock returns, which only reflect the unexpected component. I investigate adjustments based on the patent’s technology class and quality in Appendix C.3 and find that neither reduces measured R&D return dispersion.

Table 2: R&D Return Consistency across Time and Measures

Variable	Estimate	Std. err.	R^2	Observations
<i>A. 5-Year Autocorrelation</i>				
Patent valuations	0.697***	(0.020)	45.7%	7,623
Δ Sales	0.564***	(0.024)	29.6%	7,455
Δ Employment	0.552***	(0.026)	27.8%	6,447
<i>B. Correlation with Baseline R&D Returns</i>				
Δ Sales	0.597***	(0.033)	25.1%	11,688
Δ Employment	0.551***	(0.038)	14.1%	10,870

Note: Each row reports the regression coefficient of a separate regression with dependent and independent variable in logs. Panel A reports 5-year autocorrelation coefficients for alternative measures of R&D returns. The respective R&D return is calculated as the ratio of the variable indicated in column 1 and R&D expenditure at the 5-year level. Panel B reports contemporaneous correlations with alternative measures of R&D returns as dependent variables and the primary measure of R&D returns as the independent variable. All regressions control for NAICS3 \times Year fixed effects and standard errors are clustered at the NAICS6 level. See text and Appendix for additional details.

A third and final concern is measurement error related to the distinctions between expected returns, which should be equalized in the frictionless model, and realized returns, which might not be. Under rational expectations, the

gap between realized and expected R&D returns should not be predictable with information available to the firm at the time of the investment. Under the assumption that expected returns are equalized across firms, it should then also not be possible to predict differences in realized returns across firms. In contrast, and as discussed above, I find that R&D returns are highly persistent and are therefore predictable. I show that predictability persists at long time horizons in Appendix C.2. These results suggest that R&D dispersion is not primarily driven by expectation-realization gaps. I further investigate whether “superstar patents”, which we could be important given the famously fat-tailed distribution of innovation outcomes, contribute to R&D return dispersion. I find that measured dispersion is essentially unaffected by winsorizing patent valuations at the 95th percentile as reported in Appendix C.3.

In summary, my robustness exercises suggest that neither measurement error arising from the use of patents and patent valuations nor classical measurement error appear to be significant drivers of R&D return dispersion.¹⁰ Table C.1 further investigates robustness with respect to the specification and find that neither expanding the aggregation window nor changing the timing leads to lower R&D return dispersion. Focusing on observations with significantly more patents reduces measured dispersion marginally.

3.2 Economic Drivers

I investigate economic drivers of R&D return dispersion in Table 3.¹¹ First, I investigate general investment frictions in Panel A. Following the idea that investment may be distorted across multiple margins, I investigate whether measures of capital investment frictions correlate with R&D returns and find mixed results. On the one hand, R&D returns are uncorrelated with the return on capital, which is considered a summary measure for investment frictions (David et al., 2016). On the other hand, R&D returns are highly

¹⁰Appendix C.3 also considers misspecification arising from fixed costs, acquisitions of innovative firms, and knowledge capital.

¹¹See Online Appendix E for the associated theoretical foundations.

correlated with Tobin’s Q , which is also an established measure of investment frictions (Peters and Taylor, 2017). I also find mixed results for financial frictions. Firms with more liquidity, which might be less constrained by cash flow concerns, tend to have lower returns, in line with the idea that they are less constrained. On the other hand, firms with large dividend payments, which presumably are not very constrained either, have larger rather than smaller R&D returns. Finding inconclusive results for measures of financial frictions is surprising as a growing literature argues that intangible investments, including R&D, are particularly constrained by them (Brown et al., 2009).

Second, R&D return dispersion could be driven by firm-specific risk-premia due to differential exposure to aggregate risk (David et al., 2022). As reported in Panel B, I find no evidence that stock market β s, a measure of systematic risk, explain R&D returns. However, firms with volatile patent valuations tend to have higher R&D returns. Such a “risk-premium” could arise if firms’ decision makers cannot fully diversify the 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 (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, a growing literature finds that monopsony power plays an important role in shaping the allocation of workers (Card et al., 2018; Lamadon et al., 2022). The literature also finds that high-skilled workers, a group likely including inventors, are more affected by monopsony power and that larger firms tend to have more thereof (Seegmiller, 2023; Berger et al., 2022; Yeh et al., 2022). Monopsony power over inventors could drive R&D return dispersion as firms with more of it restrict their hiring more aggressively to keep wage low and, as a result, create more value per unit of cost (Lehr, 2024). In

Table 3: Correlations with R&D Returns

Variable	Estimate	Std. err.	R^2	Observations
<i>A. Frictions</i>				
Return on Capital	0.043	(0.068)	0.1%	11,844
Tobin's Q	0.202***	(0.030)	6.3%	10,471
Liquidity	-0.048**	(0.022)	0.3%	10,568
Dividend rate	36.499***	(7.142)	1.5%	11,499
<i>B. Risk</i>				
CAPM β	0.001	(0.065)	0.0%	6,799
Valuation risk	0.493***	(0.132)	1.4%	10,961
<i>C. Taxation</i>				
R&D user cost $1 - \tau$	-0.527	(0.617)	0.1%	11,247
Public patent involvement	1.392	(1.240)	0.2%	11,845
<i>D. Inventors</i>				
Inventors	0.228***	(0.032)	7.2%	11,845
Firm dominance	0.142***	(0.045)	1.4%	10,477
Inventor specialization	0.233***	(0.083)	0.4%	11,828
<i>E. Dynamics</i>				
Long-term R&D growth	0.424***	(0.054)	10.2%	6,525
Long-term TFP growth	0.451***	(0.050)	3.7%	5,421
Prior excess stock return	0.260***	(0.031)	1.1%	10,087
Prior TFP growth	0.316***	(0.041)	1.1%	7,277

Note: Each row reports the regression coefficient of a separate regression with dependent variable log R&D returns. All regressions control for NAICS3 \times Year fixed effects and standard errors are clustered at the NAICS6 level. See text and Appendix for details.

line with this idea, I find that firms hiring more inventors have larger R&D returns in Panel D.¹² In addition, I find that firms dominating their 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., 2023).

