Department of Business Administration

UZH Business Working Paper Series

Working Paper No. 336

Does the Mobility of R & D Labor Increase Innovation?

Ulrich Kaiser, Hans Christian Kongsted, Thomas Ronde

June 2013

University of Zurich, Plattenstrasse 14, CH-8053 Zurich,
http://www.business.uzh.ch/forschung/wps.html

Does the Mobility of R&D Labor Increase Innovation?

Ulrich Kaiser* Hans Christian Kongsted † Thomas Rønde‡

June 25, 2013

Abstract

We investigate the effect of mobility of highly skilled workers in Denmark on the total patenting activity of the firms involved for the population of R&D active Danish firms observed between 1999 and 2004. Our study documents how workers joining increase firms' patenting activity. The effect is strongest if workers join from patent-active firms. We also find evidence of a positive feedback effect on patenting from workers who have left for another patent-active firm. Summing up the effects of joining and leaving workers, we show that labor mobility increases the total innovative activity of the new and the old employer. Our study thus provides firm-level support for the notion that labor mobility stimulates overall innovation of a country or region.

Keywords: labor mobility, innovation, research and development, patent.

*University of Zurich, Department of Business Administration, Plattenstr. 14, 8032 Zurich, Switzerland, ulrich.kaiser@business.uzh.ch; Copenhagen Business School, Department of Innovation and Organizational Economics and Institute for the Study of Labor, Bonn.

†Copenhagen Business School, Department of Innovation and Organizational Economics, Kilevej 14A, 2000 Frederiksberg, Denmark, hck.ino@cbs.dk, and Centre for Applied Microeconometrics, University of Copenhagen.

‡Copenhagen Business School, Department of Innovation and Organizational Economics, Kilevej 14A, 2000 Frederiksberg, Denmark, thr.ino@cbs.dk; Centre for European Economic Research, Mannheim, and Centre for Economic Policy Research, London.

1 Introduction

Knowledge transfer resulting from labor mobility constitutes an important source of innovation and growth. One stream of literature focuses on the effects of mobility on the innovation activities of individual firms—typically, the firm hiring the worker—and explores many contingencies that moderate this relationship (Boeker, 1997; Palomeras and Melero, 2009). From a more macro-perspective, studies within economics, management strategy, and economic geography argue that a high level of labor turnover spurs regional innovation performance (Almeida and Kogut, 1999; Fallick et al., 2006). In this paper, we integrate these levels of analysis to study whether labor mobility increases the total R&D output of the firms involved. In other words, we investigate whether the notion that labor mobility stimulates overall innovation has a firm-level micro-foundation.

Various evidence from surveys, patent files, and court cases, shows that labor mobility is an important channel of knowledge transfer between firms (Almeida and Kogut, 1999; Hoti et al., 2006; Mansfield, 1985). The implications of this observation for firm strategy have been extensively investigated. The findings show that firms exploit the knowledge and skills brought by new recruits, to enter distant technological areas (Palomeras and Melero, 2009; Rosenkopf and Almeida, 2003; Tzabbar, 2009), to introduce new types of products (Boeker, 1997; Rao and Drazin, 2002), and to boost R&D output (Ejsing et al., 2013). Another strand of the literature adopts the perspective of worker exits and how firm performance suffers from interruption to routines and loss of knowledge (Campbell et al., 2012; Wezel et al., 2006). However, the departure of a worker also represents an opportunity for the firm to access the knowledge available at the worker's new employer (Agrawal et al., 2006; Corredoira and Rosenkopf, 2010).

At the regional level, localized knowledge sharing has long been recognized as a major benefit of agglomeration (Jacobs, 1969; Marshall, 1920). Saxenian (1994) pointed to the particular importance of labor mobility for regional innovation performance. In her comparative analysis of Silicon Valley in California and Route 128 in Massachusetts, she argued that the "job-hopping" culture in Silicon Valley creates tightly coupled social networks through which knowledge flows, causing rapid innovation and growth in that region. Consistent with this view, subsequent studies document the co-existence of high labor turnover and localized knowledge sharing among firms in the semiconductor industry in Silicon Valley (Almeida and Kogut, 1999; Breschi and Lissoni, 2005; Fallick et al., 2006). The importance of labor mobility is underlined also by more recent studies showing that regions characterized by strong enforcement of trade secrecy laws and covenants not to compete experience lower rates of labor turnover but also less patenting and entrepreneurship (Marx et al., 2009; Png, 2012; Samila and Sorenson, 2011).

Although the literature documents the importance of labor mobility for innovation in firms and in regions, to our knowledge, there are no empirical studies providing quantitative evidence of the critical link between these levels of analysis. Existing firm-level studies look at the effect of labor mobility on the innovation activities of either the new or the old employer, and in regional-level studies the central assumption is that labor mobility increases total innovation for the firms involved in the labor exchanges. It is not a priori clear whether this condition holds since worker mobility may hurt the old employer more than it benefits the new employer. In this paper, we investigate the effect of mobility of highly skilled workers in Denmark on the total patenting activity of the firms involved. We find a significant and positive effect of labor mobility on total patenting if either of the firms involved has a history of patenting activity. Mobility from and to patenting firms, that is, firms with a positive stock of existing patents, has the largest marginal effect, namely 0.023 additional patent applications in the following year per mobile worker. For the average firm in our sample, this implies a 36 percent increase in the number of patent applications.

Our empirical findings derive from an extensive data set that combines patent applications from Danish firms to the European Patent Office (EPO) with matched employer-employee registry data which essentially contain a complete record of mobility in the Danish labor market. This contrasts with most existing studies which trace mobility via patent files, meaning that mobility is observed only if an inventor applies for a patent also at the new firm. We differentiate among workers who joined the firm in the focal year ("joiners"), workers who have been in the firm since the previous year ("stayers"), and workers who left in the previous year ("leavers"). The focus is on R&D workers who we define as individuals i) holding a university degree in natural sciences, engineering and other technical fields, and ii) employed in positions classified as using or producing knowledge at an advanced level. 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 different types of labor, capital, and other observed firm characteristics onto patent counts. To control for unobserved permanent differences in firms' patent productivity, we employ two different estimation approaches: The pre-sample mean estimator (Blundell et al. 1995) and the dynamic fixed-effect GMM estimator (Blundell et al. 2002).

We find that a joiner coming from a patenting firm is associated with a significant increase in the number of the new employer's patent applications. In relative terms, the joiner in this case contributes six times more than a comparable stayer, while a joiner from a firm with no patent activity is no more productive than a stayer. Our interpretation of this result is that, on average patent-active firms are more R&D active, and thus constitute a richer source of valuable knowledge.

The relative magnitude of the contribution of joiners from patenting firms is sizable given that we consider the effect of mobility of an average worker occupying a job function and in possession of the formal qualifications required to conduct R&D, whereas most existing studies consider "star scientists" with at least two patented inventions. In the case of leavers, we find that a worker who left to join a patenting firm, is associated with a significant increase in the present number of patent applications by the previous employer. In the case of a worker who left to join a non-patenting firm there is no significant concomitant effect on patenting. To sum up, in relation to the effects of joiners and leavers, our analysis provides strong support for the view that mobility of high-skilled workers stimulates total firm-level invention conditional on at least one of the firms having been patent active in the past.

The analysis of the effect of mobility on total invention relies on two important advantages of our data set. First, compared to most previous studies that use patent data to track inventor mobility, we have a complete record of labor mobility. In addition to avoiding possible biases arising from unregistered moves, this allows us to estimate the patent productivity of workers joining a firm, leaving, and staying with a firm. Second, our dependent variable—number of patent applications—lends itself more naturally to aggregation than other dependent variables which are employed in existing studies such as entry into a new technology class or product area. A potential weakness of our analysis is that it does not include variables for whether a mobility event resulted in actual knowledge transfer. This raises the concern that the correlation between mobility and patenting observed in the data might be explained by factors other than knowledge transfer. To address this issue, in an extension we show that labor mobility is positively associated with the propensity of the firms involved to cite each other's patents, which we take as an indication that the mobility event resulted in knowledge transfer. Moreover, we corroborate our causal interpretation of the effect of mobility on invention by showing that the positive correlation between these variables is not significantly related to any large change in R&D employment. Finally, we apply instruments for firm mobility based on industry averages in GMM estimation to take account of firm-specific shocks to patent productivity. These additional analyses suggest that the relationship is not caused by some unobserved factor—such as a firm ramping up its R&D activity—which would increase both hiring and patenting.

