Examiner Characteristics and the Patent Grant Rate*

Mark Lemley †
Bhaven N. Sampat‡

July 13, 2008

1 Introduction

What determines whether a patent applicant can obtain a patent? In prior work, we studied patent grant rates and determined that, while the U.S. Patent and Trademark Office (PTO) grants patents on a large majority of the applications it receives, it appears to block patenting of a small but non–trivial number of applications. (Lemley and Sampat, 2008). Our conclusion in that paper was that the PTO was not a "rubber stamp." But without knowing how many of those applications deserved a patent in some absolute sense, we could not determine whether the relatively high PTO grant rate was too high, too low, or just right.

In this paper, we use the same data set to examine a related question: What causes certain applications to be allowed and others to be rejected? In particular, we explore the effects of patent examiner characteristics on

*We thank Ernie Beffel, John Golden, Iain Cockburn, Matt Neidell, Sherry Glied, Josh Graff Zivin and participants at the NBER Productivity Lunch for comments on the project or a prior draft, and Sarah Craven, Reuven Garrett, and Matthew Wasserman for research assistance. Patrick Crosby at xblabs.com helped obtain some the data used in this analysis. Sampat thanks the Ford Foundation for financial support.

†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

the probability an application is granted. We find, among other things, 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 the l that 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.

Only a few 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 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.[1]

[1]Unfortunately, changes to the PAIR interface make it difficult to update the prosecu-

As reported in Lemley and Sampat (2008), 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 examiner's rejection, but could also be abandoned for other reasons (e.g. the applicant's firm went out of business).[2] The remaining 2.2 percent of the applications were pending as of June 2008.

In this paper, we focus on the set of these application that have received a 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. One concern about our analysis is selection: if specific types of applications are assigned to specific examiners, this could confound our results. Section 6 examines this possibility, and suggests that our results are not driven by selection bias. Section 7 concludes with a discussion of policy implications.

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

---

tion data beyond April 2006.

[2]There 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.

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 (2007) argues 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. 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 of prior art 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 examiner's 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.[3]

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 2007; Alcacer, Gittleman, and Sampat 2008).

The large number of applications facing the PTO means that examiners are subject to significant 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 (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 currently the subject of considerable policy debate (Jaffe and Lerner 2004; Lemley and Moore 2004). Of direct relevance to our analysis below, note that junior examiners face more scrutiny than more senior examiners. Their 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 until they are promoted and themselves obtain signatory authority, or the right to sign off on an application independently.

3 Examiner Characteristics

The main problem we confronted in evaluating the effects of examiners on patentability was determining who the examiners were. Since the PTO employs its examiners, examiner names are listed in PAIR, and primary and secondary 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 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

[3] They may also file one of a variety of "continuation" applications to continue to argue for patentability.

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.[4]

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. Linking information in these directories to the PAIR data allowed us to determine the experience as of 2001 (right truncated at 9 years) for each of the 2,823 examiners who were assigned the January 2001 applications.[5]

We also 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 1976--2006 period. Collectively, the 2,803 examiners in our sample were listed on issued patents under 13,744 name variants. To gauge the accuracy of this match, we compared the experience measure derived from the examiner roster to an experience measure based on the front page data (the application date of the first patent he or she examined and issued), and found the two to be highly correlated (r=.93).[6]

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. Using front-page data from issued patents, we constructed several additional examiner-specific variables. First, using data

[4] *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 Punnathapu, Achutamurtry Ponnathapu, Achutamurty Ponnathapu, Achuthamurthy Ponnathapu, Achutmurthy Ponnathapu, Achutyamurthy Ponnathapu, Murthy Ponnathapu, Murthy Ponnathapu Achuta, Murthy Ponnathapuachut, and Ponnathapuachuta.

[5] We also used an outside service to determine the gender of examiners listed on the roster. However, gender was unrelated to any outcome of interest, and was dropped from the analyses reported below.

