IRLE WORKING PAPER
#115-14
August 2014

Shedding Light on Inventors' Returns to Patents

Domenico Depalo and Sabrina Di Addario

Cite as: Domenico Depalo and Sabrina Di Addario. (2014). "Shedding Light on Inventors' Returns to Patents". IRLE Working Paper No. 115-14. http://irle.berkeley.edu/workingpapers/115-14.pdf

Shedding Light on Inventors' Returns to Patents

Domenico Depalo Sabrina Di Addario *

August 1, 2014

Abstract

We estimate individual returns to patents using a unique longitudinal administrative dataset on patents and earnings, following individuals and firms for 20 years (1987-2006). We find that inventors' wages steadily increase before patent applications are submitted to the EPO, peak in the year preceding their filing, and then decrease again. We take the fact that earnings peak at *t-1* instead of at *t* as a bureaucratic delay between the time the invention really takes place and the time when the firm submits the application. We also find that the applications that will eventually lead to a granted patent receive a greater wage increase than those who will not. Finally, we use an event study framework to distinguish among inventor-types, and we find that the "star-inventors" (the employees submitting at least three times in their life) receive a lasting wage premium, while the employees with one or two submissions stop receiving the premium after the application date.

JEL classification: O31, J31.

Keywords: Patents; Wages; Incentives; Inventors; Performance pay; Return.

*Domenico Depalo and Sabrina Di Addario, Bank of Italy. We are grateful to David Card, Ernesto Dal Bo, Patrick Kline, Jesse Rothstein, Steven Tadelis, Eliana Viviano and Noam Yuchtman for their excellent comments. We also thank INPS for having provided the employer-employee matched dataset and the team that helped us matching it to Patstat (Carla Bertozzi, Michela Carlana, Fabrizio Marini, Carlo Menon, Giovanna Marcolongo, Michel Serafinelli and Gian Maria Tomat). Finally, we thank the participants to: the Labor Lunch Seminars at the University of Berkeley (November 16, 2012, Berkeley), the Innovation Seminars at the University of Berkeley (November 21, 2012, Berkeley), and the Bank of Italys Seminar (March 20, 2013, Rome). Some of the work for this paper was undertaken while the authors were visiting the Center for Labor Economics and the Institute for Research on Labor and Employment at the University of Berkeley; they thank the Center and I.R.L.E for their hospitality. All the routines will be available at the webpage: http://sites.google.com/site/domdepalo/. The views expressed herein are those of the authors and not necessarily those of the Bank of Italy. Corresponding author: Sabrina Di Addario, Bank of Italy, Rome Branch, Economic Research Unit, Via XX Settembre, 97/E - 00187 Rome, Tel.: +39-06-4792 3458, Fax: +39-06-4792 5626, e-mail: `sabrina.diaddario@bancaditalia.it`

1 Introduction

While the benefits firms derive from patenting have been analyzed in the literature (see, for instance, Hall, Jaffe, and Trajtenberg (2000); Geroski, Machin and Van Reenen (1993); Balasubramanian and Sivadasan, 2011), very little is known about the incentives for workers to invent. Understanding inventors' motivation may be important to increase firms' innovation rates and, ultimately, the micro-foundations of countries' technological progress. Although monetary incentives may not be the primary source of motivation for pure researchers, who are perhaps spurred more by a passion for research or a search for fame, wage premia acknowledge good work and may motivate some inventor types. However, in most countries employees are requested to cede all the property rights of their inventions to their employers, in which case patenting does not necessarily provide inventors with direct monetary compensation. In fact, informal talks with some inventors and the (scant) existing empirical evidence indicate that, on average, workers do obtain a premium after an invention. Accordingly, on the basis of a sample of almost 600 firms observed over 6 years Van Reenen (1996) finds that 4 years after commercialization a major innovation raises firms' aggregate wages by 2 percent. Toivanen and Väänänen (2012) estimate inventors' returns to patents by using individual data on almost $2,000$ Finnish inventors whose inventions were patented in the U.S. over a period of 9 years $(1991-1999)$. They show that inventors obtain a $1-2$ percent premium in the year the firm is granted a patent and a $3-4$ percent wage increase four years after.

In this paper we estimate individual returns to patents by using a unique longitudinal dataset on patents and earnings, following individuals and firms for 20 years $(1987-2006)$. In particular, we build a new concordance between the European Patent Office (EPO) database on patents (Patstat) and the Italian employer-employee matched dataset on individual wages from the Italian Social Security Institute (INPS). To the best of our knowledge we are the first to link the EPO patent data to individual inventors' wages and demographic and job characteristics. Our linking procedure, which is based on an exact matching of inventors' names, location and employer, enables us to match a higher number of inventors than we would have obtained had we used only firms' names, like standard matching procedures with these data.

We estimate individual returns to patents with a Mincerian wage function augmented with patent application indicators. We add to Toivanen's and Väänänen's work, based on a 9-year panel dataset, in many respects. First, the greater length of our panel enables us to study the dynamics of inventors' earnings. Second, we are able to verify empirically whether on average inventors obtain a one-off bonus at the time of a patent submission (as "unobserved effort" theories would predict), a permanent wage increase (in line with the "unobserved ability" literature), or do not obtain any premium (as the "intrinsic motivation" studies would suggest). Third, we are able to discern among patents of different quality (those that will eventually be granted and those that will not) and different inventor types, distinguishing between the premia associated to the first and second patent applications and those accruing to "star-inventors" (which we define as the employees submitting at least three times in our 20-year observational period). Fourth, while in our regressions we always include worker-fixed effects, we are also able to test whether our results persist after controlling for both individuals' and firms' unobserved heterogeneity.

Our results show that on average inventors' wages start increasing a few years before patent applications are submitted to the EPO, peak in the year preceding their filing, and then decrease again. We take the fact that earnings peak at *t-1* instead of at t as a bureaucratic delay between the time the invention really takes place and the time when the firm submits the application. When we test whether there is a difference between the applications that will eventually be granted and those that will not, we find that the inventors who actually obtain a grant earn a larger premium than those who do not (although the patterns are similar in the two cases). Finally, using an event study framework we test whether firms' patenting compensation policies are affected by the number of patent applications each inventor is able to submit in her/his working life. We find that the premia received by the employees with one or two submissions do not persist after the application date. In contrast, firms are willing to provide a permanent wage increase after the third submission. Our results are robust to different sample restrictions and to a double fixed-effect estimation, which makes our estimates consistent even in the presence of omitted time-invariant characteristics pertaining to individuals and/or firms.

Summarizing, the average wage setting scheme for patenting inventors in Italy is the following.

Employers pay inventors a small and increasing premium during the time they undertake research to avoid that they leave the firm or to ensure that they keep putting enough effort each year before the patent submission is completed. Firms are willing to provide a permanent premium, in line with the unobserved ability theories, only to the employees who have proven to be high-ability types, that is the inventors who contributed to at least three patent applications. In contrast, the premia of the workers who have invented just once or twice do not persist after the application date, probably because the employer is still not sure of whether the invention was due to their ability or just to luck.

The paper is structured as follows. The next section presents the related literature, Section 3 the dataset, the variables and the descriptive statistics. Section 4 reports the empirical model and discusses the estimation results. Finally, the last section concludes.

2 Related literature

The majority of the literature analyzes the impact of patents at the national, sectoral, or firm level; very few studies examine the benefits accruing to inventors, largely because of lack of data. At the country level, Eaton and Kortum (1999) find that patents raise R&D and growth. Studies at the firm level show that patents increase market value (Griliches, Pakes and Hall (1986) and Hall, Jaffe, and Trajtenberg, 2000), profitability (Geroski, Machin and Van Reenen, 1993), R&D (Arora, Ceccagnoli, and Cohen, 2008), and are positively correlated to companies' size, skill and capital intensity, labor productivity and TFP (Balasubramanian and Sivadasan, 2011). In contrast, Boldrin and Levine (2013) argue that the evidence in favor of any effect of patents on innovation or productivity is not conclusive.

The scantness of data on inventors' earnings has certainly limited the research on returns to patents. Indeed, most studies link patent data to information at the firm level. For instance, Van Reenen (1996) finds that firms' average wages increase by 2 percent 4 years after they commercialize a major innovation, using aggregate earnings data on 600 companies. Linking compensation survey data to the NBER patent information from the US Patent and Trademark Office (USPTO), Lerner and Wulf (2007) obtain that offering long-term incentives (i.e. stock options) to corporate R&D heads has a positive effect on firm's patent citations, awards and original patents. Toivanen and Väänänen (2012) link a Finnish employer-employee matched dataset to the NBER patent grant data over the period $1991 - 1999$. They show that, after controlling for individuals' unobservable heterogeneity, inventors obtain a $1 - 2$ percent premium in the year the firm is granted a patent and an average $3 - 4$ percent wage increase four years after. However, in their case the existence of a premium is somewhat expected because in Finland firms are legally obliged to reward the employees who take out a patent.

From a theoretical point of view it is not obvious that patents produce any individual return, since in most countries applicant firms retain entirely the inventions' property rights and are not legally obliged to reward their employees for patenting (except for "occasional inventions"; see Sabbatini, 2011). Thus, if, how and why firms choose to reward inventors even if they do not have to is an interesting issue, hardly examined before.

The agency literature is a useful framework to analyze individuals' incentives to patent: firms (principals) have to decide whether to hire potential inventors (agents) without knowing their ability and/or willingness to exert effort, and inventors have to choose a level of research activity without knowing whether they will be successful. Although employers cannot observe inventors' innate characteristics or behavior, they can design compensation contracts that induce employees to either reveal their type or to behave in the interest of the firm. Contracts on patentable inventions would not be difficult to design because patents are observable outcomes, unlike most of workers' output. They could be explicitly set or discretionary, but in Italy they would be firm-specific, as the Italian law does not establish whether companies should reward the employees who contribute to an invention (unlike the Finnish law, for instance). Firms may or may not commit to a prize structure for the employees who contribute to a patent application or for the inventors who obtain a grant.

Like most datasets, ours does not include information on individual contracts; thus, we will only be able to infer empirically the existence, the structure and size of inventors' patent premia on the average of patenting firms. While patent premia could potentially be in the form of piece rates, bonuses, salary revisions, deferred compensation, promotions, as well as options and profit-sharing, we will be able to assess only whether inventors' earnings increase, on average, in proximity of an invention and whether this increase is temporary or permanent, as our data does include information on earnings from capital. In this light, in this section we will broadly classify theories according to whether they would predict permanent returns to patents, one-off bonuses or no premia at all.

A) Unobserved ability and permanent wage premia.

