

Examiner Characteristics and the Patent Grant Rate^{*}

Mark Lemley[†] and Bhaven N. Sampat[‡]

[†]William H. Neukom Professor, Stanford Law School; of
counsel, Keker & Van Nest LLP

[‡]Assistant Professor, Department of Health Policy and
Management Mailman School of Public Health, Columbia
University

January 3, 2009

Abstract

In this paper, we show that there are important differences across patent examiners at the U.S. Patent and Trademark Office (USPTO), and that these relate to the most important decision made by the USPTO: whether or not to grant a patent. We find that more experienced examiners, and those who systematically cite less prior art, are more likely to grant patent applications. These results are not encouraging as a matter of public policy. But they do point to human resource policies as potentially important levers in patent system reform. (JEL Keywords: O34 - Intellectual Property Rights; O38 - Government Policy)

^{*}We thank Iain Cockburn, Wes Cohen, Ernie Beffel, Chris Cotropia, Al Engleberg, Sherry Glied, John Golden, Stuart Graham, Matt Neidell, Arti Rai, Scott Stern, and participants at the NBER Productivity Lunch, the 2008 NBER Summer Institute, the Georgia Tech Roundtable on Entrepreneurship Research (REER), and UNC Law School Conference on Empirical Patent Scholarship for useful comments. Sarah Craven, Reuven Garrett, Mayank Dabral, and Matthew Wasserman provided excellent research assistance. Patrick Crosby at xblabs.com helped obtain some of the data used in this analysis. Sampat thanks the Ford Foundation for financial support.

1 Introduction

Under the standard economic rationale, patents provide incentives to conduct research and to disclose inventions, in return for a limited right to exclude others from making, selling, or using these inventions (Nordhaus, 1969). Under optimally designed patent systems the static social costs from the market power conferred by patents are offset by the dynamic welfare gains from increased invention and disclosure. From this perspective, if patents were granted on inventions that were already known, they would create harms without concomitant social gains (Lemley and Shapiro, 2005). In the United States, examiners at the Patent and Trademark Office (USPTO) are charged with serving as gatekeepers to ensure that only novel and non-obvious inventions are granted patents.¹

In this paper, using information from several novel datasets, we show that there are important differences across patent examiners, and that these relate to the most important decision made by the USPTO: whether or not to grant a patent. In particular we find that examiners differ in significant and important ways in their experience and the depth of their prior art searching, and that these examiner characteristics have qualitatively and statistically significant impacts on whether a patent application is granted.

The results are not encouraging as a public policy matter, because they suggest that the decision to allow or reject a patent application may not be driven by the merits of that application, but rather by the luck of the draw. At the same time, they suggest that human resource policies and incentive structures at the USPTO could affect patent grant rates, an important finding amidst growing concern that it grants too many "low quality" patents and is in need of significant reform (Jaffe and Lerner 2004; NAS 2004; FTC 2005).

Several scholars have previously studied examiner heterogeneity and its effects on patenting. Cockburn, Kortum and Stern studied the effect of particular examiners on the characteristics of issued patents and their survival in litigation. (Cockburn et al. 2003) They find that examiner fixed effects explain a significant percentage of the variation in the characteristics of issued patents, and that examiner differences affected litigation outcomes ----- patents issued

¹ Or, as Thomas Jefferson--who as Secretary of State was among the nation's first patent examiners--suggested (in considering the originality of an 1813 invention on "Elevators, Conveyors, and Hopper--Boys,") patent examiners have the difficult job of "drawing a line between the things which are worth to the public the embarrassment of an exclusive patent, and those which are not." (The Writings of Thomas Jefferson, p. 355).

by certain examiners were more likely to be upheld in court than those issued by others. They conclude that "there may be as many patent offices as there are patent examiners." In another study, Doug Lichtman studied the role of the PTO in compelling amendments during the prosecution process, an issue directly relevant to prosecution history estoppel and the application of the doctrine of equivalents (Lichtman 2004). Lichtman identified the extent to which issued claims differed from the claims as originally drafted. This study, too, found examiner effects to be important: whether an applicant amended its claims depended in significant measure on which examiner reviewed the application. Lichtman concluded that different examiners had different "styles," with some examiners systematically more likely than others to compel applicants to narrow their claims.

In this paper, rather than looking at issued patents, we identify a group of patent applications and follow them through the process to examine the impact of examiner characteristics on patent application outcomes. Specifically, we collected every original utility patent application filed in the month of January 2001 and published before April 2006. After eliminating plant, design, and reissue patents, PCT applications directed at foreign filing, and continuations, divisionals, and CIPs based on earlier applications, we were left with 9,960 applications. We then collected information on the status of these applications as of June 2008, and other information about the prosecution process as of April 2006.² Evaluating applications rather than just issued patents allows us to obtain richer data on the behavior of examiners, data that was not available before 2001.

As reported in Lemley and Sampat (2009), 70.5 percent of the applications had resulted in patents by June 2008. About a quarter of the applications, 27.3 percent, were not patented because they were abandoned by the applicant. These abandonments could be because the applicant could not overcome an examiners rejection, but the applications could also be abandoned for other reasons (e.g. the applicants firm went out of business).³ The remaining 2.2 percent of the applications were pending as of June 2008.

In this paper, we focus on the set of these applications that have received a

²Unfortunately, changes to the PAIR interface make it difficult to update the prosecution data beyond April 2006.

³There is no actual way for the PTO to finally reject a patent (Lemley and Moore, 2004). We consider a patent application to have been finally abandoned if the applicant has filed a notice of abandonment or has not responded to a PTO rejection or filed an appeal within six months, the time limit for doing so.

final disposition, i.e. the 98% of the applications that are patented or abandoned. We relate whether these 9744 applications were granted by June 2008 to examiner--specific characteristics, including examiner experience and prior art searching tendencies. In the next section, we describe the patent prosecution process, both to provide context for our analysis and to motivate construction of the variables we examine. In Section 3, we describe how we collected the examiner data. In Section 4, we provide descriptive statistics on, and explore relationships between, the key independent variables. Section 5 presents results from linear probability models relating examiner and application characteristics to the probability an application is granted. There, we also relate examiner characteristics to another measure of examination rigor, whether applications were amended during prosecution.

One concern about our analysis is selection: if specific types of applications are assigned to specific examiners, this could bias our results. Section 6 examines this possibility, and suggests that our results are not driven by selection bias. Section 7 explores some possible mechanisms driving our results. Section 8 concludes.

2 The Patent Prosecution Process

Roughly 450,000 new patent applications are filed every year. Each of these applications contains, among other things, written descriptions of the invention and specific claims defining the boundaries of the property right the applicant hopes to obtain. In the United States, applicants also have a duty of candor to disclose any previous patents and publications, or "prior art", that are material to patentability of the claimed invention. Accordingly, most (but not all) patent applicants also disclose patent and non--patent prior art as part of their patent application. Though there is a duty of disclosure in the United States, there is no affirmative requirement that applicants conduct prior art searches; Sampat (2009) and Alcacer et al. (2009) argue that incentives for applicants to do so vary across inventions and industries.