[6] The correlation wouldn't be perfect even if our matching were, since examiners may not examine a patent immediately after hire, and our experience measure based on PAIR data is right censored at 9 years.

on citations in all patents issued by an examiner over the 2001-2006 period, we contructed 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 2007; Alcacer, Gittleman, Sampat 2008). Here, we use the average examiner share of references in patents he/she examined as proxies for an examiner's propensity to search for 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, we also control for the total number of patents issued by an examiner over this period, and in other models, estimate regressions excluding applications where the examiner had greater than ten patents issued over the 2001-6 period.[7]

4 Descriptive statistics on examiner characteristics

Table 1 shows descriptive statistics on each of the independent variables, calculated at the examiner level.[8] On average, the examiners were employed by the USPTO for 3.8 years, though this measure is a lower bound since our experience measure is truncated at 9 years. (Median experience is 3 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 (2007), 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.

[7] In current work we are also exploring another indicator of an examiner's intensity of search, the extent to which he or she tends to cite unique prior art for each application, as opposed to "favorite" or "pet" patents cited against all applications (Cockburn et al., 2003; Breitzman and Thomas 2005).

[8] Note that a small number of examiners issued no patents over the 2001-06 period, meaning we could not calculate the front-page-based measures for them. And even among examiners who did issue patents, some cited no non-patent prior art.

The relationships between the independent variables are also interesting. Table 2 shows the bivariate correlations between experience, propensity to search for patents (as proxied by the share of patent citations in an examiner's patents inserted by the examiner), and propensity to search for non--patented prior art. Each of the measures of search intensity are negatively related to experience, and positively related to one another. Put another way, examiners who are better at finding one type of prior art are better at finding other types as well. And the more experienced the examiner, the less detailed a search they seem to do. More experienced examiners have lower average shares of citations to both patented and non-patented prior art.

There is also considerable variation both within and across fields. Figure 1 shows experience by broad art units. In the computer industry, examiners are overwhelmingly new hires. There are also a large proportion of new hires in the communications industry. 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 in those industries. 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 nine 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 reveals significant inter-industry 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 difference across fields. In general, the share is higher in the computing and communication arenas; Sampat (2007) and Alcacer, Gittleman, and Sampat (2008) suggest this may reflect lower incentives for applicants in these fields to search for prior art. 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, Figure 4 shows the relationship between our two measures of search intensity and experience. Specifically, it shows the average of an examiners average citations to patented prior art and non-patented prior art, by cohorts of examiner experience. Strikingly, more experienced examiners appear to contribute a lower share of references to both patents and publications. This 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.

5 Results

5.1 Examiner characteristics and the grant rate

In our empirical analyses, we estimated linear probability models relating the probability a application filed in January 2001 is granted by June 2008 to examiner characteristics. Each of the models includes fixed effects for each of the 301 art units, and robust standard errors clustered on examiners.[9]

Table 3 shows our results. The first column shows results from a model where we include our experience measure non-parametrically, with dummy variables for each experience category. (The left out category is zero years.) The second column shows results from instead including our (right censored) experience variable. Strikingly, in both models, the probability an application results in a patent increases steadily with the experience of the examiner. In each model, the most experienced examiners (9 or more years of experience) have a grant rate that is about 12 percentage points greater than that for the least experienced, after controlling for art unit effects.

The descriptive statistics above showed that more experienced examiners conduct less thorough prior art searches, which could explain these results. In Table 4, we show the effects of our first citation-based measure, an examiner's propensity to search for patented prior art (as proxied by the share of citations he/she accounts for in patents issued in 2001-06). The first column shows results based on the full sample of applications. In that model,

[9] Results from probit models are qualitatively similar, and available on request from the authors.

we control for the total count of 2001-06 patents the examiner issued. The second column shows results for the subset of applications associated with examiners who issued more than 10 patents between 2001 and 2006, and for whom the "share of patent citations" variable is more precisely measured.

In each of the models, experience continues to have a positive and statistically significant effect on the probability is granted. In the full sample (Model 1) the effect is about half as large as that from the models without the patent citation variable, suggesting that the experience results from that model may in part have been reflecting differences between experienced and non-experienced examiners in prior art search tendencies. Moreover, the direct effect of our measure of how deeply an examiner searches for previous patents—her average share of citations to previous U.S. patents calculated across all over her 2001-06 patents—is negative and statistically significant in both models. Rough conterfactual simulations suggest that the effects are also qualitatively significant: Model 1 predicts that 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 5.5 percentage points; Model 2 predicts a 7.3 percent difference.

