Ideas and Firm Dynamics
When It Takes Two to Tango

Seula Kim* Jane Olmstead-Rumsey† Honghao Wang‡

August 3, 2026

Abstract

Teams of inventors now produce over 80% of U.S. patents, up from just 45% in 1976. Over the same period, inventor employment has become increasingly concentrated in large firms, rising six times more than overall employment concentration. To explain these trends, we develop a new model of endogenous team formation within firms and embed it in a general equilibrium model of the inventor labor market. Teams increase firms' returns to scale by increasing the likelihood of productive matches between inventors. The value of these matches rise over time, possibly driven by increasing knowledge specialization. This causes large firms to expand and small firms to shrink in equilibrium, accounting for the observed increase in inventor employment concentration. Because firms do not internalize knowledge spillovers, there is misallocation of inventors across firms. The rise of team production amplifies this misallocation by a factor of six by reallocating inventors away from small firms that generate disproportionately large knowledge spillovers.

Keywords: innovation, R&D teams, inventor allocation, firm dynamics, economic growth

*Pennsylvania State University; Email: seulakim@psu.edu

†LSE and CEPR; Email: j.olmstead-rumsey@lse.ac.uk

‡Northwestern University; Email: honghao.wang@kellogg.northwestern.edu. We thank Harun Alp, Gadi Barlevy, Antonin Bergaud, Francisco Brahm, Santiago Caicedo, Maarten De Ridder, Nik Engbom, Lukas Freund, Seb Graves, John Haltiwanger, Xavier Jaravel, Kala Krishna, Yueyuan Ma, Ben Moll, Chris Moser, Ezra Oberfield, Jeremy Pearce, Marta Prato, Ben Pugsley, Chris Tonetti, and Steve Yeaple for helpful comments and discussions. We thank Dongjae Lee, Francisco Libano-Monteiro and Shiori Tanaka for outstanding research assistance.

1 Introduction

The development of new ideas is a key driver of economic growth. Ideas are produced either by individual inventors or by collaborations among teams of inventors. Teams have become increasingly important in the production of ideas in recent decades, both in patenting and in scientific research (Wuchty, Jones, and Uzzi 2007). Existing models of inventor team formation study important dimensions of collaboration—such as specialization, sorting, human capital accumulation, and declining communication costs—but abstract from firm boundaries.

The central premise of this paper is that firm boundaries impose meaningful constraints on inventor collaboration. Changes to the idea production function that favor teams such as increasing specialization (Jones 2009) or falling communication costs (Pearce 2023) may therefore reshape the allocation of inventors to firms in order to exploit collaborations among inventors. When teams are the primary unit of knowledge production and team formation is constrained by the boundary of the firm, modeling the two jointly is essential for studying trends in the allocation of inventors to firms and its implications for the rate of innovative output. Inventors are increasingly concentrated at large firms, raising concerns about the implications for growth (Akcigit and Goldschlag 2023; Fernandez-Villaverde, Yu, and Zanetti 2026).

Our main contribution is to propose a novel and tractable framework for modeling team-based invention within firms in which additional inventor colleagues raise expected inventor productivity only in firms that choose to adopt a team production model. We then measure the impact of team production on innovative output (patents) using a panel of the universe of patenting firms in the U.S. Consistent with the theory, we show that team-based firms have less-decreasing returns to scale in inventor labor and the impact of choosing a team production model on returns to scale has increased over time, likely reflecting greater specialization of knowledge (Jones 2009; Freund 2025). In our model, inventors become more concentrated at large firms in response to

this shock. Finally, we explore whether this reallocation of inventors is efficient in the presence of knowledge spillovers. A planner would choose to increase concentration in response to the measured changes in the idea production function favoring teams, but not nearly as much as happens in the competitive equilibrium. The changes in the idea production function increase misallocation (the gap between innovative output in the socially optimal allocation and the competitive equilibrium) by a factor of six.

Our analysis precedes in several steps. We begin by using patent data to document several stylized facts about inventor teams. First, patents developed by teams now account for 80% of all private patents in the U.S., compared with 45% in 1980. Second, we use new data on inventor employment histories from Revelio Labs to document that nearly 90% of team collaborations in patenting occur among inventors who work at the same firm. This underscores the importance of firm boundaries in shaping collaborative idea production. Third, inventor employment concentration in large firms has risen dramatically over time, and this rise is more pronounced than the increase in overall employment concentration (Autor et al. 2020b) over the same period by a factor of six.

Motivated by these facts, we build a general equilibrium model of the inventor labor market with heterogeneous firms in the spirit of Hopenhayn (1992) and Melitz (2003), augmented with two new ingredients. First, we explicitly model the process of inventor team formation within firms. Each inventor draws a bilateral productivity for collaborating with every other inventor in the firm and can alternatively work alone with some fixed productivity. As a result, expected inventor productivity is an increasing but concave function of the number of inventor coworkers in the firm, as inventors are increasingly likely to find a productive match with each additional coworker.[1] Second, we incorporate knowledge spillovers across firms that depend on both firm productivity and employment size, following Chikis, Kleinman, and Prato (2026). The presence of spillovers implies that the equilibrium firm size distribution need not be efficient.

Next, we use a combination of direct and indirect inference to estimate how key model parameters, particularly those governing the idea production function, have

[1] We support these modeling assumptions with empirical evidence from the U.S. patent data.

evolved over time, calibrating two steady states of the model (1976-1999 and 2000-2018). We find that the productivity of both solo inventors and inventor teams has risen over time, particularly the returns to scale benefit of team production. Our model predicts that these changes in the idea production function generate a substantial increase in employment concentration and average firm size, consistent with the data. At the same time, the model predicts higher innovative output due to the increase in returns to scale, but an endogenous decline in knowledge spillovers.

Finally, we assess the welfare implications of the changes in idea production associated with the rise of patenting teams. Inventor misallocation increases by a factor of six. The underlying mechanism is that the measured shifts in the idea production function favor large, highly productive firms, and their expansion increases the inventor wage, leading to excessive firm exit and a contraction of smaller firms who generate disproportionately large knowledge spillovers compared to their innovative output (Chikis, Kleinman, and Prato 2026). These results have implications for the design of R&D policy. In particular, the welfare gains from subsidizing small firms and innovative firm entry are larger in the current environment than in the past.

Related Literature. A growing body of literature examines the match between individual inventors and firms (Manera 2022; Akcigit and Goldschlag 2023; Babalievsky 2023; Liu 2023; Guccione and Roldan-Blanco 2026). A separate strand of literature models firms as collections of coworkers or co-inventors, without incorporating additional dimensions of firm heterogeneity (Herkenhoff et al. 2024; Pearce 2023; Freund 2025). This paper bridges the gap between these two approaches by exploring how inventor team formation varies with firm-level productivity and relates to inventor (mis)allocation across firms.[2] This connection is important for understanding how changes in the idea production function might have played a role in the recent rise in employment concentration of inventors.

[2] Jarosch, Oberfield, and Rossi-Hansberg (2021) and Boerma, Tsyvinski, and Zimin (2025) study models of sorting into teams within firms, but not specifically for inventors. Acabbi et al. (2026) study sorting into entrepreneurial teams.

Our model of teamwork within the firm contributes a new mechanism to theories of the boundary of the firm and optimal scale (Lucas 1978). Recent work finds evidence of differences in returns to scale at the firm level within industries (Chan et al. 2025) and evidence that the supply of managers (Engbom et al. 2025), standardization (Argente et al. 2025), and economic development (Bassi et al. 2023; Chen 2025) all influence optimal firm scale. We validate the existence of our new mechanism using firm-level data on inputs (inventors) and output (patents). The teams mechanism may extend beyond patenting to other settings in which workers increasingly specialize along the development path. However, patent data provide a uniquely rich environment to observe team collaborations within firms, which is more challenging in other industries.

This paper is also related to the literature on how interactions between individuals shape economic growth (Jones 2009; Lucas and Moll 2014; Akcigit et al. 2018; Prato 2025). Relative to these papers, we abstract from the dynamics of individual-level knowledge accumulation and instead focus on the organizational features that determine how knowledge is produced within firms (Antràs, Garicano, and Rossi-Hansberg 2006; Garicano and Rossi-Hansberg 2015). Our model delivers the novel prediction that R&D workers increasingly concentrate in large firms, which highlights the importance of understanding how human capital accumulation for R&D workers operates. If exposure to a diverse set of employers over an inventor's career is a key source of knowledge accumulation, rising concentration in larger firms may pose challenges. On the other hand, the shift from solo invention toward collaboration may accelerate on-the-job human capital accumulation through peer learning within firms (Akcigit et al. 2018).[3] We leave the study of these dynamics for future work.

Finally, our findings contribute to understanding trends in research productivity of large and small firms (Andrews, Criscuolo, and Gal 2016; Cavenaile, Celik, and Tian 2025; Olmstead-Rumsey 2026), and the recent trends of declining business dynamism and missing high-growth young firms (Decker et al. 2014; Decker et al. 2016; Sterk,

[3] Several empirical studies highlight the role of teams for inventor career dynamics (Jaravel, Petkova, and Bell 2018; Bhaskarabhatla et al. 2021).

Sedláček, and Pugsley 2021). Our model suggests that ideas getting harder to find (Bloom et al. 2020) may be driven in part by the changing allocation of inventors to firms driven by team production. This can endogenously reduce knowledge spillovers across firms, providing a microfoundation for Akcigit and Ates (2023)'s observation that knowledge diffusion from large firms to small firms seems to have fallen over time.[4]

2 Empirical Motivation

2.1 Data Description

Our primary data sources are the United States Patent and Trademark Office (USPTO) PatentsView, Revelio Labs LinkedIn data, and S&P's Compustat. The USPTO PatentsView database tracks the universe of patents granted by the U.S. Patent and Trademark Office (USPTO) from 1976 onward. It provides detailed patent-level information, including application and grant dates, technology class classifications, inventor identifiers and addresses, citations, and the names and addresses of patent assignees (firms).

The Revelio Labs LinkedIn data offer worker-level employment histories based on LinkedIn profiles, which is linked to the USPTO inventor data. This allows us to match inventors to firms in years that they do not patent and hence do not appear in the USPTO data.

S&P Compustat provides detailed firm-level data for publicly listed firms, maintained by S&P Global. It includes information such as firm size, industry scope, balance sheets, income and cash flow statements, and stock prices.

We collect utility patents filed between 1976 and 2018, granted by the USPTO. To identify the patents assigned to U.S. public firms and track them by different firm characteristics, we link the USPTO assignee identifier to the Compustat GVKEYs (the firm identifiers in Compustat) using the bridge developed by Ding, Jo, and Kim (2022)

[4] Berkes and Gaetani (2026) provide a complementary explanation for declining knowledge diffusion based on increasing knowledge specialization that abstracts from the role of firms.

and Braguinsky et al. (2023).[5] When a patent is assigned to multiple entities (around 3% of patents in our sample), we keep the first assignee following the sequence order recorded by USPTO. To measure the quality of patents, we calculate five-year forward citations with vintage fixed effects removed, and obtain patent value measures from Kogan et al. (2017) (KPSS values, hereafter) for patents filed by publicly held firms.

2.2 Empirical Findings

Next, we summarize the core set of empirical findings that emerge from our dataset.

Fact 1: Team patents have increased from 45% (1976) to 80% (2018) of patents

First, team-based patenting has become increasingly prevalent over the past five decades. We define a *team patent* as a patent with multiple inventors, and refer to the set of inventors listed on the patent as a *team*.

In our data, the share of team patents has risen from 45% in 1976 to above 80% in 2018. This upward trend holds either measured in patent counts, market-based patent value (KPSS values), or five-year forward citations. These patterns are illustrated in Figure 1.[6]

Fact 2: 88% of team patents are by inventors at the same firm

Second, the vast majority of team collaborations in innovation occur within firm boundaries rather than across firms. This highlights the role of firms as the critical platform for team production.

In the Revelio Labs data, we compute the probability that a given pair of inventors within a patenting team is ever employed by the same firm within a ten-year window

[5] This bridge is constructed using the standard name and address matching process, complemented by an internet search-aided algorithm following Autor et al. (2020a) as well as additional manual matching. The detailed methodology is described in Ding, Jo, and Kim (2022) and Braguinsky et al. (2023).

[6] In this figure, we restrict our sample to Compustat firms with patent filings given the availability of KPSS values and to maintain consistency of the sample across series. The results are robust when using the full set of USPTO assignees.

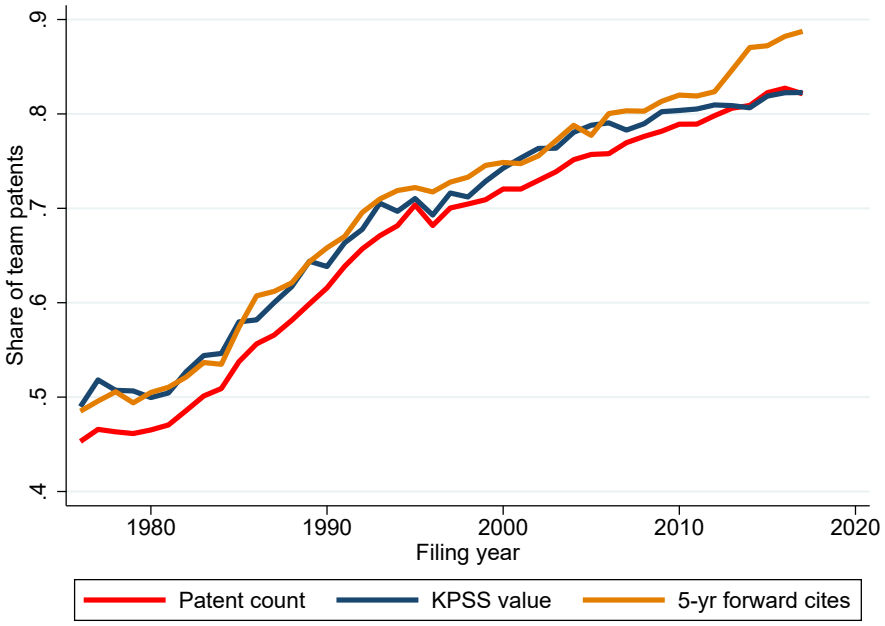


Figure 1: Team Patents Shares

Note: The figure depicts the share of team patents at Compustat firms from 1976 to 2018. A team patent is defined as one assigned by multiple inventors. The number of patents is counted unweighted (red line), weighted by KPSS values from Kogan et al. (2017) (navy line), and weighted by five-year forward citations with vintage fixed effects removed (orange line).

centered on the patent filing date. This probability is 88% (86% if restricting attention to the five years before the patent was filed). Additional details are provided in Appendix A.1. Appendix Figure A1 shows that the probability of two inventors from a patenting team working together peaks around 10 months before the filing date.[7]

Fact 3: Inventor employment concentration has increased much more than overall employment concentration

A growing literature analyzes the causes and consequences of rising employment concentration in the United States (Autor et al. 2020b; Rossi-Hansberg, Sarte, and Trachter 2020). Akcigit and Ates (2023) were the first to point out the increase in employment

[7]As a robustness check, in the USPTO data, we identify an inventor's affiliation in a given year as the assignee (firm) with which the inventor files the largest number of patents in that year. In the event of a tie, we assign the affiliation to the firm associated with the patent with the latest filing date. Using this definition, 89.5% of team patents in our sample are produced by inventors affiliated with the same firm.

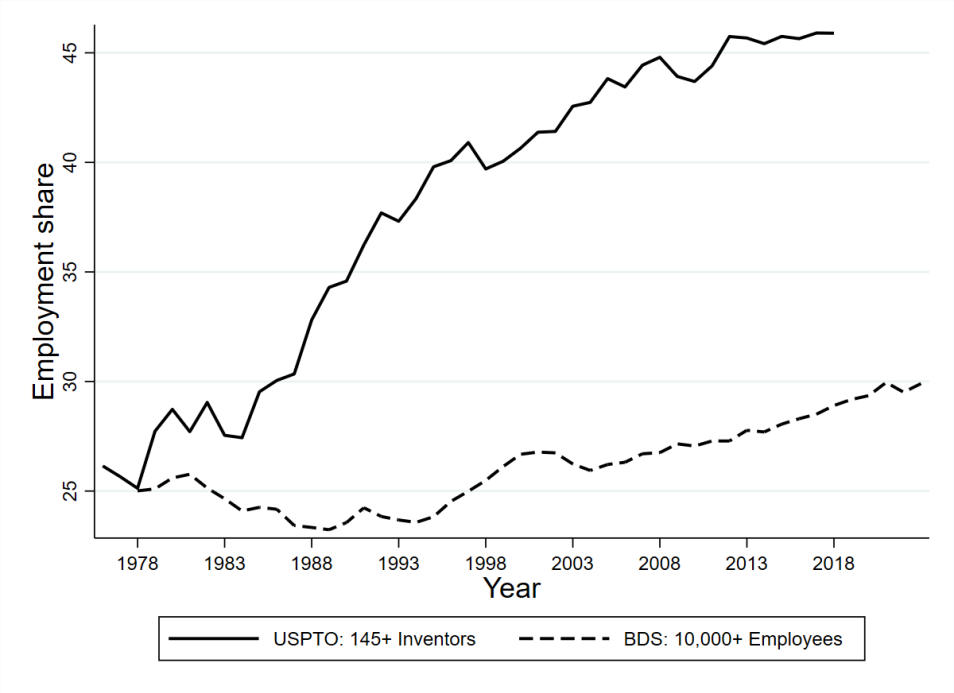


Figure 2: Employment Concentration: Inventors and Overall Employment

Note: The figure plots the employment share of inventors at the largest firms (ranked by the number of inventors), indexed to match the employment share of overall employment at the largest firms (ranked by total employment) in 1978. Inventor employment is measured using patent data in USPTO to count inventors at the firm–year level, while overall employment shares are constructed from the Census Bureau's Business Dynamics Statistics (BDS).

concentration specifically among inventors and Akcigit and Goldschlag (2023) provide additional supporting evidence. To our knowledge, we are the first to provide a direct comparison between overall employment concentration and inventor employment concentration.

To do this, we use the USPTO data to match inventors to firms (assignees in the patent data).[8] A standard measure of overall employment concentration is the share of total employment accounted for by "large firms" (firms with more than 10,000 employees) reported by the Census Business Dynamics Statistics (BDS) (Autor et al. 2020b). We define "large innovative firms" as firms with more than 145 inventors, where 145 is

[8]This dataset includes inventors who patent in a given year only. The Revelio Labs data provides a broader coverage of inventors, including those who do not patent, and we use it as a supplementary source. However, this data does not span a sufficiently long time horizon to study long-run trends in concentration.

chosen to match the share of inventor employment at large innovative firms in 1978 (the first year of the BDS data) to the share of overall employment at firms with more than 10,000 employees in 1978.

Holding fixed this threshold over time, the inventor employment share of large innovative firms has risen about six times more than the corresponding overall employment share of large firms (Figure 2). The model in the next section explores inventor employment concentration can arise from changes in the idea production function that favor knowledge production in teams.[9]

3 Theoretical Framework

To explain the trends documented in Section 2 and assess their implications for the allocation of inventors across firms, we develop a theoretical framework to study the problem of heterogeneous innovative firms hiring inventors in general equilibrium.

The model has the following three key ingredients: first, we introduce firm heterogeneity in productivity to generate dispersion in firm size as is standard in the firm dynamics literature (Hopenhayn 1992; Melitz 2003); second, we provide a novel microfoundation for the firm's choice of either solo or team-based idea production and endogenize the economy's fraction of team patenting; and third, we allow for research productivity spillovers across firms that depend both on a firm's intrinsic productivity and its employment size following Chikis, Kleinman, and Prato (2026).

