DISCUSSION PAPER SERIES

Labor Mobility, Social Network Effects, and Innovative Activity

Ulrich Kaiser
Hans Christian Kongsted
Thomas Rønde

April 2011

I Z A

Labor Mobility, Social Network Effects, and Innovative Activity

Ulrich Kaiser
University of Zurich,
CEBR, ZEW, CIE and IZA

Hans Christian Kongsted
University of Copenhagen,
CIE, CAM and CEBR

Thomas Rønde
Copenhagen Business School,
CEBR, ZEW, CIE and CEPR

Discussion Paper No. 5654
April 2011

ABSTRACT

Labor Mobility, Social Network Effects, and Innovative Activity[*]

We study the mapping between labor mobility and industrial innovative activity for the population of R&D active Danish firms observed between 1999 and 2004. Our study documents a positive relationship between the number of workers who join a firm and the firm's innovative activity. This relationship is stronger if workers join from innovative firms. We also find evidence for *positive* feedback from workers who leave for an innovative firm, presumably because the worker who left stays in contact with their former colleagues. This implies that the positive feedback ("social network effects") that has been found by other studies not only exists but even outweighs the disruption and loss of knowledge occurring to the previous employer from the worker leaving. Summing up the effects of joining and leaving workers, we find ample evidence for mobility to be associated with an increase in *total* innovative activity of the new and the old employer.

JEL Classification: O33, O34, C23

Keywords: labor mobility, innovation, social network

Corresponding author:

Ulrich Kaiser
University of Zurich
Department of Business Administration
Plattenstr. 14
8032 Zurich
Switzerland
E-mail: ulrich.kaiser@business.uzh.ch

1   Introduction

There is by now evidence from surveys (Mansfield, 1985; Zander and Kogut, 1995), patent files (Almeida and Kogut, 1999; Song et al., 2003), and litigation (Hoti et al., 2006) showing that labor mobility is an important source of inter-firm knowledge transfer. Firms are aware of the opportunities that mobility creates and hire inventors to acquire technological competencies and to enter distant technological areas (Palomeras and Melero, 2009; Rosenkopf and Almeida, 2003; Singh and Agrawal, 2011). They also hire expert managers to introduce new types of products (Boeker, 1997; Rao and Drazin, 2002). Firms experiencing worker exit may, however, encounter a loss of skills and knowledge — or knowledge exclusivity — that reduces their competitiveness, sometimes to the extent of threatening firm survival (Wezel et al., 2006).

A related body of literature has studied the regional effects of labor mobility. The focus has been primarily on the world's most successful high-tech cluster, Silicon Valley, where Almeida and Kogut (1999) and Fallick et al. (2006) have documented the co-existence of high labor turnover and localized knowledge sharing among firms.[1] Saxenian (1994) stresses how the "job-hopping" culture in Silicon Valley creates tightly coupled social networks through which knowledge flows. This view is supported by recent works showing that firms learn from workers who have left, presumably because they stay in contact with their former colleagues (Agrawal et al., 2006; Corredoira and Rosenkopf, 2010).[2]

The research on labor mobility is motivated by the notion that knowledge sharing is a powerful source of innovation and economic growth at the regional level. Codified knowledge is non-rival in nature, and one firm using a piece of knowledge does not prevent other firms from using it at the same time (Arrow, 1962). While the positive effects of knowledge sharing certainly arise if the knowledge flows are pure externalities, as envisioned in economic growth literature, this is less obviously the case if knowledge travels embodied within workers. The disruptive effect of worker exit together with the loss of skills and knowledge might decrease the level of innovation by the old employer to such an extent that it outweighs the positive effect on innovation by the new employer. A key question in this debate is therefore whether labor mobility is associated with an increase in total innovation by the firms involved. And if it is, the question is how strong these effects are. This paper sets out to answer both of these questions.

Previewing our results, we find a significant and positive effect of labor mobility on the firms' joint level of invention.[3] The effect is composed of a strictly positive effect on the level of invention by the new employer and either

---

[1]Breschi and Lissoni (2005, 2009) show that knowledge sharing being localized is primarily explained by workers being geographically immobile.

[2]In a different context, Singh (2007) finds bidirectional knowledge flows between multinational companies (MNCs) and host country organisations.

[3]As our measure of R&D output is patent applications, we will use the word inventions rather than

a positive or no effect - depending on the new employer's type - on the level of invention by the old employer. This represents to our knowledge the first firm-level evidence showing that labor mobility increases the aggregate level of invention.

Our empirical findings derive from an extensive data set that combines patent applications by Danish firms to the European Patent Office (EPO) with matched employer–employee register data that contain an essentially complete record of mobility in the Danish labor market. This data set allows us to estimate the change in the number of inventions, measured by patent applications, associated with labor mobility. We consider the effect of worker mobility on the levels of invention of the two firms involved, the new and the old employer, and allow for bidirectional knowledge flows. Summing up these effects, we obtain the net effect of labor mobility on invention.

The point of departure of our empirical approach is a standard firm–level patent production function (Hall et al., 1986; Hausman et al., 1984) that maps the different types of labor, capital and other observed firm characteristics into patent counts. We focus on workers with a university degree in natural sciences, engineering and other technical fields and analyze the effects of mobility of such workers when employed in positions classified as using or producing knowledge at an advanced level.[4]

Our results are consistent with the previous literature showing that workers bring knowledge and skills that increase the level of invention by the new employer. Firms are divided into "patenting" and "non-patenting" types according to their stock of previous patent applications being positive or not. A worker coming from a patenting firm is, according to our estimates, associated with an increase in the number of patent applications by the new employer of 0.034. A worker coming from a non-patenting firm is associated with an increase of around two thirds of that (0.022).[5] The economic significance of those numbers can be assessed by relating them to an average number of 0.76 patents per year by firms that have patented prior to the beginning of our sample period.

Although the positive effect of hirings on innovation is well-established by now, a couple of things must be noted regarding these results. Firstly, most existing studies use patent files to trace mobility, and the subjects of the analysis are therefore by construction "star scientists" who invent repeatedly. Instead of this, we consider the effect of mobility by an average worker who possesses the formal qualifications necessary to perform R&D. Thus, the magnitude of the effects that we find should be interpreted in view of the fact that our definition of R&D workers is fairly inclusive. Furthermore, as we have a complete record

---

innovations when referring to our empirical results.

[4]According to the International Standard Classification of Occupations by the International Labour Organization (ILO).

[5]The marginal effects reported in the introduction are evaluated for the average firm with at least one pre-sample patent.

of mobility, we avoid the possible biases arising from unregistered moves.[6] Secondly, with the notable exception of Hoisl (2007), this is the only study on the effects of mobility on invention output rather than invention input as measured by citations in later patent applications of the new or the old employer.

An important novelty of this paper is that we are able to estimate the effects of leavers. We find that a worker leaving for a patenting firm is associated with an *increase* in the number of patent applications by the old employer of 0.016. A worker who left for a non-patenting firm has no significant concomitant effect on patenting. It is an interesting and somewhat surprising result that a worker leaving for a patenting firm is associated with an increase in the level of invention of their old firm. This suggests that reverse knowledge flows not only exist, as Agrawal et al. (2006) as well as Corredoira and Rosenkopf (2010) have shown looking at patent citations, but are sufficiently important to outweigh negative effects resulting from worker exit.

Our analysis provides strong support to the view that mobility of high-skilled workers stimulates the total level of invention of the firms involved in the move, and the effect is stronger if the firms are of the patenting types. Hence, mobility between two patenting firms has the largest marginal effect $(0.050 = 0.034 + 0.016)$ on patent counts. Mobility between two non-patenting firms has a slightly positive effect, albeit statistically insignificant.

The remainder of this paper is organized as follows. The following section reviews the related literature. Section 2 details the hypotheses tested and the theory underlying them. Section 3 describes the data and outlines the definitions used in the analysis. Section 4 characterizes our econometric approach and provides descriptive statistics. The main results are reported in Section 5, along with certain robustness checks. Section 6 concludes.

2 Related Literature

Several studies have used patent data to analyze the importance of mobility for innovation. Song et al. (2003) as well as Rosenkopf and Almeida (2003) show that firms are more likely to cite patents produced by the inventors' former employers in their patent applications, which is evidence of forward knowledge flows. Agrawal et al. (2006) as well as Corredoira and Rosenkopf (2010) demonstrate the existence of reverse knowledge flows by showing that the old employer is also more likely to cite the patents of the worker's new employer in its later patent applications. An important difference to these papers is that they measure the effects of knowledge sharing from the input side of the invention process - the knowledge that goes into the creation of new inventions - whereas we measure it from the output side. Hoisl (2007) is to our knowledge

---

[6]If patent files are used to track mobility, a move that did not result in a patent application at the new employer is unregistered.