Lastly, I investigate the link between R&D returns and firm dynamics in Panel E. Factor return dispersion arises naturally in models with adjustment costs (Asker et al., 2014). With such frictions, a positive shock to R&D pro-

¹²This finding is also in line with a size-dependent R&D scale elasticity assuming that up-scaling is costlier for already large companies.

ductivity leads to a temporary rise in R&D returns as R&D output, which captures R&D inputs and productivity, adjusts faster than R&D inputs alone, leading to a positive correlation between R&D growth and returns. In line with this prediction, I find that long-term growth in R&D, i.e., its growth rate from $t-1$ to $t+6$, can account for around 10% of R&D return dispersion. Similarly, I find a strong correlation with long-term TFP growth, which can be rationalized by the same mechanism. Finally, I also find that prior TFP growth and stock returns are predictive for R&D returns, suggesting, again, that they are associated with firm expansion and other positive events for the firms. Such results may also be in line with information frictions or behavioral biases leading to a slow response to positive news (Hirshleifer et al., 2013).

Overall, the explanatory power of the mechanisms considered here is low, echoing similar results in the return on capital literature (David et al., 2016). This finding makes the interpretation of measured R&D wedges difficult, since we do not fully know their source. Nonetheless, the documented dispersion marks a stark deviation from the predictions of a frictionless model and can be interpreted using the formulae developed in Section 1.

3.3 Relationship with Impact-Value Factors

I find inconclusive evidence on the link between R&D wedges and the impact-value factor in Table C.5. While markup-based measures suggest slightly positive correlation, patent-based measures suggest a negative correlation.

4 Growth, Wedges, and Policy

I next turn to estimating the macroeconomic impact of R&D wedges.

4.1 Combining Data and Model

Measurement. Per Proposition 2, I estimate R&D Allocative Efficiency as

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

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 discuss sample selection further below and highlight that, within my sample, it is indeed the case that there appears to be more “misallocation” among smaller firms.

I consider two scenarios for adjustment factor $\hat{\eta}_t$. In the first case, I assume that R&D wedges and impact-value factors are independent and, thus, set $\hat{\eta}_t = 1$. In the second case, I estimate its value using citations divided by sales growth as a proxy for the impact-value factor. The adjustment factor is $\hat{\eta}_t = \sqrt{1 + 2 \cdot \hat{\beta}_t}$, where $\hat{\beta}_t$ is the regression coefficient when regressing R&D returns on the impact-value factor over a centered rolling 10-year window.

To get a sense of longer-run developments, I collapse annual estimates using geometric averages. I consider the average over the full sample from 1975 to 2014 as well as the early and late periods, 1975–90 and 2000–14, respectively. Comparing the early and late period gives us a window into long-run changes in R&D Allocative Efficiency and their ability to shed light on declining economic growth.

Finally, I calculate bootstrapping standard errors for the estimates. For each year, I sample firm 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.

Counterfactuals. 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 technologically feasible. Even when this is not the case, 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 first 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 $g^C = 1.5\% \cdot \Xi^{-1}$ to match the long-run US growth rate. I refer to this scenario as the endogenous growth case since misallocation reduces the long-run growth rate.

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

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

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 responds to changes in the environment as does the long-run productivity level.¹³ In the counterfactual, I assume that the economy is on its long-run growth path before 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%.

4.2 The Long-run View

Consider the case of unrelated R&D wedges and impact-value factors first. The blue line in Figure 3 plots the annual estimates of R&D Allocative Efficiency Ξ_t , while long-run values are reported in Panel A of Table 4. The table also reports the welfare cost in consumption-equivalent terms.

Frictions have a significantly negative impact on economic growth as measured through *R&D Allocative Efficiency*. I estimate an average growth impact of -21.3% for the full sample, which suggests a growth rate of 1.9% in absence of R&D wedges based on a realized annual productivity growth rate of 1.5%. Unsurprisingly, such a stark slowdown of economic growth has large welfare consequences. The model suggest that welfare would be 12% higher in absence of R&D wedges. For comparison, Berger et al. (2022) estimate that

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


Table 4: The Impact of R&D Wedges on Economic Growth and Welfare

Time Horizon	Growth Impact $\Xi - 1$			Welfare Cost	
	Est.	Std. Err.	95% CI	End.	Semi-End.
<i>A. Baseline</i>					
1975–2014	-21.3%	(0.41%)	[-21.9% -20.5%]	12.3%	11.7%
1975–1990	-14.7%	(0.49%)	[-15.4% -13.8%]	7.6%	7.4%
2000–2014	-24.0%	(0.68%)	[-25.0% -22.8%]	14.5%	13.7%
Δ Change	-10.9%			5.4%	5.3%
<i>B. Adjusted</i>					
1975–2014	-17.9%	(0.36%)	[-18.4% -17.3%]	9.8%	9.4%
1975–1990	-12.0%	(0.40%)	[-12.6% -11.2%]	6.0%	5.9%
2000–2014	-21.7%	(0.62%)	[-22.6% -20.6%]	12.6%	12.0%
Δ Change	-11.0%			5.5%	5.3%

Notes: Table reports estimates for impact of R&D wedges across samples together with their implications for welfare. Welfare changes are in consumption equivalent terms. Standard errors and confidence intervals are calculated using bootstrapping.

monopsony in the production sector reduces US output by 21% and welfare by 8%, while [Hsieh and Klenow \(2009\)](#) estimate 30%–40% larger US output in absence of production factor misallocation.

Figure 3: R&D Allocative Efficiency Has Declined



Notes: Annual estimates for R&D Allocative Efficiency $\Xi_t - 1$. Adjusted values estimate the adjustment factor over a 10-year rolling window. The shaded area covers the 90% confidence interval.

The annual estimates further suggest that declining R&D Efficiency contributed to the long-run growth slowdown. Estimated R&D Efficiency declined throughout the sample period with a sharp downturn and partial recover during the 1990s, potentially reflecting dynamics during the Dot-Com boom. Comparing the estimates for the 1975–90 and 2000–14 period, I find that R&D Allocative Efficiency declined from -15% to -24% with an associated welfare loss of around 5%. This decline implies an $\frac{24\%-15\%}{1-15\%} \approx 11\%$ slower short-run growth rate, which accounts for $\frac{11\%}{40\%} \approx 25\%$ of the growth slowdown.¹⁴

Adjusting for the impact-value factor reduces the cost of R&D returns marginally, but leaves their evolution essentially unaffected. As reported in Panel B, the long-run estimated impact of R&D wedges is -18%, which is slightly better than the unadjusted estimate. Nonetheless, the estimated welfare cost of 9%–10% remain large. Finally, the change in the economic growth rate implied by the evolution of the R&D Efficiency remains -11%.

4.3 Discussion

Table 5 report a range of robustness exercises.

First, **not all inventions are patented**, making patent valuations a potentially incomplete measure of the private value created from R&D, even if is the most reliable one. I consider alternative measures of R&D output in Panel A and find that using patent valuations yields the largest estimate of R&D Allocative Efficiency. Its decline is more (less) pronounced when using sales (employment) growth instead of patent valuations.