We begin by specifying inventor-level productivity, which depends on whether the inventor works alone or in a team. We then describe the firm's problem of hiring inventors and deciding whether to organize inventors into solo or team production. Finally, we define and characterize the equilibrium. Throughout, we refer to Appendix A.2, which provides additional empirical evidence supporting the key modeling assumptions and

[9] In other sectors, rising team production and knowledge specialization may have contributed to the rise in overall employment concentration, but team production is much harder to measure outside of research, so we leave this to future work.

results.

3.1 Inventor Productivity

In the initial “brainstorming" phase that determines an inventor’s productivity in the second idea production stage, an inventor can either brainstorm alone (“solo”) or collaborate with another inventor within the same firm (“team”).

Inventor i has the following productivity:

$$\begin{cases} z_{ii} = z & \text{if working alone} \\ z_{ij} \sim Pareto(x, \frac{1}{\gamma}) & \text{if working with another inventor } j. \end{cases} \tag{3.1}$$

When brainstorming alone, an inventor has productivity z.[10] Under team production, inventor i can potentially brainstorm with any other inventor j within the firm. The productivity of each potential collaboration z_{ij} is drawn from a Pareto distribution with scale and shape parameters $(x, \frac{1}{\gamma})$.

A few remarks are in order. First, inventors find potential team members only within the boundary of the firm, consistent with Fact 2. Second, while we do not model team communication costs explicitly, we maintain the assumption throughout that $x < z$, that is, the least productive team collaboration yields lower productivity than working alone, suggesting some costs associated with poorly matched teams.[11] Third, for simplicity, we assume that each potential team consists of only two inventors, supported by the data evidence in Appendix A.2.[12,13]

[10] A simple extension, though one we do not pursue here, would be to allow stochastic productivity for solo inventors with mean z. The firm’s decision of whether to adopt team production would be unchanged.

[11] We also abstract from serial collaboration or team-specific human capital accumulation since our model is static, but this is also consistent with evidence in Appendix A.2 that 80% of teams only patent together once.

[12] Appendix Figure A2 shows that even as firm size grows large, the mean team size within the firm is stable at around three inventors.

[13] Appendix B.1 sketches an extension that allows productivity draws for teams of three. This extension strengthens our mechanism by giving inventors in large firms even more draws over which the firm selects the maximum productivity.

Finally, the pairwise productivity draws reflect potential gains from collaboration: inventors may combine their expertise in highly productive ways, for instance through complementarities in skills or knowledge. The stochastic structure of pairwise productivity implies heterogeneity in match quality within the firm: not all inventors pair equally well, and not all potential collaborations are equally productive. Appendix B.2 contains a further microfoundation for this setup based on Freund (2025): developing a plan for an idea requires many tasks or fields of knowledge over which each inventor has a vector of skills. A parameter χ governs within-inventor skill dispersion across tasks, with a higher value of χ capturing greater specialization (Jones 2009; Freund 2025). The quality z_{ij} of the plan developed by a randomly drawn pair of inventors increases with the distance between i and j's skill vectors. Greater specialization χ amplifies the effect of skill distance on output by amplifying the efficiency gains from specialization: the distribution of team productivity has a fatter tail when specialization increases (Appendix Figure B1). Increased γ in our model can be thought of as greater specialization of inventors in particular tasks or fields of knowledge.

Suppose the firm's objective is to maximize the expected productivity of its inventors, which corresponds to profit maximization in the following section. Since workers are ex ante homogeneous, expected productivity is identical across inventors within a firm of size n. Before pairwise productivities are realized, the firm chooses between solo and team idea production. Under solo idea production, each inventor works independently with productivity z. Under team idea production, the firm observes all realized pairwise productivities among its workers and assigns each worker to the collaboration that yields the highest productivity.[14]

We now characterize expected inventor productivity under team production, which is the relevant object for the firm. Let z_i^{team} denote the maximal team productivity available to inventor i within the firm. In a firm with a discrete number n of inventors,

[14]Note that collaborations are not mutually exclusive: we envision teamwork as a brainstorming activity that determines inventor productivity in the second stage of idea production. Thus, if the same inventor is the most productive match for multiple coworkers in the brainstorming phase they can all be matched with that inventor.

the realized team productivity for inventor i is

$$z_i^{team} = \max\{z_{i1}, z_{i2}, ..., z_{in}\}. \tag{3.2}$$

Given the Pareto distribution of pairwise productivity, the expected team productivity of an inventor in a firm of size n is

$$\mathbb{E}(z_i^{team}) = x\Gamma\left(1-\gamma\right)n^{\gamma}. \tag{3.3}$$

See Appendix B.3.1 for the derivation. Note that this property also holds with a continuum of workers, as we will have in our general equilibrium model, where n is interpreted as the mass of inventors employed by the firm, and bilateral productivities are drawn from an inhomogeneous Poisson point process. See Appendix B.3.2 for details.

Equation (3.3) implies that expected team productivity is increasing in firm size n, the minimum team productivity parameter x, and in the thickness of the right tail of the Pareto distribution of pairwise productivity draws γ which reflects the degree of knowledge specialization of inventors.

Because of the assumption that $z > x$, expected team productivity exceeds solo productivity only in sufficiently large firms. In particular, there exists a firm-size threshold above which inventors achieve higher expected productivity under team production than under solo production. Figure 3 illustrates this relationship and the corresponding cutoff. Consistent with this prediction, Appendix Table A1 documents a positive, economically and statistically significant relationship between innovative firm size (measured by the number of inventors) and the share of team patents in the firm's total patent count using within-firm variation in employment size.

3.2 Firm-Level Innovative Output

In the second stage, firms transform inventor efficiency units—determined during the brainstorming stage—into innovative output using a decreasing-returns production tech-

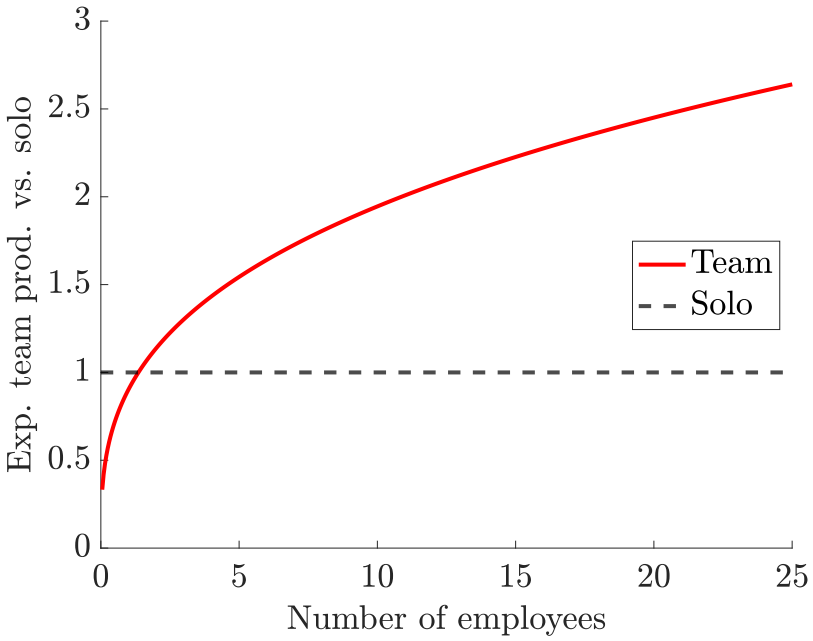


Figure 3: Team vs Solo Productivity

Note: The figure compares worker-level expected productivity under team and solo production as a function of firm size n. The red solid line indicates team production, while the black dashed line shows solo production.

nology (Lucas 1978).

Firms are heterogeneous in productivity, denoted by ε_f, which is drawn from a Pareto distribution, $\varepsilon_f \sim \text{Pareto}(x^f_{\min}, \frac{1}{\eta})$, with cumulative distribution function G. A firm's total innovative output is given by:

$$y_f = \begin{cases} \zeta \varepsilon_f \left(z\, n_f\right)^{\alpha}, & \text{if solo,} \\ \zeta \varepsilon_f \left(x\, \Gamma(1-\gamma)\; n_f^{\gamma}\, n_f\right)^{\alpha}, & \text{if team.} \end{cases} \tag{3.4}$$

The key prediction of the model is that team-based firms exhibit less-decreasing returns to scale in inventor labor than solo-based firms, reflecting a within-firm agglomeration effect through which additional coworkers raise inventor productivity. With a single input, the returns-to-scale parameter coincides with the output elasticity with respect to inventor labor. For a solo-based firm, this is simply α, whereas for a team-based

firm, this is $\alpha(1+\gamma)$. Throughout, we assume $\alpha(1+\gamma) < 1$ to ensure decreasing returns to scale for both types of firms (and thus an interior solution for labor demand). Our empirical estimates of the firm-level idea production function in Section 5.1 are consistent with this assumption.

Finally, firms learn from one another through the knowledge spillover term ζ, which affects the productivity of all firms symmetrically. This assumption is motivated by the finding that the social returns to innovation are substantially larger than the private returns (Bloom, Schankerman, and Van Reenen 2013; Dyvere 2024). We formalize spillovers in the model following Chikis, Kleinman, and Prato (2026). In particular, the spillovers generated by a particular firm depend jointly on its productivity ε_f and size n_f, and aggregate across firms as follows:

$$\zeta = \left(\int_{f \in F} \varepsilon_f n_f^{\beta} df \right)^{\theta}, \tag{3.5}$$

where F is the measure of active firms in the economy. $\theta \geq 0$ governs the importance of knowledge spillovers for firm output, and $\beta \in [0,1]$ governs the relative importance of firm productivity (ε_f) and firm size (n_f) for generating spillovers. Firm size may matter because ideas are embodied in inventors, making other inventors more likely to learn from a firm when it employs a larger inventive workforce.

When $\beta = 0$, knowledge spillovers depend only on the set of active firms–their productivity and mass–but not on how inventors are allocated across them. For $\beta \in (0,1]$, large firms contribute more to knowledge spillovers than small firms, with β governing the curvature of spillovers with respect to inventor labor. A lower value of β implies greater concavity and, consequently, lower marginal social returns to an additional inventor in a given firm.[15]

[15] Chikis, Kleinman, and Prato (2026) use data on citations patterns to the same firm in different locations where the firm has plants of different sizes and estimate $\beta = 0.25$. In Section 5.2, we estimate β following Chikis, Kleinman, and Prato (2026) with our data for 1976-1999 and 2000-2018 separately.

3.3 Firm Problem

Given their idiosyncratic productivity draw, innovative firms solve a static profit maximization problem. Producing requires a fixed operating cost wc (overhead inventor labor needed to run a lab). These firms simultaneously choose whether to produce or exit, and conditional on operating, choose the number of inventors to hire and whether to organize these inventors into solo or team production. The following characterizes the firm's profit function:

$$\Pi(\varepsilon_f) = \max\Big\{0, \Pi^{solo}(\varepsilon_f), \Pi^{team}(\varepsilon_f)\Big\}, \tag{3.6}$$

where

$$\Pi^{solo}(\varepsilon_f) = \max_{n_f^{solo}} \zeta\varepsilon_f z^{\alpha}(n_f^{solo})^{\alpha} - wn_f^{solo} - wc, \tag{3.7}$$

$$\Pi^{team}(\varepsilon_f) = \max_{n_f^{team}} \zeta\varepsilon_f \left(x\Gamma(1-\gamma)\right)^{\alpha} (n_f^{team})^{(\gamma+1)\alpha} - wn_f^{team} - wc. \tag{3.8}$$

Throughout, we normalize the "price" of innovative output to 1. The model could be extended to include a production sector where ideas generated by the firm's R&D lab are used to produce physical output.

Profit maximization creates a cutoff for ε_f above which firms choose team production. Let $\tilde{\varepsilon}^{team}$ denote this cutoff. The fixed cost creates another cutoff for firm productivity, $\tilde{\varepsilon}^{exit}$, below which firms endogenously exit. Let $d = \{0, 1\}$ denote the exit choice, which equals to 1 if firm exits.

The optimal inventor employment size of firms under solo and team production can be derived as follows:

$$n^*(\varepsilon_f) = \begin{cases} \left(\frac{\alpha\zeta\varepsilon_f z^{\alpha}}{w}\right)^{\frac{1}{1-\alpha}} & \text{w/ solo (if } \tilde{\varepsilon}^{exit} < \varepsilon_f < \tilde{\varepsilon}^{team}) \\ \left(\frac{(\gamma+1)\alpha\zeta\varepsilon_f(x\Gamma(1-\gamma))^{\alpha}}{w}\right)^{\frac{1}{1-(\gamma+1)\alpha}} & \text{w/ team (if } \tilde{\varepsilon}^{team} < \varepsilon_f), \end{cases} \tag{3.9}$$

where $\tilde{\varepsilon}^{exit}$ and $\tilde{\varepsilon}^{team}$ are determined by the following equations, respectively,

$$\tilde{\varepsilon}^{exit} = \frac{c^{1-\alpha}}{(1-\alpha)^{(1-\alpha)}(\alpha z)^{\alpha}} \frac{w}{\zeta}, \tag{3.10}$$

$$\tilde{\varepsilon}^{team} = \left(\frac{(1-\alpha)\,(\alpha z)^{\frac{\alpha}{1-\alpha}}}{(1-(\gamma+1)\alpha)\,((\gamma+1)\alpha)^{\frac{(\gamma+1)\alpha}{1-(\gamma+1)\alpha}}\,(x\Gamma(1-\gamma))^{\frac{\alpha}{1-(\gamma+1)\alpha}}} \right)^{\frac{(1-\alpha)}{\alpha}\frac{1}{\gamma}(1-(\gamma+1)\alpha)} \frac{w}{\zeta}. \tag{3.11}$$

3.4 Firm Entry

A large mass of potential entrants can enter by paying an entry cost $w\kappa$ (innovative entrepreneurial labor) before drawing their productivity. The following free entry condition must hold in equilibrium:

$$\int \Pi(\varepsilon_f) dG(\varepsilon_f) - w\kappa = 0, \tag{3.12}$$

which pins down the mass of firms M (including potential entry). Entrants draw their idiosyncratic productivities from the same distribution G of incumbent productivities.

3.5 Labor Market Clearing

There is a fixed supply of inventors with measure 1. Labor market clearing requires that the aggregate inventor labor demand is equal to the supply:

$$M \left(\int_{\tilde{\varepsilon}^{team}} \left(n^{team}(\varepsilon_f; w) + c \right) dG(\varepsilon_f) + \int_{\tilde{\varepsilon}^{exit}}^{\tilde{\varepsilon}^{team}} \left(n^{solo}(\varepsilon_f; w) + c \right) dG(\varepsilon_f) + \kappa \right) = 1. \tag{3.13}$$

This pins down the equilibrium wage w in the economy.

3.6 Equilibrium

An equilibrium is a collection of $\left\{\{n_f^{solo}, n_f^{team}, d_f, \Pi_f\}_f, \tilde{\varepsilon}^{exit}, \tilde{\varepsilon}^{team}, w, M\right\}$, where firms solve (3.6)-(3.9); the free-entry condition (3.12) holds; and the labor market clearing condition (3.13) is satisfied.

4 Social Planner's Problem

Before turning to quantification, we consider analytically how the efficient allocation of inventors across firms differs from the allocation in competitive equilibrium. For expositional simplicity, we focus on polar cases where firms operate exclusively under either solo or team production and abstract from firm entry and exit by assuming a fixed measure of firms. This setting provides intuition for the differences between the efficient allocation and the competitive equilibrium under solo and team production and how the magnitude of misallocation varies across the two production schemes. In the quantitative model (Section 6.4), we solve the full planner's problem, which jointly determines firm size n_f, the choice between solo and team production for each firm, the mass of firms M, and the firm exit cutoff (equivalently, the mass of active firms).

The planner allocates inventors across firms to maximize total innovative output, internalizing spillovers, the idea production technology, and the resource constraint. The planner's problem can be written as follows:

$$\max_{n_f} Y = \int_0^1 y_f df, \tag{4.1}$$

subject to (3.4), (3.5), and

$$\int_0^1 n_f df = 1. \tag{4.2}$$

We define a firm-level wedge τ_f as the gap between the social marginal product of inventor labor at firm f (MP_f^{SP}) and the private marginal revenue product (MP_f^{CE})

perceived by the firm:

$$MP_f^{SP} = (1+\tau_f) MP_f^{CE}. \tag{4.3}$$

4.1 Solo Production

We first solve for the optimal allocation of inventors in the case where all firms operate under solo production. Under solo production, the planner faces the following marginal product of inventor labor:

$$MP_f^{solo,SP} = \theta\beta\zeta^{-\frac{1}{\theta}}\varepsilon_f n_f^{\beta-1} Y + \alpha\zeta z^{\alpha}\varepsilon_f n_f^{\alpha-1}, \tag{4.4}$$

and the counterpart in the competitive equilibrium with all solo production is

$$MP_f^{solo,CE} = \alpha\zeta z^{\alpha}\varepsilon_f n_f^{\alpha-1}. \tag{4.5}$$

Using (4.3), (4.4) and (4.5), the wedge under solo production is:

$$\tau_f^{solo} = \frac{\theta\beta Y}{\alpha\zeta^{1+\frac{1}{\theta}}} z^{-\alpha} n_f^{\beta-\alpha}. \tag{4.6}$$

The wedge includes a common component for all firms and a firm-specific component that depends on idiosyncratic productivity ε_f. In the spirit of Hsieh and Klenow (2009), we evaluate this wedge at the competitive equilibrium firm size distribution and obtain the following result.

Proposition 1. *If $0 < \beta < \alpha < 1$, firm-level wedges are decreasing in firm productivity ε_f and size n_f under solo production.*

Proof. See Appendix B.4.1. □

Proposition 1 implies misallocation of inventor labor across firms due to knowledge spillovers, and the planner values additional inventors at small firms the most.

4.2 Team Production

Under team production, the planner faces the following marginal product of labor:

$$MP_f^{team,SP} = \theta\beta\zeta^{-\frac{1}{\theta}}\varepsilon_f n_f^{\beta-1}Y + \alpha\big(1+\gamma\big)\zeta\left(x\Gamma(1-\gamma)\right)^{\alpha}\varepsilon_f n_f^{\alpha(1+\gamma)-1}, \tag{4.7}$$

and the counterpart in the competitive equilibrium is

$$MP_f^{team,CE} = \alpha\big(1+\gamma\big)\zeta\left(x\Gamma(1-\gamma)\right)^{\alpha}\varepsilon_f n_f^{\alpha(1+\gamma)-1}. \tag{4.8}$$

Using (4.3), (4.7) and (4.8) delivers the wedge under team production as follows:

$$\tau_f^{team} = \frac{\theta\beta Y}{\alpha(1+\gamma)\zeta^{1+\frac{1}{\theta}}}\left(x\Gamma(1-\gamma)\right)^{-\alpha} n_f^{\beta-\alpha(1+\gamma)}. \tag{4.9}$$

The following proposition shows the evaluation of the wedges at the competitive equilibrium firm size distribution under team production.

Proposition 2. *If $0 < \beta < \alpha(1+\gamma) < 1$, firm-level wedges are decreasing in firm productivity ε_f and size n_f under team production.*

Proof. See Appendix B.4.2. □

Proposition 2 establishes that team production similarly generates misallocation of inventor labor across firms. The following proposition shows that team production increases the dispersion of wedges as a function of firm productivity compared to the all-solo economy. By increasing firms' private returns to scale, team formation causes wedges between social and private returns to decline more rapidly with firm productivity, thereby increasing the relative distortion between small and large firms.

