THE MACROECONOMIC IMPLICATIONS OF R&D OWNERSHIP*

Shaoshuang Yang†

Click for Update

Abstract

This paper examines how R&D ownership affects inventor incentives, innovation allocation, and aggregate productivity. I develop a general equilibrium model of endogenous growth addressing how inventors self-select into firms of different sizes because the equilibrium equity-wage contract depends on how much firm value is tied to the inventor's own effort. The model predicts that (1) when firms own R&D, incentives are diluted and inventors exert less effort, especially in large firms; (2) when inventors retain a larger share of R&D, they work harder; and (3) the inventor distribution across firm sizes depends systematically on the ownership regime. A counterfactual exercise shows how alternative U.S. innovation ownership policies would reallocate innovation and alter long-run growth. Quantitatively, R&D ownership has a nonlinear impact on the aggregate growth rate. When inventor R&D ownership improves from the baseline 2% to 10%, the growth rate increases by 0.66%. When the inventor holds 50% of the R&D ownership, the growth rate improves by 53.54%.

JEL codes: O31, O34, O41, L11.

Keywords: R&D ownership, inventor incentive, endogenous growth, innovation allocation, firm dynamics, principal-agent friction

*I am deeply indebted to my advisors, Ayse Imrohoroglu, Pablo Kurlat, David Zeke, and Gerard Hoberg for enormous support. I am also grateful to Yueyuan Ma, Thomas Chaney, Monica Morlacco, David Lagakos, Francisco Buera, Erzo Luttmer, and participants at AEA, SEA and Dynare conference for helpful comments This paper is previously circulated as "The Distribution of Innovations across Firms"

†Affiliation: The Chinese University of Hong Kong, Shenzhen, zip code 518172, China. Tel: +86 755 84218898. Email: yangshaoshuang@cuhk.edu.cn.

1 Introduction

> *"[Offering inventors exclusive rights to their invention outcome] added the fuel of interest to the fire of genius, in the discovery and production of new and useful things."*
>
> — Abraham Lincoln (Basler 1953, p. 363)

Innovation is a driving force of economic growth, yet R&D production inherently suffers from a principal-agent problem (Merges, 1999), which arises between the legal reality of employer-owned R&D and the economic importance of employee-inventor incentives. Its severity depends on the ownership of the R&D outcome (Aghion and Tirole, 1994). The tension raises questions that remain understudied: How does the allocation of R&D ownership between employers and inventors affect the aggregate growth rate of the economy, and how does it shape the distribution of innovation across heterogeneous firms? This study provides a general equilibrium framework connecting the ownership of R&D outcomes to firm dynamics and aggregate growth. I use this framework to examine how policies governing employee-inventors' R&D ownership affect innovation through their interaction with inventor incentives.

This paper is motivated by two stylized facts regarding innovation output and inventor payoff. First, inventors' statutory R&D ownership rights positively affects their output. When institutional environments shift and inventors lose ownership of their inventions, their productivity declines. Second, offering inventors compensation that ties their income to firm performance increases innovative output.

Motivated by these facts, I build an endogenous growth model in which the innovation ownership regime dictates the equilibrium contract between inventors and firms. This, in turn, determines the employee-inventor's incentive to exert effort, drives the self-selection of inventors into firms of different sizes, and alters aggregate growth and the distribution of innovations. I calibrate the model to match empirical moments on innovation ownership, inventor output heterogeneity, and the growth rate. Using this calibrated framework to analyze counterfactual policies, I show that the interaction between R&D ownership and inventor incentives is essential for understanding firm dynamics and economic growth. The counterfactual exercise shows that, when inventors are granted 50% direct ownership of their innovation output, inventors move to large firms and the growth rate rises from 2% to 3%, implying the importance of micro-level contracting frictions. Both the aggregate and the distributional effects found in the counterfactual exercise are consistent with their empirical counterparts.

In the model, each inventor is endowed with an idea of idiosyncratic quality (a la Abowd et al. (1999)) and chooses whether to participate in the R&D sector or to provide labor as a worker. If she chooses to participate as an active R&D inventor, she next chooses to join a firm of her preferred size. The inventor allocation across firms determines the distribution of innovation and the aggregate growth rate.

At the core of the theoretical framework is a principal-agent problem driven by the unobservable nature of an inventor's effort. The intensity of the incentive problem increases in firm size. The total incentive for an inventor to exert effort is a combination of two elements: her legal share of direct R&D ownership and the indirect share from the equity-wage contracts offered by her employer. However, offering an equity share forces the firm to navigate a trade-off between improving the inventor's incentives and exposing the inventor to the broader volatility of the business. The effectiveness of equity contracts decreases with firm size, and hence so does the optimal equity share offered. In a small firm, a single inventor's idea accounts for a large fraction of the firm's total value, so the firm's equity is tightly tied to the outcome of her innovation. This makes equity a highly effective incentive device, and the firm optimally offers a larger equity share to the inventor. As a result, inventors are more strongly incentivized in smaller firms. In contrast, the equity of a large firm is tied to many unrelated activities. Consequently, large firms are structurally less effective at using equity to motivate talent, leading to weaker overall incentives and lower innovative effort.

Counterbalancing this incentive disadvantage, large firms possess an advantage in the form of synergy. The expected economic value of an inventor's idiosyncratic idea is amplified when implemented by a larger firm, as they can leverage their established production scale and broader knowledge base to maximize the commercial impact of the new technology. The synergy is particularly valuable for implementing higher-quality ideas, though at the same time, working in a larger firm increases the inventor's management cost. Consequently, the framework features an endogenous selection effect: inventors self-select into firms of different sizes trading off the high-powered equity incentives of small firms, the synergy of large firms, and the management cost. Inventors evaluate this tension alongside the prevailing legal ownership regime. When the ownership regime is more pro-employer, inventors rely more on equity to align incentives, so they opt for smaller firms where equity is more effective. Additionally, inventors make participation decisions by trading off the payoff from supplying labor against the payoff from supplying ideas. When facing greater agency friction, the payoff of acting as an active R&D employee drops. Hence, more potential inventors choose to supply labor instead. This agency friction drives an endogenous self-selection of inventive talent across different firm types, dictating not only the aggregate rate of endogenous growth

but also the concentration of innovative activity across the firm-size distribution.

The model yields three key predictions. First, when the firm retains 100% R&D ownership, compared with the optimal R&D ownership split schedule, the firm optimally offer equity stakes to the inventor to compensate for the inventor's lost direct ownership. However, this offset is incomplete. Incentives are diluted and the inventor's effort decreases, especially in large firms. Second, as the inventor's direct R&D ownership rises, the firm reduces the equity stake it offers. Because equity is a diluted substitute for direct ownership, this offset is incomplete, and the inventor's effective incentives strengthen on net, raising effort. Third, the assignment of inventors across firm types depends systematically on the ownership regime. When firm ownership increases, inventors opt for smaller firms, where equity provides stronger incentives that compensate for the loss of direct ownership. Collectively, these predictions rationalize the stylized facts described above and yield testable implications.

The model's central mechanism operates through a specific contract space: firms incentivize inventors using a combination of fixed wages and equity stakes. This structure is not directly observable in standard datasets, but the assumption generates testable implications. When statutory R&D ownership is exogenously shifted toward employers, inventor output should decline. This decline should be smaller in small firms, where equity remains a potent incentive device because an individual inventor's innovation accounts for a large share of firm value. I test this prediction by exploiting a quasi-natural experiment that exogenously transferred R&D ownership from employee-inventors to employers. Using a difference-in-differences design, I find that the ruling reduced inventors' patenting output, with the negative effect systematically mitigated in small firms across measures of patent quantity, quality, and quality-weighted output, consistent with the implication of the contract space assumption.

To quantify how R&D ownership affects economic growth and the distribution of innovation across firms, I calibrate the model and run counterfactual experiments that reassign R&D ownership between inventors and firms. The model benchmark is calibrated to the United States, where 98% of R&D ownership is assigned to employers under the baseline statutory regime. In the counterfactuals, I hold all parameters fixed except the inventor ownership share. I then compute the change in the annual growth rate and in the share of innovations developed in small firms, defined as firms with no mroe than 500 employees. I find that shifting R&D ownership from firms back to inventors leads to higher economic growth, consistent with the data. This effect on the growth rate increases nonlinearly with the magnitude of the shock. In the baseline economy (characterized by 98% firm ownership), the growth rate is calibrated to 2%. When inventors are assigned a 10% R&D ownership

share, the growth rate increases to 2.03%, translating to a 0.66% proportional increase in aggregate growth relative to baseline. When 50% of R&D ownership is assigned to inventors, the growth rate increases to 3.09%, translating to a 53.54% proportional increase in aggregate growth relative to baseline.

Meanwhile, this aggregate gain is accompanied by the reallocation of innovations between small and large firms. When inventors hold 10% ownership of their R&D outcome, the share of innovations produced by small firms drops by 38.81%. When inventors hold 50% of R&D ownership, all innovations shift to large firms. This quantitative dynamic aligns with the direction of the empirical facts: as the exogenous legal shift toward employer ownership disproportionately depressed innovation within large firms in the data, restoring inventor ownership in the model disproportionately benefits large firms. The quantitative analysis thus highlights a trade-off governed by legal institutions: while higher inventor ownership promotes aggregate growth through better incentives and greater commercialization synergy, it heavily concentrates innovative activities within large firms, suppressing small-firm innovation dynamics.

I further decomposes the results sequentially into three channels: a contract channel, an inventor reallocation channel, and a participation channel. In the exercise, I sequentially add one channel at a time. The contract channel isolates the endogenous adjustment of equity-wage contracts and the corresponding change in unobservable effort. It always has a positive impact on the aggregate growth rate, regardless of the direction of the R&D ownership shift. The inventor reallocation channel allows inventors to endogenously relocate across the firm-size distribution. It captures how ownership regimes alter the sorting of inventors across firms to maximize the expected value of their utility. In the 10% ownership case, the contract channel accounts for 20.45% of the aggregate growth improvement, while adding the reallocation channel contributes an additional 37.12%. The participation channel captures inventors' occupation choice, where an enhanced R&D payoff induces an expansion of the active R&D workforce. Adding the participation channel contributes 42.42% in the 10% ownership case. This decomposition demonstrates that R&D ownership regimes affect the growth rate and the distribution of innovation by altering individual motivation, shifting inventive talent across firms, and changing the size of the R&D workforce.

Additionally, I also explore the optimal R&D ownership allocation when both firms and inventors can credibly commit to transfer R&D ownership ex ante. In this case, the R&D ownership allocation is determined endogenously between each firm-inventor pair. Inventors trade off between incentive, synergy and management cost. The aggregate growth doubles, increasing from 2% to about 4%. This result underscores the importance of R&D ownership for the aggregate growth rate. The findings imply that R&D policies, such as R&D subsi-

dies directed at employers, may be blunted if the underlying legal architecture fails to pass incentives through to the actual inventors.

Related Literature This paper develops a macroeconomic framework of endogenous growth that explicitly models the internal incentive structures driving the allocation of innovation across firms and the aggregate growth rate. Traditionally, the endogenous growth literature treats the firm as a unitary decision-maker when modeling innovation dynamics (Klette and Kortum, 2004; Phillips and Zhdanov, 2013; Cunningham et al., 2021; Buera and Oberfield, 2020; Perla et al., 2021). While an emerging strand of research explores the market for ideas (Silveira and Wright, 2010; Chatterjee and Rossi-Hansberg, 2012; Chiu et al., 2017) and the specific role of inventors (Bhaskarabhatla et al., 2021; Akcigit et al., 2025), the frictions arising from incomplete contracting between inventors and their employers remain underexplored in general equilibrium. By highlighting this incentive channel both empirically and theoretically, this paper offers a new perspective on how micro-level contracting frictions aggregate to determine economy-wide growth.

The core mechanism of this paper is an agency friction affecting individual innovative productivity. This dynamic intersects with the well-documented decline in innovation productivity (Bloom et al., 2020). While recent studies rationalize this decline through the strategic, defensive behavior of incumbent firms protecting their market power (Akcigit and Goldschlag, 2023), this paper proposes a complementary channel rooted in agency frictions. This reconciles two stylized facts: declining R&D productivity and the rise of superstar firms (Autor et al., 2020). This paper shows that when firms retain the majority of R&D ownership, the difficulty of providing equity-based incentives intensifies with firm size, dampening inventors' productivity in large firms.

To empirically and theoretically ground the agency frictions, this study formalizes the role of ownership in innovation. The theoretical literature establishes that assigning residual control rights to the inventor is crucial for incentivizing inventor effort, because employment contracts are incomplete (Grossman and Hart, 1986; Aghion and Tirole, 1994; Hart, 1995). The importance of R&D ownership is also supported by empirical studies examining the relationship between inventor incentives and innovation outcomes (Lerner and Wulf, 2007; Onishi, 2013; Giarratana et al., 2018; Hvide and Jones, 2018; Alok and Subramanian, 2023; Suh, 2023). However, this microeconomic literature has primarily focused on partial equilibrium outcomes without tracing the macro-level consequences. This paper embeds this micro-mechanism into a macroeconomic model, providing a framework to evaluate the aggregate implications of R&D ownership.

Building on this ownership structure, the model introduces a mechanism through which inventors self-select into firms of different sizes. This is related to the classic literature on firm

boundaries (Coase, 1937; Hart and Moore, 1990, 2008) and the literature on the allocation of R&D. Recent macroeconomic contributions have explored how firms' size and scope affect their innovation strategies (Akcigit et al., 2016; Akcigit and Kerr, 2018; Ma, 2022; Jo and Kim, 2024). This paper provides an additional angle by demonstrating that the distribution of innovation across the economy is also driven by inventors sorting between small and large firms, depending on the R&D ownership regime.

Finally, this study informs the debate regarding R&D policy. Standard growth models typically treat firms as the representative innovating agent determining R&D inputs (Romer, 1990; Klette and Kortum, 2004), leading to predictions that R&D subsidies yield large effects (Akcigit et al., 2025). However, empirical evidence suggests that the actual effects of R&D policy are often minimal (Goolsbee, 1998; Wilson, 2009). This paper provides an explanation for this discrepancy. This paper argues that the limited effectiveness of R&D subsidies stems from a misalignment between policy and inventor incentives. When subsidies are directed to employers under pro-employer ownership regimes, the lack of incentive pass-through severely blunts the macroeconomic effectiveness of innovation policy.

The rest of the paper is organized as follows. Section 2 reviews the institutional background governing R&D ownership. Section 3 summarizes stylized facts in the literature. Section 4 constructs a model where inventors make decisions about joining firms of specific sizes conditional on the ownership regime. I derive predictions about inventor productivity and the allocation of inventors across firms. Section 5 calibrates the model. Section 6 conducts quantitative exercises to evaluate the growth and distributional effects of shifts in the R&D ownership regime. Section 7 discusses the assumptions and extensions of this paper. Section 8 concludes.

2 Institutional Background

The incompleteness of inventor employment contracts makes the allocation of R&D ownership rights crucial for determining inventors' innovative output. Given the complexity of innovative activity, employment contracts cannot fully anticipate or describe the various dimensions of inventor effort. Furthermore, the rich interactions between a specific innovation's implementation and all of a firm's other activities makes it difficult for a contract to fully capture the possible outcomes of an inventor's effort ex ante. These frictions imply that the allocation of R&D ownership plays a crucial role in determining the severity of the agency problem. This section outlines the relevant institutional background.

To evaluate how R&D ownership shapes innovation output, it is first necessary to define the economic boundaries of an R&D outcome. The knowledge generated during the

innovation process is embedded both in patents and in the inventor herself, in the form of tacit know-how. Consequently, an inventor's incentive to exert effort depends on the joint ownership of both the patent and human capital, a composite I define as *R&D ownership*. In the United States, this effective R&D ownership is governed by two sets of legal instruments. The formal ownership of patents is determined by the interplay between common law, preinvention assignment agreements (PIAAs), and state-level statutes (Alok and Subramanian, 2023). The ownership allocation of human capital is governed by non-compete agreements (NCAs), which restrict inventor mobility and effectively determine whether the inventor can realize the human-capital component through mobility (Garmaise, 2011).[1]

Common law clearly assigns patent ownership only at two extremes—explicit employer mandates and fully independent inventions — leaving the vast majority of innovative activity in a legal gray area (Parker, 1983). PIAAs resolve this ambiguity by obligating employees to assign all rights and interests in any inventions made or conceived during the employment term to the employer (Dratler Jr, 1979). In the United States, PIAAs are widely adopted, especially among innovating firms, which require prospective employees to sign one as a precondition for employment (Pisegna-Cook, 1993; Merges, 1999).[2] The enforcement of a PIAA implies that essentially all inventions belong to the employer, granting the firm the legal apparatus to appropriate even those innovations that are entirely unrelated to the employment relationship. The combination of its expansive scope and universal adoption dictates that the PIAA governs the ownership of patents.

Recognizing the expansive reach of PIAAs, eight U.S. states enacted pro-employee "free-to-create" statutes in the late 1970s and early 1980s (Gullette, 1980; Coolley, 1985), permitting inventors to retain ownership of inventions made on their own time and unrelated to the employer's business. Separately, NCAs restrict inventor mobility, transferring effective control over the human capital generated by the innovation process from the inventor to the firm. Together, these four legal instruments define the inventor's effective R&D ownership, which is the joint control over both the patent and the human capital embedded in the inventor. PIAAs and free-to-create statutes govern the patent component; NCAs govern the inventor's residual claim on the tacit know-how component. The remainder of the paper takes this effective share as a single statutory primitive, denoted ϕ in the model of Section 4.

