Journal of Empirical Finance 74 (2023) 101435

Contents lists available at ScienceDirect

Journal of Empirical Finance

journal homepage: www.elsevier.com/locate/jempfin

The role of human capital: Evidence from corporate innovation☆

Tong Liu[a], Yifei Mao[b,*], Xuan Tian[c]

[a] MIT, USA
[b] Cornell University, USA
[c] Tsinghua University, China

ARTICLE INFO

Keywords:
Human capital
Inventors
Innovation strategy
Innovation performance

ABSTRACT

This paper examines the distinct roles played by inventors and firms in contributing to corporate innovation. Inventors are six to eight times as important as firms in contributing to innovation performance as measured by patent and citation counts, but their importance is about the same in innovation strategies as captured by patent exploratory and exploitative scores. Furthermore, when labor mobility is reduced, the relative importance of firms to inventors in contributing to innovation strategy increases. Additional tests suggest that our main findings are unlikely driven by endogenous matching between firms and inventors.

1. Introduction

Since Coase (1937), there has been a longstanding debate on what constitutes a firm. In the Hart-Moore framework, nonhuman assets are the glue that brings a firm together (Hart, 1995), but Zingales (2000) argues that "human capital is emerging as the most crucial asset" in today's world. While a firm consists of both organizational capital—physical assets, organizational structure, corporate culture, and access to resources—and human capital, the roles played by human capital and organizational capital in explaining a firm's long-run success are still unclear to researchers.[1] Is a firm's growth engine embedded in its human capital? How important is a firm's organizational capital in determining its growth strategy and performance? This paper attempts to answer these questions and sheds new light on our understanding of a firm's long-run growth.

To this end, we focus on the engine driving a firm's competitive advantage and long-run success, namely, technological innovation. Specifically, we decompose the contribution of the firm and its inventors to inventor-level innovation output. We investigate two dimensions of a firm's innovation output. The first dimension is innovation performance as measured by both the quantity (patents) and quality (citations) of the firm's innovation output. The second dimension is innovation strategy as measured by the exploratory and exploitative scores of the firm's innovation output. As defined in the existing literature (e.g., Manso et al., 2017; Gao et al., 2018; Fitzgerald et al., 2021), *exploratory innovation* is radical innovation that requires knowledge outside of the existing knowledge domain, and *exploitative innovation* is incremental innovation that builds on existing knowledge and improves existing skills, processes, and

☆ We are grateful for the comments and suggestions from Linda Canina, Daniel Ferreira, Andrew Ellul, Gerard Hoberg, Andrew Karolyi, Pam Moulton, Veronica Pool, Noah Stoffman, Scott Yonker, Alminas Zaldokas, and seminar and conference participants at Cornell University, Indiana University, the 2016 SFS Cavalcade conference, the 2016 European Finance Association annual meeting, the 2016 China International Conference in Finance, and the 2018 American Finance Association annual meeting. We alone are responsible for any errors or omissions.

* Corresponding author.
E-mail address: ym355@cornell.edu (Y. Mao).

[1] Following Ewens and Rhodes-Kropf (2015), we use the term "organizational capital" to refer to a firm's nonhuman capital.

https://doi.org/10.1016/j.jempfin.2023.101435
Received 7 December 2021; Received in revised form 2 October 2023; Accepted 4 October 2023
Available online 5 October 2023
0927-5398/

structures. Therefore, exploration and exploitation represent a firm's strategy in long-term growth.

While there have been many studies investigating the various market and firm characteristics that contribute to innovation output, little is known about the roles played by inventors as compared to the roles of firms.[2] A major challenge of quantifying the roles played by inventors in contributing to innovation is the need to separate the contribution of organizational capital from that of human capital. Patent generation provides a clean, rich, and unique setting for dealing with this empirical challenge because we can track individual inventors' patent filings and the citations of those patents. Further, we are able to observe inventors' moves from one firm to another using the Harvard Business School (HBS) Inventor Patent Citation database. Intuitively, if outputs of individual inventors do not change as they move, the inventor is likely to be the main driver of the associated output. If, however, an individual inventor's innovation output changes significantly after a move between firms, then the change in innovation output can largely be attributed to the new firm.

We employ two methods to isolate the roles played by firms and inventors. The first method focuses on a panel of inventors who have changed their affiliated firms; inventor, firm, and year fixed effects are included in the regressions. We refer to this approach as the "mover dummy variable" (henceforth, MDV) method, which has been commonly used in the existing literature (e.g., Bertrand and Schoar, 2003; Graham et al., 2012). However, because the MDV approach is limited to movers only, who account for 16 % of all inventors in our sample, we use an alternative method that includes both movers and stayers in the sample, as long as the stayers are in firms that employ at least one mover. This method was developed by Abowd et al. (1999) (henceforth AKM) and later refined by Abowd et al. (2002). The AKM method extends a rather small sample of movers to a connectedness sample, which includes 98.4 % of all inventors in our setting. The key results and their economic implications from our paper are similar regardless of the method we use.

Intuitively, using the MDV and AKM methods, we want to quantify how much of the observed variation in innovation output can be attributed to inventor fixed effects and firm fixed effects, respectively. Our results reveal that inventor fixed effects are consistently more important than firm fixed effects in explaining innovation performance as measured by patent counts and citations per patent, but they are approximately equal in importance in explaining innovation strategy as measured by patent exploratory and exploitative scores. Specifically, our estimates suggest that inventor fixed effects are about six to eight times as important as firm fixed effects in explaining the firm's innovation performance, while inventor fixed effects are of about the same importance as firms fixed effects in explaining the firm's innovation strategy. The results suggest that, while human capital is crucial in determining a firm's innovation performance, its effect on a firm's innovation strategy is much more moderated.

Furthermore, we explore how a policy shock to inventor mobility might affect firms' and inventors' relative contributions to innovation. In particular, we use the staggered adoption of the inevitable disclosure doctrine (IDD) by US state courts as plausibly exogenous shocks to the mobility of inventors. The adoption of IDD prevents employees with knowledge of trade secrets held by incumbent firms from moving to other firms, and thus lowers inventor mobility. Our results show that, when inventor mobility is reduced, the relative importance of firms versus that of inventors in terms of innovation performance does not vary, but the relative importance of firms versus that of inventors for innovation strategy increases. The results highlight the importance of firms' organizational capital in determining innovation strategy, especially when inventors' outside options are reduced due to restrictions on their mobility.

A drawback to the AKM approach is that it estimates inventor fixed effects and firm fixed effects over the entire sample period. Therefore, it is difficult to investigate how the innovation capability of inventors and firms varies over time and how it affects the match between them. To address this concern, we follow Bhaskarabhatla et al. (2021) to conduct a rolling-window estimation of inventor and firm fixed effects and study the sorting patterns between inventors and firms. We find that inventors with high-innovation performance tend to move to firms that have low-innovation performance but high-innovation performance coworkers. In addition, inventors tend to move to firms with different innovation strategies but to work with other inventors with similar innovation strategies.

In the final part of the paper, we conduct a series of tests to evaluate how endogenous matching between inventors and firms might affect our results. Endogenous matching between inventors and firms could cause us to wrongly attribute the change in the inventor's output to firm fixed effects, when the change actually comes from an interaction between the inventor and the firm. In other words, matching might contaminate our estimation of the relative importance of the inventor and the firm. We first follow Card et al. (2013) and conduct direct tests to assess how matching affects our results. The test results show that endogenous matching carries little explanatory power to innovation. We then examine a variety of subsamples that have different exposures to the endogenous moving of inventors.

We first compare the relative importance of inventors and firms when inventors move to similar firms (i.e., firms with similar past patenting outcomes or in the same industry) and when inventors move to firms that are very different from their current firms (i.e., firms with different past patenting outcomes or in a different industry). The concern is that firm fixed effects might be underestimated if inventors endogenously choose to move to similar firms in which they are less likely to experience changes in innovation output. However, we find that inventors' relative importance to firms in explaining innovation performance and strategy remains robust in the subsample where inventors move to similar firms and the subsample where inventors move to different firms. We also examine the relative importance of inventors and firms in a subsample that is less likely to be subject to the matching concern; we do so by focusing on a subsample of movers who exhibit few changes in their innovation output as they move across firms. The rationale of this test is that these movers are less likely to be subject to endogenous matching issues. Our main findings continue to hold in this test.

[2] See, for example, Aghion, Bloom, Blundell, Griffith, and Howitt (2005), Acharya and Subramanian (2009), Manso (2011), Acharya, Baghai, and Subramaniam (2014), Seru (2014), and Cornaggia, Mao, Tian, and Wolfe (2015), Chemmanur, Ertugrul, and Krishnan (2019). One exception is Bhaskrabhatla, Cabral, Hegde, and Peeters (2021), who study the contribution of inventors and firms to innovation performance.