Proposition 3. *Suppose $0 < \beta < \alpha < \alpha(1+\gamma) < 1$. Then the wedges between social and private returns are more sensitive to firm productivity under team production than under*

solo production. In particular,

$$\left|\frac{\partial \log \tau_f^{team}}{\partial \log \varepsilon_f}\right| > \left|\frac{\partial \log \tau_f^{solo}}{\partial \log \varepsilon_f}\right|.$$

Equivalently, wedges decline more rapidly with firm productivity under team production than under solo production.

Proof. See Appendix B.4.3. □

This proposition provides intuition for the result in Section 6.4 that an increase in team production increases misallocation by increasing wedges for small firms.

5 Calibration

We calibrate two sets of model parameters, one for the early period, 1976-1999 (denoted "E"), and another for the late period, 2000-2018 (denoted "L"). Several parameters are obtained by directly estimating the model's idea production function. The knowledge spillover function is also directly estimated for the two periods. Other parameters are internally calibrated to match the early-period team patent share, the entry rate into patenting, and the average firm size. The remaining parameters are externally determined, either through normalization or taken from the literature.

5.1 Estimating the Idea Production Function

Idea production function identification. First, to pin down the idea production function parameters (α, γ, z, x), we estimate the model's firm-level idea production function in the data. Specifically, we implement the following triple-interaction regression using the universe of patenting firms in the USPTO data over the period 1976–2018:

$$\log(y_{ft}) = \mu_f + \phi_1 \log n_{ft} + \phi_2 \log n_{ft} \times \mathbb{I}(team_{ft}) + \phi_3 \log n_{ft} \times \mathbb{I}(team_{ft}) \times Post_t$$

Table 1: Mapping from Estimated Coefficients to Structural Parameters

Parameter	Interpretation	Period	Mapping
Panel A. Idea Production Function			
α^E	RTS, solo firms	Early	$\alpha^E = \phi_1$
α^L	RTS, solo firms	Late	$\alpha^L = \phi_1 + \phi_6$
γ^E	Higher RTS, team firms	Early	$(\gamma^E + 1)\alpha^E = \phi_1 + \phi_2$
γ^L	Higher RTS, team firms	Late	$(\gamma^L + 1)\alpha^L = \phi_1 + \phi_2 + \phi_3 + \phi_6$
z^E	Solo prod.	Early	$\alpha^E \ln z^E + \ln \zeta^E = Cons.$
z^L	Solo prod.	Late	$\alpha^L \ln z^L + \ln \zeta^L = Cons. + \phi_5$
x^E	Min. team prod.	Early	$\alpha^E \ln \left(x^E \Gamma(1-\gamma^E)\right) + \ln \zeta^E = Cons. + \phi_4$
x^L	Min. team prod.	Late	$\alpha^L \ln \left(x^L \Gamma(1-\gamma^L)\right) + \ln \zeta^L = Cons. + \phi_4 + \phi_5 + \phi_7$
Panel B. Knowledge Spillover Function			
β^E	Scale for n in spillovers	Early	$\beta^E = \tilde{\phi}_1$
β^L	Scale for n in spillovers	Late	$\beta^L = \tilde{\phi}_1 + \tilde{\phi}_2$

Notes: This table summarizes the mapping from the regression coefficients in (5.1) (Panel A) and (5.2) (Panel B) to the model's structural parameters. The regression estimates are in Tables 2 and 3 and the resulting parameter values are in Table 4.

$$+ \phi_4 \mathbb{I}(team_{ft}) + \phi_5 Post_t + \phi_6 \log n_{ft} \times Post_t + \phi_7 \mathbb{I}(team_{ft}) \times Post_t + Cons. + \nu_{ft}. \tag{5.1}$$

The baseline measure of firm output is patent counts, where y_{ft} denotes the number of patents filed by firm f in year t.[16] The term μ_f captures a firm fixed effect, and n_{ft} is the number of (patenting) inventors employed by firm f in year t. The indicator $\mathbb{I}(team_{ft})$ equals one if the firm's share of team patents in its total patent count $\left(\frac{\text{team patents}_{ft}}{\text{total patents}_{ft}}\right)$ in year t exceeds the sample median (0.88), and zero otherwise. $Post_t$ is a dummy equal to one for the late period (2000-2018). Identification of ϕ_2 and ϕ_3, which pin down the returns to scale benefits of teams in the early and late periods, relies on random variation within a firm in the choice to use solo versus team production conditional on size n_{ft}. The regression coefficients map directly into the model parameters. This mapping is summarized in Table 1.

[16] To enter the data, a firm-year observation must have at least one patent, so taking logs does not drop any data.

Table 2: Returns to Scale in Team Production

	(1) All	(2) Public Firms
$\log n_{ft}$	0.757***	0.818***
	(0.002)	(0.004)
$\log n_{ft} \times \mathbb{I}(team_{ft})$	0.039***	0.0875***
	(0.004)	(0.006)
$\log n_{ft} \times \mathbb{I}(team_{ft}) \times Post_t$	0.031***	0.011*
	(0.004)	(0.006)
$\mathbb{I}(team_{ft})$	-0.677***	-0.732***
	(0.004)	(0.008)
$Post_t$	-0.019***	-0.058***
	(0.003)	(0.009)
$\log n_{ft} \times Post_t$	0.003	0.019***
	(0.002)	(0.004)
$\mathbb{I}(team_{ft}) \times Post_t$	-0.049***	-0.015
	(0.005)	(0.011)
$\log K_{ft}^{intan.}$		-0.012***
		(0.002)
$Constant$	0.065***	0.081***
	(0.002)	(0.014)
R-squared	0.704	0.817
Obs.	1105437	131166
Firm FE	Yes	Yes

Notes: Source: USPTO, 1976–2018. Column (1) includes all firm–year observations. Column (2) includes public firms with available intangible capital data from Peters and Taylor (2017). See equation (5.1) for details. Standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Idea production function results. The regression estimation results are reported in Column 1 of Table 2, and the implied structural parameter estimates are summarized in Table 4. First, the estimated value of $\alpha^E = 0.757 < 1$ supports the assumption of decreasing returns to scale in inventor labor for solo firms in the early period. There is no evidence of a change in returns to scale for solo firms over time: the coefficient ϕ_6 on $\log n_{ft} \times Post_t$ is close to zero and not statistically significant.

The estimated value of γ, the shape of the team productivity distribution governing the returns to scale benefit of teams, is $\gamma^E = 0.052$ in the early period, indicating that team production increases returns to scale, consistent with the model. Together,

the estimates of α^E and γ^E imply that team-based firms have decreasing returns to scale, but less-decreasing than solo-based firms (0.796 instead of 0.757 for the early period). The positive coefficient ϕ_3 on the triple interaction indicates that the returns to scale benefit of teams rises in the late period, possibly due to increasing knowledge specialization. Consistent with rising complementarities across inventors over time due to specialization, Appendix A.3 documents that patenting teams increasingly combine inventors from different technological fields—defined by the technology class in which each inventor patents most frequently—while individual inventors become increasingly specialized in their respective fields (that is, move less frequently between technology fields). The estimated returns to scale of team-based firms in the late period is 0.827.

To infer solo inventor productivity z from the regression, we first need to partial out the effect of knowledge spillovers ζ, where $\zeta^E = \left(\int_{f \in F} \varepsilon_f n_f^{\beta^E} df\right)^{\theta}$, using the value of β^E estimated in Section 5.2 below.[17] The resulting value for the early period is $z^E = 1.247$. The negative coefficient ϕ_5 on $Post_t$ indicates a broad decline in patenting productivity in the later period. Spillovers decline endogenously in the late period of the model; after partialling spillovers out, we estimate an increase in solo productivity, $z^L = 1.366$. This suggests that ideas getting harder to find (Bloom et al. 2020) may be driven by declining knowledge spillovers and not by the intrinsic quality of inventors.

The early period value of the minimum team productivity x^E is internally calibrated to match the team patent share in the early period ($x^E = 1.065$). Replacing $Cons.+\phi_4$ in the mapping in Table 1 with x^E to calibrate the change in x over time yields $x^L = 1.062$, indicating a slight decrease in the productivity of the least productive teams.

Robustness. Table 2, Column 2 report estimates from the same specification, limiting the sample to public firms, which makes it possible to control for capital inputs to R&D using the Peters and Taylor (2017) measures of intangible capital (capitalized past R&D expenses). Controlling for knowledge capital addresses concerns that omitted inputs (capital) that are positively correlated with inventor employment or team production

[17] Practically, the implementation involves guessing a set of parameter values and computing ζ in the model, then using the equations in Table 1 to arrive at the structural parameter values for z and x.

Table 3: Knowledge Spillovers: Returns to Scale in Inventor Labor

	(1)	(2)
$\log n_{fmt}$	0.429***	0.467***
	(0.011)	(0.015)
$\log n_{fmt} \times Post_t$		-0.061***
		(0.016)
R-squared	0.147	0.148
Obs.	106,362	106,362
Firm-Year FE	Yes	Yes

Notes: The dependent variable is the log number of five-year forward external citations received by firm f's patents from other firms in market m (commuting zone) divided by the number of firms in market m. The key explanatory variable is the log number of inventors at firm f in market m. Column (2) allows the elasticity of knowledge spillovers with respect to inventor employment to differ after 2000 through an interaction with $Post_t$. All specifications include firm-year fixed effects. See (5.2) for details. Standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

could bias the estimated returns to scale upward. Our main findings remain qualitatively robust.

We report additional robustness tests in Appendix A.4. In Appendix Table A2, we re-estimate equation (5.1) using forward citations instead of patent counts as the measure of firm output, adjusted for technology-class fixed effects for citations. In this case, the returns-to-scale advantage of team production is statistically significant only in the post-2000 period but the general patterns remain similar to the baseline results.

Another concern is that inventor employment is measured with error, since the USPTO data capture only inventors who successfully patent in a given year. To address this issue, we use Revelio Labs data on inventor employment histories to construct a better measure of inventors per firm. However, these data are reliable only from around 2008 onward. Appendix Table A3 reports the corresponding estimates. While the baseline returns-to-scale parameter (α) is lower in this sample, the estimated team effects (γ) in the pre- and post-periods (2008–2014 versus 2014–2018) remain qualitatively consistent with our main findings.

5.2 Estimating the Knowledge Spillover Function

Spillover function identification. We use local citation flows to estimate the concavity of the model's knowledge spillover function in firm size (β) with the following regression:

$$log(\iota_{fmt}) = \mu_{ft} + \tilde{\phi}_1 \log n_{fmt} + \tilde{\phi}_2 \log n_{fmt} \times Post_t + \nu_{fmt}, \tag{5.2}$$

where ι_{fmt} is the number of five-year forward external citations that firm f's patents in year t receive from other firms in market m (commuting zone) divided by the number of firms in market m in year t, n_{fmt} is the number of inventors at firm f in market m, μ_{ft} is a firm-year fixed effect, and $Post_t$ is an indicator for the post-2000 period. While our model has no notion of space, this strategy allows us to control for firm by year fixed effects to cleanly identify the role of firm size for spillovers separately from productivity (which is assumed to be the same across locations).[18]

Our approach is therefore identical to Chikis, Kleinman, and Prato (2026), except that we allow β to vary between the two time periods since our focus is on changes in spillovers over time.[19] The coefficient $\tilde{\phi}_1$ associated with the number of inventors captures the returns to scale of R&D labor in creating spillovers in the early period, and the coefficient $\tilde{\phi}_2$ associated with the interaction term measures the change in these returns during the post-2000 period.

Spillover function results. Table 3 presents the estimation results. Column 1 estimates the returns to scale in knowledge spillovers using the pooled sample and yields an estimate of $\beta = 0.429$. Column 2 allows it to differ between the pre- and post-2000

[18] We expect that local spillovers scale more than national spillovers in firm size, since physical location matters for inventors bumping into and learning from each other, so the estimated local concavity might be an upper bound for overall spillovers. Lower values of β strengthen the misallocation results in Section 6.4.

[19] While the main data (USPTO) is the same in both studies, Chikis, Kleinman, and Prato (2026) use Dun and Bradstreet data to identify establishments whereas we rely on inventor location in the patent data as a proxy for establishment location.

periods. The estimates imply that the returns to scale of spillovers with respect to firm size were $\beta^E = 0.467$ before 2000 and declined to $\beta^L = 0.406$ after 2000.

Allowing for differences in β over time is important for separating in the model how much of the decline in knowledge spillovers and the increase in misallocation is due to changes in the idea production function that favor teams versus a technological change in the knowledge spillover function that would have reduced spillovers and increased misallocation without any changes to the idea production function.

5.3 Other Parameters: Early Period

Based on the estimation results, we calibrate the model as follows. Table 4, Panel A reports the calibrated parameters for the early period. We first choose a functional form for the exogenous firm productivity distribution G: a Pareto distribution with shape $x^f_{\min}$ and scale η. Then we externally calibrate six parameters, $(\alpha, \gamma, z, x^f_{\min}, \beta, \theta)$ as follows. We use the estimation results in Table 2 to pin down the returns to scale in production, α, the returns to scale in team productivity, γ, and the productivity of solo inventors, z. We normalize the lower bound of the firm-level productivity distribution, $x^f_{\min}$, to one. We calibrate β using the estimates reported in Table 3. For θ, we use the estimate from Chikis, Kleinman, and Prato (2026), who identify values in the range $[0.04, 0.12]$.[20]

Next, four parameters (x, η, κ, c) are calibrated internally to match four empirical moments. The fit for these targets is reported in Table 5. The share of team patents over 1976–1999, equal to 57.8 percent (Figure 1), pins down the lower bound of team productivity, x. Second, the Pareto tail parameter of the firm productivity distribution, η, is chosen to match the employment share of large firms (with 145 inventors or more), which was 32.8% in the early period (Figure 2). Third, we target the entry rate of patenting firms, defined as the share of entrants (age 0) among all patenting firms using Business Dynamics Statistics (BDS) data over 1978–1999, which yields an average

[20] This range is also consistent with estimates reported in (Greenstone, Hornbeck, and Moretti 2010; Moretti 2021) using different samples and time periods.

Table 4: Calibration

Panel A: Early Period (1976-1999)			
Parameter	Definition	Value	Target/Source
α	RTS, solo firms	0.757	Table 2
γ	Higher RTS, team firms	0.052	Table 2
z	Solo prod.	1.247	Table 2
x	Min. team prod.	1.065	Share of team patents
$x^f_{\min}$	Location for ε	1	Normalization
η	Tail for ε	0.164	Emp. share of large firms
β	Scale for n in spillovers	0.467	Table 3
θ	Overall spillovers	0.080	Chikis, Kleinman, and Prato (2026)
κ	Entry cost	1.456	Patenting entry rate
c	Fixed cost	0.512	Average firm size
Panel B: Late Period (2000-2018)			
Parameter	Definition	Value	Target
α	RTS, solo firms	0.757	Fixed to early period
γ	Higher RTS, team firms	0.093	Table 2
z	Solo prod.	1.366	Table 2
x	Min. team prod.	1.062	Table 2
$x^f_{\min}$	Location for ε	1	Normalization
η	Tail for ε	0.164	Fixed to early period
β	Scale for n in spillovers	0.406	Table 3
θ	Overall spillovers	0.080	Chikis, Kleinman, and Prato (2026)
κ	Entry cost	1.456	Fixed to early period
c	Fixed cost	0.512	Fixed to early period

Notes: This table reports the parameter values for the calibrated model in the early (Panel A) and late period (Panel B) and the target or source for each parameter value.

annual entry rate of 2.7 percent.[21] This moment pins down the entry cost, κ. Finally, the fixed operating cost, c, is chosen to match the average inventor employment size of patenting firms over 1976–1999, equal to 7.1 inventors.

5.4 Other Parameters: Late Period

Our central question is whether changes in the idea production function can account for the observed rise in inventor employment concentration. We hold fixed the parameters

[21] The BDS data are available starting in 1978.

Table 5: Targeted Moments For the Early Period (1976-1999)

Target Moment	Data	Model
Share of team patents	57.8	57.8
Employment share of firms with $n > 145$	32.8	32.8
Patenting entry rate (annual)	2.7	2.7
Average firm size (inventors)	7.1	7.1

Notes: This table reports the targeted moments in the data (in the first column) and their corresponding values implied by the model (in the second column). The data sources are the USPTO (1976-1999) and the U.S. Census Bureau Business Dynamics Statistics (BDS, 1978-1999).

related to industry primitives (the firm productivity distribution parameter η, the entry cost κ and the fixed cost c) as well as the returns-to-scale parameter α. We then vary the parameters directly governing idea production and knowledge spillovers (γ, z, x, and β) using the estimates from equations (5.1) and (5.2). Panel B of Table 4 reports the resulting parameter values for the late period.

6 Quantitative Results

This section describes the model fit for targeted and non-targeted moments, then decomposes the role of the idea production parameters and the knowledge spillover parameter in explaining changes in the firm size distribution, spillovers across firms, and innovative output over time. The final part of this section compares the laissez faire economies in the early and late period to their efficient counterparts and shows that misallocation of inventors has risen over time due to changes in the idea production function that favor teams.

6.1 Model Fit and Validation

Table 5 shows that the calibrated model exactly matches the targeted moments for the early period. Table 6 reports the model predicted levels of and changes in various aggregate moments over time due to the estimated changes in the idea production function and knowledge spillover parameters, including various non-targeted moments.

Table 6: Model Predicted Changes in Aggregate Moments

	Early	Late	Change
Inventor wage w	0.724	0.878	21.3
Prod. cutoff (exit)	1.004	1.241	23.6
Prod. cutoff (team)	1.580	1.740	10.1
Share of solo-based firms	93.7	87.2	-6.5
Share of team patents	57.8	90.4	32.7
Employment share of firms with $n > 145$	32.8	78.9	46.1
Average firm size	7.1	27.1	282.3
Spillovers ζ	0.903	0.827	-8.4
Output (normalized)	1.000	1.213	21.3

Notes: This table reports the model-implied changes in aggregate moments between the early and late periods. The first two columns are the moments for the early and late periods, respectively, and the third column reports the corresponding percent or percentage point change relative to the early-period benchmark.

The model predicts a 21.3% increase in inventor wages, consistent with the evidence in Ekerdt and Wu (2025), who document a rise of roughly 15 percentage points in the wage premium for researchers since 1970. The increase in the inventor wage raises both the equilibrium exit cutoff and the cutoff for using teams by making hiring inventors more expensive.

The model also implies a decline in the share of firms operating with solo production of 6.5 percentage points. Consistent with this prediction, the data show a decline of approximately 12 percentage points in the share of innovative firms with only one inventor (from 39% to 27% of firm–year observations). Using our empirical definition of solo-based firms from the regression (those with below-median team patent share), the share of solo-based firms declined by 15 percentage points, from 60% to 45%.

As an additional validation of the model's key prediction that team patent share and concentration are correlated in the aggregate, we investigate this correlation at the technology class (CPC group) level in Appendix A.5.[22] Inventor employment concentration is positively associated with the share of team patents at the CPC technology-class level even after controlling for technology class and year fixed effects (Appendix Table A4). In changes, an increase in the team patenting share between the early and late

[22]An example of a CPC group is B60L 53: Methods or charging stations for charging electric vehicles.

periods is positively associated with a rise in inventor employment concentration at the technology class level (Appendix Figure A10).

