

2012

A New Look at Patent Quality: Relating Patent Prosecution to Validity

Ronald J. Mann
Columbia Law School, rmann@law.columbia.edu

Marian Underweiser
munderw@us.ibm.com

Follow this and additional works at: https://scholarship.law.columbia.edu/faculty_scholarship



Part of the [Intellectual Property Law Commons](#)

Recommended Citation

Ronald J. Mann & Marian Underweiser, *A New Look at Patent Quality: Relating Patent Prosecution to Validity*, JOURNAL OF EMPIRICAL LEGAL STUDIES, VOL. 9, P. 1, 2012; COLUMBIA LAW & ECONOMICS WORKING PAPER No. 381 (2012).

Available at: https://scholarship.law.columbia.edu/faculty_scholarship/1728

This Working Paper is brought to you for free and open access by the Faculty Publications at Scholarship Archive. It has been accepted for inclusion in Faculty Scholarship by an authorized administrator of Scholarship Archive. For more information, please contact cls2184@columbia.edu.



Journal of Empirical Legal Studies
Volume 9, Issue 1, 1–32, March 2012

A New Look at Patent Quality: Relating Patent Prosecution to Validity

*Ronald J. Mann and Marian Underweiser**

The article uses two hand-collected data sets to implement a novel research design for analyzing the precursors to patent quality. Operationalizing patent “quality” as legal validity, the article analyzes the relation between Federal Circuit decisions on patent validity and three sets of data about the patents: quantitative features of the patents themselves, textual analysis of the patent documents, and data collected from the prosecution histories of the patents. The article finds large and statistically significant relations between ex post validity and both textual features of the patents and ex ante aspects of the prosecution history (especially prior art submissions and the existence of internal patent office appeals before issuance). The results demonstrate the importance of refocusing analysis of patent quality on replicable indicators like validity, and the value that more comprehensive collection of prosecution history data can have for improving the output of the patent prosecution process.

I. INTRODUCTION

What can justify another article about patent quality? No respected observer would deny the significance of the difficulties facing the patent system or that the central problem is a decline in the quality of patents (Jaffe & Lerner 2004; National Research Council of the National Academies 2004; Bessen & Meurer 2008). Most would agree that the system has failed recently; the conventional view is that Congress overshot the mark when it gave the Federal Circuit a mandate to ease the standard applied to assess patent validity (Jaffe & Lerner 2004). Nobody believes the problems are trivial or academic; it is a given that poor-quality patents cause a drag on the competitiveness of the national economy (Bessen & Meurer 2008; Burk & Lemley 2009). This consensus has legislators poised to make substantial changes in the patent system to limit the costs that bad patents impose on

*Address correspondence to Ronald Mann, Columbia Law School, 435 W. 116th St., New York, NY 10027; email: rmann@law.columbia.edu. Mann is Professor of Law and Co-Director of the Center for Transactional Studies, Columbia Law School; Underweiser is Intellectual Property Counsel at IBM.

We acknowledge comments on presentations at workshops at the law schools at Columbia and Michigan, as well as the 2008 American Law and Economics Association Annual Meeting, a 2009 RSA Conference held at Columbia Law School, presentations at the Patent and Trademark Office, and a workshop at the law firm of Woodcock Washburn. We thank Rochelle Dreyfuss, Marc Ehrlich, Hal Edgar, Clarisa Long, Manny Schecter, David Schwartz, and John Weresh for comments of noteworthy value, and Jim Carlson, Arka Chatterjee, and Greg Odegaard for exemplary assistance with the collection and organization of the prosecution history data set. This article does not reflect the views of IBM.

innovators (H.R. 1260 and S. 515, 111th Cong., 1st Sess.), the Patent and Trademark Office (PTO) moving aggressively to take a larger role in the system (Long 2009), and leading academics proposing a lengthy menu of systemic changes (Benjamin & Rai 2007; Lichtman & Lemley 2007; Nard & Duffy 2007).

Yet there is a stark mismatch between the burgeoning academic literature about patent quality and the legal issues motivating calls for system reform. Thus, a methodologically intricate econometrics literature has used data about references in and to patents to measure the extent to which patents reflect important innovations or disseminate information to future researchers (Jaffe & Trajtenberg 2002; Owen Smith & Powell 2003; Thompson & Kean 2005; Singh 2005; Alcácer & Chung 2007). Other scholars have developed multivariable indicators that attempt to measure patent value, generally building on the intuition that litigated patents are likely to be more valuable than nonlitigated patents (Allison et al. 2004; Lanjouw & Schankerman 2004; Allison & Mann 2007), but the existing scholarship contributes little to debates about system reform because it rests on conceptions of quality that say little or nothing about the strengths and weaknesses of the patenting process: the PTO has neither control nor responsibility for the frequency with which issued patents are cited or involved in litigation.

This article takes a different approach to quality. Instead of examining the information content or economic value of patents, we investigate legal validity as a metric of quality. To operationalize that metric, we collect all patents for which the Federal Circuit has considered validity since 2003. To build a robust understanding of the features that might relate to validity, we do not limit our analysis to the features apparent on the face of the patents (the focus of the existing literature); rather, we also use the textual features of the patent (how well the specification and claims are aligned) and detailed information from the prosecution histories.

The research design allows us to analyze the empirical relation between the two parts of the patent system: the interaction between the work of the applicant and examiner and the reliability of the patent that issues. If the heart of the policy concern about patents is a suspicion that the applicant's work in preparing applications or the PTO's work in examining them is not sufficiently rigorous, the data presented here speak to the issue much more directly than analyses of economic value or inventiveness that have characterized prior work on patent quality.

Section II of the article discusses our conception of patent quality and its precursors in the prosecution process. Section III discusses the data. Section IV presents the results. Section V discusses the implication of the results for patent reform. Section VI briefly concludes.

II. PATENT QUALITY AND ITS PRECURSORS

The first step in the study of patent "quality" is to settle on an understanding of patent quality that can be operationalized and empirically tested. Because the term "quality" is itself so general, it should not be surprising that different groups of scholars have used the term to examine distinct concepts relevant to their own interests.

The most advanced literature about patent quality, characterized by the seminal work of Jaffe, Trajtenberg, and Henderson, has analyzed the extent to which patents reflect and facilitate the diffusion of knowledge, as evidenced by citations to and in patents. Important topics in that literature have included differences between patents issued to university and commercial researchers (Henderson et al. 1998), the relative importance of geographical and network ties (Jaffe et al. 1993; Jaffe & Trajtenberg 2002; Thompson & Kean 2005; Henderson et al. 2005; Singh 2005; Thompson 2006; Singh et al. 2010; Sorenson & Singh 2007), and the relative importance of lone inventors as compared to teams of inventors (Wuchty et al. 2007; Singh & Fleming 2010).

The basic premise of that literature resonates with the fundamental role of the patent system in driving the development and disclosure of new technology. Hence, if the primary purpose of the patent system is to force the public disclosure of new technology, then patent quality is a function of the importance of the technology disclosed in issued patents and the extent to which the disclosure facilitates use of the technology by subsequent researchers. At the same time, many aspects of legal validity¹ are largely irrelevant to this perspective. For example, the existence of overbroad claims does nothing to undermine the quality of a patent that discloses an important technology.


A second literature, which bridges finance and legal scholarship, has analyzed quality in the sense of economic value. The premise of this literature is that the patent system is doing a better job if the patents are worth more and a worse job if they are worth less. Related to that point, increases and decreases in the value of patents have important implications for corporate finance and investment (Rivette & Kline 1999). The fundamental difficulty that literature has faced has been the problem of valuing patents. It is difficult to observe patent value directly because sales of patents on an open market are quite uncommon, and because the terms of licensing transactions normally are proprietary and often involve nonmonetary considerations (like cross-licensing).

Hence, scholars in that literature have used various proxies for value. For example, the most prominent papers from legal academics examine the differences between litigated patents and unlitigated patents (Allison et al. 2004) or between patents for which maintenance fees are paid and those for which maintenance fees are not paid (Moore 2005; Bessen 2008). Others, building on that work, have attempted to identify particular types of patents that have the characteristics of value identified in the studies of litigated and maintained patents (Lanjouw & Schankerman 2004; Allison & Mann 2007). Finance scholars, studying patents held by publicly traded firms, have examined the relation between patent portfolios and firm value (Hall & MacGarvie 2006). Still others (exclusively in the European Union) have surveyed patentees or their attorneys (Sapsalis et al. 2006; Harhoff et al. 2003; Reitzig 2004).

Like the diffusion metric, the value metric is only loosely related to validity. Thus, the economic value of the patent will depend both on the likelihood that the patent is valid and

¹Technically, Federal Circuit decisions do not hold a patent valid or invalid, they hold only that the challenged portions of the patent are invalid or not invalid. For convenience, we refer loosely to such patents as valid and invalid. We discuss in detail below the implications of this for statistical analysis of those decisions.

Figure 1: Precursors of patent quality.



on the value of the market in which the patent is deployed. For example, a poorly drafted patent of dubious validity might be worth tens (or hundreds) of millions of dollars if it purports to claim rights to a valuable product (like the Blackberry or Microsoft Word). Conversely, a patent drafted with sterling clarity and undoubted novelty might be worth little or nothing if the product that it describes is unmarketable. This is important because an emphasis on that metric would suggest no basis for concern if we were subsequently to learn that the patents at issue in the Blackberry and Word cases—each of which was worth tens if not hundreds of millions of dollars when issued—were invalid. Because most policy analysts focus their attention on the issuance of patents that are legally dubious, a focus on economic value seems at best incomplete.