The papers closest to ours are Hoisl (2007) and Ejsing et al. (2013) which study the effect of labor mobility on patenting. Hoisl (2007) combines data on mobility constructed from patent files, with background information on inventors obtained from questionnaires. She shows that mobile inventors are on average more productive, and that mobility increases inventor productivity. However, since she does not measure the previous employer's patent productivity she is unable to address the effect of mobility on the total level of invention. Using similar data to ours, Ejsing et al. (2013) show that hiring researchers from universities has a large effect on firm patenting but they do not address the issue of how worker mobility affects total invention.

In a related literature stream, registry data is used to test a prediction of human capital theory 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 but Maliranta et al. (2009) find that workers are not able to capitalize on the knowledge acquired from participating in R&D activities. Toivanen and Väänänen (2012) find a significant and potentially long-lasting wage premium for inventors of granted patents, indicating that these workers are perceived by their firms to possess valuable knowledge and skills. However, our aim is to measure the importance of mobility for total invention.

The rest of the paper is organized as follows: Section 2 provides the theoretical background to the analysis and develops the argument that mobility occurs only if it increases the total invention of the firms involved. Section 3 describes the data and outlines the definitions used in the analysis. Section 4 disucsses our econometric approach and provides descriptive statistics. The main results are reported in Section 5 with some corroborating evidence and robustness checks. Section 6 concludes.

2 Theoretical Background

2.1 The Effect of Labor Mobility on Firm-Level Invention

We briefly outline the main effects of labor mobility on firms' R&D capabilities identified in the literature.

Reallocation of skills and abilities: R&D workers possess technical skills and problem-solving abilities which are important inputs for the production of inventions. Since these skills and abilities can only be applied in one firm at a time, they are rival in nature (Arrow, 1962). That is, when a worker moves from one firm to another, the R&D capability of the new and the old employer are respectively increased and decreased.

Immediate knowledge transfer: Different pieces of knowledge are not necessarily equally widely distributed within the organization. Some knowledge is "private", resides within a single individual, and is available only to the current employer. Other pieces of knowledge are "social" and shared among several employees (Spender, 1996). So whether a particular piece of social knowledge can be transferred by a single individual is a key issue. Knowledge regarding how a technical process works is an example of social knowledge which can be individually transferable (Liebeskind, 1997) whereas implicit knowledge embedded in the routines, culture, and norms of the firm is difficult for a single worker to transfer (Spender, 1996).

If a worker switches firms, the new employer gets access to the worker's private knowledge and the part of the worker's social knowledge that is individually transferable. We refer to this as "forward knowledge transfer" since knowledge and labor flow in the same direction. The old employer loses only the worker's private knowledge. Hence, mobility leads to sharing of social knowledge, which is the fundamental reason why labor mobility is perceived as an important source of aggregate innovation (Saxenian, 1994; Franco and Mitchell, 2008).

Social ties and attention: Agrawal et al. (2006) observe that mobility results in the old and the new employers citing each other's patents more frequently. While the citing behavior of the new employer can be explained by its use of the knowledge that the worker brings, the apparent existence of "reverse knowledge transfer"—i.e. knowledge that flows in the opposite direction to labor—is striking. Two explanations for this phenomenon have been proposed by Corredoira and Rosenkopf (2010). First, the worker may stay in contact with former co-workers, resulting in knowledge exchange among the firms' employees. Second, the old employer's awareness of the worker's new employer as a source of knowledge may be heightened and may result in closer attention to the latter's patents and other R&D activities. Although our data do not allow us to disentangle these explanations, the theoretical predictions are clear. The increased focus of the involved firms on each other's activities, and the stronger personal ties among employees reinforce the forward knowledge transfer for the new employer. Furthermore, reverse knowledge transfer alleviates the loss of knowledge that the old employer experiences.

Putting together these three effects of labor mobility suggests that the new employer gains access to new skills and knowledge. This should increase the firm's R&D capability, and in turn its invention output based on new opportunities for knowledge recombination and possibilities to pursue new lines of research (Grossman and Helpman, 1991; Schumpeter, 1934). The old employer experiences loss of the worker's skills and private knowledge but might benefit from reverse knowledge flows. Thus, the overall effect of labor mobility on the old employer's R&D capability and invention output is a priori unclear. Our main interest in this paper is how labor mobility affects aggregate invention. The above arguments suggest that the invention output of the new employer increases while the invention output of the old employer may decrease, leaving the total effect—the sum of the effects on the new and the old employers—indeterminate.

2.2 The Effect of Labor Mobility on Total Invention

In order to gain theoretical predictions, we treat the market for R&D workers as competitive and assume that mobility occurs as the result of wage competition among firms. Assuming that a sufficiently large proportion of moves happens endogenously due to wage competition—rather than for exogenous family-related reasons for example—we would expect to find empirical evidence consistent with the predictions of this model.

Following Pakes and Nitzan (1984), we consider two firms competing for a worker currently employed by one of the firms. If the current employer keeps the employee, it earns profit π. If the worker moves, the old and the new employer earn $\theta_O\pi$ and $\theta_N\pi$, respectively. The profits $\theta_O\pi$ and $\theta_N\pi$ include all the costs and benefits that arise from labor mobility. Retaining the worker has value $(1-\theta_O)\pi$ to the current employer and hiring the worker has value $\theta_N\pi$ to the potential new employer. Hence, wage competition implies that mobility occurs if and only if $(1-\theta_O)\pi < \theta_N\pi \Leftrightarrow \pi < (\theta_N+\theta_O)\pi$. In other words, mobility occurs if and only if it increases the joint profit of the two firms (Pakes and Nitzan, 1984).

There are opposing effects of labor mobility on the joint profit of the firms (Combes and Duranton, 2006; Fosfuri and Rønde, 2004). First, social knowledge that was the sole possession of the old employer before the mobility event is shared with the new employer. This leads to increased competition over some commercial uses of the knowledge, reducing the profit of the old employer. Since competition destroys rents, the new employer gains less from entry to these commercial uses than the old employer loses from increased competition. Increased competition tends to reduce the joint profit of the firms, and thus deters labor mobility. Similarly, more mundane costs related to labor turnover, such as hiring and training costs and costs resulting from interruptions to the workflow, also tend to reduce joint profit and prevent mobility. Second, labor mobility may increase the firms' joint profit through its effect on invention. Firms have different R&D capabilities and strengths, and knowledge sharing increases the likelihood that a piece of knowledge will serve as an input to a new invention (Bessen and Maskin, 2009; Scotchmer, 1991). Therefore, knowledge sharing through labor mobility has the potential to stimulate invention, whether in the form of greater variety, value, or speed of invention. This effect increases the firms' joint profit and tends to facilitate labor mobility.

Taking these arguments together, the theoretical prediction is that labor mobility occurs if and only if the positive effect from an increase in joint invention outweighs the negative effect of stronger competition and other costs associated with labor mobility. Hence, we would expect the labor mobility that we observe in our data to cause an increase in total invention.

3 Data

The core of our data set is patent applications to the EPO filed between 1978 and 2006 with at least one applicant with Danish residency. These data were retrieved from EPO's "PATSTAT" database.[1] We consider in our analysis patent applications up to and including 2004, since the database for the years after 2004 is incomplete. Our data set includes 12,873 patent applications.

We use patent applications rather than patent grants because the average grant time (4-5 years) for the patents in our data set (Kaiser and Schneider, 2005) implies that a substantial number of patents applied for during the time period considered for our estimation (2000-2004) would be lost were patent grants used.[2] The "time stamp" of the patent application is the "priority date", that is, the date of first filing of the invention for patent protection at the EPO or any national patent office.

EPO data do not have 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—mostly manually—our EPO data to Statistics Denmark's firm identifiers. We were able to assign firm identifiers to 11,280 patent applications. The unmatched applications primarily refer to firms that went out of business before 1999. The corresponding information would in any case have been lost in our analysis, since our firm-level data start in 1999. After matching these data with our firm-level information, we were left with 11,031 patent applications applied for by 2,278 unique firms.