The contribution of our paper is threefold. First, we contribute to the literature on the economics of organization. This literature proposes different hypotheses for the existence of the firm and distinguishes physical capital from human capital (Coase, 1988; Klein, 1988; Williamson and Winter, 1993). An empirical study by Kaplan et al. (2009) examines startup companies and suggests that the business (nonhuman capital) is more important than its management team (human capital). Other studies examine how firms strategically hire human capital to benefit the firm (e.g. Cohen et al., 2012). Providing a fresh view into the literature, our findings highlight the importance of human capital in innovation performance and the importance of organizational capital in innovation strategy.

Second, out paper contributes to the literature on corporate innovation. Prior studies examined various determinants including the legal environment (Acharya et al., 2014; Lin et al., 2021), CEO compensation contracts (Bazjak et al., 1993; Baranchuk et al., 2014), banking competition (Cornaggia et al., 2015), financial market development (Hsu et al., 2014), institutional ownership (Bushee, 1998, 2011; Aghion et al., 2013), government subsidy (Howell, 2017), and hedge fund interventions (Brav et al., 2018). These studies, however, focus only on extensive margins that affect firm-level innovation. In this paper, we delve further into the intensive margin and decompose innovation drivers into human-capital-related and organization-capital-related components, which allows us to further understand the relative importance of these two components in promoting firm innovation.

Third, this paper contributes to the expanding literature that attempts to separate human capital and organizational capital. The findings in the existing literature highlight the importance of human capital. Abowd et al. (1999) find that individual effects are more important than firm effects in explaining wage variations in France. Bertrand and Schoar (2003) and Graham et al. (2012) show that manager fixed effects explain a significant extent of firm policy heterogeneity and that managers with higher performance fixed effects receive higher compensation. Ewens and Rhodes-Kropf (2015) find that venture capitalists have repeatable skills and VC partners' human capital is more important than VC firms' organizational capital in explaining performance. Berk et al. (2017) stress the role of mutual fund firms in efficiently allocating capital to their managers. Gao et al. (2017) examine the relative contribution of loan officers and banks to loan lending. Chemmanur et al. (2019) compare investment banks and bankers in the M&A setting. Cho et al. (2016) find that firms are more important than managers in driving patent generation.[3]

Our paper is closely related to the work of Bhaskarabhatla et al. (2021) who study the contribution of inventors and firms to innovation performance. Different from their paper, our study examines innovation strategy as well as innovation performance. We find a lower relative importance of inventors to firms in driving innovation strategy than in driving innovation performance. We also examine how the restrictions of labor mobility affect the contribution of inventors, and find that the importance of inventors in driving innovation strategy decreases when labor mobility is restricted.

2. Hypotheses development

As shown in Bhaskarabhatla et al. (2021), inventor-specific skills are more important than firm-specific capabilities in explaining the variance in the inventive performance of inventors. This is consistent with prior studies, which argue and provide evidence that inventors possess tacit knowledge and can take this knowledge with them to new firms (Zucker et al., 2002; Palomeras and Melero 2010; Singh and Aggrawal 2011). Given that knowledge is held by individuals, it is intuitive that innovation performance is influenced more by inventors than by firms.

Though many studies on innovation have focused on innovation performance, more have started to examine innovation strategy, motivated by particularly successful outcomes or breakthroughs (e.g., Taylor and Greve, 2006; Fleming, 2007). We hypothesize that, for innovation strategy, the relative importance of firms compared to that of inventors is greater in terms of innovation performance. To provide conceptual foundations for the hypothesis, we follow the literature and borrow insights from a stylized evolutionary model of innovation (e.g., Campbell, 1960; Weitzman, 1998; Singh and Fleming, 2010).

As described in the model, innovation is an evolutionary search process with three phases: variation, selection, and retention. In the "variation" phase, inventors generate new ideas through combinatorial thought trials. The idea that novel technologies can almost always be traced back to combinations of existing technologies has been established in classic studies (Schumpeter, 1934; Weitzman, 1998; Olsson, 2000). In the "selection" phase, ideas are reviewed. Poor outcomes are rejected, and the most promising novelties are identified. In the "retention" phase, members of a larger creative community evaluate the selected ideas and go on to adopt a few of them.

We argue that in both the "variation" and "selection" phases, team collaborations—particularly those embodied in the collaboration at the firm level—plays a more important role than individuals in determining how novel the ideas can be.[4] First, in the "variation" phase, there would be greater diversity within team collaboration as each inventor brings a different set of knowledge, skills, and past experiences. As a result, team collaboration brings more recombinant opportunities to generate relatively novel ideas. Given that the new combinations are more uncertain (Fleming, 2001), collaboration increases the variance of the invention (i.e., both good and bad outcomes). Second, in the "selection" phase, collaborative selection is usually more effective than individual selection. This is because individuals are likely to be bad evaluators of their own ideas (Runco and Smith, 1991). In contrast, a team with a more diverse background is more likely to have members with experience in the new, breakthrough field and to be able to evaluate its

[3] The findings of Cho et al. (2016) are likely driven by the fact that managers are not the main drivers of patent generation. While executives may set the boundaries for the R&D group, the exact innovation directions and outputs are determined by the knowledge and expertise possessed by inventors in the firm.

[4] The "retention" stage could be thought of as the aggregation of "selection" choices made by a larger community outside the team.

quality.

Taking the above arguments together, to the extent that a firm can be viewed as a natural extension of its teams in the creation of new knowledge (Kogut and Zander, 1992), we suggest that a firm should play a larger role in innovation strategy than in innovation performance. This leads to our first hypothesis.

Hypothesis 1. The relative importance of firms compared to that of inventors is greater for innovation strategy than for innovation performance.

We also examine the dynamic shift of the relative importance of firms compared to inventors when labor mobility is reduced. As shown in the literature, when labor mobility is restricted, it is more difficult for inventors to move to other firms. For example, Jeffers (2019) shows that stronger restriction in labor mobility leads to a substantial decline in employee departures, especially in knowledge-intensive occupations. With restricted mobility, inventors would have fewer outside options, which would increase firms' bargaining power against the inventors. Prior studies have consistently shown that—as labor mobility declines—wages decrease, firm investment increases, and firm investment per employee drops (e.g., Garmaise, 2011; Jeffers 2019; Starr, 2019).

When their bargaining power increases, firms likely have an even greater impact on innovation strategy, mainly through their impact on the selection phase. This is because, when ideas are reviewed in the selection phase, firms with greater bargaining power relative to inventors have a stronger say as to which ideas are promising and which are poor. In contrast, an increase in firms' bargaining power does not affect their impact on the innovation performance, given that knowledge is still held by individuals.

The above argument leads to our second hypothesis.

Hypothesis 2. When labor mobility is reduced, firms are relatively more important than inventors for increasing innovation strategy.

3. Data

3.1. Sample construction

We begin with the latest version of the Harvard Business School (HBS) patent and inventor database available at http://dvn.iq.harvard.edu/dvn/dv/patent.[5] This database provides information both on inventors (the individuals who receive credit for producing the patent) and on the assignee (the entity—a government, a firm, an organization, or an individual—that owns the patent).

For the purpose of our study, we track the employers of inventors as they move from one firm to another. Since the patent database does not indicate the inventor's employment directly, we assume that the employer is the company to which the patent (filed by the inventor) is assigned. If a patent is only assigned to one assignee, the employer is clearly identified. However, when there are multiple assignees, the HBS patent and inventor database reports only the primary assignee. This issue confounds identification of the inventor's employer. To overcome this problem, we match the HBS patent and inventor database with the National Bureau of Economics Research (NBER) patent citation data set that contains precise patent and assignee information.[6] Appendix A describes in more detail how we pin down the employer for each inventor when there are multiple assignees.

To obtain time-varying firm characteristics, we merge the inventor-year patent sample with the firm-level annual accounting variables obtained from Compustat. We require all firms to have non-missing financial records across their life cycle as they appear on Compustat records. Finally, we omit observations occurring before 1970 (i.e., 345 observations), which are only a small portion of our final sample. Our final sample consists of 204,678 inventors (1246,951 inventor-year observations) who have worked for 5722 firms from 1970 through 2003.[7]

3.2. Variable measurement

3.2.1. Measuring innovation

We construct two sets of patent-based metrics to gauge an inventor's innovation output. The first set measures innovation performance and the second set captures innovation strategy. Following the innovation literature, we include, in the first set, the total number of patents filed and eventually granted in a given year by an inventor, which captures the quantity of the inventor's innovation output. When a patent has multiple inventors, we assign count one to each inventor; in this case, all results are robust if we split the count evenly among all co-inventors. We use the patent application year instead of the grant year to determine an inventor's innovation output because the application year is closer to the actual time when innovation activities would have occurred (Griliches et al., 1986, 1990). Although the intuition is straightforward and easy to construct, a simple measure of patent counts hardly distinguishes breakthrough innovations from incremental technological f tnd ofich of□

innovation output, the total number of non-self-citations of each patent in subsequent years. We use this measure to capture the quality (or impact) of an inventor's innovation output.

