

NBER WORKING PAPER SERIES

LOW-QUALITY PATENTS IN THE EYE OF THE BEHOLDER:  
EVIDENCE FROM MULTIPLE EXAMINERS

Gaétan de Rassenfosse  
William E. Griffiths  
Adam B. Jaffe  
Elizabeth Webster

Working Paper 22244  
<http://www.nber.org/papers/w22244>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
May 2016, revised November 2019

The authors are grateful to seminar and conference participants at the New Zealand Economic Association Conference, Asia Pacific Innovation Conference, European Intellectual Property Policy Conference, Toulouse School of Economics, ETH Zurich, and the third International Meeting in Law & Economics (Paris) for valuable feedback. Joachim Henkel and Sonia Jaffe provided valuable comments. T'Mir Julius provided excellent research assistance and her contribution is gratefully acknowledged. This study was financed by the Australian Research Council Discovery Grant ARC LP110100266 “The Efficiency of the Global Patent System” with partners IP Australia and the Institute of Patent and Trademark Attorneys. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Gaétan de Rassenfosse, William E. Griffiths, Adam B. Jaffe, and Elizabeth Webster. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Low-quality Patents in the Eye of the Beholder: Evidence from Multiple Examiners  
Gaétan de Rassenfosse, William E. Griffiths, Adam B. Jaffe, and Elizabeth Webster  
NBER Working Paper No. 22244  
May 2016, revised November 2019  
JEL No. K41,L43,O34

### **ABSTRACT**

A low-quality patent system threatens to slow the pace of technological progress. Concerns about low patent quality are supported by estimates from litigation studies suggesting that the majority of patents granted by the U.S. patent office should not have been issued. This paper proposes a new Bayesian method for measuring patent quality, based on twin patent applications granted at one office but refused at another office. Our method allows us to distinguish whether low-quality patents are issued because an office implements a (consistently) low standard, or because it violates its own standard. The results suggest that quality in patent systems is higher than previously thought. In particular, relative to the own standard of each office, the percentage of mistakenly granted patents is under 10 percent for all offices. The Japanese patent office has a greater percentage of mistakenly granted patents than those of Europe, the United States, Korea and China, largely because it has a higher standard.

Gaétan de Rassenfosse  
Ecole polytechnique fédérale de Lausanne  
College of Management of Technology  
Odyssey Station 5  
CH-1015 Lausanne  
gaetan.derassenfosse@epfl.ch

William E. Griffiths  
Department of Economics  
University of Melbourne  
3010, Victoria  
Australia  
b.griffiths@unimelb.edu.au

Adam B. Jaffe  
188 Brookline Avenue  
Apartment 26A  
Boston, MA 02215  
and Brandeis University  
and Queensland University of Technology  
and also NBER  
adam.jaffe@motu.org.nz

Elizabeth Webster  
Swinburne University of Technology  
Mail H25, PO Box 218 Hawthorn  
VIC 3122, Australia  
emwebster@swin.edu.au

## 1. Introduction

Concern that the patent system inhibits rather than encourages innovation has become a staple of the business and technology press (*e.g.*, *The Economist*, 2015). A major source of concern is that patent offices may grant too many low-quality patents, whose existence can chill the R&D investment and commercialization processes, either because of background uncertainty about freedom to operate or because of implicit or explicit threats of litigation.

Concern about patent quality is by no means new. The recent *Economist* article quoted itself from 1851 saying that the granting of patents “begets disputes and quarrels betwixt inventors, provokes endless lawsuits [and] bestows rewards on the wrong persons.” But in the last few decades, significant increases in the number of patent applications granted and the frequency of patent litigations, as well as media attention such cases have received, have given these concerns new force in the academic literature. Major patent offices are well aware of the problem and several of them have initiatives underway aimed at improving the quality of patent review. For example, the U.S. Patent and Trademark Office (USPTO) now has an Office of Patent Quality Assurance and has recently initiated an ongoing online ‘patent quality chat.’<sup>i</sup>

We interpret concern about low-quality patents as corresponding to concern that patents are being granted whose contribution is too small to deserve patent protection. Conceptually, there are two pathways by which this may be occurring. A first source of low quality in a patent system relates to the fact that patent offices might *systematically apply a standard that is too lenient*, relative to some conception of optimal stringency about, *e.g.*, non-obviousness or disclosure standards. Some of the discussion of the patent quality problem, particularly in the United States, has this flavor. Jaffe and Lerner (2004), for example, argue that changes in the incentives of the USPTO, the U.S. courts, and U.S. patentees over the 1980s and 1990s led to a systematic lowering of the standard for a U.S. patent grant.

A conceptually distinct source of low quality in patent system is mistakes—*granting patents that do not meet the office’s own implicit standard*, however high or low that standard may be. Observers of the patent system also discuss this issue. For example, Lemley and Shapiro (2005:83) write: “There is widespread and growing concern that the Patent and Trademark Office issues far too many ‘questionable’ patents that are unlikely to be found valid based on a thorough review.” Although there are clear patentability requirements, subjective elements in the examination process (Meurer, 2009; Lemley and Sampat, 2012; Frakes and Wasserman, 2017; Nagaoka and Yamauchi, 2015) and in the governance of patent offices (de Saint-Georges and van Pottelsberghe, 2013; Schuett, 2013; Picard and van Pottelsberghe, 2014) affect the consistency of the examination decision. More generally, the grant decision rests ultimately on a person’s (or team’s) comparison of the application’s inventive merit and level of disclosure and the office’s standard for patentability (including patentable subject matter, obviousness and disclosure). Perfect consistency of decision-making across examiners seems unlikely to be the outcome of such a process.

The practical and normative consequences of these different sources of low quality are different. Systematically low standards create monopoly power and transfer rents in situations where the triviality of the invention does not justify the reward. Consistently low standards are not, however, a source of uncertainty about which patents are truly valid—so long as the patent office and the courts are applying exactly the same standard. On the other hand, uncertainty arising from the perceived probability that mistaken decisions will be made dampens downstream patent development as patent owners are unsure if their patent will survive a validity challenge.

We propose a formal model that separately estimates office standards from office ‘mistakes’. The premise of our model is that a refusal by an examiner in one jurisdiction raises doubts as to the legitimacy of any patent grant secured elsewhere. We then use novel data on multiple examination outcomes for the same invention in different patent offices to estimate how many patents granted, by office, are either the result of a low office standard or an office ‘mistake’. Our data are derived from a population of 408,133 inventions with linked patent applications that have been examined in at least two of the five major patent offices, covering in total more than a million applications. We estimate a Bayesian statistical model of the grant process that reflects the underlying data generation process, and captures parametrically the effect of observable application attributes on the grant probability, the effect of systematic differences in office standards, and the possibility of personal (*i.e.*, examiner) mistakes.

To foreshadow the results, we find that systematic differences in standards *across* offices are quantitatively more important than the inconsistency (*i.e.*, mistakes) *within* offices. Up to 10 percent of granted patents, depending on the office, have dubious validity in the specific sense that they appear to be inconsistent with the country’s own standard for patent grant. On the other hand, up to 19 percent, depending on the office, would not have been granted by the strictest office (Japan). These estimates help define the boundaries in policy discussions over international patent office agreements and improvements for national patent systems.

The rest of the paper is organized as follows. Background discussion on patent quality is presented in Section 2. The empirical strategy is presented in Section 3; Section 4 contains a description of the data. The econometric results are discussed in Section 5 and Section 6 concludes.

## **2. Background**

Most of the existing literature looks at the issue of low quality by measuring the fraction of litigated patents that are found by a court to be invalid. Such studies provide valuable insights on the prevalence of invalidity. It is unclear, however, how invalidation in court relates to the two possible sources of invalidity. If one assumes that the courts are implicitly applying the same standard as the patent office, and that courts make perfect decisions, then a court invalidity finding corresponds to a case in which the office did not correctly apply its own standard. In practice, it is also possible that the court is applying a more stringent standard—and that it makes mistakes (Lemley, 2001). Thus, invalidity rates

from such studies are difficult to relate to the quality of the examination process or the stringency of the office.

It should be noted, however, that patent litigation studies report high rates of ‘invalidity’. Allison and Lemley (1998) reviewed final validity decisions of 299 litigated patents and found an invalidity rate of half. Cremers *et al.* (2014) report that about 30 percent of appealed patent suits have their initial decision overturned. More recently, scholars have also studied the outcome of *inter partes* reviews, which are post-grant reviews conducted by USPTO Patent Trial and Appeal Board (Wallach and Darrow, 2016). These studies suggest that invalidity rates might be quite high. However, given that a mere 0.1 percent of patents are litigated to trial (Lemley and Shapiro, 2005), it remains unclear what these statistics tell us about the overall prevalence of invalidity.

