THE MACROECONOMIC IMPLICATIONS OF R&D OWNERSHIP *

Shaoshuang Yang†

Click for Update

Abstract

This paper examines how R&D ownership affect inventor incentives, innovation allocation, and aggregate productivity. Exploiting variation in U.S. state-level non-compete law enforcement as exogenous shocks, I show empirically that losing R&D ownership reduces inventor output, with larger negative effects in big firms. I develop a model addressing how inventors self-select into firms of different sizes because the equilibrium contract depends on how much firm value is tied to the inventor's own effort. The model predicts that (1) when inventors retain a larger share of R&D, they work harder; (2) when firms own R&D, incentives are diluted and inventors exert less effort, especially in large firms; (3) the inventor distribution across firm sizes depends systematically on the ownership regime. A counterfactual simulation quantifies how alternative U.S. innovation ownership policies would reallocate innovation effort and alter long-run growth. Results highlight the role of intellectual property institutions in shaping both micro-level incentives and macro-level innovation outcomes.

JEL codes: O31, L11.

Keywords: innovation ownership, firm dynamics, inventor incentive, non-compete law, aggregate growth rate.

*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, while its production, R&D, inherently suffers from a principal-agent problem (Merges, 1999). The tension of this principal-agent friction, which between the legal reality of employer-owned R&D and the economic importance of employee-inventor incentives, depends critically on the ownership of the R&D outcome **(TBD: cite)**. 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. I begin by documenting stylized facts regarding innovation output across firm sizes and R&D ownership regimes. When institutional environments shift and inventors lose ownership of their inventions, their productivity declines; crucially, this negative effect is more severe in large firms. Motivated by these facts, I build an endogenous growth model featuring that 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 fundamentally 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.In the counterfactual exercise, the elasticity of the aggregate growth rate with respect to employee-inventors' R&D ownership is 0.3. This implies that a 1% increase in employee-inventors' R&D ownership raises the aggregate growth rate by 0.3%.

This paper first documents the empirical impacts of the R&D ownership assignment. I exploit the 2008 ruling by the U.S. Court of Appeals for the Federal Circuit in *DDB Technologies LLC v. MLB Advanced Media, LLP*. This decision abruptly invalidated pro-employee state statutes regarding preinvention assignment agreements (PIAAs) in eight states. This

generated an exogenous shift in R&D ownership from employee-inventors toward employers (Baniak and Dawson, 2008). In the treated states, average inventor output drops by 4.09% relative to the control group, as measured by quality-adjusted patents. In terms of aggregate economic magnitude, this shock is associated with a total loss of $0.32 billion (in 2025 dollars). On the intensive margin—examining a balanced panel of inventors who innovate every year—the ownership shift induces a severe 12.4 percent decline in output per unit of exposure. On the extensive margin, the growth rate of active inventors also drops. Crucially, I document a heterogeneity among firms of different sizes: all negative effects are much stronger in large firms. For example, relative to their counterparts in small firms, inventors in large firms experience an additional 2.66 percent decline in output. For the average innovative public firm, which generates $37.1 million in annual innovation value, this differential implies that each large firm incur an additional loss of at least $0.97 million per year. The results support that the offering inventor R&D ownership 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.

Motivated by the empirical findings, I develop an endogenous growth framework featuring R&D ownership. The R&D ownership is governed by an exogenous parameter, which measures the share that belongs to employers. Each inventor is endowed with an idea of idiosyncratic quality (a la Abowd et al. (1999)) and chooses to join a firm of the selected size. To attract inventors and elicit unobservable innovative effort, firms must design an optimal equity-wage contract. When solving this contracting problem, a firm faces a trade-off: providing a larger equity stake sharpens the inventor's incentive to exert effort, but simultaneously exposes the risk-averse inventor to greater variance in consumption. In each firm, the employee-inventor's compensation scheme consists of a flat wage, equity-based incentives (Onishi, 2013) and the inventor-share of R&D. The incentive, a combination of the latter two types of compensation, determines how much effort an inventor exerts. The counterfactual exercise on R&D ownership shows that when a decrease in inventor ownership suppress innovative output, especially in large firms. Moreover, both the aggregate and the distributional effects are consistent with their empirical counterparts.