3

the only other paper that measures the effects of labor mobility on invention output. She combines data on mobility from patent files with background information about the inventors gained from questionnaires. Her findings show that mobile inventors are on average more productive, and that mobility increases inventor productivity. However, as she does not measure the patent productivity of the previous employer, she cannot address the effect of mobility on the total level of invention.

In a study that investigates the effects of both labor mobility and membership in a large and partly publicly-sponsored research joint venture in nanotechnology, Cassiman et al. (2011) show that hiring workers who have previously been employed at this institution does not generally lead to significant changes in the innovativeness of the hiring firms. Instead, active participation in the research joint venture ("buying a spot in the lab") does indeed lead to economically and statistically significant positive effects on industrial innovation.

Another approach to documenting knowledge sharing through labor mobility has been to test different theoretical predictions. Kim and Marschke (2005) show that firms have a higher propensity to patent in regions with high labor mobility. This is consistent with a theory in which firms patent their inventions to prevent misappropriation by former employees. Human capital theory predicts that workers who acquire valuable knowledge on the job receive a wage premium but pay for this through an initial wage discount. Møen (2005) finds evidence of such a wage profile, whereas Maliranta et al. (2009) find that workers are not able to capitalize on the knowledge acquired as R&D workers.

Finally, Toivanen and Väänänen (2008) combine Finnish patent data with linked employer–employee data. They find a significant and potentially long–lasting wage premium for inventors of granted patents, indicating that these workers are perceived by firms as possessing valuable knowledge and skills. Again, our approach differs from this line of research as the aim is not to demonstrate that labor mobility leads to knowledge sharing among firms. Rather, we take this as a well-established fact and try instead to measure the importance of this phenomenon for invention.

3   Theoretical Motivation and Hypotheses

In this section, we derive the hypotheses that we subsequently test in the empirical analysis. We consider the effect of labor mobility on the level of invention and refer to the parties involved in a move as the "old employer", the "new employer" and the "worker". We denote knowledge that travels from the old to the new employer as "forward knowledge flow", as it moves in the same direction as the worker, and knowledge that travels in the opposite direction is denoted as "reverse knowledge flow". It is useful to make the distinction between two types of knowledge that a worker can possess, namely "explicit knowledge" and "intellectual human capital". Explicit knowledge is codifiable

and can be transmitted to others at low costs. It is thus non-rival in nature, and one firm using a piece of explicit knowledge does not prevent other firms from using it at the same time (Arrow, 1962; David, 1992). Following Zucker et al. (1998), intellectual human capital is defined as knowledge that is embedded in a worker's human capital and that — due to its complexity or tacitness — cannot be easily communicated and shared.

Most economic analyses of knowledge sharing through labor mobility have conceptualized workers as depositories of knowledge. A worker acquires explicit knowledge and intellectual human capital while working for the old employer and brings it to the new employer; see, e.g., Pakes and Nitzan (1983) as well as Kim and Marschke (2005). Labor mobility results therefore in a forward knowledge flow from the old to the new employer. The old employer experiences a loss of the worker's intellectual human capital as well as any explicit knowledge that is private to the worker.

Inter-organizational social networks are another important channel for the diffusion of knowledge (Coleman et al., 1957; Dahl and Pedersen, 2004; Powell et al., 1996), and labor mobility also increases knowledge sharing by changing the structure of social networks in which the firm is embedded. Applying a network perspective, firms can be seen as the nodes of the social network and the enduring social relationships between the workers as the ties. Assuming that a worker stays in contact with her former colleagues, labor mobility establishes new ties between two nodes of the network, the new and the old employer. Since relationships between former colleagues are likely to be characterized by trust and reciprocity, the ties represent channels through which ideas, problem solutions and other types of knowledge are exchanged (von Hippel, 1987). The new ties formed by the labor move may also reduce the distance to other nodes in the network, increasing both the new and the old employer's knowledge exchange with other firms in the network (Breschi and Lissoni, 2005; Singh, 2005).

The social ties provide the new employer with access to knowledge that is located at the old employer, or at other firms, and that the worker moving does not possess. The forward knowledge flow arising from labor mobility is therefore reinforced by the establishment of social ties. Interestingly, the social network perspective also introduces the possibility of reverse knowledge flows, because the new ties formed allow the employees of the old employer to tap into the knowledge available at the new employer (Agrawal et al., 2006; Corredoira and Rosenkopf, 2010).

To sum up, labor mobility results in an inflow of knowledge to the new employer. Since knowledge is a key input in the production of new inventions, we would therefore expect labor mobility to be associated with greater invention by the new employer. Furthermore, the forward knowledge flows should be greater the more new knowledge is available at the old employer. In the empirical analysis, invention performance is measured by the number of patent

applications.[7] Moreover, we proxy the amount of new knowledge available at a firm by dividing firms into two groups: "patenting firms" with a positive stock of previous patents and "non-patenting firms" without such previous patent applications. Measuring the theoretical variables in this way, the first hypothesis is as follows:

Hypothesis 1: *A. Labor mobility is associated with a higher number of patent applications made by the new employer.*
B. The effect is larger if the old employer is a patenting firm.

The net effect of labor mobility on the old employer's invention performance is *a priori* ambiguous, because the old employer loses the worker's intellectual human capital and any explicit knowledge private to the worker but receives a reverse knowledge flow. Still, arguing as above, we would expect the reverse knowledge flow to be larger when the worker leaves for a patenting firm than when she leaves for a non-patenting firm:

Hypothesis 2: *The effect of labor mobility on the number of patent applications made by the old employer is larger if the new employer is a patenting firm.*

In order to derive our hypothesis regarding the effect of labor mobility on the joint level of invention of the firms, it is useful to start from an assumption of a competitive and frictionless labor market. In such a market, labor mobility occurs if and only if it increases the joint profits of the firms involved; see, e.g., Fosfuri et al. (2001). The reason for this is that the firm that values the worker the most will offer the highest wage. Hence, mobility occurs if the gain of the potential new employer from hiring the worker outweighs the loss of the current employer from losing the worker.

There are two opposing effects of labor mobility on the joint profits of the firms (Combes and Duranton, 2006; Fosfuri and Rønde, 2004). Firstly, a certain amount of knowledge that was in the sole possession of one of the firms before mobility occurred is now shared. Therefore, the profits of the firm that loses knowledge exclusivity is reduced. It will face more competition for some commercial uses of the knowledge. Since competition destroys rents, the firm that receives the knowledge gains less from entering into these commercial uses than what is lost by the firm that loses knowledge exclusivity. This effect tends to reduce the joint profits of the firm and to prevent labor mobility from occurring. Secondly, firms have different R&D capabilities and strengths, and knowledge sharing increases the likelihood that a piece of knowledge will serve as an input in the creation of new knowledge and inventions (Bessen and Maskin, 2009).

---

[7]As part of a robustness check, we also weight patent applications by the number of forward citations that they receive in order to control for quality differences. We expect knowledge sharing to increase both the number and quality of the inventions made. Therefore, the hypotheses developed below also hold if the dependent variable is citation-weighted patent applications.

Indeed, this is the main argument as to why knowledge sharing fosters invention and sustained economic growth. Therefore, knowledge sharing through labor mobility has the potential to stimulate invention, be it in the form of more variety, higher success probability, or greater speed of invention. This second effect increases the joint profits of the firms and tends to facilitate labor mobility.

These arguments imply that labor mobility occurs if the positive effect from a higher level of invention is sufficiently strong enough to outweigh the negative effect from more competition. Put differently, an increase in the joint invention of the firms is a necessary condition for labor mobility to occur. Again, we expect the forward and the reverse knowledge flow to be greater if the old employer and the new employer, respectively, are patenting firms. This leads to our third and main hypothesis:

Hypothesis 3: *A. Labor mobility is associated with an increase in the total number of patent applications made by the new and the old employer.*
B. The effect is larger if the new employer and if the old employer are patenting firms.

Some comments relating to Hypothesis 3 must be noted. To begin with, note that mobility could possibly reduce invention if it would be very disruptive for the old employer, or if it would result in a less efficient allocation of the worker's intellectual human capital. However, labor mobility would then unambiguously reduce the firms' joint profits, and we should not observe labor mobility occurring. Secondly, the argument underlying the hypothesis does not rely on mobility occurring exactly if it maximizes the firms' joint profits, a result that is sensitive to, e.g., the introduction of labor market frictions. Rather, the central and robust argument is that mobility takes place if the negative effect from tougher competition is dominated by the positive effect on the joint level of invention. Finally, strong protection of intellectual property (IP) shields firms from product market competition that would otherwise arise from knowledge sharing by preventing imitation. Hence, we expect strong IP protection to increase labor turnover, because it reduces the downside of mobility. The argument behind Hypothesis 3 holds, however, irrespective of the strength of the IP protection.

