informs
http://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
*Articles in Advance*, pp. 1–22
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

# Are Inventors or Firms the Engines of Innovation?

**Ajay Bhaskarabhatla,[a] Luis Cabral,[b] Deepak Hegde,[b] Thomas Peeters[a]**

[a] Erasmus School of Economics, Erasmus University, 3062 PA Rotterdam, Netherlands; [b] Stern School of Business, New York University, New York, New York 10012

**Contact:** bhaskarabhatla@ese.eur.nl, https://orcid.org/0000-0002-0445-1373 (AB); luis.cabral@nyu.edu, https://orcid.org/0000-0003-2763-2422 (LC); dhegde@stern.nyu.edu, https://orcid.org/0000-0003-3888-7447 (DH); peeters@ese.eur.nl, https://orcid.org/0000-0001-8507-5164 (TP)

**Received:** February 24, 2018
**Revised:** April 29, 2019; December 21, 2019; February 18, 2020; March 2, 2020
**Accepted:** March 11, 2020
**Published Online in Articles in Advance:** October 7, 2020

https://doi.org/10.1287/mnsc.2020.3646



**Abstract.** In this study, we empirically assess the contributions of inventors and firms for innovation using a 37-year panel of U.S. patenting activity. We estimate that inventors' human capital is 5–10 times more important than firm capabilities for explaining the variance in inventor output. We then examine matching between inventors and firms and find highly talented inventors are attracted to firms that (i) have weak firm-specific invention capabilities and (ii) employ other talented inventors. A theoretical model that incorporates worker preferences for inventive output rationalizes our empirical findings of negative assortative matching between inventors and firms and positive assortative matching among inventors.

**History:** Accepted by Ashish Arora, entrepreneurship and innovation.
**Supplemental Material:** The data files and online appendices are available at https://doi.org/10.1287/mnsc.2020.3646.

**Keywords:** human capital • innovation • matching

## 1. Introduction

Innovation is a key driver of firms' productivity and competitive advantage (Griliches 1984, Jaffe 1986, Kogan et al. 2017). Yet, we know little about the sources of innovation at firms. In particular, we do not know to what extent two key inputs to innovation—human capital and firm-specific capabilities—influence the inventiveness of inventors. Aspects of modern innovation, such as the deep complexity of new products and the burden of knowledge required to invent them, suggest organizational capital, embedded in firms' structure, processes, and culture, are essential for innovation (Cohen et al. 2000, Jones 2009). At the same time, the tacit nature of breakthrough knowledge and the well-documented achievements of heroic inventors suggest human capital is crucial for innovation (Arora 1995, Zucker et al. 2002, Agarwal et al. 2004, Fallick et al. 2006, Singh and Agrawal 2011, Palomeras and Melero 2010, Schilling 2018).

How much do the human capital of inventors and the innovation capabilities of firms contribute to inventors' productivity? Are inventors who appear highly productive matched with firms that have the best innovation capabilities or is their *secret sauce* embedded in their own human capital? In this paper, we empirically and theoretically study these questions. Indeed, if the ability to invent rests with firm-specific structure, culture, and routines that are hard to transfer across organizational boundaries, then inventors may be substitutable, and firms should develop capabilities that enhance innovation. If, on the other hand, human capital is critical for innovation, then firms' innovativeness will depend on their ability to screen, attract, and retain talented workers. Thus, shedding light on two key inputs to innovation—human capital and firm capabilities—and the relationship between the inputs has direct implications for managers and for theories of innovation and competitive advantage.

Disentangling the contributions to inventor productivity of human capital and firm capabilities poses two main challenges. First, firms deploy a combination of human capital and firm capabilities to tasks, and the two factors' contributions to outcomes are difficult to separate. Second, worker productivity is a consequence of endogenously matched human capital and firm capabilities, complicating its identification through standard regression techniques. We tackle these challenges by assembling data on all U.S. patents granted by the U.S. Patent and Trademark Office (USPTO) between 1973 and 2010. Patents record the identity of their inventors and firm assignees, allowing us to construct a 37-year panel of each patenting firm's and patenting inventor's annual patenting output—our proxy for inventive performance. To tease apart the relative contributions of firms and inventors for innovation, we apply the identification strategy of Abowd et al. (1999) (henceforth AKM). This strategy requires that we use a subsample of inventors who work at firms connected to one another by moving inventors. After merging in the Compustat data and excluding continuation patents, we retain about 709,000 inventors at 2,500 U.S. publicly listed firms, with detailed information on time-varying firm characteristics that influence inventor performance from 1976 to 2010. We leverage this AKM sample to tease apart the contributions of inventors and firms to innovation and thus address the first empirical challenge described above. Unlike a traditional fixed effects approach (Bertrand and Schoar 2003), which can only estimate worker fixed effects for workers who move (only about 26% in our sample), the AKM approach allows us to estimate employer and employee effects even for the nonmoving inventors. Conceptually, we first pin down the fixed effects of moving inventors and firms connected by these movers and then use this information to identify the fixed effects of nonmovers.

Applying the AKM method to our sample reveals inventor-specific fixed effects explain 18%–37% of the observed variance in inventors' patenting performance. In contrast, only 2%–7% of the overall variance in inventor productivity is explained by firm fixed effects. The observed firm-level variables, such as age, size, patent stock, and research and development (R&D) intensity, explain another 1%–8% of the overall variance. These results suggest inventor-specific human capital explains much of the variance in inventor output, which echoes the relative importance of workers documented in prior literature using the AKM method and its extensions to estimate wage equations (Abowd et al. 1999, Andrews et al. 2008, Gruetter and Lalive 2009, Bonhomme et al. 2019). Firm-specific innovation capabilities add to the innovation performance of inventors, but their contribution pales in comparison with that of the human capital of inventors.

The AKM method extracts the time-invariant effects of employer and employee capabilities on worker output (see Iranzo et al. 2008, Graham et al. 2012, Card et al. 2013, Ewens and Rhodes-Kropf 2015, and Peeters et al. 2020 for applications of the approach to other contexts). A drawback of this approach is that it estimates worker and firm capabilities using worker performance over the entire sample period, making it difficult to study matching between employers and employees as a function of their capabilities (Bonhomme et al. 2019). We address this challenge by implementing a *rolling window* strategy for AKM estimations (as in Card et al. 2013). That is, we estimate standard AKM fixed effects for inventors and firms in progressive time windows, allowing the estimates to vary across the windows. For example, we first limit the AKM estimation sample to a 10-year window from 1978 (the first year of our estimation sample after including lagged explanatory variables for periods $t-2$ and $t-1$) through 1987 and estimate the firm and inventor effects based on movements within this window. These estimates are not contaminated by changes to inventors and firms after 1987, and we use them to examine how the fixed effects predict movements in 1988. Next, we draw a new subsample of 10 years—by rolling the window one year—from 1979 through 1988 and estimate AKM firm and inventor effects based on moves within this new window. These estimates are used to examine how the fixed effects predict movements in 1989 and so on. Conditional on the inventor moving, we find a negative correlation between inventor human capital and the innovation capability of her destination firm but a positive correlation with the average human capital of the inventor's coworkers at the destination firm. These patterns hold in the larger sample (including movers and nonmovers): high human capital inventors appear more likely to be placed at firms with low innovation capabilities but other high human capital workers.

This finding of negative assortative matching between worker and firm innovation capabilities is difficult to square with traditional models that suggest positive assortative matching between workers and firms (Becker 1973). One possible explanation is that our estimates have a downward bias, either caused by the presence of match quality and endogenous inventor mobility or as the result of limited mobility among moving inventors in the connected worker-firm network (Andrews et al. 2008). We implement a battery of empirical checks to examine the implications of endogenous mobility or limited mobility bias (Lazear et al. 2015, Abowd et al. 2019, Jochmans and Weidner 2019). We find that these potential downward biases are unlikely to explain the negative assortative matching between inventors and firms we uncover.

An alternative explanation for our results is that the standard assumption of supermodular match payoff functions fails to apply in the present case. To explore this possibility, we develop a formal model in which the logarithm of innovation rate is additive in firm and inventor type. This functional form would suggest positive assortative matching between innovators and firms, as the innovation rate is a supermodular function of firm and inventor type.[1] However, a critical assumption of our model is that inventors care about wages and the immediate outcome of their efforts, namely innovation output. Inventors may have intrinsic preferences over innovation output (e.g., publications or patents as in Stern 2004) or value outputs as a signal of their productivity to labor markets (Kline et al. 2019, Melero et al. 2020). Moreover, we assume that inventors' marginal utility of innovation output is diminishing: an additional patent matters more for an inventor with fewer patents than for an inventor with more patents under her belt. Specifically, we assume that the marginal utility of innovation declines at a higher rate than a logarithmic function. We also assume that the firms' profit function is additive in the log of innovation output and wages paid to inventors.

We show that the equilibrium of the matching game (i.e., the core of the game) features negative assortative matching between firms and inventors. This results from the fact that the innovator-firm match function is submodular (even though the innovation function itself is supermodular). In other words, our theoretical result implies that low human capital inventors match with high innovation capability firms. The idea is that the innovation boost provided by firms with high innovation capabilities is particularly valuable for low human capital inventors, who are willing to accept lower wages for the *amenity* of higher innovation rates. By contrast, high human capital inventors cluster at firms with low innovation capabilities. Intuitively, these inventors have less to gain from their employer's innovation capabilities, placing a relatively greater weight on their financial compensation.[2] In terms of the matching literature, the critical step leading to our result is the assumption that workers care about financial compensation and their innovative output and that the marginal utility of innovative output is decreasing. Together, this implies that the firm-inventor matching function is effectively submodular, which in turn leads to negative assortative matching.

Our study makes several contributions to the study and practice of innovation and strategic management. First, research that investigates how firm characteristics, such as incentive schemes or organizational culture, affect innovation, and performance suggests that most variance in performance remains unexplained even after accounting for the effects of firms' characteristics such as industry, business segment, and corporate structure (Nelson and Winter 1982; Rumelt 1984; Wernerfelt 1984; Barney 1986; McGahan and Porter 1997, 2002; Bloom et al. 2013; Martinez et al. 2015). Our results suggest that differences in the human capital embedded in firms accounts for a substantial portion of the unexplained variance, at least in innovation performance. We thus complement the growing body of research on worker characteristics and human resource management by providing one of the first assessments of the relative importance of human capital for innovation at organizations (Hatch and Dyer 2004, Huckman and Pisano 2006, Groysberg et al. 2008, Jones 2009, Lazear 2009, Campbell et al. 2012, Mayer et al. 2012). Indeed, our results call for increased attention of management scholars and managers, focused as yet on firm-level process and routines, toward shifting the focus of analysis to individual workers and human capital (Zucker et al. 2002, Agarwal et al. 2004, Wuchty et al. 2007, Azoulay et al. 2010, Groysberg 2010, Singh and Fleming 2010, Mollick 2012, Agarwal and Ohyama 2013).

Second, we extend important research that considers how workers and firms match to the context of innovation. Consistent with recent studies of routine workers that record departures from the standard positive assortative matching framework (Becker 1973), we provide evidence for negative assortative matching between firm-specific innovation capabilities and human capital (see also Lindenlaub 2017, Eeckhout 2018, Eeckhout and Kircher 2018). We document positive assortative matching among workers—highly talented workers prefer to work at firms with other highly talented workers. Our theoretical framework proposes that these nuanced sorting patterns can arise from the distinct preferences of inventors for innovation output, which they tradeoff against preference for wages (Stern 2004). Our theoretical model thus provides a unique explanation—one rooted in the unique preferences of inventors—to understand how innovative workers choose their employers.

Finally, our empirical and theoretical results suggest that inventor human capital and firm-level innovative capability are strategic substitutes, whereas inventor and coworker human capital are strategic complements. Somewhat counterintuitively then, firms with superior innovation capabilities can profit more by hiring low-type inventors at lower wages than expensive high-type inventors. Nevertheless, our empirical analysis uncovers that much of the variation in inventor output in the data are explained by human capital rather than firm-specific capabilities, suggesting strategies based on strong firm-specific capabilities are rather rare.

## 2. Data Description

### 2.1. Sample Construction

We start with the population of U.S. patents granted during the years 1973–2010, obtained from the USPTO. Limiting the last patent grant year to 2010 allows us at least five years to observe forward citations, our measure of innovation impact, to the latest patents without truncation. We disambiguate inventor names recorded by the USPTO using the procedure outlined in Li et al. (2014) and standardize assignee names using the procedure in Hall et al. (2001). Treating unproductive years of inventors as missing can lead to an overestimation of the contribution of inventors relative to firms. To address this, we include inventor-firm-year observations for unproductive years between an inventor's first and last patenting year in our sample and assign zero patents for them. To further improve the accuracy of inventor moves, following Ge et al. (2016), we link data on employment histories obtained from LinkedIn with inventor histories from the USPTO data (see Online Appendix B). The procedure improves the timing of moves for about 23,000 inventors. Approximately 20% of the successfully granted patent applications during 1973–2010 emerge from continuations of previous applications. We exclude these continuation patents from our sample, as they may be similar to their previously granted parent and may overstate innovation productivity.