At the core of the theoretical framework is a principal-agent problem driven by the unobservable nature of an inventor's effort. The total incentive for an inventor to exert effort is a combination of two elements: their legal share of direct R&D ownership and the equity claim offered by their employer. Because the legal baseline for R&D ownership is set by laws and applies uniformly across all firms, any variation in incentives arises entirely from the equity-wage contracts designed by individual employers. However, offering an equity share forces the firm to navigate a fundamental trade-off between sharpening the inventor's

incentives and exposing them to the broader volatility of the business. The effectiveness of these equity contracts depends on firm size. In a small firm, a single inventor's idea accounts for a large fraction of the firm's total value, allowing the firm to provide high-powered equity incentives. In contrast, the equity of a large firm is tied to many unrelated activities, which dilutes the incentive for any single inventor and exposes them to risks outside their control. Consequently, large firms are structurally less effective at using equity to motivate talent, leading to weaker overall incentives and lower innovative effort when direct R&D ownership is assigned to employers.

Counterbalancing this inherent incentive disadvantage, large firms possess an advantage in the form of synergy. In the model, the expected economic value of an inventor's idiosyncratic idea is amplified when implemented by a larger firm, as incumbents 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 larger firm increases the management cost of the innovation. Consequently, the framework features an endogenous selection effect: inventors self-select into firms of different sizes based on the fundamental trade-off between the high-powered equity incentives of small firms and the superior synergy of large firms. Because ideas are heterogeneous in quality, inventors evaluate this tension alongside the prevailing legal ownership regime. Ultimately, this trade-off 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 that are consistent with the empirical findings. First, when inventors retain ownership rights, they act as residual claimants; this naturally incentivizes higher effort and reduces the necessity for equity compensation. Second, when the firm's share of R&D ownership increases, the firm optimally increases the equity stake offered to the inventor to substitute direct ownership with equity incentives. However, this offset is incomplete. Incentives are diluted and inventor's effort decreases, especially in large firms. Third, the assignment of inventors across firm types depends systematically on the ownership regime. When the firm ownership increases, inventors opt for smaller firms to compensate for the loss of direct ownership through stronger equity incentives. Collectively, these predictions rationalize the empirical patterns and structurally link the R&D ownership regime to aggregate productivity.

To quantify how R&D ownership affects economic growth and the distribution of innovation across firms, I calibrate the model and run counterfactual analysis that reassign R&D ownership between inventors and firms. The model benchmark is calibrated to the United States, assuming 98% R&D ownership is assigned to employers. In the counterfactual, I

hold all parameters fixed except the inventor ownership share. I then compute the change in the annual growth rates and decompose it into (1) change driven by increase in effort level holding the assignment of inventors across firms unchanged, and (2) change driven by the shift of inventor distribution across firms. I also compute the change in two important moments: the productivity of inventors, and the share of innovations developed in small firms relative to large firms.

I find that shifting R&D ownership from firms back to inventors leads to a higher economic growth, consistent with data. In the baseline economy (characterized by 98 percent firm ownership), the relevant growth rate component is 1.01 percent. When an additional 5 percent of R&D ownership is assigned to inventors, the growth rate increases to 1.03 percent, translating to a 1.4 percent relative increase in aggregate growth, with an implied elasticity of approximately 0.3. Meanwhile, this aggregate gain is accompanied with the reallocation of innovations between small and large. The share of innovations produced by small firms drops dramatically by 40.5 percent (from 0.201 to 0.121 TBD). This quantitative dynamic aligns with the direction of the empirical findings: 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 by narrowing their structural incentive disadvantage. The quantitative analysis thus highlights a trade-off governed by legal institutions: while higher inventor ownership promotes aggregate growth through better incentives and synergy utilization, it heavily concentrates innovative activity within large incumbent firms, suppressing small-firm innovation dynamics.