4 Data

Data on all patent applications to the EPO that were filed for between 1978 and 2006 by at least one applicant with Danish residency constitute the core of our data set. This data was retrieved from EPO's "PATSTAT" database.[8] We consider patent applications up to and including 2004 in our analysis, since the

---

[8]For information on this data set, refer to http://www.epo.org/patents/patent-information/raw-data/test/product-14-24.html.

database for the years following this date is not complete. This includes 12,873 patent applications.

Patent applications are used rather than patent grants because the average grant time at the EPO of four to five years (Kaiser and Schneider, 2005) implies that a substantial number of patents applied for during the time period considered for estimation (2000–2004) would be lost if patent grants were used.[9] The "time stamp" of the patent applications is the "priority date", the date on which the invention was first filed for patent protection at the EPO or any national patent office.

The EPO data do not come with a unique firm identification number of the type used by Statistics Denmark, the provider of our firm–level and employee–level data. Therefore, we attached our EPO data — mostly manually — to Statistics Denmark's firm identifiers. We could assign firm identifiers to 11,280 patents. The unmatched applications primarily refer to firms that went out of business before 1999. In any case, the corresponding information would have been lost in our analysis, since our firm–level data starts in 1999.

These 11,280 patent applications were applied for by 2,448 unique patent applicants that are neither private independent inventors nor foreign firms. After having matched this data with our firm–level data, we are left with 11,031 patent applications accounted for by 2,278 unique firms.

Statistics Denmark provided us with firm registry data, most importantly the sectoral and regional affiliation of a firm and its book value of physical capital, and with registry data on employee characteristics, most importantly the end–of–November number of employees and their highest level of education.[10] We discard sectors with no EPO patent applications between 1978 and 2004. Sectors are defined according to the three–digit NACE Rev. 1 industrial classification level. Firms that did not file for an application at the EPO in a particular year are assigned a zero for the number of patent applications in that year. In a final step, we merge the firm–level data with employee–level data, which allows us to track the employment history of each worker across firms. We leave out firms that were founded during the estimation period 2000–2004. This last restriction is caused by the choice of estimation method. As further described in Section 5.2, we use a "pre–sample" estimator that requires information on firms' patenting behavior prior to 2000. Finally, we discard firms from the public sector, since their patenting behavior is likely to be very different from that of private firms.

When delimited according to these criteria, there are 349,595 observations on 93,725 firms among the population of private firms in Denmark. Our main

---

[9]There exists a reporting lag between the date of application and the date on which the application is published in the EPO database. This implies that not all patents applied for after 2004 are registered in the database at the time of data collection. We do not include such patents in order to avoid biases.

[10]As the firm affiliation of a worker is registered only once a year in November, we do not observe within–year mobility.

sample consists of observations on firms that employ at least one worker in an R&D–related occupation.[11] This data represents the focus of our attention, since firms with employees in R&D–related occupations are much more likely to patent than firms without workers of this nature. Of the total number of 2,861 patent applications during 2000–2004 that could be definitively assigned to a firm, 2,728 — or 95 percent — can be assigned to firms with positive R&D employment. By excluding firms with very little or no current R&D activity, we attempt to compare different varieties of apples rather than apples and oranges. Our main estimation results thus include 42,507 observations on 14,516 unique firms, and 2,728 patents over the period 2000–2004.

5   Empirical approach

This section describes the patent production function and outlines the econometric approach that we employ in estimating the relationships between worker mobility and firms' inventive output. We provide details on the treatment of unobserved firm heterogeneity and state dependence in the analysis. In addition to our worker mobility terms, our econometric specification also includes controls for firm size (total R&D employment, capital stock), sectoral affiliation (15 sectors), five different geographical regions, and time effects. We lag all explanatory variables except for the time, region and sector dummies by one year. As part of a robustness check, we also use a two–period lag.

5.1   Patent production function

We assume a Cobb–Douglas functional form as it is standard procedure within the literature (Hausman et al. 1984; Blundell et al. 1995). We differentiate between R different types of workers in terms of mobility and in terms of having (not) been employed by a patenting firm, or leaving a firm to join a patenting/not patenting firm. We denote each share of workers $s_r = L_r/L$, where L is the total R&D workforce. Our dependent variable is the total number of patent applications by firm i in a given year t which we denote by P. It is a count variable that either takes on a value of 0 or a positive integer. We will hence use count data models in the estimations which link the explanatory variables to the dependent count variable in an exponential (or log–linear) way:

$$P = \exp\Big(\ln(A) + \beta\ln(K) + \alpha\ln(L) + \sum_{r=2}^{R}\alpha_r s_r\Big). \qquad (1)$$

Labor category $r = 1$ constitutes our base category of R&D workers. Appendix A provides details on the derivation of Equation (1).

---

[11]We provide details on our definition of R&D–related occupations in Section 5.3.

The coefficients in Equation (1) do not directly translate into marginal effects as in a linear model. A positive (negative) coefficient on worker group r indicates that its estimated marginal contribution to patenting activity of firm i is greater (smaller) than that of the base category. We present both coefficient estimates and marginal effects in our empirical analysis.

5.2 Econometric model

Count data models

Equation (1) can be directly estimated using count data models. The most popular model for count data is the Poisson regression (Cameron and Trivedi 1998; Winkelmann 2008). The basic Poisson regression model assumes equality of the conditional mean of the dependent variable and its conditional variance. A common empirical observation in patent data is, however, that the conditional variance is greater than the conditional mean, which implies over–dispersion (Cincera 1997). The variance restriction does not have to be fulfilled to obtain consistent estimates of the parameters in the Poisson model, but more efficient estimates can be found by specifying a more flexible relationship between the mean and the variance. We hence consider a Negative Binomial (NegBin) model which does allow for over–dispersion.[12]

The pre–sample mean estimator

In panel data of the type we study, the presence of firm–specific variables that the econometrician does not observe but affect patenting activity like e.g. permanent differences in appropriability conditions of R&D investments or different technological opportunities is not uncommon.

In a series of papers, Blundell et al. (1995, 1999, 2002) propose the "pre–sample mean estimator" (PSME), a count data method used to directly account for unobserved permanent heterogeneity through additional historic information on the dependent variable. We shall apply their estimator to our data. The idea is to approximate firm-specific, time–invariant heterogeneity by using information on a firm's patenting behavior prior to the start of the estimation period. In our particular context, we possess data on any firm's patent activity from 1978 onwards, while our explanatory variables (allowing for lags) are observed after 1999 only.

Specifically, the Blundell et al. pre–sample mean estimator uses the average of the dependent variable over the pre–sample period as a proxy for the correlated fixed effect for each firm. Since a prominent feature of our data is an overall increase in the level of patenting during the pre–sample period from 1978 to 1999, we extend the estimator by normalizing a firm's number of patents in a pre–sample year by the total number of patents applied for during that year. We provide details on this normalization in Appendix B.

Many of the firms in our data never applied for a single patent. We again

---

[12]We implement this estimator in Stata by the nbreg procedure with the dispersion (constant) option.

follow Blundell et al. (1995, 1999) and include a dummy variable for firms having applied for at least one patent during the pre–sample period. This variable also acts as a remedy for the so–called "zero–inflation problem" that is common to many analyses of economic count data (Mullahy, 1997).[13] Our final specification of the fixed effect includes the dummy variable for pre–sample patenting and the natural logarithm of the trend-corrected number of pre–sample patents.[14]

State dependence

In addition to unobserved permanent differences between firms, it is also commonly argued that past patenting outcomes may have an effect on present patenting activities. Blundell et al. (1995), for example, argue that a firm's stock of past patents represent knowledge from which future patentable ideas can be derived. This would be evidence of positive state dependence in patenting. Crépon and Duguet (1997), on the contrary, point to the fact that applying for a patent in year t will prevent the firm from applying for the same patent in year $t + 1$, essentially diluting its stock of patentable ideas. This would suggest negative state dependence.

To capture state dependence, Blundell et al. introduce a measure of a firm's patenting history into the estimating equation, namely the discounted stock of patents which gives a lower weight to more distant periods. Due to the relative short time span of our estimation sample, this measure is highly correlated with the pre-sample mean of the number of patents. Like Crépon and Duguet (1997), we therefore use a dummy variable that indicates whether or not a firm patented at $t - 1$ as our control for state dependence.

