NBER WORKING PAPER SERIES

ARE PATENT FEES EFFECTIVE AT WEEDING OUT LOW-QUALITY PATENTS?

Gaétan de Rassenfosse
Adam B. Jaffe

Working Paper 20785
http://www.nber.org/papers/w20785

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2014

de Rassenfosse was supported by an Early Career Researcher grant from the Faculty of Business and Economics at the University of Melbourne. Jaffe was supported by Motu Economic and Public Policy Research. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Are Patent Fees Effective at Weeding out Low-quality Patents?
Gaétan de Rassenfosse and Adam B. Jaffe
NBER Working Paper No. 20785
December 2014, Revised June 2016
JEL No. K2,O31,O34,O38

ABSTRACT

The paper investigates whether patent fees are an effective mechanism to deter the filing of low-quality patent applications. The study analyzes the effect on patent quality of the Patent Law Amendment Act of 1982, which resulted in a substantial increase in patenting fees at the U.S. Patent and Trademark Office. Results from a series of difference-in-differences regressions suggest that the increase in fees led to a weeding out of low-quality patents. About 14 per cent of patents in the lowest quality decile were filtered out, and the effect was especially visible for companies with a large patent portfolio. The study has strong policy implications in the current context of concerns about declines in patent quality.

Gaétan de Rassenfosse
EPFL - CDM - ITPP - IIPP
Odyssea Station 5
CH-1015 Lausanne
gaetan.derassenfosse@epfl.ch

Adam B. Jaffe
Motu Economic and Public Policy Research
PO Box 24390
Wellington 6142
New Zealand
and Queensland University of Technology
and also NBER
adam.jaffe@motu.org.nz

1. Introduction

There is much current concern among both scholars and the business press that the issuance of large numbers of low-quality patents is increasing litigation costs and harming innovation incentives (e.g., Barton, 2000; Bessen and Meurer, 2008; Feldman 2014; The Economist, 2015). Concerns about low-quality patents are particularly acute in the United States and China (Giacopello, 2012; Liang 2012) but the issue is very much a global one. Data by the Organisation for Economic Co-operation and Development suggest that patent quality has declined in the 2000s compared to the 1990s in all advanced economies (OECD, 2011:190). The main patent offices around the world acknowledge the importance of delivering high-quality patents and are committed to improving quality standards.[1]

A potential decline in patent quality raises several inter-related policy concerns. First, the basic welfare tradeoff inherent in patents—incentives to innovate and reveal information balanced against static and dynamic monopoly distortions—is problematic when monopoly is granted for non-novel or obvious inventions. It might seem that the harm from such patents is limited by the likelihood that they are invalid and hence unenforceable, but, Lemley and Shapiro (2005) point out that patents are seldom (in)valid with certainty. This creates a deadweight loss and distorts ex-ante incentives to engage in research (Farrell and Shapiro, 2008). Bessen and Meurer (2008:145) argue along these lines that the patent system has turned from a source of net subsidy to R&D to a net tax. Second, as patents become easier to obtain, the patenting of marginal inventions increases, which may lead to a fragmentation of intellectual property (IP) rights. Fragmentation significantly raises the cost of access to and use of knowledge and may ultimately lower R&D investment (Heller and Eisenberg, 1998). This loss is exacerbated by the cumulative process that is prominent in complex technology industries (Hunt, 2006; Bessen and Maskin, 2009). Finally, the decrease in patent quality creates self-reinforcing operational challenges to patent offices. A perception that marginal applications are likely to be successful encourages more such applications; the resulting increase in the rate of application strains examination resources, likely resulting in continued decline in quality—if the patent office does a rush job to maintain reasonable levels of backlog and pendency (Caillaud and Duchêne, 2011; van Pottelsberghe, 2011).

Most fixes for the quality problem, such as more rigorous examination, or implementing additional options for post-grant review of examiner decisions, require increased resources devoted to maintenance of quality. But the interactions among application decision, the resource cost of examination, and patent quality suggest that reduction in the incentive to apply for low-quality patents through an appropriate fee schedule might significantly reduce the social cost of achieving the desired level of patent quality. Further, there are theoretical reasons to expect that an increase in patent fees would

[1] See, for example, the statements about quality in the 'Four Office Statistics Report 2010 Edition', October 2011, JPO, Tokyo, 82p; USPTO "Patent Quality Chat" http://www.uspto.gov/patent/initiatives/patent-quality-chat.

disproportionately discourage low-quality applications. The objective of this paper is to investigate whether patent fees act effectively ex-ante to screen patent quality, by testing whether a fee increase caused a reduction in the proportion of low-quality patents. This research question fits into the broad literature on the optimal design of patent systems (see, for example, DeBrock, 1985; Matutes et al., 1996; Gallini, 2002; Farrell and Shapiro, 2008; Stiglitz, 2014), and more particularly on the use of fees as a policy tool (see, for example, Scotchmer, 1999; Gans et al., 2004; Caillaud and Duchêne, 2011).

To answer the research question this paper exploits a quasi-natural experiment that occurred in the United States in 1982. To address the declining financial situation of the U.S. Patent and Trademark Office (USPTO) in an era of increasingly tight budgets for federal agencies, Congress passed the Patent Law Amendment Act (PLAA), which resulted in a substantial increase in overall patenting costs. We postulate that the effective fee increase was smaller for foreign firms seeking to extend to the U.S. patent protection already sought for in other countries. This allows us to compare the change in the proportion of low-quality foreign applications to the change for domestic applications in a difference-in-differences (DiD) formulation. We build on the latent quality model by Lanjouw and Schankerman (2004), but we expand on that model to allow for separate latent variables for patent quality and patent value, based on four commonly-used patent metrics.

To anticipate the results, we find evidence of a significant trimming of low-quality patents after the reform. Estimates suggest that 14 per cent of patents in the lowest quality decile were filtered out. The figure reaches 15 per cent for patents in the lowest value decile and 19 per cent for patents in the lowest value quintile. The result is robust to a range of alternative specifications. The effect on quality was particularly strong for companies with a large patent portfolio.

The rest of the paper is organized as follows. Section 2 provides background information on patent quality, the use of fees to screen quality, and the reform. Section 3 presents the data and Section 4 presents the econometric framework and results. The last section offers policy implications.

2. Background

Patent quality and patent value

The concept of patent quality is widely discussed but not always precisely defined. In a theoretical model, patent quality can be defined as the size of the inventive step, i.e. the distance along a technology line that separates the invention from the best pre-existing invention. In a world of certainty and perfect information, the statutory standard for patentability can be represented as a minimum qualifying inventive step: inventions that embody a step of that size or higher are granted and otherwise not. If the inventive step of an invention is uncertain and/or imperfectly observed, then inventions with a bigger inventive step are *more likely* to be granted, but some small inventive-step patents are also granted. Conditional on being granted, a patent on an invention with a larger inventive step is more likely to be upheld as valid if challenged. From this perspective, the problem of patent quality is some combination of patent offices' systematically using too low an inventive step threshold, and their making too many mistakes in evaluating the inventive step of patent applications (de Rassenfosse et al*.,* 2016).