This point is well made by Marco (2004), who emphasizes the importance of accounting for selection effects in patent validity adjudications. Recognizing this problem, Miller (2013) attempts to correct for selection into an invalidity hearing. Using 980 adjudicated and 1960 control patents at the USPTO, he estimates a population-wide invalidity rate of 28 percent. However, the selection into Miller’s sample is twofold: selection into a patent being disputed, and selection into parties choosing trial over settlement. The first selection is not accounted for, suggesting that the 28-percent figure may still be biased, though the direction of bias is unclear. Zischka and Henkel (2014) have also studied carefully the presence of selection bias in their data but did not identify statistically significant selection covariates. They find a 75 percent invalidity rate of appeals at the German Federal Patent Court between 2000 and 2012.

As illustrated by the litigation studies, the basic approach to assessing the level of quality in the system is to investigate what happens when another qualified decision maker (but ideally many) takes a fresh look at the question of whether an asserted invention qualifies for patent protection. Paradise *et al.* (2005) manually examine the validity of 1167 claims of 74 U.S. patents on human genetic material. They find that 448 claims (38%) were problematic. The ‘second-pair-of-eyes review’ program at the USPTO, which began in the year 2000 but has been discontinued since, aims at assessing examination quality by re-examining patent applications related to business methods. However, data are not publicly available and Allison and Hunter (2006:737–8) comment that this review is a “subjective, in-house process metric guided by no apparent standards that may fall victim to unconscious bias or external influence.”

In contrast with these studies, Palangkaraya, Jensen and Webster (2011) use a revealed behavior method to estimate rates of patent invalidity. They analyze the population of all 34,000 patent applications that were granted by the USPTO and examined at both the EPO and JPO during the 1990s. If the number of forward citations at the USPTO is a proxy for the real size of the inventive step, they estimate that 9.8 percent of patents are incorrectly granted.

Finally, other studies have empirically examined the issue of patent quality at the USPTO by exploiting data on examiner decisions at other offices. Lemley and Sampat (2012) link 2176 USPTO

applications to decisions at the EPO (European Patent Office) to assess the effect of examiner experience on grant outcomes. Lei and Wright (2017) measure the weakness of U.S. patents using outcomes of related applications at the EPO. Frakes and Wasserman (2017) use outcomes of EPO and JPO (Japanese Patent Office) applications to assess whether the time allocated to review applications at the USPTO affects the validity of patents issued. This line of work offers the possibility of important nuanced understanding of the determinants of possibly heterogeneous decisions within patent offices. As far as we can ascertain, no study has exploited variations in outcomes across offices to study the quality of patent systems in a systematic manner.

### 3. Empirical strategy

Our research seeks to implement the second-pair-of-eyes approach with a much larger set of inventions and with more pairs of eyes. Our context allows each patent office to have its own *de facto* standard, and every decision-maker to make mistakes. We do so by analyzing the grant outcome of ‘twin’ patent applications submitted to multiple jurisdictions. Twin applications are applications covering the same technical content in different jurisdictions (Palangkaraya, Jensen and Webster, 2011; Webster, Jensen and Palangkaraya, 2014; Sampat and Shadlen, 2015).<sup>ii</sup> We estimate an index of the probability that each patent application is granted under the differing circumstances of the different patent offices, and then use the resulting estimates to predict the overall ease of obtaining a patent and the proportion of low quality patents in each office.

The sample for the analysis is the population of 408,133 inventions described in patent applications filed between 2000–2006 in at least two of the EPO, the USPTO, the JPO, the KIPO (Korean Intellectual Property Office) and the CNIPA (National Intellectual Property Administration of China, formerly SIPO). We use this time period in order to ensure that the applicant has had a chance to pursue protection in as many countries as he or she chooses, and to allow sufficient time to reach a grant decision. These five offices, known collectively as the ‘IP5 Offices’, attract about 80 percent of worldwide patenting activity.<sup>iii</sup>

#### 3.1. The model

Our model of the actual examination decision assumes that each invention  $i$  has a unique but unobservable inventive merit  $c_i$ , which is shared by all the applications to different offices. The probability of granting patent application  $i$ , by an examiner in office  $j$  is a function of this inventive merit  $c_i$ ; the office-specific *de facto* standard required for a grant and related institutional incentives  $\tau_j$ ; a set of covariates  $\mathbf{x}_{ij}$ , capturing observed heterogeneity at the patent-patent office level (*e.g.*, differences in the number of claims, filing route); and examiner-specific factors that are not systematic to the office,  $\varepsilon_{ij}$ .<sup>iv</sup> These elements combine to give an index  $y_{ij}^*$ , which maps into the probability of a grant for each application in each office. To formally specify this model, we begin by noting that patent  $i$  is accepted by office  $j$  if

$$c_i \geq \tau_j + (\mathbf{x}_{ij} - \bar{\mathbf{x}}_j) \boldsymbol{\beta}_j - \varepsilon_{ij} \quad (1)$$

That is, the inventive merit  $c_i$  is greater than or equal to the office standard  $\tau_j$  plus deviations from that standard attributable to specified covariates  $(\mathbf{x}_{ij} - \bar{\mathbf{x}}_j)\boldsymbol{\beta}_j$ , minus a non-systematic or random component  $\varepsilon_{ij}$ . In equation (1), the vector  $\bar{\mathbf{x}}_j = \sum_{i=1}^{N_j} \mathbf{x}_{ij} / N_j$  is the sample average of the covariates associated with applications to office  $j$ ;  $N_j$  is the number of applications considered by office  $j$ ; and  $\boldsymbol{\beta}_j$  is a vector of unknown parameters that quantify the impact of the covariates on the acceptance decision. Under the assumptions we outline below, the sign in front of the random component  $\varepsilon_{ij}$  could be positive or negative; we have made it negative to be consistent with the development that follows. Expressing the covariates in terms of deviations from their office means makes  $\tau_j$  the average acceptance standard for office  $j$ . Differences in office standards can be assessed by comparing estimates of the  $\tau_j$ 's.

Office “mistakes” occur when the random element  $\varepsilon_{ij}$  is such that the outcome of a patent application is contrary to that predicted by the relative magnitudes of  $c_i$  and  $\tau_j + (\mathbf{x}_{ij} - \bar{\mathbf{x}}_j)\boldsymbol{\beta}_j$ . We have a mistakenly granted or low-quality patent if an application is successful when  $c_i < \tau_j + (\mathbf{x}_{ij} - \bar{\mathbf{x}}_j)\boldsymbol{\beta}_j$  and a mistakenly refused patent if an application is not successful when  $c_i \geq \tau_j + (\mathbf{x}_{ij} - \bar{\mathbf{x}}_j)\boldsymbol{\beta}_j$ .

To specify the model in a form suitable for estimation, the latent variable or index  $y_{ij}^*$  is written as

$$y_{ij}^* = c_i - \tau_j - (\mathbf{x}_{ij} - \bar{\mathbf{x}}_j)\boldsymbol{\beta}_j + \varepsilon_{ij} \quad (2)$$

The condition  $c_i \geq \tau_j + (\mathbf{x}_{ij} - \bar{\mathbf{x}}_j)\boldsymbol{\beta}_j - \varepsilon_{ij}$  is equivalent to  $y_{ij}^* \geq 0$ . Let the binary variable  $y_{ij}$  take the value  $y_{ij} = 1$  if the patent application for the  $i$ -th invention is granted by the  $j$ -th office ( $y_{ij}^* \geq 0$ ) and  $y_{ij} = 0$  otherwise ( $y_{ij}^* < 0$ ). Then, when the random variable  $\varepsilon_{ij}$  is assumed to have a standard normal distribution, equation (2) becomes a probit model. The probability that an application is granted is given by

$$\begin{aligned} \Pr(y_{ij} = 1) &= \Pr(y_{ij}^* > 0) \\ &= \Pr\left[c_i \geq \tau_j + (\mathbf{x}_{ij} - \bar{\mathbf{x}}_j)\boldsymbol{\beta}_j + \varepsilon_{ij}\right] \\ &= \Pr\left[\varepsilon_{ij} < c_i - \tau_j - (\mathbf{x}_{ij} - \bar{\mathbf{x}}_j)\boldsymbol{\beta}_j\right] \\ &= \Phi\left(c_i - \tau_j - (\mathbf{x}_{ij} - \bar{\mathbf{x}}_j)\boldsymbol{\beta}_j\right) \end{aligned} \quad (3)$$

where  $\Phi(\bullet)$  is the standard normal cumulative distribution function.

We assume  $(\varepsilon_{ij} | \mathbf{X}) \sim N(0,1)$ , where  $\mathbf{X}$  denotes all observations  $\{\mathbf{x}_{ij}\}$  and the  $(\varepsilon_{ij} | \mathbf{X})$  are independent over  $i$  and  $j$ . An implicit identifying assumption is that  $E_j(\varepsilon_{ij} | \mathbf{X}) = 0$ , *i.e.*, examiners at office  $j$  take correct decisions on average. (Any systematic deviation from the ‘correct’ outcome is captured by the office-specific component—the *de facto* standard.) Likewise,  $E_i(\varepsilon_{ij} | \mathbf{X}) = 0$ , *i.e.*, every invention is treated fairly on average across offices. The error term  $\varepsilon_{ij}$  captures the subjectivity of interpretation of the patent law or the ‘mood’ of the examiner among other things. That is, if the same application were examined in the same office, under the same office procedures and covariates, but by a different examiner, any difference in the decision would be explained by  $\varepsilon_{ij}$ .