This paper further decompose the results into intensive margin (inventor effort) and extensive margin (inventor allocation). The intensive margin captures the growth change driven exclusively by the increase in optimal inventor effort, holding the existing distribution of inventors across firms fixed. A 5 percent inventor R&D ownership increment lifts the aggregate growth rate by 1.5 percent. The small firm share of innovation drops by 40.5 percent. The extensive margin isolates the growth generated by the endogenous shift of inventors across the firm-size distribution. When I shift 5 percent of R&D ownership to inventors, it increases the aggregate growth rate by 1 percent. The small firm share of innovation drops by 35 percent. This decomposition demonstrates that the macroeconomic benefit of pro-employee ownership regimes stems jointly from improved individual motivation and a systemic shift of inventive talent.

Related Literature This paper develops a macroeconomic framework of endogenous growth that explicitly models the internal incentive structures driving aggregate innovation. Traditionally, endogenous growth literature treats the firm as a unitary decision-maker when modeling innovation dynamics, market power, and resource reallocation (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 naturally intersects with investigations into the innovation productivity, which has been declining (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. Reconciling the two facts of declining R&D productivity and the rise of superstar firms (Autor et al., 2020), this paper shows that as firms grow larger, the inherent difficulty of providing high-powered incentives systematically dampens the R&D productivity of their inventors.

To empirically and theoretically ground the agency frictions, this study formalizes the role of ownership in R&D innovation. Theoretical literature establishes that assigning residual control rights to the inventor is crucial for incentivizing inventor effort (Grossman and Hart, 1986; Aghion and Tirole, 1994; Hart, 1995), a dynamic supported by empirical studies examining the relationship between inventor incentive and innovation outcomes (Lerner and Wulf, 2007; Onishi, 2013; Giarratana et al., 2018; Hvide and Jones, 2018; Alok and Subramanian, 2023; Suh, 2023). For example, Lerner and Wulf (2007) shows that compared with a flat wage, paying by equity increase inventor R&D incentive and improve innovation outcomes. However, this microeconomic literature has primarily focused on partial equilibrium outcomes without tracing the macro-level consequences. My 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 how inventors self-select into firms of different sizes. This approach enriches the classic literature on firm boundaries (Coase, 1937; Hart and Moore, 1990, 2008) and the allocation of R&D. Recent macroeconomic contributions have explored how firm's size and scope affect its innovation strategy (Akcigit et al., 2016; Akcigit and Kerr, 2018; Ma, 2022; Jo and Kim, 2024). This paper provide an additional angel 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 agent determining innovative input (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. I argue 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-firm ownership regimes, the lack of incentive pass-through severely blunts the macroeconomic effectiveness of innovation policy.

2 Empirical Patterns

This section establishes the empirical motivation for the incentive channel. The effective allocation of R&D ownership in the United States is jointly determined by two legal institutions: preinvention assignment agreements (PIAAs) and non-compete agreements (NCAs). To isolate the causal effect of transferring this ownership from inventors to employers, the analysis exploits a quasi-natural experiment generated by an abrupt shift in federal case law. Empirical results demonstrate that inventor innovative output respond negatively to the pro-employer judgment. The magnitude of this response is heavily mediated by firm size, providing motivating evidence for a structural incentive channel that will be formally embedded into the general equilibrium model in Section 3.

2.1 Institutional Background

To systematically evaluate how R&D ownership shape 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 dual-faceted: 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 intersecting legal institutions. Non-compete agreements (NCAs) restrict inventor mobility and capture the ownership of human capital, 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 our 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 a firm (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 (**?**). 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 contractual instrument is widely used in the United States, especially among innovating firms, which require prospective employees to sign PIAAs as a precondition for employment (**??**). 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 prior to and past 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). TBD. Under these statutes, inventors in these states enjoy a legally protected degree of residual ownership over a subset of their human capital and the associated inventive output.

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* [1]. 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—an innovation entirely unrelated to his employer's line of business as confirmed by Schlumberger's general counsel and the lab director when he invented the technology. 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 because patents are federal grants, questions of patent assignment are governed by federal law, overriding state-level regulations. This ruling effectively caused 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 patent rights toward employers by invalidating state-level PIAA exemptions. In the following section, I explain the identification using this empirical setting in regression analysis.