As discussed further below, empirical analysis of patent quality is sometimes unclear on the relationship between patent quality and patent value. One of our goals is to establish an empirical framework for relating these two concepts. Starting first from a theoretical perspective, a higher quality invention is likely to be more valuable, all else equal. First, it is more likely to be upheld as valid if challenged. In addition, though this is somewhat more speculative, it seems likely that all else equal a high-quality invention is likely to generate greater market demand due to its intrinsic degree of technical superiority over pre-existing substitutes. In a similar vein, an invention with a higher inventive step may also be more challenging to invent around, providing a longer monopoly period to its owner.

But of course all else is not equal, and there are several factors that affect value that are not related to invention quality. Some markets are inherently bigger than others, e.g. as between two therapeutic advances of equal technological magnitude, the one that relates to a widespread disease will be more valuable than the one that relates to a rare disease. Further, even for markets of a given size, the elasticity of demand with respect to product attributes can differ, so that a given increment of technological superiority may translate to a larger market share in some markets than in others. Finally, patent value depends on the value of patent protection, as distinct from the value of the patented invention, and this can vary because of the relative efficacy of patent protection and other appropriation mechanisms in different contexts.

In summary, patent quality and patent value are distinct concepts, but likely to be positively correlated. The empirical magnitude of the correlation will depend on the variance of the aspects of value that do not relate to quality relative to the variance in patent quality. We will return below to the empirical challenge of distinguishing quality and value.

Patent quality, patent value and patent fees

As far as we can ascertain, there has been no empirical study of any relationship between fees and quality. Empirical studies on patent fees have focused mainly on estimating the price elasticity of demand for patent applications. Estimates performed on patent filing fees typically vary around -0.3, meaning that a ten per cent increase in fees results in a three per cent decrease in the number of patent applications (de Rassenfosse and van Pottelsberghe, 2012).

Starting with Scotchmer (1999) and Cornelli and Schankerman (1999), economists have simply *assumed* that patent fees can be used to screen quality. Caillaud and Duchêne (2011) explicitly look at patent filing fees in the context of congested patent offices with imperfect examination. They show that there exists a range of values of application fees that lead to a unique high-R&D equilibrium in which firms self-select in their decision to apply for a patent. Picard and van Pottelsberghe (2013) study how the mode of governance of patent offices affects the setting of fees and the quality of the examination process. In their model, the willingness to pay the fees increases with the inventiveness of the patent.

It is theoretically likely that both low-quality and low-value inventions would be disproportionately reduced by an increase in fees. Consider the following. The owner of an invention applies for a patent when the expected benefit from patent protection (probability of grant times economic value of receiving patent protection) exceeds the patenting cost.[2] When fees are increased, some applications for which the expected value was marginally positive will no longer be profitable to pursue. The marginal applications deterred by a fee increase could be marginal for different reasons. Some of them are marginal because the applicant understands that their chances of grant are low, i.e. they are low-quality. Some of them are marginal because the *value* of the patent, even if granted, is not very high, and this could be true in some cases despite the *quality* of the invention being high, for the reasons discussed above.

While it is theoretically expected that a fee increase would screen out both low-quality and low-value inventions, the relative empirical significance of the two is unclear. Further the policy implications are quite different. It is the job of a patent office to screen out low-*quality* inventions, so a fee increase that accomplishes that is normatively helpful. But the stated mission of the patent office is not to screen out inventions that just happen to have low (private) value.[3] Hence from a policy perspective, fees are only successful as a screening device to the extent that they reduce the number of low-quality patents, not just low-value patents. Thus the importance of the empirical challenge of distinguishing quality and value, to which we return below.

The U.S. Patent Law Amendment Act of 1982

Implementation of the PLAA, which resulted in a significant increase in patent fees, provides a useful policy-change framework for studying the effect of fees on patent quality. It led to the largest increase in fees in the history of the USPTO (de Rassenfosse and van Pottelsberghe, 2013) and it occurred sufficiently long ago for ex-post patent quality indicators to be available without truncation. It was also implemented for reasons that are not related to

[2] We return in Section 4.2 to the possibility that patent applications are filed for reasons other than the maximization of expected value.
[3] However, such a policy goal would be desirable if low value inventions resulted from R&D investments that the firm would have made anyway, that is, absent of the patent system.

concerns about quality.[4] At that time, indeed, patent quality at the USPTO was not yet an issue. The fee increase became effective on October 1, 1982.

Patenting costs were affected as follows. Official fees from filing to grant rose from an estimated $239 before the reform (H.R. 96-1307) to $800 after the reform. In addition, the reform also introduced renewal fees, which are due 3.5, 7.5, and 11.5 years from the date of the original patent grant. Renewal fees increase linearly with age, from $400 in year 3.5 to $1,200 in year 11.5. Thus the fees for maintaining a patent to full term rose from $239 to $3,200.

Skeptics may advance several reasons why filing fees would have only a limited impact on patent quality. A first, commonly-heard argument is that patent fees represent only a fraction of patenting cost, which also includes attorney fees. Based on a survey of patent attorneys, Helfgott (1993) reported that patent attorney fees in the United States averaged around $635 in 1992. Assuming that attorney fees followed the evolution of the consumer price index would give a 1983 figure of $440. Thus, while the precise share of patent fees in total application cost cannot be determined, the available evidence suggests that the PLAA fee changes represented a significant increase in total application cost.

Second, patenting costs are usually modest in comparison with R&D costs, such that they would only marginally affect the decision to apply for a patent. Sunk R&D costs should not be relevant for the patent application decision, so this argument must really boil down to an assertion that the expected economic return from patent application is large relative to the patent fees. But, by definition, this is not true for marginal patent applications, so even if for most applications the fees are not a consideration, theory still suggests that they should have an effect on potential applications near the margin of economic viability. These are precisely the low-quality applications for which we wish to measure the impact.

A final argument is that patents are usually filed early in the life of innovation projects, so it is possible that the evaluation error of invention owners at that time is so large as to make the application decision nearly random.[5] In a seminal article, Griliches (1990:1699) discusses precisely this issue. He argues that under perfect information about invention quality, a rise in the cost of patenting would deter the marginal, low-quality inventions. At the other extreme, too large a degree of uncertainty at the time of filing would limit the effectiveness of patent fees as a screening device. Griliches' opinion is explicit:

[4] The PLAA was largely adopted to strengthen the financial resources of the USPTO in an era of increasingly tight budgets for federal agencies. According to a 1980 House Report (H.R. 96-1307), patent fees had not been adjusted since 1967. At that time, the fee structure provided revenue which met 67 per cent of the costs of operating the USPTO. By 1980 inflation had reduced the real value of patent fees, which were estimated to cover a mere 27 per cent of the operating costs (Public Law 97-247).

[5] For instance, Kondo (1999) analyzes the dynamic mechanism of the R&D-patent relationship of Japanese industry and shows that R&D expenditure leads to patent applications with a 1.5 year time-lag. Pakes (1986) estimates an option model of patents and finds evidence of a learning effect early in patent life.

"The truth, I believe, is somewhere in the middle, but closer to the first case, with some definite knowledge about the potential importance of the particular invention."

Thus, at the end of the day, it is an empirical question whether fees are significant enough to affect the application decision. Theory suggests that if there is any effect, it should be seen most clearly for applications at the margin of being worth patenting. The PLAA provides an opportunity to test for this effect.