Once applications arrive at the USPTO, they are divided by the PTO into technology classes, or Art Units. Supervisory Patent Examiners (SPEs) within each of the art units assign particular applications to particular examiners based on a rather loose set of rules (MPEP sec. 903.08(b)). One issue relevant to our empirical analyses below is whether there is sorting, i.e. whether particular types of applications are assigned to particular types of examiners.

ers. Merges (2001) suggests that while sorting could be good from a policy perspective, there is a strong "all patents are created equal" tradition at the PTO militating against this. Our interviews with SPEs suggest that there is some sorting, but that familiarity with particular technologies and docket flow management, rather than judgments of an application's quality or patent-worthiness, are the dominant considerations.

The patent examiner assigned the application reviews it and conducts his or her own search for prior art that might make the application unpatentable. This involves searching databases of previous U.S. patents, either manually or through algorithms available to examiners. Examiners may also search foreign patents and the non-patent literature (e.g., scientific and technical journals) to find prior art that might compromise patentability. Searching the non-patent prior art in particular may be more difficult: Thomas (2001) argues that "[i]n comparison to much of the secondary literature [non-patent prior art], patents are readily accessible, conveniently identified, and printed in a common format. Identification of a promising secondary reference, and full comprehension of its contents, often prove to be more difficult tasks." (318)

Examiners then assess the novelty and non-obviousness of the claims in the application, relative to what is disclosed in the complete list of prior art, i.e. the prior art references from the applicant plus any discovered through the examiners own search. Examiners generally issue an initial rejection of the application, setting out the problems they find in one or more of the claims. (Lemley & Sampat, 2008). The applicant responds, generally by amending the patent claims or by disputing the rejection. After the response, the examiner may then allow the application or issue a "final rejection." Even a final rejection is not really final, however; applicants can respond by amending their application, or by requesting an interview with the examiner to press their case. They may also file one of a variety of "continuation" applications to continue to argue for patentability.

If the patent issues, the front page includes a range of bibliographic data, including the final claims, and information on the applicant and examiner involved in examination. Issued patents also list all of the prior art references considered during the prosecution process, and, since 2001, indications of whether these references came from applicant disclosures or examiner searches (Alcacer and Gittleman 2006; Sampat 2009; Alcacer, et al. 2009).

The large number of applications facing the PTO means that examiners are subject to sharp time constraints; the entire process of reading and evaluating

an application, searching for prior art, writing a rejection, responding to an amendment with a second office action, having an interview, and fulfilling various formal requirements can take 3--4 years on average (Allison & Lemley, 2000), but the examiner spends an average of only 18 hours over those years working on any given application. (Lemley, 2001).

The incentives facing examiners are complicated, and the object of considerable policy debate (Jaffe and Lerner 2004; GAO 2005). Examiners' promotion and bonus decisions are tied to the number of "counts" they accumulate. These counts can be earned by first office actions, or by disposal of cases, which occur when an application is granted or abandoned. Other examiner activities, including searching for prior art, compelling further amendments after a first office action, or issuing final rejections, are not rewarded by the count system. Some have argued that these rules create a bias towards granting rather than rejecting applications (Jaffe and Lerner 2004).

Moreover, these incentives operate differently over an examiner's career. Examiners are hired at different grades (GS--5, GS--7, GS--9, or GS--11) on the government employee pay schedule, depending on their technical experience and educational background. The counts per hour they are required to attain to achieve satisfactory performance reviews (and thus further promotion) and annual bonuses increase sharply with pay grade.⁴ Another difference between more junior and senior examiner is the scrutiny they receive. Junior examiners work is subject to review from more senior examiners; indeed, though they do the bulk of the examination, they are listed as secondary examiners on applications they examine. Once promoted to the GS--14 level (generally between 4 and 6 years of hire, depending on what level they were hired at) they receive signatory authority, or the right to sign off on an application independently.

⁴ A recent GAO report provides this illustration: "a GS--12 patent examiner working on data processing applications is expected to achieve two counts in 31.6 hours, whereas a GS--12 patent examiner working on plastic molding applications is expected to do so in 20.1 hours. GS--7 patent examiners working on those types of applications, however, are expected to achieve two counts in 45.1 and 28.7 hours, respectively." <http://www.gao.gov/htext/d08527t.html>

3 Examiner Characteristics

This discussion suggests that an examiner's propensity to search for prior art, and her experience, might affect how she handles patent applications they are assigned. To construct data on prior art search patterns, we began by linking the examiners in our January 2001 sample to patents they issued over the 2001--2006 period. We used information on whether the citations in these 2001--2006 patents (to previous patents and to non--patent literature) were provided by the examiner or by applicants to construct measures of examiner search intensity, as we discuss in detail below.

Since the PTO employs its examiners, examiner names are listed in PAIR, and examiner names are listed on the front page of issued patents, it might seem straightforward to identify the examiners associated with each application, and to link these to other examiner characteristics constructed from data on the front page of subsequently issued patents. However, while the PAIR examiner name data are cleanly linked to the standardized names in the official USPTO employee directory, the front page examiner data are reported in a haphazard format and are rife with errors, as Cockburn et al. (2003) point out. (One source of error is that the large--sample front page patent data provided by the USPTO appears to be constructed via optical character recognition of the patent images.) To take just one example, we identified one examiner whose name was spelled no less than 20 different ways on the front page of issued patents.⁵

To solve this problem, we acquired the USPTO Employee Directories from 1992 to the present. These Directories list the examiner name in the same standardized format as in the PAIR data. We used a combination of programming and manual correction to match each of these clean examiner names to the noisy names listed on patents issued over the 2001----2006 period. Collectively, the 2,803 examiners in our sample were listed on these later issued patents in 13,772 name variants. Access to the employee directories also allowed us to determine exactly when the examiner first began working at the

⁵ *Ponnathapura Achutamurthy's* name was listed as Achutamurthy Donnathapu, Achutamurthy P., Achutamurthy Ponnathapau, Achuta--murthy Ponnathapu, Achutamurthy Ponnathapua, Achutamurthy Ponnathapuea, Achutamurthy Ponnathapur, Achutamurthy Ponnathapura, Achutamurthy Ponnathaput, Achutamurthy Ponnathupura, Achutamurthy Punathapu, Achutamurtry Ponnathapu, Achutamurty Ponnathapu, Achuthamurthy Ponnathapu, Achutmurthy Ponnathapu, Achutyamurthy Ponnathapu, Murthy Ponnathapu, Murthy Ponnathapu Achuta, Murthy Ponnathapuachut, and Ponnathapuachuta.

USPTO, which we used to construct experience measures. These measures are right--censored, since we lack access to Employee Directories before 1992.

