THE MACROECONOMIC IMPLICATIONS OF R&D OWNERSHIP *

Shaoshuang Yang†

Very preliminary - do not circulate

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, patent trade, firm dynamics, incentive constraint, firms' boundaries.

*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)

There is little consensus regarding policies governing employee-inventors' intellectual property rights. The U.S. Congress has repeatedly debated strengthening these rights (Hovell, 1982), and legal frameworks vary significantly across countries.In the United States, for example, pre-invention assignment agreements are standard practice, effectively assigning ownership to firms (Pisegna-Cook, 1993). Conversely, in countries like Germany, such agreements are prohibited, and inventors retain ownership. Other nations, such as Japan, fall between the U.S. and German regimes (Onishi, 2013). How does the allocation of intellectual property rights affect the economy? Some answers have been provided. A literature on intellectual property rights protection usually focus on the firm size, namely the trade-off between R&D investment incentive and monopoly power (**?**Acemoglu and Akcigit, 2012). In this paper, I highlight a novel mechanism: the inventor incentive channel. Providing both empirical evidence and a theoretical framework, I show that innovation ownership 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 R&D outcomes and growth.

This study provides a 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' intellectual property affect innovation through their interaction with inventor incentives. I begin by documenting stylized facts regarding individual innovation output across firm sizes and R&D ownership regimes. I show that 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 with a micro-foundation based on inventor incentives. The theory focuses on the inventor's joint decision of how much effort to exert and which firm to join, conditional on the ownership regime. The central mechanism relies on the self-selection of inventors into firms of different sizes: the equilibrium inventor-firm contract depends heavily on the extent to which firm value is tied to the inventor's own effort. Because of the complexity and uncertainty inherent in R&D, contracts are incomplete, and I assume that firms can only use equity-based incentive. Consequently, startups can offer high-powered

incentives because their value depends almost entirely on the inventor's idea, while large firms suffer from incentive dilution as their equity value is driven by many unrelated factors. I calibrate the model to match empirical moments on innovation ownership, inventor output heterogeneity, and innovation value. 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 doing so, this paper complements the endogenous growth literature, which has previously focused primarily on the firm side of innovation, by explicitly micro-founding the inventor's choices.

I first examine how the 2008 federal ruling, which weakened inventors' ownership rights over RD outcomes, affected innovation. I find that when RD ownership shifts from employee-inventors to employers, average inventor output drops by 19.0 percent 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 $132.5 billion (in 1982 dollars). I also test the extensive margin by analyzing the change in active inventors at the firm level. The results are negative, indicating that treated firms experience significantly slower growth in their stock of innovative personnel.

I then investigate heterogeneous effects by introducing firm size as an additional dimension. A triple-difference regression analysis reveals that the negative impact of firm-held R&D ownership is significantly more severe for large firms. Relative to their counterparts in small firms, inventors in large firms experience an additional 8.93 percent decline in output. For the average innovative public firm, which generates $11.1 million in annual innovation value (in 1982 dollars), this differential implies that each large firm incur an additional loss of at least $0.99 million per year. Furthermore, patent-level regressions show that the economic value of a given innovation increases with firm size, suggesting a strong complementarity between inventors and firm assets. Consequently, the disproportionate drop in patent quantity in large firms translates into an even larger loss in total economic value. Finally, on the extensive margin, the slower growth in innovative personnel is more pronounced for large firms, implying that these firms become relatively less attractive to inventors under the new regime. Why does R&D output depend negatively on employee-inventor's ownership of innovation, and why is this effect heterogeneous across firm sizes? I argue that the primary channel is the inventor's incentive to exert innovative effort. Motivated by the empirical findings, I develop a novel endogenous growth framework micro-founded on the inventor's innovative effort and affiliation choice. In the model, the R&D ownership regime is governed by an exogenous parameter, which measures the share of ownership that belongs to the inventor. Each inventor is endowed with an idea and idiosyncratic productivity (a la Abowd, Kramarz, and Margolis, 1999) and chooses to join the optimal firm, who differs in

size; I model the decision to found a startup as choosing a firm of size zero. The employee-inventor's compensation scheme consists of a flat wage and equity-based incentives (Onishi, 2013). Crucially, while equity is tied to total firm value, the inventor's effort affects only their specific R&D project. As firm size increases, the marginal contribution of a single project to total firm equity diminishes. Consequently, when inventors retain ownership, they act as residual claimants and face no incentive problem. However, when ownership shifts to the firm, an agency friction emerges. This friction is particularly severe in large firms due to the dilution of equity incentives. Thus, for a given inventor-firm pair, the inventor exerts less effort under firm ownership, with the decline being most pronounced in large firms. Furthermore, shifting ownership to employers alters the allocation of inventors, inducing more inventors to self-select into startups. These micro-level choices aggregate to determine firm entry, R&D productivity, and overall economic growth.

The model features the interaction between two types of heterogeneous agents: inventors and firms. Inventors are endowed with idiosyncratic R&D productivity, while firms operate at varying sizes that are endogenously determined by the R&D process. The allocation of inventors across firms is driven by two key forces: the inventor-firm contract and commercialization synergy. The inventor-firm contract governs the inventor's compensation scheme and is designed to induce optimal effort, which is crucial for R&D success but unobservable to the firm. Due to the inherent complexity and uncertainty of innovation, employment contracts are incomplete and cannot specify all dimensions of effort ex ante (Alok and Subramanian 2023). Consequently, the model focuses on a standard compensation structure: a mix of fixed wages and equity. Equity is highly effective at eliciting effort in startups, where firm value depends almost entirely on the inventor's specific innovation. In larger firms, however, equity value is driven by numerous stochastic factors unrelated to the inventor's effort (e.g., the success of other product lines), thereby diluting the power of equity-based incentives.

Counteracting this incentive friction is commercialization synergy, which captures the value added by a firm's production and distribution expertise. Large firms typically provide greater synergy, enhancing the economic value of a successful innovation. Together, the contract structure and synergy determine the equilibrium allocation of inventors. On one hand, smaller firms offer better incentive alignment, making them particularly attractive for inducing effort from highly productive inventors. On the other hand, larger firms can extract more value from an innovation once it is delivered. Thus, the inventor faces a fundamental trade-off between maximizing the probability of success (via high-powered incentives in small firms) and maximizing the value of success (via synergies in large firms). This trade-off dictates both the sorting of inventors into firms and their subsequent effort levels.