Among unique assignees of patents in this USPTO sample, 45% are U.S. firms and 45% are foreign. The remaining are individual inventors and other assignee types (universities, nonprofits, and government institutions). To incorporate firm characteristics, we match the USPTO sample to Compustat data on publicly listed U.S. firms using the procedure described in Bessen (2009). The matching procedure is based on patent grants from 1976 and accounts for changes in patent ownership because of mergers, acquisitions, and spinoffs as of 2006, which we extend to 2010. This yields a sample of active inventors over 1976–2010, for whom we have information on their firm's Compustat variables, including firm age, R&D expenditures, capital intensity, sales, changes in operating income, and the number of employees. However, we lose the first two years of observations because we use the moving average of the values from $t-2$ to $t$ for R&D intensity, capital intensity, sales, operating income change, and patent stocks.

From this sample, we construct the baseline AKM estimation sample by identifying firms connected through mobile inventors. The subsample of the largest connected network has not only mobile inventors but also inventors who do not change firms in the network. This encompasses more than 99% of all inventor-firm-year observations formed by inventors with at least one patent-year observation assigned to a Compustat firm.[3] As Table 1 shows, the baseline AKM sample has information on 708,547 unique inventors working at 2,511 firms, resulting in more than 2.5 million inventor-firm-year observations. In robustness analyses, we restrict the sample to include inventors with careers of at least 6 or 10 years in this sample, because these restrictions may improve the reliability of the estimated inventor fixed effects. Table 1 summarizes these samples as well.

### 2.2. Sample Selection

In an ideal world, we would apply our identification strategies to a sample that includes data on all inventors and their full employment and invention histories. Instead, we are forced to work with a second-best: U.S. patenting records of inventors and their employment histories inferred from the records, potentially updated using LinkedIn data. The AKM sample restricts U.S. patentees to those working at Compustat firms, and the requirement that inventors belong to firms connected through the inventor movement further narrows the sample.

We examine potential selection biases induced by our sample construction procedure by comparing the AKM sample to the overall USPTO sample with all inventors active in our sample period. The USPTO sample is comprehensive but precludes including observable firm characteristics drawn from Compustat. Consequently, the AKM sample has no inventors that are unconnected to firms, whereas unconnected inventors make up 13% of the USPTO inventors. Given Compustat includes some of the largest firms with thousands of employees, one can also expect AKM firms to have a larger number of movers. Indeed, 94% of the 2,511 AKM firms have at least two movers, whereas only 45% of the 248,198 overall USPTO firms (which includes AKM firms) have two or more movers. Nevertheless, the AKM sample has a higher fraction of inventors who have never moved firms (about 91% of all inventors vs. 83% in USPTO), likely because Compustat firms have a longer lifespan. A considerable number of inventors moved once in both samples (6% and 11% in AKM and USPTO, respectively) but moving twice or more is rare in the AKM sample (3% of inventors). Online Appendix A, Table A1, tabulates the corresponding descriptive statistics.

### 2.3. Variable Description

Table 2 describes the variables we use to measure innovation performance, inventor characteristics, and firm characteristics. For each inventor $i$ at firm $j$ in year $t$, we measure the total number of patents weighted by forward citations excluding self-citations over the first five years after patent publication. We correct for teamwork by dividing the measure by the number of coinventors on each patent. To measure extreme innovation outcomes such as breakthroughs, we construct two additional measures. Following Singh and Fleming (2010), the first measure counts patents in the fifth percentile of citations for a given cohort of patents by grant year for an inventor, known as breakthroughs. The second is the log of the number of patents produced, which fail to obtain a single citation, referred to as useless inventions.

The AKM samples include an array of variables that control for correlates of inventors' performance (Table 2). Following prior literature (Hall and Ziedonis 2001), we control for firm age, the existence of R&D expenditures, R&D intensity, capital intensity, sales, changes in operating income, and the number of employees. We also control for the effects of firms' knowledge stocks on inventor output with a measure

**Table 1.** Sample Description and Summary Statistics

| Column | 1 | 2 | 3 |
|---|---|---|---|
| AKM estimation subsamples | Full | 6+ observations | 10+ observations |
| Inventor-firm-year observations | 2,566,626 | 1,675,784 | 1,174,268 |
| Unique inventors | 708,560 | 146,391 | 76,862 |
| Unique firms | 2,511 | 2,273 | 2,000 |
| Mean number of inventors per firm | 282.4 | 64.4 | 38.5 |
| Mean and standard deviation AKM variables | Full | 6+ observations | 10+ observations |
| *Citation Weighted Patents* | 3.22 | 3.33 | 3.19 |
| | (75.1) | (88.1) | (65.5) |
| *Experience* | 5.81 | 7.96 | 9.44 |
| | (5.67) | (5.92) | (6.29) |
| *Firm Age* | 17.7 | 17.9 | 17.8 |
| | (8.80) | (8.37) | (8.04) |
| *Dummy R&D* | 0.98 | 0.98 | 0.98 |
| | (0.13) | (0.13) | (0.14) |
| *R&D Intensity* (%) | 0.37 | 0.23 | 0.26 |
| | (61.5) | (45.3) | (53.5) |
| *Capital Intensity* (%) | 0.55 | 0.43 | 0.47 |
| | (55.0) | (40.0) | (47.6) |
| *Firm Sales* (m$) | 42.1 | 43.6 | 44.4 |
| | (45.1) | (44.9) | (44.8) |
| *Operating Income Change* (%) | 0.52 | 0.47 | 0.39 |
| | (5.33) | (4.97) | (4.39) |
| *Employees* | 119.1 | 126.3 | 123.7 |
| | (121.1) | (119.0) | (119.8) |
| *Patent Stock* | 4,985.1 | 5,081.9 | 4,921.9 |
| | (7,494.2) | (7,424.5) | (7,141.2) |

*Notes.* The top panel of the table describes the observations in the datasets of the AKM analyses. The estimation subsamples correspond to the connectedness sample or the sample of firms connected to each other by inventor mobility. The first column for the AKM estimation subsamples describes all inventors who filed at least one patent during 1978–2010, the second column describes inventors with at least 6 years of patenting experience, and the third describes the subsample of inventors with at least 10 years of experience. In both columns 2 and 3, we include intermediate unproductive years, when counting an inventor's career length. Because the AKM estimation subsamples contain observations for each of the years during which the inventors were active, the number of inventor-firm-year observations are strictly greater than the number of unique inventors. In the bottom panel, we detail the mean and standard deviation (in parentheses) of the variables included in the regression model of Equation (1). All summary statistics refer to the variable in levels, even when they are included in logarithmic scale in the model. Please refer to Table 2 for detailed descriptions of these variables.

of a firm's patent stock in a given year. Table 1 compares the various AKM subsamples with respect to these firm characteristics. A last set of variables pertains to the overall financial performance of the organizations in our data set. Here we consider a firm's net income and Tobin's Q as calculated from its Compustat data for a given year.

## 3. Contributions of Firms and Inventors to Innovation

### 3.1. The AKM Model

Assessing whether persistence in innovation performance is driven by human capital or high ability firms requires disentangling the contributions of inventor and firm-specific capabilities. To accomplish this, following Abowd et al. (1999), we model inventor's inventive output $y$ following the function $y = e^{(\alpha+\phi)}$, where $\alpha$ is the inventor's human capital and $\phi$ is the firm's innovation capability. The output function is supermodular in inventor and firm ability, reflecting their complementarity in output. Our theoretical analysis, in Section 6, shows that such complementarity need not imply complementarity in match surplus and positive assortative matching between inventors and firms. The reason is that match surplus also includes inventor utility, and if the latter is sufficiently concave in innovation output, we obtain a submodular match function.

Taking logs, we derive the log-additive innovation production function, $\log y = \alpha + \phi$. We include the set of time-varying contributors to innovation defined in Table 2 and estimate a model of the form

$$\log y_{ijt} = \beta_x X_{it} + \beta_z Z_{jt} + \omega_{jt} + \gamma_t + \alpha_i + \phi_j + \epsilon_{ijt}. \quad (1)$$

**Table 2.** Variable Descriptions

| Variable | Description |
|---|---|
| Innovation output measure | |
| *Citation-weighted patents* | Number of patents $p$ multiplied by the five-year forward citations (excluding self-citations) to these patents, filed by inventor $i$ at firm $j$, in year $t$, divided by number of coauthors on patent $p$. |
| *Breakthrough patents* | Number of patents $p$ in the fifth percentile of citations for a cohort of patents, by inventor $i$ at firm $j$, in grant year $t$. |
| *Useless patents* | Number of patents $p$ which fail to obtain a single citation, filed by inventor $i$ at firm $j$, in grant year $t$. |
| Inventor characteristics (continuous variables enter estimation in logarithmic scale) | |
| *Past "X"* | Average value output measure "X" in previous one, five, or nine years (depending on specification) for inventor $i$. |
| *Experience* | Number of years between first and current patent in data set for inventor $i$. |
| *Coworkers' citation-weighted patents* | Average of "*Citation Weighted Patents*" by other inventors at firm $j$ in year $t$ excluding focal inventor $i$. |
| Firm characteristics (continuous variables enter estimation in logarithmic scale) | |
| *Firm age* | Firm $j$'s age in year $t$ in years. |
| *Dummy R&D* | Dummy whether firm $j$ reports R&D expenditure in year $t$. |
| *R&D intensity* | R&D Expenditures/Sales averaged over years $t-2$ to $t$. |
| *Capital intensity* | PP&E/Sales averaged over years $t-2$ to $t$, where PP&E is Property, Plant and Equipment expenditure. |
| *Sales* | Firm $j$'s averaged sales over years $t-2$ to $t$. |
| *Operating income change* | Change in operating income of firm $j$ averaged over years $t-2$ to $t$. |
| *Employees* | Number of employees for firm $j$ in year $t$. |
| *Patent stock* | Sum of patents at firm $j$ in years $t$-2 to $t$. |
| Firm performance measures | |
| *Tobin's Q* | Tobin's Q for firm $j$ in year $t$ computed using the formula: $\frac{AT+(CSHO*PRCC\_C)-CEQ}{AT}$, where AT is total assets, CSHO is common outstanding shares, PRCC_C is the annual closing stock price, and CEQ is common equity. |
| *Net income* | Net income for firm $j$ in year $t$ minus minimum of net income in year $t$ over all firms. |

Here, $y_{ijt}$ refers to the number of citation-weighted patents of inventor ($i$) at firm ($j$) in year ($t$). The vectors $X_{it}$ and $Z_{jt}$ represent time-varying inputs related to the inventor ($X_{it}$) and firm ($Z_{jt}$). The vectors $\gamma_t$, $\alpha_i$, and $\phi_j$ contain sets of year, individual inventor, and firm fixed effects, respectively, and $\omega_{jt}$ is a dummy variable set to one if firm $j$ reports R&D expenditure in year $t$ and zero otherwise. $\epsilon_{ijt}$ denotes an inventor-firm-year-specific error term. A key identifying assumption of the AKM method is exogenous mobility, implying that the mobility decisions of inventors may be driven by components of the model, such as their own fixed effect or the firm fixed effect of their current and future firms but not by (components of) the error term.

### 3.2. Baseline AKM Results

To adjudicate the contributions of inventor- and firm-specific effects on inventors' performance, we calculate the covariance of annual innovation output with the inventor-, firm-, and year-fixed effects, divided by the variance of the dependent variable, that is, $\frac{\mathrm{Cov}(y, inventorFE)}{\mathrm{Var}(y)}$, $\frac{\mathrm{Cov}(y, firmFE)}{\mathrm{Var}(y)}$ and $\frac{\mathrm{Cov}(y, yearFE)}{\mathrm{Var}(y)}$. These ratios obtained from AKM regression estimates can be interpreted as the fraction of the total $R^2$ attributable to inventor-specific, firm-specific, and year-specific factors, respectively. We are also interested in the joint significance of the inventor and firm effects, which we assess with a joint $F$-test of the estimated coefficients.

We estimate the AKM model in Equation (1) using the user-written STATA command FELSDVREG (Cornelissen 2008) and report the results in Table 3. Column (1) reports the results from the full AKM sample for the regression with all firm characteristics ($Z_{jt}$) and inventor observables ($X_{it}$) included. Our results indicate that the contributions of inventor and firm effects to innovation performance are highly significant. Inventor heterogeneity explains 34.1% and firm heterogeneity 3.2% of the total variance in inventors' innovation performance.[4] In relative terms, inventor effects are by far the most important factor contributing to the variance in innovation performance among inventors. The covariance between the inventor and firm effects in column (1) is negative, which anticipates our result on negative assortative matching between inventor and firm fixed effects in Section 5. Although not immediately relevant to our objective, we note that year-effects subsume the influence of factors such as the macroeconomic environment or patent law changes that commonly affect the patenting intensity of all inventors in the sample and account for about 2.1% of the explained variance in patent performance in our panel. Inventor fixed effects, firm fixed effects, and year fixed effects are all jointly significant at $p < 0.01$.

