Does Monopsony Matter for Innovation?*

Nils H. Lehr†
International Monetary Fund

April 5, 2026

Abstract

This paper examines how monopsony power in the market for inventors—firms' ability to depress wages by restricting employment—affects U.S. innovation and economic growth. Using an instrumental variable strategy, I estimate firm-level inventor labor supply elasticities and find that firms face less than perfectly elastic inventor labor supply, with larger employers wielding greater monopsony power. I develop and quantify a heterogeneous-firm growth model with size-dependent monopsony power that matches this evidence. Quantitatively, monopsony power reduces annual U.S. economic growth by 0.12 percentage points and welfare by 8% in consumption-equivalent terms through depressed R&D employment and misallocation.

*I thank Stephen J. Terry, Pascual Restrepo, Tarek A. Hassan, David Lagakos, Stefania Garetto, Adam Guren, James Bessen, Titan Alon, Dimitris Papanikolaou, Fil Babalievsky, and many seminar and conference participants for their insightful comments. The views expressed in this paper are my own and do not necessarily reflect those of the IMF, its Executive Board, or its Management. All errors are my own.

†Email: nilshlehr@gmail.com

1 Introduction

The growing dominance of large firms has fueled an active debate over its causes and consequences for the U.S. economy (Grullon et al., 2019; Autor et al., 2020). A central concern is that rising concentration may be closely linked to declining competition and greater market power (Wu, 2018; Philippon, 2019; Meagher, 2020). Concerns about insufficient competition increasingly extend to labor markets, where large firms may have the power to suppress wages. This concern is reflected, for example, in the White House's 2021 executive order on "Promoting Competition in the American Economy" and in the revised 2023 Horizontal Merger Guidelines, both of which drew on the recent monopsony literature (House, 2021; Berger et al., 2023).

Labor market power, commonly referred to as monopsony power, was considered most prevalent for "low-skilled" workers in rural communities with limited employment opportunities; however, recent evidence suggests that it extends to "high-skilled" workers and may be especially important in settings where workers are costly to replace and more specialized (Goolsbee and Syverson, 2023; Seegmiller, 2025; Jäger et al., 2024). One interpretation of these novel findings focuses on a perhaps previously less emphasized source of monopsony power: human capital specificity. For example, registered nurses provide invaluable services to hospitals, but their significant human capital—as indicated by the required graduate degree—is only valuable within the profession. Consequently, hospitals can suppress nurses' compensation in the face of limited competition for their services (Prager and Schmitt, 2021).

Building on these findings, I study the macroeconomic consequences of monopsony power over inventors, a group of workers with an outsized impact on productivity and growth. Monopsony may be especially relevant in this market because inventors possess highly specialized skills, while their output—new inventions—is a key driver of long-run growth and welfare. The compounding nature of innovation also implies that monopsony over inventors may generate dynamic losses that go beyond static inefficiencies. Anecdotal evidence suggests that firms have at times recognized and exploited such power. For example, several large technology firms entered into anti-poaching agreements aimed at suppressing wages for engineers and related workers (Edwards, 2014). Apple, Adobe, Intel, and Google were sanctioned by the Department of Justice in 2010 for such agreements, while Microsoft recently announced that it would no longer enforce non-compete agreements (of Justice, 2010; Reuters, 2022).

This paper concludes that monopsony power in the market for corporate inventors has a sizable negative impact on innovation and economic growth. I find in the data that firms with a large inventor workforce appear to have significant monopsony power, while smaller firms face more competitive conditions. Monopsony power generates a markdown—a wedge between wages and marginal products—that depresses inventor demand and, when monopsony power rises with firm size, disproportionately reduces inventor employment at larger employers, distorting the allocation of inventors across firms. Interpreted through the lens of a quantitative endogenous growth model, my results indicate that monopsony power over inventors reduces long-run economic growth by 0.12 pp, leading to a welfare loss of 8%.

I reach these conclusions in three steps. First, I introduce monopsony power over inventors into a heterogeneous-firm endogenous growth model. Inventors choose their employer based on idiosyncratic preference shocks and wages offered as in Card et al. (2018). As a result, firms face an upward-sloping labor supply curve and can lower wages marginally without losing their entire inventor workforce, as would be the case in the standard competitive model. Similar to Berger et al. (2022), I allow for size-dependent monopsony power such that large employers of corporate inventors may have more power over them. Monopsony power depresses the aggregate demand for corporate inventors, resulting in lower R&D employment and lower economic growth. Size-dependence of monopsony power further induces misallocation as larger firms disproportionately depress their demand for inventors, leading to an additional drag on innovation and economic growth. This paper isolates and quantifies this channel, while recognizing that other forces—such as differences in firms' ability to appropriate rents from innovation—may operate simultaneously with potentially offsetting effects.

In the second step, I present novel evidence on firms' monopsony power over corporate inventors. I estimate the average firm-level elasticity of inventor wages with respect to their employment, i.e., the inverse labor supply elasticity, in a sample of U.S.-listed firms by regressing inventor wage growth on employment growth. I construct inventor employment and their wages by combining firms' financial statements with their patent records. The literature has long recognized the potential identification challenges in this setup (Manning, 2011). Labor supply shocks, such as preference shocks over firms, can lead to a downward bias in the estimated elasticity. In particular, a positive labor supply shock reduces the wage a firm needs to pay in order to maintain a given level of employment, which is informative about the wage level of a firm, but not its local labor supply elasticity. I propose to address this identification challenge by using stock market returns as an instrument for shocks to firms' labor demand as in Seegmiller (2025). The instrument is relevant if returns partly reflect shocks that induce firms to hire more inventors, such as demand shocks for their products. It satisfies the exclusion restriction if there is no link between stock market returns and inventor wages other than their employment. I confirm robustness using firm-level productivity shocks from Imrohoroglu and Tuzel (2014).

My estimates suggest that monopsony power is sizable and size-dependent. I estimate an average inverse labor supply elasticity of 0.38, which implies that a firm would lose about 26% of its inventors if it were to reduce their wages by 10%. For comparison, Seegmiller (2025) estimates an inverse elasticity of 0.84 for high-skilled workers using a similar research design, while Kline et al. (2019) find that workers at small innovative firms capture roughly 30% of patent-generated rents; though the frameworks differ, both imply stronger monopsony power than my baseline, suggesting that my estimates are conservative. Importantly, I find that firms with an above-median R&D workforce face an inverse labor supply elasticity of 0.67 compared to approximately 0 for smaller firms, consistent with evidence that larger employers face less elastic labor supply in the production sector (Berger et al., 2022; Yeh et al., 2022).

In the final step, I calibrate the model by moment matching and use it as a laboratory to study the impact of monopsony in the market for inventors on innovation and economic growth. I discipline monopsony power in the model by targeting the empirical estimates of the inverse labor supply elasticity, computed identically in model-simulated data. The calibration accounts for non-listed R&D firms, which tend to be smaller and thus attenuate the aggregate costs of size-dependent monopsony.

The calibrated model suggests that monopsony power over inventors slows down innovation and economic growth significantly due to a combination of insufficient R&D employment and misallocation of inventors across firms. Inducing firms to act as price takers in the market for inventors increases economic growth by 0.12 pp per year, leading to a sizable welfare improvement of 8%, comparable to, for example, the welfare cost of production-sector monopsony estimated by Berger et al. (2022). The acceleration in economic growth is driven by a 5% rise in R&D employment and a significant improvement in aggregate R&D productivity due to a more productive allocation of inventors. The reallocation channel alone accounts for two-thirds of the growth acceleration, highlighting the importance of the misallocation channel. I show that the baseline results are robust to several extensions including wage discrimination, misalignment between private and social R&D returns, creative destruction, firm entry, and alternative subsidy design. Wage discrimination and misalignment between private and social R&D returns tend to dampen the cost of monopsony power; however, their extent is likely limited in practice. Allowing for entry raises the gains from tackling monopsony power when it is mitigated through targeted subsidies. Finally, the results are robust to financing the targeted R&D subsidies through higher corporate or labor income taxes, but are moderated when financed by lowering untargeted R&D subsidies.

Literature. This paper is closely connected to three strands of the literature. First, I contribute to the literature on monopsony power by providing novel evidence in the market for corporate inventors and linking size-dependent monopsony power to R&D investments and, thus, economic growth. The literature documents that monopsony power is pervasive in the production sector and stronger for larger employers (Azar et al., 2020; Arnold, 2021; Kroft et al., 2021; Lamadon et al., 2022; Yeh et al., 2022). Furthermore, there is growing evidence of monopsony power in labor markets for "high-skilled" workers (Prager and Schmitt, 2021; Goolsbee and Syverson, 2023; Seegmiller, 2025). I complement this literature by documenting monopsony power over an important group of skilled workers: corporate inventors. This group is crucial due to its close link to R&D investments, which in turn are commonly identified as a main driver of long-run productivity growth in the U.S. My model builds on the literature microfounding monopsony power via preferences over employers. An alternative approach focuses on a lack of outside options for workers as a microfoundation of monopsony power (Shi, 2023; Schubert et al., 2025; Bagga, 2023). Relatedly, recent work studies how firms may shape the composition of innovation in order to influence labor mobility and retention in the presence of labor market frictions (Ma et al., 2025).

I complement the literature by introducing preference-based monopsony power into a heterogeneous-firm general equilibrium endogenous growth model and estimating that tackling monopsony power could significantly accelerate U.S. economic growth. Relatedly, Berger et al. (2022) introduce a structural general equilibrium model of the production sector with monopsony power and Barr and Roy (2008) study how monopsony can reduce growth by reducing human capital investment.

Second, I contribute to the literature on resource allocation in the R&D sector. The literature focuses primarily on the misalignment of private and public benefits of R&D, which can also lead to misallocation, rather than misalignment of marginal costs as in my case. Existing work highlights a range of potential mechanisms for such misalignment including knowledge and business stealing externalities, and differences in firms' ability to profit from their inventions or protect their intellectual property. Romer (1990) and Aghion and Howitt (1992) first argued that this misalignment can lead to underinvestment in R&D, while the more recent literature is focused on heterogeneous misalignment across firms that leads to misallocation of R&D resources (Acemoglu et al., 2018; Cavenaile et al., 2021; Mezzanotti, 2021; König et al., 2022; Terry, 2023; Aghion et al., 2025). I complement this literature by instead focusing on a misalignment in the marginal costs perceived by the firm and a planner due to monopsony power. Interestingly, this mechanism leads to the conclusion that large firms might not do enough R&D relative to small firms, while the literature typically finds that they might do too much (Akcigit et al., 2022; Manera, 2022; de Ridder, 2024). These findings suggest that both types of mechanisms might partly offset each other in practice. My paper is also related to the literature on talent (mis-)allocation in the R&D sector (Akcigit et al., 2020; Prato, 2022; Celik, 2023). I complement this literature by focusing on market power as a source of talent misallocation.

Finally, my paper falls within the larger literature on the macroeconomic implications of factor misallocation, which has mostly focused on the production sector. Restuccia and Rogerson (2008) and Hsieh and Klenow (2009) first argued that misallocation of production factors may be significant and could have a large impact on productivity and output. The subsequent literature investigated a range of potential sources of misallocation including financial frictions, government intervention, information frictions, and adjustment costs (Asker et al., 2014; Midrigan and Xu, 2014;

David et al., 2016, 2022). More recently, the literature has considered market power in product and labor markets as a driver of resource misallocation in the production sector that may reduce aggregate productivity and depress output levels (Loecker et al., 2020; Berger et al., 2022). I contribute to this literature by focusing on misallocation in the R&D sector, which may lead to slower innovation and economic growth rather than lower output. This focus coincides with Lehr (2025), who studies misallocation in the R&D sector in general. This paper is complementary as it studies and provides evidence for a particular mechanism of misallocation: monopsony power.

2 A Growth Model with Inventor Monopsony

This section introduces preference-based monopsony power into a general equilibrium growth model in the tradition of Romer (1990) to investigate the potential impact of monopsony power on innovation and economic growth. I calibrate the model in Section 4, targeting the evidence on labor market power presented in Section 3.

2.1 Model Description and Decentralized Equilibrium

Time is discrete and indexed by t. The economy is populated by a representative household that supplies production and research labor and allocates its income between consumption and savings. The final good is produced by a representative firm from production labor and intermediate inputs. The latter are produced by a mass N of profit-maximizing research firms that own their exclusive production rights. Research firms hire researchers and materials to invent new types of intermediate goods. Table 1 presents an overview of the main model equations.

The **representative household** with a population of $\bar{L} = 1$ has KPR preferences over consumption C_t and employment disutility L_t, which aggregates labor supply to R&D $L_{R,t}$ and production $L_{P,t}$ (King et al., 1988).

$$U(C_t, L_{R,t}, L_{P,t}) = \frac{(C_t\, V(L_t))^{1-\varsigma} - 1}{1-\varsigma}$$
$$\text{with } \ln V(L_t) = -\frac{\varphi}{1+\varphi}\, L_t^{1+1/\varphi} \text{ and } L_t = \left(\alpha_P^{-\frac{1}{\epsilon}}\, L_{P,t}^{1+1/\epsilon} + \alpha_R^{-\frac{1}{\epsilon}}\, L_{R,t}^{1+1/\epsilon}\right)^{\frac{\epsilon}{1+\epsilon}} \tag{1}$$

Table 1: The Economic Environment

(1)	$U = \frac{(C_t\, V(L_t))^{1-\varsigma}-1}{1-\varsigma}$ with $\ln V = -\frac{\varphi}{1+\varphi} L_t^{1+1/\varphi}$ with $L_t = \left(\alpha_P^{-\frac{1}{\epsilon}}\, L_{P,t}^{1+1/\epsilon} + \alpha_R^{-\frac{1}{\epsilon}}\, L_{R,t}^{1+1/\epsilon}\right)^{\frac{\epsilon}{1+\epsilon}}$	Household utility (KPR)
(2)	$L_{R,t} = \left(\underline{\ell}^{\xi} + \frac{1}{1+\xi}\right)^{-1} \int_N \ell_{kt} \left(\underline{\ell}^{\xi} + \frac{1}{1+\xi}\left(\frac{\ell_{kt} N}{L_{R,t}}\right)^{\xi}\right) dk$	R&D labor aggregator
(4)	$Y_t = L_{P,t}^{1-\alpha} \int_{\mathcal{Q}_t} z_{jt}^{1-\alpha} x_{jt}^{\alpha}\, dj$	Final good production
(7)	$Q_{kt}^N = Z_t A_k \Big(\alpha_L^{1/\theta} \ell_{kt}^{\frac{\theta-1}{\theta}} + (1-\alpha_L)^{1/\theta}\left(\frac{M_{kt}}{Z_t}\right)^{\frac{\theta-1}{\theta}}\Big)^{\gamma\frac{\theta}{\theta-1}}$	R&D production
(8)	$Z_{kt+1} = Q_{kt}^N\, z_{kt} + Z_{kt}$ and $\mathcal{Q}_{kt+1} = \mathcal{Q}_{kt} + \mathcal{Q}_{kt}^N$ $Z_t = \int_N Z_{kt} dk$ and $\mathcal{Q}_{\sqcup} = \int_N \mathcal{Q}_{kt} dk$	Blueprint accumulation and aggregation
(9)	$\ln z_{kt} = (1-\rho)\mu_z + \rho \ln z_{kt-1} + \sigma_z \varepsilon_{kt}^z$	Idea quality (stochastic)
(13)	$Y_t = C_t + \psi \int_{Q_t} x_{jt} dj + \psi_M \int_N M_{kt}\, dk$	Resource constraint

Notes: N is the fixed mass of R&D firms. All labor supply parameters are weakly positive, i.e., $\xi, \underline{\ell}, \varphi, \epsilon \geq 0$.

The elasticities satisfy $\epsilon, \varphi \geq 0$. Sectoral labor supply is independent if $\epsilon = \varphi$.

Labor supply for researchers is potentially differentiated, capturing the idea that firms are imperfect substitutes from the perspective of workers due to, e.g., amenities, management styles, company cultures or visions. I denote the supply of researchers for firm $k \in N$ by ℓ_{kt} and assume that labor supply $L_{R,t}$ satisfies

$$L_{R,t} = \left(\underline{\ell}^{\xi} + \frac{1}{1+\xi}\right)^{-1} \left(\int_N \int_{\ell_{kt}} \left(\underline{\ell}^{\xi} + \left(\frac{\ell\ N}{L_{R,t}}\right)^{\xi}\right) d\ell\ dk\right) \text{ with } \xi \geq 0.$$

The aggregator integrates over firms as well as over workers within firms, where the marginal disutility is increasing in the number of workers hired as long as $\xi > 0$. This formulation captures the idea that workers have idiosyncratic preferences over employers such that firms hiring more researchers face ever less enthusiastic workers at the margin. For $\underline{\ell} = 0$, the aggregator is of the CES-type and, thus, has a constant elasticity with respect to labor supply for an individual firm. For $\underline{\ell} > 0$, the aggregator is non-homothetic with a rising elasticity with respect to ℓ_{kt} such that larger employers face more inelastic R&D labor supply. Evaluating the integral, we have

$$L_{R,t} = \left(\underline{\ell}^{\xi} + \frac{1}{1+\xi}\right)^{-1} \left(\int_N \ell_{kt} \left(\underline{\ell}^{\xi} + \frac{1}{1+\xi} \left(\frac{\ell_{kt}\, N}{L_{R,t}}\right)^{\xi}\right) \, dk\right). \tag{2}$$

The household's budget constraint balances expenditure on consumption C_t and savings B_{t+1} in a zero-net-supply riskless bond with income from labor supply, from firm ownership, and interest from prior savings. Labor income is derived from production sector wages $W_{P,t}$ and firm-specific wages from R&D work $W_{R,kt}$, and taxed at rate τ_L. The household further receives dividends from profits Π_t net of corporate taxes τ_π, government transfers T_t and interest R_t on bond holdings B_t.

$$C_t + B_{t+1} = (1-\tau_L)\left(W_{P,t}\, L_{P,t} + \int_N W_{R,kt}\, \ell_{kt}\, dk\right) + (1-\tau_\pi)\Pi_t + T_t + R_t\, B_t. \tag{3}$$

A representative firm hires production labor $L_{P,t}$ and buys intermediate inputs $\{x_{jt}\}_{j\in\mathcal{Q}_t}$ to produce **final output** Y_t with production function

$$Y_t = L_{P,t}^{1-\alpha} \int_{\mathcal{Q}_t} z_{jt}^{1-\alpha}\, x_{jt}^{\alpha}\, dj, \tag{4}$$

where z_{jt} is an intermediate good-specific demand-shifter. The firm takes the wage $W_{P,t}$ and intermediate input prices $\{p_{jt}\}_{j\in\mathcal{Q}_t}$ as given and maximizes its profits:

$$\max_{L_{P,t},\{x_{jt}\}_{j\in\mathcal{Q}_t}} \left\{Y_t - W_{P,t}\, L_{P,t} - \int_{\mathcal{Q}_t} p_{jt}\, x_{jt}\, dj\right\}. \tag{5}$$

The production rights to **intermediate goods** are owned by a fixed mass N of innovative firms, which can produce them at constant marginal cost ψ and take into account the final goods sector's demand curve when maximizing profits:

$$\pi_{jt} = \max_{x_{jt}} \left\{p_{jt}(x_{jt})\, x_{jt} - \psi\, x_{jt}\right\}. \tag{6}$$

To add to their portfolio of intermediate goods, **innovative firms** hire inventors

ℓ_{kt} and materials M_{kt} to produce new blueprints $\mathcal{Q}_{kt}^N$ with production function

$$Q_{kt}^N = Z_t\ A_k\ \left(\alpha_L^{\frac{1}{\theta}}\ \ell_{kt}^{\frac{\theta-1}{\theta}} + (1-\alpha_L)^{\frac{1}{\theta}} \left(\frac{M_{kt}}{Z_t}\right)^{\frac{\theta-1}{\theta}}\right)^{\gamma\ \frac{\theta}{\theta-1}}. \tag{7}$$

The parameter A_k is a time-invariant firm characteristic that captures permanent differences in R&D productivity across firms. In the quantitative analysis, I use A_k to distinguish between listed and non-listed firms.

The mass of (quality-adjusted) blueprints for a firm evolves according to

$$Z_{kt+1} = Q_{kt}^N\ z_{kt} + Z_{kt} \quad \text{and} \quad \mathcal{Q}_{kt+1} = Q_{kt}^N + \mathcal{Q}_{kt}. \tag{8}$$

I denote the respective aggregate levels of (quality-adjusted) blueprints by dropping the firm subscript, i.e., $Z_t = \int_N Z_{kt}\ dk$ and $\mathcal{Q}_t = \int_N \mathcal{Q}_{kt}\ dk$.