The model yields three primary predictions that map directly to 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 ownership shifts to the firm, the contract must substitute direct ownership with equity incentives. However, because equity is imperfect—especially in large firms where it is diluted by unrelated factors—inventors exert less effort. This theoretical result mirrors the empirical decline in the inventors' output (intensive margin), particularly the disproportionate drop observed in large firms. Third, the ownership regime fundamentally alters the allocation of talent. When R&D ownership shifts to employers, the agency costs in large firms outweigh their commercialization synergies, making them less attractive to high-ability inventors. This matches the empirical evidence on the inventor allocation (extensive margin), where large firms suffer from a relative decline in innovative personnel. Collectively, these predictions rationalize the empirical patterns and structurally link the R&D ownership regime to aggregate productivity.

To quantify the effect of policies on R&D ownership, I calibrate the model and run counterfactual analysis that reassign ownership of the residual rights to innovations. The model benchmark is calibrated to the United States, where all R&D ownership is assigned to employers. In the counterfactual, I hold all parameters fixed except the inventor ownership share. I compute the productivity of inventors by firms, the allocation of innovations between firms, and the aggregate growth rate. I then decompose the effect into extensive margin and intensive margin, and compare the relative magnitude.

Related Literature

First, this paper contributes to the endogenous growth literature, which has traditionally focused on firm-level decision-making. Prominent examples include Klette and Kortum (2004); Phillips and Zhdanov (2013); Cunningham et al. (2021); Buera and Oberfield (2020); Perla et al. (2021). More recently, an emerging strand of literature has explored the market for ideas (Silveira and Wright, 2010; Chatterjee and Rossi-Hansberg, 2012; Chiu et al., 2017) and the specific role of inventors in the innovation process (Bhaskarabhatla et al., 2021; Akcigit et al., 2025). However, the specific channel of inventor incentives, particularly those stemming from incomplete contracts, remains under-explored within this framework. This paper highlights this incentive channel both empirically and theoretically, offering a new perspective on the determinants of innovation productivity and aggregate growth.

This paper also relates to the ongoing discussion regarding trends in inventor productivity. As documented by Bloom et al. (2020), aggregate R&D productivity has been declining. While Akcigit and Goldschlag (2023) rationalize this decline by arguing that incumbent firms strategically stifle innovation to protect their market position, this paper proposes a complementary channel rooted in agency frictions. Specifically, I argue that as firms grow

larger—consistent with the rise of 'superstar firms' (Autor et al., 2020)—it becomes increasingly difficult to effectively incentivize inventors due to the dilution of equity-based compensation.

This paper focuses on the relationship between incentives and R&D productivity. Theoretical literature has long underscored the importance of residual rights in determining incentives in the presence of incomplete contracts (Grossman and Hart, 1986; Hart, 1995). Similarly, Aghion and Tirole (1994) theoretically examines the implications of ownership on innovation outcomes. Recent empirical work (Onishi, 2013; Lerner and Wulf, 2007) provides supporting evidence that assigning ownership to inventors is crucial for innovation success. For instance, Hvide and Jones (2018) focuses on a university setting, while Giarratana et al. (2018) leverages an inventor survey dataset. Complementing these studies, I exploit legal shocks to provide direct, inventor-level evidence that holding greater R&D ownership enhances an inventor's incentive to exert effort and improves R&D outcomes. This finding aligns with recent firm-level (Suh, 2023) and state-level (Alok and Subramanian, 2023) results. My paper further contributes to this literature by proposing a macroeconomic growth model featuring this micro-mechanism of employee incentives, providing a framework to evaluate the aggregate implications of ownership on R&D productivity.

Understanding the role of incentives in aggregate innovation enriches the discussion on R&D allocation and the boundary of the firm, building on the foundational work of Coase (1937), Hart and Moore (1990), and Hart and Moore (2008). Recent contributions in this area include Akcigit et al. (2016); Kogan et al. (2017); Akcigit and Kerr (2018); Ma (2022); Jo and Kim (2024). This paper distinguishes itself by identifying the inventor's choice as the primary driver of firm boundaries. I propose a novel mechanism determining the allocation of innovation across firms: a fundamental trade-off between incentive alignment (strongest in startups) and commercialization synergy (strongest in large firms). I show that this mechanism is empirically consistent with responses to a natural experiment in the legal environment.

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. However, empirical evidence suggests that the actual effects of R&D policy are often minimal (Goolsbee, 1998; Wilson, 2009). This paper provides a structural explanation for this discrepancy. I argue that the limited effectiveness of R&D subsidies stems from a misalignment between policy and inventor incentives: while policies directly subsidize employers, there is limited pass-through to the inventors who actually generate the ideas.

2 INSTITUTIONAL BACKGROUND, DATA, AND Empirical FACTS

In the United States, it is common practice for firms to require employees to sign pre-invention assignment agreements as a condition of employment. In contrast, German employee invention law explicitly prohibits such pre-invention agreements, mandating that all inventions initially belong to the inventors, albeit with certain associated costs and procedural requirements.

Within the U.S. legal framework, the ownership of innovation output—specifically patents and human capital—is governed by patent law and non-compete regulations, respectively. During the late 1970s to early 1980s, eight states enacted 'freedom-to-create' legislation designed to protect employee-inventors' ownership of innovations created outside the scope of their employment. However, in 2008, the Federal Circuit's decision in DDB Technologies LLC v. MLB Advanced Media, LP effectively overruled these state protections. Following Suh (2023), I exploit this ruling as a natural experiment. I utilize state-level variation in enforcement stringency to measure prior exposure to this legal change (Garmaise, 2011), noting that even in states with freedom-to-create laws, enforcement was often constrained by non-compete covenants. The patent output is lagged two year behind, to accommodate the process time of the patent system following (Suh, 2023). The main results are presented in the tables below.

Table 1 presents the baseline results, showing that the treatment has a significant negative impact on inventor productivity. In column 4, the coefficient is -0.19, indicating that, post-treatment, average inventor output declines by approximately 19 percent. In Table 2, I employ an alternative specification that directly utilizes the non-compete law enforcement index, as this index also influences the enforcement of freedom-to-create laws (Suh, 2023). In this specification, the treatment implies an average increase in the stringency index of 5.5; multiplying this by the estimated coefficient yields a predicted output drop of approximately 19 percent. This magnitude is consistent with the benchmark estimation, confirming the robustness of the result. Crucially, the interaction term between the treatment and the small-firm indicator is positive (0.12), implying that the negative shock is significantly mitigated for inventors in small firms.