In our empirical analyses, we treat the examiner for each application or patent as the examiner who did the most direct work on that application: the secondary examiner if there was one, or the primary examiner if there was no secondary assigned.⁶ Using front--page data from issued patents, we constructed two measures of examiner search intensity. Using data on citations in all patents issued by an examiner over the 2001--2006 period (cf. Thompson 2006; Alcacer and Gittelman 2006) we constructed a measure of the share of all citations to (a) patents, and (b) non--patent prior art that came from the examiner rather than the applicant. Previous research has used patent citations by applicants as proxies for how well they are searching prior art (Sampat 2009; Alcacer et al. 2009). Here, we use the average examiner share of references (taken as an average, across all patents issued by that examiner) as proxies for an examiners propensity to search for prior art. Specifically, we construct the variable *US_SHARE* as a measure of how thoroughly the examiner searches U.S. patents, and *NP_SHARE* defined analogously for non--patent prior art.

The theory here is that examiners who conduct more thorough searches are, on average, likely to contribute a greater share of prior art references in patents they issue over the 2001--2006 period, relative to other examiners in their art units. Because it is an average, this measure is more precise for examiners who issue more patents. To account for this, in the models with the search intensity measures we exclude applications where the examiner had fewer than ten patents issued over the 2001--6 period, which leads us to drop 89 examiners and 714 of the January 2001 patent applications.⁷

We also used information in the examiner directory to construct experience measures. Though the applications were filed in 2001, there can be lags before

⁶Given the high turnover rate at the USPTO, it is not surprising that there are some applications for which the examiner changed over the prosecution period. Our data are on the last examiner associated with an application, the one associated with patent issuance or with rejection/abandonment. Based on manual inspection of a sample of 1,637 applications for which the image file wrappers were available, we determined that there were examiner switches for about 244 of these, or slightly less than 15 percent, and there were no qualitatively or statistically significant differences in the characteristics of original versus final examiners who handled these applications.

⁷ We also estimated models including these examiners/applications and controlling for the count of patents they issued over the 2001--6 period. These models produce qualitatively similar results, and are available on request.

they are docketed to an examiner. (Ninety percent of the applications in our sample were docketed by the end of 2001, and 98 percent by the end of 2002.) Accordingly, in our analyses we calculate experience as the difference between the first year an examiner appears in the Employee Directory and the year the application is docketed to the examiner.⁸

This experience measure is also right-censored, since we lack employee directories prior to FY1992. In view of the censoring, in our empirical analyses we examine the effects of examiner experience non-parametrically, including dummy variables for *ZERO YEARS EXPERIENCE*, *2--4 YEARS EXPERIENCE*, *5--7 YEARS EXPERIENCE*, AND *8+ YEARS EXPERIENCE*, though the main results are similar if we use different cohort groupings or treat "experience" as a continuous variable.

4 Descriptive statistics -- examiner citation patterns

Table 1 shows descriptive statistics on each of the main independent variables, calculated at the examiner level. The modal experience cohort is the first, i.e. most examiners on our January 2001 patents are relatively new. Also notable is that there are few examiners in the 5--7 year cohort, consistent with concerns that numerous examiners leave the USPTO relatively early in their careers (GAO 2005). But over a quarter of the examining corps are veterans, having stayed with the Office for over 8 years.

On average, examiners accounted for 43.6 percent of citations to patented prior art in their issued patents, but only 9.5 percent of all citations to non-patent prior art. These figures are consistent with the figures reported in Sampat (2009), and with the arguments that examiners are better at searching patented prior art than non-patent prior art. But for each of these measures, there is considerable variation around the mean, suggesting heterogeneity across patent examiners. The bivariate correlation between these two measures is positive and significant but weak ($r=.29$, $p<.001$), indicating they move together but that each may be picking up different information about search behavior.

⁸ For a non-trivial share of the applications, 16.95 percent, this experience measure was negative, in most of these cases taking on a value of -1. This is because the Employee Directories are by fiscal rather than calendar years, and also that they are occasionally updated with lags. In these cases, we recoded experience to zero. Note also that all of our main results are unchanged if we instead use 2001 minus year first in directory as our experience measure.

There is also considerable within-- and across--field variation in these measures. Figure 1 shows experience by art unit. (Here, for expositional convenience, we use six broad art units, rather than the 300--plus detailed art units employed in the regressions.) In the computer art units examiners are overwhelmingly new hires. There are also a large proportion of new hires in communications. This may reflect either greater turnover in those art units (a function in part of other job opportunities in those industries) or the growth in the number of applications (and thus hiring) in those fields. By contrast, examiners in the chemical and the mechanical art units have substantially more tenure, and indeed the largest number of them have been at the PTO for more than eight years. The biotechnology and organic chemistry art unit also includes a large number of examiners with a long tenure at the PTO, though not as many as in the mechanical and chemical industries.

Figure 2 significant cross--field variation in the pattern of prior art citations. In most industries, the average examiner share of citations to previous patents is roughly normally distributed, though the means differ across fields. In general, the share is higher in the computing and communication arenas; Sampat (2009) and Alcacer et al. (2009) suggest this may reflect lower incentives for applicants in these fields to search for prior art rather than additional work by examiners in those fields. By contrast, in biotechnology the average examiner share of citations is lower, suggesting that applicants in that industry conduct prior art searches and that examiners either rely on the submitted results or are unable to find relevant art beyond that submitted by the applicant. But even within fields there is considerable variation around the mean, indicating examiner heterogeneity.

Figure 3 shows that the distribution of citations to non--patent prior art are heavily skewed toward applicant submission in every industry. This may be a function of the time constraints under which examiners operate, and that is it harder for examiners to find for non--patent prior art for the reasons discussed above. Here again there is variation in the examiner averages within fields, suggesting that some examiners are systematically more active in searching non--patent prior art than others.

Finally, Table 2 shows regression results relating our two measures of examiner search intensity to experience, after controlling for detailed art unit effects. The first column shows that the examiner's average share of patent references decreases monotonically with experience, with the most experienced examiners' share 10.2 percentage points less than that of the least experienced examiners (the left out category), and the differences between each of the co-

hort categories statistically significant at the 1 percent level. For non--patent references, however, there are no qualitatively or statistically significant differences between over the first 7 years. But the most experienced examiners have a significantly lower share than each of the other cohorts. Relative to the least experienced examiners, their average share of non--patent references is about 1.9 percentage points lower.

These descriptive statistics already suggests that more experienced examiners behave differently than less experienced ones. We explore this in more detail below, relating examiner experience, and the prior art search measures conditional on experience, to the probability an examiner grants a given application, and conditional on grant, the probability an examiner granted the patent without issuing any rejections.

5 Empirical approach and results

5.1 Main results

In our main empirical analyses, we estimated linear probability models relating the probability that an application filed in January 2001 is granted by June 2008 to examiner characteristics, including the experience variables and the two measures of examiner search intensity. Each of the models includes fixed effects for each of over 300 art units, and robust standard errors clustered on examiners.

Table 3 shows our results. The first column shows that the grant probability increases with experience, though this effect plateaus after 7 years. The most experienced examiners have an 11.3 percentage point greater grant rate than the youngest examiners. Column Two shows that, conditional on experience, examiners who on average account for a higher share of citations to previous patents (in their 2001--6 patents) have a significantly lower grant rate. The results are also qualitatively significant. Replacing examiners in the 5th percentile of the distribution (accounting for 14 percent of citations to patents, averaged across art units) with those from the 95th percentile (accounting for 75 percent of citations) would reduce the grant rate by about 8 percentage points. Column 3 shows our measure of examiner search intensity based on non--patent references is also negatively related to the grant rate, with an almost 11 percentage point difference in the grant rate between examiners in the 5th and 95th percentiles of the distribution (accounting for 0

percent and 33 percent of references to non-patent prior art, respectively).