The firm-specific demand-shifter is stochastic and follows a log-normal AR(1):

$$\ln z_{kt} = (1-\rho)\ \mu_z + \rho\ \ln z_{kt-1} + \sigma_z\ \varepsilon_{kt}^z \quad \text{with} \quad \varepsilon_{kt}^z \sim N(0,1). \tag{9}$$

Innovative firms choose materials M_{kt} at constant unit costs ψ_M and hiring ℓ_{kt}, internalizing their upward sloping labor supply curve and hence the impact of their hiring on the wage $W_{R,kt}$, to maximize their value:

$$\begin{aligned} V_{kt}(Z_{kt}, z_{kt}) = \max_{\ell_{kt}, M_{kt}} \Bigg\{ \int_{Q_{kt}} \pi_{jt}\ dj - (1-\tau_{kt})(W_{R,kt}\ \ell_{kt} + \psi_M\ M_{kt}) \\ + R_{t+1}^{-1}\ \mathbb{E}_t[V_{kt+1}(Z_{kt+1}, z_{kt+1}) | Z_{kt}, z_{kt}] \Bigg\} \end{aligned} \tag{10}$$

Profits from the corporate sector are transferred to the household as dividends:

$$\Pi_t = Y_t - \left(W_{P,t}\ L_{P,t} + \psi\ \int_{\mathcal{Q}_t} x_{jt}\ dj\right) - \int_N (1-\tau_{kt})\ \left(W_{R,kt}\ \ell_{kt} + \psi_M\ M_{kt}\right) dk \tag{11}$$

The **government** budget constraint determines transfers T_t as the gap between tax income and tax expenditure:

$$\begin{aligned} T_t =& \tau_\pi \int_N \left(\int_{Q_{kt}} \pi_{jt} \, dj - (1-\tau_{kt})(W_{R,kt} \, \ell_{kt} + \psi_M \, M_{kt}) \right) \, dk \\ &+ \tau_L \left(W_{P,t} \, L_{P,t} + \int_N W_{R,kt} \, \ell_{kt} \, dk \right) - \int_N \tau_{kt} \left(W_{R,kt} \, \ell_{kt} + \psi_M \, M_{kt} \right) \, dk \end{aligned} \tag{12}$$

Market clearing in the product market requires that

$$Y_t = C_t + \psi \int_{\mathcal{Q}_t} x_{jt} \, dj + \psi_M \int_N M_{kt} \, dk. \tag{13}$$

The decentralized equilibrium definition is standard and formalized in Definition 1. I focus on a Balanced Growth Path (BGP) equilibrium.

Definition 1 (Decentralized Balanced Growth Path Equilibrium)**.** *A sequence of quantities and prices such that (a) households maximize utility* (1) *subject to budget constraint* (3)*, (b) firms maximize profits by solving* (5) *and* (6)*, and firm value by solving* (10)*, (c) markets clear* (13)*, (d) quantities grow at a constant rate* $g = Z_{t+1}/Z_t - 1$*, except for labor supply, which remains constant at the aggregate level.*

2.2 Planner's Problem

To study optimal policy, it is useful to introduce the planner problem. The planner chooses quantities to maximize expected utility:

$$\begin{aligned} \max_{L_{P,t}, \{M_{kt}, \ell_{kt}\}, \{x_{jt}\}} \quad \sum_{t=0}^{\infty} \beta^t \, U(C_t, L_{P,t}, L_{R,t}) \text{ s.t. } (1), (2), (4), (7), (9), (13), \text{ and} \\ Z_{t+1}/Z_t - 1 = \int_N z_{kt} \, (Q_{kt}^N / Z_t) \, dk \end{aligned} \tag{14}$$

Definition 2 (Planner BGP Allocation)**.** *A sequence of quantities that solve the planner problem* (14) *such that productivity* Z_t *grows at a constant rate* g*.*

2.3 Monopsony in R&D and Growth

To highlight the direct impact of monopsony power on economic growth, I make two simplifying assumptions in this section that focus the attention on the main mechanism studied in this paper:

1. Labor is the only input to R&D, i.e., $\alpha_L = 1$.

2. There is no taxation in the decentralized equilibrium, i.e., $\tau_L = \tau_\pi = \tau_{kt} = 0$.

For aggregate implications, I further assume $\varphi = \epsilon > 0$, which eliminates cross-sectoral labor supply substitution and implies a positive aggregate labor supply elasticity. I relax these assumptions in the quantitative exploration. Derivations and proofs are provided in Appendix A.

The characteristic feature of monopsony power is that firms' wages respond to their demand for labor and that firms take this effect into account. Lemma 1 highlights the first property in the model by showing that firms' R&D wages respond to their demand for R&D workers. Furthermore, this sensitivity is stronger for firms that are already larger when $\underline{\ell} > 0$, i.e., in the case of log-concave labor supply. The response of R&D wages and employment to small changes in the firm-specific R&D subsidy τ_{kt} or R&D productivity z_{kt} identifies this elasticity.

Lemma 1 (Wages in the R&D sector)**.** *The elasticity of the firm's R&D wage with respect to a change in its R&D employment is given by*

$$\epsilon_{kt} \equiv \frac{\partial \ln W_{R,kt}}{\partial \ln \ell_{kt}} = \xi \, \frac{(\ell_{kt} \; N/L_{R,t})^{\xi}}{\underline{\ell}^{\xi} + (\ell_{kt} \; N/L_{R,t})^{\xi}}. \tag{15}$$

Note that $\epsilon_{kt} \in [0, \xi]$, with the lower bound attained by infinitesimally small firms and the upper bound in the limit as firm size grows large. As a result, firms' demand for R&D workers, and thus R&D employment, is less sensitive to R&D productivity shocks or targeted subsidies at larger firms, as discussed in Proposition 1. For $\underline{\ell}, \xi > 0$, this further implies that the responsiveness of R&D employment to R&D productivity is decreasing in firm size. When monopsony power is log-linear ($\underline{\ell} = 0$), the sensitivity of R&D employment to productivity or subsidies is identical in the decentralized equilibrium and planner allocation. Under log-concave labor supply, however, the decentralized equilibrium exhibits lower sensitivity. The curvature of the labor supply schedule causes the markdown to vary with the level of hiring: as firms expand in response to a productivity shock, the markdown rises, dampening the employment response. Since the planner faces no markdown, this attenuating force is absent from the first-best allocation.

Proposition 1 (Employment in the R&D sector)**.** *Compared to a model without monopsony power, i.e.* $\xi = 0$*, a model with* $\xi > 0$ *features a lower sensitivity of R&D employment with respect to R&D productivity shocks. This sensitivity is unambiguously declining with R&D employment if* $\underline{\ell} > 0$*. These properties are preserved if R&D uses some material, i.e.,* $\alpha_L \in (0,1)$*, under limited substitutability:* $\theta \in [0,1]$*. Relative to a planner allocation, firms' R&D employment is equally sensitive to productivity shocks in the decentralized equilibrium for* $\underline{\ell} = 0$ *and becomes less sensitive for a given initial* ℓ_{kt} *as inventor employment increases if* $\underline{\ell} > 0$ *.*

Proposition 2 highlights two effects of monopsony power on equilibrium R&D employment. First, monopsony power reduces aggregate R&D effort relative to the competitive benchmark, provided the aggregate supply of inventors is not perfectly inelastic. Since the decentralized equilibrium already features inefficiently low R&D, owing to the monopoly distortion in the product market and intertemporal knowledge externalities, monopsony power widens this gap further. Second, when labor supply is log-concave, the allocation of R&D workers across firms is skewed toward small firms: large firms exploit their greater monopsony power by disproportionately reducing demand. Thus, monopsony distorts not only the aggregate level of R&D but also its allocation across firms, lowering growth beyond what the level effect alone would imply. I refer to this second channel as misallocation.

Proposition 2 (The Impact of Monopsony Power)**.** *Denote quantities in the* D*ecentralized equilibrium and* P*lanner allocation by superscripts. Under homogeneous monopsony power,* $\xi > 0$ *and* $\underline{\ell} = 0$*, the decentralized equilibrium features lower effective R&D employment than the planner allocation* ($L_R^D < L_R^P$) *and this gap is larger compared to a model without monopsony power. However, the relative R&D employment across firms coincides in the decentralized equilibrium and planner allocation, i.e., for any firms* k *and* k' $\frac{\ell_{kt}^D}{\ell_{k't}^D} = \frac{\ell_{kt}^P}{\ell_{k't}^P}$*. In contrast, under heterogeneous monopsony power,* $\underline{\ell} > 0$*, R&D employment increases faster with R&D productivity in the planner allocation than the decentralized equilibrium, i.e., for any* k, k' *such that* $\ell_{kt}^D > \ell_{k't}^D$ $\frac{\ell_{kt}^D}{\ell_{k't}^D} < \frac{\ell_{kt}^P}{\ell_{k't}^P}$*.*

What are the policy implications? In the case of common monopsony power, the planner allocation can be achieved by a general subsidy to firms' R&D activity or, alternatively, by subsidizing R&D workers. Such a subsidy becomes ever more important the more elastic the supply of R&D workers in the economy. In the case of heterogeneous monopsony power, general R&D subsidies are insufficient and targeted interventions become necessary. The optimal (marginal) R&D subsidy rate is larger for firms hiring more inventors.

Corollary 1 (Optimal Policy)**.** *Under homogeneous monopsony power, the planner allocation in the R&D sector can be achieved through a uniform R&D subsidy. Under heterogeneous monopsony power, targeted subsidies are necessary, with effective subsidy rates increasing in firm size.*

Optimal policy under size-dependent monopsony power calls for large employers of inventors to expand their R&D workforce, implying that these firms underinvest in R&D in the decentralized equilibrium. This contrasts with the recent literature arguing that large firms overinvest in R&D relative to small firms (de Ridder, 2024; Aghion et al., 2025). The two results are reconciled by the source of heterogeneity in innovation activity. In my model, heterogeneity is driven by productivity differences that a planner would also value when allocating R&D workers. In the aforementioned papers, heterogeneity arises from differences in the ability to charge markups, which inflates the private return to R&D at large firms relative to the social return. In practice, both forces likely operate simultaneously with potentially offsetting effects. This paper isolates and quantifies the monopsony channel.

Finally, size-dependent monopsony power has a directly observable implication that does not require structural estimation of the labor supply elasticity. As shown in Lemma 2, the R&D return — the ratio of R&D output to its costs — is increasing in R&D employment if and only if monopsony power varies with firm size. The intuition is straightforward: firms with greater monopsony power suppress wages below marginal products, raising output relative to costs. Under price-taking, by contrast, all firms equalize wages to marginal products and thus earn identical R&D returns. A positive correlation between R&D returns and R&D employment in the data is therefore a strong signal of size-dependent monopsony power.

Lemma 2 (R&D Returns). *Let the R&D return of a firm be the ratio of the expected value created from innovation to the total cost. Its equilibrium value is given by*

$$\textit{R\&D Return}_{kt} \equiv \frac{Q_{kt}^{N}\, z_{kt}\, V_{\pi,t+1}/R_{t+1}}{W_{R,kt}\, \ell_{kt}} = \frac{1}{\gamma}\left(1+\frac{1}{\epsilon_{kt}}\right),$$

where $V_{\pi,t+1}$ is the present discounted value of profits from a blueprint of quality 1. The R&D return is constant across firms for $\underline{\ell}=0$, and increasing in ℓ_{kt} for $\underline{\ell}>0$.

3 Evidence

This section presents evidence on inventor monopsony power, which disciplines the parameters $\{\xi, \underline{\ell}\}$ in the model calibration (Section 4). The estimates serve as directly targeted moments in the structural calibration, where I replicate the empirical regressions identically on model-simulated data.

3.1 Data

My data combine information on the financial performance and innovation activity of US listed firms. Firm-level data are essential because the firm-level labor supply elasticity, which governs monopsony power, differs from the market-level elasticity. Intuitively, individual firms can expand R&D employment by hiring from competitors or from non-employment, but if all firms expand simultaneously, only the latter option is feasible, yielding a potentially different elasticity.

I obtain financial data from WRDS Compustat, which collects data based on mandatory filings by listed firms. Variables of interest include R&D expenditure (`xrd`), employment (`emp`), and stock market returns. I combine these data with information on firms' patenting activity using the crosswalk between firms and patents developed in Kogan et al. (2017). The patent data from Kogan et al. (2017) and the USPTO's PatentsView database include information on firms' granted patents, including application date and technology classification, and the inventors listed on each patent. Patents provide a direct, though incomplete, measure of R&D output (Cohen et al., 2000; Mezzanotti and Simcoe, 2025).

The primary variables of interest are inventor employment and wages. I measure inventor employment using patent records, linking inventors across patents using the USPTO's inventor disambiguation and assigning them to firms based on whether they appear on a patent application filed by the firm in a given year that is ultimately granted. I aggregate to the firm level by summing over all inventors. This measure may be incomplete, e.g., because not all active researchers at a firm are listed on a patent within a given period; however, it provides a readily available measure of innovators contributing to firms' patent output. I similarly construct a measure of inventor productivity based on lifetime patenting and calculate firm-level averages. See Appendix B.1 for additional details.

I measure inventor wages as R&D expenditure divided by inventor employment. This measure is subject to three potential issues. First, not all R&D expenditure is on labor inputs, as R&D also requires materials and equipment, though NSF statistics suggest a labor share of 79%.[1] Second, my measure of inventors may be incomplete as discussed above, adding further measurement error. Third, measuring inventor employment from patents implicitly assumes that R&D projects result in a patent application within a given year. Projects with longer time horizons could generate a misalignment between R&D expenditure and recorded patents. My analysis must therefore account for potential measurement error in R&D wages.

As discussed in the previous section, the R&D return can be informative about monopsony power. I measure it as the ratio of patent valuations to the previous year's R&D expenditure at a 5-year horizon:

$$\text{R\&D Return}_{it} \equiv \frac{\sum_{s=0}^{4} \text{Patent Valuations}_{it+s}}{\sum_{s=0}^{4} \text{R\&D Expenditure}_{it-1+s}}. \tag{16}$$

I restrict the sample to 1975–2019 and drop firms with consistently low R&D expenditure (less than \$2.5m in 2012 USD per year), low patenting (fewer than 2.5 patents per year), or fewer than 5 sample-years. The final sample contains approximately 17,200 observations for 810 firms and covers more than 85% of R&D expenditure in Compustat and 80% of patent valuations in Kogan et al. (2017), and more

[1] I calculate this figure using Table 10 in the NSF's Business Enterprise Research and Development Survey statistics for 2019. I exclude "other" R&D expenditure and "other purchased services" and add 1/3 of depreciation expenditure to capture the cost of capital, assuming a 5% interest rate and 15% depreciation rate. Total R&D labor expenditure includes "salaries, wages, and fringe benefits," "stock-based compensation," and "temporary staffing." See Supplemental Appendix G.3.

than 40% of R&D recorded in BEA accounts. On average, firms employ about 130 inventors and produce 85 patents per year. See Appendix B.1 for further details.

3.2 Estimation Approach

The inverse labor supply elasticity for inventors is the key monopsony parameter in the model and the primary object of estimation. It can be estimated by regressing log changes in inventor wages on log changes in inventor employment, as shown in equation (17) (Manning, 2003). Estimating in differences eliminates time-invariant firm heterogeneity. In my baseline, I select $t-2$ as the reference period and investigate changes up to $t+3$. The estimate should therefore be interpreted as a medium-run labor supply elasticity rather than as a contemporaneous impact elasticity.

$$\Delta \ln \text{Inventor Wage}_{kt} = \bar{\epsilon} \times \Delta \ln \text{Inventors}_{kt} + \alpha_{j(k)\times t} + \delta X_{kt} + \varepsilon_{kt}, \tag{17}$$

where $\alpha_{j(k)\times t}$ are industry-year fixed effects and X_{kt} control variables. OLS estimation may be biased if labor supply shocks simultaneously affect wages and employment. For example, if workers become more attracted to a firm, it can lower wages while hiring more workers. Identifying a supply elasticity thus requires demand shocks.

To address this concern, I estimate (17) by 2SLS, using stock market returns in t as the excluded instrument for changes in inventor employment, following Seegmiller (2025)'s identification strategy for the overall labor supply elasticity. The instrument is relevant if stock market returns reflect changes in firm productivity or demand that incentivize the firm to expand production, which in turn increases the market size for new products and raises the incentive to expand R&D. The exclusion restriction requires that stock market returns do not affect inventor wage growth other than through their impact on inventor employment growth. In the baseline specification, I control for prior wage and employment growth as in Seegmiller (2025).

I allow the inverse labor supply elasticity to vary with firm size by interacting inventor employment growth with an indicator for prior above-median inventor employment. Under size-dependent monopsony power, the interaction coefficient should be positive. I follow a similar approach using above-median R&D returns, which are also linked to the labor supply elasticity as discussed in the previous section.

$$
\begin{aligned}
\Delta \ln \text{Inv. Wage}_{kt} = \quad & \epsilon_l \times \Delta \ln \text{Inv.}_{kt} \\
& + (\epsilon_h - \epsilon_l) \times \Delta \ln \text{Inv.}_{kt} \times \{\text{Above Median Inventors}\}_{kt} \\
& + \beta \{\text{Above Median Inventors}\}_{kt} + \alpha_{j(k)\times t} + \delta X_{kt} + \varepsilon_{kt}
\end{aligned} \tag{18}
$$

The exclusion restriction for the interaction term requires that stock returns do not differentially affect R&D wage growth across firm size groups other than through their impact on inventor employment growth.

There are several potential identification challenges. First, stock market returns may partly reflect labor supply shocks that also raise firm value. The estimated elasticity would then be downward biased, as supply shocks reduce wages while raising employment. Importantly, the relevant shocks must apply to the market for inventors specifically; a shock that lowers required wages for non-inventor workers without affecting inventor wages does not violate the exclusion restriction. More generally, the identifying variation in the empirical analysis comes from firms' dynamic responses to shocks rather than from permanent cross-sectional differences across employers. One concrete version is that positive productivity shocks raise both stock returns and a firm's attractiveness as an employer—for example, through improved prestige or career prospects—shifting inventor supply outward.[2] If this channel is stronger at larger firms, it would attenuate both the average inverse elasticity and the size gradient, reinforcing the conservative interpretation of the baseline estimates.

Second, equity-linked compensation could induce a direct correlation between stock returns and inventor wages unrelated to inventor employment, a concern that is particularly relevant for high-skilled workers who increasingly receive equity-linked pay (Eisfeldt et al., 2023). While stock-based pay accounts for roughly 12% of R&D labor compensation (NSF BERDS, 2019), the concern is substantive because equity grants may co-move with firm performance. The long differencing window ($t-2$ to $t+3$) mitigates this, as one-off grants around the time of the return shock wash out of the wage growth measure. I address this threat in detail in Supplemental Appendix E.2, where I extend the quantitative model to incorporate stock-based compensa-

[2] For example, larger employers may rely more on their reputation to attract and retain inventors, exposing them more to such shocks (see also Kim, 2026).

tion and show that the estimates are robust. Finally, R&D wages include non-labor expenditure, introducing potential measurement error that I address in the model calibration (see also Supplemental Appendix E.1).

3.3 Results

My estimation results, as reported in Table 2, reveal three main findings.[3] First, the estimated inverse labor supply elasticity is positive and significant: a 10% increase in inventor employment requires 3.8% higher wages, implying that inventors receive roughly $1/(1+\hat{\epsilon}) \approx 72\%$ of their marginal product. For comparison, Seegmiller (2025) estimates a larger inverse elasticity of 0.84 for high-skilled workers using LEHD data, while Kline et al. (2019) estimate that workers at small, innovative firms receive closer to 30% of firm-specific rents, though the two objects are not directly comparable.[4] Second, the heterogeneity analysis across firm size suggests that this effect is driven entirely by firms with a large inventor workforce. A firm with above-median inventor employment faces an inverse elasticity of 0.67, while the elasticity for smaller firms is not statistically distinguishable from zero.[5] Third, firms with above-median R&D returns face more inelastic inventor supply, as predicted by the model. Quantitatively, these estimates are closely aligned with the results based on inventor employment; remaining differences likely reflect the fact that R&D returns capture frictions beyond labor market power. In summary, smaller firms face approximately competitive labor markets for inventors, while larger firms exercise meaningful monopsony power.

I consider several robustness exercises reported in Appendix Table G.1. First, observed wage growth could reflect changing inventor quality composition; however, the table confirms that the estimates are robust to controlling for inventor quality. Second, the estimates are robust to using TFP shocks from Imrohoroglu and Tuzel (2014) instead of stock market returns as an instrument for inventor demand; however, the first stage is weaker and, as a result, the standard errors larger. Third,

[3] Supplemental Appendix G.1 reports the first-stage results. Stock market returns are associated with a significant subsequent expansion of inventor employment. The baseline estimate ($t = 3$) suggests that a 10% increase in firm valuation is associated with a 2% expansion of R&D employment. The estimate is highly statistically significant with an F-statistic comfortably above 10.

[4] I revisit the comparison through the lens of the quantitative model in Supplemental Appendix F.

[5] These estimates do not imply that wage levels are higher at smaller R&D employers, as marginal products may differ substantially.

Table 2: Inverse R&D Labor Supply Elasticity Estimates

	(1)	(2)	(3)
	R&D Wage Growth		
R&D Employment Growth	0.382**	0.003	0.098
	(0.162)	(0.129)	(0.194)
— × Above Median R&D Employment		0.666***	
		(0.222)	
— × Above Median R&D Return			0.469**
			(0.190)
First stage F stat. (Main)	68	40	28
First stage F stat. (Inter.)	–	27	50
Observations	12,729	12,729	12,729

Note: R&D employment and wage growth are log differences between $t-2$ and $t+3$. R&D employment growth is instrumented with stock market returns in t. Regressions control for inventor wage and employment growth from $t-2$ to $t-1$ and NAICS3 × year fixed effects. F statistics reported are based on Sanderson and Windmeijer (2015). Standard errors are clustered at the NAICS6 level.

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

dropping the control variables reduces the estimates only marginally. Fourth, capping or extending the time window by a year does not materially alter the estimates. Lastly, the estimates are essentially unaffected when alternatively measuring inventor employment via (1) a full-time equivalent measure, (2) only US inventors, or (3) only verified inventors as identified by Kaltenberg et al. (2021), and are also robust to replacing industry-year fixed effects with an alternative technology-based firm classification derived from patent CPC codes.