3. Data and Descriptive Statistics

Figure 1 shows the monthly evolution of the number of USPTO utility patents granted, by application date. The effect of the reform is clearly visible, with a peak in patenting activity in September 1982, immediately followed by a drop in October. This suggests that applicants rushed to file their patent applications before the fee increase, providing a first sign that fees matter. The reform also seems to have had a lasting effect on the demand for patents. The total number of patent applications fell from 116,052 in 1982 to 96,847 in 1983 and 109,010 in 1984. At the same time, total funds for industrial R&D grew by 9 per cent annually, from US 93,496 million in 1982 to US 110,553 million in 1984 (in constant 2000 dollar terms).[6]

Figure 1. Number of utility patents granted by the USPTO, by application month (1981–1984)

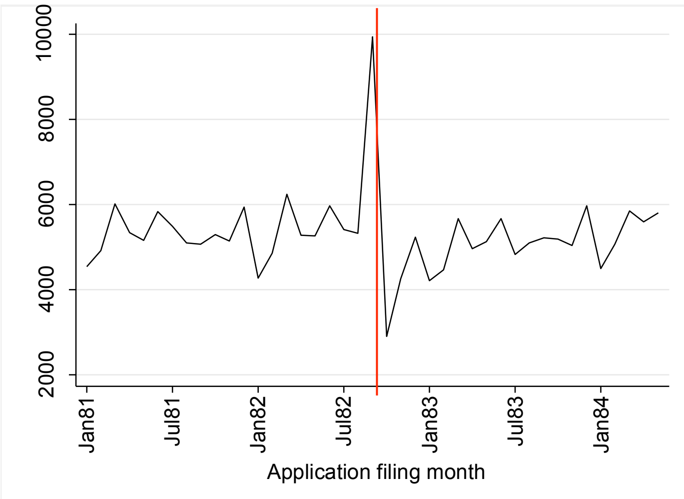

Notes: The vertical line indicates the time at which the fee increase became effective.

[6] See 'USPTO Annual Report FY 1993', Table 6: Patent applications filed (FY 1973–1993); National Science Foundation's '2005 Survey of Industry Research and Development', Table 2: Industrial Research and Development performed in the United States, by source of funds (1953–2005).

3.1 Patent quality indicators

Our theoretical expectation is that low-quality and low-value potential patent applications would have been differentially screened out by the fee increase. We expect that the ex post distribution of quality or value was shifted to the right, and shifted more for domestic patents than for foreign patents. To test for this effect, we need an empirical proxy for patent quality and patent value.

There is an existing literature on indicators of patent quality, but the existing literature often does not distinguish whether the proposed indicator is measuring quality or value. As noted above, anything that contributes to quality can be expected to contribute to value, but the converse is not true: some measures of value may not necessarily reflect quality, to the extent they are driven by other factors such as market size. We use four indicators that have been identified in the literature with patent quality or value: the number of citations received by the patent (Y_1); the number of independent claims at grant (Y_2); the size of the patent family (Y_3); and the number of times the patent was renewed (Y_4). As explained below, we consider citations and claims to be indicative of both quality and value, but we treat family size and renewal as indicative of value but not quality.

The *number of citations* received by a patent has been shown to be a good measure of its technological importance (e.g., Carpenter et al., 1981; Narin et al., 1987; Albert et al., 1991) as well as its economic value (Trajtenberg, 1990). Other authors have also used citation data to estimate the probability that a patent should be granted (Palangkaraya et al., 2011). Recent criticisms have questioned the use of patent citations as measures of knowledge flows, because many citations are added by examiners and not by applicants themselves (Alcácer and Gittelman, 2006). As far as patent quality is concerned, however, there is evidence that examiner citations actually increase the informational content of citation counts (Hegde and Sampat, 2009). See Jaffe and de Rassenfosse (forthcoming) for a review.

The *number of claims* has been used as an indicator of the breadth and the profitability of an invention (Tong and Frame, 1994; Lanjouw and Schankerman, 2004). Claims are the substance of a patent. They codify the description of the invention and constitute the scope of protection in case of grant. In estimates of the values of U.S. patents Bessen (2008) finds that each additional claim increases value by about 2 per cent (as revealed by renewal data).

Patent claims can be *independent* or *dependent*. An independent claim stands by itself in defining a technology that the inventor "claims" has not been previously implemented. A dependent claim represents a further specification or description of subject matter already covered by an independent claim in the same patent.[7] As far as we are aware, all previous

[7] To illustrate with a clearly non-novel example, an independent claim could be "Claim 1: A device for holding stuff up, consisting of a horizontal surface supported by four vertical members of equal length." A dependent

empirical work using claims as a proxy for patent quality have simply counted the total number of claims, independent and dependent. However, by definition dependent claims cover subject matter that is a proper subset of that covered by the parent independent claim, and so the count of independent claims alone should be a better indicator of patent quality in the sense of the magnitude of the inventive step. Hence we utilize the count of independent claims.

The *family size* is the number of jurisdictions in which patent protection is sought. It was first used by Putnam (1996) and Lanjouw et al. (1998). The rationale is that inventions protected by a large international family are of high value given the many costs incurred in the international patent application process. Using data from a survey of patent holders in Germany, Harhoff et al. (2003) report that patents representing large international families correlate particularly well with estimates of patent values. There is, however, no direct path by which family size relates to patent quality, so we consider it an indicator of value but not quality.

The *number of times the patent was renewed* is also a useful indicator of patent value. Most patent offices require the regular payment of renewal fees in order to keep the patent in force. The use of patent renewal data rests on the premise that inventions for which patent protection is more valuable will tend to be protected by payment of renewal fees for longer periods (Schankerman and Pakes, 1986). Renewal fees at the USPTO are due 3.5, 7.5, and 11.5 years from the date of the original patent grant. As explained in section 2, renewal fees were introduced with the reform. However, the change with respect to renewal fees was applied retroactively, such that renewal fees had to be paid for all the patents applied for on or after December 12, 1980. This feature allows us to track the number of renewals for patents both before and after the fee change. As with family size, there is no direct pathway by which renewals are connected to patent quality, so we consider it only as an indicator of value.

The population includes 223,408 utility patents filed in a 21-month period before and after the PLAA, that is, from January 1, 1981 to June 30, 1984. The choice of the start date is motivated by the fact that January 1981 is the first month for which all the patent quality indicators can be constructed. Data sources and technical details are relegated to the Data Appendix.

Table 1 reports the descriptive statistics for the four metrics. Patents in the sample received an average of 5.81 citations in the ten-year period following grant and have 2.10 independent claims. They have an average family size of 4.24, meaning that they were extended to 3.24 jurisdictions besides the United States. The lowest value for the family size

claim in the same patent could be "Claim 2: A device as described in Claim 1, in which the vertical members are made of wood".

is 1, corresponding to a patent that is filed in the United States only. These patents remained valid for 2.80 periods on average, corresponding to an average life of 10.7 years after grant.

Table 1. Descriptive statistics

	Min	Mean	Max	Std. Dev.
Y_1 (citations)	0	5.81	619	7.71
Y_2 (indep. claims)	1	2.10	79	1.78
Y_3 (family size)	1	4.24	49	4.39
Y_4 (renewals)	1	2.80	4	1.10

Notes: N = 222,886.