5.3 Variable definitions

Definition of R&D workers

We define R&D workers as those workers within a firm who are likely to be engaged in R&D related tasks. Specifically, we apply two main criteria to identify the relevant group of workers.[15] The first criterion is that the person must hold a bachelor, master, or Ph.D. degree in technical or natural sciences, veterinary

---

[13]We prefer our approach over the alternative zero inflation model since our model belongs to the family of linear–exponential models which are shown to produce consistent estimates under a large set of circumstances (Winkelmann, Ch. 3, 2008). Further, Staub and Winkelmann (2009) demonstrate that zero–inflation models are, unlike the Poisson and the NegBin model, not robust to mis–specification of the data-generating process.

[14]The number of pre–sample patents is replaced by a small number if the prior patent count is zero.

[15]Additional and somewhat straightforward criteria are that the individual must not be retired, be between 20 and 75 years old and be employed by a Danish firm (since we only have data on Danish firms at our disposal).

and agricultural sciences, or health sciences.[16] This criterion originates from the idea that knowledge flows are mainly associated with the mobility of high–skilled workers. The definition corresponds closely to the findings of Kaiser (2006) who uses patent inventor survey (PATVAL) data to show that Danish inventors are likely to hold a bachelor's degree or higher. More precisely, 30.5 percent of the inventors hold a bachelor's degree, 40.8 percent a master's degree and 17.4 percent a Ph.D. degree. We intend to capture all persons who possess the formal skills necessary to perform R&D related activities within the firm.

Some high-skilled workers may, however, never perform actual R&D tasks. We introduce the additional criterion that a person's job function must involve use or production of knowledge at an advanced level. This information is included in our data by way of the International Standard Classification of Occupations (ISCO) code, which is prepared by the International Labour Organization.[17] At its first–digit level, ISCO classifies different occupations depending on their knowledge content. In particular, we can distinguish between "professionals" (level 2) and "technicians and associate professionals" (level 3).[18] Individuals are categorized in the former group if they work in a position in which they "increase the existing stock of knowledge, apply scientific or artistic concepts and theories, teach about the foregoing in a systematic manner, or engage in any combination of these three activities." We denote this group "R&D professionals". They will be the focus of our analysis of mobility as they are most likely to be directly involved in the creation of new knowledge. Individuals categorized as technicians and associate professionals occupy support positions that are more directed towards utilizing already existing knowledge. We refer to this group as "R&D support workers". As they are not directly engaged in developing new knowledge, they are not expected to be the main carriers of knowledge between firms. The share of support workers in a firm is therefore included in our model as a control variable only.

To summarize, we define R&D professionals as individuals who hold a technical or scientific degree and perform job functions with an advanced knowledge content. R&D support workers have similar formal skills but are currently employed in positions with less emphasis on the creation of new knowledge. Jointly, these two groups constitute the current stock of R&D workers within the firm.

Mobility

We next characterize different categories of R&D professionals according to their mobility status. We differentiate between three categories of joiners to firm i in

---

[16]The health sciences category includes many general practitioners and hospital doctors who *a priori* are not expected to perform R&D related activities. Most of these will be excluded in our estimations since we exclude the public sector.

[17]http://www.ilo.org/public/english/bureau/stat/isco/intro.htm

[18]We include R&D managers (ISCO 1237) in the group of professionals. The codes are very detailed but a change in the way individuals were classified in 2003 prevents us from using more narrowly defined occupations consistently over time.

year t: The main group, which we simply term *joiners*, are workers who have been employed at another firm l in year $t-1$ (and hence were not employed at firm i in year $t-1$). *Graduates* are workers who meet the criteria for being R&D professionals at time t and graduated between $t-1$ and t. *Other joiners* are workers whose job market status in year $t-1$ is unknown but who are employed by firm i in year t. *Stayers* are R&D professionals who are employed by firm i both at time $t-1$ and t. They constitute the base category in our model.

A final group of R&D professionals as distinguished by mobility status in our model is the group of *leavers*. They were employed at firm i in year $t-1$ and are now employed at a different firm j in year t. In order to test for social network effects as discussed in Section 3, we further differentiate this group according to the level of innovativeness of their new employer. We proxy the level of innovativeness by the previous patenting success of the firm. Specifically, we interact the leaver variable with a dummy variable for the new employer having a non-zero historic patent count.

Using the same logic, we introduce two categories of joiners to firm i according to the previous patenting status of the old employer l.

5.4   Descriptive statistics

Table 1 displays descriptive statistics for the variables involved in our estimations. It differentiates between observations (firm–years) related to firms with and without patents prior to the beginning of the estimation period. They are denoted as pre–sample patenters and non-patenters, respectively. Around 6.6 percent of the sample comprises observations on pre–sample patenters.

Firms in our data are selected from the general population of firms by having a positive number of R&D workers. Nevertheless, the average firm is fairly small, having about eight R&D employees and a capital stock of about 78 mio. DKK (the median, however, is 2.7 mio. DKK only).[19] The overall level of patenting is fairly low. The average firm applies for 0.06 patents per year within the sample period. As expected, there are pronounced differences between pre–sample patenters and non-patenters in terms of firm size and their within-sample patent output. While firms in the former group employ 39 R&D workers on average and produce 0.76 patent applications per year, the corresponding figures for firms without pre-sample patents are five workers and 0.02 applications . The composition of the R&D workforce in terms of professionals and support workers is fairly similar. Pre-sample patenters employ on average 42 percent of their R&D workforce in supporting positions against 46 percent for non-patenters.

The main focus of our analysis is the mobility of R&D professionals. The overall level of mobility is high, with the four groups of joiners (from patenting firms, from non-patenting firms, graduates, other joiners) constituting about 20 percent of the current year's employment of R&D professionals (joiners and

---

[19]One U.S. dollar corresponds roughly to 5.3 DKK.

stayers). Analogously, this also holds true for pre-sample patenters and non-patenters. A high mobility rate is in keeping with the finding that Danish labor mobility rates are among the highest in OECD conutries, as documented by Eriksson and Westergård-Nielsen (2009).

Whereas the overall level of mobility is comparable between pre-sample patenters and non-patenters, there are clear differences in terms of the origins of firm joiners and the destinations of leavers. More than one in every four firm joiners come from a patenting firm in the sample of pre-sample patenters: for pre-sample non-patenters, this figure is less than one in seven. A similar difference can be observed on the leaver side, with approximately one in three (one in seven) of the leavers from pre-sample patenters (non-patenters) going to a patenting firm.

Table 1 does not display pronounced differences between firms with and without pre–sample patents in terms of their regional distributions. The sectoral distributions, on the other hand, show that certain sectors such as chemicals (which includes biotech), machinery, electronics and instruments are slightly overrepresented in the sub-sample of pre-sample patenters.

Insert Table 1 about here.

Appendix C displays a correlation table for the variables involved in our estimations. It shows that our explanatory variables are moderately correlated. This is confirmed by a variance inflation factor of 1.86, which is well below the critical value of 10 (Besley et al. 1980).

6 Results

Table 2 presents our NegBin PSME estimation results. The table contains two sets of results. The first specification is without controlling for the patenting status of the old firm for joiners or the new firm for leavers. In the second specification, we differentiate both groups according to patenting status. This extension enables us to directly address the existence of social network effects and the net effect of mobility, as laid out in Section 3.[20]

6.1 Principal results

From the results contained in Table 2, we infer the validity of our main theoretical hypotheses regarding mobility. In doing so we control for other important

---

[20]For hypotheses with a definite sign prediction, such as Hypothesis 1, 2 and 3, we apply one-sided t-tests and present the respective p-values. If there is no sign prediction, we apply two-sided t-tests and report the corresponding p-values.

determinants of patenting, including the size and composition of the R&D work-force of the firm, its capital stock, sectoral and regional characteristics as well as unobserved heterogeneity and possible state dependence in patenting.

The results without differentiation by patenting status show that the share of R&D joiners in a firm is positively and statistically significantly related to patenting activity when compared to the base category of workers, R&D stayers. The same holds for R&D graduates. The effects for other R&D worker groups are statistically insignificantly different from that of R&D stayers.

The results for the differentiated specification show that it is both the joiners from patenting firms and from non–patenting firms who are positively related to patenting activity. We find a positive difference between the effects of moves that involve a patenting firm over moves from a non-patenting firm. A similar difference is found when differentiating by patenting status of the new firm for leavers. Simplifying the model to the specification with no differentiation by patenting status is strongly rejected (p-value 0.012), which is why we prefer the more general model in the following. There is still no statistically signifi-cant difference between either R&D support workers or joiners with unknown employment history compared to R&D stayers.

Our results regarding joiners from firms immediately lead to our first main conclusions. There is a positive effect on patenting of joiners in general (p-value 0.000, one-sided). Moreover, there also is an additional positive effect when joiners come from a previously patenting firm (p-value 0.048, one-sided). These findings support our Hypothesis 1.