In summary, the evidence suggests that large firms exercise significant monopsony power in the market for inventors, while smaller firms face approximately competitive conditions. The estimated elasticities serve as the key calibration targets for the quantitative model in the next section.

4 Quantification

The evidence presented in the previous section suggests a potentially meaningful role for monopsony power in the market for inventors. This section quantifies its impact on innovation and economic growth by calibrating the model to match the evidence. I solve the model numerically by discretizing the ergodic R&D productivity distribution over 100 grid points. For simulated micro-moments, I discretize the AR(1) process over 50 grid points using the Rouwenhorst method and simulate a listed firm over 100,000 periods with 10,000 periods of burn in, where the firms' policy function relies on aggregates calculated over the larger grid. I then calculate all moments identically to their empirical counterparts, e.g., I define R&D wages as the ratio of R&D expenditure to inventors when replicating the empirical evidence from the previous section in the model. See Appendix A.4 for additional details.

4.1 Calibration

I calibrate the model with a combination of external parameters from the literature and moments matching. An important consideration in this process is that my data only covers listed firms, which tend to be larger and cover about 60% of U.S. R&D expenditure.[6] Since monopsony power increases with firm size in the model, calibrating it to match listed firms only could artificially inflate the cost of monopsony. I thus assume that only a fraction ζ of firms is listed and differentiate non-listed firms via the productivity shifter A_{nl}, setting $A_k = 1$ for listed and $A_k = A_{nl}$ for non-listed firms. For all empirical moments, I match using simulated data for listed firms only; all parameters other than A_k are common across firm types.

For the external calibration, I pick a standard value for discount factor $\beta = 0.97$ and intertemporal elasticity of substitution $\varsigma = 2$, which together with a targeted growth rate of 1.5% imply an annual risk-free interest rate of 6%. I calibrate the demand parameter α to achieve a markup of $1/\alpha - 1 = 25\%$ and set $\psi = \alpha$ to normalize the price of intermediates to 1. Following Chetty et al. (2012), I set the aggregate labor supply elasticity to $\varphi = 0.5$, such that an exogenous 1% rise in wages raises aggregate employment by 0.5%. For the elasticity of substitution across sectors, I rely on the estimate in Ekerdt and Wu (2025) and set $\epsilon = 0.3$, implying limited substitutability. I set the elasticity of substitution between materials and labor in R&D to $\theta = 0.5$, adopting estimates for the production sector in Oberfield and Raval

[6]Total R&D expenditure in the Compustat sample in 2019 is 340 billion USD, while the NSF reports a total expenditure on R&D for all firms of 564 billion USD, implying that listed firms account for 60% of R&D expenditure. For 2000, this share is slightly higher at 72%.

(2021) since direct estimates for the R&D sector are not available.[7] I set the share of listed firms ζ to 5% based on the firms in my sample relative to the NSF R&D surveys.[8] Finally, I set labor taxes to 30%, the corporate tax rate to 15%, and the R&D subsidy rate to 7%, broadly in line with averages for the estimation period.[9]

The remaining parameters are jointly calibrated to match a set of macro and micro moments. At the macro level, I target an annual growth rate of 1.5% and a relative size of listed to non-listed firms of 35 in terms of R&D expenditure, in line with the relative size of firms in my sample and in the NSF aggregate statistics. These moments are particularly informative about the R&D productivity intercept and the relative productivity of non-listed firms $\{\mu, A_{nl}\}$. I target a total labor supply of 1/3, equivalent to 8 hours per day, of which 14.6% work in R&D as in Acemoglu et al. (2018), to pin down the labor disutility shifters $\{\alpha_P, \alpha_R\}$. I target a labor share of 79% in R&D to pin down the labor-intensity of the R&D production function α_L, while ψ_M is left free for computational convenience.[10] At the micro-level, I target the standard deviation of R&D expenditure growth and the autocorrelation of R&D expenditure to pin down the demand process parameters $\{\sigma, \rho\}$. Following Acemoglu et al. (2018), I target a unit cost elasticity of R&D expenditure of -1, computed from firms' partial equilibrium response to an idiosyncratic subsidy increase. Lastly, I target the regression evidence in columns (1) and (2) of Table 2 to discipline the monopsony parameters $\{\xi, \underline{\ell}\}$, constructing R&D wages and employment identically to the data and running the same 2SLS regressions. I then minimize a weighted absolute distance between model and data moments.

Table 3 reports the calibrated parameters alongside the targeted moments and their model counterparts. The model fits well, with the largest deviation being the wage elasticity for small firms, which the model overestimates.

[7] The innovation process typically relies on highly specialized labor leveraging specialized machines such that opportunities for substitution may be limited.

[8] Compustat in 2000 has 1,068 firms conducting R&D while the NSF reports 17,757 firms, implying a share of listed firms among R&D-conducting firms of 5%.

[9] The value for the Corporate Income Tax is based on a GAO report, while I draw on estimates from the CBO for labor income taxation. Average R&D subsidies are based on OECD indicators.

[10] I calculate this figure based on NSF data. See footnote 6 and Supplemental Appendix G.3 on the labor share. On ψ_M, I first solve the labor market equilibrium and then scale the R&D productivity intercept, R&D wage intercept, and ψ_M jointly by the ratio of targeted to actual growth to achieve $g = 1.5\%$. This scaling preserves the first-order conditions and, thus, the equilibrium allocation, avoiding an outer loop over the growth rate. Consequently, ψ_M is not independently targeted.

Table 3: Calibrated Parameters and Targeted Moments

A. Parameters	Symbol	Value	Method
A.1. External			
Intertemporal elasticity of substitution	ς	2.00	External
Discount factor	β	0.97	External
Aggregate labor supply elasticity	φ	0.50	External
Occupational labor supply elasticity	ϵ	0.30	External
Profit Parameter	α	0.80	External
Share of firms listed	ζ	0.05	External
Elasticity Inventors vs Materials	θ	0.50	External
Labor tax rate	τ_L	0.30	External
Corporate Income Tax rate	τ_Π	0.15	External
R&D Subsidy Rate	τ	0.07	External
A.2. Internal			
Production labor supply shifter	α_P	0.29	Direct
R&D Labor supply shifter	α_R	0.06	Direct
R&D productivity intercept	μ	0.27	Direct
Non-listed firms R&D prod.	A_{nl}	0.08	Moment matching
R&D productivity variation	σ	0.68	Moment matching
R&D productivity autocorrelation	ρ	0.93	Moment matching
Returns to scale in R&D	γ	0.51	Moment matching
Labor Weight in R&D	α_L	0.92	Moment matching
Materials Price	ψ_M	0.27	Moment matching
Firm-level labor supply elasticity	ξ	4.36	Moment matching
Firm-level labor supply elasticity	$\underline{\ell}$	20.8	Moment matching
B. Moments	Target	Model	Source
Growth rate (%)	1.50	1.50	Standard value
Relative R&D size	35.0	35.0	NSF/ Compustat
SD of growth rate	0.29	0.29	Compustat
Autocorrelation of R&D	0.93	0.93	Compustat
Avg. Estimate	0.38	0.39	Table 2, Column (1)
Low Estimate	0.00	0.14	Table 2, Column (2)
High-Low Estimate	0.67	0.67	Table 2, Column (2)
Labor share	0.79	0.79	NSF
R&D Labor Force	0.05	0.05	Acemoglu et al. (2018)
Production Labor Force	0.29	0.29	Acemoglu et al. (2018)
Unit cost Elasticity	-1.00	-1.00	Acemoglu et al. (2018)

Notes: Table reports values for the main calibration. Panel A reports calibrated parameters. Panel B reports moments together with target values. See text for details.

Table 4 reports a set of untargeted moments from the data and calibrated model. Panel A shows that the model understates the dispersion in sales growth, TFP growth, and stock market returns. This is expected, as the model's only source of firm-level variation is the demand shifter, while firms face additional shocks unrelated to R&D in practice.[11] Nonetheless, the model closely matches moments related to R&D returns, which are tightly linked to monopsony power as discussed in Lemma 2. Panel B shows that the elasticity of R&D returns with respect to inventor employment is virtually identical in model and data, while the elasticities with respect to stock market returns and TFP growth also closely align. The latter two moments arise because stock returns and TFP growth are correlated with R&D productivity shocks, which raise R&D returns through the rising markdown that accompanies the expansion in inventor employment, consistent with interpreting stock returns as shocks relevant for firms' demand for inventors. Panel B further shows that R&D returns are highly persistent in both model and data, driven by persistence in R&D productivity. Note that the model could not match these moments if firms acted as price takers, since R&D returns would be constant in that case. Finally, Panel C reports additional moments related to the product market. The model qualitatively matches the R&D-to-sales elasticity and the negative relationship between sales growth and firm size, and generates highly persistent sales, all broadly in line with the data.

4.2 Counterfactuals

Monopsony power over inventors has substantial costs for growth and welfare in the calibrated model, as documented in Table 5. The "Baseline" column reports the main counterfactual, which induces firms to effectively act as price takers in the R&D labor market through targeted subsidies. This should be interpreted as a corrective-policy experiment offsetting the labor-market distortion from monopsony, rather than as a regulatory intervention that forces firms to act competitively without compensation. In the absence of monopsony power, growth accelerates by 0.12 pp to 1.62 percent per year, yielding a welfare improvement of 8% in consumption-equivalent terms.

[11] The simulation includes an iid log-normal demand shock with a standard deviation of 0.1, which captures short-run uncertainty in R&D project outcomes. This value is not targeted but chosen to provide a realistic dispersion of short-run outcomes and does not affect the targeted moments.

Table 4: Untargeted Moments in the Data and Calibrated Model

Outcome	Data		Model
A. Measures of Dispersion			
Sales Growth	0.302		0.142
TFP Growth	0.282		0.142
Stockmarket Returns	0.472		0.032
B. R&D Return Elasticities			
R&D Employment	0.282	(0.023)	0.290
Stockmarket Returns	0.278	(0.021)	0.214
TFP Growth	0.232	(0.013)	0.107
Lagged R&D Return	0.733	(0.020)	0.780
C. Additional Moments			
R&D to Sales Elasticity	0.680	(0.066)	0.677
Sales Growth to Log Sales	−0.050	(0.003)	−0.066
Autocorrelation of Sales	0.908	(0.006)	0.934

Notes: Table reports untargeted moments from the data and the calibrated model. Standard errors for data moments reported in brackets.Firms face iid log-normal demand shocks with mean normalized to 1 in levels and a standard deviation of 0.1 in logs. Demand shocks enter multiplicatively with firm technology stocks. In Panel B, stock market returns also include an iid normal shock with a standard deviation of 0.1. Regressions control for NAICS3×Year fixed effects and standard errors are clustered at the NAICS6 level.

These costs are comparable in magnitude to other market distortions studied in the literature: Berger et al. (2022) estimate that monopsony power in the production sector reduces output by 21% and welfare by 8%, while Aghion et al. (2025) estimate that resolving misallocation from heterogeneous product market power would boost growth by 31% with a welfare gain of 9%. Similarly, de Ridder (2024) finds that rising fixed costs reduced growth by 23% and welfare by 9%. Thus, monopsony in the R&D sector has sizable welfare implications, comparable to other prominent distortions.

Faster TFP growth is driven by an expansion and reallocation of R&D activity, as documented in Panel A of Table 5. Scaling R&D inputs to their counterfactual aggregate levels while holding the allocation across firms fixed accounts for 0.04 pp of the growth acceleration, while reallocation of R&D toward more productive firms at constant aggregate input levels accounts for 0.08 pp. Thus, two-thirds of the growth impact of monopsony is due to misallocation. Panel B confirms that R&D activity

Table 5: The Impact of Inventor Monopsony on Economic Growth and Welfare

Outcome	Calibration	
	Baseline	Conservative
A. Growth & R&D		
TFP Growth	1.62	1.59
Δ *Levels Only*	1.54	1.53
Δ *Reallocation Only*	1.58	1.56
$\Delta\%$ R&D Employment	4.51	4.13
$\Delta\%$ R&D Materials	5.91	4.74
Δ Avg. R&D Subsidy	37.3	34.3
$\Delta\%$ Firm Values	9.81	7.17
B. Δ *Top 1% Share*		
Innovation	4.43	3.80
R&D Employment	7.28	6.16
R&D Expenditure	17.6	17.2
C. Consumption and Welfare		
$\Delta\%$ Consumption	0.40	0.36
$\Delta\%$ *Production Employment*	0.46	0.41
$\Delta\%$ Welfare	7.75	5.80

Notes: Table reports counterfactuals for the baseline and conservative calibrations. The conservative calibration incorporates wage discrimination ($\alpha_D = 0.25$), misalignment between private and public incentives ($\chi = 0.95$), and Schumpeterian displacement ($\delta^S = 0.25$). Panel A reports the counterfactual growth rates and changes in aggregate R&D moments ($\Delta\%$: percent changes; Δ: level changes). Panel B reports changes in the share of R&D activity accounted for by the top 1% of R&D firms. Panel C reports percent changes in aggregates beyond the R&D sector. See text for details.

concentrates toward the top firms in the counterfactual. Finally, R&D employment and material use increase, though the impact on consumption is partially offset by an expansion in production employment driven by sectoral labor complementarity.

The targeted subsidies are sizable, with an implied increase in the average subsidy rate of 36 percentage points, highlighting the cost of inducing the largest firms to expand R&D.[12] Consequently, the average firm value increases by 10 percent. These

[12] The average subsidy rate is the ratio of total fiscal expenditure to total pre-subsidy R&D expenditure. I assume firm-specific linear subsidy rates as discussed below.

results raise the question of how to finance such subsidies and whether they would induce firm entry, as discussed below.

4.3 Robustness and Discussion

Next, I discuss some natural extensions to the model and various robustness checks, in each case analyzing the impact of these changes on my key counterfactual magnitudes.

Wage discrimination. The baseline model assumes uniform wages within firms. With perfect wage discrimination, firms could hire efficiently despite upward-sloping labor supply, eliminating the allocative cost of monopsony. I extend the model to allow for partial wage discrimination in Supplemental Appendix C.1. Higher levels of wage discrimination reduce the growth impact of monopsony; however, several pieces of evidence point toward limited discrimination in practice: many employers set wages nationally (Hazell et al., 2025); within-firm residual wage dispersion is flat or declining in firm size, the opposite of what discrimination would predict (Song et al., 2019); and labor supply elasticities are larger for new recruits than incumbents, again inconsistent with discriminatory wage setting (Seegmiller, 2025). Even under a calibration that allows wage discrimination to play a substantial role, offsetting residual monopsony power still improves growth by 0.10 pp and welfare by 7%, as reported in Supplemental Table D.2.

Incentive misalignment. Monopsony power unambiguously generates misallocation in the baseline model, since the planner and firms agree on the relative value of ideas but disagree on the cost of producing them. If instead the planner valued idea quality differently from firms, e.g., because knowledge spillovers are less sensitive to idea quality than private returns, then monopsony power could be less costly if it reallocated activity toward firms whose ideas the planner values more highly. I formalize this possibility in Supplemental Appendix C.3 through knowledge externalities that load less than proportionally on idea quality. While the data suggest some misalignment between private and social valuations of idea quality, as measured by the elasticity of patent citations to valuations, variation in idea quality explains little of the cross-firm variation in R&D activity. Thus, effective alignment is likely close to one and production efficiency is the primary driver of R&D heterogeneity across firms, though I cannot rule out other unmeasured sources of misalignment.

Wage discrimination and incentive misalignment each attenuate the baseline welfare cost individually. Additionally, creative destruction associated with faster growth may reduce the net-present value of new inventions and thus dampen the growth acceleration (Supplemental Appendix C.4). To provide a single conservative benchmark, I combine these three forces into a joint calibration reported in the "Conservative" column of Table 5. Jointly, they reduce the estimated growth and welfare impacts only modestly, confirming the robustness of the baseline.

Firm entry. The exercise of monopsony power raises firm values and may therefore encourage entry, partially offsetting the innovation losses. I extend the model to allow the mass of firms N to be determined by a free entry condition in Supplemental Appendix C.6. Under free entry, however, monopsony becomes even more costly (Column 1, Panel B of Supplemental Table D.3). Offsetting monopsony power through subsidies raises firm values further, inducing additional entry that amplifies the innovation gains. Inducing firms to effectively act as price takers increases the number of active firms by 12% and raises growth by 0.24 pp and welfare by 17%. By contrast, under a regulatory counterfactual that prevents firms from exercising monopsony power without fiscal compensation, firm values fall and free entry leads to substantial exit, reducing growth by 0.14 pp and welfare by 9% relative to the monopsony baseline (Column 2).

Financing and targeting subsidies. The counterfactual finances R&D subsidies through lump-sum taxation. Supplemental Table D.3 alternatively considers funding by adjusting the untargeted component of R&D subsidies or by raising labor or corporate taxes.[13] Adjusting untargeted R&D subsidies dampens the growth effects significantly and reduces the welfare impact to 2%. Labor taxes are more promising and yield a comparable growth acceleration to the baseline and a 6% welfare gain, partly reflecting lost consumption due to lower production employment. With a fixed number of firms, corporate taxation replicates the lump-sum counterfactual since R&D is expensed; however, under free entry corporate taxes dampen the benefits of offsetting monopsony power by reducing firm values.

[13] An important consideration is the fiscal cost of R&D subsidies. Implementing them through linear, type-specific taxes incurs significant inframarginal expenditure. Alternatively, one could consider a take-it-or-leave-it subsidy offer just large enough to induce firms to ignore the markdown. I discuss this further in Supplemental Appendix C.5 and assume an intermediate case here.

Another important question is whether imperfectly targeted subsidies could achieve similar improvement to the policy counterfactual. Supplemental Table D.4 suggests that this is not the case. Subsidies at identical fiscal costs but targeted uniformly to the top 50%, 5%, or 1% only achieve a fraction of the welfare gains of the main policy. Under type-specific linear costs untargeted subsidies actually achieve stronger growth and welfare improvements at the same fiscal costs. This result is driven by the large inframarginal fiscal costs of incentivizing the largest R&D firms and is overturned when subsidies are instead implemented under the most fiscally efficient scheme.

Competition and innovation. The baseline analysis focuses on imperfect competition in the R&D labor market. In Supplemental Appendix C.7, I investigate other forms of imperfect competition and show that their impact on innovation depends on the market in which they occur. Monopsony power in the production labor market reduces innovation by compressing employment and shrinking the market for blueprints, while increased fringe competition in the blueprint market also reduces innovation by eroding profits, despite raising welfare through lower static markups. These results highlight that the mechanism studied in this paper—input market power distorting R&D hiring—is distinct from, and potentially counteracted by, output market distortions that sustain innovation incentives.

Measurement error. The R&D wage measure used in estimation may be subject to three sources of measurement error: the inclusion of non-labor expenditure in R&D costs, performance-linked compensation such as stock grants and bonuses, and the imperfect measurement of inventor employment via patents. I investigate all three quantitatively in Supplemental Appendix E using the calibrated model. The materials component generates a modest bias whose direction depends on the elasticity of substitution between labor and materials in R&D; however, since I compute moments identically in model and data, the calibration absorbs this bias by construction. For bonus and equity-linked pay, I augment the model's wage process with three alternative compensation structures calibrated to match the stock-based pay share of 12.7% reported in NSF BERDS. All three specifications generate bias at short horizons, but the bias vanishes by the targeted three-year horizon, confirming that the long differencing window insulates the estimates. Finally, I simulate the patent-based measurement of inventor employment by modeling a discrete project-success process calibrated to match the observed patent and inventor counts. The key determinant of bias is whether the measurement error covaries with the instrument, which depends on how patents scale with firm size. The data strongly support a one-for-one scaling of patents with inventor employment, under which detection probability is approximately constant across firms and the resulting bias is small and downward — implying that the baseline estimates are, if anything, conservative.

Alternative estimates. I compare my estimates with two additional sources of evidence on labor supply elasticities in Supplemental Appendix F. First, Seegmiller (2025) estimates larger inverse elasticities for high-skilled workers using LEHD data and a similar research design. Re-calibrating the model to his estimates raises the welfare cost of monopsony significantly, suggesting that my baseline is conservative. Second, I compare the implied markdown schedule with estimates from Kline et al. (2019), who exploit quasi-random patent allowance decisions to estimate pass-through rates at small, first-time patentees. While the mapping between their bargaining framework and my monopsony model is necessarily indirect, the implied markdowns from their estimates exceed those in my calibration across most of the firm-size distribution, again suggesting that the baseline results represent a lower bound.

Calibration robustness. Supplemental Appendix D.3 reports the sensitivity of the baseline results to 10% perturbations of each calibrated parameter (Table D.5). The counterfactual growth and welfare effects are stable across perturbations, with welfare costs ranging from 8% to 10% and the largest sensitivity arising from the R&D returns to scale γ and the intertemporal elasticity of substitution ς.

I also re-calibrate the model under several alternative parametrizations in Supplemental Appendix D.1 (Table D.2). Higher substitutability between labor and materials in R&D ($\theta = 1.5$) reduces the welfare cost to 4%, partly because materials substitute for misallocated labor, but also because the measurement error from including materials in R&D wages is substantially larger under this parametrization. This specification also implies a sharp decline in the R&D labor cost share with firm size, a pattern not supported by NSF data. A lower R&D labor share or a higher unit cost elasticity target similarly attenuate the results. In the other direction, raising the sectoral labor supply elasticity to $\epsilon = 1.3$ substantially increases the welfare cost to 16%, as greater substitutability between R&D and production labor makes aggregate R&D employment highly responsive to the removal of monopsony distortions. Finally, replacing the log-normal productivity distribution with a Pareto yields results close to the baseline. Across all alternative calibrations, the welfare cost of monopsony ranges from 4% to 16%, bracketing the baseline estimate of 8%.