### 3.2. Estimation

For estimating the parameters of the model, we adopt a Bayesian approach with a hierarchical prior on the  $c_i$ . The hierarchical prior for  $c_i$  provides scope for improving the precision of estimation given that the number of observations on each patent family (the offices considering each patent) is relatively small ( $\leq 5$ ). The starting point for Bayesian analysis is the specification of prior distributions for  $(c_i, \tau_j, \beta_j)$ . These prior distributions are combined with the likelihood function that corresponds to equation (3) to form a joint posterior distribution for  $(c_i, \tau_j, \beta_j)$ . The means from the posterior distribution are the estimates for  $c_i$ ,  $\tau_j$  and  $\beta_j$  typically reported by Bayesian studies, although arguments can be made for using other values such as medians or modes. Along with the posterior means, the posterior standard deviations are usually reported to gauge the level of uncertainty in the estimates.

While this procedure is conceptually simple, complications can arise when it is being implemented. For many models, the integrals that define the posterior means and standard deviations are too difficult to evaluate. This problem has led to an explosion of Markov chain Monte Carlo (MCMC) techniques for sampling observations from posterior densities, with these observations being used to estimate the posterior means and standard deviations. See, for example, Brooks *et al.* (2011). For the probit model defined by (2) and (3), the obvious choice of an MCMC technique is Gibbs' sampling, which draws observations from the joint posterior by sampling sequentially from the conditional posterior densities for  $y_{it}^*$ ,  $c_i$ ,  $\tau_j$  and  $\beta_j$  (Albert and Chib 1993). Note that, in this set up, the latent variable  $y_{it}^*$  is introduced as an extra unknown parameter. An alternative to Gibbs sampling, and one which we adopt in this paper, is to derive approximate marginal posterior densities from which the means and standard deviations can be readily calculated. The technique we use for deriving approximate marginal posteriors is called variational Bayes. Details of how it works can be found in Ormerod and Wand (2010) or Gelman *et al.* (2014: 331–338). Variational Bayes was chosen in preference to Gibbs sampling because of the daunting computational task implied by Gibbs sampling. With over one million observations on  $y_{it}^*$ , and over 400 thousand observations on  $c_i$ , drawing and storing repeated samples is computationally demanding.

In the Appendix, we describe the prior distributions we placed on the parameters and the essential results needed to obtain variational Bayes estimates for our model. Proofs of these results can be obtained in supplementary material available from the second author upon request.

## 4. Data and variables

### 4.1 A dataset of one-to-one equivalents across offices

Our empirical set up follows the same logic as studies of identical twins. We endeavor as far as possible to collate a dataset comprising the same invention that was the subject of multiple applications (at different offices). Accordingly, we exclude patent applications that lead to divisionals and



continuations. As with identical twin studies, we infer that our results hold for out-of-sample observations but cannot formally prove this.

To construct this ‘twin’ dataset, we combined data from seven offline and online sources. The main data source is the EPO-OECD PATSTAT database (October 2014 release) for the backbone of the dataset. We start from the universe of priority patent applications filed anywhere in the world over the period 2000 to 2006 (de Rassenfosse *et al.*, 2013) and track their one-to-one equivalents in any of the five offices.<sup>v</sup> A priority filing is the first patent application describing an invention. Application  $P_B$  in country  $B$  is a one-to-one equivalent of application  $P_A$  in country  $A$  if  $P_B$  claims  $P_A$  as sole priority (*i.e.*, no merged patent applications) and  $P_A$  is only claimed by  $P_B$  in office  $B$  (*i.e.*, no split patent applications). In this sense,  $P_A$  and  $P_B$  cover the same technical content and are ‘twin’ applications. We also extract from PATSTAT information on applicants’ country of residence, patents technological fields as identified with the International Patent Classification (IPC) codes and filing route (either the ‘Paris Convention’ route or the PCT route).<sup>vi</sup>

Data on the application legal status (granted/refused/withdrawn) come from: the EPO’s INPADOC PRS table for PATSTAT for European and Chinese applications; from JPO’s public access on-line Industrial Property Digital Library Database (IPDL) for Japanese applications; from KIPO public access on-line IPR Information Service (KIPRIS) for Korean applications; and from the USPTO’s Public Pair on-line database for U.S. applications.

Data on the number of claims of published patent applications come from: PATSTAT for European applications; CNIPA’s on-line patent search platform for Chinese applications; IPDL for Japanese applications; KIPRIS for Korean applications; and lens.org for U.S. applications. We developed specific web-crawlers to collect online information.

#### 4.2 Variables

Our main dependent variable,  $y_{ij}$ , is the binary outcome that takes the value of 1 if patent application  $i$  was granted by an examiner in patent office  $j$  and 0 if refused. Our measure of refusal includes applications that were examined and refused by the patent office plus all quasi-refusals. Quasi-refusals include patent applications that were withdrawn at the EPO following a negative search report containing  $X$  or  $Y$  citations, which challenge the inventive step of an application. Indeed, many applications at the EPO are withdrawn after a (negative) office communication, which Lazaridis and van Pottelsberghe (2007) interpret as quasi-refused applications.

There are three fundamental sources of heterogeneity with respect to the grant outcome in the data: systematic office differences ( $j$ ), systematic invention differences ( $i$ ), and application-patent office differences ( $ij$ ). The first two sources are accounted for by estimation of the office ( $\tau_j$ ) and invention ( $c_i$ ) parameters, respectively. Concerning the third source, we control for four variables  $\mathbf{x}_{ij}$  that are likely to induce heterogeneity in the grant decision across offices for the same invention. On average the examiners from the different offices make a true assessment of the inventive merit in the

application, which is measured with the invention parameter  $c_i$ . The four variables influence the examiner's decision over and above the objective quality of the invention.

The first of these controls is a dummy variable, *local applicant<sub>ij</sub>*, which equals 1 if there is at least one applicant with an address in the same jurisdiction as the examining patent office, and 0 otherwise. There is empirical evidence that patent offices give preferential treatment to domestic applications (Webster, Jensen and Palangkaraya, 2014, de Rassenfosse *et al.* forthcoming). This home bias may reflect prejudice, but it may also reflect the fact that domestic applicants have stronger incentives to push the patent application in their home market or that they may be more familiar with their home patent system.

The second is the dummy variable *priority filing<sub>ij</sub>*, which takes the value 1 if application  $i$  is a priority filing in office  $j$  and 0 otherwise. By the construction of our data (using one-to-one equivalents), there can be only one priority filing per family. Firms usually file a priority filing in the office they know best, which may affect the likelihood that they receive a grant in that office. The country of the priority office may also be the most important market, where incentives to push for a grant are stronger.

The third is the dummy variable *PCT<sub>ij</sub>*, which indicates whether the patent application was filed through the Patent-Cooperation Treaty route.<sup>vii</sup> There are non-trivial administrative implications of using the PCT route that may affect the consistency of examination outcome (*e.g.*, search report shared between all the offices, extension of priority right from 12 to 31 months).

Finally, we control for the number of claims (*claims<sub>ij</sub>*), which is the number of claims articulated in the patent application at the time of lodgment. Although twin applications in our sample cover the same technical content, there might be slight differences in the construction of the applications across offices. The number of claims is a proxy for these differences.

Table 1 presents a summary of the characteristics of the patent applications at each office. The sample includes 1,064,513 patent applications associated with 408,133 inventions (or patent families). Thus, one invention is submitted for patent protection in 2.6 offices on average. Overall, the JPO, at 72.23 percent, recorded the lowest grant rate and the CNIPA, at 96.3 percent, the highest. More than half of applications at the JPO had at least one local applicant compared with only 3.1 percent at CNIPA.<sup>viii</sup> CNIPA had also the smallest rate of priority filings and JPO the highest. (Indeed, except for the EPO, there is a strong correlation between the office of priority filing and whether the applicant is local to that office.) Use of the PCT was highest for the EPO but lowest for KIPO. Finally, the average number of claims at the time of application varies between 10.3 at the JPO and 17.8 at the USPTO.<sup>ix</sup>

Table 2 provides an overview of the number of equivalents (*i.e.*, twins) between offices. There are 125,704 direct equivalents between the USPTO and the EPO. The lowest number of equivalents is reached between the EPO and the KIPO (32,082 patent applications) and the highest number is reached

between the USPTO and the JPO (212,673 applications). As far as the CNIPA is concerned, it is most integrated with the USPTO, closely followed by the JPO.

**Table 1. Descriptive statistics, 2000–2006**

	All offices	EPO	USPTO	KIPO	JPO	CNIPA
Number of patent applications	1,064,513	162,803	325,068	127,314	278,760	170,568
Number of inventions	408,133	162,803	325,068	127,314	278,760	170,568
Proportion of accepted patents	0.8410	0.7677	0.9139	0.8444	0.7223	0.9634
Proportion of local applicants (LA)	0.3305	0.4414	0.2000	0.4151	0.5626	0.0311
Proportion of priority filing (PF)	0.2850	0.0984	0.2226	0.4098	0.5643	0.0325
Proportion of LA × PF	0.2643	0.0706	0.1828	0.4058	0.5521	0.0283
Proportion of PCT patents	0.2423	0.4521	0.2283	0.0225	0.2646	0.1962
Average number of claims	14.747	15.549	17.777	14.928	10.299	15.340

Notes: Data relate to patent applications filed between 2000 and 2006. See main text for data sources.

