Melero, E., Palomeras, N. & Wehrheim, D. (2020). The Effect of Patent Protection on Inventor Mobility. *Management Science*, 66 (12), pp. 5485-5504.

DOI: 10.1287/mnsc.2019.3500

The effect of patent protection on inventor mobility*

Eduardo Melero[†a], Neus Palomeras[a], and David Wehrheim[b]

[a]Universidad Carlos III de Madrid
[b]IESE Business School

Forthcoming, Management Science

ABSTRACT

This article investigates the effect of patent protection on the mobility of early-career employee-inventors. Using data on patent applications filed at the US Patent and Trademark Office between 2001 and 2012 and examiner leniency as a source of exogenous variation in patent protection, we find that one additional patent granted decreases the likelihood of changing employers, on average, by 23 percent. This decrease is stronger when the employee has fewer coinventors, works outside the core of the firm, and produces more basic-research innovations. These findings are consistent with the idea that patents turn innovation-related skills into patent-holder-specific human capital.

KEYWORDS: inventors, patents, mobility, specific human capital, examiner leniency

*We are grateful to Ashish Arora, Manuel F. Bagüés, Stefano H. Baruffaldi, Bruno Cassiman, Emre Ekinci, Andrea Fosfuri, Karin Hoisl, Peter Mueser, Pablo Ruiz-Verdú, Evan Starr, Giovanni Valentini, and Natalia Zinovyeva, as well as the conference participants at the SMS Special Conference on Strategic Human Capital (Milan, 2017); the Barcelona GSE Summer Forum (2017); the SEI Faculty Workshop (Madrid, 2017); the Munich Summer Institute (2017); the 10th Annual Conference on Innovation Economics (Chicago, 2017); the INSEAD Workshop on Innovators Mobility (Fontainebleau, 2018); and seminar participants at Universidad Carlos III de Madrid (2017) for their useful comments and suggestions. All errors are our own. The authors acknowledge financial support from the Spanish Ministry of Science, Innovation and Universities through projects ECO2015-65599-P, ECO2015-69615-R, PGC2018-098767-B-C21, PGC2018-096316-B-I00 and PGC2018-094418-B-I00 and the Madrid regional government through project EARLYFIN-CM #S2015/HUM-3353.

†Corresponding author: eduardo.melero@uc3m.es

1. Introduction

The inter-firm mobility of skilled workers is a key factor behind the dissemination of knowledge, as first noticed by Arrow (1962). Especially relevant in this respect is the case of R&D workers, given the tacit technological knowledge that they frequently embody. Accordingly, past research has been interested in understanding the determinants of this type of mobility at the individual (Hoisl, 2007; Palomeras and Melero, 2010; Ganco, 2013), organizational (Paruchuri et al., 2006) and institutional (Marx et al., 2009; Hombert and Matray, 2016; Akcigit et al., 2016) levels. A good understanding of these determinants may allow firms to design innovation strategies to profit from (and not be damaged by) the knowledge spread by moving workers.

The role of institutional factors on the mobility of R&D workers is of particular interest, as its consequences for knowledge diffusion concern both managers and policy-makers. Previous research has paid particular attention to the laws that govern noncompete agreements. These agreements expressly prevent departing employees from working for competitors and have a direct negative impact on the interfirm mobility of skilled talent (Gilson, 1999; Fallick et al., 2006; Marx et al., 2009; Marx, 2011). Nevertheless, the literature has overlooked other legal mechanisms that may pose implicit barriers to mobility by preventing R&D workers from transferring knowledge across employers.[1]

This paper contributes to filling this gap in the literature by investigating the effect of patent protection on inventor mobility. The premise of our study is that inventors possess tacit knowledge that can contribute to the implementation (or further development) of the innovations that they have created, as documented by Zucker et al. (1998). At the same time, firms' proprietary technologies usually operate as complementary assets to employees' human capital, turning it less redeployable in other companies (Campbell et al., 2012). As a consequence, we suggest that patent protection turns the innovation-related skills of the R&D workers into patent-holder-specific human capital and therefore decreases inventor mobility. We formalize these ideas in a simple model and derive two additional implications concerning this negative effect of patent protection on mobility: (i) it inversely depends on the existence of other firm-specific human capital (e.g., complementarities between the inventor and the firm's resources), and (ii) it is more intense when the role of the inventor in the implementation of the innovation

[1] An exception is Ganco et al. (2015), who investigate the negative relationship between a firm's reputation for IP enforcement and outbound mobility. They implicitly assume the instrumentality of patent protection in generating a retention effect at high-litigiousness firms. However, they do not explicitly analyze the role of patents.

is more relevant (e.g., for innovations that build on basic research).

We empirically investigate the effect of patenting on inventor mobility by comparing the trajectories of inventors with different numbers of granted patents among those with a given number of applications. This analysis poses an important methodological difficulty. Inventors with high (unobservable) talent are arguably more likely to produce innovations that satisfy legal patentability requirements than inventors with less talent. Thus, a simple comparison of inventors with approved applications and those without them may lead to inferences that mirror only their underlying dissimilarities. We address this challenge by using differences in leniency across patent examiners as a source of exogenous variation in the number of granted patents (Sampat and Williams, 2019) to estimate the causal relationship. The validity of this instrument is supported by qualitative evidence regarding the allocation of patent applications to examiners at the United States Patent and Trademark Office (USPTO) (Cockburn et al., 2003; Lemley and Sampat, 2012) and by our own exogeneity tests.

Our analysis is based on the career trajectories of 67,775 inventors who filed their first patent application with the USPTO between 2001 and 2012. By identifying inventors' early career moves from patent application data, our results show a negative effect of patent protection on mobility. One additional patent granted (due to a "lucky" examiner assignment) leads, on average, to a 23 percent decrease in the probability of changing employers. As expected, this negative effect is stronger for inventors with fewer coauthors, those who work outside their firms' core technologies, and those working on innovations in the realm of basic research.

Our findings provide new insights into several domains. First, we contribute to the analysis of the impact of institutional factors on the mobility of high-skilled workers. We suggest that patents preclude some inventor moves that would otherwise occur. To the extent that reduced mobility leads to worse employer-employee matches and lower diffusion of tacit knowledge, this result has potentially important policy implications for the patent system as an institution intended to foster innovation. Second, in terms of managerial implications, our results suggest that patents may work as devices to retain inventive talent. Thus, patent-intensive innovative firms may find it particularly profitable to implement long-term-oriented human resources practices. Finally, our findings have a methodological implication by documenting that patent grant events have a significant impact on the behavior of firms and inventors. This implies that any research in the field of inventor career dynamics exclusively using data from granted patent applications may suffer from selection bias.

2. Theoretical framework

2.1. A model of patent protection and inventor mobility

Inventor employees are responsible for the overwhelming majority of patented innovations, but most of them do not own the corresponding property rights. Employment contracts usually stipulate that the employer, not the inventor, holds the property rights on innovations made on the job (Merges, 1999). This generates interesting implications for the relationship between patent protection and inventor mobility, which we develop below in a simple model.

We consider an inventor-employee who creates an innovation for which her firm applies for patent protection. We distinguish two additive components of the inventor's marginal productivity with a given employer (both of them fully observable): (i) *inventive productivity*, which reflects the match between the inventor's inventive abilities and the employer, and (ii) *innovation-related skills*, which reflect the tacit knowledge the inventor has acquired on the specific innovation while developing it.

First, we assume that the value of the inventive productivity (not related to the focal innovation) of a given inventor with each potential employer follows a random process. Each combination of inventor i and employer k generates some match-idiosyncratic productivity $\phi_{ik} \geqslant 0$.[2] To study mobility, the relevant inventive productivities to be compared are ϕ_{ij}, her match with the current employer (j), and $\bar{\phi}_{ij} = \max_{k \neq j}\{\phi_{ik}\}$, her best match among the rest of the (otherwise identical) alternative employers. Only the difference between these two components is relevant for the moving behavior. Hence, we define $\theta_i = \phi_{ij} - \bar{\phi}_{ij}$ as the *relative quality of the match* of the inventor with her current employer, and we assume that it follows a strictly increasing cumulative distribution function $G(\theta_i)$.

The second (and key) component of the inventor's marginal productivity concerns her (focal) innovation-related skills. In line with the evidence provided by Zucker et al. (1998), we assume that the development of an innovation into an actual product ready for market launch, further development, or internal implementation is not a trivial process and benefits considerably from the involvement of its creators. To characterize the value of these skills, we assume that (i) the new technology is public knowledge by the time of the patent decision due to the earlier publication of the application and (ii) patent rights provide their owners with an effective

[2] Although we analyze the moving decision in a given moment of time, ϕ_{ik} changes over time. Thus, a given inventor is not necessarily employed by her best match in that focal moment, and a change of employer may increase her inventive productivity.

monopoly. Therefore, the patent office's decision determines whether the applicant firm receives monopoly rights on the innovation or any firm in the market can implement it. We abstract from technology markets by assuming that transactions in them are prohibitively costly.

To compute the value of innovation-related skills, we first define the profits that a firm obtains from exploiting an innovation generated by inventor i at firm j as $\pi^R_{m,n}$, where $m = \{i, -i\}$ denotes whether the firm employs the inventor, $n = \{j, -j\}$ denotes whether the firm is the patent applicant (and current employer), and $R = \{M, C\}$ indicates whether the innovation is exploited in monopoly or competition, according to whether the patent is granted.

Importantly, the innovation-related skills of the inventor grant the firm that employs her an efficiency advantatge in the exploitation of the innovation. We allow for these skills to be partly specific to the current employer since they may be complementary to some of its resources, such as technological assets or coinventors. Hence, in a setting without patent protection, the innovation-related skills of inventor i generate profits $\pi^C_{i,j}$ for her current employer and profits $\pi^C_{i,-j}$ for alternative employers, with $\pi^C_{i,j} > \pi^C_{i,-j}$. Because of competition, firms lacking the efficiency advantage contributed by the inventor would make zero profits. If the patent is granted, only the patent holder can implement the innovation and will obtain monopoly profits $\pi^M_{i,j}$ if the inventor stays and monopoly profits $\pi^M_{-i,j}$ if she leaves, with $\pi^M_{i,j} > \pi^M_{-i,j}$. Thus, the inventor's set of innovation-related skills generates rents of $\pi^M_{i,j} - \pi^M_{-i,j}$ if employed by the patent holder and has zero value for alternative employers.

We assume that each firm makes a wage offer to the inventor that equals her marginal productivity in the company.[3] This offer is the result of adding the two components described above. The inventor will move away from her current employer if she receives a better offer. In the absence of patent protection, the inventor will switch employers if and only if $\theta_i \leqslant \pi^C_{i,-j} - \pi^C_{i,j}$. This occurs when the added value of the firm-specific component of her innovation-related skills in competition is not large enough to compensate for a bad relative match with the current employer. With patent protection, the full set of innovation-related skills is patent-holder-specific. Moreover, the patent allows the current employer to obtain some monopoly profits from the innovation even if the inventor leaves. Therefore, the inventor will move if and only if $\theta_i \leqslant \pi^M_{-i,j} - \pi^M_{i,j}$. This occurs when the added value of the full set of her innovation-related skills in monopoly is not large enough to compensate for a bad relative match with the current

[3] Note that we only make this assumption for exposition purposes. With public information on all the components of productivity, the inventor always ends up being employed by the firm where she generates the highest value. The wage determination process may affect how the inventor and her final employer share the surplus with respect to her best employment alternative but not the mobility outcome.

employer.

Combining the conditions for mobility with and without patent protection, we find that patents have a negative effect on employee mobility if and only if the following is true:

$$\pi^M_{i,j} - \pi^M_{-i,j} > \pi^C_{i,j} - \pi^C_{i,-j} \tag{1}$$

The left-hand side of the inequality represents the added value in a monopoly setting of the full set of innovation-related skills. The right-hand side represents the value in a competitive setting of the subset of those skills that is specific to the current employer. Inequality (1) implies that some inventors with $\pi^M_{-i,j} - \pi^M_{i,j} < \theta_i < \pi^C_{i,-j} - \pi^C_{i,j}$ move to an employer with a better match in the absence of patent protection but are retained otherwise. Formally, Inequality (1) implies that patent grants will reduce the probability of inventor mobility, as $G(\pi^M_{-i,j} - \pi^M_{i,j}) - G(\pi^C_{i,-j} - \pi^C_{i,j}) < 0$ for any strictly increasing cumulative distribution function G. Based on this framework, we next develop a set of testable implications. The Online Appendix, Section A.1, provides further details on their mechanisms and formal derivation.

2.2. Testable implications

The key insight behind Inequality (1) is that patent protection affects innovation-related skills in two important ways. First, it converts them into fully patent-holder-specific human capital, which leads to a decrease in inventor mobility (the *human capital effect*). Second, patents change the market structure from competition to monopoly. Since competitive pressures lead companies to more intensely exploit efficiency gains, this change in market structure reduces the marginal value of innovation-related skills.[4] In particular, it reduces the value of the firm-specific component of these skills, leading to an increase in mobility (the *product-market effect*). If innovation-related skills are mostly general, the product-market effect will only be of second order, and the net impact of patent protection on mobility will be negative. On the contrary, if these skills are to a large extent specific to the current employer even in the absence of patents, then the product-market effect of patent protection could dominate the human capital effect and lead to a positive effect of patents on mobility. Thus, the *first implication* of our analysis is that patent protection will have a negative effect on mobility as long as the firm-specific component of innovation-related skills is small enough.