5 Conclusion

Monopsony power in U.S. labor markets has become an important policy concern. This paper suggests that such concerns are especially relevant in the market for inventors, who possess highly specialized skills and play an outsized role in productivity growth. Because innovation is cumulative, static inefficiencies in this market translate into dynamic losses for growth and welfare.

The paper reaches this conclusion in three steps. First, I develop a heterogeneous-firm endogenous growth model with monopsony power in the market for inventors. The model highlights two channels: monopsony reduces aggregate inventor employment by depressing wages, and size-dependent monopsony distorts the allocation of inventors across firms by disproportionately compressing labor demand at larger employers. Second, using an instrumental-variables strategy to estimate firm-level inventor labor supply elasticities, I find substantial monopsony power for large firms. While smaller firms appear to face competitive conditions, larger employers lose only 15% of their inventor workforce following a 10% wage reduction, implying that inventors at those firms are paid only 60% of their marginal product. Third, I calibrate the model to these elasticities and find that eliminating monopsony power raises growth by 0.12 pp and welfare by 8% in consumption-equivalent terms.

These findings point to at least two directions for future research. First, monopsony power in the corporate sector may affect inventors' entrepreneurship, a particularly relevant margin given the importance of startup acquisition by large technology firms. Second, monopsony power may shape human capital investment by depressing returns and creating differential exposure across skills, with consequences for both the allocation of inventors across firms and the accumulation of inventive human capital.

References

Acemoglu, Daron, 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," 2025.

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

_ , Jeremy Pearce, and Marta Prato, "Tapping into Talent: Coupling Education and Innovation Policies for Economic Growth," 2020.

Arnold, David, "Mergers and Acquisitions, Local Labor Market Concentration, and Worker Outcomes," 2021.

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.

Azar, José, Ioana Marinescu, Marshall Steinbaum, and Bledi Taska, "Concentration in US labor markets: Evidence from online vacancy data," *Labour Economics*, 2020, *66.*

Bagga, Sadhika, "Firm Market Power, Worker Mobility, and Wages in the US Labor Market *," *Journal of Labor Economics*, 2023, *41*, S205–S256.

Barr, Tavis and Udayan Roy, "The effect of labor market monopsony on economic growth," *Journal of Macroeconomics*, 12 2008, *30*, 1446–1467.

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

Berger, David W, Thomas Hasenzagl, Kyle F Herkenhoff, Simon Mongey, and Eric A Posner, "Comments on the 2023 Draft Merger Guidelines: A Labor Market Perspective," Technical Report 2023.

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.

Cavenaile, Laurent, Murat Alp Celik, and Xu Tian, "The Dynamic Effects of Antitrust Policy on Growth and Welfare," *Journal of Monetary Economics*, 7 2021, *121*, 42–59.

Celik, Murat Alp, "Does the Cream Always Rise to the Top? The Misallocation of Talent in Innovation," *Journal of Monetary Economics*, 1 2023, *133*, 105–128.

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

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

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

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

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

Edwards, Jim, "Emails From Google's Eric Schmidt And Sergey Brin Show A Shady Agreement Not To Hire Apple Workers," 3 2014.

Eisfeldt, Andrea L., Antonio Falato, and Mindy Z. Xiaolan, "Human Capitalists," *NBER Macroeconomics Annual*, 1 2023, *37*, 1–61.

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

Goolsbee, Austan and Chad Syverson, "Monopsony Power in Higher Education: A Tale of Two Tracks," *Journal of Labor Economics*, 2023, *41*, S257–S290.

Grullon, Gustavo, Yelena Larkin, and Roni Michaely, "Are US industries becoming more concentrated?," *Review of Finance*, 7 2019, *23*, 697–743.

Han, Chirok and Peter C.B. Phillips, "GMM estimation for dynamic panels with fixed effects and strong instruments at unity," *Econometric Theory*, 2010, *26*, 119–151.

Hazell, Jonathon, Christina Patterson, Heather Sarsons, and Taska Bledi, "National Wage Setting," 2025.

House, The White, "Executive Order on Promoting Competition in the American Economy," 7 2021.

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

Imrohoroglu, Ayse and Selale Tuzel, "Firm-Level Productivity, Risk, and Return," *Management Science*, 2014, *60*, 2073–2090.

Jäger, Simon, Jörg Heining, and Nathan Lazarus, "How Substitutable Are Workers? Evidence from Worker Deaths," 2024.

Kaltenberg, Mary, Adam Jaffe, and Margie E Lachman, "Invention and the Life Course: Age Differences in Patenting," *NBER Working Paper 28769*, 2021.

Kim, Seula, "Workers' Job Prospects and Young Firm Dynamics," 2026.

King, Robert G., Charles I. Plosser, and Sergio T. Rebelo, "Production, Growth and Business Cycles," *Journal of Monetary Economics*, 1988, *21*, 195–232.

Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar, "Who Profits from Patents? Rent-Sharing at Innovative Firms," *The Quarterly Journal of Economics*, 2019, *134*, 1343–1404.

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.

Kroft, Kory, Yao Luo, Magne Mogstad, and Bradley Setzler, "Imperfect Competition and Rents in Labor and Product Markets: The Case of the Construction Industry *," 2021.

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 Haakon, "Did R&D Misallocation Contribute to Slower Growth?," 2025.

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.

Ma, Song, Wenyu Wang, and Yufeng Wu, "Steering Labor Mobility Through Innovation *," 2025.

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

Manning, Alan, *Monopsony in Motion: Imperfect Competition in Labor Markets*, Princeton University Press, 2003.

— , “Imperfect competition in the labor market,” *Handbook of Labor Economics*, 2011, *4*, 973–1041.

Meagher, Michelle, *Competition is Killing Us*, Penguin Random House UK, 2020.

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

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

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

Oberfield, Ezra and Devesh Raval, “Micro Data and Macro Technology,” *Econometrica*, 2021, *89*, 703–732.

of Justice, The Department, “Justice Department Requires Six High Tech Companies to Stop Entering into Anticompetitive Employee Solicitation Agreements,” 9 2010.

Philippon, Thomas, *The Great Reversal – How American Gave up on Free Markets*, Belknap Press, 2019.

Prager, Elena and Matt Schmitt, “Employer Consolidation and Wages: Evidence from Hospitals,” *American Economic Review*, 2 2021, *111*, 397–427.

Prato, Marta, “The Global Race for Talent: Brain Drain, Knowledge Transfer, and Growth,” 2022.

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

Reuters, “Microsoft says it will not enforce non-compete clauses in U.S. employee agreements,” 6 2022.

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

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

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

Shi, Liyan, "Optimal Regulation of Noncompete Contracts," *Econometrica*, 2023, *91*, 425–463.

Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter, "Firming Up Inequality," *The Quarterly Journal of Economics*, 2 2019, *134*, 1–50.

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

Wu, Tim, *The Curse of Bigness*, Columbia Global Reports, 2018.

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

Appendix

A Model Appendix

A.1 Solving the Firm Problem

The firm's value maximization problem can be solved by guessing and verifying that the value function can be expressed as $V_{kt}(z_{kt}, Z_{kt}) = Z_{kt}\, V_{\pi,t} + V_{kt}(z_{kt})$, where $V_{\pi,t}$ is the present discounted value of profits from an invention of quality 1 and $V_{kt}(z_{kt})$ is the solution to

$$V_{kt}(z_{kt}) = \max_{\ell_{kt}, M_{kt}} \Big\{ z_{kt}\, Q_{kt}^N \frac{V_{\pi,t+1}}{R_{t+1}} - (1-\tau)\,(w_{R,kt}\, \ell_{kt} + \psi_M\, M_{kt}) \\ + R_{t+1}^{-1} \mathbb{E}_t[V_{kt+1}(z_{kt+1}) | z_{kt}] \Big\}.$$

Consequently, firms' R&D choices balance the profits and costs of innovation. Along the Balanced Growth Path, the value of profits simplifies to $V_\pi = \frac{1}{R-1}\pi$, while the value function normalized by Z_t is the solution to

$$V_k(z) = \max_{\ell_{kt}, m_{kt}} \Big\{ z\, q_{kt}^N \frac{V_\pi}{R} - (1-\tau)\,(w_{R,kt}\, \ell_{kt} + \psi_M\, m_{kt}) + \frac{1+g}{R} \mathbb{E}_t[V_k(z') | z] \Big\},$$

where I denoted next period's values with a prime.

A.2 Balanced Growth Path Equations

This section reports the equations describing the Balanced Growth Path for the decentralized equilibrium and balanced growth path allocation. Variables in lower case are normalized by Z_t. I maintain the notation of denoting firms by subscript k, where each firm is a defined by its demand shifter and (normalized) stock of knowledge $\{z_k, Z_k\}$ and define $l_R = L_R/N$. Both model solutions rely on definitions (1), (2), and (7), and the implied ergodic distribution over z_k from (9).

The **decentralized equilibrium** can be described by a system of $8+3N$ equations for $8+3N$ unknowns:

$$(1+g)^{\varsigma} = \beta\ R \tag{A.1}$$

$$(1-\alpha)\ \alpha^{\frac{\alpha}{1-\alpha}} = \frac{c}{1-\tau_L}\ L^{\frac{1}{\varphi}}\ \left(\frac{L_P}{\alpha_P\ L}\right)^{\frac{1}{\epsilon}} \tag{A.2}$$

$$c + \psi_M\ \int_N m_k\ dk = (1-\alpha^2)\ \alpha^{\frac{\alpha}{1-\alpha}}\ L_P \tag{A.3}$$

$$V_\pi = \frac{1}{R-1}\ (1-\alpha)\ \alpha^{\frac{2-\alpha}{1-\alpha}}\ L_P \tag{A.4}$$

$$w_R = \frac{c}{1-\tau_L}\ L^{\frac{1}{\varphi}}\left(\frac{L_R}{\alpha_R\ L}\right)^{\frac{1}{\epsilon}}\ \frac{L_R}{\int_N \ell_k\ \left(\underline{\ell}^{\xi} + (\ell_k/l_R)^{\xi}\right)\ dk} \tag{A.5}$$

$$(1-\tau)\ \psi_M = z_k\left(\alpha_L^{\frac{1}{\theta}}(\ell_k/m_k)^{\frac{\theta-1}{\theta}} + (1-\alpha_L)^{\frac{1}{\theta}}\right)^{\frac{\gamma\ \theta}{\theta-1}-1}(1-\alpha_L)^{\frac{1}{\theta}}\ m_k^{\gamma-1}\ V_\pi \tag{A.6}$$

$$\frac{m_k}{\ell_k} = \frac{1-\alpha_L}{\alpha_L}\left(\frac{w_R}{\psi_M}\ \left(\underline{\ell}^{\xi} + (\ell_k/l_R)^{\xi}\right)\left(1+\xi\ \frac{(\ell_k/l_R)^{\xi}}{\underline{\ell}^{\xi} + (\ell_k/l_R)^{\xi}}\right)\right)^{\theta} \tag{A.7}$$

$$g = \int_N z_k\ q_k^N\ dk \tag{A.8}$$

Similarly, the **planner allocation** can be described by

$$(1+g)^{\varsigma} = \beta\ R \tag{A.9}$$

$$(1-\alpha) = c\ L^{\frac{1}{\varphi}}\ \left(\frac{L_P}{\alpha_P\ L}\right)^{\frac{1}{\epsilon}} \tag{A.10}$$

$$c + \psi_M\ \int_N m_k\ dk = (1-\alpha)\ L_P \tag{A.11}$$

$$V_Q = c\ \frac{\beta(1+g)^{-\varsigma}}{1-\beta(1+g)^{1-\varsigma}} \tag{A.12}$$

$$\tilde{w}_R = c\ L^{\frac{1}{\varphi}}\left(\frac{L_R}{\alpha_R\ L}\right)^{\frac{1}{\epsilon}}\ \frac{L_R}{\int_N \ell_k\ \left(\underline{\ell}^{\xi} + (\ell_k/l_R)^{\xi}\right)\ dk} \tag{A.13}$$

$$\psi_M = z_k\left(\alpha_L^{\frac{1}{\theta}}(\ell_k/m_k)^{\frac{\theta-1}{\theta}} + (1-\alpha_L)^{\frac{1}{\theta}}\right)^{\frac{\gamma\ \theta}{\theta-1}-1}(1-\alpha_L)^{\frac{1}{\theta}} m_k^{\gamma-1}\ V_Q \tag{A.14}$$

$$\frac{m_k}{\ell_k} = \frac{1-\alpha_L}{\alpha_L}\left(\frac{\tilde{w}_R}{\psi_M}\ \left(\underline{\ell}^{\xi} + (\ell_k/l_R)^{\xi}\right)\right)^{\theta} \tag{A.15}$$

$$g = \int_N z_k\ q_k^N\ dk \tag{A.16}$$

A.3 Proofs

Proof of Lemma 1. The equation follows immediately from the labor supply curve, i.e., the household's first-order condition with respect to R&D labor supply. □

Proof of Proposition 1. Consider the first order condition for ℓ_{kt} under $\alpha_L = 1$:

$$\gamma \; z_{kt} \; V_\pi \; \ell_{kt}^{\gamma-1} = (1-\tau_{kt}) \; W_{R,kt} \; (1+\epsilon_{kt}) \quad \text{with} \quad \epsilon_{kt} = \xi \; \frac{(\ell_{kt}/l_{R,t})^\xi}{\underline{\ell}^\xi + (\ell_{kt}/l_{R,t})^\xi}$$

Taking a full derivative with respect to a small shock $\Delta \ln(1-\tau_{kt})$ or $\Delta \ln z_{kt}$, we have

$$\frac{\Delta \ln \ell_{kt}}{\Delta \ln z_{kt}} = -\frac{\Delta \ln \ell_{kt}}{\Delta \ln(1-\tau_{kt})} = \frac{1}{1-\gamma+\epsilon_{kt}\,\left(1+\frac{1}{1+\epsilon_{kt}}\,\frac{\Delta \ln \epsilon_{kt}}{\Delta \ln \ell_{kt}}\right)}.$$

Since ϵ_{kt} is a constant for $\underline{\ell} = 0$, it follows that the employment elasticity is as well in this case. Furthermore, one can show that the denominator is strictly increasing in ℓ_{kt}, such that the elasticity is lower for higher employment levels.[14]

The planner faces the same first order conditions except for dropping the markdown term. As a result, the derivations are equivalent to setting $\frac{\partial \ln \epsilon_{kt}}{\partial \ln \ell_{kt}} = 0$.

Next, in the case with materials, we can define

$$f(m_{kt}/\ell_{kt}) = \left(\alpha_L^{\frac{1}{\theta}} + (1-\alpha_L)^{\frac{1}{\theta}} \left(\frac{m_{kt}}{\ell_{kt}}\right)^{\frac{\theta-1}{\theta}}\right)^{\gamma - \frac{\theta-1}{\theta}}$$

and derive

$$\frac{\Delta \ln \ell_{kt}}{\Delta \ln z_{kt}} = -\frac{\Delta \ln \ell_{kt}}{\Delta \ln(1-\tau_{kt})} = \frac{1}{1-\gamma+\left(1-\theta\,\frac{\partial \ln f(m_{kt}/\ell_{kt})}{\partial \ln(m_{kt}/\ell_{kt})}\right)\epsilon_{kt}\,\left(1+\frac{1}{1+\epsilon_{kt}}\,\frac{\Delta \ln \epsilon_{kt}}{\Delta \ln \ell_{kt}}\right)},$$

where I took advantage of the relative demand (A.7). Next, we can derive

$$\theta\,\frac{\partial \ln f(m_{kt}/\ell_{kt})}{\partial \ln(m_{kt}/\ell_{kt})} = (\gamma\theta - (\theta-1))\;\; s_{kt}^m \quad \text{with} \quad s_{kt}^m = \frac{(1-\alpha_L)^{\frac{1}{\theta}}(m_{kt}/\ell_{kt})^{\frac{\theta-1}{\theta}}}{\alpha_L^{\frac{1}{\theta}} + (1-\alpha_L)^{\frac{1}{\theta}}(m_{kt}/\ell_{kt})^{\frac{\theta-1}{\theta}}}.$$

[14] For the derivation, define $x = \frac{(\ell_{kt}/L_{R,t})^\xi}{\underline{\ell}^\xi + (\ell_{kt}/L_{R,t})^\xi}$ s.t. $\epsilon_{kt} = \xi\; x$. Then, one can show that $\frac{\partial \ln \epsilon_{kt}}{\partial \ln \ell_{kt}} = \xi(1-x)$. Consequently, the derivative of the entire term involving elasticities is $\xi(1+\xi)\;(1+\xi\; x)^{-2}$, which is strictly positive.

Using the equilibrium materials-labor ratio in R&D, we have

$$s_{kt}^{m} = \frac{(1-\alpha_L)\,\psi_M^{1-\theta}}{\alpha_L\left(w_{R,kt}\,(1+\epsilon_{kt})\right)^{1-\theta} + (1-\alpha_L)\psi_M^{1-\theta}}$$

It follows that the term $1-(1-\theta(1-\gamma))\,s_{kt}^{m}$ is increasing in the marked-down wage as long as $\theta \leq 1$ and, thus, the overall elasticity is declining in firm size. □

Proof of Proposition 2. I will denote values for the planner with an $*$. Combining equations for both solution concepts from Section A.2 and imposing $\alpha_L = 1$ and $\underline{\ell} = 0$, we can express the relative R&D employment in the planner allocation and decentralized equilibrium as

$$\left(\frac{L_R^*}{L_R}\right)^{1+\frac{1}{\epsilon}\frac{1}{1-\gamma}} = (1+\xi)^{\frac{1}{1-\gamma}}\left(\frac{\alpha^2}{1+\alpha}\right)^{-\frac{1}{1-\gamma}}\left(\frac{\beta^{-1}(1+g)^{\varsigma}-1}{\beta^{-1}(1+g^*)^{\varsigma}-(1+g^*)}\right)^{\frac{1}{1-\gamma}}$$

Furthermore, relative growth depends on R&D employment only:

$$\frac{g^*}{g} = \left(\frac{L_R^*}{L_R}\right)^{\gamma},$$

Inverting the equation for relative R&D employment and plugging in above, one can show that $g^* > g$, which in turn implies $L_R^* > L_R$ as long as $\epsilon > 0$. Furthermore, this gap is increasing in ξ due to the markdown term.

Under heterogeneous monopsony power, the structure of the first-order conditions only differs due to the markdown. Since the markdown itself is increasing in ℓ_{kt}, this yields a rising wedge such that the planner's allocation is steeper in productivity than the decentralized equilibrium for a given L_R. □

Proof of Lemma 2. The expression follows immediately from rearranging the first-order conditions. □

A.4 Numerical details

Model solution. I solve the model on a discretized grid for the productivity process with 100 grid points. For the **calibration**, I first bisect over the wage intercept to create a model solution matching the targeted mass of inventors and treating $\underline{\ell}\ L_{R,t}$ as a joint parameter. I then scale this solution to achieve the targeted growth rate, multiplying the productivity intercept, wage intercept, and cost of materials by the ratio of targeted to actual growth. This method immediately yields the desired growth rate and preserves the policy functions in the R&D sector. Using these policy functions, I then back out the remaining model aggregates and parameters. Note that the profits are pinned down by L_P, which is also directly targeted. For the **counterfactuals**, I instead run a solver for the relevant aggregates for a given parameter combination, solving the firm problem for a given guess and iterating until convergence.

Simulation. I discretize the AR(1) process over 50 grid points using the Rouwenhorst method and solve the related policy function based on aggregates from the balanced growth path. Then, I simulate a listed firm for 100,000 periods, whereof I discard the first 10,000 periods for warm-up.

Calibration. I calibrate the model through moment matching as discussed in the main text. I construct all aggregates using the model solution. For dynamic moments, I simulate the model and construct all variables as in the data, including using R&D expenditure per worker as my proxy for R&D wages. For autocorrelation, I use the estimator proposed in Han and Phillips (2010) to control for permanent firm differences. I calculate the moments, including the instrumental variable regressions, as in the empirical section. I then calculate a weighted absolute distance of target and model moments, which I minimize using standard solvers.

B Empirical Appendix

B.1 Variable Construction

R&D Employment. I calculate R&D employment based on inventors listed on firms' granted patents at the time of application. I link patents to firms in Compustat using the crosswalk in Kogan et al. (2017). In the baseline, I assign each firm all their inventors. Alternatively, I assign each firm a share of a given inventor in each year based on the share of patents assigned to the firm. In my baseline, I use the inventors identified by USPTO's Patentsview and confirm robustness with either only U.S.-based inventors or inventors identified by Kaltenberg et al. (2021).

R&D Wages. I construct a proxy for R&D wages as the ratio of R&D expenditure to R&D employment.

R&D Returns. I construct R&D returns as the ratio of patent valuations estimated in Kogan et al. (2017) to R&D expenditure over a 5-year window. I focus on observations with at least 50 underlying patents when computing R&D returns. See Lehr (2025) for a further discussion of the measure.

Stock-market returns. I construct annual stock market returns from monthly returns reported in CRSP. I then construct excess returns using the S&P500 index returns from the same data source. When constructing returns, I line up the month with the fiscal year of the company.