Attempting to avoid the limitations of the diffusion and value metrics, this article conceives of quality as legal validity. A focus on legal validity has several advantages for present purposes. First, many of the problems attributed to the patent system relate to patents of dubious validity. Second, although legal validity is itself difficult to predict and subject to change over time, it is still more objective than an inquiry into the importance or innovative flair of the invention. As discussed above, those inquiries are likely to dissolve into an investigation into loose proxies for the monetary value of the patent or some post hoc consideration of its societal impact. Although it is important for businesses to understand the value of their patents, monetary value is of relatively little use in a project focused on improving the patent system. Similarly, although motivating invention and diffusion of knowledge are central to the success of the patent system, the focus of present concerns is on the dubious and uncertain validity of issued patents. For empirical analysis to inform those concerns it must analyze them directly. In sum, to improve the process through which patents are issued, it is important to understand the relation between characteristics of that process and the validity of the patents it produces.²

Having settled on validity as the relevant aspect of quality, the next step is to consider what features of the application and prosecution process might relate to better or worse quality (Figure 1). We analyze the validity of the patent as a function of three distinct sets of inputs: the invention, effort by the applicant, and effort by the examiner. The nature of the invention is important to the validity of the patent because the “inventiveness” that may

²The conceptions of quality we discuss here relate to the public benefits of the patent system. Thus, we are not considering the “quality” of the patent in the sense of the likelihood that it secures for the applicant the broadest possible monopoly over the disclosed invention.

(or may not) distinguish the invention from prior art is a central component of patent validity. A patent based on an invention that lacks inventiveness will have poor quality because it lacks the requisite novelty and nonobviousness required for legal validity; conversely, a patent based on an invention with inventiveness at least has the chance of issuing with high quality. To be sure, an invention that makes only a small advance over prior art still could produce a patent of high quality if the distinction from prior art is not obvious and is carefully specified and delineated. But the point is that legal validity depends on an inherent feature of the invention—not within the control of the applicant or the examiner—a sort of “technological truth” about the difference between the invention and the prior art.

The second step of the process is the application. We distinguish two distinct concepts that inform the quality of the application: the drafting effort of the applicant, and the search effort of the applicant. Thus, the quality of the application depends on the effort the applicant puts into accurately specifying the nature of the invention and defining the claims that will set the patent apart from prior art. The quality of the application also depends on the extent to which the applicant searches for prior art because the applicant’s knowledge of prior art will affect the accuracy and effectiveness of the applicant’s discussion of the patent’s inventiveness. An applicant that does not fully understand the prior art related to the patent cannot produce an application of the same quality as one that does.

The third and most important step of the process is the prosecution process. We distinguish four distinct concepts that inform the quality of the prosecution process. The examiner’s effort in criticizing the way the patent is drafted is important because it enhances the likelihood that the final patent will be well crafted, so that (for example) the claims will be definite in light of the specification. The examiner’s search effort affects the likelihood that the claims accurately distinguish the disclosed invention from prior art. The applicant’s effort in providing information about the invention and in responding to the examiner’s input is also important to the quality of the final product. For example, even if the applicant understands the prior art, the quality of the process is undermined if the applicant fails to disclose the prior art to the examiner. There is significant potential for cooperation; because both parties are searching for the same prior art, a better search effort by the applicant lowers the necessary search effort for the examiner, and vice versa. Similarly, there is a great deal of room for more (and less) cooperative efforts by the applicant, amending claims readily or grudgingly in response to newly discovered prior art. Finally, where applicable, the review process of the PTO Board provides an opportunity to oversee the work of the examiners and thus to improve the final quality of the office’s output.

Several things about this discussion warrant emphasis. First, the conception is dynamic: the quality of the issued patent is likely to differ from the quality of the application initially filed, and the extent of this difference will depend on what happens during the prosecution process. Second, quality (validity) depends on the joint effort that the examiner and applicant bring to the process: the better that effort, the more the process will improve the application before issuing the resulting patent, and the greater the likelihood that the resulting patent will be valid. Neither party standing alone can ensure that the outcome of the process is a patent of high quality and the successful fulfillment of the

responsibilities of each party depends on input from the other. To put it another way, because each party has unique and important capabilities, a system that does not motivate both parties to work toward the production of quality patents cannot be expected to succeed.

III. DATA COLLECTION

As the discussion above suggests, the research design relates quantitative information about the patents and the process that led to their issuance to the ultimate validity of the patents. This requires us first to define the population of valid and invalid patents, and then to collect data about those patents that relate to the concepts. Specifically, as we discuss in detail in the succeeding sections, the data set includes three types of information about the patents in the population. First, following prior researchers (Allison et al. 2004; Moore 2005; Lanjouw & Schankerman 2004), we collected as much automated information about quantitative features evident on the face of the patents (claims, references, classes, etc.) as was practicable. Second, using content analysis, we estimated the degree of alignment between the claims and specification of each of the patents. Third, we hand collected detailed information about the prosecution histories of each of the patents. In general, the idea was to estimate the following equation:

$$\Pr(v_i = 1) = \text{logit}^{-1}(\alpha + \beta X_i + \gamma Y_i + \delta Z_i + \varepsilon),$$

where v_i is the validity of the patent i , X is a vector of variables from the prosecution history, Y is the alignment of the claims and the specification in the issued patent, and Z is a vector of variables from the face of the patent, standardized by matching where practicable. The overriding hypothesis is that the prosecution history and alignment data will relate to validity in important ways.

A. Federal Circuit Decisions

For validity, the most direct measure is the results of litigation,³ and in particular a decision of the Federal Circuit, available only for the small subset of patents that are the subject of litigation before that court. Except in the rare case of U.S. Supreme Court review, Federal Circuit decisions on validity definitively establish that claims of the patent in question are (or are not) invalid. Even there, the Supreme Court has never reversed a Federal Circuit decision holding a patent invalid. To begin, then, we collected Federal Circuit decisions

³It is not necessarily the only measure. For example, Lei and Wright (2009) use success of parallel applications in the European Union as a proxy for the quality of applications filed with the PTO, and consider whether the share of relevant prior art located before the patent is issued correlates positively with success in the prosecution of E.U. applications. Although the research design is intriguing, marked differences between validity considerations in the European Union and the United States make us unsure about how to generalize from those findings to the U.S. population.

that definitively resolve a claim about the validity of an issued patent.⁴ Because decisions that come from district courts and decisions that come directly from the PTO (when the PTO refuses to issue a patent) raise quite different considerations, and because there are far more cases from district courts than from the PTO, this portion of the project uses only cases that come from the district courts. Working from the *United States Patent Quarterly*, supplemented by LEXIS searches for unreported decisions, we include all cases decided from January 1, 2003 through December 31, 2009, for a total data set of 266 cases adjudicating the validity of 366 patents.

The balance between validity and invalidity was relatively even: the court held 147 of the 366 patents (40.2 percent) valid.⁵ We collected several variables about the Federal Circuit's disposition of the case. For example, to judge the clarity of the patent's validity or invalidity, we collected information about whether there was a dissent, whether the decision was published or unpublished, and whether the Federal Circuit affirmed the decision of the trial court. We also categorized all the decisions into three broad categories based on the legal issues that the court addressed: prior art (usually anticipation or obviousness), drafting (usually definiteness or enablement, but occasionally double patenting or recapture), and prior use (usually a statutory bar). The data set includes no cases on patentable subject matter, which tend to reach the Federal Circuit by appeal from the Board's refusal to issue a patent. Several of those variables appear to vary with validity (Figure 2). Reversing decisions, unanimous decisions, and decisions involving drafting issues are more likely to find patents valid; affirming decisions, dissenting decisions, and decisions involving prior art questions are more likely to find patents invalid.


B. Facial Characteristics of Patents (*Delphion*)

We next collected about 15 variables that are susceptible of automated collection, largely because they are apparent on the face of the patent and thus available from *Delphion*. These are the variables related to the number of claims, independent claims, references in the patent, and the like that have dominated the existing literature. As discussed above, it was important to take account of the apparent variations in those features over time and technology type. For each of the 366 nodal patents, we selected a series of matched patents, permitting us to quantify the difference on the relevant variable between the nodal patent and the typical values of the matching patents. The matched patents were designed to provide a robust control for the time of issuance and type of technology.

⁴Thus, we exclude cases in which the Federal Circuit vacates a validity determination of a lower court and remands for further consideration. At the same time, we include summary decisions on validity, reasoning that cases in which validity is so clear that it does not require extended explanation are useful data points for a data set seeking to identify objective differences between patents that are and are not valid. This distinguishes our work from Allison and Lemley (1998), who also included district court determinations and nonfinal Federal Circuit decisions. The high rate of reversal in cases appealed to the Federal Circuit suggests it would not be useful to rely on district court determinations that were not appealed as conclusive evidence of validity or invalidity (Burk & Lemley 2009:28).

⁵Although several of the patents were repeatedly litigated, none had inconsistent outcomes. The 40 percent validity rate is of course a substantial drop off from the early days of the Federal Circuit, when about 55 percent of patents were found valid (Dunner 1985; Lemley 1994).

Figure 2: Decision attributes and validity (2003–2009).



Specifically, we identified for each of the nodal patents all the patents with an issue date within six months of the issue date of the nodal patent that included in their classifications both the same main International Patent Classification (IPC) class and the same main U.S. class. We then randomly selected 30 of the matching patents (except in one case where there were only 13 matched patents, all of which we used). To standardize the relation between the nodal patent and the matched patents, we calculated on each of the variables the mean and standard deviation for the matched patents. We used a standardized value equal to the difference between the feature on the nodal patent and the mean of that feature on the matched patent, divided by the standard deviation of that feature on the matched patents.

Table 1 displays the means for representative standardized variables, by validity group. For the most part the differences between litigated patents and their matches point in the directions we expected. Thus, all the variables for claims and references are larger for both classes of litigated patents (valid and invalid) than their matches. This replicates the findings in the existing literature that associates those features with patents likely to be selected for litigation.

On the other hand, a different pattern emerges when we compare the valid patents to the invalid patents. Here, the relation between validity and the claims variables is, if anything, inverse: the difference from the matches is less for the valid patents than it is for the invalid patents. This could suggest either that those patents are more exposed to invalidation (because they have more claims) or that they are more likely to overclaim. At

Table 1: Standardized Facial Patent Characteristics (2003–2009)

	<i>Valid Patents</i>	<i>Invalid Patents</i>	<i>Difference</i>
Claims	0.63	0.80	(0.17)
Independent claims	0.58	0.83	(0.25)
Words (abstract)	0.07	0.16	(0.09)
References	1.44	0.87	0.57*
U.S. references	0.99	0.68	0.31
Foreign references	0.59	0.13	0.46**
NPPA	1.87	1.60	0.27
Classes	0.14	-0.05	0.19
Priority countries	0.76	1.43	(0.67)***
U.S. priority patents	0.78	1.36	(0.58)***
U.S. patents in family	1.20	3.81	(2.61)
Foreign patents in family	0.73	0.96	(0.23)

NOTES: Table reports mean characteristics for Federal Circuit patents in terms of standard deviations from the mean of a sample of nonlitigated patents matched on time of issuance and technology class. Differences are calculated as mean of valid–mean of invalid. $N=363$. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$ (based on two-group test).