[4]Any cost advantage contributed by the inventor, for example, will have a stronger impact in competition than in monopoly (see Online Appendix, Section A.1).

A *second implication* of this double-edged effect of patents is that, as the firm-specific component of innovation-related skills increases, the human capital effect becomes less relevant, while the product-market effect gains importance. Thus, the hypothesized negative effect of patent protection on mobility is attenuated. In terms of Inequality (1), this component has a stronger impact on the right-hand side (i.e., competition) than on the left-hand side (i.e., monopoly) and makes the inequality less likely to hold. In empirical terms, we capture the relevance of firm-specific skills with the number of coauthors with whom the inventor works and her position within the technological core of the company. Hence, we expect that the effect of patent protection on inventor mobility should be positively moderated by these two factors.

The practical relevance of Inequality (1) depends on the importance of the skills of the inventor for the exploitation of the innovation. Accordingly, a *third implication* from our analysis indicates that the effect of patent protection on inventor mobility should be more intense when these skills are particularly useful for the implementation or further development of the innovation. This is likely to be the case for innovations with a more basic-research orientation. The development of such innovations into potentially marketable products is usually a long and uncertain journey, in which the role of tacit knowledge is critical (Zucker et al., 2002). In comparison, the implementation of applied innovations is a more straightforward process. Hence, we expect that the negative effect of patent protection on mobility should be more intense for basic-research innovations than for applied ones.

2.3. The role of inventor ability

The analysis of the relationship between patenting and inventor mobility has thus far been framed in terms of the value of the skills connected to the focal innovation. However, the existence of general inventor ability that cannot be observed by the econometrician may affect our capacity to detect that relationship in practice. In this section, a simple extension of our model provides some insights into how this unobserved inventor ability affects our empirical analysis.

Let us assume that the value the inventive productivity of inventor i with employer k, defined above as ϕ_{ik}, is a complementary combination between the inventor's general ability $\delta_i \geqslant 0$ and a *pure* match-idiosyncratic component of productivity $\epsilon_{ik} \geqslant 0$, so that $\phi_{ik} = \delta_i \epsilon_{ik}$. Accordingly, the relative quality of the match with the current employer is $\theta_i = \delta_i(\epsilon_{ij} - \bar{\epsilon}_{ij})$, where ϵ_{ij} is the *pure* match-idiosyncratic component of inventor i's productivity with her current employer, and

$\bar{\epsilon}_{ij} = \max_{k \neq j}\{\epsilon_{ik}\}$ is her best ϵ_{ik} among the other employers. Inventor ability δ_i amplifies any underlying differences in the quality of the match across employers. Crucially, it also increases the heterogeneity in the relative match with the current employer θ_i. In addition, our analysis in Section 2.1 indicates that only inventors with values of θ_i sufficiently below zero will move. Hence, some inventors with a negative difference in the pure match component $\epsilon_{ij} - \bar{\epsilon}_{ij}$ and high ability levels δ_i will switch firms, whereas otherwise similar inventors with less ability will remain with their current employer.

This simple extension indicates a positive relationship between inventor ability and mobility, as suggested by previous empirical studies (Palomeras and Melero, 2010; Akcigit et al., 2016; Moretti and Wilson, 2017). In our context, this possibility is an empirical concern because high-ability inventors may also enjoy a higher probability of obtaining patent protection for their innovations. The patentability requirements of novelty, usefulness and nonobviousness used by the USPTO suggest that this would be the case. Thus, if inventor ability can be observed by the market but not by the econometrician, the estimates of the effect of patenting on mobility will be upward biased unless some source of exogenous variation in patent grants is used.

2.4. Patents as signals of human capital

A related issue emerges when firms cannot observe inventor ability δ_i (and therefore also ignore ϕ_{ik}) and, given the patentability requirements, rely on patent grants as signals of such ability. It follows from the above discussion that expected inventor ability increases the heterogeneity in the expected relative quality of the match. Some inventors with a moderately bad pure match with their current employer will leave if the patent is granted and remain otherwise. Hence, if patent grants serve as signals of ability, they will have a positive impact on mobility. Analogously, if the pure match component of productivity with the current employer ϵ_{ij} cannot be observed, patent grants may be used as signals of the quality of that match. In that case, patents will generate a positive update in the expected value of θ_i and reduce mobility. Consequently, a signaling effect of patent grants could drive the estimated effect of patenting on mobility up or down, depending on whether they are used as signals of ability (δ_i) or signals of the match with the current employer (ϵ_{ij}), respectively. These signaling mechanisms cannot be circumvented with the use of exogenous variation in patent grants unless we make rather extreme assumptions on the observability of such variation.[5] Therefore, we acknowledge

[5]In particular, when using examiner leniency as an instrument for patent grants, we would need to assume that the market observes (and discounts) leniency to rule out the signaling argument.

the signaling argument as an alternative mechanism behind our estimated effects of patenting on mobility, and in Section 6, we discuss the extent to which it is compatible with the rest of evidence presented in the article.

3. Data and methods

3.1. Description of the data

Our study combines data from several sources. We begin with the USPTO Patent Examination Research (PatEx) dataset, which sources its information from the public Patent Application Information Retrieval (PAIR) database. PAIR contains detailed information on published patent applications filed with the USPTO. This includes standard data, such as the filing date, technological classification and application type, as well as data about the examination process, such as its current application status and the identity and art unit of the examiner (see Graham et al., 2018, for details). We identify from this dataset every original utility patent application filed between 2001 and 2012. We are constrained to this time period due to the availability of data on applications. Applications only began to be made public regardless of whether the patent was granted after the implementation of the American Inventors Protection Act (AIPA) in November 2000.[6]

To identify the firm responsible for filing the patent application, we use the information on ownership reassignments from inventors to their (presumed) employers available in the USPTO Patent Assignment Dataset.[7] Our sample effectively runs until September 2012, when the process of reassignments was made unnecessary. From our original set of approximately 3.6 million patent applications, we identify approximately 2.8 million that were reassigned from the inventors to their employers.

Next, we use data from the PatentsView initiative (www.patentsview.org) to identify the inventors listed in our sample of applications and compile their career histories. This dataset contains the results of the disambiguation algorithm specific to the inventor data provided in Li et al. (2014), which allows the robust identification of individual inventors across patent

[6]PAIR covers approximately 95% of the regular utility filings from 2001 to 2012 (Graham et al., 2018). The remaining 5% of applications correspond to: (i) applications that were abandoned before the 18-month publication lag and (ii) applications that opted out of pre-grant publication (thus relinquishing the possibility of international protection) and for which patents were not eventually granted. According to Graham et al. (2018), the applications covered by PatEx are very similar to the population of USPTO applications.

[7]Before September 2012, the USPTO considered the inventor to be the owner of the application. Since inventors typically have contractual obligations to transfer ownership to their employer, it was necessary to submit to the USPTO a chain of title from the inventor to the assignee (i.e., the firm) (Marco et al., 2015).

applications (since 2001) and granted patents (since 1976).[8] Through these data, we can identify 2.1 million disambiguated individual inventors.

We impose the following three restrictions on the merged dataset. First, we focus our analysis on inventors who filed their first patent application between 2001 and 2012. This sample represents a subset of inventors who are in their early careers, that is, in their first ten years (at most) of inventive activity. We select these inventors for the following reasons. Primarily, the initial steps of an inventor's career are more likely to be affected by the outcome of one or a few patent applications. Thus, if patents have an effect on mobility, it will be more clearly detected among inventors at the beginning of their careers. Moreover, regression to the mean in random processes eliminates the variation in the average patent examiner leniency among inventors with a large number of applications. Thus, our identification strategy would be less effective for the subpopulation of more experienced inventors. Second, we focus on inventors who receive at least one decision on an application during our sample period, given our aim of identifying the impact of the patent office's decisions. In particular, we require that they receive a first decision on their application prior to 2012 to ensure that we have at least a nine-month window within which to observe mobility for the last cohort in the data. Third, given our interest in employee inventors, we further restrict our sample to inventors who started their careers at a company. To capture them, we select applications assigned to originating firms included in the Standard and Poor's (S&P) Capital IQ database, which provides the names and transactions (such as mergers and acquisitions) for the most extensive set of public and privately held U.S. firms. To match the firm names from Capital IQ with the assignee names, we first apply the name standardization procedure used in the NBER patent data project.[9] We then run the Jaro-Winkler algorithm to correct for typos and misspellings, grouping together records with an overlap of 90% or higher. Finally, we apply exact string matching to select the final list of standardized assignee names that coincide exactly with firm names in Capital IQ.

To detect inventor mobility, we infer employer changes from changes in the assignee between two consecutive applications, following Hoisl (2007). An analogous method based on granted patents is common in the literature (see, e.g., Marx et al., 2009; Palomeras and Melero, 2010; Singh and Agrawal, 2011; Ganco et al., 2015). Because our approach is based on applications and

[8] We are grateful to the PatentsView team for sharing these data with us. PatentsView is supported by the Office of the Chief Economist at the USPTO, with additional support from the U.S. Department of Agriculture (USDA). The PatentsView platform was established in 2012 and is a collaboration among the USPTO, USDA, Center for the Science of Science and Innovation Policy at the American Institutes for Research, the University of California, Berkeley, Twin Arch Technologies, and Periscopic.

[9] See https://sites.google.com/site/patentdataproject.

not granted patents, it suffers less from the concern that the sample may be biased towards high-skilled inventors (i.e., those who have met the USPTO patentability requirements). However, it shares with them a number of acknowledged limitations (Palomeras and Melero, 2010; Ge et al., 2016). We discuss below those with implications for our current exercise.

First, an inventor's career can only be tracked if she repeatedly appears in patent applications. Otherwise, she is censored out of the sample. One potential concern of this approach is that inventors whose applications are not granted may be less likely to apply again in the future. This would be problematic if the attrition were not randomly distributed across movers and stayers. For example, a moving inventor's likelihood of changing research areas and filing new patent applications may depend (positively or negatively) on whether her previous innovations are patent-protected. Furthermore, applying for patents is largely a firm-level decision, and there may be substantial heterogeneity among firms in the intensity of patent use. All the inventors in our population have, by definition, been included in an application by their initial employers. Thus, it is natural to expect that they tend to be initially employed by firms with relatively high patent intensity and that movers will tend to switch to less patent-intensive employers. Therefore, measuring inventor mobility with patent applications poses an attrition problem that might affect the validity of our results. We address this issue in Section 4.3.

Second, identifying inventors through the names that appear in patent documents may lead to misclassification errors, which we mitigate using the name disambiguation algorithm provided by Li et al. (2014). As remarked by Ge et al. (2016), other potentially important sources of misclassification in tracking mobility with patent data are the inability to detect the exact point in time at which a move takes place (not relevant for our study) and the recording as moves of contract R&D, collaborations, or mergers and acquisitions. To address this last problem, we do not consider as actual moves the following changes in assignee: (i) those that are followed by a return to the initial employer within less than one year, as these cases most likely reflect contract research or collaborations; (ii) those due to mergers and acquisitions, which are detected through information provided by Capital IQ. The existence of some remaining misclassification error is unavoidable due to the nature of the large-scale representative sample used in our study.

Our final sample consists of inventors who started their careers at one of our identified Capital IQ firms during the period of 2001–2011 and who received a first decision on a patent application prior to 2012. We track inventors from their first application until they move or until their last application with the initial firm. The resulting dataset comprises 67,775 first-time inventors employed by 2,231 originating firms and who filed 320,874 applications during the sample period. We detect 13,984 first-employer changes for those inventors, suggesting that approximately 20% of them moved at least once over the 11-year sample period. This number is not far from the 37% over a 25-year period detected by Hoisl (2007).

3.2. Econometric specification

To identify how patent grants affect subsequent inventor mobility, we estimate the probability that an inventor changes her employer between application year t and application year $t+1$, conditional on not having moved at t. We use a linear probability model to estimate the hazard that an inventor moves as follows:

$$Probability_{ijt}\ (Y_{ij,t+1} = 1) = \alpha + \beta\ Patents\ granted_{it} + \gamma\ Z_{it} + \kappa\ D_i + \eta_t + \tau_j + \varepsilon_{ijt}, \qquad (2)$$

where i indexes inventors, j indexes their employers, and t is an ordinal index of the application year. The dependent variable, $Y_{ij,t+1}$, is an indicator that equals one if an inventor moves between t and $t+1$. When inventors file several applications in a given year, we consolidate the information on an annual basis.[10] Hence, our measure of mobility records whether the inventor changed employers at least once during that observation window.[11]

Our main variable of interest, $Patents\ granted_{it}$, is the total number of patents issued to inventor i up to t (included). The vector Z_{it} contains a range of time-variant covariates. First, and most important, Z_{it} contains the total number of applications filed by inventor i up to t, as we want to compare inventors that filed the same number of applications. Second, we condition on the time elapsed since the inventor's first decision year (in spells of 365 days), allowing mobility decisions to be shaped by seniority. Third, to account for technological differences in mobility rates, we include indicators for the six nonexclusive NBER categories in which the applications are classified. D_i represents the year in which the inventor receives her first decision from the patent office. This cohort indicator controls for the fact that inventors entering later in the panel have less time to move than do those entering earlier. Finally, η_t are year fixed

[10] Note that this implies considering inventors with at least 2 applications in 2 different years.