Table 5 shows that the effect is also negative and statistically significant, and qualitatively larger, when using an examiner's share of references to non-patent prior art as a measure of examiner search intensity. (The effects of the experience measure are qualitatively similar to those from the previous model.) Here, the difference between 5th and 95th percentiles of examiners (accounting for 0 percent and 33 percent of references to non--patent prior art, respectively) maps to a 8.6 percentage point change in the grant rate in the full model, and a 10 percentage point difference using estimates from the restricted sample.

Table 6 shows results from models simultaneously introducing each of the two citation measures. The effects of the average non-patent citation share is nearly four times larger in magnitude than the effects of the patent citation share in each of the models. In both models, we can reject the hypothesis that the coefficients on of the two citation measures are equal at the 1 percent level. This may reflect that searching for non-prior art is a stronger signal of an examiner's search intensity, i.e. that it is easier to identify and cite patented prior art, making this a noisier signal of an examiner's thoroughness than propensity to cite non-patent prior art.

Taken together, the 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 inexperienced 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. Examiners who tend to cite more non-patent prior, on average, are much less likely to grant an application.

The heterogeneity across patent examiners illustrated in our descriptive statistics relates thus 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 certainly suggestive that examiners are doing more work, and rejecting applications with more rigor, at early stages in their career, and doing less work 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 narrowing 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 in the following section.

5.2 Examiner characteristics and rejections during patent prosecution

To measure whether a patent's scope was narrowed during the patent prosecution process, we would ideally look carefully at the final claims in the patent and compare to claims in the application. 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.

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.

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. Table Seven shows he about 81 percent of these granted applications received a non-final rejection, and 26 percent a final rejection, during patent prosecution. In Table 8, we relate these outcomes to experience. Models 1 and 2 show the effects of the continuous experience measure on the probability that a granted patent had a non-final rejection, or a final rejection, respectively. In each, an additional year of examiner experience is associated with a 1.4 percentage point reduction in the probability of rejection. Models 3 and 4 show results from models with experience dummies, and show similar results. After controlling for art unit effects, the most experienced examiners have a 13 percentage point lower probability than the least experienced examiners of issuing a patent with non-final rejections, and a 14 percentage point lower probability of final rejection. Table 9 shows results from models including the examiner's average share of patent citations, for the whole sample (Models 1 and 2) and for the subset associated with examiners who issued more than 10 patents (Models 3 and 4). Interestingly, conditional on experience, this measure of examiner search propensity is not significantly related to the probability that an application had non-final rejections or final rejections in the full sample. But in the restricted sample, where the search intensity proxy is measured more precisely, examiner propensity to cite patented prior art is positively and significantly related to the probability that an issued patent had rejections, and thus was likely to have had amendments before issued.

Table 10 shows analogous results for the non-patent citation measure of search intensity. Here, there is no effect on the probability of non-final rejections (which, as Table 7 showed, are quite common), but a positive and significant effect on the probability of final rejections in both the full and restricted sample. Moreover, the magnitude of the effect of non-patent citation propensity on final rejections is substantially larger than that that for patent citation propensity, nearly twice as large in Model 4. This is again consistent with the notion that the extent to which examiner's search and cite patented prior art is less related to their rigor or thoroughness than the extent to which they search for and cite non--patent prior art.

The fact that experience reduces the likelihood of both non-final and final rejections tends to reject a possible explanation for the relationship between examination and grant: 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. Instead, the data seem to be consistent with what we refer to below as the "tenure effect" – examiners begin by devoting more time and energy to examination, but after years in the office they simply don't work as hard.

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 policy 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 examiner's 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. 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 judgment (ideally, by a person having ordinary skill in the art) about whether they are novel and non-obvious in light of the prior art.[10]

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 application level (for applications that were granted or not): 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 11 shows the effects of these variables on our three examiner characteristics: years of experience, the examiner's average share of citations to patents, and the examiner's average share of citations to non--patent references. In none of the models do these application level variables have qualitatively or statistically significant relationships to the examiner characteristics, consistent with our impression that there is no application level sorting at the USPTO.