SOURCE: Author calculations based on Delphion and hand-collected data set.

least in this bivariate analysis, the latter hypothesis draws support from the strong relation to validity of the length of the abstract: where the abstracts of the valid patents were shorter than the abstracts of their matches, the abstracts of the invalid patents were substantially longer than the abstracts of their matches. Conversely, all the reference variables (total references, U.S. references, foreign references, and NPPA) relate positively to validity: the difference from the matches is greater for the valid patents than it is for the invalid patents. This resonates with the idea that patents issued after a more thorough canvassing of prior art are more likely to be valid.⁶ The number of classes into which a patent is classified has a slight positive relation to validity, perhaps because the breadth of classes reflects a more careful search. Finally, the variables measuring the number of domestic “family member” patents have strong positive relations to litigation (they are positive for both valid and invalid patents) and strong inverse relations to validity. This well might reflect more aggressive patenting strategies for more important and controversial inventions. The multivariate analysis discussed below considers more carefully the extent to which any of those differences have an important relation to validity.

C. Text-Based Features

Because the validity of patents ultimately relates to the actual text of the patents, we also analyzed text-based features of the nodal patents. These allow us to quantify the differences

⁶We note the contrary hypothesis in Lei and Wright (2009), that examiners search prior art more carefully in the case of patents they believe to be weak, but find that hypothesis unsupported in our data set. Our working supposition is that their findings are related to systematic differences between the concepts of validity that govern prosecution in the European Union and before the PTO.

in these features between the nodal patents and the set of matched patents discussed above. Patent experts point regularly to the degree of alignment of the claims and the patent body as an important quality component. Accordingly, we analyze the text of the patent to capture the degree of alignment of the text of the various parts of the patent.⁷ Our approach follows the basic principles in the rapidly developing area of content analysis. First, we remove words too common for useful analysis (“the,” “and,” “for,” etc.). Second, we “stem” remaining words, a process to reduce inflected (or sometimes derived) words to their stem, base, or root form (Porter 1980). Once we have a reduced universe of words W , we represent the claims and the specification of the patent p as an indicator vector that is 1 if a given word in the universe is present; 0 otherwise. We then calculate the measure of alignment between the claims and the specification as the Euclidean norm on the difference of the two indicator vectors, as indicated in this equation:

$$\text{Alignment}(p) = PCV(p) - DV(p)P = \left\{ \sum_{i \in W} (cv(p)_i - dv(p)_i)^2 \right\}^{1/2}.$$

One of the most challenging problems was accommodating the reality that the specification is almost invariably much longer than the claims, and that a difference in relative length is not an indicator of poor alignment. To respond to that problem, we rely on an indicator rather than a count representation of alignment. Thus, our textual analysis considers only whether the specification mentions the salient terms in the claims, not necessarily if they occur with similar frequency.⁸

The alignment variable measures the distance between the content of the specification and the content of the claims. Thus, it increases if the claims and description are poorly aligned. Accordingly, we would expect this measure to correlate inversely with validity. This expectation is borne out by a bivariate comparison: the mean vector for valid patents is 22.7 (with a standard deviation of 5.1); the mean vector for invalid patents is 24.2 (with a standard deviation of 6.6). We discuss below the robustness of that finding in multivariate analysis.

D. Prosecution History Data

Because a major goal of our project is to identify relations between the application and prosecution process, on the one hand, and the ultimate question of validity, on the other hand, we hand collected 30 additional variables from the prosecution histories of the applications and from the subsequently issued patents, generally falling into six distinct

⁷Our data collection resembles Lichtman (2004) in its reliance on an automated textual assessment of the patent documents. Where his goal, though, was to consider the quantity of changes during the prosecution process, our goal is to develop an automated tool that can assess the textual features of the document that relate closely to its validity.

⁸Given the novelty of this measure, it is difficult to be sure that we have selected the best method of analyzing alignment. In the end, we constructed a number of different methods (using alternative Euclidean and cosine functions of alignment). The different measures were highly correlated and we report the Euclidean measure here largely because it is the simplest.

categories.⁹ First, we collected the number of classes and subclasses searched during the prosecution process, reasoning that those items would provide information about the rigor of the prosecution process.

Second, we collected information about the number of claims and independent claims in the application. This information should shed light on two separate questions. By allowing us to quantify the extent to which the prosecution process reduces the number of claims in the patent, these data should provide a measure of the rigor of the prosecution process. Similarly, by allowing us to examine data about independent claims in addition to total claims, we have a second and arguably more precise method for measuring the aggressiveness of drafting reflected in the application.

Third, we collected information from the Information Disclosure Statements (IDS) often filed with patent applications (typically on a PTO Form 1449). The intuition here is that applications in which the inventor did file an IDS and included a greater share of the prior art ultimately located during the prosecution are of higher quality than those applications that did not include an IDS or that disclosed a lower share of the prior art ultimately located during the prosecution.

Fourth, we calculated two separate measures of the time in examination—the time from application to issuance, and the time from the first office action to issuance. Because of differential backlogs in art units in the PTO, we hypothesized that a measure of examination time running from the first substantive action by the examiner was an alternative and potentially more accurate measure of the rigor of the prosecution process.

Fifth, we collected information about the steps of interaction by the examiner and the applicant. On the part of the examiner, we counted the number of times that the examiner rejected the patent before ultimately granting the application, as well as a number of data points about the basis for the rejections. On the part of the applicant, we counted the number of separate times that the applicant amended the claims in a substantial way (with separate calculations for amendments of dependent and independent claims).

Finally, we collected several miscellaneous items that are likely to appear rarely but that might correlate with quality *ex post*. The most important of these relate to counterpart applications in other countries and appellate processes inside the PTO. On the first point, counterpart applications filed in the European Union or Japan should indicate that the applicant had a high *ex ante* estimation of the value of the patent and thus should correlate with applicant effort in preparing the application. Moreover, the parallel search and examination effort should be reflected in a higher level of applicant effort on the nodal

⁹Two separate coders examined each prosecution history, and one of the principal investigators resolved all inconsistencies in coding by personal examination of the relevant files. Tests of intercoder reliability find an accuracy rate of 94.6 percent. Because other researchers (e.g., Alcácer et al. 2009) have used data from the PTO's Patent Application Information Retrieval System (PAIR) as a source for information about prosecution histories, we separately collected PAIR data on the features that are available from PAIR. It appears that the PAIR data substantially undercount the items we collected. So, for example, the PAIR data find 219 IDS filings in the 333 patents from our data set for which PAIR data are available, while our hand-collected data located 275 IDS filings for the 362 patents for which we obtained prosecution histories. Because of the undercounting, we do not use any PAIR data in the models we present in this article.

Table 2: Prosecution History Variables (2003–2009)

	<i>Valid Patents</i>	<i>Invalid Patents</i>
Mean claims in application	24	26
Mean independent claims in application	4.0	4.6
IDS (%)	79	75
IDS filings (#)**	1.9	1.4
IDS references	40	28
Rejections (%)	90	85
Rejections based on unlisted references***	0.54	0.29
Examiner-added U.S. references*	4.0	5.6
Days of prosecution time (Issue date–First-action date)	634	628
PTO decisions (%)*	0.7	3.7
E.U. priority in family (%)	10.2	13.7
Mean continuations/RCEs/CPAs***	0.6	1.0

NOTES: Table reports means or percentages, as indicated, for Federal Circuit patents, for data hand collected from prosecution histories. $N = 360$. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$ (based on two-group tests of mean or proportion, as appropriate).

patent application in the data set. On the second, information about patents issued after a notice of appeal by the applicant has ambiguous connotations. On the one hand, it indicates a patent issued after an unusually careful process, which might signal ex post quality. On the other hand, it also indicates a patent as to which the examiner came to a definitive conclusion of invalidity, which might relate to a lack of quality ex post. Table 2 provides summary statistics by validity group for selected prosecution history variables.

IV. ANALYSIS

A. Analytical Strategy

The data collection rests on a group of patents for which the Federal Circuit has made final rulings of validity or invalidity. Thus, the research design assumes that the difference between those groups is dichotomous, so that it is appropriate to compare the characteristics of the two groups of patents with multivariate logistic regression. That assumption confronts several challenges. For one thing, the Federal Circuit does not often rule on the validity of an entire patent. Most cases (and all the cases in the data set) involve challenges to the validity of some particular claim or subset of claims. Thus, when the Federal Circuit upholds a claim of invalidity, it leaves intact the claims in the patent that were not challenged. Conversely, when it rejects a claim of invalidity, it says nothing at all about the validity of claims not before it. In each case, we code the case based on what the Federal Circuit decides: a case invalidating even a single claim is treated as a case of invalidity, while a case that invalidates no claims is treated as a case of validity.

Similarly, the legal significance of decisions finding claims valid and invalid is not precisely obverse. A decision invalidating a claim is much more complete than a decision validating a claim: once the Federal Circuit invalidates a claim, the matter is for all practical

purposes finally resolved. A decision that the patent is valid, on the other hand, technically means nothing more than that the particular legal question presented to the Federal Circuit was not an adequate basis for rejecting the patent.¹⁰

Although those problems undermine the dichotomous nature of validity in a technical way, in practice we find them unimportant. Given the costs and delay involved in litigation to the Federal Circuit, it is likely that litigants challenging a patent before the Federal Circuit will challenge claims that have economic significance. It is thus not unreasonable to suppose that a decision invalidating (or validating) those claims says something important about the failure (or success) of the prosecution process in adequately demarcating the invention the patent discloses. Similarly, litigants that press validity challenges to a final decision in the Federal Circuit are likely to raise the strongest challenges to validity in the first instance. The costs of litigation are quite high, and litigants have no incentive to retain their best challenges for later appeals; indeed, doctrines related to preclusion and the law of the case well might limit their ability to present subsequent challenges to the same claims.