[11] Because t indexes application years and not calendar years, the time window between two consecutive values of t may be longer than one year if the inventor takes longer than that to apply for a patent again. Thus, an inventor with a first patent application in 2001, a first decision in 2002, and subsequent applications in 2004, 2005, 2008 and 2010 contributes to our sample with four time windows of unequal length.

effects that account for period differences, and τ_j are firm fixed effects, as we want to eliminate firm-level unobserved heterogeneity. We cluster standard errors at the inventor level, although clustering at the firm level produces almost identical results.

One important challenge in the estimation of Equation (2) is that β may be upward biased if it captures the combined impact of patent grants and omitted inventor characteristics. As explained in Section 2.3, inventors of higher quality may be both more likely to obtain granted patents and more likely to switch employers. Failing to account for this would lead us to mistakenly attribute a causal impact of patents on mobility due to the confounding effect of inventor ability. A second source of bias stems from the fact that patent applications cannot be definitively rejected, only "abandoned" by the applicants. After a rejection from the patent office, applicants can always "resubmit" an updated patent application, accommodating the USPTO examiners' criticism, typically by narrowing down the scope of the claims (Sampat and Williams, 2019). Therefore, although the feedback from the patent office is potentially important in driving abandonment, unobserved innovation characteristics may also affect the applicant's prosecution efforts and, consequently, the probability of an eventual grant decision. To overcome these identification challenges, one needs some exogenous source of variation in the likelihood of patent grants that does not belong directly in the mobility equation. We next discuss the extent to which patent examiner leniency meets these requirements.

3.3. Identification strategy

To construct an instrument that allows us to identify the effect of patent grants on mobility, we build on the previous research on the patent examination procedures at the USPTO. Although the process is relatively structured and standardized, patent examiners have considerable discretion in how they conduct the examination. This influences the outcome of otherwise similar patent applications. Lemley and Sampat (2012) find an 11-percentage-point difference in the grant rate between the least and the most experienced examiners in a given technology area. Frakes and Wasserman (2016) report differences in the odds of patent approval by examiner cohort, which they attribute to differences in the training received at hiring. Accordingly, we exploit differences in leniency across patent examiners as a source of variation in patent grants, following a similar approach to recent studies that investigate the impact of patent grants on firm and technology-level outcomes (Gaulé, 2018; Sampat and Williams, 2019). Our main identification assumption is that examiner leniency is only related to the mobility of an inventor through her expected number of patents granted. While our exogeneity checks in Section 4.2 and Online Appendix A.2 suggest that this is the case, a description of the process for allocating applications to examiners may help to understand the rationale behind this assumption.

Upon arriving at the USPTO, patent applications are received by a central office, where they are assigned an application number, patent class and subclass codes and are then allocated accordingly to one of the art units in charge of the examination process. Once in the art unit, a supervisory patent examiner (SPE) assigns the application to a specific examiner. Each art unit has discretion over how work is organized, including how applications are allocated to examiners by SPEs. Interviews with patent examiners conducted by Lemley and Sampat (2012) reveal that this assignment process is done according to either arbitrary rules (such as the last digit of the application number), docket management needs or/and familiarity with the technology. There is no evidence from these interviews that SPEs engage in any substantive evaluation of applications that could inform about their patent-worthiness. Therefore, it is unlikely that they assign applications to examiners according to this dimension. The empirical evidence is consistent with these observations. The work by Lemley and Sampat (2012), Righi and Simcoe (2019) and Sampat and Williams (2019) shows that patent applications assigned to lenient and strict patent examiners within an art unit have similar observable characteristics at the time of application.

3.3.1. Average examiner leniency as an instrumental variable

Our objective is to obtain an instrumental variable for the number of applications granted to an inventor up to a given point in time. We start by operationalizing examiner leniency at the application level. In the spirit of Gaulé (2018), we compute time-varying measures of leniency as follows:

$$E_{jkat} = \frac{Grants_{kat} - 1(Grant_j = 1)}{Reviews_{kat} - 1} \tag{3}$$

and

$$U_{jat} = \frac{Grants_{at} - 1(Grant_j = 1)}{Reviews_{at} - 1}, \tag{4}$$

where E_{jkat} is the approval rate of examiner k in art unit a assigned to review application j submitted at time t. $Reviews_{kat}$ and $Grants_{kat}$ represent the number of applications in art unit a and application year t that examiner k has reviewed and granted, respectively.[12] Similarly,

[12] To avoid overstating variations in leniency when the number of applications per examiner is small, we consider only examiners who reviewed at least ten applications within an art-unit-year (Sampat and Williams, 2019).

U_{jat} is the approval rate of art unit a and is constructed as the share of reviewed applications filed in year t that were granted by art unit a, excluding the focal patent.[13] The difference between E_{jkat} and U_{jat} represents the relative leniency faced by an inventor who files patent application j in year t assigned to examiner k within art unit a. For a single patent application, the corresponding examiner's relative leniency, $E_{jkat} - U_{jat}$, is a suitable instrument for whether that application is granted. However, we are interested in obtaining an instrument for the inventor's total number of applications granted up to a given point in time. Accordingly, we average $E_{jkat} - U_{jat}$ across the n_{it} patents applied for by inventor i up to year t, and obtain the inventor-level average examiner leniency at t as follows:

$$L_{it} = \frac{1}{n_{it}} \sum_{j=1}^{n_{it}} (E_{jkat} - U_{jat}). \tag{5}$$

4. Results: Patent protection and inventor mobility

4.1. Descriptive statistics

Table 1 provides the descriptive statistics for the main variables used in this study. The unit of observation in our analysis is the inventor-application year. The figures indicate that 85% of our sample of early-career inventors have at least one granted patent, with an average of 2.2 granted patents and 7.7 applications per inventor. On average, 10% of the inventors change employers between two application years. Because the average duration of such a window is of 1.7 years, this amounts to an annual mobility rate of approximately 6%.

[13] Note that our instrument differs from that proposed by Gaulé (2018) in two respects. First, while he considers the overall approval rate of an examiner, we follow Sampat and Williams (2019) in adjusting for art unit. We do this because 39% of examiners in our sample reviewed patent applications for multiple art units in the same year. Second, our equations differ in the denominator, as we use the number of applications reviewed rather than the number of applications filed. However, nothing hinges on the use of the leniency measure of Gaulé (2018).

Table 1: Descriptive statistics

Variable	Mean	SD	p10	p50	p90	Observations
Move	0.10		0	0	1	129,740
Any patent granted	0.85		0	1	1	129,740
# of patents granted	2.2	3.6	0	1	5	129,740
# of basic patents granted	0.45	1.2	0	0	1	74,430
# of applied patents granted	2.3	3.4	0	1	5	94,737
Average patent scope	0.95	0.85	0	0.83	1.7	116,766
Examiner leniency	0.005	0.088	-0.095	0.008	0.103	129,740
Examiner scope leniency	0.0002	0.0014	-0.0014	0.0001	0.0020	116,766
# of applications filed	7.7	10.6	1	4	17	129,740
Years since 1st decision	3.1	2.0	1	3	6	129,740
# of coinventors	11.1	12.3	2	7	25	129,740
# of USPC classes	2.9	2.3	1	2	6	129,740
% of applications in firm's core	0.30	0.39	0	0	1	129,740
# of applications per firm	1,269	1,779	15	463	4,320	129,740
Sales ($m)	47,340	45,748	1,082	34,589	106,916	68,213
# of expired patents	0.64	1.7	0	0	2	33,813
# of post self cites	3.8	67.7	0	0	5	33,813
# of post cites	6.7	21.7	0	2	16	33,813
Enforceability index	-0.68	1.6	-3.8	-0.07	0.90	33,813
Female	0.09		0	0	0	33,813
# of pre self cites	9.2	186.8	0	1	11	33,813
# of pre cites	11.4	38.5	0	3	27	33,813

4.2. Patent grants and inventor mobility

Table 2 provides the main results of our analysis. In Column 1, we present the OLS estimates of the baseline specification relating mobility to the number of patents granted and our main set of controls. We find a small, positive and significant relationship between patent grants and mobility. This relationship cannot be interpreted as causal, however. As argued above, there are reasons for which we should expect unobservable factors to affect both the extent to which an inventor's patent applications are approved and her subsequent mobility.

Turning to the instrumental variable approach, Column 2 presents the first stage, in which we regress the number of patents granted on average examiner leniency. As expected, the instrument is positive and highly significant; a one-standard-deviation increase in the average leniency of the examiners assigned to review an inventor's patent applications is associated with a six-percentage-point increase in the number of patents issued for the average inventor-year. The first-stage F-statistic of the excluded instrument is large and well above the rule of thumb for weak instruments proposed by Stock and Yogo (2005), indicating that the instrument explains a substantial part of the variation in granted patents. Column 3 reports the result from the second-stage regressions estimating Equation (2), with the main variable of interest replaced with the predicted number of patents granted from the first stage. The coefficient is strongly negative and significant at the 1% level. The point estimate implies that an exogenous increase in one successful patent application reduces the probability of moving by 2.3 percentage points, which represents a 23% decrease over the sample probability of 10%. This result is economically significant. For context, Marx et al. (2009) find that Michigan's adoption of enforceable noncompete agreements in 1985 reduced inventor mobility by approximately 8%, and Balasubramanian et al. (2018) document that the 2015 noncompete ban for technology workers in Hawaii increased mobility by 11%. Since Starr et al. (2019) report a noncompete signing rate for high-skilled workers (e.g., computer, mathematical, engineering) of 36%, the relative effect of noncompetes on mobility can be approximately estimated to be between 22% and 30%. These numbers are within the same range as our results for the effect of patent protection.

For the ease of estimation and interpretation, we use linear probability models as our main specification throughout the paper. In Column 4 of Table 2, however, we report the results from a probit model where we implement the instrumental variable estimator by using the control function method (see Blundell and Powell, 2004). This leads to qualitatively and quantitatively similar results for the estimated effect of the number patents granted, supporting the notion that the linear 2SLS model is a reasonable approximation to obtain average partial effects.

Table 2: Patent grants and inventor mobility

Estimation method	OLS	OLS (1^{st} st.)	2SLS (2^{nd} st.)	Probit (2^{nd} st.)
Dependent variable	Move	# of pats granted	Move	Move
	(1)	(2)	(3)	(4)
Examiner leniency		1.577***		
		(0.062)		
# of patents granted	0.001***			
	(0.000)			
# of patents granted *(instr.)*			-0.023***	-0.020***
			(0.007)	(0.005)
Years since 1^{st} decision (L)	0.029***	0.131***	0.033***	0.031***
	(0.003)	(0.029)	(0.003)	(0.003)
Firm FE	Yes	Yes	Yes	Yes
First decision year FE	Yes	Yes	Yes	Yes
Calendar year FE	Yes	Yes	Yes	Yes
# of applications filed FE	Yes	Yes	Yes	Yes
Technological class FE	Yes	Yes	Yes	Yes
F-statistic		430		
Exogeneity test (p-value)			0.000	

$N = 129,740$. Number of inventors: 67,775. Number of firms: 2,231. The estimation period is 2001–2011. Robust standard errors are clustered by inventor (in parentheses). The exogeneity test is a Hausman-based test. The test for weak instruments is the Cragg-Donald Wald F-statistic. Column 4 displays the average marginal effect from a probit model with endogenous regressors.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Overall, these instrumental variable specifications provide evidence suggesting that patent grants cause, on average, a decrease in the subsequent mobility of early-career inventors. Since the consistency of these results depends on the validity of average examiner leniency as an instrument for patent grants, we provide a series of supporting exogeneity tests. We first test whether the application portfolios of the inventors assigned to tough and lenient examiners are similar in terms of the observable characteristics that predict granting probability at the time of filing. Following Sampat and Williams (2019), we use the variables *patent family size* and *number of claims* to predict the probability of patent approval. Based on these predictions, we compute inventor-level predicted grant rates. Figure 1 shows that average examiner leniency is visually orthogonal to the predicted grant rate, although it is clearly related with the actual grant rate (see details of this test in Online Appendix A.2.1).

