RETURNS TO INVENTORS

Otto Toivanen and Lotta Väänänen*

Abstract—A key input to inventive activity is human capital. Hence, it is important to understand the monetary incentives of inventors. We estimate the effect of patented inventions on individual earnings by linking data on U.S. patents and their inventors to Finnish employer-employee data. Returns are heterogeneous: inventors get a temporary reward of 3% of annual earnings for a patent grant and for highly cited patents a longer-lasting premium of 30% in earnings three years later. Similar medium-term premium's accrue to inventors who initially hold the patent rights, although they forgo earnings at the time of the grant.

I. Introduction

A key input to inventive activity is human capital. Thus, incentives for individuals to invest in acquiring appropriate human capital, select to engage in inventive activity, and put effort into it may be important. Therefore, it is of substantial interest and significance to shed light on the monetary incentives to inventors provided by the labor market. While the incentives of firms to invest in innovation have been established in numerous studies on the returns to R&D and the value of patents to firms, very little is known about how the inventive individuals behind these patents are rewarded for their effort. We take a step toward filling this gap by empirically examining the financial returns to patent inventors and the heterogeneity of those returns. We construct a data set where U.S. patents and their inventors from the NBER patents and citations data file (Hall, Jaffe, & Trajtenberg, 2002) are linked to Finnish employee-employer data containing detailed information on personal characteristics and earnings, as well as information on the employers from 1988 to 1999. These data give us the opportunity to explore, ex post, the existence of monetary incentives for individuals to invent by estimating the effect of patent grants on their earnings.

The sources of financial returns to employee inventors include incentive schemes that are directly tied to patenting, as well as indirect effects on earnings through an improved position of the individual in the labor market through a patent (or its quality) signaling the ability of the inventor. Compensation schemes that are tied to signals of effort and successful outcomes are particularly important when tasks are hard to monitor, including R&D activities. A nascent literature (Lerner & Wulff, 2007; Lach & Schankerman, 2008) exists on the effect of different incentive schemes on innovative performance. Once granted, a patent provides a publicly observable signal about the ability of the inventor and may lead to wage increases, for example, through increased outside offers and bargaining.[1] We therefore expect wages to respond to patenting through direct compensation or labor market effects. Taking an ex post view and measuring the level of rewards provided by the labor market and existing incentive schemes is a key input in the design of new ones. Our research therefore complements the literature studying the effects of incentives on innovation output.

We have access to a unique data set of individuals, where we observe individuals' earnings and the USPTO patent grants for their inventions from 1991 to 1999 (together with citations received up to 2002). This enables us to take a novel and holistic approach to evaluating the existence of monetary incentives of invention and estimate the effect of patent grants and patent citations on individuals' earnings. Patents offer a convenient, if not trouble-free, window on individual inventiveness and have been exploited in economic research at least since the 1950s (Schmookler, 1957; Griliches, 1990). Using citations to patents improves this measure by accounting for patent value (see Trajtenberg, 1990). Monetary rewards to patenting may take various forms, including one-time bonuses, value-contingent payments, stock options, and wage raises. Thus, the returns can be both temporary and long term and may be realized some time after the patent grant. Therefore, we adopt a flexible specification including up to six lags of granted patents to allow us to identify the timing of the returns. Given what is known about the heterogeneity in the value of patents and about the time it takes to learn this value, we expect the signal to become more informative as some time passes from the patent grant and that citations to a patent (shown to be a good indicator of patent value) are a more informative signal than a simple patent count.

With panel data at the individual level and variation over time in patent grants, we can control for unobserved individual heterogeneity with fixed effects and remove the ability bias, which is often a problem in exercises of a similar nature, such as in estimating the returns to schooling (Card, 2001). The lag between the time of an invention and the patent grant enables us to treat granted patents as predetermined variables. Our specification also survives the strict exogeneity test, and the fact that first-differences and fixed-effects results are very similar lends further confidence that strict exogeneity is satisfied. This provides some indication that what we measure is the causal effect of patent grants on wages.

Received for publication March 4, 2010. Revision accepted for publication May 1, 2011.

* Toivanen: Katholieke Universiteit Leuven; Väänänen: Nokia.

We thank two anonymous referees, Ari Hyytinen, Josh Lerner, Bill Kerr, Suzanne Scotchmer, Kathryn Shaw, Konrad Stahl and Manuel Trajtenberg for discussions, and seminar participants at the Thirtieth Annual Meeting of the Finnish Society for Economic Research, HECER, EARIE 2008, EEA-ESEM 2008, Third ZEW Conference on the Economics of Innovation and Patenting, EALE 2008, Uppsala, Mannheim, NHH Bergen, and Copenhagen Business School for comments. We also thank Satu Nurmi, Margit Lahtinen, Villiina Hellsten, and Jouko Verho at Statistics Finland for their work with the data. We thank the Yrjö Jahnsson Foundation for financial support. L.V. also gratefully acknowledges financial support from the FDPE and the Deutsche Forschungsgemeinschaft through SFB/TR 15. The usual caveat applies.

[1] To the extent that such revelation of information leads to better employee-employer/task matches, those returns are also partly social returns.

The Review of Economics and Statistics, November 2012, 94(4): 1173–1190

We find that inventors get a temporary reward of about 3% of their gross annual earnings in the year of the patent grant, presumably corresponding to a one-time bonus for being awarded a patent. In addition, patent grants result in a 4% to 5% average premium in annual earnings three to four years after the patent grant. This premium remains for at least the following two years, possibly representing a permanent wage increase. The results are robust to (a) including any number of lags from four to six; (b) including a firm-level measure of invention (the number of patent grants and its lags, or alternatively, the expected lifetime citations to firm patent grants and their lags) to control for possible firm-level wage effects that are due to invention; (c) controlling for the average wages in the firm; (d) excluding the year 1999, which may be affected by the IT boom of the turn of the millennium; (e) controlling for firm R&D intensity; (f) removing the employees at the three largest patenting firms; and (g) including a large control group of noninventors.

We also find that behind the average effect we identify, there are interesting differences in returns, depending on the quality or value of the patent as measured by expected lifetime citations. These quality-dependent returns are first realized three years after the granting of the patent, coinciding with the time it typically takes to learn the value of a patent (Pakes, 1986; Lanjouw, 1998). Similar to the value of patents to firms and in line with the findings of Harhoff and Hoisl (2007), the returns to inventors seem heavily skewed and linked to citations (see Trajtenberg, 1990; Hall, Jaffe, & Trajtenberg, 2005). Indeed, it is only the highest-quality patents that yield positive returns for inventors. When we include the quality of patent output as a categorical variable depending on the expected lifetime citations received, we find that patents with twenty to thirty citations generate a premium of about 20% in annual earnings in the sixth year and patents with over thirty citations generate a premium of over 30% from the fourth year onward. In contrast, patents with fewer than twenty citations seem to generate no returns. Returns to inventors are thus very heterogeneous and tied to observable signals of the quality of the patent.

It seems natural to think of these rewards to patenting as part of "pay for performance," the increase in which has recently been shown to explain a large part of growth in male wage inequality in the United States from the 1970s to the 1990s (Lemieux, MacLeod, & Parent, 2009). One can also view patenting and the citations a patent receives as observable signals of an employee's ability or productivity. In models of learning about worker ability (Farber & Gibbons, 1996), the job market obtains signals of worker ability over time (usually assumed to be unobservable to the econometrician) and wages respond to these signals. In our application, these signals, represented by patents and their quality, are public information and thus observable also to the researchers. Our results indicate that they do play a large role in determining inventive individuals' remuneration. These results are also in line with survey evidence on the incentive schemes for inventors in Finnish firms, explaining, for example, the immediate reward due to a patent being granted and the information we collected by interviewing the director of patenting (Harri Honkasalo) at Nokia.[2] Furthermore, the results indicate compliance with the law on employee inventions in Finland, which states that inventors are entitled to compensation that depends on the value of the invention. While we are unable to identify whether the measured returns are due to incentive (hidden action) or signaling (hidden type) reasons, they do provide a motivation for economists' interest in emphasizing and studying pecuniary incentives to invent.

Most of the patents in our sample are assigned to corporations, but there is a small fraction where rights are unassigned or assigned to individuals, presumably the inventor.[3] We also analyze the dependence of the returns on the ownership of the intellectual property by comparing the returns to inventors who initially own the patent to the returns to those whose patent is assigned to an organization. Individuals and firms may have different capabilities to internalize the revenues from an invention, and the overall private returns from an invention may be greater for patents assigned to firms (Grönqvist, 2009). Individuals whose patents are assigned to firms receive only a share of the rents. We find that inventors who initially own their patents first forgo some of their earnings but eventually earn substantially higher rewards than inventors who do not have the intellectual property rights over their invention. The returns to inventors who initially own their patents is of the order of a 15% to 30% premium in gross earnings in the fifth to sixth year after patent grant (in contrast to the 4% to 5% average premium for those who do not). This difference is not explained by the higher quality of inventor-owned patents: the number of citations to inventor-owned patents is on average lower than to company-owned patents.

[2] Pekari (1993) examines employee inventions through case studies and interviews of sixteen actively patenting companies (six large, five medium, five small) in Finland. In eleven of the sixteen companies (and in all of the large companies), there were explicit rules for rewarding employees for their inventions. In large companies, the reward structure typically had three phases: at the time of the notice of invention, a fixed reward of 1,000–2,500 FIM (160–420 euros); at the time of the patent grant, a fixed reward of 2,200–10,000 FIM (360–1,700 euros); and as the value of the invention is revealed over time, a special value-contingent reward. The fixed fees were designed so that in most of the cases, they would represent reasonable compensation for the inventor and no special reward would be paid. However, if the invention later proved to be of exceptional value, the inventor would have been entitled to a special reward. This special reward is determined by the fraction of the value of the invention that the inventor is entitled to, depending on the employee's overall role. The value-contingent reward was typically paid two to three years after the patent grant. In small and medium-sized companies, while fixed rewards for invention and patenting were less common than in large firms, special value-contingent rewards for all patented inventions were used.

[3] These could be patents granted to university employees, for example, who at the end of the observation period are working in private sector. Some could be patents invented while working in the private sector but whose rights the employer for some reason gives away to the employee (e.g., when the invention falls outside the activities of the employer). Finally, they could be patents granted to individuals who own the company they work for (but are not classified as entrepreneurs, whom we exclude).