To be specific, the occurrence that would undermine our analytical frame would be that a patent that endured a challenge in the Federal Circuit was invalid in an important way (either with respect to the claims at issue in the first case or with respect to some other set of claims). However, there is no case in the data set in which the Federal Circuit rejected the first challenge to validity and accepted a second challenge; for each of the seven patents that returned to the Federal Circuit after an initial holding of validity, the Federal Circuit rejected subsequent validity challenges. The reverse is equally true; for the only patent to return to the Federal Circuit after an initial finding of invalidity (the claimant raising a broader invalidity claim), the Federal Circuit accepted the second invalidity challenge as well. Similarly, buttressing our hypothesis that litigation in the Federal Circuit is likely to focus on the core claims of the patent, we note considerable overlap in the claims involved in repeat litigation: for four of the eight repeat litigation patents, the same claims were at issue in each Federal Circuit case. For three of the others there was substantial overlap, so that one common claim was involved in each case. For only one patent did the repeated appearances involve different claims. In sum, recognizing the technical concern, we believe as a practical matter that the difference between patents that stand and fall in the Federal Circuit is sufficiently important to justify analyzing that determination in a dichotomous way. At worst, it suggests a measurement error that would make it more difficult for the statistical analysis to discern patterns that exist in the data.

Because the collection of data from the prosecution histories required individualized analysis of those files, it was not feasible to provide the robust matching for those data points that the standardization process provided for the variables that could be collected through automated processes. Thus, the data set includes both a large group of variables for which we have no matching information and also a set of variables that use matching; comparison of valid and invalid patents on the matched variables provides a difference-in-difference look at the relation between those variables and validity.

¹⁰Our concern on this point motivated our decision to exclude nonfinal decisions from our data set.

To take advantage of all the information in the data set, it is necessary to estimate models that include both the standardized variables and the unstandardized variables. This raises a number of concerns, the most obvious of which is the possibility that some combination of technology class and periodicity might explain relations in the unstandardized variables. In particular, controls for those concepts might be inadequate if they require extrapolation beyond the range of the data.¹¹

To assess how problematic it is to use both the standardized and unstandardized variables in the data set, we estimated the logistic regression models described below separately—once using the standardized variables where available and once using the unstandardized forms of the same variables. Although the coefficients shifted somewhat, for the most part the results were quite similar. Accordingly, we conclude that the technology and time controls function adequately. Nevertheless, as a conservative analytical approach we use standardized variables for all the concepts for which they are available.

Because our project relates directly to the prosecution process, we differentiate patents based on the technology centers in which the PTO examines them.¹² Recognizing that the implicit premise of the statistical analysis below is that patent validity is a static and time-invariant characteristic of patents, it is problematic to recognize that the Federal Circuit's perspective of what types of patents are valid and invalid doubtless has shifted during the period between the issuance of the patents in the data set (the oldest of which was issued in 1975) and the points in time at which the Federal Circuit assessed their validity. As we discuss in more detail below, the apparent variation related to time and technology was an important factor in the design of the data collection and analysis.

Because the data set contains a large number of variables (188) in comparison to the number of observations (366), and because many of the variables are closely related to each other, it was not useful for our regressions to include all the variables in the data set. Accordingly, we started by organizing the variables into a number of distinct concepts that should relate to validity. Using multiple variables from any one of those categories would pose obvious problems of collinearity: the variables for total claims, independent claims, and dependent claims, for example, correlate closely among themselves.¹³ Accordingly, we endeavored to identify a single representative variable or small number of variables for each

¹¹To enhance our ability to understand how the relevant variables relate to different technology centers, we also estimated multilevel models that allow us to observe how the effects of the explanatory variables differ by technology center. Because those models did not reveal significant variation by technology center, we do not report them here. They do, however, buttress our perspective that the technology controls are adequate. The data set also includes PTO data about features of the process that differ by technology center (number of examiners, backlog, etc.). None of those data appear to relate significantly to validity.

¹²Prior scholarship has used ad hoc classifications loosely based on the top-level IPC codes (Jaffe & Trajtenberg 2002; Lanjouw & Schankerman 2004). We attempted to replicate the classification from Lanjouw and Schankerman (2004), but neither the authors nor the journal have retained the data necessary to do so.

¹³We also used principal component analysis to extract factors from the multiple variables for each concept, but generally the use of factors did not improve the results relative to using representative variables for each concept, and the difficulty of interpreting the coefficients on those factors ultimately persuaded us to report models that use representative variables.

separate concept, based largely on a theoretical understanding of the likely relations of the variables to validity.

Finally, as controls for robustness, we repeated our analysis on distinct subsets of the data set, using the features of the Federal Circuit decision to segment the data set. Thus, where the results below are based on all the observations in the data set, we repeated that analysis omitting all the patents in which there was a dissent; because judges on the Federal Circuit disagreed about the validity of those patents, we reasonably might regard them as less clearly marked valid or invalid. Similarly, we repeated the analysis for the subset of cases in which the decision was not published; unpublished decisions might be regarded as involving patents in which the validity issues were clearer. The results on those segments were not substantively different from the results we report.¹⁴ Although we thought separate analysis of those segments of the data set was useful as a robustness check, we did not include any of the attributes of the Federal Circuit decisions in our statistical analysis directly. Because the premise of our statistical analysis is that validity is a function of the prosecution process (and thus that validity is fixed at the moment of issuance), attributes of the Federal Circuit decision cannot properly be used to explain validity; logically, postissuance variables are irrelevant to analysis of the decision to issue.¹⁵

A related issue is the circumstance that many of the Federal Circuit decisions in the data set addressed the validity of multiple patents within the same opinion. Because those opinions showed a slight tendency to conclude that all (or almost all) patents in the group were valid or invalid, we were concerned that those observations were not independent. To respond to that concern, the results we summarize below include two sets of standard errors. The first are conventional standard errors clustered on the separate judicial decisions. The second result from a simulation that used a single randomly selected patent from each of the Federal Circuit decisions for 10,000 repetitions; we report bootstrapped standard errors from the results.

A final segmentation of the data took account of subject matter coding of the issues addressed by the Federal Circuit. Specifically, we repeated our analysis separately for the prior art and drafting cases, reasoning that the covariates should relate differently to those two categories of cases. None of those robustness checks suggested instability in the results that we present below.

¹⁴Results are available on request. We considered the possibility of identifying a broader set of patents for which validity might be even clearer, such as patents pressed in litigation for which the defendants did not challenge validity. In the end, however, considering the strategic reasons that might motivate defendants to forego a challenge to validity, we decided that the absence of a judicial decision on validity made those patents a poor choice for reliable statistical analysis of validity.

¹⁵That is not to say that events after issuance in fact have no relation to the Federal Circuit's decision. Of course, many postissuance decisions have an important effect on the Federal Circuit's decision—the quality of the parties' attorneys, strategy related to claims construction, and the identity of the district court judge and the judges on the Federal Circuit panel, but for purposes of statistical analysis of the PTO's decision to issue, those events are better regarded as “noise” that might (or might not) obscure the clarity of the relations among the applicant's conduct, the prosecution process, and the quality of the issued patent.

B. Analyzing Patent Validity

Because the purpose of our analysis is to understand the extent to which the activities of the applicant and the examiner interact to foster (or undermine) the validity of the patents that issue, we present three distinct models, reflecting three different points in the process. The first set is limited to information about the initial submission of the applicant to the PTO: the characteristics of the application itself and any IDS filings that the applicant submits in support. The second set includes both information about the applicant's input to the process and the activities of the applicant, examiner, and Board during the process itself. The third set adds information apparent only from the issued patent (such as the standardized variables that compare the issued patent to contemporaneous patents in similar technology areas).

Before we describe those results, it is useful to note the empirical challenge we face. There is much that might influence a Federal Circuit decision on the validity of a patent that is not determined at the moment that the PTO issues the patent. The value of the patent may influence the vigor with which the patentee will defend it and with which an infringer will attack it. The perspective of the trial judge and jury will shape factual findings and the record that will be available to the appellate court. Expert witnesses may influence the juries' understanding of the invention and the state of the art. The quality of the attorneys in the trial and appellate courts will significantly affect how the courts perceive the relevant legal questions. All would agree that the identities of the judges on the Federal Circuit panel chosen to resolve the appeal will have some influence on how the case is decided.¹⁶ To put it more directly, even an observer deeply committed to the integrity and objectivity of the legal process would be surprised if quantitative information about a patent could provide enough information to accurately predict the likelihood that the Federal Circuit would find any particular patent valid or invalid. Recognizing those difficulties, we think that the statistical analysis we present below was surprisingly successful. For example, the version that uses data from all three stages correctly classifies 74 percent of the patents and has a generalized R^2 of 22 percent.¹⁷

Table 3 summarizes the logistic regression analysis. At a global level, the most interesting thing is that even with the inclusion of the standardized variables, many of the most important data points are hand-collected features from the prosecution history. This suggests that a reliable understanding of patent validity cannot rest on information from the face of patents alone, but instead depends on features of the process that produced the patents in question. On the other hand, even data available from the applicant alone provide considerable information, correctly classifying 68 percent of the patents. Turning from the analysis as a whole to the individual concepts, it is useful to discuss in turn what the analysis suggests about the relation between each of those concepts and the ultimate question of patent validity.

¹⁶For example, Judge Gajarsa ruled in favor of validity for 73 percent of the patents he considered, while Judge Moore ruled in favor of validity only 21 percent of the time. Similarly, reversal rates ranged from 35 percent (Judge Dyk) to 11 percent (Judge Mayer).

¹⁷We use generalized R^2 here as a rough measure of the explanatory power of the models.