[1] 517 F.3d 1284, 1290 Fed. Cir. 2008

2.2 Identification

Following Suh (2023), the empirical strategy employs a difference-in-differences framework that exploits the exogenous shift in R&D ownership generated by the 2008 *DDB* ruling. The treatment group comprises individuals working for firms located in the eight states directly affected by the CAFC's invalidation of pro-employee statutes.[2] While the 2008 ruling uniformly shifted statutory patent rights toward employers across all eight treated states, the practical severity of this shock varied cross-sectionally depending on the inventor's pre-existing control over her human capital, governed by state-level NCA enforcement. Specifically, in states with strict NCA enforcement, inventors already faced severe mobility restrictions and possessed weak outside options; the marginal impact of losing patent rights via the ruling shift was relatively muted. Conversely, in states with weak NCA enforcement, inventors retained high mobility and valuable outside options to commercialize independent ideas. For these inventors, the sudden invalidation of free-to-create statutes constituted a severe destruction of their effective R&D ownership. To accurately capture this heterogeneous treatment intensity, the empirical specification defines exposure inversely to NCA stringency

$$\text{Exposure}_{is} = \max_{s'}(\text{NCA}_{s'}) - \text{NCA}_s, \tag{1}$$

where NCA_s is the NCA Enforcement Index for state s as measured by Garmaise (2011), and the maximum is taken over all states. A higher value of Exposure_{is} indicates that inventor i in state s enjoyed weaker NCA protection prior to the shock, meaning that the DDB ruling had a larger practical impact on her effective ownership. This construction implies that within the eight treated states, inventors in lower-NCA-enforcement states experienced the largest ownership shift. To identify the causal impact of R&D ownership on innovation, the empirical specification exploits this heterogeneous exposure, interacting the abrupt 2008 legal shock with the pre-existing variation in NCA enforceability.

To estimate the impact of R&D ownership on innovation output, the analysis employs a difference-in-differences (DID) specification that exploits the 2008 *DDB Technologies LLC v. MLB Advanced Media, LLP* decision. The baseline empirical model is estimated at the inventor-firm-state-year level:

$$y_{ijs,t+2} = \alpha + \beta\,(\text{Exposure}_{is} \times \text{Post}_t \times \text{Treat}_s) + \mathbf{X}'_{ijt}\boldsymbol{\omega}_{\boldsymbol{X}} + \eta_t + \gamma_s + \theta_j + \varepsilon_{ijst}, \tag{2}$$

where $y_{ijs,t+2}$ is the quality-weighted patent output[3] of inventor i, employed at firm j, in state s, in year $t+2$. I measure the number of granted patents at $t+2$ to account for the patent system process Hall et al. (2001). Post_t is an indicator equal to one for years 2008

[2] 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).

[3] $y_{ijs,t+2} = \sum_{p \in P(i,t)} \ln(1 + \text{citation}_p)$, where $P(i,t)$ is the sample of patents granted to inventor i at year t. The citation is truncated at the 5-year window to ensure that cohorts are compared on common ground, mitigating right-censoring issues.

and after. Treat_s is an indicator equal to one for the eight treated states. The exposure measure, Exposure_{is}, is defined in equation (1). The specification includes year fixed effects η_t, state fixed effects γ_s, and firm fixed effects θ_j. The vector of observables $\mathbf{X}_{ijt}$ controls for inventor gender, time-varying firm size[4], and inventor-level technology-class (CPC) fixed effects[5]. For the balanced-panel specification, I replace the CPC fixed effects with inventor fixed effects α_i, providing a within-inventor estimate of the ownership effect. The coefficient of interest, β, captures the differential change in innovative output for inventors who were more exposed to the DDB shock, relative to less-exposed inventors in the same treated states and relative to inventors in control states.

The central identifying assumption of this empirical framework is an inventor-level parallel-trends condition. Specifically, the causal interpretation of β requires that, absent the 2008 *DDB* decision, the changes in innovative output ($y_{ijs,t+2}$) for inventors in treated states with higher exposure would have evolved similarly to those with lower exposure, as well as relative to inventors in unaffected control states, conditional on the specified fixed effects and observable characteristics. Furthermore, this identification strictly requires that there was no differential anticipation of the federal legal shock prior to 2008, and no contemporaneous state-level economic shocks that were differentially correlated with the pre-existing NCA exposure within the treated states. Because the treatment shock and the continuous exposure measure are both structurally derived 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 (2) is augmented with an interaction term:

$$\beta_2 \left(\text{Small}_{ijt} \times \text{Exposure}_{is} \times \text{Post}_t \times \text{Treat}_s\right), \tag{3}$$

where Small_{ijt} is an indicator equal to one if firm j employs 500 or fewer workers. The coefficient β_2 isolates the moderating effect of firm scale. If the theoretical incentive channel holds, the decline in innovative output following the pro-employer shift should be systematically attenuated in smaller firms, generating a positive estimate for β_2.