[1] A detailed discussion of each legal instrument is provided in Appendix A.2.

[2] Although one might naturally expect that the ideal ownership setting depends on the inventor-firm pair, the standard PIAA is widely adopted. As a result, nearly all innovations are automatically assigned to the employer. Section D provides a discussion of this.

3 Stylized Facts

Before developing the general equilibrium framework, this section synthesizes a set of stylized facts from the empirical literature that establish a relationship between property rights and inventor productivity. Across diverse institutional and geographic settings, empirical evidence consistently demonstrates a pervasive incentive channel: stripping inventors of residual control rights systematically depresses their R&D effort and innovative output. These stylized facts motivate the macroeconomic model of endogenous growth developed in this paper, which incorporates the incentive structures governing innovation.

Fact 1: Inventors' statutory R&D ownership rights positively affect their output.

The legal allocation of R&D ownership to employee-inventors directly shapes innovation outcomes. Evaluating an institutional reform in Norway, Hvide and Jones (2018) document that a legal shift in invention ownership from researchers to universities caused a severe decline in innovative output. Following a policy in 2003 that transferred two-thirds of intellectual property rights to universities, the patenting rates of researchers fell by approximately 50%, accompanied by significant drops in patent quality and university startup rates, the latter being a primary channel for commercializing innovation.

Beyond the academic context, Alok and Subramanian (2023) evaluate variation in state-level legislation across the United States and find consistent state-level evidence within the corporate sector. Exploiting the passage of state laws that restricted employers' claims on employee-inventors' inventions and strengthened inventors' residual property rights, the authors find that pro-employee legal shocks boost corporate innovation. Patent counts rise by 14%, while patent quality also improves. Acharya et al. (2014) also corroborates this statutory incentive channel.

Fact 2: Compensation schemes that tie inventors' income to firm performance increase innovative output.

When statutory regimes favor employers, firms can rely on alternative compensation mechanisms to bridge the incentive gap. However, the inventor employment contracts are incomplete. First, inventor effort cannot be fully described. Second, firms cannot credibly commit ex ante that they will compensate an inventor ex post. Therefore, contracts contingent on effort or specific innovation outcomes are not feasible. Employers instead rely on indirect incentive tools such as revenue-based compensation and stock grants to improve employee-inventors' performance. Analyzing Japanese firm data, Onishi (2013) finds that while Japanese law requires firms to offer "reasonable remuneration" in exchange for appropriating an employee's invention, corporate surveys reveal that firms deploy these

compensation plans primarily to "enhance inventor motivation" beyond what a fixed wage can achieve. Indeed, implementing revenue-based compensation significantly increases an inventor's patenting output, and this effect is more pronounced in small firms.

This effect is confirmed by firm and inventor survey data. Lerner and Wulf (2007) use European survey data and reach a similar conclusion. The paper demonstrates that compensating R&D heads with incentives which ties the employees' income to firms' performance, rather than fixed wages, systematically increases quality-weighted patent output. Furthermore, Giarratana et al. (2018) highlight the importance of tying compensation directly to an inventor's own innovative outcome. This compensation structure increases inventor effort, measured by self-reported interactions with other researchers and time devoted to research and development. This higher effort, in turn, drives the increase in realized innovative outcomes.

Taken together, these stylized facts indicate that both statutory R&D ownership and compensation tied to innovation value positively affect innovative output, and this relationship holds across diverse datasets and environments. The consistent direction of these findings motivates the theoretical framework developed in Section 4, which formalizes the incentive channel through which statutory R&D ownership and endogenous employee-employer contracts jointly determine innovative effort and aggregate economic growth.

4 Model

I build a continuous-time model of endogenous growth in which firms compete to attract inventors for innovation. The goal of this model is to study how R&D ownership affects the distribution of innovation across firms and aggregate growth. The economy features a representative household, short-lived inventors, a competitive final good sector, and a continuum of heterogeneous intermediate firms. Upon drawing an innovative idea of idiosyncratic quality, an inventor optimally selects a firm of her preferred size to develop her idea into an invention. The central mechanism of the model hinges on a principal-agent friction: because innovative effort is unobservable, the effective incentive to innovate is jointly governed by the exogenous statutory allocation of R&D ownership and the firm's optimal design of an equity-wage contract.

4.1 Preferences

A representative household provides labor and consumes final goods. The household is long-lived and supplies one unit of labor inelastically with one unit of human capital. The

household's preferences are given by a utility function:

$$U_H = \int_0^\infty e^{-\rho t} \log(C_H(t)) dt, \tag{1}$$

where $\rho > 0$ is the discount rate and $C_H(t)$ denotes the consumption of final goods of the household at time t.

A continuum of risk-averse inventors of unit measure is born at each instant. Inventors are short-lived, existing only for an infinitesimal interval dt. An inventor exerts unobservable effort $e_I(t)$ to produce innovations within intermediate firms and consumes final goods. Inventor preferences are specified using a mean-variance utility function defined over stochastic consumption $c_I(t)$ and effort $e_I(t)$:

$$U_I\left(c_I(t), e_I(t), q(t)\right) = \mathbb{E}\left(c_I(t)\right) - A\frac{\mathrm{Var}\left(c_I(t)\right)}{V(t|\bar{q}(t))} - R\left(e_I, q, \bar{q}(t)\right) V(t|\bar{q}(t)), \tag{2}$$

where c_I is consumption, $A > 0$ is the degree of risk aversion, and $R(e_I, q, \bar{q}(t))V(t|\bar{q}(t))$ represents the associated disutility cost of effort, where

$$R(e_I, q, \bar{q}(t)) = \frac{1}{1+\delta} e_I^{\delta+1} + \eta \frac{q}{\bar{q}(t)}. \tag{3}$$

It combines two types of cost: an effort cost that increases in effort, and a management cost that increases in firm size[3]. $\bar{q}(t)$ (defined below) is the average quality in the economy at time t, and $V(t|\bar{q}(t))$ is the value of a firm whose size equals the average quality. The variance, the effort, and management costs are normalized by $V(t|\bar{q}(t))$ to ensure the problem is stationary.

In addition to engaging in risky innovation, an inventor has the outside option of entering the labor market alongside the representative household, supplying ξ units of effective human capital. This alternative yields a reservation utility of $\bar{U}_I(t) = c_I^L(t)$, where $c_I^L(t)$ denotes the inventor's consumption when supplying labor. In equilibrium, a share $l_I(t)$ of inventors will opt to supply labor, while $(1 - l_I(t))$ of inventors develop R&D. Finally, aggregating across the unit measure of agents, the total consumption of the inventor cohort is given by $C_I(t) = \int_0^1 c_{Ii}(t) di$, where i indexes individual inventors.

4.2 Technology

The economy features two types of firms: intermediate goods producers and final goods producers, as in Akcigit and Kerr (2018). Both types of firms are owned by the household. Intermediate goods producers hire inventors to develop innovations and produce intermediate goods. Final goods producers assemble intermediate goods into final goods.

The final goods producers use a continuum of intermediate goods of measure 1 to pro-

[3]In equilibrium, firm size is linear in its quality q. The proof is in Section 4.7.

duce final goods, with a production technology similar to that in Akcigit and Kerr (2018):

$$Y(t) = \frac{1}{1-\beta} \int_0^1 q_j^{\beta}(t) y_j^{1-\beta}(t) dj. \tag{4}$$

In this function, $q_j(t)$ is the quality of the intermediate good j for $j \in [0,1]$, and $y_j(t)$ is its quantity. I normalize the price of the final good to one in every period. The final good producers are perfectly competitive, taking input prices as given. Henceforth, I will drop the time index t when it does not cause confusion.

The final goods are consumed by the household and inventors. The resource constraint of the economy is:

$$Y = C_H + C_I. \tag{5}$$

The intermediate goods producers form a continuum of risk-neutral firms with unit measure. Each firm produces one type of good, with a using a linear, labor-only technology:

$$y_j(t) = \bar{q}(t) l_j(t), \tag{6}$$

where $l_j(t)$ is the firm's total human-capital-adjusted labor input. Each household member contributes 1 unit of effective human capital, while each inventor who supplies labor contributes ξ units. $\bar{q} = \int_0^1 q_j dj$ is the average quality level of the economy. The labor cost per firm is $w(t) l_j(t)$, where the wage $w(t)$ is taken as given. Labor market clearing requires

$$\int_0^1 l_j dj \leq 1 + \xi l_I(t). \tag{7}$$

where $\xi l_I(t)$ is the total effective human capital supplied by inventors at time t, implying that a share of $l_I(t)$ inventors choose to supply labor instead of working in the R&D sector.

This paper focuses on the balanced growth path. I normalize the variables using the average quality $\bar{q}$. I denote the normalized variables using a tilde:

$$\tilde{q}_j \equiv \frac{q_j}{\bar{q}}. \tag{8}$$

4.3 Innovation

The innovation production function follows the standard formulation in the endogenous growth literature (Romer, 1990; Klette and Kortum, 2004; Akcigit and Kerr, 2018). In the literature, the firm both creates and implements innovations; in this paper, however, the inventor creates the innovation while the legal instruments determine who captures the commercial value.

Each inventor is born with one idiosyncratic idea of quality $z \sim \Psi(z)$. To transform this idea into an innovation, the inventor must affiliate with an intermediate goods producer of her choosing. Innovation raises the intermediate goods producer's quality q_j. It is costly to work, and the flow cost of choosing effort e_I is $R(e_I, q, \bar{q}) V(\bar{q})$, which combines an effort cost and a management cost. Thus, without incentives, the inventor chooses $e_I = 0$. Moreover, the inventor can always costlessly produce a sham innovation that does not improve the

firm's quality. Whether an innovation is genuinely useful is not verifiable to outsiders.

The innovation process is stochastic. The inventor's chosen effort level governs the arrival rate of a successful innovation according to a Poisson process:

$$\lambda(e_I) = \mu e_I, \tag{9}$$

where $\mu > 0$ is the baseline efficiency of the R&D technology. The inventor's effort e_I is unobservable and unverifiable by the firm (Aghion and Tirole, 1994; Merges, 1999).

If the inventor's effort successfully yields an innovation, the allocation of the resulting property rights is governed by the statutory regime.[4] The statutory regime therefore *exogenously* determines the share ϕ of total innovation value that the inventor captures, with the remaining $(1 - \phi)$ share retained by the firm. Furthermore, the framework assumes no commitment regarding effort or innovation outcomes. The combination of unobservable effort and statutory split generates a principal-agent friction. In response to this friction, this paper assumes that all firms employ a common contract structure that combines a fixed wage for risk-sharing with an equity component for incentive alignment.

Once a useful innovation is realized, the quality improvement $x(z, \tilde{q})\bar{q}$ is jointly determined by the idiosyncratic quality of the idea, z, and the size of the commercializing firm, $\tilde{q}(t)$. Specifically, it takes the following form:

$$x(z, \tilde{q})\bar{q} = \gamma z \bar{q}, \tag{10}$$

where γ can take two values, $\gamma \in \{\gamma_H, \gamma_L\}$ with $\gamma_H > \gamma_L$, corresponding to a high and low realization, respectively. The probability of the high state takes the logistic form

$$h(\tilde{q}(t)) = \frac{1}{1 + \exp\left(-b_0 - b_1 \log(\tilde{q})\right)}, \tag{11}$$

where $b_1 > 0$. Denote the expected quality improvement by $x_0(z, \tilde{q}) = \mathbb{E}\left[x(z, \tilde{q})\right]$.

An innovation generates a higher expected value when implemented within a larger firm. This captures a synergy effect: large firms possess better commercialization scale, complementary assets, and broader production capacities, allowing them to extract a larger aggregate payoff from a given technological improvement than a smaller firm could. The model assumes that each inventor makes decisions individually. When a firm hires more than one inventor, the inventors act sequentially so that their innovation outcomes do not interact. The quality improvement is scaled by $\bar{q}$ to ensure that innovation step sizes grow proportionally with the aggregate technology stock, preserving balanced growth. It captures how the total technology stock in the economy affects new innovations.

[4]The institutional determinants of patent assignment are PIAAs and free-to-create statutes, and the determinants of the inventor's residual claim on tacit know-how are NCAs. Both are discussed in Section 2.

4.4 Information Structure

After an inventor is born, her type z is publicly known. Then the inventor chooses which firm to work for and her effort level e_I. The effort e_I is unobservable and unverifiable. Hence, contracts cannot be contingent on the effort level.

For each innovation, whether it is genuinely useful is not verifiable to outsiders, meaning that the quality improvement $x(z, \tilde{q})\bar{q}$ is private information. This reflects the difficulty outsiders face in assessing an innovation's exact value (Chatterjee and Rossi-Hansberg, 2012). Therefore, hiring contracts cannot be contingent on either effort or innovation outcomes.

4.5 Employment Contracting Problem

Intermediate firms engage in Bertrand competition to hire inventors by offering employment contracts. The employment contract consists of two parts: a wage rate $T_I V(t|\bar{q}(t))$, and a share of equity $a \in [0, 1]$. In addition to her contractual equity stake a, the inventor retains a share ϕ of the R&D output she produces which is mandated by the legal regime and does not require a separate contract.

4.5.1 Discussion on Contract Space

One natural question is whether an employment contract can be contingent on ϕ? This is not feasible in the model, since ϕ is set jointly by laws and other legal instruments. In the baseline model, neither the firm nor the inventor can credibly promise to transfer a share of ownership to the other party, since they can always obtain the statutory share through a lawsuit ex post. Section 6.3 considers a counterfactual in which the firm and the inventor can negotiate over ϕ. Section D discusses why this negotiation does not occur in practice.

One may think that a better tool for providing incentives would be to offer a bonus contingent on the delivery of an innovation, since it would avoid exposing inventors to unrelated shocks. However, it is not feasible. Bonus offers require payoffs to be contingent on the delivery of an innovation, which is not feasible due to lack of commitment. On the one hand, a firm cannot credibly commit to pay a bonus ex post. The firm's statutory $(1 - \phi)$ share is already enforceable. The firm has no incentive to make additional payments to the inventor once an innovation has been delivered. On the other hand, the inventor cannot commit to deliver a genuinely useful innovation, since the innovation's value cannot be proven to an outsider.[5] Inventors have the incentive to earn bonuses effortlessly by providing useless

[5]Neither the stock reaction, nor the patent transaction can be used to evaluate innovation value. First, public firms only account for less than 5% of all innovative firms in my sample period, which means that most firms do not have access to the stock market, let alone using it to evaluate innovation value. Even for

innovations. Therefore, using a bonus cannot solve the moral hazard problem.

A firm might attempt to draft a detailed contract that links the inventor's innovation to changes in firm profit. I will follow the incomplete contract literature (Grossman and Hart, 1986), which assumes that it is very difficult to specify all innovation outcomes in contracts, and, in practice, it is usually not contractible (Acemoglu, 1996; Frésard et al., 2020); I assume that firms can only use equity and wages in employment contracts, as commonly seen in the real world (Brickley and Hevert, 1991).

4.5.2 Contracting Problem

The timing is as follows. At birth, each inventor observes the initial quality z of her idea. Firms observe the distribution of value $x(z, \tilde{q})$ that they would generate from a successful innovation developed by the inventor z. All firms simultaneously extend personalized employment offers to inventors. After evaluating all offers, the inventor either joins her preferred firm with relative quality $\tilde{q}^*(z)$ or takes the outside option and supplies labor.

Upon recruiting an inventor, the firm's quality follows a stochastic path:

$$dq = -qdJ_{\text{exit}} + x(z, \tilde{q})\bar{q}dJ_{R\&D}(e_I), \tag{12}$$

where J_{exit} and $J_{R\&D}(e_I)$ represent two Poisson jump processes with arrival intensities τ and $\lambda(e_I) = \mu e_I$, respectively. The former captures exogenous exit shocks, while the latter represents successful internal R&D developed by the inventor. The magnitude of an innovation's impact is given by $x(z, \tilde{q})\bar{q}$, which denotes the stochastic value added by an inventor of quality z within a firm of relative size $\tilde{q}$ scaled by the aggregate technology level $\bar{q}$.

The inventor's compensation $c_I(a, T_I, q, z, \phi, \bar{q}$ is defined as:

$$\begin{aligned} c_I(a, T_I, q, z, \phi, \bar{q})dt =& T_I V(\bar{q})dt + \phi X(q, z)dJ_{R\&D}(e_I) \\ &+ a\left(dV(j|q, z) - \phi X(q, z)dJ_{R\&D}(e_I)\right) \end{aligned} \tag{13}$$

where $dV(j|q, z)$ is the firm equity value change from t to $(t + dt)$ when hiring an inventor of quality z. In the second element, $X(t|q, z)$ denotes the realized value of an innovation produced by an inventor of type z in firm q at time t:

$$X(q, z) = V(j|q + x(z, \tilde{q})\bar{q}) - V(j|q),$$

which is not verifiable for an outsider. Here, $V(j|q)$ denotes the firm's value if the innovation does not arrive. Consistent with a principal-agent structure, beyond the statutory ϕ-share that the inventor retains, the contract bundles a fixed wage with an incentive components

public firms, it is not easy to use stock market returns to identify innovation value. One reason is that many events, other than innovations, also affect stock prices. Another reason is that, for most public firms, multiple innovations are granted on the same day, which makes it impossible to assign value to each patent separately. Admittedly, for traded patents, their market value can be estimated using the price information. However, most firms do not innovate for the purpose of selling—they innovate to improve their own productivity. That is, many innovations are never sold, and there is no price information available.

which allows the inventor to acquire a fraction a of the firm's equity at market value.