Models of averse selection, whereby firms cannot access important information on inventors' permanent characteristics (e.g. ability, productivity, basic skills, or competence) are best suited to explain the existence of permanent returns to patents. In an efficiency-wages framework, employers might be willing to offer inventors greater than market-clearing wages for two main reasons.

First, to select the best workers (Lazear, 1986). The idea is that the more able employees benefit more from greater compensation contracts linked to output than the less able ones, so by offering higher pay-for-performance contracts employers are able to select a better distribution of workers. One of the first systems put in place in history to attract the best inventors was created in the Venetian Republic in 1474 by assigning artisans exclusive rights on their inventions (Moser, 2013).

Second, to minimize turnover costs. Since the cost of losing an inventor could be particularly high to the firm, employers might offer economic-rents to retain the key-workers in their company. To have a better understanding of the results that we will present in this paper it is useful to distinguish between two phases: (a) the period inventors are engaged in the research process leading a patent application (which might take a long time), and (b) the period following the patent submission.

(a) Even though initially employers cannot observe inventors' innate ability, after a worker is hired they gradually update their information set by observing his/her output realizations over time. During the research time leading to a patent application current employers are more able to learn about the worker's ability than prospective firms ("private learning"; Farber and Gibbons, 1996). Although in this period rival firms cannot yet observe the employee's performance, current employers might be willing to pay inventors efficiency wages from the time they update their beliefs on their ability to the application date, to prevent that these key-workers leave the firm before completing the patent submission.[1]

(b) When inventors take out a patent they reveal their true ability also to rival firms ("public learning") and thus increase the value of their outside options. In this case returns to patents reflect current employers' attempt to prevent poaching from rivals. Patent compensation would timely occur after the patent submission and would permanently shift inventors' wage curve upwards, in the form of either a salary revision or a promotion.

B) Unobserved effort and temporary wage premia.

Moral hazard models are best suited to explain the existence of temporary returns to patents. Although employers cannot observe workers' research effort, they can induce effort exertion by conditioning payments to the inventor's performance. In particular, they may commit to a contract that provides inventors with a temporary bonus any time the worker contributes to a patent application and/or obtains a grant.

There are a few theoretical frameworks in which firms might find it profitable to pay a premium to inventors to raise their effort.

First, efficiency wages models postulate that workers' effort increases with the size of the rent received (Prendergast, 1999). The idea is that firms encourage employees' good performance by raising the value of the job, which increases the worker's cost of being caught shirking (Krueger and Summers, 1988). Second, in a relational contract context (see Malcomson (2013) for a review) firms may be willing to share their expected surplus from patents to prevent that inventors perceive their lack of benefits from patenting as unfair and will consequently reduce their future creative effort.[2] Third, repeated interactions may create implicit incentives that generate patent premia. In repeated-games firms build a reputation for rewarding inventors, and employees exert enough effort in research activity with the expectation of the reward. Employers would not renege their

[1]Employers might fear more that inventors do not complete their invention than that they disclose their knowledge to rival firms, which is often prevented by non-competing agreements. Note that in the absence of firm competition there would be no enforcement mechanism guaranteeing that employers adjusted wages once they change their beliefs on inventors' true productivity.

[2]In Fehr's and Schmidt's (1999) model the presence of inequity-averse workers induces (selfish) firms to pay a rent in the fear that otherwise workers would retaliate by reducing their effort. In Rabin's (1993) approach firms reward employees to increase their loyalty, fearing that without a compensation workers would feel mistreated and would sabotage them.

obligations (especially if expected rents from patents high), otherwise they would lose credibility and workers would innovate anymore.

C) Intrinsic motivation and no patent premia.

Moral hazard models assume that agents are effort averse. However, inventors may invest in research activity just because they enjoy it. According to some non-economic literature (see Prendergast (1999) for a review and Kreps (1997) for a critique to this approach) offering workers a monetary premium to carry out research activity would *reduce* their intrinsic enjoyment of the task, and thus would *lower* their propensity to patent. The idea is that without compensation inventors attribute the reason of their patenting efforts to their own enjoyment, while when they are paid to undertake the task they unconsciously attribute it to the monetary reward, which, in turn, lowers their intrinsic desire to invent. Thus, according to this view, pay-for-performance contracts inhibit creativity. In this framework, we should observe no wage variation after a patent submission or grant.[3]

A lot of empirical evidence (not necessarily on the market of inventions) shows that pay-for-performance contracts increase output in settings where performance can be measured. Lazear (2000) finds that a switch from hourly wages to piece-rate pay raises average productivity of labor; almost half of the increase is due to incentive effects, the rest to an improved selection of workers. Moreover, firms share the productivity gains with their workforce. Since attempting to innovate necessarily involves failures, Manso (2011) argues that the compensation schemes best tailored at motivating innovation should tolerate early failing and reward long-term success. In this light, standard pay-for-performance schemes punishing failures with low rewards and/or termination may in fact discourage innovation. Ederer and Manso (2012) suggest that the compensation schemes best aimed at fostering innovation should be long-term, should provide job security and timely feedback on performance, and should also be path dependent (i.e. inventors doing well initially but poorly later should earn less than those who perform badly initially but well later).

[3] Note that the absence of patent premia in the short run can be explained also in an economics framework. In career concerns models (see Holmstrom (1982), for instance), workers exert enough effort to build a reputation that enables them to renegotiate future contracts. In this case inventors would exert effort to affect prospective employers' perceptions even without explicit compensation schemes based on performance (because outside offers matter). However, in this case we should observe higher earnings for inventors than for non-inventors in the long-run, while in the non-economic models described above we should not observe any premium even in the long-run.

Although the goal of this paper is not the selection of one of these alternative explanations of the existence of returns to patents, it is useful to verify which of the different compensation schemes is more consistent with our data (Section 4). We argue that one-off premia are more compatible with effort-related theories, while permanent wage increases are more consistent with theories related to ability. However, observing patent premia for a continuous but limited period of time would be consistent with both. Indeed, a temporary premium lasting for a few years could be due to the fact that firms learn about the true ability of an inventor researching on ideas about to be patented and want to retain the employee until he/she completes the work, but it is also compatible with a dynamic moral hazard model, whereby firms may provide a yearly compensation flow to motivate effort for the duration of the research period leading to a patent.

3 The data and descriptive statistics

We link the employer-employee matched data from the Italian Social Security Institute (Istituto Nazionale di Previdenza Sociale, INPS) to Patstat, the European Patent Office (EPO) Worldwide Patent Statistical Database.

INPS is an administrative dataset following all private-sector workers and firms over time. Being an administrative database, it offers many advantages: it is not affected by systematic measurement errors (see Abowd and Card, 1989) nor by systematic unit- or item- non-response; moreover, earnings are not top-coded, in contrast to wages from employer-employee matched datasets in other countries (e.g. Germany). Like other administrative datasets, its disadvantages are related to the scarcity of the available information, which, at the individual level comprises: age, gender, municipality of residence and municipality of birth, work status (blue collar; white collar; manager; other), type of contract (full-time versus part-time) and gross yearly earnings. The information on firms includes: average gross yearly earnings, yearly number of employees, industry, plant location (at the municipality level), date of plant opening and closure.

Patstat contains the universe of patent applications and grants presented at the EPO by any Italian firm since 1978 (when Italian companies started applying at the EPO). The database provides a detailed description of each patent submission, including its title, abstract and technological field, the name and address of residence of all its inventors and applicants (i.e. the firms submitting a patent application and retaining the relative property rights), the dates of application filing, publication and grant obtainment. The EPO releases a new version of Patstat twice a year, but it takes about three years to update its records, thus the most recent data are always incomplete. For this reason we dropped all the patent applications filed in 2007-2009 (the last three available years, as our Patstat version was released in April 2009). In 2009 the stock of Italian firms' applications at the EPO was about 52,500. Besides dropping the 4,523 applications filed after 2007 (of which only 157 had already been granted), we also excluded those presented before 1987 (amounting to 5,342, of which 2,788 granted) because we lack of INPS data for that period. After excluding the applications missing relevant information (e.g. applicants' or inventors' names, application dates, etc.) our EPO dataset comprises 42,699 patent submissions pertaining to 90,743 firm-inventor-patent records, that is about 12,000 firms and 33,895 inventors (an average of 3.6 submissions per firm and 1.6 application per inventor).

Like the USPTO dataset (its American counterpart), Patstat lacks of a firm identifier. Thus, we did a challenging matching-name work to link Patstat companies and inventors to INPS firms and employees. Accomplishing this task is not easy, because of homonymy, spelling errors in the EPO registry, and different filing of business names in the two datasets. The details of our matching procedure are as follows. We merged the datasets in several steps. We first attributed VAT codes to Patstat firms on the basis of the company name and location. To minimize the errors, we verified the code using four alternative datasets (Cebi, Infocamere, INPS, Orbis). We were able to associate a VAT code to 70 percent of Patstat firms (86 percent of which have been assigned a unique VAT code). Subsequently INPS linked Patstat companies to all possible INPS firms that had either the same VAT identifier or the same name and location (at the municipality level). In particular, INPS was able to match 80 percent of the firms to its records (i.e. 9,748 firms, 19,022 plants), distributed across all the Italian regions. To determine which of these were correct matches, we verified whether the inventors appearing in each patent submission were actually employed in the corresponding applicant firm matched by INPS. To this aim, INPS searched for all the inventor names recorded in its registry of employees, and found 23,542 employees with the same name of a Patstat inventor (about 70 percent of the total). For half of these there was also a correspondence between Patstat's and INPS' employer (at the time of the patent submission). Thus, we were able to match one-third of the universe of inventors employed in an Italian firm with at least one submission to the EPO between 1987 and 2006. The 10,353 unmatchiable inventors are either misreported, or not employed (e.g. consultants, self-employed or non-formally employed), or employed in an institution not registered with INPS, most likely in the public sector.[4] Both because private research accounts for about 55 percent of total research in Italy (Istat, 2013) and because INPS deals almost exclusively with the private sector, the majority of our unmatchiable inventors must be public sector employees, rather than mis-reported and self-employed workers. Since this paper regards the private sector[5] we are confident that selection plays a minor role in our matching procedure. The possibility to exploit the information on employees improves the precision of our matching procedure with respect to the methods used previously (e.g. the NBER patent database), matching patenting firms only on the basis of companies' names and location (see Thoma et al. (2010) for an extensive discussion on how to combine large patent datasets at the firm level).

Our employer-employee matched dataset covers the years 1987-2009. It includes the full work-history of the employees working in a patenting firm that INPS was able to match, even if they moved from / to a non-patenting firm. In total it comprises 13,545 inventors working in 3,633 patenting firms (4,615 establishments). The unbalanced panel includes more than 160,200 observations (on average, about 8,000 inventors per year). When we include all the inventors' co-workers our sample size rises to almost 25,000,000 observations: inventors are 0.7 percent of the patenting firms' workforce.