A number of other findings are also of potential interest. We find that employer changes after the patent grant do not affect the returns; regardless of whether the inventor stays with the firm where the invention is made or changes employers, the same returns accrue. Looking at gender differences, we find a male-female wage gap of 20%, even conditional on being an inventor. Regarding the difference in returns for males and females, we find the same immediate reward but no long-term returns for females. This could, of course, be due to a number of factors (e.g., females working in different industries and different firms), and we find some evidence of heterogeneity between firms (e.g., no significant long-term returns in the pharmaceutical sector). We also find that inventors have particularly high returns to age (experience), of the order of 10%–12%, possibly mirroring the results of Møen (2005), but the returns to tenure are low (less than 1%).

We proceed as follows. In section II, we review the related literature. In section III, we describe the data. In section IV, we present the empirical framework and in section V the results. We conclude in section VI.

II. Literature Review and Theoretical Framework

There is evidence that inventors are driven by profit motives. In the nineteenth century, individual inventors reaped the benefits of their inventions through own commercialization and through licensing (Khan & Sokoloff, 1993). As inventions became more technical and capital intensive, inventive activity (and individuals) moved into larger organizations (Lamoreaux & Sokoloff, 2005). Today inventors mostly invent as a part of their job, and inventive activity is to a large extent organized in R&D laboratories in firms and other R&D-performing organizations. Thus, for today's employee inventors, the financial incentives to invent are provided through employer compensation and through the labor market.

Incentives to invent are not limited to financial gains; they can also be nonpecuniary. Rossman (1931) reports the survey responses of a group of over seven hundred inventors, including the most prominent inventors of the time, who were asked for their motives and incentives to invent. The most commonly cited reason was "love of inventing," followed by "the desire to improve existing devices." "Financial gain," although clearly important, was only the third most frequently mentioned motive. More recent findings from Sauermann and Cohen (2010), who study the link between individual motives and innovative output for PhD scientists and engineers, indicate that intellectual challenge, independence, and money are motives that have a positive relationship to innovative output. There is clearly an element of current satisfaction ("on-the-job-consumption") that research activity provides, in addition to any financial rewards, as also noted by Levin and Stephan (1991) and emphasized in biographies of past inventors (Rossman, 1931). Similar evidence is provided by Stern (2004), who finds that scientists employed by firms in fact "pay to be scientists," that is, accept lower earnings in return for being able to pursue individual research agendas and publish in scientific journals.

The importance of nonpecuniary incentives notwithstanding, economists have studied the role of monetary incentives in the innovative process. Aghion and Tirole's (1994) incomplete contracts analysis, for example, normalizes the nonmonetary incentives to a constant and studies the effects of monetary incentives. The standard theoretical foundation for providing employees with (monetary) incentives comes from principal-agent models. In such a setup, the employee chooses the level of effort devoted to inventive activities. The assumption would be that the probability of coming up with a patentable invention, as well as the quality of the patent conditional on obtaining one, are increasing functions of effort. These models suggest that compensation should be tied to an informative signal of the level of effort (Holmström, 1979). Because both patents and measures of their quality (citations) are observable, it is useful for the principal to tie the agent's rewards to them in a context where unobservable effort is important for invention.

While incentive schemes have been subject to empirical research (Bandiera, Rasul, & Barankay, 2005; Lazear, 2000), they have been less studied in the context of innovation. An important exception is Lerner and Wulf (2007), who analyze how corporate R&D managers' compensation affects innovation in firms. Their key finding is that when the corporate R&D head has substantial firmwide authority over R&D decisions, long-term incentives such as stock options are associated with a higher level of innovation (more heavily cited patents, patents of greater generality, and more frequent awards). Another important exception is Lach and Schankerman (2008), who study the effect of university royalty-sharing schemes on university patenting in order to understand the importance of monetary incentives for university inventors. They find a positive correlation between the royalty share granted to faculty scientists (inventors) and university patenting. These papers differ from ours in that they use direct measures of monetary incentives where we use outcomes, and aggregate (firm or university level) data where we use individual-level panel data. These authors study the effects of incentive schemes on inventors' performance, whereas we seek to measure the aggregate level of all incentive schemes. Thus, our work complements theirs.

The provision of incentives is not the only reason that the labor market would reward inventors. For example, being a patent inventor may work as a signal of the individual's ability and productivity and so result in a wage premium. In an adverse selection framework, an employer could resort to offering a menu of contracts that induces individuals with a low ability to invent (useful or high-quality inventions) to self-select to other occupations. This could be accomplished by punishing (e.g., through lower wage growth) for "bad" invention outcomes, such as no patents or patents of low quality. The employer would have to reward for the high-quality output in order to induce individuals with a high ability to invent to become inventors. Thus, rewarding employees for patents and their quality may be due to signaling instead of incentive reasons.

Furthermore, such signaling can lead to improved firm-worker matches, thus raising earnings. Additionally, an invention represents knowledge, some of which is tacit and embedded in the individual, and this knowledge should earn a return in the labor market. A related point concerns knowledge spillovers: if firms want to prevent such spillovers, they may have to pay a wage premium to inventors in order to retain them. Evidence for this is provided by Møen (2005), who finds that while the technical staff in R&D-intensive firms first pay for the knowledge they accumulate on the job through lower earnings at the beginning of their career, they later earn a return on these implicit investments through higher earnings. At the same time, in order to benefit from incoming spillovers,[4] firms are willing to pay a premium in order to acquire such workers, as shown by Andersson et al. (2009), who find that firms with high-potential payoffs from innovation pay more in starting salaries than other firms in order to attract star workers (workers with a history of higher earnings and wage growth) and, furthermore, that such firms also reward these workers for loyalty. Van Reenen (1996) finds that technological innovation leads to higher average earnings in innovating firms and interprets the result in the light of theories of rent sharing.

Both the incentive and the signaling theories lead to an employee being rewarded for patenting or the quality of her patents. We are thus largely unable, in this paper, to separate between the two possible explanations.

Finally, as in many other countries, there is a legal framework that provides a basis to expect inventors to earn a return on the inventions they produce while employed (the law on employee inventions in Finland, 29.12.1967/656). While giving the right to the invention to the employer (in most cases), the law also rules that the employee has the right to reasonable compensation from the employer for the invention, taking into account the value of the invention.[5] Similar legal provisions exist (e.g., in Germany) and have been studied by Harhoff and Hoisl (2007). They address a question that is closely related to ours: using survey data on German inventors of EPO patents, they study how the characteristics of the surveyed patent affect the share of the inventor's salary received as compensation for that patent.[6] The survey responses from the inventors indicate that the average compensation for one patent is 1.8% of annual gross income and for all patents an average of 8.3%.

III. Data

A. *Matching USPTO and FLEED Data*

Our source of information on inventions and inventors is the NBER patents and citations data file (Hall, Jaffe, & Trajtenberg, 2002) on U.S. Patent Office (USPTO) patents. In the past few years, there have been some research projects making use of large-scale inventors' data. Most notably, Trajtenberg, Shiff, and Melamed (2006) have developed a computerized matching procedure to identify inventors in the NBER patent data. Some studies have used smaller-scale data. Kim, Lee, and Marschke (2004) use matched firm-inventor data from the pharmaceutical and semiconductor industries to study the relationship between firm size and inventor productivity. We go a step further than the previous studies and match inventor data to the employee records in a longitudinal employer-employee data set of the Finnish working-aged population (FLEED) that resides at Statistics Finland. The FLEED is a register-based data set that contains detailed information on all Finnish individuals in the working-age population and their characteristics, in particular their annual earnings, as well as firm-level information on their employers.[7]

The NBER patent data contain the names of all inventors of a given patent and information on their address (at a minimum, the municipality of residence). In Finland, each resident is given a unique identifier (the personal identity code), which is contained in the Finnish Population Information System (FPIS), together with basic personal information, including the address and municipality of residence. With the aid of the FPIS inventor information from the NBER patent data can be linked to personal identity codes. These personal identity codes are also contained in the FLEED (in encrypted form), enabling the linking of inventor information with it.[8] Those Finnish patents from

[4] Kaiser, Kongsted, and Rønde (2009) provide evidence that firms hiring R&D workers from patenting firms can benefit from incoming knowledge spillovers. They show that such workers contribute to an increase in patenting in the recipient firm and also that the labor mobility of R&D workers increases the joint patenting activity of the firms involved.

[5] Finnish law divides inventions into four groups in this respect. Inventions in group A came about through a close relation with the job of the inventor, and utilization of the invention fits into the activities of the employer or came about as part of the job of the inventor (no matter whether the utilization fits into the activities of the employer). In this case, the employer owns the invention if it so chooses. Inventions in group B came about in a different relation to the job as those in group A but fit into the activities of the employer. For these inventions, the employer has user rights, but must negotiate over any larger rights. Inventions in group C came about without a connection to the job of the inventor, but the utilization falls into the activities of the employer. The employer then has the right to negotiate over use rights first. Inventions in group D came about without a connection to the job of the inventor, and the utilization does not fall into the activities of the employer. The employer has no rights in this case (Mansala, 2008).

[6] Their survey contains a question about this share, but apparently no questions on levels of monetary compensation. Harhoff and Hoisl also offer a good discussion of legal compensation schemes for inventors in various countries.

[7] Given the richness of the FLEED data, it is impossible to detail its contents here. See Korkeamäki and Kyyrä (2000) for a description of the FLEED.

[8] The process of linking the inventor records to personal identification codes was done at Statistics Finland by its own personnel under strict confidentiality. We never had access to any information that would have enabled the identification of individual people from the data.

the NBER data that are assigned to Finnish companies have also been linked to their assignee firms in the FLEED. This provides us with an additional link we can use to help identify the inventors. In cases where the name and residence information in the inventor data matches more than one personal identity code from the FPIS, we also use this link between the patent inventor and the patent assignee, allowing us to search for the correct personal identity code from among the employees of the assignee firm. This information helps us solve a key issue that has hampered progress in studying inventors: the matching of inventors from patent documents to other data.

We use USPTO patents rather than Finnish patents because they should be more valuable. Grönqvist (2009) has estimated that the average value of a Finnish patent is on the order of only 5,000 euros, reflecting the small size of the Finnish market. Using USPTO data will also make our results comparable to other studies using the same data.

The data construction proceeded as follows. Using the full name and the municipality of residence on the inventor record (as well as the full address where available), together with the patent application year, the FPIS was searched for matching records, and all matching personal identity numbers were linked to the inventor record. For some, this resulted in a unique match, while for others, a number of potential identity numbers matched the inventor information. In order to determine the right identity for the inventor, we used the link between the patent inventor and the assignee firm to search the personal identity codes of all employees in the assignee for matches with those linked to the inventor record.