The second line of the compensation equation captures the inventor's net equity payoff. Because the inventor already captures the ϕ share of innovation, the equity share a applies only to the firm's residual claim. Equity thus exposes the inventor to firm-level risk (including exit shocks and the firm's $(1-\phi)$ share of innovation outcomes) in exchange for incentive alignment. To fund consumption, the inventor liquidates her equity and receives her innovation holdings at the conclusion of every dt interval. In doing so, the inventor bears part of the risk from exit shocks and innovation shocks. For a firm j with quality q, the optimal contract for an inventor of type z solves:

$$\begin{aligned}\Omega_F(z,q) = \max_{a,T_I} &(1-a)(1-\phi)\mu e_I X(q,z) - T_I V(\bar{q}) \\ \text{s.t. } e_I = \arg\max &\left\{U_I\left(c_I(a,T_I,q,z,\phi,\bar{q}),e_I\right\}\right. \\ &U_I\left(c_I(a,T_I,q,z,\phi,\bar{q}),e_I\right) \geq \max\{\bar{u}(z),\bar{U}_I\} \\ &(1-a)(1-\phi)\mu e_I X(q,z) - T_I V(\bar{q}) \geq 0\end{aligned} \tag{14}$$

where $\Omega_F(z,q)$ represents the firm's expected return. The instantaneous payoff to the firm's existing shareholders is the residual innovation value, after accounting for the inventor's ϕ share, minus the costs of the wage and the equity portion a granted to the inventor. The optimization is subject to three conditions. First, the incentive-compatibility constraint ensures that the inventor selects an effort level e_I that maximizes their expected utility $U_I(c_I, e_I)$ given the contract. Second, the participation constraint requires the inventor's utility to meet a reservation threshold $\bar{u}(z)$, which is determined endogenously through Bertrand competition and the outside option. Finally, the firm's participation constraint ensures that hiring the inventor remains profitable for the original shareholders.

Inventor's utility function can be expressed as:

$$\begin{aligned}U_I\left(c_I(a,T_I,q,z,\phi,\bar{q}),e_I\right) =& \mathbb{E}\left(c_I(a,T_I,q,z,\phi,\bar{q})\right) - A\frac{\mathrm{Var}\left(c_I(a,T_I,q,z,\phi,\bar{q})\right)}{V(\bar{q})} \\ &- R\left(e_I,q,\bar{q}\right)V(\bar{q}).\end{aligned} \tag{15}$$

An inventor then chooses an optimal firm size $\tilde{q}^*(z)$ to maximize her utility, conditional on working in the R&D sector.

Then the inventor decides whether to work in the R&D sector or supply labor. Her decision rule is:

$$u(z) = \max\{U_I\left(c_I(z,q^*(z),\phi),e_I(z,q^*(z),\phi),\right),\bar{U}_I\}. \tag{16}$$

4.6 Entry and Exit

Intermediate firms face an exogenous exit rate τ, which is independent of their size and is a risk unrelated to innovation.[6] After a firm exits the market, the technology it owns is distributed randomly to new entrants. Entry of new intermediate firms occurs exogenously at rate τ. The new firm incurs a cost equal to the value of the technology associated with its drawn quality, q_I. At any instant, there is always a unit measure of firms in the economy.

4.7 Equilibrium

I now characterize the equilibria of the economy. I focus on the balanced growth path on which aggregate variables $(Y, C_H, C_I, w, \bar{q})$ grow at the constant rate g, and other variables, such as l_I, stay constant.

4.7.1 Household's and Firms' Choices

The household maximization problem reveals the relationship between the growth rate, g, and the time discount factor,

$$\frac{\dot{C}_H}{C_H} = r - \rho. \tag{17}$$

The final goods producer's problem can be written as:

$$\max_{\{y_j\}} \frac{1}{1-\beta} \int_0^1 q_j^{\beta} y_j^{1-\beta} dj - \int_0^1 y_j p_j dj. \tag{18}$$

The first-order condition

$$p_j = q_j^{\beta} y_j^{-\beta}$$

yields the demand function for goods produced by intermediate firms, $j \in [0, 1]$.

Intermediate good producers engage in monopolistic competition, optimizing their profit by choosing l_j, p_j, y_j, given the wage level w:

$$\begin{aligned} &\max_{l_j, p_j, y_j} y_j p_j - w l_j. \\ &s.t.\ y_j = \bar{q} l_j \\ &\quad p_j = q_j^{\beta} y_j^{-\beta} \end{aligned} \tag{19}$$

The FOC yields:

$$y_j = q_j \left(\frac{\bar{q}(1-\beta)}{w} \right)^{\frac{1}{\beta}}, l_j = \frac{y_j}{\bar{q}}, p_j = \frac{w}{\bar{q}(1-\beta)}. \tag{20}$$

[6] A modest negative correlation between exit rates and firm size would not affect the model's qualitative predictions. Allowing the exit risks in small firms to be mitigated by R&D outcomes also does not affect the model's predictions.

The labor market clears, which derives that

$$\frac{\int_0^1 q_j \left(\frac{\bar{q}(1-\beta)}{w}\right)^{\frac{1}{\beta}} dj}{\bar{q}} = 1 + \xi l_I.$$

Therefore, the wage is:

$$w = \frac{(1-\beta)}{(1+\xi l_I)^{\beta}} \bar{q}. \tag{21}$$

Both the production output y_j and profit π_j are linear in quality:

$$y_j = (1+\xi l_I) q_j, \pi_j = \beta (1+\xi l_I)^{1-\beta} q_j. \tag{22}$$

Using the equilibrium quantity from Equation (20) and the wage rate from Equation (21), the aggregate output is a linear function of the average quality $\bar{q}$:

$$Y = \frac{(1+\xi l_I)^{1-\beta}}{1-\beta} \bar{q} \tag{23}$$

The HJB equation of a firm's value is

$$\begin{aligned} rV(j|q) - \dot{V}(j|q) &= \pi_j + \Omega_F(z,q) - \tau V(j|q) \\ &= \pi_j - \tau V(j|q) \end{aligned} \tag{24}$$

The second line follows from the scarcity of inventors and the assumption that competition for their ideas ensures that inventors capture the full value of innovations. Since $\pi_j = \beta(1+\xi l_I)^{1-\beta} q_j$, the value of an intermediate firm is linear in its quality, i.e.,

$$V(j|q) = \nu q. \tag{25}$$

Solving the innovation problem shows that the inventor's effort, e_I, is independent of time when (25) holds. The inventor's consumption, $c_I(a, T_I, q, z, \phi, \bar{q})$, and utility, $U_I(c_I, e_I)$, are both linear in ν. The inventor's optimal firm choice is invariant to time and depends only on parameters.[7]

4.7.2 Aggregate Variables

The economy's final output is allocated to the consumption of both households and inventors. By the law of large numbers, the household's total consumption $C_H(t)$ satisfies:

$$C_H = w + \Pi, \tag{26}$$

where Π denotes aggregate profit from the production sector.

Similarly, aggregating across active inventors yields total inventor consumption, C_I:

$$C_I = \nu \bar{q} \int_{z \in \{z | U_I(c_I, e_I) > \bar{U}_I\}} \mu e_I\left(z, \tilde{q}^*\right) x_0\left(z, \tilde{q}^*\right) d\Psi(z) + w \xi l_I. \tag{27}$$

Innovation is the sole source of growth in this economy. The aggregate growth rate is given by:

$$g = \frac{\dot{\bar{q}}}{\bar{q}} = \int_{z \in \{z | U_I(c_I, e_I) > \bar{U}_I\}} \mu e_I\left(z, \tilde{q}^*\right) x_0\left(z, \tilde{q}^*\right) \, d\Psi\left(z\right) \tag{28}$$

[7] See Appendix B.8 for a formal proof.

4.7.3 Equilibrium

This section characterizes an equilibrium of the economy in which aggregate variables $(Y, C, w, \bar{q})$ grow at a rate g.

Definition An equilibrium of this economy is the mapping between $\tilde{q}^*$ and z, the allocation $\left(\{y_j^*\}_j, Y^*, C_I^*, C_H^*\right)$, the prices $\left(w^*, r^*, \{p_j^*\}_j\right)$, the growth rate g^*, and the share of inventors supplying labor l_I, such that (1) for any $j \in [0,1]$, y_j^* and p_j^* satisfy Equation (20); (2) the wage w^* satisfies Equation (21); (3) the interest rate r^* satisfies Equation (17); (4) the measure of human capital supplied by inventors l_I and the mapping between $\tilde{q}^*$ and z come from inventors' optimal choices; (5) R&D spending C_I^* satisfies Equation (27); (6) the aggregate output Y^* satisfies Equation (23); (7) the aggregate consumption C_H^* satisfies Equation (5); and (8) the growth rate g satisfies Equation (28).

4.8 Closed form

To provide analytical clarity and highlight the primary economic mechanisms, this section introduces a tractable framework that yields a closed-form characterization of the inventor's decision-making process. A key refinement in this setup concerns the realized innovation quality, $x(z, \tilde{q})\bar{q}$, which is assumed to follow a distribution with mean $x_0(z, \tilde{q})\bar{q}$ and second moment $\mu e_I^{-1} \left(\mathbb{E}(\gamma) z\bar{q}\right)^2$, rather than the distribution described in Equation (10). This modification effectively represents a setting where increased inventor effort mitigates innovation-related risk. The variance of the inventor's consumption is given by:

$$\mathrm{Var}\left(c_I\left(a, T_I, \tilde{q}, z, \phi, \bar{q}\right)\right) = \nu^2 \bar{q}^2 \left(\left(a(1-\phi)+\phi\right)^2 x_0\left(z, \tilde{q}\right)^2 + a^2 \tau \tilde{q}^2\right).$$

The first element is the uncertainty related to the innovation, while the second is the uncertainty caused by the firm's exit risk. In this special case, innovation-related uncertainty no longer depends on the effort level.

The parameter δ in the effort cost function, $R(e_I, q, \bar{q})$, is set to 1, consistent with empirical estimates in the literature (e.g., Akcigit and Kerr (2018)). Applying backward induction, firms anticipate the inventor's optimal effort response:

$$e_I^* = \mu\left(a(1-\phi)+\phi\right) x_0\left(z, \tilde{q}\right).$$

Stronger incentives either through higher equity share, a, or larger inventor R&D ownership share ϕ, induces greater inventor effort. Similarly, a higher expected innovation quality, $x_0(z, \tilde{q})$, also induces greater inventor effort. In both cases, the marginal return to effort rises, making additional effort more attractive.[8]

[8] The derivation of the optimal compensation package is provided in Appendix B.1.

The firm's problem in Equation (14) yields:

$$a^*(\tilde{q}, z, \phi) = \max\left\{ \frac{1 - \frac{2A\phi}{(1-\phi)}}{1 + 2A(\frac{\tau\tilde{q}^2}{(1-\phi)^2\mu^2 x_0(z,\tilde{q})^2} + 1)}, 0 \right\}. \tag{29}$$

When inventor ownership satisfies $\phi < \frac{1}{1+2A}$, the optimal equity share a^* is always positive. Otherwise, the optimal equity share is zero. It means that the inventor has sufficiently large statutory ownership. Additional incentives are not necessary.

Lemma 1. [9] *When $b_1 \in (0, 1]$, the optimal equity decreases in firm size,*

$$\frac{\partial a^*(\tilde{q}, z, \phi)}{\partial \tilde{q}} < 0.$$

When $b_1 \in (0, 1]$, the optimal equity share a^* decreases monotonically with firm size $\tilde{q}$. This relationship is governed by the tension between two competing effects: the commercialization potential $x_0(z, \tilde{q})$ and the non-innovative shock, represented by the exit risk $\tau\tilde{q}^2$. Firm size influences these two components in opposing ways. Larger firms offer greater synergy that boosts $x_0(z, \tilde{q})$, making it more valuable for the firm to incentivize effort, pushing a^* up. Meanwhile, in larger firms, the exit risk takes a larger share of the total variance. Thus, reducing the equity stake helps insulate the inventor from shocks unrelated to her innovative performance, pushing a^* down. The overall effect depends on the dominant channel. When $b_1 \leq 1$, as in the calibration, the exit-risk channel dominates, leading larger firms to offer lower equity shares.

Lemma 2. [10] *The optimal equity a^* decreases with the inventor R&D ownership ϕ. When $\phi < \frac{1}{1+2A}$, the optimal equity a^* strictly decreases with the inventor R&D ownership ϕ. When $\phi \geq \frac{1}{1+2A}$, the optimal equity level is always zero.*

The reason firms compensate inventors with equity is to provide incentives. The equity transfers an a^* fraction of the firm-owned $(1 - \phi)$ R&D ownership to the inventor, which complements the share ϕ that the inventor directly owns. When ϕ increases, two forces jointly affect a^*. First, the inventor already owns a larger share of R&D, which decreases the demand for additional incentives. Second, the effectiveness of equity also drops, since the firm retains a smaller share of R&D. Both channels reduce the optimal equity share a^* when ϕ increases.

Lemma 3. [11] *The inventor effort level e_I^* increases in her effective ownership $(a^*(1-\phi)+\phi)$.*

The total incentives an inventor receives combine direct R&D ownership (ϕ) with the equity-implied R&D ownership $(a^*(1 - \phi))$, giving an effective ownership of $(a^*(1 - \phi) + \phi)$.

[9] See Appendix B.2 for the proof.
[10] See Appendix B.3 for the proof.
[11] See Appendix B.4 for the proof.

An inventor optimally chooses her effort level based on the total incentive. When the inventor is better incentivized, she chooses a higher effort level.

Proposition 1. [12] *When firms retain ownership, compared with the ownership case which generates the optimal incentive ($\phi = \frac{1}{1+2A}$)*[13]*, incentives are diluted and inventors exert less effort, especially in large firms.*

When firms retain ownership, without equity incentive, the inventor's output is disconnected from her unobservable effort input. To incentivize inventors, the firm must offer them equity. However, the equity incentive is diluted, compared to the direct ownership of the R&D outcome, since it is related to both an inventor's own R&D result and all other activities in the firm. Thus, the inventor optimally exerts less effort. In large firms, when the synergy does not grow too fast, the innovation takes a smaller share of firm value. Equity, as an incentive tool, is even more diluted. Therefore, this effect is more salient in large firms.

Proposition 2. [14] *Suppose $\phi < \frac{1}{1+2A}$, so that the optimal equity share is strictly positive ($a^* > 0$). There exists a cutoff $\bar{\phi}$:*

(i) If $\phi \geq \bar{\phi}$, the optimal effort level $e_I(z, \tilde{q})$ is monotonically increasing in the inventor R&D ownership ϕ for all $\tilde{q}$.

(ii) If $\phi < \bar{\phi}$, there exists a firm-size threshold $\tilde{q}_0(z, \phi)$ such that $e_I(z, \tilde{q})$ is increasing in ϕ for $\tilde{q} \geq \tilde{q}_0(z, \phi)$ and decreasing in ϕ for $\tilde{q} < \tilde{q}_0(z, \phi)$.

When the inventor's direct R&D ownership ϕ drops, the firm's optimal equity offer a^* increases (Lemma 2). The net effect on the effective share depends on which adjustment dominates. When the inventor owns a sufficiently large share of R&D ($\phi \geq \bar{\phi}$), the inventor is already partially incentivized. Therefore, risk-sharing considerations dominate incentive provision. Since the firm's optimal choice of a^* is constrained by the inventor's risk-bearing capacity, the equity-implied share $a^*(1-\phi)$ rises by less than the fall in ϕ. As a result, the inventor's effective ownership drops as ϕ falls, her payoff is less tightly linked to her own effort, and she optimally reduces her effort level.

When inventors own a smaller share ($\phi < \bar{\phi}$), the result depends on the composition of firm-level risk. In larger firms ($\tilde{q} \geq \tilde{q}_0$), non-innovative risk takes a sufficiently large share of total risk, so equity loads substantial unrelated variance on the inventor and the logic above still applies: both the effective share and effort move with ϕ. In small firms ($\tilde{q} < \tilde{q}_0$), non-innovative risk is small relative to innovation risk, so equity loads almost only

[12] See Appendix B.5 for the proof.
[13] The optimal R&D ownership is when the equity incentive becomes redundant.
[14] See Appendix B.6 for the proof.

innovation variance on the inventor and is a near-perfect substitute for direct ownership. The firm therefore has room to push a^* up aggressively when ϕ falls — both to compensate for the lost direct stake and because the larger retained share (1-ϕ) makes equity more potent (Lemma 2). The combined adjustment more than offsets the fall in ϕ, raising the inventor's effective ownership. The negative relationship between firm ownership and inventor effort is therefore mitigated, and may reverse, at the bottom of the firm-size distribution.