### 3.3. Robustness Checks

Columns (2)–(10) of Table 3 report the results of robustness checks. Columns (2) and (3) repeat the analysis

**Table 3.** Contributions of Inventor and Firm Fixed Effects for Inventors' Performance

| | Log number of citation weighted patents (AKM) | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| Dependent variable | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
| Sample | All | 6+ observations | 10+ observations | Movers | No observable covariates | Coworker output | No merger | Established by state/ country | Established by technical field | Includes self-cites |
| Cov($y$, *inventor FE*)/Var($y$) | 0.341 | 0.202 | 0.183 | 0.179 | 0.368 | 0.338 | 0.340 | 0.329 | 0.367 | 0.359 |
| Cov($y$, *firm FE*)/Var($y$) | 0.032 | 0.029 | 0.026 | 0.040 | 0.035 | 0.023 | 0.031 | 0.043 | 0.071 | 0.033 |
| Cov($y$, *year FE*)/Var($y$) | 0.021 | 0.025 | 0.029 | 0.029 | 0.032 | 0.007 | 0.021 | 0.021 | 0.120 | 0.021 |
| Cov(*inventor FE, firm FE*)/Var($y$) | −0.109 | −0.082 | −0.059 | −0.074 | −0.043 | −0.104 | −0.106 | −0.095 | −0.011 | −0.129 |
| $F$-test on inventor and firm FE ($p$ value) | <0.01 | <0.01 | <0.01 | <0.01 | <0.01 | <0.01 | <0.01 | <0.01 | <0.01 | <0.01 |
| Log experience | −0.31* | −0.33* | −0.33* | −0.30* | – | −0.30* | −0.31* | −0.31* | 0.02* | −0.37* |
| | (0.002) | (0.002) | (0.002) | (0.003) | | (0.002) | (0.002) | (0.002) | (0.003) | (0.002) |
| Firm characteristics | Yes | Yes | Yes | Yes | No | Yes | Yes | Yes | Yes | Yes |
| No. firms | 2,511 | 2,273 | 2,000 | 2,511 | 2,511 | 2,283 | 2,485 | 7,830 | 3,130 | 2,511 |
| No. movers | 60,926 | 45,706 | 30,099 | 60,926 | 60,926 | 59,918 | 59,373 | 52,548 | 157,147 | 60,926 |
| No. inventors | 708,560 | 146,391 | 76,862 | 60,926 | 708,560 | 702,261 | 698,429 | 644,436 | 666,768 | 708,560 |
| Year FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Firm FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Inventor FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Observations | 2,566,626 | 1,675,784 | 1,174,268 | 676,164 | 2,566,626 | 2,541,551 | 2,533,024 | 2,412,643 | 1,451,072 | 2,566,626 |
| $R^2$ | 0.447 | 0.295 | 0.273 | 0.287 | 0.435 | 0.450 | 0.447 | 0.446 | 0.569 | 0.445 |

*Notes.* The table reports AKM regression estimates where log number of citation weighted patents (proxy for inventive performance) is the dependent variable. Each observation is at the inventor-firm-year level. Both years with and without patents are included. All estimations, except column 6, include the firm characteristics described in Table 2: Firm age, Dummy R&D, R&D intensity, Capital intensity, Sales, Operating income change, Employees, and Patent stock. Column 1 reports on the full AKM sample, columns 2 and 3 use a subsample of AKM inventors with at least 6 and 10 yearly observations, column 4 uses the subsample of inventors who moved at least once, column 5 excludes all time-varying covariates, column 6 inserts coworker output in the firm characteristics, column 7 only uses observations from firms that did not have mergers or acquisitions in the sample, column 8 splits firms into establishments by .U.S state and (non-U.S.) country, column 9 splits firms by the technology class of the patent filed, and column 10 includes self-citation in the patent output measure. Cov($y$, *inventor FE*)/Var($y$), Cov($y$, *firm FE*)/Var($y$), and Cov($y$, *year FE*)/Var($y$) report the contribution of inventor, firm, and year fixed effects toward explaining the observed variance in the inventor's output of citation weighted patents. Cov(*inventor FE, firm FE*)/Var($y$) reports the covariance between inventor and firm fixed effects scaled by the variance of inventor's output of citation weighted patents. Robust standard errors in parentheses. FE, fixed effect.

*$p < 0.01$.

in Column (1) using a subsample of inventors with at least 6 and 10 years of experience, respectively, where we count both years with and without patents toward an inventor's experience. This allows more observations, and hence more degrees of freedom, to identify each individual effect. It also addresses the concern that inventor effects may be noisy as a result of *overfitting* the fixed effect on a short inventor career. The relative contribution of inventor fixed effects goes down to 20.2% and 18.3% in these analyses compared with a 3% contribution of firm fixed effects. Hence, the inventor effects still explain a far larger portion of the variance in innovation output than the firm effects. The decline in the explanatory power of the inventor effects could indicate we are overfitting the fixed effects of inventors with a very small number of yearly observations. If we estimate the AKM model for a subsample where we require inventors to have at least two observations, we find that this leads to a sharp decrease in the proportion of variance explained by the inventor fixed effects (from 34% to 25%). On the other hand, restricting the sample to inventors with two or more observations reduces the number of inventors in the sample by about 60%, which clearly poses a problem in terms of the representativeness of our estimation sample.

Our AKM sample in columns (1)–(3) are based on firms connected through a network formed by inventor moves. This allows us to compute the fixed effects of both mobile and immobile inventors in connected firms. Of course, mobile inventors may be systematically different from inventors who have never changed firms and one may question AKM's imputation of fixed effects for nonmobile inventors. Therefore, column (4) reports the importance of inventor and firm effects obtained by estimating Equation (1) on a subsample of mobile inventors alone (as in Bertrand and Schoar 2003). This subsample of movers yields estimates of inventor-fixed effects quite close (17.9%) in importance to the ones obtained in columns (2) and (3). Column (5) repeats the estimation in column (1), but without any observed firm or inventor characteristics. This assures that the estimated importance of inventor-fixed effects relative to firm-fixed effects reported in Column (1) is not because we included several firm characteristics and only a few inventor characteristics as controls. The results we obtain in column (5) are similar to those in column (1).

Another potential concern is that the productivity of coworkers may affect the inventor's own output over and above the firm's time-invariant impact (Jones 2009, Jaravel et al. 2018). In our baseline model, the firm and year effects may both partially account for this, depending on whether coworker productivity evolves through common time-varying shocks or mainly varies across firms. Our baseline model may, therefore, overestimate the relative importance of the firm and year effects in this case. Because we cannot estimate a full set of coworker effects, which would lead to strict multicollinearity, we introduce the contemporaneous output of inventors at the firm to probe the severity of this issue. Although this approach raises endogeneity concerns, we show in column (6) that it does not dramatically affect the importance of the firm effects. Finally, we exclude observations of firms that change ownership because of mergers and acquisitions of entities in column (7), and these results are similar to those in column (1). Again, the inventor-specific effects explain 34% of the variance in their inventiveness.[5]

In the AKM estimation, we estimate one fixed effect for each firm regardless of its degree of centralization. Treating a large multinational firm such as IBM with multiple product and geographic units as a single entity can bias our results (Arora et al. 2014). To address this concern, we conduct two robustness checks where we split each Compustat entity into branches. In the first check, we group all inventors located in the same state (if located in the United States) or country (if located outside of the United States) into one firm entity (column 8). In the second analysis, we group all patent activity in the same National Bureau of Economic Research (NBER) technology class within the Compustat firm into one entity (column 9). These robustness checks more likely reflect independent divisions and departments within firms. A concern in this exercise is the translation of the explanatory variables to a suitable level of disaggregation. We therefore recreated the patent stock variable by calculating through patents at this lower level of aggregation in each case. We assume that the capital and R&D intensity are uniform across subunits within firms because these cannot be apportioned in a meaningful way. We find that the results remain qualitatively similar in these cases. In column (10), we examine the robustness of our results to including self-citations in the dependent variable and again find similar results as in column (1).

### 3.4. Distribution of Inventor and Firm Effects

Here we examine heterogeneity among inventors and firms in their estimated inventive capabilities. Because a given inventor (firm) fixed effect should be interpreted relative to all other inventor (firm) fixed effects in the sample, we follow the common practice of rescaling the estimated effects by the distribution mean. Rescaling centers the distributions of fixed effects at zero. Figures 1 and 2 display the distribution of inventor- and firm-fixed effects obtained from the regression model in column (1) of Table 3. After rescaling, the average fixed effect equals 0 for both firms and inventors. The standard deviations of the

**Figure 1.** (Color online) Distribution of Inventor Fixed Effects Drawn from AKM Estimation

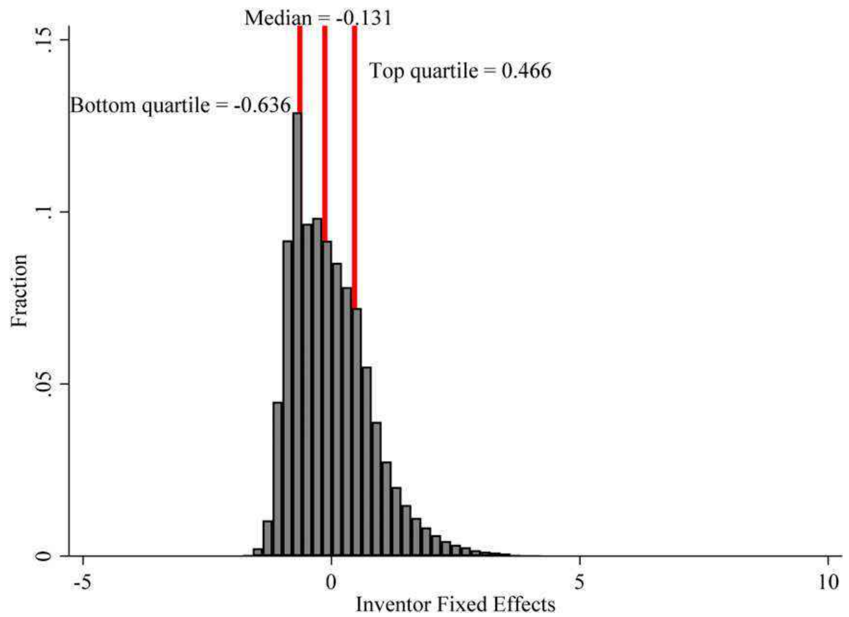


*Notes.* The figure plots the distribution of the 708,560 inventor fixed effects estimated using the AKM specification and sample corresponding to column (1) of Table 3. The estimated inventor fixed effects have been standardized by subtracting the population mean from the estimates. The vertical lines indicate the top quartile, median, and bottom quartile of the estimated inventor fixed effects.

inventor- and firm-fixed effects are 0.81 and 0.59, respectively. The median inventor has an estimated effect of −0.131, that is, slightly below the population average, whereas the first and third quartile stand at −0.636 and 0.466, respectively. This leftward shift with respect to the population average is caused by

**Figure 2.** (Color online) Distribution of Firm Fixed Effects Drawn from AKM Estimation

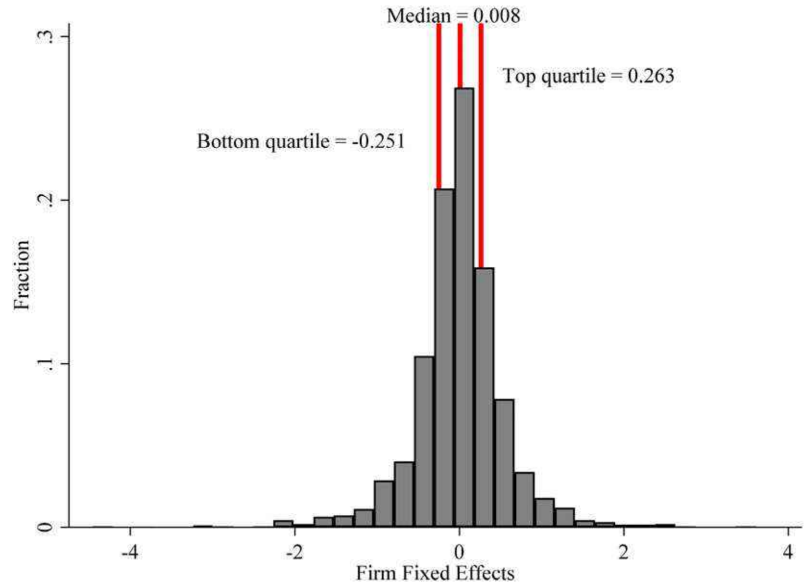


*Notes.* The figure plots the distribution of the 2,511 firm fixed effects estimated through the AKM specification and sample corresponding to column (1) of Table 3. The estimated firm fixed effects have been standardized by subtracting the population mean from the estimates. The vertical lines indicate the top quartile, median, and bottom quartile of the estimated firm fixed effects.

the relatively long right tail of the distribution. As can be seen in Figure 1, the left tail of *underperforming* inventors is fairly short relative to the right tail, suggesting the presence of star inventors. By comparison, the distribution of firm effects is more balanced. Here the median estimated effect is 0.008, with the first and third quartiles at −0.251 and 0.263, respectively. Figure 2 confirms this observation, because it shows no apparent skewness in the distribution of firm effects. Hence, star firms seem less common than star inventors.