Figure 1: Probability of patent grants by average examiner leniency

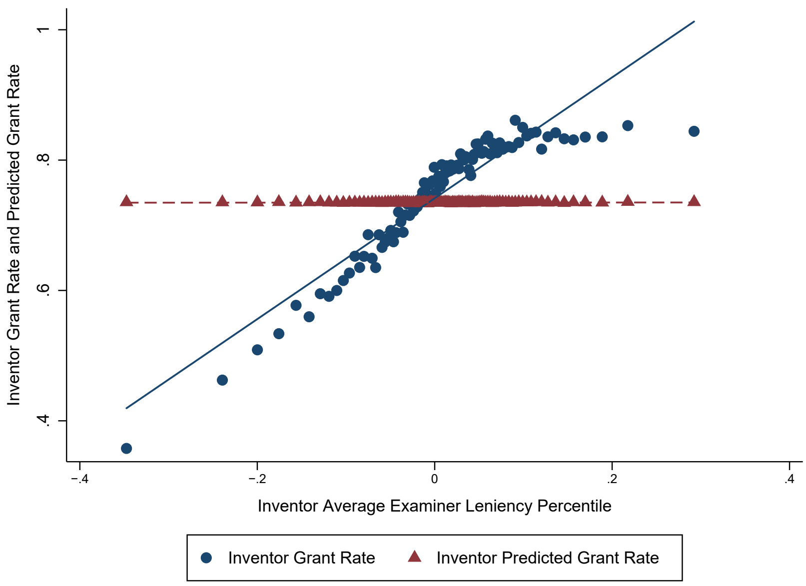

The figure relates the average examiner leniency with the inventor-level grant rate and predicted grant-rate. The predicted grant rate is based on two observable proxies for patent value at the time of filing (patent family size and claims count). The graph was generated using Michael Stepner's binscatter for Stata.

We also perform the following set of additional tests, which are fully discussed in Online Appendix A.2. First, we provide evidence suggesting that average examiner leniency is also orthogonal to inventor characteristics that impact mobility. Second, we document that our approach of taking deviations in examiner grant rates from art-unit-year averages [see Equations (3) and (4)] is equivalent to running the analysis with raw grant rates and controlling for art-unit-year fixed effects (as in Sampat and Williams, 2019). Third, we investigate whether the existence of examiner specialization within art units (as reported by Righi and Simcoe, 2019) affects our results. We do so by replicating our analysis when computing examiner leniency at the art-unit-subclass level. Fourth, we also replicate the analysis for the subsample of art units where, according to Feng and Jaravel (forthcoming), the SPEs assign patent applications to examiners based on the last digit of the application's serial number. These are art units for which the "next-to-random" application allocation assumption is particularly plausible. Our results remain robust to such considerations. Other potential concerns generated by the nature of our empirical analysis (not strictly related to the use of our instrumental variable) are addressed in the next section.

4.3. Robustness checks

In this section, we provide a variety of additional tests that confirm the robustness of our baseline findings. We also address the concern that our estimates could be affected by some attrition effect generated by our approach of detecting mobility with patent applications. We summarize the results of both kinds of checks in Table 3.[14]

First, our baseline model incorporates fixed effects for six broad technology fields (NBER categories). To account for technological effects at a more disaggregated level, in Column 1, we estimate our main specification with technology effects that include 414 patent-class fixed effects. The point estimate from this more stringent specification is similar to and not statistically distinguishable from the baseline point estimate presented in Table 2, suggesting that the impact of patents on mobility is not driven by characteristics of technology sub-fields not captured by our six broad technology category dummies.[15]

Second, we address the concern that our estimated effect may be driven by the resolution of the uncertainty surrounding an application. In Equation (2), the control for the total number of applications filed by an inventor includes granted, abandoned and pending patent applications. The estimated effect of patent grants may therefore be related to the release of the decision regarding the application rather than to the granting of the patent itself. For this reason, in

[14]See Online Appendix, Table A.3 for a summary of the relative effect sizes of our estimates of patenting across the different specifications and subsamples reported in the paper.

[15]Patent classes are defined according to the United States Patent Classification (USPC). We choose to report the rest of the analyses in the article while accounting for NBER category fixed effects to remain consistent with existing studies on inventor mobility and with previous research using examiner leniency as an instrumental variable. Nonetheless, we replicate all of them using USPC class fixed effects. The corresponding results are qualitatively and quantitatively similar to those presented here.

Column 2, we drop pending applications and re-estimate the model with dummies controlling for the number of applications reviewed. The coefficient of interest is slightly weaker in this case but remains statistically significant.

Third, we collapse our panel data into a cross-sectional structure by considering inventor mobility in any moment after the first decision. This strategy reduces the concerns of potential oversampling of frequent applicants. It may also offer a more precise identification of our parameter of interest, as we can more neatly exploit variations caused by examiner leniency. The results, reported in Column 3, show that the estimate from this procedure is qualitatively similar to the baseline estimate. The point estimate of -0.027 implies that obtaining a positive first answer decreases the probability of moving in the next years by approximately 18% on average. A comparison of this effect size with our baseline of 23% allows us to evaluate whether the impact of patenting on mobility operates at the intensive and/or extensive margin. The results indicate that the effect is present at both margins, as the effect of the first patent grant on inventor mobility is similar to the average effect of having one additional patent.

Fourth, to rule out the possibility that the effect we detect is driven by outlier inventors, we exclude superstar inventors (those at the top 5% of the distribution according to the number of applications submitted to the USPTO) from our sample. In Column 4, we again observe qualitatively and quantitatively similar results to those of our baseline estimation.

Finally, we run three robustness checks to examine the extent to which the sample attrition inherent to measuring inventor mobility with subsequent patent applications is an issue in our study. As suggested in Section 3.1, this would be the case if moving inventors had different probabilities of filing new applications and those differences were contingent on previous patent grants. We first address this issue by considering one potentially important underlying reason behind the differential attrition of movers and stayers, i.e., changes in technological areas. It is reasonable to expect that moving inventors are more likely to switch areas than stayers. A change in the area of research, in turn, may make the moving inventors more or less likely to apply for patents again (and be included in our sample) in the near future. In Column 5, we deal with this possibility by restricting the sample to pharmaceutical inventors (based on U.S. patent classes 424 and 514). Pharmaceutical scientists make a good testing field for this concern because they are very unlikely to switch areas (Paruchuri et al., 2006). Scientists researching cures for cancer, for instance, typically spend a lifetime in this area. Moreover, the research output in this field usually leads to patent filings, thus further reducing attrition concerns. The estimated coefficient for inventors in this group is -0.030, which implies an average reduction of 25% in inventor mobility for an additional patent (average mobility rate of pharmaceutical inventors between two patenting years is approximately 12%). This is only slightly larger than the average effect in the overall sample.

More generally, we investigate the extent to which sample attrition due to lack of reapplications is likely to drive our results. In Column 6, we restrict the sample to inventors who filed patents before 2007, so that we have a complete 5-year time window to observe another application. This approach implies dropping more than half of the observations of the sample, but it substantially increases the reapplication rate.[16] The coefficient of interest for this subsample is negative and significant. Although the -0.033 point estimate is larger than that of our baseline in absolute terms, it is very similar in relative terms (i.e., 21% vs. 23%). As a complementary approach, we examine how our main results are altered by the variation in patenting intensity across industries. Intuitively, if inventor sample attrition due to lack of reapplications generates some bias, it will be particularly present in industries with less intensive patenting activity. To test this, we first calculate patenting intensities using all firms in the Compustat/CRSP database by finding the (time-varying) average number of patents per R&D dollar at the two-digit SIC code level.[17] Next, we merge this data with the subsample of firms in our dataset for which we have industry classification data (i.e., those listed in Compustat) and then split the sample into observations with high and low patenting intensities based on the sample median. In Columns 7 and 8 we find, in two split sample regressions, that the point estimate is -0.023 in the low patenting-intensity group and -0.034 in the high patenting-intensity group. Evaluated at their respective sample means, however, these point estimates imply economically identical effects. Having one additional patent leads to an average 26% reduction in mobility in both types of sectors. In unreported tests, we further confirm these results by using the full sample and introducing the (instrumented) interaction between patents and a patenting-intensity group dummy. This generates a nonsignificant coefficient of -0.003 on the interaction term.

[16]While 54% of first-time inventors reappear at least once in our data, this percentage increases to 73% when we impose this restriction.

[17]Patent data for publicly traded firms in the Compustat/CRSP-merged universe come from Kogan et al. (2017).

Table 3: Robustness tests

Estimation method: 2SLS Sample Dependent variable: Move	All (1)	Reviewed applications (2)	First decision (3)	Non-superstars (4)	Pharma (5)	Filing year $\leqslant 2006$ (6)	Low patent intensity (7)	High patent intensity (8)
Patent granted *(instr.)*			-0.027**					
			(0.012)					
# of patents granted *(instr.)*	-0.027***	-0.017***		-0.026***	-0.030*	-0.033**	-0.023*	-0.034*
	(0.008)	(0.005)		(0.008)	(0.017)	(0.015)	(0.014)	(0.019)
Redefined (USPC class) Tech class FE	Yes							
First-stage F-statistic	400	2,559	4,912	828	177	352	132	81
Exogeneity test (p-value)	0.000	0.005	0.378	0.000	0.036	0.007	0.035	0.022
N	129,721	127,248	66,285	123,596	23,069	58,351	34,567	33,563
# of inventors	67,764	66,660	66,285	67,240	12,921	42,897	18,600	18,342
# of firms	2,231	2,209	2,052	2,228	1,241	1,682	380	551

Robust standard errors are clustered by inventor (in parentheses). The exogeneity test is a Hausman-based test. The test for weak instruments is the Cragg-Donald Wald F-statistic. All regressions control for the number of years since the inventor's first decision (log), firm fixed effects, fixed effects for the number of applications filed by the inventor, the technology field, the year of the inventor's first decision and the calendar year. Nonsuperstar inventors are defined as belonging to the bottom 95% of the patent application distribution. Pharmaceutical scientists are identified using USPC classes 424 and 514, which are both defined as drug, bio-affecting, and body-treating compositions. We calculate the industry patenting intensity for Compustat firms by finding the average number of patents per \$m of R&D at the two-digit SIC code level in the last five years before the focal year. High (low) patenting-intensive industries are those where this ration is above (below) the sample median of 0.08. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

4.4. Extension: The scope of patent protection

Because of our reliance on examiner leniency as an instrumental variable for patent grants, it is likely that our baseline results represent the local average treatment effect of an additional granted patent for the group of inventors producing innovations around the margin of approval and rejection. In this section, we extend the notion of patent protection beyond the granted/nongranted dichotomy and focus on patent scope as a more fine-grained measure of the concept (Lerner, 1994; Lanjouw and Schankerman, 2004). This analysis allows us to evaluate the effect of a marginal increase in the scope of protection, which may occur at any point in the distribution of patentability.

To operationalize the scope of patent protection, we use a metric based on patent claims. The claims of a patent document describe, in technical terms, the different elements of the underlying technology, delimiting the technological space covered by the patent. Kuhn and Thompson (2019) show that the length of the first claim in a patent document is the best available measure of the patent scope. The rationale behind this measure is that shorter claims result in a broader scope because each added word introduces more conditions that a competitor must meet to be considered an infringer. Thus, the applicant has incentives to write the claims as short as possible in the application to aim for the broadest possible protection. On the other hand, examiners typically amend the claims during examination adding words (Kuhn and Thompson, 2019; Marco et al., 2019). The scope of patent protection is, therefore, affected by the examination process and influenced by individual examiners, as further supported by anecdotal evidence (Cockburn et al., 2003; Lemley and Sampat, 2012). Consequently, how strict or lenient examiners are when evaluating claims can also be used as an instrumental variable to estimate the effect of the patent scope on mobility.

To construct our measure of patent scope, we start with the number of words in the first claim of the patents issued to an inventor. To be able to use the full sample of patent applications and given that nongranted applications have no approved claim, we compute the inverse of the first-claim length and assign a value of zero to nongranted applications. Finally, we average the indicator at the inventor level.[18] Higher values of this indicator (*average patent scope*) denote that the inventor has a patent portfolio of broader scope. Second, to construct the instrument for patent scope, we compute the average examiner scope leniency by following the process

[18]Kuhn and Thompson (2019) note that this measure of scope is not suitable for the analysis in biotechnology. Therefore, we exclude filings from this technology center. Further, to make claim length comparable across different technology fields, we scale it by the respective mean at the art-unit-year level.

described in Equations (3) to (5) but using the inverse of the number of words approved in the first claim instead of patent grant rates.

Table 4 presents the results. Column 1 reports the simple OLS estimates, which are small, but negative and significant. Column 2 shows the first stage where we regress the inventor's average patent scope on the average examiner scope leniency. The estimated effect of the instrument is positive, as expected, and statistically significant at the 1% level. Column 3 reports the second stage. In line with our findings for patent grants, we estimate a negative and significant effect of the scope of patent protection on mobility. The point estimate implies that a one standard deviation increase in the average of the inverse of the claim length leads to a 29% reduction in the probability of leaving.

In Columns 4 to 6, we examine the extent to which these results are driven by the fact that we include nongranted patents, which have zero scope. To do so, we repeat the analysis of the effect of patent scope considering granted patents only. This poses a selection problem that we address with a Heckman correction process. We use the average examiner leniency in patent grants to address the sample selection problem and the examiner leniency in patent scope for the subsequent 2SLS instrumental variable estimation. The estimated effect of patent scope on mobility, available in Column 6, is almost identical to that of Column 3.