Proposition 3. [15] *The assignment of inventors across firm types depends systematically on the ownership regime.*

When inventors make their decision to join a firm, they trade off synergy, incentives, and cost. Large firms provide better synergy, but offer weaker incentives and impose higher cost. The ownership regime determines the degree of incentive, the intensity of the agency friction, and the importance of equity compensation. Since the effectiveness of equity incentives differs across firms, a shift in the ownership regime also changes the allocation of inventors across firms. Specifically, compared with the optimally incentivized case ($\phi = \frac{1}{1+2A}$), under firm ownership ($\phi \to 0$), it is easy to show that inventors opt for smaller firms.

4.9 Size-dependent reaction to an ownership shift

The assumption that the contract space is consist of a fixed wage and equity implies a testable prediction: if statutory R&D ownership is shifted from inventors to employers, aggregate innovative output will decline, and the decline will be systematically mitigated or even reversed within small firms[16]. Appendix A examines this prediction using an exogenous legal shock. The analysis exploits a quasi-natural experiment generated by a shift in U.S. federal case law that exogenously transferred R&D ownership from employee-inventors to employers. By estimating the heterogeneous response of inventor output across the firm-size distribution, the empirical results are consistent with the model prediction.

5 Calibration

I calibrate the full model to target the average US economy from 1980 to 2020. Patents are used as a proxy for innovations. An innovation's idea quality, denoted by z, and the realized value, $x\nu\bar{q}$, correspond to the patent's citations (scientific importance) and the patent's economic value, respectively. Additionally, I assume that the idea quality z follows the log-normal distribution characterized by parameters (z_m, σ).

[15] See Appendix B.7 for the proof.

[16] A quantitative example of this heterogeneous effect is shown in Section 6.

5.1 Identification

The parameters of the quantitative model are divided into two categories. The first group comprises parameters calibrated *a priori* based on direct aggregate statistics or established values from the macroeconomic literature. The second group consists of parameters that are jointly estimated internally to match specific moments in the microdata. Table 1 summarizes the first group of parameters: $\{\rho, \beta, \tau, \mu, A, \delta, \phi\}$. The household's subjective discount rate, ρ, is set to 0.02 to match the average real interest rate observed during the sample period. The elasticity of intermediate goods quality in final production, β, is set to 0.109, following Akcigit and Kerr (2018). The exogenous firm exit rate, τ, is calibrated to 0.09, targeting the average annual firm exit rate reported in the U.S. Census Bureau's Business Dynamics Statistics (BDS). The scale parameter of the innovation arrival rate, μ, is normalized to 1.[17] Regarding the curvature of the utility and cost functions, the coefficient of absolute risk aversion, A, is set to 0.5, and the effort cost elasticity, δ, is set to 1, both of which align with standard calibrations in the macro-labor literature (Hall and Van Reenen, 2000).

The employee-inventors' statutory ownership share is calibrated to $\phi = 0.02$, implying that firms retain a residual baseline ownership of 98%. This parameterization is grounded in comparative employee-invention law across countries. In Germany, a jurisdiction renowned for its statutory protection of employee-inventors, the legal minimum for inventor ownership is formally set at 3%. The U.S. lies below this threshold because of the widespread enforcement of pre-invention assignment agreements (PIAAs) and non-compete agreements (NCAs), which prevent the inventor from capturing the patent and jointly compress her ability to realize the human-capital share. Consequently, assigning an inventor ownership share of 2% places the U.S. baseline below the German statutory minimum.[18]

Table 1: Parameter Values from a Priori Information

Parameter	Description	Value	Identification
ρ	Discount rate	0.02	Interest rate
β	Production function quality share	0.109	Firm profitability
τ	Exo. exit rate	0.09	BDS
μ	Innovation arrival rate	1	Normalization
A	Risk aversion	0.5	Risk aversion
δ	Effort cost elasticity	1	Effort cost elasticity
ϕ	Inventor ownership share	0.02	Legal ownership

Notes: This table shows parameter values from the literature or direct estimation.

[17] In the model, μ enters only multiplicatively, either as $\mu x_0(z, \tilde{q})$ or as $\mu \mathbb{E}\left[x_0(z, \tilde{q})^2\right]$, which implies that μ is not separately identified from $x_0(z, \tilde{q})$.

[18] Robustness checks demonstrate that the model channels remain stable across alternative baseline calibrations of ϕ ranging between 0 and 0.05.

This paper jointly calibrates the eight remaining parameters, $(b_0, b_1, \gamma_H, \gamma_L, z_m, \sigma, \xi, \eta)$, using the minimum distance method inspired by Lentz and Mortensen (2008). The parameters, along with their corresponding moments, are shown in Table 2.

Table 2: Parameters from the Minimum Distance Estimation

Para.	Description	Identification
b_0	Constant of logistic function	Share of high realization in an average-sized firm
b_1	Scale factor of logistic function	Coefficient of regression (30)
γ_L	Low realization of innovation value	Normalize to 1
γ_H	High realization of innovation value	High-to-low economic value ratio
z_m	Log mean of idea quality distribution	Aggregate growth rate
σ	Dispersion of idea quality distribution	SD/Mean of inventor wages
ξ	Inventor human capital	Share of innovations in small entities
η	Scale of management cost	SD/Mean of inventor total earnings

Notes: Parameters in this table are jointly calibrated to minimize the distance between the model and data moments.

Probability of a high realization and synergy in an average-sized firm—In the model, the probability that an innovation achieves a high-value realization is governed by a logistic synergy function, parameterized by a baseline constant b_0 and a scale factor b_1. The realization value of an innovation is inferred from stock market reactions to patent grants, using a linkage between U.S. patents and CRSP–Compustat data, following the methodology of Kogan et al. (2017). Empirically, a high realization is defined as an innovation whose economic value exceeds the median value within its specific grant-year and citation-decile cohort. To identify the scale factor b_1, which captures how rapidly commercialization synergy improves with firm size, the following logistic regression is estimated:

$$\text{logit}(P(y_{pjt} = 1)) = \alpha + b_1 \ln(\text{size}_{jt}) + \mathbf{X}_t + \varepsilon_{pjt}, \tag{30}$$

where y_{pjt} is an indicator equal to one if patent p, granted to firm j in year t, achieves a high realization value. $\text{logit}(P(y_{pjt} = 1))$ is the log-odds of the associated probability. Firm size, size_{jt}, is measured separately using total revenue and total assets in two specifications, while $\mathbf{X}_t$ absorbs year fixed effects. Across all specifications, the estimated coefficient b_1 is positive, yielding values of 0.138 and 0.171, respectively. The calibration targets the scale factor of $b_1 = 0.17$.

The baseline constant, b_0, determines the probability of a high realization in an average-sized firm, given by $(1 + \exp(-b_0))^{-1}$. In the patent–firm matched public-firm data, the raw share of high-realization patents for the mean firm is 0.62. However, publicly traded firms in the Compustat universe are right-skewed and systematically larger than the average innovating firm in the broader economy, so using the public-firm mean would overstate baseline synergy. To correct for this size distortion, this paper also calculates the high-

realization share of the median public firm, which ranges from 0.46 to 0.51 across the two firm-size measures. The baseline calibration targets a 50% high-realization share for the average firm.[19]

High to low economic value ratio—In the model, the low and high realizations of innovation value cannot be separately identified, as they always appear in multiplicative form with z. The low-realization state γ_L is normalized to 1. The high state is calibrated to the value ratio of high-realization patents and low-realization patents. Using the stock-market-reaction-based innovation values described above, the relative value of a high-realization patent compared to a low-realization patent within the same grant-year and citation-decile cohort is 17.84.

Growth rate—Innovation is the sole driver of growth in the model. The log mean of the idea quality distribution, z_m, is central to the distribution of inventor ideas and thus to aggregate growth. A higher z_m increases the step size of innovation, raising the overall growth rate. I calibrate z_m so that the model-implied aggregate growth rate equals 2.0%, matching the average annual U.S. growth rate in GDP per capita.

Wage dispersion—In the model, the inventor wage, T_I, is determined in equilibrium as a share of the expected innovation payoff. Wage dispersion is mainly driven by the dispersion of idea quality, parameterized by σ. When idea quality is more dispersed, the wage dispersion also increases. Therefore, σ is identified by targeting the empirical standard deviation-to-mean ratio of inventor wages, whose average is 0.95 (Toivanen and Väänänen, 2012).

Share of innovations in small firms[20]—The parameter ξ defines the inventor's baseline human capital and dictates the value of the outside option. This reservation utility is more likely to bind for an inventor with a low idea quality. In the calibrated baseline, the model features positive assortative matching between inventor quality and firm size. Hence, a higher ξ raises the reservation utility, inducing more low-quality inventors to exit the R&D sector. Since low-quality inventors are matched to small firms, the share of innovations conducted in small firms decreases as ξ rises. Therefore, ξ is identified by matching the aggregate share of innovations occurring within small entities. In the data, this is calculated using the share of patents granted to small entities, which is 0.197.

Earning dispersion—In the model, the inventor's earnings, c_I, includes both wage and equity payouts. The management cost scale parameter, η, introduces an additional cost for inventors who work in larger organizations. While the baseline variance of fixed compen-

[19]The model mechanisms remain robust when targeting alternative high-realization shares ranging from 25% to 62%.

[20]A small firm is defined as a business entity with no more than 500 employees, following the USPTO.

sation is driven by the underlying dispersion of talent, the management cost η scales the management cost that grows with firm size. By altering the cost of joining large organizations, η determines the sorting of inventors, which reshapes the variance of stochastic equity payouts. Hence, η is identified by targeting the earnings dispersion, which is 1.03 (Akcigit and Goldschlag, 2025).[21]

5.2 Calibration Results

Table 3 reports the model-generated moments and their counterparts in the data. Overall, the model matches the targeted moments closely. The resulting parameter values are reported in Table 4. The quantitative analysis shows that, on average, active inventors in the baseline economy accept offers with an average equity share of 58.37%. This indicates that firms provide substantial additional incentives when the statutory stake is at the baseline level. However, the effectiveness of this equity compensation varies across the firm-size distribution. The calibration demonstrates that the optimal equity stake offered to the average inventor strictly decreases with firm size, as shown in Figure 1. In large firms, equity is heavily diluted by unrelated activities and therefore becomes a less efficient incentive device for the risk-averse inventor. Because the calibrated management cost parameter η is small, inventors' firm-size choice is governed primarily by the trade-off between incentives and synergy.

Table 3: Moments

Identification moment	Data	Model
Share of high realization in an average-sized firm	0.50	0.50
Coefficient b_1 in regression (30)	0.17	0.17
High-to-low economic value ratio	17.84	17.84
Aggregate growth rate	0.020	0.020
SD/Mean of inventor wage	0.95	0.95
Share of innovations in small entities	0.20	0.20
SD/Mean of inventor total earning	1.03	1.01

Notes: This table compares the moments generated from the calibrated model and the data. In general, the model-generated moments match the data well.

[21] The estimation is explained in Appendix C.1.

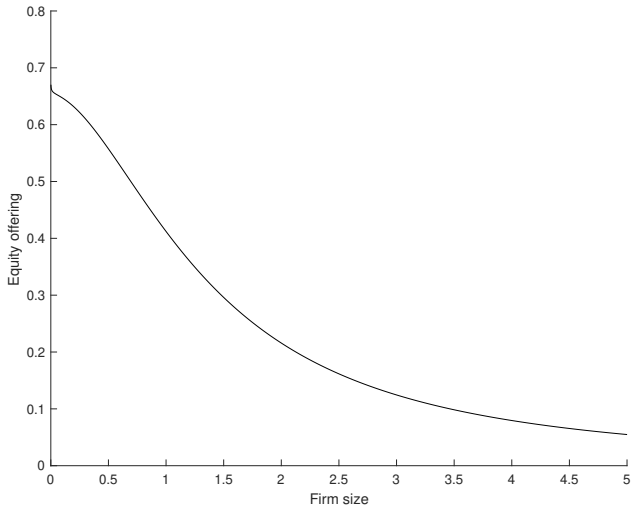


Figure 1: Equity offered to an average inventor

Table 4: Estimated Parameter Values

Parameter	Description	Value
b_0	Constant of logistic function	0
b_1	Scale factor of logistic function	0.17
γ_L	Low realization of innovation value	1
γ_H	High realization of innovation value	17.84
z_m	Log mean of idea quality distribution	-3.89
σ	Dispersion of idea quality distribution	0.52
ξ	Inventor human capital	0.012
η	Scale of management cost	5.6e-4

Notes: All parameters except γ_L (normalized to 1) in this table are jointly calibrated to minimize the distance between the model and data moments.

6 Quantitative Analysis

This section conducts counterfactual experiments to evaluate the macroeconomic implications of alternative R&D ownership regimes. Holding all other parameters constant, the analysis simulates economies where the statutory property rights of inventors are exogenously shifted from the baseline (2%) to 0%, 10%, and 50%. In the first case, the firm holds all intellectually property right. The remaining two counterfactual thresholds are calibrated to reflect the statutory guidelines of the German Employee Inventions Act (Gesetz

über Arbeitnehmererfindungen), a legal framework renowned for explicitly mandating R&D ownership allocation based on an employee's role and reliance on corporate resources. For instance, a 10% inventor share characterizes a research employee executing assigned tasks using substantial firm resources, whereas a 50% share reflects a highly educated employee (e.g., holding a university or engineering degree) who independently identifies and develops an innovation with minimal firm support. The results of these policy experiments are summarized in Table 5, with Panel A reporting absolute equilibrium levels and Panel B detailing percentage changes relative to the baseline.

Table 5: Counterfactuals under Alternative R&D Ownership Regimes

	Baseline	**Counterfactual Inventor Ownership Regimes**		
Moment	(2% Ownership)	0% Ownership	10% Ownership	50% Ownership
Panel A: Absolute Levels				
Aggregate Growth Rate	2.01%	2.01%	2.03%	3.09%
Average Equity Share Offered	58.37%	59.40%	53.67%	0.26%
Small Firm Innovation Share	19.68%	21.19%	12.04%	0.00%
Inventor Participation Rate	40.19%	40.11%	40.44%	47.83%
Panel B: Percentage Changes Relative to Baseline				
Aggregate Growth Rate	—	-0.15%	+0.66%	+53.54%
Average Equity Share Offered	—	1.76%	-8.06%	-99.56%
Small Firm Innovation Share	—	+7.68%	-38.81%	-100.00%
Inventor Participation Rate	—	-0.20%	+0.63%	+19.01%

Notes: This table reports the counterfactual outcomes under alternative R&D ownership regimes. Column 1 reports the baseline economy, where employee-inventors retain a 2 percent statutory ownership share. Columns 2 through 4 simulate exogenous shifts in this statutory allocation, holding all other parameters constant. Panel A reports the levels of the aggregate moments. Panel B reports the percentage change of each moment relative to its baseline value, defined as $(x_{\text{counterfactual}} - x_{\text{baseline}})/x_{\text{baseline}}$. The baseline column in Panel B serves as the reference numeraire, denoted by dashes (—).

The counterfactual exercises demonstrate that shifting statutory R&D ownership from inventors to firms depresses innovative output and moves innovations to smaller firms, consistent with the empirical observations. On the other hand, shifting ownership from firms back to inventors increases aggregate growth. In the counterfactual where inventors retain 50% of the invention's value, the aggregate growth rate accelerates from the baseline 2.01% to 3.09%, as shown in the first row of both panels. In terms of magnitude, reversing the comparison (moving from 50% inventor ownership down to the 2% baseline) reduces the aggregate growth rate by roughly one-third. This is broadly consistent with Hvide and Jones (2018), who document a halving of researcher patenting following a two-thirds ownership shift from inventors to universities in Norway, though the institutional settings and the units of measurement differ.

The change in ownership regime also affects household welfare. When statutory ownership increases from 2% to 50%, the consumption-equivalent welfare improvement for the representative household is 71.60%. In other words, the representative household requires a 71.60% permanent increase in consumption to be indifferent between the original path and the higher-growth path.

Allocating more R&D ownership to inventors stimulates growth through three channels. First, the contract adjustment channel induces inventors to exert higher effort on average, because more direct statutory ownership provides more efficient incentives. The contract adjustment channel raises the average effort level by 0.12% in the 10% counterfactual case and 5% in the 50% case, respectively, conditional on working in the same firm. Second, the relocation channel shifts the equilibrium assignment of inventors toward larger, high-synergy firms, which increases the expected value of innovation. This endogenous reallocation occurs because the statutory increase in R&D ownership directly mitigates the principal-agent friction, thereby diminishing the inventor's reliance on high-powered firm equity. This effect is reflected in the firms' optimal contracting response: when the legal regime guarantees inventors a 10% statutory stake, the average equity share offered by firms drops from the baseline of 58.37% to 53.67%. Figure 2 illustrates this shift in optimal equity schedule across inventor types. As statutory ownership, ϕ, increases, the optimal equity offered by firms decreases across all inventor types, reflecting a reduced structural demand for supplementary incentives. As a result, as shown in Figure 3, the optimal firm size chosen by any given inventor increases. Inventors sort into larger firms. Third, on the extensive margin, the enhanced direct payoff induces an expansion of the active R&D workforce. For example, when statutory ownership increases to 10% and 50%, the inventor participation rate rises from 40.19% to 40.63% and 47.83, respectively, as shown in the last row. Notably, all three channels exhibit nonlinearity in ϕ with responses disproportionately more significant at higher statutory ownership shares.