A potential caveat in the baseline analysis is that many inventors do not patent every year, which complicates the use of inventor fixed effects. To address this, Table 3 restricts the sample to inventors who appear in the dataset annually. In this high-productivity subsample, the treatment effect is notably larger, with a coefficient of -1.00 compared to -0.19 in the full sample. This suggests that the adverse effect of losing ownership rights is particularly

	(1)	(2)	(3)	(4)
treat	0.363***	-0.176***	0.355***	-0.190***
	(0.111)	(0.051)	(0.121)	(0.051)
treat#small biz			0.051	0.120***
			(0.073)	(0.035)
small biz	-0.825***	-0.488***	-0.824***	-0.519***
	(0.057)	(0.067)	(0.068)	(0.074)
prior exposure	-0.085**	-0.001	-0.081**	-0.002
	(0.034)	(0.039)	(0.039)	(0.042)
small biz#prior exposure			-0.026	0.008
			(0.033)	(0.039)
gender	0.260***	0.218***	0.261***	0.218***
	(0.022)	(0.025)	(0.022)	(0.025)
Constant	2.946***	3.010***	2.945***	3.014***
	(0.132)	(0.056)	(0.134)	(0.058)
Observations	1,101,607	1,046,212	1,101,607	1,046,212
R-squared	0.006	0.148	0.006	0.148
Year FE	No	Yes	No	Yes
State FE	No	Yes	No	Yes
Industry FE	No	Yes	No	Yes
Firm FE	No	Yes	No	Yes

Table 1: Inventor productivity - all sample

*Treatment using free-to-create state and non-competence law enforcement as prior exposure. this is conducted on all observations. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1*

	(1)	(2)	(3)	(7)
exforcement	-0.070***	-0.034**	-0.073***	-0.036**
	(0.007)	(0.017)	(0.008)	(0.016)
small biz	-0.813***	-0.488***	-0.904***	-0.591***
	(0.058)	(0.066)	(0.076)	(0.070)
small biz#exforcement			0.013*	0.016**
			(0.008)	(0.007)
gender	0.258***	0.218***	0.258***	0.218***
	(0.022)	(0.025)	(0.022)	(0.025)
Constant	3.464***	3.193***	3.478***	3.203***
	(0.102)	(0.120)	(0.100)	(0.117)
Observations	1,101,607	1,046,212	1,101,607	1,046,212
R-squared	0.008	0.148	0.008	0.148
Year FE	No	Yes	No	Yes
State FE	No	Yes	No	Yes
Industry FE	No	Yes	No	Yes
Firm FE	No	Yes	No	Yes

Table 2: Inventor productivity - non-competence law enforcement

*Treatment using non-competence law enforcement. Change the enforcement index to the highest level after the Federal ruling. this is conducted on all observations. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1*

	(1)	(2)	(3)	(4)
treat	0.213	-0.943***	0.203	-0.998***
	(0.385)	(0.271)	(0.400)	(0.269)
treat#small biz			0.237	1.341***
			(0.659)	(0.268)
small biz	-1.905***	-0.397	-1.926***	-0.633
	(0.525)	(0.373)	(0.634)	(0.461)
prior exposure	-0.097	-0.030	-0.095	-0.021
	(0.143)	(0.150)	(0.149)	(0.154)
small biz#prior exposure			-0.073	-0.197
			(0.284)	(0.241)
Constant	9.354***	9.517***	9.355***	9.527***
	(0.364)	(0.091)	(0.367)	(0.093)
Observations	27,112	27,111	27,112	27,111
R-squared	0.001	0.568	0.001	0.568
Year FE	No	Yes	No	Yes
State FE	No	Yes	No	Yes
Inventor FE	No	Yes	No	Yes

Table 3: Inventor productivity - balanced panel

*Treatment using free-to-create state and non-competence law enforcement as prior exposure. this is conducted on a balanced panel with inventors who innovate every period. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1*

severe for highly productive inventors. Consistent with the previous results, the interaction term remains positive, reinforcing the finding that employment in a small firm mitigates the negative impact of the shock.

	(1)	(2)	(3)	(4)
treat		-0.107***		-0.093
		(0.029)		(0.099)
small biz	-0.077	-0.005	-0.124	0.073
	(0.051)	(0.041)	(0.217)	(0.110)
treat#small biz		0.073***		0.177**
		(0.018)		(0.070)
prior exposure		0.031***		0.050**
		(0.009)		(0.025)
small biz#prior exposure		-0.007		-0.045
		(0.016)		(0.045)
Enforcement	-0.023***		-0.032**	
	(0.006)		(0.015)	
small biz#Enforcement	0.013**		0.034	
	(0.006)		(0.021)	
Constant	0.522***	0.384***	0.700***	0.489***
	(0.041)	(0.017)	(0.113)	(0.043)
Observations	38,856	38,856	5,935	5,935
R-squared	0.151	0.151	0.124	0.124
Year FE	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes
Balanced Panel	No	No	Yes	Yes

Table 4: Change in personnel
*Treatment using free-to-create state and non-competence law enforcement as prior exposure. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1*

Finally, Table 4 reports the impact on the extensive margin, measured by the change in innovative personnel by firm. Columns 1 and 2 utilize the full sample of firms (conditional on having at least one observation in the shock year to ensure the personnel change is well-defined), while Columns 3 and 4 restrict the analysis to a subsample of firms that innovate every year. Columns 1 and 3 use the continuous enforcement index, while Columns 2 and 4 use the binary freedom-to-create state indicator as the treatment. Across all four specifications, the treatment effect is consistently negative, indicating a slower accumulation of innovative personnel when ownership rights shift from employees to employers. Further-

more, the interaction term is consistently positive, confirming that small firms are relatively less affected by this decline in talent acquisition.

3 A Theoretical Model (In progress)

I build a theoretical model where firms compete to attract inventors for in-house innovation. The goal of this model is to study the innovation distribution among firms and how innovation market affects firms' growth. The model addresses two problems: In the primary market, inventors make strategic decisions regarding the size of the firm they prefer to work for, while in the secondary market, firms engage in transactions to buy and sell innovations among themselves. In the following sections, I first introduce the modeling environment and then describe the equilibrium.

3.1 Environment

Time is continuous. Agents are categorized into four primary groups: households, inventors, intermediate goods producers, and final goods producers. A representative household provides labor and consumes final goods. Inventors use effort to produce innovations in intermediate firms and consume final goods. There is a continuum of intermediate firms, each producing one unique type of product and leveraging innovations to improve product quality. The final good producers assemble intermediate goods to produce final goods.

3.1.1 Preferences

The household is long-lived. It supplies one unit of labor in-elastically. The utility function is

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

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

An inventor is short-lived and lives for dt time periods. There is a continuum of inventors of measure 1 in every period. An inventor provides effort e_I to produce inventions. She is risk averse and has a mean-variance utility:

$$U_I(c_I, e_I) = \mathbb{E}(c_I) - \frac{var(c_I)}{2\bar{q}} - R(e_I)\bar{q}, \tag{2}$$

where c_I is the consumption, e_I is the effort level, and $R(e_I)\bar{q}$ is the associated cost. $\bar{q}$ (defined below) is the average quality in the economy. The cost is scaled by $\bar{q}$ to keep the problem stable over time. Denote the inventors' aggregate consumption using C_I.