6.2 Aggregate Implications of the Rise of Teams

The model predicts a substantial rise in employment concentration between the early and late periods. The employment share of large firms (defined as those with 145 or more inventors) increases from 32.8% to 78.9%, compared with an increase to 44.1% in the data. Average firm size in the model also increases from 7.1 inventors to 27.1, whereas in the data the average size rose to 10.8 inventors per firm. The team patent share rises to 90.4%, compared to 77.6% in the data.[23]

Together, these changes reduce knowledge spillovers by 8.4%. This suggests a possible microfoundation for the declining knowledge diffusion hypothesis of Akcigit and Ates (2023): shifts in the idea production function generate greater misallocation of inventors across firms by concentrating inventors in large firms, which in turn dampens spillovers since the returns to additional inventors in large firms in terms of additional knowledge spillovers are low. However, the rise in returns to scale from team production increases output despite the decline in spillovers across firms (it has a similar effect on output to increasing returns to scale α in a standard Lucas (1978) model). Section 6.3 decomposes the role of the individual idea production parameters in more detail, and we return to the issue of declining knowledge spillovers and misallocation in Section 6.4.

Robustness. Given that the baseline exercise overpredicts some of the changes in the data, Appendix B.5.1 explores an alternative exercise in which γ^L is calibrated internally (rather than taken directly from the regression) to match the late period team patent share and (x^L, z^L, β^L) are calibrated in the same way as before. In this exercise, employment concentration rises to 60.7% and the average firm size rises to

[23]The robustness section below shows that the main results are robust to an alternative calibration that fits these moments more tightly.

13.1, both closer to the changes in the data.

Appendix B.5.2 reports two additional experiments where the exogenous firm productivity distribution, $\varepsilon_f \sim \text{Pareto}(x^f_{min}, \frac{1}{\eta})$, is allowed to vary between the early and late periods, either jointly with the idea production and spillover parameters or on its own. In particular, we recalibrate η^L so that the model matches the rise in employment concentration. Jointly changing industry primitives and the idea production function and spillover parameters allows the model to jointly fit changes in the team patent share and employment concentration, but varying the firm productivity distribution alone explains less than half of the rise in team patenting in the data, suggesting that changes in the idea production function are needed to fully explain the rise in team patenting.

6.3 Decomposition of Aggregate Outcomes

This section decomposes the contribution of each estimated change in the idea production function parameters (z, x, γ) individually and isolates the role of idea production function changes from the role of the spillover function parameter β to the evolution of key aggregate moments. Table 7 shows the results for each of the idea production parameters one at a time. Table 8 presents the results for changes in all three idea production parameters ("IP only") with all else (including β) fixed to the early period. Below, we discuss the effects of these parameter changes on each aggregate variable.

Share of team patents. The increase in solo productivity z on its own reduces the share of team patents because it benefits solo production relative to team production. Similarly, the decrease in the minimum team productivity x also predicts a decline in the share of team patents. On the other hand, the increase in the returns to scale benefit of teams, captured by the thicker right tail of the team productivity distribution (higher γ), increases the predicted team patent share substantially, to 97.8% of all patents. The joint changes in the idea production parameters explain the entire rise in the team patent share between the early and late periods (Table 8).

Table 7: Decomposing the Role of Idea Production Function Changes

	Early	Late	z only	Change
Inventor wage w	0.724	0.878	0.757	4.6
Prod. cutoff (exit)	1.004	1.241	1.000	-0.4
Prod. cutoff (team)	1.580	1.740	2.358	49.2
Share of solo-based firms	93.7	87.2	99.5	5.8
Share of team patents	57.8	90.4	30.8	-26.9
Employment share of firms with $n > 145$	32.8	78.9	27.3	-5.6
Average firm size	7.1	27.1	6.6	-7.7
Spillovers ζ	0.903	0.827	0.908	0.5
Output (normalized)	1.000	1.213	1.046	4.6
	Early	Late	x only	Change
Inventor wage w	0.724	0.878	0.724	-0.1
Prod. cutoff (exit)	1.004	1.241	1.003	-0.1
Prod. cutoff (team)	1.580	1.740	1.596	1.0
Share of solo-based firms	93.7	87.2	94.1	0.4
Share of team patents	57.8	90.4	56.9	-0.9
Employment share of firms with $n > 145$	32.8	78.9	32.7	-0.1
Average firm size	7.1	27.1	7.1	-0.6
Spillovers ζ	0.903	0.827	0.903	0.0
Output (normalized)	1.000	1.213	0.999	-0.1
	Early	Late	γ only	Change
Inventor wage w	0.724	0.878	0.877	21.0
Prod. cutoff (exit)	1.004	1.241	1.327	32.1
Prod. cutoff (team)	1.580	1.740	1.455	-8.0
Share of solo-based firms	93.7	87.2	42.8	-50.9
Share of team patents	57.8	90.4	97.8	40.1
Employment share of firms with $n > 145$	32.8	78.9	80.4	47.6
Average firm size	7.1	27.1	41.0	477.7
Spillovers ζ	0.903	0.827	0.827	-8.4
Output (normalized)	1.000	1.213	1.210	21.0

Notes: This table reports the counterfactual exercises in which we vary one parameter at a time—z (top panel), x (middle panel), and γ (bottom panel)—while holding all other parameters fixed at their early-period values. The first two columns are the baseline model moments for the early and late periods, respectively. The third column presents the model-implied value in a counterfactual when only the given parameter is adjusted to its late-period estimate, and the last column reports the corresponding percent or percentage change relative to the early-period benchmark.

Inventor employment concentration. The increase in solo productivity z and the decline in the minimum team productivity both increase the productivity cutoff for adopting team production. This reduces the average firm size slightly and results in

Table 8: Decomposing the Role of Idea Production Function vs. Spillover Changes

	Early	Late	IP only	Change
Inventor wage w	0.724	0.878	0.898	24.0
Prod. cutoff (exit)	1.004	1.241	1.270	26.4
Prod. cutoff (team)	1.580	1.740	1.634	3.4
Share of solo-based firms	93.7	87.2	78.5	-15.2
Share of team patents	57.8	90.4	93.3	35.5
Employment share of firms with $n > 145$	32.8	78.9	79.7	46.9
Average firm size	7.1	27.1	31.2	339.5
Spillovers ζ	0.903	0.827	0.833	-7.7
Output (normalized)	1.000	1.213	1.240	24.0

Notes: This table reports the counterfactual exercises in which we vary the three parameters (z, x, γ) in idea production function while holding all other parameters fixed, including β, at their early-period values. The first two columns are the baseline model moments for the early and late periods, respectively. The third column presents the model-implied value in a counterfactual when only the three parameters are adjusted to its late-period estimate, and the last column reports the corresponding percent or percentage change relative to the early-period benchmark.

a decline in the employment share of large firms. The increase in the right tail of team productivity draws (higher γ), however, generates a substantially larger rise in average firm size and the employment share of large firms. The change in employment concentration between the early and late periods in the model is entirely driven by the idea production parameters (Table 8).

Innovative output and knowledge spillovers. The rise of solo productivity increases innovative output but has only minor effects on knowledge spillovers, while the decrease in the minimum team productivity reduces innovative output slightly with almost no effects on knowledge spillovers. The rise in solo productivity raises knowledge spillovers by reallocating inventors toward smaller firms. On the other hand, the increase in the returns to scale benefit of using teams (higher γ) raises output substantially but reduces knowledge spillovers significantly since inventors are reallocated toward larger firms. In Table 8, changes in the idea production parameters account for 92% of the decline in knowledge spillovers (suggesting the decline in β is responsible for the remaining 8%).

Table 9: Competitive Equilibrium vs. Social Planner (CE to SPP)

Variable	Early	Late	IP Only
Total output	-0.20%	-1.72%	-1.25%
Spillovers (ζ)	-0.60%	-6.04%	-5.58%
Emp. share of firms with $n > 145$	7.43 pp	21.60 pp	22.35 pp

Notes: This table reports the deviation of the competitive equilibrium (CE) from the social planner (SPP) for aggregate innovative output, knowledge spillovers, and the employment share of large firms. The first two columns report the early- and late-period economies. The final column reports a counterfactual in which only idea production function parameters change from their early-period values, while all other parameters, including β, remain fixed at its early-period value.

6.4 Effects on Misallocation

Finally, we compare the competitive equilibrium allocations in the early and late periods with the corresponding social optimum given the different idea production and knowledge spillover functions in the two periods. To do this, the planner's problem in Section 4 is extended to allow the planner to choose the mass of firms, firm entry and exit, and assign firms to use solo or team production. The full set-up is outlined in Appendix B.6. The parameter values are the same as in Table 4. We first compare aggregate outcomes and then explore the mechanisms.

Aggregate outcomes. Table 9 reports the aggregate results, measuring the deviation of the competitive equilibrium from the social optimum in aggregate innovative output, knowledge spillovers, and the employment share of large firms in each period.

The central result is that the shift to team production substantially increases the efficiency costs of inventor employment concentration. The competitive equilibrium underproduces innovative output by only 0.2% in the early period but by 1.7% in the late period, largely because the shift toward team production amplifies the distortion created by uninternalized knowledge spillovers. This result is robust to the alternative calibration in which the late-period value of γ is chosen to match the observed share of team patents, though in that version the output gap rises to 1.5% instead of 1.7%. Appendix B.5.1 provides more details.

The changes in the idea production function alone (absent the change in the spillover

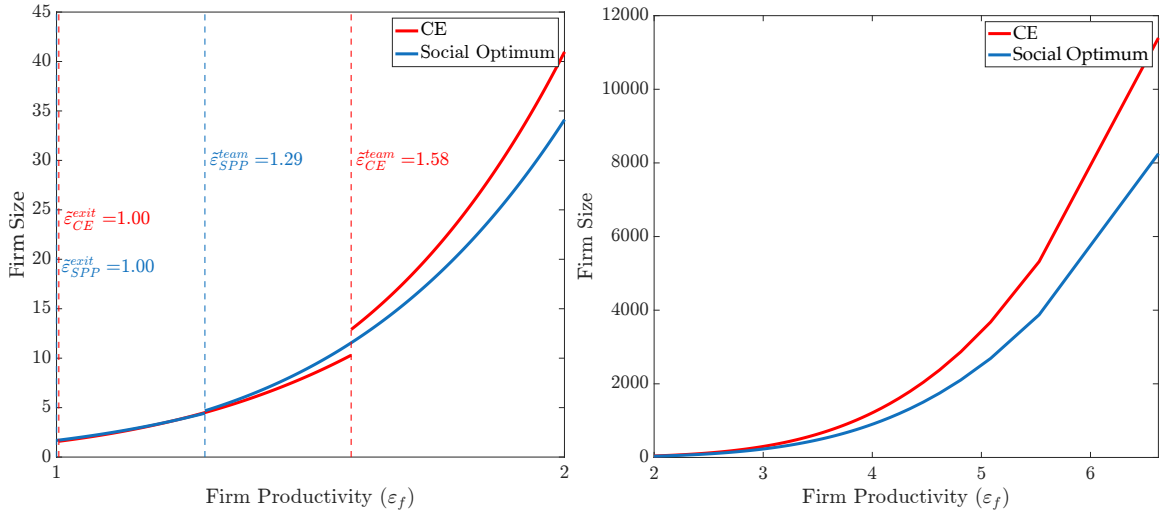
function parameter β) increase the output gap by seven times its early value, suggesting that they are responsible for around 73% of the change in the output gap between the competitive equilibrium and the planner's allocation (Table 9, column 3). The endogenous decline in knowledge spillovers in the competitive equilibrium is a major driver of this widening gap: the difference between the planner's realized spillovers and the competitive equilibrium's spillovers grows by a factor of ten in the late period.

The planner improves aggregate output by internalizing knowledge spillovers, retaining a larger mass of firms and assigning more low-productivity firms to team production. In contrast, higher equilibrium wages in the competitive equilibrium induce excessive exit and discourage low-productivity firms from expanding enough to adopt team production. Although the planner also increases employment concentration in the late period in response to changes in the idea production function that favor teams, the increase is substantially smaller than in the competitive equilibrium.

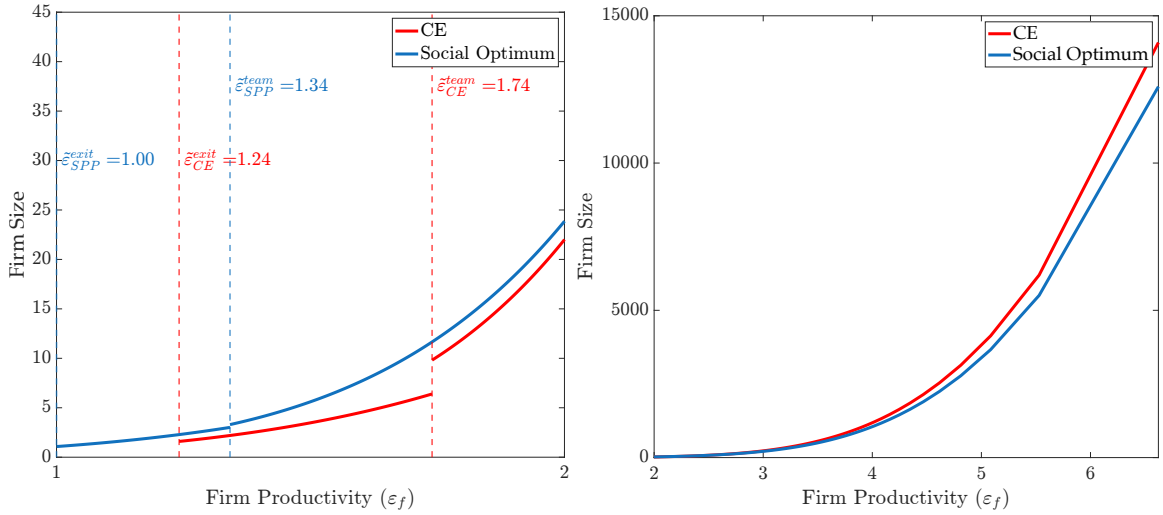
Firm size decision rules. Figure 4 compares firm-size policies under the planner and the competitive equilibrium. The planner reallocates inventors from large firms to small firms and adopts team production at lower productivity levels. The planner sustains a larger mass of active firms through a lower exit cutoff. Together these adjustments increase knowledge spillovers and mitigate employment concentration.

Firm-level wedges. Figure 5 translates the size differences in Figure 4 into the firm-level wedges τ_f, defined in (4.3) as the gap between the social marginal product and private marginal revenue product of inventor labor, evaluated at the competitive equilibrium allocation. Allowing an endogenous choice between solo and team production creates a productivity cutoff for team production in both the social planner's allocation and the competitive equilibrium. The planner also chooses a productivity threshold for exit. We evaluate the wedges accordingly, taking into account the optimal regime of production under each case.[24]

[24]Specifically, as in Figure 4, for firms below the social planner's team cutoff $\tilde{\varepsilon}^{team}_{SPP}$, both allocations involve solo production. For firms between the planner $\tilde{\varepsilon}^{team}_{SPP}$ and competitive equilibrium team cutoffs $\tilde{\varepsilon}^{team}_{CE}$, the planner adopts team production while the competitive equilibrium remains in solo produc-



(a) Early Period



(b) Late Period

Figure 4: Firm Size Decision Rule: CE vs. SPP in the Early and Late Periods

Note: These figures illustrate firm-size decision rules under the competitive equilibrium and the social planner's allocation for the early period (top panel) and the late period (bottom panel). In each panel, the left subfigure focuses on the low-productivity (small-firm) region, while the right subfigure focuses on the high-productivity (large-firm) region to highlight the underlying patterns clearly. The red lines correspond to the competitive equilibrium, and the blue lines indicate the planner's equilibrium. The dashed vertical lines denote the two productivity cutoffs for adopting team production and firm exit in each equilibrium for each period.

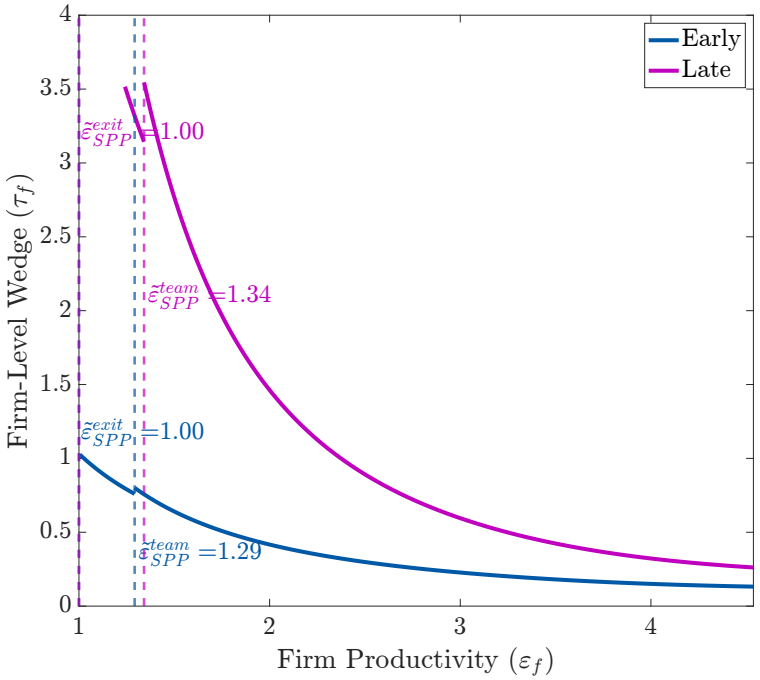


Figure 5: Firm-level Wedges: CE vs. SPP in the Early and Late Periods

Note: The figure illustrates the firm-level wedges τ_f between the competitive equilibrium and the social optimum, defined in equation (4.3), as a function of firm productivity ε_f. The blue lines correspond to the early period, and the purple lines correspond to the late period. The dashed vertical lines denote the two productivity cutoffs in the planner's problem for adopting team production and firm exit in each period.

Comparing wedges across the two periods, the level of wedges is uniformly higher in the late period (purple line) across the entire productivity distribution. In both periods, the wedges are decreasing in firm productivity (and therefore in firm size), consistent with the analytical results in Propositions 1 and 2, and the wedges are steeper with respect to productivity in the late period, consistent with Proposition 3. Implementation of the first best would therefore involve larger R&D subsidies to small firms and greater entry subsidies in the late period.

tion. For firms above the competitive equilibrium cutoff $\tilde{\varepsilon}^{team}_{CE}$, both allocations involve team production. Furthermore, there are no active firms below the competitive equilibrium exit cutoff, $\tilde{\varepsilon}^{\text{exit}}_{CE}$, unlike in the planner's allocation, so we cannot evaluate wedges for these firms.

7 Conclusion

This paper studies how the rise of team-based idea production interacts with firm boundaries and shapes the allocation of innovative labor across firms. We document a dramatic shift toward team production in U.S. patenting and a concurrent rise in inventor employment concentration at large firms. Using employment history data from Revelio Labs, we provide new evidence that firms play a central role in determining inventors' collaborators and ultimately the realized set of innovative teams.

To interpret these patterns, we develop a new model in which team formation occurs within firms and knowledge spillovers operate across firms. In the model, firms that use team production have endogenously higher returns to scale in inventor labor. We estimate the model's idea production function using U.S. patent data and show that team production increases returns to scale in inventor labor, consistent with the model. We further find that the idea production function has changed over time to increasingly favor teams and offer a potential explanation based on increased knowledge specialization of inventors. In the model, these shifts in the idea production function increase the concentration of inventor employment, while endogenously reducing knowledge spillovers across firms and increasing the misallocation of innovative labor.

The welfare implications of rising team production therefore depend critically on firm boundaries and the structure of knowledge spillovers. One promising direction for future research is to examine how this shift shapes the long-run dynamics of knowledge accumulation for R&D workers.

References

Acabbi, Edoardo Maria, Andrea Alati, Luca Mazzone, and Marta Morazzoni. 2026. "Sorting into Entrepreneurial Teams."

Akcigit, Ufuk, and Sina T. Ates. 2023. "What Happened to US Business Dynamism?" *Journal of Political Economy* 131 (8): 2059–2124.