**Table 2. Cross-country number of equivalents, 2000–2006**

	EPO	USPTO	KIPO	JPO
USPTO	125,704			
KIPO	32,082	87,228		
JPO	91,878	212,673	79,757	
CNIPA	59,388	119,841	64,925	113,561

Next, Table 3 presents the breakdown of ‘inconsistent’ decisions by patent family size. For instance, inventions filed in two countries (family size of 2) represent 56.7 percent of the cases. The inventions are usually granted patent protection in both jurisdictions (67.8 %), although 27.7 percent of them are granted patent protection in one jurisdiction and refused in the other. These differences in grant outcome can be legitimate, resulting from different standards for a grant, or they could signal a mistake by one office. The empirical analysis will seek to quantify these two dimensions.

**Table 3. Overview of inconsistent decisions**

Number of Applications	2	3	4	5
Proportion (among total)	0.56715	0.28396	0.12237	0.02652
Proportion with 0 acceptances	0.04544	0.01007	0.00144	0.00000
Proportion with 1 acceptance	0.27664	0.07905	0.01734	0.00388
Proportion with 2 acceptances	0.67792	0.27494	0.08283	0.02615
Proportion with 3 acceptances		0.63594	0.26914	0.09619
Proportion with 4 acceptances			0.62924	0.26206
Proportion with 5 acceptances				0.61172

Not only is it possible to use the results to decompose *all* inconsistent decisions across offices into those attributable to different standards and those attributable to office mistakes, but it is also possible to focus on those applications *granted* by each office and refused elsewhere, and to decompose these inconsistencies into those attributable to a more lenient standard in the office being considered (the focal office), those attributable to a mistake by the focal office, and those attributable to a mistake by another office. This decomposition is particularly relevant for the issue of low-quality patterns. In Table 4 we examine the ‘raw’ rates of mixed grant-refusal patent families, that is, before correcting for office-specific differences and the influence of examiners’ subjective assessments. The raw rates show that 21.3 percent of the patents that were granted by the EPO were refused in at least one other office. The JPO has the lowest rates (13.9%), and the CNIPA the highest (26.9%).

**Table 4. Mixed grant-refusal patent families, 2000–2006**

Office	Number of granted patents	Proportion refused elsewhere
EPO	124,988	0.21299
USPTO	297,072	0.25236
KIPO	107,501	0.25666
JPO	201,335	0.13948
CNIPA	164,321	0.26899

## 5. Results

### 5.1 Parameter estimates

The parameter estimates (posterior means) were obtained by solving the equations described in the Appendix. Convergence of the iterations used to solve these equations was achieved after 1330 iterations, the parameter estimates changing very little in the last 300 iterations. The posterior standard deviations that are presented in Table 5 alongside the posterior means were calculated from the standard deviations of the variational Bayes approximate posterior densities. Of particular interest are the estimates of the  $\tau_j$ , the office standards for grants. They show that, relative to the CNIPA which is normalized at 0.0000, the USPTO, KIPO, EPO and JPO have more stringent standards. They are estimated as 0.8970, 1.2109, 1.8388 and 2.4018, respectively, with JPO having the strictest standard.

To use Table 5 to assess the estimated impact of the covariates on the probability of granting a patent, note that, from equation (3), a negative sign suggests a positive effect on the probability, and vice versa. There are exceptions, but the signs of the coefficients are generally uniform across offices, with *local applicant* (*LA*), *priority filing* (*PF*), and *PCT* having a positive effect on the probability, and *claims* having a negative effect. When both *LA* and *PF* are present, the conclusion about their effects needs to be modified in some instances because of the presence of the interaction term  $LA \times PF$ .

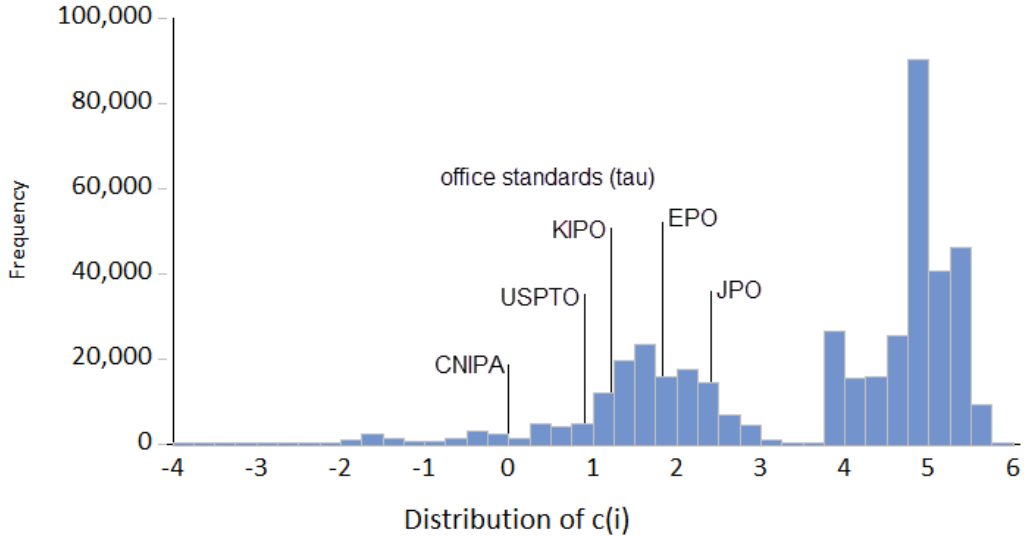
Estimates of the average probabilities of acceptance for each office are given by  $N_j^{-1} \sum_{i=1}^{N_j} \Phi(\bar{c}_i - \bar{\tau}_j - (\mathbf{x}_{ij} - \bar{\mathbf{x}}_j) \bar{\boldsymbol{\beta}}_j)$ . They match the raw data proportions of acceptance to the first two decimal places.

**Table 5. Parameter estimates**

	EPO	USPTO	KIPO	JPO	CNIPA
$\bar{\beta}_{j1}$ <i>local applicant (LA)</i>	-1.3619 (0.0055)	-0.0292 (0.0135)	-1.3581 (0.0292)	-0.1996 (0.0188)	-0.8928 (0.0463)
$\bar{\beta}_{j2}$ <i>priority filing (PF)</i>	0.0673 (0.0154)	-0.4743 (0.0093)	-1.6407 (0.0444)	-0.6838 (0.0177)	-0.8720 (0.0375)
$\bar{\beta}_{j3}$ <i>LA × PF</i>	1.3766 (0.0184)	0.1075 (0.0166)	1.4071 (0.0532)	0.1494 (0.0255)	1.3714 (0.0613)
$\bar{\beta}_{j4}$ <i>PCT</i>	-1.0266 (0.0054)	0.9725 (0.0044)	-2.1387 (0.0190)	-0.3962 (0.0055)	-1.5998 (0.0061)
$\bar{\beta}_{j5}$ <i>ln(claims)</i>	0.1944 (0.0039)	0.4422 (0.0029)	-0.0901 (0.0038)	0.0702 (0.0026)	0.1053 (0.0034)
$\bar{\tau}_j$ <i>office standard</i>	1.8388 (0.0025)	0.8970 (0.0018)	1.2109 (0.0028)	2.4018 (0.0019)	0.0000 (0.0000)
average probability of acceptance	0.7689	0.9130	0.8443	0.7237	0.9632

Note: Posterior standard deviations in parentheses. For the  $\beta$ 's, a negative sign implies a positive effect on the probability of acceptance and vice versa. A low office standard (higher probability of acceptance) implies that an office is more lenient. Average probabilities of acceptance match the raw data proportions of acceptance to the first two decimal places.

Figure 1 depicts the distribution of estimates of patent merit along with the estimated thresholds that depict the office standards. The distribution is bimodal, with a first group of inventions spread across the office thresholds and a second group of high inventive merit applications. Further investigation (not reported) suggests that the bimodal nature of the distribution can be explained by the proportion of offices that accepted each patent application. Values for  $\bar{c}_i$  above approximately 3.8 correspond to inventions that were accepted by all offices who considered them. The middle range corresponds to applications accepted by one or more offices and rejected by one or more offices. There is also a long tail to the left that corresponds to applications rejected by all offices who considered them.