3.1.2 Technology

Firms are owned by household. The final good producers produce final goods using a continuum of intermediate goods $j \in [0, N_F]$ with production technology which is similar to Akcigit and Kerr (2018)[1]:

$$Y(t) = \frac{1}{1-\beta} \int_0^{N_F} q_j^{\beta}(t) y_j^{1-\beta}(t) dj. \tag{3}$$

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:

$$Y = C_H + C_I. \tag{4}$$

There is a continuum of measure N_F risk neutral firms producing intermediate goods. Each firm produces one kind of good, using a linear technology using only labor:

$$y_j = \bar{q} l_j, \tag{5}$$

where l_j is the labor input, $\bar{q} = \frac{1}{N_F} \int_0^{N_F} q_j dj$ is the average quality. It means that innovations have positive externality, similar to Romer (1986). The cost is linear in wage w, which firms take as given. The labor market satisfies the constraint:

$$\int_0^{N_F} l_j dj \leq 1. \tag{6}$$

The production technologies, together with the market setting on innovation, ensure that a firm's value $V(q_j)$ is linear in quality q_j

$$V(q_j) = \nu q_j, \tag{7}$$

where ν is endogenous. The proof is in Section 3.2.1.

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:

$$\tilde{q_j} \equiv \frac{q_j}{\bar{q}}, \tilde{Q} \equiv \frac{Q}{\bar{q}}, \tilde{V}(\tilde{q}) \equiv \frac{V(q_j)}{\bar{q}} = \nu \tilde{q_j}, \tag{8}$$

where $Q \equiv \int_0^{N_F} q_j dj$ is the total technology stock in the economy.

A firm's quality q is affected by both business shocks and innovations. The business shock δ follows a Poisson arrival rate normalized to 1 and is randomly drawn from $\Xi(\delta)$. When hit by a shock, the firm's quality becomes $q + \delta q$.

In intermediate firms, innovations are produced by inventors using effort. Each inventor is born with one innovative idea type θ, that measures the idea's effort sensitivity and is independent of how good an innovation is$(\theta \in (0,1), \theta \sim \Psi(\theta))$. She joins a firm, either

[1]The difference is that in my specification, the final good producers only use intermediate goods, not labor, to assemble the final good. My model yields similar results if the final good producers use labor as well.

an incumbent or a start-up, and exerts unobservable effort e_I to transform the idea into an innovation. Given e_I, an innovation arrives with the instantaneous Poisson flow rate:

$$\lambda_\theta(\theta, e_I) = \mu e_I^\theta. \tag{9}$$

The innovation production function is based on the growth theory literature (Romer, 1990; Klette and Kortum, 2004; Akcigit and Kerr, 2018). The literature usually treats θ as a parameter, whereas here, θ is heterogeneous across innovations. It allows me to consider the mapping between innovations and firms. θ is the effort-elasticity in the innovation production function and I refer to inventors with high θ as effort-sensitive inventors. Additionally, in the literature, a unified firm produces and implements innovations; in my model, however, it is the inventor who creates innovations, and the firm only enjoys the outcome. It is costly to work, and the flow cost of choosing effort e_I is $R(e_I)\,\bar{q}$. Thus, without incentives, the inventor chooses $e_I = 0$. Meanwhile, she can always effortlessly make up a useless innovation that does not improve quality. The usefulness is not verifiable.

If the inventor creates a useful innovation, a patent is granted. An innovation's step size z, which is the patent owner's private information, is drawn from a distribution $\Phi(z)$. By implementing an innovation, a firm's quality increment is Δq, where

$$\Delta q = \begin{cases} \gamma_L z Q & \text{with probability } (1 - h(\tilde{q})) \\ \gamma_H z Q & \text{with probability } h(\tilde{q}), \gamma_H > \gamma_L. \end{cases} \tag{10}$$

The increment is assumed to be scaled up with Q to capture how the total technology stock in the economy affects new innovations. $h(\tilde{q})$ is the probability that a firm has complementarity with the innovation. I assume that $h'(\tilde{q}) > 0$; larger firms are more likely to have complementarity with an innovation. The complementarity includes that a firm can apply the technology to a larger production scope, the innovation suits its business operation, or the firm has enough liquidity to meet potential requirements in commercialization. γ_L and γ_H capture the different efficiency in applying the technology. When there is complementarity, the firm is the efficient user of the innovation. In this case, the quality increment is $\Delta q = \gamma_H z Q$. Otherwise, the quality improvement is lower. There are always some firms that have complementarity.[2] All firms know whether it complements with the innovation, since it is easy to observe the innovation's technology class.

Because the firm value is linear in its quality, given an innovation step size, its value to any given firm only depends on the complementarity. If the firm is an efficient user, then there is no gain from trade. Otherwise, there is positive gain from trade. All firms try to sell innovations with which they don't have complementarity, and buyers are the firms who

[2] Akcigit et al. 2016 and Ma (2022) also consider the technology mismatch, but in a different setting. In their papers, if there is a mismatch, the firm cannot use the innovation while in my model, the firm can still apply it.

complement with the innovation. For each innovation, assume that there are at least two firms have complementarity.

3.1.3 Information Structure

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

For each innovation, the step size z is private information and which firms complement with the innovation is public information. This is because though it is very difficult for an outsider to know an innovation's exact value (Chatterjee and Rossi-Hansberg, 2012), it is easy to observe the technology class. For example, if there is a new technology about compressing digital files, all firms in computing know the technology has synergy with their business, yet only the technology owner can tell the pecuniary value. Since the step size is private information and unverifiable, the existence of a useful innovation cannot be proved. Therefore, hiring contracts cannot be contingent on either. Meanwhile, when firms sell and buy innovations, they trade under asymmetric information.

3.1.4 Employment Contracting Problem

In the primary market, a firm hires an inventor by offering an employment contract. Firms compete with each other following Betrand competition. I assume that the employment contract has two dimensions: a constant wage T and a share of equity $a \in [0, 1]$.

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 payoff to be contingent on the delivery of an innovation, which is not achievable due to lack of commitment. On the one hand, a firm cannot commit to pay bonus after an innovation is created, since an inventor cannot prove an innovation is developed by her. On the other hand, an inventor cannot commit to deliver a useful innovation, since the innovation value cannot be proved to an outsider (i.e. the firm knows the actual value, but others don't [3]). Inventors have the

[3]Neither the stock reaction, nor the patent transaction can be used to evaluate innovation value. First of all, public firms only account for less than 13% 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 public firms, it is not easy to use stock market return to identify innovation value. It is difficult to implement. One reason is that lots of 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

incentive to earn bonuses effortlessly by providing useless innovations. Therefore, using a bonus cannot solve the moral hazard problem.