3.2 Proportion of lower-tier patents before and after the fee increase

Table 2 shows that the proportion of lower-tier patents decreased after the PLAA. For each metric, a cut-off value between 'low' and 'high' was chosen somewhat arbitrarily but in such a way that a relatively small fraction of all patents are classified as 'low': those that have no citation; three or less claims; a family size of one; or never been renewed. The proportion of lower-tier patents is reported before and after the fee change for the population of U.S. patents and for a more restricted sample that is used for further analysis below.

Table 2. Difference in the proportion of low quality patents before and after the fee increase (per cent), four quality indicators

	Low quality (Y/N) according to:			
	Y_1 (citations)	Y_2 (indep. claims)	Y_3 (family size)	Y_4 (renewals)
Population of US patents – all patents by all entities (N=210,041)				
Before	13.39	49.64	37.09	16.10
After	11.31	46.63	33.80	14.71
Δ	2.08*	3.01*	3.29*	1.39*
	(0.14)	(0.22)	(0.21)	(0.16)
DiD sample – priority filings by US large entities and second filings by foreign large entities, entities active in both periods (N=107,490)				
Before	11.99	48.08	30.96	11.27
After	10.23	45.73	25.83	10.21
Δ	1.76*	2.34*	5.13*	1.05*
	(0.19)	(0.30)	(0.27)	(0.19)

Notes: Quality thresholds used: citations = 0; independent claims = 1; family size = 1; renewals = 1. The samples exclude patents filed in the first month before and after the PLAA. See Table A-1 in Data Appendix for sample construction. Standard errors in parentheses. * indicates that the difference in proportion is significantly different from zero at the 0.5 per cent probability threshold.

The proportion of lower-tier patents was markedly lower after the fee increase, for every metric and in both samples. For example, there are 13.39 per cent of patents in the population of U.S. patents that have no citation before the fee increase down to 11.31 after the increase. The difference is 2.08 percentage points and is significantly different from zero at the 0.5 per cent probability threshold.[8] The decrease in the proportion of lower-tier patents is in the range 1.05–5.13 percentage points across all indicators and samples.

Table 2 presents some *prima facie* evidence that the fee increase was associated with a filtering out of low-quality patents. There are, however, three important limitations with the figures presented. First, for any one of the metrics, there may be non-screening reasons for the observed decrease in the number of low-quality patents. For example, the fee increase may have induced firms to cram more claims into each patent application. Second, these metrics may to some extent capture invention value instead of quality. More generally, as with any change over time, there may be other temporally changing factors that are driving the observed effect (such as an inflation of citation rate), rather than the fee increase itself. To deal with these issues, next section presents a framework that optimally combines the information from the multiple indicators to recover an implicit 'true' quality indicator. The framework also utilizes a DiD formulation to isolate the effect of the fee change from other temporal changes.

4. Econometric results

4.1 Non-linear latent variable model of patent quality and value

We build on the latent quality model of Lanjouw and Schankerman (2004). In its original version, the model estimates a linear latent variable that captures variations in observed 'quality' indicators, where 'quality' captures the technological importance of the innovation and market opportunities. We seek to improve the original model in two substantial ways. The first improvement is the use of a *non-linear* latent variable model, which allows better capturing the distributional features of the data. The second improvement is the use of *two latent variables*, $V*$ and $Q*$, capturing the economic value and the technical quality of the inventions, respectively. Our model is more flexible than the original model but flexibility comes at the cost of a significantly greater computational complexity.

We estimate the following model:

$$E[Y_k|Q^*, V^*] = G(C\beta_{k,1} + L\beta_{k,2} + V^*\beta_{k,3} + Q^*\beta_{k,4})$$

where Y_k (k=1,…,4) is the N×1vector of values for the k-th quality indicator, $G(.)$ is a link function, C is the vector 1 (with all entries equal to 1), L is a dummy vector that captures whether the assignee is local, $V*$ is the vector of latent economic value with factor loading $\beta_{k,3}$ with $\beta_{1,3} = 1$ to allow for identification, $Q*$ is the vector of latent technical quality with factor loading $\beta_{k,4}$ with $\beta_{2,4} = 1$ to allow for identification. We further impose identification restrictions in the form $\beta_{k,4} = 0$ for some k as discussed next.

[8] We rely on conservative evidence thresholds for the declaration of significant coefficients (p-value of 0.005) following Johnson (2013). The author shows that commonly-used levels of significance represent only weak evidence in favour of hypothesised effects and argues for the use of more stringent thresholds.

The four indicators are: the number of citations (Y_1); the number of independent claims (Y_2); the size of the patent family (Y_3); and the number of times the patent was renewed (Y_4). We note that the use of the number of independent claims instead of the total number of claims is another improvement over existing studies. Independent claims are directed to the essential features of the invention and, therefore, are better suited to capture information about the scope of the invention.

We assume that Y_k has a negative binomial distribution for k=1,2,3 to account for overdispersion in count data and use a log link function. We assume that Y_4 has an ordinal distribution and use a probit link function. We set two identification restrictions, namely $\beta_{3,4} = \beta_{4,4} = 0$. In concrete terms, these restrictions mean that the technical quality of the invention has no effect on the patent family size and the number of times the patent was renewed—these dimensions being predominantly associated with the economic value of an invention.

Figure 2 depicts the distribution of the latent variables. Both variables are skewed to the right, which conforms to our expectations. The skewness of the latent quality variable is particularly strong, as indicated in the left panel.

Figure 2. Density estimates of the latent quality variables

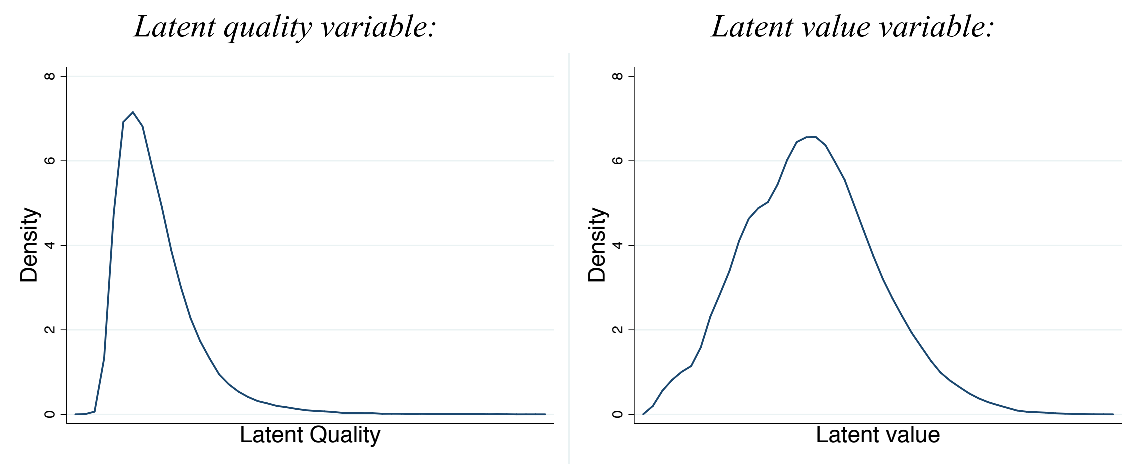

Notes: Epanechnikov kernel with bandwidth ≈ 0.0055. DiD sample used, excluding observations with more than 45 citations or more than 25 family members or more than 20 independent claims (each corresponding to the 0.5 percentile of the respective distributions) for convergence issues (n=107,490). The unit of the latent variables has no meaningful interpretation.