**Figure 1. Distribution of estimates of patent merit and estimated office thresholds**

### 5.2 Office mistakes

The model predicts that the  $i$ -th application to the  $j$ -th office will be granted if the estimates  $(\bar{c}_i, \bar{\tau}_j, \bar{\beta}_j)$  are such that  $\bar{c}_i \geq \bar{\tau}_j + (\mathbf{x}_{ij} - \bar{\mathbf{x}}_j) \bar{\beta}_j$ , and rejected otherwise. When a decision to grant or refuse a patent conflicts with the decision predicted by the model, we say an office mistake has occurred. Using this definition, observations can be classified into correct acceptances, correct rejections, incorrect acceptances and incorrect rejections. The numbers and proportions of observations in each of these categories for each office are given in Table 6. Incorrect decisions can only be identified for those applications that were accepted by at least one office and rejected by at least one office. The model correctly predicted the outcome for all applications where there were consistent decisions across offices.

**Table 6. Predicted versus actual grant outcomes**

Predicted Outcome	Realised outcome			
	Granted	$y_{ij} = 1$	Refused	$y_{ij} = 0$
<b>EPO</b>				
Grant	117,280	(0.720)	5,685	(0.035)
Refuse	7,708	(0.047)	32,130	(0.197)
<b>USPTO</b>				
Grant	293,398	(0.903)	16,638	(0.051)
Refuse	3,674	(0.011)	11,358	(0.035)
<b>KIPO</b>				
Grant	104,762	(0.823)	7,960	(0.063)
Refuse	2,739	(0.022)	11,853	(0.093)
<b>JPO</b>				
Grant	182,189	(0.654)	6,269	(0.022)
Refuse	19,146	(0.069)	71,156	(0.255)
<b>CNIPA</b>				
Grant	163,978	(0.961)	4,570	(0.027)
Refuse	343	(0.002)	1,677	(0.010)

Notes: Number of patent applications falling in each category (proportion in parenthesis for each office).

Regarding the EPO, our model predicts that 91.7 percent of decisions were accurate (72% correct grants and 19.7% correct refusals), in the sense that the EPO applied its own standard in a consistent manner. The percentages of accurate decisions for the other officers were 93.8 percent, 91.6 percent, 90.9 percent and 97.1 percent for USPTO, KIPO, JPO and CNIPA, respectively. While these percentages are relatively consistent across offices, the breakdowns between correctly granted and refused, and incorrectly granted and refused, are not. Offices with low thresholds (CNIPA and USPTO) have relatively large percentages of correct acceptances (96.1% and 90.3%, respectively), whereas for JPO and EPO, who have high thresholds, the situation is reversed (65.4% and 72.0%, respectively). With respect to incorrect decisions, for CNIPA, USPTO and KIPO, we find that the proportions of mistakenly granted patents (0.2%, 1.1% and 2.2%) are smaller than the proportion of mistakenly refused patents (2.7%, 5.1% and 6.3%), an outcome which is contrary to that expected based on evidence from earlier studies. In this case, because we have a common set of applications, and because CNIPA, USPTO and KIPO have relatively low thresholds, their probabilities of not reaching those thresholds are low. For JPO and EPO, who have relatively high thresholds, the opposite is true. The probability of a grant not reaching their thresholds is relatively high.

### *5.3 Econometric decomposition of the mixed grant-refusal patent families*

In this Section we consider a decomposition of applications where decisions across offices were inconsistent into those where different decisions can be attributed to different standards, and those where there has been an office mistake. Two decompositions are presented. In the first decomposition we examine all patents with inconsistent decisions. In the second, we focus on the subset of granted patents refused elsewhere.

**Table 7. Decomposition of applications with inconsistent decisions**

	Total number	Inconsistent decisions	Correct by office standard	Violation of office standard
All	408,133	127,706 (0.313)	84,896 (0.208)	42,810 (0.105)
EPO	162,803	58,098 (0.357)	36,840 (0.226)	21,258 (0.131)
USPTO	325,068	99,085 (0.305)	64,028 (0.197)	35,057 (0.108)
JIPO	127,314	43,238 (0.340)	21,177 (0.166)	22,061 (0.173)
JPO	278,760	96,296 (0.345)	61,655 (0.221)	34,641 (0.124)
CNIPA	170,568	49,216 (0.289)	27,832 (0.163)	21,384 (0.125)

Notes: Each column contains the numbers of patent applications in each category with the proportion of the total number of applications given in parentheses.

Table 7 contains the decomposition of all applications with inconsistent decisions. Approximately 30–35 percent of applications fall into this category. After accounting for different office standards, we find that between 10.8 percent (USPTO) and 17.3 percent (KIPO) of applications have conflicting outcomes that can be attributed to an office mistake.

Of particular interest are the reasons why applications granted by one office might be refused by another. In Table 8 we limit our analysis to applications that have been granted in each office. We have decomposed these patents according to the origin of the discrepancy in treatment across offices. Reported are the total number of patents in each office (column 1); of these, the number that were refused in at least one other office (column 2, similar to Table 4); and of these, the number we estimate that were refused because other offices had a higher standard (column 3); or because the office mistakenly applied its own standards (column 4); or because another office incorrectly applied its standard (column 5). The results show that 19.2 percent of patents issued by the USPTO are most likely refused elsewhere because other offices have higher granting standards. We estimate that only 1.2 percent of USPTO patents would be deemed invalid if examined again. This contrasts with Japan. Only 1.7 percent of JPO patents are refused elsewhere. Because JPO has the most stringent standards, this non-zero percentage occurs because of differences in how offices react to foreignness, use of the PCT route, the number of claims and the priority office.<sup>x</sup> On the other hand, 9.5 percent of JPO patents are likely to be found invalid if re-examined under the same set of JPO rules.

The other offices fall between these two extremes. At the EPO, the office standard accounts for 10.5 percent of mixed grant-refusal outcomes and office mistakes account for 6.2 percent. At KIPO, different office standards are 14.2 percent and office mistakes are 2.5 percent. At CNIPA, different office standards are 16.8 percent and office mistakes are 0.2 percent.

Note that we are not justifying the stringency of each office standard, nor the legitimacy of how each office reacts to features such as foreignness, use of the PCT route, the number of claims and the priority office. We are merely asking: if the patent application was assessed again within the same jurisdiction under the same set of rules, would it be deemed valid?



**Table 8. Number (proportion) of mixed grant-refusal patent families by estimated reason, 2000–2006**

	(1)	(2)	(3)	(4)	(5)
	Granted patents	Mixed grant-refusal patent families (proportion)	Due to more lenient focal office standard (proportion)	Due to mistake applying focal office standard (proportion)	Due to a mistake applying another office standard (proportion)
EPO	124,988	26,621 (0.213)	13,151 (0.105)	7,708 (0.062)	5,762 (0.046)
USPTO	297,072	74,968 (0.252)	57,155 (0.192)	3,674 (0.012)	14,139 (0.048)
KIPO	107,501	27,591 (0.257)	15,216 (0.142)	2,739 (0.025)	9,636 (0.090)
JPO	201,335	28,083 (0.139)	3,514 (0.017)	19,146 (0.095)	5,423 (0.027)
CNIPA	164,321	44,201 (0.269)	27,576 (0.168)	343 (0.002)	16,282 (0.099)

Note: Column (1) is the total number of granted patents; Column (2) shows applications accepted by the focal office and rejected by one or more of the other offices; Column (3) is the number (proportion) of offices who made a correct decision given their office standards; Column (4) is the number (proportion) of patents that were incorrectly accepted by the focal office and rejected elsewhere; Column (5) is the number (proportion) of patent applications correctly accepted by the focal office and incorrectly rejected by at least one of the other offices.

#### 5.4 Self-selection bias

One potential concern with the estimates relates to self-selection into specific offices. Typically, applicants with weak inventions may purposefully avoid filing a patent at the strictest offices to avoid a rejection. This is potentially a concern because the patenting process is known to be costly and applicants strategically chose where to file for patent protection. However, we believe that selection is not a concern in our sample. Indeed, the analysis focuses on high-value patent applications that are filed in at least two countries. Empirical evidence suggests that applicants are inelastic to patenting costs (de Rassenfosse and van Pottelsberghe 2012). Given that applicants in our sample are active internationally, they are presumably less elastic than the average patent applicant.

To spot the presence of selection effects, we look at differences in the distribution of the inventive merit of inventions across offices. Can we see more low inventive merit inventions in more lenient offices, or are distributions similar for all offices? Table 9 presents key summary statistics for the offices. The estimates are remarkably similar across offices, suggesting that selection is not a concern.

**Table 9. Summary statistics for the distributions of the estimates of the  $c_i$** 

Statistic	All	EPO	USPTO	KIPO	JPO	CNIPA
Mean	3.6969	3.6850	3.8104	3.5750	3.7685	3.8373
Minimum	-3.8255	-3.4052	-3.8255	-3.8255	-3.5920	-3.5282
0.025 quantile	2.0958	2.1697	2.2346	2.1250	2.0859	2.3742
Median	4.6051	4.4712	4.6280	4.0443	4.8367	4.6244
0.75 quantile	4.9683	5.2128	4.9723	4.9801	5.1345	4.9878
Maximum	5.8542	5.8542	5.8542	5.8091	5.8542	5.8456
Std. Dev.	1.7179	1.7712	1.6105	1.7544	1.7725	1.5984
Observations	408133	162803	325068	127314	278760	170568

## 6. Conclusion

There is significant concern around the world that patent offices are issuing patents that should not have been granted. Studies based on litigation outcomes suggest that this problem is quantitatively significant, with the overall fraction of dubious patents ranging from a quarter to three-quarter of all patents. Our analysis of patent applications examined by multiple offices around the world suggests that the overall prevalence of low-quality patents might be smaller.