Table 3: Logistic Regression Models of Patent Validity (2003–2009)

<i>Variables</i>	<i>Applicant Input Model</i>	<i>Prosecution Model</i>	<i>Final Model</i>
Technological Breadth			
Standardized classes	—	—	0.35 (0.12) [0.06]***
Scope of Protection			
Standardized claims	—	—	-0.07 (0.07) [0.04]
Independent claims (application)	-0.071 (0.044) [0.01]	-0.061 (0.038) [0.01]	—
Patents in EU/JP	—	—	-0.40** (0.19) [0.06]
Applicant Search			
# of IDS filings	0.20*** (0.07) [0.02]	0.17** (0.069) [0.03]	0.22*** (0.08) [0.03]
IDS references	0.004* (0.002) [0.001]	0.005** (0.002) [0.00]	—
Office Search Effort			
# of examiner-added U.S. references	—	-0.068** (0.031) [0.01]	-0.10*** (0.03) [0.01]
Office Actions			
# of rejections based on unlisted references	—	0.37** (0.15) [0.07]	0.37** (0.18) [0.10]
Postexaminer Actions			
Continuations	—	-0.24* (0.14) [0.08]	-0.22* (0.15) [0.09]
Board decision (y/n)	—	-2.0* (1.2) [0.22]	-2.12** (0.99) [0.23]
Rule 312 amendment (y/n)	—	-0.67* (0.37) [0.19]	-0.90** (0.38) [0.21]
Total Research			
Standardized U.S. references	—	—	0.12 (0.09) [0.05]
Standardized foreign references	—	—	0.17** (0.09) [0.03]
Standardized nonpatent prior art	—	—	-0.02 (0.02) [0.02]
Drafting Quality			
Vector (specification to claims)	—	—	-0.04* (0.02) [0.01]
Controls			
Application date (years)	-0.14*** (0.03) [0.01]	-0.14*** (0.03) [0.01]	-0.15*** (0.03) [0.01]
{Technology controls omitted}			
Constant	4.23*** (0.89) [0.23]	4.70*** (0.97) [0.25]	6.90*** (1.38) [0.41]
Observations (clusters)	359 (254)	359 (254)	356 (252)
Log pseudo-likelihood	-214	-201	-188
Generalized R^2	12%	17%	22%
Correctly classified	67%	70%	74%

NOTES: Table displays logistic coefficients, with robust standard errors in parentheses (clustered on decisions) and bootstrapped standard errors in brackets; *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$ (based on clustered errors).

1. Technological Breadth

Several variables in the data set relate to the technological breadth of the invention. Among other things, we have data about the number of inventors and several different metrics of the number of classes into which the patent is classified (both under the U.S. system and under the IPC system). Although existing literature suggests that the number of inventors relates to the diffusion of knowledge (Wuchty et al. 2007), it appears not to relate in any significant way to validity. Given the ease with which additional individuals can be listed as inventors on a patent, this was not surprising to us. We do not include any inventor variables in the analysis that we report here.

At the same time, the number of classes into which the patent is classified has a significant and positive relation to validity. The coefficient suggests (at the reference values of the other variables) that each additional class into which the patent falls (the mean is four classes) increases the likelihood of validity by about 9 percentage points. The results on that variable initially surprised us. Our initial hypothesis had been that an invention spanning multiple classes would be a more ambitious invention and thus more susceptible of invalidation because of the multiplicity of technologies from which relevant art might be found. On reflection, however, our initial hypothesis seems to reflect a failure to distinguish between the ambition of what the inventor claims and the ability of the prosecution process to identify the significance of the disclosed invention. If we hold other important features of the patent and the process constant, the number of classes into which the invention is classified might better be understood as reflecting the success of the applicant and examiner in understanding the entire range of technology over which the invention operates. This perspective treats the variable as a proxy for quality of the joint search. Another explanation would be that these highly cross-classified patents often are issued in arts in which the classifications are poorly organized. Poorly organized classifications, in turn, are likely to reflect rapid growth in the technology. From this perspective, cross-classification of patents might be a proxy for innovation in a cutting-edge area.

2. Scope of Protection

Given the focus of the existing literature on claims (largely because of the relative ease of obtaining that information), we explored several different variables that might capture the scope of protection. We explored variables for total claims and independent claims, in both the patent and the application, for the length of the patent (in pages) and of the abstract (in words). For most of those variables we examined both standardized and unstandardized forms. We also explored quadratic forms of several of these variables to capture the possibility that these variables might be important only at the extremes. The only claims-related variable that differs substantially from zero is independent claims in the application, which is marginally negative and appears at the applicant and prosecution stages, with a coefficient that suggests each additional independent claim in the application reduces the likelihood of final validity by 1–2 percentage points.

As discussed above, the existing literature has emphasized a positive connection between the number of claims and value. The data suggest, however, that the existing literature reflects little more than the reality that patents with more claims are more likely

to be selected into litigation than patents with fewer claims. It is easy to suggest hypotheses that would relate the number of claims or complexity of the patent to validity. For example, a patent with more claims necessarily has more places in which mistakes could have been made. Conversely, a patent with a more complex claim structure is more likely to include a claim just broad enough to successfully attack a competitor, decreasing the likelihood that the alleged infringer can repel an infringement claim by showing that the asserted claim is overly broad. Perhaps there is some truth to both stories. In this data set, at least, none of those variables relates strongly to validity.

Given the weakness of the relation between independent claims and validity, the final model omits that variable in preference of a variable that counts counterpart patents in the European Union and Japan (a variable available only for the final model). We collected those data based on the notion that patents in those offices might reflect the most economically significant extensions of the invention. As suggested above, this variable and the related variables appear to relate positively to litigation and negatively to validity. In the final model, the coefficient suggests that each major additional patent reduces the likelihood of validity by about 10 percentage points.

We originally collected the data on patent families because of an expectation that patents subjected to the prosecution processes of the European and Japanese patent offices would display a higher quality, both because of the benefits of a parallel examination and search and because of the reputedly high capabilities of those offices.

The family variables, however, suggest an inverse relation to validity. Our best explanation is that it reflects different national standards for patent drafting. Attorneys experienced in multinational patent prosecution have explained to us that strategies for patent drafting in the European Union and Japan are quite different from those in the United States. If companies first draft their patents for prosecution in an E.U. or Japanese office, they well might produce applications for the U.S. patent office that are less well designed for the U.S. patenting process, and thus less likely to result in valid patents within that system. Another possibility suggested to us by early readers is that companies seek multinational protection only for their most important patents, and that in those cases they (or their counsel) might be more interested in aggressively broad patenting in each country, losing sight of their interest in restricting their claims to a line of clear validity.¹⁸

3. Applicant Search

As discussed above, several variables related to the applicant's search effort have significantly positive bivariate correlations with validity. We include the general number of IDS filings and the number of references in those filings (both reflecting the applicant's contribution to the patent search effort). We omit the references variable in the final model because it seems to be collinear with other variables included in that model that are not

¹⁸Several readers of early drafts have suggested that validity problems can be ascribed to conflicting motivations of the patenting company (with a more conservative motivation in validity) and their outside counsel (with a slightly different aggressive interest in displaying zeal). Although the idea is intriguing, the data structure available to us does not make it practical to explore those issues in this article.

available for the earlier-stage models. All those variables have positive coefficients significantly different from zero in all three models. The coefficients suggest that each additional IDS filing (the mean is 1.6 filings) relates to an increased likelihood of validity of about 5 percentage points, and that each 10 additional IDS references increases the likelihood of validity by about one-half of a percentage point.

4. Office Search Effort

As discussed above, we collected data about changes from the references in the IDS to the references in the final patent.¹⁹ At the second and third stages we include the number of U.S. references added by the examiner. That variable is inversely related to validity, with a coefficient that suggests that each reference added (the mean is seven examiner-added U.S. references) is associated with a two-point decrease in the likelihood of validity. Generally, patents issued after the applicant does a poor job of identifying the relevant prior art (either because it fails to submit relevant art or because it submits irrelevant art) are of lower quality. Even though the examiner might expend considerable effort eliminating the references provided by the applicant and adding new references based on the examiner's own search, the data suggest that the weakness of the application persists. This buttresses the idea, discussed above, that the robustness of the prior art search of the applicant has independent value for the quality of the application, even apart from the likelihood that the application would be subject to challenge as anticipated or obvious. It is possible, to be sure, that the substantial change in references between IDS and patent signal only that the examiner chose to work harder on the application because of an apparent defect. The point of importance for us, however, is that whatever signaled to the examiner the need for unusual effort, it appears to correlate with a poor identification of prior art by the applicant.

5. Office Actions

Building on that concern, we include at the second and third stages a variable for office actions based on unlisted references—not submitted by the applicant but found by the examiner (a cross-cutting descriptive category determined by individualized examination of the office action document). The coefficient is positive and significantly different from zero, suggesting that each such rejection increases the likelihood of validity by 9–10 percentage points. Juxtaposed with the data summarized above (about the adverse implications of a process in which an examiner adds a large number of references), this finding suggests that it is not enough for the examiner to locate relevant prior art the applicant fails to cite; what seems to be valuable in the process is that the examiner locate invalidating prior art that, by hypothesis, forces the applicant to reshape the claims. The feature of the process that relates to validity is the office action rejecting the application.

¹⁹In recent years, patents indicate on their face the references added by the examiner. We are not aware of any data source that previously has examined patents cited by the applicant and omitted from the final patent by the examiner.

6. Postexaminer Actions

The final two stages include three distinct variables to account for activity by which applicants can alter the results of the decisions of the examiners: seeking continuations, seeking review from the Board, and seeking amendments of the patent under Rule 312. All three of those variables have large and significant coefficients²⁰ that correlate negatively with validity. The estimates suggests that each continuation is associated with a six-point decrease in the likelihood of validity, that the issuance of patents only after Board decisions is associated with a 50-point decline in the likelihood of validity, and that Rule 312 amendments are associated with a 22-point decline in the likelihood of validity.