Table 3 presents the correlation coefficients between the latent variables and the patent indicators. The latent quality variable is most strongly correlated with the number of independent claims, followed by the number of forward citations. By construction, the correlation between the quality variable and the family size and renewal indicators is about zero. Note also the positive correlation coefficient of 0.35 between the latent quality variable and the latent value variable. It suggests that higher quality patents are also more valuable. The latent value variable is most strongly correlated with the number of times the patent was renewed, followed by the number of forward citations, family size and the number of independent claims (in that order).

Table 3. Correlation coefficients

	$Q*$	$V*$	Y_1	Y_2	Y_3	Y_4
Q*	1.00					
$V*$	0.35	1.00				
Y_1 (citations)	0.49	0.62	1.00			
Y_2 (independent claims)	0.89	0.23	0.13	1.00		
Y_3 (family size)	-0.05	0.51	0.04	0.01	1.00	
Y_4 (renewals)	0.03	0.71	0.18	0.07	0.12	1.00

Notes: n=107,490. All coefficients are significantly different from zero at the 0.1 per cent probability threshold.

4.2 Difference-in-differences (DiD) estimates

The empirical analysis seeks to quantify the intensity of the weeding out of low-quality patents. We ask: how many low-quality patents did the fee increase weed out, where we define low-quality patents as patents situated in the lowest deciles of the quality distribution. The gist of the identification strategy is to use difference in the sensitivity to fees across groups of applicants. Indeed, foreign applicants face greater overall patenting cost than U.S. applicants: in addition to application cost at home, foreign applicants have to incur application cost at the USPTO and possibly translation cost of their patent document into English. Hence, if patent application fees affect quality, the increase in observed quality should be higher for priority filings by U.S. applicants compared to second filings by foreign applicants. For example, we observe that ten percent of patents by local assignees before the fee increase have a quality $Q* \leq -0.084$, and we seek to estimate the proportion of patents with $Q* \leq -0.084$ after the fee increase for both local and foreign applicants.

Table 4 presents econometric estimates for two thresholds of the latent variables: patents below the first decile in column (1)–(3); and patents below the first quintile (i.e., first and second deciles) in columns (4)–(6). Results in Panel A relate to the latent quality variable whereas results in Panel B relate to the latent value variable. Focusing on Panel A, results in columns (1) and (4) show that fewer patents fell below the set quality thresholds after the reform. DiD estimates are reported in columns (2)–(3) and (5)–(6). Results in column (2) read as follows. By construction, approximately 10 per cent of patents by locals before the PLAA were in the low-quality group. After the fee increase, only 8.8 per cent (=0.097+0.006-0.001-0.014) of patents by locals were in the low-quality group, and we estimate that the reform led to at least a 1.4 percentage points decrease in the proportion of low-quality patents (coefficient associated with variable *post* × *local*). This figure corresponds to a trimming of 14 per cent of low-quality patents. It is a lower bound estimate of the true effect given that the control group of second filings by foreign applicants was also affected by the fee increase. The upper bound estimate is 15 per cent (=0.014+0.001). Column (3) controls for technology and year effects, as well as for changes in the composition of patented technologies after the

PLAA. Again, the treatment effect (*post* × *local*) is negative and highly significant. However, it seems that only the lowest quality patents were affected: the coefficient associated with the treatment effect is negative but not statistically different from 0 when we use the first quintile of the quality variable in columns (5)–(6).

Turning to panel B, the results suggest that the reform also had an effect on the value of patent applications submitted, leading to a 1.5 percentage points decrease in the first decile of value and a 1.9 percentage points decrease in the first quintile. That the PLAA had a stronger and more persistent effect on the value of patents compared to the quality of patents conforms with our expectation, increasing confidence in the relevance of the overall latent variables.

For an external observer the fact that patents are filed for reasons other than maximization of expected value (e.g., some patent applications are used as metric for staff scientist performance) has similar effects on ex-post quality and value metrics than applicants not being able to tell ex-ante which patents are valuable: it limits the effectiveness of fees as a screening device. Thus a fraction of patents in the first quality decile before and after the PLAA could be patents for which the threshold model did not apply. This conclusion supports the staged nature of the schedule of fees, which is characterized but entry fees in the form of application and grant fees and maintenance fees once the patent is granted (Cornelli an Schankerman, 1999; Scotchmer, 1999). However, the result that a non-trivial proportion of low-quality patents was filtered out suggests that at least some firms do have some definite knowledge about the potential importance of their inventions early on.

The approach adopted in this paper is similar in spirit to a quantile regression. The difference is that it does not estimate the value of the first decile for each group but the proportion of patents by each group that lie in the first decile of a reference group (locals before the fee increase). Quantile regressions give qualitatively similar conclusions, namely qualified evidence of increase in quality and strong evidence of increase in value (not reported). However, they are more difficult to interpret in economic terms.

Table 4. DiD estimates of the effect of the fee increase on the latent quality and value

Panel A. Impact on latent quality

	(1)	(2)	(3)	(4)	(5)	(6)
	First decile ($Q^* \leq -0.084$)			First quintile ($Q^* \leq -0.069$)		
local	0.091*	0.097*	0.098*	0.095*	0.098*	0.097*
	(0.001)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)
post	-0.009*	-0.001	0.016	-0.023*	-0.020*	0.010
	(0.001)	(0.002)	(0.012)	(0.002)	(0.003)	(0.019)
local x post		-0.014*	-0.014*		-0.007	-0.007
		(0.003)	(0.003)		(0.004)	(0.004)
IPC			Yes			Yes
IPC x post			Yes			Yes
Year			Yes			Yes
Constant	0.010*	0.006*	0.039*	0.105*	0.103*	0.203*
	(0.001)	(0.001)	(0.008)	(0.002)	(0.002)	(0.012)
R-squared	0.039	0.039	0.058	0.019	0.019	0.051

Panel B. Impact on latent value

	(1)	(2)	(3)	(4)	(5)	(6)
	First decile ($V^* \leq -0.088$)			First quintile ($V^* \leq -0.061$)		
local	0.038*	0.045*	0.043*	0.059*	0.068*	0.065*
	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)
post	-0.013*	-0.004	0.016	-0.014*	-0.003	0.026
	(0.002)	(0.002)	(0.015)	(0.002)	(0.003)	(0.021)
local x post		-0.015*	-0.014*		-0.019*	-0.019*
		(0.003)	(0.003)		(0.005)	(0.005)
IPC			Yes			Yes
IPC x post			Yes			Yes
Year			Yes			Yes
Constant	0.061*	0.057*	0.115*	0.148*	0.143*	0.237*
	(0.001)	(0.002)	(0.009)	(0.002)	(0.002)	(0.013)
R-squared	0.006	0.006	0.023	0.006	0.007	0.031

Notes: N=107,490. Dependent variable is a dummy = 1 if patent quality/value falls in the reference decile/quintile, 0 otherwise. Quantiles estimated for the group of patents by locals before the PLAA. Econometric method is OLS. '*': significantly different from zero at the 0.5 per cent probability threshold.