As known, the distribution of patents per inventor is very skewed. In our dataset almost 40 percent of the inventors contributes to just one patent submission in their life, less than one-fifth to two; just 1.5 percent of the inventors applies more than 30 times (Figure 1, upper panel). The distribution of granted patents is similar, with 30 percent of the inventors not obtaining any grant

[4]In period analyzed there were two main social security institutions for employees Italy: INPS and INPDAP. The former dealt with most of the private sector and the latter with most of the public sector (including universities).

[5]Note that we eliminated the small fraction of the matched inventors employed in the public sector (less than 2 percent).

within our observational period, and almost one-third being granted just one (Figure 1's lower panel). The average number of years for an application to be granted is about 4.4 for the whole period, although the time to obtain a grant decreased in the last decade (see Table 1).

Table 2 shows the distribution of patent applications by sector. As expected, patents are mostly concentrated in the industrial sector, especially in terms of inventors (98 percent of total inventors are employed in industry). The retail sector accounts for 1.8 percent of submissions; handcraft industries, which tend to be smaller and more traditional than the others, hardly apply. Geographically, most of the applications are concentrated in the North of Italy: Lombardy accounts for more than 40 percent of total submissions (see Table 3 and Figure 2).

Table 4 shows the descriptive statistics for the whole sample, the sub-sample of inventors and a 10-percent random sample of their co-workers. The last column reports the p-values of the t-tests of equality in means between inventors and non-inventors. As expected, individual wages are significantly (much) higher for inventors than for non-inventors. Figure 3 reports the average gross yearly earnings for the two groups of workers (inventors' wages have been re-scaled for an easier comparison). The figure shows that inventors' wages have been increasing over time more markedly than those of the rest of the sample. From a comparison across the second and the third columns of Table 4 it is apparent that inventors tend to be older than the rest of the workers. The inventor's sub-sample exhibits a much lower share of blue-collar workers, slightly less women, and a higher percentage of white-collar workers. Finally, the share of workers with a full-time contract is very similar across the inventor and the non-inventor sub-groups.

4 The empirical model and results

We estimate individual returns to patents with a Mincerian wage function (see Mincer, 1976) augmented with patents:

$$w_{ijt} = \sum_{k=-K}^{K} \delta_{t-k} Patent_{ij,t-k} + X'_{ijt}\beta + \alpha_i + \gamma_t + \epsilon_{ijt}, \quad (1)$$

where w_{ijt} is the gross yearly income (including social security contributions, taxes, overtime work, Christmas bonuses) of firm j's employee i in year t; *Patent* is an innovation proxy, X_{ijt} is a vector of individual observable characteristics; α_i are worker fixed effects; γ_t are year dummies; and ϵ_{ijt} is an error component with zero mean. Errors are always clustered both at the individual and at the firm level (Cameron, Gelbach and Miller, 2011) and always exploit the longitudinal dimension of the data. In what follows we suppress the index j unless we explicitly deal with firm characteristics.

The main objective of this paper is to test whether employees obtain a premium when they contribute to a patent submission, and whether this takes the form of a one-off bonus at the time of invention, in line with the "unobserved effort" theoretical framework, or whether it permanently increases wages, as "unobserved ability" theories would predict. Consistently with this aim we augment the set of standard Mincerian co-variates with *Patent*, a proxy of innovation. Following a large body of the literature we measure innovation with patent submissions (see, for instance, Griliches, 1990). In contrast to Toivanen and Väänänen (2012) we use submitted rather than granted patents for two main reasons. First, because the former are a better and more timely indicator of whether individuals are engaged in research activity (Lotti and Schivardi, 2005). Second, because using submissions instead of grants increases the robustness of our results to endogenous mobility (i.e. the bias arising from inventors switching from one firm to another as a consequence of a successful patent).

For these reasons, in our main model we define $Patent_{ijt}$ as a dummy variable equal to one if the employee i contributed to a patent submitted by firm j at time t. Nevertheless, we also recognize that grants may provide a better quality and a more objective signal on inventors' ability than applications. Therefore, in another specification of the model we distinguish between the submission that will eventually be granted a patent in our observational period and those that will not.[6] Since this paper aims at discerning between competing theories of inventors' wage schedules, in particular between whether they obtain a one-off or a lasting premium after patenting, in all specifications we include 19 forward lags and 19 backward lags of the patent variable ($Patent_{t-k}$ and $Patent_{t+k}$, with k=0,1, . . . 19). Test results on these parameters will support one or the other

[6] Note that to ensure that the EPO had the time to examine all the submissions received we restricted our sample to the period 1987-2006. We are thus confident that our sample is only marginally affected by a truncation problem.

theory. In addition, we control for the number of co-authors of the patent application, because informal talks with some inventors suggest that firms may pay a fixed premium per submission to be split among all co-authors. Finally, we add the number of patent applications that inventors submit each year to test whether more submissions provide higher premia (Hamermesh and Pfann, 2012).

Our observable co-variates include a quadratic form of age, work-status (blue-collar, white-collar, manager or other) and type of contract (full-time or part-time). Ideally, we would need to control for labor market experience and education (see Mincer (1974) and Card, 1999), which are not available. However, this turns out to be a minor issue in our data, because we include workers' fixed effects in all the regressions, exploiting the longitudinal dimension of the dataset. Thus, our estimates are consistent even in the presence of unobservable / unavailable individual time-invariant characteristics, such as personal ability, intelligence, motivation, or education (to the extent that most workers do not improve their education after entering the labor market).[7]

4.1 Main specification

Tables 5–Table 7 report the results. In Table 5 we progressively enrich the model specifications to test whether our coefficients of interest change as the set of covariates increases. In Table 6 we focus on our benchmark specification to evaluate systematic differences across sub-samples, and in Table 7 we use an event study framework to take into account the stock of patent applications each inventor has over our entire observational period.

The patent proxy we use in Table 5 is the dummy variable denoting whether the individual contributed to a patent application in the current year. We normalize the first lead (*t-1*) to zero, so that, all else being equal, the marginal returns to patents at time t with respect to *t-1* is just the coefficient on earnings at t. We test six specifications, progressively adding the observable co-variates. In the first column (5.5) we just control for the patent proxies (including the number of co-authors and the yearly number of patent submissions). In the second specification we add the time dummies to control for the business cycle, in the third one we include a quadratic form

[7] Although we recognize that part of these unobservable characteristics may in fact be time-varying, we are confident that our approach eliminates most of this pitfall (Vella, 1998).

of inventor's age, to proxy human capital accumulation. In the fourth column we control for contract type and work status (white collar, manager, other, versus blue collar) as a proxy for the individual's skills; in specification (5.5) we add some characteristics of the worker's plant, to control for regional and sectoral structural differences in employees' wage compensation (the sectors are: services, handcraft, retail, agriculture; industry is the omitted category). Finally, in our benchmark specification (5.6), we also include firm's size (namely, the number of employees and the number of plants).

Results show that controlling for the time dummies lowers the wage premium, while adding the covariates other than the business cycle has little impact on our coefficients of interest. In the period before submission, the average wage growth amounts to $1-2$ percent, with a small year-on-year increase up to 3 years before the application date and a 3.5 percent premium at *t-1* (with respect to *t-2*; column (5.6)). We run a formal test to verify whether earnings steadily increase in the period before submission and we always reject the null hypothesis that returns are equal across the years preceding submission. A possible interpretation of this result is that employers are aware of their inventors' research progress and thus pay them a yearly premium to prevent them from leaving the company before the patent is completed. Figure 4 shows the wage dynamics (note that we added a constant equal to the wage average in the whole period to the premia reported in column (5.6)). After submission earnings drop, but they never return to their previous level (we always reject the null hypothesis that $Patent_{i,t}$=$Patent_{i,t+1}$=...$Patent_{i,t+k}$). In particular, inventors' earnings fall by 1.9 percent the year of submission, keep dropping for the two following years, and then remain stable. Moreover, the wage curve is not symmetrical: k years before an invention the premium is always lower than k years after (the difference in the premium size between -k and k increases with $\mid k \mid$ and it amounts to $2-5$ percentage points for k$\geq$ 5). More precisely, a formal t-test always rejects the null hypothesis that premia at time *t-k* are equal to premia at time *t+k* (for any k) at standard confidence levels. Inventors' wages peak at time *t-1* rather than at t (after controlling for individual/firms characteristics). A possible reason is the time lag between the time inventors complete an invention (which is when employers reward them) and the firm's application filing date. Indeed, completed inventions must first be analyzed by the firm's legal office, which might take a few months before deciding whether it is worth submitting the patent application to the EPO. Moreover, submissions to the EPO through a national authority might delay further the filing date.

Finally, the variables number of co-authors and number of applications per year have the expected sign but are never statistically significant.

4.2 Granted versus non-granted patent submissions

Our dataset enables us to distinguish between the applications that will eventually be granted a patent from those that will not. This could be a relevant distinction because employers might be able to discern ex-ante the most promising submissions and may decide to reward them accordingly (Table 6). Column (6.1) reports the outcome on the whole sample, corresponding to our benchmark specification ((5.6)). Results show that although the pattern is similar to the benchmark's, the inventors who contribute to a patent application that will eventually be granted experience a higher wage increase than those who do not obtain a grant (the difference between the two types of inventors' premia is about 2 percent large two years after submission). Moreover, the pattern between the two is similar before submission occurs, but afterwards it differs: while the earnings of the employees whose invention will not be granted decrease, those who will eventually be granted a patent exhibit relatively stable wages (see Figure 5 (a) and (b)). We formally test whether the size of premium is the same at time *t-k* and at time *t+k* and we cannot reject the null hypothesis that it is only for the non-granted applications; it is always rejected for the submissions that will eventually be granted. We also test whether k years from/to submission granted and non-granted premia are the same, and we always reject the null hypothesis that they are. This suggests that the applications who will be granted a patent obtain a lasting premium, consistently with the theories of unobserved ability. In contrast, the premium accruing to the submissions who do not obtain the grant does not persist after the application is submitted.

The premium for granted submissions is about 3 percent both at time *t-1* and 3-4 years after the application filing date. The magnitude of these parameters is consistent with Toivanen's and Väänänen's findings except that they are measured at the time of submission, rather than at the time inventors receive a grant, as the authors do.