To examine this further, we also looked at application level characteristics

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

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 12 shows results from regressing these variables on the examiner characteristics. Models 1 through 3 show that none of these measures have a qualitatively or statistically significant relationship with examiner experience. Models 4 and 5 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.

In any event, while statistically significant these effects are qualitatively small: essentially very precisely estimated zeroes. For example, increasing the applicant share of references to patents from its 5th to its 95th percentile (an increase from 0 to 20 applicant citations) would be associated with only a .6 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 .3 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 examiner's 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 this directly. In the next section, we examine a proxy variable, whether the same application was granted in by the EPO.

6.2 Evidence from EPO Outcomes

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. 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.

Of these U.S. applications, 2,694 were filed at the EPO, of which 43.5 percent had been granted as of June 2008. Table 13 shows a cross-tabulation of EPO and USPTO outcomes as of June 2008. Of the applications that were granted in the U.S., slightly less than half (48 percent) have been granted by the EPO.[11] By contrast, of those granted by the EPO, the vast majority (nearly 90 percent) were also granted by the USPTO.[12]

Table 14 relates each of our examiner characteristics to whether the application was applied for in the EPO (Models 1 to 3) and whether granted in the EPO, conditional on application (Models 4 to 6). Neither of these shows a significant relationship with examiner characteristics, supporting our argument that there is no sorting across examiners based on either perceived importance of the patent or likelihood of patentability.

In Table 15, we explicitly include the EPO patent status variable to our baseline models, i.e. control for this measure of the patent-worthiness of the

[11] This number differs somewhat 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.

[12] Note that we currently lack data on which of the 1522 applications that are unpatented in Europe were rejected or withdrawn, or are still pending. (Recall that by construction, our U.S. sample excludes the less than 2 percent of January 2001 filed applications that remained pending as of June 2008.) Based on a 5 percent sample of the applications filed but not patented in Europe, we determined, using data from the European Register, that the vast majority (73 percent) were rejected or withdrawn, rather than pending. Since we are primarily interested in controlling for the patent-worthiness of the application, it is unclear that we need to distinguish exactly why the application has not been patented in Europe. Nonetheless, for future work we plan to collect the detailed EPO status information for all of our applications that were filed but not patented in Europe.

invention. Models 1 and 2 replicate our previous findings for the subset of applications that have counterpart EPO applications, and show the results are qualitatively similar. Models 3 and 4 include the EPO patent status variable, which is both positive and highly significant in each. All else equal, if an application that is filed at the EPO has been granted by the EPO, the U.S. grant probability is 15.3 percentage points higher. But neither the direction, statistical significance, or relative magnitude of the experience or patent search propensity variables change, again suggesting that selection is not driving our main results.

7 Conclusions and Policy Implications

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.

We emphasize that our data cannot answer the question of what the "right" grant rate is at the PTO, and therefore whether (to take one example) junior or senior examiners, or those who identify more or less prior art, have it right. Nonetheless, we think it is potentially troubling that the examiner grant rate is so strongly related to examiner experience. One possible explanation for this result --- though one we cannot prove --- is that new examiners come into the PTO with a certain mindset about their job, and that they begin by rejecting a fairly large percentage of applications. As they become inculcated into the culture of the PTO, which as a whole grants patents to the large majority of applicants (Lemley and Sampat 2008), they relent somewhat and are more likely to approve applications. It may be that once they are promoted from secondary to primary examiner, they are subject to significantly less scrutiny, and what we might call the "tenure effect" takes hold. Since it is in most respects easier for examiners to allow patents than to reject them (Lemley and Moore, 2004), primary examiners can ease the burden of their job by granting rather than rejecting applications in doubtful cases. This hypothesis is also consistent with data that show that primary examiners issue patents more quickly than secondary examiners. (Crouch, 2008)

Some have suggested to us that this might simply mean that more senior examiners 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 fact 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.