Second, we do not have a convincing measure for the **impact-value factor**. My preferred specification estimates adjustment factor η by regressing citations over sales growth on R&D returns. Alternatively, I consider using the profit-based measure of the impact-value factor to estimate the adjustment term in Panel B. I find that the resulting estimates are slightly larger, while changes over time continue to hover around a growth impact of -11%.

¹⁴Total factor productivity growth declined from 0.5% for the 1976–1995 period to 0.3% for the 2005–2018 period, a $\frac{0.5\%-0.3\%}{0.5\%} \approx 40\%$ reduction. See Figure 1.

Third, I investigate the impact of **entry and exit** by focusing on continuing firms only.¹⁵ The estimates continue to suggest a significantly negative impact of R&D wedges on growth of -16%, while changes over time among continuing firms predict a 9% reduction in economic growth.

Fourth, **measurement error** is an important consideration regardless of the precise measure of R&D wedges, however, I find little evidence thereof in practice. I consider two sources in detail in Appendix F. First, the outcome of each innovation effort is uncertain and, thus, we might be concerned that some of the variation in measured R&D wedges is due to firms being more or less lucky in their research projects. I propose to estimate the contribution of this channel in a bootstrapping approach in which I first redraw firms' patent valuations, and then calculate how far aggregated values are from the true expectation as measured by the firms actual patent valuation. Naturally, these differences are smaller for firms with more patents by the law of large numbers. Second, firms might be subject to ex-post firm-level shocks that have a uniform effect on their R&D output. Such variation is not accounted for by the bootstrapping approach as it is common across all inventions within a given period. I propose an estimation methodology for this source of variation using a GMM estimator in Appendix F. The main idea is to exploit the persistence of R&D returns to estimate the contribution of non-persistent variation, such as one-off luck or firm-level measurement error, to the overall dispersion in R&D returns. My results suggest almost no contribution of these two sources of "measurement error" for my main estimates as reported in Panel D. An alternative adjustment is proposed in [Bils et al. \(2021\)](#), where the authors assume that measurement error is additive in levels rather than logs. I implement this adjustment procedure and find that it indeed improves the estimated R&D Allocative Efficiency. As reported in Panel D, the adjusted measure declines by 5% compared to 10% in the baseline, however, the associated welfare costs remain sizable. See Appendix F.3 for details.

¹⁵See Appendix C.5 on how I calculate estimates for continuing firms.

Fifth, the joint measurement of **observation weights** and R&D returns could yield mechanical correlation. Panel E confirms that equalizing weights or using sales growth instead of patent valuations yields smaller estimates of R&D Allocative Efficiency while holding the decline approximately equal.

Finally, and as discussed above, the **sample** is skewed towards larger firms due to the data sources and selection criteria. A representative sample would include a larger share of smaller firms with, presumably, smaller R&D expenditure per firm as well. I explore size-heterogeneity in Panel F and find that R&D Allocative Efficiency is smaller and declining faster among firms with lower R&D expenditure is lower. This finding is in line with a world in which smaller firms face more frictions, which is commonly assumed (Brown et al., 2009), and suggest that a representative sample of firms might feature lower estimated R&D Allocative Efficiency as well as a faster decline therein, such that the estimates reported are conservative as to the impact of frictions in the R&D sector. See Appendix C.7 for details.

5 Conclusion

This paper presents evidence that frictions, and their impact on the allocation of R&D resources, contributed to the recent decline in aggregate R&D productivity and U.S. economic 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 US-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. I measure these as the ratio of patent valuations divided by R&D expenditure and show that there are large and persistent differences therein. In contrast, the frictionless model predicts return equalization and associates R&D return dispersion with frictions. R&D return dispersion persists in a large set of robustness exercises and measure-

Table 5: R&D Wedges, Economic Growth and Welfare — Robustness

Specification	Growth Impact $\Xi - 1$				Δ Welfare	
	1975–2014	1975–90	2000–14	Δ	End.	Semi-E.
Baseline	-17.9%	-12.0%	-21.7%	-11.0%	5.5%	5.3%
<i>A. Value of Innovation</i>						
Δ Sales	-25.5%	-19.6%	-32.0%	-15.4%	8.1%	7.8%
Δ Employment	-40.3%	-39.0%	-43.9%	-8.0%	3.8%	3.7%
<i>B. Impact-Value Adjustment</i>						
Profit-Based	-22.9%	-15.7%	-25.4%	-11.5%	5.8%	5.6%
<i>C. Entry & Exit</i>						
Continuing Firms	-16.3%	-11.5%	-19.4%	-8.9%	4.3%	4.2%
<i>D. Measurement Error</i>						
Direct adjustment	-18.0%	-12.1%	-21.7%	-10.9%	5.4%	5.3%
Bils et al. (2021)	-13.7%	-10.5%	-15.1%	-5.1%	2.3%	2.3%
<i>E. Observation Weights</i>						
Unweighted	-21.6%	-16.1%	-25.1%	-10.7%	5.3%	5.1%
Sales growth	-23.9%	-18.2%	-26.3%	-9.8%	4.8%	4.7%
<i>F. Firm Size</i>						
Small R&D	-25.2%	-16.8%	-30.7%	-16.7%	9.0%	8.7%
Larger R&D	-16.3%	-10.3%	-19.9%	-10.7%	5.3%	5.1%

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.

ment error adjustments. Lastly, 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 suggests that frictions reduce US economic growth significantly and increasingly so. I estimate for the full sample that economic growth was 18% slower due to frictions, implying a welfare cost of 11% in

consumption-equivalent terms. Furthermore, I find that rising frictions can account for an 11% lower growth rate for 2000–14 compared to 1975–90, which accounts for 25% of the observed productivity slowdown.

These findings suggest important avenues for future research. Most importantly, more research is needed to understand the underlying forces driving rising frictions. A thorough understanding of the variation in R&D wedges and impact-value factors may allow for the development of potentially targeted policies and could thus be essential for improving U.S. R&D productivity and economic growth.