TFP estimates. I use the updated firm-level TFP estimates from Imrohoroglu and Tuzel (2014) to construct TFP growth as an alternative instrument for firm-level R&D demand shocks.

Table B.1: Summary Statistics & Sample Coverage

	1975–2019	1975–1995	1996–2019
A. Summary Statistics			
Sales	6,022 (19,243)	3,877 (9,266)	7,313 (23,174)
Stock market returns	0.10 (0.56)	0.10 (0.39)	0.10 (0.64)
(Residual) TFP Growth	-0.00 (0.32)	0.00 (0.21)	-0.00 (0.38)
R&D Expenditure	307.5 (1,022.5)	117.6 (312.4)	421.8 (1,257.4)
Patents	84.3 (304.1)	51.8 (103.3)	103.9 (375.1)
Inventors	128.5 (432.1)	70.7 (154.9)	163.3 (530.5)
Inventors (FTE)	126.3 (423.4)	70.2 (154.1)	160.0 (519.4)
Inventors (U.S.)	109.9 (348.2)	64.6 (139.2)	137.1 (424.9)
Inventors (Verified)	75.0 (233.9)	52.1 (113.3)	88.8 (281.8)
R&D Wages (log)	0.52 (1.14)	0.17 (1.06)	0.73 (1.13)
B. Observations			
Observations	17,231	6,473	10,758
Firms per Year	383	308	448
Unique Firms	811	510	722
C. Sample Coverage			
Patents	61.5%	64.3%	59.1%
Patent Valuations	81.7%	86.5%	77.4%
R&D vs Compustat	87.4%	88.8%	86.2%
R&D vs BEA	44.9%	37.2%	51.7%

Note: Table reports mean and standard deviation over different subperiods. Standard deviations are reported in brackets. The sample is restricted to firm-year observations with at least 10 patents in the five-year forward window and firms with an average of at least 5 patents per year. The Compustat sample used as denominator in the sample coverage section excludes utilities, financials, and public sector firms.

Supplemental Appendix to "Does Monopsony Matter for Innovation?"

Nils H. Lehr

C Extensions and their Calibration

This Appendix collects extensions and discusses the calibration of related parameters. Additional quantitative results and calibrations are reported in Appendix D.

C.1 Wage Discrimination

The inability of firms to set discriminatory wages across employees is crucial to generating monopsony power. This section considers the case of wage discrimination and highlights the challenges of disentangling it from monopsony power empirically. For expositional purposes, I focus on the case without materials.

C.1.1 Theoretical Exploration

We can write the labor disutility for R&D workers equivalently as

$$L_{R,t} = \left(\underline{\ell}^{\xi} + \frac{1}{1+\xi}\right)^{-1} \int_{N_t} \left(\int_0^{\ell_{kt}} \left(\underline{\ell}^{\xi} + \left(\frac{\ell \ N_t}{L_{R,t}} \right)^{\xi} \right) \, d\ell \right) \, dk,$$

which highlights that the marginal disutility differs among the employees of a given firm. Evaluating the integral we see that the 0th worker has a disutility proportional to $\underline{\ell}^{\xi}$, while the ℓ_{kt}-th worker has a disutility proportional to $\underline{\ell}^{\xi} + (\ell_{kt} \ N_t / L_{R,t})^{\xi}$. If a firm can impose perfectly discriminatory wage, then it will pay a lower wage to the former than to the latter. As a result, the wage for the ℓth worker at any company needs to satisfy

$$W_{R,kt}(\ell) = W_{R,t} \ \left(\underline{\ell}^{\xi} + \left(\frac{\ell \ N_t}{L_{R,t}} \right)^{\xi} \right).$$

where $W_{R,t}$ is defined as in the main text. Total labor cost for the firm, C_{kt}, is then the integral over all employees, and marginal cost is the wage of the last employee:

$$C_{kt} = \int_0^{\ell_{kt}} W_{R,kt}(\ell)\, d\ell \propto \ell_{kt} \left(\underline{\ell}^{\xi} + \frac{1}{1+\xi} \left(\frac{N\, \ell_{kt}}{L_{R,t}} \right)^{\xi} \right) \text{ s.t. } \frac{\partial C_{kt}}{\partial \ell_{kt}} \propto \underline{\ell}^{\xi} + \left(\frac{\ell_{kt}\, N}{L_{R,t}} \right)^{\xi}.$$

Consequently, firms' marginal costs are the true marginal costs of hiring the last worker, so that the planner allocation and the decentralized equilibrium agree on the relative marginal cost of R&D workers across firms. Thus, there is no misallocation of R&D workers across firms nor insufficient demand due to firms' gaming of the labor market. Naturally, other sources of insufficient demand persist.

A natural question is then whether we can distinguish between both models empirically. Unfortunately, this task is difficult as average wages behave quite similarly in both models. In particular, one can verify that the elasticity of the average wage with respect to employment, $W_{R,kt} = C_{kt}/\ell_{kt}$, remains positive:

$$\frac{\partial \ln W_{R,kt}}{\partial \ln \ell_{kt}} = \xi \, \frac{\frac{1}{1+\xi} \left(\frac{\ell_{kt}\, N_t}{L_{R,t}} \right)^{\xi}}{\underline{\ell}^{\xi} + \frac{1}{1+\xi} \left(\frac{\ell_{kt}\, N_t}{L_{R,t}} \right)^{\xi}}.$$

This phenomenon occurs as rising wages at the margin also push up the average wage, even though inframarginal wages are unaffected.

In practice, firms' power to discriminate is likely limited due to information asymmetries and/or fairness considerations. I therefore consider a model in which workers are paid a fraction α_D of the fully discriminatory wage and a fraction $1 - \alpha_D$ of the required marginal wage given total hiring. Total labor costs then satisfy

$$C(\ell_{kt}) \propto \ell_{kt} \left(\underline{\ell}^{\xi} + \frac{1 + (1-\alpha_D)\, \xi}{1+\xi} \left(\frac{\ell_{kt}\, N_t}{L_{R,t}} \right)^{\xi} \right)$$

Consequently, marginal costs become proportional to

$$\frac{\partial C_{kt}}{\partial \ell_{kt}} \propto \left(1 + (1-\alpha_D)\, \xi \, \frac{\left(\frac{\ell_{kt}\, N_t}{L_{R,t}} \right)^{\xi}}{\underline{\ell}^{\xi} + \left(\frac{\ell_{kt}\, N_t}{L_{R,t}} \right)^{\xi}} \right) \left(\underline{\ell}^{\xi} + \left(\frac{\ell_{kt}\, N_t}{L_{R,t}} \right)^{\xi} \right)$$

Marginal costs are proportional to marginal disutility for $\alpha_D = 1$ and to marginal average disutility for $\alpha_D = 0$.

C.1.2 The Cost of Monopsony Under (Partial) Wage Discrimination

Figure C.1 explores the impact of wage discrimination quantitatively in the main calibration. For reference, I also report values for the baseline monopsony case and the case of no monopsony implemented through subsidies. As shown in Panels (a) and (b), growth and welfare converge to the no-monopsony case as the model approaches full wage discrimination; however, the gap remains large at intermediate values. Panel (c) further highlights that wage discrimination dampens the model-implied labor supply elasticities, especially for larger firms. Thus, an appropriate comparison across models requires a recalibration to match the empirical estimates of the inverse labor supply elasticities.

Figure C.1: The Effect of Wage Discrimination on Efficiency

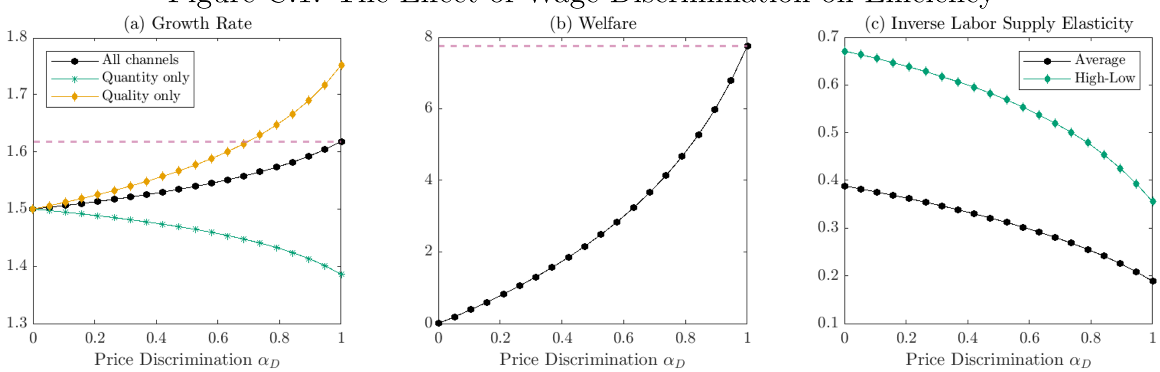

Notes: The figure shows the impact of increasing firms' ability to discriminate among workers' wages. All models employ the baseline calibration and then impose alternative values for α_D. The dotted red and blue lines show the outcomes under the calibrated model and the counterfactual with no monopsony power respectively.

C.2 Calibration

The baseline assumes uniform wage setting within the firm ($\alpha_D = 0$). In practice, firms may have some ability to differentiate pay among their workers, though the evidence suggests this is limited. First, Hazell et al. (2025) find that 40–50 percent of postings for the same job within a firm carry identical wages across locations, a pattern they attribute to fairness norms and institutional rigidities.[1] Second, a large

[1] See also Dube et al. (2019), who document that even small arbitrary wage differences among peers generate large quit responses, rationalizing why firms compress pay.

literature documents significant dispersion in wages within firms. While it would be tempting to interpret this evidence as suggesting large discrimination, it could equally reflect compositional differences—larger firms hire across a broader range of skills and occupations. Indeed, the model with discrimination links wage dispersion to firm size.[2] To this end, Song et al. (2019) find that dispersion in within-firm wages that is not accounted for by worker characteristics is essentially unaffected by firm size.[3] The evidence in Song et al. (2019) and Hazell et al. (2025) thus suggests limited scope for wage discrimination, motivating the baseline of uniform wage setting.

Motivated by this evidence, I re-calibrate the model setting $\alpha_D = 0.25$—a value that is deliberately generous to the discrimination hypothesis—and matching the same moments. The resulting calibration is reported in Table D.1. I repeat the baseline counterfactual exercise of offsetting (the remaining) monopsony power and report the results in Table D.2. As expected, the impact of monopsony power is moderated, but remains quantitatively important.

C.3 Incentive Misalignment and the Cost of Monopsony

This appendix investigates the robustness of the main results to relaxing a key assumption of the baseline model, namely that public and private valuations of R&D are aligned.

C.3.1 Model Extension

In the baseline model, public and private values of ideas both scale with z_{kt}, so the equilibrium is independent of whether variation in effective R&D productivity $A_k \; z_{kt+1}$ reflects idea quality or production efficiency. This equivalence breaks if the planner values idea quality differently from firms—for instance, because knowledge spillovers are less sensitive to idea quality than private returns. In that case, differences in R&D driven by idea quality may be inefficient, while those driven by

[2] To see this, note that under full wage discrimination, within-firm wage dispersion is governed by the range of marginal disutilities $[\underline{\ell}^{\xi}, \underline{\ell}^{\xi} + (\ell_{kt} N_t / L_{R,t})^{\xi}]$. For small firms, ℓ_{kt} is small and wages are approximately uniform at $\underline{\ell}^{\xi}$; for larger firms, the gap widens.

[3] Table V in Song et al. (2019) reports this component excluding firms above different thresholds. For 1980–86 and 2007–13, their estimates excluding firms above 10,000 employees (0.163, 0.141) or firms above 1,000 employees (0.171, 0.145) are *higher* than the overall estimates from Table IV (0.154, 0.136), implying that there is less unexplained wage variation within the largest firms, not more—the opposite of what wage discrimination would predict.

production efficiency remain efficient. I formalize this distinction by decomposing effective R&D productivity into a quality and a production component and allowing for misalignment in the social value of the former.

Formally, let ω measure the degree to which effective R&D productivity is driven by idea quality. Then, the value of a given quantity of ideas is given by $(A_k \; z_{kt+1})^{\omega}$, while the R&D production function is given by

$$Q^N_{kt+1} = Z_t \; (A_k \; z_{kt+1})^{1-\omega} \; \left(\alpha_L^{\frac{1}{\theta}} \; \ell_{kt}^{\frac{\theta-1}{\theta}} + (1-\alpha_L) \left(\frac{M_{kt}}{Z_t} \right)^{\frac{\theta-1}{\theta}} \right)^{\gamma \; \frac{\theta}{\theta-1}} .$$

Profits from a given invention scale with the idea quality, i.e., $(z_{kt+1} \; A_k)^{\omega}$, leaving the firm's optimization problem unaffected by ω. Finally, the quality adjusted number of ideas evolves according to

$$Z_{t+1} = \int_{N_t} (z_{kt+1} \; A_k)^{\omega} \; Q^N_{kt+1} \; dk + Z_t.$$

By construction, the growth rate is unaffected by ω, confirming that the source of effective R&D productivity does not matter in this model.

I introduce misalignment by allowing the planner to weight idea quality differently from firms. Let $\tilde{Z}_t$ be the state of knowledge, evolving according to

$$\tilde{Z}_{t+1} = \int_{N_t} \bar{z}^{\omega} \; \left(\frac{z_{kt+1} \; A_k}{\bar{z}} \right)^{\omega \, \chi} \; Q^N_{kt+1} \; dk + \tilde{Z}_t \text{ with } \chi \in [0,1] \text{ and } \bar{z} = \mathbb{E}[(z_{kt+1} \; A_k)].$$

Knowledge loads differentially on idea quality relative to Z_t when $\chi < 1$, and feeds into the R&D production function as follows:

$$Q^N_{kt} = \tilde{Z}_t \; (A_k \; z_{kt+1})^{1-\omega} \; \left(\alpha_L^{\frac{1}{\theta}} \; \ell_{kt}^{\frac{\theta-1}{\theta}} + (1-\alpha_L) \left(\frac{M_{kt}}{Z_t} \right)^{\frac{\theta-1}{\theta}} \right)^{\gamma \; \frac{\theta}{\theta-1}} .$$

Firms do not internalize their impact on knowledge, making knowledge creation a pure externality. Solving for a balanced growth path with $g_Z = g_{\tilde{Z}}$ pins down the ratio of both stocks:

$$\tilde{z} \equiv \frac{\tilde{Z}_t}{Z_t} = \frac{\int_{N_t} \bar{z}^{\omega} \; \left(\frac{z_{kt+1} \; A_k}{\bar{z}} \right)^{\omega \, \chi} \; Q^N_{kt+1} \; dk}{\int_{N_t} (A_k \; z_{kt+1})^{\omega} \; Q^N_{kt+1} \; dk}.$$

The normalized R&D production function is then given by the original one adjusted for $\tilde{z}$. If $\chi < 1$ and $\omega > 0$, there is misalignment between planner incentives, which take into account the knowledge externalities, and the firm incentives for R&D, which do not. This follows immediately, since a planner would maximize the growth rate, which can be expressed as

$$g = \int_{N_t} \bar{z}^{\omega} \left(\frac{z_{kt+1} \; A_k}{\bar{z}} \right)^{\omega \, \chi} \; \frac{Q_{kt+1}^N}{\tilde{Z}_t} \, dk,$$

while (implicitly) firms maximize just the stock of quality adjusted ideas Z_t, ignoring their actions' impact on $\tilde{Z}_t$.

C.3.2 Calibration

To assess the importance of incentive misalignment, one needs information both on the misalignment of quality assessments, χ, and the importance of quality differences in effective R&D productivity, ω. In the following, I argue that there is plausible misalignment in the valuation of ideas; however, idea quality drives little of the overall variation in R&D expenditure. Consequently, the effective alignment, $1 - \omega(1 - \chi)$, is likely very high.

Consider idea alignment first. I propose to estimate the degree of alignment using the elasticity of patent forward-citations to valuations. Patent valuations are intended to be a proxy for the profit impact of the patent, scaling with $A_k \; z_{kt+1}$, while citations measure the knowledge creation, scaling with $(A_k \; z_{kt+1})^{\chi}$. The estimates reported in Table C.1 suggest an elasticity in the range of $[0.6, 0.7]$, depending on the inclusion of firm fixed effects. While this points to some misalignment, the key question is whether idea quality drives sufficient variation in R&D to make this misalignment consequential.

Table C.1: Patent Citation-Valuation Elasticity

	(1)	(2)	(3)	(4)
	Log Patent Citations			
Log Patent Valuations	0.612***	0.691***	0.739***	0.696***
	(0.013)	(0.014)	(0.013)	(0.014)
Citation Window	Life-time	Life-time	5-year	20-year
Firm FEs		Yes	Yes	Yes
Observations	17,964	17,964	17,964	17,964

Note: All regressions control for NAICS3 × year fixed effects. Sample restricted to observations with at least 25 patents in the 5 years from t to $t+4$. Standard errors are clustered at the firm level.

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

Next, consider the degree to which variation in idea quality drives R&D decisions. Table C.2 shows that average patent valuations—a proxy for the private value of ideas—explain very little of the variation in R&D expenditure or total patents across firms. Once controlling for permanent firm characteristics, which may conflate idea quality with production efficiency, the within-firm R^2 is 0–2%; even the across-firm R^2 never exceeds 40%. This suggests a value of ω close to zero: the component of effective R&D productivity orthogonal to idea quality drives most R&D variation. Measurement error in patent valuations biases these R^2 values downward, but even so, a within-firm R^2 of at most 2% implies that idea quality plays a negligible role in time-series R&D variation.

Table C.2: R&D-Patent Value Elasticity

	(1)	(2)	(3)	(4)
	Log R&D Expenditure/ Patents			
Log Avg. Patent Valuations	0.818***	0.150***	0.442***	-0.069**
	(0.028)	(0.028)	(0.028)	(0.028)
Outcome	R&D	R&D	Patents	Patents
Firm FEs		Yes		Yes
Within R2	0.37	0.02	0.12	0.00
Observations	17,445	17,441	17,964	17,964

Note: All regressions control for NAICS3 $\times$ year fixed effects. Sample restricted to observations with at least 25 patents in the 5 years from t to $t+4$. Standard errors are clustered at the firm level.

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

In summary, even taking the upper end of the citation-valuation elasticity ($\chi \approx 0.6$) and a generous reading of the within-firm R^2 as implying $\omega \leq 0.05$, effective alignment would be no less than 0.98. Nonetheless, I present estimates for the growth impact of monopsony power under alternative values for effective alignment $\tilde{\chi} = 1 - \omega(1 - \chi)$ below.

C.3.3 Results

Figure C.2 reports the impact of offsetting monopsony power for different values of public-private alignment. As expected, lower levels of public-private alignment reduce the impact of offsetting monopsony power due to the impact on R&D productivity via the knowledge externality. Indeed, the impact turns slightly negative at $\tilde{\chi} = 0.5$. Nonetheless, the impact remains large for the empirically relevant range of $\tilde{\chi}$ close to one.

C.4 Schumpeterian Forces

In Schumpeterian growth models, faster innovation induces displacement of incumbents, which in turn reduces the value of incumbency and, thus, innovation (Aghion and Howitt, 1992). A potential concern in my counterfactual could thus be that it

Figure C.2: The Cost of Monopsony under Public-Private Misalignment

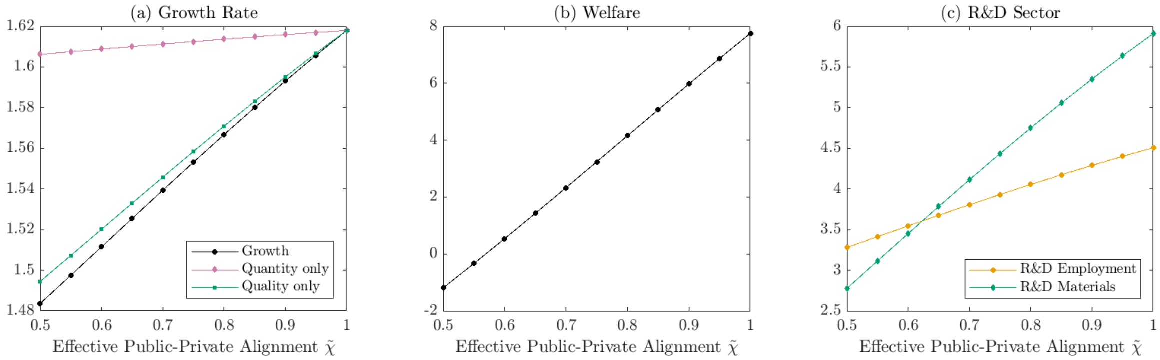

Notes: Figure reports the impact of offsetting monopsony power on economic growth, welfare, and R&D employment. For all values of $\tilde{\chi}$, I first recalibrate the model to match 1.5% growth with monopsony power by adjusting the R&D productivity intercept μ and then offset monopsony power via subsidies as in the baseline. Panel (a) reports the impact on growth and decomposes it into the R&D productivity impact from $\tilde{Z}$ and the quantity impact induced by lower R&D productivity. Panel (b) reports the welfare impact in consumption-equivalent terms, while Panel (c) reports impacts on R&D employment and materials.

overstates growth effects since my model lacks such Schumpeterian forces. To account for this channel, I introduce an extension incorporating Schumpeterian displacement, which is reflected in the "conservative" counterfactual in the main text.