For individuals for whom more than one personal identity number was found from the population register, the identification of the correct individual was based on the assumption that he or she is an employee of the patent assignee firm. While we expect this to hold true for the majority, in some cases this may lead to misidentification of the inventor. Thus, we may have assigned a patent to some noninventors or failed to assign the patent to its proper inventor. If this is the case, it introduces some measurement error into our patent variable and biases our estimates downward.

Unfortunately, though not surprisingly, we were unable to identify and link all the patent inventor records to employee records for two reasons. First, for some inventor records, the search from the population register produced no match. This could be due to misspellings in the names or incorrect information for some other reason. Second, for some of those inventor records for which several matching identity numbers were obtained from the population register, more than one of these identity numbers was also found among the employees of the patent assignee firm. Without a unique match, we failed to identify and link the patent to any individual, so these inventors are not included in our sample.

Taking from the NBER patents data all the patents whose country code is FI and were applied for between 1988 and 1999, and linking these patents to their inventors, whose country code is FI, we ended up with 8,065 inventor patent records. From these, we are able to find a matching individual in the FLEED for 5,905 of them. These records correspond to 3,253 individuals (many individuals have more than one patent over the time period). Although there is good grounds to believe that the (in)ability to match patent records to individuals is random (e.g., spelling mistakes, lack of information on the inventor's address), we investigated whether there is any significant selection bias by comparing the patent records that we are able to match to those that we cannot. We compared the matched and unmatched patent records in terms of patent fields, number of inventors, number of citations, application year, and assignee code and found no substantial differences in any of these respects.[9]

For our empirical analysis, we limited the sample to observations from 1991 onward, because the linking of inventors and patents to the FLEED is based on the application year of the patent but our analysis uses the grant year of the patent. The typical lag from patent application to grant is one to three years, so for most of the cases, we were able to match a patent inventor to a granted patent from 1991 on. About 90% of the individuals were observed for all the years for the time period of analysis (1991–1999), resulting in close to 28,000 observations.

B. *Samples and Descriptive Statistics*

The process we have described generates our data on inventors: individuals listed as inventors in at least one USPTO patent during our observation period.[10] We restricted our analysis to individuals who were full-time employees at the end of the years in which we measured their earnings (i.e., removed those classified as entrepreneurs, unemployed, students, retired, in military service, or otherwise out of the labor market). This left us with just over 24,000 observations. We lost some observations due to missing firm (employer) information, leaving us with about 17,000 observations. There is no firm (financial statements) information available for public sector organizations, so the majority of the individuals for whom no firm information is available are employed in the public sector (e.g., in universities).[11] Finally,

[9] There is no statistically significant difference in the expected number of lifetime citations. The difference in the mean of inventors is statistically significant at the 10% level, but the difference is very small (3.05 for the unmatched versus 2.97 for the matched). There are slight differences in the distributions of patent fields, assignee types, and application years.

[10] Some 11% of the individuals listed as inventors are in a managerial position and could be managers of a laboratory. It could be that they have their names on the patent as a matter of policy rather than through having been involved hands-on in the inventive process. Our view on this is that to the extent that the laboratory manager is responsible for creating an environment that is conducive to invention, he or she should be rewarded, and we want to include them in our sample as inventors.

[11] For the years before 1994, we also lose some employees working in smaller companies, as firm financial statements information is then available for a limited number of firms (mainly large firms), but from 1994 onward (which is what matters for our main specification), it is available for almost the full population of Finnish firms.

we removed from the sample individuals who have patent grants only before the year 1991. We were left with a sample of 15,996 observations on 2,156 individuals. For our full specification, which includes six lags of the patent variable, the sample consisted of about 5,000 observations on 1,789 individuals who worked as full-time employees in the private sector.

Table 1A presents some descriptive statistics for this sample for the years 1991 and 1995–1999. The individuals in this sample are predominantly male (92%), on average 39 years old in 1991 (45 years old in 1999), and employed by their current employer (tenure) for 8 years on average in 1991. Mean annual earnings in the sample, about 37,000 euros (median 34,400) in 1991, increases throughout the period, reaching over 50,000 euros (median 44,900) in 1998 (all converted to 1999 money). The mean earnings in 1999 are at 80,000 euros, with a very high variance (median 44,900).

Table 1B presents the descriptive statistics conditional on having been granted a patent that year: The number of individual inventors almost tripled over the 1990s, from 196 to 560; the mean number of patents per inventor ranged from 1.2 to 1.4. The patent quality—the mean number of expected lifetime citations (see section VD for an explanation) received per patent—varies around 13 and shows no particular trend.

Table 1C presents the levels and fields of education for the sample. The inventors are fairly highly educated, with more than half of the inventors having a master's degree or a doctorate. Most of the inventors have an engineering degree (78%). Table 1D shows the number of observations in the main industry sectors represented in the sample, with 70% of the individuals working in the following five sectors: manufacturing of chemicals and chemical products; machinery and equipment; radio, TV, and communication; medical, precision, and optical instruments; and provision of business services.

The number of firms represented in the data is 224 in 1991 and 528 in 1999, with 936 different firms over the time period. The distribution of the number of individuals per firm is skewed, with (in 1999) over 350 firms employing just one inventor, 60 firms employing two inventors, 30 firms with three inventors, and only three firms with more than one hundred inventors.

Figure 1 presents a histogram of the number of patents per inventor over our sample period. The great majority of them (60%) have just 1 patent, about 20% have 2 patents, and the most inventive of them as many as 23 patents. To gain further insight into this, figure 2 presents a histogram displaying the frequency of observations with n patents. This distribution is also heavily skewed with a mass at 0 patents: almost 12,993 observations with 0 patents in a given year (not shown in the figure), 2,422 observations with one patent, and 409 with two patents. Figure 3 shows the distribution of citations for observations with at least one patent. This distribution is also heavily skewed to the left with a long right tail.

We have 127 inventor patent grant observations where the assignee is an individual; the rest are observations where the patent is assigned to an organization (mostly companies, so we refer to these as corporate-owned patents). Assuming that the patents assigned to individuals are owned by the inventor at the time of the patent grant, we refer to these as inventor owned. These could be patents granted to university employees, for example, who at the end of the observation period are working in the private sector. Some could be patents invented while working in the private sector but whose rights the employer for some reason gives away to the employee (e.g., when the invention falls outside the activities of the employer). Finally, they could be patents granted to individuals who own the company they work for but are not classified as entrepreneurs, whom we exclude. Comparing the number of citations by ownership, we find that inventor-owned patents receive fewer citations than those owned by organizations: the mean number of citations for inventor-owned patents is 7.32 and that for corporate-owned patents 10.27.

IV. The Empirical Framework

We estimate equations of the following form,

$$\ln(w_{it}) = X_{it}\beta + \sum_{j=0}^{\tau} \gamma_{j+1} Patent_{i(t-j)} + \alpha_i + \mu_t + \varepsilon_{it}, \quad (1)$$

where $\ln(w_{it})$ refers to the log of annual wage income, X_{it} is a vector of person- and firm-level characteristics, α_i is an individual-specific unobservable fixed effect, possibly correlated with the patent variables, μ_t is a vector of year dummies, and ε_{it} is the error term. Personal characteristics include the person's age and its square, a vector of 42 dummy variables for the level and field of education, gender, tenure with the current employer, and the number of months employed during the year. Firm characteristics include the sector of the firm, the number of employees in the firm, and its location regionally (NUTS2: 5 location dummies).[12]

The variable $Patent_{it}$ captures individual i's inventions in period t. The simplest measure of invention we use is a patent count: the number of patents granted in a given year in which the individual is listed as an inventor. Because inventions can affect earnings in subsequent years, not just in the year of the patent grant, we include τ lags of the patent variable in order to estimate any long-term wage effects of innovation. We experiment with as many lags as the data enable.

We also explore the implications of patent value or quality on the inventors' earnings by using forward citations to the patent. A number of studies have shown substantial het-

[12] NUTS 2 is a five-level regional classification system of the European Union. In Finland the five major regions are Southern Finland, Western Finland, Eastern Finland, Northern Finland, and Åland.

TABLE 1.—DESCRIPTIVE STATISTICS

A. Descriptive Statistics by Year

Variable	1991	1995	1997	1998	1999
Annual wage income					
Mean	37,468	41,280	46,215	52,287	79,556
s.d.	16,299	18,427	36,234	44,612	260,253
Number of patents					
Mean	0.15	0.19	0.26	0.33	0.42
s.d.	0.44	0.51	0.54	0.68	0.81
Citations-weighted patents					
Mean	1.54	2.64	2.50	3.45	3.84
s.d.	5.86	11.95	8.60	13.19	14.16
Age					
Mean	37.7	40.9	42.7	43.5	44.3
s.d.	7.8	8.2	8.2	8.0	7.9
Female dummy					
Mean	0.08	0.07	0.07	0.07	0.07
s.d.	0.27	0.26	0.26	0.26	0.26
Tenure with current employer					
Mean	8.6	10.4	11.3	11.8	12.3
s.d.	7.4	8.0	8.3	8.4	8.5
Number of months employed per year					
Mean	11.9	11.9	11.9	11.9	11.7
s.d.	0.75	0.79	0.68	0.70	1.57
Firm size (number of employees/100)					
Mean	26.4	23.6	28.2	28.5	28.0
s.d.	22.3	25.3	34.8	35.3	38.8
Observations	1,567	1,877	1,896	1,866	1,825

B. Descriptive Statistics Conditional on Having a Patent Grant That Year

Variable	1991	1995	1997	1998	1999
Annual wage income					
Mean	43,446	43,825	49,080	53,577	72,322
s.d.	20,718	20,343	22,558	48,189	167,175
Number of patents					
Mean	1.22	1.25	1.18	1.28	1.38
s.d.	0.51	0.62	0.50	0.75	0.91
Citations-weighted patents					
Mean	12.3	17.4	11.3	13.5	12.5
s.d.	12.0	26.2	15.3	23.3	23.3
Age					
Mean	41.7	41.8	42.4	42.7	42.8
s.d.	8.3	7.9	7.7	7.9	8.4
Female dummy					
Mean	0.06	0.07	0.04	0.08	0.10
s.d.	0.24	0.25	0.20	0.26	0.30
Tenure with current employer					
Mean	11.5	11.4	11.7	10.9	11.3
s.d.	8.0	7.8	7.9	7.8	8.3
Number of months employed per year					
Mean	12	12	12.0	11.9	11.7
s.d.	0	0	0.4	0.6	1.5
Firm size (number of employees/100)					
Mean	27.5	25.7	31.8	34.9	34.7
s.d.	24.3	23.4	36.9	38.9	43.0
Observations	196	284	421	478	560

C. Levels and Fields of Education[a]

Levels of Education	Percent	Fields of Education	Percent
Upper secondary	8.7	General	2.8
Lowest tertiary	9.7	Humanities and arts	0.4
Lower degree (bachelor)	21.8	Social science and business	1.3
Higher degree (master)	43.3	Natural sciences	11.1
Doctorate	12.2	Engineering	77.3
		Agriculture and forestry	0.9
		Health and welfare	1.8
		Services	0.1
Unknown	4.2	Unknown	4.2
Total	100.0	Total	100.0

TABLE 1.—(CONTINUED)

D. Main Industry Sectors	Class[b]	Observations	Percent
Manufacturing			
Machinery and Equipment	29	3,741	23.4
Radio, TV, and Communication	32	2,992	18.7
Chemicals and Chemical Products	24	1,907	11.9
Medical, Precision, and Optical Instruments	33	1,173	7.3
Other business activities (services)	74	1,328	8.3
All remaining sectors		4,855	30.4
Total		15,996	100

[a]The classification is based on the Finnish Standard Classification of Education 2007.
[b]Refers to Finnish National Standard Industrial Classification TOL1995.