Taken together, these results in this section provide strong evidence that (a) more experienced examiners are more likely to grant patents; (b) this effect does not simply reflect differences in prior art search tendencies between experienced and unexperienced examiners, i.e. experience matters independently of these measures; and (c) even conditional on experience, the extent to which examiners search for prior art (as measured by the average share of citations they account for in their issued patents) is strongly related to the probability an application is granted.

The heterogeneity across patent examiners illustrated in our descriptive statistics thus relates to what is arguably the most important decision the USPTO makes, whether or not to grant a patent. As we discuss detail in the conclusion, we cannot conclude from these results whether the more or less experienced examiners, or even the more or less thorough examiners, have it right, absent priors on whether the grant rate is currently too high or too low. But it is suggestive that examiners are doing more work, and rejecting applications with more rigor, at early stages in their career, and both doing less work and allowing more patents as their tenure increases.

It may also be that whether a patent is granted is too blunt a measure of the quality of examination. As discussed in Section Two, claims can change over the course of patent prosecution. Amendments to the application compelled by examiners (in the face of prior art) can narrow the scope of the property right granted by a patent. Another measure of the rigor or quality of patent prosecution is the extent to which this narrowing occurs. We explore this below.

5.2 Examiner characteristics and changes during patent prosecution

Following an examiner office action (a non-final or final rejection), applicants can amend the offending claims in an application. These narrowing amendments can take different forms, including adding and subtracting claims, combining claims, and changing claim language. While amendments could either broaden or narrow patent claims, it is reasonable to expect that amendments made in response to an examiner rejection are more likely to narrow rather than broaden the claim.

To measure whether a patents scope was narrowed during the patent prose-

cution process, we would ideally look carefully at the final claims in the patent and compare to claims in the application. Determining the scope of the claims in a patent or application is done via careful and generally contentious interpretation of the language of the claim, the meaning of words, and even the history of the patent prosecution process. Accordingly, it is impossible to properly measure the extent of narrowing during patent prosecution process for a sample as large as ours. Indeed, even in a small sample it would be difficult if not impossible to do so in an objective way; lawyers and judges regularly disagree over whether particular amendments broaden or narrow patents even after expending years and millions of dollars assessing a particular patent.

One candidate measure we could use is the change in the number of claims between the application and issued patent. The economics literature commonly uses the number of claims in a patent as a measure of its scope; following this logic, the change in number of claims between the granted patent and the patent application might serve as a measure of narrowing of scope. However, this measure is problematic. Moore (2003), for example, questions the use of number of claims as evidence of broad scope, noting that applicants often obtain many narrow claims because they cannot get one broad claim. Our data allow us to directly test whether the number of claims is a useful measure of narrowing during patent prosecution. Specifically, for a random 5 percent sample of applications that granted patents before April 2006, or 301 applications, we used data from PAIR to determine whether there was ever an examiner rejection (non-final or final) during the patent prosecution process, and data from Delphion to determine the number of claims in the application and in the issued patent.

Figure 5 shows the histogram of net claim changes, i.e. the number of claims in the patent minus the number of claims in the application, by whether there was a rejection during the patent prosecution process. Applications granted after rejections are more likely to have been narrowed (via amendment) before grant. If the number of claims measures scope, we would expect to see applications that had rejections----those in the right panel----have a significantly lower number of claims in the final patent than in the application, i.e. all of the values should be less than zero. However, nearly half (49 percent) of the applications with examiner rejections have either no changes in the number of claims or an increase in number of claims in the patent. (By contrast, the left panel shows that 81 percent of the applications that were granted with no rejections during patent prosecution had zero changes in the number of claims.) These data suggest that the number of claims can increase,

decrease, or stay the same when an application is narrowed in response to an examiner rejection: the number of claims alone doesn't provide useful information about scope.

Instead, in our analyses we use the PAIR data on whether there was ever a rejection during patent prosecution for the 6459 applications in our sample that resulted in a granted patent by April 2006. Of these applications, 81 percent of these granted applications received a non-final rejection, and 26 percent a final rejection, during patent prosecution. But 19 percent did not receive any rejections before they were issued, i.e. they were issued on the first office action, as is. For these patents, there was no narrowing of scope during patent prosecution.

Table 4 relates the probability that these applications ever had a rejection to the experience and search intensity measures. In each of the models, the probability that the patents were granted with at least one rejection decreases steadily with experience. The most experienced examiners have 13 percentage point higher probability of issuing on first office action, i.e. without any rejections or narrowing amendments, than the most recently hired examiners. But, conditional on experience the two measures of search propensity are not significantly related to the probability of issuing without rejections.

The finding that experience increases the probability of having had no rejections, conditional on issuing a patent, rejects one possible explanation for the relationship between examination and grant rates reported above: that experienced examiners are better at getting quickly to the patentable piece of an application by negotiating with the applicant to amend his claims appropriately. Were that the explanation, experience would not be correlated with a reduction in non-final rejections, since those rejections generally occur before any amendment or negotiation opportunity.

6 Selection

Taken together, the data from the previous two sections show that not only are more experienced examiners and those who are less active in identifying prior art (conditional on experience) more likely to grant an application, but that they are also less likely to have rejected claims (and thus compelled amendments to the application) in the patents they ultimately grant. We discuss potential implications of these findings in the conclusion, but note for now that at the very least they suggest that the heterogeneity in patent examiners

we identified above is strongly related to important outcomes.

One potential threat to the validity of these results is selection. If different examiners are systematically assigned to "easier" or "harder" applications, or those more or less patent-worthy, our estimates would be biased. As discussed above, our interviews with Supervisory Examiners suggest that an examiners familiarity with a specific technology appears to play a role in assignment decisions. But for selection bias to explain our results, a different type of selection effect would be necessary. It would require, first, that the SPEs be in a position at the outset to assess whether an application was more or less likely to be patentable, and second, that they would be motivated to give those more likely to be patentable to the most experienced examiners and those who are less thorough in their prior art searching.

But our interviews reveal no evidence that SPEs do any kind of substantive evaluation of the applications before assigning them to particular examiners, and the press of work makes it implausible that they could do enough of a review to make a judgment as to whether an incoming application was likely to make it through the office. Nor would such a selection bias be logical; if SPEs were in fact engaged in some sort of sorting, we would expect it to work in the opposite direction than our data suggest, with the toughest applications rather than the easiest being assigned to the more seasoned examiners.

Nonetheless, in this section we explore selection more systematically. First, using both the application data and granted patent data, we look for evidence of selection on observables such as the length of the patent, family size, or number of claims or prior art citations. Second, we directly control for how "patent-worthy" an application is by looking at whether the same applications were granted in Europe (at the European Patent Office, or EPO).