Statistics Denmark provided us with firm registry data, including most importantly firms' sectoral and regional affiliations and physical capital book value, and registry data on employee characteristics including, most importantly the end-of-November number of employees and their highest level of education.[3] We discarded sectors with no EPO patent applications between 1978 and 2004. Sectors are defined according

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

[2]There is a reporting lag between date of application and date of publication of the application in the EPO database. This implies that not all patents applied for after 2004 had been registered in the database at the time of data collection. We excluded these patents in order to avoid biases.

[3]As workers' firm affiliations are registered only once a year, in November, we do not observe within–year mobility.

to the three-digit NACE Rev. 1 industry classification. In a final step, we merged the firm-level data with employee-level data, which allows us to track the employment history of individual workers. We excluded firms founded during the estimation period 2000-2004. As described further in Section 4, our main estimation results are based partly on an estimator that requires information on firms' patenting behavior prior to 2000. Finally, we discarded firms from the public sector, since its patenting behavior is likely to be very different from that of industry.

When delimited according to these criteria, we have 349,595 observations for 93,725 firms in the population of private firms in Denmark. Our main sample consists of observations for firms that employ at least one worker in an R&D-related occupation. We focus on these since firms with employees in R&D-related occupations are much more likely to patent than firms with no R&D workers. Of the 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 between different varieties of apples rather than between apples and oranges. Our main estimation results thus include 42,507 observations for 14,516 unique firms, and 2,728 patents over the period 2000-2004.

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.[4] First, the person must have a Bachelor's, a Master's, or a Doctoral degree in technical or natural sciences, veterinary and agricultural sciences, or health sciences.[5] This criterion is based on the idea that knowledge flows are associated mainly with mobility of high-skilled workers. The definition corresponds closely to the findings in Kaiser (2006) who uses patent inventor survey (PATVAL) data for Denmark to show that 30.5 percent of inventors have Bachelor's degrees, 40.8 percent have Master's degrees and 17.4 percent have Doctorates. We intend to capture all individuals possessing the formal skills necessary to perform R&D related activities within the firm. Since some high-skilled workers may never conduct R&D 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 through the International Standard Classification of Occupations (ISCO) coding applied by the International Labour Organization.[6] At its first-digit level, ISCO classifies occupations according to their knowledge content. In particular, we can distinguish between "professionals" (level 2) and "technicians and associate professionals" (level 3).[7] 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 are the focus of our analysis of mobility since they are most likely to be directly involved in the creation of new knowledge. Individuals categorized as technicians and associate professionals occupy support positions which more likely utilize already existing knowledge. We call this group "R&D support workers". Since they are not directly engaged in developing new knowledge, they are not expected to be the main carriers of knowledge between firms. Therefore the share of the firm's support workers is included in our model as a control

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

[5] 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 not be included in our estimations since we exclude the public sector.

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

[7] 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.

variable only.

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

We next characterize categories of R&D professionals according to their mobility status. We differentiate between a main group simply termed joiners, who were employed in another firm in year $t-1$ (and hence were not employed in firm i in year $t-1$), and other joiners, who are workers whose job market status in year $t-1$ is unknown or who graduated between time t and time $t-1$. Stayers are R&D professionals who were employed by firm i at time $t-1$ and time t. Finally, leavers are workers employed in firm i in year $t-1$ who are employed in a different firm in year t. We also differentiate mobile workers, joiners, and leavers according to the patent activity of their old and new employers. Specifically, we distinguish between joiners who previously were employed by a "patenting firm", which we define as a firm with a positive patent stock at $t-1$, and joiners who previously were employed by a firm with no patents, a "non-patenting firm". We distinguish also between leavers who joined patenting vs. non-patenting firms. We do this to account for the inherent differences between firms that are patent active (and most likely also R&D active) and those that are not. Patent active firms are likely to possess a workforce that is endowed with a deeper and broader knowledge base than firms that do not patent. When a worker from a patent active firm joins a new employer, she may bring in a set of knowledge that is more valuable for invention.

Appendix A displays descriptive statistics for the variables involved in our estimations. The descriptive statistics for the overall sample show that the firms in our sample are generally small: The average firm has around eight R&D employees and a capital stock of about DKK78 mill. (the median is DKK2.7 mill.).[8] The overall level of patenting is fairly low. The average firm applied for 0.06 patents per year during the sample period.

We also provide descriptive statistics for the subset of firms which patented at least once before the beginning of our sample period, so-called pre-sample patenters. These firms can be expected to patent more regularly than the average firm for several reasons, including state dependence (Blundell et al. 1995) and the likely presence of unobserved firm-specific factors that favor patenting. In addition, observable firm characteristics are conducive to patenting for firms in this sub-sample compared to the average firm in the full sample. We find that firms with one or more pre-sample patents employ an average of 39 R&D workers, employ a stock of capital of DKK400 mill. on average, and produce 0.76 patent applications per year.

In relation to mobility and the composition of the R&D work force, our three groups of joiners (from patenting firms, from non-patenting firms, other joiners) jointly constitute more than 20 per cent of the current year's total employment of R&D professionals (joiners plus the reference category of stayers). The overall level of mobility of R&D professionals is high compared for example to the annual turnover rate of scientists and engineers of 13 per cent reported by Kim and Marschke (2005). The group of R&D supporters amounts to 45.7 per cent of current R&D employment.

When comparing the subsamples of firms with or without any pre-sample patents, the share of support workers is lower for pre-sample patenters (42.3 per cent) than for the other firms in our sample (46 per cent). Pre-sample patenters also attract a larger proportion of their joiners from other patenting firms (2.6 per cent of the current R&D work force compared to 1.3 per cent for firms without pre-sample patents). This is consistent with higher in-sample R&D intensity among "pre-sample patenters". The overall level of

[8] \$US1 corresponds roughly to DKK5.7.

mobility is comparable between the two sub-samples with 20 per cent of R&D professionals having joined within a year in the case of pre-sample patenters against 23 per cent for firms without any pre-sample patents.

Appendix B provides the correlations for the variables in our estimations. The table 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 (Belsley et al. 1980).

4 Empirical approach

This section describes our patent production function and outlines the econometric approach employed to estimate the relationship between worker mobility and firms' inventive output.

For the patent production function we assume a Cobb-Douglas specification as it is standard in the literature (Blundell et al. 1995; Hausman et al. 1984; Kim and Marschke 2005). Our dependent variable is the total number of a firm's patent applications in a given year, which we denote by P.[9] It is a count variable that takes the value zero or a positive integer. Hence we use count data models in the estimations. The mean of the count variable is linked to the explanatory variables exponentially:

$$E(P) = \exp\Big(\ln(A) + \alpha\ln(QL) + \beta\ln(K)\Big) \tag{1}$$

where QL denotes quality-adjusted R&D labor input and K denotes capital input. Our measure of labor input is defined to be specific to a firm's R&D activities. In the case of capital, our data do not allow us to measure the specific input of capital to R&D, hence we interpret capital stock as a general measure of firm size. The variable A summarizes factors other than capital and labor that affect patent production such as sectoral, geographical, and time effects which we also include in our empirical model.[10]

We choose an additive specification for quality-adjusted labor QL following Griliches (1967). We differentiate between four main types of R&D labor currently employed in the firm, namely stayers (denoted by St), joiners from firms (J), other joiners (O), and support workers (Su). Our specification for quality-adjusted labor is:

$$\begin{aligned} QL &= L_{St} + \gamma_J L_J + \gamma_O L_O + \gamma_{Su} L_{Su} + \gamma_X L_X \\ &= L\left(1 + (\gamma_J - 1)\frac{L_J}{L} + (\gamma_O - 1)\frac{L_O}{L} + (\gamma_{Su} - 1)\frac{L_{Su}}{L} + \gamma_X\frac{L_X}{L}\right), \end{aligned} \tag{2}$$

where current employment ($L = L_{St} + L_J + L_O + L_{Su}$) does not include leavers. We normalize the effect of stayers to unity. The coefficients γ_r measure the contribution of the rth worker type to quality-adjusted labor QL relative to the contribution of stayers.