Akcigit, Ufuk, Santiago Caicedo, Ernest Miguelez, Stefanie Stantcheva, and Valerio Sterzi. 2018, March. "Dancing with the Stars: Innovation Through Interactions." Working paper 24466, National Bureau of Economic Research.

Akcigit, Ufuk, and Nathan Goldschlag. 2023, March. "Where Have All the Creative Talents Gone? Employment Dynamics of US Inventors." Working paper 31085, National Bureau of Economic Research.

Andrews, Dan, Chiara Criscuolo, and Peter Gal. 2016. "The Global Productivity Slowdown, Technology Divergence and Public Policy: A Firm Level Perspective." *The Future of Productivity: Main Background Papers*, pp. 1–50.

Antràs, Pol, Luis Garicano, and Esteban Rossi-Hansberg. 2006. "Offshoring in a Knowledge Economy." *Quarterly Journal of Economics* 121:31–77.

Argente, David, Sara Moreira, Ezra Oberfeld, and Venky Venkateswaran. 2025. Scalable Expertise: How Standardization Drives Scale and Scope.

Autor, David, David Dorn, Gordon H Hanson, Gary Pisano, and Pian Shu. 2020a. "Foreign competition and domestic innovation: Evidence from US patents." *American Economic Review: Insights* 2 (3): 357–374.

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2020b. "The Fall of the Labor Share and the Rise of Superstar Firms." *Quarterly Journal of Economics* 135 (2): 645–709.

Babalievsky, Fil. 2023. "Misallocation in the Market for Inventors." pp. 1–45.

Bassi, Vittorio, Jung Hyuk Lee, Alessandra Peter, Tommaso Porzio, Ritwika Sen, and Esau Tugume. 2023. "Self-employment within the firm." Technical Report, National Bureau of Economic Research.

Berkes, Enrico, and Ruben Gaetani. 2026. The Decline in the Transmission of Scientific Ideas.

Bhaskarabhatla, Ajay, Luis Cabral, Deepak Hegde, and Thomas Peeters. 2021. "Are Inventors or Firms the Engines of Innovation?" *Management Science* 67 (6): 3899–3920.

Bloom, Nicholas, Charles Jones, John Van Reenen, and Michael Webb. 2020. "Are Ideas Getting Harder to Find?" *American Economic Review* 110 (4): 1104–1144.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen. 2013. "Identifying Technology Spillovers and Product Market Rivalry." *Econometrica* 81 (4): 1347–1393.

Boerma, Job, Aleh Tsyvinski, and Alexander P. Zimin. 2025. "Sorting with Teams." *Journal of Political Economy* 133 (2): 421–454.

Braguinsky, Serguey, Joonkyu Choi, Yuheng Ding, Karam Jo, and Seula Kim. 2023. "Mega firms and recent trends in the us innovation: Empirical evidence from the us patent data." Technical Report, National Bureau of Economic Research.

Cavenaile, Laurent, Murat Alp Celik, and Xu Tian. 2025. "Are Markups Too High? Competition, Strategic Innovation, and Industry Dynamics." *Review of Economic Studies, forthcoming*.

Chan, Mons, Guangbin Hong, Joachim Hubmer, Serdar Ozkan, and Sergio Salgado. 2025. Scalable versus Productive Technologies.

Chen, Zhang. 2025. "Economic Growth and the Rise of Large Firms." *Econometrica, forthcoming*.

Chikis, Craig, Benny Kleinman, and Marta Prato. 2026. The Geography of Innovative Firms.

Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda. 2014. "The role of entrepreneurship in US job creation and economic dynamism." *Journal of Economic Perspectives* 28 (3): 3–24.

Decker, Ryan A, John Haltiwanger, Ron S Jarmin, and Javier Miranda. 2016. "Where has all the skewness gone? The decline in high-growth (young) firms in the US." *European Economic Review* 86:4–23.

Ding, Yuheng, Karam Jo, and Seula Kim. 2022. "Improving Patent Assignee-Firm Bridge with Web Search Results." Working papers, U.S. Census Bureau, Center for Economic Studies.

Dyvere, Arnaud. 2024. "Public RD Spillovers and Productivity Growth."

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

Engbom, Nik, Hannes Malmberg, Tommaso Porzio, Federico Rossi, and Todd Schoellman. 2025. Economic Development According to Chandler.

Fernandez-Villaverde, Jesus, Yang Yu, and Francesco Zanetti. 2026. "Defensive Hiring and Creative Destruction."

Freund, Lukas. 2025. "Superstar Teams." Technical Report, Cambridge University.

Garicano, Luis, and Esteban Rossi-Hansberg. 2015. "Knowledge-Based Hierarchies: Using Organizations to Understand the Economy." *Annual Review of Economics* 7 (Volume 7, 2015): 1–30.

Greenstone, Michael, Richard Hornbeck, and Enrico Moretti. 2010. "Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings." *Journal of Political Economy* 118 (3): 536–598.

Guccione, Andrea, and Pau Roldan-Blanco. 2026. "Economic Growth when Knowledge is Concentrated." UAB.

Herkenhoff, Kyle, Jeremy Lise, Guido Menzio, and Gordon Phillips. 2024. "Production and Learning in Teams." *Econometrica*.

Hopenhayn, Hugo A. 1992. "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica* 60 (5): 1127–1150.

Hsieh, Chang-Tai, and Peter J. Klenow. 2009. "Misallocation and Manufacturing TFP in China and India*." *The Quarterly Journal of Economics* 124 (4): 1403–1448 (11).

Jaravel, Xavier, Neviana Petkova, and Alex Bell. 2018. "Team-Specific Capital and Innovation." *American Economic Review* 108 (4-5): 1034–73 (April).

Jarosch, Gregor, Ezra Oberfield, and Esteban Rossi-Hansberg. 2021. "Learning From Coworkers." *Econometrica* 89 (2): 647–676.

Jones, Benjamin F. 2009. "The Burden of Knowledge and the 'Death of the Renaissance Man': Is Innovation Getting Harder?" *Review of Economic Studies*.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. "Technological Innovation, Resource Allocation, and Growth." *Quarterly Journal of Economics*.

Liu, Jingnan Jane. 2023. "Worker Mobility, Knowledge Diffusion, and Non-Compete Contracts." Technical Report, UW-Madison.

Lucas, Robert E. 1978. "On the Size Distribution of Business Firms." *The Bell Journal of Economics* 9 (2): 508–523.

Lucas, Robert E., and Benjamin Moll. 2014. "Knowledge Growth and the Allocation of Time." *Journal of Political Economy* 122 (1): 1–51.

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

Melitz, Marc. 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71 (6): 1695–1725.

Moretti, Enrico. 2021. "The Effect of High-Tech Clusters on the Productivity of Top Inventors." *American Economic Review* 111 (10): 3328–75 (October).

Olmstead-Rumsey, Jane. 2026. "Market Concentration and the Productivity Slowdown."

Pearce, Jeremy. 2023. "Idea Production and Team Structure." Technical Report, University of Chicago.

Peters, Ryan H., and Lucian A. Taylor. 2017. "Intangible Capital and the Investment-q Relation." *Journal of Financial Economics* 123 (2): 251–272.

Prato, Marta. 2025. "The Global Race for Talent: Brain Drain, Knowledge Transfer, and Growth." *The Quarterly Journal of Economics* 140 (1): 165–238 (02).

Rossi-Hansberg, Esteban, Pierre-Daniel Sarte, and Nicholas Trachter. 2020. "Diverging Trends in National and Local Concentration." *NBER Macroeconomics Annual*, vol. 35.

Sterk, Vincent, Petr Sedláček, and Benjamin Pugsley. 2021. "The nature of firm growth." *American Economic Review* 111 (2): 547–579.

Wuchty, Stefan, Benjamin F. Jones, and Brian Uzzi. 2007. "The Increasing Dominance of Teams in Production of Knowledge." *Science* 316 (5827): 1036–1039.

A Empirical Appendix

A.1 Probability of Inventors Working at the Same Firm

We use patent data and Revelio Labs data to measure how often inventors listed on the same team patent are employed by the same firm around the time of invention. Specifically, we identify all pairs of inventors listed on each team patent and construct a monthly panel spanning five years before and after the patent filing date. For each inventor pair and month, we define an indicator equal to one if both inventors are employed by the same firm in that month and zero otherwise. We extract at most two firms per inventor per month from the Revelio Labs data.

Based on this indicator, we compute the probability that a pair of inventors is ever employed by the same firm over a five-year window before and after the patent filing date. This probability is 87.5%. If we restrict attention to teams working together up to five years before the filing date, the value is 85.9%. Furthermore, in Figure A1, we plot the average value of this indicator by month relative to the patent filing date. The probability that the two inventors are employed by the same firm peaks at approximately 72% twelve months before the patent filing date and declines steadily thereafter, reaching approximately 45% three years after the filing date. This pattern suggests that inventor teams are more likely to be formed within the same firm as the invention approaches and frequently separate following the completion of a joint invention.

A.2 Empirical Support for Modeling Assumptions

One model assumption that deserves additional discussion is the assumption that there is no serial collaboration or team-specific human capital. Our model is a static, one period model, so there is no natural way of modeling accumulation of team-specific capital. Moreover, somewhat surprisingly, patenting teams are quite short lived in the data: 80% of teams in our sample only patent together once. Even for pairs of inventors

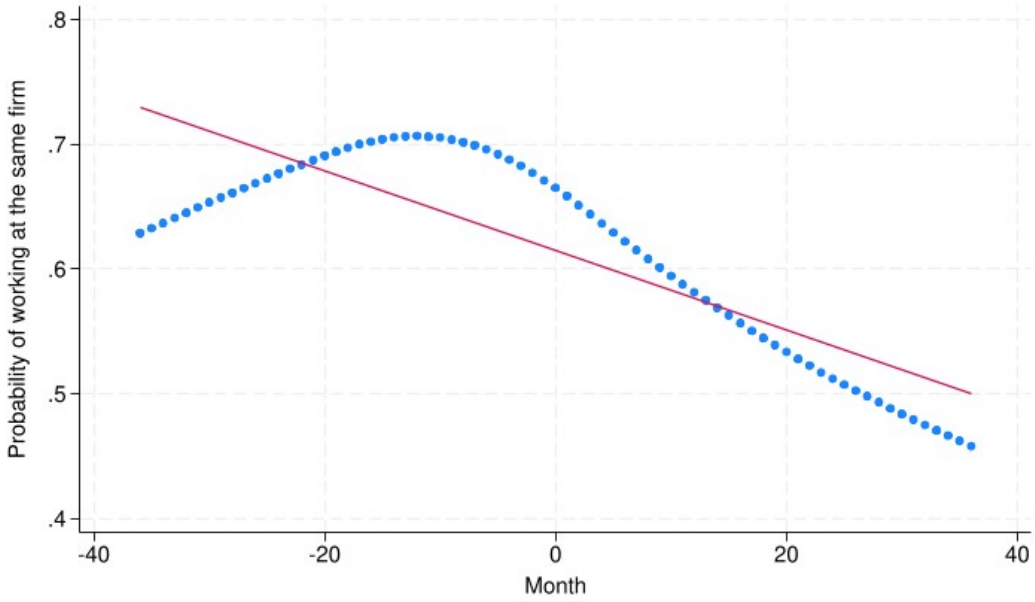


Figure A1: Probability that Inventor Pairs Work at the Same Firm around Patent Filing
Note: The figure illustrates the probability for a pair of inventors on a patent working at the same firm 36 months before and after the filing month. $t = 0$ corresponds to the filing month of the patent.

within larger teams, 65% of pairs never collaborate again.

A second assumption is that inventors only work in teams of two. This assumption is stylized, since we also observe many teams of three to five inventors in the data (teams of two constitute 37% of team patents, and teams of less than 6 team members together constitute 80%). Importantly, the number of inventors per patent is not increasing one for one in the number of inventors at the firm (that is, firms with many inventors still divide them into smaller teams rather than pooling them into a single large team), leveling off at around three inventors (Figure A2). Section B.1 considers a richer model with teams of three.

In partial equilibrium the model predicts that output per worker is increasing in the number of coworkers. Figure A3, which plots patenting output per worker as a function of the number of (successful) inventors, shows that this is the case separately for two subperiods using USPTO data. Consistent with our finding that returns to scale increase, there is a level shift downward in output per worker in the later period. One

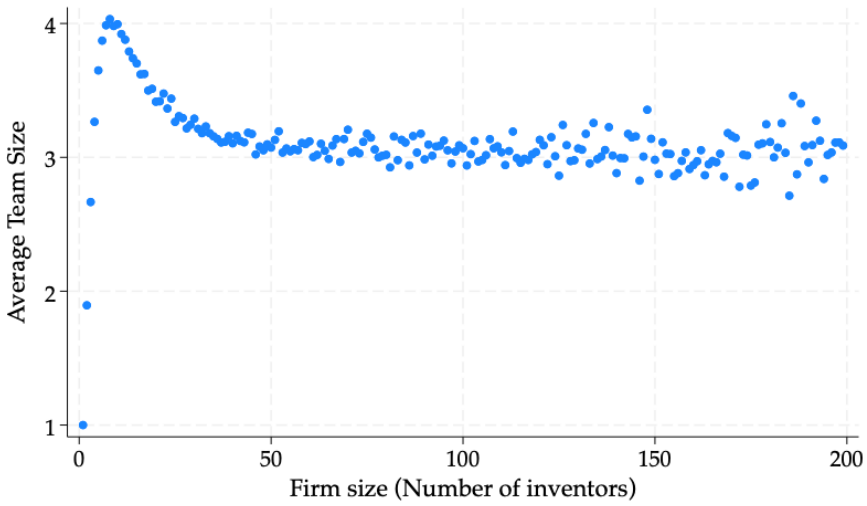


Figure A2: Average Team Size By Number of Inventors

Note: The figure illustrates the average team size within firms, as a function of the total number of patenting inventors at the firm in that year pooling all years in USPTO, 1976-2018.

concern with the figure is that output per worker looks highest for firms with one inventor. Recall that our measure of inventors is the number of successfully patenting inventors, which mechanically ensures that output per worker for single-inventor firms is at least one. Figure A4 corrects for this using employment histories of inventors from Revelio Labs to get a more accurate measure of inventors per firm (2008-2018), enabling us to include firm-year pairs with inventors but no patents, and shows that output per worker at single inventor firms is actually quite low. Otherwise the pattern is qualitatively similar to the USPTO figure.

Finally, the model predicts that larger firms will have a higher share of team patents in total patents. Defining firm size by the number of inventors, we estimate the following regression to examine the relationship between firm size and the propensity to engage in team patenting:

$$TeamPatShare_{ft} = \eta_f + \gamma_t + \log(Num.Inventors)_{ft} + \epsilon_{ft} \tag{A.1}$$

where $TeamPatShare_{ft}$ is the share of team patents out of total patents assigned to firm f in year t; $\log(Num.Inventors)_{ft}$ is the log number of inventors within firm f in

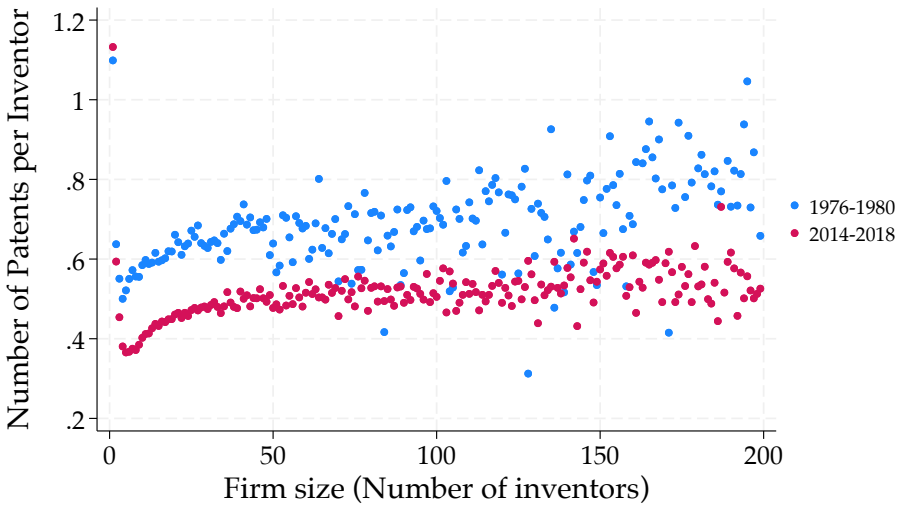


Figure A3: Number of patents per worker as a function of firm size in the USPTO data for two subperiods: 1976-1980 and 2014-2018.

Note: The figure illustrates the number of patents per worker within firms, as a function of the total number of patenting inventors at the firm in that year. The blue dots pool all years in the early period, 1976-1980, and the red dots pool all years in the late period, 2014-2018.

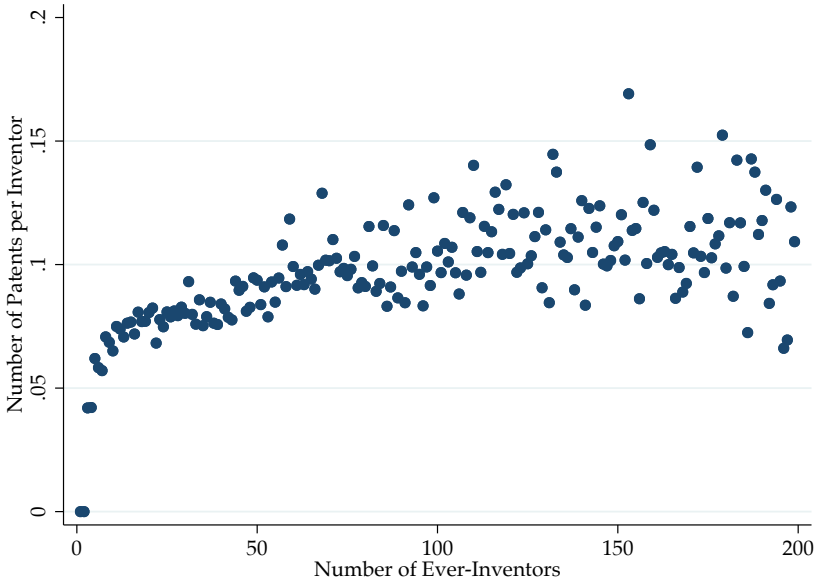


Figure A4: Number of patents per worker as a function of firm size in the Revelio Labs data, 2008-2018.

Note: The figure illustrates the number of patents per worker within firms, as a function of the total number of inventors at the firm in that year. The data is sourced from the Revelio Labs data, pooling all years between 2008 and 2018.

t; and η_f and γ_t are firm and year fixed effects, respectively.

Table A1 presents the results, where the first column uses all USPTO assignees, and the second column restricts the sample to public firms in Compustat. Across both samples, we find that larger firms exhibit a higher share of team patents than smaller firms. The effect is economically meaningful: a 10% increase in the number of inventors within a firm is associated with a 2-percentage-point increase in the share of team patents for USPTO firms, and a 1.2-percentage-point for public firms.

Table A1: Team Patent Share and Firm Size

	(1) All Firms	(2) Public Firms
Log num. inventors	0.204***	0.122***
	(282.52)	(22.46)
N of Obs.	717,111	74,501
Firm FE	Yes	Yes
Year FE	Yes	Yes
Adj. R^2	0.491	0.375

Note: The table presents the regression result from equation (A.1), where the dependent variable is the share of team patents within a firm. The coefficient of interest is on the log number of inventors in the firm. Firm and year fixed effects are included, respectively. Column (1) covers all USPTO assignees, and Column (2) includes Compustat firms only. *** p<0.01, ** p<0.05, * p<0.1

A.3 Team Diversity and Inventor Specialization

Another salient trend over the past several decades is rising knowledge diversity within teams alongside increasing inventor-level specialization.