### 3.5. Impact of Technology Field, Firm Size, and Alternative Output Measures

Previous research suggests that the production of innovation output differs significantly among technology fields (Cohen et al. 2000, Malerba 2005). We explore how these differences impact the relative contribution of firms and inventors to innovation production by estimating the baseline AKM specifications, with the complete set of covariates, for each of the six technology fields defined in Hall et al. (2001). In this exercise, we only consider patents within the technology field to calculate an inventor's output, as well as a firm's stock of patents. As shown in Online Appendix A, Table A3, inventor effects explain between 25% and 30% of the variance of innovation output in each technology field. Firm capabilities are significantly less important and explain between 3% and 7.5% of the variance in inventor output for each field.

The relative contributions of firms and inventors may also be sensitive to firm size. We, therefore, repeat the AKM estimation in subsamples of firms with varying numbers of inventors. Our results (Table A3) indicate that firm effects explain a larger share of the variance for smaller firms with less than 50 inventors compared with firms with more than 1000 inventors (10% vs. 1%). As discussed before, the firm effect in smaller firms may take up some of the effect of group-level human capital. Even so, the contribution of inventor effects is much larger across the entire firm size distribution, with values ranging from 34.5% to 40%.

If firms contribute more to the generation of high-impact patents rather than to the number of patents or citations, then our previous analyses do not capture such nuance. To address this concern, we repeat the estimations using measures developed in Singh and Fleming (2010) for breakthroughs and useless inventions as the dependent variables. The results are reported in Online Appendix A, Table A4. Both in terms of high-end and low-end patents, the contribution of inventors outweighs the contribution of firms toward explaining variation in innovation performance. This is in line with the results we obtain for the log of citation-weighted patents produced per year.

## 4. Inventor-Firm Matching

### 4.1. Time-Varying Inventor and Firm Effects

If human capital is the most important contributing factor for inventor performance, then how do inventors choose the firms they work for? A deeper understanding of the matching process between firms and high-skilled workers is essential to address this question. In this section, we focus our attention on matching between human capital and firm capabilities.

The standard AKM estimates reported in Section 4 poses some challenges to study matching between inventors and firms. To illustrate why, suppose we are interested in relating an inventor's movement between two employers in year $t$ to her individual ability, as estimated by AKM. When individual effects are estimated on the full sample, an inventor's effect is constructed from her average innovation output across all her employers, net of observable inputs, and firm capabilities. This includes observations both before and after year $t$, and as such, these estimates are *contaminated* by the firms to which the inventor has not yet moved in year $t$ (but will do so in a later year in the sample period). If we were to use these estimates to analyze the inventor's move in year $t$, it would be impossible to disentangle whether an inventor with a high estimate moved to a firm with greater ability or whether the inventor's estimate is high, because it is partly derived from her time working at a firm with greater capabilities. The same holds true for estimates of firm capabilities. The presence of such match dynamics undermine the assumption of our baseline analysis (and theory we present in Section 6) that firm and inventor types are constant through the sample.

With this caveat in mind, we propose a *rolling window* procedure that derives time-varying estimates of inventor and firm effects through the AKM methodology. To implement the procedure, we begin by limiting the sample to a 10-year period from 1978 through 1987. Then, we estimate Equation (1) on the largest network in this subsample to obtain firm and inventor effects. Crucially, these estimates are not contaminated by how the inventor and firm effects change as a result of inventor moves after 1987. Next, we draw a new subsample of 10 years by rolling the window by one year, from 1979 through 1988. We again estimate Equation (1) on this sample. We continue this rolling procedure until we arrive at the end of our main sample in 2010. Because the effects in different windows may be estimated in comparison with different benchmark inventors (firms), we standardize the estimated inventor (firm) effects by subtracting the mean and dividing by the standard deviation of all inventor (firm) effects in the same subsample. We thus end up with a set of standardized time-varying estimates for firm and inventor effects which we can leverage to informally check matching patterns. These are, in our view, best interpreted as time-varying measures of an inventor's (firm's) relative innovation ability, compared with the distribution of contemporary inventor (firm) abilities (i.e., those active in the past 10 years). To examine the robustness of our results, we implement the same procedure for rolling windows that span five years.

## 4.2. Which Inventors Move?

We first use the individual effects estimated through the rolling window algorithm to investigate which inventors are likely to move to another firm in the future as a function of their current human capital. To this end, we define a mobility indicator $y_{it}$, which equals 1 in year $t$ for inventor $i$ at firm $j$ if the next patent by inventor $i$ is filed at a different firm than firm $j$. We set this indicator equal to 0 if the next patent filed by inventor $i$ is filed at firm $j$ and code the variable value as missing if the inventor does not reappear in the sample after year $t$. We then estimate a regression model to relate this indicator to the estimates of inventor effects obtained from the window ending in year $t$. In light of the literature on star inventors, we allow for nonlinear effects along the distribution of the estimated inventor capabilities. Hence, we do not include the estimated inventor effect directly but construct a vector of indicator variables, which classify the estimated inventor effects into deciles. As such, we obtain a profile of effects along the deciles of the estimated inventor capabilities rather than a single estimate for the average effect. Our model takes the form

$$y_{ijt} = \beta_0 + \beta_a \hat{\alpha}_{it} + \beta_f \hat{\varphi}_{jt} + \beta_c E[\hat{\alpha}_{ct}]_{jt} + \beta_x x_{ijt} + \gamma_t + \sigma_j + \varepsilon_{it}. \quad (2)$$

In Equation (2), $\hat{\alpha}_{it}$ denotes the vector of indicators for the decile of the estimate of inventor $i$'s individual effect obtained from the 10- (or 5-) year rolling window ending in year $t$. The model further includes $\hat{\varphi}_{jt}$, which represents a vector of indicators for the decile of the estimated firm effect, and $E[\hat{\alpha}_{ct}]_{jt}$, which stands for the average estimated inventor effect for all inventors $c$, filing patents at firm $j$ in year $t$. We construct these variables based on estimates drawn from the same rolling window as those used for $\hat{\alpha}_{it}$. We also control for the current tenure of inventor $i$ at firm $j$ ($x_{ijt}$), and a set of year ($\gamma_t$) and North American Industry Classification System two-digit level fixed effects ($\sigma_j$). We estimate Equation (2) using a linear probability model with bootstrapped standard errors, such that the coefficients of the inventor effects can be interpreted as marginal probabilities relative to the lowest decile.[6]

To allow a clear interpretation of the estimation results, we plot the coefficients of the decile indicators with their confidence interval in Figure 3.[7] We find that inventors in higher deciles are more likely to move and the probability to move increases along with the inventor decile. Inventors in the eighth and ninth decile of estimated inventor ability are most likely to move (about 4% more likely than the lowest decile). By contrast, inventors in the very top decile of the ability distribution move less often than those just below the top category. Although they are still more likely to move than the bottom decile (about 2% more likely), this may indicate that firms fight harder to retain the very top performers compared with those just below.

## 4.3. Human Capital and Inventor-Firm Matching

Next, we study how the inventor's estimated capabilities derived from the rolling window procedure

**Figure 3.** Estimates of Inventor Mobility by Decile of Inventor Effect

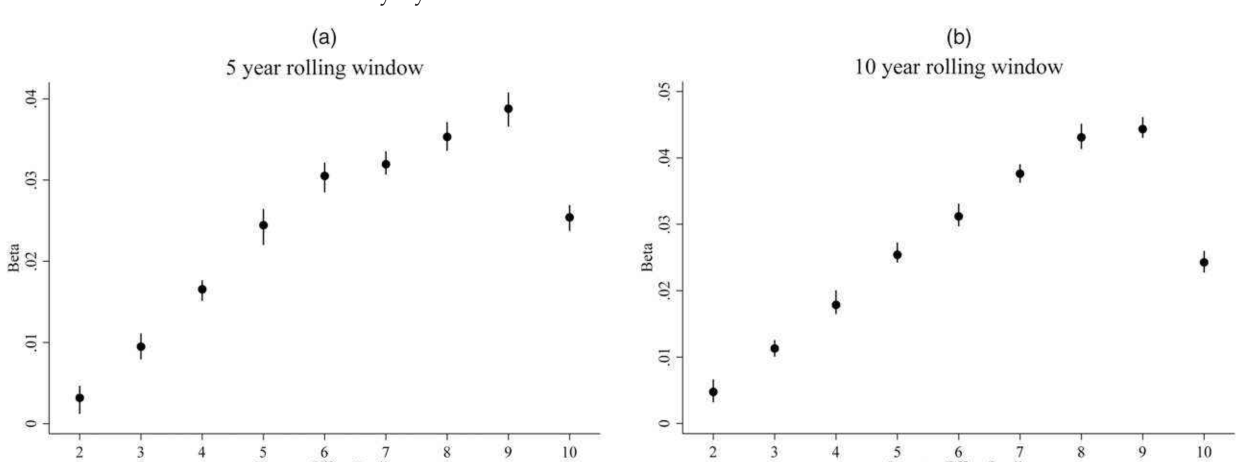


*Notes.* This figure plots the estimated coefficients for the inventor effect decile estimates in the inventor-firm matching models reported in Table A5 of the online appendix. We depict the point estimate relative to the bottom decile (dot) together with its bootstrapped 95% confidence interval (whiskers). Panel (a) refers to column (2) and panel (b) to column (4) of Online Appendix A, Table A5.

correlate with characteristics of the firm, to which the mobile inventor moves. In particular, we test whether more high-skilled inventors are attracted to (a) firms with superior firm-specific inventive capabilities or (b) firms with high-skilled coworkers. Our sample for this analysis consists of all movements by an inventor $i$ from a firm $j$ to a new firm $k$, for which we are able to obtain an estimate of the inventor effect from the rolling window ending in the year $t$, that is, the last year inventor $i$ is observed at firm $j$. Formally, we estimate the following regression model

$$\hat{\alpha}_{it} = \beta_0 + \beta_k x_{kt} + \beta_i x_{it} + \gamma_t + \varepsilon_{it}. \tag{3}$$

In Equation (3), the dependent variable $\hat{\alpha}_{it}$ refers to the mobile inventor's estimated effect before the move. The vector $x_{kt}$ refers to characteristics of the next firm $k$ at time $t$, that is, before inventor $i$ has joined firm $k$. These characteristics primarily include (a) firm $k$'s estimated firm capability and (b) the average estimated ability of inventors active at firm $k$. We further control for firm $k$'s size, measured as log of assets, age, and profitability, proxied by the Tobin's Q, and log net income. We measure these firm characteristics as described by the variable definitions in Table 2. We standardize each characteristic by subtracting the average across all active firms in year $t$ and dividing by the standard deviation among firms in year $t$. In each specification, we further add a constant term ($\beta_0$), the log of inventor experience in years ($x_{it}$), and a set of year fixed effects, $\gamma_t$.