FIGURE 1.—TOTAL NUMBER OF USPTO PATENTS GRANTED PER INVENTOR, 1991–1999

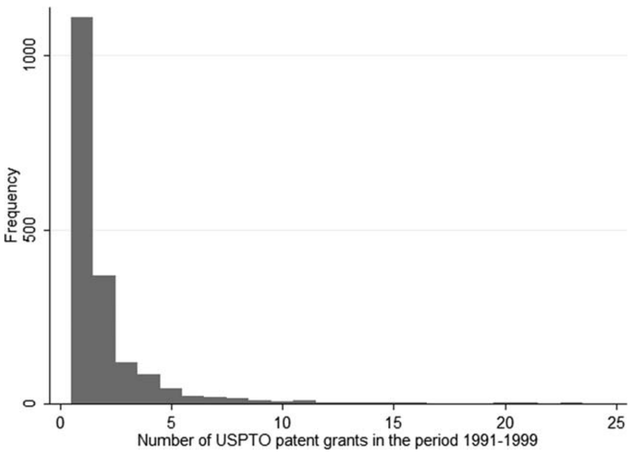

FIGURE 2.—NUMBER OF USPTO PATENT GRANTS PER INVENTOR PER YEAR

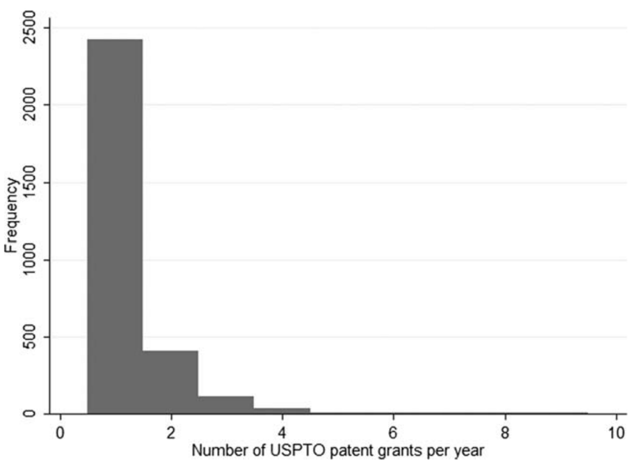

Observations with no patents (12,993) are excluded from the graph.

FIGURE 3.—EXPECTED LIFETIME CITATIONS RECEIVED BY PATENTS

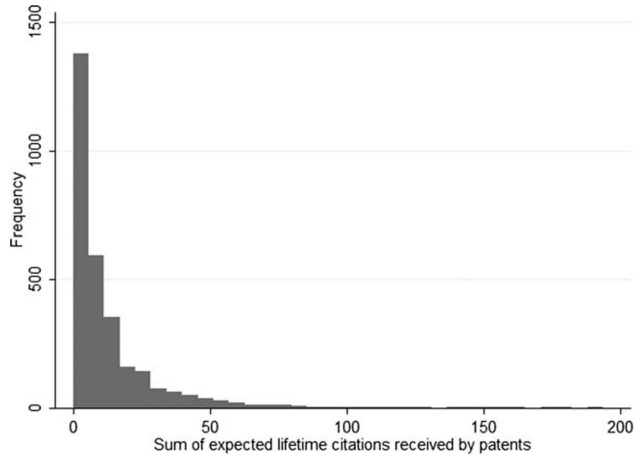

Observations with 0 patents (12,993) are excluded from the graph.

erogeneity in the value of innovations and that this distribution is highly skewed, for example, by using patent counts and renewal decisions (Pakes, 1986; Lanjouw, 1998; Grönqvist, 2007), and survey questions on patent value (Harhoff et al., 1999) and from patent citations (Trajtenberg, 1990; Hall et al., 2005). Given that the returns to firms from patents are highly variable, one might expect that the rewards that employers pay to inventors are also based on the value of the innovation.

We use both the within and the first-differencing transformations to identify the effect of patenting on an individual's wage. The key aspect is that any unobservable individual time-invariant factors are removed by these transformations. Importantly, this relieves us of the ability bias typically encountered in the returns-to-schooling studies (Card, 2001, for a review of the schooling studies). Both the within and first-differenced estimators are consistent under the assumption of strict exogeneity: $E[\varepsilon_{it}|Z_{i1}, \ldots, Z_{iT}, \alpha_i] = 0$.[13] We expect no contemporaneous correlation between the error term and the patenting variable, because a patent granted in year t has in effect been (pre)determined before year t. The lag between the years of patent application and granting of the patent is on average two years in our data. Therefore, the effort put into developing the innovation has been put in at least a couple of years, probably more, before the granting of the patent. One possible worry about the strict exogeneity condition is that future wage

[13] First-differencing does not necessitate strict exogeneity. For a discussion, see Wooldridge (2002).

TABLE 2.—EFFECT OF CONTEMPORANEOUS PATENT GRANTS ON ANNUAL EARNINGS

	OLS	FE	FD
Patent grants in year *t*	0.0354***	0.0161**	0.0129**
	0.0076	0.0072	0.0061
Age	0.110***	0.129***	
	0.008	0.008	
Age^2	−0.0011***	−0.0011***	−0.0014***
	0.0001	0.0001	0.0002
Tenure with current employer	0.0068***	0.0093***	0.0018
	0.0014	0.0013	0.0016
Female dummy	−0.213***		
	0.0228		
Number of months employed per year	0.114***	0.0901***	0.0870***
	0.009	0.007	0.009
Firm size (number of employees/100)	0.0008***	0.0023***	0.0009**
	0.0003	0.0003	0.0003
Constant	6.724***	5.853***	0.166***
	0.22	0.219	0.0157
Observations	15,996	15,996	13,419
Individuals	2,156	2,156	2,077
R^2	0.33	0.23	0.06

The dependent variable is log annual wage income. All regressions include dummies for the field and level of education, dummies for the sector of the firm, dummies for the firm's regional location, and year dummies. OLS are the results from pooled OLS estimations with clustered standard errors, FE are the results from using the within (fixed-effects) estimator, and FD are the results from the first-differenced regressions. Robust standard errors below. Statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

shocks may be correlated with the current period value of the patent variable, for example, through labor markets treating patenting as a signal of permanent or at least long-lasting productivity. However, this is part of the effect we estimate and is captured by the inclusion of the lagged values of the patent variable. If the realization of patents in the future is correlated with the contemporaneous error term in the wage equation, the strict exogeneity condition would be violated. This could happen (e.g., through changes in jobs either within or between firms) if a job change results in a better match between inventor and firm and also improves the patent productivity of the inventor. We apply a test of strict exogeneity and do not reject it. Under this assumption, the individual fixed effects also take care of selection into the sample and thus make the use of a control sample of noninventing individuals unnecessary. As one of our robustness tests, we include a large control group of noninventors in our estimation sample. Our results are robust to this.

V. Results

A. Patent Counts

We start with a specification that includes only the contemporaneous patent count and no lagged terms, that is, the variable $Patent_{it}$ being the number of patents granted to individual i in year t. The results are presented in table 2. Our preferred estimation results are those from the fixed-effects and first-differenced specifications, but we also report the results from the pooled OLS for comparison. The pooled OLS estimate of the patent grant reward to inventors is 0.035 (significant at the 1% level), the fixed-effects estimate is 0.016, and the first-difference estimate is 0.013 (both significant at the 5% level). The magnitude of the OLS estimate reflects the upward bias generated from unobserved individual heterogeneity, as expected.[14]

We also estimate a specification that includes the number of patent applications in addition to the number of patent grants in year t in order to test whether inventors are rewarded at the time of the patent application.[15] In the OLS estimations, patent applications obtain a positive significant coefficient in addition to the patent grant, but we find no statistically significant effect in the FE estimations, and the coefficient on the patent grants remains statistically significant and of the same magnitude as before.

We then move on to investigate whether the effect of patents on the inventor's wages is a permanent increase in the wage level (a raise) or a temporary one (a bonus) by including lags of the patent variable. This specification also allows us to investigate whether there are rewards to patenting that accrue only some years after the patent grant. Including lags is also important because patent grants may be correlated over time and thus introduce an omitted variable bias when not included in the estimations (in other words, a violation of the strict exogeneity).

We run a series of regressions where we include lagged values of the patent variable, experimenting with one to six

[14] Some of the control variable coefficients are of interest. The age premium (the return to experience) is relatively high (coefficient on age circa 0.1 and that of squared age −0.001) compared to the coefficient on tenure (measured in years), which is only 0.002 to 0.009. We also tried specifications including the square of tenure, which was mostly insignificant and did not affect our results. The coefficient on the female dummy is −0.21 (OLS coefficient). Firm size has a positive effect on earnings (large firms pay higher wages). Most of the year dummy coefficients are significant, as are many of the coefficients on the education dummies and the firms' sector dummies.

[15] For these regressions, we are forced to exclude the most recent years of our data (1997–1999), because these years are likely to suffer the most from the truncation problem (i.e., that we do not observe the patent applications for patents granted after 1999).