7. Total Research

The prior literature has emphasized a strong relation between the number of references in the final patent and the patent's value. Because, holding other things constant, the number of references is more likely to reflect the quality of the search than the value of the underlying invention, it would be natural for this variable (unlike the claims variable) also to relate positively to validity. At the same time, it is reasonable to expect that the variables gleaned from the prosecution history would measure that concept more accurately. Thus, it was not surprising that variables measuring total references, U.S. references, and non-patent prior art have little force in the multivariate analysis. Accordingly, we include only the variable for foreign references, which is positive and significantly different from zero. The coefficient suggests that each additional foreign reference (the mean is three) relates to an increase in the likelihood of validity of almost 2 percentage points. The intuition here is that the greatest variation in search quality relates not to U.S. patents (as to which routine examiner skill should be adequate) but to foreign patents, where it may be more difficult to locate the relevant art.

8. Drafting Quality

We include in the final model the vector that compares the text of the specification to the text of the claims. After the decision in *Phillips v. AWH Corporation*, 415 F.3d 1303 (Fed. Cir. 2005) (en banc) accentuated the importance of using the specification to interpret the claims, conformity between the claims and the specification will be especially important in assessing challenges to the enablement and definiteness of the patent. Also, given that validity against prior art is determined by comparing the prior art to what is claimed, a precise determination of claim meaning is critical to clear validity determinations. Thus, it is not surprising that the distance between the concepts articulated in the specification and the concepts articulated in the claims is negatively correlated to validity (albeit only at the 0.10 level).

²⁰As the table indicates, some of those coefficients are significant only at the 0.10 level.

9. Time

The time control is significant, underscoring the importance of accounting for the age of the patent. We have experimented with different measures of time (application date, first-action date, and issue date),²¹ but have settled on application date because that seems to us the date that fits best with a theory that emphasizes the process from application to issuance. Scatterplots suggest that the effect of time is roughly linear. The coefficients suggest that the likelihood of validity declines by about 3.5 percentage points for each later calendar year of the application date (1998 applications would result in patents almost 4 percentage points less likely to be valid than 1997 applications).

V. IMPLICATIONS

Building specific policy recommendations on the analysis above would be a delicate matter. Because the data set uses a population of validity decisions from the Federal Circuit, it is reasonable to treat the results as a reliable description of the characteristics associated with favorable and adverse decisions by that court. This is, of course, an important question in its own right. However, if the ultimate goal is to understand the quality of the work done by the PTO, we must consider whether the Federal Circuit data are sufficiently representative to make it fair to generalize from data about those patents to conclusions about the much larger population of issued patents.

Although the population of issued patents is indeed orders of magnitude larger than the group of Federal Circuit patents we examine here, basic principles of statistical inference suggest that it is not unreasonable to make findings about the larger population using the approach we present here. This is not because there are no systematic differences between the patents that do and do not reach the Federal Circuit. Of course there are, as the large coefficients in Table 1 demonstrate. Nor is it because none of the variables that affect selection into the Federal Circuit relate to validity because validity surely is relevant to that process.

That is not the relevant question. The relevant question is whether one of the variables in our statistical analysis has a different relation to validity for the patents that do and do not reach the Federal Circuit. In causal terminology, we make no strong claims about the internal validity of the inferences that might be drawn from the data. We have no experimental data, and it would be quite difficult to construct an experiment that would permit identification of treatment and control groups (effectively examining the variation in a single patent characteristic while randomizing everything else). A separate concern about external validity (generalizability) arises if there is an interaction between selection into Federal Circuit review and the effects of the independent variables.

Given the infrequent nature of Federal Circuit review, and its dependency on factors wholly unrelated to the patent prosecution process (such as the value of the market on which

²¹These measures of the time when the patent was examined are distinct from measures of the length of time for which a patent was examined. As discussed above, we do not find any significant relations between the duration of examination and validity.

the patent writes and the aggressiveness of the firm that holds it), it is not easy to see obvious reasons why any of the variables should interact with selection into the Federal Circuit. Accordingly, although we recognize that caution is appropriate, we do believe that the findings are at least suggestive with respect to the larger universe of patents issued by the PTO.

Turning to that point, we view the data analysis as supporting a three-fold shift in the conception of patent reform. First, if the goal of patent reform is to increase the likely validity of the patents the PTO issues, then the focus of reform should be on the prosecution process, not on the procedures of postissuance litigation that dominate existing legislative proposals. Second, given the complexity of the PTO process (only glimpsed in the data discussed above) and the volume of applications the PTO examines, the only reliable way to assess the strengths and weaknesses of the process is through empirical observation of that process and its results. Third, given the distinct capabilities that applicants and the PTO possess, the optimal system for maximizing the likely validity of issued patents is one in which both applicants and the PTO have incentives to cooperate in developing and improving applications of high quality.

Though they do lead to a shift in attention from existing scholarship, the first two points are hardly controversial: Who can contend that reforms should ignore the patenting process or that reforms will work better if they are not informed by data? The third point, however, warrants some elaboration because the conception of quality that animates our discussion calls for a fundamental cultural shift in the perspective through which specific procedures and reforms are evaluated. Thus, we acknowledge a stark tension between the joint production conception of our analysis and the perspective of entitlement that permeates the system. It is perhaps only a small exaggeration that patent applicants and their attorneys often view patents like welfare payments: something to which they are entitled unless the agency can identify a serious substantive defect with the request. Conversely, the agents of the public within the PTO traditionally have been motivated by a point system and procedural framework that gives them only limited and indirect incentives to improve the quality of applications, and far too little power to reject poor applications.

To be sure, the system does impose some obligations on the applicant and leaves open a real possibility that a patent might be unenforceable because of misconduct during the application process (Merges & Duffy 2007:1102–40). However, in general, the applicant's duties are complete when it submits a document on the requisite form and pays the appropriate fee (Burk & Lemley 2009:23). At that point, the burden of effort on the application shifts almost entirely to the PTO (Rai 2009:2075–76). The entitlement perspective is if anything even more damaging during the prosecution process, where examiners as a practical matter have no way to dismiss weak applications (Lemley & Moore 2004; Benjamin & Rai 2007), and for years have worked under incentive structures that provide powerful incentives to grant dubious applications pressed by motivated applicants (Rai 2009:2062–65; Office of Inspector General, U.S. Department of Commerce 2004).²² The

²²The PTO has responded to this problem in several ways in recent years, attempting to reward examiners more for accuracy than for disposal of cases. It remains to be seen how successful those efforts will be. See U.S. Patent and Trademark Office (2010a).

problem is only exacerbated by the ethical obligations zealous attorneys face to obtain the broadest possible scope of protection for their inventive clients.

In that vein, the findings in Section IV provide two distinct types of support for the shift in perspective that we propose. First, at the global level, the predictive power of the analysis underscores the importance of the prosecution process to the ultimate legal question of validity. The existing patent literature proceeds on the static assumption that applications are good or bad, and that little can be done to improve the process except to encourage the PTO to more accurately determine which applications should be rejected. The findings of Section IV reflect a more complex and dynamic conception of the application, under which most, if not all, applications could be improved through the application of more effort by the applicant, the examiner, or both working together.

Most importantly, the predictive power of the first-stage analysis shows that some of what is good and bad about the applications can be discerned readily at the moment they reach the patent office, and a great deal of what is good and bad about the applications can be discerned from what has happened to them before they leave the office. This strongly suggests the need for continuing research in the area, taking advantage of the unique data-collection opportunities available only to the PTO, analyzing more information about different sets of patents (reexamined patents, patents appealed to the Federal Circuit directly from the Board, etc.). Of central significance for this article, it underscores the possibility that improvements in that process could increase the share of issued patents that are valid even if those improvements did nothing to alter the pool of incoming applications.

More specifically, the analysis in Section IV suggests a variety of approaches that could be considered to improve the effort of the applicant, the examiners, and the PTO as a whole. Our goal in all cases is to improve the incentives of both the examiners and the applicant. In general, we argue that the PTO should devise internal criteria to monitor “red-flag” applications, and then create incentives for applicants and examiners to revise or terminate such applications. Our suggestions at this point are only illustrative, building on the relations evident from the analysis in Section IV. Still, they do suggest that the PTO could implement data-driven reforms that would improve the likely validity of the patents that issued from the process.

A. Applicant Effort

The most important change in the prosecution process would be to alter the ethic of entitlement to one of collaboration. The applicant should have discernible incentives to provide the best possible application and to respond in the most effective way to concerns expressed by the examiner during the prosecution process, and the examiner should have parallel incentives to provide effective feedback leading toward an issued patent that accurately defines the innovation. Several distinct issues are apparent.

1. Care in Drafting

More than a quarter of the patents held invalid in the data set suffered from drafting problems, which generally reflect a failure to include a specification that adequately describes and enables an invention that is delineated with definiteness in the claims. The

text-based analysis we present in this article suggests that it is well within the reach of existing technology to design simple software-based tools that can measure the alignment between the specification and the claims in a way that matches reasonably well to the ultimate legal question the PTO faces when it evaluates an application. If this is so, then without the expenditure of any examiner time, applications could be objectively evaluated and subjected to differential treatment based on how well they have been drafted.

Several types of administrative responses are apparent. Perhaps the PTO could reject poorly aligned applications out of hand under Section 112. This has the advantage of removing applications from the PTO's docket and of balancing the workload: if the applicant has not met an objective baseline of aligning the claims and the specification, then the application does not yet warrant a thorough examination. Another approach would subject poorly aligned applications to higher fees, reflecting an anticipated need for extra examiner time (Rai 2009).²³ Or they could be shunted to a slow track in which they would receive extra layers of review.²⁴ In any case, the expected response would be for applicants to improve their applications before filing them with the PTO to ensure that they avoid whatever adverse treatment is contemplated. At that point, the quality of applications submitted to the PTO might measurably improve. So long as the premise holds—that the software tool accurately can test whether the specification adequately conforms to the claims—the tool cannot be gamed because the only response that satisfies the tool is a response that in fact improves the quality of the application.²⁵

2. Overclaiming

To the extent the existing academic literature has discussed the number of claims in patents, it has emphasized the positive correlation between the number of claims and the likelihood that the patent will be litigated, and reasoned from that correlation to the view that claims correlate with quality (Allison et al. 2004). If that relation were taken seriously, it would be difficult to justify the PTO's proposed restrictions on patents with large numbers of claims.²⁶

As discussed above, however, the data that we present here suggest no positive correlation between the number of claims and the validity of the patent: patents with more claims are more likely to be selected into litigation, but when they are selected into litigation they are, if anything, less likely to be held valid than those with fewer claims.