Table 4 reports the results obtained by estimating Equation (3) for the 5- and 10-year rolling windows. We find that firm-specific innovation capability correlates negatively with the ability of inventors moving into the firm. This suggests negative assortative matching between the innovation capability of the firm and the human capital of inventors. Thus, firms with lower estimated firm-specific innovation capabilities attract inventors with higher estimated human capital. Our results for the hiring firms' average inventor ability lead us to the opposite conclusion. Moving inventors with higher human capital join firms where their future coworkers also have, on average, higher

**Table 4.** Inventor-Firm Matching

| | Moving inventor's estimated inventor effect prior to move | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| Dependent variable | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
| *Rolling window* | 5 year | 10 year | 5 year | 10 year | 5 year | 10 year | 5 year | 10 year | 5 year | 10 year |
| *Destination firm's firm effect* | −0.035* | −0.033* | | | −0.023* | −0.013* | −0.026* | −0.017* | −0.022* | −0.012* |
| | (0.004) | (0.007) | | | (0.006) | (0.007) | (0.006) | (0.008) | (0.006) | (0.007) |
| *Destination firm's mean inventor effect* | | | 0.100* | 0.086* | 0.063* | 0.070* | 0.059* | 0.062* | 0.064* | 0.070* |
| | | | (0.008) | (0.007) | (0.013) | (0.013) | (0.013) | (0.013) | (0.013) | (0.013) |
| *Destination firm's log assets* | | | | | 0.008* | 0.012* | 0.009* | 0.013* | 0.006* | 0.009† |
| | | | | | (0.002) | (0.004) | (0.002) | (0.004) | (0.002) | (0.004) |
| *Destination firm's log age* | | | | | −0.011* | −0.012* | −0.011* | −0.012* | −0.011* | −0.012* |
| | | | | | (0.002) | (0.003) | (0.002) | (0.004) | (0.002) | (0.003) |
| *Destination firm's Tobin's Q* | | | | | | | 0.006† | 0.005 | | |
| | | | | | | | (0.002) | (0.004) | | |
| *Destination firm's log net income* | | | | | | | | | 0.002* | 0.003* |
| | | | | | | | | | (0.001) | (0.001) |
| *Log inventor experience* | 0.205* | 0.404* | 0.204* | 0.403* | 0.205* | 0.403* | 0.204* | 0.402* | 0.205* | 0.404* |
| | (0.016) | (0.024) | (0.016) | (0.024) | (0.016) | (0.024) | (0.016) | (0.023) | (0.016) | (0.024) |
| Constant | 0.057 | −0.537 | 0.057 | −0.537 | −0.103 | −0.551 | −0.104 | −0.551 | −0.102 | −0.544 |
| Year FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Observations | 47,697 | 46,090 | 47,697 | 46,090 | 47,181 | 46,084 | 45,938 | 44,851 | 47,036 | 46,084 |
| $R^2$ | 0.256 | 0.274 | 0.256 | 0.275 | 0.260 | 0.275 | 0.259 | 0.273 | 0.260 | 0.275 |

*Notes.* The table reports ordinary least squares regression results, where the dependent variable is the estimated inventor effect of moving inventors in the year before the move obtained from the moving window AKM regression. The main explanatory variables refer to characteristics of the destination firm, also obtained in the year before the inventor's arrival at this firm. We include the destination firm's estimated firm effect from the AKM moving window, the mean of the fixed effects of inventors at the destination firm, the destination firm's log assets, log firm age, log net income, and Tobin's Q. We standardize all inventor and firm effects and firm characteristics by subtracting the mean and dividing by the standard deviation of the respective year. We control for the inventor's premove experience and a set of year fixed effects in all models. We use a bootstrap procedure to correct for the use of estimated explanatory variables. We first resample the original observations 100 times and create 100 sets of estimates from the moving window estimation. We then run the second stage regressions separately for each of these 100 sets of inventor and firm effects to obtain confidence intervals. We report the median of the bootstrapped regression coefficients as the estimated coefficient. Bootstrapped errors are reported in parentheses. FE, fixed effect.

$^{*}p < 0.01$; $^{\dagger}p < 0.05$.

human capital. This suggests positive assortative matching among coworkers. In particular, if a mobile inventor's estimated human capital is one standard deviation higher, the estimated capability of her destination firm is around 3% of a standard deviation lower, whereas the average estimated human capital of her future coworkers is between 8% and 10% of a standard deviation higher. These results are robust when we control for the other firm characteristics detailed above. The contrasting matching patterns among inventors and coinventors on the one hand and inventors and firm capabilities on the other suggests that rather than treating firms as having one dimension of quality, as previous literature on matching has, it may be important to consider firms as multidimensional entities and examine the quality of each dimension separately.[8]

To examine whether our findings of negative assortative matching between inventors and firms based on innovation capability and positive assortative matching based on human capital characterize the stock of inventors at firms, not just movers, we return to the baseline AKM estimates (of Equation (1)) presented in column (1) of Table 3. We plot the firm-fixed effects and the mean estimated inventor-fixed effects at the firms in Figure 4. These estimates are derived from the AKM sample of all connected firms and incorporate information on all employees at the firms. The figure shows a large negative correlation (−0.676) between firm-fixed effects and mean inventor fixed effects. In contrast, Figure 5 shows a positive correlation (0.429) between the estimated inventor and coworker fixed effects at the firm. Thus, even considering a snapshot of inventor-firm assignments, high-skilled inventors are more likely to be matched with firms that have other high-skilled inventors, but low firm-specific innovation capabilities.

In a final analysis based on the rolling window estimates, we relate the firm's profitability in year $t$ to (a) its own estimated firm effect from the rolling window in year $t-1$ and (b) the average inventor effect of its active inventors in $t-1$. As reported in Online Appendix A, Table A7, we find a significantly positive relation between the current firm and average inventor effect and Tobin's Q and the firm's net income. Although we do not attach any causal interpretation to these findings, they support the notion that innovative capabilities both at the firm and inventor level may be valuable for the firm's bottom line. As such, firms may indeed face a choice to either develop their inventor workforce or attempt to build up firm capabilities, which may both lead the firm to profitability.

## 5. Potential Estimation Biases

### 5.1. Limited Mobility Bias and Sparseness of the Connected Network

In Figure 3, we depict the scatterplot of inventor and firm fixed effects drawn from the AKM estimation and report a negative correlation between both types

**Figure 4.** Inventor Fixed Effects and Firm Fixed Effects Drawn from AKM Estimation

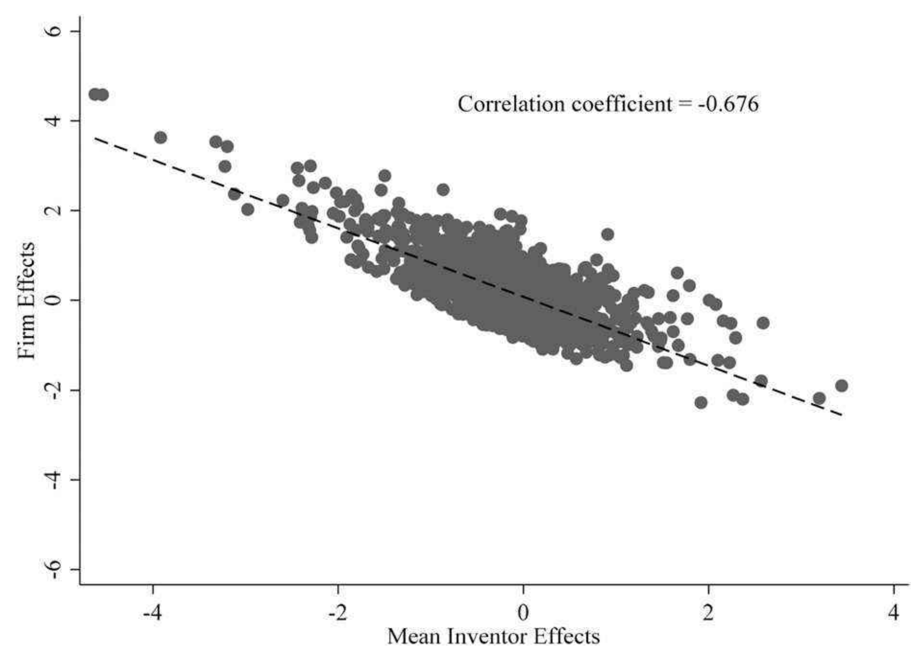


*Notes.* The figure plots estimated firm fixed effects against the mean estimated inventor fixed effect at the firms. The figure is based on 2,273 estimated firm fixed effects and the same number of inventor fixed effects obtained by averaging estimated inventor fixed effects of all inventors at each firm. Fixed effects are obtained using the AKM specification and sample corresponding to column (1) of Table 3. The estimated fixed effects have been standardized by subtracting the population mean from the estimates.

**Figure 5.** Inventor and Average Coworker Fixed Effects Drawn from AKM Estimation

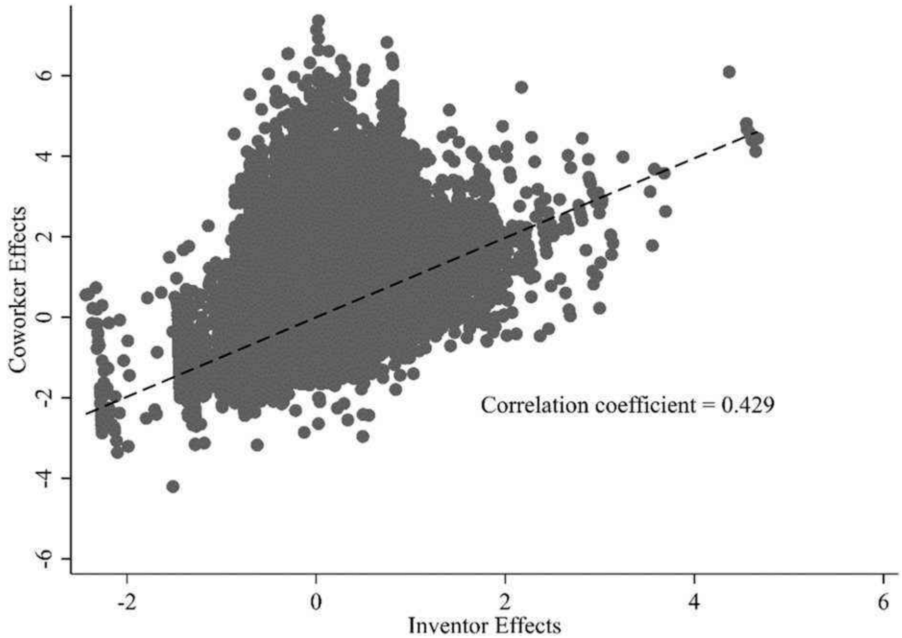


*Notes.* The figure plots estimated inventor fixed effects against estimated coworkers' fixed effects. The figure is based on 788,051 estimated inventor and coworker fixed effects (greater than the number of inventor effects—708,560—because some inventors change employers and thus coworkers). All fixed effects are obtained using the AKM specification and sample corresponding to column (1) of Table 3. The estimated fixed effects have been standardized by subtracting the population mean from the estimates.

of fixed effects (–0.676). However, this number may present an overly negative estimate of the correlation between inventor and firm fixed effects, if our data have *limited mobility* of workers across firms (Andrews et al. 2008). The limited mobility bias arises because worker and firm effects are characterized by estimation error, but at the same time, worker and firm effects are usually estimated based on the same observations. This implies for each data point that an overestimate of the effect in one dimension (e.g., worker effect) on average results in an underestimate of the effect in the other dimension (e.g., firm effect). In the asymptotic case, the bias approaches zero, because both the estimation error declines and every firm has many moving workers, which ensures the estimation error can be averaged out over many worker-firm combinations. However, in samples with short time dimensions and few movers, the correlation between the estimation error of the worker and firm effects will be factored into the calculated correlation of the estimated fixed effects. This leads to a potential downward bias in the correlation between worker and firm effect estimates, such as the one we report in Figure 3.

A first solution to counteract this potential bias is to calculate the correlation of worker and firm fixed effects, which have not been obtained from the same observations. For example, in the analysis reported in Table 4, we relate the firm fixed effect of an inventor's future firm to the inventor fixed effects estimated on observations before their move to this firm. This implies that there cannot be a mechanical correlation between the estimation error in the firm and inventor fixed effects because both are drawn from independent data points. As can be seen in Table 4, we also find a significantly negative correlation in this analysis.

A second solution, proposed by Andrews et al. (2008), is to assess the empirical relevance of the limited mobility bias in the data under investigation. Here, we follow the approach of Jochmans and Weidner (2019), who expand the work of Andrews et al. (2008). Their approach measures the connectedness of the underlying worker-firm network to characterize the estimation error in the worker and firm fixed effects. For a given network, this then allows us to evaluate the precision of the estimated fixed effects and how this (im)precision biases the variance and covariance of the estimated fixed effects.

To implement this approach, we derive the adjacency matrix, $\boldsymbol{A}$, for the largest network of connected firms in the samples examined in columns (1), (2), and (3) of Table 3. We weigh the importance of each connection (i.e., each moving inventor) by the number of observations we have for the inventor at each firm in the connection. Using the notation of Jochmans and Weidner (2019), we then calculate the Laplacian $\boldsymbol{L}^*$ and normalized Laplacian $\boldsymbol{S}$ of $\boldsymbol{A}$. These matrices characterize the large sample properties for fixed effects estimated on the connected network. In particular, $\lambda_2$, the first nonzero Eigenvalue of $S$, should be significantly larger than 0 for the network to be sufficiently connected. As shown in Table 5, we find values of 0.030, 0.030, and 0.023. These values are clearly above those in the teacher value-added example reported by Jochmans and Weidner (2019) as an example of a weakly connected network. We gauge the implied bias in the variance of the fixed effects (see Jochmans and Weidner 2019 for relevant formulas), by calculating the distribution of $S^{\dagger}$, which measures the scale of the variance approximation (as calculated from the fixed effects) to its exact value. Ideally, this value should be close to unity for most of its distribution. As shown in Table 5, the median and mean of the variance approximations are close to 1, with the 6+ observations sample yielding the most favorable results.[9] By comparison, the teacher value-added example in Jochmans and Weidner (2019) yielded a mean value of 2.5 with 1.29 in the first decile. Finally, we calculate the weighted trace of $L^{*}$ to get the bias in the plug-in variance estimator measured as a proportion of the error variance. Because the sample of all inventors contains more inventors with short careers, who are by default less mobile, the potential bias is higher for the analysis in column (1) (approximately 10% of the error variance) than in the other samples (approximately 2.5%–3%). Overall, these values are in line with the occupational network analysis in Jochmans and Weidner (2019), which serves as an example of a reassuringly well-connected network. This indicates that our conclusions are less likely to be the result of biases in the variances and covariances reported in Table 3.