The average treatment effect can hide important disparities across patentees. In particular, we suspect that the intensity of the response function varies with the size of the patent portfolio held by patentees. To test this hypothesis, we estimate the DiD regression model on four subsamples constructed according to the number of patents held by assignees in the pre-reform period: patents by assignees that had no more than 5 patents; patents by assignees with 6 to 20 patents; patents by assignees with 21 to 100 patents; and patents by assignees with a portfolio size greater than 100 patents. Table 5 only reports the treatment effects for ease of readability.

Table 5. Treatment effect by size of the pre-reform patent portfolio

	(1a)	(1b)	(1c)	(1d)	(2a)	(2b)	(2c)	(2d)
	Latent quality				Latent value			
Portfolio size:	1–5	6–20	21–100	> 100	1–5	6–20	21–100	> 100
Treatment effect at first decile								
	-0.006	-0.007	-0.016*	-0.011*	-0.012	-0.035*	-0.022*	-0.008
	(0.011)	(0.008)	(0.006)	(0.004)	(0.011)	(0.010)	(.007)	(0.004)
Treatment effect at second decile								
	0.016	0.017	-0.009	-0.011	-0.015	-0.054*	-0.034*	-0.001
	(0.017)	(0.014)	(0.009)	(0.006)	(0.016)	(0.014)	(0.010)	(0.006)
N	7959	12,206	22,392	61,543	7959	12,206	22,392	61,543

Notes: coefficients associated with variables *post* × *local* reported. Standard errors in parentheses. '*': significantly different from zero at the 0.5 per cent probability threshold.

The intensity of trimming with respect to patent quality is particularly strong for the very low quality patents (first decile) of large to very large patent owners (more than 20 patents in the study period). The fee increase had no effect on the quality of patents by small and medium patent owners. Regarding the value indicator, the reform had virtually no impact for small and very large patent owners. It had the largest effect for medium and large owners.

4.3 Confounding factors

An important assumption of the DiD regression model is that the quality of patents by locals follows the same trend in the pre-treatment period as that of patents by foreign assignees (the so-called 'parallel-trend assumption'). It ensures that the control group provides an adequate basis for the counterfactual case. We briefly discuss two tests that we performed in order to ensure that the assumption holds, although we do not report them for the sake of brevity. First, we have tested for the presence of lag effects in the DiD regression model by including dummy variables that take the value 1 for patents assigned to locals and filed in the few months preceding the reform and 0 otherwise. Significant coefficients typically provide evidence that the trend of the treatment group started departing from that of the control group before the reform. Coefficients associated with the various lag variables were largely not statistically different from zero, suggesting no significant change in trend before the reform. Second, we have also performed a placebo DiD before the reform. The time window used for the placebo test goes from January 1981 to August 1982 and the variable *post* takes the value 1 for patents filed on or after November 1, 1981. The interaction term was not significantly different from zero, providing additional evidence that controls and treatments did not differ before the reform.

The DiD setting does not follow conventional textbook practice because both treatments and controls were subject to a policy change. Yet, this particular setting does not compromise the validity of the findings: as long as the intensity of the behavioral response to a change in fees differs between groups, this setting provides information about whether patent fees affect quality. However, since the control group was also subject to a change in fees, the treatment effect may underestimate the true effect on quality. In order to gauge the sensitivity of the results to the control group we have estimated the DiD regression model using control patents that were probably among the least sensitive to fees, namely patents by German applicants.[9] Treatment effects were of similar magnitude.

The patent landscape was also changed in 1982 by the creation of the Court of Appeals for the Federal Circuit (CAFC) and its assignment as the sole U.S. appeals court in patent cases. This change eventually affected incentives to apply for patents by lowering the standards for patentability and increasing the value of patent protection. But these changes do not undermine our results because: (1) the ultimate effects of the CAFC on patent practice and enforcement were not seen until at least 1985 (Bender et al., 1986; Strawbridge et al., 1987); and (2) once these effects were known, the effect was to *increase* rather than decrease the incentive to apply for low-quality patents (Hall, 2005; Quillen, 2006)—so if the effect of the CAFC was somehow anticipated during our data period that would cause a conservative bias in our estimates of the impact of the fee change.[10]

Finally, one may question whether the results do not reflect intertemporal substitution rather than a genuine weeding out effect. The main regression model excludes patents filed the month directly before and after the PLAA, where intertemporal substitution is most likely to occur. We have also estimated the regression model on a sample that excludes patents filed two months before and after the reform, with very similar results (not reported). While it seems possible that patentees may rush to file applications before the fee increase, it seems unlikely that they will be able to rush by more than two months in light of the short R&D gestation lag. More generally, results of the specification with lags and the placebo test discussed above indicate that the trends between groups started changing only after the reform, which further suggests the presence of a genuine effect.

5. Concluding remarks

This paper investigates the effect of the U.S. Patent Law Amendment Act of 1982 (which involved a substantial increase in patenting costs) on the quality of patents at the USPTO. The empirical analysis suggests a positive answer to the question asked in the title of the

[9] The group of U.S. second filings granted to German applicants is a strong control group because the German patent system is usually seen as a high-quality system involving a high inventive step (see, for example, Michel and Bettels, 2001, p. 189). In addition, having been substantially changed in 1976 (Mueller and Wegner, 1977), German patent law did not undergo any major reform in the early 1980s. Finally, the total patenting cost for German applicants willing to protect an invention in the United States is much higher than that for U.S. applicants. Helfgott (1993) estimated that the cost of translating a typical patent application from German to English was $2,000 in 1992, equivalent to $ 1,400 in 1983 using the CPI deflator (i.e., more expensive than U.S. attorneys' fees and application fees combined).

[10] Atkinson et al. (2009) provide some evidence that forum-shopping practice may have anticipated the formation of the CAFC.

paper. It presents evidence that applications perceived to be weaker were weeded out. Results from a series of DiD regressions indicate that 14 per cent of patents in the lowest quality decile were filtered out. The figure reaches 15 per cent for patents in the lowest value decile and 19 per cent for patents in the lowest value quintile. However, the increase in quality was not constant across the board. The effect was particularly strong for the largest patentees.

We note that our analysis is guided by but not embedded in a behavioral model of the firm and its environment. As such, the analysis is subject to Lucas' critique (Lucas, 1976), meaning that the results may have limited predictive power. The legal environment has changed since the early 1980s, and the increased emphasis on the alternative, strategic uses of patents has modified patenting practices. It would be erroneous to directly transpose the estimates to the current situation. However, the qualitative message of the empirical analysis is likely to remain: higher patent fees reduce applications for low-quality patents disproportionately.

The study has important implications for intellectual property policy. A fee reform, much like any tax reform, must be studied from three perspectives: revenues, efficiency, and equity. It is well established empirically that an increase in fees will increase the total revenues collected by patent offices because the price elasticity of demand for patents is lower than unity. In the traditional tax context, this revenue benefit must be balanced against an assumed efficiency loss associated with reducing quantities below their supposedly optimal level. But for the reasons discussed in the introduction, to the extent that the reduction in patent applications effectively screens out low-quality patents, the resulting reduction in processing time, and increase in the average quality of patents, likely improves the overall functioning of the patent system. Regarding the equity dimension, the results presented in this paper show that the fee increase had no impact on the quality (to a large extent) and value (to some extent) of patents by small patentees. The intensity of trimming with respect to patent quality was particularly strong for the very low quality patents of large to very large patent owners, which is a desirable policy outcome as current concern about low quality largely targets this group.