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**, “Good Rents versus Bad Rents: R&D Misallocation and Growth,” 2022.
- , —, —, **Peter J. Klenow**, and **Huiyu Li**, “A Theory of Falling Growth and Rising Rents,” *Review of Economic Studies*, 2023, *forthcoming*.
- , **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,” 2022.
- 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.
- Bils, Mark**, **Peter J. Klenow**, and **Cian Ruane**, “Misallocation or mis-measurement?,” *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.
- 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*, 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*, 2023, *forthcoming*.
- 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.
- Ewens, Michael, Ryan Peters, and Sean Wang**, “Measuring Intangible Capital with Market Prices,” 2022.
- Fons-Rosen, Christian, Pau Roldan-Blanco, and Tom Schmitz**, “The Effects of Startup Acquisitions on Innovation and Economic Growth,” 2023.
- 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.
- Hirshleifer, David, Po Hsuan Hsu, and Dongmei Li**, “Innovative efficiency and stock returns,” *Journal of Financial Economics*, 3 2013, *107*, 632–654.
- 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.
- Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler**, “Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market,” *American Economic Review*, 1 2022, *112*, 169–212.
- Lehr, Nils H.**, “Does Monopsony Matter for Innovation?,” 2024.
- 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**, “Do R&D Tax Credits Create Jobs?,” 2019.
- 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 and Appropriability: Revisiting the Role of Intellectual Property,” *NBER Working Paper No. 31428*, 2023.

- Midrigan, Virgiliu and Daniel Yi Xu**, “Finance and misallocation: Evidence from plant-level data,” *American Economic Review*, 2014, 104.
- Olmstead-Rumsey, Jane**, “Market Concentration and the Productivity Slowdown,” 2022.
- Peters, Ryan H. and Lucian A. Taylor**, “Intangible capital and the investment-q relation,” *Journal of Financial Economics*, 2017, 123.
- 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, and Bledi Taska**, “Employer Concentration and Outside Options,” 2023.
- Seegmiller, Bryan**, “Valuing Labor Market Power : The Role of Productivity Advantages,” 2023.
- Terry, Stephen J.**, “The Macro Impact of Short-Termism,” *Econometrica*, 2023, 91, 1881–1912.
- , **Konrad B. Burchardi, Thomas Chaney, Lisa Tarquinio, and Tarek A. Hassan**, “Immigration, Innovation, and Growth,” 2023.
- , **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 $\theta_{it} = \varphi_{it} \cdot V_{it}$ we can solve for firms' demand for R&D inputs as

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

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

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

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

$$g_t = \int_0^1 \zeta_{it} \cdot \ell_{it} \cdot \frac{W_t}{\gamma} \cdot di.$$

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

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

Some rearrangement yields the formulae in the proposition. \square

Proof of Corollary 1. The formula follows immediately since the terms in the nominator 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. \square

Proof of Proposition 3. The planner problem is given by

$$\begin{aligned} \max \quad & g_t = \int_0^1 \zeta_{it} \cdot z_{it} \cdot V_{it} \cdot di \\ \text{s.t.} \quad & L_t = \int_0^1 \ell_{it} \cdot di \quad \text{and} \quad z_{it} = \varphi_{it} \cdot \ell_{it}^\gamma \end{aligned}$$

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

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

where λ_t^W is the shadow wage.

One can confirm immediately, that the implied allocation coincides with the competitive equilibrium iff $\zeta_{it} \cdot (1 + \Delta_{it})$ 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 $1 + \Delta_{it} = 1/\zeta_{it}$. \square

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

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

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. The approximation is precise if all variables are jointly log-normal and, in this case, 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. \square

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

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

In turn, it is straight-forward to show that a second order approximation of the formula in Proposition 2 yields the same expression. \square

B Data Appendix

Inventor employment. Let $\mathcal{P}_{it \rightarrow t+4}$ be the set of successful patent applications for firm i between t and $t+4$ and $\mathcal{I}_{it \rightarrow t+4}$ be the set of associated inventors. I denote the number of patents assigned to firm i and listing j and inventor 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

$$\text{Inventors}_{it \rightarrow t+4} = \sum_{j \in \mathcal{I}_{it \rightarrow t+4}} \frac{\sum_{s=0}^4 P_{ijt+s}}{\sum_{s=0}^4 P_{jt+s}}. \quad (\text{B.1})$$

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. As for the R&D return, I construct the measure at the 5-year level:

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

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 ($\text{prcc_f} \times \text{csho} + \text{dltt} + \text{dlc} - \text{act}$), to physical capital (ppeqgt).

Liquidity. I define liquidity as cash holdings divided by assets ch/at .

Dividend rate. I define the dividend rate as dividends over assets dvt/at .

Public patent involvement. I classify patents as connected to public actors either if they are assigned to a government entity, research lab, or university, or if they have 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 patent within a 5-year window, I calculate the share of inventors working for the firm among those that worked on patents in the exactly same technology classification. For the latter, I use the complete CPC classification of the patent, which has more than 600 technology classes, which 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 in 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 vice versa.

Inventor specialization is constructed in two steps. First, I calculate inventor specialization for a given 5-year window as the average cosine similarity between patent classifications in an inventors portfolio of new patents. For each patent I then create an indicator vector over the set of available patent classification with individual categories weighted by their inverse frequency. I then calculate the average cosine similarity across all patents in the portfolio and take the simple average across all patents. This measure is between 0 and 1 by construction with 0 implying completely different patents and 1 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.

Online Appendix

Not for publication

C Additional Empirical Results

C.1 Additional Robustness for R&D Return Dispersion

Table C.1: R&D Return Dispersion Across Specifications

Specification	Standard Deviation	Observations
<i>A. Aggregation horizon</i>		
1-year	1.00	11,083
5-year	0.93	11,845
10-year	0.92	11,845
20-year	0.91	11,845
<i>B. Realization horizon</i>		
Same year	0.86	10,885
1-year	0.93	11,845
2-year	0.97	10,852
5-year	1.08	8,377
<i>C. Minimum Patents</i>		
50 patents	0.93	11,845
100 patents	0.84	7,846
200 patents	0.79	4,859

Note: All returns in logs and residualized with respect to NAICS3-Year fixed effects. Aggregation horizon is the number of years over which valuations and R&D expenditure are summed. Realization horizon is the difference between the patent application year and the year of R&D expenditure considered. Unless otherwise specified, R&D returns are measured with a 5-year aggregation horizon, 1-year realization horizon, and 50 minimum patents.

C.2 The Realization-Expectation Gap

The standard, frictionless firm investment model adapted for R&D equalizes the expect return on R&D across all firms. Variation in realized R&D returns could then arise simply due to the stochastic nature of innovation:

Table C.2: R&D Return Dispersion Across Industries

Industry	Standard Deviation	Observations
All industries	0.93	11,844
Life Science	0.82	1,630
IT	1.07	4,732
Manufacturing	0.83	5,108
Other	0.83	374

Note: R&D returns residualized with respect to NAICS3×Year fixed effects. Returns are measured in logs. Industries are defined as in [Mezzanotti and Simcoe \(2023\)](#).