Suppose that instead of adding new products, a share δ of inventions improve upon an existing invention by the quality increment z_{kt}. As a result, the previous owner of the invention is displaced. To simplify the exposition, I will assume any improvements occur only among inventions of the same quality level z. Since any type's share in new inventions is the same as its share in all inventions, the replacement rate still scales with the aggregate innovation rate $\bar{m}$. Thus, a firm's value of an invention with $n-1$ subsequent improvements is given by

$$V_\pi(z,n) = \pi\ z\ n\ \frac{1}{r + \delta\ \bar{m}}.$$

Since each invention is an improvement with probability δ and original with probability $1-\delta$, the number of cumulated improvements follows a geometric distribution (Klette and Kortum, 2004):

$$s_n = (1-\delta)\ \delta^{n-1}.$$

Consequently, the expected value of an invention is given by

$$V_\pi(z) = \sum_{n=1}^{\infty} V_\pi(z,n)\; s_n = \pi\; z\; \frac{1}{(r+\delta\;\bar{m})(1-\delta)}.$$

It follows that for $\delta \in (0,1)$, a higher innovation rate $\bar{m}$ reduces the value of inventions. All other elements of the model, however, are unaffected due to the additive nature of improvements, including the growth rate. For convenience, I implement a shorthand for this extension in the model by setting the firms' discount rate to $R/(1-\delta(g-\bar{g}))$, where $\bar{g}$ is the growth rate under monopsony. This specification is exact at the calibrated equilibrium and captures the first-order effect of growth deviations on firm value.

C.5 The Fiscal Impact of Alternative Policies

The baseline counterfactual offsets monopsony power through targeted subsidies that offset the markdown. An important consideration when evaluating the resulting fiscal impact and the resulting financing need is thus how these subsidies are implemented. For the purpose of this exercise I will focus on the cost of targeted subsidies and ignore the fiscal costs of untargeted R&D subsidies τ. Consider two alternatives. In the first alternative, monopsony power is offset through a set of firm-specific subsidies $(1-\tau_{kt})$ that perfectly offset the markdown. The resulting fiscal burden from offsetting monopsony power alone is

$$F_t^L = (1-\tau)\int_{N_t} \tau_{kt}\;\ell_{kt}\;W_{R,kt}\;dk.$$

In the second alternative, I assume that the government makes each firm a take-it-or-leave-it offer T_{kt} for setting R&D employment ignoring the markdown. Denote with a prime values set by the firm in absence of a subsidy, then this offer needs to satisfy the participation constraint

$$T_{kt}+z_{kt}\;q_{kt}^{N}V_\pi-(1-\tau)\left(\ell_{kt}\;W_{R,kt}+\psi_M\;M_{kt}\right) \geq z_{kt}\;q_{kt}^{N'}V_\pi-(1-\tau)\left(\ell'_{kt}\;W'_{R,kt}+\psi_M\;M'_{kt}\right).$$

Let this condition hold with equality. Then, the associated fiscal costs are

$$F_t^D = \int_{N_t} T_{kt} \, dk.$$

It is straightforward to prove that $F_t^D \leq F_t^L$ as the next lemma shows.

Lemma 1. The fiscal burden under the take-it-or-leave-it offer is strictly lower than the fiscal burden under linear taxes as long as $\tau_{kt} > 0$ and all production factors are used.

Proof. Taking the definition of T_{kt} and subtracting the fiscal burden under linear taxes yields

$$\begin{aligned} &\Big(z_{kt}\, q_{kt}^{N'} V_\pi - (1-\tau)\left(\ell_{kt}'\, W_{R,kt}' + \psi_M\ M_{kt}'\right)\Big) \\ &- \left(z_{kt}\, q_{kt}^{N} V_\pi - (1-\tau)\left((1-\tau_{kt})\, \ell_{kt}\, W_{R,kt} + \psi_M\ M_{kt}\right)\right) \end{aligned}$$

The first term is the firm's maximized profit absent any subsidy, while the second is the firm's maximized profit under the linear subsidy $(1-\tau_{kt})$, which strictly reduces marginal costs. Since the firm could always choose (ℓ_{kt}', M_{kt}') under the subsidy and earn weakly more than without it, the second term must be strictly larger whenever $\tau_{kt} > 0$ and the firm adjusts its inputs. The difference is thus strictly negative, implying $T_{kt} < \tau_{kt}\, \ell_{kt}\, W_{R,kt}$. □

Thus, a highly sophisticated government might face a much lower burden than what linear taxes imply. To allow for intermediate cases, I parameterize the effective fiscal burden as $F_t = \lambda\, F_t^L + (1-\lambda)\, F_t^D$, where $\lambda \in [0,1]$ captures the sophistication of the subsidy scheme. In practice, existing R&D subsidies in the U.S. appear closer to the linear case. Analysis from the Congressional Research Service on the U.S. R&D tax credit finds that "[...] the credit is no more efficient at targeting marginal investment than a flat credit for all research."[4] Motivated by this evidence, I report results at both extremes of λ in the targeting exercise reported in Table D.4 and set $\lambda = 0.5$ when evaluating different tax reforms to fund the R&D subsidies.

[4]See Congressional Research Service, "The Federal Research and Development (R&D) Tax Credit," CRS Report R48848, 2026.

C.6 Entry

The baseline model assumes an exogenous number of firms. In an extension, I instead allow N_t to be determined endogenously via a free entry condition. I assume that there is a pool of potential entrants that can pay entry costs $\phi_E\, Z_t$ to draw an R&D productivity level from the ergodic distribution over z_{kt} and enter without any existing blueprints, i.e., $Z_{kt} = 0$. Normalizing by Z_t, the free entry condition is thus given by

$$\phi_E = \mathbb{E}[V(0, z_{kt})].$$

The condition pins down N_t since a larger number of firms increases the demand for R&D workers and reduces the available resources for consumption, both of which increase wages and thus reduce $V(0, z_{kt})$. For $\varsigma > 1$, larger entry also increases the discount rate through its effect on growth, providing an additional stabilizing force.

C.7 Competition and Innovation

The main focus of this paper is the impact of monopsony power on innovation. A natural question is then the impact of other sources of imperfect competition on growth in this model. I discuss two additional channels and argue that the link between competition and innovation crucially depends on the market in which it occurs. Imperfect competition in R&D input markets reduces innovation by incentivizing firms to restrain hiring, while imperfect competition in the output market for blueprints increases innovation by raising profits. Finally, imperfect competition downstream of the output market reduces innovation by limiting firms' market size.

C.7.1 Model

Fringe Competition. I consider two forms of fringe competition. First, for each intermediate good, a fraction $1 - f_e \in [0, 1]$ is immediately imitated by a competitor with identical unit costs. Bertrand competition then compresses the equilibrium price for these goods to marginal costs ψ. Second, the remaining intermediate goods are subject to Bertrand competition from imitators with marginal costs ψ/α^{f_i} with $f_i \in [0, 1]$. The firm then engages in limit pricing, setting the price at the marginal costs of the fringe competitor. The resulting expected profits are given by

$$\mathbb{E}_t[\pi_{jt}] = z_{jt}\ L_{P,t}\ (1-\alpha^{f_i})\ \left(\frac{\psi}{\alpha^{1+f_i}}\right)^{-\frac{\alpha}{1-\alpha}}\ f_e.$$

Furthermore, net-output is given by

$$Y_t - X_t = Z_t\ L_{P,t}\ \left((1-f_e)(1-\alpha) + f_e(1-\alpha^{1+f_i})\ \alpha^{\frac{f_i\ \alpha}{1-\alpha}}\right)\left(\frac{\psi}{\alpha}\right)^{-\frac{\alpha}{1-\alpha}}$$

Note that f_i and f_e measure the degree of competition, ranging from perfect competition at 0 to monopoly at 1. As is standard in endogenous growth models, perfect competition reduces profits to 0 and, thus, incentives to innovate. Finally, the resulting wage is given by

$$W_{P,t} = (1-\alpha)\ Z_t\ \left((1-f_e) + f_e\ \alpha^{\frac{f_i\ \alpha}{1-\alpha}}\right)\left(\frac{\psi}{\alpha}\right)^{-\frac{\alpha}{1-\alpha}}.$$

Production Labor Wedge. I introduce imperfect competition in the production labor market in a reduced form through a wedge Δ_L in the production firm's first-order condition, so that labor demand is given by

$$\frac{\partial Y_t}{\partial L_{P,t}} = (1+\Delta_L)\ W_{P,t}.$$

A wedge above zero can be interpreted as capturing the impact of monopsony power over production workers in reduced form.

C.7.2 Results

Table C.3 reports the counterfactual impact of imperfect competition. The first column repeats the baseline R&D monopsony counterfactual, while the second and third columns report the results of introducing a production labor wedge and increasing fringe competition in the blueprint market, respectively. A production labor wedge reduces both growth and welfare (column 2): by compressing production employment, it shrinks the market for blueprints and thus weakens innovation incentives. Increasing fringe competition in the blueprint market also reduces growth, since lower profits discourage innovation, but raises welfare through lower static markups (column 3). Thus, whether competition raises or lowers innovation depends on the market in which it occurs: it raises innovation by correcting input market distortions or expanding demand for blueprints, but lowers it by eroding blueprint profitability. Welfare and growth can diverge when lower innovation is accompanied by reduced static distortions.

Table C.3: The Impact of Competition on Growth and Welfare

Outcome	No R&D Monop.	$\Delta_L = 1.5$	$f_e = 0.75$
A. Growth & R&D			
TFP Growth (%)	1.62	1.46	1.43
Δ% R&D Employment	4.51	-4.31	-7.51
Δ% R&D Materials	5.91	-9.45	-15.4
Δ% Firm Values	9.81	-12.9	-20.5
Δ% Flow Profits	0.46	-12.3	-19.6
B. Consumption & Welfare			
Δ% Consumption	0.40	-12.4	17.1
Δ% *Production Empl.*	0.46	-12.3	7.20
Δ% Welfare	7.75	-10.6	9.19

Notes: This table shows the effects of different competition channels. The three counterfactual columns report the impact of removing R&D monopsony through R&D subsidies, introducing a production labor wedge ($\Delta_L = 1.5$), andreducing markups in the intermediate goods sector through competition from imitators ($f_e = 0.75$).Panel A reports the counterfactual growth rate and changes aggregate moments for the R&D sector. Δ% refers to percent changes, while Δ refers to level changes. Panel B reports percent changes in aggregate moments beyond the R&D sector. See text for details.

C.8 Alternative R&D Productivity Distribution

The baseline assumes a log-normal distribution for z_{kt}. I confirm the robustness of the main conclusions by alternatively exploring a calibration with a Pareto distribution. I assume that the stochastic properties of the idea quality are determined by a persistent random normal variable $\tilde{z}$:

$$\tilde{z}_{kt+1} = \rho\, \tilde{z}_{kt} + (1-\rho^2)^{1/2}\, \nu_{kt+1} \quad \text{with} \quad \nu_{kt+1} \sim N(0,1).$$

The actual values are derived by projecting the values onto a Pareto:

$$z_{kt} = z_0 \left(1 - \Phi(\tilde{z}_{kt})\right)^{-\frac{1}{\sigma}},$$

where $\Phi(\tilde{z}_{kt})$ is the cumulative distribution function of the normal distribution.

Table D.1 reports the resulting calibration, which undershoots the High-Low estimate but provides a closer match for the Low estimate; the opposite of the log-normal calibration. Nonetheless, the counterfactual analysis reported in Table D.2 finds a slightly stronger growth impact of offsetting monopsony power, accelerating growth to 1.71 instead of 1.62 percent per year, and a slightly stronger welfare effect of 15% compared to 8%. As in the baseline, growth acceleration occurs by raising idea quality at the expense of quantity.

D Extensions and Additional Quantitative Results

D.1 Alternative Calibrations

Table D.1: Calibration under Alternative Parameterizations

A. Parameters	Baseline	Alt. targets II	$\theta =$ 1.5	High γ	Low α_L	$\alpha_D =$ 0.25	Pareto
θ	0.50	0.50	1.50	0.50	0.50	0.50	0.50
μ	0.27	0.26	0.28	0.28	0.27	0.27	0.06
A_{nl}	0.08	0.09	0.11	0.12	0.09	0.09	0.06
σ	0.68	0.66	0.56	0.59	0.67	0.69	1.43
ρ	0.93	0.94	0.93	0.93	0.94	0.94	0.93
γ	0.51	0.51	0.51	0.61	0.51	0.51	0.51
α_L	0.92	0.92	0.70	0.92	0.87	0.92	0.91
ξ	4.36	4.76	4.32	4.50	4.26	4.90	7.14
$\underline{\ell}$	20.8	23.3	25.8	20.3	21.2	20.5	31.4
α_D	0.00	0.00	0.00	0.00	0.00	0.25	0.00
B. Moments	Baseline	Alt. targets II	$\theta =$ 1.5	High γ	Low α_L	$\alpha_D =$ 0.25	Pareto
Growth rate (%)	1.50	1.50	1.50	1.50	1.50	1.50	1.50
Relative R&D size	35.0	35.0	35.0	35.0	35.0	35.0	35.0
SD of growth rate	0.29	0.29	0.29	0.29	0.29	0.29	0.29
Autocorr. of R&D	0.93	0.93	0.92	0.92	0.93	0.93	0.93
Avg. Estimate	0.39	0.29	0.38	0.39	0.39	0.34	0.43
Low Estimate	0.14	0.09	0.11	0.15	0.14	0.11	0.01
High-Low Estimate	0.67	0.56	0.67	0.67	0.67	0.67	0.52
Labor share	0.79	0.79	0.79	0.79	0.67	0.79	0.79
R&D Labor Force	0.05	0.05	0.05	0.05	0.05	0.05	0.05
Prod. Labor Force	0.29	0.29	0.29	0.29	0.29	0.29	0.29
Unit cost Elasticity	-1.00	-1.00	-1.00	-1.50	-1.00	-1.00	-1.01

Notes: Table reports calibrated parameters (Panel A) and targeted moments (Panel B) for alternative calibrations. Parameters not reported are held at baseline values. Alt. targets II uses the elasticity estimates without controls from Table 2. Low α_L targets a labor share of 2/3, High γ targets a unit cost elasticity of 1.5, and Pareto uses a Pareto R&D productivity distribution. See text for details.

D.2 Additional Counterfactuals

Table D.2: Monopsony and Growth under Alternative Calibrations

Outcome	Baseline	Alt. Targets	Alt. Targets II	$\theta =$ 1.5	Low α_L	High γ	$\epsilon =$ 1.3	$\alpha_D =$ 0.25	Pareto
A. Growth & R&D									
TFP Growth	1.62	1.66	1.60	1.56	1.61	1.64	1.74	1.60	1.71
Δ *Levels Only*	1.54	1.53	1.53	1.52	1.53	1.54	1.65	1.53	1.55
Δ *Reallocation Only*	1.58	1.63	1.56	1.54	1.57	1.60	1.58	1.57	1.65
Δ% R&D Employment	4.51	3.49	4.50	3.30	4.35	4.50	22.4	4.33	6.51
Δ% R&D Materials	5.91	7.38	5.22	0.20	5.54	6.19	15.3	5.24	7.95
Δ Avg. R&D Subsidy	37.3	44.2	34.1	25.6	32.0	38.6	37.0	34.3	52.5
Δ% Firm Values	9.81	8.93	10.2	9.25	8.21	15.2	9.25	7.70	9.51
B. Consumption and Welfare									
Δ% Consumption	0.40	0.53	0.34	0.26	0.32	0.47	-1.18	0.36	0.66
Δ% *Production Empl.*	0.46	0.60	0.40	0.26	0.42	0.55	-0.99	0.42	0.72
Δ% Welfare	7.75	10.9	6.48	3.91	6.86	9.65	15.9	6.59	14.6

Notes: Table reports counterfactuals for alternative calibrations. Alt. targets calibrates the model targeting the labor supply elasticity estimates in Seegmiller (2025). Alt. targets II uses the elasticity estimates without controls from Table G.1. Low α_L targets a labor share of 2/3, High γ targets a unit cost elasticity of 1.5, and Pareto uses a Pareto R&D productivity distribution. Panel A reports counterfactual growth rates and changes in aggregate R&D moments (Δ%: percent changes; Δ: level changes). Panel B reports percent changes in aggregates beyond the R&D sector. See text for details.

Table D.3: Monopsony and Growth under Alternative R&D Subsidy Financing

Outcome	Alternative Funding Mechanisms				
	Lump-Sum	Regulate	R&D Tax	Labor Tax	Corp. Tax
A. Fixed Mass of Firms					
A.1. Growth & R&D					
TFP Growth (%)	1.62	1.62	1.54	1.61	1.62
Δ% R&D Employment	4.51	4.51	-3.94	3.00	4.51
Δ% R&D Materials	5.91	5.91	-14.6	4.22	5.91
Δ% Firm Values	9.81	-22.4	7.86	8.18	6.19
A.2. Consumption and Welfare					
Δ% Consumption	0.40	0.40	0.22	-0.96	0.40
Δ% Welfare	7.75	7.75	2.36	6.06	7.75
A.3. Subsidies and Taxes					
Avg. R&D Subsidy	32.3	7.00	9.07	32.3	32.3
Labor Tax	30.0	30.0	30.0	32.8	30.0
Corporate Tax	15.0	15.0	15.0	15.0	17.8
B. Free Entry					
B.1. Growth & R&D					
TFP Growth (%)	1.74	1.36	1.62	1.70	1.69
Δ% R&D Employment	10.6	-7.48	0.16	7.78	8.17
Δ% R&D Materials	10.4	-3.74	-12.6	7.73	8.66
Δ% Active Firms	11.8	-24.6	9.57	9.70	7.16
B.2. Consumption and Welfare					
Δ% Consumption	0.16	1.32	0.07	-1.27	0.24
Δ% Welfare	17.4	-8.68	8.58	13.5	13.4
B.3. Subsidies and Taxes					
Avg. R&D Subsidy	33.0	7.00	8.40	32.9	32.7
Labor Tax	30.0	30.0	30.0	33.0	30.0
Corporate Tax	15.0	15.0	15.0	15.0	18.0

Notes: This table compares mechanisms for financing the subsidies required to offset firm markdowns. Lump-sum is the baseline and assumes non-distortionary financing. The remaining counterfactuals achieve budget neutrality through the average R&D subsidy rate τ, the labor income tax τ_L, or the corporate income tax τ_Π. The specification assumes an intermediate level of targeting sophistication such that the government only partly pays the inframarginal costs implied by type-specific linear tax rates. Panel A holds the number of firms fixed, while Panel B allows for free entry. Within each panel, sub-panels 1–3 report R&D moments (Δ%: percent changes; Δ: level changes), broader aggregates, and tax or subsidy rates, respectively. See text for details.

Table D.4: Monopsony and Growth under Alternative R&D Subsidy Targeting

Outcome	Perfect	Untargeted		Targeting Top X%		
	Targeting	High	Low	50%	5%	1 %
A. Growth & R&D						
TFP Growth	1.62	1.62	1.54	1.54	1.54	1.54
Δ% *Levels Only*	1.54	1.62	1.54	1.54	1.52	1.52
Δ% *Reallocation Only*	1.58	1.50	1.50	1.50	1.51	1.53
Δ% R&D Employment	4.51	13.9	4.15	4.01	2.05	0.80
Δ% R&D Materials	5.91	39.3	10.9	10.9	10.4	10.8
Δ Avg. R&D Subsidy	37.3	34.1	12.4	12.4	12.2	12.2
Δ% Firm Values	9.81	2.40	0.75	0.87	3.11	4.14
B. Δ Top 1% Share						
Innovation	4.43	-0.24	-0.08	-0.04	0.46	2.29
R&D Employment	7.28	-0.13	-0.04	0.02	0.83	2.88
R&D Expenditure	17.6	-0.42	-0.14	-0.02	1.61	8.11
C. Consumption & Welfare						
Δ% Consumption	0.40	0.21	0.05	0.05	0.01	-0.01
Δ% *Production Empl.*	0.46	0.65	0.17	0.17	0.13	0.12
Δ% Welfare	7.75	7.74	2.27	2.25	2.20	2.58

Notes: This table evaluates the importance of targeting R&D subsidies. Perfect targeting refers to subsidies perfectly targeting the monopsony wedge. Untargeted involves no targeting but set the subsidy rate to achieve the same fiscal expenditure as under the perfect targeting scheme. High refers to the scenario where the government expenditure reflects linear type-specific taxes rates, while low assumes the budget cost assuming that the government makes a take-it-or-leave-it offer compensating firms for additional R&D expenditure. All other schemes target the top X% of firms by R&D productivity achieving the same budget as the untargeted low specification. Panel A reports the counterfactual growth rate and changes aggregate moments for the R&D sector. Δ% refers to percent changes, while Δ refers to level changes. Panel B reports changes in the share of R&D activity accounted for by the top 1% of R&D firms. Panel C reports percent changes in aggregate moments beyond the R&D sector. See text for details.