Nevertheless, both innovation measures suffer from severe truncation problems. Our matched sample includes only patents that were eventually granted by the end of 2006; patents filed in the last few years of our sample period may still be under review and may not have been granted by 2006. (This truncation problem is caused mainly by using the NBER database, which is updated through 2006). To adjust truncation bias in patent counts, we calculate the number of patents filed in a given year (and eventually granted) by each inventor in the HBS database, which contains patents granted through 2010. To the extent that the patent application outcomes were announced by 2010 for patents filed by 2006, this approach greatly alleviates the truncation concern. While patents tend to be cited over a long period after their grant date, we observe only the citations occurring through 2010. To deal with this truncation bias, we correct the citation data by using "weight factors" following Hall et al. (2001, 2005) and estimate the shape of the citation-lag distribution.

Consistent with the innovation literature, the distribution of patent grants in our final sample is right skewed, with its median at zero. Due to the right skewness of patent counts and citations per patent, we use the natural logarithm of one plus patent counts (*LnPatent)* and the natural logarithm of one plus the number of citations per patent (*LnCitePat)* as the main innovation metrics to measure innovation performance in our analysis. We also winsorize all our dependent variables at the 99th percentile.

The second set of metrics includes the exploitative (*Exploit*) and exploratory (*Explore*) scores of patents to reflect an inventor's innovation strategy. We follow existing literature (e.g., among others, Sørensen and Stuart, 2000; Katila and Ahuja, 2002) to categorize an inventor's patenting activity as exploratory innovation or exploitative innovation. The basic idea is that inventors who concentrate on existing knowledge are expected to produce more exploitative patents while inventors who explore new ideas are expected to create more exploratory patents. We define an inventors' existing knowledge as their previous patent portfolio and the set of patents cited by their own patents over the past five years. Proxies are constructed to classify a patent as exploitative if at least 60 % of its citations are based on existing knowledge, which includes all the patents that the firm filed and all the patents that were cited by the firm's patents filed over the past five years. A patent is classified as being exploratory if at least 60 % of its citations are not based on existing knowledge.[8] We then set *Exploit* equal to the ratio of the number of exploitative patents filed by an inventor in a given year to the total number of patents filed by the same inventor in the same year. Similarly, we define *Explore* as the ratio of the number of exploratory patents for a given year to the total number of patents filed by the inventor in the same year.

Note that the patent databases used in our study are unlikely to be affected by survivorship bias. As long as a patent application is eventually granted, the patent is attributed to the applying firm at the time of application even if the firm is later acquired or goes bankrupt. Moreover, because patent citations are attributed to the patent rather than to the assignee, the patent granted to a firm that is later acquired or goes bankrupt may continue to be cited long after the firm disappears.

For firm characteristics, we compute all variables for firm *i* in fiscal year *t*. Our control variables include firm size (the natural logarithm of book value assets), firm age (the natural logarithm of a firm's age since its IPO year), profitability (ROA), investments in intangible assets (R&D expenditures over total assets), asset tangibility (net PPE scaled by total assets), leverage, capital expenditures, growth opportunities (Tobin's Q), financial constraints (the KZ index of Kaplan and Zingales [1997]), and industry concentration (the Herfindahl index based on sales). Aghion et al. (2005) point out non-linear effects of product market competition on innovation output. Hence, we include the squared Herfindahl index in our regressions. We provide detailed variable definitions in Appendix B.

3.2.2. Measuring inventor mobility

Table 1 presents information on the inventors in our sample who moved and did not move. Panel A shows that during the sample period, 15.9 % (32,561) of inventors were movers who worked in more than one firm in the sample, while the rest 84.1 % were non-movers who worked in the same sample firm throughout our sample period. Panel B provides information on the number of movers per firm during the sample period. Of the total, 24.7 % of the sample firms had no movers while the remaining 75.3 % had at least one mover.

3.2.3. Measuring control variables

Following the innovation literature, we control for a vector of inventor and firm characteristics that may affect innovation output. For inventor time-varying characteristics, we create proxies for inventors' prior innovation experience. Two variables, Ln*Expnum* and Ln*Expcit*, are defined as the logarithm of one plus the innovation metrics (adjusted patent count and citations per patent, respectively) over the past three years of a given year. We use a three-year rolling window because recent experience is a good indicator that the inventor is actively innovating (Chemmanur et al., 2019). The construction of these variables requires information on the past three years' invention experience of a given year; hence we exclude the first three-year observations for all inventors because their prior innovation experience is missing. To keep the mobility structure intact, we also exclude inventors whose moving happens in the first three years of our sample period.

3.3. Summary statistics

We construct the baseline AKM estimation sample by identifying firms in a connectedness sample, which includes not only

[8] We also use 80% and the results are robust when using variables defined by this alternative cutoff.

inventors who moved but also non-movers who work in firms that have hired at least one mover. To define a connectedness sample, we use graph theory to determine groups of inventors and firms that are connected. Detailed procedures are as follows: We start with an arbitrary inventor and track all firms where this inventor ever worked. Then we include in our connectedness sample all inventors who worked in these firms and continue tracking all firms for which these inventors have ever worked. We repeat the procedure until all data are exhausted.

We examine whether the connectedness sample is representative of the full sample and provide summary statistics of the variables in both the full sample and the connectedness sample in Table 2 (Brav et al., 2005).[9] Panel A summarizes the representativeness of these variables for inventors. To minimize the effect of outliers, we winsorize all control variables at the 1st and 99th percentiles. On average, an inventor in the sample files 0.9 patents, which are eventually granted, each year. Each patent is cited 6.5 times, which is comparable to the full sample in the NBER database where each inventor on average has 0.934 patents per year and each patent is cited 8.009 times. In the connectedness sample, inventors have similar numbers of patents granted per year, 0.9, and each patent is cited 6.5 times. Other variables that measure an inventor's exploratory or exploitative characteristics are also close in both the full sample and the connectedness sample.

Panel B of Table 2 summarizes the representativeness of these variables for firms. In the full sample, an average firm has book value assets of $7.12 billiof of□0

Table 2
Summary statistics
This table reports summary statistics for the full sample and for the connectedness sample at the inventor and firm levels when using patent counts and citations per patent as dependent variables in our baseline regression. Panel A presents the summary statistics of patent counts and citations per patent as well as the time-varying inventor characteristic measures for inventor-year observations. Panel B presents the summary statistics of the time-varying firm characteristics measures for firm-year observations. Definitions of variables are listed in Appendix B.

Panel A: Summary Statistics for Inventors

Variable	Mean	Median	SD	25th	75th	N (inventor-year)
Patent Counts						
Full Sample	0.910	0	1.484	0	1	1246,951
Connectedness Sample	0.912	0	1.486	0	1	1236,561
Citations per Patent						
Full Sample	6.463	0	13.077	0	7.457	1246,951
Connectedness Sample	6.476	0	13.091	0	7.481	1236,561
Exploit						
Full Sample	0.126	0	0.307	0	0	555,592
Connectedness Sample	0.127	0	0.307	0	0	548,233
Explore						
Full Sample	0.745	1	0.401	0.5	1	555,592
Connectedness Sample	0.745	1	0.401	0.5	1	548,233
Panel B: Summary Statistics for Firms						
Variable	Mean	Median	SD	25th	75th	N (firm-year)
Assets (million)						
Full Sample	7120.888	691.71	21,416.62	110.87	4084.982	46,177
Connectedness Sample	7569.452	844.164	22,253.39	137.228	4512	40,047
RDAssets						
Full Sample	0.059	0.032	0.143	0.010	0.068	46,177
Connectedness Sample	0.060	0.035	0.140	0.012	0.069	40,047
Age						
Full Sample	21.932	21.000	13.151	10.000	32.000	46,177

include time-varying controls for inventor and firm. The time-varying controls at the firm level include firm asset size, R&D expenditure, age, profitability (measured by ROA), tangible assets (measured by PPE), leverage, capital expenditure, Tobin's Q, financial constraints, and product market competition. The time-varying controls at the inventor level include the inventors' previous innovation performance (measured by patent counts and citations per patent in the three previous years). The term μ_t captures the year fixed effects and the term γ_{it} controls for the functional area effects that inventor i belongs to year t. We define the functional area for an inventor as the primary functional class of most patents for which inventor i applies in year t.[10] Our focus is on retrieving both inventor and firm fixed effects ϕ_i and θ_j using the movements of inventors across firms and then evaluating the contribution of inventor and firm fixed effects to the R-squared. Hence the treatment of the standard errors in the estimation would not affect the conclusion on the contribution from inventors relative to that from firms.

We use a method first proposed by Abowd et al. (1999) (hereafter referred to as the AKM method) and later refined by Abowd et al. (2002).[11] Abowd et al. (2002) show that connections make the estimation of inventor and firm fixed effects for each connected group relative to a within-group benchmark computationally feasible. To make inventor and firm fixed effects directly comparable across groups, we follow the normalization procedure suggested by Cornelissen (2008): First, we normalize the mean firm fixed effects for each group to zero and add the group mean firm fixed effects to inventor fixed effects; Second, we subtract the grand mean of inventor fixed effects from each inventor fixed effect and then add this grand mean inventor fixed effect to the intercept.