### 5.2. Match Effects and Endogenous Mobility

The AKM technique uses inventor movements to pin down inventor and firm-fixed effects, but it is unlikely that inventor movement across firms is random. Our analysis in Section 5, where we introduce a rolling window estimation strategy to investigate the relationship between inventor human capital and inventor-firm matching, indeed confirms that inventor and firm effects affect job mobility. This raises concerns about the maintained assumption of exogenous mobility in standard AKM models, but in itself does not imply that our estimates of inventor and firm effects are biased. In fact, as long as inventor mobility is a function of the fixed inventor or firm effects or other components included in Equation (1), there is no reason to expect biases in the estimated individual effects (Abowd et al. 2019). However, if certain qualities of individuals or firms not captured by the fixed effects or other included terms drive mobility and thus worker-firm matching, our estimates will be biased.

A potential channel for the error term to influence mobility is through the existence of *match effects* between firms and workers. Both Lazear et al. (2015) and Abowd et al. (2019) examine the correlation between the average residual of a match and the fixed effects of future employers. The test of Abowd et al. (2019) partitions moving workers by the decile of their origin firm fixed effect and their destination firm fixed effect. If there exists a systematic pattern in these average residuals, a proxy for the match effect, this indicates mobility may be endogenous. After all, sorting based on match effects would imply that inventors moving to firms with higher firm effect should see systematically higher match effects before their move. The top panel of Figure 6 shows a graph of the average match effect of mobile inventors at the origin firm partitioned by the origin firm effect decile (left horizontal axis) and the destination firm effect decile (right horizontal axis) for the model of column (1) of Table 3. Unlike in Abowd et al. (2019), there appears to be no systematic pattern, which suggests that mobility is not endogenous with respect to match effects. However, the bottom panel of Figure 6 depicts the number of moving inventors in each decile to decile cell. It is clear that some cells are not sufficiently populated in our sample to draw firm conclusions and conduct the $\chi^2$ test described in Abowd et al. (2019).

**Table 5.** Diagnostic Statistics for the Sparseness of Connected Firm Network

| Column | 1 | 2 | 3 |
|---|---|---|---|
| Sample | All | 6+ observations | 10+ observations |
| $\lambda_2$ | 0.030 | 0.030 | 0.023 |
| Median ($S^{\dagger}$) | 1.0205 | 1.0162 | 1.0177 |
| Mean ($S^{\dagger}$) | 1.2626 | 1.1771 | 1.1877 |
| Standard deviation ($S^{\dagger}$) | 1.4517 | 0.8896 | 0.9978 |
| Weighted trace ($L^{*}$) = bias in variance as % error variance | 9.86% | 3.21% | 2.45% |

*Notes.* Table shows the statistics developed in Jochmans and Weidner (2019) to test the bias in two-way fixed effects estimated on bipartite connected networks. The samples tested refer to columns (1)–(3) of Table 5. We refer to Section 5.1 of the paper for further details and Jochmans and Weidner (2019) for detailed definitions, formulas, and more explanation on these statistics.

**Figure 6.** Average and Number of Match Effects by Current and Next Firm Fixed Effect Decile

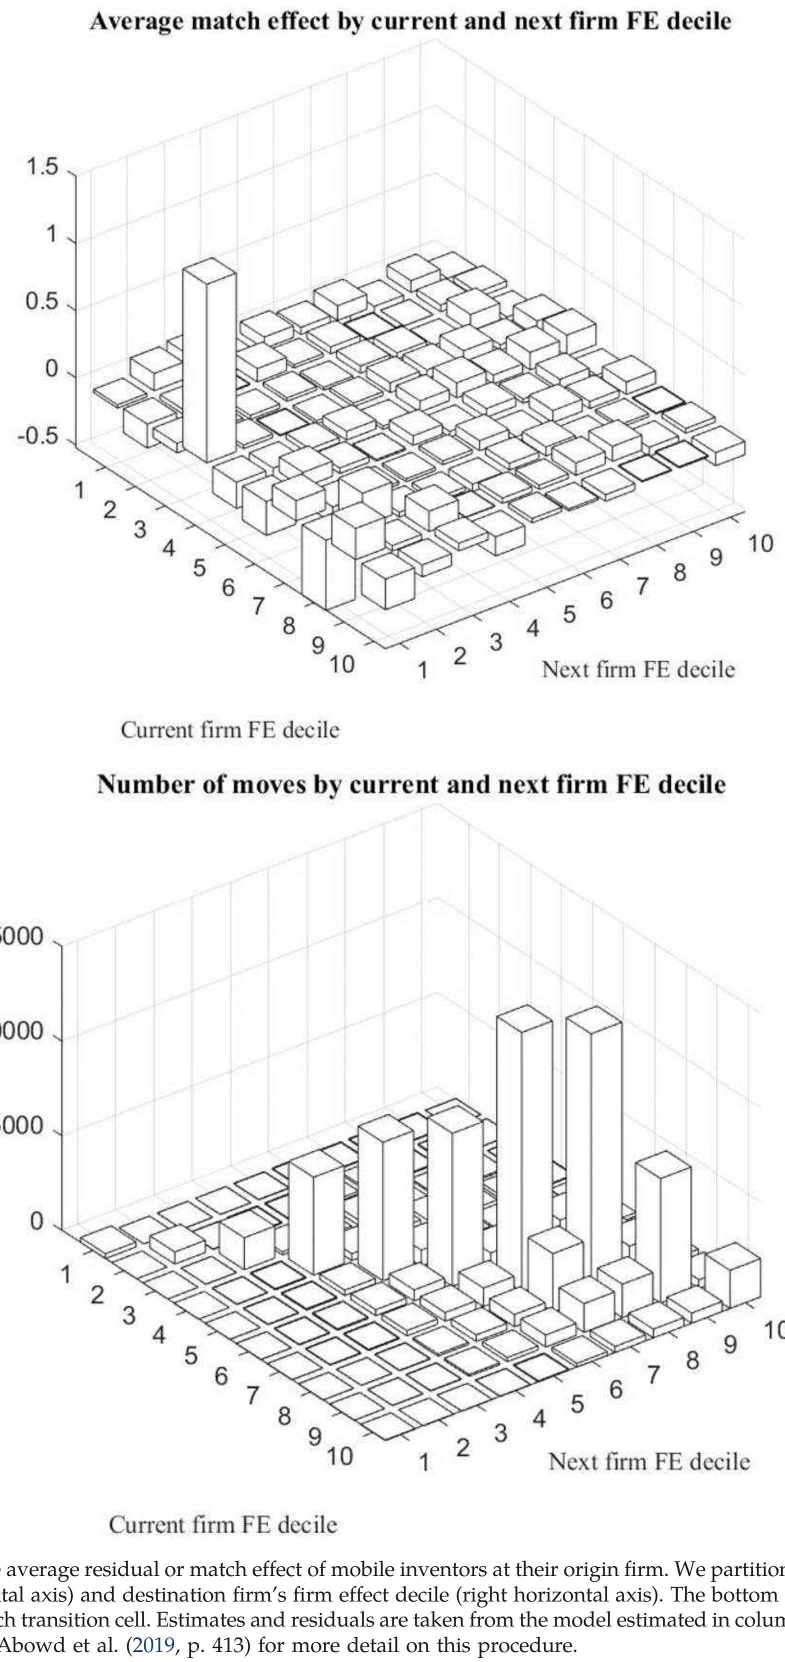


*Notes.* The top panel plots the average residual or match effect of mobile inventors at their origin firm. We partition these by their origin firm's firm effect decile (left horizontal axis) and destination firm's firm effect decile (right horizontal axis). The bottom panel shows the number of individual match effects in each transition cell. Estimates and residuals are taken from the model estimated in column (1) of Table 3. We refer to Section 6.3 of the paper and Abowd et al. (2019, p. 413) for more detail on this procedure.

An alternative test, suggested by Lazear et al. (2015), takes the subsample of mobile inventors and regresses the destination firm's fixed effect on the average residual during employment at the origin firm. We perform this regression on the samples used in columns (1)–(3) of Table 3. The results in columns (1)–(3) of Table 6 show that the match effect does not predict the destination firm fixed effect in our data. The coefficient estimate is insignificant and the model explanatory power ($R^2$) is practically zero. A second test proposed by Lazear et al. (2015) examines whether the match effect in a worker's initial employment spell explains the quartile of the next boss's fixed effect. Translating this to our context, we regress the match effect in an inventor's first employment spell on the decile of her second employer. The results in columns (4)–(6) of Table 6 indicate that the estimated effects are not significantly different from one another. Moreover, they jointly explain only a tiny proportion of the overall variance in the model. We find similar results when we repeat this analysis including all moves in an inventor's career. Taken together, these results suggest that, as in Card et al. (2013) and Lazear et al. (2015), endogenous mobility is not likely to be a critical concern for our analysis.

### 5.3. Robustness of Matching Results across Industries

If inventor and industry effects are positively correlated between industries but negatively correlated within industries, then the addition of these two opposite effects can lead to no, or as in our case, negative, correlation between the inventor and firm effects (Abowd et al. 2000, Postel-Vinay and Robin 2002). That is, industry-specific correlation patterns obscure potential positive correlation between the inventor and firm effects. Related, the extent of the sparse network problem and the bias it introduces can be different across industries. To examine whether our results are robust, we examine them by industry. We show that the matching correlations we observe

**Table 6.** Regression Results for Endogenous Mobility Tests

| | Destination firm FE | | | Origin match effect | | |
|---|---|---|---|---|---|---|
| Dependent variable | 1 | 2 | 3 | 4 | 5 | 6 |
| Sample | 1+ observations | 6+ observations | 10+ observations | 1+ observations | 6+ observations | 10+ observations |
| *Origin match effect* | −0.00 | −0.00 | −0.00 | | | |
| | (0.002) | (0.002) | (0.003) | | | |
| *Postmove firm FE* | | | | | | |
| *Second decile* | | | | −0.02 | 0.03 | 0.05 |
| | | | | (0.023) | (0.019) | (0.029) |
| *Third decile* | | | | 0.00 | 0.03 | 0.03 |
| | | | | (0.021) | (0.017) | (0.025) |
| *Fourth decile* | | | | 0.00 | 0.01 | 0.02 |
| | | | | (0.020) | (0.017) | (0.025) |
| *Fifth decile* | | | | −0.01 | 0.01 | 0.02 |
| | | | | (0.021) | (0.017) | (0.025) |
| *Sixth decile* | | | | −0.02 | −0.00 | −0.00 |
| | | | | (0.020) | (0.017) | (0.025) |
| *Seventh decile* | | | | −0.03 | −0.01 | −0.01 |
| | | | | (0.020) | (0.017) | (0.026) |
| *Eighth decile* | | | | −0.05$^{\dagger}$ | −0.03 | −0.02 |
| | | | | (0.021) | (0.019) | (0.027) |
| *Ninth decile* | | | | −0.01 | −0.01 | −0.01 |
| | | | | (0.023) | (0.024) | (0.032) |
| *Tenth decile* | | | | −0.04 | −0.02 | −0.01 |
| | | | | (0.023) | (0.021) | (0.037) |
| Constant | 0.07 | 0.09 | 0.22 | −0.04 | −0.06 | −0.06 |
| Observations | 89,743 | 73,698 | 54,017 | 59,682 | 44,963 | 29,592 |
| $R^2$ | 0.000 | 0.000 | 0.000 | 0.001 | 0.002 | 0.003 |

*Notes.* This table shows results for the endogenous mobility tests described in Lazear et al. (2015, table 5). In columns (1)–(3), we regress the destination firm fixed effect on the match effect in the origin firm. In columns (4)–(6), we regress the match effect in the inventor's first employment spell on the next firm's decile in the firm fixed effects distribution, using the first decile of firm fixed effects as the reference category. All match effects and firm fixed effects relate to the full AKM model reported in columns (1)–(3) in Table 3. See Lazear et al. (2015) for more information on these tests. Standard errors clustered at the firm level are reported in parentheses.

$^{*}p < 0.01$; $^{\dagger}p < 0.05$; $^{\ddagger}p < 0.1$.

previously are not merely aggregate results but borne out in each of the six NBER patent-based technology categories (Online Appendix A, Figures A2 and A3).