Turning to the evidence on Hypothesis 2 regarding leaver effects, the coeffi-cient of leavers who have left for a firm that has previously patented is *positive* and even significantly larger than the effect of the base category of stayers (p-value 0.006). Leavers to non-patenting firms, on the other hand, have no significant effect on current patent output of their previous employer (p-value 0.193). We take the finding of a positive and significant (p-value 0.004, one-sided) difference between the effects of leavers to patenting and non-patenting firms as evidence for the existence of a positive reverse knowledge effect, thus confirming our Hypothesis 2.

Finally, Hypothesis 3 combines the joiner and leaver effects to determine the effect of mobility on the joint level of invention of the firms involved. There are four possible combinations depending on the patenting status of the old firm of joiners and the new firm of leavers, respectively. We determine the statistical significance of the joint effect for each case by adding the coefficients of the respective types of joiners and leavers. For the case where a non-patenting firm is involved on both sides, we estimate a joint invention effect that is positive albeit insignificant (p-value 0.305, one-sided). If the joiner comes from a patent-ing firm whereas the leaver goes to a non-patenter, the effect becomes larger and borderline significant (p-value 0.044, one-sided). For those cases where leavers move to a patenting firm, we find large positive and strongly significant effects (p-value 0.000, one-sided), irrespective of the origin of joiners. These results are

strongly suggestive of an overall positive joint effect of mobility on the level of invention of the firms involved. We find the positive joint effect to be strengthened if either of the firms involved have previously patented. These findings support our main Hypothesis 3.

To sum up, Table 2 shows that increases in the number of joiners from either type of firm, in the number of graduates and in the number of leavers who have joined a patenting firm, are positively correlated with own patenting activity. It also indicates the importance of both forward and reverse knowledge flows. We also find strong evidence for overall positive effects of labor mobility on the level of invention of the firm.

Insert Table 2 about here.

6.2   Marginal effects

To study the absolute magnitude of effects, we interpret our empirical findings in terms of marginal effects in Table 3. The marginal effect is the absolute change in the number of patents due to an increase in the number of workers from skill group r by one worker. It depends both on the coefficient estimates, the number of workers in each of the worker groups, and the number of patents. The marginal effects are larger the more patents are applied for per R&D worker (compare Appendix D). We evaluate the marginal effects at the averages across (i) all observations and (ii) firms with at least one pre–sample patent.

We find that an increase in the number of joiners from patenting firms by one is related to an increase in the number of patents by 0.015 when evaluated for the average of all observations in our sample. This seems like a relatively large effect when compared to the average number of patent applications per year in our sample of 0.064 (compare Table 1). However, the average number of joiners from patenting firms is 0.1 and as many as 79 percent of firms do not employ any worker of this type. Hiring one worker of this type hence implies a large change for most firms in our data. When evaluated for an average pre–sample patenter, the corresponding marginal effect of joiners from patenting firms is 0.034 patents per additional worker against a sample average of 0.76 patent applications per year.

Insert Table 3 about here.

We finally evaluate the joint effects of leavers to and joiners from firms in terms of marginal changes in Table 4. The idea here is to consider the thought experiment of substituting a worker who has left the firm in the previous period with one who joins the firm. This is somewhat similar to our discussion of Hypothesis 3 at the end of Section 6.1, with the main difference being that the section referred to relative sizes of coefficient estimates while the following paragraphs provide quantitative estimates for the mapping between exchanges of different types of mobile workers.

We calculate the net effect for joiners who join from patenting and from non–patenting firms in combination with leavers who have left for patenting firms or for non–patenting firms. Specifically, we add the marginal effect of a joiner and the marginal effect of a leaver. For example, the total effect of replacing one worker who leaves for a patenting firm by one joiner from a patenting firm is 0.015+0.007=0.022 patents (compare Table 3) across all firms. For firms with at least one pre–sample patent, the same type of substitution yields 0.05 additional patents.

For all leaver/joiner combinations, and for both the full sample and for pre-sample patenters, we find positive mobility effects: all types of switches between different types of labor lead to additional patents. The marginal effects are statistically significant except for the combination of joiners from and leavers to non-patenting firms.

Insert Table 4 about here.

6.3   Other results

Table 2 also shows that there is substantial positive state dependence in patenting activity. Firms with patenting activity in the previous period have a much higher probability to patent again in the current period. This may reflect sunk costs associated with learning about conducting successful research and, more practically, filing a successful patent application.

Our correction for unobserved heterogeneity also has a significantly positive impact. An increase in the number of pre–sample patents by one percent is associated with an increase in the number of current patents by around 0.3 percent. Both the capital stock and the number of R&D workers are positively associated with the number of patents in year t. The respective elasticities are 0.135 and 0.314.

The year dummies and the sector dummies are statistically highly significant, while the region dummies are statistically insignificant: this implies that there are no regional differences in patenting activity once it is controlled for state dependence, unobserved heterogeneity and other patent production factors.

6.4   Alternative interpretations

We find our empirical results to be consistent overall with the hypotheses that we forwarded in Section 3. This of course does not rule out alternative explanations. One potential concern would be that causality ran in the opposite direction of what is assumed in our theoretical discussion. On the side of the joiners, if ideas were perceived inside the firms, and then workers were hired to transform - more or less mechanically - the ideas into patentable inventions, a positive correlation between the number of joiners and firms' patenting could result without any flow of knowledge taking place between firms. Furthermore,

the workers hired to implement the ideas would acquire knowledge that they could share with their former colleagues, explaining the positive leaver effect. It must be noted that the leaver effect would still represent a knowledge flow that would increase the joint level of invention of the firms involved in the mobility event.

The equivalent interpretation on the side of the leavers seems less plausible. Suppose, say, that a firm reached the end of its technology life-cycle and stopped patenting while laying off R&D workers. This would result in a negative correlation between the number of leavers and patenting, not a positive one as found in our analysis.

In our theory development, we abstract from potential selection of R&D workers with different unobserved ability or human capital endowment into different types of firms. The main concern here is that firms with the best conditions for research may attract the best R&D workers, so-called "positive assortative matching". In order to assess the empirical importance of this argument, one could - in the spirit of using pre-sample patenting as a proxy of unobserved differences in firms' ability to innovate - consider the previous patenting activity of a firm to indicate its inherently unobservable "research conditions". Joiners from firms with past patenting activity would be on average of higher ability than joiners from firms with no past patenting activity. This could explain at least part of the difference that we observe between joiners from these two types of firms under this interpretation. A similar argument should then apply to the leavers' side as well, which would reinforce our finding of reverse knowledge flows. Leavers to firms with previous patenting activity would then be on average of higher ability and the firm would suffer a greater loss of human capital for this group than for leavers to firms with no previous patenting. Given that we find a positive effect on current patenting in the former case and a zero or slightly negative effect for the latter, the estimated difference would then constitute a lower bound on the effect of reverse knowledge flows. Hence, selection may upwardly (downwardly) bias our results on joiners (leavers) from (to) firms with previous patenting activity. Nevertheless, even if selection would account for the full difference between the two types of joiners, the sum of the joiner and the leaver effects remains positive, resulting in a positive effect of labor mobility on the joint level invention as posited by Hypothesis 3.

6.5   Robustness checks

We conduct two robustness checks. First, we weight our dependent variable, the count of the number of patents of firm i at time t, by the number of citations a patent receives. By doing this, we attempt to account for the value of patents. As a second check, we lag the labor mobility variables by two periods instead of one, since it may take more than one year for joiner and leaver effects to materialize.

Citation weights

The distribution of the economic and technological value of patents is heavily skewed in the sense that few patents have a very high value while the bulk of patents has very little value, as discussed, e.g., by Harhoff et al. (1999); Lanjouw et al. (1998), and Hall et al. (2005). Trajtenberg (1990) shows that there is a close relationship between the number of citations a patent receives ("forward citations") and the social value of the invention in the computer tomography industry. Thus, he suggests approximating value by patent forward citations since they capture the enormous heterogeneity in the "quality" or "importance" of patents. Like Trajtenberg (1990), we weight each patent by one plus the number of citations the patent received within a three–year period after the EPO publication. Our patent citation data stem from the "EPO/OECD patent citations database" which is available from the OECD (Webb et al., 2005) and covers the period 1978–2006.

The left-hand part of Appendix E presents estimation results with citation weights. The coefficient estimates differ only slightly between the citation–weighted results and results for the unweighted specifications in Table 2.

Two year lags

The right-hand part of Appendix E contains estimation results with two lags on the R&D worker share variables as well as on the total number of R&D workers. Using two lags instead of one lag only implies that we lose one year of data for each firm, which leaves us with 32,182 observations on 12,481 unique firms.