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

where the expectations are taken with respect to the firm’s information set at the time of the investment decision.

Assuming that firms optimize and have rational expectation, the realization-expectation gap should be an i.i.d. random variable.¹⁶ Thus, one should not be able to predict it using any information that in the firm’s information set at the time of the investment decision. Otherwise, the firm would chose a different investment level that would make the R&D return unpredictable beyond its expected level. Furthermore, this property should extent to the realized R&D return overall as long as expected R&D returns are indeed equalized.

As I show next, this prediction is not borne out by the data. In particular, I show that past R&D returns are a consistent predictor of future R&D returns by estimating a simply autoregressive model:

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

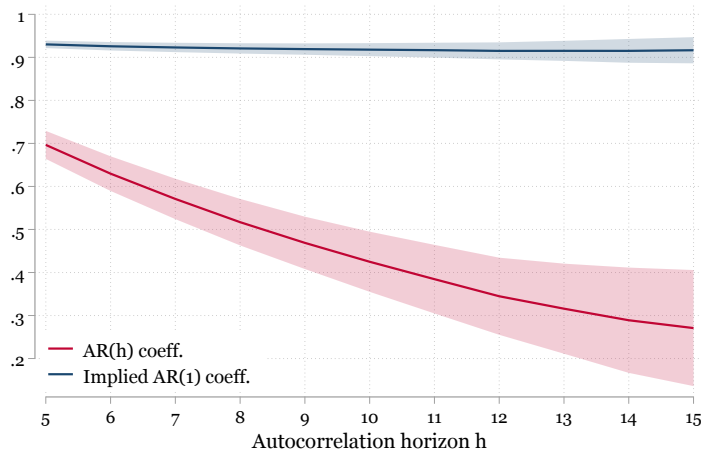
¹⁶This statement is general and can be seen immediately by assuming innovation outcomes are log-normal. In this case, the realization-expectation gap is a normal random variable.

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 C.1 suggest that R&D returns are highly predictable, even at longer horizons, using their own 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 using their h 's 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, thus, reject the hypothesis that variation in R&D returns is primarily driven by the stochastic nature of realization vs expectation.

Figure C.1: 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.

C.3 Measurement Robustness

Table C.3: Return Dispersion with Adjustments

Adjustment	Standard Deviation	Observations
<i>A. Acquisitions</i>		
Unadjusted	0.923	11,829
Adjusted ($s = 6.3\%$)	0.910	11,829
Adjusted ($s = 8.5\%$)	0.909	11,829
Adjusted ($s = 100\%$)	0.982	11,829
<i>B. Fixed-costs</i>		
Unadjusted	0.924	11,807
Adjusted	0.937	11,807
<i>C. Knowledge capital</i>		
R&D Expenditure	0.925	11,845
Knowledge capital	0.961	11,845
Organizational capital	0.985	11,845

Note: See text for description of measures. All return measures residualized with respect to NAICS3×Year fixed effects. Second column reports standard deviation of log R&D returns.

Acquisitions, which are common in the innovation economy, might lead to measurement error in R&D returns due to misattribution, i.e., by not counting all R&D costs associated with the measured patents (Phillips and Zhdanov, 2013; Fons-Rosen et al., 2023). Suppose that the firm acquires some inventions that are subsequently patented and added to total value created $z_{it} \cdot V_{it}$, however, the costs are recorded as acquisition cost $Aqc_{.it}$ instead of R&D expenditure $R\&D_{it}$. Assuming that the firm is otherwise unconstrained, the measured R&D return then becomes

$$\frac{V_{it} \cdot z_{it}}{R\&D_{it}} = \frac{1}{\gamma} \cdot \left(1 + \frac{Aqc_{.it}}{R\&D_{it}} \right),$$

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

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 expenditure on innovative products such that $\text{Aqc.}_{it} = s \cdot \text{Total aqc.}_{it}$. Total acquisition expenditure is 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 adjust R&D returns as

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

Panel A of Table C.3 reports the associated results. Adjusting for acquisitions marginally reduces R&D return dispersion, however, the magnitudes are small. Adjusting by 8.5% of total acquisitions reduces measured R&D return dispersion by 1.5%. Counting all acquisitions as R&D expenditure increases R&D return dispersion.

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

$$\frac{V_{it} \cdot z_{it}}{(f_i + \ell_{it}) \cdot W_t} = \frac{1}{\gamma} \cdot \frac{\ell_{it}}{f_i + \ell_{it}}. \quad (\text{C.3})$$

Resultingly, 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} and we have R&D return dispersion that is unrelated to frictions. Note, however, that the average R&D return for very large firms, i.e. $\ell_{it} \gg f_i$, 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 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 inverse of the average R&D return for firm 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 25th percentile of the R&D cost distribution. I can then estimate fixed costs and adjusted R&D returns as

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

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 be equalized across firms.

The fixed-costs adjustment increases measured R&D return dispersion marginally as reported in Table C.3, Panel B. Thus, fixed-costs do not appear to be a significant source of measured R&D return dispersion.

Knowledge capital. R&D expenditure is often interpreted as a cumulative investment in the firms knowledge base (Peters and Taylor, 2017). Under this alternative view, R&D capital, rather than expenditure, is the appropriate denominator for the R&D returns. I explore the robustness of my findings with respect to the input measure using the knowledge capital and organizational capital measures developed in Ewens et al. (2022). The knowledge capital measure is built up from R&D investments only, while organizational capital focuses on other overhead expenses. I refer to the sum of both as organizational capital. R&D return dispersion is strictly higher when using either the knowledge or organizational capital as reported in Table C.3, Panel C.

Outlier patents. Innovation outcomes are famously fat-tailed: While most inventions have moderate impacts, some transform entire industries (Akcigit and Kerr, 2018). This consideration raises the question as to whether variation in R&D returns is driven by “outlier-patents” with extremely large valuations. I investigate this question by creating winsorized measures of patent valuations that ignore value above the top 1% or top 5% of the annual patent valuation distribution and recalculate R&D return dispersion. As reported in Table

C.4, Panel A, winsorizing patent valuations at the top 1% (5%) reduces R&D return dispersion by 1% (3.6%). Thus, only a small fraction of the dispersion in R&D returns is potentially attributable to outlier patents.

Low value patents. Jaffe and Lerner (2007), among others, argue that changes in patent law, grant procedures, and enforcement have led to an onslaught of low quality patents with questionable economic value. The methodology in Kogan et al. (2017) takes into account low quality patents, however, we might still wonder whether their presence adds more noise to measured R&D returns.¹⁷ I investigate this question by constructing measures excluding valuations below 250k (500k) in 2010 USD and recalculating R&D return dispersion.¹⁸ As reported in Table C.4, Panel B, excluding low quality patents from the measure increases measured dispersion in R&D returns slightly.