6.1 Selection on Observables

There are two difficulties in examining potential selection on observable variables. A first is that for patent applications, much of the front-page data available for issued patents (including citations and assignee information) is not available. Accordingly, in addition to examining selection on observables for our applications, we also do supplementary analyses on the subset that eventually issued as patents.

More generally, for both applications and patents, it is difficult to identify variables that would measure the "patent-worthiness" of an invention. Patent-worthiness is ultimately based on a reading of the claims and a judg-

ment (ideally, by a person having ordinary skill in the art) about whether they are novel and non-obvious in light of the prior art.⁹

Lacking the ability to make that determination for each of the patents in our sample, we instead collected data on things we could measure to test for selection on observable characteristics. We started with examining the two variables on which we do have data at the time of application: the number of pages in the application and the patent family size, i.e. the number of countries in which an application was filed. The latter has been used as a measure of patent value in the applied literature. Table 5 shows the effects of these variables on our three examiner characteristics: years of experience (Models 1 and 2), the examiners average share of citations to patents (Models 3 and 4) and the examiners average share of citations to non-patent references (Models 5 and 6). In none of the models do these application level variables have a qualitatively and statistically significant relationship to the examiner characteristics. This is consistent with our impression that there is no application level sorting at the USPTO.¹⁰

To examine this further, we also looked at application level characteristics for granted patents, focusing again on the patents that were granted by 2006 for which we have comprehensive data. This analysis requires characteristics of the application that could plausibly influence patent-worthiness, but which would not themselves be associated with the effects of examiner characteristics on patent prosecution or features of the granted patents. Accordingly, we do not examine measures that could reflect the impact of patent prosecution process itself (like forward citation counts to an issued patent); we instead focus on "time zero" measures that reflect characteristics of the application as filed, rather than the patent as granted.

Specifically, we examine the number of references to patented prior art provided by the applicant as part of the application, the number of references to non-patented prior art, and the number of patents the applicant was issued in the previous year (2000), a measure of the applicant's experience. Table 6 shows results from regressing examiner characteristics on these variables. Models 1 through 3 show that none of these measures have a qualitatively or statistically significant relationship with examiner experience. Models 4 and 5

⁹ If there were a ready variable or set of variables that proxied for this, the USPTO's task would be much easier.

¹⁰ Model 5 does show a statistically significant relationship between number of pages in the application and the examiner's average share of cites to non-patent prior art, but the magnitude of this relationship is tiny.

do show negative and statistically significant relationships between the volume of citations (to both patented and non-patent prior art) provided by an applicant and our measure of examiner propensity to cite U.S. patents. This could suggest that applications with less patented prior art are assigned to examiners who tend to be more thorough in their searches, providing some evidence of selection. But while statistically significant these effects are qualitatively small: essentially very precisely estimated zeroes. For example, increasing the number of applicant references to patents from its 5th to its 95th percentile (an increase from 0 to 20 applicant citations) would be associated with only a 0.4 percentage point difference in the examiners average share of patent citations. A similar increase for the number of applicant non-patent references (from 0 to 6 references) would be associated with a 0.36 percentage point difference. Moreover, Models 7 through 9 show that we don't see similar effects for our other measure of examiner search propensity, an examiners average share of citations to non-patent literature.

The analyses thus show little evidence of selection on observables, using either the limited observables we have for applications overall, or the more comprehensive set of application-level measures we have for applications that eventually mature into patents. But none of these variables is a perfect measure of the patent-worthiness of the application. For reasons discussed above, it is likely impossible to measure patent-worthiness directly.¹¹ In the next section, we examine a proxy variable, whether the same application was granted in by the EPO.

6.2 Selection: Evidence from the EPO

We collected information from Delphion on whether the applications in our sample were also filed at the European Patent Office (EPO), which examines and grants European patents for the 32 States in the European Patent Convention. For applications that were not granted by the EPO, we supplemented the Delphion data with information from the European Register on whether they were rejected/withdrawn or currently pending. Following Webster et al. (2007), we limited our sample to the 8905 applications filed in January 2001 that had a single priority application, to be certain that the corresponding EPO applications were on the same invention.

¹¹For the same reason it is unlikely that SPEs are judging patent-worthiness either, in deciding how to distribute applications across examiners.

Of these U.S. applications, 2,761 were filed at the EPO, which had granted patents on 46 percent of them as of June 2008. Table 7 shows a cross-tabulation of EPO and USPTO outcomes as of June 2008. Of the applications that were granted in the U.S., slightly more than half (52.1 percent) have been granted by the EPO.¹² By contrast, of those granted by the EPO, the vast majority (88 percent) were also granted by the USPTO.

Table 8 shows the effects of the experience and search intensity measures on probability of U.S. grant, conditional on EPO status (As in the main analyses above, we drop the small number of applications still pending in the U.S.) Model 1 shows that relative to applications patented at the EPO (the left out category), applications rejected at the EPO have a 27 percentage point lower probability of being patented in the U.S. But controlling for EPO status, more experienced examiners have a statistically and qualitatively higher likelihood of grant than less experienced examiners. Similarly, Models 2 and 3 show that conditional on experience and EPO status, our two measures of examiner search intensity are negatively and significantly related to probability of grant in the U.S.

These models control for EPO status. In Table 9, we also assess the effects of experience on likelihood of U.S. patent grant separately for applications that were (a) patented at the EPO and (b) rejected at the EPO. Model 1 shows that examiner experience doesn't affect the likelihood that applications patented at the EPO are patented in the U.S.: there is no statistically significant difference between older and newer examiners probability of approving applications that were also approved by the EPO. By contrast, Model 2 shows that more experienced examiners are more likely to grant applications that were rejected by the EPO.

We are reluctant to treat EPO decisions as a gold standard to which U.S. decisions should be compared, given that the EPO too faces concerns about lax patent quality. But it is interesting the experience effect does not reflect younger examiners being more likely to reject applications patented at the EPO, but instead that older examiners are more likely to patent applications rejected by the EPO. And our finding that the experience effect persists when we control for EPO decisions is further reason to believe there is not a selection

¹²This number differs somewhat from that in Jensen et al. (2006), who show that 63 percent of granted applications in the U.S. are also granted by the EPO in the most recent cohort for which they have data (priority year 1995). But this figure is trending downward over time (from 1990 to 1995) in their data, and, based on our data from 2001, appears to have continued to do so.

effect at work.

7 Other explanations for the experience effect

The analyses in the previous section suggest that the experience effect does not reflect sorting of applications across examiners. What then, is its source?

One possible explanation is selective retention. As we pointed out above, the PTO faces significant employee attrition, particularly among examiners who have been with the agency less than 5 years (GAO 2008). If examiners who were more diligent, more thorough, more technically sophisticated, or more highly educated were more likely to leave the USPTO earlier in their careers, perhaps because they have better job opportunities, this could provide one explanation for our results. Under this interpretation, the experience effect would reflect different features of examiners who stay versus leave, not anything about changing incentives or capabilities facing a given examiner over her career at the USPTO.