We model the patent grant process in a way in which imperfect decision-makers compare their assessment of the quality of an invention to an internal standard of quality necessary for grant. This method allows us to decompose differences in the decisions of multiple decision-makers into those that are due to a mistake by a decision-maker and those that are due to systematic differences in the policies and practices at each office. The kind of decomposition that we have undertaken requires repeated observations on each invention and each decision-making unit.

Our analysis of about 400,000 inventions considered for patent protection by multiple patent offices suggests that both sources of inconsistent decisions occur. The strength of our analysis is to compare various offices using the same invention. To push the point, it allows us to conclude that differences in grant outcomes across offices are primarily driven by office policy choices and practices rather than subjectivity of the examination process.

Specifically, we find that the fraction of patents that should not have been granted given the offices own implicit grant standard and systematic treatment of foreignness, use of the PCT route, the number of claims and the priority office, is less than 10 percent in all offices. Most of the inconsistency in the grant-refusal decision across offices is due to difference in office standards. These differences are greatest between the USPTO and the JPO; and between the CNIPA and the JPO.

Whereas the sample used for the analysis is large, it is not randomly drawn. Patents examined in multiple international jurisdictions are likely to be of higher economic value than the average patent. Our analysis of the selection problem suggests, however, that rates of low quality patents for the population of all applications to each office are unlikely to be much higher than our estimates for this

IP5 sample. Thus, even allowing for selection bias, our results imply rates of an order of magnitude lower than the rates found by litigation studies.

Teasing out the source of inconsistent examiner decisions across offices is important for understanding investor confidence and for designing policies to improve the quality of patent systems. If patents issued by one office are thought to be of a highly variable validity, then patent protection will not offer investors proposing to develop the invention much reduction in their uncertainty in that jurisdiction. On the other hand, if the inconsistency in examiners decisions across offices is due to systematic differences in office standards and procedures, then investors have a better idea of where they stand. This may increase their confidence about developing the invention in a specific jurisdiction. With respect to policy design, if uneven application of office standards is the main source of inconsistency, the appropriate policy would be to get the examiners from the said office on the same page (or to investigate examiners whose own reports vary widely). If systematic office differences are the main source of inconsistency, then the policy is either to review the legal standards or to investigate the institutional procedures that are believed to lead to a variable standard.

One of our remarkable findings is our much lower rate of ‘mistaken’ patents compared to litigation studies. We do not know the reason for this finding but suggest four possibilities. First, litigated patents are highly selected toward those with greater uncertainty about their validity. Clearer cases can be expected to be settled out of court. Second, litigation studies implicitly assume that courts apply the same standard as that of the office whose grant is being reviewed, and do not make mistakes themselves. In practice, it is possible that courts systematically apply a stricter standard for validity than the patent office—and make mistakes themselves. Third, although patent applications in our sample are examined by up to five examiners from very different cultures and language groups, every examiner spends considerably less time than if the patent were re-examined in litigation. It is possible that all examiners have mistakenly granted the patent. Finally, review by a court is fundamentally different from review by another examiner because the court review is an adversarial proceeding. It is possible that there is prior art that no patent examiner will ever find, but which the adverse party is able to bring to the court’s attention. Thus, overall our results provide a different perspective on patent quality and should be viewed as complementary to those of litigation studies rather than directly comparable.

The magnitude of the difference between the figures presented in this paper and the figures obtained using patent litigation data bear important implications for discussions about patent quality. One difficulty in interpreting the difference is that we do not know how much of it might be due to selection bias in the litigation studies. But if we assume for the sake of argument that invalidity in the view of the courts is truly significantly higher than invalidity in the view of the offices, we can make four general points.

First, much of the debate around quality focuses on improving examination. Our results suggest that examiner error is low in most jurisdictions, and therefore this effort will be only marginally useful. Second, some of the debate has a flavor of the United States, in particular, having a low standard. Our

results suggest that it is true that the U.S. standard is somewhat low, corresponding to about 19 percent more patents granted than the JPO if both offices were acting consistently. Thus, raising the U.S. standard to the level of the highest country would eliminate some low-quality patents, but perhaps not as many as some commentators believe. Third, more generally, the tone of the debate is frequently that the uncertainty around validity is the patent offices' fault. Our results suggest rather that it is inherent in the examination process that a non-trivial number of invalid patents will be approved. Finally, we bring into sharp focus the question of *why* courts are more likely to invalidate than examiners. To the extent that it is because of the adversarial nature of litigation, the finding brings the question of how to best to organize re-examination processes that are undertaken within offices. But if it is because judges are fundamentally tougher than examiners, the finding raises deeper questions about administrative law, since judges are not supposed to apply different standards.

The findings presented in this paper are interesting in light of concerns about patent quality, but they also contribute to current policy discussions on patent prosecution highway (PPH) agreements. PPH designates a set of initiatives for providing accelerated prosecution procedures by sharing information between patent offices. Our results show that there is considerable heterogeneity across offices. The PPH agreements intend to increase the harmonization of decision. However, they may also propagate a wrong decision into the whole patent family, further weakening patent rights. Our results further illustrate that some offices are more accurate than others, which may create additional tensions in the context of PPH agreements.

Finally, our analysis is silent on the optimal level of ambiguity that the patent system should tolerate. On the one hand, low quality patents hurt businesses and may slow down the pace of technological progress. On the other hand, ensuring high-quality examination is costly, especially because the majority of patents have limited economic potential. In our analysis we have controlled for how each office reacts systematically to foreignness, use of the PCT route, the number of claims and the priority office. We do not justify these differences but note them. Future research should investigate whether delivering more harmonized outcomes for businesses is likely to improve welfare. Our results provide a useful starting point in that regard.

## Acknowledgments

Dan L. Burk, David Card, Annamaria Conti, Dietmar Harhoff, Joachim Henkel, Karin Hoisl, Sonia Jaffe, Keld Laursen, Yann Ménière, Alfons Palangkaraya, Arti Rai, Ben Roin and Carl Shapiro provided valuable comments. The authors are also grateful to seminar and conference participants at the NBER Summer Institute, Duke Law School, Congress of the European Economic Association, European Patent Office, European Association for Research in Industrial Economics Conference, European Policy for Intellectual Property Conference, Toulouse School of Economics, ETH Zurich, the third International Meeting in Law & Economics (Paris), the Munich Summer Institute, New Zealand Economic Association Conference, and Asia Pacific Innovation Conference, T’Mir Julius provided excellent research assistance and her contribution is gratefully acknowledged. This study was financed by the Australian Research Council Discovery Grant ARC LP110100266 ‘The Efficiency of the Global Patent System’ with partners IP Australia and the Institute of Patent and Trademark Attorneys.

## References

- Albert, J.H. & Chib, S. (1993). Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association*, 88, 669–679.
- Albert, M. B., Avery, D., Narin, F., & McAllister, P. (1991). Direct validation of citation counts as indicators of industrially important patents. *Research policy*, 20(3), 251–259.
- Allison, J., & Lemley, M. (1998). Empirical evidence on the validity of litigated patents. *American Intellectual Property Law Association Quarterly Journal*, 26(3), 185–275.
- Allison, J., & Hunter, S. (2006). On The Feasibility Of Improving Patent Quality One Technology At A Time: The Case Of Business Methods. *Berkeley Technology Law Journal*, 21, 729–794.
- Brooks, S., Gelman, A., Jones, G.L. & Meng Xiao-Li, editors (2011). *Handbook of Markov Chain Monte Carlo*. Boca Raton: CRC Press.
- Brunner, J. (2014). Patent Prosecution as Dispute Resolution: A Negotiation Between Applicant and Examiner. *Journal of Dispute Resolution*, 3(1), 7–21.
- Cremers, K., Ernicke, M. Gaessler, F. Harhoff, D. Helmers, C. McDonagh, L. Schliessler, P. Van Zeebroeck, N. (2013) Patent litigation in Europe, *ZEW Discussion Paper* 13-072.
- Choi, J. P., & Gerlach, H. (2016). Patent pools, litigations, and innovation. *RAND Journal of Economics*, forthcoming.
- de Rassenfosse, G., & van Pottelsberghe, B. (2012). On the price elasticity of demand for patents. *Oxford Bulletin of Economics and Statistics*, 74(1), 58–77.
- de Rassenfosse, G., Dernis, H., Guellec, D., Picci, L., & van Pottelsberghe, B. (2013). The worldwide count of priority patents: A new indicator of inventive activity. *Research Policy*, 42(3), 720–737.
- de Rassenfosse, G., Jensen, P., Julius, T., Palangkaraya, A., Webster, E. (forthcoming). Are foreigners treated equally under TRIPS? *Journal of Law & Economics*, accepted.
- de Saint-Georges, M., & van Pottelsberghe, B. (2013). A quality index for patent systems. *Research Policy*, 42(3), 704–719.