A firm could try to use a detailed contract, and try to directly link the inventor's innovation to the firm's profit increase and use this as a basis for a bonus. I will follow the incomplete contract literature (Grossman and Hart, 1986), who assumes that it is very hard 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 the employee contracts, as commonly seen in the real world (Brickley and Hevert, 1991).

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 $q = 0$, we 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.

3.1.5 Secondary Market Setup

Firms can trade unwanted innovations on secondary markets. Sellers are the ones who don't have complementarity with their innovations. Each innovation is unique and not substitutable by others—there is one market for each innovation. Buyers are the firms that have complementarity with the innovation. For each innovation for sale, a firm becomes a potential buyer with a probability proportional to the normalized quality $\tilde{q}$, corresponding to the fact that big firms are more likely to be a buyer (Figueroa and Serrano, 2019). For an innovation, there are at least two potential buyers.

Firms sell and buy under incomplete information because the step size is private information. Firms compete to buy innovations following Bertrand competition. Each buyer offers a price p_z simultaneously, and the seller chooses whether to accept one. It is a one-shot game—neither sellers nor buyers keep track of past trades. The settings lead to a standard lemons market (Akerlof, 1978).

3.1.6 Entry and Exit

There are two types of entry. One is innovation-related, which happens when an inventor chooses to work in $\tilde{q} = 0$ firm, successfully creates an innovation and decides to keep it. Denote the amount of innovation-related entrants with λ_I. The other is the exogenous entry where firms enter due to reasons other than innovations. Denote the amount of exogenous entrants with λ_0. In both cases, upon entry, the firm draws a quality $\tilde{q}$ from a distribution

available.

$\tilde{F}_{q0}(\tilde{q})$ and incurs a cost which equals the firm value. This represents the spillover from incumbents to entrants. Assume that $\tilde{f}_{q0}(\tilde{q}/\kappa) = \tilde{f}_q(\tilde{q})/\kappa$. When $\kappa < 1$, entrants are on average smaller. When the new entrant owns an innovation, then its quality q is boosted by Δq because of the innovation.

Firms face an exogenous exit rate τ. I will focus on a balanced growth path such that entry equals exit

$$\tau N_f = \lambda_I + \lambda_0. \tag{11}$$

3.2 Equilibrium

I now characterize the equilibria of the economy in which aggregate variables $(Y, C, R, w, \bar{q})$ grow at the constant rate g, while other variables (such as N_F) are constant.

3.2.1 Production

The final good producer chooses $\{y_j\}_j$ to maximize its profit using the technology described in Section 3.1.2:

$$\max_{\{y_j\}} \frac{1}{1-\beta} \int_0^{N_F} q_j^{\beta} y_j^{1-\beta} dj - \int_0^{N_F} y_j p_j dj. \tag{12}$$

The first-order condition yields the demand function of intermediate firms $p_j = q_j^{\beta} y_j^{-\beta}$. The intermediate goods are produced by corresponding firm $j \in [0, N_F]$ using only labor $y_j = \bar{q} l_j$, where $\bar{q} = \frac{1}{N_F} \int_0^{N_F} q_j dj$ is the average quality, and l_j is the labor input. Intermediate good producers are in monopolistic competition and choose l_j, p_j, y_j to maximize their profit, given the wage level w:

$$\begin{aligned} &\max_{l_j, p_j, y_j} y_j p_j - w l_j. \\ &s.t.\ y_j = \bar{q} l_j \\ &\quad p_j = q_j^{\beta} y_j^{-\beta} \end{aligned} \tag{13}$$

Therefore, the FOC yields

$$y_j = q_j \left(\frac{\bar{q}(1-\beta)}{w} \right)^{\frac{1}{\beta}}, l_j = y_j/\bar{q}, p_j = \frac{w}{\bar{q}(1-\beta)}. \tag{14}$$

In each period, the labor market clearing satisfies $\int_0^{N_F} l_j dj = 1$, which gives that $\frac{\int_0^{N_F} q_j \left(\frac{\bar{q}(1-\beta)}{w}\right)^{\frac{1}{\beta}} dj}{\bar{q}} = 1$. The wage w can be solved

$$w = N_F^{\beta} (1-\beta) \bar{q}. \tag{15}$$

Plug it back into the intermediate firm's problem. Both production y_j and profit π_j are linear in quality

$$y_j = \frac{q_j}{N_F}, \pi_j = \frac{\beta q_j}{N_F^{1-\beta}}. \tag{16}$$

I drop the subscript j from the firm-level variable when it does not cause confusion. In this model, the firm size is linear in its quality and I will use q to denote firm size when it does not cause confusion. Intermediate firms are the ones who hire inventors to innovate. Because of competition, any value from innovations is captured by inventors, and any value from acquisitions is captured by the seller. The discounted value of being a firm of quality q is the same as the net present value under the assumption that the firm never innovates. Thus, the value function of intermediate firm q at time t can be written as

$$V(q,t) = \int_t^\infty e^{-(r+\tau)(s-t)} \beta q / N_F^{1-\beta} ds = \nu q. \tag{17}$$

where $\nu = \frac{\beta}{(r+\tau)N_F^{1-\beta}}$. The value function is linear in q and does not depend on time. This result implies that for any firm, the value of a same quality improvement Δq is the same.

The aggregate production is linear in average quality $\bar{q}$. The resource constraint of the economy is $Y = C_H + C_I$. The Euler equation is:

$$g = \frac{\dot{Y}}{Y} = \frac{\dot{C}_H}{C_H} = \frac{\dot{\bar{q}}}{\bar{q}} = r - \rho. \tag{18}$$

3.2.2 Secondary market

The secondary innovation market is where firms buy and sell innovations. For each innovation, there are at least two buyers. Buyers follow Betrand competition and offer prices to the seller. A seller accepts an offer only when it is profitable, which means $\gamma_L z \nu \tilde{Q} \leq \tilde{p}_z$. The left-hand side is the value of the innovation if the seller keeps it, whereas the right-hand side is the value if the seller sells it. The threshold $\hat{z}$ satisfies

$$\gamma_L \hat{z} \nu \tilde{Q} = \tilde{p}_z. \tag{19}$$

Buyers set the price $\tilde{p_z}$ according to the zero-profit condition

$$\int_0^{\hat{z}} \left(\gamma_H \nu z \tilde{Q} - \tilde{p}_z \right) \phi(z) dz = 0, \tag{20}$$

where $\gamma_H \nu z \tilde{Q}$ is the innovation value to any buyer with complementarity; $\tilde{p}_z$ is the price of the innovation; $\hat{z}$ is the step size threshold. If $z \leq \hat{z}$, an innovation is sold.