Our results are robust to restricting the sample to: a) industry, which is the most patent-intensive sector (column (6.2)); b) the North of Italy, which is the most patent-intensive area of the country (specification (6.3)); and c) the industrial sector in the North (specification (6.4)). In specifications (6.5)-(6.6) we split the sectors on the basis of their patent stock in the period 1987-2006. In particular, we define as "high frequency" the sectors accounting for about three quarters of total patents (namely, the chemical sector, buildings, and mechanics; the pharmaceutical sector is analyzed separately) and as "low frequency" the rest. Results indicate that in both cases the structure of the premia is quite similar to that of the full sample. In the pharmaceutical industry (column (6.7)) the number of years preceding submission in which inventors obtain wage increases is greater than the average, probably because in this sector the research process leading to a patent is longer than in other industries.

Our interpretation of the results relies on the fact that inventors move randomly across firms, otherwise any observed wage increase may just be the consequence of the job change, reflecting, for instance, different bargaining conditions or the price of a different employer-employee match. Moreover, if workers moved across firms because of patent submissions, our variable of interest would not be exogenous and our estimates would not be consistent. In the absence of an instrumental variable, we tested whether results are confirmed within the "stayers" (i.e. the workers who never change firm over the studied period), who are 70 percent of the sample (column (6.8)). Results are qualitatively similar to those reported in column (6.1), although the stayers' premium is higher than that of the full sample (and thus of the movers). Thus, if anything, after a firm switch the movers are subject to a wage reduction with respect to what they would have earned had they stayed in the previous company, which attenuates the possible concern that our results are driven by job changes.

Finally, adding to Toivanen's and Väänänen's work, we run a double fixed-effect estimation (see Abowd, Kramarz and Margolis, 1999) to ensure that our outcome is not affected by unobservable variables at the firm level (e.g. firm-specific compensation policies aimed at inventors). The latter might be particularly relevant, as firm-specific heterogeneity has become increasingly important

(Card, Heining and Kline, 2013). The variance explained by the firms' unobservable heterogeneity is only 2 percent of the total variance; thus, previous results are confirmed also in this case (specification (6.9)).

Summarizing, the inventors who do not obtain a grant stop receiving a premium after the application date. This result is more consistent with the "unobserved effort" theories, which predict a one-off bonus as an incentive to increase inventors' productivity. Conversely, the fact that the inventors whose submission will be granted a patent continue receiving a premium after submission is in line with the "unobserved ability" framework, whereby employees' wages permanently increase after their ability is revealed. This suggests that employers make a distinction between different inventors' types. We will investigate this hypothesis in the next section.

4.3 An event study

A striking result from Section 4.2 is that different types of submissions provide different premia. This novel result leads us to explore along other dimensions of heterogeneity, such as the quality of inventors' types, which we proxy with the number of patent applications each individual contributes to during her/his entire working life. The high skewness of the distribution of patents per inventor (see Section 3) suggests that inventors' productivity is very heterogeneous and might be rewarded differently. In particular, it is possible that employers pay higher patent premia to the "stars-inventors" than to the employees who contribute just to a few inventions, because the former have proven to be high-ability types and might have acquired more bargaining power. Conversely, the amount of information embodied in the first patents may not be sufficient for the employer to judge whether the inventor is high-ability or just lucky. The aim of this section is to test whether inventors of different types (i.e. with a different number of submissions in their life) obtain different patent premia.

The length of our panel enables us to differentiate among the employee types according to the number of patents they submitted in the twenty years we observe, and to measure the premium k years before/after each invention. Thus, we now analyze whether the effect of each inventor's *last* patent on earnings differs between who invents n times in the years 1987-1996 (for n=1,2, . . . 10).

We use an "event study" framework (see, for instance, Jacobson, LaLonde and Sullivan (1993) or Kline, 2012) to test whether the premium changes in the neighborhood of the main event (which we define as the individual's last patent application date), by comparing the earnings developments of the employees who submit at time t to those of the inventors who apply in a different period. To this aim, we reorder the panel in event time. In particular, we now assume that the effect of the n-th invention on an inventor's wage in year t depends on the distance between t and the application date (after controlling for the year dummies). In other words, we assume that after a period of k years, the effect of submitting a patent at time t_1 or t_2 is the same.

We now estimate:

$$w_{it} = \sum_{n} \sum_{k=-19}^{19} \delta_n^k E_{n,it}^k + X_{it}'\beta + \alpha_i + \gamma_t + \epsilon_{it}, \tag{2}$$

where $E_{n,it}^k$ are a set of dummy variables equal to 1 if the employee i applied for the n-th patent (for n=1, 2, 3, and $n \geq 5$) k periods before the current year t (or k periods after, if k is negative), with $k = -19, -18, \ldots, 0, 1, 2, \ldots 19$. δ_n^k is the effect of the n-th submission on wages k periods away from the application date. As in equation (1), X_{it} are the employees' observable time-varying characteristics, α_i the worker fixed effects, γ_t the year dummies and ϵ_{it} the error term.

Table 7 reports the results. In the first column we restrict the sample to the employees who submitted only once in the 1987-1996 period. In the second specification we keep the sub-sample of the inventors with just two patent applications, and we report the premia relative to the second submission (although we also control for the first one). In the third column the event is the third submission; in the fourth specification it is the fifth, estimated on the sample of the inventors with *at least* five applications (while controlling for the previous four submissions).

Figure 6 plots the δ_n^k coefficients over time for each n we tested (with 10 percent confidence intervals). They represent the inventors' mean wage premia after taking out the effect of the other co-variates. Each quadrant of the figure represents a specification of Table 7. The path of the δ coefficients reported in Figure 6 (a) (corresponding to (7.1)) is quite similar to that in Figure 4 (corresponding to specification (5.6)) until the application date (the difference between the two estimates is below 1 percentage point). However, after submission the two curves differ: the

δ coefficients estimated with the event study approach decline more rapidly than those estimated with the panel approach: the difference between the two becomes sizable two years after submission and increases over time (it widens to 10 percentage points from the seventh year). This pattern suggests that the results observed in Table 5 are largely driven by the star inventors.[8] To support this intuition we now investigate whether patent premia differ substantially between inventor types. We thus turn to analyze the employees who invent more than once.

Column (7.2) and Figure 6 (b) report the development of the earning increases of the inventors who submit twice in their life in the neighborhood of the second invention. Similarly to the previous case, wage premia grow up to time t and decrease after, although earnings remain on a higher level than they were before submission (by 2 percentage points after five years). To test whether employers behave differently towards the employees who have already proven to be able to submit a patent, we now compare the size of the premium these individuals obtain at their first and second application and we find that the size of the wage increases is similar (both before and after the application date). Moreover, we compare the premia obtained by these employees the first time to those received by the individuals who submit just once in their life, and find that the difference is negligible before the application date (less than 2 percent), but larger afterwards (increasing from 1 percent one year after to about to 10 percent after 7 years). Thus, even though the employers cannot (obviously) forecast the number of inventions each employee is going to undertake over her/his life, they are able, to some extent, to discern the inventors types ex-ante.

The earnings developments of the employees applying three times differ from the previous cases (after controlling for the previous submissions, all their lags, and the other co-variates reported in equation 2), especially after the submission date (column (7.3) and Figure 6 (c)). In particular, we find that the employees who invent three times earn a permanent wage increase after the application date (the δ coefficients are not statistically different from zero after t, while in specifications (7.1)-(7.2) they were significantly negative). The difference of the premia size between the first, second and third submission of this group of employees is generally negligible (2 percent or less) before the application date; afterwards the wage-gap gradually increases. These results suggest that on average

[8] This is because in column (7.1) the δ coefficients are estimated on the sample of whom submits just once, whereas in specification (5.6) they are computed on the whole sample of inventors.

employers start updating their priors on their inventors' ability at around the third invention.

In contrast to the previous cases, the inventors with five applications or more continue receiving earning increases after the application date (the δ coefficients are statistically positive from t onwards; specification (7.4) and Figure 6 (d)), possibly because of the presence of outstanding inventors in this group.[9]

Thus, our results suggest that employers cannot discern the inventor's type of the employees inventing just once or twice, probably because they are not able to judge whether the patent was the outcome of luck or ability. However, after the third submission they update their priors, recognizing that these inventors are of the high-ability type, and thus provide them with a lasting premium, in line with the unobserved ability theories.

5 Conclusions

This paper analyzes the impact of patents on inventors' wages, using a unique matched employers-employees dataset on patents and earnings, over a time span of 20 years (1987-2006). We test whether a patent application premium exists, and if so, whether it is permanent (as "unobserved-ability" theories would predict) or temporary (in line with the "unobserved-effort" literature).

By exploiting the longitudinal dimension of our data, we estimate a Mincerian equation where most of the unobservable components determining the inventors' wages are absorbed by the employees fixed-effects. Applying standard techniques, our results show the presence of an inversely U-shaped profile centered at the time of application. In particular, our estimates indicate that inventors' wages steadily increase before patent applications are submitted to the EPO, peak in the year preceding their filing, and then decrease again. To further refine our conclusions we also undertake an event study analysis. Because in our case "events" (i.e., inventions) may occur more than once in each inventor's in life, we split the sample on the basis of the total number of patents each employee submits in our entire observational period. We find that the earnings of the inventors contributing to at least three submissions increase up to the application date, after which they

[9] To increase the robustness of our results we also run separate regressions for n=5, 6, . . . 10 (not shown here because of the small sample sizes, but available upon request), confirming the previous findings.

remain flat, in contrast to what occurs to the workers who invent just once or twice in their lives, whose wages lower immediately after submission. A few more points are worth noticing.

First, firms pay a yearly premium up to the time the invention is completed, possibly to retain their key workers until they are able to submit the patent application to the EPO. An alternative explanation of this result is within dynamic moral hazard models, whereby employers cannot observe effort but know that the employee is undertaking (promising) research (which takes time). In this context, firms might pay a yearly premium during the research period to ensure that inventors exert enough effort for the whole duration of the project.

Second, the fact that earnings peak at *t-1* instead of at *t* is probably the result of bureaucratic delays between the time the invention really takes place and the time when the EPO registers it.

Third, while the wage developments before submission are quite homogenous among different groups of inventors, afterwards they become heterogeneous. For instance, the inventors who will eventually be granted a patent receive a greater wage increase than those who will not. Moreover, employers are willing to provide a lasting premium, in line with the unobserved ability theories, only to the employees who have proven to be high-ability types (i.e., on average, the employees who contributed to at least three patent applications). In contrast, the employees who have invented just once or twice do not receive any premium after the application date, probably because the employer is still not sure of whether the invention was due to their ability or just to luck.

References

Abowd, John M., and David Card. 1989. "On the Covariance Structure of Earnings and Hours Changes", *Econometrica*, 57, 411–445.