- The Economist (2015). A question of utility. August 8 2015, pp 50–52.
- Encaoua, D., & Lefouili, Y. (2009). Licensing ‘weak’ patents. *Journal of Industrial Economics*, 57(3), 492–525.
- Farrell, J., & Shapiro, C. (2008). How strong are weak patents? *American Economic Review*, 98(4), 1347–1369.
- Frakes, M. & Wasserman, M. (2017). Is the Time Allocated to Review Patent Applications Inducing Examiners to Grant Invalid Patents? Evidence from Micro-Level Application Data, *Review of Economics and Statistics*, 99(3), 550–563.
- Gelman, A., J. B. Carlin, H.S. Stern, D.B. Dunson, A. Vehtari and D.B. Rubin (2014), *Bayesian Data Analysis*, 3rd edition, Boca Raton FL: CRC Press.
- Graham, S. & Mowery, D. (2004). Submarines in software? continuations in US software patenting in the 1980s and 1990s. *Economics of Innovation and New Technology*, 13, 443–456.
- Hall, B., Jaffe, A., & Trajtenberg, M. (2001). The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools. *NBER Working Paper* 8498.
- Hall, Bronwyn H., Adam Jaffe, and Manuel Trajtenberg (2005). Market value and patent citations. *RAND Journal of Economics*, 36(1), 16–38.
- Henkel, J. and Zischka, H. (2014). Why most patents are invalid – Extent, reasons, and potential remedies of patent invalidity. mimeo, TUM School of Management, Technische Universität München, 29 September 2014.
- Jaffe, Adam B. and Josh Lerner (2004). *Innovation and Its Discontents: How Our Broken Patent System is Endangering Innovation and Progress, and What to Do About It*. Princeton: Princeton University Press.
- Lazaridis, G., & van Pottelsberghe, B. (2007). The rigour of EPO’s patentability criteria: An insight into the ‘induced withdrawals.’ *World Patent Information*, 29(4), 317–326.
- Lei, Z., & Wright, B. (2017). Why weak patents? Testing the examiner ignorance hypothesis. *Journal of Public Economics*, 148, 43–56.
- Lemley, M. A. (2001). Rational ignorance at the patent office. *Northwestern University Law Review* 95.4.
- Lemley, M. A., & Sampat, B. (2012). Examiner Characteristics and Patent Office Outcomes. *Review of Economics and Statistics*, 94(3), 817–827.
- Lemley, M. A., & Shapiro, C. (2005). Probabilistic Patents. *Journal of Economic Perspectives*, 19(2), 75–98.
- Marco, A. (2004). The selection effects (and lack thereof) in patent litigation: Evidence from trials. *The B.E. Journal of Economic Analysis & Policy*, 4(1), 1226.
- Merges, R., & Nelson, R. (1990). On the complex economics of patent scope. *Columbia Law Review*, 90(4), 839–916.
- Meurer, M. J. (2009). Patent Examination Priorities. *William & Mary Law Review*, 51(2), 675–709 .
- Miller, S. (2013). Where’s the innovation? An analysis of the quantity and qualities of anticipated and obvious patents. *Virginia Journal of Law and Technology*, 18(1), 1–58.

- Nagaoka, S. and Yamauchi, I. (2015). Information constraint of the patent office and examination quality: Evidence from the effects of initiation lag. Mimeo.
- Narin, F. (1995). Patents as indicators for the evaluation of industrial research output. *Scientometrics*, 34(3), 489–496.
- OECD (2003) Science, Technology and Industry Scoreboard, Organization for Economic Cooperation and Development, Paris.
- Ormerod, J.T. and M.P. Wand (2010), Explaining variational Approximations, *American Statistician* 64(2), 140–153.
- Palangkaraya, A., Webster, E. & Jensen, P. (2011). Misclassification between patent offices: Evidence from a matched sample of patent applications. *Review of Economics and Statistics*, 93(3), 1063–1075.
- Paradise, J., Andrews, L., & Holbrook, T. (2005). Patents on Human Genes: An analysis of Scope and Claims. *Science*, 307, 1566–1567.
- Picard, P., & van Pottelsberghe, B. (2013). Patent office governance and patent examination quality. *Journal of Public Economics*, 104, 14–25.
- Quillen, C., & Webster, O. (2001). Continuing patent applications and performance of the US patent and trademark office. *Federal Circuit Bar Journal*, 11, 1–21.
- Sampat, B. & Shadlen, K. (2015). TRIPS implementation and secondary pharmaceutical patenting in Brazil and India. *Studies in Comparative International Development*, 50 (2), 228–257.
- Schuett, F. (2013). Patent quality and incentives at the patent office. *The RAND Journal of Economics*, 44(2), 313–336.
- Trajtenberg, M. (1990). A penny for your quotes: Patent citations and the value of innovations. *RAND Journal of Economics*, 21(1), 172–187.
- Trajtenberg, M., Henderson, R., & Jaffe, A. (1997). University versus corporate patents: A window on the basicness of invention. *Economics of Innovation and New Technology*, 5(1), 19–50.
- Wallach, E. J., & Darrow, J. J. (2016). Federal Circuit Review of USPTO Inter Partes Review Decisions, by the Numbers. *Journal of the Patent & Trademark Office Society*, 98, 105.
- Webster, E., Jensen, P., & Palangkaraya, A. (2014). Patent examination outcomes and the national treatment principle. *RAND Journal of Economics*, 45(2), 449–469.

## Appendix: Details for estimation

The first step for estimation is the assignment of priors for the parameters  $(c_i, \tau_j, \boldsymbol{\beta}_j)$  in the model  $y_{ij}^* = c_i - \tau_j - (\mathbf{x}_{ij} - \bar{\mathbf{x}}_j) \boldsymbol{\beta}_j + \varepsilon_{ij}$ .

Adding a constant to all the  $c_i$  and adding the same constant to all the  $\tau_j$  leaves the model unchanged, implying all  $c_i$  and  $\tau_j$  are not identified. To anchor the location, we set  $\tau_5 = 0$ . For the remaining  $\tau_j$ ,  $j=1,2,3,4$ , we assign the independent priors  $\tau_j \sim (\tau_0, a^{-1})$ , with the settings  $\tau_0 = 0$  and  $a = 0.01$ . The relatively large prior variance of  $a^{-1} = 100$  is designed to be noninformative, making the estimates of the  $\tau_j$  largely determined by the data. For the same reasons, the prior we assign to  $\boldsymbol{\beta}_j$  is  $\boldsymbol{\beta}_j \sim N(\boldsymbol{\beta}_0, a^{-1} \mathbf{I})$ ,  $j=1,2,\dots,5$  with  $\boldsymbol{\beta}_0 = \mathbf{0}$ . For  $c_i$ , we assign independent hierarchical priors  $c_i \sim N(\mu_c, h^{-1})$ ,  $i=1,2,\dots,M$ , where  $M = 408,133$  is the number of patents, and with  $\mu_c \sim N(c_0, w_0^{-1})$ ,  $c_0 = 0$  and  $h_0 = w_0 = 0.01$ . These settings are designed to provide scope for improving the efficiency of estimation for the  $c_i$ , and, at the same time, allowing the data to provide sufficiently wide variation in the  $c_i$ . We also assume all  $c_i$ ,  $\tau_j$  and  $\boldsymbol{\beta}_j$  are a priori independent.

To describe details of the variational Bayes estimation procedure, it is convenient to write the observations for the  $j$ -th office in matrix notation as

$$\mathbf{y}_j^* = \mathbf{L}_j \mathbf{c} + \begin{bmatrix} -\mathbf{1}_{N_j} & -\mathbf{X}_j \end{bmatrix} \begin{bmatrix} \tau_j \\ \boldsymbol{\beta}_j \end{bmatrix} + \boldsymbol{\varepsilon}_j \quad j=1,2,3,4$$

$$\mathbf{y}_5^* = \mathbf{L}_5 \mathbf{c} - \mathbf{X}_5 \boldsymbol{\beta}_5 + \boldsymbol{\varepsilon}_5$$

Here,  $\mathbf{y}_j^*$  is a  $N_j \times 1$  vector measuring propensities of the  $j$ -th office to accept its patent applications,  $\mathbf{c}$  is an  $M \times 1$  vector containing the merit associated with each of the  $M$  patent applications,  $\mathbf{L}_j$  is a  $N_j \times M$  selection matrix that chooses from  $\mathbf{c}$  the patents considered by the  $j$ -th office; each of its rows contains one 1 and  $M-1$  zeros,  $\mathbf{1}_{N_j}$  is a  $N_j$ -dimensional vector of 1's,

$$\mathbf{X}'_j = \left( (\mathbf{x}_{1j} - \bar{\mathbf{x}}_1)', (\mathbf{x}_{2j} - \bar{\mathbf{x}}_2)', \dots, (\mathbf{x}_{N_j j} - \bar{\mathbf{x}}_{N_j})' \right) \quad \text{and} \quad \boldsymbol{\varepsilon}'_j = [\varepsilon_{1j}, \varepsilon_{2j}, \dots, \varepsilon_{N_j j}].$$

For  $j=1,2,3,4$ , we define  $\mathbf{Z}_j = \begin{pmatrix} -\mathbf{1}_{N_j} & -\mathbf{X}_j \end{pmatrix}$  and  $\boldsymbol{\gamma}'_j = (\tau_j \quad \boldsymbol{\beta}'_j)$ . For  $j=5$ , we write  $\mathbf{Z}_5 = -\mathbf{X}_5$  and  $\boldsymbol{\gamma}_5 = \boldsymbol{\beta}_5$ , so that the model becomes

$$\mathbf{y}_j^* = \mathbf{L}_j \mathbf{c} + \mathbf{Z}_j \boldsymbol{\gamma}_j + \boldsymbol{\varepsilon}_j$$

Let  $k_i$  be the number of offices who consider patent application  $i$ . Then, the total number of observations is  $N = \sum_{j=1}^5 N_j = \sum_{i=1}^M k_i = 1,064,513$ . Let  $\mathbf{y}$  be an  $(N \times 1)$  vector containing all the observations on the  $y_{ij}$ .