D.3 Parameter Sensitivity

Table D.5: Parameter Robustness

Outcome	Base	ξ	$\underline{\ell}$	ϵ	φ	A_{nl}	γ	ρ	σ	ζ	ς
A. Targeted Moments											
TFP Growth	1.50	1.50	1.52	1.50	1.48	1.54	1.45	1.50	1.50	1.58	1.49
Relative R&D size	35.0	35.0	38.2	35.0	35.0	31.0	41.2	35.0	35.0	36.9	35.1
SD of growth rate	0.29	0.29	0.30	0.29	0.29	0.30	0.30	0.45	0.33	0.30	0.29
Autocorr. of R&D	0.93	0.93	0.93	0.93	0.93	0.93	0.93	0.83	0.93	0.93	0.93
Avg. Estimate	0.39	0.38	0.34	0.39	0.39	0.35	0.44	0.37	0.36	0.35	0.39
Low Estimate	0.14	0.13	0.11	0.14	0.14	0.12	0.16	0.16	0.12	0.12	0.14
High-Low Estimate	0.67	0.70	0.60	0.67	0.67	0.62	0.76	0.49	0.67	0.63	0.67
Labor share	0.79	0.79	0.79	0.79	0.79	0.80	0.79	0.79	0.79	0.79	0.79
R&D Labor Force	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05
Prod. Labor Force	0.29	0.29	0.29	0.29	0.28	0.29	0.29	0.29	0.29	0.29	0.29
Unit cost Elasticity	-1.00	-1.00	-1.08	-1.01	-0.95	-1.13	-1.03	-1.00	-1.00	-1.25	-0.96
B. Counterfactual Impact											
TFP Growth	1.62	1.62	1.64	1.63	1.60	1.65	1.60	1.62	1.62	1.71	1.60
Δ% R&D Empl.	4.51	4.97	4.53	5.32	4.61	4.12	4.91	4.51	4.51	4.77	4.44
Δ% R&D Materials	5.91	6.10	5.69	6.35	6.04	5.37	6.98	5.91	5.91	5.93	5.77
Δ% Consumption	0.40	0.40	0.38	0.32	0.49	0.37	0.49	0.40	0.40	0.41	0.38
Δ% Welfare	7.75	7.84	7.62	8.15	7.70	7.55	9.46	7.75	7.75	8.65	9.69

Notes: Table reports model sensitivity with respect to selected parameters. The first column replicates the baseline calibration; each subsequent column varies the indicated parameter by 10%. The ρ column reduces the parameter, while all other columns increase it. Panel A reports targeted moments under monopsony power.Panel B reports results after offsetting monopsony power through subsidies. See text for details.

E Measurement Error in Labor Supply Elasticities

In this Appendix, I discuss three sources of potential measurement error for the IV estimates of the labor supply elasticity of inventors. Section E.1 discusses the presence of cost for materials in the R&D per inventor measure, Section E.2 discusses bonus payments to inventors, and Section E.3 discusses the imperfect measurement of R&D employment via patents. None of the three materially affect the main estimates. The materials bias is quantitatively relevant in principle but is already absorbed by the moment matching strategy used in calibration. Bonus payments generate negligible bias at the targeted three-year horizon. Finally, accounting for imperfect inventor measurement would imply marginally stronger monopsony power, so that the baseline estimates can be considered conservative.

E.1 Materials in R&D

In the empirical section, I define the R&D wedges as the ratio of R&D expenditure to inventors. If R&D expenditure also includes materials, this assumption implies a proportional measurement error equal to the inverse labor share in R&D. Let $\widetilde{W}_{R,kt}$ be the imperfect measure of R&D wages and s_{kt} the labor share, then:

$$\widetilde{W}_{R,kt} = \frac{W_{R,kt}}{s_{kt}} \qquad \text{with} \quad s_{kt} = \frac{W_{R,kt}\ \ell_{kt}}{W_{R,kt}\ \ell_{kt} + \psi_M\ M_{kt}}.$$

Using the relative demand for R&D inputs, the inverse labor share can be expressed as

$$\frac{1}{s_{kt}} = 1 + \frac{1-\alpha_L}{\alpha_L}(1+\epsilon_{kt})^{\theta}\left(\frac{w_{R,kt}}{\psi_M}\right)^{\theta-1}.$$

Thus, the elasticity of the imperfect wage measure with respect to R&D employment as induced by a factor-neutral productivity shock is

$$\frac{\partial \ln \widetilde{W}_{kt}}{\partial \ln \ell_{kt}} = \epsilon_{kt}\left(1 + (1-s_{kt})\left((\theta-1) + \theta\,\frac{1}{1+\epsilon_{kt}}\,\frac{\partial \ln \epsilon_{kt}}{\partial \ln \ell_{kt}}\right)\right).$$

The expression implies that the elasticity estimate could be biased, but the direction is ambiguous. Under homogeneous monopsony power, $\epsilon_{kt} = \epsilon$ and $\frac{\partial \ln \epsilon_{kt}}{\partial \ln \ell_{kt}} = 0$. Consequently, the direction of the bias depends on whether material and labor are substitutes or complements. It is negative for complements, 0 for Cobb-Douglas, and positive for substitutes.[5] For heterogeneous monopsony, the same concern applies, but the bias is shifted towards positive, since $\frac{\partial \ln \epsilon_{kt}}{\partial \ln \ell_{kt}} > 0$.

In the calibration, I follow the empirical data treatment and compute moments using $\widetilde{W}_{kt}$ in both the data and the model. Since the materials component enters symmetrically on both sides, the calibrated parameters are unaffected by this bias. The remaining question is whether the labor supply elasticities implied by the calibrated model differ from those that would obtain under a pure wage measure. Figure E.1 investigates this by comparing IV estimates computed from $\widetilde{W}_{kt}$ with those computed from the true wage W_{kt} in simulated data from the same calibrated model. Under the baseline calibration, materials generate a modest upward bias of 0.02–0.05 for the 3-year coefficient, which does not affect the qualitative findings. This bias is larger when materials and labor are more substitutable or when the R&D labor share is lower, but remains marginal under the baseline calibration.[6]

Figure E.1: Materials and Estimated Labor Supply Elasticities

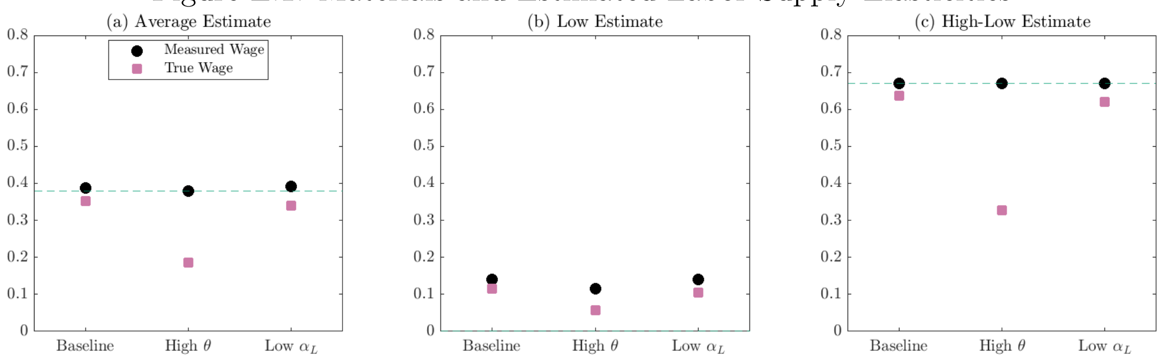

Notes: Figure reports estimates of the labor supply elasticity for different specifications of the quantitative model. Panel (a) follows specification (17), while (b) and (c) report the coefficients corresponding to specification (18). The "Wage" coefficients are computed using true wages $W_{R,kt}$, while "Baseline" reports coefficients using the measured wages $\widetilde{W}_{R,kt}$. "High θ" and "Low α_L" report coefficients using measured wages from calibrations with larger θ and a lower labor share target.

E.2 Bonus Payments

Apart from wages, inventors may be compensated in stocks, or stock options, or bonuses based on the firm's performance (Lerner and Wulf, 2007). These types of

[5] The size of the bias changes with the labor share such that larger and smaller firms may have different estimates even if the underlying labor supply elasticity is constant.

[6] The model with $\theta = 1.5$ also implies a strong decline in the labor share in R&D for large firms, which is not present in the NSF data.

compensation could in principle link R&D wages to the company's overall performance and, thus, its stock returns, which may constitute a violation of the exclusion restriction when using stockmarket returns as an instrument to estimate the labor supply elasticity. I next investigate this concern quantitatively by augmenting the model's wage process with alternative specifications of bonus payments or stock-based compensation and studying their impact on the estimated firm-level labor supply elasticity.

E.2.1 Specification

I proceed in two steps. First, I add i.i.d. log-normal demand shocks ν_{kt} to the firm's innovation stock Z_{kt} that are realized after the firm's R&D decision and that could be tied to bonus payments for R&D workers. This shock affects stockmarket returns through profits and could simultaneously affect measured R&D wages through the bonus component, constituting a potential violation of the exclusion restriction to the instrument. Second, I consider three separate specifications for alternative pay structures:

1. **Direct exposure.** In the first specification, I introduce the simplest form of performance pay and allow for a direct link between R&D wages and the idiosyncratic R&D productivity shock:

$$w_{R,kt} \, \frac{\nu_{kt}^{\phi}}{\mathbb{E}[\nu_{kt}^{\phi}]} \, \ell_{kt} + \psi_M \, m_{kt}.$$

   The parameter ϕ controls the degree to which R&D wages respond to the idiosyncratic shocks and, thus, the degree of potential bias.

2. **Vested shares.** In the second specification, I assume that a share ω_S of wages is paid in stocks, which vest over the t_S periods starting from the grant date, s.t., $t_S = 1$ is immediate vesting. Per GAAP accounting rule, such vested stocks are valued at the grant-date fair value and recognized as R&D expenditure ratably over the vesting period. As a result, realized R&D expenditure is

$$(1 - \omega_S) \, w_{R,kt} \, \ell_{kt} + \frac{\omega_S}{t_S} \sum_{s=0}^{t_S - 1} w_{R,kt-s} \, \ell_{kt-s} + \psi_M \, m_{kt}.$$

3. **Bonus payments.** In the final specification, I assume that a share ω_S of wages is base-pay, while the expected rest takes the form of bonus payments that are only paid under positive shocks. As a result, realized R&D expenditure is

$$(1-\omega_S)\, w_{R,kt}\, \ell_{kt} + \omega_S\, \frac{\max\{\nu_{kt}-1,0\}}{\mathbb{E}[\max\{\nu_{kt}-1,0\}]}\, w_{R,kt}\, \ell_{kt} + \psi_M\, m_{kt}.$$

I calibrate the second and third specifications to achieve a bonus share of 12.7% of total payments, matching NSF data, and set the vesting period for the second specification to three years, a value covering more than 85% of executive pay packages analyzed in Cadman et al. (2013). In the first specification I set $\phi = 0.5$, deliberately chosen to overstate the plausible degree of performance sensitivity: a 10% unexpected increase in R&D productivity raises the wage bill by approximately 5%. Finally, I set the standard deviation of $\ln \nu_{kt}$ to 0.1, consistent with the baseline calibration, and normalize its mean in levels to 1. This calibration is intended to highlight the qualitative implications for the main estimates without being a quantitative assessment.

E.2.2 Results

Figure E.2 reports the results for the different specifications across estimation horizons. The first specification leads to an upwards bias of all three coefficients of interest on impact, but no meaningful bias at longer horizons. The second specification instead leads to a significant downwards bias over the first two years, i.e., the period over which vesting was only partial, but not beyond. This occurs as realized wages respond less, since they partly reflect vesting from previous periods. Finally, the third specification directionally yields a similar bias as the first specification. Notably, the bias from each specification vanishes by year 2. Since any mixture of these channels that matches the observed bonus share of 12.7% inherits this property, the targeted coefficients at the three-year horizon are robust to these bonus specifications or any convex combination thereof.

E.3 Patent-Based Measurement Error

In the empirical section, I measure R&D employment based on the number of inventors observed in patents. Such measured R&D employment (unique inventors on patents) could be subject to systematic measurement error that could bias the es-

Figure E.2: Main Coefficient for Different Bonus Schemes

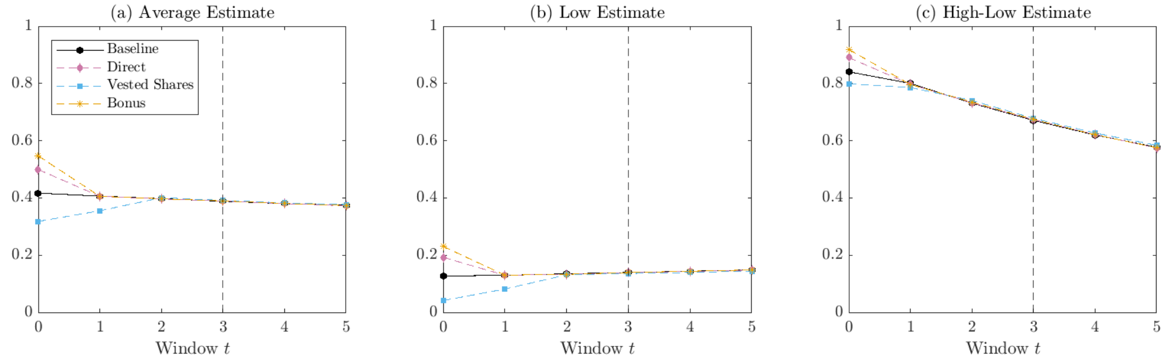

Notes: Figure reports main regression estimates for labor supply elasticities for bonus or stock-based compensation specifications. Panel (a) reports estimates for the specification without interaction, while (b) and (c) report coefficients from the interaction regression. The vertical lines indicate the targeted coefficients.

timates of the R&D labor supply elasticity. I propose to use the calibrated model together with a simulation of the patenting process therein to assess this concern. My results suggest that the estimates are conservative.

I mimic the data generating process by explicitly simulating a discrete set of projects with a discrete number of inventors per firm-year. Inventors can participate in multiple projects and each project has an ex-ante identical chance of success π. Projects are only recorded when they are successful and, as a result, inventors only appear in the data if at least one of their projects was successful. I simulate this process using simulated model data mapped to projects and inventors with scaling factors $\{\kappa, \kappa_L\}$ chosen to match an average of 72 patents and 108 inventors per firm-year as in the data. I first describe the simulation process followed by a description of the mapping between model and simulation and a presentation of the main results. In the final section, I discuss alternative mappings.

E.3.1 Algorithm

Consider a firm k with a set of P_{kt} projects and N_{kt} inventors for a set of periods indexed by t. Each project is successful with probability π resulting in a patent on which the associated inventors are listed. For each firm-year (P_{kt}, N_{kt}):

1. Assign each inventor $j \in \{1, \dots, N_{kt}\}$ to a home project $p(j)$ drawn uniformly from $\{1, \dots, P_{kt}\}$. Let $H_p = \{j : p(j) = p\}$.

2. Draw target team size $m_p \sim F_M$ for each project, where F_M is the empirical

distribution of inventors per patent.

3. Draw project success $S_p \sim \text{Bernoulli}(\pi)$ for each project.
4. Initialize observed inventors $\mathcal{O} = \{j : S_{p(j)} = 1\}$.
5. For each successful project with $|H_p| < m_p$, draw $m_p - |H_p|$ additional inventors uniformly from the set not already assigned to project p and add to $\mathcal{O}$.
6. Measured employment: $\tilde{N}_{kt} = |\mathcal{O}|$.

R&D expenditure is perfectly observed from the data.

E.3.2 Mapping Model to Algorithm and Calibration

I discretize the number of inventors as $N_{kt} = \text{round}(\ell_{kt}\ \ \kappa_L)$, where κ_L is a scaling parameter. Next, I similarly set the number of projects to scale with R&D employment $P_{kt} = \text{round}(\ell_{kt}\ \ \kappa/\pi)$, where κ is a scaling parameter and dividing by π ensures that the expected number of projects is unaffected by π itself. Scaling the number of patents with R&D employment is motivated by the observation, documented below, that patents and inventor counts scale 1-for-1 in the data. In the discussion below, I consider two alternatives: scaling with the innovation rate $\tilde{z}_{kt}$, i.e., $P_{kt} = \text{round}(z_{kt}\ \ q^N_{kt}\kappa/\pi)$ or scaling with the number of blueprints, i.e., $P_{kt} = \text{round}(q^N_{kt}\kappa/\pi)$. As I highlight below, these alternative specifications imply empirical patent-inventor relationships that are not borne out by the data.

I calibrate $\{\kappa, \kappa_L\}$ to match an average of 72 patents and 108 inventors per firm-year, as in the data. I experiment with different values of the project success rate π, since it is fundamentally unobserved. Lower values imply greater uncertainty and, thus, potentially more measurement error.

E.3.3 Main Results

Figure E.3 reports the resulting estimates for different levels of π. For specifications with measurement error, i.e., $\pi < 1$, coefficients are marginally smaller, but the quantitative implication is limited. These results suggest that calibrating the model without measurement error is (marginally) conservative, since adding measurement error decreases the simulated regression coefficients.

Figure E.3: Main Coefficient for Different Degrees of Project Uncertainty π

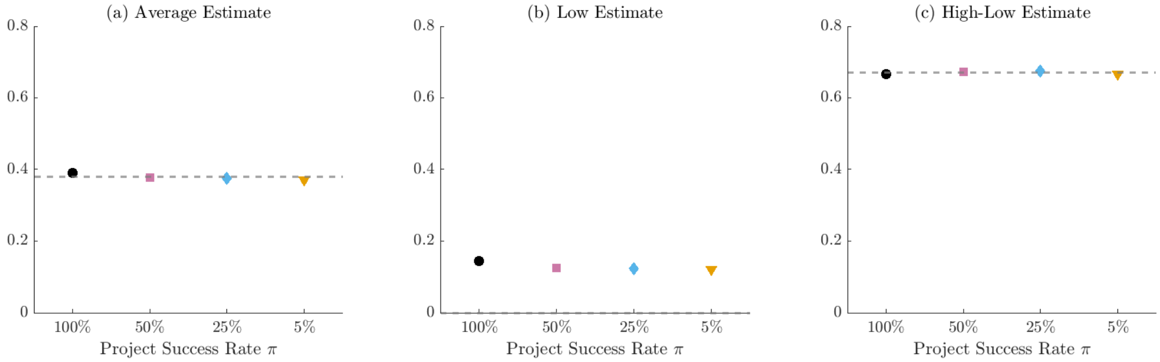

Notes: Figure reports main regression estimates for labor supply elasticities for different values of project success rate π. Panel (a) reports estimates for the specification without interaction, while (b) and (c) report coefficients from the interaction regression.

E.3.4 Discussion

The conservative finding above hinges on the scaling assumption governing how patents map to R&D employment in the simulation. Choosing the right scaling is ultimately an empirical question.[7] Table E.1 estimates a 1-for-1 relationship between observed patents and inventors in my data, and Panel (b) of Figure E.4 confirms that this elasticity is only consistent with the R&D employment scaling used in the main specification. The two alternatives—scaling with the number of blueprints q_{kt}^N or with quality-adjusted inventions $z_{kt}\, q_{kt}^N$—imply patent-to-inventor elasticities that are too small or too large relative to the data.

To understand why the scaling choice matters for the estimates, it is useful to review the basic estimation idea. Let X denote log R&D employment, Y the log wage bill, and Z the stock return instrument. Denote by $\tilde{X} = X + \eta$ log measured R&D employment including a measurement error component η. Maintaining the validity of the IV strategy absent measurement error, the IV with true employment identifies the coefficient β:

$$\beta^{IV} = \frac{\text{Cov}(Y - X, Z)}{\text{Cov}(X, Z)} = \frac{\text{Cov}(Y, Z)}{\text{Cov}(X, Z)} - 1 = \beta.$$

[7] It is straightforward to show that the model is equivalent whether differences in z_{kt} are interpreted as idea quality, i.e., product demand shifters, or productivity shifters in the R&D production function. As a result, the calibrated model is silent on what would appropriately constitute a patent.

Table E.1: Patent-Inventor Elasticity

	(1)	(2)	(3)	(4)
		Log Patents		
Log Inventors	0.995***	1.023***	1.024***	1.009***
	(0.007)	(0.009)	(0.009)	(0.010)
Inventor Measure	FTE	FTE	Raw	US
Firm FEs		Yes	Yes	Yes
Observations	15,885	15,885	15,885	15,885

Note: All regressions control for NAICS3 $\times$ year fixed effects. All variables at 5-year horizon. Sample restricted to observations with at least 25 patents in the 5 years from t to $t+4$. Standard errors are clustered at the firm level.

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

The IV regression with measurement error in X then yields:

$$\tilde{\beta}^{IV} = \frac{\text{Cov}(Y - \tilde{X}, Z)}{\text{Cov}(\tilde{X}, Z)} = (1+\beta)\frac{\text{Cov}(X, Z)}{\text{Cov}(X, Z) + \text{Cov}(\eta, Z)} - 1.$$

Thus, as long as the instrument is uncorrelated with the measurement error, $\tilde{\beta}^{IV} = \beta$. Otherwise, the estimate is biased downward for $\text{Cov}(\eta, Z) > 0$ and upward otherwise. The key question is therefore whether the measurement error covaries with the instrument, which depends on how detection probability varies with firm size.