TABLE 3.—EFFECT OF PATENT GRANTS AND ITS SIX LAGS ON ANNUAL EARNINGS

	OLS	FE	FD
Patent grants			
In year t	0.0494***	0.0235	0.0275*
	0.0126	0.0144	0.0148
In year $t-1$	0.0005	−0.0052	0.0035
	0.0167	0.0218	0.0232
In year $t-2$	−0.0033	−0.0237	−0.0252
	0.0143	0.0225	0.0249
In year $t-3$	0.0050	0.0126	0.0080
	0.0206	0.0196	0.0214
In year $t-4$	0.0328**	0.0427**	0.0421*
	0.0144	0.0212	0.0218
In year $t-5$	0.0203	0.0552***	0.0468**
	0.0148	0.021	0.0199
In year $t-6$	0.0126	0.0493***	0.0522**
	0.0125	0.0176	0.0206
Age	0.113***	0.202***	
	0.0206	0.0458	
Age^2	−0.0012***	−0.0017***	−0.0016***
	0.0002	0.0005	0.0006
Tenure with current employer	0.0063***	0.0079***	0.0067***
	0.0017	0.0022	0.0021
Female dummy	−0.225***		
	0.0348		
Number of months employed per year	0.0177***	0.0067*	0.0044
	0.0065	0.0037	0.0045
Firm size (number of employees/100)	0.0007	0.0042***	0.0035***
	0.0005	0.0009	0.0010
Constant	7.768***	4.578***	0.186***
	0.446	1.177	0.057
Observations	4,938	4,938	3,126
Individuals	1,789	1,789	1,639
R^2	0.23	0.08	0.04

The dependent variable is log annual wage income. All regressions include dummies for the field and level of education, dummies for the sector of the firm, dummies for the firm's regional location, and year dummies. OLS are the results from pooled OLS estimations with clustered standard errors, FE are the results from using the within (fixed-effects) estimator, and FD are the results from the first-differenced regressions. Robust standard errors below. Statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

lags. We also test the strict exogeneity assumption by including the lead of the patent variable in the fixed-effect model and by including the levels of the patent variables in the first-differenced model (Wooldridge 2002). We cannot reject the null (of strict exogeneity) in either case. In table 3 we present the results from the estimations with six lags of the patent grant.[16] In all the estimations, the coefficient of the contemporaneous reward to a patent grant remains positive and in fact goes up (0.050 in OLS, 0.022 in FE, and 0.028 in FD). This suggests that there indeed is an omitted variable bias in the specification without lagged terms.[17] However, the statistical significance is somewhat weakened (the coefficient is significant only at 11% level in the FE and at 10% in the FD). In addition, interesting results emerge from the estimates on the lagged patent grants: the fourth, fifth, and sixth lags obtain a positive significant coefficient in the fixed-effects and first-differenced regressions, and they are of similar magnitude for these three lags (ranging from 0.042 to 0.055). These results indicate that in addition to a temporary reward in the year of the patent grant, there appears to be a longer-lasting, possibly permanent, effect that results in earnings increases of about 4% to 5% four years after the invention is patented.[18]

The fact that this wage increase comes a few years after the patent grant may be related to the fact that it typically takes three to four years to learn the value of the patent (see Pakes, 1986, and Lanjouw, 1998, for German, U.K., and French patents and Grönqvist, 2007, for Finnish patents). For example, Pakes (1986) finds that only 1.2 (0.5)% of French patent owners learn that their patent has no value in the third (fourth) year of patent life, and that the probability of learning a better use of the patent is only 0.1 (0.0)% in the third (fourth) year of patent life. His respective numbers for German patents are even lower.

[16] Results from specifications including up to four and five lags produce qualitatively similar results; the estimated coefficients are slightly smaller in magnitude, indicative of the omitted variable bias generated when ignoring longer lags. Specifications with fewer than four lags miss the positive effect captured by the fourth to sixth lags. Including a seventh lag reduces the significance of the estimates below 10%, although the signs are the same and magnitudes are close to those with fewer lags. This may be due to the fact that using seven lags leaves only two years of data at our disposal.

[17] Intuitively, what happens in the base specification is that the fourth to sixth years after the patent grant are wrongly allocated into the control group of "no patent grant" years, raising the average wage earned while in the control group and thereby inducing a downward bias in the base specification patent coefficient.

[18] It seems very unlikely that the patent lags would capture nonlinear experience or tenure effects, as our specification includes a quadratic function of age and tenure. In addition, we have estimated a specification with squared tenure included, with identical results.

TABLE 4.—EFFECT OF CITATIONS-WEIGHTED PATENT GRANTS ON EARNINGS

	OLS	FE	FD
Patent grants in year t	0.0398***	0.0286**	0.0270*
	0.0125	0.0136	0.0145
Citations-weighted patents			
In year $t-1$	0.0009	−0.0006	−0.00002
	0.0012	0.0015	0.0017
In year $t-2$	0.0012	0.0011	0.0003
	0.0008	0.0017	0.0021
In year $t-3$	0.0025	0.0035*	0.0023
	0.0015	0.0018	0.0021
In year $t-4$	0.0026*	0.0033*	0.0029
	0.0014	0.0018	0.0021
In year $t-5$	0.0014	0.0042**	0.0033*
	0.0013	0.0018	0.0019
In year $t-6$	0.0020	0.0050**	0.0042
	0.0020	0.0024	0.0026
Age	0.111***	0.179***	
	0.0207	0.0457	
Age^2	−0.0011***	−0.0015***	−0.0014**
	0.0002	0.0005	0.0006
Tenure with current employer	0.0062***	0.0076***	0.0064***
	0.0017	0.0021	0.0021
Female dummy	−0.224***		
	0.0348		
Number of months employed per year	0.018***	0.0052	0.0027
	0.0065	0.0044	0.0052
Firm size (number of employees/100)	0.0007	0.0043***	0.0035***
	0.0005	0.0009	0.0010
Constant	7.801***	5.114***	0.170***
	0.4490	1.185	0.0565
Observations	4,938	4,938	3,126
Individuals	1,789	1,789	1,639
R^2	0.24	0.08	0.04

The dependent variable is log annual wage income. All regressions include dummies for the field and level of education, dummies for the sector of the firm, dummies for the firm's regional location, and year dummies. OLS are the results from pooled OLS estimations with clustered standard errors, FE are the results from using the within (fixed-effects) estimator, and FD are the results from the first-differenced regressions. Robust standard errors below. Statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

B. *Patent Citations*

As the discussion at the end of the previous section suggests, the effect of patent grants on the inventors' earnings potentially depends also on the quality or value of the patent. Existing studies show that the number of citations received by a patent is a fairly good proxy for its value, so we make use of citations to better account for the quality of the inventors' patent output.[19] Using citations suffers from the problem of truncation, as citations to a patent arrive over long periods of time, but we observe them only until the last year of the available data.[20] We take the number of citations received by a patent by the end of the year 2002 and adjust these citation counts using the results in Hall et al. (2002) to remove the effects of truncation (based on the application year and the technology field of the patent). These adjustments provide an estimate of the total number of citations a given patent will receive in its lifetime. These estimates will be somewhat noisy; for the patents in our data, we observe citations only for the subsequent three to fifteen years. Typically the prime citation years for a patent are roughly three to ten years after the grant (Hall et al., 2005). The fewer citation years we observe for a patent, the noisier these estimates are.

Our specification now includes, instead of the number of patent grants, the expected lifetime citations that the patents granted in year t receive. These are the total (expected) citations to inventors' patents granted in year t. Thus, it combines both the citations per patent and the number of patents, and in essence describes the total quality of an inventor's patent output in a given year. Again, we include six lags of this quality-adjusted patent output measure to examine the timing of the rewards.[21] The results of these estimations are presented in table 4. We find that starting from the third year and lasting up to six years after the patent grant (and possibly permanently), the number of citations received has a positive effect on the inventor's earnings, with every ten citations received increasing the inventor's wage by around 3% to 5% (the estimates from the FD estimation are slightly lower than from the FE and less sta-

[19] For example, Hall et al. (2005) find that both the patents to R&D and the citations to patent ratios have a significant impact on the market value of the patenting firm.

[20] Here we make use of the updates to the NBER patent data, available from Bronwyn H. Hall's Web site (http://elsa.berkeley.ed./~bhhall/patents.html), allowing us to observe the number of citations received by the patents up until 2002.

[21] In the reported estimations, we use the patent count variable for the identification of the contemporaneous effect because we find no evidence for quality-dependent rewards in the year of the patent grant.

TABLE 5.—EFFECT OF CITATION-WEIGHTED PATENTS CATEGORIZED INTO FIVE BINS

Citations	0	0–9	10–19	20–29	30 or more
In year t	−0.019	−0.039	0.016	0.014	0.097
	0.033	0.026	0.052	0.093	0.088
In year $t-1$	−0.056	−0.027	−0.010	−0.078	0.034
	0.045	0.033	0.056	0.131	0.095
In year $t-2$	−0.127**	−0.045	−0.038	−0.131	0.086
	0.064	0.039	0.054	0.125	0.135
In year $t-3$	−0.077	−0.023	0.013	0.001	0.228*
	0.054	0.030	0.064	0.129	0.125
In year $t-4$	−0.101*	−0.015	−0.025	0.075	0.315***
	0.056	0.031	0.061	0.124	0.118
In year $t-5$	−0.073	0.038	0.044	0.131	0.327***
	0.055	0.028	0.057	0.124	0.112
In year $t-6$	−0.073	0.019	0.007	0.175**	0.363**
	0.053	0.023	0.030	0.078	0.158
Constant	4.800***				
	1.242				
Observations	4,938				
Individuals	1,789				
R^2, within	0.09				

The dependent variable is log annual wage income. The regressions include all the same control variables as before: age, gender, tenure, months employed, firm size, dummies for the field and level of education, dummies for the sector of the firm, dummies for the firm's regional location, and year dummies. Robust standard errors below. Statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

tistically significant). These results lend support to the notion that the returns to inventors depend on the value of the patent and are realized three years after the patent grant once the value of the invention is learned. The contemporaneous effect of the patent grant remains. Similar to the value of patents to firms and in line with the findings of Harhoff and Hoisl (2007), the returns to inventors thus seem heavily skewed. These findings lend further support for the claim, originating from Trajtenberg (1990), that citations are a measure of patent value.[22]

To study the link between patent quality and returns further, we categorize patents according to the expected lifetime citations they receive. Nineteen percent of the patent observations have no citations, 43% have one to ten citations, 18% have ten to twenty citations, 8% have twenty to thirty citations, and 12% of them have more than thirty citations. The results from the estimations with the five citation categories and their six lags are displayed in table 5. They offer further evidence that the returns to inventors are highly tied to patent quality: patents in the two highest quality categories (20–30 citations and over 30 citations) receive high positive returns. Those in the top category start earning significant returns in the third year after patenting. The coefficients are 0.23, 0.32, 0.33, and 0.36 for lags 3 to 6, respectively, implying returns in the range of 30% to 40%. Those in the category of twenty to thirty citations have a statistically significant effect of 0.175 in the sixth year after the patent grant. The point estimates for patents that receive no citations indicate that inventors with such patents earn a negative premium throughout. Two of these (for the second and fourth years) are statistically significant.