Two important conclusions emerge from the scope analysis. First, the results suggest that the negative estimated effect of patenting on mobility is not driven by the subset of inventors whose innovations lie around the approval threshold, but that it is a more general phenomenon. Second, these findings reinforce the idea that the impact of patent protection on inventor mobility operates not only on the extensive margin (through obtaining a patent for an innovation) but also on the intensive margin (in this case, through a broader protection for the granted patents).

Table 4: Patent scope and inventor mobility

Sample	All	All	All	Patents	Patents	Patents
Estimation method	OLS	OLS (1^{st} st.)	2SLS (2^{nd} st.)	OLS	OLS (1^{st} st.)	2SLS (2^{nd} st.)
Dependent variable	Move	Average patent scope	Move	Move	Average patent scope	Move
	(1)	(2)	(3)	(4)	(5)	(6)
Average patent scope	-0.007***			0.004**		
	(0.001)			(0.002)		
Examiner scope leniency		70.322***			104.178***	
		(3.377)			(6.761)	
Average patent scope *(instr.)*			-0.034***			-0.035***
			(0.009)			(0.009)
Inverse mills ratio				-0.142***	-0.352***	-0.194***
				(0.035)	(0.121)	(0.037)
First-stage F-statistic			1,681			3,587
Exogeneity test (p-value)			0.002			0.001
N	116,766	116,766	116,766	99,264	99,264	99,264
# of inventors	60,928	60,928	60,928	50,860	50,860	50,860
# of firms	1,873	1,873	1,873	1,671	1,671	1,671

The estimation period is 2001–2011. Robust standard errors are clustered by inventor (in parentheses). The exogeneity test is a Hausman-based test. The test for weak instruments is the Cragg-Donald Wald F-statistic. All regressions control for the number of years since the inventor's first decision (log), firm fixed effects, fixed effects for the technology field, the year of the inventor's first decision and the calendar year. The words present in the first independent claim are standardized at the level of the art unit and filing year. Because of how language is used in their claims, we exclude patents examined by the "Biotechnology and Organic Chemistry" center (see Kuhn and Thompson, 2019). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5. Heterogeneous impact of patent grants

Thus far, we have presented robust evidence consistent with our leading hypothesis that patent protection decreases subsequent inventor mobility. In this section, we test the additional implications generated by the idea that this negative effect is, at least in part, driven by innovation-related skills that become patent holder specific.

First of all, we have argued that the effect of patent protection on mobility will be less intense when complementarities with other resources of the company make the inventor's innovation-related skills inherently firm specific (thus weakening the human capital effect and strengthening the product market effect of patents). In empirical terms, we capture the inventor's complementarities with her employer's resources through the following: (i) the number of coinventors collaborating with her at the company and (ii) an indicator of whether she works in the technological core of the company. The extent to which an inventor works in teams is a natural measure of the complementarities between her and her employer. As documented by Singh and Agrawal (2011), the tacit knowledge embodied by an individual inventor tends to remain contained among close collaborators (i.e., coinventors) rather than diffusing widely inside the firm. At one extreme, if the inventor is a sole author, she will be valuable in isolation for the efficient implementation of the innovation, as she will likely be the only person who embodies the necessary tacit know-how. Conversely, when the innovation is the product of teamwork, the implementation knowledge contributed by an individual will depend on the extent to which she collaborates with other the team members. Consequently, the negative effect of patents on mobility should be less intense for inventors with multiple coauthors.

Analogously, the extent to which the inventor works in a technological area in which her employer is particularly research active captures the potential complementarity between the inventor and her employer's technological resources. In these areas, usually referred by the innovation literature as "core technology areas," firms are likely to have well-established trajectories of research (Song et al., 2003). This implies that routines, procedures and know-how in those areas are deeply embedded in the firm, and the skills of an individual inventor can hardly be implemented in isolation. Thus, we also expect that the negative impact of patent grants on mobility is larger for inventors working outside the firm's core areas than for those employed in the firm's core.

We investigate the effects of these two sources of heterogeneity using the following variables: (i) the natural logarithm of (one plus) the number of unique coinventors with whom an inventor has worked in her applications up to t and (ii) the percentage of her patent applications that fall in the firm's core technology areas. Following Song et al. (2003), we consider a technological area as part of the core if its corresponding patent class appears with a frequency greater than 10% in the firm's application portfolio (over the entire sample period). In this analysis, we also control for the inventor's degree of specialization (captured by the number of different patent classes into which her applications fall) and firm size (proxied by the number of applications filed by the firm in that year), which are necessary controls to consistently estimate the effects of coinventors and core areas.

Panel A of Table 5 presents the results of specifications that include the (instrumented) interactions between the number of patent grants an inventor received and her complementarities with other resources of the organization. The first column reproduces the baseline model to which we add the above-mentioned variables. In Columns 2 and 3, we separately add the (instrumented) multiplicative terms. As expected, the estimates from the interaction effects show that the negative impact of patent grants on mobility is most intense for solo inventors and for inventors working outside the firm's core technologies. As the inventor has a larger group of coauthors and her work falls more in the core areas of the firm, this negative effect is attenuated. In Column 4, both interactions are considered simultaneously. The positive estimated interaction effect of the number of patents and coinventors decreases to nonsignificant levels, but the estimated moderating effect of the core remains virtually unchanged. Overall, these results support the implication that the negative impact of patent protection on inventor mobility is stronger in the absence of other sources of specific human capital.

The second source of heterogeneous effects from our theoretical framework concerns the importance of the inventor's skills for the implementation or further development of her innovation. If, as argued, patent protection turns innovation-related skills into patent-holder-specific human capital, the impact of a patent grant on mobility should be more intense when these skills are more relevant. We capture the relative importance of the inventor's implementation skills by distinguishing between whether her innovations are closer to basic or applied research. Basic research involves early stage technologies for which the returns are difficult to identify and appropriate; therefore, the knowledge behind them is typically less codified and more difficult to transfer (Cassiman et al., 2018). Thus, ensuring the correct translation of this complex and more tacit knowledge into commercializable products is likely to require the involvement of the individuals that closely work with it (Zucker et al., 2002). Innovations based on applied research, on the other hand, are more likely to involve explicit, codified and transferable knowledge, which makes the involvement of the individual inventor less important for their implementation. Hence, we expect that the negative effect of patent protection on mobility should be more intense for innovations in the sphere of basic knowledge.

To distinguish between innovations that build on basic vs applied research, we rely on the presence of nonpatent literature (NPL) references in the list of citations included in patent documents (Fleming and Sorenson, 2004; Arora et al., 2017). We define patented innovations as basic-research oriented if the majority of NPL references consist of citations to scientific documents, and applied-research oriented otherwise.[19] We instrument our variables of interest, the number of basic (applied) patents granted, with the average examiner leniency among basic (applied) applications.[20] As the results of Panel B of Table 5 show, the estimated negative effect

[19]NPLs references include scientific publications, as well as, e.g., books, abstract meetings and abstract services, trade journals, presentations at trade fairs, patent office actions, and search reports. To identify scientific publications, we linked the list of NPL references to Thomson Reuters Master Journal List which comprises all active titles that are eligible for inclusion in the Essential Science Indicators.

[20]Note that backward citations information is available for granted patents only. Therefore, our distinction between basic and applied patent applications is based on the USPC class-subclass combinations to which they belong. We define applications as basic (applied) if they belong to a USPC class-subclass where the share of NPL references is above (below) the sample median.

of patent protection on mobility is, as expected, stronger for basic innovations than for applied ones (and the difference is statistically significant at the 10% level).

In summary, the heterogeneity analysis presented in this section supports the idea that patent protection has a negative effect on inventor mobility by converting innovation-related skills into patent-holder-specific human capital. First, the impact of patents is stronger in the absence of other sources of firm-specific human capital – i.e., when the innovation is developed with fewer coauthors and fewer firm resources are devoted to the focal technology. Moreover, the effect is particularly intense for patents based on basic research, which arguably require a more active role of their inventors for implementation.

Table 5: Patent grants, interaction with firm-specific skills, research direction and inventor mobility

	Panel A: Interaction with firm specific skills				Panel B: Research direction	
Estimation method: 2SLS Dependent variable: Move	(1)	(2)	(3)	(4)	(5)	(6)
# of patents granted *(instr.)*	-0.023***	-0.061***	-0.045***	-0.062***		
	(0.007)	(0.023)	(0.011)	(0.023)		
# of patents granted × # of coinventors (L) *(instr.)*		0.018**		0.008		
		(0.008)		(0.009)		
# of patents granted × % of applications in firm's core *(instr.)*			0.066***	0.062***		
			(0.020)	(0.021)		
# of coinventors (L)	-0.002	-0.041**	-0.000	-0.018		
	(0.002)	(0.018)	(0.002)	(0.020)		
% of applications in firm's core	-0.031***	-0.033***	-0.154***	-0.149***		
	(0.003)	(0.003)	(0.037)	(0.038)		
# of USPC classes (L)	-0.019***	-0.014***	0.002	0.004		
	(0.004)	(0.003)	(0.007)	(0.006)		
# of applications per firm (L)	-0.050***	-0.046***	-0.051***	-0.049***		
	(0.003)	(0.003)	(0.003)	(0.003)		
# of basic patents granted *(instr.)*					-0.100**	
					(0.045)	
# of applied patents granted *(instr.)*						-0.025**
						(0.010)
First-stage F-statistic	407	98	97	32	47	199
Exogeneity test (p-value)	0.000	0.000	0.000	0.000	0.007	0.003
N	129,740	129,740	129,740	129,740	74,430	94,737
# of inventors	67,775	67,775	67,775	67,775	37,398	46,644
# of firms	2,231	2,231	2,231	2,231	1,579	1,655

The estimation period is 2001–2011. Robust standard errors are clustered by inventor (in parentheses). The exogeneity test is a Hausman-based test. The test for weak instruments is the Cragg-Donald Wald F-statistic. All regressions control for the number of years since the inventor's first decision (log), firm fixed effects, fixed effects for the number of applications filed by the inventor, the technology field, the year of the inventor's first decision and the calendar year. Panel B further includes fixed effects for the number of basic or applied applications filed by the inventor. Basic (applied) patents are defined as those where the majority of backward citations are nonpatent (patent) literature citations.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6. Discussion of alternative interpretations

In principle, patent protection may reduce inventor mobility through mechanisms other than the transformation of innovation-related skills into patent-holder-specific human capital. In this section, we discuss two prominent alternative explanations: the existence of financial constraints that can be mitigated with patent rights and the potential signaling role of patents. We also provide additional evidence in support of our suggested mechanism by showing that inventor turnover negatively affects the use of the protected technology by the patent holder.

6.1. Financial constraints

Patents may serve as a value certificate and collateral for investors, facilitating access to debt and venture capital financing (Mann, 2018; Gaulé, 2018). In that case, firms that secure patent protection for their innovations would differ from firms that do not in their ability to raise funds to continue with their activity in a specific technology. Consequently, frictions in financial markets may prevent firms without granted patents from retaining their inventors. This financial constraint interpretation could explain the overall negative effect of patents on mobility that we document. We next explore its plausibility.

First, we investigate whether our results could be driven by the subsample of inventors whose employers went out of business. If patents affect firm survival by loosening financial constraints, a substantial number of inventor moves following patent rejections may be motivated by their employers' bankruptcies. We address this possibility by estimating our main specification considering only moves from source firms that have at least one patent filed in the years after the detected employer change. The results of this analysis are available in Column 1 of Table 6. They are nearly identical to the results of our baseline estimation.

Second, we repeat our main analysis separately for the subsamples of inventors working in small and big firms. If financial constraints were behind the negative effect of patent grants on mobility, the effect should be larger for those inventors working in smaller firms, which are more likely to suffer from them (Beck et al., 2005). On the other hand, if our main result is driven by the specificity of the inventor's implementation skills to the patent holder, the effect should be more important in larger organizations. This is because, in comparison with small start-ups, large firms are more likely to possess the downstream specialized complementary assets necessary for the commercialization of an innovation. This increases the likelihood that their technology is developed in-house instead of being licensed out (Arora and Ceccagnoli, 2006).

To classify companies according to their size, we define as small firms those in the lowest quartile of our sample in terms of the number of patent applications filed with the USPTO in a given year.[21] Columns 2 and 3 of Table 6 show the results when we split the sample accordingly. The estimated effect of patent grants on mobility is close to zero and insignificant in the case of small firms, whereas it is large, negative, and significant for large firms. As shown in Column 4, however, the differenceen between the two estimated effects is not statistically significant. For robustness purposes, we use the firms' net sales as an alternative measure of size for the subsample of firms included in Compustat and define small firms analogously. The analysis using this alternative size distinction, presented in the last three columns of Table 6, shows very similar results. Overall, the evidence displayed in Table 6 is inconsistent with the financial constraint interpretation of our main findings and well aligned with our preferred explanation.