One counterintuitive implication of our results is that the high turnover rate at the PTO, long lamented as a problem that reduces examination quality, may actually improve the quality (or at least the rigor) of examinations by ensuring that a relatively high percentage of examiners are new hires who are subject to more scrutiny and therefore do a more careful job of searching for prior art and evaluating patent applications. This could be a good or a bad thing, depending on whether one believes PTO examination today is too lenient, too strong, or just right. But the relationship between turnover and examination rigor is a surprising one with consequences for PTO hiring and promotion decisions. The PTO should take steps to eliminate the "tenure effect," either by engaging in closer scrutiny of senior examiners or by moving to a team-review system in which the individual predilections of a single examiner have less impact on patenting decisions.

Second, our findings may 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. For example, we reported the surprising result that the computer industry had the lowest grant rate of any industry. 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

[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.

paper, we found that the computer industry had a surprisingly low grant rate. 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. Exploring the relative importance of art units effects and examiner effects in patent outcomes is an important task for future research.

Third, we show--we believe for the first time in a large sample analysis--that identification of prior art—and, it appears, particularly non-patent prior art—matters for patent office outcomes. The welfare implications of these findings depend on the costs of identifying additional prior art, one's beliefs on whether the PTO currently makes too many Type I or Type II errors (rejecting patents that merit patent protection, or approving patents that do not, respectively), and on the social costs of each of these types of errors. But our findings do 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.

8 Bibliography

Alacacer, J. and M. Gittleman (2006). "How do I know what you know? Patent examiners and the generation of patent citations." *Review of Eco-*

[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 or to a third party may not produce the same result. Our data do not provide evidence on this issue.

nomics and Statistics 84(4): 774-779.

Alcacer, Gittleman, and Sampat. (2008). "Applicant and Examiner citations to patents; an exploration and overview.'' *Working paper.*

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

Breitzman and Thomas (2005). "The Pet Patent Theory - Myth or Reality?" *ipFrontline* 12 August 2005.

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." *Patently–O Patent Law Blog, posting on 14 March 2008*

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 B. Sampat (2008). "Is the Patent Office a Rubber Stamp." *Emory Law Journal*, forthcoming.

Litchtman, D. (2004). “Rethinking Prosecution History Estoppel.” *University of Chicago Law Review* 71(151).

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).

Patel, R. P., D. R. Brownstone, et al. (2003). “Understanding After Final and After Allowance Patent Practice.” Fundamentals of Patent Prosecution, Practising Law Institute, No. G0-01EV, 2003-2005.

Sampat, Bhaven (2007). “When do applicants search for prior art?” *Working paper.*