Taking logs and using the approximation $\ln(1 + z) \approx z$ for small z we can plug the expression for $\ln(QL)$ into the patent production function. This leads to our basic estimating equation which differentiates between different R&D worker types:

$$E(P) = \exp\left[\ln(A) + \alpha\ln(L) + \alpha_J\frac{L_J}{L} + \alpha_O\frac{L_O}{L} + \alpha_{Su}\frac{L_{Su}}{L} + \alpha_X\frac{L_X}{L} + \beta\ln(K)\right],$$

[9] For brevity, we exclude explicit indices in what follows.

[10] Our econometric specification controls for 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 1 year to allow for time lags in the R&D process and to alleviate concerns about reverse causality.

where $\alpha_r = \alpha(\gamma_r - 1)$ for worker group r currently in the firm and $\alpha_X = \alpha\gamma_X$ for leavers. Our estimations identify the α-coefficients from which we shall back out the productivity ratios γ_r. As discussed in Section 3, our main specification also differentiates between mobile workers who join from or leave to go to patenting vs. non–patenting firms. This introduces α_r^P and α_r^N as a straightforward extension where the superscripts P and N denote patenting and non-patenting firm, respectively.

The econometric models we apply will account for both state dependence in patenting activity and unobserved firm heterogeneity. We account for patenting dynamics since existing firm-level studies show that previous patenting activity has substantial positive effects on current patenting (Blundell et al. 1995, 1999, 2002; Crépon and Duguet 1997). Arguing that a firm's stock of past own patents represents knowledge from which future patentable ideas can be derived, Blundell et al. include the lagged discounted stock of patents as a regressor. Due to the relative short time span of our estimation sample, we follow Crépon and Duguet (1997) and use a dummy variable that indicates whether or not a firm patented in $t-1$ as our control for state dependence.

Likewise, our model allows for fixed effects in order to capture unobserved firm-specific permanent differences such as appropriability conditions for R&D investments or different technological opportunities that might affect current patenting. Estimating a simple model including a dummy variable for each firm would produce consistent estimates only in count data models where all the regressors are strictly exogenous, which clearly does not apply to a dummy variable related to past patenting activity, see Blundell et al. (2002). This is similar to the bias introduced by the simultaneous inclusion in a linear model of fixed effects and a lagged dependent variable which renders the lagged dependent variable endogenous by construction (Nickell 1981). To solve this problem, we consider two different fixed effect approaches for dynamic count data models: The Pre-Sample Mean (PSM) estimator of Blundell et al. (1995) and the GMM estimator of Blundell et al. (2002). We discuss each model in turn.

The idea behind the PSM estimator is to approximate unobserved time-invariant heterogeneity by using information on the firm's patenting behavior prior to the start of the estimation period. This is exactly the setting in our data: We possess information on all firms' patenting activity from 1978 onwards, while our explanatory variables (allowing for lags) are observed after 1999 only. The PSM estimator approximates the "true" fixed effect by the pre-sample patenting history of each firm, summarized by the term FE, which in our case consists of patents applied for during the period 1978 to 1998. Specifically, the PSM estimator uses the average of the dependent variable over the pre-sample period as a proxy for the correlated effects for each firm. Since a prominent feature of our data is an overall increase in the level of patenting during the pre-sample period, we normalize the firm's number of patents in a pre-sample year by the total number of patents applied for during that year.[11]

Since many of the firms in our data did not apply for patents during the pre-sample period, we follow Blundell et al. (1995, 1999) and include a dummy variable for firms that applied for at least one patent during the pre-sample period. This variable also serves to correct for the (arbitrary) small constant added to the number of pre-sample patents to make log-transformation of FE feasible. It also remedies the so-called "zero-inflation problem" (Mullahy, 1997). It is common in analyses of economic count data to find excess numbers of zeros, a problem often treated by a zero-inflated model. This model treats the selection of firms into patenting as separate from the determination of the non-zero number of applications by actual

[11] Our approach hence allows for trends in patenting at the general level such as business cycle effects, changes in the propensity of firms to patent rather than to opt for secrecy, or changes in the propensity of Danish firms to patent at the EPO rather than the national patent office.

patenters. In section 5.3, we consider a zero-inflated model as a robustness check and find that our main results hold qualitatively across alternative empirical approaches.

As an alternative to PSM estimation, Blundell et al. (2002) consider a GMM estimator which also accounts for both fixed effects and lagged dependent variables. This estimator is best compared to the more popular dynamic panel data estimators for linear models (Arellano and Bond 1991; Arellano and Bover 1995). We follow Kim and Marschke in applying a quasi-differencing transformation to correct for fixed effects as suggested by Wooldridge (1991). It essentially removes the fixed effects by a non-linear transformation, much like the standard "within transformation" in linear models. The zero inflation problem is accounted for by the transformation as any time-invariant explanatory variable—including variables that relate to the selection into patenting—are swamped by the fixed effects. As in Blundell et al. (2002) we use longer lags of our dependent variable as instruments for the lagged dependent variable.[12]

The GMM estimator accounts also for other variables apart from the lagged dependent being endogenous. One potential concern is that causality may run not only from mobile workers to patent applications (as it is assumed in our theoretical discussion) but also in the reverse direction. We discuss this concern further in subsection 5.2 where we present evidence to corroborate our interpretation. Our GMM estimator addresses this concern by instrumenting worker shares. As instruments we use the firms' own lagged shares and the average share of each type of worker in other firms in the same sector and in the same region. The intuition is that sector-specific and region-specific supply and demand shocks to other firms will affect the demand for each skill group from the focal firm. At the same time, the average shares of the skill groups of other firms are unlikely to be correlated with the error term of our equation of interest—unobserved firm-specific factors that affect the firm's patent production.

The GMM estimator comes with a set of data requirements that considerably reduces the number of our observations. The most binding requirement is that there must be at least three subsequent observations per firm. The GMM estimator uses 23,769 observations for 6,751 firms while the PSM estimator uses 42,507 observations for 14,516 firms.

5 Results

5.1 Main results

Our estimation results are presented in Table 1. We report results for different count data specifications (Poisson and negative binomial) and for different estimators (PSM and GMM). We comment mainly on the negative binomial PSM results which are our preferred results.

Looking at the effects of joiners, the statistical significance of the α-coefficient of a group of R&D joiners is relative to the reference group of R&D stayers. The sign tells us whether this R&D worker type contributes more or less to patenting than stayers. The share of joiners from patenting firms has the largest effect on patent productivity. However, joiners from non-patenting firms are not significantly more productive than stayers. We interpret the finding of stronger effects for joiners from patenting compared to non-patenting firms as reflecting that in the former case workers transfer knowledge valuable for invention but not in the latter case.[13] Hence our results suggest that any negative effects of intellectual

[12] To estimate the GMM model we use Windmeijer's (2002) "ExpEnd" program that runs under the software package GAUSS.

[13] The reverse argument may hold that patenting firms have established a reputation for strict enforcement of patents in

property protection and strategic litigation of departing employees are outweighed by positive knowledge transfer effects. The heterogeneous group of "other joiners" also has a positive effect on patenting, which is statistically highly significant although smaller than the effect of joiners from patenting firms. This effect is most likely due to the presence within this group of firm expatriates and graduates for example, an issue analyzed in detail in Ejsing et al. (2013).

For the leaver groups, their α-coefficients show whether R&D workers of this type contribute to the focal firm's patenting activity even though they are no longer employed by that firm. Our results show a positive and significant effect of leavers who left for a patenting firm but no significant effect of leavers to non-patenting firms. Again, this suggests that patenting firms are richer sources of knowledge.

Comparing the models in Table 1, we see the results are qualitatively very similar. The GMM Poisson is estimated with less precision which is unsurprising given that it instruments the main mobility variables. The only major difference in the estimates is that the coefficients of joiners from and leavers to patenting firms are bigger in the GMM than the PSM. For our GMM approach to be valid, we require our instruments to be strongly correlated with the mobility terms and simultaneously uncorrelated with the error term in the patent production function. We test the first property by running "first stage" regressions of our instruments and our exogenous variables on the endogenous variables. F-tests of joint significance of the instruments should be above 10 for them to be "sufficiently correlated" with the endogenous variables (Stock et al. 2002). In our case, all F-statistics are substantially above 10. We consider the second property by Hansen J-tests and find no evidence against orthogonality of the instruments. The null hypothesis that our instruments are orthogonal cannot be rejected at a marginal significance level of more than 70 percent.