An analogous method (i.e., the MDV method) proposed by Bertrand and Schoar (2003) uses the standard least square dummy variable (LSDV) approach (as in Graham et al., [2012]) and employs a mobility sample, consisting only of movers and the firms for which they work, to separate firm fixed effects from individual fixed effects. One disadvantage of the MDV method compared to the AKM method is a potential sample selection bias resulting from restricting the sample to only movers, who may be different from non-movers. In addition, there are other important benefits from adopting the AKM framework. First, the AKM method uses information on both movers and non-movers, which provides a larger sample size and higher statistical power. Second, this approach can significantly reduce the computational work involving the large data set used in our study. Nonetheless, we follow the MDV approach in our analysis as a robustness check.

We now provide a detailed discussion on how the AKM method identifies inventor and firm fixed effects separately in the connectedness sample. We define the variable D_{ikt} as a dummy that equals one if inventor i works at firm k at time t and zero otherwise. Then we can rewrite Eq. (1) as:

$$Y_{ij(t+3)} = \beta_1 X_{it} + \beta_2 Z_{jt} + \gamma_{it} + \phi_i + \sum_{k=1}^{J} D_{ikt}\theta_k + \mu_t + \epsilon_{ijt}. \tag{2}$$

In the first step, the AKM approach sweeps out the inventor fixed effects by averaging over all inventor i's innovation performance to obtain:

$$\bar{Y}_i = \beta_1 \bar{X}_i + \beta_2 \bar{Z}_i + \bar{\gamma}_i + \phi_i + \sum_{k=1}^{J} \bar{D}_{ik}\theta_k + \bar{\mu}_t + \bar{\epsilon}_i. \tag{3}$$

Here $\bar{Y}_i$ is inventor i's average innovation performance across the full sample period. Then we begin to demean (2) with (3) to get:

$$Y_{ij(t+3)} - \bar{Y}_i = \beta_1 (X_{it} - \bar{X}_i) + \beta_2 \left(Z_{jt} - \bar{Z}_i\right) + (\gamma_{it} - \bar{\gamma}_i) + \sum_{k=1}^{J} (D_{ikt} - \bar{D}_{ik})\theta_k + (\mu_t - \bar{\mu}_t) + \left(\epsilon_{ijt} - \bar{\epsilon}_i\right). \tag{4}$$

The demeaning process removes the inventor fixed effects. Now it is clear that we can exploit movers' information to identify firm fixed effects since $D_{ikt} - \bar{D}_{ik} \neq 0$ for a mover, which can be estimated by the LSDV method. Finally, using the estimates in the above regression, we can recover the inventor fixed effects by the following equation:

$$\hat{\phi}_i = \bar{Y}_i - \hat{\beta}_1 \bar{X}_i - \hat{\beta}_2 \bar{Z}_i - \bar{\gamma}_i - \sum_{k=1}^{J} \bar{D}_{ik}\hat{\theta}_k \tag{5}$$

and here $\bar{\mu}_t$ is often treated as the benchmark in estimating time effects and is thereby assumed to be zero.

As Abowd et al. (2004) and Andrews et al. (2008) note, an estimation bias may emerge when inventor mobility is limited; this could lead to imprecise estimation of inventor and firm fixed effects. Consequently, we need to exercise caution when interpreting the results in both the MDV and AKM methods. However, this issue is not severe in our study because movers represent about 16 % of our sample. This proportion is relatively high compared to percentages reported in previous literature (e.g., Graham et al., 2012). Another property of the AKM estimator is that fixed effect estimates themselves have properties that are similar to other estimators. As shown by Wooldridge (2010), estimates of time-varying variable coefficients are both unbiased and consistent, while fixed effect estimates are only unbiased.

4.2. Hypothesis 1 testing: the relative importance of firms compared to that of inventors

In this section, we test our first hypothesis by investigating whether the relative importance of firms compared to that of inventors is greater for innovation strategy than for innovation performance. To this end, we first examine the relative importance of firms compared to inventors separately for innovation performance and for innovation strategy; we then compare the two.

Our testing sample is the connectedness sample used in the AKM method. This sample excludes firms with no movers during our sample period. Based on this procedure, the connectedness sample for innovation performance has 201,461 inventors (32,561 movers), 4310 firms, and 1239,614 inventor-year observations; this accounts for 98 % of all inventors, 75 % of all firms, and 99 % of all

[11] In most of our analyses, we use Cornelissen's (2008) Stata command "felsdvreg" to implement the AKM method and estimate both inventor and firm fixed effects. This command facilitates the estimation of a linear model with two high-dimensional fixed effects (i.e., inventor and firm fixed effects) by using a memory-saving decomposition of the design matrix. It also provides useful summary statistics. In some cases with tremendous data size, we switch to the Stata command "reghdfe" proposed by Correia (2014), which is more efficient when dealing with data that requires large memory.

observation units.[12] We estimate Eq. (1) using the AKM method in the connectedness sample. Following prior literature, we select the observable characteristics of inventors and firms that may affect an inventor's future innovation output (e.g., He and Tian, 2013; Seru, 2014; Cornaggia et al., 2015). Specifically, in our full fixed effects model we regress the proxy of inventors' innovation performance on both firm and time-varying variables, such as firm size, age, profitability, and intangible assets and on inventor time-varying variables such as the inventor's prior experience. Additionally, we include year fixed effects to capture the impact of economic conditions.

Table 3 reports the estimation results. Following Graham et al. (2012) and Ewens and Rhodes-Kropf (2015), we use $\frac{cov(Y,\ Inventor\ FE)}{var(Y)}$ to capture the contribution of inventor fixed effects to the total variation in inventors' innovation output. The ratio $\frac{cov(Y,\ Inventor\ FE)}{var(Y)}$ reports the covariance of the dependent variable with inventor fixed effects, scaled by variance of the dependent variable. These normalized covariance terms represent the fractions of total variations attributable to particular factors, which can effectively capture the relative importance of individual fixed effects in explaining the dependent variable for a given regression model. In addition, in Table 3 we report adjusted R-squared across four different model specifications for each dependent variable: the first specification includes all control variables and year effects; the second specification includes firm fixed effects in addition to the variables in the first specification; the third specification includes inventor fixed effects in addition to the variable in the second specification; the fourth specification includes all control variables and both firm fixed effects and inventor fixed effects. We adopt adjusted R-squares in this case because the number of explanatory variables changes across models.

In Table 3, Columns 1 and 2, we examine innovation performance. For patent counts in Column 1, inventor fixed effects account for 53.1 % of the total variation while firm fixed effects contribute 8.4 % of the total variation (the remaining portion is attributable to all other controls). The relative importance of inventor fixed effects compared to firm fixed effects is measured by the ratio of the contribution of these two fixed effects, which is around six times as reported in Column 1. For citations per patent shown in Column 2, 62.2 % of the total variation corresponds to inventor fixed effects while 7.2 % of the total variation corresponds to firm fixed effects. Inventor fixed effects are about eight times as important as firm fixed effects. Overall, the stark differences in explanatory power between inventor and firm fixed effects for patent counts and citations per patent indicate that innovation performance is largely driven by inventor fixed effects. These results are consistent with the findings of Bhaskarabhatla et al. (2021), who also show the importance of inventors' inherent ability or time-invariant characteristics—compared to firms' time-invariant characteristics—in shaping innovation output.

In Table 3, Columns 3 and 4, we examine exploitative and exploratory scores to gauge innovation strategy. The results reported in Column 3 show that the relative importance of inventor and firm fixed effects is about 1.4 and those reported in Column 4 are about 1.3; this indicates that the explanatory powers of inventor fixed effects and firm fixed effects are very close in explaining innovation strategy. Table 3 also reports the F-statistics for the joint significance of both fixed effects and reports the individual significance of fixed effects for inventor or firm. The F statistics all consistently reject the null hypothesis that these coefficient estimates are jointly zero.

The explanatory power of all control variables—except that of inventor fixed effects and year fixed effects—is different with different dependent variables: 35.4 % for patent counts, 21.3 % for citations per patent, 11.7 % for the exploitative score and 15.6 % for the exploratory score. When inventor (or firm) fixed effects is included, the adjusted R-squared is increased. For example, when the dependent variable is Ln*Patent,* the adjusted R-squared is increased by 1.9 percentage points if firm fixed effects are added and by 8.4 percentage points if inventor fixed effects are added. The extent of the adjusted R-squared increment corresponding to inventor (or firm) fixed effects is consistent with our results above on the relative importance of inventor and firm fixed effects. In the example of Ln*Patent*, the ratio of the increment of adjusted R-squared when including inventor fixed effects to the increment of adjusted R-squared when including firm fixed effects is about 4.4, which is close to our estimates above.