These results are qualitatively in line with models that suggest that the job market learns an employer's ability over time and rewards it. While such learning is often (Farber & Gibbons 1996) modeled as unobservable to the econometrician, one could view patenting and citations as observable measures of learning, available to the job market, public as they are.

C. *Controlling for Firm Effects*

While the previous estimations control for some differences among firms, including firm size and firm industry sector (at the two-digit level) as control variables, there could be other firm effects that affect the results. We first test whether our results remain once we control for potential firm-level rent sharing due to patenting, as found by van Reenen (1996). We also include the average wage in the firm in which the individual works to control for firm-level heterogeneity in wages. Finally, we redo all the estimations while controlling for the R&D intensity of the firm. Table 6 reports the results of these estimations using the patent count, table 7 reports the results using the expected lifetime citations, and table 8 reports the results using the categorical variable for citations.

Given the result of van Reenen (1996) that innovation in a firm leads to higher average wages (interpreted as higher wages for all employees), and given that our goal in this paper is to estimate the returns to individuals who make the inventions, we want to remove the possible concern that the returns we estimate are a reflection of firm-level rent sharing.[23] To accomplish this, we include a variable for the number of patents granted to the firm in year t, together with six lagged values of it. While some of the firm-level patent variables obtain significant (positive) coefficients (the contemporaneous term and the fourth lag), our other

[22] Trajtenberg (1990) found that citations reflect the social value of inventions. We find that they reflect the private (inventor) value of inventions.

[23] Although given that the sample contains individuals from the same firm who receive patents that year and those who do not, the results are not likely to be merely the result of firm-level wage effects.

TABLE 6.—EFFECT OF PATENT GRANTS ON EARNINGS, CONTROLLING FOR FIRM EFFECTS

	Firm Patents	Firm Wage	RD Sample	RD Intensity
Patent grants				
In year t	0.017	0.0276*	0.0352**	0.0353**
	0.015	0.015	0.014	0.014
In year $t-1$	−0.012	−0.009	0.008	0.008
	0.023	0.025	0.028	0.028
In year $t-2$	−0.027	−0.030	0.001	0.001
	0.023	0.025	0.019	0.019
In year $t-3$	0.011	0.008	0.005	0.005
	0.021	0.023	0.020	0.020
In year $t-4$	0.0343*	0.0413*	0.0367*	0.0366*
	0.021	0.023	0.021	0.021
In year $t-5$	0.0518**	0.0488**	0.0356**	0.0356**
	0.021	0.023	0.015	0.015
In year $t-6$	0.0503***	0.0453**	0.0463**	0.0463**
	0.019	0.020	0.019	0.019
Observations	4,938	4,450	3,709	3,709
Individuals	1,789	1,656	1,622	1,622
R^2, within	0.10	0.08	0.09	0.09

The dependent variable is log annual wage income. The results are from a fixed-effects estimation. The first column shows the results when controlling for firm patents and their six lags, the second column controlling for firm average wage, the third column for the sample where R&D intensity is observed, and the last column when controlling for firm R&D intensity. In addition to these firm-level control variables, the regressions include all the same control variables as before: age, gender, tenure, months employed, firm size, dummies for the field and level of education, dummies for the sector of the firm, dummies for the firm's regional location, and year dummies. Robust standard errors below. Statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

TABLE 7.—EFFECT OF CITATIONS-WEIGHTED PATENTS ON EARNINGS, CONTROLLING FOR FIRM EFFECTS

	Firm citations	Firm wage	RD sample	RD intensity
Patent grants in year t	0.0225*	0.0339**	0.0336***	0.0336***
	0.0136	0.0146	0.0121	0.0121
Cits-weighted patents				
In year $t-1$	−0.0008	−0.0005	0.0006	0.0006
	0.0016	0.0016	0.0015	0.0015
In year $t-2$	0.0008	0.0010	−0.0006	−0.0006
	0.0018	0.0018	0.0014	0.0014
In year $t-3$	0.0028	0.0031	0.0009	0.0009
	0.0019	0.0020	0.0013	0.0013
In year $t-4$	0.0026	0.0029	0.00345**	0.00345**
	0.0019	0.0019	0.0015	0.0015
In year $t-5$	0.00352**	0.00361**	0.0025**	0.0025**
	0.0016	0.0017	0.0010	0.0010
In year $t-6$	0.00447*	0.0044	0.0044*	0.0044*
	0.0025	0.0027	0.0024	0.0024
Observations	4,938	4,450	3,709	3,709
Individuals	1,789	1,656	1,622	1,622
R^2, within	0.10	0.08	0.10	0.10

Cits = citations. The dependent variable is log annual wage income. The results are from a fixed-effects estimation. The first column shows the results when controlling for firm citations and their six lags, the second column controlling for firm average wage, the third column for the sample where R&D intensity is observed, and the last column when controlling for firm R&D intensity. In addition to these firm-level control variables, the regressions include all the same control variables as before: age, gender, tenure, months employed, firm size, dummies for the field and level of education, dummies for the sector of the firm, dummies for the firm's regional location, and year dummies. Robust standard errors below. Statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

results remain as before (column 1 of table 6). Alternatively, in the specification with the expected lifetime citations, we control for the expected lifetime citations received by the firm's patents, again including six lags, and our results remain very similar (column 1 of table 7). The estimated coefficients are slightly smaller, and their significance is weakened a little. For the specification with citations as a categorical variable, we present the results for the top citation category (more than thirty citations relative to no patents), and the results are very close to the ones obtained earlier (column 1 of table 8). Similarly, when we control for the average wage in the firm, our main results on the earnings premiums remain (column 2 of tables 6, 7, and 8). Overall, these results reduce concerns that it is the firm-level rent-sharing effect that is driving our results.

Controlling for the firms' R&D intensity also has little effect on our main results. When we include firm R&D intensity, we have to restrict the sample to those individuals who are employees in companies that are in the R&D data, and thus our sample size falls slightly (now 3,709 observations, 1,622 individuals). Fortunately most of the firms that perform systematic R&D or have patents are included in the R&D data. We present the results for this smaller sample also without the R&D intensity control (column 3) as well as with the control (column 4). These two specifications produce identical results. There is one notable difference that arises when restricting attention to the R&D sample: for patents in the top citation category, the wage premiums appear at the time of the patent grant and the following year (see table 8, columns 3 and 4). This may be an

TABLE 8.—EFFECT OF PATENTS IN THE TOP CITATIONS CATEGORY, CONTROLLING FOR FIRM EFFECTS

	Firm Citations	Firm Wage	RD Sample	RD Intensity
Cits-weighted patents > 30				
In year t	0.101	0.129	0.228**	0.228**
	0.092	0.098	0.098	0.098
In year $t-1$	−0.026	0.007	0.158*	0.158*
	0.094	0.099	0.086	0.086
In year $t-2$	0.054	0.088	0.111	0.111
	0.133	0.139	0.082	0.082
In year $t-3$	0.183	0.23*	0.211***	0.211***
	0.121	0.127	0.073	0.073
In year $t-4$	0.275**	0.314**	0.338***	0.338***
	0.119	0.122	0.102	0.102
In year $t-5$	0.285**	0.319***	0.218**	0.218**
	0.117	0.123	0.086	0.086
In year $t-6$	0.347**	0.339**	0.304**	0.304**
	0.160	0.170	0.152	0.153
Observations	4,938	4,450	3,709	3,709
Individuals	1,789	1,656	1,622	1,622
R^2, within	0.11	0.09	0.11	0.11

Cits = citations. The dependent variable is log annual wage income. The results are from a fixed-effects estimation with citations as a categorical variable. The results are reported for the top citation category (cits > 30). The first column shows the results when controlling for firm citations and their six lags, the second column controlling for firm average wage, the third column for the sample where R&D intensity is observed, and the last column when controlling for firm R&D intensity. In addition to these firm-level control variables, the regressions include all the same control variables as before: age, gender, tenure, months employed, firm size, dummies for the field and level of education, dummies for the sector of the firm, dummies for the firm's regional location, and year dummies. Robust standard errors below. Statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

indication that firms with systematic R&D activities are quicker to learn (and appropriate) the value of their patents and subsequently reward the inventors.

D. Reward Mechanisms

To extend our analysis from the level of returns to inventors to the sources of returns, we do four things. First, we examine if the returns are affected when accounting for the number of co-inventors in a patent. Second, we study whether it is changing of employer that yields the estimated returns. Because patents are public information, the granting of a patent may make the inventors more visible or more valuable to other employees, and returns to inventors could then be realized through job changes. Third, patents are not just a measure of invention: they also dictate who has the intellectual property over a given invention at the time of the patent grant, and (not) owning the intellectual property may affect the return to inventors, keeping the value of the patent constant. These returns may be realized through a variety of mechanisms, such as licensing fees, or through the sale of the intellectual property rights, or simply by increasing the value of the individual in the job market. We therefore study whether the rewards to patenting are different for those who own the intellectual property at the time of the patent grant. Concentrating on ownership of intellectual property at the time of patent grant allows us to capture the returns to inventors generated through subsequent sale of the intellectual property rights. Finally, we change our dependent variable to include capital income. As discussed in section I, if patents are valuable to the employer and producing patents requires effort (which is hard to monitor or measure), the employer may resort to providing incentives that generate capital income as well. It should be noted that since 1995 in Finland, stock options have been taxed as income and not as capital gains and thus are included in the dependent variable in our earlier regressions.[24]

Number of co-inventors. We first investigate whether the returns are dependent on the number of co-inventors in the patents. A quarter of the patents have a single inventor, 25% have two inventors, 19% have three inventors, 13% have four, 7% have five, and the remaining 11% have six to twelve inventors. We calculate the share of the patents granted each year that an inventor is responsible for by dividing each patent by the number of its inventors (i.e., if, in a given year, the individual is an inventor in two patents, one alone and one with one co-inventor, his or her patent count takes the value of 1.5). We do the same for the expected lifetime citations received by the patents. The results are presented in table 9 for both the patent count and the expected lifetime citations. The contemporaneous effect is not statistically different from 0 in either of the specifications, possibly indicating that the contemporaneous reward does not depend on the number of inventors (and that the measure of the inventor's patent share is a poorer measure than the simple patent count). The estimated coefficients on the fourth to sixth lags go up as expected if the returns to a patent are divided among the inventors. However, for the patent count, the coefficients are now significant only at the 10% and 15% levels. This could be due to measurement error, as the true contribution of each inventor is not observed and the measure is based on an equal division of the patent among the inventors. For the specification with the inventors' shares of the expected lifetime citations, the coefficients for the third to fifth lags are almost twice as

[24] Unfortunately, we cannot observe stock options that have not been exercised.