We base our further analysis on the PSM Negbin results since this model is indicated by a selection test for overdispersion and incorporates a larger number of observations and firms than the GMM results.

Table 1 shows also that we find substantial evidence of state dependence. This may reflect sunk costs associated with learning to conduct successful research or more practical knowledge related to patent application filing. Our correction for unobserved heterogeneity in the PSM model has a significantly positive impact on current patenting activity. An increase in the number of pre-sample patents by 1 percent is associated with an increase in the number of current patents of around 0.3 percent.

Insert Table 1 about here.

The coefficients in Table 1 do not translate directly into marginal effects as in a linear model. To facilitate interpretation of the magnitude of these effects, we convert the PSM NegBin estimates into productivity ratios, the γ terms discussed in section 4. The productivity ratios displayed in Table 2 show that joiners from patenting firms are 6.3 times more patent productive than R&D stayers, a figure that is statistically highly significant. The related figures for other joiners from and for leavers to patenting firms are 4.6 and 3.5, respectively. The remaining ratios are statistically insignificantly different from 1 indicating that these groups of R&D workers are no more productive than stayers.

Insert Table 2 about here.

Finally we evaluate the *total* effect of labor mobility on patenting which is the main focus of the paper. We conduct a thought experiment designed to increase the rate of turnover while keeping R&D

order to reduce the risk of knowledge leaks due to worker exit (Agarwal et al. 2009). This would suggest a *smaller* effect of mobility flows involving patenting rather than non-patenting firms, since the firm receiving the knowledge transfer might be reluctant to use proprietary knowledge.

employment unchanged. Going from period $t-1$ to t, our experiment replaces an incumbent worker with a joiner, keeping total employment constant. Compared to the situation of no mobility, the effective labor input QL will include an additional joiner effect, γ_J , and an additional leaver effect γ_X, while there will be one less stayer in the firm in period t. The total effect of mobility then is calculated as the marginal effect making this substitution. Partially differentiating our patent production function, Equation (1), we obtain the total effect of mobility as:

$$\text{Total effect} = \frac{\partial E(P)}{\partial L_J^i} + \frac{\partial E(P)}{\partial L_X^j} - \frac{\partial E(P)}{\partial L_{St}} = \frac{E(P)}{L}(\alpha_J^i + \alpha_X^j), \quad i, j \in \{N, P\}. \tag{3}$$

The α-coefficients are found in Table 1 for different types of workers. The expected number of patents, $E(P)$, and total R&D employment, L, are evaluated for an average firm in our sample as well as the average of firms with at least one pre-sample patent. The results are shown in Table 3. The strongest effect is found for the total effect of one worker leaving for a patenting firm and one worker joining from a patenting firm while keeping total R&D employment constant. This results in an additional 0.020 patents for the average firm in our sample, a 31 percent increase compared to the average number of patents. When evaluated for the average for firms with at least one pre-sample patent, the same type of substitution yields 0.046 additional patents, an increase over the average number of patents of 6 percent.

For combinations of leavers and joiners that involve at least one patenting firm, we find a positive and statistically significant effect of mobility on total patenting while mobility between non-patenting firms has no significant effect on total patenting. Overall, these findings are consistent with our theoretical prediction in section 2 that labor mobility increases the total innovative output of the firms involved.

Insert Table 3 about here.

5.2 Identification issues

Here we discuss our findings and provide additional evidence to corroborate our interpretation of the main results as being driven by knowledge spillovers from mobility. First, we establish a "paper trail" of patent citation links between firms that are connected by labor flows. Second, we examine the relative importance of another potential driver of our results, the "innovative thrust" created by a firm ramping up its R&D activities to further exploit an already existing knowledge base in the firm. Third, we address the argument in Kim and Marschke (2005) of preemptive patenting in the face of labor mobility. Finally, we discuss the extent to which our results could be affected by positive assortative matching between workers and firms.

First, we want to verify that the probability of citation links between firms increases if there is movement of labor between the firms. The presumption is that if a worker joins another firm and transfers knowledge, there will be an increased likelihood of patent citations between the firms. Mobility between the firms can go in either direction, or there might be a bi-directional exchange of labor. For the event that firm A, say, cites firm B, we distinguish between (i) a forward "joiner" effect if one or more R&D workers join firm A from firm B, and (ii) a reverse "leaver" effect if one or more R&D workers left firm A for firm B in the previous period.

We construct a dyadic data set of all possible combinations of firms that patent within the present period ("firm A") and firms which hold a positive patent stock at the beginning of the period ("firm B"). We define indicator variables for i) the event of one or more workers joining firm A from firm B, and ii) one or more R&D workers leaving firm A for firm B. For the case of bi-directional mobility, we define a separate indicator variable coded 1 if such bi-directional exchange occurred (and 0 otherwise). We set the forward and reverse mobility indicators to 0 if the bi-directional variable is coded 1. Finally, we define our dependent variable as an indicator of the existence of one or more citations in firm A's current patent application, to a patent in firm B's patent stock.

Table 4 shows the results of linear regressions of the citations link variable on our three labor mobility indicators (Probit regressions show very similar results and even stronger significance.) The positive and statistically significant coefficients of our three mobility dummy variables show that labor mobility is positively associated with the probability of firm A citing a firm B patent. This holds for all three types of mobility: A joiner's link, a leaver's link, and the bi-directional link. The key finding holds for the base specification which includes the mobility terms only (1st column), for a specification where we control also for industry and year (2nd column), and for firm characteristics (total number of R&D workers and size of the patent stock of the cited firms, 3rd column) in addition to industry and year controls. The results displayed in Table 4 strongly corroborate our interpretation of the results from the main empirical analysis: Both forward and reverse labor mobility appear to be positively associated with knowledge flows.

Insert Table 4 about here.

The second issue we investigate is the relative importance of the "innovative thrust" interpretation of our empirical results. The idea is that firms prepare to patent by investing heavily in R&D, investing in a laboratory and filling it with R&D workers. If such an alternative interpretation holds, our estimated joiner effect could simply be picking up massive R&D investments possibly unrelated to knowledge flows.

For the GMM results reported in Table 1, we note that instruments for firm mobility based on industry averages are applied. Therefore these results are not sensitive to firm-specific increases in innovative thrust. For the PSM estimation results, we provide an empirical check based on the presumption that innovative thrust will be related to the establishment of new projects or entire new lines of business which, in turn, are likely to be associated with large net additions to the stock of R&D workers in the firm. We re-run the PSM estimations in Table 1 and interact the shares of joiners from patenting and non-patenting firms and "other" joiners with a dummy variable that is coded 1 (and 0 otherwise) if the firm underwent a "large" change in R&D employment. Statistical significance of these interactions would provide evidence in favor of the innovative thrust interpretation.

Our definition of a "large" change has to take into account that most firms are small and there are only a few very large firms, measured by R&D employment. To do this we consider three different definitions of a "large" change: The case where we require R&D employment to change by more than 100 percent or by more than 10 workers; the case of R&D employment changing by more than 100 percent or by more than 5 workers; and the case of R&D employment changing by more than 50 percent or by more than 10 workers. The second definition is more inclusive in terms of relative changes; the third definition is more inclusive in terms of absolute changes.

Insert Table 5 about here.

Table 5 displays the estimation results. We find the interaction terms neither individually nor jointly to be significantly different from zero. Hence we find no evidence of the alternative innovative thrust interpretation mattering significantly for our findings. Our main PSM estimation results in Table 1 remain unchanged except for the effect of leavers which has the same sign and magnitude but is more imprecisely estimated.

A third issue related to the interpretation of our results is the knowledge protection argument proposed by Kim and Marschke (2005). It suggests that firms patent in order to prevent workers from transferring knowledge to other firms. However, the leaver effect we identify, on average materializes one year after the worker has left the firm which makes it unlikely that the patenting activity is related to an attempt to protect a specific invention that the departing worker had knowledge of. Furthermore, if we re-estimate the model using two lags instead of just one for all R&D worker related variables, the estimated leaver effect is only slightly lower. Although the order of magnitude is unaffected, the leaver effect is less significant (p-value of 11 percent) than in the one-lag specification. This is likely due to a substantial reduction in the sample size caused by the additional lag. It suggests that the effect is not driven primarily by any protective measures taken by the old firm since they would need to be put in place soon after the worker departed in order to secure priority over the invention.