Comparing the findings in Table 3, Columns 1 and 2 with the findings in Table 3, Columns 3 and 4, we find the evidence to be consistent with Hypothesis 1. The result suggests that the firm's organizational capital has a relatively more important impact on innovation strategy than on innovation performance. While inventors possess tacit knowledge and can take this knowledge to new firms, it appears that the collaborative team environment of a firm has a significant impact on the novelty of an idea is and on how ideas are selected. That is, firms will eventually have a greater impact on innovation strategy.

Additional tests conducted to check the robustness of our baseline findings are reported in the Online Appendix: (1) Table A1 exhibits the results of robustness checks for Table 3 when implementing the MDV and AKM methods on the largest group of the connectedness sample. All these results are qualitatively similar to what we obtained in Table 3. (2) We redo our main tests by using alternative specifications without covariate controls for firms and inventors and report the results in Table A2; (3) To alleviate the concern about mechanic truncation of the career lengths of inventors—given that inventors will not enter the sample until they start to file patents—we reexamine our main results by arbitrarily adding years (e.g., 1, 3, 5 years) back to an inventor's career preceding or following our original sample, as shown in Table A3; (4) To alleviate the concern about using log transformation as the dependent variable (Cohn et al., 2022), we redo our main tests by using an alternative dependent variable—raw counts for patent numbers and citations, as shown in Table A4; (5) We reexamine our main regressions on subsamples of industries with different levels of secrecy, as shown in Table A5; (6) To exclude the threat of employee synergy, we conduct a robustness check on the subsample of inventors who are solo authors, as shown in Table A6; (7) To alleviate the concern that firms may patent across diverse functional areas, we redo our tests on a subsample of firms with low functional area diversity, as shown in Table A7; (8) We redo our main tests by dropping

[12] The final connected sample entered into AKM e9

continuation patents in the sample (Bhaskarabhatla et al., 2021) and exhibit the results in Table A8. Our main results survive all the above robustness tests. In the Online Appendix we also report tests conducted to investigate the heterogeneity of inventor fixed effects. As shown in Figs. A1 and A2 and in Table A9, there is high variation in inventor fixed effects both in the connected sample and in the largest group. In Table A10, we examine the subsample of hi ennothX

Table 4

Impact of IDD ruling

This table reports the differences in the mean ratios of *Inventor FE/Firm FE* before and after the IDD rulings in subsamples of states that ever had passed the IDD rulings. The test specifications follow Table 3, but we focus on a subsample of states that had positive IDD rulings passed. Specifically, we follow Contigiani et al. (2018) and Gu et al. (2022) to classify states that ever had positive IDD rulings before 2003: Illinois (1995), New York (1997), Washington (1997), Utah (1998), and Iowa (2002). We use a bootstrap procedure to compute the mean and standard errors of the estimates. We resample the original observations 300 times and then re-estimate the inventor FE/Firm FE before and after IDD rulings using the AKM estimation. We exclude draws with negative covariances because it becomes difficult to interpret the numbers as shares (Cornelissen et al. 2008). Column 1 uses the natural logarithm of one plus the adjusted number of patents as the dependent variable and Column 2 uses the natural logarithm of one plus the adjusted number of citations per patent as the dependent variable (zero if no patents filed by an inventor of a year). Columns 3 and 4 use the *Exploit* and *Explore*

bootstrapped sample for the ratios. Therefore, there are no significant differences in the ratio between inventor fixed effects and firm fixed effects before and after the passage of IDD rulings.

Table 4, Columns 3 and 4 focus on innovation strategy. The results indicate a drop in the ratio of inventor fixed effects over firm fixed effects after the adoption of the IDD. The differences in the ratios pre and post the IDD shock are significant. The findings are consistent with Hypothesis 2, which suggests that firms' role in innovation strategy increases when there is a reduction in labor mobility.[14]

In the Online Appendix Table A13, we conduct robustness checks of Table 4 and find consistent results when we remove the controls, focus on a subsample period after 1990, or winsorize the regression variables in the bootstrap procedures.

5. Inventor-firm matching

So far, we have shown that the importance of firms relative to that of inventors is greater for innovation strategy than for innovation performance, and that the importance of firms relative to that of inventors for innovation strategy further increases when labor mobility is reduced. Given the importance of firms in determining innovation strategy, how do inventors choose the firms where they work? To better understand this, in this section we focus our attention on the match between human capital and firm capabilities.

Using the standard AKM approach to study the match between inventors and firms has its drawbacks. That is, inventor fixed effects are estimated as an average for the full sample under the AKM method. As a result, it is difficult to extract inventor fixed effects that are affected by the firms to which they move. To solve this, we follow Bhaskarabhatla et al. (2021) and use a rolling-window approach for the years 1985 through 2006. We construct a 10-year window by including the previous nine years and use the AKM method to re-estimate Eq. (1) with the 10-year window subsample. We obtain the time-varying inventor fixed effects using this approach.

5.1. Which inventors move?

After obtaining the individual effects estimated through the rolling-window algorithm, we first use them to investigate which inventors are more likely to move to another firm. Following Bhaskarabhatla et al. (2021), we define a mobility indicator y_{ijt} as the dependent variable that equals one if inventor i of firm j in year t will depart the current firm in future years. The key independent variable is the vector of indicators for the decile of the estimate of inventor i's individual effect estimated from the 10-year rolling window ending in year t. In the regression, we control for a vector of indicators for the decile of the estimated firm effect in the 10-year rolling window, the average estimated individual effect for all employees in firm j of year t, the current tenure of individual i as well as the year and two-digit SIC code fixed effects.

Fig. 1 plots the estimated coefficients of inventors' tendency to leave a firm on the inventor effect decile. Consistent with Bhaskarabhatla et al. (2021), we show in Panels A and B that inventors in high deciles for innovation performance (patents and citations) are more likely to move and that their probability of moving increases along with the inventor deciles. However, in terms of innovation strategy, inventors' incentives to move are a bit different. In Fig. 1, Panel C we find a sharp drop in the moving likelihood for inventors in the top decile of exploitativenss but little difference in the incentive for inventors in other deciles to move. Potentially this is because inventors who are particularly exploitative have their human capital tied closely to the firms' organizational capital, and therefore are less likely to leave the firm.

In Fig. 1, Panel D we find that the likelihood of moving is an inverted U shape with regard to inventors' tendency to explore. Specifically, inventors at the bottom tercile of exploratory scores are least likely to move. Given that the least exploratory inventors are the most exploitative, this finding is consistent with the results in Panel C: the most exploitative inventors are least likely to move from the firm. In addition, among inventors above the median exploratory score level, those who are more exploratory seem less likely to move. This may indicate that firms try to retain inventors who are highly novel, as indicated by their high exploratory scores.

5.2. *How are inventors matched to firms?*

After investigating which inventors tend to move, we study how the inventors' estimated capabilities, derived from the rolling-window procedure, correlate with firm characteristics. To do so, we run a regression with the dependent variable being the estimated individual effect from the 10-year rolling window in the year before moving. The key independent variable is either the fixed effect of the firm to which the inventor moves, obtained from the 10-year rolling-window AKM regression, or the average inventor fixed effects of the firm to which the inventor moves. Control variables include firm characteristics.

Table 5 reports the regression results. The findings about innovation performance are consistent with Bhaskarabhatla et al. (2021). We find that firm-specific innovation performance correlated negatively with the innovation performance of inventors moving into the firm. However, those inventors who move—and who also have higher innovation performance— join firms where their coworkers also have, on average, higher innovation performance. Our findings about innovation strategy are similar. We show that the innovation strategy of inventors correlates negatively with the innovation strategy of firms they join, but correlates positively with the innovation strategy of inventors in those firms.

Our findings are consistent with the negative assortative matching model proposed by Bhaskarabhatla et al. (2021). Three key assumptions in the model lead to negative assortative matching. First, inventors care about both wage and innovation output. Second, the marginal utility from innovation is decreasing. That is, the increase in innovation matters more for an inventor (a firm) with higher innovation performance than for an inventor (a firm) with lower innovation performance. Third, inventors' innovation is boosted more in firms with higher innovation performance. As a result, in equilibrium, inventors with high innovation performance are matched to firms with low innovation performance where innovation is boosted less but wages are higher. Inventors with low innovation performance are matched to firms with high innovation performance where innovation is boosted more but wages are lower.[15]

6. Endogenous matching between inventors and firms

The results in Section 5 indicate that a subset of inventors and firms exhibits a pattern of sorting in the market. But to what extent does it impact our main conclusions? Unfortunately, the AKM approach assumes the separability of inventor effects and firm effects, and it may not capture the endogenous matching component. To partially alleviate this concern, we follow Card et al. (2013) to undertake a test to assess whether the AKM method is appropriate in our innovation setting and provide evidence on whether the inventor-firm matching issue jeopardizes our main results. Table 6 reports added explanatory power from adding inventor-firm matching effects. The results suggest that inventor-firm-matching fixed effects explain a very small portion of innovation performance and style, thus validating the applicability of the AKM method here.