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

[5]More than 50% inventors are only granted with one patent in the sample period. To avoid the selection issue caused by adding an inventor fixed effect, I control for inventor observables (gender and CPC) instead. The same analysis is implemented on a dataset with inventors who develop patents every year (balanced-panel), in which I substitute the observables with individual fixed effects.

2.3 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 exclusively 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 economics of innovation literature, forward citations serve as a robust proxy for technological quality and economic value, 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 qualifies for this statutory designation if it employs 500 or fewer workers. This binary indicator is utilized throughout the analysis to precisely distinguish between small, highly-equitized firms and large incumbent organizations.

Small Entity Indicator.— To explore how the response to ownership shocks varies across firm scale—a dimension of heterogeneity that directly motivate the theoretical mechanism developed in the next section—the empirical specification requires a measure of firm size. This is captured using the USPTO's small entity status indicator. A firm qualifies for this statutory designation if it employs 500 or fewer workers. This binary indicator is utilized throughout the analysis to distinguish between small firms and large incumbent organizations, allowing us to document the stylized facts regarding firm-size heterogeneity.

Our baseline sample covers the period 2006 to 2013, spanning four years before and four years after exogenous legal shock, consistent with Suh (2023). 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. To examine the intensive margin of inventor effort, the analysis also constructs a balanced panel by restricting attention to inventors who generate patents in every single year of the sample window. The resulting dataset contains 1,046,256 inventor-year observations in the full, unbalanced sample, and

24,860 inventor-year observations in the balanced panel.

2.4 Results

Table 1: Relationship between Innovation Output and Pro-employer Ruling Treatment

	Innovation Output				g_{inventor}	
	All Inventor Sample		Balanced Sample		All Firm Sample	
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{Exposure}_{is} \times \text{Post}_t \times \text{Treat}_{is}$	-0.020***	-0.022***	-0.124***	-0.131***	-0.005	-0.007**
	(0.007)	(0.008)	(0.040)	(0.040)	(0.003)	(0.003)
$\text{Small}_{ijt} \times \text{Exposure}_{is} \times \text{Post}_t \times \text{Treat}_{is}$		0.013*		0.205***		0.009***
		(0.007)		(0.04)		(0.002)
Small_{ijt}	-0.505***	-0.526***	0.031	-0.236	-0.049	-0.062
	(0.059)	(0.065)	(0.557)	(0.487)	(0.044)	(0.040)
gender	0.217***	0.217***	0.745**	0.745**	-0.077***	-0.077***
	(0.026)	(0.026)	(0.362)	(0.362)	(0.028)	(0.028)
Observations	1,046,256	1,046,256	24,860	24,860	55,582	55,582
R-squared	0.148	0.148	0.150	0.150	0.198	0.198
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Inventor CPC FE	Yes	Yes	No	No	Yes	Yes
Inventor FE	No	No	Yes	Yes	No	No

Notes: Standard errors are clustered at the state level. All depend variables are measured at $t + 2$. Columns (1)–(2) use the full inventor sample with quality-weighted patent output as the outcome. Columns (3)–(4) use the balanced panel (inventors present in every year) with quality-weighted patent output as the outcome. Columns (5)–(6) use the full firm sample and the growth rate of active inventor number in each firm. All specifications control for year, state and firm fixed effect. The inclusion of inventor's technology class and inventor fixed effects is indicated in the last two rows. *** Significant at the 1 percent level; ** significant at the 5 percent level; * significant at the 10 percent level.