If a firm is not the efficient user of an innovation, it chooses to sell the innovation if its step size is lower than the threshold $\hat{z}$. Otherwise, the firm uses the innovation to improve its own technology.

For each innovation, the probability of being sold is $(1 - h(\tilde{q}))\,\Phi(\hat{z})$, where $(1 - h(\tilde{q}))$ is the probability that the innovation does not have complementarity with the firm $\tilde{q}$ where the inventor works; $\Phi(\hat{z})$ is the cumulative distribution function of z at $\hat{z}$. The total amount of innovations sold per unit of time is

$$\Phi(\hat{z}) \int_0^1 \lambda_\theta^* \left(1 - h\left(\tilde{q}^*(\theta)\right)\right) \psi(\theta)\, d\theta, \tag{21}$$

where λ_θ^* is the equilibrium innovation arrival rate of type θ. $\tilde{q}^*$ denotes the firm for which

inventor θ works in equilibrium. Recall that the probability of buying an innovation is $\lambda_b \tilde{q}$, because the probability of buying an innovation is assumed to be proportional to $\tilde{q}$ and the constant λ_b is what I solve for. The secondary market clearing gives

$$N_F \int \lambda_b \tilde{q} f_q(\tilde{q})\, d\tilde{q} = \Phi(\hat{z}) \int_0^1 \lambda_\theta^* \left(1 - h\left(\tilde{q}^*(\theta)\right)\right) \psi(\theta)\, d\theta, \tag{22}$$

which implies that

$$N_F \lambda_b = \Phi(\hat{z}) \int_0^1 \lambda_\theta^* \left(1 - h\left(\tilde{q}^*(\theta)\right)\right) \psi(\theta)\, d\theta. \tag{23}$$

It pins down the endogenous variable λ_b, and the probability of buying an innovation.

The innovation value for the original firm $\tilde{q}$ satisfies

$$x(z, \tilde{q}) = \begin{cases} \gamma_L \nu \hat{z} \tilde{Q} & \text{no complementarity and } z \leq \hat{z} \\ \gamma_L \nu z \tilde{Q} & \text{no complementarity and } z > \hat{z} \\ \gamma_H \nu z \tilde{Q} & \text{with complementarity.} \end{cases} \tag{24}$$

When there is no complementarity, and the step size z is lower than the threshold $\hat{z}$, the firm sells the innovation on the innovation market at the price $\gamma_L \nu \hat{z} Q$. If the step size z is higher than the threshold, the firm keeps it, and it is worth $\gamma_L \nu z Q$. When there is complementarity, the firm is the efficient user of the innovation. Then the firm also implements it in-house, and its value is $\gamma_H \nu z Q$. In a first-best allocation, all innovations would be sold to the efficient user. The asymmetric information on the secondary market, however, prevents this from happening.

The expected innovation value $\mathbb{E}(x(\tilde{q}))$ can be written as

$$\begin{aligned} \mathbb{E}(x(\tilde{q})) = {} & \gamma_L \left[\Phi(\hat{z})\, \hat{z} + (1 - \Phi(\hat{z}))\, \mathbb{E}(z | z > \hat{z})\right] (1 - h(\tilde{q}))\, \nu Q \\ & + \gamma_H \mathbb{E}(z)\, h(\tilde{q}) \nu Q. \end{aligned} \tag{25}$$

It increases with the firm size $\tilde{q}$ because $h\tilde{q}$ is increasing in $\tilde{q}$. Therefore, conditional on an innovation being created, its value is higher if it is in a large firm.

3.2.3 Primary market

The primary market is where firms compete to hire inventors using a combination of equity a and wage $\tilde{T}$. The setup yields a principal-agent framework. Firms are risk neutral, and inventors are risk averse. Firms enjoy the innovations produced by inventors, but it is costly for inventors to work and firms cannot monitor the effort. Thus, firms want to split the surplus with the inventor by offering a constant wage; meanwhile, firms need to incentivize the inventor to exert effort by offering equity. Because of the competition, the firm's problem is the same as maximizing the utility it offers to an inventor, subject to the zero profit

condition:

$$\begin{aligned}
&\max_{a,T,e_I,c_I} U_I\left(a,\tilde{T};\tilde{q},\theta,e_I\right)=\mathbb{E}\left(\tilde{c_I}\right)-\frac{1}{2}\text{var}\left(\tilde{c_I}\right)-R\left(e_I\right)dt.\\
&\quad s.t.\lambda_\theta\left(\theta,e_I\right)\mathbb{E}\left(x\left(\tilde{q}\right)\right)dt-\mathbb{E}\left(\tilde{c_I}\right)\geq 0\\
&\qquad e_I=\arg\max\left\{U_I\left(a,\tilde{T};\tilde{q},\theta,e_I\right)\right\}\\
&\qquad \mathbb{E}\left(\tilde{c_I}\right)=a\left(\mathbb{E}\left(\tilde{V}_0\left(\tilde{q}\right)\right)+\lambda_\theta\left(\theta,e_I\right)\mathbb{E}\left(x\left(\tilde{q}\right)\right)dt\right)+\tilde{T}\\
&\qquad \text{var}\left(\tilde{c_I}\right)=a^2\left(\lambda_\theta\left(\theta,e_I\right)\sigma_x^2(\tilde{q})+\sigma_0^2(\tilde{q})\right)dt
\end{aligned}\tag{26}$$

The optimal contract is $\left\{a^*\left(\theta,\tilde{q}\right),T^*\left(\theta,\tilde{q}\right)\right\}$ [4]. The first line is the objective function $U_I\left(a,\tilde{T};\tilde{q},\theta,e_I\right)$, representing the inventor θ's utility when working for firm $\tilde{q}$ with effort level e_I; $\left(a,\tilde{T}\right)$ characterize the contract offered by the firm, and $\tilde{c_I}$ is the inventor's consumption level, while $R\left(e_I\right)dt$ is the cost of effort. The second line is the firm's individual rationality constraint, where $x\left(\tilde{q}\right)$ is the value of an innovation to a firm $\tilde{q}$. It multiplies by the probability of creating an innovation $\lambda_\theta\left(\theta,e_I\right)dt$, which gives the expected payoff of hiring an inventor. The cost of hiring is $(\tilde{c_I})$. The firm will only participate if the expected gain is nonnegative. The third line is the inventor's incentive compatibility constraint. The inventor chooses an effort level e_I to maximize the expected utility $U_I\left(a,\tilde{T};\tilde{q},\theta,e_I\right)$ given firm quality $\tilde{q}$ and the wage scheme $\left\{a,\tilde{T}\right\}$. The fourth line describes the inventor's expected consumption. The first part, $a\left(\tilde{V}_0\left(\tilde{q}\right)+\lambda_\theta\left(\theta,e_I\right)\mathbb{E}\left(x\left(\tilde{q}\right)\right)dt\right)$, is expected value of the firm equity a, which is the sum of the value without innovations and the expected value of an innovation. The other component, $\tilde{T}$, is the constant transfer from the firm to the inventor. The fifth line shows the variance of the inventor's consumption. The exposure is the equity a. The uncertainty comes from two sources: innovation-related part $\lambda_\theta\left(\theta,e_I\right)\sigma_x^2(\tilde{q})$ and the innovation-independent part $\sigma_0^2(\tilde{q})$. In start-ups, the second term is zero. I explain both in detail later.