Furthermore, given that we rely on inventors' moves across firms to estimate the relative importance of firms' organizational capital and inventors' human capital, we gauge how our results vary in subsamples that are subject to different levels of endogenous moving. We perform several sets of tests to alleviate the matching concern and show our main results are robust.

First, if inventors move to firms with similar performance, we could underestimate the contribution of firms' organizational capital to the inventors' innovation output. This is because we would attribute little innovation to firm fixed effects if inventors only move to firms with similar performance, given that it is less likely for us to observe a change in the inventors' patenting around the moves. However, firms could have contributed more to innovation had the inventors moved to the firms with different performance. To illustrate that this endogenous matching would have minimal impact on our estimation, we compare two different groups: inventors who move to firms with *similar* innovation performance and inventors who move to firms with *different* innovation performance.

Table 7, Panels A and B exhibit the results. Panel A reports inventors who move to firms with similar performance, which we define

[15] A similar rationale applies to innovation strategy. We can make three parallel assumptions in a model that also leads to negative assortative matching of innovation strategy. First, inventors care about both wage and innovation output. Second, for a inventor who scores high in a certain innovation strategy, the marginal utility from innovation decreases in this type of innovation. For example, a high-exploratory inventor or firm has marginal utility that decreases in exploratory innovation. Third, inventors' innovation is boosted more when they are matched with the same type of firm in terms of innovation strategy. As a result, high-exploratory (high-exploitative) inventors will match to low-exploratory (low-exploitative) firms that are less likely to boost exploratory (exploitative) patents but provide higher wages. Similarly, low-exploratory (low-exploitative) inventors will match with high-exploratory (high-exploitative) firms that provide a greater boost to exploratory (exploitative) patents but provide lower wages.

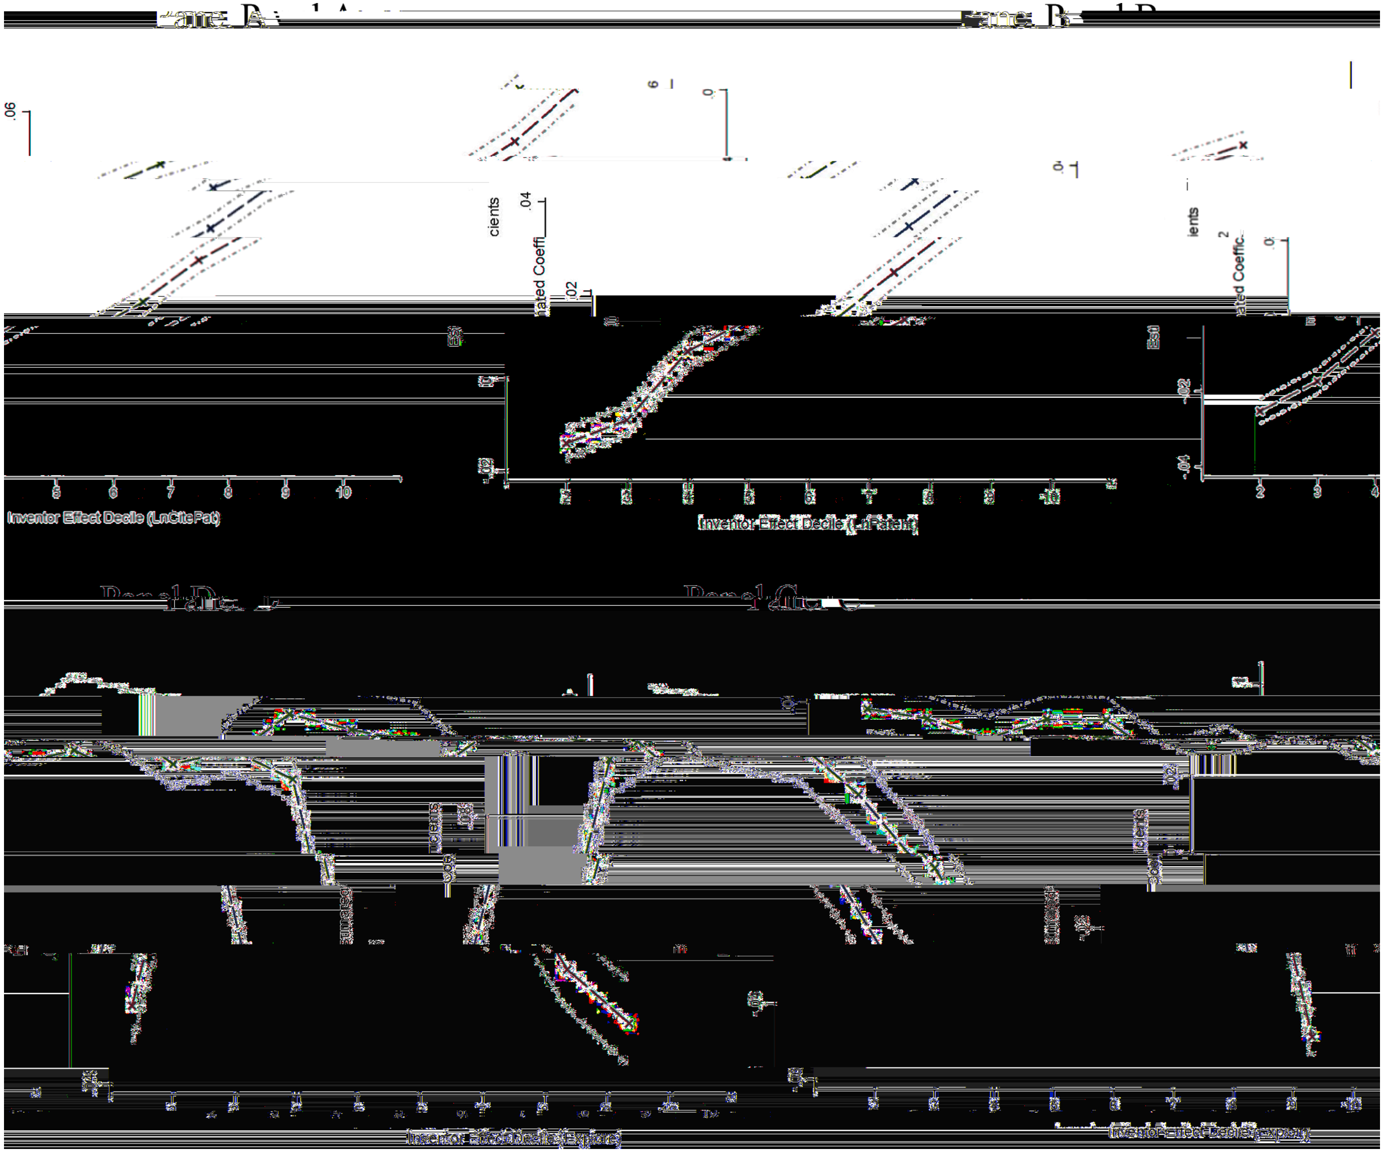

Fig. 1. Inventor mobility by decile of inventor effects. These figures plot the estimated coefficients for the inventor effect decile estimates in the inventor-firm matching models following Bhaskarabhatla et al. (2021). Four panels correspond to four different dependent variables: the log of one plus the adjusted number of patents (Panel A), the log of one plus the adjusted number of citations per patent (Panel B), the exploitative score (Panel C), and exploratory score (Panel D).

Table 5
Inventor-Firm Matching
This table reports the OLS regression results with the dependent variable

as those in the bottom quartile of the difference in patenting outcome between the inventor's current and previous firms. Panel B reports inventors who move to firms with different performance, which we define as those in the top quartile of the difference in patenting outcome between the inventor's current and previous firms. We compute a firm's patenting outcome by dividing the total number of filed-and-granted patents in the last three years by the firm's total assets in year *t*.

The results reported in Table 7, Panels A and B show that an inventor moving between firms with similar performance is 9.4 times more important than the firm in explaining patent counts. When moving between firms with different performance, an inventor is 8.9 times more important than the firm in explaining patent counts. For citations per patent, the inventor is also more important than the firm in terms of explanatory power. In line with our main results, firm fixed effects have a similar magnitude to that of inventor fixed effects in explaining the score variation in both subsamples for exploratory and exploitative scores, The minor difference in the explanatory power between inventor fixed effects and firm fixed effects across the subsamples suggests that endogenous moving could not completely account for our main results.