Measurement details of patent valuations. Kogan et al. (2017) measure patent valuations using 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 of the patent being granted, then the change in firm valuation ΔM at the moment that the patent is granted should equal

$$\Delta M = (1 - \pi) \cdot V \quad \text{or, equivalently,} \quad V = \frac{\Delta M}{1 - \pi}. \quad (\text{C.4})$$

They measure the nominator using stock market returns and assume that the probability that a patent is granted is constant across all patents. The latter assumption is quite stringent for at least two reasons. First, patents of different patent classes might have different probabilities of being granted. For example, during the 1991–2014 period, 85% of patents applications classified as semiconductor memory devices (CPC subclass G11C) were granted within

¹⁷Patent valuations from Kogan et al. (2017) are strictly positive and monotonically increasing in the stock market return. Thus, even if the stock market is unresponsive, because the patent is worthless, the patent is assigned a positive value.

¹⁸I find that 10% (15%) of patents are valued less than 250k (500k) from 1975–84, while 17% (22%) are for the 2005–14 period.

Table C.4: Return Dispersion with Patent Valuation Adjustments

Adjustment	Std. dev. of R&D return	Observations
<i>A. Outlier patents</i>		
Unadjusted valuations	0.925	11,845
Winsorized at top 1%	0.916	11,845
Winsorized at top 5%	0.892	11,845
<i>B. Low value patents</i>		
All valuations	0.925	11,845
Valuations > 250k	0.942	11,845
Valuations > 500k	0.959	11,842
<i>C. Class grant rate</i>		
Unadjusted	0.925	11,845
Adjusted	0.933	11,845
<i>D. Value-dependent grant rate</i>		
$\eta = 0$ (Unadjusted)	0.925	11,845
$\eta = .05$	0.968	11,845
$\eta = .25$	1.163	11,845
$\eta = .5$	1.740	11,845
$\eta = 1.5$	2.584	11,845
$\eta = 2$	1.648	11,845

Note: See text for description of measures. All return measures residualized with respect to NAICS3× Year fixed effects. Second column reports standard deviation of log R&D returns.

3.5 years compared to 30% of those classified as healthcare informatics (CPC subclass G16H). Second, patent grant decision are assumed to be independent of the value of the patent. Such an assumption would not hold, e.g., if higher quality patents are more valuable, but also more likely to be granted.

I investigate whether either of these possibilities contributes to R&D return dispersion as follows. First, I use data on patent application and grant decision from the USPTO for the 1991–2014 period to calculate the grant probability by patent class 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 that were granted within 3.5 years of the application. Second, I assume that the probability of patent rejection takes the form $1 - \pi_p = \pi_0 \cdot V_p^{-\eta}$, where η measures the degree to which higher value patents are also more likely to be granted. The adjusted patent valuation is

$$\tilde{V}_p = \left(V \cdot \frac{1 - \bar{\pi}}{\pi_0} \right)^{\frac{1}{1-\eta}}. \quad (\text{C.5})$$

I calibrate π_0 at the annual level to keep the average patent valuation constant and experiment with alternative values for η . Optimally, one would want to estimate this value, however, we only observe valuations for granted patents.

I find that neither adjustment reduces the measured dispersion as reported in Table C.4. Adjustment for class grant probabilities, as reported in Panel C, increase R&D return dispersion marginally. Panel D considers value-dependent grant rates and finds that the resulting R&D return dispersion can be significantly larger. For example, it increases by 25% when assuming that a 10% larger patent valuation translated into a 2.5% higher grant probability.

C.4 Relationship with Impact-Value Factors

I investigate the link between R&D wedges and the impact-value factor in Table C.5 and find mixed results. On the one hand, markup-based measure suggest slightly positive correlation. For example, I find a significant positive correlation of the R&D return with the profit-implied markup, i.e., revenue divided by revenue minus cost, implying that firms with higher R&D returns also tend to have larger markups. The same relationship, although smaller in absolute magnitude, holds when using the markup measure developed in [Loecker et al. \(2020\)](#). To the degree that markup differences are primarily driven by persistent differences in the quality of innovation, as in a model with limit pricing and heterogeneous innovation quality, these results suggest that the impact-value factor might amplify the misallocation due to R&D return dispersion.

On the other hand, patent-based measures suggest a negative correlation. For example, I find a slightly negative correlation when measuring the impact-value factor as citations over sales growth. I also find a robust negative correlation when using citations over valuations to measure the impact value factor and sales growth over R&D expenditure to measure the R&D wedge. Finally, using the text-based patent quality measure developed in Kelly et al. (2021), I find a strong negative correlation with the R&D return. Thus, if these patent-based measures provide a good proxy for the impact-value factor, then they might partly offset misallocation due to R&D wedges.

Table C.5: The Relationship of R&D Wedges and Impact-Value Factors

Impact-Value Factor	Estimate	Standard Error	R^2	Observations
<i>A. Markup-based Measures</i>				
Estimated Markup	0.030***	(0.007)	2.7%	10,615
Profit-implied Markup	0.066***	(0.016)	4.6%	11,845
<i>B. Patent-based Measures</i>				
Citations/Δ Sales	-0.077	(0.054)	0.3%	11,688
Citations/Valuations*	-0.201***	(0.032)	3.8%	11,688
Text-Impact/ Δ Sales	-0.184***	(0.053)	1.9%	7,481

Note: Each coefficient stems from a separate regression with the R&D wedge as the independent variable and a measure of the impact-value factors as the dependent variable. The R&D wedge is measured as the ratio of patent valuations over R&D expenditure excepts for the third row, where it is measured as changes in sales over R&D expenditure. All variables are in logs. Regressions control for NAICS3 \times Year fixed effects and standard errors are clustered at the NAICS6 level.

C.5 Constructing Estimates for Continuing Firms

I investigate whether entry and exit contributed to the evolution of R&D Allocative Efficiency by constructing a measure thereof solely for continuing firms. I construct the baseline measure in 1975 and annual changes in the Impact of R&D wedges for all subsequent years, which I accumulate over time. For the year 1976, I first filter to firms active in 1975 and 1976. For these firms, I estimate of R&D efficiency for both 1975 and 1976. I then take the ratio to get the rate of change and apply it to my original estimate for 1975. Subsequent years are calculated accordingly.