Considering two lags instead of one makes some of our estimated effects stronger, most notably the effect of joiners from patenting firms. The leaver effects to patenting and non-patenting firms become less positive and more negative, respectively. The joint mobility effect for joiners from and leavers to non-patenters is now slightly negative, although insignificantly so.

All qualitative results including our conclusions on the three main hypotheses therefore remain when we extend the model using longer lags. The results also further reinforce our argument that reverse causality may not be a major concern here.

7   Conclusions

This paper assesses the quantitative importance of inter-firm mobility of labor for invention, using a data set that combines patent applications by Danish firms to the European Patent Office with matched employer–employee register data that track the employment history of R&D workers across time. We estimate the effect of labor mobility on the inventive activity by both the new and the old employer.

For the average firm with at least one patent prior to 1999, when our data set begins, a worker joining from a patenting firm is associated with an increase in the number of patent applications of 0.034 by the new employer, and a

worker coming from a non-patenting firm is associated with an increase of 0.022 patent applications. Turning to worker exit, we find that a worker leaving for a patenting firm is associated with an increase in the number of patent applications of 0.016 by the old employer. We explain this by mobility creating new ties in the social network: employees who have left stay in touch with their former colleagues and exchange information about their present research. There is no significant effect on invention by the old employer associated with a worker leaving for a non-patenting firm.

Summing up these effects, we find that labor mobility is related to a positive and both economically and statistically significant increase in the joint invention by the old and the new employer. The effect on joint invention is strongest for mobility between two patenting firms (0.05) and insignificantly positive for mobility between two non-patenting firms (0.011). These are notable results that provide a "missing link" in the literature between the firm-level and regional effects of labor mobility on invention.

Saxenian (1994) has forcefully argued that "job-hopping" is the key to the success of Silicon Valley by spurring knowledge sharing, innovation and, thus, the competitiveness of the local firms. Nevertheless, the reasons for the unusual high labor turnover in Silicon Valley remain controversial. Saxenian attributes it to an open business culture whereas Gilson (1999) stresses the weak enforcement of covenants not to compete in California. Marx et al. (2009) use an - apparently inadvertent - reform of the legal enforcement of covenants not to compete in Michigan as a natural experiment to study the effect of such covenants on labor mobility. In line with the arguments of Gilson, Marx et al. find that covenants not to compete do indeed restrict labor mobility. Our results would suggest that the Michigan reform is likely to have reduced invention, at least in the short-term. At the same time, our results add to this debate by showing that covenants not to compete are likely not only to reduce the aggregate level of invention, but also the invention of the firms imposing them. The reason for this is that firms restricting the outward mobility of their workers benefit less from reverse knowledge flows, suggesting that covenants not to compete should be used selectively and to protect key knowledge only. Franco and Mitchell (2008), as well as Kräkel and Sliwka (2009), argue that covenants not to compete may in fact increase innovation by stimulating firms' R&D investments - an issue that touches upon a limitation of our work. In the empirical analysis, firms are compared that are observationally equivalent, except with respect to the mobility of their workforce. Since these firms are likely to face very similar labor market conditions, it is not possible to identify how differences in the general level of labor turnover - e.g., due to different enforcement of covenants not to compete - affect firms' R&D investments. Here, there are likely to be opposing effects: high mobility increases appropriability hazards (reducing the return on R&D investment) but facilitates access to knowledge inputs (increasing the return on R&D investment). Interestingly, Samila and Sorensen (2011) find that the former effect dominates and that covenants not to compete tend to

impede entrepreneurship and employment growth.

There are several interesting research questions that could be pursued in future works. Our analysis looks at all workers who possess sufficient formal qualifications (and job descriptions) to perform R&D tasks, whereas most other studies have looked at highly productive inventors in the semiconductor industry. Several countries have register data similar to the data that we use, and it would be very interesting to link inventors to register data in order to track the mobility of inventors and to measure their importance for invention relative to other workers. We have only looked at mobility between private firms in this analysis, but another interesting question is how the mobility between university and private firms affects the knowledge production in these two sectors, as measured by patents and academic publications.

References

Agrawal, A. I. Cockburn and J. McHale (2006), Gone but not forgotten: knowledge flows, labor mobility, and enduring social relationships, Journal of Economic Geography 6, 571–591.

Almeida, P. and B. Kogut (1999), Localization of knowledge and the mobility of engineers in regional networks, Management Science 45, 905–916.

Arrow, K.J. (1962), Economic welfare and the allocation of resources for innovation. In: Nelson, R.R. Editor, The rate and direction of inventive activity, Princeton University Press, Princeton, NJ, 609–626.

Besley, D.A., E. Kuh and R.E. Walsh (1980), Regression diagnostics: identifying influential data and sources of collinearity, New York: John Wiley and Sons.

Bessen, J. and E. Maskin (2009), Sequential innovation, patents, and imitation, RAND Journal of Economics 40, 611–635.

Blundell, R., R. Griffith and J. van Reenen (1995), Dynamic count data models of technological innovation, The Economic Journal 105, 333–344.

Blundell, R., R. Griffith and J. van Reenen (1999), Market share, market value and innovation in a panel of British manufacturing firms, Review of Economic Studies 66, 529–554.

Blundell, R., R. Griffith and F. Windmeijer (2002), Individual effects and dynamics in count data models, Journal of Econometrics 108, 113–131.

Boeker, W. (1997), Strategic change: the influence of managerial characteristics and organizational growth, The Academy of Management Journal 40, 152–170.

Breschi, S. and F. Lissoni (2005), Cross-firm inventors and social networks: localized knowledge spillovers revisited, Annales d'économie et de statistique 79/80, 189–209.

Breschi, S. and F. Lissoni (2009), Mobility of skilled workers and co-invention networks: an anatomy of localized knowledge flows, Journal of Economic Geography 9(4), 439–468.

Cameron, C. and P. Trivedi (1998), The analysis of count data, Cambridge University Press, New York.

Cassiman, B., R. Veugelers and S. Arts (2011), Tracing the effect of links between science and industry: the role of researcher interaction and mobility between firms and research organizations, IESE working paper.

Cincera, M. (1997), Patents, R&D, and technological spillovers at the firm level: some evidence from econometric count models for panel data, Journal of Applied Econometrics 12, 265–280.

Coleman, J., E. Katz and H. Menzel (1957), The diffusion of an innovation among physicians 20, 253–270.

Combes, P.-P. and G. Duranton (2006), Labour pooling, labour poaching, and spatial clustering, Regional Science and Urban Economics 36, 1–28.

Corredoira, R. and L. Rosenkopf (2010), Should auld acquaintance be Forgot? The reverse transfer of knowledge through mobility ties, Strategic Management Journal 31, 159–181.

Crépon, B. and E. Duguet (1997), Estimating the innovation function from patent numbers: GMM on count panel data, Journal of Applied Econometrics 12, 243–263.

Dahl, M.S. and C.Ø.R. Pedersen (2004), Knowledge flows through informal contacts in industrial clusters: myth or reality?, Research Policy 33, 1673–1686.

David, P. A. (1992), Knowledge, property, and the system dynamics of technological change. Proceedings of the World Bank Annual Conference on Development Economics, 215–248.

Eriksson, T. and N. Westergård-Nielsen (2009), Wage and Labor Mobility in Denmark, 1980-2000, in Lazear, E. and K. Shaw (eds): The Structure of Wages, National Bureau of Economic Research, Cambridge/Mass.

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

Fosfuri, A., M. Motta and T. Rønde (2001), Foreign direct investments and spillovers through workers' mobility, Journal of International Economics 53, 205–222.

Fosfuri, A. and T. Rønde (2004), High-tech clusters, technology spillovers, and trade secret laws, International Journal of Industrial Organization 22, 45–65.

Franco, A. M. and M. F. Mitchell (2008), Covenants not to compete, labor mobility, and industry dynamics, Journal of Economics and Management Strategy 17, 581–606.

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, 575–629.

Hall, B. H., J. A. Hausman and Z. Griliches (1986), Patents and R&D: is there a lag?, International Economic Review 27, 265–283.

Hall B. H., A. B. Jaffe and M. Trajtenberg (2005), Market value and patent citations, RAND Journal of Economics 36, 16–38.

Harhoff, D., F. Narin, F.M. Scherer, and K. Vopel (1999), Citation frequency and the value of patented inventions, Review of Economics and Statistics 81, 511–515.

Hausman, J.A., B.H. Hall and Z. Griliches (1984), Econometric models for count data with an application to the patents–R&D relationship, Econometrica 47, 909–938.

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

Hoti, S., M. McAleer, D. Slottje (2006), Intellectual property litigation in the USA, Journal of Economic Surveys 20, 715–729.