Table 6: Patent grants, firm survival, firm size and inventor mobility

Estimation method: 2SLS Sample Dependent variable: Move	Active firms (1)	Small firms (2)	Big firms (3)	All (4)	CS: Small firms (5)	CS: Big firms (6)	CS: All (7)
# of patents granted *(instr.)*	-0.025***	-0.005	-0.028***	-0.030***	-0.008	-0.036***	-0.039***
	(0.007)	(0.018)	(0.008)	(0.008)	(0.022)	(0.013)	(0.014)
# of patents granted				0.027			0.046
× # Small firm *(instr.)*				(0.021)			(0.032)
Small firm				-0.009			-0.116
				(0.044)			(0.079)
First-stage F-statistic	425	181	312	119	75	142	35
Exogeneity test (p-value)	0.000	0.731	0.000	0.000	0.621	0.000	0.002
N	128,556	31,622	98,070	129,740	16,837	51,334	68,213
# of inventors	66,917	19,750	50,171	67,775	10,243	27,037	36,214
# of firms	2,020	2,111	330	2,231	685	140	790

The estimation period is 2001–2011. Robust standard errors are clustered by inventor (in parentheses). The exogeneity test is a Hausman-based test. The test for weak instruments is the Cragg-Donald Wald F-statistic. All regressions control for the number of years since the inventor's first decision (log), firm fixed effects, fixed effects for the number of applications filed by the inventor, the technology field, the year of the inventor's first decision and the calendar year. The sample in Columns 5–7 is restricted to Compustat (CS) firms. In Columns 2–4 and 5–7, a firm is classified as small (big) if its application portfolio or net sales are below (above) the 25^{th} sample percentile in a given year, which corresponds to 68 applications or 7,305 \$m in net sales. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

[21] Using alternative quantiles to classify small firms, such the lowest tercile or quintile, does not alter our results. Note that accounting for firm fixed effects in the analysis requires that even the smallest firms in our sample have at least two patent applications.

6.2. Patents as a signal of match quality

Patent grants may serve as signals of an inventor's human capital, as discussed in Section 2.3. If patents signal inventor ability, they should lead to increased mobility. This, however, appears to be inconsistent with our findings for early inventors. Patents can also signal a good employer-employee match, as they indicate that the inventor and the firm have successfully completed a discovery process together. If the match quality is not easily observable, the signaling effect of patents would lead to a decrease in mobility. Furthermore, the signal would be stronger when the inventor has fewer coauthors and the innovation falls outside the core of the firm. Therefore, match signaling is a potential alternative driver of our results.

Nevertheless, there are some arguments that question this alternative interpretation. First, the signaling mechanism is difficult to reconcile with the finding that patents have a stronger negative effect on mobility for innovations closer to basic research. Basic research is inherently less focused and more uncertain than applied research (Rosenberg, 1990). Accordingly, the value of innovations in this realm should be more difficult to identify, and patents covering them should provide noisier signals. Second, our focus on post-AIPA patent filings means that the vast majority of applications are already published by the grant date. This implies that much of the key technological information about patented innovations is released at the publication date (Hegde and Luo, 2018). Thus, although the grant decision may provide the market with some additional information on the quality of the match, we expect that the majority of such information is conveyed in the publication of the application.

6.3. Evidence on postmobility knowledge usage

To shed further light on our proposed mechanism, we analyze the impact of inventor mobility on the extent to which the patent holder continues to use her innovations. If the inventor's tacit knowledge plays an important role in the implementation and further development of these innovations, her departure will negatively impact the firm's ability to exploit them. We use patent expiration (i.e., the decision to not renew a patent) as our main measure of the firms' knowledge usage behavior. Since renewal fees have to be paid at regular intervals (every four years until year 12) to keep patent protection active, firms have to periodically determine whether patent renewal is worthwhile. Thus, if the premise of our analysis holds, we should expect an increase in the firms' propensity to let patents expire as a result of a mobility event.

We operationalize *patent expiration* with a dummy variable capturing whether at least one of the inventor's patents that had to go through its first renewal during the last spell has been left to expire.[22] As an alternative proxy for the degree of technology usage, we draw on the citations received by the inventor's patent(s) from future patents by the patent holder, commonly referred to as *self-citations*. If innovation-related knowledge is important for retention decisions, inventor departures should curtail the firm's ability to build on her patented innovations, which should be reflected fewer in self-citations.

The sample for this analysis is restricted to inventors with at least one granted patent and is, therefore, not a random subsample of the population of inventors. As in Section 4.4, we address this problem by using Heckman's two-step sample selection correction under the identifying assumption that the selection model contains one covariate – i.e., examiner leniency – that is correlated with patent grants but not with patent use (renewal decisions and self-citations). A more general concern in this analysis is that mobility events cannot be considered exogenous determinants of patent renewals (or citation patterns). Unobservable technology characteristics can simultaneously affect mobility and patent renewal decisions (and citation patterns). We address this issue by using plausibly exogenous instrumental variables of inventor mobility based on the level of enforcement of noncompete covenants in the inventor's state. In particular, we use the enforceability index proposed by Starr (2019) and its interaction with the sex of the inventor.[23] We implement the instrumental variable estimator using 2SLS and introduce an extensive set of inventor-level and technology-level control variables.

In Table 7, we first present the results of the analysis with the expired patent dummy as the dependent variable. In Column 3, we observe a positive and significant estimate of the coefficient of mobility. After an inventor moves, there is a 53% increase in the probability that her former employer allows at least one of her patents to expire. This result is consistent with our argument that inventor mobility causes a loss of implementation know-how. The idea is reinforced by the analysis of self-citations received by the inventor's patents. Column 5 reports a negative and significant coefficient of mobility on the number of self-citations, suggesting that inventor departures lead to less follow-on innovation by the originating firm.

[22] Information on patent maintenance fee events since September 1981 is available on the USPTO website.

[23] Sex as an instrument for mobility has been used in the past, for example, by Kim and Marschke (2005). Inventor sex was coded by matching the first name of inventors listed on the patent with common male and female first names collected from Gender API (see https://gender-api.com).

Table 7: Postmobility, patent expiration and knowledge usage

Estimation method	OLS	OLS (1^{st} st.)	2SLS (2^{nd} st.)	OLS	2SLS (2^{nd} st.)
Dependent variable	Patent expired	Move	Patent expired	# of post self cites (L)	# of post self cites (L)
	(1)	(2)	(3)	(4)	(5)
Move	0.022***			-0.073***	
	(0.008)			(0.011)	
Enforceability index		-0.006***			
		(0.001)			
Enforceability index × Female		-0.009**			
		(0.004)			
Move *(instr.)*			0.532**		-0.652*
			(0.258)		(0.379)
Female	0.026**	-0.001	0.023*	-0.023	-0.022
	(0.012)	(0.006)	(0.013)	(0.017)	(0.017)
# of post cites (L)				0.174***	0.172***
				(0.008)	(0.008)
# of pre cites (L)	-0.041***	0.014***	-0.048***	-0.069***	-0.061***
	(0.004)	(0.002)	(0.005)	(0.007)	(0.008)
# of pre self-cites (L)	-0.023***	-0.024***	-0.011	0.498***	0.477***
	(0.004)	(0.002)	(0.008)	(0.011)	(0.012)
# of patents granted (L)	0.230***	-0.003	0.216***	-0.011	0.001
	(0.007)	(0.003)	(0.009)	(0.011)	(0.014)
% of patents in firm's core	-0.046***	-0.016***	-0.037***	0.094***	0.085***
	(0.009)	(0.004)	(0.010)	(0.013)	(0.014)
Average patent scope	-0.013***	0.004*	-0.017***	0.001	0.006
	(0.003)	(0.002)	(0.004)	(0.006)	(0.007)
# of coinventors (L)	0.005	-0.044***	0.009	0.024***	0.032***
	(0.004)	(0.002)	(0.006)	(0.006)	(0.009)
# of applications per firm (L)	0.039***	-0.013***	0.046***	0.013***	0.005
	(0.002)	(0.001)	(0.004)	(0.003)	(0.006)
Inverse mills ratio	0.212**	0.041	0.155	-0.450***	-0.350**
	(0.092)	(0.058)	(0.104)	(0.128)	(0.141)
F-statistic		40			
Exogeneity test (p-value)			0.000		0.000

$N = 33,813$. Number of inventors: 18,043. Number of firms: 1,109. The estimation period is 2001–2011. Robust standard errors are clustered by inventor (in parentheses). The exogeneity test is a Hausman-based test. The test for weak instruments is the Cragg-Donald Wald F-statistic. All regressions control for the number of years since the inventor's first decision (log), fixed effects for the technology field, the year of the inventor's first decision and the calendar year. The enforceability scores for each state are from Starr (2019). Since noncompetes are enforceable in the state where the employee is located, we use information on the inventor's location from the last patent application prior to observing the outcome. *# of pre self-cites (pre cites)* are the citations received from patents owned by the focal (other) assignees before year t (i.e., the year in which an inventor is observed to stay or exit). *# of post self-cites (post cites)* are the citations received from patents owned by the focal (other) assignees in a 5-year window after year t. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

To sum up, the overall negative effect of patents on mobility documented in this article for early-career inventors does not have a unique possible interpretation. In particular, it is plausible that the role of patents as signals of a good inventor-firm match partially accounts for it. Nevertheless, given the post-AIPA context of our analysis (with pre-grant publication of applications), the results of our heterogeneity analysis, and the evidence of a decrease in post-mobility usage of patented knowledge, our preferred explanation for the main findings is the existence of innovation-related skills specific to the patented technology. Finally, although our analysis does not explicitly test the possibility that patents act as signals of inventor ability, our findings suggest that such an ability-signalling process would be dominated by other mechanisms that operate in the opposite direction in the relationship between patenting and mobility.

7. Conclusion

In this study, we investigate the effect of patent protection on inventor mobility. We argue that patent rights decrease interfirm mobility to the extent that inventors have tacit knowledge that is especially valuable in the implementation or further development of their patented innovations. Patents convert these innovation-related skills into patent-holder-specific human capital, thus discouraging mobility.

To test this idea, we examine the impact of obtaining a patent on the mobility patterns of early-career inventors involved in patent applications. Since inventions that result in granted and not granted patents are expected to be inherently different, we adopt an instrumental variables approach to estimate the effect of patenting on inventor mobility. In particular, we use the variations in the granting rates among patent examiners within an art unit and filing year as an exogenous source of variation in patent grants. We analyze the early careers of employee inventors who apply for patents at the USPTO for the first time between 2001 and 2011. Our results indicate that receiving a patent grant causes a substantial decrease in mobility. Moreover, we find support for the following additional predictions that characterize our proposed mechanism: (i) patents have a weaker negative effect on mobility in the presence of other sources of firm-specific human capital (collaboration with coinventors and working in the technological core of the firm), and (ii) patents have a stronger negative effect on mobility when the involvement of the inventor is particularly important for implementation (innovations built on basic research). We also provide complementary evidence showing that an inventor move negatively affects the subsequent use of her patented innovations, which is consistent with the idea that inventors hold innovation-related skills useful for implementation. Overall, this combination of findings suggests that, while other mechanisms may play a role in explaining the negative effect of patent protection on inventor mobility, an interpretation based on patent-holder-specific human capital is the most plausible in our setting.

These results have some relevant implications. First, they suggest that patents, despite disclosing codified technical knowledge, may hamper the diffusion of tacit know-how that comes with interfirm mobility. There is an ongoing debate among scholars about whether patent rights help or hinder the spread of knowledge and, ultimately, the generation of further innovations. On the one hand, patents encourage the disclosure of information that would otherwise be kept secret and promote technology trade (Moser, 2011). On the other hand, the recombinant nature of new knowledge, coupled with the fragmentation of ownership rights and steep transaction costs, may turn the patent system into an obstruction for the transfer of technological information (Heller and Eisenberg, 1998). Recent research on the impact of patents on subsequent innovation has produced mixed evidence (Galasso and Schankerman, 2015; Sampat and Williams, 2019). In this context, our examination of the effect of patents on inventor mobility adds a new perspective to the debate, allowing for a more comprehensive view of the relationship between patents and the diffusion of knowledge. In particular, our results suggest that patent protection, by inducing lower mobility, may as well reduce the spread of noncodified knowledge associated with the protected technology and, more generally, of other know-how not related to the replicability of a specific innovation.

Furthermore, the reduction in mobility induced by patents may also contribute to an inefficient allocation of inventive skills. To the extent that inventors' career moves are motivated by employer-employee match improvements, patents, by discouraging mobility, will inhibit these efficiency improvements. On the bright side, our results suggest that patents, by making some inventor skills patent-holder-specific, shift the incentives to invest in human capital from the employees (i.e., the inventors) to their employers (i.e., the patent holders). This shift may encourage some efficient investments in training that might not have been otherwise made by the inventors themselves because of financial constraints or risk considerations. More generally, by contributing to the retention of their R&D workers, the patent system may provide incentives to companies to invest in long-term-oriented human resources practices. Finally, our findings bring to light an important methodological issue. If patent protection affects the mobility of inventors, tracking their careers through their issued patents (as most studies have done thus far) introduces a downward bias in the detection of mobility. This bias may also extend to the analysis of any research question related to inventor career dynamics, from mobility studies to analyses of knowledge spillovers or coinventor networks. To avoid this problem, future studies in the field should rely on information from all, not only granted, patent applications.