Further work is needed to understand the net welfare gains of an increase in fees. While the benefits are fairly obvious—especially in today's context of low quality—the costs are more difficult to evaluate. A possible cost of dearer patents is the overall reduction of the return to patenting and hence a potential reduction of the incentives to innovate. (Note that for high-quality applications that are undeterred by higher fees, the net benefit of patenting is still reduced.) There are, however, reasons to believe that this effect is negligible. The price elasticity of demand for patents is inelastic (de Rassenfosse and van Pottelsberghe, 2012) and the decision to invest in inventive activities should be *a fortiori* even less sensitive to patent fees. Empirical evidence by Nicholas (2011) supports this view. The author shows that a dramatic lowering of patent filing fees in Britain in 1883 had no effect on the level of innovation of the British economy. Other side effects of higher fees could include a higher prevalence of secrecy, and the exclusion of cash-poor players. But these risks seem modest in the current context of concerns about patent quality, large backlogs and financial vulnerability of patent offices, suggesting that the option of fee increases should be considered seriously.

Acknowledgements

The research was supported by an Early Career Researcher grant from the Faculty of Business and Economics at the University of Melbourne. We are grateful to two anonymous referees, a Co-Editor, Justus Baron, Jérôme Danguy, Chris Dent, Paul Jensen, Anne Leahy, Alan Marco, Elisabeth Mueller, Mark Schankerman, Scott Stern, Russell Thomson, Bruno van Pottelsberghe, and Elizabeth Webster for valuable comments, as well as to conference and seminar participants at IP Australia (Canberra), the University of Melbourne (Melbourne), the USPTO (Washington, D.C.), the IPTS (Seville), LUISS Guido Carli University (EARIE conference, Rome), Seoul National University (APIC Conference, Seoul), MINES ParisTech (Paris), ESADE Business School (DRUID conference, Barcelona) and the EPFL (Lausanne). Antonio De Marco provided valuable research assistance.

References

Albert, M., D. Avery, F. Narin, and P. McAllister, 1991. 'Direct validation of citation counts as indicators of industrially important patents'. *Research Policy*, 20(3), 251–259.

Alcácer, J. and M. Gittelman, 2006. 'Patent citations as a measure of knowledge flows: The influence of examiner citations'. *Review of Economics and Statistics*, 88(4), 774–779.

Atkinson, S., A. Marco and J. Turner, 2009. 'The Economics of a Centralized Judiciary: Uniformity, Forum Shopping, and the Federal Circuit'. *Journal of Law and Economics*, 52(3), 411–443.

Barton, J., 2000. 'Reforming the patent system'. *Science*, 287(5460), 1933–1934.

Bender, K., S. Griffin and C. Lipsey, 1986. 'Patent decisions of the United States Court of Appeals for the Federal Circuit: The Year 1985 in Review'. *American University Law Review*, 35, 995-1036.

Bessen, J., 2008. 'The value of U.S. patents by owner and patent characteristics'. *Research Policy*, 37(5), 932–945.

Bessen, J. and E. Maskin, 2009. 'Sequential innovation, patents, and imitation'. *RAND Journal of Economics*, 40(4), 611–635.

Bessen, J. and M. Meurer, 2008. *Patent failure: How judges, bureaucrats, and lawyers put innovators at risk*. New Jersey: Princeton University Press.

Caillaud, B., A. Duchêne, 2011. ‘Patent office in innovation policy: Nobody’s perfect’. *International Journal of Industrial Organization*, 29(2), 242–252.

Carpenter, M., F. Narin and P. Woolf, 1981. ‘Citation rates to technologically important patents’. *World Patent Information*, 3(4), 160–163.

Cornelli, F. and M. Schankerman, 1999. ‘Patent renewals and R&D incentives’. *RAND Journal of Economics*, 30(2), 197–213.

DeBrock, L., 1985. ‘Market structure, innovation, and optimal patent life’. *Journal of Law and Economics*, 28(1), 223-244.

de Rassenfosse, G., A. Jaffe and E. Webster, 2016. “Low-quality patents in the eye of the beholder: Evidence from multiple examiners,” National Bureau of Economic Research Working Paper No. 22244.

de Rassenfosse, G. and B. van Pottelsberghe, 2012. ‘On the price elasticity of demand for patents’. *Oxford Bulletin of Economics and Statistics*, 74(1), 58–77.

de Rassenfosse, G. and B. van Pottelsberghe, 2013. ‘The role of fees in patent systems: Theory and evidence’. *Journal of Economic Surveys*, 27(4), 696–716.

The Economist (2015). A question of utility. August 8 2015, pp 50-52.

Farrell, J. and C. Shapiro, 2008. ‘How strong are weak patents?’. *American Economic Review*, 98(4), 1347–1369.

Feldman, R., 2014. Slowing the Patent Trolls. *The New York Times*, March 29 2014, Op-Ed, p. A21.

Gallini, N., 2002. ‘The economics of patents: Lessons from recent US patent reform’. *The Journal of Economic Perspectives*, 16(2), 131-154.

Gans, J., S. King and R. Lampe, 2004. ‘Patent renewal fees and self-funding patent offices’. *Topics In Theoretical Economics*, 4(1), Article 6.

Giacopello, F., 2012. ‘Rise in quantity, fall in quality: assessing China’s patent filings’. *Intellectual Property Magazine*, October issue, 63–64.

Griliches, Z., 1990. ‘Patent statistics as economic indicators: A survey’. *Journal of Economic Literature*, 28(4), 1661–1707.

Hall, B., 2005. ‘Exploring the patent explosion’. *Journal of Technology Transfer*, 30(1/2), 35-48.

Hall, B., S. Graham, D. Harhoff and D. Mowery, 2004. ‘Prospects for improving U.S. patent quality via postgrant opposition’. In A. Jaffe, J. Lerner and S.Stern (Eds), *Innovation Policy and the Economy*, Vol. 4. Cambridge: MIT Press.

Harhoff, D., F. Scherer and K. Vopel, 2003. ‘Citations, family size, opposition and the value of patent rights’. *Research Policy*, 32(8), 1343–1363.

Hegde. D and B. Sampat, 2009. ‘Examiner citations, applicant citations, and the private value of patents’. *Economics Letters*, 105(3), 287–289.

Helfgott, S., 1993. ‘Patent filing costs around the world’. *Journal of the Patent Office and Trademark Society*, 75(7), 567–580.

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

Hunt, R., 2006. ‘When do more patents reduce R&D?’. *The American Economic Review*, 96(2), 87–91.

Jaffe, A. and G. de Rassenfosse, Forthcoming. ‘Patent citation data in social science research: Overview and best practices’. *Journal of the Association for Information Science and Technology*, accepted.

Johnson, V., 2013. ‘Revised standards for statistical evidence’. *Proceedings of the National Academy of Sciences*, 110(48), 19313-19317.

Kondo, M., 1999. ‘R&D dynamics of creating patents in the Japanese industry’. *Research Policy*, 28(6), 587–600.

Lanjouw, J., A. Pakes and J. Putnam, 1998. ‘How to count patents and value intellectual property: The uses of patent renewal and application data’. *The Journal of Industrial Economics*, 46(4), 405–432.