6.1 Size-dependent reaction

While enhanced inventor ownership generates robust aggregate growth, it induces a reallocation of innovative activity across the firm-size distribution. As shown in Section 4.9, the model predicts that the response of innovation output to an ownership shift is size-dependent. This section examines this prediction quantitatively.

Innovation output is defined by a quality-weighted patent count, which is the same object that determines economic growth in the model. Normalizing the baseline innovation efficiency to 1, the economy-wide innovation efficiency drops by 0.15% when inventor own-

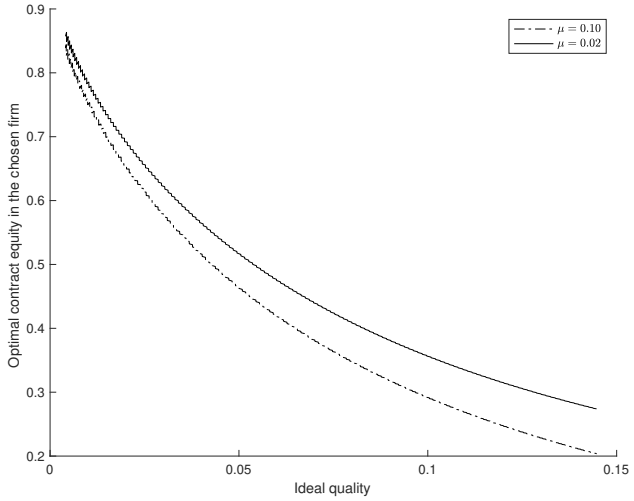


Figure 2: Optimal Equity by Inventor Type

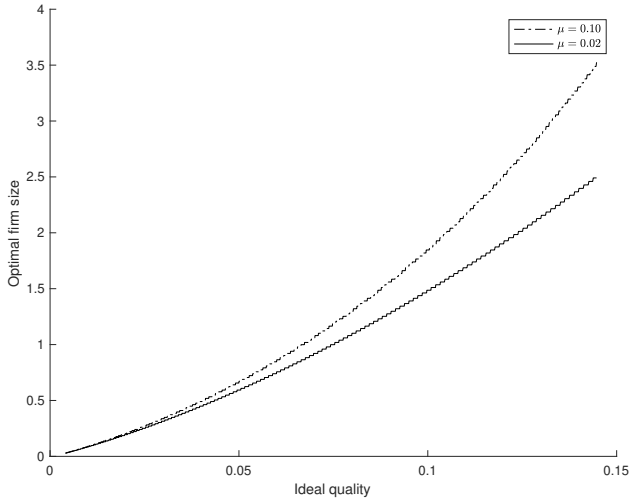


Figure 3: Optimal Firm Size by Inventor Type

ership drops from 2% to 0%. This drop is mitigated, and even reversed, in small firms. For the same change in ownership allocation, on average, the innovation efficiency increases by 0.28% in small firms, matching the sign and rough magnitude of the untargeted size-dependent response documented in Appendix A.4. This gain in small firm is driven by both the contract channel and the reallocation channel. First, by Proposition 2, when inventors hold only a small share of R&D ownership, the contract adjustment channel has a size-dependent effect. In small firms, incentive provision dominates risk-sharing considerations, so equity rises by enough to raise effective ownership and effort. Second, following the shock, more high-quality inventors are matched to small firms in equilibrium, raising small-firm innovation productivity. Together, these two forces offset the negative impact of the shift in small firms.

The reallocation also changes the share of innovations conducted in small firms, as illustrated by the Small Firm Innovation Share row of Table 5. When all R&D ownership is assigned to firms, small firms account for a larger share of total innovation output, rising by 7.68% relative to baseline. Conversely, when inventors hold more ownership, small firms become less attractive. In the baseline economy, small firms account for 19.68% of all innovations, because their equity is more closely tied to inventors' own performance, allowing them to offer more equity incentives which help to overcome the pro-employer legal regime. As statutory inventor ownership rises, the inventor's reliance on firm equity diminishes. When the inventor holds 10% ownership, the small-firm innovation share falls to 12.04%, a 38.81% drop relative to baseline. In the 50% ownership counterfactual, the small-firm innovation share drops to 0%. When inventors are legally guaranteed a sufficiently large direct share, the original trade-off between incentive alignment and commercialization synergy is resolved in favor of synergy. Inventors no longer need to join small firms for incentives; instead, the trade-off becomes one between the high synergy of large firms and the low management cost of small firms. Thus, the quantitative model highlights a policy trade-off: strengthening inventor property rights raises aggregate growth and leverages large-firm synergy, but it also reduces innovative activity in small firms.

6.2 Decomposition

To isolate the mechanisms through which statutory R&D ownership drives aggregate growth, the analysis employs a sequential decomposition of the counterfactual[22]. The results are reported in Table 6. Starting from the baseline equilibrium, the first step isolates the contract channel: inventor participation and firm sorting are held fixed at baseline levels, while only

[22]This paper conducts another decomposition exercise. This exercise starts with the baseline counterfactual and shuts down one channel at a time. The results are reported in Appendix C.2.

the optimal contract and unobservable effort adjust. In the 0% ownership case, this adjustment has a positive direct effect on growth (+0.02%), which goes against the overall direction of the general-equilibrium response (−0.15% total). This is because of the force described by Proposition 2. When firms hold all the R&D ownership, the share of firm-level risk that is innovation-related is particularly large in small firms, where non-innovative variance is small relative to innovation variance. When the inventor-firm allocation is fixed at baseline levels, a non-trivial mass of inventors remain in small firms. In these firms, the firm's equity adjustment overshoots: equity rises by enough to more than offset the lost direct ownership, so the inventor's effective stake rises slightly above the baseline. The average equity share rises from 58.37% to 59.23%, which means that although the direct R&D ownership drops, the effective ownership increases slightly by 0.03% overall, which is mainly drive by small firms. As a result, the aggregate effective ownership rises slightly from 59.20% to 59.23%, with the bulk of the change concentrated in small firms. This non-monotonic response of effective ownership in small firms depends on the calibration assumption that the exit rate is constant across firm sizes. If small firms faced higher exit risk, the non-innovative variance would not vanish in small firms, and the size-dependent flip would weaken or disappear. The qualitative aggregate-growth predictions of the model are unaffected; only the small-firm reversal would attenuate. In the 10% and 50% ownership cases, the contract channel contributes 20.45% and 9.32%, respectively. The contract channel itself contributes less when the regime shift is more significant. When facing a larger ownership shift, the reallocation and participation decisions are pushed further from their unconstrained optima, so the interaction between channels becomes more important.

The second step relaxes the sorting constraint, permitting the endogenous reallocation of inventors across the firm-size distribution (the inventor reallocation channel). In the 0% ownership case, this adjustment reverses the contract channel's positive contribution. The cumulative effect of the contract and reallocation channels combined is negative. This is because, when sorting is unlocked, inventors sort further toward small firms, dampening commercialization synergy and reducing aggregate growth. When ownership shifts more to inventors, the reallocation channel works in the same direction as the contract channel, further raising growth. In the 10% and 50% ownership cases, the combination of the contract channel and the inventor reallocation channel accounts for 57.58% and 82.54% of the aggregate growth improvement, respectively. This indicates that the combination of these two channels contributes more when the shock is larger.

Finally, the third step allows inventors to change their occupation decision, activating all three channels. Although the absolute change attributable to the participation channel grows as the shock becomes larger, its relative contribution to the total growth change declines from

over 60% (in the 0% counterfactual) to below 20% (in the 50% counterfactual).

This decomposition reveals non-linearities and complementarities among channels. The marginal contribution of the contracting channel diminishes as the regime shift becomes larger falling from 20.45% (in the 10% case) to 9.32% (in the 50% case). Because a massive ownership shock pushes the economy far from its baseline equilibrium, isolated adjustments in effort are insufficient; the interaction between optimizing incentive scheme and efficient firm-size sorting becomes increasingly important. This tension is better illustrated by the 0% ownership counterfactual, where the contracting channel actually operates in the opposite direction of the general equilibrium effect. When statutory ownership is revoked, firms compensate for the inventor's lost direct stake by offering more equity, which may lead to a raise in effective ownership above the baseline in small firms. If inventors are locked into their baseline firms (Step 1), this equity injection temporarily induces higher effort and boosts growth slightly. However, once the sorting constraint is lifted (Step 2), risk-averse inventors sort toward smaller firms. This endogenous reallocation undermines the commercialization synergy of the economy, reversing the initial growth bump and driving down the aggregate growth rate.

Table 6: Sequential Decomposition of Counterfactual

	Baseline	**Counterfactual Inventor Ownership Regimes**		
Moment	(2% Ownership)	0% Ownership	10% Ownership	50% Ownership
Growth Rate	**2.01%**	**2.01%**	**2.03%**	**3.09%**
Panel A: Incremental Percentage Change				
Contract Channel	—	+0.02%	+0.13%	+4.99%
+ Reallocation Channel	—	-0.07%	+0.24%	+39.21%
+ Participation Channel	—	-0.09%	+0.28%	+9.35%
Total Percentage Changes	—	-0.15%	+0.66%	+53.54%
Panel B: Marginal Percentage Contribution				
Contract Channel	—	-13.33%	+20.45%	+9.32%
+ Reallocation Channel	—	+50.00%	+37.12%	+73.22%
+ Participation Channel	—	+63.33%	+42.42%	+17.46%
Total Percentage Changes	—	+100.00%	+100.00%	+100.00%

Notes: This table reports the sequential decomposition of the aggregate growth rate under counterfactual statutory R&D ownership regimes. The first row reports the counterfactual growth rate as in Table 5. Panel A reports the marginal contribution to the percentage change in aggregate growth from sequentially activating each channel; the sum of these contributions equals the total percentage change relative to baseline. Panel B reports each channel's contribution as a share of the total percentage change. For example, the marginal contribution of the reallocation channel is computed as the difference between the cumulative growth change with only the contract channel active and the cumulative growth change with both the contract and reallocation channels active.

6.3 Optimal Ownership: full commitment

The counterfactual exercises presented above evaluate shifts in the uniform statutory allocation of R&D ownership (ϕ). This section evaluates a case where firms and inventors both have full commitment over the division of R&D ownership, ϕ. Therefore, both parties can negotiate an optimal division of property rights ex ante. Specifically, firms and inventors endogenously divide R&D ownership under full commitment, trading off between incentive alignment and risk-sharing.

Table 7 compares the results of this optimal ownership setting against the baseline. When agents can commit to the agreed property rights allocation, the aggregate growth rate nearly doubles, surging from 2.01% to 3.79%. This expansion is driven by all three channels. First, under the counterfactual setting, firms and inventors jointly choose the optimal ownership allocation, which improves the inventor effort level. As shown in the second row, in the optimal ownership case, the average level of effort increases from 13.78% to 14.61%, conditional on inventors working in the same firm. Secondly, fully committed ownership division means that firms can provide customized incentives tailored to inventor type, and they do not need to use equity anymore. The third row shows the average inventor ownership share. In the baseline case, on average an inventor holds an effective 59.20% share of R&D, which is a combination of 2% direct ownership and 58.37% ownership through equity holding. In the full-commitment case, the inventors' average R&D ownership share is 62.39%, all of which is held directly. The comparison of the effective inventor R&D ownership share by inventor type is shown in Figure 4. The figure confirms that for all inventor types, effective inventor R&D ownership share is always higher in the full-commitment case. Because large firms no longer suffer from the incentive-dilution problem, the share of innovations conducted in small firms drops to 0. Third, inventor participation also increases by more than one-fourth in the optimal ownership case. Since all inventors have a customized incentive schedule, more inventors are willing to participate in the R&D sector. All three channels also contribute to aggregate growth rate.

7 Discussion

This paper assumes that the exit rate is constant across firms. One extension is to allow the exit rate to be higher in small firms. One can show that, as long as $\frac{\tau(\tilde{q})\tilde{q}^2}{x_0^2(z,\tilde{q})}$ is increasing in $\tilde{q}$, the optimal equity offer decreases in firm size. This condition holds whenever the exit rate $\tau(\tilde{q})$ does not diminish too fast in firm size. See Appendix D for discussions on other assumptions and extensions.

Table 7: Comparison of Baseline vs. Optimal Full-Commitment Ownership

Macroeconomic Moment	**Baseline**	**Optimal Ownership**
Aggregate Growth Rate	2.01%	3.79%
Average effort (Within same firm)	13.78%	14.61%
Average Inventor R&D Ownership Share	59.20%	62.39%
Small Firm Innovation Share	19.68%	0.00%
Inventor Participation Rate	40.19%	51.11%

Notes: This table compares the moments of the baseline against the first-best optimum where firms and inventors negotiate R&D ownership shares under full commitment. The average effort is computed holding the inventor-firm assignment fixed at the baseline.

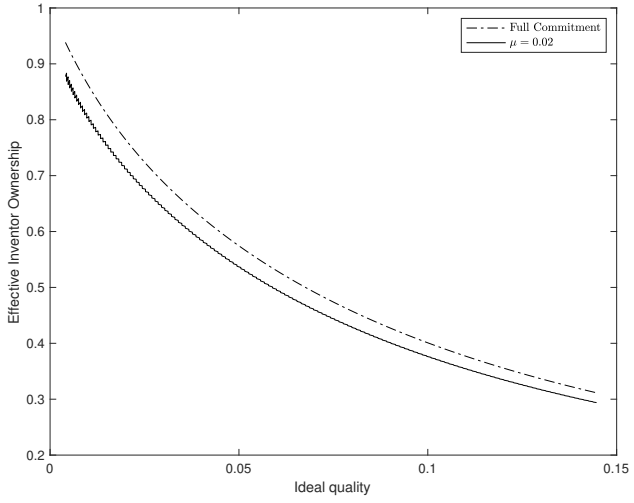


Figure 4: Effective Inventor R&D Ownership Share by Inventor Type

8 Conclusion

This paper develops a general equilibrium framework of endogenous growth to investigate how the statutory allocation of R&D property rights affects the aggregate growth rate. Departing from standard macroeconomic models that treat innovating firms as the sole R&D decision-makers, this analysis explicitly incorporates the principal-agent friction between inventors and employers. When the legal regime appropriates R&D ownership from inventors, firms must endogenously compensate for this loss of statutory ownership through equity. This framework generates a trade-off between incentive alignment and commercialization synergy: while large firms offer superior synergistic scale, their equity payouts are more diluted, rendering them less effective at eliciting unobservable effort. Small firms, in contrast, are better at incentiving inventors using equity but lack scale.

The quantitative counterfactuals reveal that restoring statutory property rights to inventors not only accelerates aggregate growth but also induces a reallocation of inventive talent. In an optimal contracting environment where ownership is jointly determined by firms and inventors, the growth rate nearly doubles and the share of innovations in small firms drops to zero. Optimally incentivized inventors no longer need to rely on equity; instead, they sort into large firms to maximize the commercialization synergy of their ideas. This indicates that a significant fraction of small-firm innovation operates as a response to a pro-employer legal regime. These findings yield policy implications: macroeconomic interventions, such as standard R&D subsidies directed at employers, may be severely blunted if the underlying legal architecture fails to pass incentives through to inventors themselves. Future research could extend this framework by incorporating dynamic on-the-job search or team-based production, to further illuminate how micro-level contracting frictions shape economic growth.

References

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999): "High wage workers and high wage firms," *Econometrica*, 67, 251–333.

Acemoglu, D. (1996): "A Microfoundation for Social Increasing Returns in Human Capital Accumulation," *The Quarterly Journal of Economics*, 111, 779–804.

Acharya, V. V., R. P. Baghai, and K. V. Subramanian (2014): "Wrongful discharge laws and innovation," *The Review of Financial Studies*, 27, 301–346.

Aghion, P. and J. Tirole (1994): "The Management of Innovation," *The Quarterly Journal of Economics*, 109, 1185–1209.

Akcigit, U., M. A. Celik, and J. Greenwood (2016): "Buy, Keep, or Sell: Economic Growth and the Market for Ideas," *Econometrica*, 84, 943–984.

Akcigit, U. and N. Goldschlag (2023): "Where have all the" creative talents" gone? Employment dynamics of us inventors," Tech. rep., National Bureau of Economic Research.

——— (2025): "Measuring the characteristics and employment dynamics of US inventors," *Journal of Economic Growth*, 30, 237–269.

Akcigit, U. and W. R. Kerr (2018): "Growth through Heterogeneous Innovations," *Journal of Political Economy*, 126, 1374–1443.

Akcigit, U., J. Pearce, and M. Prato (2025): "Tapping into talent: Coupling education and innovation policies for economic growth," *Review of Economic Studies*, 92, 696–736.

Alok, S. and K. Subramanian (2023): "Does Strengthening the Property Rights of Employee-Inventors Spur Innovation? Empirical Evidence on Freedom-to-Create Laws Passed by US States," *The Journal of Law and Economics*, 66, 369–408.

Autor, D., D. Dorn, L. F. Katz, C. Patterson, and J. Van Reenen (2020): "The Fall of the Labor Share and the Rise of Superstar Firms," *The Quarterly Journal of Economics*, 135, 645–709.

Bhaskarabhatla, A., L. Cabral, D. Hegde, and T. Peeters (2021): "Are inventors or firms the engines of innovation?" *Management Science*, 67, 3899–3920.