Finally, although our data are very detailed about individual characteristics there is some concern related to the selection of R&D workers with different unobserved ability or human capital endowment into different types of firms. The thought experiment underlying our estimation of the total effect of mobility assumes homogenous unobserved quality of joiners, leavers, and stayers. Following Becker (1973) we might suspect that firms with the best conditions for conducting research may attract the best workers, so-called "positive assortative matching". In relation to matching R&D workers to firms this would suggest that workers in patenting firms are of higher quality on average than workers in firms with no previous patenting activity. This could explain at least part of the difference that we observe between the effects of joiners from these two types of firms. However, a similar argument applies to the leavers' side. Leavers to firms with previous patenting activity are on average of higher ability and the old firm suffers a greater loss of human capital for this group than for leavers to firms with no previous patenting. In this interpretation of our results, selection may upwardly (downwardly) bias the effect on joiners (leavers) from (to) firms with previous patenting activity. If the biases are approximately equal, the total effect of mobility in (3) will remain unaffected by the presence of matching. More importantly, even if the estimated effects of joiners from patenting and non-patenting firms were entirely due to unobserved differences in worker quality, our key finding of a positive effect of mobility on total patenting would remain due to the (possibly downward biased) leaver effect.

5.3 Robustness checks

We conduct five different checks of the robustness of our results: (i) Accounting for patent heterogeneity by weighting them according to the number of citations received, (ii) discarding the top 20 patenting firms, or alternatively all the biotechnology firms, to check whether our main results are driven by selected firms, (iii) applying a narrow definition of R&D worker by considering only workers with a Master's or Doctoral degree, (iv) re-running the regressions without correcting for trends in overall patenting behavior, and (v) using alternative estimators such as probit, Tobit, and zero-inflated models. For reasons of space, the results of these checks are not displayed here but can be found in the working paper version of this paper.

First there might be some concern that our estimates do not account for patent value heterogeneity. It is well known that the distribution of the economic and technological value of patents is heavily skewed in the sense that a few patents are very high value, while most are of very little value (see discussion in e.g. Hall et al., 2005, Harhoff et al., 1999 and Lanjouw et al,. 1998). Trajtenberg (1990) shows that, in the computer tomography industry, there is a close relationship between the number of citations a patent receives ("forward citations") and the social value of an invention. He therefore suggests using forward citations to approximate patent value since they capture the wide heterogeneity in the "quality" or "importance" of patents. Like Trajtenberg (1990), we weight each patent by 1 plus the number of citations the patent received within the three years after EPO publication. Our patent citation data are from the "EPO/OECD patent citations database" which is available from the OECD (Webb et al., 2005) and covers the period 1978-2006. The citations-weighted and citations-unweighted estimation results show only slight differences. The coefficients in the estimates referring to joiners become slightly bigger, while the coefficients of leavers remain almost unchanged. Citation-weighting hence generates results that provide even stronger support for our main result that mobility enhances total innovation.

There is a second concern related to whether our results are driven primarily by selected industries or firms that are very patent active. We estimate our main specification again, excluding the biotechnology sector or the 20 largest firms measured by their patent stocks in 2002. We find that the results of the estimations on these restricted samples differ very little qualitatively and quantitatively from our main results based on the full data.

A third issue is related to our R&D worker definitions. The main worry is that their definition might be too broad if it includes groups of workers who are unlikely to be engaged in research. The effect, if any, would be to bias our main results downwards. To assess the likelihood of this, we apply a less inclusive definition that selects only workers with a Master's or Ph.D. degree. While this restriction leads to results that are qualitatively very similar to our main results, the coefficients are generally smaller for the narrower definition of R&D workers than for our main definition. In view of the survey evidence reported by Kaiser (2006), we interpret this finding as meaning that workers with a Bachelor's level degree constitute a significant fraction of actual inventors in Denmark.

A fourth robustness check relates to our trend correction of correlated effects as discussed in section 4. Leaving out normalization for the general upward trend in patenting activity leads to results that are qualitatively identical and quantitatively very similar.

As a fifth robustness check we re-estimate our main model using alternative estimators. We run probit and Tobit regressions as well as a zero-inflated Poisson model. Although none of these estimators accounts for firm fixed effects, the positive effects of mobility reappear across the alternative estimation methods. The share of joiners from patenting firms, as well as the share of "other joiners", remain strongly significant in the probit and Tobit models. For the count part of the zero-inflated model we find positive effects of joiners although they are partly insignificant. The effect of leavers who left for patenting firms again is strongly positive and significant in the probit and Tobit models, but negative although insignificant for the count part of the zero-inflated model. The effects of the firm control variables (total R&D employment, capital stock, and previous patenting) have the same signs as in our main results and are significant overall. The main difference between the alternative estimators and our main results lie in the selection part of the zero-inflated model. This model estimates two equations, a probit model for having at least one patent application, and a count model for number of patents conditional on a non-zero patent count. Unsurprisingly, we find that the share of a firm's R&D professionals who apply and develop advanced knowledge, is a strong predictor of non-zero patent activity—reflected by a negative sign on the share of R&D support workers. The mobility terms show little significance in the selection equation. However, identification of the selection equation in this model is based entirely on functional form which may cause difficulties in separating out the effects on applying for one patent and the effects on applying for more than one patent. It also explains a general lack of significance of the results of the zero-inflated model. Overall, we find that the results of the alternative estimators for the effects of mobility are in general agreement with our main results based on the PSM and GMM estimators.

6 Conclusions

This paper assesses the quantitative importance of inter-firm labor mobility for invention, using a unique data set that combines patent applications by Danish firms to the European Patent Office with matched employer-employee registry data that track the employment history of R&D workers across time. We estimate the effect of labor mobility on the total patenting activity of the firms involved in the mobility event.

In line with the results in the literature, we show that an inflow of workers is associated with an increase in the firm's patenting activity. A worker joining from a patenting firm has a six times higher patenting productivity than a worker who stays with the firm. Interestingly, worker departure is not associated with a decrease in patenting. A worker who left to join a patenting firm contributes three times more to the original firm's patenting activity than a worker that stays, while a worker leaving for a non-patenting firm has no significant concomitant effect on patenting. Most importantly, we show that firms are not involved in a zero-sum game when competing for R&D workers to increase their R&D output. Worker mobility is related to a positive and statistically significant increase in total invention by the old and the new employer. The effect on total invention is strongest for mobility between two patenting firms where a mobility event increases the total patenting of the firms involved by 0.023. While this number might seem low, it compares to an average number of 0.064 patents per year for the average firm in our data which implies an increase of 36 per cent. Mobility between firms with no history of patenting is not associated with a significant increase in total patenting.

These results, to the best of our knowledge, provide the first quantitative support for the notion that inter-firm mobility stimulates total innovation. Saxenian (1994) in her study of Silicon Valley argues that "job-hopping" is crucially important for the innovation performance of the firms in that region, and our results confirm the importance of labor mobility in a much more representative setting covering all types of industries in Denmark.

A key issue is whether labor mobility causes the observed increase in patenting. We provide several pieces of evidence supporting such a causal interpretation. First, we show that mobility is associated with an increase in the probability of the old and the new employer citing each other in subsequent patents, which suggests that mobility does lead to knowledge transfer between the firms. Second, we find very similar results both qualitatively and quantitatively when we instrument labor mobility to reduce concern that our results might be driven by unobserved factors affecting both hiring and patenting. Finally, we show that the positive correlation between mobility and patenting is not significantly related to any large changes in R&D employment. This again suggests that the positive relationship between these variables is not caused by some unobserved factor—such as a firm ramping up its R&D activity—that would increase both.