Since we lack any additional demographic information about our examiners, this is difficult to assess directly. Moreover, since we only have a cross-section of applications at a given point in time, we cannot examine changes in an examiner's grant behavior over her career.

However, the USPTO Employee Directories do give us information on when and if the examiners in our January 2001 sample left the agency, assuming this happened before 2007. Using these data, we determined that 20.1 percent of the examiners in our sample had left by the end of 2006. Attrition is more pronounced for the newest examiners (consistent with the figures in GAO 2008) with 27.1 percent of the examiners hired in 2001 having left by the end of 2006.

To assess whether examiners who would exit are different from others, we estimated models relating the two measures of search intensity and grant rates to a dummy variable indicating whether the examiner would leave the USPTO by 2006. Table 10 shows the results. Models 1 and 2 respectively regress an examiner's average share of patent references and non-patent references on whether the examiner would exit, for examiners hired in 2001. Model 3, estimated at the patent application level, relates whether an application is granted to whether the examiner would exit, again only for examiners hired in 2001. In none of these models do examiners who would exit the USPTO within five years look different from the others hired at the same time. Models

4 through 6 are similar, but include all examiners (not just those hired in 2001), and indicator variables for each of the experience categories (not reported). The estimated coefficients show that those who would leave have indeed have significantly higher shares of patent and non--patent citations (Columns 4 and 5), but that their grant rates are similar to those who did not leave (Model 6). These results suggest that selective retention is unlikely to be the source of the experience effect we documented above.

Another possibility is examiner tenure. After promotion, examiners are not subject to the same level of scrutiny. Among other things, with full signatory authority they can sign off on their own applications without review. This could plausibly cause them to be more lax.

It is difficult to test for a tenure effect directly. A first problem is data: we lack direct information about whether examiners are tenured or untenured. But we can determine which examiners were definitely tenured by 2001, using data from issued patents. Specifically, we can infer an examiner's probable tenure year as the first year after which she was listed as primary examiner on all the patents she issued that were applied for in that year. In other words, we determined the first application year after which an examiner consistently had signatory authority.

Not surprisingly, the likelihood an examiner is tenured increases sharply with experience, with 9.6 percent of examiners with 2--4 years of experience tenured by 2001, 41.6 percent of examiners with 5--7 years of experience, and 90.7 percent of examiners with 8 or more years experience. The strong relationship between tenure and experience also make it difficult to separately identify a tenure effect separate from an experience effect. More importantly tenure is non--random: for a given level of experience, those that are tenured are likely to be systematically different from those who are not.

While we cannot rule out a tenure effect, our finding that there are differences in prior art searching tendencies, grant probabilities, and likelihood of granting on first office action even between the 2--4 year cohort (the bulk of whom are not tenured) and brand new examiners suggest that it alone is unlikely to drive the experience effect.

Another potential explanation is the quota system, which imposes increasing demands on examiners, and decreased time per count, as they accumulate experience at the USPTO. It may be that more experienced examiners are doing less prior art searching and granting more patents because they are required to process more applications in the same time frame. Still another possible explanation is that more senior examiners are less familiar with current

trends in their field, and therefore more likely to find applications patentable. Whether the experience effect documented in the main models reflects the quota system, technological obsolescence, or tenure, it suggests that human resource policies at the USPTO could affect outcomes, as we discuss in more detail below.

8 Discussion

Consistent with the qualitative and quantitative results of Cockburn et al. (2003), our data show considerable examiner--level heterogeneity both within and across art units. Moreover, this variation is related to the most important decision made by the USPTO: whether or not to grant a patent. In particular, we find that more experienced examiners are significantly more likely to grant, and, conditional on experience, examiners that conduct more intensive prior art searches are least likely to grant.

Taken alone, the result that more senior examiners are more likely to grant could suggest that they can more quickly figure out what is patentable in an application. But our data on prior art citation patterns do not support that conclusion. The finding that more senior examiners systematically cite less prior art reinforces the inference that senior examiners are doing less work, rather than that they are merely getting it right more often than junior examiners.¹³ And the fact that seniority is correlated with more first--action allowances is also inconsistent with the idea that more experienced examiners are simply negotiating the applicant to a narrower, patentable outcome; in the first--action allowance cases there is no negotiation at all.

We show that these results are unlikely to reflect either sorting of applications across examiners or selective retention of examiners. But our data cannot separately distinguish between a technological obsolescence effect, an experience effect driven by the count/quota requirements over an examiner's career, or a tenure effect, each of which would have different policy implications. But under any of these explanations, our results suggest that human resource policies have important effects on USPTO outcomes. The tenure system, the count system, and examiner recruitment and retention policies should

¹³ Similarly, while one might have sought to explain the lower citation patterns by more experienced examiners as greater parsimony learned from experience, the fact that those reduced citations are accompanied by a greater propensity to grant patents undermines that explanation.

be a more prominent part of current patent reform deliberations.

Our findings may also have significant implications for the industry--specific results we discussed in our prior paper (Lemley and Sampat 2008). While there is no question that there are industry--specific differences throughout patent prosecution, some of the differences we identified in the prior article may turn out to owe their origin to differences in examiners. In this study, we determined that the computer industry had by far the highest percentage of new examiners: more than 60 percent of examiners in that art unit had less than a year of experience, compared with less than 20 percent in mechanics and chemistry. In our prior paper, we found that the computer industry had a surprisingly low grant rate: lower than any other industry. At least some of that result may be explained by the prevalence of new examiners in that industry. If it is generalizable, this result may have another effect: booms in patenting in new industries may be self--limiting. The more applications are filed in an art unit, the more new examiners the art unit will have to hire. And because new examiners are more likely to reject patents, this will drive down the grant rate in that art unit, limiting the number of patents that result.

We also show, we believe for the first time in a large sample analysis, that identification of prior art matters for patent office outcomes. Examiners that tend to identify more prior art reject more applications. While the welfare implications of these findings depend on the costs of identifying additional prior art, and ones beliefs on whether the PTO currently makes too many Type I or Type II errors, our findings provide support for the feasibility of current initiatives (e.g. the Peer to Patent initiative, or post--grant opposition) aimed at affecting the grant rate by bringing more prior art to the attention of examiners. Indeed, they may also support more dramatic proposals, such as the idea (currently under consideration in Congress) that applicants should be forced to conduct a diligent search for prior art when they file a patent application.¹⁴

Finally, our data suggest that whether the PTO grants or rejects a patent is significantly related to the happenstance of which examiner is assigned the application. That is not an encouraging result if our goal is a system that rewards deserving patent applicants while denying patents to the rest. Nor

¹⁴ This policy proposal presupposes that applicant searches and examiner searches are substitutes. That may not be so. If it is the act of searching that engages an examiner with the application, increasing the rigor of the examination process, outsourcing search to the applicant, to another patent office, or to a third party may not produce the same result. Our data do not provide evidence on this issue.

is it clear that, in this environment, all issues patents should enjoy the same presumption of validity.

9 Sources

Alcacer, J. and M. Gittleman (2006). "How do I know what you know? Patent examiners and the generation of patent citations." *Review of Economics and Statistics* 84(4): 774-779.