Bloom, N., C. I. Jones, J. Van Reenen, and M. Webb (2020): "Are ideas getting harder to find?" *American Economic Review*, 110, 1104–1144.

Brickley, J. A. and K. T. Hevert (1991): "Direct employee stock ownership: An empirical investigation," *Financial Management*, 70–84.

Buera, F. J. and E. Oberfield (2020): "The Global Diffusion of Ideas," *Econometrica*, 88, 83–114.

Chatterjee, S. and E. Rossi-Hansberg (2012): "Spinoffs and the Market for Ideas," *International Economic Review*, 53, 53–93.

Chiu, J., C. Meh, and R. Wright (2017): "Innovation and Growth with Financial, and Other, Frictions," *International Economic Review*, 58, 95–125.

Coase, R. H. (1937): "The Nature of the Firm," *Economica*, 4, 386–405.

Coolley, R. B. (1985): "Recent changes in employee ownership laws: Employers may not own their inventions and confidential information," *The Business Lawyer*, 57–75.

Cunningham, C., F. Ederer, and S. Ma (2021): "Killer Acquisitions," *Journal of Political Economy*, 129, 649–702.

Dratler Jr, J. (1979): "Incentives for people: The forgotten purpose of the patent system," *Harv. J. on Legis.*, 16, 129.

Frésard, L., G. Hoberg, and G. M. Phillips (2020): "Innovation Activities and Integration through Vertical Acquisitions," *The Review of Financial Studies*, 33, 2937–2976.

Garmaise, M. J. (2011): "Ties that truly bind: Noncompetition agreements, executive compensation, and firm investment," *The Journal of Law, Economics, & Organization*, 27, 376–425.

Giarratana, M. S., M. Mariani, and I. Weller (2018): "Rewards for patents and inventor behaviors in industrial research and development," *Academy of Management Journal*, 61, 264–292.

Goolsbee, A. (1998): "Does government R&D policy mainly benefit scientists and engineers?" .

Grossman, S. J. and O. D. Hart (1986): "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," *Journal of Political Economy*, 94, 691–719.

Gullette, R. L. (1980): "State legislation governing ownership rights in inventions under employee invention agreements," *J. Pat. Off. Soc'y*, 62, 732.

Hall, B. and J. Van Reenen (2000): "How effective are fiscal incentives for R&D? A review of the evidence," *Research policy*, 29, 449–469.

Hall, B. H., A. B. Jaffe, and M. Trajtenberg (2001): "The Nber Patent Citation Data File: Lessons, Insights and Methodological Tools," .

Hall, B. H. and R. H. Ziedonis (2001): "The Patent Paradox Revisited: An Empirical Study of Patenting in the Us Semiconductor Industry, 1979-1995," *Rand Journal of Economics*, 101–128.

Hart, O. (1995): *Firms, contracts, and financial structure*, Clarendon press.

Hart, O. and J. Moore (1990): "Property Rights and the Nature of the Firm," *Journal of Political Economy*, 98, 1119–1158.

——— (2008): "Contracts As Reference Points," *The Quarterly Journal of Economics*, 123, 1–48.

Hovell, W. P. (1982): "Patent Ownership: An Employer's Rights to His Employee's Invention," *Notre Dame L. Rev.*, 58, 863.

Hvide, H. K. and B. F. Jones (2018): "University innovation and the professor's privilege," *American Economic Review*, 108, 1860–1898.

Jo, K. and S. Kim (2024): "Heterogeneous Innovations and Growth Under Imperfect Technology Spillovers," .

Klette, T. J. and S. Kortum (2004): "Innovating Firms and Aggregate Innovation," *Journal of Political Economy*, 112, 986–1018.

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman (2017): "Technological Innovation, Resource Allocation, and Growth," *The Quarterly Journal of Economics*, 132, 665–712.

Lentz, R. and D. T. Mortensen (2008): "An Empirical Model of Growth through Product Innovation," *Econometrica*, 76, 1317–1373.

Lerner, J. and J. Wulf (2007): "Innovation and incentives: Evidence from corporate R&D," *the Review of Economics and Statistics*, 89, 634–644.

Ma, Y. (2022): "Specialization in a Knowledge Economy," Working Paper 4052990, SSRN.

Merges, R. P. (1999): "The law and economics of employee inventions," *Harv. JL & Tech.*, 13, 1.

Onishi, K. (2013): "The effects of compensation plans for employee inventions on R&D productivity: New evidence from Japanese panel data," *Research Policy*, 42, 367–378.

Parker, H. D. (1983): "Reform for rights of employed inventors," *S. Cal. L. Rev.*, 57, 603.

Perla, J., C. Tonetti, and M. E. Waugh (2021): "Equilibrium Technology Diffusion, Trade, and Growth," *American Economic Review*, 111, 73–128.

Phillips, G. M. and A. Zhdanov (2013): "R&d and the Incentives from Merger and Acquisition Activity," *The Review of Financial Studies*, 26, 34–78.

Pisegna-Cook, E. D. (1993): "Ownership rights of employee inventions: The role of preinvention assignment agreements and state statutes," *U. Balt. Intell. Prop. LJ*, 2, 163.

Romer, P. M. (1990): "Endogenous Technological Change," *Journal of Political Economy*, 98, S71–S102.

Silveira, R. and R. Wright (2010): "Search and the Market for Ideas," *Journal of Economic Theory*, 145, 1550–1573.

Suh, P. (2023): "Intellectual property rights and debt financing," *The Review of Financial Studies*, 36, 1970–2003.

Toivanen, O. and L. Väänänen (2012): "Returns to inventors," *Review of Economics and Statistics*, 94, 1173–1190.

Trajtenberg, M. (1990): "A penny for your quotes: patent citations and the value of innovations," *The Rand journal of economics*, 172–187.

Wilson, D. J. (2009): "Beggar thy neighbor? The in-state, out-of-state, and aggregate effects of R&D tax credits," *The Review of Economics and Statistics*, 91, 431–436.

Appendix A Empirics and Data

A.1 Data

The empirical analysis relies primarily on patent records from the United States Patent and Trademark Office (USPTO). This dataset, accessible via the PatentsView database, contains detailed information on all patents issued from 1926 to the present. To isolate the introduction of novel products and processes, the sample is restricted to utility patents granted to U.S. corporations. For each patent, the raw data document the grant year, the names and locations of all listed inventors and assignees, the technology classification code (CPC), and citation linkages. Based on this raw data, the core empirical variables are constructed as follows.

Forward Citations.— The primary measure of innovation output relies on forward citations, defined as the number of citations a focal patent receives from subsequent patents. Following standard conventions in the literature on the economics of innovation, forward citations serve as a proxy for technological quality, capturing the extent to which a given invention forms the foundation for subsequent R&D (Trajtenberg, 1990; Hall and Ziedonis, 2001). To avoid the right-censoring issue, the analysis calculates the cumulative number of forward citations each patent receives within a fixed five-year window following its issuance.

Small Entity Indicator.— To test the firm-size heterogeneity, the empirical specification requires a measure of firm scale. This is captured using the USPTO's small entity status indicator. A firm is classified as a small entity if it employs 500 or fewer workers. This binary indicator is utilized throughout the analysis to precisely distinguish between small firms and large firms.

Our baseline sample covers the period 2002 to 2017, spanning eight years before and after the exogenous legal shock. The baseline sample includes all inventors who file at least one patent during this window and whose assignee's geographic state can be definitively identified from the patent record. The resulting dataset contains 2,065,225 inventor-year observations in the full sample.

A.2 Institutional Background

To evaluate how R&D ownership shapes innovation output, it is first necessary to define the economic boundaries of an R&D outcome. The knowledge generated during the innovation process is inherently two-sided: it is embedded formally in patents and tacitly within the inventor herself, thereby enhancing her human capital. Consequently, an inventor's incentive to exert effort depends on the joint ownership of both the patent and her human capital—a

composite I define as *R&D ownership*. In the United States, this effective R&D ownership is governed by two sets of legal instruments. Non-compete agreements (NCAs) restrict inventor mobility and transfer the ownership of human capital from inventor to employer, while the formal assignment of patents is determined by the interplay between common law, preinvention assignment agreements (PIAAs), and state-level statutes (Alok and Subramanian, 2023). This subsection outlines this legal architecture, laying the institutional groundwork for the empirical identification.

Non-Compete Agreements. A non-compete agreement (NCA) restricts an inventor-employee from leaving her current employer to work for a competitor, or from founding a competing firm, for a specified period after the employment relationship ends. Because the threat of departure is a key source of bargaining power for inventor-employees, tighter NCA enforcement effectively transfers control over the inventor's human capital from the inventor to the employer. NCAs are widely used in the United States and represent one of the most important mechanisms binding employees to firms (Garmaise, 2011), thereby determining the division of human capital ownership.

Common Law. State law determines the ownership of an invention until a patent application is filed (Hovell, 1982). Under common law, the ownership of inventions is clearly allocated only at two extremes (Parker, 1983): if an inventor is hired explicitly to develop a specific technology, the rights belong to the employer; conversely, if an inventor creates an innovation entirely independently, without utilizing any of the employer's resources or proprietary information, the rights remain with the inventor. However, this leaves a middle ground of inventive activity in a legal gray area. For example, when an inventor is hired for general research and development and utilizes the firm's resources, the ownership of the resulting inventions is inherently ambiguous.

Preinvention Assignment Agreements (PIAA). To eliminate this uncertainty, employers rely on the PIAA, an employment contract that obligates an employee to assign to the employer all rights and interests in any inventions made or conceived during the employment term (Dratler Jr, 1979). Acting as a complement to common law, this contract instrument is widely used in the United States, especially among innovating firms, which require prospective employees to sign PIAAs as a precondition for employment (Pisegna-Cook, 1993; Merges, 1999). In practice, the scope of a PIAA may be sufficiently broad to cover categories of inventions far beyond the employer's immediate business, and the agreement may extend both before and after the terms of employment (Alok and Subramanian, 2023).

The enforcement of a PIAA implies that essentially all inventions can potentially belong to the employer, granting the firm the legal apparatus to appropriate even those innovations that are entirely unrelated to the employment relationship. The combination of its expansive scope and universal adoption dictates that the PIAA governs the ownership of patents.

Free-to-Create Statutes. In response to the expansive reach of PIAAs, several states enacted pro-employee "free-to-create" statutes in the late 1970s to early 1980s. These laws set statutory limits on PIAAs, permitting employees to retain ownership of inventions made using their own time and resources, provided the inventions did not relate to the employer's current or reasonably anticipated business. Eight states passed such legislation: California, Minnesota, Washington, North Carolina, Illinois, Delaware, Kansas, and Utah (Gullette, 1980; Coolley, 1985). Under these statutes, inventors in these states enjoy a legally protected degree of residual ownership over a subset of their human capital.

The 2008 DDB Ruling. This state-level legal protection was abruptly dismantled by the 2008 Court of Appeals for the Federal Circuit (CAFC) decision in *DDB Technologies LLC v. MLB Advanced Media, LLP*[23]. The case centered on David Barstow, a software engineer employed by Schlumberger Technology Corporation to develop oil-well software. During his employment, Barstow invented and patented a baseball statistical tracking technology. This innovation is entirely unrelated to his employer's line of business, as confirmed by Schlumberger's general counsel and the lab director at the time of the invention. Barstow formed DDB Technologies and assigned the patents to it. However, when Barstow's newly formed entity sued MLB for patent infringement, MLB successfully defended itself by securing a competing license from Schlumberger. MLB argued that Barstow's PIAA automatically assigned ownership of the patents to his employer, regardless of the invention's subject matter. The CAFC ruled in MLB's favor, establishing the precedent that questions of patent assignment are governed by federal law, overriding state-level regulations. This ruling effectively rendered the pro-employee 'free-to-create' statutes in the eight affected states unenforceable. Consequently, the *DDB* decision generated a legally mandated exogenous shift in R&D ownership from inventors to employers within these eight formerly protected states.

In summary, the effective ownership of an inventor's R&D output is shaped simultaneously by these intersecting legal instruments. The 2008 *DDB* ruling exogenously shifted statutory R&D ownership toward employers by invalidating state-level PIAA exemptions. In sections below, I show empirical patterns using this shock as the identifying variation in regression analysis.

[23] 517 F.3d 1284, 1290 Fed. Cir. 2008.

A.3 Identification

Following Suh (2023), the empirical strategy employs a difference-in-differences framework with continuous treatment. This design exploits the exogenous shift in R&D ownership generated by the 2008 *DDB* ruling, utilizing cross-sectional variation in treatment intensity. The treatment group comprises individuals working for firms located in the eight states directly affected by the invalidation of pro-employee statutes.[24]

While the 2008 ruling uniformly shifted statutory R&D rights toward employers across all eight treated states, the practical severity of this shock varied depending on the inventor's pre-existing control over her human capital, governed by state-level NCA enforcement. In states with strict NCA enforcement, inventors already faced severe mobility restrictions and possessed weak outside options; the marginal impact of losing R&D right via the ruling shift was relatively muted. Conversely, in states with weak NCA enforcement, inventors retained high mobility. For these inventors, the invalidation of free-to-create statutes constituted a substantial reduction of their effective R&D ownership.

To accurately capture this heterogeneous treatment intensity, the empirical specification defines a treatment dose that is inversely proportional to NCA stringency:

$$\text{Dose}_s \;=\; \mathbb{1}_{s \in S_{\text{treat}}} \left(\max_{s'} (\text{NCA}_{s'}) \;-\; \text{NCA}_s \right), \tag{31}$$

where S_{treat} is the set of eight treated states with pro-employee statutes, and NCA_s is the NCA Enforcement Index for state s as measured by Garmaise (2011). A higher value of Dose_s indicates that the DDB ruling had a larger practical impact on effective ownership of inventor i in treated state s. For all control states, Dose_s takes a value of zero.

Utilizing this exposure measure, the baseline regression is estimated at the inventor level:

$$y_{ijs,t+2} = \alpha + \beta \left(\text{Dose}_s \times \text{Post}_t \right) + \mathbf{X}'_{ijst} \boldsymbol{\omega}_{\boldsymbol{X}} + \varepsilon_{ijst}, \tag{32}$$

where y_{ijst} is the innovative output of inventor i, who is employed at firm j in state s in year t. I measure the innovative output at $t+2$ to account for the patent system process (Hall et al., 2001). Post_t is an indicator equal to one for all years following the ruling. The vector of observables $\mathbf{X}_{ijst}$ controls for time-varying firm size[25], year fixed effects, state fixed effects, firm fixed effects, and inventor-level technology-class (CPC) fixed effects[26]. Because more than 50 percent of inventors in the sample are granted only one patent, the inclusion of inventor fixed effects induces selection by dropping single-patent observations. The baseline specification therefore controls for inventor-level technology-class fixed effects instead. Both

[24] Generally, the state where the employer is located or where the job duties are performed constitutes a reasonable choice of law and is likely to be honored by courts (American Bar Association 2014).

[25] I use a small firm indicator to denote firm size. $\text{Small}_{ijt} = 1$ if firm j employing inventor i at time t satisfies the criteria of being a small firm (no more than 500 employees), as defined by USPTO.

[26] An inventor-level technology-class is defined as the class field that the inventor most often innovates in.

Dose_s and Post_t are absorbed by the state and year fixed effects. The coefficient of interest, β, captures the differential change in innovative output for inventors subjected to a higher treatment dose.

The analysis evaluates innovation output using three measures. The primary one is based on a patent count, defined as:

$$y_{ijs,t} = \text{Patent Count}_{i,t}. \tag{33}$$

The second measure captures individual technological output based on total patent quality, measured by citations:

$$y_{ijs,t} = \sum_{p \in P(i,t)} \text{citation}_p. \tag{34}$$

where $P(i,t)$ is the sample of patents granted to inventor i in year t. Citations are truncated at the 5-year window to mitigate right-censoring issues. The third metric is quality-weighted patent output, defined as:

$$y_{ijs,t} = \sum_{p \in P(i,t)} f(\text{citation}_p), \tag{35}$$

where $f(\cdot)$ is an increasing function.[27] The first two specifications are evaluated using a Poisson regression, while the third one is evaluated using OLS. In Appendix A.5, I report a firm-level regression using the change in the number of active inventors as the dependent variable.

The identifying assumption is a parallel-trends condition. Suh (2023) validates the parallel trends assumption for this legal shock at the aggregate firm level. Extending the identification to individual inventors requires that, prior to 2008, the distribution of inventors across firms of different sizes was evolving in parallel across treated and control states. Because the treatment dose derives from state-level legal environments, standard errors are clustered at the state level throughout all specifications to account for arbitrary serial correlation within jurisdictions.

This paper further explores the firm-size heterogeneity in response. The baseline specification in (32) is augmented with a three-way interaction term, alongside all relevant lower-order interactions:

$$\gamma \left(\text{Small}_{ijt} \times \text{Dose}_{is} \times \text{Post}_t \right), \tag{36}$$

The coefficient γ captures the heterogeneous treatment effect by firm size. If there is a positive estimate for γ, it suggests that small firms can mitigate the negative effect of pro-employer ownership.

[27] Here I assume that $f(\text{citation}_p) = \ln(1 + \text{citation}_p)$. The notation $(1 + \text{citation}_p)$ here captures the fact that the impact of an innovation includes both the innovation itself and the subsequent citations it receives.