²³For discussion of the importance (and difficulty) of enhancing PTO fee authority, see Rai (2009:2067–72).

²⁴The PTO already promises accelerated examination to applicants that meet certain disclosure requirements (U.S. Patent and Trademark Office 2006, 2010b).

²⁵Because application of such a tool would increase the likelihood that the specification matches well to the claims, it should have the ancillary salutary effect of decreasing the uncertainty of claims interpretation that led to the Federal Circuit's en banc decision in *Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed. Cir. 2005). Clarity of patent boundaries is an important consideration even apart from validity (Bessen & Meurer 2008).

²⁶37 C.F.R. § 1.75(b)(1) (2008) (discussing examination support documents for applications with more than five independent claims or 25 total claims) (withdrawn after *Tafas v. Doll*, 559 F.3d 1345 (Fed. Cir. 2009)).

Indeed, the multivariate analysis suggests that there is no robust relationship between the number of claims and validity. There are any number of obvious reasons, totally unrelated to validity, why patents with more claims might more likely be the subject of litigation. For example, they might reflect broader inventions, and thus be more likely to write onto the goods and services of competitors. Or the variation of claims might lead their holders to perceive a greater chance that they will be able to identify a claim that is broad enough to write on the competing product but narrow enough to withstand validity challenges.

To be sure, it is possible to suggest reasons why patents with more claims might more likely be invalid. Most obviously, the broader number of claims might increase the likelihood that at least one of the claims is invalid, but, given the discussion above, it is not at all clear that the marking of validity in this data set would capture those patents as invalid; the relevant question in Federal Circuit litigation is likely to turn on whether claims of central economic importance can be upheld, and patents with more reticulated claims structures well might fare better in that venue than those with simpler claims structures.

The fact remains, however, that the data presented here do not suggest a strong relation between the number of claims and the ultimate question of validity. As discussed there, this is true despite our efforts to analyze the number of claims in a variety of functional forms.²⁷ In our view, then, the data presented here cast doubt on efforts to use claims limits as a method to improve patent quality. That is not to say that the PTO could not identify other data sets that would link particularly large numbers of claims to problems with validity. However, given the difficulty of identifying a robust reason why the number of claims should relate positively or negatively to validity, the absence of strong support in the data makes us skeptical that this should be an important part of the reform agenda. We do not doubt that applications with large numbers of claims impose larger burdens on the PTO, and we believe it makes sense for the PTO to pass a share of the examination costs back to applicants in the form of larger fees for such applications. But those fees should be regarded as based on workload. For now, the case that they improve quality is unimpressive.

3. Applicant Search

Previous scholars have documented in detail the dysfunction of the patent system's existing structure of incentives for applicants to search and identify relevant prior art (Benjamin & Rai 2007; Lemley 2001). Among other things, the desire to avoid liability for willful infringement, combined with a duty of candor limited only to prior art of which the applicant actually is aware, combine to ensure that some companies perceive a strong incentive not to discover or disclose prior art related to their invention (Burk & Lemley 2009:23; Lemley & Tangri 2003).

²⁷In addition to the variables for claims and independent claims in the patent and application, we explored variables for dependent claims, for various ratios between independent claims and total claims for changes in the number of claims between the application and the patent. We also explored quadratic forms of the variables to test the possibility that relations with validity occur only at the extreme ends of the distribution. None of those efforts revealed a robust relation to validity.

The data presented here underscore the adverse effect this has on the prosecution process. The most obvious effects are the strong relations between validity on the one hand and IDS filings (a positive relation) and examiner-added references on the other (a negative relation). The significance of those findings is underscored by the evidence about the high share of examiner-added references presented in Cotropia et al. (2010). The applicant is much better situated to locate and understand the relevance of prior art related to the applicant's invention. This is true in part for the simple reason that the applicant has much more time to spend on the invention; we doubt many applicants file applications without spending far more than the 20 hours that would be a great superfluity of time on the part of the examiner. It also is true because the applicant or its lawyer should know a great deal more about the relevant technical field than the examiner studying the application.

A robust search effort by the applicant also should improve the quality of the patent application because it should enhance the crispness with which the applicant can identify the core inventiveness that justifies the patent: only by remarkable luck can an application written without detailed knowledge of the prior art include claims that distinguish the claimed invention from existing technology with precision. Creating an incentive to use skilled legal counsel likely to integrate the prior art capably into the application is an important goal here. To be sure, it is possible for the examiner to force redrafting of the application once the examiner identifies the prior art, but editing of claims in response to specific challenges is a poor substitute for informed drafting of the entire application in the first instance.

Again, we are agnostic about the question of regulatory design, which is indeed a difficult one. In this particular case, the variety of legal rules that interact to undermine the incentives of the applicant suggest that substantial statutory reform is required to ensure a robust and affirmative obligation on the part of the applicant to provide information about the prior art.²⁸ We note one possibility discussed prominently in the existing literature: to remove the statutory presumption of validity with respect to art not cited to the examiner (Lichtman & Lemley 2007). At bottom, however, it seems clear that applications in which the applicant does not undertake to provide a reasonable assessment of prior art should be treated more cautiously than those in which the applicant does provide such an assessment. At the same time, any system that obligates applicants to provide prior art must deal with the obvious problems of designing such a proposal to avoid large amounts of irrelevant information; a duty to provide information could lead to irrelevant information either out of an abundance of caution or, less benignly, as a species of prior art flooding designed to confuse the examiner.²⁹ Reliance on the peer-to-patent process (www.peertopatent.org)

²⁸One of the hardest problems is limiting the risks that inadvertent errors will render the patent invalid in subsequent litigation, a realistic concern under current Federal Circuit rules for inequitable conduct (Rai 2009:2075–76).

²⁹For what it is worth, our data do not suggest that such a problem exists at the present time—there is no greater likelihood of invalidity for the patents with unusually large numbers of IDS submissions or dropped references. However, such a problem might emerge rapidly if applicants become affirmatively obligated to provide relevant information. Any reform in this area should monitor data continuously to identify any such adverse outcomes.

also could be constructive as a way to respond to systematic weaknesses in the PTO's capabilities to discover prior art, and it might be particularly useful for nonpatent prior art.

One promising idea here, though more ambitious, would be to establish a system of qualified search authorities. Applications accompanied by search certificates from qualified search authorities could receive preferential treatment, including an exemption from liability for inequitable conduct.³⁰ To be sure, there is the risk that the details of such a proposal might disadvantage smaller or less liquid applicants, especially if they effectively shift the costs of examination to the applicant. We emphasize, however, that our principal goal is not to press any particular reform; rather, our goal is to emphasize that the combination of our data with the implications of theory and practice all suggest a major flaw in the existing system.

B. Respecting the Results

Another set of issues relates to finality of the examiner's process. One of the most recognized problems with the present process is the limited practical ability of examiners to finally close the consideration of an application. Even "final" rejections are final only for those applicants that do not wish to press the matter further. A coherent structure for the prosecution process necessarily includes some mechanism for the examiners to terminate further examination of the invention, leaving disgruntled applicants to their remedies before the Board and the Federal Circuit. The data speak to this problem in three distinct ways.

First, the strong inverse relation between continuations and validity buttresses the widely held perception that continuation practice affords inventors too many bites at the apple (Lemley & Moore 2004). This increases both the likelihood that frustrated examiners will approve weak applications and that strategic applicants will be able to amend claims before issuance to cover inventions that others have made independently during the pendency of the application. The PTO of course has attempted to rein in this practice, but courts to date have been unsympathetic.³¹ The data afford an empirical basis that justifies continued pressure on this point, either regulatory or statutory.

Second, it is important to ensure that review by the Board improves the quality of the patents that are issued, rather than the reverse. Given the increased expertise and time for

³⁰For discussion of a similar proposal, see Statement of Jon W. Dudas, Under Secretary of Commerce for Intellectual Property Before the Committee on the Judiciary, U.S. Senate (June 7, 2007), available at <<http://www.uspto.gov/news/speeches/2007/2007jun06.jsp>>; USPTO Contracts International Patent Application Searches to Commercial Firms (PTO Press Release 05 48, Sept. 21, 2005), available at <<http://www.uspto.gov/news/pr/2005/0548.jsp>>.

³¹The continuation rules formerly codified at 37 C.F.R. § 1.78(d) (1) (2008) were held invalid in *Tafas v. Doll*, 559 F.3d 1345 (Fed. Cir. 2009), and then withdrawn shortly after the Obama Administration took office. To be sure, many informed applicants offered persuasive criticisms of the particular approach taken in the rules invalidated in *Doll*. It remains to be seen whether a second effort will either exhibit more nuance or withstand judicial review.

deliberation by the Board, we expected that those patents would be much more likely to be valid than patents that had not gone through Board review. In the data, however, a better explanation seems to be that the Board's quasi-judicial decision-making process gives an advantage to the law firm presentations characteristic of sophisticated applicants, as opposed to the presentations from those within the PTO. It is perverse that the patents issued after Board review are so much more likely (50 percentage points more likely by our estimation) to be invalid than those issued without Board review. Given the size of the data set, and the relative rarity of Board review, the number of patents in question is so small that other explanations might exist. Still, the data do suggest that further inquiry into the reliability of the Board's determinations is appropriate. Among other things, analysis of decisions appealed to the Federal Circuit (where affirmance seems much more likely) might illuminate contrasts between Board decisions in favor of and against the applicant.

Finally, from the opposite perspective—what happens when examiners approve applications?—the data suggest an odd but strong relation between Rule 312 amendments and invalidity. Rule 312 permits amendments after a patent application has been approved, with the purpose of correcting errors before the patent actually is issued.³² In concept, the relevant amendments should be minor and nonsubstantive. Yet, in the data set, the patents issued after those amendments are much less likely to be valid than those without those amendments. What this suggests is that the defects repaired by those amendments are papering over other more serious problems with the patents.³³ Accordingly, it might be appropriate to scrutinize patents subject to those amendments much more carefully than currently is done. To be sure, additional scrutiny might deter those motions. However, if that is true, it might result in those patents being issued with even more glaring defects. If so, then perhaps those patents would be less costly in litigation because they could be dealt with more expeditiously. In any event, the relations apparent from the data suggest the need for empirical inquiry into the types and consequences of Rule 312 amendments under existing practice.