Alcacer, Gittleman, et al. (2009). "Applicant and Examiner citations to patents; an exploration and overview" Forthcoming, *Research Policy*.

Allison, J. R. and M. A. Lemley (2000). "Who's Patenting What? An Empirical Exploration of Patent Prosecution" *Vanderbilt Law Review* 53(2099).

Cockburn, I., S. Kortum, et al. (2003). "Are All Patent Examiners Equal? Examiners, Patent Characteristics, and Litigation Outcomes" *Patents in the Knowledge-Based Economy*: 46-47.

Crouch, D. (2008). "Extending the Patent Term: Most Patents are Extended Due to PTO Delay," <http://www.patentlyo.com/patent/2008/03/extending-the-p.html>.

Federal Trade Commission. (2003) *To promote innovation: The proper balance of competition and patent law and policy*.

Jaffe, A. B. and J. Lerner (2004). *Innovation and Its Discontents: How our broken patent system is endangering innovation and progress, and what to do about it*. Princeton, Princeton University Press.

Jensen, P. H., A. Palangkaraya, et al. (2006). "Disharmony in international patent office decisions." *Federal Circuit Bar Journal* 15 (4) 679(704).

Lanjouw and Schankerman (2001). "Characteristics of Patent Litigation: a window on competition." *RAND Journal of Economics* 32(1): 129-151.

Lemley, M. A. (2001). "Rational Ignorance at the Patent Office." *Northwestern University Law Review* 95(1495).

Lemley, M. A. and K. A. Moore (2004). "Ending Abuse of Patent Continuations." *Boston University Law Review* 82(77).

Lemley, M.A. and Carl Shapiro (2005) "Probabilistic Patents" *Journal of Economic Perspectives* 19(2), 2005.

Lemley, M. A. and B. Sampat (2009). "Is the Patent Office a Rubber Stamp." *Emory Law Journal* forthcoming

Lichtman, D. (2004). "Rethinking Prosecution History Estoppel" *University of Chicago Law Review* 71(151).

A. A. Lipscomb and A. E. Bergh, editors. *The Writings of Thomas Jefferson. The Writings of Thomas Jefferson*. Thomas Jefferson Memorial Association, Washington.

Merges (2001). "As Many as Six Impossible Patents Before Breakfast: Property Rights for Business Concepts and Patent System Reform" *Berkeley Technology Law Journal* 14: 577-615.

Moore (2003). "Xenophobia in American Courts." *Northwestern University Law Review* 97(1497).

Nordhaus, W. (1969) *Invention, Growth, and Welfare: A Theoretical Treatment of Technological Change*. Cambridge, MA.

Patel, R. P., D. R. Brownstone, et al. (2003). "Understanding After Final and After Allowance Patent Practice."

Sampat (2009). "When do applicants search for prior art?" Mimeo, Columbia University.

Thomas (2001). "Collusion and Collective Action in the Patent System: A Proposal for Patent Bounties." *University of Illinois Law Review* 2001(1): 305-353.

Thompson, Peter. (2006) "Patent Citations and the Geography of Knowledge Spillovers: Evidence from Inventor and Applicant Added Citations" *Review of Economics and Statistics*. Vol. 88, No. 2: 383-388.

U.S. General Accountability Office. (2005) "USPTO Has Made Progress in Hiring Examiners, but Challenges to Retention Remain"

Webster, E., A. Palangkaraya, et al. (2007). "Characteristics of international patent application outcomes." *Economic Letters* 95 362(368).

Figure 1: Examiner Experience, by Art Unit




Figure 2: Distribution of Examiner's Share of Patent Cites, by Art Unit




Figure 3: Distribution of Examiner's Share of Non-Patent Cites, by Art Unit

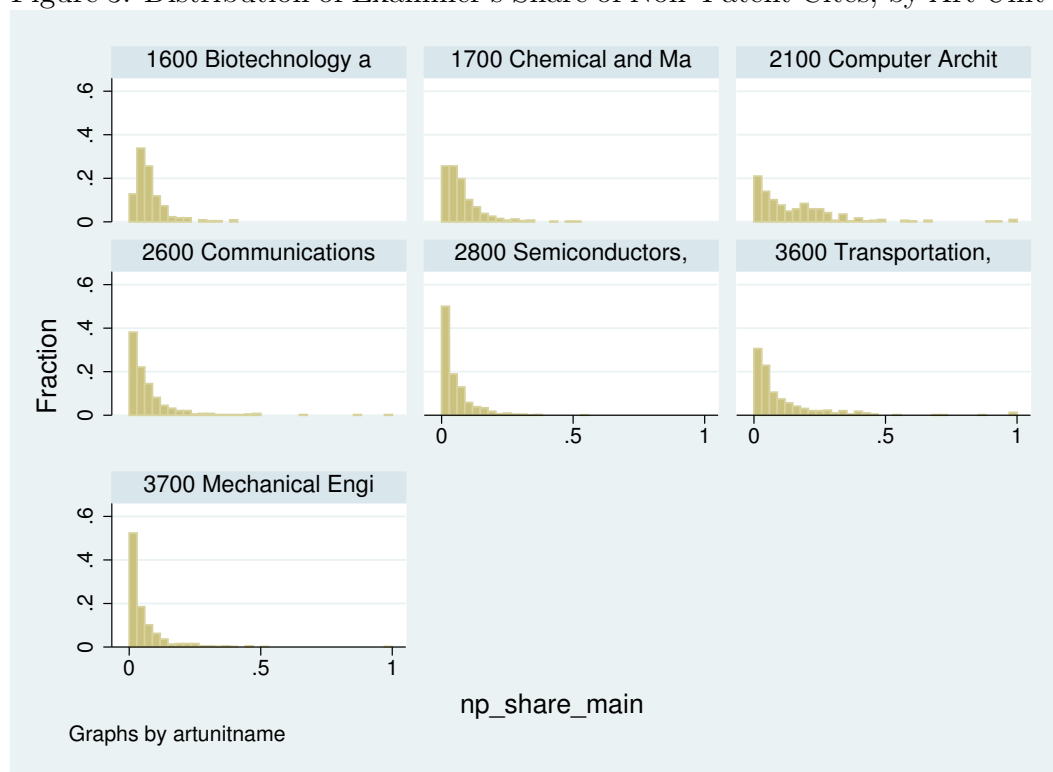


Figure 4: Claim Changes Versus Rejections




Table 1: Summary Statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
0-1 Yrs Experience	0.358	0.479	0	1	2797
2-4 Yrs Experience	0.274	0.446	0	1	2797
5-7 Yrs Experience	0.095	0.293	0	1	2797
8+ Yrs Experience	0.274	0.446	0	1	2797
count0106_main	218.679	155.471	1	842	2708
us_share_main	0.436	0.188	0.029	1	2708
np_share_main	0.081	0.112	0	1	2670

Table 2: Linear probability model: search intensity measures vs. examiner experience

	USShare	NPSHare
	(1)	(2)
2-4 Yrs Experience	-.035*** (.007)	-.005 (.004)
5-7 Yrs Experience	-.076*** (.009)	-.004 (.005)
8+ Yrs Experience	-.102*** (.007)	-.019*** (.004)
Obs.	2533	2533