## 6. Theoretical Analysis

### 6.1. Potential Explanations for AKM and Matching Results

If our empirical findings are not caused by biased estimation, then they raise the question of what economic mechanism generates negative assortative matching between the inventor and firm effects, as well as positive assortative matching among inventors, both of which are observed in the data. Recent labor economics literature seeks to explain evidence for weak positive assortative matching, seemingly random matching, or negative assortative matching between firms and workers (Abowd et al. 1999, 2002; Goux and Maurin 1999; Andrews et al. 2008; Gruetter and Lalive 2009). We extend this literature to the case of innovation, noting that inventors are different from *normal* workers.

The classical model of matching (Becker 1973) is associated with the prediction of positive assortative matching. As mentioned in the previous paragraph, the empirical literature is not always consistent with this prediction. Eeckhout (2018) considers three classes of explanations to reconcile theory and empirical evidence. The first is job search frictions: the idea that the formation of new matches between workers and firms is costly and thus worker and firm types fail to match perfectly. The second explanation rests on the inability of firms to perfectly observe worker productivity, which also makes it more costly to obtain positive assortative matching (PAM). Finally, the third explanation points out that workers and firms may match on more than one dimension, and the multidimensionality of matching can produce both PAM and NAM (negative assortative matching) (Lindenlaub 2017, Eeckhout and Kirchner 2018).

Our model presents a variation on the third explanation. We assume that matching is driven not only by the innovation production function but also by inventor utility, which in turn depends on job characteristics other than wage compensation (specifically, innovation output). In this sense, our model formalizes the idea put forward by Bonhomme et al. (2019), who note that workers may be willing to sacrifice wages in exchange for better nonwage job characteristics or *amenities*. Lamadon et al. (2019) show the relative importance of amenities (e.g., proximity to work, flexible work schedules, preference for the type of tasks performed) for worker sorting in the U.S. labor market. In our case, we posit that, for inventors, an important amenity is given by innovation output. Specifically, our theoretical model considers a worker utility function with two inputs: wage and innovation output.

### 6.2. The Model

In this section, we develop a theoretical framework to formally illustrate how the presence of inventors' taste for amenities, specifically innovation output, may yield the surprising pattern of negative assortative matching between innovating firms and workers suggested by the data. Accordingly, we focus primarily on this feature and then extend the framework to address inventor-inventor matching as well.

Similar to our empirical section, an inventor's innovation function is given by

$$\log(y) = \alpha + \phi, \tag{4}$$

where $y$, as before, is the number of patents of an individual inventor-firm pair, $\alpha$ an indicator of the inventor's ability, and $\phi$ a measure of the firm's innovation strength. The inventor's utility function, in turn, is given by

$$u = w + f(y), \tag{5}$$

where $w$ is wage earnings. We assume that $f(y)$ can be written as $f(y) = g(\log(y))$ and that $g'(y) > 0$ and $g''(y) < 0$, that is, the inventor's utility is increasing and concave in the inventor's output and the marginal utility of innovation output $y$ declines at a faster rate than $1/y^2$. This is not an innocuous assumption. In fact, it drives much of the results that follow. That said, we believe that Equation (5) describes well the reality of inventor motivation. First, in addition to monetary compensation, inventors care about the result of their efforts. This may result from self-esteem considerations, career concerns, or other factors. Second, consistent with standard models of agent preferences, we posit that the inventor's marginal utility from innovation is decreasing: an extra patent matters a lot more for an inventor with a low number of patents than for an inventor with a large number of patents under her belt.[10] We do, moreover, make the stronger assumption that the marginal utility of innovation declines at a fast rate (specifically, faster than the logarithm function).

The simplest model that is complex enough to address the issue of matching between inventors and firms is a model with two firms of different type and two inventors of different type (at the end of the section we consider a variety of extensions).

Specifically, we assume that one of the inventors is type $\alpha_L$ and the other $\alpha_H$. Similarly, one of the firms is type $\phi_L$ and the other $\phi_H$. $\alpha_L$ and $\alpha_H$ correspond, respectively, to low-human capital and high-human capital inventors. Similarly, $\phi_L$ and $\phi_H$ correspond to firms with low and high firm-specific innovation specific capabilities, respectively.

We assume the firm's profit function is given by

$$\pi = \log(y) - w. \tag{6}$$

Similarly to the inventor utility function, the assumption that firm profits are convex in y can be understood as a reflection of decreasing marginal product.[11] For example, suppose that the firm is trying to solve a specific problem (e.g., to cure a disease) and that revenues are proportional to the probability of achieving that goal. Suppose moreover that each patent corresponds to a try at solving the problem. If these tries are independent draws, then the probability of achieving the desired goal is given by $h(y) = 1 - (1 - \rho)^y$, where $\rho$ is the probability that a patent leads to a solution to the problem. As can be seen, $h(y)$ is a concave function of $y$.[12] To conclude the model's description, we assume that a firm can hire only one worker.[13]

The process of matching firms with workers, as well as the process of wage setting, can be complex and highly idiosyncratic when the worker is an inventor. In other words, many of these inventors are *superstars* who are paid a negotiated wage rather than a $w$ from a salary scale. Given the complexity of the process, it makes sense to analyze the problem as a coalitional or cooperative game. The terms *coalitional* and *cooperative*, common as they may be in the literature, are probably not as appropriate as *protocol free*, which better describes the idea: instead of assuming a specific extensive form (i.e., rigid rules regarding who does what and when in the game), we simply address the identity of who does business with whom (matches) and what payoff they get (more specifically, bounds on what their payoffs are). Specifically, following a common approach in this type of problems we look for the core of the game in question.

**Definition 1.** *The core of the firm-inventor matching game is a set of matches and payoff values such that no firm-inventor pair can increase their payoffs by forming a deviating coalition.*

We are now ready to present our central theoretical result.

**Proposition 1.** *The core of the firm-inventor matching game includes a unique element. It is characterized by negative assortative matching, that is, a low-type inventor is matched with a high-type firm and vice-versa.*

The proofs of our theoretical results are presented in Online Appendix C.[14] At first, the result may seem a little counterintuitive. One might expect the high-type firm to be matched with the high-type inventor. However, the equilibrium is not determined by the maximization of the surplus of any particular match—as the $HH$ match would be—but rather by the maximization of the sum of surplus levels over all matches. In the present case, the maximum joint value of a particular match corresponds to the high-type inventor working at a low-type firm.

To understand the intuition for Proposition 1, note that, although the innovation function is supermodular, the inventor's utility function is very concave (i.e., the inventor's marginal utility of innovation is decreasing at a high rate). This implies that the match function is submodular and that maximizing total value calls for negative assortative matching. In fact, a boost in innovation is worth a lot more for a low-type inventor than it is for a high-type inventor, and a high-type firm provides that boost better than a low-type firm.

In general, the core does not indicate the exact equilibrium payoff received by each player. However, it does provide bounds on the payoff received by each player. In the present context, we are able to prove the following result regarding equilibrium inventor wages.

**Proposition 2.** *The wage paid by the low-type firm, $w_L$, is greater than the wage paid by the high-type firm, $w_H$.*

Together, Propositions 1 and 2 suggest that inventors trade off innovation and wage (the two inputs into their utility function) when they choose what firm to work for. At the high-type firm, inventors get a bigger boost to their innovation output but receive a lower wage. By contrast, at the low-type firm, inventors get a lower boost to their innovation output but receive a higher wage. Note that Proposition 2 does not exactly imply the tradeoff just described because $w_L$, for example, is the wage paid by a low-type firm to a high-type inventor and not to a low-type inventor. However, this intuition stands.

### 6.3. Model Extensions

One natural extension of this model is to consider an arbitrary number of firms, $n > 2$. If we keep the number of inventors at $n$ too, then the extension is relatively straightforward. That said, we do not think there is really any additional intuition gained by analyzing the $n$ firm case in this way.

A more relevant extension would be to allow for more inventors than firms. This possibility opens the question of sorting among inventors. The simplest case is when there are two firms of different type, as before, but now four inventors, two of each type. And we assume that a firm operates with exactly two inventors. If innovation functions are additive in inventor outputs, and if inventor utility is the same as in Equation (5), then it is straightforward to show that the core corresponds to low-type inventors working for the high-type firm and high-type inventors working for the low-type firm, just as in Proposition 1. The idea is that, because the profit function is additive in inventor output, the problems of matching each of the inventors are separable. In other words, Proposition 1 applies to each of the inventors in each firm. The result of this multi-inventor matching game is that, in equilibrium, we observe positive assortative matching among inventors. However, this PAM is a result of the NAM between inventors and firms.

## 7. Concluding Thoughts

In this paper, we establish that inventor-specific skills are 5–10 times more important than firm-specific capabilities for explaining the variance in the inventive performance of inventors. The relatively small effect of firm-specific capabilities, which include capabilities such as corporate culture and organizational routines, and take several years to build, may explain why several decades of research has not uncovered a clear advantage for established firms in innovation. Our findings make the case for a more central role for human capital in theories of the firm and studies of competitive advantage.

We also study the matching of inventors to employers—a topic of central importance for labor economists and human resource management. We find that high human capital inventors match with firms that (i) have weak firm-specific invention capabilities and (ii) employ other talented inventors. Our theoretical analysis rationalizes these empirical patterns by incorporating preferences for innovation outputs into workers' utility function. This analysis also suggests that firms can seek to enhance their competitive advantage either by employing inventors with lower human capital and contributing to their innovativeness (through firm capabilities) or by serving as a platform for highly talented, but expensive, workers.

Of course, our analysis has several limitations. First, we use a specific measure of innovation based on patent counts—indeed not all inventors and firms have the same propensity to patent their inventions. Second, our AKM analyses explain around 44% of the variation in innovation output, leaving a significant role for unobserved inventor-firm-year specific idiosyncratic factors. These factors could include, for example, job-fit, learning, and experience (both individual and firm-level), and changes in firm-level patenting policies, legal personnel, or leadership. Although well suited to account for unobserved inventor and firm heterogeneity, the standard AKM cannot be leveraged to examine complementarities in innovation arising from interactions between worker and firm attributes. Third, the AKM model is static, in the sense that it does not take into account worker mobility's likely dependence on innovation performance (Bonhomme et al. 2019). We address this limitation in multiple ways described in Sections 4 and 5 by using the rolling window technique and tests to rule out different types of biases. These methods allow us to partially, but not perfectly, deal with time-varying omitted variables such as firm leadership or governance that may influence inventor output and firm capability and matching. Fourth, our analyses establish the importance of inventor-specific ability for explaining variance in innovation performance, but we do not know what drives the fixed effects. Inventor fixed effects likely subsume the influences of a variety of intrinsic traits (e.g., innate ability and persistence) and acquired experiences (e.g., education). Fifth, working at a firm may have a persistent effect on inventors' human capital that may be inaccurately attributed to the worker. Unpacking and identifying changes in firm capabilities and the ingredients of human capital presents promising avenues for future research.

Finally, although overall human capital appears more important than firm capabilities in explaining variance in inventor productivity, firms can improve their innovation output by investing in capabilities. Our results leave open the question of whether firms can improve their overall innovation performance through a strategy that builds firm capabilities or focusses on hiring and retaining top inventors. The answer to the question is likely to depend on the relative costs of innovation inputs.

### Acknowledgments

The authors are grateful to Ashish Arora (department editor), the anonymous associate editor, and four referees for their excellent suggestions throughout the review process. They thank Pierre Azoulay, Stephane Bonhomme, Ines Black, Gino Cattani, Francis Kramarz, Thibaut Lamadon, Franco Malerba, Jan van Ours, and Enrico Pennings for helpful comments and suggestions. The authors also thank conference and seminar participants at the INFORMS Annual Meeting; Academy of Management Meetings; SMS Conference; EARIE Conference; R&D Management Conference; The 10th Workshop on The Organization, Economics and Policy of Scientific Research; the 2018 BYU-Utah Winter Strategy Workshop; Erasmus School of Economics; Indian Institute of Management, Bangalore; Saunders College of Business at Rochester Institute of Technology; Cox School of Business at Southern Methodist University; and New York University's Stern School of Business. All errors are the authors' own.

### Endnotes

[1] Specifically, we assume that the innovation rate $y$ is an exponential function of $\alpha + \phi$, where $\alpha$ is the inventor's type and $\phi$ is the firm's type. This implies that the cross–partial derivative of $y$ with respect to $\alpha$ and $\phi$ is positive, so the innovation function is supermodular in its inputs.

[2] This, per se, would not suffice for submodularity. However, we effectively assume that the inventor's utility function is more concave than the firm's payoff function is convex.

[3] We exclude from AKM analyses observations not part of the largest connected network (about 1%) because there is no basis to normalize

the estimated fixed effects to a reference firm or inventor across unconnected networks.