TABLE 9.—EFFECT OF PATENTS AND CITATIONS ON ANNUAL EARNINGS, ACCOUNTING FOR CO-INVENTORS

	Patent Share	Citations Share
Patents granted		
In year t	0.005	0.0008
	0.028	0.0023
In year $t-1$	−0.039	−0.0009
	0.038	0.0028
In year $t-2$	−0.063	0.0017
	0.050	0.0028
In year $t-3$	0.014	0.0061*
	0.038	0.0033
In year $t-4$	0.0606*	0.0085***
	0.034	0.0033
In year $t-5$	0.067	0.0072**
	0.047	0.0032
In year $t-6$	0.049	0.0063
	0.031	0.0042
Observations	4,938	4,938
Individuals	1,789	1,789
R^2, within	0.07	0.08

The dependent variable is log annual wage income. The results are from a fixed-effects estimation. The results on the left are from a specification where patents are measured as the inventor's share of patents where he or she was an inventor. The results on the right are from a specification where citations are measured as the inventor's share of citations where he or she was an inventor. The regressions include the same control variables as in the main specification: age, gender, tenure, months employed, firm size, dummies for the field and level of education, dummies for the sector of the firm, dummies for the firm's regional location, and year dummies. Robust standard errors below. Statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

high as for the specification ignoring the number of co-inventors, and they are also statistically significant. Altogether these results indicate that while the contemporaneous reward seems independent of the number of co-inventors, there is some division of the longer-term rewards among co-inventors. An inventor who is alone responsible for a patent gets a wage premium of 5% to 7% to a patent (6% to 8% for every ten citations), while the average return to being part of a patented invention is around 4% to 5% (and 3% to 5% for every ten citations).

Employer changes. Turning then to the question of returns due to employer changes, the data show that about 4% of the individuals change employers in a given year and that over six years (from 1993–1999), 22% of the individuals changed employers at least once. To study the possibility that the returns to inventors are generated through changes in jobs, we include a series of indicator variables and interactions between them and the patent variables to capture the effect of job changes between the year of the patent grant and the year when income is measured. To illustrate, consider an individual who obtained one patent three years ago and changed her job last year. For her, the interaction between the job change indicator and the count of patents obtained three years ago would take the value 1. This interaction allows us to separately identify the returns coming from patents obtained three years ago to individuals who have subsequently changed jobs and those who have not. Adding these variables into the specification containing lags of patent counts, we find that only one of the interactions obtains a significant coefficient; the sixth lag of the interaction carries a negative coefficient. In the specification using citations, none of the interactions obtains a significant coefficient. Furthermore, our point estimates for the patent count and citation variables are virtually unchanged. While this result suggests that actual job changes do not generate any extra returns to inventors, it does not mean that the existence of the possibility of changing jobs would not be a causal factor behind the returns we estimate. Furthermore, some of the coefficients for the direct effects of job change are positive and significant, indicating that job changes (in the past three years) are associated with a wage premium of about 10%.

Assignee type. Analyzing the returns conditional on the type of patent assignee, we find significant differences that indicate that the ownership of intellectual property rights is a mechanism through which the returns to inventors are generated. We separate the patents into two classes: those owned by a company (whether the employer of the inventor or some other) at the time of the patent grant and those owned by the inventor (the patent assignee is an individual). We then reestimate the model with lags of patent counts for both types of patents and for quality-weighted patent counts. The coefficients of the patent variables from a fixed-effects specification are reported in table 10. From that table, it is obvious that the reward structures are different when we condition for ownership: inventors who initially own the patent first forgo some of their earnings (possibly due to efforts in developing and commercializing the invention) but later earn returns higher than those earned by inventors of patents owned by a firm. Patents initially owned by the inventor yield negative returns in the year of the patent grant and the year after that (inventors forgo 7% and 15% of their annual earnings in these years), but later yield returns of circa 30% in the sixth year. The coefficients for the patent count variables when the inventor is not the initial owner are very close to those we obtained earlier (see table 3), with returns in years 4 to 6 after the patent grant lying between 3.5% (sixth year) and 5.1% (fifth year). These results should be interpreted with caution, as time-varying unobservables may be correlated with who becomes the assignee.

These differences in the returns are not explained by the individual-assigned patents being of higher quality. In fact, the mean number of citations for patents assigned to individuals is lower than that for patents assigned to companies, and the distribution is also more skewed; 28% have no cites (19% for corporate owned), and the share of highly cited patents (with more than 20 cites) is only 8% compared to the 20% for company-owned patents. Also, estimating the specification with expected lifetime citations (table 10), the results indicate that the returns to inventor-owned patents are less strongly dependent on the quality (less statistically significant, and coefficients not nearly as large as would be suggested by the average returns to a patent with the average number of citations), although the initial losses appear strongly (negatively) tied to quality.

Table 10.—Effect of Patent Grants on Annual Earnings, by Assignee Type

	Patent Grants		Citations-Weighted Patents	
	Individual	Firm	Individual	Firm
In year t	−0.076*	0.0251*	−0.0088**	0.0020*
	0.042	0.015	0.0035	0.0011
In year $t-1$	−0.127**	−0.003	−0.0202***	0.0000
	0.064	0.022	0.0061	0.0016
In year $t-2$	−0.057	−0.023	−0.0164*	0.0015
	0.080	0.023	0.0090	0.0018
In year $t-3$	0.109	0.011	0.0107	0.0039*
	0.121	0.020	0.0099	0.0020
In year $t-4$	0.039	0.043**	−0.0060	0.0036*
	0.103	0.022	0.0053	0.0019
In year $t-5$	0.204	0.051**	0.0061	0.0043**
	0.125	0.021	0.0040	0.0018
In year $t-6$	0.306**	0.040**	0.0088	0.0047*
	0.149	0.016	0.0065	0.0027
Constant	4.561***		4.561***	
	1.18		1.18	
Observations	4,938		4,938	
Individuals	1,789		1,789	
R^2, within	0.08		0.08	

The dependent variable is log annual wage income. The results are from fixed-effects estimations. The results on the left are from a specification where patents are measured as the number of patent grants. The results on the right are from a specification where patents are weighted by expected lifetime citations. All regressions include dummies for the field and level of education, dummies for the sector of the firm, dummies for the firm's regional location, and year dummies. Statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A possible explanation for the initial negative returns to inventors who own their patents is that after obtaining a patent, they invest in increasing the value of the patent. Such investments could include development of the technology, spending time informing potential buyers about the technology, or organizing the licensing or sale of the patent. Such activities could lead to a short-term decrease in earnings.

Capital income. Finally, turning to the question of whether inventors are rewarded through capital income-generating mechanisms, we reestimate our model by changing the dependent variable to be the logarithm of the sum of wage and capital income (instead of being the logarithm of the former only). Estimating the model with lagged patenting variables (and a fixed-effects estimator), we find that the coefficients of the lags for the fourth to sixth year are significant (the fourth year only at the 7% level, others at the 1% level) with point estimates of 0.038, 0.052, and 0.04. These are all slightly lower than those reported in table 3. Converting these percentage returns to monetary rewards, we find that the monetary rewards at the wage level are almost exactly the same as when including both wage and capital income. When we use the mean wage and capital income over the years 1997 to 1999 as our base, the estimated monetary returns at the wage level are 2,550 euros, in the fourth year after the patent, 3,260 euros, in the fifth and 2,900 euros, in the sixth. These compare to monetary returns of 2,560 euros, 3,500 euros, and 2,700 euros when capital income is included in the dependent variable. It thus seems that the job market does not reward inventors through capital income. One reason we find no extra returns in capital income is probably that stock options are in fact taxed (and reported) as annual wage income.

E. *Robustness*

We perform a number of additional estimations to check the robustness of our results and examine some alternative explanations for them. We examine whether the results are solely due to the IT boom of the late 1990s, affecting only some of the firms and sectors in the sample. Finally, we check that our results are robust to including a random sample of controls; we append a large random sample of individuals who are employed at R&D-performing firms but do not invent and perform all of our estimations for this sample.

Given that the late 1990s (and particularly 1999) was a period characterized by sharply rising market values in the IT sector, it is worth checking whether only these years, or these particular sectors, are the ones when and where inventors earned returns. In order to allow us to keep our specification with all six lags, we remove only the year 1999 from the sample. Doing so hardly affects our results: the coefficients (standard errors) in the patent count specification on the fourth, fifth, and sixth lags are 0.06 (0.03), 0.04 (0.02), and 0.05 (0.02). In the specification with expected lifetime citations, the statistical significance of the coefficients is somewhat reduced, but the magnitudes remain as before. With citations as a categorical variable, patents with twenty or more expected lifetime citations are rewarded with a premium of about 20%.

We also test whether the returns are different for different sectors of the economy. We interact the patent count and its lagged values with variables for the main manufacturing industries in our sample: Machinery and Equipment; Chemicals and Chemical Products; "Radio, TV and Communication; Medical, Precision, and Optical Instruments"; and Fabricated Metal Products. The direct effects of the patent counts remain as before, and only a few significant differences emerge between the sectors. In particular, the medical instruments sector stands out as not providing any medium-term returns to patenting (negative significant coefficients on the interactions with the fourth to sixth lags of the same magnitude as the direct effects). And the IT sector does not stand out as being different from the average.

To check whether the returns we estimate are driven by the few firms that are the largest patenting firms in Finland, we also perform our estimations removing from the sample the observations from the largest three patenting firms. Again, we find that our results do not change much.[25] In fact, in all three specifications, the magnitude of the estimated coefficients goes up.

We also check whether men and women earn similar returns for their inventions. When we allow the returns to be different for women and men (by taking interactions of gender with the patent count and its lags), we find that while the contemporaneous bonus is not significantly different for the genders, the estimated long-run returns are driven by returns to men, not women. The interactions for females are negative and significant and of the same magnitude as the direct patent count effects. In the specification using citations, the interaction terms are not statistically significant.