Kaiser, U. (2006), The Value of Danish Patents — Evidence From a Survey of Inventors, Centre for Economic and Business Research Discussion Paper 2006-01.

Kaiser, U. and C. Schneider (2005), The CEBR matched Patent–employer–employee data set, Centre for Economic and Business Research mimeo, Copenhagen.

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

Kräkel, M. and D. Sliwka (2009), Should you allow your employee to become your competitor? On noncompete agreements in employment contracts, International Economic Review 50, 117–141.

Lanjouw, J., A. Pakes and J. Putnam (1998), How to count patents and the value of intellectual property: the use of patent renewal and application data, Journal of Industrial Economics 46, 405–432.

Maliranta, M., P. Mohnen and P. Rouvinen (2009), Is inter-firm labor mobility a channel of knowledge spillovers? Evidence from a linked employer–employee panel, Industrial and Corporate Change 18, 1161–1191.

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

Mansfield, E. (1985), How rapidly does new industrial technology leak out, The Journal of Industrial Economics 34, 217–223.

Møen, Jarle (2005), Is mobility of technical personnel a source of R&D spillovers?, Journal of Labor Economics 23, 81–114.

Mullahy, J. (1997), Heterogeneity, excess zeros, and the structure of count data models, Journal of Applied Econometrics 12, 337–350.

Pakes, A. and S. Nitzan (1983), Optimal contracts for research personnel, research employment and the establishment of 'rival' enterprises, Journal of Labor Economics 1, 345–365.

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

Powell, W. W., K. W. Koput, and L. Smith-Doerr (1996), Interorganizational collaboration and the locus of innovation: networks of learning in biotechnology, Administrative Science Quarterly, 41, 116–145.

Rao, H. and R. Drazin (2002), Overcoming resource constraints on product innovation by recruiting talent from rivals: a study of the mutual fund industry, 1986-94, Academy of Management Journal 45, 491–507.

Rosenkopf, L. and P. Almeida (2003), Overcoming local search through alliances and mobility, Management Science 49, 751–766.

Samila S. and O. Sorensen (2011), Noncompete covenants: incentives to innovate or impediments to growth, forthcoming Management Science.

Saxenian, A. (1994), Regional advantage: culture and competition in Silicon Valley and Route 128. Cambridge, Massachusetts: Harvard University Press.

Singh, J. (2005), Collaborative networks as determinants of knowledge diffusion patterns, Management Science 51, 756–770.

Singh, J. (2007), Asymmetry of knowledge spillovers between MNCs and host country firms, Journal of International Business Studies 38, 764–786.

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

Song, J., P. Almeida, and G. Wu (2003), Learning–by–hiring: when is mobility more likely to facilitate interfirm knowledge transfer?, Management Science 49, 351–365.

Staub, K. E. and R. Winkelmann (2009), Robust estimation of zero–inflated count models, University of Zurich, Working Paper No. 0908.

Toivanen, O. and L. Väänänen (2008), Returns to inventors, Helsinki School of Economics mimeo.

Trajtenberg, M. (1990), A penny for your quotes: patent citations and the value of innovations, The Rand Journal of Economics 21, 172–187.

von Hippel, E. (1987), Cooperation between rivals: informal know-how trading, Research Policy 16, 291–302.

Webb, C., H. Dernis, D. Harhoff and K. Hoisl (2005), Analysing European and international patent citations: a set of EPO patent database building blocks, STI Working Paper 2005/9, OECD.

Wezel, F. C., G. Cattani and J. M. Pennings (2006), Competitive implications of interfirm mobility, Organization Science 17, 691–709.

Winkelmann, R. (2008), Econometric analysis of count Data, Heidelberg, New York: Springer.

Zander, U. and B. Kogut (1995), Knowledge and the speed of the transfer and imitation of organizational capabilities: an empirical test, Organization Science 6, 76–92.

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

	All observations		Observations with pre–sample patent		Observations without pre–sample patent	
	Mean	**SD**	**Mean**	**SD**	**Mean**	**SD**
Dependent variable						
# patent applications t	0.064	1.143	0.761	3.900	0.015	0.535
Dummy patent $t-1$	0.019	—	0.209	—	0.005	—
R&D worker shares						
Joiners from patenting firms	0.012	0.082	0.020	0.078	0.011	0.082
Joiners from non–patenting firms	0.067	0.205	0.056	0.138	0.068	0.209
Graduates	0.022	0.115	0.021	0.078	0.022	0.117
Other joiners	0.026	0.137	0.017	0.086	0.027	0.140
Stayers	0.416	0.422	0.463	0.323	0.412	0.428
Support workers	0.458	0.441	0.423	0.337	0.460	0.447
Leavers to patenting firms	0.015	0.091	0.032	0.110	0.013	0.090
Leavers to non–patenting firms	0.077	0.245	0.076	0.188	0.077	0.249
Capital and R&D labor						
Total R&D workers	7.568	44.523	38.895	140.269	5.349	25.602
Capital stock (in mio. DKK)	78	1,280	399	2,520	55	1,140
Year dummies (base: 2000)						
2001	0.203	—	0.206	—	0.202	—
2002	0.196	—	0.202	—	0.196	—
2003	0.187	—	0.190	—	0.187	—
2004	0.183	—	0.181	—	0.183	—
Sector dummies (base: wholesale and retail trade)					—	—
Farm & food	0.016	—	0.019	—	0.016	—
Textiles & paper	0.041	—	0.036	—	0.041	—
Chemicals	0.014	—	0.054	—	0.011	—
Plastic & glass	0.026	—	0.072	—	0.023	—
Metals	0.049	—	0.084	—	0.047	—
Machinery	0.069	—	0.233	—	0.057	—
Electronics	0.030	—	0.067	—	0.028	—
Instruments	0.018	—	0.063	—	0.015	—
Vehicles	0.007	—	0.021	—	0.006	—
Furniture	0.016	—	0.021	—	0.016	—
IT & telecom	0.070	—	0.035	—	0.072	—
Technical services	0.140	—	0.127	—	0.141	—
Business–related services	0.095	—	0.044	—	0.099	—
Other	0.180	—	0.023	—	0.191	—
Region dummies (base: Capital region)					—	—
Zealand	0.097	—	0.088	—	0.098	—
Southern	0.224	—	0.237	—	0.223	—
Central	0.207	—	0.196	—	0.208	—
Northern	0.074	—	0.090	—	0.073	—
Pre–sample variables						
# pre–sample patents	0.031	0.557	0.468	2.120	0	0
Dummy pre–sample patents	0.066	—	1	—	0	—
# obs.	42,507		2,811		39,696	

Table 1 displays descriptive statistics for the entire set of observations, for observations with a pre–sample patent and for those without a pre–sample patent. "SD" denotes the respective standard deviation.

Table 2: NegBin pre-sample mean estimation results

	Without controls for patenting status of old/new employer		With controls for patenting status of old/new employer	
	Coefficient	p–value	Coefficient	p–value
R&D worker shares				
Joiners	1.010	0.000	—	—
... from patenting firms	—	—	1.465	0.000
... from non–patenting firms	—	—	0.852	0.002
Graduates	1.084	0.004	1.101	0.003
Other joiners	0.673	0.122	0.685	0.115
Support workers	-0.119	0.572	-0.133	0.526
Leavers	-0.082	0.781	—	—
... to patenting firms	—	—	0.808	0.006
... to non–patenting firms	—	—	-0.588	0.193
Capital and total R&D workers				
ln(R&D workers)	0.307	0.000	0.314	0.000
ln(capital stock)	0.135	0.000	0.135	0.000
Lagged dependent and pre–sample variables				
Dummy patent $t-1$	1.493	0.000	1.482	0.000
$ln(FE)$	0.321	0.000	0.315	0.000
FE dummy	0.340	0.145	0.334	0.150
Tests for joint significance	**Test statistic**	p–**value**	**Test statistic**	p–**value**
Worker shares	30.55	0.000	46.97	0.000
Year dummies	9.77	0.044	9.97	0.041
Sector dummies	75.85	0.000	75.48	0.000
Region dummies	2.17	0.705	2.03	0.730
Pre–sample variables	55.54	0.000	52.69	0.000
Pseudo R^2	0.3581		0.3590	

Table 2 displays NegBin PSME regression results. The estimations involve 42,507 observations on 14,516 unique firms, 2,728 patents and 390 unique patenting firms. Patent citation weights have not been applied. p-values are based on clustered standard errors.