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

Lemley, M. and C. Shapiro, 2005. ‘Probabilistic patents’. *Journal of Economic Perspectives*, 19(2), 75–98.

Liang, M., 2012. ‘Chinese patent quality: Running the numbers and possible remedies’. *John Marshall Review of Intellectual Property Law*, 11, 477–522.

Lucas, R., 1976. ‘Econometric policy evaluation: A critique’, in K. Brunner and A. Meltzer, *The Phillips Curve and Labor Markets*, Carnegie-Rochester Conference Series on Public Policy, 1, New York, pp. 19–46.

Martínez, C., 2011. 'Patent families: when do different definitions really matter?'. *Scientometrics*, 86(1), 39–63.

Matutes, C., P. Régibeau and K. Rockett, 1996. 'Optimal patent design and the diffusion of innovations'. *RAND Journal of Economics*, 27(1), 60–83.

Michel, J. and B. Bettels, 2001. 'Patent citation analysis. A closer look at the basic input data from patent search reports'. *Scientometrics*, 51(1), 185–201.

Mueller, F., and H. Wegner, 1977. 'The 1976 German patent law'. *Journal of the Patent Office Society*, 59(2), 89–137.

Narin, F., E. Noma and R. Perry, 1987. 'Patents as indicators of corporate technological strength'. *Research Policy*, 16(2–4), 143–155.

Nicholas, T., 2011. 'Cheaper patents'. *Research Policy*, 40(2), 325–339.

Organisation for Economic Co-operation and Development (OECD), 2011. *OECD Science, Technology and Industry Scoreboard 2011*, 204 p. Paris: OECD Publishing.

Pakes, A., 1986. 'Patents as options: Some estimates of the value of holding European patent stocks'. *Econometrica*, 54(4), 755–784.

Palangkaraya, A., E. Webster and P. Jensen, 2011. 'Misclassification between patent offices: Evidence from a matched sample of patent applications'. *Review of Economics and Statistics*, 93(3), 1063–1075.

Picard, P. and B. van Pottelsberghe, 2013. 'Patent office governance and patent system quality'. *Journal of Public Economics*, 104, 14-25.

Putnam, J., 1996. *The Value of International Patent Protection*. PhD Thesis. Yale.

Quillen, C., 2006. 'Innovation and the US patent system'. *Virginia Law & Business Review*, 1(2), 207–237.

Schankerman, M. and A. Pakes, 1986. 'Estimates of the value of patent rights in European countries during the post-1950 period'. *Economic Journal*, 96(384), 1052–1076.

Schmoch, U., F. Laville, P. Patel and R. Frietsch, 2003, 'Linking technology areas to industrial sectors'. Final report to the European Commission DG Research: Fraunhofer Institut fur Systemtechnik und Innovationsforschung: Karlsruhe.

Scotchmer, S., 1999. 'On the optimality of the patent renewal system'. *RAND Journal of Economics*, 30(2), 181–196.

Stiglitz, J., 2014. 'Intellectual property rights, the pool of knowledge, and innovation'. *NBER Working Paper* 20014.

Strawbridge, D., D. McDonald and R. Moy, 1987. 'Patent law developments in the United States Court of Appeals for the Federal Circuit during 1986'. *American University Law Review*, 36, 861-902.

Tong, X. and J. Frame, 1994. 'Measuring national technological performance with patent claims data'. *Research Policy*, 23(2), 133-141.

Trajtenberg, M., 1990. 'A penny for your quotes: Patent citations and the value of inventions'. *RAND Journal of Economics*, 21(1), 172–187. New Technology, 5(1), 19–50.

van Pottelsberghe de la Potterie, B., 2011. 'The quality factor in patent systems'. *Industrial and Corporate Change*, 20(6), 1755–1793.

Data Appendix

Data sources

The central data source is the October 2012 edition of Patstat, the worldwide patent statistical database by the European Patent Office (EPO). It provides a listing of all patents granted by the USPTO around the time of the reform.[11] We include all utility patents filed 21 months before and after the reform, that is, from January 1, 1981 to June 30, 1984. There are 223,393 such patent documents. These patent documents all correspond to granted patents, because the USPTO at that time did not publish patent applications that were rejected. The quality indicators related to the number of citations, the number of claims and the family size are computed from Patstat. Information on the number of claims is missing for 545 patents. The data is complemented with the USPTO Patent Maintenance Fee Events (PMFE) database in order to compute the number of renewals. A total of 350 patents could not be matched with the USPTO PMFE database. The full sample contains 222,434 patents with all quality indicators available.

A second sample is used for the econometric analysis. A total of 143 patents are excluded from the sample due to missing IPC codes. The DiD sample also excludes 50,496 patents by small entities. Although the reform led to an increase in fees for all patentees, the increase was much smaller for assignees that could claim the small entity fee reduction. Thus, excluding patents assigned to small entities allows for an intensity of treatment that is homogenous across all local applicants. Second, patents in the control group are overwhelmingly owned by large entities, and the exclusion of patents by small entities therefore increases the homogeneity between treatments and controls. Finally, this filter leads to the exclusion of university-owned patents, thereby mitigating the potential effect of the Bayh-Dole Act. The treatment group is composed of priority filings by local applicants, and the control group is composed of second filings by foreign applicants. The DID sample therefore excludes 38,577 second filings by local applicants and 3,547 priority filings by foreign applicants. Finally, we keep firms active both before and after the PLAA in order to increase homogeneity between pre and post samples. Table A-1 provides an overview of samples' composition.

The data is also complemented with the OECD Applicant Harmonized Name (HAN) data table, which provides a clean listing of assignees (used for computing patent portfolio size).

[11] The MySQL source code is available upon request from the authors.

Table A-1. Overview of data sources

Description	Source	Full sample size	DiD sample size
All patents granted around PLAA	Patstat	223,408	223,408
Missing number of citations	Patstat	(79)	(0)
Missing number of claims	Patstat	(93)	(93)
Missing family size	Patstat	(0)	(0)
Missing renewal data	USPTO PMFE	(350)	(350)
Missing IPC codes	Patstat	-	(144)
Excluding small entities	USPTO PMFE	-	(50,560)
Excluding second filings by local applicants	Patstat	-	(38,646)
Excluding priority filings by foreign applicants	Patstat	-	(3,555)
Excluding patents by firms active only before or after PLAA	Patstat	-	(14,720)
Excluding outliers	Patstat	-	(1272)
Final sample		*222,886*	*113,989*
Final sample, excl. patents filed 1 month around PLAA		*210,041*	*107,490*

Construction of patent quality indicators

The number of citations received by a patent is computed by counting the number of times the patent document was cited ten years after grant. We consider only citations from USPTO patents; we exclude citations from patents filed in other jurisdictions. The indicator is subject to an inflation bias because the number of citing patents increases over time, with the growing number of patent applications.

The number of claims is directly available from Patstat and corresponds to the number of claims listed in the publication associated with the first granted document.

The size of the patent family is computed by counting the number of jurisdictions covered by the patent documents in the same DOCDB family. The DOCDB family is constructed by examiners and covers patent documents protecting the same technical content. See Martínez (2011) for additional information on patent families.

The number of renewals is computed from the USPTO PMFE database. Every patent that has expired is associated with a code EPX. and a corresponding expiration date in the database. Patents that are not associated with an expiration code were maintained to full term (17 years after grant at that time).