Figure E.4: The Implication of Different Patent Scaling Relationships

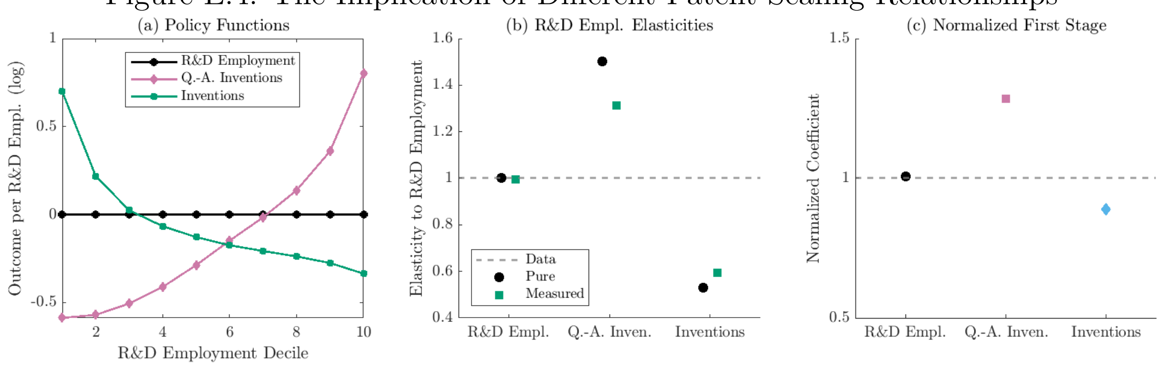

Notes: The left panel reports the number of quality-adjusted inventions $z_{kt}\, q_{kt}^N$, blueprints q_{kt}^N, and R&D employment ℓ_{kt} per R&D employee in the model against log R&D employment. The middle panel reports the elasticity of patents with respect to R&D employment in simulated data under each scaling assumption; the dashed red line reports the empirical value as documented in Table E.1. The right panel reports the first-stage coefficient normalized by its value without measurement error for each scaling assumption.

Figure E.4 illustrates the mechanisms. The left panel shows that the number of projects per inventor differs across scaling specifications, directly affecting each inventor's detection probability. If patenting increases one-for-one with R&D employment, detection probability per inventor is approximately constant across firm sizes, so $\text{Cov}(\eta, Z) \approx 0$ and the bias is small. Under blueprint scaling, inventions per inventor decrease with employment leading to a negative correlation between the measurement error and returns, deflating the first stage and biasing the IV estimates upward; under quality-adjusted scaling, the reverse holds. Panel (c) confirms this pattern: the first stage is approximately unaffected by measurement error under employment scaling but shifts significantly under the alternatives.

Figure E.5 reports the resulting coefficient estimates. Scaling with blueprints q^N_{kt} leads to an upward bias, while quality-adjusted scaling $z_{kt}\ q^N_{kt}$ yields a downward bias. Employment scaling implies only a small downward bias, reinforcing the conservative interpretation of the baseline results.

Figure E.5: Main Coefficient for Different Patent Scaling Relationships

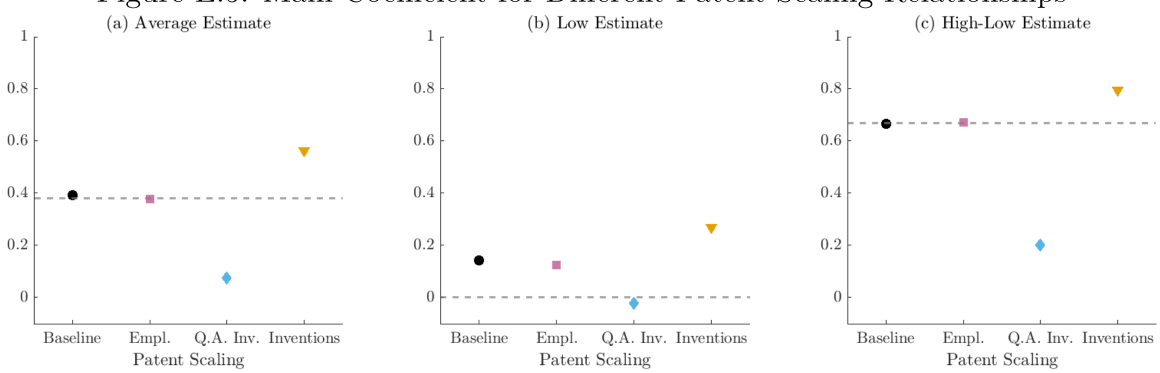

Notes: Figure reports main regression estimates for labor supply elasticities for different specifications of the patent scaling. Baseline values set the project success rate π to 1, implying no measurement error, while all other specifications set $\pi = 0.5$. Panel (a) reports estimates for the specification without interaction, while (b) and (c) report coefficients from the interaction regression.

F Alternative Labor Supply Elasticity Estimates

In this section, I compare my estimates with two additional sources of evidence on labor supply elasticities. Section F.1 calibrates the model to the estimates in Seegmiller (2025), who uses a similar research design with LEHD data, and shows that the baseline results are conservative. Section F.2 compares the implied markdown schedule with estimates derived from Kline et al. (2019), who use patent grants to infer labor supply elasticities, and reaches a similar conclusion.

F.1 Comparison with Seegmiller (2025)

Seegmiller (2025) exploits a similar research design to estimate the labor supply elasticity for all workers using LEHD data. While all workers are likely a flawed comparison group for inventors, he also reports estimates for highly skilled workers in Appendix Table IA.8, a group likely most comparable to inventors.[8] As discussed in the text, his estimates are consistently above the ones reported in the main table, i.e., they suggest more monopsony power. Nonetheless, I explore using these estimates as alternative targets to calibrate the model. An important distinction in the estimates is firm sorting: Seegmiller (2025) sorts by average product of labor, while my main estimates sort by R&D employment size. However, these sorting criteria are conceptually equivalent in my framework when applied to the market for inventors, since the return on R&D depends on the markdown, which is increasing in firm size.[9] I thus maintain his sorting, but target his estimates for 4 quartiles instead of my estimates for above and below median R&D employment. Table F.1 reports the calibrated parameters and targeted moments. The calibration matches Q1 precisely, but implies slightly more monopsony power than Seegmiller's targets for Q2 and Q3 while slightly undershooting Q4.

I next explore the impact of monopsony power on growth. The second column in Table D.2 reports the effect of setting monopsony power to zero in the Seegmiller-calibrated model, while the first column reports the baseline for comparison. As anticipated, monopsony power is much more detrimental to growth compared to the baseline: growth increases to 1.66 in its absence compared to 1.62 in the baseline. Similarly, the welfare costs rise significantly. Interestingly, the calibrated model suggests an even stronger role for misallocation: levels alone only boost growth by 0.03

[8]Seegmiller (2025) defines skill groups using AKM worker fixed effects, so the highly skilled category captures high earners regardless of occupation and may include workers outside of R&D.

[9]It is straightforward to show that the average product of labor in the model is given by

$$\frac{z_{kt}\; q_{kt}^{N}\; V_{\pi}}{W_{R,kt}\; \ell_{kt} + \psi_M\; M_{kt}} = \frac{1}{\gamma}\left(1 + s_{kt}\; \epsilon_{kt}\right),$$

where s_{kt} is the firm's labor share in R&D. Under the baseline calibration, this value is strictly increasing in ℓ_{kt}.

Table F.1: Calibrated Parameters and Targeted Moments — Alternative Targets

A. Internal Parameters	Symbol	Value	Method
Production labor supply shifter	α_P	0.29	Direct
R&D Labor supply shifter	α_R	0.07	Direct
R&D productivity intercept	μ	0.30	Direct
Non-listed firms R&D prod.	A_{nl}	0.07	Moment matching
R&D productivity variation	σ	0.69	Moment matching
R&D productivity autocorrelation	ρ	0.90	Moment matching
Returns to scale in R&D	γ	0.51	Moment matching
Labor Weight in R&D	α_L	0.93	Moment matching
Materials Price	ψ_M	0.30	Moment matching
Firm-level labor supply elasticity	ξ	3.74	Moment matching
Firm-level labor supply elasticity	$\underline{\ell}$	14.7	Moment matching
B. Moments	Target	Model	Source
Growth rate (%)	1.50	1.50	Standard value
Relative R&D size	35.0	35.0	NSF/ Compustat
SD of growth rate	0.29	0.29	Compustat
Autocorrelation of R&D	0.93	0.90	Compustat
Estimate Q1	0.49	0.35	S25, Table IA.8
Estimate Q2	0.68	0.68	S25, Table IA.8
Estimate Q3	0.80	0.93	S25, Table IA.8
Estimate Q4	1.35	1.35	S25,Table IA.8
Labor share	0.79	0.79	NSF
R&D Labor Force	0.05	0.05	AAABK18
Production Labor Force	0.29	0.29	AAABK18
Unit cost Elasticity	-1.00	-1.00	AAABK18

Notes: Table reports values for an alternative calibration targeting the labor supply elasticity estimates for high skilled workers in Seegmiller (2025), referred to as S25. AAABK refers to Acemoglu et al. (2018). Panel A reports calibrated parameters. Panel B reports moments together with target values.See Table 3 for external parameter values. See text for details.

pp, while reallocation yield a 0.13 pp increase. Thus, these results suggest that the baseline specification is conservative compared to the values implied by the estimates in Seegmiller (2025).

F.2 Comparison with Kline et al. (2019)

Kline et al. (2019), henceforth KPWZ, analyze how patent-induced shocks to firm performance propagate into worker compensation using a linkage of U.S. patent applications to Treasury business and worker tax records. Exploiting quasi-random variation in initial patent allowance decisions, they estimate how much of patent-generated surplus flows to workers at innovative firms and use this estimate to back out the implied labor supply elasticity faced by the firm. I find that my calibration is conservative relative to their estimates.

While KPWZ do not report pass-through rates separately by firm size, their Supplemental Appendix Table D.4 reports reduced-form impacts of initial patent allowances on both surplus per worker and wage bill per worker, split at median baseline employment (17 employees). For firms with ex ante valuable patents, an initial allowance raises surplus per worker by \$10,670 and wage bill per worker by \$2,160 at above-median firms, compared to \$13,280 and \$5,150 at below-median firms. Taking the ratio of these reduced-form impacts yields Wald estimates of the pass-through rate: $\hat{\pi} \approx 0.20$ at above-median firms and $\hat{\pi} \approx 0.39$ at below-median firms. KPWZ then show that one can back out the implied labor supply elasticity ϵ with an additional assumption on the demand elasticity η by inverting $\pi = (1 - 1/\eta)\,(1 + 1/\epsilon)^{-1}$ (see p. 1352). Following their approach and calibration of $\eta = 6$, their pass-through rate estimates imply labor supply elasticities of 0.32 and 0.87 for above and below median size firms.

This comparison is necessarily indirect for two reasons. First, KPWZ estimate a bargaining model in which incumbent workers extract rents proportional to their replacement costs; my framework features no such bargaining channel. Moreover, their identification strategy—exploiting quasi-random patent allowance decisions—would not identify the firm-level inventor supply elasticity in my setting: the grant of a patent does not alter the first-order condition for inventor hiring, so this variation would instead identify labor supply elasticities on the production side of the firm, which I model as perfectly competitive. The mapping from their pass-through estimates to monopsony elasticities therefore relies entirely on the structural equivalence that KPWZ themselves draw between η and a firm-specific labor supply elasticity. Second, KPWZ focus on small, first-time patentees drawn from the universe of U.S. tax-filing firms (median employment of 17), whereas my sample consists of Compustat-listed firms that are substantially larger. Even their "above-median" firms are small relative to typical firms in my sample. If the labor supply elasticity is decreasing in firm size, as both frameworks predict, the relevant comparison for firms in my sample would imply even lower elasticities than those derived above.

Figure F.1 plots the markdown schedule implied by my calibrated model against the two Kline et al. (2019) benchmarks. For completeness, I also report the value implied by their estimate for stayer earnings, which does not exhibit a size-link in their estimates. Compared to the full distribution of firms, my calibration implies markdown rates well above those suggested by the estimates in KPWZ across all firm-size deciles. For listed firms, which were not in their sample, markdowns in the top four deciles fall below the KPWZ stayer estimate. The top decile also falls below their estimate for below-median firms, while remaining above their estimate for above-median firms. These comparisons reinforce the interpretation that my welfare estimates represent a lower bound.

Figure F.1: Markdown in the Model and KPWZ

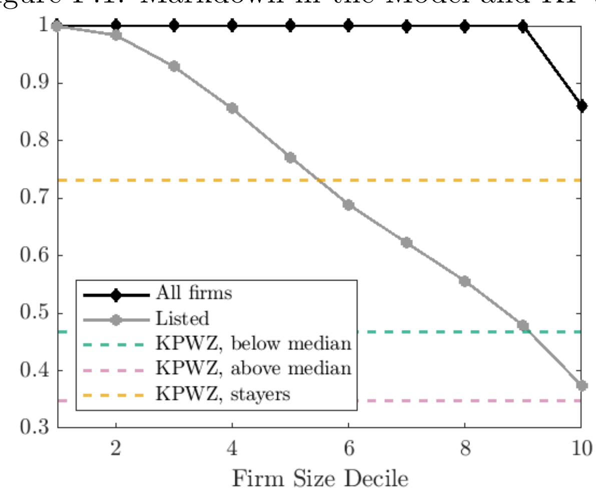

Notes: Figure reports markdowns $\epsilon/(1+\epsilon)$ by deciles of the firm size distribution for all and listed firms together with the implied estimates reported in Kline et al. (2019), referred to as KPWZ, assuming a demand elasticity of $\eta = 6$. See text for details.

G Additional Empirical Results

G.1 First Stage Empirical Results

I report the first stage results with and without controls for prior inventor wage and employment growth in Figure G.1. The baseline estimate ($t = 3$) suggests that a 10% increase in the firm valuation is associated with a 2% expansion of R&D employment for both specifications. R&D employment rises gradually over time, echoing estimates for regular workers in Seegmiller (2025).

Figure G.1: First Stage Estimates Across Specifications

(a) Baseline (b) Without Controls

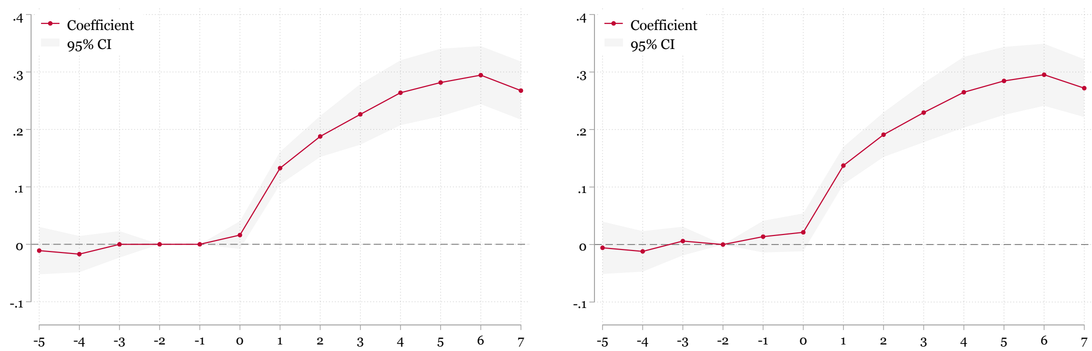

Notes: Panel (a) controls for inventor wage and employment growth from $t-2$ to $t-1$. Both control for NAICS3 $\times$ year fixed effects.

G.2 Robustness for Empirical Estimates

Table G.1 reports alternative estimates for specifications (17) and (18). Column (1) reports the baseline, while column (2) drops controls for prior inventor wage and employment growth. Column (3) replaces NAICS3 codes in the NAICS3$\times$Year fixed effects with an LDA-based firm classification leveraging their patents' CPC codes. Column (4) uses TFP growth in $t-1$ constructed from the TFP estimates in Imrohoroglu and Tuzel (2014) as an alternative instrument for R&D employment growth, while column (5) controls for inventor productivity growth. Columns (6)-(8) use alternative measures of inventor employment and columns (9) and (10) cap and extend the estimation window by 1 year.

Table G.1: Robstuness for Inverse R&D Labor Supply Elasticity Estimates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
A. Baseline	**R&D Wage Growth**									
R&D Employment Growth	0.382**	0.294**	0.441**	0.251	0.389**	0.381**	0.385**	0.378**	0.286**	0.226*
	(0.162)	(0.142)	(0.177)	(0.255)	(0.158)	(0.160)	(0.162)	(0.159)	(0.126)	(0.116)
R&D Prod. Growth					0.393***					
					(0.091)					
First stage F stat.	68	76	63	18	72	74	63	60	93	78
B. Interaction with Size	**R&D Wage Growth**									
R&D Employment Growth	0.003	-0.055	-0.030	-0.139	0.000	-0.033	-0.004	-0.169	0.086	-0.079
	(0.129)	(0.112)	(0.129)	(0.198)	(0.131)	(0.130)	(0.130)	(0.142)	(0.096)	(0.120)
— × Above Median R&D Empl.	0.666***	0.560***	0.667***	0.463***	0.673***	0.655***	0.670***	0.516**	0.630***	0.571***
	(0.222)	(0.200)	(0.206)	(0.173)	(0.218)	(0.219)	(0.220)	(0.213)	(0.192)	(0.174)
R&D Prod. Growth					0.309***					
					(0.073)					
No Controls		✓								
Alt. Instr.				✓						
LDA FEs			✓							
Alt. inventors	All	All	All	All	All	US only	FTE	Verified	All	All
End Year	t+3	t+3	t+3	t+3	t+3	t+3	t+3	t+3	t+2	t+4
First stage F stat. (Main)	40	43	42	9	42	44	37	51	53	44
First stage F stat. (Inter.)	27	26	22	13	29	27	26	34	32	29
Observations	12,729	12,755	12,916	10,593	12,566	12,692	12,729	11,905	12,712	11,953

Note: Robustness table for alternative specifications. Unless stated otherwise, regressions control for inventor wage and employment growth from $t-2$ to $t-1$ and NAICS3 $\times$ year fixed effects. F statistics reported are based on Sanderson and Windmeijer (2015). Standard errors are clustered at the NAICS6 level.

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

G.3 Calculating the Labor Share in R&D

I calculate the labor share in R&D for the US in 2000 and 2019 using the "All industries" data reported in the 2000 Survey of Industrial Research and Development (SIRD), which was conducted by the Division of Science Resources Statistics within the National Science Foundation (NSF), and the 2019 Business Enterprise Research and Development Survey (BERDS), which was conducted by the National Center for Science and Engineering Statistics (NCSES) and Census Bureau. In both cases, I first calculate the attributable R&D costs, which excludes undefined costs and includes imputed opportunity cost for capital, and then report the share of labor costs. For the 2000 figures I make a range of adjustment to capture costs that are reported in detail in 2019, but lumped into an "Other" category in 2000. These adjustments are based on the 2019 values reported for these categories and detailed in the footnotes of Table G.2.

As reported in Table G.2, the labor share of attributable R&D costs was 79% in 2019. The remainder of the costs is split between "materials and equipment" and capital, where the former tends to be more important. Notably, the labor share in R&D costs is significantly higher than the labor share in the US overall, which is typically reported around 67% (Autor et al., 2020). Hence, R&D is a very labor-intensive task, justifying the focus on labor markets in R&D.

Table G.2: National Labor Share in R&D

	2019
A. Raw R&D costs [% thereof]	
Raw R&D cost	493.0
R&D wages and benefits	268.0 [54.4%]
Stock-based compensation	39.0 [7.9%]
Temporary staffing	21.4 [4.3%]
Materials and supplies	34.4 [7.0%]
Royalties and licensing fees	9.2 [1.9%]
Expensed equipment	7.2 [1.5%]
Lease and rental payments	8.2 [1.7%]
Depreciation	18.9 [3.8%]
Other	86.6 [17.6%]
B. Attributable R&D cost	
Raw R&D costs	493.0
– Other	- 86.6
+ Imputed cost of capital	9.4
Attributable R&D costs	415.8
C. Attributable cost shares [% thereof]	
Materials and equipment	50.9 [12.2%]
Capital	36.5 [8.8%]
Labor	328.4 [79.0%]

Notes: Values in Panel A are taken from the source noted in the text except those marked with *, which are imputed. Labor-related values are imputed to keep constant their relative size to R&D wages and benefits. Other values are imputed to keep constant their relative size to overall R&D. The "Other" category is adjusted such that the individual items add up to raw R&D cost. Panel B calculates attributable R&D costs as raw R&D cost minus other cost plus cost of capital. The latter are imputed as 50% of depreciation, which is in line with an interest rate of 7.5% and depreciation rate of 15%. The final panel categorizes R&D costs into materials and equipment, capital, and labor. Materials and equipment includes materials and supplies, royalties and licensing fees, and expensed equipment. Capital includes depreciation, lease and rental payments, and imputed cost of capital. Labor includes R&D wages and benefits, stock-based compensation, and temporary staffing.

References

Acemoglu, Daron, 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.

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.

Cadman, Brian D., Tjomme O. Rusticus, and Jayanthi Sunder, "Stock option grant vesting terms: Economic and financial reporting determinants," *Review of Accounting Studies*, 12 2013, *18*, 1159–1190.

Dube, Arindrajit, Laura Giuliano, and Jonathan Leonard, "Fairness and frictions: The impact of unequal raises on quit behavior," *American Economic Review*, 2 2019, *109*, 620–663.

Hazell, Jonathon, Christina Patterson, Heather Sarsons, and Taska Bledi, "National Wage Setting," 2025.

Imrohoroglu, Ayse and Selale Tuzel, "Firm-Level Productivity, Risk, and Return," *Management Science*, 2014, *60*, 2073–2090.

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

Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar, "Who Profits from Patents? Rent-Sharing at Innovative Firms," *The Quarterly Journal of Economics*, 2019, *134*, 1343–1404.

Lerner, Josh and Julie Wulf, "Innovation and Incentives: Evidence from Corporate R&D," *The Review of Economics and Statistics*, 11 2007, *89*, 634–644.

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

Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter, "Firming Up Inequality," *The Quarterly Journal of Economics*, 2 2019, *134*, 1–50.