[4] Because the model explains 44.7% of the overall variation in inventor performance, inventor heterogeneity and firm heterogeneity account for 76% (34.1/44.7) and 7% (3.2/44.7) of the variance explained by our model, respectively. See Cornelissen (2008), p.183, for more discussion on the interpretation of these numbers.

[5] Online Appendix A, Table A2, also reports pairwise correlation coefficients among inventor- and firm-fixed effects obtained from the different estimations described above. The coefficients are all higher than 0.77, suggesting the robustness of our findings to the different specifications and samples.

[6] We construct these and all further bootstrapped standard errors by re-estimating the entire 5-year and 10-year rolling window procedure 100 times with replacement. See table notes for more details.

[7] See Online Appendix A, Table A5, for the full estimation results.

[8] As noted in Section 4.1, the presence of match dynamics among firms and inventors run counter to the assumption of our baseline AKM that firm and inventor types are constant through the sample. We examine whether the mobility/matching results obtained through the rolling window procedure (reported in Table 4) differ from mobility/matching results obtained from the baseline AKM estimates that are time invariant (column (1) of Table 3). The estimates, shown in Online Appendix A, Table A6, confirm our findings of PAM among inventors and coinventors and NAM between inventors and firms obtained from the rolling window analyses. The effect sizes pertaining to inventors-firm matching (i.e., NAM) are very close to those obtained through the rolling window estimation, and those pertaining to coinventor matching (i.e., PAM) are stronger when we do not use rolling windows.

[9] The distributions for our three samples, depicted in Online Appendix A, Figure A1, show that the variance approximation is indeed very close to 1 for all but the highest deciles of the distribution.

[10] A similar consideration applies to academic publications as well.

[11] We could also place a coefficient in front of $\log(y)$ to compare units of $y$ to units of $w$. However, the qualitative nature of our results would not change.

[12] Although this does not correspond precisely to a log function, the assumption of independence across draws is also unlikely to hold.

[13] We could explicitly model this in the form of decreasing returns to scale but that would unnecessarily complicate the analysis. That said, some assumption is required to avoid the outcome of one large firm that hires all of the inventors in the economy.

[14] The proof of Proposition 1 basically adapts features of results in Gale and Shapley (1962) and Becker (1973). In other words, there is nothing novel about Proposition 1 other than a slightly different proof strategy and the fact that we derive the result in the specific context of a firm-inventor game, that is, with a particular set of players and payoff functions. Moreover, although most of the literature has focused on super-modular payoff functions and positive assortative matching, Proposition 1 deals with negative assortative matching.

## References

Abowd J, Creecy R, Kramarz F (2002) Computing person and firm effects using linked longitudinal employer-employee data. Technical Report 2002-06, U.S. Census Bureau, Washington, DC.

Abowd J, Kramarz F, Margolis D (1999) High wage workers and high wage firms. *Econometrica* 67(2):251–333.

Abowd J, McKinney K, Schutte I (2019) Modeling endogenous mobility in earnings determination. *J. Bus. Econom. Statist.* 37(3): 405–418.

Abowd J, Kramarz F, Lengermann P, Perez-Duarte S (2004) Are good workers employed by good firms? A test of a simple assortative matching model for France and the United States. Working paper, Cornell University, Ithaca, NY.

Abowd JM, Kramarz F, Lengermann P, Roux S (2000) Inter-industry and firm-size wage differentials: New evidence from linked employer-employee data. Working paper, Cornell University, Ithaca, NY.

Agarwal R, Ohyama A (2013) Industry or academia, basic or applied? Career choices and earnings trajectories of scientists. *Management Sci.* 59(4):950–970.

Agarwal R, Echambadi R, Franco A, Sarkar M (2004) Knowledge transfer through inheritance: Spin-out generation, development and performance. *Acad. Management J.* 47(4):501–522.

Andrews ML, Gill L, Schank T, Upward R (2008) High wage workers and low wage firms: negative assortative matching or limited mobility bias? *J. Royal Statist. Soc. A* 171(3):673–697.

Arora A (1995) Licensing tacit knowledge: Intellectual property rights and the market for know-how. *Econom. Innovative New Tech.* 4(1):41–59.

Arora A, Belenzon S, Rios LA (2014) Make, buy, organize: The interplay between research, external knowledge, and firm structure. *Strategic Management J.* 35(3):317–337.

Azoulay P, Graff Zivin J, Wang J (2010) Superstar extinction. *Quart. J. Econom.* 125(2):549–589.

Barney J (1986) Organizational culture: Can it be a source of sustained competitive advantage? *Acad. Management Rev.* 11(3):656–665.

Becker GS (1973) A theory of marriage: Part I. *J. Political Econom.* 81(4):813–846.

Bertrand M, Schoar A (2003) Managing with style: The effect of managers on firm policies. *Quart. J. Econom.* 118(4):1169–1208.

Bessen J (2009) Matching patent data to Compustat firms. Accessed April 22, 2020, http://users.nber.org/~jbessen/matchdoc.pdf.

Bloom N, Eifert B, McKenzie D, Mahajan A, Roberts J (2013) Does management matter: Evidence from India. *Quart. J. Econom.* 128(1):1–51.

Bonhomme S, Lamadon T, Manresa E (2019) A distributional framework for matched employer employee data. *Econometrica* 87(3):699–739.

Campbell B, Coff R, Kryscynski D (2012) Rethinking sustained competitive advantage from human capital. *Acad. Management Rev.* 37(3):376–395.

Card D, Heining J, Kline P (2013) Workplace heterogeneity and the rise of West-German wage inequality. *Quart. J. Econom.* 128(3):967–1015.

Cohen WM, Nelson RR, Walsh JP (2000) Protecting their intellectual assets: Appropriability conditions and why U.S. manufacturing firms patent (or not). NBER Working Paper No. 7552, National Bureau of Economic Research, Cambridge, MA.

Cornelissen T (2008) The Stata command felsdvreg to fit a linear model with two high-dimensional fixed effects. *Stata J.* 8(2):170–189.

Eeckhout J (2018) Sorting in the labor market. *Annual Rev. Econom.* 10(2018):1–29.

Eeckhout J, Kircher P (2018) Assortative matching with large firms. *Econometrica* 86(1):85–132.

Ewens M, Rhodes-Kropf M (2015) Is a VC partnership greater than the sum of its partners? *J. Finance* 120(3):1081–1113.

Fallick B, Fleischman C, Rebitzer JB (2006) Job-hopping in Silicon Valley: Some evidence concerning the microfoundations of a high-technology cluster. *Rev. Econom. Statist.* 88(3):472–481.

Gale D, Shapley LS (1962) College admissions and the stability of marriage. *Amer. Math. Monthly* 69(1):9–15.

Ge C, Huang KW, Png IP (2016) Engineer/scientist careers: Patents, online profiles, and misclassification bias. *Strategic Management J.* 37(1):232–253.

Goux D, Maurin E (1999) Persistence of interindustry wage differentials: A reexamination using matched worker-firm panel data. *J. Labor Econom.* 17(3):492–533.

Graham J, Li S, Qiu J (2012) Managerial attributes and executive compensation. *Rev. Financial Stud.* 25(1):144–186.

Griliches Z, ed. (1984) *R&D, Patents and Productivity* (University of Chicago Press, Chicago).

Groysberg B (2010) *Chasing Stars: The Myth of Talent and the Portability of Performance* (Princeton University Press, Princeton, NJ).

Groysberg B, Lee L-E, Nanda A (2008) Can they take it with them? The portability of star knowledge workers' performance. *Management Sci.* 54(7):1213–1230.

Gruetter M, Lalive R (2009) The importance of firms in wage determination. *Labour Econom.* 16(2):149–160.

Hall BH, Ziedonis RH (2001) The patent paradox revisited: An empirical study of patenting in the U.S. semiconductor industry, 1979–1995. *RAND J. Econom.* 32(1):101–128.

Hall BH, Jaffe AB, Trajtenberg M (2001) The NBER patent citation data file: Lessons, insights and methodological tools. NBER Working Paper 8498, National Bureau of Economic Research, Cambridge, MA.

Hatch NW, Dyer JH (2004) Human capital and learning as a source of sustainable competitive advantage. *Strategic Management J.* 25(12):1155–1178.

Huckman RS, Pisano GP (2006) The firm specificity of individual performance: Evidence from cardiac surgery. *Management Sci.* 52(4):473–488.

Iranzo S, Schivardi F, Tosetti A (2008) Skill dispersion and firm productivity: An analysis with employer-employee matched data. *J. Labor Econom.* 26(2):247–285.

Jaffe A (1986) Technological opportunity and spillovers of R&D: Evidence from firms' patents, profits, and market value. *Amer. Econom. Rev.* 76(5):984–1001.

Jaravel X, Petkova N, Bell A (2018) Team-specific capital and innovation. *Amer. Econom. Rev.* 108(4-5):1034–1073.

Jochmans K, Weidner M (2019) Fixed effects regressions on network data. *Econometrica* 87(5):1543–1560.

Jones BF (2009) The burden of knowledge and the 'death of the renaissance man': Is innovation getting harder? *Rev. Econom. Stud.* 76(1):283–317.

Kline P, Petkova N, Williams H, Zidar O (2019) Who profits from patents? Rent-sharing at innovative firms. *Quart. J. Econom.* 134(3):1343–1404.

Kogan L, Papanikolaou D, Seru A, Stoffman N (2017) Technological innovation, resource allocation, and growth. *Quart. J. Econom.* 132(2):665–712.

Lamadon T, Mogstad M, Setzler B (2019) Imperfect competition, compensating differentials and rent sharing in the U.S. labor market. Working paper, University of Chicago, Chicago.

Lazear EP (2009) Firm-specific human capital: A skill-weights approach. *J. Political Econom.* 117:914–940.

Lazear EP, Shaw KL, Stanton C (2015) The value of bosses. *J. Labor Econom.* 33(4):823–861.

Li G-C, Lai R, D'Amour A, Doolin DM, Sun Y, Torvik VI, Yu AZ, Fleming L (2014) Disambiguation and co-authorship networks of the U.S. patent inventor database (1975–2010). *Res. Policy* 43(6):941–955.

Lindenlaub I (2017) Sorting multidimensional types: Theory and application. *Rev. Econom. Stud.* 84(2):718–789.

Malerba F (2005) Sectoral systems of innovation: A framework for linking innovation to the knowledge base, structure and dynamics of sectors. *Econom. Innovations New Tech.* 14(1-2): 63–82.

Martinez E, Beaulieu N, Gibbons R, Pronovost P, Wang T (2015) Organizational culture and performance. *Amer. Econom. Rev.* 105(5):331–335.

Mayer KJ, Somaya D, Williamson IO (2012) Firm-specific, industry-specific and occupational human capital, and the sourcing of knowledge work. *Organ. Sci.* 23(5):1311–1329.

McGahan A, Porter M (1997) How much does industry matter, really? *Strategic Management J.* 18(S1):15–30.

McGahan A, Porter M (2002) What do we know about variance in accounting profitability? *Management Sci.* 48(7):834–851.

Melero E, Palomeras N, Wehrheim D (2020) The effect of patent protection on inventor mobility. *Management Sci.* Forthcoming.

Mollick E (2012) People and process, suits and innovators: The role of individuals in firm performance. *Strategic Management J.* 33(9):1001–1015.

Nelson R, Winter SG (1982) *An Evolutionary Theory of Economic Change* (Belknap Press, Cambridge, MA).

Palomeras N, Melero E (2010) Markets for inventors: Learning-by-hiring as a driver of mobility. *Management Sci.* 56(5):881–895.

Peeters T, Salaga S, Juravich M (2020) Matching and winning? The impact of upper- and middle managers on team performance in Major League Baseball. *Management Sci.* Forthcoming.

Postel-Vinay F, Robin JM (2002) Equilibrium wage dispersion with worker and employer heterogeneity. *Econometrica* 70(6): 2295–2350.

Rumelt R (1984) Toward a strategic theory of the firm. Lamp R, ed. *Competitive Strategic Management* (Prentice-Hall, Englewood Cliffs, NJ).

Schilling MA (2018) *Quirky: The Remarkable Story of the Traits, Foibles, and Genius of Breakthrough Innovators Who Changed the World* (Public Affairs, New York).

Singh J, Agrawal A (2011) Recruiting for ideas: How firms exploit the prior inventions of new hires. *Management Sci.* 57(1): 129–150.

Singh J, Fleming L (2010) Lone inventors as sources of breakthroughs: Myth or reality? *Management Sci.* 56(1):41–56.

Stern S (2004) Do scientists pay to be scientists? *Management Sci.* 50(6):835–853.

Wernerfelt B (1984) A resource-based view of the firm. *Strategic Management J.* 5(2):171–180.

Wuchty S, Jones BF, Uzzi B (2007) The increasing dominance of teams in production of knowledge. *Science* 316(5827):1036–1039.

Zucker L, Darby MR, Torero M (2002) Labor mobility from academe to commerce. *J. Labor Econom.* 20(3):629–660.