Table 1 reports the estimation results of the regression specifications in section 2.2. The panels are organized by sample and outcome variable. Columns (1) and (2) utilize the full inventor sample to evaluate overall quality-weighted patent output. Columns (3) and (4) restrict the sample to a balanced panel of continuously active inventors to isolate the intensive margin. Finally, columns (5) and (6) shift the unit of observation to the firm level to analyze the extensive margin, measured as the annual growth rate in the number of active inventors [6]. For each set of the results, I report the average estimates first, followed by the exploration of firm-size heterogeneity. Across all specifications, the evidence demonstrates

[6] I use the annual growth rate, instead of the total amount of active inventors, because the amount of inventors in a firm can be serially correlated.

that shifting statutory R&D ownership toward employers systematically depresses innovative output, and this distortion is heavily mitigated within small firms.

Columns (1) and (2) present the baseline results, showing that shifting R&D ownership toward employers reduces innovation. The coefficient on the triple interaction Exposure × Post × Treat is −0.020. This estimate implies that, in terms of patents of average-quality, for an inventor in a treated state at the average level of exposure, the DDB ruling reduced output by approximately 4.09% of an average inventor's output. In terms of dollar value, on aggregate level, this drop is about $1 million, per year per firm relative to inventors in control states [7]. The coefficient of the small-firm interaction in column (2) is 0.013, meaning that the working in a small firm offsets more than half of the baseline negative impact.

This dampening effect is substantially amplified among high-productivity inventors. Because approximately 50 percent of inventors in the full sample appear only once during the sample period, I further isolate the intensive margin by analyzing a balanced panel restricting the sample to inventors who appear in the dataset annually. The results are reported in columns (3) and (4), where inventor fixed effects are included. In this high-productivity subsample, the treatment effect is larger. The coefficient is −0.124 in column (3), approximately six times larger in magnitude than the all sample average estimate in column (1). It implies that the high-productivity inventors are much more exposed and negative affected than average ones. Consistent with the baseline result, column (4) shows that this effect is significantly mitigated in small firms, suggesting that small firms are better at coping the decline in employee R&D ownership.

The ownership shock also distorts the extensive margin of innovation, discouraging the retention and participation of inventive talent. Columns (5) and (6) shift the unit of observation to the firm level and study the growth rate of active inventors [8], which implies the extensive margin of innovation activity. The coefficient on the triple interaction in column (5) is −0.005, which is not significant, meaning that on average there is no change in the trend of inventor workforce. Once control for small firm, the coefficient becomes significant and the small-firm interaction coefficient is positive, indicating that the ownership shift discouraged inventor participation especially in large firms. This result suggests that the ownership effects operate not only by reducing the output of existing inventors but also by decreasing the growth rate of active inventors, mainly for large firms.

To sum up, when innovation ownership shifts from inventors to employers, overall inno-

[7] The economic value of a patent is measured following Kogan et al. (2017)

[8] The growth rate of active inventors is defined as $g_{\text{inventor},t+1} = \sum N_{j,t+1}/N_{j,t}$, where $N_{j,t+1}$ measures the amount of inventors in firm j with a patent in year $t+1$. Alternatively, one can define the active inventor as someone has at least one patent in a five-year window. However, this definition is not feasible due to the data availability constraint.

vative outputs are affected negatively, especially for large firms. The results suggest that the R&D ownership matters for innovative output via an incentive channel. Additionally, the observed size heterogeneity suggests that, small firms are better able to incentivize inventors even when baseline direct ownership is stripped away. One potential explanation is that when employers hold more ownership, firms increase equity compensation to partially offset the reduction in inventors' direct stakes. But equity is a less powerful incentive device in large firms, where the link between any individual inventor's effort and the overall firm value is diluted by the many unrelated activities and contributors 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. I formalize this mechanism in the model in Section 3.

2.5 Related Stylized Facts

This section uses an exogenous shift in the statutes regarding invention ownership and shows that when innovation ownership shifts from inventors to employers, individual-level innovative outputs are affected negatively, especially the ones who work in large firms. The result is consistent with the findings in the literature. The inventor-level findings established in this section align with existing empirical evidence on the relationship between employee-inventor invention ownership rights, inventor incentive and innovation outputs.