Given contracts, the inventor with idea θ chooses firm $\tilde{q}^*\left(\theta\right)\in[0,\infty)$ and exerts effort e_I to maximize her utility by solving the following problem:

$$\max_{\tilde{q}} U_I\left(a\left(\tilde{q},\theta\right),\tilde{T}\left(\tilde{q},\theta\right),e_I\left(\tilde{q},\theta\right),\tilde{q},\theta\right).\tag{27}$$

Here $a\left(\tilde{q},\theta\right),\tilde{T}\left(\tilde{q},\theta\right)$ and $e_I\left(\tilde{q},\theta\right)$ are the solutions to the firm's problem in Equation (26). If $\tilde{q}^*\left(\theta\right)>0$, the inventor works in an incumbent firm; $\tilde{q}^*\left(\theta\right)=0$ means the inventor joins a new firm. The equilibrium effort is $e_I^*\left(\theta\right)$ and the corresponding arrival rate is $\lambda_\theta^*=\lambda_\theta\left(\theta,e_I^*\left(\theta\right)\right)$.

The firm value is affected by four factors: in-house innovation, purchased innovation, business shocks, and exogenous death. The firm value without in-house innovations $\tilde{V}_0\left(\tilde{q}\right)$ is

[4] The numerical results are in Figure **??**.

$$\tilde{V}_0(\tilde{q}) = \nu\tilde{q} + \begin{cases} -\nu\tilde{q} & \tau dt \\ \delta\nu\tilde{q} & dt \\ \left(\gamma_H \nu z \tilde{Q} - \tilde{p}_z\right)|_{z\leq\hat{z}} & \lambda_b \tilde{q}\Phi(\hat{z})\, dt \\ 0 & 1 - (\tau + 1 + \lambda_b \tilde{q}\Phi(\hat{z}))\, dt. \end{cases} \tag{28}$$

The first line is when the firm dies exogenous. The second line is when hit by a business shock δ that arrives with an arrival rate of 1. The third line is when the firm purchases an innovation. Though the expected gain from purchasing an innovation is zero, the realized gain usually is not zero—it is the gap between the price paid and the firm value improvement. The quality improves by $\gamma_H z\tilde{Q}$, where z is a random draw from the $\Phi(z)$ distribution truncated at $\hat{z}$. The cost of buying an innovation is $\nu\gamma_L\hat{z}\tilde{Q}$. The last line is the firm value otherwise.

The uncertainty in the firm equity value leads to the variance in the inventor's consumption c_I. Because the in-house innovation is independent of other activities, the firm-level risk can be decomposed into two parts: in-house innovation-related risk and the rest. The former one is $\lambda_\theta(\theta, e_I)\,\sigma_x^2(\tilde{q})dt$, where $\sigma_x^2(\tilde{q})$ is the second moment of the innovation value. The latter one is $\sigma_0^2(\tilde{q})dt$ which contains all shocks unrelated to innovations, including the risk from death, business shocks, and unknown step size in innovation trade. It can be written as

$$\sigma_0^2(\tilde{q}) = \nu^2\left\{\tilde{q}^2\left(\tau + \sigma_\delta^2\right) + \lambda_b\tilde{q}\Phi(\hat{z})\,\mathbb{E}\left[(\gamma_H z - \gamma_L\hat{z})^2\,|z\leq\hat{z}\right]\right\}. \tag{29}$$

This risk increases with the firm size $\tilde{q}$. Therefore, in a larger firm, the innovation-related risk takes a smaller share in the total equity variance.

The firm determines how much equity to offer to the inventor, which depends on the firm size. In start-ups, the equity value depends solely on the inventor's innovation. Thus, start-ups can offer compensation in the form of equity and wages merely restricted by the inventor's risk aversion—the inventor can be well-incentivized to put effort into innovating. In larger firms, by contrast, the inventor's innovation only makes up a small fraction of the equity value; much of the equity value depends on stochastic factors unrelated to the innovator's effort (e.g. success of other product lines). If larger firms were to offer the same incentive as start-ups for the same expected compensation, inventors would be exposed to many unrelated risks and face a highly levered compensation package; this is not attractive to risk averse inventors. Hence, both the optimal equity a and the resulting incentive decrease with firm size. This is consistent with the empirical findings in Akcigit and Goldschlag (2023).

After seeing all contracts, the inventor picks the contract offered by firm $\tilde{q}^*(\theta)$, which

gives her the highest utility. Equation (27) can be written as:

$$\begin{aligned}\max_{\tilde{q},e_I} U_I\left(\tilde{q};\theta,e_I\right) &= \lambda_\theta\left(\theta,e_I\right)\mathbb{E}\left(x\left(\tilde{q}\right)\right)dt - \frac{1}{2}a^2(\tilde{q})\left[\lambda_\theta\left(\theta,e_I\right)\sigma_x^2(\tilde{q})+\sigma_0^2(\tilde{q})\right]dt \\ &\quad - R\left(e_I\right)dt. \\ s.t. e_I &= \arg\max\left\{U_I\left(\tilde{q};\theta,e_I\right)\right\}.\end{aligned} \tag{30}$$

The main trade-off an inventor facing is a higher chance to create an innovation versus a better value after an innovation is created, which is represented by the first term of Equation (30): $\lambda_\theta\left(\theta,e_I\right)\mathbb{E}\left(x\left(\tilde{q}\right)\right)$. The two elements work in opposite directions. A firm with higher quality $\tilde{q}$ on average can implement an innovation more efficiently, which means that the value-added of an innovation $\mathbb{E}\left(x\left(\tilde{q}\right)\right)$ is higher; meanwhile, a smaller firm can offer better incentive and hence higher arrival rate $\lambda_\theta\left(\theta,e_I\right)$. The inventor's decision then depends on the relative strength of the two forces, which, in turn, is determined by the inventor type. For an effort-sensitive inventor with high θ, the incentive channel is more important. As a result, she chooses to work for a small firm. Similarly, when an inventor's idea has low effort-elasticity, the value channel is more crucial and she works for a large firm.