To measure knowledge diversity, we first identify the "core" CPC group and subgroup associated with an inventor in a given year as a proxy for the inventor's type or area of expertise.[25,26] The core CPC group/subgroup of an inventor is defined as the group/-subgroup with the highest number of patents filed by that inventor up to that year. For

[25]For example, G06N 20/00 is "Machine learning" or A23B 11/00 is "Preservation of milk or dairy products."

[26]For a patent associated with multiple CPC groups/subgroups, we pick the main CPC groups/subgroups based on the sequence variable provided by the USPTO.

each team, we calculate the number of core CPC groups/subgroups associated with its inventors, which serves as a measure of the team's knowledge diversity. We then average this knowledge diversity measure across all teams within a firm for that year. Finally, we compute the average of these firm-level measures across all firms in a given year. The results are plotted in Figures A5 and A6, which show an upward trend in the likelihood that teams include inventors with different technological specialties.

To capture inventor-level specialization, we measure the propensity to move across core CPC groups over one- and five-year horizons. For each inventor-year, we construct an indicator equal to one if the inventor's current core CPC group differs from a core CPC group observed during the preceding one or five years, respectively. We then average this indicator across inventors in each year. Figures A7 and A8 show a downward trend in jump intensity, indicating that inventors have become less likely to diversify across technological fields and increasingly specialized in their core fields over time. Consistent with this pattern, Figure A9 reports the average number of new CPC groups that have not appeared in an inventor's prior patents. This measure also declines over time, suggesting that inventors have become less likely to enter new technological fields.

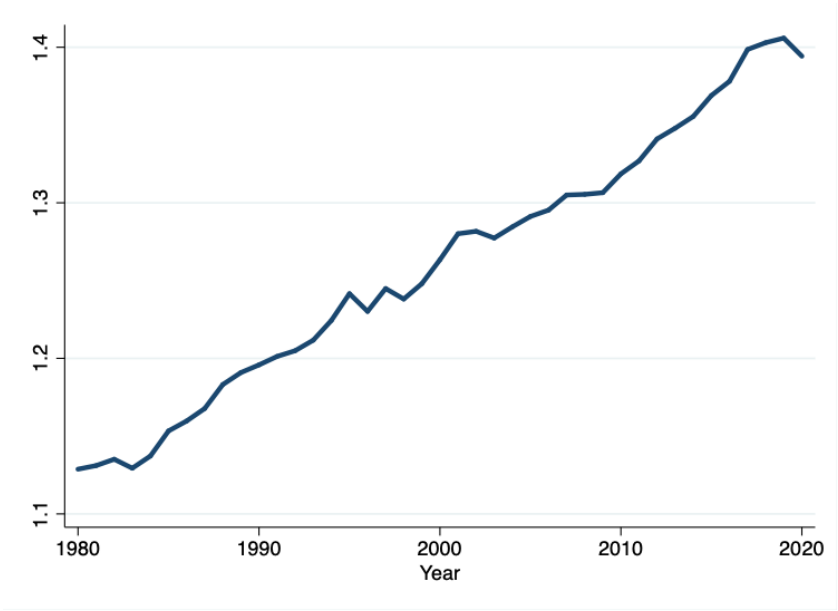


Figure A5: Average Number of Inventor CPC Groups in Teams in a Firm

Note: The figure illustrates the average number of core CPC groups associated with inventors on team patents in USPTO. A core CPC group is defined as the modal CPC group in an inventor's patenting history up to a given year. For each firm, we calculate the average number of core CPC groups for inventors on team patents and then compute the mean across all firms in a given year.

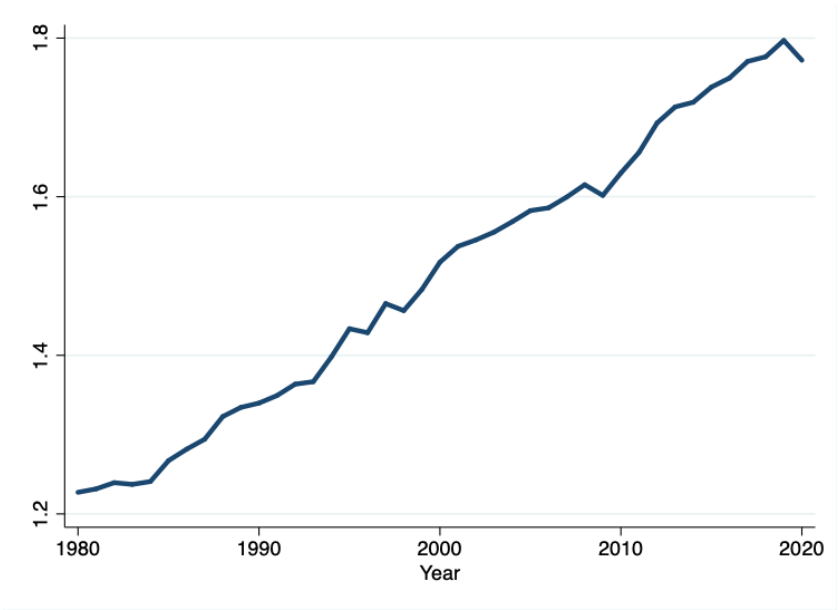


Figure A6: Average Number of Inventor CPC Subgroups in Teams in a Firm

Note: The figure illustrates the average number of core CPC subgroups associated with inventors on team patents in USPTO. A core CPC subgroup is defined as the modal CPC subgroup in an inventor's patenting history up to a given year. For each firm, we calculate the average number of core CPC subgroups for inventors on team patents and then compute the mean across all firms in a given year.

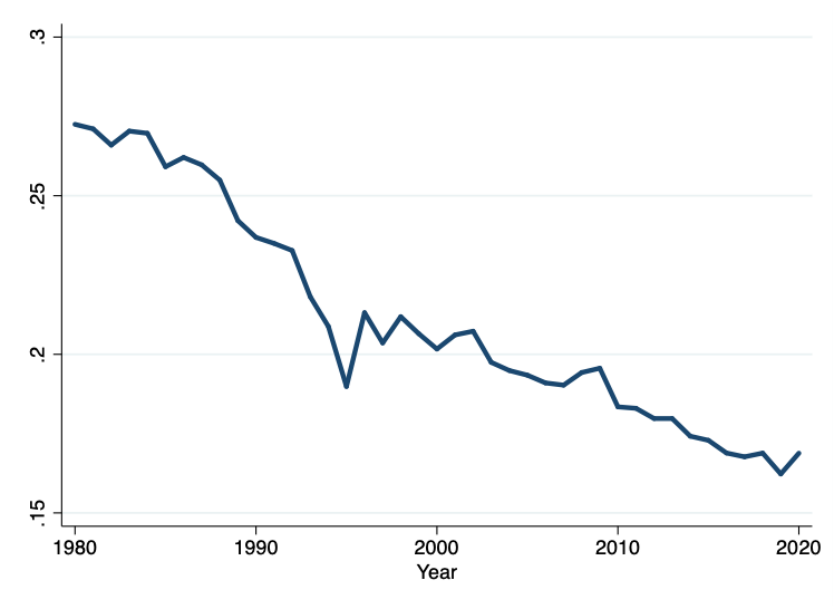


Figure A7: Average Jump Intensity in Inventors' Core CPC Groups (1-year window)

Note: The figure reports the average jump intensity in inventors' core CPC groups based on patenting activity in the previous year in the USPTO data. An inventor's core CPC group is defined as the modal CPC group in the inventor's patenting history up to a given year. For each inventor-year, the jump indicator equals one if the current core CPC group differs from the core CPC group observed in the previous year. The figure reports the annual mean of this indicator across inventors.

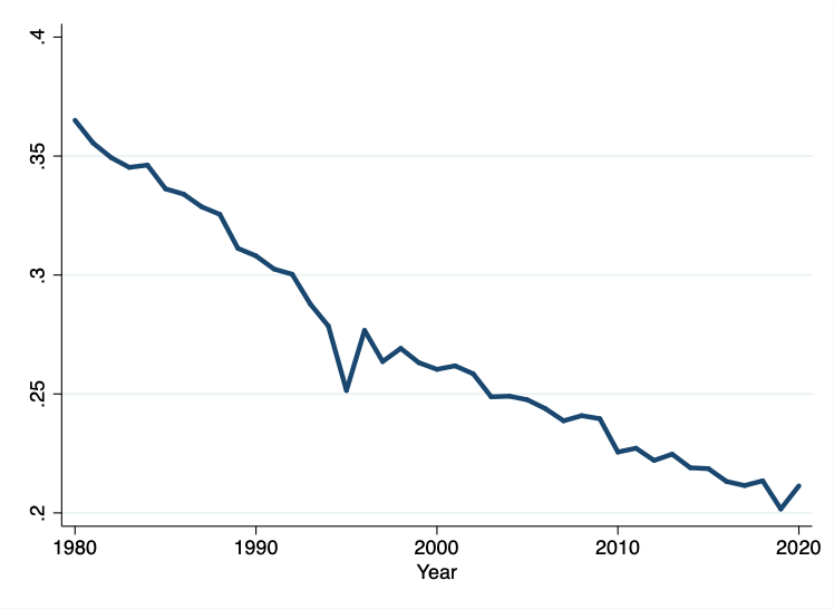


Figure A8: Average Jump Intensity in Inventors' Core CPC Groups (5-year window)

Note: The figure reports the average jump intensity in inventors' core CPC groups based on their patenting activity over the previous five years in the USPTO data. An inventor's core CPC group is defined as the modal CPC group in the inventor's patenting history up to a given year. For each inventor-year, the jump indicator equals one if the current core CPC group differs from at least one core CPC group observed for the inventor during the previous five years. The figure reports the annual mean of this indicator across inventors.

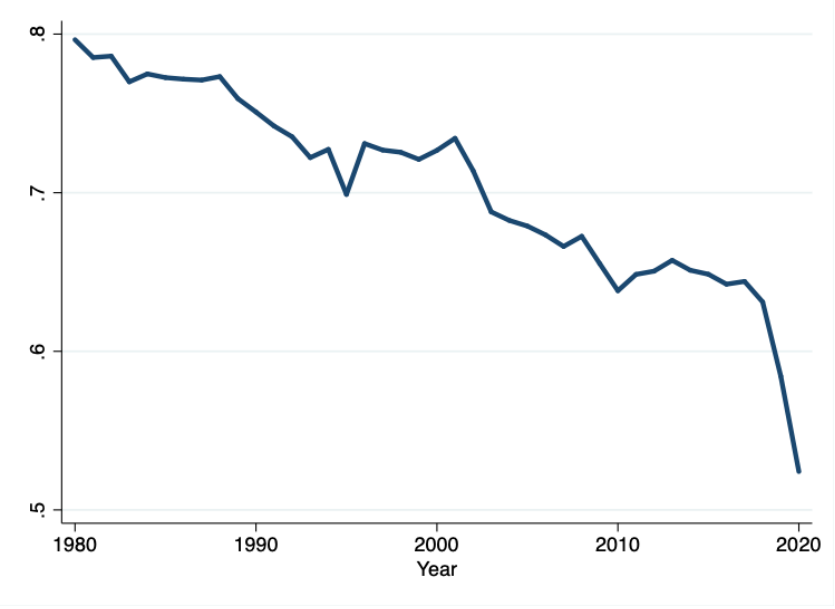


Figure A9: Average Number of New CPC Groups in Inventor Patent Portfolio

Note: The figure reports the average number of new CPC groups added to inventors' patent portfolios in each application year. For each inventor-year, we count the CPC groups that have not appeared in the inventor's prior patents and then average this count across inventors.

A.4 Robustness for Idea Production Function Estimation

Table A2: Returns to Scale in Team Production: Citations

	(1) All	(2) Public Firms
$\log(n_{ft})$	0.575***	0.686***
	(0.003)	(0.006)
$\log(n_{ft}) \times \text{Team}_{ft}$	0.001	0.041***
	(0.004)	(0.007)
$\log(n_{ft}) \times \text{Team}_{ft} \times \text{Post}_t$	0.022***	0.003
	(0.005)	(0.008)
Team_{ft}	-0.429***	-0.472***
	(0.005)	(0.011)
Post_t	-0.095***	-0.105***
	(0.003)	(0.012)
$\log(n_{ft}) \times \text{Post}_t$	0.009***	0.028***
	(0.003)	(0.005)
$\text{Team}_{ft} \times \text{Post}_t$	-0.042***	-0.014
	(0.006)	(0.015)
$\log(K_{ft}^{intan.})$		-0.040***
		(0.003)
Constant	0.613***	0.774***
	(0.003)	(0.021)
R-squared	0.452	0.616
Obs.	1105437	131166
Firm FE	Yes	Yes

Notes: Column (1) includes all firm–year observations. Column (2) includes public firms with available intangible capital data from Peters and Taylor (2017). Outcome is log of 5-year forward citations, purged of technology-class fixed effects, plus one. See equation (5.1) for details. Standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A3: Returns to Scale in Team Production: Revelio Labs Data

	(1) All	(2) $n > 1$
$\log(n_{ft})$	0.353***	0.457***
	(0.010)	(0.012)
$\log(n_{ft}) \times \text{Team}_{ft}$	0.169***	0.161***
	(0.007)	(0.009)
$\log(n_{ft}) \times \text{Team}_{ft} \times \text{Post}_t$	0.041***	0.055***
	(0.010)	(0.014)
Team_{ft}	-0.071***	-0.096***
	(0.014)	(0.019)
Post_t	0.006	0.003
	(0.013)	(0.025)
$\log(n_{ft}) \times \text{Post}_t$	0.019**	0.009
	(0.009)	(0.014)
$\text{Team}_{ft} \times \text{Post}_t$	0.023	-0.009
	(0.018)	(0.027)
Constant	0.026*	-0.254***
	(0.016)	(0.025)
R-squared	0.244	0.275
Obs.	169,800	135,406
Firm FE	Yes	Yes

Notes: Column (1) includes all firm–year observations. Column (2) restricts the sample to firm–year observations with at least two inventors. Outcome is number of patents. Data is Revelio Labs, 2008-2018, post-period is 2014-2018. See equation (5.1) for details. Standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.5 Inventor Concentration and Team Patent Share

We examine the relationship between inventor concentration and team innovation using patent data. We construct these measures at the CPC technology group level and estimate the following regression:

$$Concentration_{ct} = \beta TeamPatentShare_{ct} + X'_{ct}\gamma + \mu_c + \delta_t + \epsilon_{ct}, \tag{A.2}$$

Table A4: Inventor Concentration and Team Patenting

	(1)	(2)	(3)	(4)	(5)	(6)
Team patent share	0.311***	0.212***	0.260***	0.187***	0.0584***	0.0541***
	(0.00739)	(0.00726)	(0.0208)	(0.0177)	(0.0107)	(0.0103)
log(firms)		-0.247***		-0.246***		-0.151***
		(0.00426)		(0.0150)		(0.0131)
log(patents)		0.254***		0.253***		0.148***
		(0.00347)		(0.0125)		(0.0114)
Year FE	No	No	Yes	Yes	Yes	Yes
CPC FE	No	No	No	No	Yes	Yes
Controls	No	Yes	No	Yes	No	Yes
R^2	0.080	0.339	0.092	0.343	0.118	0.167
Obs.	25961	25961	25961	25961	25961	25961

Notes: The dependent variable is inventor concentration at the CPC technology-group level. Team patent share is the average fraction of patents that are produced by teams of inventors at the firm level within a CPC technology group. Controls include the log number of firms and the log number of patents in the technology class. Standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

where $Concentration_{ct}$ denotes the share of inventors at firms with more than 145 inventors in CPC technology group c and year t, $TeamPatentShare_{ct}$ is the average firm-level share of team patents, and X_{ct} is a vector of controls, including the log number of firms and the log number of patents, all measured at the CPC technology group–year level. μ_c and δ_t denote CPC technology-group and year fixed effects, respectively. Firms are included in a CPC technology group in a given year if they have at least one patent in that CPC technology group in that year.[27] The coefficient of interest, β, measures the association between inventor concentration and the prevalence of team-based innovation within a technology class.

Table A4 reports the estimation results. We find a positive and statistically significant association between inventor concentration and the share of team patents across all specifications. The estimated relationship is robust to the inclusion of controls, including year and CPC technology-class fixed effects, and is consistent with our model predictions.

[27] For patents with multiple assigned CPC technology groups, we use the main CPC group based on the sequence variable provided by the USPTO.

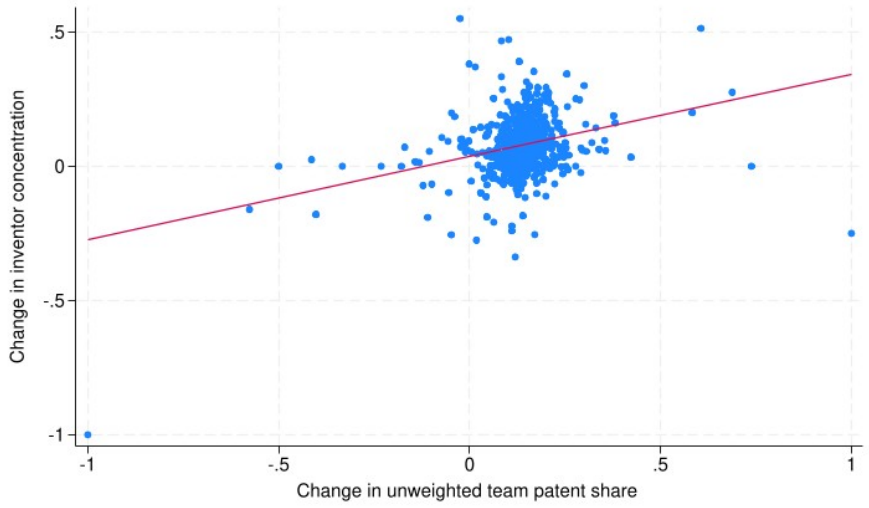


Figure A10: Inventor Concentration and Team Patent Share

Note: The figure illustrates the long-run changes in the average inventor concentration and the average firm-level team patent share between 1976-1999 and 2000-2018 across CPC technology groups in USPTO.

Furthermore, Figure A10 provides complementary evidence by examining changes in inventor concentration and team patenting across CPC technology groups over time. For each CPC technology group, we compute the average inventor concentration and the average firm-level team patent share separately for the periods 1976–1999 and 2000–2018. We then measure the change in each variable as the difference between its average in the later period and its average in the earlier period. The figure plots the change in inventor concentration against the corresponding change in the team patent share across CPC technology groups. The figure shows a positive relationship, indicating that technology groups experiencing larger increases in inventor concentration also tend to exhibit larger increases in the prevalence of team-based innovation.

B Model Appendix

B.1 Teams of Three

In the baseline model each pair of inventors within a firm gets a bilateral productivity draw and an inventor's productivity is the maximum over all their pairwise draws. With teams of three, every triple of inventors gets a draw, and an inventor's productivity is the maximum over all triples they belong to. With n workers, each worker belongs to $C(n-1,2)$ distinct triples ("choose 2 out of $n-1$ coworkers"). $C(n-1,2) = (n-1)(n-2)/2$, which grows like $\frac{n^2}{2}$ for large n. So instead of n^γ scaling, teams of three would deliver $(n^2)^\gamma = n^{2\gamma}$ scaling, making inventor productivity scale faster in the number of coworkers than in the baseline model.

B.2 Microfounded Team-Level Idea Production Function

This section shows how a version of the team production model in Freund (2025) generates stochastic team productivity with a distribution resembling the Pareto distribution specified in equation (3.1).