Second, firms that appear very different in patenting performance, may be operationally similar, which still causes the firm fixed effect to be underestimated. For example, if an inventor moves to a firm in an industry different from that of the inventor's previous firm, the new firm is more likely to affect the inventor's patenting by providing a vastly different environment and different access to resources. In contrast, if inventor moves between firms in the same industry, the new firm is less likely to affect the inventor's patenting. Hence, we are more likely to underestimate firm fixed effects in the former situation than in the latter. To address this concern, we classify inventor-moving into two groups, one containing inventors who move within the same industry and one containing inventors who move across industries. Table 7, Panels C and D repeat the main tests on these two subsamples. We find that inventor fixed effects are 6.1 times as important as firm fixed effects in explaining patent counts when inventors move across different industries, and 12.9 times as important as firm fixed effects in explaining patent counts when inventors move within the same industry. In terms of citations per patent, the inventor also appears to be much more important than the firm. But for exploitative and exploratory scores, the explanatory power of inventor fixed effects is comparable to that of firm fixed effects across both subsamples. As a matter of fact, the ratio of the explanatory power of inventor fixed effects and firm fixed effects for inventors who move across different industries is very close to our baseline findings. Thus, the results suggest that inventors who move within the same or different industries have limited effects on our results.

Third, if an inventor moves to a firm that is a better match, then part of the change in innovation output would come from the matching effect, rather than from the firm-specific fixed effects in innovation. However, under the AKM method, we would attribute the change in innovation output to firm-specific fixed effects. In other words, this matching possibility may lead to an overestimation of firm fixed effects. To address this concern we first control for the previous firm and inventor performance in our regression, under the assumption that part of the matching effect is reflected in firm and inventor quality. We then examine a subsample of movers who are less likely to be subject to the matching effect. Table 7, Panel E considers inventors who moved between firms and who, within three years, had a change in patent output that was within 25 %.[16] While the maximum change was 25 %, the mean change in patent output was only 0.5%, and the median change was 0 %. Given such a small change in patent output, the matching issue should bear very little effect in this subsample. The results are consistent with our main finding that the inventor plays a more important role than does the firm in explaining innovation performance, and plays a role that is about equally important in explaining innovation strategy.

Overall, the test results in this section suggest that our main findings are not driven completely by inventors' endogenous moving. We also note that the matching issue is not specific to our method but is present in any data that matches labor with capital. Finally, we acknowledge that the need for caution in interpreting and generalizing our results, because endogenous moving by inventors appears to play some role in our findings and cannot be ruled out completely.

7. Conclusion and discussion

In this paper, we study the role of inventor- and firm-specific heterogeneities in promoting a firm's innovation output. We find that

Table 7

Different types of movers

This table reports the subsample analysis of results from using the AKM method to estimate both inventor and firm fixed effects in subsamples containing different types of movers. The connected group is constructed after restricting to this set of movers. The subsample of inventors who move to firms with "similar innovation performance" considers those who moved to a firm in the bottom quartile of the absolute difference in patent output between the new firm and the previous firm. The subsample of inventors who moved to firms with "different innovation performance" considers those who moved to a firm in the top quartile of the absolute difference in patent output between the new firm and the previous firm. We compute the patent output of a firm by the total number of filed-and-granted patents in the last three years divided by the total assets of the firm in year *t*. The "different industries" subsample considers inventors who moved across different industries. The "same industry" subsample considers inventors who moved within the same industries. Industry is defined based on 3-digit SIC codes. The subsample of inventors who moved with a "minor change of output" includes those who moved between firms with a change in patent output within 25 % in three years. Panels A–E correspond to the estimation with the five different subsamples described above. The connected group is constructed after restricting to this set of movers. Inventors who moved more than once were excluded from the subsamples. The estimation is implemented by using the Stata command "felsdvreg" proposed by Cornelissen (2008). Column 1 uses the natural logarithm of one plus the adjusted number of patents as the dependent variable and Column 2 uses the natural logarithm of one plus the adjusted number of citations per patent as the dependent variable (zero if no patents were filed by an inventor in a given year). Columns 3 and 4 use the *Exploit* and *Explore* indices as the dependent variables (missing value is assigned if no patents were filed by an inventor of a year). Definitions of variables are defined in Appendix B. ***, ** and * indicate significance at the 1, 5, and 10 percent levels, respectively.

Panel A: Moving to firms with similar innovation performance				
Dependent Variable	LnPatent (1)	LnCitePat (2)	Exploit (3)	Explore (4)
Relative importance of inventor and firm fixed effects (percentage of R-sq. explained)				
Contribution of Inv. FE	0.299 (54.34 %)	0.282 (62.59 %)	0.372 (60.91 %)	0.347 (52.96 %)
Contribution of Firm FE	0.032 (5.76 %)	0.024 (5.23 %)	0.195 (32.00 %)	0.227 (34.76 %)
Inventor FE / Firm FE	9.431	11.960	1.904	1.524
F-test on fixed effects				
Joint F-statistic	2.00***	1.95***	2.19***	2.49***
Inventor FE F-statistic	1.76***	1.72***	1.24***	1.33***
Firm FE F-statistic	2.39***	2.20***	2.08***	3.01***
Adj. R-squared	0.454	0.333	0.392	0.458
# Movers	7351	7351	5131	5131
# Stayers	159,089	159,089	151,457	151,457
# Firms	1762	1762	1435	1435
Year FE	Yes	Yes	Yes	Yes
Functional Area FE	Yes	Yes	Yes	Yes
Observations	954,201	954,201	440,857	440,857
Panel B: Moving to firms with different innovation performance				
Dependent variable	LnPatent (1)	LnCitePat (2)	Exploit (3)	Explore (4)
Relative importance of inventor and firm fixed effects (percentage of R-sq. explained)				
Contribution of Inv. FE	0.301 (54.10 %)	0.281 (61.95 %)	0.345 (56.53 %)	0.332 (50.89 %)
Contribution of Firm FE	0.034 (6.11 %)	0.021 (4.52 %)	0.223 (36.49 %)	0.239 (36.59 %)
Inventor FE / Firm FE	8.855	13.719	1.549	1.391
F-test on Fixed Effects				
Joint F-statistic	1.96***	1.91***	2.24***	2.53***
Inventor FE F-statistic	1.78***	1.73***	1.26***	1.35***
Firm FE F-statistic	3.34***	3.00***	2.79***	3.75***
Adj. R-squared	0.461	0.337	0.399	0.463
# Movers	6769	6769	5117	5117
# Stayers	157,624	157,624	149,602	149,602
# Firms	1517	1517	1270	1270
Year FE	Yes	Yes	Yes	Yes
Functional Area FE	Yes	Yes	Yes	Yes
Observations	944,291	944,291	444,064	444,064
Panel C: Moving across different industries				
Dependent Variable	LnPatent (1)	LnCitePat (2)	Exploit (3)	Explore (4)
Relative importance of inventor and firm fixed effects (percentage of R-sq. explained)				
Contribution of Inv. FE	0.289 (52.64 %)	0.273 (61.35 %)	0.339 (55.57 %)	0.326 (50.00 %)
Contribution of Firm FE	0.047 (8.56 %)	0.032 (7.19 %)	0.227 (37.21 %)	0.246 (37.73 %)
Inventor FE / Firm FE	6.149	8.531	1.493	1.325
F-test on fixed effects				
Joint F-statistic	1.94***	1.88***	2.16***	2.46***
Inventor FE F-statistic	1.76***	1.70***	1.24***	1.33***
Firm FE F-statistic	2.52***	2.61***	2.45***	3.47***
Adj. R-squared	0.452	0.326	0.391	0.458
# Movers	14,443	14,443	7473	7473
# Stayers	165,613	165,613	152,929	152,929
# Firms	2997	2997	1903	1903
Year FE	Yes	Yes	Yes	Yes

(*continued on next page*)

Table 7 (*continued*)

Panel A: Moving to firms with similar innovation performance				
Dependent Variable	LnPatent (1)	LnCitePat (2)	Exploit (3)	Explore (4)
Functional Area FE	Yes	Yes	Yes	Yes
Observations	1036,775	1036,775	453,325	453,325
Panel D: Moving within the same industry				
Dependent Variable	LnPatent (1)	LnCitePat (2)	Exploit (3)	Explore (4)
Relative importance of inventor and firm fixed effects (percentage of R-sq. explained)				
Contribution of Inv. FE	0.310 (55.96 %)	0.285 (62.91 %)	0.430 (69.47 %)	0.450 (68.70 %)
Contribution of Firm FE	0.024 (4.33 %)	0.019 (4.19 %)	0.147 (23.75 %)	0.123 (18.78 %)
Inventor FE / Firm FE	12.917	15.000	2.925	3.659
F-test on fixed effects				
Joint F-statistic	1.97***	1.92***	2.32***	2.56***
Inventor FE F-statistic	1.78***	1.73***	1.29***	1.35***
Firm FE F-statistic	2.24***	2.23***	2.99***	3.97***
Adj. R-squared	0.457	0.335	0.408	0.463
# Movers	11,107	11,107	7033	7033
# Stayers	156,462	156,462	138,270	138,270
# Firms	2135	2135	1532	1532
Year FE	Yes	Yes	Yes	Yes
Functional Area FE	Yes	Yes	Yes	Yes
Observations	958,953	958,953	412,490	412,490
Panel E: Moving with a minor change of output				
	(1)	(2)	(3)	(4)
Dependent Variable	LnPatent	LnCitation	Exploit	Explore
Relative importance of inventor and firm fixed effects (percentage of R-sq. explained)				
Contribution of Inv. FE	0.290 (51.69 %)	0.269 (58.35 %)	0.368 (59.84 %)	0.332 (50.00 %)
Contribution of Firm FE	0.029 (5.17 %)	0.026 (5.64 %)	0.197 (32.03 %)	0.248 (37.35 %)
Inventor FE / Firm FE	10.000	10.346	1.868	1.339
F-test on fixed effects				
Joint F-statistic	2.01***	1.97***	2.16***	2.54***
Inventor FE F-statistic	1.82***	1.79***	1.28***	1.36***
Firm FE F-statistic	2.02***	2.08***	1.15**	1.95***
Adj. R-squared	0.470	0.349	0.403	0.478
# Movers	2938	2938	910	910
# Stayers	115,790	115,790	101,990	101,990
# Firms	1314	1314	581	581
Year FE	Yes	Yes	Yes	Yes
Functional Area FE	Yes	Yes	Yes	Yes
Observations	700,903	700,903	292,440	292,440