Table 3: Marginal effects

	...all firms		...firms with pre–sample patent	
	ME	*p*–**value**	**ME**	*p*–**value**
Joiners from patenting firms	0.015	0.000	0.034	0.000
Joiners from non–patenting firms	0.010	0.000	0.022	0.000
Graduates	0.012	0.000	0.027	0.000
Other joiners	0.008	0.020	0.019	0.023
Support workers	0.001	0.180	0.003	0.245
Stayers	0.002	0.018	0.005	0.018
Leavers to patenting firms	0.007	0.006	0.016	0.006
Leavers to non–patenting firms	-0.005	0.193	-0.011	0.193

Table 3 displays the marginal effects of different types of workers and across alternative types of firms. It is based on the estimation results for the specification with controls for patenting status of old/new employer presented in Table 2. Marginal effects are evaluated at the sample means of the involved variables. **Reading example:** one additional joiner from a patenting firm is related to 0.015 additional patents (marginal effect). The figure is statistically highly significant (*p*–value 0.000).

Table 4: Joint mobility effects

	Left for patenting firm		Left for non–patenting firm	
	ME	*p*–**value**	**ME**	*p*–**value**
All observations				
Joiners from patenting firms	0.022	0.000	0.010	0.012
Joiners from non–patenting firms	0.017	0.000	0.005	0.223
Observations with pre–sample patent				
Joiners from patenting firms	0.050	0.000	0.023	0.014
Joiners from non–patenting firms	0.038	0.000	0.011	0.251

Table 4 displays marginal effects for joint mobility under different scenarios. These calculations are based on the estimation results displayed in Table 2. **Reading example:** the replacement of one R&D worker who has left for a patenting firm by a joiner previously employed by a patenting firm is associated with an increase in the number of patents by 0.022 across all firms. For firms with at least one pre–sample patent, this relates to an increase in the number of patents by 0.050.

Appendix A: Production function derivation

The Cobb–Douglas function implicitly assumes that firms never have zero quantities of an input, because all inputs are multiplied. Dividing labor into different types makes this assumption unlikely to hold in the present sample with many small firms included. To be able to apply a Cobb–Douglas function without forcing output to zero for many firms, we need to assume perfect substitution between the different types of labor inputs. However, as long as we assume separability between a relative homogenous set of labor inputs, here R&D employment, we consider this assumption to be reasonable.

An additive specification of the quality–adjusted labor input assuming that R&D workers are perfectly substitutable inputs is given by:

$$
\begin{aligned}
QL &= L_{St} + \gamma_J L_J + \gamma_O L_O + \gamma_G L_G + \gamma_{Su} L_{Su} + \delta_P L_P \\
&= L \left(1 + (\gamma_J - 1) \frac{L_J}{L} + (\gamma_O - 1) \frac{L_O}{L} + (\gamma_G - 1) \frac{L_G}{L} + (\gamma_{Su} - 1) \frac{L_{Su}}{L} + \delta_X \frac{L_X}{L} \right)
\end{aligned}
$$

This includes both leaver effects and effects of workers presently in the firm. The worker groups are denoted St for *stayers*, J for *joiners* from firms, O for *other joiners*, Su for *support workers*, and X for *leavers*. It is straightforward to extend the model with terms that differentiate joiners and leavers according to patenting status.

The marginal productivity of L_{St} is normalized to unity in this equation. For the shares of R&D workers currently in the firm, the coefficients $(\gamma_r - 1)$, $r = J, O, G, Su$, measure the relative productivity differentials between an additional worker and an R&D stayer. For leavers, the coefficient δ_X measures their patent productivity relative to the effect of stayers.

Taking logs, using the approximation that $\ln(1 + z) \approx z$ for small z and substituting for QL, we obtain a log–linear specification of the patent production function:

$$
\begin{aligned}
P &= \exp \Big[\ln(A) + \beta \ln(K) + \alpha \ln(L) + \alpha (\gamma_J - 1) \frac{L_J}{L} \\
&\quad + \alpha (\gamma_O - 1) \frac{L_O}{L} + \alpha (\gamma_G - 1) \frac{L_G}{L} + \alpha (\gamma_{Su} - 1) \frac{L_{Su}}{L} + \alpha \delta_X \frac{L_X}{L} \Big]
\end{aligned}
$$

In a final step to derive estimating equation (1), we define $\alpha_r = \alpha(\gamma_r - 1)$, $r = J, O, G, Su$, and $s_r = L_r/L$ for the shares of R&D workers in the firm. Similarly for leavers, $\alpha_X = \alpha \delta_X$, and s_X denotes the ratio of leavers to current employment.

Appendix B: Trend Correction of the PSME

We extend the pre-sample mean estimator of Blundell et al. (1995) in allowing for aggregate trends.

Let T_P denote the number of pre–sample observations on the dependent variable. P_{it} denotes the number of patent applications. We define our *weighted* proxy variable for firm–specific fixed effects, FE_i, by

$$FE_i \;\; = \;\; \frac{1}{T_P} \sum_{t=1}^{T_P} \frac{P_{it}}{\sum_{j \in A} P_{jt}}$$

where A denotes the set of all firms potentially applying for patents.

Essentially, by using this extension we allow for general business cycle effects, the general propensity of firms to patent (vs. secrecy), the propensity of Danish firms to patent at the EPO, etc., to vary over time.

Appendix C: Correlation table

	J. from pat. firms	J. from non-pat. firms	Other joiners	Grad.	Supp.
Joiners from pat. firms	1				
Joiners from non–pat. firms	-0.019	1			
Other joiners	-0.015	-0.035	1		
Graduates	-0.008	-0.024	-0.018	1	
Support workers	-0.105	-0.266	-0.169	-0.158	1
Leavers to pat. firms	0.065	0.032	0.013	0.034	-0.082
Leavers to non–pat. firms	0.019	0.088	0.051	0.061	-0.173
ln(total R&D workers)	0.029	-0.011	-0.042	0.012	-0.078
ln(cap. stock)	0.014	-0.031	-0.053	-0.027	0.100
Dummy patent $t-1$	0.038	0.001	-0.012	0.011	-0.029
$ln(FE)$	0.030	-0.011	-0.016	0.002	-0.028
FE dummy	0.027	-0.014	-0.018	-0.001	-0.021

	L. to pat. firms	L. to non pat. firms	ln(total R&D workers)	ln(cap. stock)	Lag patent	ln
Joiners from pat. firms						
Joiners from non–pat. firms						
Other joiners						
Graduates						
Support workers						
Leavers to pat. firms	1					
Leavers to non–pat. firms	0.049	1				
ln(total R&D workers)	0.057	0.035	1			
ln(cap. stock)	0.022	-0.008	0.395	1		
Dummy patent $t-1$	0.033	-0.009	0.267	0.151	1	
$ln(FE)$	0.054	-0.003	0.358	0.226	0.4678	
FE dummy	0.052	-0.001	0.320	0.218	0.375	0

Appendix D: Marginal effects

The marginal effect — the absolute change in the number of patents related to a change in the number of workers of group r by one — is obtained by partially differentiating the patent production function (1) with respect to worker group r.

For the workers presently in the firm (except stayers), namely groups J, O, G and Su, the marginal effects are:

$$\frac{\partial E[P]}{\partial L_r} = \frac{E[P]}{L}(\alpha + \alpha_r - \sum_{j \in J,O,G,Su} \alpha_j s_j - \alpha_X s_X).$$

For stayers, St, the same expression applies noting that $\alpha_{St} = 0$ by the normalization adopted here.

For the leavers, there are no effects via total employment, and we thus have the following simple expression for the marginal effects:

$$\frac{\partial E[P]}{\partial L_X} = \frac{E[P]}{L}\alpha_X$$

Appendix E: Robustness checks

	Citations-weighted		Two-time lagged variables	
	Coefficient	*p*–value	**Coefficient**	*p*–value
R&D worker shares				
Joiners from patenting firms	1.606	0.000	2.108	0.000
Joiners from non–pat. firms	0.917	0.001	0.746	0.020
Graduates	1.116	0.003	1.044	0.007
Other joiners	0.712	0.103	1.121	0.010
Support workers	-0.096	0.634	0.158	0.495
Leavers to patenting firms	0.763	0.011	0.685	0.085
Leavers to non–patenting firms	-0.627	0.168	-0.813	0.074
Capital and R&D labor				
ln(R&D workers)	0.299	0.000	0.319	0.000
ln(capital stock)	0.140	0.000	0.137	0.001
Lagged dependent and pre–sample variables				
Dummy patent $t-1$	1.433	0.000	1.571	0.000
$ln(FE)$	0.339	0.000	0.305	0.000
FE dummy	0.361	0.161	0.466	0.095
Pseudo R^2	0.328		0.370	
Number of observations	42,507		32,182	
Number of firms	14,516		12,481	

The table displays NegBin PSME regression results for the number of patent applications weighted by citations within three years (left-hand part of table) and two-time lagged variables (right-hand part of table). The specification estimated is otherwise identical to the one in the main results table, Table 2.