A.4 Results

This section establishes the empirical validation for the contract assumption. Empirical results demonstrate that inventor innovation output responds negatively to the pro-employer ruling. The magnitude of this response is heavily mediated by firm size, consistent with the assumption of a structural incentive channel that is described in the general equilibrium model in Section 4.

Table A1: Relationship between Innovation Output and Pro-employer Ruling Treatment

	Patent Quantity		Patent Quality		Quality-weighted Quantity	
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{Dose}_s \times \text{Post}_t$	-0.005*	-0.007**	-0.026**	-0.031**	-0.042***	-0.047***
	(0.003)	(0.003)	(0.013)	(0.014)	(0.01)	(0.012)
$\text{Small}_{ijt} \times \text{Dose}_s \times \text{Post}_t$		0.013***		0.044***		0.036***
		(0.002)		(0.014)		(0.011)
Observations	2,065,225	2,065,225	2,052,616	2,052,616	2,065,225	2,065,225
R-squared					0.148	0.149
Year, state, industry and firm FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Standard errors are clustered at the state level. All dependent variables are measured at $t+2$. Columns (1)–(2) use the patent count defined in Equation (33) as the outcome. Columns (3)–(4) use the patent quality defined in Equation (34) as the outcome. Columns (5)–(6) use the quality-weighted patent output defined in Equation (35) as the outcome. Column (2), (4), and (6) also control for the small firm indicator, the interaction term between the small firm indicator and treatment dose, and interaction term between the small firm indicator and post treatment indicator. All specifications control for year, state, industry and firm fixed effect. *** Significant at the 1 percent level; ** significant at the 5 percent level; * significant at the 10 percent level.

Table A1 reports the results of regression (32). The first two specifications are evaluated using Poisson regression to account for the zeros. The columns are organized by outcome variable. Columns (1) and (2) estimate the impact on innovation output using raw patent counts (Equation (33)). Columns (3) and (4) use total forward citations (Equation (34)). Finally, columns (5) and (6) capture quality-weighted patent output (Equation (35)). For each set of results, the table first presents the average treatment effect, followed by a specification exploring firm-size heterogeneity. Across all columns, the evidence demonstrates that exogenously shifting statutory R&D ownership toward employers depresses innovative output, with this depression mitigated within small firms.

Columns (1) and (2) present the results of the first measure, showing that shifting R&D ownership toward employers reduces innovation. In column (1), the coefficient on the interaction Dose $\times$ Post is -0.005. It implies that for an inventor in a treated state at the average level of exposure, Dose $= 6.98$, the DDB ruling reduced an average inventor's output by approximately 3.43%. Translating this productivity loss into aggregate financial terms using the patent valuation methodology of Kogan et al. (2017), this drop corresponds to an

annualized loss of roughly $1.60 million (in 2025 dollars) in innovation value per firm, relative to the unaffected control states. In terms of aggregate economic magnitude, this shock is associated with a total loss of $87.55 billion. Using the second measure, for an inventor in a treated state at the average level of exposure, the DDB ruling reduced an average inventor's output by approximately 16.60%. This result implies that the shock has a stronger effect on quality, compared to quantity. Using the third measure, the effect is approximately 10.37%. This negative effect is consistent across alternative measures of innovative output as shown in column (3) and column (5). Together, these baseline estimates confirm that when statutory innovation ownership shifts from inventors to employers, overall innovative output drops.

I document heterogeneity across firms of different sizes. The negative effects of the pro-employer ownership shift are mediated by firm size, as shown in the even-numbered columns of Table A1. The treatment effects across all specifications are always negative, while the interaction coefficient for small firms is always positive. This implies that affiliation with a small firm offsets the negative impact of the ownership shift. In column (2) and (4), the net effect is reversed when working in a small firm. In column (6), the baseline treatment effect for large firms is -0.047, while the interaction coefficient for small firms is $+0.036$. This implies that affiliation with a small firm offsets more than 75 percent of the baseline negative impact induced by the legal shock. For the average innovative public firm, which generates $46.29 million in annual innovation value, this differential implies that each large firm incurs an additional loss of at least $4.11 million per year. The results support the claim that offering R&D ownership to inventors matters for innovation output, suggesting an inventor incentive channel. Additionally, small firms are better able to incentivize inventors even when baseline direct ownership is stripped away, consistent with the implications of the model's contract space assumption. This mitigation effect is replicated across both patent counts (Column 2) and total citations (Column 4), where the positive small-firm interaction term systematically dampens the primary negative shock. This mitigation effect is a combination of two channels. One is the contract adjustment channel. Firms provide more additional incentives in small firms, leading to a smaller drop in inventors' productivity. The other is the inventor reallocation channel. Inventors move to small firms. Both channels are captured in the model in Section 4.

In summary, when innovation ownership shifts from inventors to employers, overall innovative output is affected negatively, especially for large firms. In Appendix A.5, I report a firm-level regression using the change in the number of active inventors as the dependent variable. It captures inventor allocation across firms. The results are consistent with the inventor-level outcome. All results are consistent with the model's implications, confirming the assumption of the model. When direct ownership is stripped away, small firms are better

able to incentivize inventors. One potential explanation is that, to tie the inventors' payoff to innovation outcomes, firms use alternative incentives such as equity. However, equity is a less powerful incentive device in large firms, where the link between the inventor's effort and the overall firm value is diluted by unrelated activities within the organization. In small firms, equity is more tightly tied to the inventor's own contribution, so firms can more effectively compensate inventors for the loss of direct ownership.

A.5 Change in the Number of Active Inventors

Table A2 reports the change in the number of active inventors by firm size. The growth rate of active inventors is defined as

$$g_{\text{inventor},j,t} = \frac{N_{j,t}}{N_{j,t-1}} - 1,$$

where $N_{j,t}$ is the number of inventors in firm j who are granted at least one patent in year t. For a firm with no active inventors in year $t-1$, I use the most recent year s in which the firm had an active inventor ($N_{j,s} > 0$) and compute the average annual change over the $(t-s)$-year interval.

The firm-level results on inventor counts are consistent with the inventor-level findings in Table A1. The shock has a negative effect on the inventor growth rate. After the shock, the growth rate of active inventors drops by approximately 2 percentage points. This negative effect is mitigated in small firms, where the drop in growth rate is roughly half the drop in large firms. This size-dependent heterogeneity is consistent with the model's implications.

Appendix B Model and Proof

B.1 Closed-Form Model Solution

We can reformulate the firm's optimization problem from Equation (14) as follows:

$$\begin{aligned} &\max_{a,\phi} \left(\mu e_I x_0(z,\tilde{q}) - A\left((a(1-\phi)+\phi)^2 \mu^2 x_0(z,\tilde{q})^2 + a^2 \tau \tilde{q}^2 \right) - \frac{1}{2} e_I^2 - \eta \tilde{q} \right) \nu \bar{q} \\ &\text{st } e_I = \mu \left(a(1-\phi)+\phi \right) x_0(z,\tilde{q}) \end{aligned} \tag{37}$$

By substituting the expression for e_I into the objective function and evaluating the first-order condition with respect to the equity share a, we obtain:

$$\begin{aligned} a^* &= \max \left\{ \frac{(1-\phi)\left[(1-\phi)\mu^2 x_0(z,\tilde{q})^2 - 2A\phi\mu^2 x_0(z,\tilde{q})^2\right]}{(1-\phi)^2\left(2A\mu^2 x_0(z,\tilde{q})^2 + \mu^2 x_0(z,\tilde{q})^2\right) + 2A\tau\tilde{q}^2}, 0 \right\} \\ &= \max \left\{ \frac{1 - \frac{2A\phi}{(1-\phi)}}{1 + 2A\left(\frac{\tau\tilde{q}^2}{(1-\phi)^2\mu^2 x_0(z,\tilde{q})^2} + 1\right)}, 0 \right\} \end{aligned} \tag{38}$$

Table A2: Relationship between the Growth of the Number of Active Inventors and Pro-employer Ruling Treatment

	(1)	(2)
$\text{Dose}_s \times \text{Post}_t$	-0.016***	-0.020***
	(0.005)	(0.006)
$\text{Small}_{ijt} \times \text{Dose}_s \times \text{Post}_t$		0.011**
		(0.004)
Observations	121,857	121,857
R-squared	0.274	0.274
Year, state, and firm FE	Yes	Yes

Notes: Standard errors are clustered at the state level. All depend variables are measured at $t+2$. The table uses the change rate of the number of active inventors as the outcome. An active inventor is defined as an inventor who is granted with at least one patent in the given year t. I drop the extreme growth rates which are either in the top 1% or the bottom 1%. Column (2) controls for the small firm indicator, the interaction term between the small firm indicator and treatment dose, and interaction term between the small firm indicator and post treatment indicator. Both specifications control for year, state, industry and firm fixed effect. *** Significant at the 1 percent level; ** significant at the 5 percent level; * significant at the 10 percent level.

Faced with a menu of potential contracts, an inventor characterized by idea quality z selects her preferred firm size $\tilde{q}$ by maximizing her instantaneous utility:

$$\begin{aligned} \max_{\tilde{q}} u\left(c_I, e_I\right) =& \mathbb{E}\left(c_I\right) - \frac{A}{\nu\bar{q}}\text{Var}\left(c_I\right) - R\left(e_I, \tilde{q}\right)\nu\bar{q} \\ \text{st } a =& a^*\left(\tilde{q}\right) \\ T_I =& T_I^*\left(\tilde{q}\right) \end{aligned} \tag{39}$$

Putting the expression of the optimal equity level, $a^*(\tilde{q})$, the optimal onwership ϕ, and $T_I^*(\tilde{q}) = (1-\phi)\left(1-a^*\right)\mu e_I x_0\left(z,\tilde{q}\right)$ into the maximization problem and solving the first-order condition,

$$\mu^2 \frac{x_0\left(z,\tilde{q}\right)}{\tilde{q}} \frac{\partial x_0\left(z,\tilde{q}\right)}{\partial \tilde{q}} \left((1+(1-a)(1-\phi)) - 2A\left(a(1-\phi)+\phi\right)\right)\left(a(1-\phi)+\phi\right) = 2Aa^2\tau - \frac{\eta}{\tilde{q}}, \tag{40}$$

which implies the optimal firm size.

B.2 Proof for Lemma 1

Proof. When $\phi < \frac{1}{2A+1}$, $a^* > 0$, it takes the value:

$$a^* = \frac{1 - \frac{2A\phi}{(1-\phi)}}{1 + 2A\left(\frac{\tau\tilde{q}^2}{(1-\phi)^2\mu^2 x_0^2} + 1\right)},$$

which decreases in $\frac{\tilde{q}}{x_0}$. Define $f(\tilde{q}) = \frac{\tilde{q}}{x_0(z,\tilde{q})}$. To ensure that the optimal equity decreases in firm size, it requires that $f(\tilde{q})$ increases in firm size $\tilde{q}$.

Recall that $h(\tilde{q})$ is as:

$$h(\tilde{q}) = \frac{1}{1 + \exp\left(-b_0 - b_1 \log(\tilde{q})\right)} \tag{41}$$

where $b_1 > 0$. Then $x_0(z, \tilde{q}) = \gamma_L z(1 + \frac{\gamma_H}{\gamma_L} h(\tilde{q}))$.

First, we can rewrite the function $h(\tilde{q})$ as:

$$h(\tilde{q}) = \frac{1}{1 + \exp(-b_0)\tilde{q}^{-b_1}} = \frac{\tilde{q}^{b_1}}{\tilde{q}^{b_1} + \exp(-b_0)}.$$

Substitute this into $x_0(z, \tilde{q})$:

$$x_0(z, \tilde{q}) = \gamma_L z \left(1 + \gamma \frac{\tilde{q}^{b_1}}{\tilde{q}^{b_1} + \exp(-b_0)}\right) = \gamma_L z \frac{(1+\gamma)\tilde{q}^{b_1} + \exp(-b_0)}{\tilde{q}^{b_1} + \exp(-b_0)},$$

where $\frac{\partial x_0(z,\tilde{q})}{\partial \tilde{q}} = \gamma_L z b_1 \tilde{q}^{b_1 - 1} \frac{\gamma \exp(-b_0)}{\left(\tilde{q}^{b_1} + \exp(-b_0)\right)^2} > 0$. The target function $f(\tilde{q})$ becomes:

$$f(\tilde{q}) = \frac{\tilde{q}}{x_0(z, \tilde{q})} = \frac{1}{\gamma_L z} \frac{\tilde{q}(\tilde{q}^{b_1} + \exp(-b_0))}{(1+\gamma)\tilde{q}^{b_1} + \exp(-b_0)}.$$

Take the derivative of the non-constant part, $\frac{\tilde{q}(\tilde{q}^{b_1} + \exp(-b_0))}{(1+\gamma)\tilde{q}^{b_1} + \exp(-b_0)}$. Let the numerator be $N(\tilde{q}) = \tilde{q}^{b_1+1} + \exp(-b_0)\tilde{q}$ and the denominator be $D(\tilde{q}) = (1+\gamma)\tilde{q}^{b_1} + \exp(-b_0)$.

Taking the derivatives:

$$N'(\tilde{q}) = (b_1 + 1)\tilde{q}^{b_1} + \exp(-b_0),$$
$$D'(\tilde{q}) = b_1(1+\gamma)\tilde{q}^{b_1 - 1}.$$

Applying the quotient rule numerator $N'D - ND'$:

$$\begin{aligned}
N'D - ND' &= [(b_1+1)\tilde{q}^{b_1} + \exp(-b_0)][(1+\gamma)\tilde{q}^{b_1} + \exp(-b_0)] - [\tilde{q}^{b_1+1} + \exp(-b_0)\tilde{q}][b_1(1+\gamma)\tilde{q}^{b_1-1}] \\
&= (b_1+1)(1+\gamma)\tilde{q}^{2b_1} + \exp(-b_0)(b_1+1)\tilde{q}^{b_1} + \exp(-b_0)(1+\gamma)\tilde{q}^{b_1} + \exp(-b_0)^2 \\
&\quad - b_1(1+\gamma)\tilde{q}^{2b_1} - \exp(-b_0)b_1(1+\gamma)\tilde{q}^{b_1} \\
&= (1+\gamma)\tilde{q}^{2b_1} + \exp(-b_0)[(b_1+1) + (1+\gamma) - b_1(1+\gamma)]\tilde{q}^{b_1} + \exp(-b_0)^2 \\
&= (1+\gamma)\tilde{q}^{2b_1} + \exp(-b_0)(2 + \gamma - b_1\gamma)\tilde{q}^{b_1} + \exp(-b_0)^2.
\end{aligned}$$

When $b_1 \leq 1$, the expression above is strictly positive for all $\tilde{q} > 0$, meaning that the equity level a^* decreases in $\tilde{q}$. ■

B.3 Proof for Lemma 2

Proof. From the Equation 29, it is easy to see that when $\phi \geq \frac{1}{2A+1}$, $a^* = 0$. When $a^* > 0$, which is equivalent to $\phi < \frac{1}{1+2A}$, the expression of equity a^* can be written as:

$$a^* = \frac{N(\phi)}{D(\phi)}$$

where the numerator is

$$N(\phi) = 1 - 2A\frac{\phi}{1-\phi}$$

and the denominator is

$$D(\phi) = 1 + 2A\left(\frac{\tau\tilde{q}^2}{(1-\phi)^2\mu^2 x_0(z,\tilde{q})^2} + 1\right).$$

Let us analyze the derivatives of $N(\phi)$ and $D(\phi)$ with respect to ϕ. It is easy to see that $N(\phi)$ is strictly decreasing in ϕ, and the denominator $D(\phi)$ is strictly increasing in ϕ. Therefore, their ratio $a^* = \frac{N(\phi)}{D(\phi)}$ is strictly decreasing in ϕ. This concludes the proof. ■

B.4 Proof for Lemma 3

Proof. At the optimal level:

$$e_I^* = \mu \left(a^*(1-\phi)+\phi\right) x_0\left(z, \tilde{q}\right).$$

It is easy to see that $\frac{\partial e_I^*}{\partial(a^*(1-\phi)+\phi)} > 0$. ■

B.5 Proof for Proposition 1

Proof. When inventors are fully incentivized, $\phi = \frac{1}{1+2A}$. The optimal equity is $a^*(\tilde{q}, z) = 0$. The optimal effort $e_I^I\left(z, \tilde{q}\right)$ is:

$$e_I^I\left(z, \tilde{q}\right) = \mu \frac{1}{1+2A} x_0\left(z, \tilde{q}\right). \tag{42}$$

When firms retain ownership, $\phi \to 0$, the optimal equity is:

$$a^* = \frac{1}{1+2A\left(\frac{\tau \tilde{q}^2}{\mu^2 x_0(z,\tilde{q})^2}+1\right)}, \tag{43}$$

which is always lower than $\frac{1}{(1+2A)}$. The corresponding effort e_I^F is:

$$e_I^F\left(z, \tilde{q}\right) = \mu a x_0\left(z, \tilde{q}\right). \tag{44}$$

Compare the two cases, the change in effort is:

$$\frac{e_I^I\left(z, \tilde{q}\right) - e_I^F\left(z, \tilde{q}\right)}{e_I^I\left(z, \tilde{q}\right)} = 1 - (1+2A)a^*(z, \tilde{q}). \tag{45}$$

Since $(1+2A)a^*(z, \tilde{q})$ is less than 1 and decreases in firm size, $\tilde{q}$, the change in effort, the gap between the optimal R&D ownership effort and the firm-owned effort is always positive and more salient in large firms. ■

B.6 Proof for Proposition 2

Proof. At the optimal level:

$$e_I^* = \mu \left(a^*(1-\phi)+\phi\right) x_0\left(z, \tilde{q}\right).$$

where a is given by:

$$a^* = \frac{1-\frac{2A\phi}{(1-\phi)}}{1+2A\left(\frac{\tau \tilde{q}^2}{(1-\phi)^2\mu^2 x_0^2}+1\right)},$$

when $\phi < \frac{1}{1+2A}$, and $a^* = 0$ otherwise. This proof focuses on the case where equity is positive.