time-invariant inventor fixed effects explain a majority of the variation in innovation performance in terms of patent counts and citations, but inventor fixed effects play a relatively less important role in explaining innovation strategy in terms of patent exploratory and exploitive scores. We also use the staggered adoption of the inevitable disclosure doctrine as plausibly exogenous shocks to labor mobility, and show that inventors' contributions to innovation performance increase and contributions to innovation strategy decrease when labor mobility is restricted. Reduced mobility means that firms have greater control of the innovation strategy because inventors have fewer outside options.

We need to bear in mind three caveats when generalizing our results. First, similar to all other studies that use movements of individuals (e.g., executives, venture capitalists, bankers, employees, etc.) as an identification strategy, our empirical setting is subject to the concern that movement of inventors could be endogenous. We intend to show the average effect of moving across firms on inventors who actually move. We are silent regarding why inventors move, although additional tests suggest that endogenous moves do not alter our conclusion. Second, we are only able to capture the contribution to innovation from movers and stayers in firms with at least one mover. If a firm has no inventor who moves, we cannot separate the inventor's contribution from the firm's. Finally, because innovation is human capital intensive, we are likely to attribute more innovation contributions to inventors. Hence, our findings likely represent a lower bound of firms' contribution to innovation.

CRediT authorship contribution statement

Tong Liu: Conceptualization, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization. **Yifei Mao:** Conceptualization, Writing – original draft, Writing – review & editing, Visualization, Supervision. **Xuan Tian:** Conceptualization, Resources, Writing – review & editing, Supervision, Project administration.

Acknowledgment

Tian acknowledges financial support from the National Natural Science Foundation of China (Grant No. 71825002).

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.jempfin.2023.101435.

Appendix A. Details in sample construction

We match the HBS patent and inventor database with the NBER patent citation database, following four steps:

(1) We sort all patents in the NBER database into two subsets based on the number of assignees for each patent: one subset (hereafter Subset A) contains all patents owned by a single assignee while the other (Subset B) includes patents owned by multiple assignees. The company affiliations are easily identified for inventors with patents in Subset A. We use the patent number to match all patents in Subset A with the HBS database; the result is, Set A, a set of 6270,074 matched inventor-patent observations.
(2) We divide Subset B, which consists of all unmatched observations after step one, into two groups: Subset B1 is composed of inventor-patent observations in which each patent was filed by a single inventor; the other group, Subset B2, collects the remaining inventor-patent observations with multiple inventors. We then match all observations in Subset B1 with Set B by patent number. This leads to 22,555 inventor-patent-assignee observations, which correspond to 11,461 inventor-patent observations in Subset B1 as each patent may be held by several assignees. We then determine one assignee for each observation based on the matched information in Subset A, that is, we designate a unique assignee to an inventor in the year that patent is granted, if this assignee in Subset A coincides with the one assignee for which the inventor is recognized as working in the same year. If we are able to identify multiple assignees through the above method, we exploit the location information to pin down the assignee for these patent filings. In another extreme case, if we find no appropriate assignees in Subset A using the above method, we also exploit the location information to determine the assignee. If several assignees have the same location, we randomly choose one. Otherwise, we relax our search criteria and select an assignee sharing the same state, New Jersey. In this way, we can pin down all assignees for the 11,461 observations.
(3) For all observations in Subset B2, patents are filed by multiple inventors and belong to various assignees. Using the patent number, we join these with Subset B to form all pairwise combinations and then select one assignee for each inventor-patent observation. The selection procedure is identical to that in step 2. As a result, assignees for the 250,168 inventor-patent instances in Counterpart B2 can be identified.
(4) Combing all observations obtained from the three steps above, our final matched sample consists of 6531,703 inventor-patent observations whose assignee can be uniquely identified. Then we identify the company affiliations for all inventors throughout their careers with the assistance of 6531,703 matched inventor-patent observations. If all patents filed by an inventor of a year belong to a single assignee, we assume that the inventor was hired by this particular assignee during the patent-filing years. Frequently patents filed by an inventor of one year are owned by different assignees. For instance, in one year, two patents of an inventor are claimed by assignee A while, in the same year, five other patents from the same inventor are claimed by assignee B. In this case, it is reasonable to assume that the inventor was employed by the assignee that owns the majority of the inventor's patents from that year. Particularly, when an inventor files the same number of patents for both assignee A and B in certain year, we utilize the inventor's employment information of last year to help us identify — if the inventor worked for assignee A (B) last year, we presume the employer was the same this year.[17] Otherwise, we randomly pick one assignee out of the two for this inventor. This procedure leads to about 4251,546 inventor-year observations.

For our analysis, we augment our inventor-year sample in a time order by filling all year gaps for inventors who appear in the patent database but who do not have patents in the gap years. For example, if an inventor filed patents in 1986 and 1991, our sample captures only the inventor's performances in 1986 and 1991. We expand the observations between 1986 and 1991 for this inventor by assigning zero to patent counts and to citations per patent.[18] This method comes with a caveat: how can we can accurately identify inventors' employer in gap years? Following the example above, it is intuitive and easy to determine which company employed the inventor between 1986 and 1991 if the patents filed in both years are owned by the same company. It would be difficult to determine the employer if the patents filed in 1986 and 1991 belong to different companies. In other words, how do we decide the company affiliations of a mover for the transition years where we have no observations of the mover's patent filing? We assume that the inventor

[17] Admittedly, this is an ad hoc assignment. To alleviate this concern, we repeat our analysis with different assigning methods. For example, we use the inventor's employment information in the subsequent year, that is, if the inventor worked for assignee A (B) the next year, we assume the employer during this year was A (B). We also tried picking an assignee for the inventor randomly. These alternative methods do not alter the nature of the results.

[18] We assign missing values to metrics that measure inventors' innovation strategy (exploratory ratio and exploitative ratio) in years with no patent filings.

belongs to the old company in the first half of the transition years (1987 to 1988) and to the new company in the second half (1989 to 1990).[19] This procedure leads to 7445,855 inventor-year observations in our augmented sample.

Appendix B. Definition

Chemmanur, T., Ertugrul, M., Krishnan, K., 2019. Is it the investment bank or the investment banker? A study of the role of investment banker human capital in acquisitions. J. Financ. Quant. Anal. 54, 587–627.
Cho, C., Halford, J., Hsu, S., Ng, L., 2016. Do managers matter for corporate innovation? J. Corp. Finance 36, 206–229.
Coase, R.H., 1937. The nature of the firm. Economica 4, 386–405.
Coase, R.H., 1988. The nature of the firm: origin, meaning, influence. J. Law Econ. Org. 4, 3–17.
Cohen, L., Malloy, C., Frazzini, A., 2012. Hiring cheerleaders: board appointments of "independent" directors. Manag. Sci. 58, 1039–1058.
Cohn, J.B., Liu, Z., Wardlaw, M.I., 2022. Count (and count-like) data in finance. J. Financ. Econ. 146, 529–551.
Contigiani, A., Hsu, D.H., Barankay, I., 2018. Trade secrets and innovation: evidence from the "inevitable disclosure" doctrine. Strat. Manag. J. 39, 2921–2942.
Cornaggia, J., Mao, Y., Tian, X., Wolfe, B., 2015. Does banking competition affect innovation? J. Financ. Econ. 115, 189–209.
Cornelissen, T., 2008. The Stata command felsdvreg to fit a linear model with two high-dimensional fixed effects. Stata J. 8, 170–189.
Correia, S., 2014. REGHDFE: stata module to perform linear or instrumental-variable regression absorbing any number of high-dimensional fixed effects. Stat. Softw. Comp.
Ewens, M., Rhodes-Kropf, M., 2015. Is a