VI. CONCLUSION

Our article starts from a simple premise, absent from the burgeoning empirical literature on patents, that the first place to start in improving the patent system is to understand how valid patents differ from invalid patents. The apparent flood of invalid patents could be lessened substantially if the seeds of invalidity could be identified—and removed—during the prosecution process. To be sure, given the small number of patents that ever will end up in litigation there is reason to doubt that the solution to the patent problem is to increase the resources expended on patent examination by an order of

³²See 37 C.F.R. 1.312.

³³Review of the Rule 312 amendments to the patents in our data set suggests that they quite often effect substantial changes such as an alteration in the language of the claims.

magnitude. However, as long as we have a prosecution process, surely it is important to ensure that the resources we do expend work to sort valid from invalid patents as best as they can.

We acknowledge that our perspective ignores numerous serious problems in the patent system—not only substantive questions about patentable subject matter, but also important issues about allocation of decision-making authority (between courts and the PTO, between the Federal Circuit and district courts, and between judges and juries)—but if the root of the problem is the weakness of the prosecution process, empirical attention to that process must be part of any serious effort. We recognize that our effort relies on a small set of the patents that have issued from that process, and that many of the possibilities we discuss could be evaluated more thoroughly in the light of more complete information. However, that only suggests the most important of all reforms, an opening up of the PTO's records to provide ready access to more complete information about all the applications it examines.

REFERENCES

- Alcácer, Juan, & Wilbur Chung (2007) "Location Strategies and Knowledge Spillovers," 53(5) *Management Science* 760.
- Alcácer, Juan, Michelle Gittelman, & Bhaven Sampat (2009) "Applicant and Examiner Citations in U.S. Patents: An Overview and Analysis," 38(2) *Research Policy* 415.
- Allison, John R., & Mark A. Lemley (1998) "Empirical Evidence on the Validity of Litigated Patents," 26 *American Intellectual Property Lawyers Association Q. J.* 185.
- Allison, John R., Mark A. Lemley, Kimberly A. Moore, & R. Derek Trunkey (2004) "Valuable Patents," 92 *Georgetown Law J.* 435.
- Allison, John R., & Ronald J. Mann (2007) "The Disputed Quality of Software Patents," 85 *Washington Univ. Law Rev.* 297.
- Benjamin, Stuart Minor, & Arti K. Rai (2007) "Who's Afraid of the APA: What the Patent System Can Learn from Administrative Law," 95(2) *Georgetown Law J.* 269.
- Bessen, James (2008) "The Value of U.S. Patents by Owner and Patent Characteristics," 37(5) *Research Policy* 932.
- Bessen, James, & Michael Meurer (2008) *How Judges, Bureaucrats and Lawyers Put Innovators at Risk*. Princeton, NJ: Princeton Univ. Press.
- Burk, Dan L., & Mark A. Lemley (2009) *The Patent Crisis and How the Courts Can Solve It*. Chicago, IL: Univ. of Chicago Press.
- Cotropia, Christopher A., Mark Lemley, & Bhaven Sampat (2010) "Do Applicant Patent Citations Matter? Implications for the Presumption of Validity," Mimeo. Available at <http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1656568>.
- Dunner, Donald A. (1985) "The United States Court of Appeals for the Federal Circuit: Its First Three Years," 13 *American Intellectual Property Lawyers Association Q. J.* 185.
- Hall, Bronwyn H., & Megan MacGarvie (2006) "The Private Value of Software Patents," *NBER Working Paper* 12195.
- Harhoff, Dietmar, Frederic M. Scherer, & Katrin Vopel (2003) "Citations, Family Size, Opposition and the Value of Patent Rights," 32(8) *Research Policy* 1343.
- Henderson, Rebecca, Adam B. Jaffe, & Manuel Trajtenberg (1998) "Universities as a Source of Commercial Technology: A Detailed Analysis of University Patenting, 1965–1988," 80(1) *Rev. of Economics & Statistics* 119.

- (2005) “Patent Citations and the Geography of Knowledge Spillovers: A Reassessment: Comment,” 95(1) *American Economic Rev.* 461.
- Jaffe, Adam B., & Josh Lerner (2004) *Innovation and its Discontents: How Our Broken Patent System is Endangering Innovation and Progress, and What to Do About It*. Princeton, NJ: Princeton Univ. Press.
- Jaffe, Adam B., & Manuel Trajtenberg (2002) *Patents, Citations & Innovations: A Window on the Knowledge Economy*. Cambridge, MA: MIT Press.
- Jaffe, Adam B., Manuel Trajtenberg, & Rebecca Henderson (1993) “Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations,” 108(3) *Q. J. of Economics* 577.
- Lanjouw, Jean, & Mark Schankerman (2004) “Patent Quality and Research Productivity: Measuring Innovation with Multiple Indicators,” 114(495) *Economics J.* 441.
- Lei, Zhen, & Brian Wright (2009) *Why Weak Patents? Rational Ignorance or Pro “Customer” Tilt?* Available at <http://ageconsearch.umn.edu/bitstream/49279/2/Lei_Wright_AAEA_2009.pdf>.
- Lemley, Mark A. (1994) “An Empirical Study of the Twenty Year Patent Term,” 22 *American Intellectual Property Lawyers Association Q. J.* 369.
- (2001) “Rational Ignorance at the Patent Office,” 95(4) *Northwestern Univ. Law Rev.* 1495.
- Lemley, Mark A., & Kimberly A. Moore (2004) “Ending Abuse of Patent Continuations,” 84(1) *Boston Univ. Law Rev.* 63.
- Lemley, Mark A., & R. K. Tangri (2003) “Ending Patent Law’s Willfulness Game,” 18(4) *Berkeley Technology Law J.* 1085.
- Lichtman, Doug (2004) “Rethinking Prosecution History Estoppel,” 71 *Univ. of Chicago Law Rev.* 151.
- Lichtman, Doug, & Mark A. Lemley (2007) “Rethinking Patent Law’s Presumption of Validity,” 60 *Stanford Law Rev.* 45.
- Long, Clarisa (2009) “The PTO and the Market for Influence in Patent Law,” 157 *Univ. of Pennsylvania Law Rev.* 1965.
- Merges, Robert P., & John F. Duffy (2007) *Patent Law and Policy: Cases and Materials*. Charlottesville, VA: LexisNexis.
- Moore, Kimberly A. (2005) “Worthless Patents,” 20 *Berkeley Technology Law J.* 1521.
- Nard, Craig Allen, & John F. Duffy (2007) “Rethinking Patent Law’s Uniformity Principle,” 101 *Northwestern Univ. Law Rev.* 1619.
- National Research Council of the National Academies (2004) *A Patent System for the 21st Century*. Washington, DC: National Research Council.
- Office of Inspector General, U.S. Department of Commerce (2004) *Final Inspection Report No. IPE 15722, USPTO Should Reassess How Examiner Goals, Performance Appraisal Plans, and the Award System Stimulate and Reward Examiner Production*. Washington, DC: Office of Inspector General.
- Owen Smith, Jason, & Walter W. Powell (2003) “The Expanding Role of University Patenting in the Life Sciences: Assessing the Importance of Experience and Connectivity,” 32(9) *Research Policy* 1695.
- Porter, M. F. (1980) “An Algorithm for Suffix Stripping,” 14(3) *Program* 130.
- Rai, Arti K. (2009) “Growing Pains in the Administrative State: The Patent Office’s Troubled Quest for Managerial Control,” 157 *Univ. of Pennsylvania Law Rev.* 2051.
- Reitzig, Markus (2004) “Improving Patent Valuations for Management Purposes—Validating New Indicators by Analyzing Application Rationales,” 33(6–7) *Research Policy* 939.
- Rivette, Kevin G., & David Kline (1999) *Rembrandts in the Attic: Unlocking the Hidden Value of Patents*. Cambridge, MA: Harvard Business Press.
- Sapsalis, Eleftherios, Bruno van Pottelsberghe, & Ran Navon (2006) “Academic Versus Industry Patenting: An In Depth Analysis of What Determines Patent Value,” 35(10) *Research Policy* 1631.
- Singh, Jasjit (2005) “Collaborative Networks as Determinants of Knowledge Diffusion,” 51(5) *Management Science* 756.
- Singh, Jasjit, & Lee Fleming (2010) “Lone Inventors as Sources of Breakthroughs: Myth or Reality,” 56(1) *Management Science* 41.
- Singh, Jasjit, Matt Marx, & Lee Fleming (2010) “Patent Citations and the Geography of Knowledge Spillovers: Disentangling the Role of State Borders, Metropolitan Boundaries and Distance.” Available at <<http://www.insead.edu/facultyresearch/research/doc.cfm?did=43298>>.

- Sorenson, Olav, & Jasjit Singh (2007) "Science, Social Networks and Spillovers," 14(2) *Industry Innovation* 219.
- Thompson, Peter (2006) "Patent Citations and the Geography of Knowledge Spillovers: Evidence from Inventor and Examiner Added Citations," 88(2) *Rev. of Economics & Statistics* 383.
- Thompson, Peter, & Melanie Fox Kean (2005) "Patent Citations and the Geography of Knowledge Spillovers: A Reassessment," 95(1) *American Economic Rev.* 450.
- U.S. Patent and Trademark Office (2006) "Changes to Practice for Petitions in Patent Applications to Make Special and for Accelerated Examination," 71 *Federal Register* 36,323.
- (2010a) "Recently Announced Changes to USPTO's Examiner Count System Go into Effect," press release October 8.
- (2010b) "Enhanced Examination Timing Control Initiative: Notice of Public Meeting," 75 *Federal Register* 31,763.
- Wuchty, Stefan, Benjamin F. Jones, & Brian Uzzi (2007) "The Increasing Dominance of Teams in Production of Knowledge," 316(5827) *Science* 1036.