The environment follows Freund (2025), except that we abstract from talent heterogeneity across inventors. A pair of inventors, i and j, jointly produces a plan for an idea. Producing the plan requires a continuum of tasks of measure one, such as researching the state of the art, defining the problem, and outlining a solution. The quality of the resulting plan is given by

$$z_{ij} = \left(\int_0^1 q(\tau)^{\frac{\rho-1}{\rho}} d\tau \right)^{\frac{\rho}{\rho-1}},$$

where $q(\tau)$ denotes the amount of task input τ, and ρ governs the elasticity of substitution across tasks. Task input $q(\tau)$ depends on the labor input $\ell_k(\tau)$ and task-level productivity $z_k(\tau)$ of each inventor $k \in \{i,j\}$:

$$q(\tau) = z_i(\tau) l_i(\tau) + z_j(\tau) l_j(\tau).$$

Each worker k is endowed with one unit of time: $\int_0^1 \ell_k(\tau)d\tau = 1$.

For each worker k and task τ, talent $z_k(\tau)$ is drawn independently and identically from a Fréchet distribution:

$$P(z_k(\tau) < z) = exp\left(-\left(\frac{z}{\iota x}\right)^{-\frac{1}{\chi}}\right),$$

where x is a common scale parameter that shifts talent for all inventors.[28] The parameter χ governs skill specialization: a higher χ implies greater variation in worker skill across tasks. The term $\iota = \Gamma\left(1+\chi(1-\rho)\right)^{1/(1-\rho)}$ is a normalization that ensures that changes in χ do not mechanically affect average productivity. We assume $1+\chi(1-\rho) > 0$ throughout.

Freund (2025) shows that the output produced by any pair of workers can be characterized without information on task-level productivities and is instead given by:

$$z_{ij} \sim 2^{1+\chi\xi_{ij}}x, \tag{B.1}$$

where ξ_{ij} denotes the pairwise skill distance between workers i and j in Freund (2025) and is uniformly distributed, $\xi_{ij} \sim U[0,1]$, in the population of inventors.

The key insight is as follows. For a given skill distance ξ_{ij}, greater knowledge specialization, χ, increases the quality of the plan produced by the team because the gains from specialization across tasks are larger. This effect is amplified when inventors' skills are more distant from each other (that is, when ξ_{ij} is high). Talent, x, shifts the level of team output.

Figure B1 illustrates the distribution of team productivity z_{ij} using 50,000 draws of pairwise skill distance with x normalized to 1. The dashed lines show the corresponding fitted Pareto distributions, where the Pareto tail parameter is estimated using the Hill maximum likelihood estimator applied to the productivity data for each value of χ. Greater skill specialization, χ, implies a thicker right tail in the fitted Pareto distribution,

[28] Freund (2025) allows overall talent x to vary across individuals, but this heterogeneity is not needed for our results.

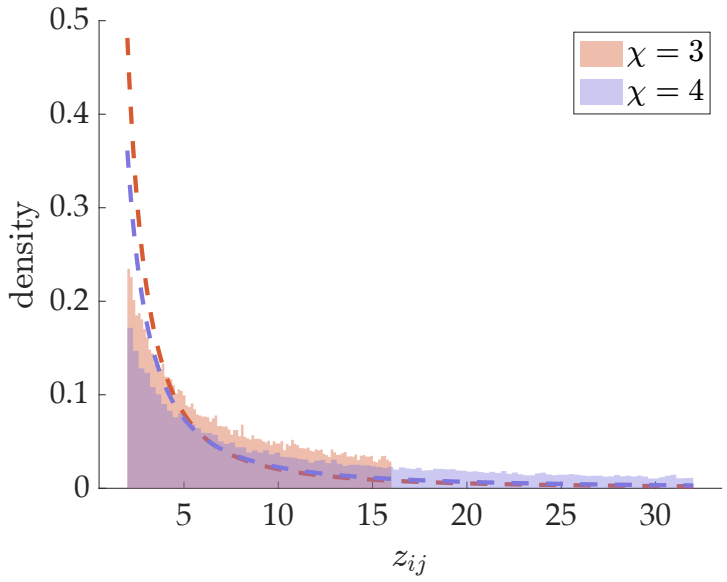


Figure B1: Distribution of Idea Quality z_{ij} and Fitted Pareto Distribution

Note: This figure illustrates the distribution of idea quality z_{ij} (bars) and fitted Pareto (dashed lines) under two different values of specialization χ. Greater specialization implies more dispersion of output for the same draws of worker skill distance ξ_{ij}.

analogous to an increase in γ in our baseline model.

B.3 Derivation of Team Productivity

B.3.1 Discrete Firm Size

Recall that bilateral productivity is:

$$z_{ij} \sim Pareto\left(x, \frac{1}{\gamma}\right): \quad P(z_{ij} \leq k) = 1 - \left(\frac{x}{k}\right)^{\frac{1}{\gamma}} \quad \forall\, k \geq x$$

Let the max of n draws be

$$M_n \equiv \max\{z_{1j}, z_{2j}, .., z_{nj}\}$$

then

$$F_{M_n}(k) \equiv P(M_n \leq k) = \left[1 - \left(\frac{x}{k}\right)^{\frac{1}{\gamma}}\right]^n \quad \forall k \geq x \tag{B.2}$$

Then, we can derive the pdf as follows:

$$f_{M_n}(k) = n\frac{1}{\gamma}\frac{x^{1/\gamma}}{k^{1+1/\gamma}}\Big[1-\Big(\frac{x}{k}\Big)^{\frac{1}{\gamma}}\Big]^{n-1} \quad \forall k \geq x.$$

Thus, the expected value is

$$\mathbb{E}(M_n) = \int_k k f_{M_n}(k)dk = xnB(n, 1-\gamma),$$

where B is Beta function.[29]

Note that the Beta function can be approximated by a power function

$$B(n, 1-\gamma) = \frac{\Gamma(n)\Gamma(1-\gamma)}{\Gamma(n+1-\gamma)} \simeq \Gamma(1-\gamma)n^{-1+\gamma},$$

where the relative error is given by

$$\frac{B(n, 1-\gamma)}{\Gamma(1-\gamma)n^{-1+\gamma}} - 1 = \frac{\gamma(1-\gamma)}{2n} + O(n^{-2}).$$

In our calibration, $\gamma < 0.1$, thus the leading term in this error does not exceed 4.5% for $n = 1$, and is only 0.63% for the calibrated average firm size $n = 7.1$. Thus, for our aim, the approximation is quite good and we will take it as equality in the following derivations for simplicity.

Therefore, the expected productivity of working in a firm with n workers becomes:

$$\mathbb{E}(M_n) = x\Gamma(1-\gamma)n^{\gamma}, \tag{B.3}$$

which increases in n.

[29] It can be derived with a transformation of variable with $t = 1-(x/k)^{\gamma}$. So that it can be rephrased as $\mathbb{E}(M_n) = nx\int_0^1 t^{n-1}(1-t)^{-\gamma}dt$.

B.3.2 Continuous Firm Size

Assume that worker bilateral productivity is drawn from an inhomogeneous Poisson point process with intensity function $\lambda(k)$ within firm of size $n \in [0, \infty)$,

$$\lambda(k) = n\frac{f(k)}{F(k)},$$

where $f(k)$ and $F(k)$ are the PDF and the CDF of the same Pareto distribution as in equation (3.1) so that the intensity function depends on the hazard rate of the Pareto distribution for team productivity. More workers (higher n) scales up the probability of productivity draws uniformly.

The intuition of this process comes from rain falling on a sidewalk. If the rain is uniform across the sidewalk this is a homogeneous Poisson point process. The number of workers n scales up how heavy the rainfall is. Our choice of the particular intensity function over the real line (sidewalk) based on the Pareto distribution in the model puts more “raindrops" on low productivity draws and ensures that even with a continuum of workers the maximum draw is finite.

Then the CDF of the maximum draw for worker i is given by the probability that no draw lies above k, $P[N(k, +\infty) = 0]$:

$$\begin{aligned} F_{z_{it}^{team}}(k) &= \mathbb{P}[N(k, +\infty) = 0] \\ &= \exp\left(-\int_k^{+\infty} \lambda(u)du\right) \\ &= \left(\exp\left(-(\log F(\infty) - \log F(k))\right)\right)^n \\ &= F(k)^n, \end{aligned}$$

and thus, everything follows as in the discrete case from equation (B.2) onward.

B.4 Proofs for Propositions

B.4.1 Proof for Proposition 1

We evaluate the wedges in equation (4.6) at the competitive equilibrium under solo production. Note that the firm decision rule in competitive equilibrium remains the same as in (3.9), except that the equilibrium wage is determined by the labor market clearing condition with solo production. The equilibrium wage is pinned down by

$$\int n^{solo}(\varepsilon_f)g(\varepsilon_f)d\varepsilon_f = \bar{N},$$

which gives

$$w = \frac{1}{\bar{N}^{1-\alpha}}\alpha\zeta z^{\alpha}x_{\min}^{f}\left(\frac{\frac{1}{\eta}}{\frac{1}{\eta}-\frac{1}{1-\alpha}}\right)^{1-\alpha}.$$

We normalize $\bar{N}=1$, and this gives a closed-from solution for firm labor demand under solo production as follows:

$$n_f^{solo} = \frac{(\frac{1}{\eta}-\frac{1}{1-\alpha})\varepsilon_f^{\frac{1}{1-\alpha}}}{\frac{1}{\eta}(x_{\min}^{f})^{\frac{1}{1-\alpha}}}.$$

This also determines spillovers and total output at the equilibrium,

$$\zeta = (x_{\min}^{f})^{\theta}\left(\frac{\frac{1}{\eta}}{\frac{1}{\eta}-(\frac{1-\alpha+\beta}{1-\alpha})}\right)^{\theta}\left(\frac{\frac{1}{\eta}-\frac{1}{1-\alpha}}{\frac{1}{\eta}}\right)^{\beta\theta}$$

$$Y = \zeta z^{\alpha}x_{\min}^{f}\left(\frac{\frac{1}{\eta}}{\frac{1}{\eta}-\frac{1}{1-\alpha}}\right)^{1-\alpha}.$$

Plugging them into (4.6), we obtain the following expression for wedges under solo production:

$$\tau_f^{solo} = \frac{\theta\beta}{\alpha}(x_{\min}^{f})^{\frac{\alpha-\beta}{1-\alpha}}\left(\frac{\frac{1}{\eta}-\frac{1-\alpha+\beta}{1-\alpha}}{\frac{1}{\eta}-\frac{1}{1-\alpha}}\right)\varepsilon_f^{-\frac{\alpha-\beta}{1-\alpha}}. \tag{B.4}$$

Since $0<\beta<\alpha<1$, (B.4) is decreasing in firm productivity, and hence in employment

size, given their positive relationship in competitive equilibrium.

B.4.2 Proof for Proposition 2

We evaluate the wedges in equation (4.9) at the competitive equilibrium under team production. As before, the equilibrium wage is pinned down by the following, which determines the firm decision rule in (3.9) under team production:

$$\int n^{team}(\varepsilon_f) g(\varepsilon_f) d\varepsilon_f = \bar{N},$$

which gives

$$w = \zeta \alpha (1+\gamma) \left(x\Gamma(1-\gamma)\right)^{\alpha} x_{\min}^{f} \left(\frac{\frac{1}{\eta}}{\frac{1}{\eta} - \frac{1}{1-\alpha(1+\gamma)}} \right)^{1-\alpha(1+\gamma)}.$$

Thus, firm labor demand under team production is

$$n_f^{team} = \frac{(\frac{1}{\eta} - \frac{1}{1-\alpha(1+\gamma)}) \varepsilon_f^{\frac{1}{1-\alpha(1+\gamma)}}}{\frac{1}{\eta} (x_{\min}^{f})^{\frac{1}{1-\alpha(1+\gamma)}}}.$$

This also determines spillovers and total output at the equilibrium,

$$\zeta = \left(x_{\min}^{f}\right)^{\theta} \left(\frac{\frac{1}{\eta}}{\frac{1}{\eta} - \frac{1-\alpha(1+\gamma)+\beta}{1-\alpha(1+\gamma)}} \right)^{\theta} \left(\frac{\frac{1}{\eta} - \frac{1}{1-\alpha(1+\gamma)}}{\frac{1}{\eta}} \right)^{\beta\theta}$$

$$Y = \zeta \left(x\Gamma(1-\gamma)\right)^{\alpha} x_{\min}^{f} \left(\frac{\frac{1}{\eta}}{\frac{1}{\eta} - \frac{1}{1-\alpha(1+\gamma)}} \right)^{1-\alpha(1+\gamma)}.$$

Plugging them into (4.9), we obtain the following expression for wedges under team production:

$$\tau_f^{team} = \frac{\theta\beta}{\alpha(1+\gamma)} (x_{\min}^{f})^{\frac{\alpha(1+\gamma)-\beta}{1-\alpha(1+\gamma)}} \left(\frac{\frac{1}{\eta} - \frac{1-\alpha(1+\gamma)+\beta}{1-\alpha(1+\gamma)}}{\frac{1}{\eta} - \frac{1}{1-\alpha(1+\gamma)}} \right) \varepsilon_f^{-\frac{\alpha(1+\gamma)-\beta}{1-\alpha(1+\gamma)}}. \tag{B.5}$$

If $0 < \beta < \alpha(1+\gamma) < 1$, (B.5) is decreasing in firm productivity, and hence in employment size, given their positive relationship in competitive equilibrium.

B.4.3 Proof for Proposition 3

From equations (B.4) and (B.5), the wedges under solo and team production can be written in the unified form:

$$\tau_f = C(\bar{\gamma}) \left(\frac{\varepsilon_f}{x^f_{\min}} \right)^{-\frac{\alpha(1+\bar{\gamma})-\beta}{1-\alpha(1+\bar{\gamma})}},$$

where $\bar{\gamma} = 0$ under solo production, $\bar{\gamma} = \gamma$ under team production, and

$$C(\bar{\gamma}) = \frac{\beta\theta}{\alpha(1+\bar{\gamma})} \left(\frac{\frac{1}{\eta} - \frac{1+\beta-\alpha(1+\bar{\gamma})}{1-\alpha(1+\bar{\gamma})}}{\frac{1}{\eta} - \frac{1}{1-\alpha(1+\bar{\gamma})}} \right)$$

is independent of firm productivity.

Taking logs gives

$$\log \tau_f = \log C(\bar{\gamma}) - \frac{\alpha(1+\bar{\gamma}) - \beta}{1-\alpha(1+\bar{\gamma})} \log \left(\frac{\varepsilon_f}{x^f_{\min}} \right).$$

Hence, the elasticity of the wedge with respect to firm productivity is

$$\frac{\partial \log \tau_f}{\partial \log \varepsilon_f} = -\frac{\alpha(1+\bar{\gamma}) - \beta}{1-\alpha(1+\bar{\gamma})}.$$

Differentiating with respect to $\bar{\gamma}$ yields

$$\frac{\partial}{\partial \bar{\gamma}} \left(\frac{\partial \log \tau_f}{\partial \log \varepsilon_f} \right) = -\frac{\alpha(1-\beta)}{(1-\alpha(1+\bar{\gamma}))^2} < 0,$$

where the inequality follows from

$$0 < \beta < \alpha < \alpha(1+\gamma) < 1.$$

Therefore, the elasticity of wedges with respect to firm productivity is more negative under team production than under solo production. Equivalently, wedges decline more rapidly with firm productivity under team production, implying a steeper relationship between wedges and firm productivity.

B.5 Additional Counterfactuals

B.5.1 Experiment with Internally Calibrated Late-Period γ

In this section, as a robustness exercise, we consider an alternative exercise in which we internally calibrate the late-period value of γ^L to match the share of team patents observed in the data (77.6% in the late period). All other aspects of the calibration strategy remain unchanged. In this case, we obtain $\gamma^L = 0.081$ in the late period.[30]

The results, reported in Table B1, are consistent with the baseline findings. In particular, a lower value of γ^L (relative to the baseline value for the late period) attenuates the increase in the benefits of team production, allowing the model to more closely match the observed increases in firm size and inventor concentration in large firms, which are 10.8 and 44% in the data, respectively (13.1 and 60.7% in the model, respectively).

We further examine how the alternative calibration of γ^L affects the magnitude of misallocation in Table B2. Under the alternative calibration, the competitive equilibrium underproduces innovative output by 1.49% relative to the social planner in the late period. Although the gaps in spillovers and inventor concentration between the competitive equilibrium and the social planner become smaller, the implied total output loss remains similar in magnitude and economically meaningful as before.

[30] Note that this alternative calibration also changes the late-period values of z^L and x^L, as both are determined endogenously along with ζ in the model. The resulting late-period values are $z^L = 1.290$ and $x^L = 1.012$. All other parameter values remain as reported in Table 4.

Table B1: Model Predicted Changes in Aggregate Moments ($\gamma^L = 0.081$)

	Early	Late	$\gamma^L = 0.081$	Change
Inventor wage w	0.724	0.878	0.782	8.0
Prod. cutoff (exit)	1.004	1.241	1.105	10.0
Prod. cutoff (team)	1.580	1.740	1.678	6.2
Share of solo-based firms	93.7	87.2	92.1	-1.5
Share of team patents	57.8	90.4	77.6	19.8
Employment share of firms with $n > 145$	32.8	78.9	60.7	27.9
Average firm size	7.1	27.1	13.1	84.8
Spillovers ζ	0.903	0.827	0.863	-4.4
Output (normalized)	1.000	1.213	1.080	8.0

Notes: This table reports aggregate moments of the competitive equilibrium. The first column presents the early-period calibration, and the second presents the baseline late-period calibration. The third uses an internally calibrated late-period value of $\gamma^L = 0.081$ (targeting the share of team patents), while all other parameters follow the calibration strategy described in Table 4, and the last column reports the corresponding percent or percentage change relative to the early-period benchmark.

Table B2: Competitive Equilibrium vs. Social Planner (CE to SPP, $\gamma^L = 0.081$)

Variable	Early	Late	$\gamma^L = 0.081$
Total Output	-0.20%	-1.72%	-1.49%
Spillovers (ζ)	-0.60%	-6.04%	-2.95%
Emp. share of firms with $n > 145$	7.43 pp	21.60 pp	15.93 pp

Notes: This table reports the deviation of the competitive equilibrium (CE) from the social planner (SPP) for aggregate innovative output, knowledge spillovers, and the employment share of large firms. The first two columns report the early- and the baseline late-period economies. The final column reports a counterfactual in which we internally calibrate the late-period value of $\gamma^L = 0.081$ (targeting the share of team patents), while all other parameters follow the calibration strategy described in Table 4.

B.5.2 Experiments with firm productivity distribution (η)

This Appendix presents results from two additional experiments where η, the shape of the Pareto distribution of firm productivities ε_f, is allowed to vary between the early and late periods. In the first experiment, we vary the idea production and spillover parameters as before in Table 4 and then choose a new η^L to exactly fit the level of employment concentration in the late period. This involves choosing a thinner tail for the firm productivity distribution with lower η^L, lowering mean productivity of firms in the economy. This is because the idea production function changes overstate the rise

Table B3: Model Predicted Changes in Aggregate Moments, $\eta^L = 0.146$

	Early	Late	$\eta^L = 0.146$	Change
Inventor wage w	0.724	0.878	0.717	-1.0
Prod. cutoff (exit)	1.004	1.241	1.014	0.9
Prod. cutoff (team)	1.580	1.740	1.394	-11.8
Share of solo-based firms	93.7	87.2	88.7	-4.9
Share of team patents	57.8	90.4	72.0	14.2
Employment share of firms with $n > 145$	32.8	78.9	44.0	11.2
Average firm size	7.1	27.1	8.8	24.7
Spillovers ζ	0.903	0.827	0.880	-2.5
Output (normalized)	1.000	1.213	0.990	-1.0

Notes: This table reports various aggregate moments of the competitive equilibrium. The first column is the early-period calibration. The second column is the baseline late-period calibration. The third column uses $\eta^L = 0.146$ in the late period (targeting employment concentration) and changes the idea production and spillover parameters as in Table 4, and the last column reports the corresponding percent or percentage change relative to the early-period benchmark.

in employment concentration, and less dispersed firm productivities in the late period are needed to lower employment concentration back down to the value observed in the data.