Since μ and $x_0(z, \tilde{q})$ are strictly positive and independent of ϕ, the monotonicity of e_I

with respect to ϕ is identical to the monotonicity of the term inside the parentheses:

$$\tilde{\phi} \equiv a^*(1-\phi)+\phi.$$

Substituting the expression for a^* into $\tilde{\phi}$, we have:

$$\begin{aligned}\tilde{\phi} &= \frac{(1-\phi)-2A\phi}{1+2A+\frac{2A\tau\tilde{q}^2}{(1-\phi)^2\mu^2x_0^2}}+\phi \\ &= \frac{1-\phi-2A\phi+\phi\left(1+2A+\frac{2A\tau\tilde{q}^2}{(1-\phi)^2\mu^2x_0^2}\right)}{1+2A+\frac{2A\tau\tilde{q}^2}{(1-\phi)^2\mu^2x_0^2}} \\ &= \frac{1+\phi\left(\frac{2A\tau\tilde{q}^2}{(1-\phi)^2\mu^2x_0^2}\right)}{1+2A+\frac{2A\tau\tilde{q}^2}{(1-\phi)^2\mu^2x_0^2}}.\end{aligned}$$

Therefore,

$$\begin{aligned}\frac{\partial\tilde{\phi}}{\partial\phi} &= \text{Constant}\times\left[\left(1+2A+\frac{2A\tau\tilde{q}^2}{(1-\phi)^2\mu^2x_0^2}\right)(1+\phi)-2\left(1+\phi\left(\frac{2A\tau\tilde{q}^2}{(1-\phi)^2\mu^2x_0^2}\right)\right)\right] \\ &= \text{Constant}\times\left[\frac{2A\tau\tilde{q}^2}{(1-\phi)\mu^2x_0^2}+\phi(1+2A)+2A-1\right],\end{aligned}$$

where the Constant $= \frac{\left(\frac{2A\tau\tilde{q}^2}{\mu^2x_0^2}\right)}{\left(1+2A+\frac{2A\tau\tilde{q}^2}{(1-\phi)^2\mu^2x_0^2}\right)^2(1-\phi)^3}$ is positive. The sign of $\frac{\partial\tilde{\phi}}{\partial\phi}$ depends on the ratio of non-innovative to innovative risk $\frac{\tau\tilde{q}^2}{\mu^2x_0^2}$, and on the inventor ownership ϕ. When the inventor holds a sufficiently large share of R&D, $\phi \geq \bar{\phi} \equiv \frac{2A-1}{1+2A}$, $\frac{\partial\tilde{\phi}}{\partial\phi} > 0$ always holds. Therefore, e_I is strictly increasing in ϕ.

Otherwise, the result depends on the sign of the first term. Following the proof of Lemma 1:

$$\frac{\partial\frac{\tau\tilde{q}^2}{\mu^2x_0^2}}{\partial\tilde{q}} \geq 0,$$

implying the effective ownership $\tilde{\phi}$ increases in ϕ when the $\tilde{q} \geq \tilde{q}_0$.[28] When $\phi < \tilde{\phi}$ and $\tilde{q} < \tilde{q}_0$, $\frac{\partial\tilde{\phi}}{\partial\phi} < 0$, an increase in ϕ leads to a decrease in effective ownership and hence a drop in effort e^*.

In the closed-form model with modified assumptions, the cutoff is $\bar{\phi} \equiv \frac{2A-1}{1+2A}$. In the full model, the ownership cutoff is higher than $\frac{2A-1}{1+2A}$. This is because that the closed-form variance specification suppresses the effort-dependence of innovation risk, which inflates the share of innovation risk in the inventor's variance and tightens the threshold. ■

[28] $\tilde{q}_0$ satisfies $\frac{2A\tau\tilde{q}_0^2}{(1-\phi)\mu^2x_0^2(z,\tilde{q}_0)}+\phi(1+2A)+2A-1=0$.

B.7 Proof for Proposition 3

Recall that the optimal firm size is determined by Equation (40):

$$\mu^2 \frac{x_0(z,\tilde{q})}{\tilde{q}} \frac{\partial x_0(z,\tilde{q})}{\partial \tilde{q}} \left((1+(1-a)(1-\phi)) - 2A\left(a(1-\phi)+\phi\right)\right)\left(a(1-\phi)+\phi\right) = 2Aa^2\tau - \frac{\eta}{\tilde{q}} \tag{46}$$

It is easy to see that the solution of $\tilde{q}$ depends on the value of ϕ.

B.8 Full Model: Innovation

For an inventor who chooses to work in the production sector, the consumption is

$$c_I^L = \xi w = \frac{\xi(1-\beta)}{(1+\xi l_I)^\beta}\bar{q}. \tag{47}$$

For an inventor who chooses to work in the R&D sector, the consumption is

$$\begin{aligned} &c_I(a, T_I, q, z, \phi, \bar{q})dt \\ &= T_I V(\bar{q})dt + adV(j|q,z), \\ &= T_I \nu \bar{q} dt \begin{cases} +a\tau\nu q dt, & Pr = 1-\tau dt - \mu e_I dt \\ -a(1-\tau dt - rdt)\nu q, & Pr = \tau dt \\ +\left(\left(a(1-\phi)+\phi\right)(1-rdt)x(z,\tilde{q})\bar{q} + a\tau q dt\right)\nu, & Pr = \mu e_I dt \end{cases} \end{aligned} \tag{48}$$

where r is the interest rate. The consumption is linear in ν. The inventor's expected consumption is given by:

$$\mathbb{E}\left(c_I(a, T_I, q, z, \phi, \bar{q})\right) = \left(\left(a(1-\phi)+\phi\right)\mu e_I x_0(z,\tilde{q}) + T_I\right)\nu\bar{q},$$

with the corresponding variance expressed as:

$$\begin{aligned} \mathrm{Var}\left(c_I(a, T_I, q, z, \phi, \bar{q})dt\right) =& \mathbb{E}((c_I dt)^2) - \mathbb{E}(c_I dt)^2 \\ =& \mathbb{E}\left(\left(a(1-\phi)+\phi\right)^2 x^2(z,\tilde{q})\mu e_I + a^2\tilde{q}^2\tau\right)\nu^2\bar{q}^2 dt \\ =& \left(a^2\tilde{q}^2\tau + \left(a(1-\phi)+\phi\right)^2 \mathbb{E}(x^2(z,\tilde{q}))\mu e_I\right)\nu^2\bar{q}^2 dt \end{aligned}$$

This variance is driven by two distinct channels: fluctuations in the baseline non-innovation-related firm value and the stochastic nature of the R&D process itself. While both components scale with firm size $\tilde{q}$, the volatility stemming from non-innovation shocks grows at a faster rate. This implies that inventors in larger firms are exposed to more intense noise unrelated to their research output. Consequently, larger firms face more severe agency frictions, as the equity stake held by the inventor becomes a less effective instrument for stimulating R&D effort. Based on the menu of available employment contracts, the inventor selects the optimal firm size $\tilde{q}$ that maximizes her objective function.

The inventor's disutility of choosing effort e_I in a firm $\tilde{q}$, $R(e_I, q, \bar{q})$, is given by Equation (3). Using the normalized $\tilde{q}$, it can be rewritten as $R(e_I, \tilde{q}) = \frac{1}{1+\delta}e_I^{1+\delta} + \eta\tilde{q}$.

The flow utility is defined as:

$$U_I(c_I, e_I) = \left(\left(a(1-\phi)+\phi\right) \mu e_I x_0 \left(z, \tilde{q}\right) + T_I - A \left(a^2 \tilde{q}^2 \tau + \left(a(1-\phi)+\phi\right)^2 \mathbb{E}(x^2(z,\tilde{q})) \mu e_I \right) - R(e_I, \tilde{q}) \right) \nu \bar{q}, \tag{49}$$

The inventor's objective is to choose an effort level that maximizes this utility:

$$\max_{e_I} U_I(c_I, e_I).$$

The first-order condition (FOC) characterizes the optimal effort:

$$e_I = \left(\left(a(1-\phi)+\phi\right) \mu x_0 \left(z, \tilde{q}\right) - A \left(a(1-\phi)+\phi\right)^2 \mathbb{E}(x^2(z,\tilde{q})) \mu \right)^{\frac{1}{\delta}} \tag{50}$$

which is invariant to ν.

The equilibrium wage rate T_I is determined via Bertrand competition:

$$\begin{gathered} (1-a)\left(1-\phi\right) \mu e_I \left[\mathbb{E}\left(V(j|q + x(z,\tilde{q})\bar{q}) - V(j|q) \right] \right) - T_I V(\bar{q}) = 0 \\ T_I = (1-a) \mu e_I x_0(z, \tilde{q}) \end{gathered}$$

which is invariant to ν.

Therefore, the instant utility is linear in $\nu\bar{q}$:

$$\begin{aligned} U_I(c_I, e_I) = & \left(\left(a(1-\phi)+\phi\right) \mu e_I x_0 \left(z, \tilde{q}\right) + T_I - A \left(a^2 \tilde{q}^2 \tau + \left(a(1-\phi)+\phi\right)^2 \mathbb{E}(x^2(z,\tilde{q})) \mu e_I \right) - R(e_I, \tilde{q}) \right) \nu \bar{q} \\ = & u_I(z, \tilde{q}) \nu \bar{q} \end{aligned} \tag{51}$$

where $u_I(z, \tilde{q})$ is unaffected by ν.

The inventor's firm selection problem can be written as:

$$q^* = \arg\max \{ u_I \left(z, \tilde{q}\right) \nu \bar{q} \} \tag{52}$$

Since $\nu\bar{q} > 0$, the decision rule is the same as:

$$q^* = \arg\max \{ u_I \left(z, \tilde{q}\right) \} \tag{53}$$

This demonstrates that the normalized decision rule $\tilde{q}^*(z)$ only depends on parameters.

Appendix C Quantitative Exercise

C.1 Earning standard deviation to mean ratio

The estimation uses a parametric approach. Assume that inventor earnings c_I follow a log-normal distribution:

$$c_I \sim \text{LogNormal}(\phi,\, \sigma_e^2), \quad \phi \in \mathbb{R},\ \sigma_e > 0.$$

The standard-deviation-to-mean ratio (CV) depends only on σ_e:

$$\text{CV} = \frac{\sqrt{\text{Var}(c_I)}}{\mathbb{E}[c_I]} = \sqrt{e^{\sigma_e^2} - 1}.$$

The Lorenz curve $L(p)$ gives the share of total earnings held by the bottom fraction p of the population. For a log-normal distribution, it has the closed-form expression

$$L(p) = F\big(F^{-1}(p) - \sigma_e\big)\,,$$

Akcigit and Goldschlag (2025) reports the empirical Lorenz curve of inventor earnings. Given this curve, the parameter σ_e can be recovered, and the implied standard-deviation-to-mean ratio follows directly from the closed-form expression above.

C.2 Counterfactual Decomposition (Leave-One-Out)

To further isolate the quantitative contribution of each adjustment margin, this appendix reports an alternative "leave-one-out" decomposition. In this exercise, all three counterfactuals are decomposed by shutting down one channel while permitting the other two margins to re-optimize. Results are reported in Table C3.

The results are consistent with the sequential exercise reported in Table 6. Because the participation channel is added last in the sequential decomposition, its marginal contribution coincides with the result of removing it in the leave-one-out exercise. The reallocation channel yields magnitudes similar to those in the sequential exercise, especially for the 0% and 10% ownership cases, implying that it has limited interaction with the participation channel. This is mainly because the inventor participation rate changes little in both cases. The contract channel has the largest discrepancy in magnitude. Shutting down the contract channel has two effects. One effect is to change inventors' incentives. The second effect is to change the equilibrium inventor allocation relative to the full counterfactual with all three channels active. In the 0% ownership case, shutting down the contract channel leads to a larger drop in the growth rate, since firms cannot compensate for inventors' lost direct ownership by offering more equity. In the 10% ownership case, shutting down the contract channel leads to a large increase in the growth rate. This is mainly due to the incentive effect. The baseline equity scheme provides a higher equity level than the new optimum. Combined with the

increase in the inventor's direct ownership, this over-incentivizes inventors, leading to higher effort. In the 50% ownership case, however, the second reallocation effect dominates. The equity provided in the baseline economy is substantially higher than the optimal level. As a result, it restricts the inventor's reallocation and leads to a smaller growth gain. Taken together, both decompositions point to the same conclusion. The interaction between contracts and reallocation, rather than any single channel in isolation, drives the aggregate growth response to ownership shifts.

Table C3: Leave-One-Out Decomposition of Macroeconomic Growth

	Baseline	**Counterfactual Inventor Ownership Regimes**		
Moment	(2% Ownership)	0% Ownership	10% Ownership	50% Ownership
Growth Rate	**2.01%**	**2.01%**	**2.03%**	**3.09%**
Percentage Change(Shutting Down Individual Channels)				
Without Participation Channel	—	-0.05%	+0.38%	+44.19%
Without Reallocation Channel	—	-0.07%	+0.41%	+7.12%
Without Contract Channel	—	-1.26%	+4.93%	+53.51%
Total Percentage Changes	—	-0.15%	+0.66%	+53.54%
Marginal Percentage Contribution				
Participation Channel	—	+63.33%	+42.42%	+17.46%
Reallocation Channel	—	+50.00%	+37.12%	+86.70%
Contract Channel	—	-746.67%	-652.27%	+0.06%

Notes: This table reports the "leave-one-out" decomposition of the aggregate growth rate. **Panel A** reports the percentage change relative to the baseline when the specified channel is rigidly fixed at its baseline optimal policy while all other margins re-optimize in general equilibrium. **Panel B** isolates the percentage contribution of that specific channel to the total change in General Equilibrium (GE) growth induced by the ownership shock. The contribution is mathematically defined as $(g_{\text{GE}} - g_{\text{shut_down}})/(g_{\text{GE}} - g_{\text{baseline}}) \times 100$.

Appendix D Discussion on Assumptions

This paper is based on several assumptions. This section discusses assumptions, extensions and future research. The baseline assumes no commitment regarding ownership allocation. An alternative setting is to allow inventors to commit to transferring part of their statutory share to the firm, but not vice versa, reflecting the disparity in legal resources between individual employees and corporate employers (Alok and Subramanian, 2023). This would alter the counterfactual predictions at high ϕ: when inventors hold substantial statutory ownership, they could voluntarily transfer part of it to improve risk-sharing. In the limit, when $\phi \to 1$, the result converges to the full-commitment case as in Section 6.3.

In the baseline model, each inventor is matched to her optimal firm. One natural extension is to incorporate a more sophisticated search and matching process. Search frictions would prevent some

inventors from reaching their preferred firm. As a result, the reallocation of inventors between small and large firms in response to ownership shifts would be less pronounced than the counterfactuals suggest.

Another important assumption is that inventors are short-lived, which abstracts from repeated interactions. This assumption is reasonable for one-time inventors, who account for more than 50% of inventors in the USPTO sample period. For serial inventors, long-lived agents would allow firms to sustain implicit incentive contracts through reputation, which could partially mitigate the agency friction. However, credit attribution within large firms is noisy and slow, so the core mechanism would be attenuated but not eliminated. A full dynamic contracting analysis of this channel is left for future research.

D.1 Discussion on Results

One natural question is why the firm doesn't promise the inventor a certain share of her innovative outcome? It seems a natural cure for the problem. In other words, why do firms enforce pre-invention assignment agreements?

There are three possible reasons. First, the gray area in common law (Section 2) creates a hold-up risk. Without a pre-invention assignment agreement, inventors could exploit the ambiguity to extract rents from employers. Though pre-invention assignment agreements lower innovative productivity, they protect firms from this hold-up risk. Second, the US lacks the institutional infrastructure required to implement a more sophisticated pre-invention assignment contract that specifies the shares belonging to the inventor and the employer. In fact, the German Employee Inventions Act (Gesetz über Arbeitnehmererfindungen) explicitly states the invention assignment rules between inventors and firms, and does not allow for private negotiation. Implementing such a system requires a careful evaluation of every innovation's value, which is extremely demanding and calls for institutional support. In Germany, the Arbitration Board under the Employee Inventions Act is in charge of settling disputes between inventor-employees and their employers. However, in the United States, such institutional infrastructure does not exist, making implementation infeasible. Third, institutional inertia may prevent reform even when the system is suboptimal. Hvide and Jones (2018) document that after Norway shifted innovation ownership from professors to universities, innovative output dropped by half, yet no policy reversal followed.