References

Akcigit, U., S. Baslandze, and S. Stantcheva (2016). Taxation and the international mobility of inventors. *American Economic Review 106*(10), 2930–81.

Arora, A., S. Belenzon, and A. Patacconi (2017). The decline of science in corporate R&D. *Strategic Management Journal 39*(1), 3–32.

Arora, A. and M. Ceccagnoli (2006). Patent protection, complementary assets, and firms' incentives for technology licensing. *Management Science 52*(2), 293–308.

Arrow, K. (1962). The economic implications of learning by doing. *The Review of Economic Studies 29*(3), 155–173.

Balasubramanian, N., J. W. Chang, M. Sakakibara, J. Sivadasan, and E. Starr (2018). Locked in? The enforceability of covenants not to compete and the careers of high-tech workers. US Census Bureau Center for Economic Studies Paper No. CES-WP-17-09.

Beck, T., A. Demirguc-Kunt, and V. Maskimovic (2005). Financial and legal constraints to growth: Does firm size matter? *The Journal of Finance 60*(1), 137–177.

Blundell, R. and J. Powell (2004). Endogeneity in semiparametric binary response models. *The Review of Economic Studies 71*(3), 655–679.

Campbell, B. A., M. Ganco, A. M. Franco, and R. Agarwal (2012). Who leaves, where to, and why worry? Employee mobility, entrepreneurship and effects on source firm performance. *Strategic Management Journal 33*(1), 65–87.

Cassiman, B., R. Veugelers, and S. Arts (2018). Mind the gap: Capturing value from basic research through combining mobile inventors and partnerships. *Research Policy 47*(9), 1811–1824.

Cockburn, I., S. Kortum, and S. Stern (2003). Are all patent examiners created equal? Examiners, patent characteristics and litigation outcomes. In W. Cohen and S. Merril (Eds.), Patents in the Knowledge-Based Economy, pp. 17–53. National Academies Press.

Fallick, B., C. Fleischman, and J. Rebitzer (2006). Job-hopping in silicon valley: Some evidence concerning the microfoundations of a high-technology cluster. *The Review of Economics and Statistics 88*(3), 472–481.

Feng, J. and X. Jaravel (forthcoming). Crafting intellectual property rights: Implications for patent assertion entities, litigation, and innovation. *American Economic Journal: Applied Economics*.

Fleming, L. and O. Sorenson (2004). Science as a map in technological search. *Strategic Management Journal 25*(8-9), 909–928.

Frakes, M. and M. Wasserman (2016). Patent office cohorts. *Duke Law Journal 65*(8), 1601–1655.

Galasso, A. and M. Schankerman (2015). Patents and cumulative innovation: Causal evidence from the courts. *The Quarterly Journal of Economics 130*(1), 317–369.

Ganco, M. (2013). Cutting the gordian knot: The effect of knowledge complexity on employee mobility and entrepreneurship. *Strategic Management Journal 34*(6), 666–686.

Ganco, M., R. Ziedonis, and R. Agarwal (2015). More stars stay, but the brightest ones still leave: Job hopping in the shadow of patent enforcement. *Strategic Management Journal 36*(5), 659–685.

Gaulé, P. (2018). Patents and the success of venture-capital backed startups: Using examiner assignment to estimate causal effects. *The Journal of Industrial Economics 66*(2), 350–376.

Ge, C., K.-W. Huang, and I. Png (2016). Engineer/scientist careers: Patents, online profiles, and misclassification bias. *Strategic Management Journal 37*(1), 232–253.

Gilson, R. J. (1999). The legal infrastructure of high technology industrial districts: Silicon valley, route 128, and covenants not to compete. *New York University Law Review 74*, 575.

Graham, S. J., A. C. Marco, and R. Miller (2018). The USPTO Patent Examination Research Dataset: A window on patent processing. *Journal of Economics & Management Strategy 27*(3), 554–578.

Hegde, D. and H. Luo (2018). Patent publication and the market for ideas. *Management Science 64*(2), 652–672.

Heller, M. and R. Eisenberg (1998). Can patents deter innovation? The anticommons in biomedical research. *Science 280*(5364), 698–701.

Hoisl, K. (2007). Tracing mobile inventors – The causality between inventor mobility and inventor productivity. *Research Policy 36*(5), 619–636.

Hombert, J. and A. Matray (2016). The real effects of lending relationships on innovative firms and inventor mobility. *The Review of Financial Studies 30*(7), 2413–2445.

Kim, J. and G. Marschke (2005). Labor mobility of scientists, technological diffusion, and the firm's patenting decision. *The RAND Journal of Economics 36*(2), 298–317.

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman (2017). Technological innovation, resource allocation, and growth. *The Quarterly Journal of Economics 132*(2), 665–712.

Kuhn, J. M. and N. C. Thompson (2019). How to measure and draw causal inferences with patent scope. *International Journal of the Economics of Business 26*(1), 5–38.

Lanjouw, J. and M. Schankerman (2004). Patent quality and research productivity: Measuring innovation with multiple indicators. *The Economic Journal 114*(495), 441–465.

Lemley, M. and B. Sampat (2012). Examiner characteristics and patent office outcomes. *Review of Economics and Statistics 94*(3), 817–827.

Lerner, J. (1994). The importance of patent scope: An empirical analysis. *The RAND Journal of Economics 25*(2), 319–333.

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

Mann, W. (2018). Creditor rights and innovation: Evidence from patent collateral. *Journal of Financial Economics 130*(1), 25–47.

Marco, A., A. Myers, S. Graham, P. D'Agostino, and K. Apple (2015). The USPTO patent assignment dataset: Descriptions and analysis. USPTO Working Paper No. 2015-2.

Marco, A. C., J. D. Sarnoff, and C. A. deGrazia (2019). Patent claims and patent scope. *Research Policy 48*(9), 103790.

Marx, M. (2011). The firm strikes back: Non-compete agreements and the mobility of technical professionals. *American Sociological Review 76*(5), 695–712.

Marx, M., D. Strumsky, and L. Fleming (2009). Mobility, skills, and the michigan non-compete experiment. *Management Science 55*(6), 875–889.

Merges, R. P. (1999). The law and economics of employee inventions. *Harvard Journal of Law & Technology 13*(1), 1–54.

Moretti, E. and D. J. Wilson (2017). The effect of state taxes on the geographical location of top earners: Evidence from star scientists. *American Economic Review 107*(7), 1858–1903.

Moser, P. (2011). Do patents weaken the localization of innovations? Evidence from world's fairs. *The Journal of Economic History 71*(2), 363–382.

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

Paruchuri, S., A. Nerkar, and D. C. Hambrick (2006). Acquisition integration and productivity losses in the technical core: Disruption of inventors in acquired companies. *Organization Science 17*(5), 545–562.

Righi, C. and T. Simcoe (2019). Patent examiner specialization. *Research Policy 48*(1), 137–148.

Rosenberg, N. (1990). Why do firms do basic research (with their own money)? *Research Policy 19*(2), 165–174.

Sampat, B. and H. L. Williams (2019). How do patents affect follow-on innovation? Evidence from the human genome. *American Economic Review 109*(1), 203–236.

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

Song, J., P. Almeida, and G. Wu (2003). Learning-by-hiring: When is mobility more likely to facilitate inter-firm knowledge transfer? *Management Science 49*(4), 351–365.

Starr, E. (2019). Consider this: Training, wages, and the enforceability of covenants not to compete. *ILR Review 72*(4), 783–817.

Starr, E., J. J. Prescott, and N. Bishara (2019). Noncompetes in the U.S. labor force. University of Michigan

Law & Economics Research Paper No. 18-013.

Stock, J. and M. Yogo (2005). Testing for weak instruments in linear IV regression. In D. Andrews (Ed.), *Identification and Inference for Econometric Models*, pp. 80–108. New York: Cambridge University Press.

Zucker, L. G., M. R. Darby, and M. B. Brewer (1998). Intellectual human capital and the birth of U.S. biotechnology enterprises. *The American Economic Review* *88*(1), 290–306.

Zucker, L. G., M. R. Darby, and M. Torero (2002). Labor mobility from academe to commerce. *Journal of Labor Economics* *20*(3), 629–660.

Online Appendix for "The Effect of Patent Protection on Inventor Mobility" by Eduardo Melero, Neus Palomeras and David Wehrheim

This Online Appendix provides additional material to the results presented in *"The Effect of Patent Protection on Inventor Mobility."* In Section A.1, we provide the details on the conditions under which patent protection leads to a reduction in inventor mobility and derive the testable implications presented in Section 2.2 of the main article. In Section A.2, we discuss and report additional robustness checks for the instrumental variables approach that supplement Section 4.2 of the paper. In Section A.3, we offer a summary of the relative effect sizes of our estimates of patenting on mobility across the different specifications and subsamples reported in the paper.

A.1. Patent protection and inventor mobility: Formal derivation of testable implications

To characterize the elements of Inequality (1) in the main article, we assume that the profits from an innovation derive from its market introduction. The demand for the innovation can be represented by the following inverse linear demand function relating price p to quantity q: $p = a - bq$, with $a, b > 0$. The implementation advantage contributed by the inventor's innovation-related skills takes the form of lower production costs. Because those skills are partially firm specific, we need to define three types of situations. First, any firm taking the innovation to market without the help of the inventor will face a constant marginal production cost of c, with $0 < c < \frac{a}{2}$. Second, if the inventor remains with the current employer, the employer can introduce the innovation at a marginal cost of 0. Finally, if a different employer hires away the inventor and implements the innovation using her contribution, this firm will face a marginal cost of dc, with $0 < d < 1$. Therefore, c captures the cost reduction due to the inventor's innovation-related skills, and d captures the degree of firm specificity of those skills.

Within this framework, if a patent is granted, the current employer of the inventor will obtain monopoly profits $\pi^M_{i,j} = \frac{a^2}{4b}$ if the inventor stays and $\pi^M_{-i,j} = \frac{(a-c)^2}{4b}$ if the inventor leaves. Alternative employers cannot make any profit from the patent-protected innovation. If the patent is not granted, any competitor can implement the innovation at a marginal cost of c. The firm employing the inventor, however, enjoys a cost advantage that allows it to set a price of $p = c$ and serve the whole market. The firm employing the inventor, therefore, will obtain profits $\pi^C_{i,j} = \frac{c(a-c)}{b}$ if it is the initial employer j and $\pi^C_{i,-j} = \frac{(c-dc)(a-c)}{b}$ if it is any other employer. Therefore, Inequality (1) in the paper – stating the conditions under which patents have a negative effect on mobility – can be rewritten as follows:

$$\frac{2ac - c^2}{4b} - \frac{dc(a-c)}{b} > 0 \tag{A.1}$$

After some rearrangement of terms, the condition established in Inequality (A.1) can be reduced to $d < \hat{d}$, with $\hat{d} = \frac{1}{4}(\frac{a}{a-c} + 1) > 0$. In other words, patent protection generates a negative effect on inventor mobility as long as the share of the inventor's innovation-related skills that is firm specific is small enough.

Our second implication from Section 2.2 in the paper is that the effect of patenting on mobility is positively moderated by the importance of the inventor's firm-specific innovation-related skills, as captured by her complementarities with the firm's resources. The negative effect of patents on mobility is attenuated or can even turn positive as d increases. This implication can be easily proved using Inequality (A.1). The left-hand side of this inequality captures the net change caused by patents in the balance of the inside vs. outside value of innovation-related skills. Hence, the probability of switching employers decreases monotonically with that expression. Consequently, stating that firm-specific skills positively moderate the effect of patent protection on mobility is equivalent to stating that the left-hand side of Inequality (A.1) decreases with d. In mathematical terms, we need to show that the derivative of such an expression with respect to d is negative. That is:

$$-\frac{c(a-c)}{b} < 0 \tag{A.2}$$

which always holds given our assumptions regarding a, b and c.

The two implications stated above follow a similar rationale. Without patent protection, monopoly pricing is not possible. Instead, the firm employing the inventor sets $p = c$, which is below the optimal monopoly price, and sells an amount of the innovative product above the optimal monopoly amount. This means that the marginal profit of a cost reduction is higher without patent protection. This is the *product market effect*: the marginal value of firm-specific skills dc is lower with patent protection, leading to a positive effect on mobility. On the other hand, patent protection turns the general component of the inventor's innovation-related skills,

$(1-d)c$, into patent-holder-specific human capital. This is the *human capital effect*, which generates an important retention effect of patents that will dominate in the determination of the total effect if d is small enough, and it becomes weaker as d increases.