Abowd, John M., Francis, Kramarz, and David N. Margolis. 1999. "High Wage Workers and High Wage Firms". *Econometrica*, 67, 251–333.

Arora, Ashish, Marco Ceccagnoli, and Wesley M. Cohen, 2008. "R&D and the patent premium". *International Journal of Industrial Organization*, 26, 1153-1179.

Balasubramanian, Natarajan and Jagadeesh Sivadasan. 2011. "What Happens When Firms Patent? New Evidence from U.S. Economic Census Data". *The Review of Economics and Statistics*, 93, 126-146.

Boldrin, Michele and David K. Levine. 2013. "The Case Against Patents". *Journal of Economic Perspectives*, 27, 3–22.

Cameron, Colin, Jonah Gelbach, and Douglas Miller. 2011. "Robust Inference with Multi-way Clustering," *Journal of Business and Economic Statistics*, 29, 238–249.

Card, David. 1999. "The Causal Effect of Education on Earnings", *Handbook of labor economics*, Vol.30, pp.1801–1863.

Card, David, Jorg Heining, and Patrick Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality," *Quarterly Journal of Economics*, 128, 967–1015.

Eaton, J., and S. Kortum. 1999. "International technology diffusion, theory and measurement". *International Economic Review*, 40, 537-570.

Ederer, Florian, and Gustavo Manso. 2012. "Is Pay-for-Performance Detrimental to Innovation?," mimeo.

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

Fehr, Ernst, and Klaus M. Schmidt. 1999. "A Theory of Fairness, Competition, and Cooperation", *The Quarterly Journal of Economics*, 817–868.

Geroski, Paul, Steve Machin and John Van Reenen. 1993. "The Profitability of Innovating Firms," *The RAND Journal of Economics*, 24, 198–211.

Griliches, Zvi. 1990. "Patent Statistics as Economic Indicators: A Survey," *Journal of Economic Literature*, 28, 1661–1707.

Griliches, Zvi, Ariel Pakes and Bronwyn Hall. 1986. "The Value of Patents as Indicators of Inventive Activity," NBER Working Paper no. 2083.

Hall, Bronwyn, Jaffe, Adam, and Manuel Trajtenbergm. 2000. "Market Value and Patent Citations: A First Look," NBER Working Paper no. 7741.

Hamermesh, Daniel, and Gerard Pfann. 2012. "Reputation and Earnings: the Roles of Quality and Quantity in Academe," *Economic Inquiry*, 50, 1–16.

Holmstrom, Bengt. 1982. "Moral Hazard in Teams," *Bell Journal of Economics*, 13, 324–340.

Jacobson, Louis S., LaLonde, Robert, and Daniel Sullivan. 1993. "Earnings Losses of Displaced Workers," *The American Economic Review*, 83, 685–709.

Kline, Patrick. 2012. "The Impact of Juvenile Curfew Laws on Arrest of Youth and Adults," *The American Law and Economics Review*, 14, 44–67.

Kreps, David M. 1997. "Intrinsic Motivation and Extrinsic Incentives," *The American Economic Review*, 87, 359–364.

Krueger, Alan, and Lawrence H. Summers. 1988. "Efficiency Wages and the Inter-Industry Wage Structure," *Econometrica*, 56, 259–293.

Lazear, Edward P. 1986. "Salaries and Piece Rates," *Journal of Business*, 59, 405–431.

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

Lerner, Josh and Julie Wulf. 2007. "Innovation and incentives: Evidence from Corporate RD," *The Review of Economics and Statistics*, 89, 634-644.

Lotti, Francesca, and Fabiano Schivardi. 2005. "Cross-country differences in patent-propensity: a firm-level investigation," *Giornale degli economisti*, 64, 469–502.

Malcomson, James M. 2013. "Relational Incentive Contracts". In Robert Gibbons and John Roberts (eds.), *Handbook of Organizational Economics*, Princeton University Press, Princeton, pp. 1014–1065.

Manso, Gustavo. 2011. "Motivating Innovation," *The Journal of Finance*, 66, 1823–1860.

Mincer, J. 1974. *Schooling, Experience and Earnings.* New York: Columbia University Press.

Moser, Petra. 2013. "Patents and Innovation: Evidence from Economic History". *Journal of Economic Perspectives*, 27, 23–44.

Prendergast, Canice. 1999. "The provision of incentives in firms". *Journal of Economics Literature*, 37, 7–63.

Rabin, Matthew. 1993. "Incorporating Fairness into Game Theory and Economics," *The American Economic Review*, 83, 1281–1302.

Sabbatini, Daniele. 2011. "La tutela delle invenzioni industriali: adeguatezza della regolazione e possibili linee evolutive," *mimeo.*

Thoma, Grid, Salvatore Torrisi, Alfonso Gambardella, Dominique Guellec, Bronwyn H. Hall and Dietmar Harhoff. 2010. "Harmonizing and Combining Large Datasets - An Application to Firm-level Patent and Account Data," NBER Working Paper no. 15851.

Toivanen, Otto, and Lotta Väänänen. 2012. "Returns to Inventors," *The Review of Economics and Statistics* 94, 1173–1190.

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

Vella, F. 1998. "Estimating Models with Sample Selection Bias: A Survey" *Journal of Human Resources*, 33, 88–126.

Figure 1: Patent applications per inventor (upper panel) and patent grants per inventor (lower panel)

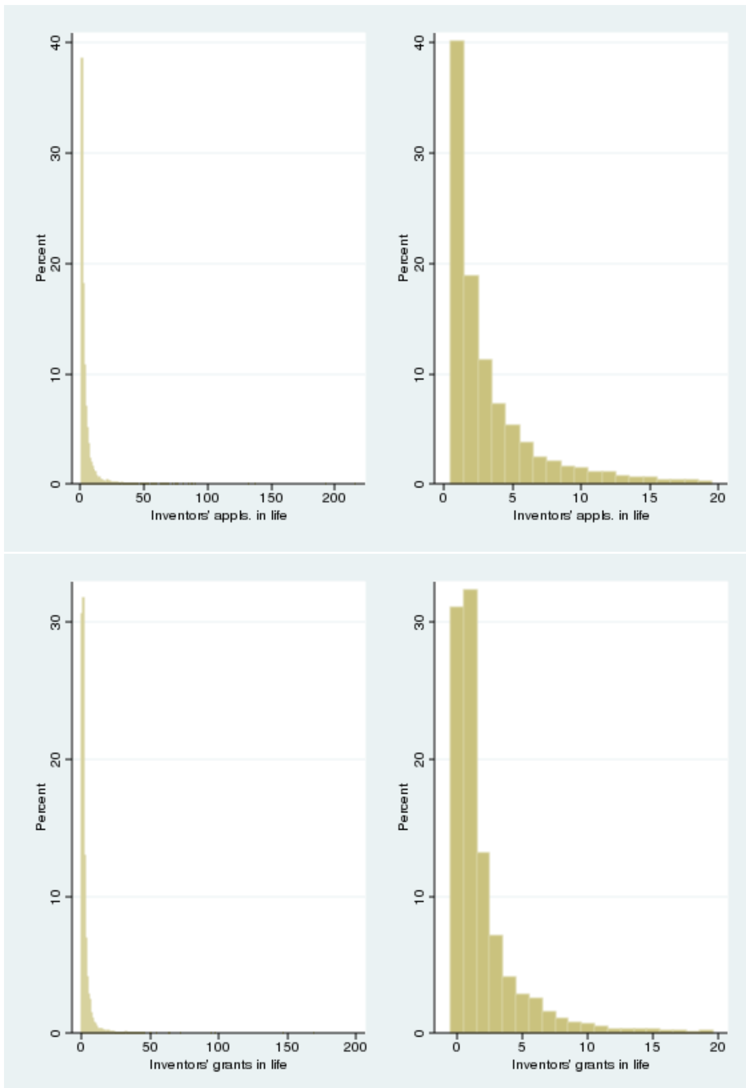

Table 1: Number of years from application to grant

	Applications			Years to grant			
Year	Total	Granted	Not yet granted	Min	Mean	Median	Max
1987	791	585	206	2	4.846154	5	16
1988	850	617	233	2	4.560778	4	9
1989	979	661	318	2	4.503782	4	16
1990	1198	832	366	2	4.497596	4	14
1991	1111	756	355	2	4.367725	4	11
1992	1344	839	505	2	4.296782	4	14
1993	1307	968	339	1	4.532025	4	14
1994	1256	905	351	1	4.701657	4	13
1995	1401	1034	367	1	5.029014	5	14
1996	1572	1177	395	1	5.169074	5	12
1997	1698	1163	535	1	5.28031	5	12
1998	1824	1174	650	1	5.261499	5	11
1999	1800	1162	638	2	5.08864	5	10
2000	2068	1267	801	2	4.797948	4	9
2001	2373	1365	1008	1	4.394139	4	8
2002	2785	1530	1255	1	3.983006	4	7
2003	3035	1434	1601	1	3.864017	4	6
2004	3084	1339	1745	1	3.275579	3	5
2005	3025	902	2123	1	2.978936	3	4
2006	3195	613	2582	1	2.401305	2	3

Table 2: Distribution of patent applications by sector

Sector	Applications (2.1)	Inventors (2.2)	Applicants (2.3)
Industry	97.3	98.0	94.9
Services	0.4	0.2	0.1
Handcraft	0.3	0.1	0.6
Agriculture and Fishing	0.2	0.1	0.2
Retail	1.8	1.6	4.1
Total	100.0	100.0	100.0

Table 3: Distribution of patent applications by region

Region	Applications (3.1)	Inventors (3.2)	Applicants (3.3)
ABRUZZI	1.2	1.1	0.9
BASILICATA	0.3	0.2	0.2
CALABRIA	0.1	0.0	0.1
CAMPANIA	0.7	0.9	1.4
EMILIA-ROMAGNA	11.8	12.4	16.0
FRIULI	5.1	4.1	3.3
LAZIO	4.7	6.2	4.6
LIGURIA	2.1	2.7	2.2
LOMBARDY	41.1	38.6	35.1
MARCHE	2.0	1.9	2.7
MOLISE	0.2	0.1	0.1
PIEDMONT	15.3	17.0	13.0
PUGLIA	0.4	0.4	0.8
SARDINIA	0.2	0.0	0.1
SICILY	2.0	1.5	0.4
TRENTINO ALTO ADIGE	0.9	0.7	1.3
TUSCANY	4.7	4.9	4.5
UMBRIA	0.5	0.3	0.7
VALLE D'AOSTA	0.2	0.1	0.2
VENETO	6.7	7.2	12.6