Let $\hat{\Xi}_t$ be the baseline estimate for the Impact of R&D wedges for year t . Let $\hat{\Xi}_t^{t,t-1}$ be the estimate when using only firms that were active in both t and $t - 1$ with $\hat{\Xi}_{t-1}^{t,t-1}$ being the respective value for $t - 1$. I then calculate the time-series for the Impact of R&D wedges for continuing firms $\hat{\Xi}_t^C$ as

$$\hat{\Xi}_t^C = \begin{cases} \hat{\Xi}_t & \text{if } t = 1975 \\ \hat{\Xi}_{t-1}^C \cdot \left(\frac{\hat{\Xi}_t^{t,t-1}}{\hat{\Xi}_{t-1}^{t,t-1}} \right) & \text{if } t = 1976, \dots, 2014 \end{cases}. \quad (\text{C.6})$$

C.6 Robustness for Aggregate Measures

Table C.6 reports estimates of R&D efficiency for alternative specifications.

Table C.6: 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.
<i>A. Fixed Effects</i>						
Year	-21.3%	-15.4%	-25.2%	-11.7%	5.8%	5.7%
NAICS3 \times Year	-17.9%	-12.0%	-21.7%	-11.0%	5.5%	5.3%
NAICS6 \times Year	-13.5%	-8.1%	-17.6%	-10.3%	5.1%	5.0%
<i>B. Minimum Patents</i>						
50 Patent	-17.9%	-12.0%	-21.7%	-11.0%	5.5%	5.3%
100 Patents	-17.3%	-10.7%	-21.7%	-12.3%	6.2%	6.0%
200 Patents	-16.5%	-9.4%	-21.5%	-13.4%	6.8%	6.7%
<i>C. Time Horizon</i>						
5-Year	-17.9%	-12.0%	-21.7%	-11.0%	5.5%	5.3%
10-Year	-18.1%	-16.3%	-19.8%	-4.2%	1.9%	1.9%
20-Year	-17.2%	-17.9%	-18.6%	-0.9%	0.4%	0.4%

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

C.7 Unobserved Firms and Size Heterogeneity

The main results assume that the sample is broadly representative of firm conducting R&D. However, firms covered in Compustat, and especially those that also patent frequently, tend to be larger than the average firm in the economy, including those that also conduct R&D. In this section, I first introduce a decomposition of aggregate efficiency for subsamples and then investigate size heterogeneity within my sample to understand the implications of adding unobserved firms with, on average, lower R&D expenditure.

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:

$$\Xi_t = \Omega_t^U \cdot \Xi_t^U + \Omega_t^O \cdot \Xi_t^O, \quad (\text{C.7})$$

where

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

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

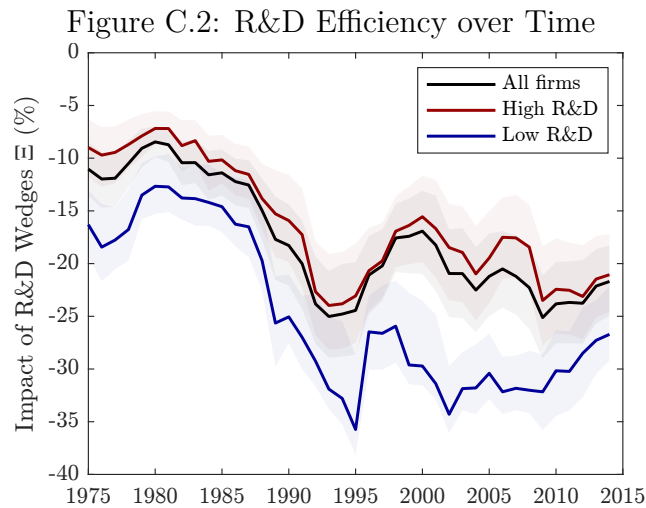
It follows immediately, that Ω_t^X is equivalent to the employment, and thus expenditure share, as long as the distribution of Δ_{it} is equivalent across both groups. Otherwise, and assuming a common level of friction, the weight is larger than the relative employment if the variation in R&D wedges is larger.

We can leverage these insights together with some additional assumptions to understand the implications of only observing a subset of firms. First, it follows immediately that Ξ_t^O is an unbiased estimator for Ξ if we assume that the distributions are comparable across unobserved and observed firms. Second, the bias in the evolution of estimated Ξ_t depends on whether we believe that the unobserved firms are subject to similar trends or not. For example, the estimates are going to be exaggerated if the distribution is constant for unobserved firms or moving in the opposite direction.

In practice, selection into the sample is well understood along one dimen-

sion: Firms in Compustat that patent heavily are larger. Thus, the “unobserved” firms here are mostly smaller firms. It is, thus, tempting to investigate size heterogeneity within the sample to get a sense for whether a heavier skew towards smaller firms would impact any conclusions about the level and evolution of R&D efficiency.

Following this logic, I investigate size heterogeneity by splitting the sample along the median R&D expenditure within a given year. Figure C.2 reports the annual estimates. Two findings emerge immediately. First, R&D Allocative Efficiency, and thus misallocation, is worse among firms with low R&D expenditure. This finding is well aligned with the common perception that such firms are subject to, e.g., tighter credit constraints. Second, R&D Allocative Efficiency declines faster for low R&D expenditure firms. While the gap between high and low R&D firms is modest in the early sample, it opens up significantly post 2000. Thus, if we were to put more emphasis on smaller firms, aggregate R&D Allocative Efficiency would be lower and declining faster, which suggest that the baseline estimates are conservative.



Notes: Figure reports annual estimates for R&D Allocative Efficiency $\Xi_t - 1$. Baseline estimates assumes that R&D wedges are independent from impact-value factors. Adjusted values estimate the adjustment factor over a 10-year rolling window. The shaded area covers the 90% confidence interval calculated using a bootstrapping procedure.