Our third implication is that the negative effect of patent protection on mobility should be more intense when the innovation-related skills of the inventor are particularly relevant for the implementation of the new technology (for example, when it is closer to basic science). In our model, those skills are captured by c. Therefore, the implication is equivalent to stating that the derivative of the left-hand side of Inequality (A.1) with respect to c must be positive, at least for the range where $d < \hat{d}$. Such a derivative can be stated as follows:

$$\frac{a-c}{2b} - \frac{d(a-2c)}{b} > 0 \tag{A.3}$$

After some rearrangement of terms, Inequality (A.3) can be expressed as $d < \hat{\hat{d}}$, with $\hat{\hat{d}} = \frac{1}{2}(\frac{c}{a-2c}+1)$. Thus, the effect of patent protection on mobility is negatively moderated by the importance of innovation-related skills only if a small enough share of those skills is firm specific. If this range includes the range for which patents have a negative effect on mobility ($d < \hat{d}$), then the negative effect of patenting on mobility is intensified by the importance of innovation-related skills. In other words, the third implication is generally true as long as $\hat{d} < \hat{\hat{d}}$, which is always the case given our assumptions regarding a, b and c.

A.2. Exogeneity checks for average examiner leniency

A.2.1. Investigating selection

For the average examiner leniency to be a valid instrumental variable for the number of patents granted to an inventor, it must satisfy the exclusion restriction. The institutional details provided in Section 3.3 in the article suggest that the assignment of applications to examiners is plausibly orthogonal to the drivers of inventor mobility, and the exogeneity tests for examiner leniency provided at the patent level by Sampat and Williams (2019) suggest that this is the case. We provide two empirical tests that support the validity of the instrument at the inventor level.[1]

First, we test whether the application portfolios of the inventors assigned to tough or lenient examiners are similar in terms of the observable characteristics that predict granting probability

[1] We thank an anonymous referee for suggesting these two pieces of analysis.

at the time of filing. As discussed in Lemley and Sampat (2012), the assessment of whether a certain type of invention is assigned to examiners with a certain leniency is challenging for the following two reasons: (i) it is difficult to identify variables that at the time of application would capture the characteristics of the underlying invention, and (ii) much of the front-page information contained in patent documents is not available in applications. Sampat and Williams (2019) propose the following two measures for such a purpose that are available at the time of filing: patent family size (# of countries in which patent equivalents have been applied for) and the number of claims listed in the application. We follow their approach and assess whether these variables are predictive of patent grants. In line with the findings in Sampat and Williams (2019), an F-test of joint significance reveals that both measures have a significant positive effect on grant decisions (F=113; p-value = 0.000). Next, we average the predicted grant probabilities across individuals and regress it on average examiner leniency. This produces an estimated coefficient of 0.0005 with a standard error of 0.0005, suggesting that there is no systematic relationship between the two variables. As displayed in the nonparametric representation in Figure 1 in the article, there is also no visual relationship between the average predicted probabilities of patent grants and our instrument of average examiner leniency.

Second, we examine the correlation between examiner leniency and an index of the likelihood of changing employers based on the inventor's set of observable characteristics. In addition to the array of fixed effects, this index includes several additional individual characteristics based on pre-patent filing measures. In particular, we follow prior patent-based studies on the determinants of inventor mobility (e.g., Song et al., 2003; Hoisl, 2007; Palomeras and Melero, 2010; Marx, 2011) and include a measure of inventor productivity (# of applications filed/tenure), inventor specialization (# of unique patent classes), the inventor's collaborative ties (# of unique coinventors), her expertise in the firm's core technology areas (as % of total applications filed), her experience (time since 1st decision) and the size of her current employer (# of applications filed by the employer). We run a probit model, which is more suitable for the purpose of obtaining specific predictions than a linear probability model (Wooldridge, 2014). Column 1 of Table A.1 reports the results for the determinants of inventor mobility, which are all statistically significant. The resulting mobility index captures approximately 16% of the variation in actual mobility. In Column 2 of Table A.1, we then regress this mobility index on examiner leniency. As the results show, the correlation between $\widehat{Move}$ and examiner leniency is close to zero and statistically insignificant. Thus, there appears to be no evidence that inventors with characteristics that make them more likely to move are particularly likely to be assigned to more lenient examiners.

Table A.1: Examiner leniency, and patent grants and mobility

	Panel A: Investigating selection		Panel B: Deviations vs. fixed effects	
Sample	All	All	1^{st} decision	1^{st} decision
Estimation method	Probit	OLS	OLS	OLS
Dependent variable	Move	$\widehat{Move}$	Patent issued	Patent issued
	(1)	(2)	(3)	(4)
Examiner leniency		-0.002	0.846***	
		(0.005)	(0.018)	
Examiner approval rate				0.846***
				(0.018)
Years since 1^{st} decision (L)	0.020***			
	(0.003)			
# of coinventors (L)	-0.009***			
	(0.001)			
% of applications in firm's core	-0.039***			
	(0.003)			
# of USPC classes (L)	-0.042***			
	(0.002)			
# of applications per firm (L)	-0.034***			
	(0.002)			
Productivity (L)	-0.122***			
	(0.005)			
Firm FE	Yes	No	No	No
First decision year FE	Yes	No	No	No
Calendar year FE	Yes	No	No	No
# of applications filed FE	No	No	No	No
Technological class FE	Yes	No	No	Art unit × Filing year
N	129,740	129,740	66,285	66,285
# of inventors	67,775	67,775	66,285	66,285
# of firms	2,231	2,231	2,052	2,052

The estimation period is 2001–2011. Robust standard errors (in parentheses) are clustered by inventor in columns 1 and 2; and by art unit in columns 3 and 4. Column 1 displays average marginal effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.2.2. Fixed effects vs. deviations

The work by Sampat and Williams (2019) supports the validity of examiner leniency as an instrument for patent grants, when accounting for art-unit-year fixed effects. Because inventors may file applications in different technological areas, we cannot explicitly control for such fixed effects in our regressions. Instead, we implicitly account for them by averaging, at the inventor level, the differences between the grant rate of the assigned examiner and the grant rate of the corresponding art unit and year [see Equations (3), (4) and (5) in Section 3.3.1 in the article]. To show that the two approaches are essentially equivalent, in Columns 3 and 4 of Table A.1, we restrict our analysis to the first application decision. In Column 3, we use the deviations from art-unit means in a given application year, and in Column 4, we use the examiner's raw approval rate in combination with a set of art-unit-by-application-year fixed effects. As can be seen, in both specifications, we obtain a coefficient of 0.846 with a standard error of 0.018. Interestingly, these point estimates are also similar in magnitude to that reported in Sampat and Williams (2019).[2]

A.2.3. Addressing examiner specialization

In this subsection, we address the concern raised by Righi and Simcoe (2019), who provide evidence of patent examiner specialization within art units. If such specialization is correlated with examination outcomes (as also suggested by these authors) and, simultaneously, with omitted factors that drive employee-inventor mobility patterns, then our instrumental-variable estimator would be flawed.

Since Righi and Simcoe (2019) note that examiner specialization disappears within subclasses for most technology centers, we replicate our instrument at the art-unit-subclass level. However, at this level of disaggregation, we are left with approximately one-third of the observations since many cells are not available for estimation. In addition, we need to exclude those technology centers from the analysis where, according to Righi and Simcoe (2019), some within-subclass specialization remains [i.e., (1600) Biotech and Organic Chemistry and (1700) Chemistry and Materials Engineering]. Thus, to obtain an instrument with sufficient variation, we build the within-subclass examiner leniency measure using a 3-year moving window when computing grant rates. Everything else remains the same as in our main analysis. Panel A of Table A.2 reports the replication of our baseline results on this subsample of 74,288 observations and the alternative instrument. As can be seen, the coefficient on the instrumented number of patents granted (in Column 3) continues to be negative and significant.

A.2.4. "Next-to-random" allocation of patent applications

In Panel B of Table A.2, we repeat our main analysis for a subsample of patent applications examined in art units where the "next-to-random" assumption on the allocation of applications to examiners is particularly likely to hold. According to qualitative evidence reported by Lemley and Sampat (2012), some art units allocate applications to examiners based on the last digit of the serial number of the application. Because the serial numbers are assigned at the

[2] For the sake of comparison with Sampat and Williams (2019), we do not include firm fixed effects in these specifications.

central office of the USPTO according to the order of arrival, their last digits are orthogonal to the application characteristics. Accordingly, Feng and Jaravel (forthcoming) use data from USPTO records to make a systematic analysis of this pattern that allows them to determine which art units tend to use a last-digit rule for the assignment of applications.[3] Relying on their classification, we repeat our baseline analysis for the subsample of art units that tend to allocate applications to examiners using the last digits of serial numbers.[4] The 2SLS coefficient presented in Column 6 of Table A.2 shows that our results are virtually unaffected by the removal of potentially problematic art-unit-years; considering that the probability of changing employers is reduced to 9% for this subsample, these estimates imply that an exogenous increase of one granted patent decreases mobility by approximately 22%.

Table A.2: Addressing examiner specialization

	Panel A: Examiner leniency within USPC-subclasses			Panel B: Art units with assignment based on last digit		
Estimation method	OLS	OLS (1st st.)	2SLS (2nd st.)	OLS	OLS (1st st.)	2SLS (2nd st.)
Dependent variable	Move	# of pats issued	Move	Move	# of pats issued	Move
	(1)	(2)	(3)	(4)	(5)	(6)
# of patents issued	0.002***			0.002***		
	(0.000)			(0.000)		
# of patents issued *(instr.)*			-0.086***			-0.020**
			(0.029)			(0.009)
Examiner leniency		0.450***			1.289***	
		(0.098)			(0.118)	
F-statistic		21			170	
Exogeneity test (p-value)			0.000			0.006
N	74,288	74,288	74,288	74,146	74,146	74,146
# of inventors	34,753	34,753	34,753	35,944	35,944	35,944
# of firms	1,573	1,573	1,573	1,432	1,432	1,432

The estimation period is 2001–2011. Robust standard errors are clustered by inventor (in parentheses). The exogeneity test is a Hausman-based test. The test for weak instruments is the Cragg-Donald Wald F-statistic. All regressions control for the number of years since the inventor's first decision (log), firm fixed effects, fixed effects for the number of applications filed by the inventor, the technology field, the year of the inventor's first decision and the calendar year. Following Righi and Simcoe (2019), the instrument at the art-unit-year-USPC-subclass level excludes the "Biotechnology and Organic Chemistry" and "Chemistry and Materials Engineering" technology centers. The art-unit-years where application assignment to examiners is determined by the last digit of the serial number of the patent application are from Feng and Jaravel (forthcoming). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

As with any instrumental variable, we cannot completely rule out the possibility that average examiner leniency is correlated with some unobserved factor affecting our outcome of interest.

[3] Feng and Jaravel (forthcoming) rely on the same divergence index and test statistics used in Righi and Simcoe (2019) to detect specialization, but they instead use it to detect the degree of concentration of patent examiners on certain last digits in comparison to a theoretical uniform distribution.

[4] We thank the authors for sharing their data with us.

However, the different exogeneity tests performed in this section place some bounds on that concern. This, combined with qualitative and quantitative evidence provided in related studies, offers us some confidence that this variable satisfies the exclusion restriction for the purpose of our study.

A.3. Summary of results of the article

Table A.3: Relative sizes of estimated effects (one additional patent)

Description	Specification	Observations	Average Mobility	Relative Impact
Baseline	Table 2, column 3	129,740	0.10	23%
Additional controls	Table 5, column 1	129,740	0.10	23%
Narrower tech classes FEs	Table 3, column 1	129,721	0.10	27%
Reviewed applications	Table 3, column 2	127,248	0.10	17%
First decision	Table 3, column 3	66,285	0.15	18%
Nonsuperstar inventors	Table 3, column 4	123,596	0.11	24%
Pharmaceutical inventors	Table 3, column 5	23,069	0.12	25%
Filing year $\leqslant$ 2006	Table 3, column 6	58,351	0.16	21%
Industries with low patenting intensities	Table 3, column 7	34,567	0.09	26%
Industries with high patenting intensities	Table 3, column 8	33,563	0.13	26%
Active firms	Table 6, column 1	128,556	0.10	25%
Larger firms (in terms of applications)	Table 6, column 3	98,070	0.09	31%
Larger firms (in terms of sales)	Table 6, column 6	51,334	0.10	36%

References

Feng, J. and X. Jaravel (forthcoming). Crafting intellectual property rights: Implications for patent assertion entities, litigation, and innovation. *American Economic Journal: Applied Economics*.

Hoisl, K. (2007). Tracing mobile inventors – The causality between inventor mobility and inventor productivity. *Research Policy 36*(5), 619–636.

Lemley, M. and B. Sampat (2012). Examiner characteristics and patent office outcomes. *Review of Economics and Statistics 94*(3), 817–827.

Marx, M. (2011). The firm strikes back: Non-compete agreements and the mobility of technical professionals. *American Sociological Review 76*(5), 695–712.

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

Righi, C. and T. Simcoe (2019). Patent examiner specialization. *Research Policy 48*(1), 137–148.

Sampat, B. and H. L. Williams (2019). How do patents affect follow-on innovation? Evidence from the human genome. *American Economic Review 109*(1), 203–236.

Song, J., P. Almeida, and G. Wu (2003). Learning-by-hiring: When is mobility more likely to facilitate inter-firm knowledge transfer? *Management Science 49*(4), 351–365.

Wooldridge, J. (2014). Quasi-maximum likelihood estimation and testing for nonlinear models with endogenous explanatory variables. *Journal of Econometrics 182*(1), 226–234.