We see our results as improving our understanding of the circumstances in which labor mobility stimulates firm-level innovation and aggregate growth. However the results in this paper should be interpreted with caution in relation to drawing conclusions regarding the optimal level of labor turnover in an industry or region. In a small country such as Denmark, firms are likely to face very similar labor market conditions. This is advantageous for the econometric identification but the results represent the association between mobility and patenting given the rate of labor turnover in Denmark. An important factor that must be considered is how labor turnover affects firms' incentives to invest in R&D. It would clearly be an important contribution if future work investigated exogenous variation in mobility rates to analyze how this affects aggregate innovation in an analysis of optimal turnover rates.

Table 1: Main estimation results

	NegBin PSM		GMM		Poisson PSM	
	Coeff.	**SE**	**Coeff.**	**SE**	**Coeff.**	**SE**
R&D worker shares						
Joiners from patenting firms	1.631***	0.266	2.389**	1.063	1.458***	0.397
Joiners from non–patenting firms	0.257	0.360	0.690	0.490	0.338	0.423
Other joiners	1.131***	0.276	1.168**	0.585	1.250***	0.352
Support	-0.149	0.207	-0.211	0.253	0.321	0.350
Leavers to pat. firms	0.720***	0.317	1.214*	0.730	0.836***	0.330
Leavers to non–pat. firms	-0.509	0.436	0.126	0.260	-0.936	0.840
Capital and R&D labor						
ln(total R&D workers)	0.309***	0.060	0.374***	0.093	0.417***	0.104
ln(capital stock)	0.138***	0.037	0.008	0.029	0.230***	0.073
Lagged dependent and pre–sample variables						
Dummy patent $t-1$	1.489***	0.148	1.450***	0.305	2.318***	0.421
ln(# pre-sample patents)	0.316***	0.079	—	—	0.135	0.129
Dummy pre-sample patent	0.325	0.231	—	—	-0.079	0.281
# of obs.	42,507		23,769		42,507	
# of firms	14,516		6,751		14,516	

Table 1 displays PSM NegBin, dynamic Poisson fixed effects GMM, and PSM Poisson estimation results. "SE" denotes the standard error. Patent citation weights have not been applied. The PSM specifications additionally include sector dummies, year dummies, region dummies and a constant term. These variables are time–invariant and drop out of the fixed effects GMM specification. The GMM specification uses Wooldridge moment conditions and additionally contains year dummies. It uses lags of worker shares and the average share of each type of worker in other firms in the same sector and in the same region as instruments for current worker shares. The asterisks' '***', '**' and '*' denote marginal significance at the one, five and ten percent level.

Table 2: Relative patent productivities

	γ_r	p–value
Joiners from patenting firms	6.277	0.000
Joiners from non–patenting firms	1.833	0.474
Other joiners	4.659	0.000
Support	0.519	0.482
Leavers to patenting firms	3.329	0.031
Leavers to non–patenting firms	-0.646	0.242

Table 2 displays the productivities of different types of R&D workers relative to the productivity of R&D stayers. "$P - value$" denotes the marginal significance level of the hypothesis that the relative productivity is one. These calculations are based on the Negbin PSM estimation results displayed in Table 1. **Reading example:** joiners from patenting firms are 6.3 times more patent–productive than R&D stayers.

Table 3: Total effects of mobility

	Left for patenting firm		**Left for non–patenting firm**	
	Coefficient	**p–value**	**Coefficient**	**p–value**
Average firm				
Joiners from patenting firms	0.020	0.000	0.009	0.023
Joiners from non–patenting firms	0.008	0.033	-0.002	0.643
Average firm with at least one pre-sample patent				
Joiners from patenting firms	0.046	0.000	0.022	0.023
Joiners from non–patenting firms	0.019	0.033	-0.005	0.643

Table 3 displays our estimates of the total change in the number of patents if one worker (of different types) left the firm while one worker (of different types) joins the firm, keeping total R&D employment constant. The upper panel displays our results across all observations and the lower panel shows results for firms with at least one pre–sample patent. These calculations are based on the Negbin PSM estimation results displayed in Table 1. **Reading example:** if one R&D worker leaves for a patenting firm and one worker previously employed by a patenting firm joins, the expected increase in the number of patents is 0.020 for the average firm.

Table 4: Linear regression results for the relationship between R&D worker mobility and citations

	Base specification		**Industry & year**		**Firm characteristics**	
	Coeff.	**SE**	**Coeff.**	**SE**	**Coeff.**	**SE**
Forward worker mobility only	0.010***	0.003	0.009***	0.003	0.009***	0.003
Reverse worker mobility only	0.016***	0.006	0.016***	0.006	0.015***	0.006
Bi–directional flow of workers	0.047*	0.026	0.047*	0.026	0.045*	0.026
Year dummies	no		yes		yes	
Industry dummies	no		yes		yes	
Firm characteristics	no		no		yes	

Table 4 displays linear regression results for a firm's probability to cite another firm's patents. The results to the left refer to the base model which includes the three worker flow terms and a constant. The specification in the middle additionally includes year and dummies for firms being in the same industry. The model to the right also includes the log number of R&D workers of the citing firm, the log number of R&D workers of the cited firm, and the lagged log stock of patent applications. Years 2000 through 2004 are included. There are 516,049 dyads, 141 non-zero citation links, 1,011 instances of firms linked by a forward mobility link only, 866 instances of reverse links, and 168 instances of bi-directional mobility links. "SE" denotes the standard error. Standard errors are clustered at the firm–level. The asterisks '***' and '*' denote marginal significance at the one and ten percent level.

Table 5: Large vs small changes in R&D employment

	Change # R&D workers >100% or >10 workers		Change # R&D workers >100% or >5 workers		Change # R&D workers >50% or >10 workers	
	Coeff.	**SE**	**Coeff.**	**SE**	**Coeff.**	**SE**
R&D worker shares						
Joiners from patenting firms	1.419***	0.465	1.303***	0.491	1.444***	0.555
Joiners from non–patenting firms	-0.271	0.816	-0.202	0.813	-0.260	0.937
Other joiners	1.465***	0.463	1.484***	0.462	1.604***	0.467
Support	-0.086	0.240	-0.087	0.240	-0.071	0.242
Leavers to pat. firms	0.471	0.421	0.474	0.427	0.456	0.440
Leavers to non–pat. firms	-0.759	0.513	-0.766	0.520	-0.766	0.509
"Large change" interactions						
Share joiners pat. firms × large change	0.859	0.994	1.172	0.943	0.570	0.831
Share joiners non–pat. firms × large change	1.734	1.097	1.477	1.133	1.051	1.133
Share other joiners × large change	-0.398	2.464	-1.254	2.446	-1.524	1.438
Capital and R&D labor						
ln(total R&D workers)	0.321***	0.064	0.322***	0.065	0.333***	0.064
ln(capital stock)	0.121***	0.043	0.122***	0.043	0.119***	0.043
Lagged dependent and pre–sample variables						
Dummy patent $t-1$	1.541***	0.163	1.539***	0.163	1.540***	0.163
ln(# pre-sample patents)	0.295***	0.085	0.295***	0.085	0.294***	0.085
Dummy pre-sample patent	0.459	0.284	0.452	0.287	0.454	0.284
Tests for joint significance	**F-test**	***p*-value**	**F-test**	***p*-value**	**F-test**	***p*-value**
Joint significance interaction terms	4.18	0.242	4.15	0.246	2.56	0.465

Table 5 displays NegBin PSM results that include interactions of the three different shares of joiners with a dummy variable that is coded one if there was a "large change" in the total number of R&D workers. A "large change" is defined in three different ways: (i) Change in R&D employment by more than 100 percent or a total change by more than ten workers (left part), (ii) change in R&D employment by more than 100 percent or a total change by more than five workers (middle part), (iii) change in R&D employment by more than 50 percent or a total change by more than ten workers (right part). The definitions of "large changes" involve two–year lags which reduces the number of observations and the number of firms to 27,199 and 9,438 respectively. The specification additionally includes sector dummies, year dummies, region dummies and a constant term. The asterisk '***' denotes marginal significance at the one percent level.

References

Agarwal, R., M. Ganco and R.H. Ziedonis (2009), Reputations for Toughness in Patent Enforcement: Implications for Knowledge Spillovers Via Inventor Mobility, Strategic Management Journal, 30, 1349-1374.

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.

Arellano, M. and S. Bond (1991), Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations, Review of Economic Studies 58, 277-297.