The patterns are seen on individual-, coorporate- and state-level, inside and outside US. At the corporate level, Lerner and Wulf (2007) use European survey data to show that compensating R&D personnel with equity rather than standard wages increases self-reported effort and quality-weighted patent output. Suh (2023) uses firm-level innovative output confirms that the presence of a principal-agent friction. Analyzing Japanese firm data, Onishi (2013) finds that revenue-based innovation compensation increases patenting, and that this effect is more pronounced in small firms. This aligns with the firm-size heterogeneity documented in Section 2.4, suggesting that smaller organizations can more readily substitute statutory property rights with equity incentives. Beyond individual and firm-level responses, the allocation of R&D property rights aggregates to broader economic outcomes. Evaluating legislative variations across the United States, Alok and Subramanian (2023) find that evidence at the state level.

Additionally, the result is not limited to a corporate setting. Hvide and Jones (2018) observe a similar dynamic in a university setting, finding that a legal shift of invention ownership from researchers to universities in Norway lowered innovative output.

Taken together, these stylized facts—spanning European university laboratories, firms

across Europe, Asia and North America, and state-level aggregates—indicate that the relationship between inventor R&D ownership and innovation output holds across different datasets and environments. To formalize these empirical regularities and evaluate their macroeconomic consequences, the subsequent section embeds this micro-level incentive channel into a general equilibrium framework of endogenous growth.

3 Model

To formalize the empirical findings and trace their general equilibrium implications, this section develops a continuous-time model of endogenous 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 an incumbent firm of a specific 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.

3.1 Preferences

A representative household provides labor and consumes final goods. The household is long-lived. She supplies one unit of labor inelastically at one unit of human capital. The household's preferences are given by a utility function:

$$U_H = \int_0^\infty e^{-\rho t} log(C_H(t)) dt, \tag{4}$$

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

A continuous unit measure of risk-averse inventors 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 $e_I(t)$:

$$U_I\left(c_I(t), e_I(t), q(t)\right) = \mathbb{E}\left(c_I(t)\right) - A\frac{var\left(c_I(t)\right)}{V(t|\bar{q}(t))} - R\left(e_I, q\right) V(t|\bar{q}(t)), \tag{5}$$

where c_I is the consumption, $A > 0$ is the degree of risk aversion, e_I is the effort level, and $R(e_I, q)V(t|\bar{q}(t))$ represents the the associated disutility cost of effort. $\bar{q}(t)$ (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 to the average quality. The variance and cost are normalized by $V(t|\bar{q}(t))$ to keep

the problem stable over time.

Additional to engaging in risky innovation, an inventor possesses an outside option to enter the labor market alongside the representative household, supplying ξ units of effective human capital. This alternative yields a reservation utility of $\bar{U}_I(t) = c_I^L(t)$. In equilibrium, a share $l_I(t)$ inventors will opt to supply labor, while $(1 - l_I(t))$ inventors develop $R\&D$. Finally, aggregating across the unit measure of agents, the total consumption of the inventor cohort is given by $C_I(t) = \int_0^1 c_{Ii}(t)di$.

3.2 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. The former hires inventors to develop innovations, and produce intermediate goods. The latter assembles intermediate goods and produces final goods.

The final good producers assemble intermediate goods to produce final goods. Firms are owned by household. The final good producers produce final goods using a continuum of intermediate goods of measure 1 with production technology which is similar to Akcigit and Kerr (2018):

$$Y(t) = \frac{1}{1-\beta} \int_0^1 q_j^{\beta}(t) y_j^{1-\beta}(t) dj. \tag{6}$$

In this function, $q_j(t)$ is the quality of the intermediate good j, 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:

$$Y = C_H + C_I. \tag{7}$$

The intermediate goods producers are a continuum of measure 1 of risk-neutral firms. Each firm produces one type of good, with a linear technology using only labor:

$$y_j(t) = \bar{q}(t) l_j(t), \tag{8}$$

where $l_j(t)$ is the human capital augmented labor input. $l_j(t) = 1$ for a regular household, while $l_j(t) = \xi$ for an inventor-type labor supplier. $\bar{q} = \frac{1}{1} \int_0^1 q_j dj$ is the average quality level of the economy. The cost is linear in wage $w(t)$, which intermediate firms take as given. The labor market satisfies the constraint:

$$\int_0^1 l_j dj \leq 1 + l_I(t). \tag{9}$$

This paper focuses on the balanced growth path. I normalize the variables using the

average quality $\bar{q}$. I denote the normalized variables using tilde:

$$\tilde{q_j} \equiv \frac{q_j}{\bar{q}}. \tag{10}$$

3.3 Innovation

Innovation takes the form that improving the quality of existing products within the intermediate goods sector. Each inventor is born with one idiosyncratic idea of quality z_0. To transform this idea into an innovation, the inventor must affiliate with an incumbent firm of her choosing and exert effort, e_I. The innovation process is stochastic. The inventor's chosen effort level governs the arrival rate of a successful innovation according to a Poisson process:

$$\lambda(e_I) = \lambda_0 e_I, \tag{11}$$

where $\lambda_0 > 0$ is the baseline efficiency of the R&D technology. The inventor's effort e_I is unobservable and unverifiable by the firm (**??**).

If the inventor's unobservable effort successfully yields an innovation, the allocation of the resulting property rights is governed by an exogenous statutory regime. The legal environment dictates that the firm captures a share $\mu \in [0, 1]$ of the generated value, while the inventor retains statutory ownership over the remaining $1 - \mu$ share. To fully commercialize the invention, the firm must acquire the inventor's residual share at fair market value[9]. Furthermore, the framework assumes an asymmetry in contractual commitment. An inventor can credibly commit to transferring a fraction of her statutory share to the firm, but the firm cannot make a symmetric, enforceable commitment to the inventor. This assumption reflects the disparity in legal resources. In the event of a contractual dispute, the firm enjoys a disproportionate advantage over the individual employee, rendering any ex-ante promise by the firm to relinquish its statutory rights dynamically inconsistent(Alok and Subramanian, 2023). The combination of unobservable effort and divided statutory ownership generates a principal-agent friction. Because the firm cannot write a contract contingent on effort, and the inventor must bear the private cost of effort while only inherently capturing a fraction of the returns, a fixed wage will fail to elicit innovation. 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 (Onishi, 2013).

Once an innovation is realized, the economic value $x(z, \tilde{q}(t))$ is jointly determined by the idiosyncratic quality of the idea, z, and, $\tilde{q}(t)$, the size of the firm that commercializes it.

[9]This setting is to be consistent with the legal requirement. Another way to understand the parameter μ is that the firm needs to pay to compensate the inventor at $1 - mu$ share of the total value, otherwise the inventor can sell this part, either in the form of invention or human capital, to others.

Specifically, it takes the following form:

$$x(z, \tilde{q}(t)) = \gamma(\tilde{q}(t))z\bar{q}(t), \tag{12}$$

where γ can take two states, high or low, $\gamma \in \{\gamma_H, \gamma_L\}, \gamma_H > \gamma_L$. The probability to be in the high state is $h(\tilde{q}(t)), h \in (0,1), h' > 0$. On average, an innovation generates a higher expected value when implemented within a larger firm. This captures a synergy effect: large incumbent 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 increment is assumed to be scaled up with $\bar{q}(t)$ to capture how the total technology stock in the economy affects new innovations.

3.4 Employment Contracting Problem (In progress)

Intermediate firms hire inventors, following Bertrand competition, by offering an employment contract. The employment contract includes two part: a wage rate $T_I V(t|\bar{q}(t))$ and a share of equity $a \in [0,1]$.

The timing is as follows. All firms simultaneously offer firm-inventor-specific contracts to inventors. After viewing all contracts, the inventor chooses her favorite one and joins the firm. If the best contract is offered by $q = 0$, I call it a start-up. Even in this case, the inventor does not bear all the risk nor holding 100% equity. I assume that $q = 0$ still (e.g. venture capital funds) offer contracts with endogenous a.

References

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

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.

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.

Baniak, M. and T. Dawson (2008): "Discussion of Employer Assignment Agreements after DDB Technologies v. MLB Advanced Media," *Nw. J. Tech. & Intell. Prop.*, 7, 298.

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.

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.

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.

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.

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.

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.

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.

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.

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.