Figure 2: Distribution of patent applications by region

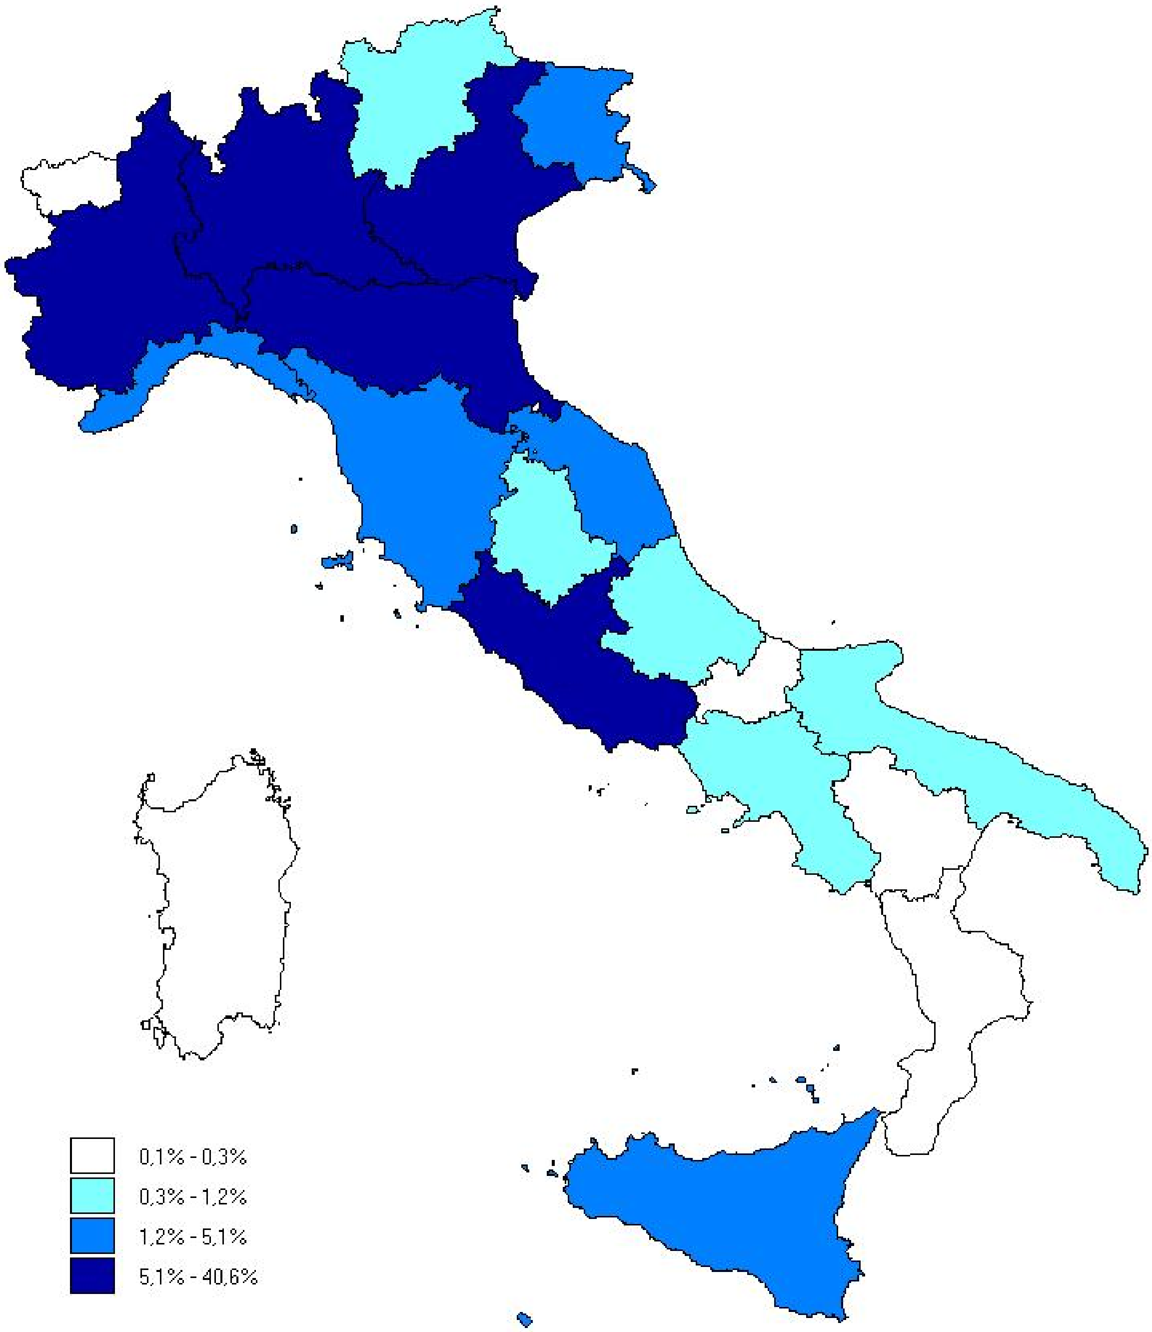

Table 4: Descriptive statistics: employees

	Whole sample (4.1)		Non-Inventors (4.2)		Inventors (4.3)		T-test (4.4)
Variable	Mean	S.D.	Mean	S.D.	P-val.		
Yearly wage	22094.0	18835.5	20712.3	16399.9	40830.6	33787.6	0.000
Female	0.2	0.4	0.2	0.4	0.1	0.3	0.000
Age	39.7	10.3	39.7	10.4	40.0	9.2	0.000
Blue collar	0.5	0.5	0.5	0.5	0.0	0.2	0.000
White collar	0.4	0.5	0.4	0.5	0.6	0.5	0.000
Manager	0.0	0.0	0.0	0.0	0.0	0.1	0.000
Other work status	0.1	0.3	0.1	0.2	0.3	0.5	0.000
Full-time	1.0	0.2	1.0	0.2	1.0	0.1	0.000
No. plants per firm	1.8	8.1					
No. workers per firm	2327140.3	4.6e+07					

T-tests for equality in means between columns (4.2) and (4.3) are reported in the last column.