Inventor intensity is endogenous. For each inventor, the equilibrium innovation arrival rate generated by herself in the firm $\tilde{q}$ is λ_θ^*. In firm $\tilde{q}$, the innovation arrival rate is affected by both *each innovation's arrival rate* and *the number of inventors per firm*. Denote the arrival rate in firm $\tilde{q}$ by $\lambda_q\left(\tilde{q}\right)$ and the firm probability distribution function by $f_q(\tilde{q})$. The total amount of innovations in a firm should be the same as the total innovations made by inventors who work for the firm. The inventor market clearing gives

$$N_F\lambda_q\left(\tilde{q}\left(\theta\right)\right)f_q\left(\tilde{q}\left(\theta\right)\right)d\tilde{q} = \lambda_\theta^*\psi(\theta)d\theta, \tag{31}$$

which pins down the innovation arrival rate per firm.

3.2.4 Entry and Exit

Recall there are two types of entries in this model. One is innovation-related. It happens when an inventor chooses to work in a firm with $q=0$, successfully gets an innovation, and decides to keep it. The number of firms that enter this way can be written as

$$\lambda_I = \int_{\tilde{q}^*(\theta)=0} \lambda_\theta^* \left[\overbrace{\underbrace{\left(1-\Phi\left(\hat{z}\right)\right)\left(1-h\left(0\right)\right)}_{\text{Non-complementary \& big step size}} + \underbrace{h\left(0\right)}_{\text{Complementary}}}^{\text{Innovations kept}}\right]\psi\left(\theta\right)d\theta. \tag{32}$$

The other type is the exogenous entry, which enters the market without an innovation. The number is denoted by a parameter λ_0.

On the exit front, the exit rate satisfies that the number of firms who exit equals the number that enters.

$$\tau N_F = \lambda_I + \lambda_0. \tag{33}$$

I use (33) to solve for the measure of firms N_F.

3.2.5 Growth rate

The growth rate of the aggregate quality $\bar{q}$ can be written as:

$$g = \frac{\mathbb{E}_t\left(\bar{q}(t+dt)\right) - \bar{q}(t)}{\bar{q}(t)dt}. \tag{34}$$

The growth rate is

$$\begin{aligned} g = &-\tau\left(1-\kappa\right) + \gamma_H \mathbb{E}\left(z\right) \int h\left(q^*\left(\theta\right)\right) \lambda_\theta^* \psi\left(\theta\right) d\theta \\ &+ \left[\gamma_L \mathbb{E}\left(z\right) + \left(\gamma_H - \gamma_L\right) \mathbb{E}\left(z | z \leq \hat{z}\right) \Phi\left(\hat{z}\right)\right] \int \left[1 - h\left(q^*\left(\theta\right)\right)\right] \lambda_\theta^* \psi\left(\theta\right) d\theta. \end{aligned} \tag{35}$$

It includes three parts. $-\tau\left(1-\kappa\right)$ is quality destruction if there were no innovations. For each unit of time, τ fraction of firms with an average quality $\bar{q}$ leave the economy, while the same fraction of firms enter the economy with an average quality κ. The rest reports the quality improvement due to the innovation. The rest of the first line is the quality improvement because of innovations that complement with the firms. The average quality is $\gamma_H \mathbb{E}\left(z\right)$. The second line measures the improvement that comes from the innovations that do not complement the initial firm. If there were no secondary market, then the average quality would be $\gamma_L \mathbb{E}\left(z\right)$. With the secondary market, some low-step-size innovations are sold to more efficient users, which increases the quality improvement by $\left(\gamma_H - \gamma_L\right) \mathbb{E}\left(z | z \leq \hat{z}\right) F\left(\hat{z}\right)$. Hence when the innovation is worth more or the arrival rate increases, the growth rate goes up.

3.2.6 Equilibrium

I end this section by summarizing the equilibrium. The in-house R&D expenditure C_I of the economy can be written as

$$C_I = \int \lambda_\theta\left(\theta\right) \mathbb{E}\left(x\left(q^*\left(\theta\right)\right)\right) \psi\left(\theta\right) d\theta. \tag{36}$$

It captures all transfers made to inventors. The total firm-level R&D expenditure is the sum of R and the spending on purchasing patents on the secondary market

$$X = C_I + \lambda_b p_z. \tag{37}$$

Based on Equation (12), the equilibrium out put level Y is linear in $\bar{q}$

$$Y = \frac{1}{1-\beta} \frac{\bar{q}}{N_F^{1-\beta}}. \tag{38}$$

and the consumption level is

$$C_H = Y - C_I. \tag{39}$$

Definition 1. A balanced growth path of this economy for any combination of t, q, is the mapping between q and θ, the allocation $\left(\left\{y_j^*\right\}_j, Y^*, X^*, C_I^*, C_H^*\right)$, the prices $\left(w^*, \left\{p_j^*\right\}_j, \hat{z}\right)$, the growth rate g^*, the entry rate λ_{eI}^*, and the measure of firms N_F^*, such that (1) for any $j \in [0,1]$, y_j^* and p_j^* satisfy Equation (14); (2) wage w^* satisfies Equation (15); (3) measure

of the intermediate producers N_F^* satisfies Equation (33); (4) the mapping is the solution of Equation (27); (5) aggregate creative destruction τ_e satisfies Equation (33); (6) the price offer $\hat{z}$ for inventions satisfies Equation (19); (7) the entry rates λ_{eI}^* satisfy Equation (32); (8) in-house R&D spending C_I^* satisfies Equation (36); (9) total R&D spending X satisfies Equation (37); (10) aggregate output Y^* satisfies Equation (38); (11) aggregate consumption C_H^* satisfies Equation (39); and (12) steady-state growth rate g^* satisfies Equation (35).

4 Calibration

4.1 Data

4.2 Moment Estimation

5 Results

5.1 Optimal Contracts

5.2 Untargeted Moments

5.3 Growth Decomposition

5.4 Counterfactual

6 Discussion and Extension

7 Conclusion

References

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

Acemoglu, D. and U. Akcigit (2012): "Intellectual property rights policy, competition and innovation," *Journal of the European Economic Association*, 10, 1–42.

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.

Akerlof, G. A. (1978): "The Market for "Lemons": Quality Uncertainty and the Market Mechanism," in *Uncertainty in Economics*, Elsevier, 235–251.

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.

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

Figueroa, N. and C. J. Serrano (2019): "Patent Trading Flows of Small and Large Firms," *Research Policy*, 48, 1601–1616.

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.

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.

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.

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.

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. (1986): "Increasing Returns and Long-Run Growth," *Journal of Political Economy*, 94, 1002–1037.

——— (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.

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.