Arellano, M. and O. Bover (1995), Another Look at Instrumental Variables Estimation of Error-component Models, Journal of Econometrics 68, 29-51.

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.

Becker, G.S. (1973), A Theory of Marriage: Part I, The Journal of Political Economy 81, 813-846.

Belsley, D. A., E. Kuh and R.E. Welsh (1980), Regression Diagnostics: Identifying Influential Data and Sources of Collinearity. New York: Wiley.

Bessen, J. and E. Maskin (2009), Sequential Innovation, Patents, and Imitation, RAND Journal of Economics 40(4), 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.

Campbell, B.A., M. Ganco, A.M. Franco and R. Agarwal (2011), Who Leaves, Where to, and Why Worry? Employee Mobility, Entrepreneurship and Effects on Source Firm Performance, Strategic Management Journal 33, 65-87.

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, 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.

Ejsing, A.K., U. Kaiser, H.C. Kongsted and K.Laursen (2013), The Role of University Scientist Mobility for Industrial Innovation, Copenhagen Business School mimeo.

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. 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.

Grossman, G. and E. Helpman (1991), Innovation and Growth in the Global Economy, Cambridge: MIT Press.

Hall, B.H., J.A. Hausman and Z. Griliches (1986), Patents and R&D: Is There a Lag?, International Economic Review 27, 265–83.

Hall B. H., Jaffe A. B. and Trajtenberg M. (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 and D. Slottje (2006), Intellectual Property Litigation in the USA, Journal of Economic Surveys 20, 715–729.

Jacobs, J. (1969), The Economy of Cities, New York: Random House.

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.

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.

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.

Liebeskind, J.P. (1997), Keeping Organizational Secrets: Protective Institutional Mechanisms and their Costs, Industrial and Corporate Change 6, 623-663.

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.

Mansfield, E. (1985), How Rapidly Does New Industrial Technology Leak Out, The Journal of Industrial Economics 34, 217–223.

Marshall, A. (1920), Principles of Economics, London: Macmillan.

Marx, M., D. Strumsky and L. Fleming (2009), Mobility, Skills, and the Michigan Non-Compete Experiment, Management Science 55, 875-889.

Mullahy, J. (1997), Heterogeneity, Excess Zeros, and the Structure of Count Data Models, Journal of Applied Econometrics 12, 337–350.

Møen, J. (2005), Is Mobility of Technical Personnel a Source of R&D Spillovers?, Journal of Labor Economics 23, 81-114.

Nickell, S. (1981), Biases in Dynamic Models with Fixed Effects, Econometrica 49, 1417-1426.

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.

Png, I. (2012), Trade Secrets, Non-Competes, and Mobility of Engineers and Scientists: Empirical Evidence, National University of Singapore mimeo.

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. Sorenson (2011), Noncompete covenants: Incentives to Innovate or Impediments to Growth, Management Science 57, 425-438.

Saxenian, A. (1994), Regional Advantage: Culture and Competition in Silicon Valley and Route 128. Cambridge, MA: Harvard University Press.

Schumpeter, J.A. (1934), The Theory of Economic Development: An Inquiry into Profits, Credit, Interest, and the Business Cycle, Cambridge, MA: Harvard University Press.

Scotchmer, S. (1991), Standing on the Shoulders of Giants: Cumulative Research and the Patent Law, Journal of Economic Perspectives 5, 29-41.

Spender, J.C. (1996), Making Knowledge the Basis of a Dynamic Theory of the Firm, Strategic Management Journal 17, 45-62.

Stock, J.H., J.H. Wright and M. Yogo (2002), A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments, Journal of Business & Economic Statistics, 518-529.

Toivanen, O. and L. Väänänen (2012), Returns to inventors, The Review of Economics and Statistics 94, 1173-1190.

Trajtenberg, M. (1990), A Penny for Your Quotes: Patent Citations and the Value of Innovations, The Rand Journal of Economics 21, 172–187.

Tzabbar, D. (2009), When does Scientist Recruitment Affect Technological Repositioning? The Academy of Management Journal 52, 873-896.

Webb, C., H. Dernis, D. Harhoff and K. Hoisl, K. (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.

Windmeijer, F. (2002), ExpEnd, A Gauss Programme for Non-Linear GMM Estimation of Exponential Models with Endogenous Regressors for Cross Section and Panel Data, Cemmap Working Paper No. CWP14/02.

Wooldridge, J.M. (1991), Specification testing and quasi-maximum-likelihood estimation, Journal of Econometrics 48, 29-55.

Appendix A: Descriptive statistics

	All obs.		Obs. without pre–sample patents		Obs. with pre–sample patents	
	Mean	**SD**	**Mean**	**SD**	**Mean**	**SD**
Dependent variable						
# patent appl. t	0.064	—	0.015	—	0.761	—
Dummy patent $t-1$	0.019	—	0.005	—	0.209	—
R&D worker shares						
Share joiners from pat. firms	0.014	0.089	0.013	0.089	0.026	0.089
Share joiners from non–pat. firms	0.062	0.199	0.063	0.203	0.046	0.122
Share other joiners	0.051	0.182	0.052	0.186	0.043	0.122
Share support	0.457	0.441	0.460	0.447	0.423	0.337
Share leavers to pat. firms	0.014	0.091	0.013	0.089	0.032	0.110
Share leavers to non–pat. firms	0.077	0.244	0.077	0.248	0.074	0.184
Capital and R&D employment						
Total R&D workers	7.693	44.570	5.473	25.694	39.043	140.253
Cap. stock (in mio. DKK)	77.50	1'280	54.80	1'140	399.00	2'520
Year dummies (base: 2000)						
2001	0.203	—	0.202	—	0.206	—
2002	0.196	—	0.196	—	0.202	—
2003	0.187	—	0.187	—	0.190	—
2004	0.183	—	0.183	—	0.181	—
Sector dummies (base: Wholesale and retail trade)						
Farm & food	0.016	—	0.016	—	0.019	—
Textiles & paper	0.041	—	0.041	—	0.036	—
Plastic & glass	0.026	—	0.023	—	0.072	—
Chemicals	0.014	—	0.011	—	0.054	—
Metals	0.049	—	0.047	—	0.084	—
Machinery	0.069	—	0.057	—	0.233	—
Electrics	0.030	—	0.028	—	0.067	—
Medical technology	0.018	—	0.015	—	0.063	—
Vehicles	0.007	—	0.006	—	0.021	—
Furniture	0.016	—	0.016	—	0.021	—
IT	0.070	—	0.072	—	0.035	—
Technical services	0.140	—	0.141	—	0.127	—
Business–related services	0.095	—	0.099	—	0.044	—
Other	0.180	—	0.191	—	0.023	—
Region dummies (base: Greater Copenhagen)						
Sjælland	0.097	—	0.098	—	0.088	—
Syd	0.224	—	0.223	—	0.237	—
Midt	0.207	—	0.208	—	0.196	—
Nord	0.074	—	0.073	—	0.090	—
Pre–sample variables						
# pre–sample patents	0.061	1.465	—	—	0.929	5.625
Dummy pre–sample pat.	0.000	—	—	—	—	—
# observations	42'507		39'696		2'811	

The Table 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 standard deviation.

Appendix B: Table of correlations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
(1) Share joiners from pat. firms	1										
(2) Share joiners from non–pat. firms	-0.032	1									
(3) Share other joiners	-0.002	-0.050	1								
(4) Share support	-0.115	-0.253	-0.240	1							
(5) Share leavers to pat. firms	0.070	0.022	0.040	-0.084	1						
(6) Share leavers to non–pat. firms	0.024	0.092	0.082	-0.172	0.043	1					
(7) ln(cap. stock)	0.016	-0.030	-0.059	0.100	0.025	-0.003	1				
(8) ln(total R&D workers)	0.031	-0.017	-0.025	-0.067	0.055	0.030	0.410	1			
(9) Dummy patent $t-1$	0.042	-0.007	0.003	-0.029	0.033	-0.009	0.151	0.264	1		
(10) ln(# pre-sample patents)	0.038	-0.020	-0.009	-0.027	0.055	-0.004	0.227	0.349	0.451	1	
(11) Dummy pre-sample patent	0.036	-0.022	-0.012	-0.020	0.053	-0.003	0.218	0.315	0.375	0.950	1