Figure 3: Inventors' and non-inventors' gross yearly wages

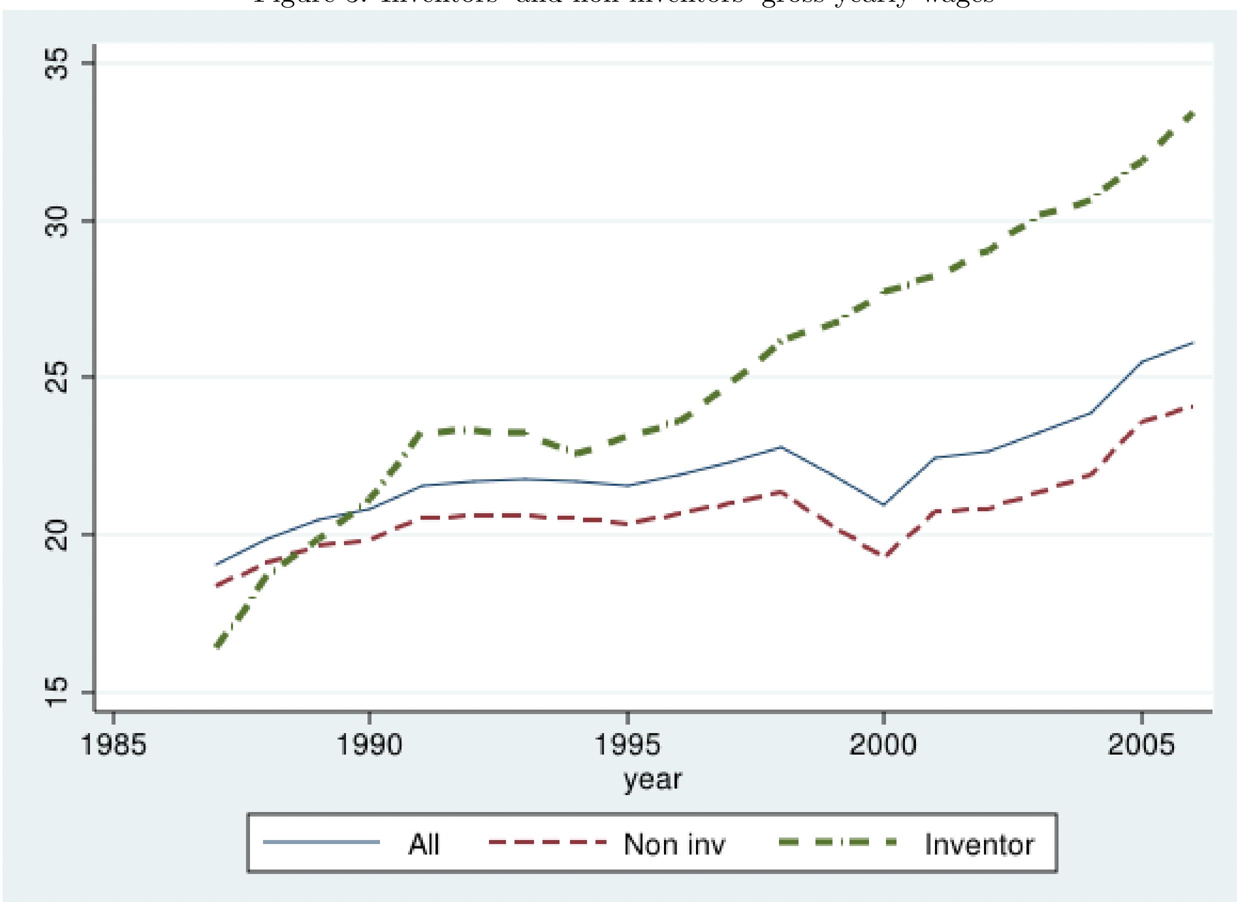

Table 5: Inventors' returns to patents

Variables	(5.1)	(5.2)	(5.3)	(5.4)	(5.5)	(5.6)
No. inventors per submis.	-0.006 ***	-0.002	-0.002	-0.003	-0.003	-0.003
Inventors' appls. per year	-0.021 ***	-0.002	0.003	0.002	0.002	0.002
Year submis.$_{t-8}$	-0.191 ***	-0.106 ***	-0.103 ***	-0.100 ***	-0.100 ***	-0.100 ***
Year submis.$_{t-7}$	-0.169 ***	-0.097 ***	-0.095 ***	-0.093 ***	-0.093 ***	-0.093 ***
Year submis.$_{t-6}$	-0.155 ***	-0.092 ***	-0.091 ***	-0.088 ***	-0.088 ***	-0.088 ***
Year submis.$_{t-5}$	-0.136 ***	-0.085 ***	-0.084 ***	-0.082 ***	-0.082 ***	-0.082 ***
Year submis.$_{t-4}$	-0.118 ***	-0.076 ***	-0.075 ***	-0.073 ***	-0.073 ***	-0.073 ***
Year submis.$_{t-3}$	-0.097 ***	-0.065 ***	-0.064 ***	-0.063 ***	-0.063 ***	-0.063 ***
Year submis.$_{t-2}$	-0.053 ***	-0.036 ***	-0.035 ***	-0.035 ***	-0.035 ***	-0.035 ***
Year submis.$_{t}$	0.046 ***	-0.014	-0.021 **	-0.020 **	-0.020 **	-0.019 **
Year submis.$_{t+1}$	-0.003	-0.034 ***	-0.035 ***	-0.035 ***	-0.035 ***	-0.034 ***
Year submis.$_{t+2}$	-0.019 ***	-0.058 ***	-0.057 ***	-0.057 ***	-0.057 ***	-0.057 ***
Year submis.$_{t+3}$	-0.024 ***	-0.069 ***	-0.065 ***	-0.065 ***	-0.065 ***	-0.065 ***
Year submis.$_{t+4}$	-0.021 ***	-0.074 ***	-0.066 ***	-0.066 ***	-0.066 ***	-0.065 ***
Year submis.$_{t+5}$	-0.030 ***	-0.085 ***	-0.072 ***	-0.072 ***	-0.073 ***	-0.072 ***
Year submis.$_{t+6}$	-0.031 ***	-0.089 ***	-0.072 ***	-0.072 ***	-0.072 ***	-0.071 ***
Year submis.$_{t+7}$	-0.024 ***	-0.087 ***	-0.065 ***	-0.064 ***	-0.064 ***	-0.064 ***
Year submis.$_{t+8}$	-0.022 **	-0.097 ***	-0.069 ***	-0.067 ***	-0.067 ***	-0.067 ***
Individual-FE	YES	YES	YES	YES	YES	YES
Time dummies	NO	YES	YES	YES	YES	YES
Age	NO	NO	YES	YES	YES	YES
Job characteristics	NO	NO	NO	YES	YES	YES
Sector; Region	NO	NO	NO	NO	YES	YES
Firm size; no. plants	NO	NO	NO	NO	NO	YES
No. obs.	160,217	160,217	160,217	160,217	160,217	160,217

Notes: The dependent variable is the yearly wage. Regressions are run on the inventors' sample and use a fixed effect estimation method. Standard errors are always clustered both at the firm and at the employee level. Patent wage premia are computed with respect to the year before submission at EPO. All specifications control for: 19 forward-lags and 19 backward lags of the patent variable (available upon request). Variables denoted with * (**) [***] indicate statistical significance at the 10 (5) [1] percent level.

Figure 4: Inventors' earnings at the time of application (Table 5, column (5.6))

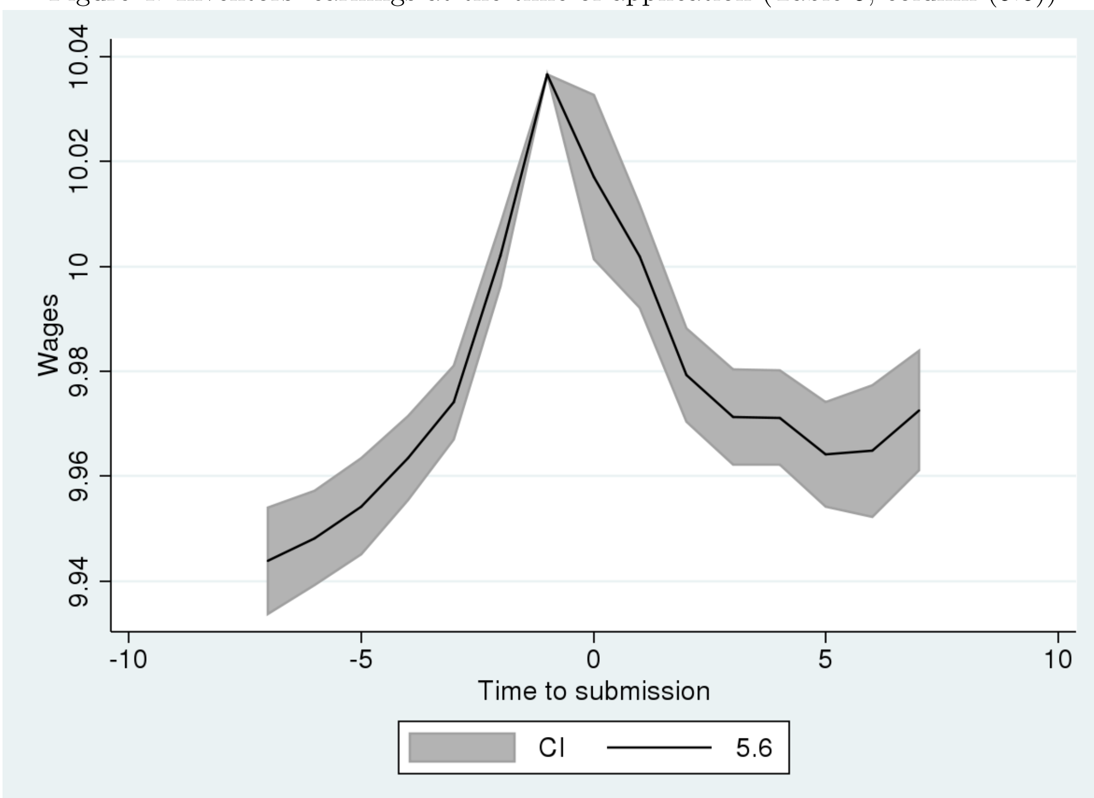

Figure 5: Inventors' earnings at the time of application (granted patents (a); non-granted patents (b))

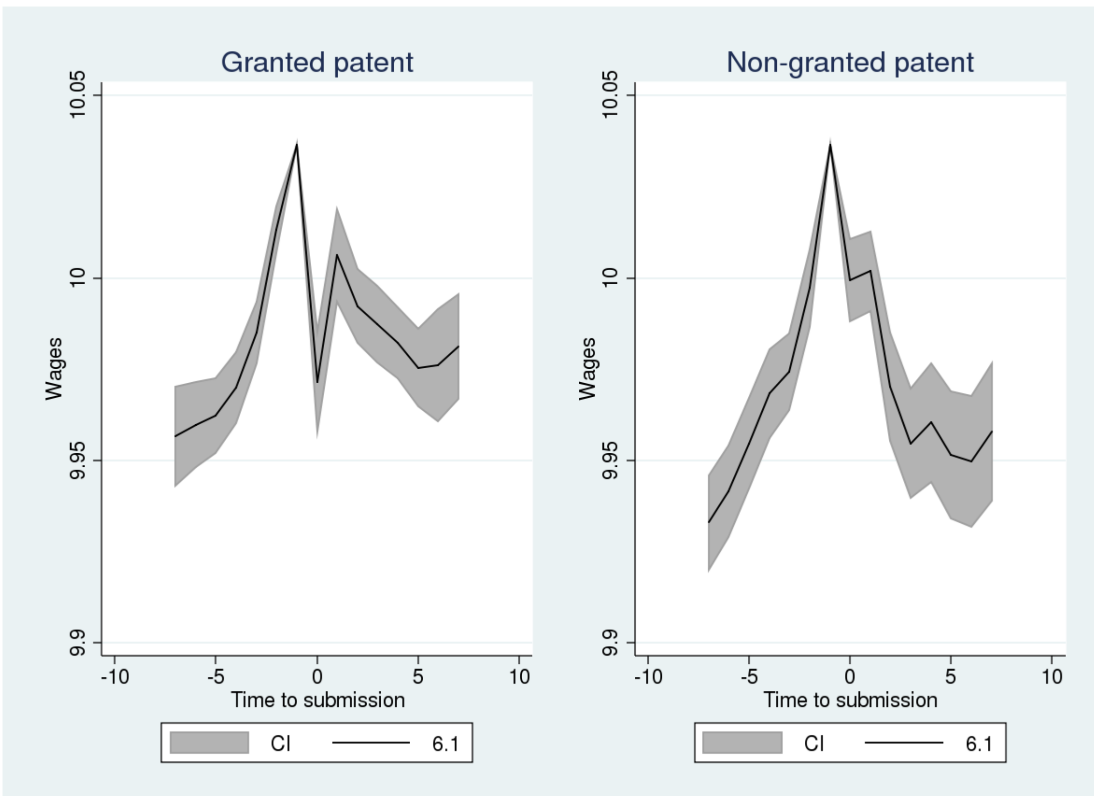

Table 6: Inventors' returns to patents (fixed effect estimation on the inventors' sample)