*Notes: All models include 301 art-unit fixed effects; Robust standard errors, are reported in parentheses. * denotes significance at the 95 percent level; ** at the 99 percent level; *** at the 99.9 percent level*

Table 3: Linear probability model: whether application granted vs. experience, search intensity measures

	Model1	Model2	Model3
	(1)	(2)	(3)
2-4 Yrs Experience	.058*** (.013)	.046*** (.013)	.049*** (.013)
5-7 Yrs Experience	.111*** (.018)	.094*** (.018)	.102*** (.018)
8+ Yrs Experience	.113*** (.013)	.094*** (.014)	.100*** (.013)
Avg. Share Pat Cites		-.126*** (.039)	
Avg. Share NonPat Cites			-.321*** (.074)
Const.	.659*** (.009)	.739*** (.021)	.702*** (.010)
Obs.	9743	9029	9029
R ²	.145	.129	.131

*Notes: All models include 301 art-unit fixed effects; Robust standard errors, clustered on examiners, are reported in parentheses. * denotes significance at the 95 percent level; ** at the 99 percent level; *** at the 99.9 percent level*

Table 4: Linear probability model: whether patented application had any rejections vs. experience, search intensity measures

	Model1	Model2	Model3
	(1)	(2)	(3)
2-4 Yrs Experience			
5-7 Yrs Experience			
8+ Yrs Experience			
Avg. Share Pat Cites		.045 (.051)	
Avg. Share NonPat Cites			-.024 (.090)
Const.	.701*** (.044)	.686*** (.047)	.702*** (.043)
Obs.	6427	6333	6285
R ²	.101	.104	.106

*Notes: All models include 301 art-unit fixed effects; Robust standard errors, clustered on examiners, are reported in parentheses. * denotes significance at the 95 percent level; ** at the 99 percent level; *** at the 99.9 percent level*

Table 5: OLS Models: Examiner Characteristics versus application characteristics

	Model1	Model2	Model3	Model4	Model5	Model6
	(1)	(2)	(3)	(4)	(5)	(6)
Pages	.002 (.003)		-.0001 (.0001)		.0001* (.00007)	
Family Size		.017 (.015)		-.0007 (.0005)		.0002 (.0003)
Const.	4.113*** (.083)	4.087*** (.085)	.442*** (.003)	.442*** (.003)	.066*** (.002)	.068*** (.002)
Obs.	9748	9748	9034	9034	9034	9034
R ²	.252	.252	.48	.48	.297	.297

*Notes: All models include 301 art-unit fixed effects; Robust standard errors, clustered on examiners, are reported in parentheses. * denotes significance at the 95 percent level; ** at the 99 percent level; *** at the 99.9 percent level*

Table 6: Linear probability model: Examiner Characteristics versus application characteristics, patented applications only

	Model1	Model2	Model3	Model4	Model5	Model6	Model7	Model8	Model9
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Applicant Patent Cites	.0008 (.003)			-.0002** (.0001)			.00002 (.00006)		
np-ap		.007 (.007)			-.0006** (.0003)			.0003 (.0002)	
applicant patent volume			7.76e-06 (6.59e-06)			2.44e-07 (2.25e-07)			1.18e-07 (1.25e-07)
Const.	4.652*** (.081)	4.639*** (.081)	4.627*** (.083)	.433*** (.003)	.432*** (.003)	.432*** (.003)	.061*** (.002)	.061*** (.002)	.061*** (.002)
Obs.	6528	6528	6532	6256	6256	6259	6256	6256	6259
R ²	.246	.246	.246	.482	.482	.482	.296	.296	.296

*Notes: All models include 301 art-unit fixed effects; Robust standard errors, clustered on examiners, are reported in parentheses. * denotes significance at the 95 percent level; ** at the 99 percent level; *** at the 99.9 percent level*

Table 7: US Status by EPO Status

US Status	EPO Status			
	Rejected	Patented	Pending	Total
Rejected	346	126	77	549
Patented	526	1,118	501	2,145
Pending	14	26	27	67
Total	886	1,270	605	2,761

Source:

Table 8: Linear probability model: whether application granted vs. experience, search measures, EPO status

	Model1	Model2	Model3
	(1)	(2)	(3)
2-4 Yrs Experience	.038* (.021)	.022 (.022)	.025 (.022)
5-7 Yrs Experience	.085*** (.028)	.068** (.029)	.079*** (.029)
8+ Yrs Experience	.073*** (.021)	.051** (.021)	.061*** (.022)
Rejected at EPO	-.268*** (.019)	-.263*** (.020)	-.263*** (.020)
Pending at EPO	-.007 (.018)	-.009 (.019)	-.010 (.019)
Avg. Share Pat Cites		-.156** (.064)	
Avg. Share NonPat Cites			-.374*** (.127)
Const.	.842*** (.016)	.924*** (.032)	.879*** (.019)
Obs.	2694	2544	2544
R ²	.29	.281	.283

*Notes: All models include 301 art-unit fixed effects; Robust standard errors, clustered on examiners, are reported in parentheses. * denotes significance at the 95 percent level; ** at the 99 percent level; *** at the 99.9 percent level*

Table 9: Linear probability model: whether application granted vs. experience, search measures, by EPO status

	Model1	Model2	Model3	Model4	Model5	Model6
	(1)	(2)	(3)	(4)	(5)	(6)
2-4 Yrs Experience	-.002 (.027)	-.023 (.027)	-.018 (.028)	.047 (.053)	.035 (.056)	.043 (.056)
5-7 Yrs Experience	.014 (.032)	-.007 (.035)	.006 (.035)	.181** (.071)	.170** (.073)	.180** (.073)
8+ Yrs Experience	.029 (.028)	.008 (.028)	.020 (.028)	.109** (.054)	.069 (.060)	.088 (.058)
Avg. Share Pat Cites		-.159* (.088)			-.244 (.165)	
Avg. Share NonPat Cites			-.147 (.171)			-.701** (.355)
Const.	.888*** (.019)	.970*** (.041)	.907*** (.023)	.538*** (.036)	.665*** (.085)	.602*** (.048)
Obs.	1244	1191	1191	872	811	811

*Notes: All models include 301 art-unit fixed effects; Robust standard errors, clustered on examiners, are reported in parentheses. * denotes significance at the 95 percent level; ** at the 99 percent level; *** at the 99.9 percent level*

Table 10: Linear probability model: Examiner Characteristics by Whether They Will Leave in 5 Years

	Model1	Model2	Model3	Model4	Model5	Model6
	(1)	(2)	(3)	(4)	(5)	(6)
Examiner Will Leave	.038 (.046)	.023 (.028)	.057 (.060)	.018** (.007)	.010* (.005)	-.004 (.013)
Const.	.447*** (.016)	.087*** (.011)	.659*** (.018)	.481*** (.006)	.095*** (.005)	.660*** (.009)
Obs.	238	237	778	2708	2670	9743
R^2	.742	.658	.327	.521	.325	.145