Finally, our results are robust to including a random sample of noninventors from the same firms. With a sample of over 70,000 individuals (nearly 200,000 observations), all of our qualitative results remain, with the additional result that the coefficient on the third lag is now significant in all of the estimations. The estimated coefficients go up in all specifications: Their magnitudes are 1.3 to 1.5 times the ones from the estimations on the sample of inventors.

F. *Qualitative Evidence*

To deepen our understanding of the reward mechanisms, we interviewed the director of patenting (Harri Honkasalo) at Nokia, the leading Finnish patenter in the United States. According to the interview and the written documents we obtained, the firm rewards its inventors in a fashion that is well in line with our results: there is both an immediate award and potentially larger awards that depend on outcomes realized after the invention.

First, there is an immediate one-time reward at the time of the patent application. If the invention is a minor one, this is the only award. This award is independent of the number of inventors involved in the patent. If the invention is not minor, the inventor receives a royalty award tied to the importance of the invention. There are two types of royalty schemes: fixed and running. The fixed award depends on the turnover of the product into which the invention is incorporated. The largest fixed royalty is twenty times the smallest one. The rules of the award scheme state that the fixed royalty is to be used in “most cases.” If the inventor and the employee are unable to agree on a fixed royalty, a running royalty may be agreed on. The running royalty is calculated based on a formula that takes into account the total number of patents incorporated into the final product, in addition to the turnover of the product (in which the invention is incorporated), the importance of the invention, intellectual property rights (IPR) costs, licensing income, the role of the inventor in the invention, and her or his position in the organization.

The number of inventors of a patent has a direct impact on the royalty (whether fixed or running), with the caveat that each inventor will receive at least 20% of what she or he would have received if she or he had made the invention alone.

VI. Conclusion

The engine of economic growth is technological progress; the engine of technological progress is human inventiveness. We address the question of the returns to individual inventors by estimating the effect of obtaining a U.S. patent on the earnings of Finnish inventors over subsequent years. We investigate the timing and nature of these returns and their dependence on the quality of the invention and the ownership of IPRs.

Our results indicate that there is a close to 3% temporary increase in earnings in the year the patent is granted, probably representing a one-time bonus, and a 4% to 5% increase in earnings three to four years after the patent grant, which remains there for at least the following two years and possibly represents a permanent wage increase. We also find that the returns to being a patent inventor depend strongly on the quality or value of the patent as measured by the expected lifetime citations it receives. Highest-quality patents generate high returns to inventors, while low-quality patents generate no or even negative returns. These quality-dependent returns are first realized three years after the granting of the patent, coinciding with the time it typically takes to learn the value of a patent. We also find that the returns to inventors depend not only on the quality of the invention but also on ownership of intellectual property. Having ownership of the intellectual property when the patent is granted is first associated with negative returns but also associated with later increases. The estimated returns in years 5 and 6 after the patent grant increase four- to sixfold, from around 4% to between 15% and 30%. This result is not explained by quality differences between inventor-owned and other patents.

Our results can thus be summarized in the following three points: First, returns to inventors are very heterogeneous, with low-quality patents yielding no returns and high-quality patents yielding high returns; second, while a patent grant is accompanied by a small bonus reward, the main part of the returns accrues to the inventors only after the quality of the patent is revealed; and third, it is not only

[25] In previous versions of the paper, we reported that the estimated returns are not robust to removing the three largest patenting firms. This result was due to a mistake in the code calculating the firm patent counts and not accounting for missing assignee firm IDs. After correcting the mistake, the result is overturned.

the act of invention that yields returns but also the ownership of the intellectual property associated with them, as the returns to those inventors who own their patents are much higher than the returns to inventors whose employer has the rights on the intellectual property.

The results are consistent with the possible explanations presented in sections I and II. First, firms' optimal design of incentive compensation schemes may be such that they give rewards for observed signals of effort, and patent grants and the revealed quality of the patent in later years work as such signals. Second, patents and, in particular, their later-revealed quality may work as important signals of individual ability, and part of the later wage premium may be a result of the labor market effect of public learning of individuals' ability and productivity. Third, the results are in line with the law on employee inventions in Finland. While we are unable to identify whether the observed returns are due to incentive (hidden action) or signaling (hidden type), our results provide a justification for economists' interest in studying the pecuniary incentives to invent.

The results indicate that incentive mechanisms for inventors in Finland are such that they promote invention and direct effort toward high-valued inventions. As Finland is one of the countries that has improved its rate of invention, measured by U.S. patents, the most over the last decades (Trajtenberg, 2001), understanding the role of monetary incentives in bringing this change about may offer lessons of more general applicability.

REFERENCES

Aghion, Philippe, and Jean Tirole, "The Management of Innovation," *Quarterly Journal of Economics* 109 (1994), 1185–1209.

Andersson, Frederik, Matthew Freedman, John Haltiwanger, Julia Lane, and Kathryn Shaw, "Reaching for the Stars: Who Pays for Talent in Innovative Industries?" *Economic Journal* 119 (2009), F308–F332.

Bandiera, Oriana, Imran Rasul, and Iwan Barankay, "Social Preferences and the Response to Incentives: Evidence from Personnel Data," *Quarterly Journal of Economics* 120 (2008), 917–962.

Card, David, "Estimating the Return to Schooling: Progress on Some Persistent Econometric Problems," *Econometrica* 69 (2001), 1127–1160.

Farber, Henry, and Robert Gibbons, "Learning and Wage Dynamics," *Quarterly Journal of Economics* 111 (1996), 1007–1047.

Griliches, Zvi, "Patent Statistics as Economic Indicators," *Journal of Economic Literature* 28 (1990), 1661–1707.

Grönqvist, Charlotta, "The Private Value of Patents by Patent Characteristics: Evidence from Finland," *Journal of Technology Transfer* 34 (2009), 159–168.

Hall, Bronwyn, Adam B. Jaffe, and Manuel Trajtenberg, "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools," in Adam Jaffe and Manuel Trajtenberg, (Eds.), *Patents, Citations and Innovations: A Window on the Knowledge Economy* (Cambridge, MA: MIT Press, 2002).

——— "Market Value and Patent Citations," *RAND Journal of Economics* 36 (2005), 16–38.

Harhoff, Dietmar, and Karin Hoisl, "Institutionalized Incentives for Ingenuity: Patent Value and the German Employees' Inventions Act," *Research Policy* 36 (2007), 1143–1162.

Harhoff, Dietmar, Francis Narin, Frederic M. Scherer, and Katrin Vopel, "Citation Frequency and the Value of Patented Inventions," this REVIEW 81 (1999), 511–515.

Holmström, Bengt, "Moral Hazard and Observability," *Bell Journal of Economics* 10 (1979), 74–91.

Kaiser, Ulrich, Hans Christian Kongsted, and Thomas Rønde, "Labor Mobility and Patenting Activity," mimeograph (2009).

Khan, B. Zorina, and Kenneth L. Sokoloff, "'Schemes of Practical Utility': Entrepreneurship and Innovation among 'Great Inventors' in the United States, 1790–1865," *Journal of Economic History* 53 (1993), 289–307.

Kim, Jinyoung, Sangjoon John Lee, and Gerald Marschke, "Research Scientist Productivity and Firm Size: Evidence from Panel Data on Inventors," mimeograph, SUNY Albany (2004).

Koreamëki, Ossi, and Tomi Kygrë, "Integrated Panel of Finnish Companies and Workers," VATT discussion paper no. 26, Government Institute of Economic Research, Helsinki (2000).

Lach, Saul, and Mark Schankerman, "Incentives and Invention in Universities," *RAND Journal of Economics* 39 (2008), 403–433.

Lamoreaux, Naomi R., and Kenneth L. Sokoloff, "The Decline of the Independent Inventor: A Schumpeterian Story," NBER working paper no. 11654 (2005).

Lanjouw, Jean O., "Patent Protection in the Shadow of Infringement: Simulation Estimations of Patent Value," *Review of Economic Studies* 65 (1998), 671–710.

Lazear, Edward, "Performance Pay and Productivity," *American Economic Review* 90 (2000), 1346–1361.

Lemieux, Thomas, W. Bentley Macleod, and Daniel Parent, "Performance Pay and Wage Inequality," *Quarterly Journal of Economics* 124 (2009), 1–49.

Lerner, Josh, and Julie Wulf, "Innovation and Incentives: Evidence from Corporate R&D," this REVIEW 89 (2007), 634–644.

Levin, Sharon, and Paula Stephan, "Research Productivity over the Life Cycle: Evidence for Academic Scientists," *American Economic Review* 81 (1991), 114–132.

Mansala, Marja-Leena, "Työsuhdekeksijälle kuuluu korvaus," IPR Info no. 4/2008 (2008). (An employed inventor is entitled to a compensation; in Finnish only.)

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

Pakes, Ariel, "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica* 54 (1986), 755–784.

Pekari, Ari-Pekka, "Työsuhdekeksintöjen merkitys ja työsuhdekeksintöjen syntymiseen vaikuttavat tekijät suomalaisissa teollisuusyrityksissä, Sitra. (The significance of employee inventions and factors affecting employee inventions in Finnish industrial firms; in Finnish only.)

Rossman, Joseph, "The Motives of Inventors," *Quarterly Journal of Economics* 45 (1931), 522–528.

Sauermann, Henry and Wesley M. Cohen, "What Makes Them Tick? Employee Motives and Firm Innovation," *Management Science* 56 (2010), 2134–2153.

Schmookler, Jacob, "Inventors Past and Present," this REVIEW 39 (1957) 321–333.

Stern, Scott, "Do Scientists Pay to Be Scientists?" *Management Science* 50 (2004), 835–853.

Trajtenberg, Manuel, "A Penny for Your Quotes: Patent Citations and the Value of Innovations," *RAND Journal of Economics* 21 (1990), 172–187.

Trajtenberg, Manuel, "Innovation in Isreal, 1968–1987: A Comparative Analysis Using Patent Data," *Research Policy* 30 (2001), 363–389.

Trajtenberg, Manuel, Gil Shiff, and Ran Melamed, "The 'Names Game': Harnessing Inventors Patent Data for Economic Research," CEPR discussion paper no. 5833 (2006).

Van Reenen, John, "The Creation and Capture of Rents: Wages and Innovation in a Panel of U.K. Companies," *Quarterly Journal of Economics* 111 (1996), 195–226.

Wooldridge, Jeffrey, *Econometric Analysis of Cross Section and Panel Data* (Cambridge, MA: MIT Press: 2002).