Variables	Whole (6.1)	Industry (6.2)	North (6.3)	Ind.-North (6.4)	Low Freq. (6.5)	High Freq. (6.6)	Pharmac. (6.7)	Stayers (6.8)	2FE (6.9)
No. inventors per submis.	-0.002	-0.002	-0.002	-0.003	-0.002	-0.000	0.004	0.003	-0.001
Inventors' appls. per year	0.005	0.004	0.005	0.010 *	0.005	0.002	0.005 **	-0.012	0.005 *
(Granted application)$_{t-8}$	-0.090 ***	-0.090 ***	-0.086 ***	-0.098 ***	-0.086 ***	-0.083 ***	-0.081 ***	-0.059 ***	-0.088 ***
(Granted application)$_{t-7}$	-0.081 ***	-0.080 ***	-0.074 ***	-0.088 ***	-0.074 ***	-0.075 ***	-0.070 ***	-0.079 ***	-0.081 ***
(Granted application)$_{t-6}$	-0.078 ***	-0.078 ***	-0.076 ***	-0.083 ***	-0.075 ***	-0.073 ***	-0.082 ***	-0.050 ***	-0.079 ***
(Granted application)$_{t-5}$	-0.076 ***	-0.077 ***	-0.072 ***	-0.076 ***	-0.073 ***	-0.070 ***	-0.069 ***	-0.082 ***	-0.075 ***
(Granted application)$_{t-4}$	-0.069 ***	-0.069 ***	-0.065 ***	-0.068 ***	-0.064 ***	-0.065 ***	-0.070 ***	-0.062 ***	-0.069 ***
(Granted application)$_{t-3}$	-0.056 ***	-0.056 ***	-0.053 ***	-0.058 ***	-0.053 ***	-0.049 ***	-0.058 ***	-0.033 **	-0.056 ***
(Granted application)$_{t-2}$	-0.027 ***	-0.026 ***	-0.026 ***	-0.028 ***	-0.026 ***	-0.025 ***	-0.041 ***	0.008	-0.026 ***
(Granted submis.)$_{t}$	-0.028 ***	-0.029 ***	-0.023 **	-0.034 ***	-0.024 **	-0.013	-0.040	-0.040	-0.031 ***
(Granted application)$_{t+1}$	-0.034 ***	-0.034 ***	-0.034 ***	-0.039 ***	-0.034 ***	-0.027 ***	-0.018	-0.035	-0.032 ***
(Granted application)$_{t+2}$	-0.051 ***	-0.050 ***	-0.047 ***	-0.046 ***	-0.046 ***	-0.049 ***	-0.058 ***	-0.048 **	-0.049 ***
(Granted application)$_{t+3}$	-0.056 ***	-0.056 ***	-0.053 ***	-0.061 ***	-0.053 ***	-0.050 ***	-0.043 ***	-0.060 ***	-0.054 ***
(Granted application)$_{t+4}$	-0.060 ***	-0.060 ***	-0.058 ***	-0.062 ***	-0.057 ***	-0.057 ***	-0.056 ***	-0.049 ***	-0.057 ***
(Granted application)$_{t+5}$	-0.067 ***	-0.067 ***	-0.069 ***	-0.064 ***	-0.069 ***	-0.062 ***	-0.072 ***	-0.057 ***	-0.063 ***
(Granted application)$_{t+6}$	-0.066 ***	-0.066 ***	-0.058 ***	-0.083 ***	-0.058 ***	-0.056 ***	-0.044 ***	-0.033 **	-0.062 ***
(Granted application)$_{t+7}$	-0.059 ***	-0.059 ***	-0.058 ***	-0.078 ***	-0.058 ***	-0.061 ***	-0.044 ***	-0.006	-0.056 ***
(Granted application)$_{t+8}$	-0.056 ***	-0.055 ***	-0.054 ***	-0.067 ***	-0.053 ***	-0.060 ***	-0.049 ***	-0.020	-0.054 ***
(Non-granted application)$_{t-8}$	-0.113 ***	-0.114 ***	-0.109 ***	-0.115 ***	-0.110 ***	-0.108 ***	-0.099 ***	-0.106 ***	-0.110 ***
(Non-granted application)$_{t-7}$	-0.108 ***	-0.110 ***	-0.100 ***	-0.110 ***	-0.101 ***	-0.102 ***	-0.104 ***	-0.104 ***	-0.103 ***
(Non-granted application)$_{t-6}$	-0.101 ***	-0.101 ***	-0.088 ***	-0.110 ***	-0.088 ***	-0.103 ***	-0.083 ***	-0.096 ***	-0.100 ***
(Non-granted application)$_{t-5}$	-0.090 ***	-0.091 ***	-0.083 ***	-0.099 ***	-0.085 ***	-0.092 ***	-0.079 ***	-0.074 ***	-0.090 ***
(Non-granted application)$_{t-4}$	-0.078 ***	-0.079 ***	-0.075 ***	-0.087 ***	-0.076 ***	-0.070 ***	-0.052 ***	-0.088 ***	-0.077 ***
(Non-granted application)$_{t-3}$	-0.071 ***	-0.072 ***	-0.071 ***	-0.079 ***	-0.072 ***	-0.065 ***	-0.056 ***	-0.057 ***	-0.069 ***
(Non-granted application)$_{t-2}$	-0.045 ***	-0.045 ***	-0.044 ***	-0.055 ***	-0.045 ***	-0.038 ***	-0.031 ***	-0.027 **	-0.043 ***
(Non-granted submis.)$_{t}$	-0.018 *	-0.018 **	-0.017 *	-0.027 **	-0.018 *	-0.001	-0.014	-0.045	-0.020 ***
(Non-granted submis.)$_{t+1}$	-0.036 ***	-0.039 ***	-0.035 ***	-0.047 ***	-0.037 ***	-0.032 ***	-0.006	-0.039 *	-0.035 ***
(Non-granted submis.)$_{t+2}$	-0.070 ***	-0.072 ***	-0.068 ***	-0.077 ***	-0.070 ***	-0.069 ***	-0.046 ***	-0.088 ***	-0.071 ***
(Non-granted submis.)$_{t+3}$	-0.085 ***	-0.086 ***	-0.080 ***	-0.092 ***	-0.081 ***	-0.087 ***	-0.060 ***	-0.106 ***	-0.084 ***
(Non-granted submis.)$_{t+4}$	-0.079 ***	-0.079 ***	-0.068 ***	-0.095 ***	-0.067 ***	-0.082 ***	-0.058 ***	-0.072 ***	-0.079 ***
(Non-granted submis.)$_{t+5}$	-0.087 ***	-0.087 ***	-0.084 ***	-0.093 ***	-0.084 ***	-0.094 ***	-0.082 ***	-0.099 ***	-0.089 ***
(Non-granted submis.)$_{t+6}$	-0.089 ***	-0.089 ***	-0.070 ***	-0.104 ***	-0.070 ***	-0.096 ***	-0.074 ***	-0.044 *	-0.089 ***
(Non-granted submis.)$_{t+7}$	-0.080 ***	-0.080 ***	-0.076 ***	-0.081 ***	-0.076 ***	-0.077 ***	-0.065 ***	-0.084 ***	-0.081 ***
(Non-granted submis.)$_{t+8}$	-0.103 ***	-0.101 ***	-0.101 ***	-0.104 ***	-0.101 ***	-0.092 ***	-0.116 ***	-0.049	-0.100 ***
Obs.	160,217	157,136	132,081	102,789	129,832	110,259	40,433	16,643	160,303

Notes: The dependent variable is the yearly wage. Regressions are run on the inventors' sample and use a fixed effect estimation method. Standard errors are always clustered both at the firm and at the employee level. Patent wage premia are computed with respect to the year before submission at EPO. All specifications control for: 19 forward-lags and 19 backward lags of the patent variable, inventors' age, work status, individual fixed effects, sector, region, firm size and time dummies. Variables denoted with * (**) [***] indicate statistical significance at the 10 (5) [1] percent level.

Table 7: Inventors' returns to patents by number of invention: Event studies

Variables	First (7.1)	Second (7.2)	Third (7.3)	Fifth (7.4)
Year submis.$_{t-8}$	-0.106 ***	-0.119 ***	-0.096 ***	-0.107 ***
Year submis.$_{t-7}$	-0.091 ***	-0.109 ***	-0.094 ***	-0.108 ***
Year submis.$_{t-6}$	-0.090 ***	-0.107 ***	-0.085 ***	-0.113 ***
Year submis.$_{t-5}$	-0.092 ***	-0.097 ***	-0.064 **	-0.090 ***
Year submis.$_{t-4}$	-0.078 ***	-0.068 ***	-0.078 ***	-0.089 ***
Year submis.$_{t-3}$	-0.073 ***	-0.056 ***	-0.051 ***	-0.077 ***
Year submis.$_{t-2}$	-0.030 ***	-0.033 ***	-0.056 ***	-0.038 ***
Year submis.$_{t}$	-0.005	-0.019	-0.025	-0.006
Year submis.$_{t+1}$	-0.050 ***	-0.050 ***	-0.007	0.030 ***
Year submis.$_{t+2}$	-0.087 ***	-0.061 ***	-0.027	0.013 ***
Year submis.$_{t+3}$	-0.108 ***	-0.104 ***	-0.090	0.042 ***
Year submis.$_{t+4}$	-0.118 ***	-0.064 **	-0.134 **	0.014 **
Year submis.$_{t+5}$	-0.140 ***	-0.079 **	-0.037	0.045 **
Year submis.$_{t+6}$	-0.164 ***	-0.146 ***	-0.006	0.081 ***
Year submis.$_{t+7}$	-0.161 ***	-0.062	-0.225 *	0.101
Year submis.$_{t+8}$	-0.164 ***	-0.023	-0.546	0.078
No. inventors per submis.	-0.010 **	0.002	-0.008	0.002
Inventors' appls. per year		-0.022	0.021	0.012
Age	0.049 ***			
No. obs.	58,722	29,255	17,871	42,584

Notes: The dependent variable is the yearly wage. Regressions are run on the inventors' sample and use a fixed effect estimation method. Standard errors are always clustered both at the firm and at the employee level. Patent wage premia are computed with respect to the year before submission at EPO. All specifications control for: 19 forward-lags and 19 backward lags of the patent variable (available upon request). Variables denoted with * (**) [***] indicate statistical significance at the 10 (5) [1] percent level.

Figure 6: Inventors' earnings in the neighborhood of submission

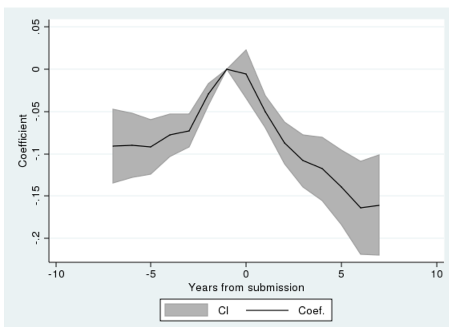

(a) n=1

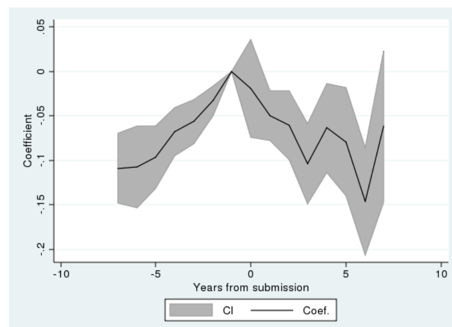

(b) n=2

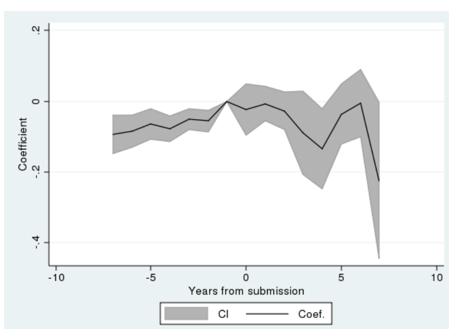

(c) n=3

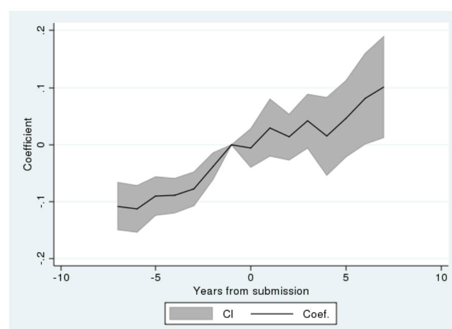

(d) n > 5