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

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

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

$$\left. \frac{\partial z_{it}}{\partial \ell_{it}} \right|_{\ell_{it}=\ell_{it}^*} \cdot V_{it} = (1 + \Delta_{it}) \cdot W_t \cdot \ell_{it}^\xi. \quad (\text{D.2})$$

Proposition 4 highlights that the main results carry over to this alternative setup, however, the effective scale elasticity is lower. Resultingly, 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:*

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

where $\tilde{\zeta}_{it} = \zeta_{it} / \left(\int_0^1 \omega_{it} \cdot \zeta_{it} di \right)$ 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

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

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

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

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

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 will index a firm by $i \in \mathcal{I}$ and a R&D line by $j \in \mathcal{J}_i$. The production function is given by

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

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

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

We can solve for the R&D wage as

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

The economic growth rate is then

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

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

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

Some algebra confirms the familiar growth rate formula

$$g = \frac{L^\gamma}{A_t^\phi} \cdot \frac{\int_0^1 \zeta_i \cdot \theta_i^{\frac{1}{1-\gamma}} \cdot (1 + \Delta_i)^{-\frac{\gamma}{1-\gamma}} \cdot di}{\left(\int_0^1 \theta_i^{\frac{1}{1-\gamma}} \cdot (1 + \Delta_i)^{-\frac{1}{1-\gamma}} \cdot di \right)^\gamma}. \quad (\text{D.9})$$

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

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

where \bar{L}_t is given exogenously and $\xi/(1-\gamma)$ is the aggregate supply elasticity. Also, let L_t^* be the supply in 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,*

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

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

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

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. \square

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

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

It is straight-forward 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

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

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

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 absence of frictions. The equilibrium wage satisfies

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

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

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

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

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

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 $\phi_t^E \cdot \frac{R \cdot (1-\gamma)}{R-1} \cdot M_t^{\frac{\gamma}{\varphi}}$. The free entry condition is

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

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

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

Note that $\varphi \rightarrow 0$ recovers the baseline model with $M_t = 1$, while $\varphi \rightarrow \infty$ 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 cost are linked to the R&D wage and given by $\phi_t^E \cdot (1-\gamma) \cdot M_t^{\frac{1}{\varphi}} \cdot \frac{W_t}{\gamma}$. The free entry condition is

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

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

$$\frac{M_t}{M_t^*} = \left(\frac{\int_0^1 \omega_{it}(1 + \Delta_{it})^{-\frac{\gamma}{1-\gamma}} di}{\left(\int_0^1 \omega_{it}(1 + \Delta_{it})^{-\frac{1}{1-\gamma}} di \right)^\gamma} \cdot \left(\int_0^1 \omega_{it}(1 + \Delta_{it})^{-\frac{1}{1-\gamma}} di \right)^{\gamma-1} \right)^{\frac{\varphi}{1+\varphi}}$$

s.t. $M_t^* = \left(\frac{L_t}{\phi_t^E} \right)^{\frac{\varphi}{1+\varphi}}.$

(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,*

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

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

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

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. \square

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 straight-forward to show that the impact of frictions is always worse in the economy with free entry holding constant the average level of R&D returns.

E Mechanisms Driving R&D Wedges

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

E.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:

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

Each input is supplied by a single monopolist with constant marginal ψ/A_{jt} . The monopolist is free to chose any price p_{jt} , however, there is a competitive fringe of firm with constant unit costs $\lambda_{jt} \cdot (\psi/A_{jt})$ 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

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

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 :

$$z_i = \varphi_i \cdot \ell_i^\gamma.$$

An invention improves technology in a random product line by λ_i such that $A_{j1} = \lambda_i \cdot A_{j0}$ 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. Resultingly, 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

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

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

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

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

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

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:

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

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

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

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

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

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

$$g^* = \max \int_0^1 z_i \cdot V_i \cdot \zeta_i \cdot di \quad \text{s.t.} \quad L = \int_0^1 \ell_i \cdot di. \quad (\text{E.1})$$

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

E.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

$$\max_{\ell_i} \{V_i \cdot z_i - (1 - \tau_i) \cdot W \cdot \ell_i\}.$$

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

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

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

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

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

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

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

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

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

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

Adjustment costs. Suppose firms face exogenous adjustment costs $\phi \cdot W \cdot$

$(\ell_i - \bar{\ell}_i)^2$. The firm problem is then given by

$$\max_{\ell_i} \{V_i \cdot z_i - W \cdot \ell_i - \phi \cdot W \cdot (\ell_i - \bar{\ell}_i)^2\}.$$

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

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

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

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

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

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

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

Firms internalize the impact 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

$$\gamma \cdot \theta \cdot \ell_i^{\gamma-1} = \left(1 + \frac{1}{N_i} \cdot \xi \right) \cdot W_i \quad (\text{E.4})$$

R&D return is given by $(1/\gamma) \cdot \left(1 + \frac{1}{N_i} \cdot \xi \right)$ with $\Delta_i = \frac{1}{N_i} \cdot \xi$. 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.

E.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 $V_i = P_i \cdot \pi_i / R$, while the public value remains $\lambda_i - 1$. Resultingly, 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 $\pi_i = Y_1 \cdot (1 - (\psi/\psi_i) \cdot \lambda^{-1})$. Resultingly, 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 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 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 straight-forward to see that any frictions in the product market that affect π_i without changing the growth impact of an invention naturally yields variation in ζ_i as well. Firms with artificially low profits under-provide innovation.

Knowledge externalities. More general knowledge externalities can also variation in the impact-value factor. For example, let the growth rate be

$$g = \left(\int_0^1 \varphi_i \cdot z_i \cdot (\zeta_i \cdot V_i) \cdot di \right)^\phi \cdot \int_0^1 z_i \cdot (\zeta_i \cdot V_i) \cdot di, \quad (\text{E.5})$$

where the first term on the right-hand side captures simultaneous knowledge externalities. Here, the marginal benefit to R&D as perceived by the firm for high φ_i firms is generally too low compared to the social planner perspective if $\phi > 0$ and vice versa.

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 ratio as the log of valuations in the alternative portfolio divided by 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 project 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 $(0.06/0.93)^2 \approx 0.4\%$ of the variance of R&D returns. Uncertainty across patents thus, does not appear to have a large contribution of dispersion therein. 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	0.059	0.056	0.058
	[0.051,0.069]	[0.049,0.066]	[0.051,0.069]
Adjustment factor	0.998	0.997	0.998
	[0.997,0.998]	[0.996,0.998]	[0.997,0.999]

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.¹⁹

¹⁹R&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.

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 $\{y_{it}\}$:

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

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

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

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

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

Proof. The results follow immediately from the assumptions. \square

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} the sample moments, then

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

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. \square

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	Δ Sales/ R&D
ρ	0.626*** (0.083)	0.787*** (0.149)
σ_ϵ^2	0.628*** (0.085)	0.404*** (0.092)
σ_μ^2	0.197 (0.140)	-0.086 (0.726)
σ_ν^2	0.007 (0.090)	0.509*** (0.091)
Observations	8,014	7,653
Adjustment factor	0.997	0.810

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 suggest the following procedure to account for additive measurement error:

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

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

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

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

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

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

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	1.477	1.112	1.154	1.025	1.088	0.919	0.861	0.834	0.803	0.524
Std. Err.	(0.12)	(0.14)	(0.09)	(0.08)	(0.09)	(0.10)	(0.11)	(0.08)	(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.