Given this background and notation, the following variational Bayes approximate posterior densities for  $(\mathbf{y}_j^*, \mathbf{c}, \boldsymbol{\gamma}_j, \mu_c)$  can be derived. Derivations are provided in the supplementary material. The densities for the  $y_{ij}^*$  are the truncated normal densities



$$q(y_{ij}^* | \mathbf{y}) = \begin{cases} I(y_{ij}^* \geq 0) \times N(\bar{c}_i + \mathbf{z}_{ij} \bar{\gamma}_j, 1) & \text{when } y_{ij} = 1 \\ I(y_{ij}^* < 0) \times N(\bar{c}_i + \mathbf{z}_{ij} \bar{\gamma}_j, 1) & \text{when } y_{ij} = 0 \end{cases}$$

For the  $c_i$ ,  $\mu_c$  and  $\gamma_j$ , we have the normal densities,

$$q(c_i | \mathbf{y}) = N \left[ \frac{\sum_{j=1}^5 \delta_{ij} (\bar{y}_{ij}^* - \mathbf{z}_{ij} \bar{\gamma}_j) + h \bar{\mu}_c}{k_i + h}, \frac{1}{k_i + h} \right]$$

$$q(\mu_c | \mathbf{y}) = N \left[ \frac{h \sum_{i=1}^M \bar{c}_i + w_0 c_0}{hM + w_0}, \frac{1}{hM + w_0} \right]$$

$$q(\gamma_j | \mathbf{y}) = N[\bar{\gamma}_j, \mathbf{V}(\gamma_j)]$$

where  $\bar{\gamma}_j = (\mathbf{Z}'_j \mathbf{Z}_j + a \mathbf{I}_j)^{-1} [\mathbf{Z}'_j (\bar{\mathbf{y}}_j^* - \mathbf{L}_j \bar{\mathbf{c}}) + a \gamma_{0j}]$ ,  $\mathbf{V}(\gamma_j) = (\mathbf{Z}'_j \mathbf{Z}_j + a \mathbf{I}_j)^{-1}$ , and the prior means are  $\gamma'_{0j} = (\tau_0, \beta'_0)$  for  $j=1,2,3,4$  and  $\gamma_{05} = \beta_0$ . Also,  $\mathbf{z}_{ij}$  is a row of  $\mathbf{Z}_j$ ,  $\delta_{ij}=1$  if the  $i$ -th patent is considered by the  $j$ -th office and zero otherwise,  $I(\cdot)$  is an indicator function equal to one if its argument is true and zero otherwise, and  $(\bar{y}_{ij}^*, \bar{c}_i, \bar{\gamma}_j, \bar{\mu}_c)$  are posterior means obtained by solving iteratively the equations given below.

$$\bar{y}_{ij}^* = \bar{c}_i + z_{ij} \bar{\gamma}_j + \frac{\phi(\bar{c}_i + z_{ij} \bar{\gamma}_j)}{\Phi(\bar{c}_i + z_{ij} \bar{\gamma}_j)} \quad \text{when } y_{ij} = 1$$

$$\bar{y}_{ij}^* = \bar{c}_i + z_{ij} \bar{\gamma}_j - \frac{\phi(\bar{c}_i + z_{ij} \bar{\gamma}_j)}{1 - \Phi(\bar{c}_i + z_{ij} \bar{\gamma}_j)} \quad \text{when } y_{ij} = 0$$

$$\bar{c}_i = \frac{\sum_{j=1}^5 \delta_{ij} (\bar{y}_{ij}^* - \mathbf{z}_{ij} \bar{\gamma}_j) + h \bar{\mu}_c}{k_i + h}$$

$$\bar{\mu}_c = \frac{h \sum_{i=1}^M \bar{c}_i + w_0 c_0}{hM + w_0}$$

$$\bar{\gamma}_j = (\mathbf{Z}'_j \mathbf{Z}_j + a \mathbf{I}_j)^{-1} [\mathbf{Z}'_j (\bar{\mathbf{y}}_j^* - \mathbf{L}_j \bar{\mathbf{c}}) + a \gamma_{0j}]$$

To solve for  $(\bar{y}_{ij}^*, \bar{c}_i, \bar{\gamma}_j, \bar{\mu}_c)$ , we iterate the above equations until the lower bound of the log of the marginal likelihood converges. Once convergence has been achieved, the variational Bayes approximate posterior densities are completely defined by the values for  $(\bar{y}_{ij}^*, \bar{c}_i, \bar{\gamma}_j, \bar{\mu}_c)$ . These values also represent the Bayesian estimates. In the supplementary material, the lower bound of the log of the marginal likelihood is shown to be

$$\begin{aligned}
\log \underline{ML} &= E_q \left[ \log p(\mathbf{y}, \mathbf{y}^*, \mathbf{c}, \boldsymbol{\gamma}, \mu_c) - \log q(\mathbf{y}^*, \mathbf{c}, \boldsymbol{\gamma}, \mu_c) \right] \\
&= C + \sum_{j=1}^5 \sum_{i=1}^{N_j} y_{ij} \log \Phi(\bar{c}_i + \mathbf{z}_{ij} \bar{\boldsymbol{\gamma}}_j) + \sum_{j=1}^5 \sum_{i=1}^{N_j} (1 - y_{ij}) \log \Phi(-\bar{c}_i - \mathbf{z}_{ij} \bar{\boldsymbol{\gamma}}_j) \\
&\quad - \frac{h}{2} (\bar{\mathbf{c}} - \bar{\mu}_c \mathbf{1}_M)' (\bar{\mathbf{c}} - \bar{\mu}_c \mathbf{1}_M) - \frac{a}{2} \sum_{j=1}^5 (\bar{\boldsymbol{\gamma}}_j - \boldsymbol{\gamma}_{0j})' (\bar{\boldsymbol{\gamma}}_j - \boldsymbol{\gamma}_{0j}) - \frac{w_0}{2} (\bar{\mu}_c - c_0)^2
\end{aligned}$$

where  $C$  is a constant.

## End notes

---

<sup>i</sup> See <<https://www.uspto.gov/patent/initiatives/2016-patent-quality-chats>>

<sup>ii</sup> Because applicants must submit twin applications to foreign jurisdictions shortly after the submission of the priority filing (up to 12 or 31 months after), the decision to submit twin applications is not driven by the outcome of examination in the office of priority. There is thus no selection on actual grant outcome. Other selection effects may be at work, as discussed in Section 6.2.

<sup>iii</sup> There were 1,821,150 patent applications filed worldwide in 2010 (priority plus second filings). Of these, 1,452,925 (79.8%) were filed in the IP5 offices (PATSTAT Autumn 2014 version).

<sup>iv</sup> For instance, the fact that some examiners might be more lenient than others will fall in the error term. By contrast, the fact that examiners might be collectively more lenient than what they should be will fall in the office-specific effect, hence our use of the term *de facto* standard.

<sup>v</sup> Thus, our sample may include a priority patent application filed, say, at the Brazilian patent office and with an equivalent at the EPO and the USPTO.

<sup>vi</sup> The ‘Paris Convention’ route is the traditional filing route for patent applications (sometimes call the national route). The term PCT stands for ‘Patent Cooperation Treaty.’ It is an international treaty that facilitates international patenting.

<sup>vii</sup> Note that some equivalents are filed partly through the PCT route and partly through the Paris route, leading to within-twin variation.

<sup>viii</sup> The low proportion of local applicants at the CNIPA reflects the fact that very few Chinese firms apply for patent protection in foreign jurisdictions, which is a pre-condition for being in the sample.

<sup>ix</sup> Differences in the number of claims across offices must be interpreted cautiously. For example, the EPO allows for claims with many alternatives and preferred embodiments compared to the USPTO so that differences in the number of claims do not necessarily indicate differences in the scope of protection. Furthermore, the number of claims is affected by other institutional factors such as claim-based fees.

<sup>x</sup> Because the office standard accounts for how the office treats foreign applicants, PCT applications, the number of claims and the priority office, this amount in the JPO is not zero.