The results are in Table B3. Lower average firm productivity reduces inventors wages, contrary to the data, and output falls to 0.990. The joint changes, though, get closer to matching the increase in firm size in the data (10.8 in the data, 8.8 in the model), and the team patent share (77.6% in the data and 72% in the model), despite not targeting these two moments.

In a second experiment, we change η^L by itself holding the other parameters fixed at their early values to target employment concentration in the late period. In this case a higher value of $\eta^L = 0.173$ is needed to match the rise in concentration by making the firm productivity distribution more skewed (and raising average firm productivity). The results are in Table B4. The change in the firm productivity distribution to match higher concentration is consistent with a rise in the team patenting share and a modest rise in firm size, consistent with the data. However, the share of solo-based firms declines by less than in the data, and this margin accounts for less than half of the observed increase in the team-patent share in the data.

Table B4: Model Predicted Changes in Aggregate Moments, $\eta^L = 0.173$

	Early	Late	$\eta^L = 0.173$	Change
Inventor wage w	0.724	0.878	0.755	4.3
Prod. cutoff (exit)	1.004	1.241	1.058	5.4
Prod. cutoff (team)	1.580	1.740	1.665	5.4
Share of solo-based firms	93.7	87.2	92.7	-1.0
Share of team patents	57.8	90.4	66.7	9.0
Employment share of firms with $n > 145$	32.8	78.9	44.0	11.2
Average firm size	7.1	27.1	9.1	28.7
Spillovers ζ	0.903	0.827	0.894	-1.0
Output (normalized)	1.000	1.213	1.043	4.3

Notes: This table reports various aggregate moments of the competitive equilibrium. The first column is the early-period calibration. The second column is the baseline late-period calibration. The third column uses $\eta^L = 0.173$ in the late period (targeting employment concentration) and holds the idea production parameters fixed to their early period values in Table 4, and the last column reports the corresponding percent or percentage change relative to the early-period benchmark.

B.6 The Full-Fledged Social Planner's Problem

This section presents the full-fledged setup of the social planner's problem with endogenous firm entry and exit. The following shows the planner's maximization:

$$\max_{M,A,n_f,\text{type}_f} M\int_A \zeta\varepsilon_f(z_f n_f)^\alpha df$$

subject to

$$M\left(\int_A (n_f + c)df + \kappa\right) = \bar{N},\ \zeta = \left(M\int_A \varepsilon_f n_f^\beta df\right)^\theta,$$

where M denotes the mass of firms, A denotes the set of active firms (equivalently, the choice of an exit cutoff), and type_f indicates whether firm f operates as a team-based or as a solo-based firm. The planner maximizes aggregate output, accounting for both the operating cost c and the entry cost κ, which are paid in units of labor.

First, the social planner chooses the mass of firms M and incurs a fixed entry cost κ per firm. After entry, each firm f draws idiosyncratic productivity ε_f. Conditional on ε_f, the planner assigns to each firm f a production mode—solo or team—with labor

efficiency given by

$$z_f^{solo} = z,$$
$$z_f^{team} = x\Gamma(1-\gamma)n_f^{\gamma},$$

and chooses its labor input n_f. The planner can instead shut down the firm. Each active firm $f \in A$ incurs a fixed cost c.

Let λ denote the Lagrange multiplier on the labor market clearing constraint. The Lagrangian associated with the social planner's problem is given by

$$\mathcal{L} = M^{1+\theta}\left(\int_A \varepsilon_f n_f^{\beta} df\right)^{\theta}\left(\int_A \varepsilon_f (z_f n_f)^{\alpha} df\right) - \lambda M\left(\int_A (n_f + c) df + \kappa\right) + \lambda \bar{N}.$$

For expositional simplicity, let

$$Z^{raw} = \int_A \varepsilon_f n_f^{\beta} df,\ Y^{raw} = \int_A \varepsilon_f (z_f n_f)^{\alpha} df,\ Y = M^{1+\theta}(Z^{raw})^{\theta}(Y^{raw}).$$

The first-order condition for a solo firm's labor choice is given by

$$\frac{Y}{M}\left(\theta\beta\frac{\varepsilon_f (n_f^{solo})^{\beta-1}}{Z^{raw}} + \alpha\frac{\varepsilon_f (z_f^{solo})^{\alpha}(n_f^{solo})^{\alpha-1}}{Y^{raw}}\right) = \lambda. \tag{B.6}$$

The first-order condition for a team firm's labor choice is

$$\frac{Y}{M}\left(\theta\beta\frac{\varepsilon_f (n_f^{team})^{\beta-1}}{Z^{raw}} + \alpha\,(1+\gamma)\frac{\varepsilon_f (z_f^{team})^{\alpha}(n_f^{team})^{\alpha-1}}{Y^{raw}}\right) = \lambda. \tag{B.7}$$

The marginal benefits from assigning firms to each production type are given by the

optimized surplus under each technology as follows:

$$\begin{cases} \dfrac{Y}{M}\left(\theta\dfrac{\varepsilon_f(n_f^{solo})^\beta}{Z^{raw}}+\dfrac{\varepsilon_f(z_f^{solo})^\alpha(n_f^{solo})^\alpha}{Y^{raw}}\right)-\lambda(n_f^{solo}+c), & \text{for solo firms,} \\ \dfrac{Y}{M}\left(\theta\dfrac{\varepsilon_f(n_f^{team})^\beta}{Z^{raw}}+\dfrac{\varepsilon_f(z_f^{team})^\alpha(n_f^{team})^\alpha}{Y^{raw}}\right)-\lambda(n_f^{team}+c), & \text{for team firms,} \\ 0, & \text{for inactive firms.} \end{cases} \tag{B.8}$$

The planner assigns firm f to the production type that yields the highest marginal benefit.

Finally, the first-order condition for choosing the mass of firms is

$$(1+\theta)\frac{Y}{M}-\lambda\left(\int_A(n_f+c)df+\kappa\right)=0. \tag{B.9}$$

The solution to the social planner's problem is jointly characterized by conditions (B.6)–(B.9) together with the labor market clearing condition.

The derivation of (B.8) is as follows. For simplicity, consider the following prototype problem:

$$J=\theta\log\int f(c,\varepsilon)d\varepsilon+\log\int g(c,\varepsilon)d\varepsilon,\ \ c(\varepsilon)\in\{0,1\}.$$

Consider a (measurable) perturbation $\delta c(\varepsilon)\in\{-1,0,1\}$ to the choice of firm types such that $\int|\delta c(\varepsilon)|d\varepsilon$ is small. Let $\Delta f(\varepsilon)$ denote the difference $f(1,\varepsilon)-f(0,\varepsilon)$. Then,

$$\delta\int f(c,\varepsilon)d\varepsilon=\int\Delta f(\varepsilon)\delta c(\varepsilon)d\varepsilon.$$

This is also small when $\Delta f(\varepsilon)$ is bounded. Thus,

$$\delta\log\int f(c,\varepsilon)d\varepsilon\doteq\frac{\int\Delta f(\varepsilon)\delta c(\varepsilon)d\varepsilon}{\int f(c,\varepsilon)d\varepsilon}.$$

Similar results hold for g. Therefore, the first-order difference is

$$\begin{aligned}\delta J &\doteq \theta \frac{\int \Delta f(\varepsilon)\delta c(\varepsilon)d\varepsilon}{\int f(c,\varepsilon)d\varepsilon} + \frac{\int \Delta g(\varepsilon)\delta c(\varepsilon)d\varepsilon}{\int g(c,\varepsilon)d\varepsilon} \\ &= \int \left(\theta \frac{\Delta f(\varepsilon)}{\int f(c,\varepsilon)d\varepsilon} + \frac{\Delta g(\varepsilon)}{\int g(c,\varepsilon)d\varepsilon}\right) \delta c(\varepsilon)d\varepsilon.\end{aligned}$$

By the standard argument, this difference is necessarily nonpositive for all valid $\delta c(\varepsilon)$ around the optimal choice of firm types. Finally, the marginal benefit of switching firm type around the optimal firm type allocation is

$$\theta \frac{\Delta f(\varepsilon)}{\int f(c,\varepsilon)d\varepsilon} + \frac{\Delta g(\varepsilon)}{\int g(c,\varepsilon)d\varepsilon}.$$

Substituting into the expressions for f and g yields (B.8). Intuitively, the marginal contribution of changing a firm's type is given by a weighted sum of its effects on the spillover component and the productivity component.

Then, we can prove the following proposition, suggesting the existence of the two productivity cutoffs in the planner's solution.

Proposition B4. *Suppose that* $0 < \alpha, \beta < (1+\gamma)^{-1}$. *When inactive, solo, and team firms coexist, there exist cutoffs* $\tilde{\varepsilon}_{spp}^{solo}$ *and* $\tilde{\varepsilon}_{spp}^{team}$ *such that*

$$x_{min}^{f} < \tilde{\varepsilon}_{spp}^{solo} < \tilde{\varepsilon}_{spp}^{team}.$$

Firms are inactive if $\varepsilon_f < \tilde{\varepsilon}_{spp}^{solo}$, *operate as solo firms if* $\tilde{\varepsilon}_{spp}^{solo} \leq \varepsilon_f < \tilde{\varepsilon}_{spp}^{team}$, *and operate as team firms if* $\varepsilon_f \geq \tilde{\varepsilon}_{spp}^{team}$. *Moreover, at the solo–team cutoff* $\tilde{\varepsilon}_{spp}^{team}$, *we have*

$$z_f^{solo} < z_f^{team} \quad \text{and} \quad n_f^{solo} < n_f^{team}.$$

Proof. Fix λ, Y, M, Z^{raw}, and Y^{raw}. Equations (B.6) and (B.7) imply that n^{solo} and n^{team} are continuously differentiable and strictly increasing in ε. The same holds for the marginal benefit functions MB^{solo} and MB^{team} in (B.8). Hence, at the solo–team cut-

off $\tilde{\varepsilon}_{\text{spp}}^{team}$, the planner must be indifferent between assigning the firm to solo or team production, which requires $MB^{solo}(\tilde{\varepsilon}_{spp}^{team}) = MB^{team}(\tilde{\varepsilon}_{spp}^{team})$.

First, we show that $z_f^{solo} < z_f^{team}$ and $n_f^{solo} < n_f^{team}$ at the solo–team cutoff $\tilde{\varepsilon}_{\text{spp}}^{team}$. Let μ denote the common marginal benefit of assigning a firm to solo and team production at this cutoff. Then (z_f^{solo}, n_f^{solo}) and (z_f^{team}, n_f^{team}) solve, respectively, the same system of equations evaluated at $\eta = 0$ and $\eta = \gamma$:

$$\frac{Y}{M}\left(\theta\beta\frac{\tilde{\varepsilon}_{spp}^{team} n_f^{\beta-1}}{Z^{raw}} + \alpha(1+\eta)\frac{\tilde{\varepsilon}_{spp}^{team} z_f^{\alpha} n_f^{\alpha-1}}{Y^{raw}}\right) = \lambda,$$
$$\frac{Y}{M}\left(\theta(1-\beta)\frac{\tilde{\varepsilon}_{spp}^{team} n_f^{\beta}}{Z^{raw}} + (1-\alpha(1+\eta))\frac{\tilde{\varepsilon}_{spp}^{team} z_f^{\alpha} n_f^{\alpha}}{Y^{raw}}\right) = \lambda c + \mu.$$

The second equation is obtained by substituting (B.6) and (B.7) into the marginal benefit expressions in (B.8). Differentiating this system with respect to η yields

$$\beta\Psi_Z(\beta-1)\frac{n_f'}{n_f} + \alpha_\eta\Psi_Y\left(\alpha\frac{z_f'}{z_f} + (\alpha-1)\frac{n_f'}{n_f}\right) = -\alpha\Psi_Y,$$
$$(1-\beta)\Psi_Z\beta\frac{n_f'}{n_f} + (1-\alpha_\eta)\Psi_Y\left(\alpha\frac{z_f'}{z_f} + \alpha\frac{n_f'}{n_f}\right) = \alpha\Psi_Y,$$

where,

$$\Psi_Z = \frac{Y}{M}\theta\frac{\tilde{\varepsilon}_{spp}^{team} n_f^{\beta}}{Z^{raw}},\ \Psi_Y = \frac{Y}{M}\frac{\tilde{\varepsilon}_{spp}^{team} z_f^{\alpha} n_f^{\alpha}}{Y^{raw}},\ \alpha_\eta = \alpha(1+\eta),\ z_f' = \frac{dz_f}{d\eta},\ n_f' = \frac{dn_f}{d\eta}.$$

By Cramer's rule, we find that

$$\frac{z_f'}{z_f} = \frac{\eta}{D}(\alpha\Psi_Y)^2 > 0,\ \frac{n_f'}{n_f} = \frac{1}{D}(\alpha\Psi_Y)^2 > 0,$$

where, the determinant is given by

$$D = \alpha\Psi_Y\left(\beta(1-\beta)\Psi_Z + \alpha_\eta(1-\alpha_\eta)\Psi_Y\right) > 0.$$

Therefore, $z_f^{solo} < z_f^{team}$ and $n_f^{solo} < n_f^{team}$.

Next, we show that the solo-versus-team cutoff $\tilde{\varepsilon}_{spp}^{team}$ is unique whenever it exists, i.e., MB^{solo} and MB^{team} can cross at most once. A sufficient condition is that whenever $MB^{solo}(\tilde{\varepsilon}_{spp}^{team}) = MB^{team}(\tilde{\varepsilon}_{spp}^{team})$, it holds that

$$\left.\frac{dMB^{solo}}{d\varepsilon}\right|_{\varepsilon=\tilde{\varepsilon}_{spp}^{team}} < \left.\frac{dMB^{team}}{d\varepsilon}\right|_{\varepsilon=\tilde{\varepsilon}_{spp}^{team}}.$$

We establish this by directly evaluating the derivatives at the cutoff $\tilde{\varepsilon}_{spp}^{team}$. The derivatives are, respectively,

$$\begin{aligned}
\frac{dMB^{solo}}{d\varepsilon} &= \left(\frac{Y}{M}\left(\theta\beta\frac{\varepsilon(n_f^{solo})^{\beta-1}}{Z^{raw}} + \alpha\frac{\varepsilon(z_f^{solo})^{\alpha}(n_f^{solo})^{\alpha-1}}{Y^{raw}}\right) - \lambda\right)\frac{dn^{solo}}{d\varepsilon} \\
&\quad + \frac{Y}{M}\left(\theta\frac{(n_f^{solo})^{\beta}}{Z^{raw}} + \frac{(z_f^{solo})^{\alpha}(n_f^{solo})^{\alpha}}{Y^{raw}}\right) \\
&= \frac{Y}{M}\left(\theta\frac{(n_f^{solo})^{\beta}}{Z^{raw}} + \frac{(z_f^{solo})^{\alpha}(n_f^{solo})^{\alpha}}{Y^{raw}}\right), \\
\frac{dMB^{team}}{d\varepsilon} &= \left(\frac{Y}{M}\left(\theta\beta\frac{\varepsilon(n_f^{team})^{\beta-1}}{Z^{raw}} + \alpha\,(1+\gamma)\frac{\varepsilon(z_f^{team})^{\alpha}(n_f^{team})^{\alpha-1}}{Y^{raw}}\right) - \lambda\right)\frac{dn^{team}}{d\varepsilon} \\
&\quad + \frac{Y}{M}\left(\theta\frac{(n_f^{team})^{\beta}}{Z^{raw}} + \frac{(z_f^{team})^{\alpha}(n_f^{team})^{\alpha}}{Y^{raw}}\right) \\
&= \frac{Y}{M}\left(\theta\frac{(n_f^{team})^{\beta}}{Z^{raw}} + \frac{(z_f^{team})^{\alpha}(n_f^{team})^{\alpha}}{Y^{raw}}\right).
\end{aligned}$$

Here, we substitute (B.6) and (B.7) and cancel out the first terms. At the cutoff $\tilde{\varepsilon}_{spp}^{team}$, the marginal benefits coincide. Let μ denote the common value, i.e.,

$$\begin{aligned}
\mu &= \frac{Y}{M}\left(\theta\frac{\tilde{\varepsilon}_{spp}^{team}(n_f^{solo})^{\beta}}{Z^{raw}} + \frac{\tilde{\varepsilon}_{spp}^{team}(z_f^{solo})^{\alpha}(n_f^{solo})^{\alpha}}{Y^{raw}}\right) - \lambda(n_f^{solo} + c), \\
&= \frac{Y}{M}\left(\theta\frac{\tilde{\varepsilon}_{spp}^{team}(n_f^{team})^{\beta}}{Z^{raw}} + \frac{\tilde{\varepsilon}_{spp}^{team}(z_f^{team})^{\alpha}(n_f^{team})^{\alpha}}{Y^{raw}}\right) - \lambda(n_f^{team} + c).
\end{aligned}$$

Therefore, evaluating the derivatives at the cutoff $\tilde{\varepsilon}_{spp}^{team}$, we obtain the following:

$$\left.\frac{dMB^{solo}}{d\varepsilon}\right|_{\varepsilon=\tilde{\varepsilon}_{spp}^{team}} = \frac{1}{\tilde{\varepsilon}_{spp}^{team}}(\mu + \lambda c + \lambda n_f^{solo}),$$
$$\left.\frac{dMB^{team}}{d\varepsilon}\right|_{\varepsilon=\tilde{\varepsilon}_{spp}^{team}} = \frac{1}{\tilde{\varepsilon}_{spp}^{team}}(\mu + \lambda c + \lambda n_f^{team}).$$

Since at this cutoff $n_f^{solo} < n_f^{team}$, it follows that

$$\left.\frac{dMB^{solo}}{d\varepsilon}\right|_{\varepsilon=\tilde{\varepsilon}_{spp}^{team}} < \left.\frac{dMB^{team}}{d\varepsilon}\right|_{\varepsilon=\tilde{\varepsilon}_{spp}^{team}}.$$

As previously shown, this implies that the marginal benefit functions MB^{solo} and MB^{team} can cross at most once. Therefore, the solo–team cutoff is unique whenever it exists.

Finally, we show that the exit cutoff $\tilde{\varepsilon}_{spp}^{solo}$ is unique whenever it exists. Substituting (B.6) and (B.7) into the marginal benefit expressions in (B.8) yields that

$$MB_f^{solo} = \frac{Y}{M}\left(\theta(1-\beta)\frac{\varepsilon_f (n_f^{solo})^{\beta}}{Z^{raw}} + (1-\alpha)\frac{\varepsilon_f (z_f^{solo})^{\alpha}(n_f^{solo})^{\alpha}}{Y^{raw}}\right) - \lambda c,$$
$$MB_f^{team} = \frac{Y}{M}\left(\theta(1-\beta)\frac{\varepsilon_f (n_f^{team})^{\beta}}{Z^{raw}} + (1-\alpha\,(1+\gamma))\frac{\varepsilon_f (z_f^{team})^{\alpha}(n_f^{team})^{\alpha}}{Y^{raw}}\right) - \lambda c.$$

Regarding both as functions of ε_f, each is strictly increasing in ε_f. Hence, their pointwise maximum can cross the horizontal line $MB = 0$ at most once. This proves that the exit cutoff is unique whenever it exists. □