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

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

9 Figures and Tables

Figure 1: Examiner Experience, by Art Unit

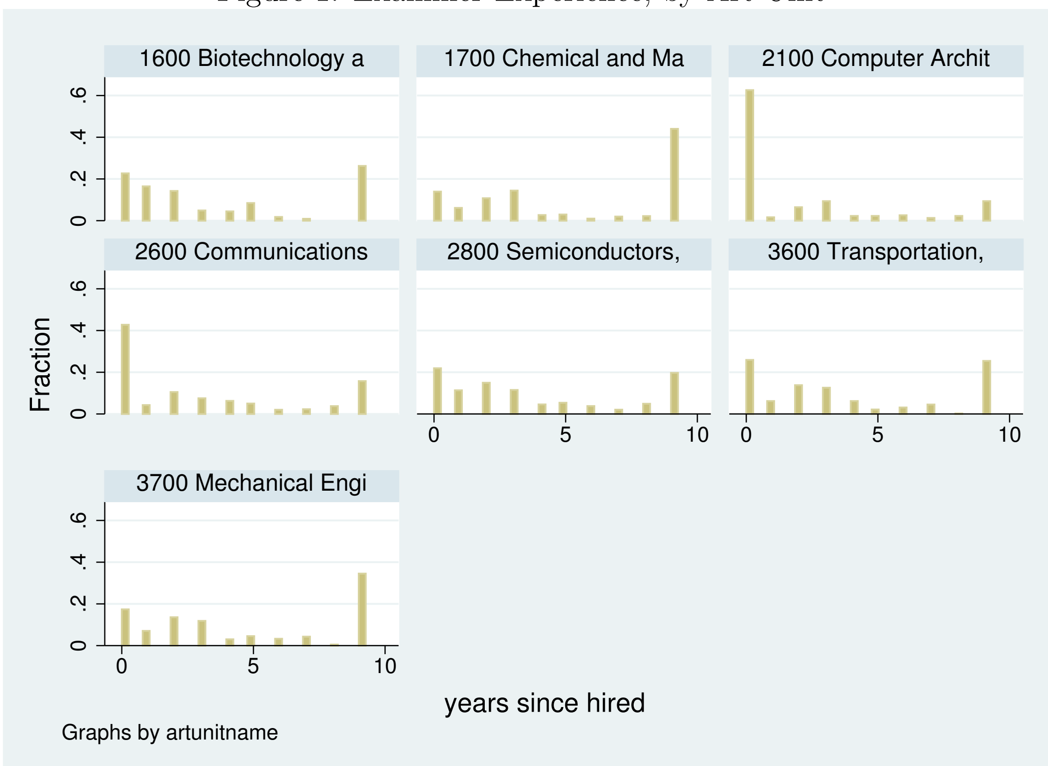

Figure 2: Distribution of Examiner's Share of Patent Cites, by Art Unit

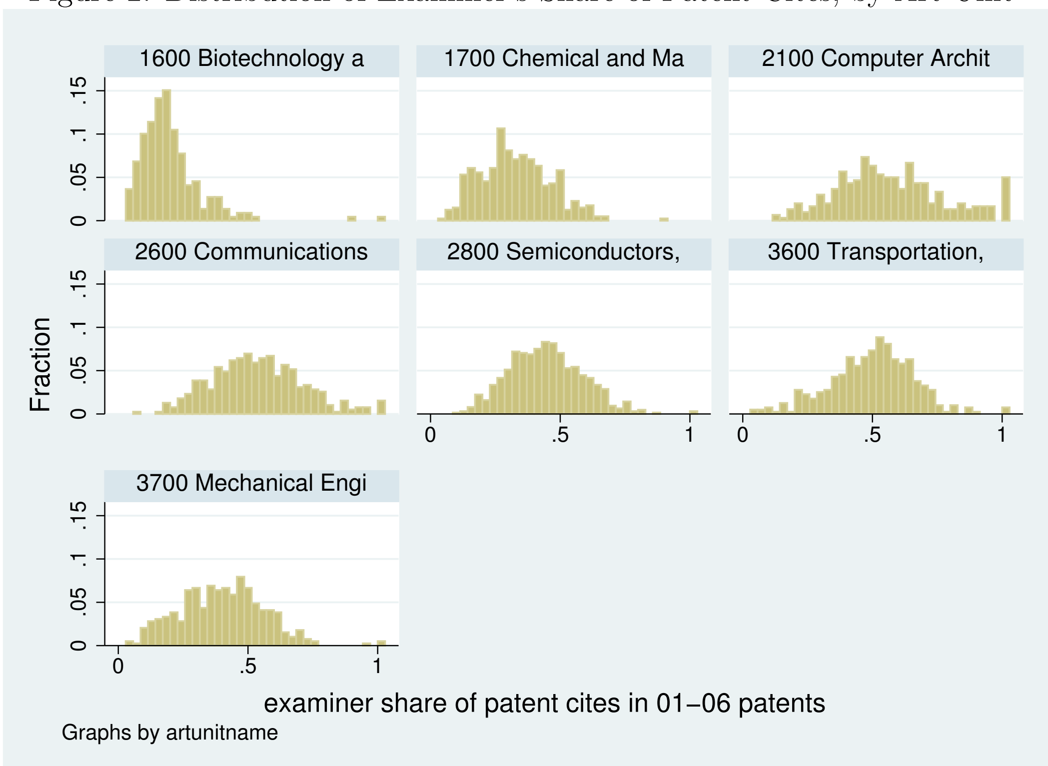

Figure 3: Distribution of Examiner's Share of Non–Patent Cites, by Art Unit

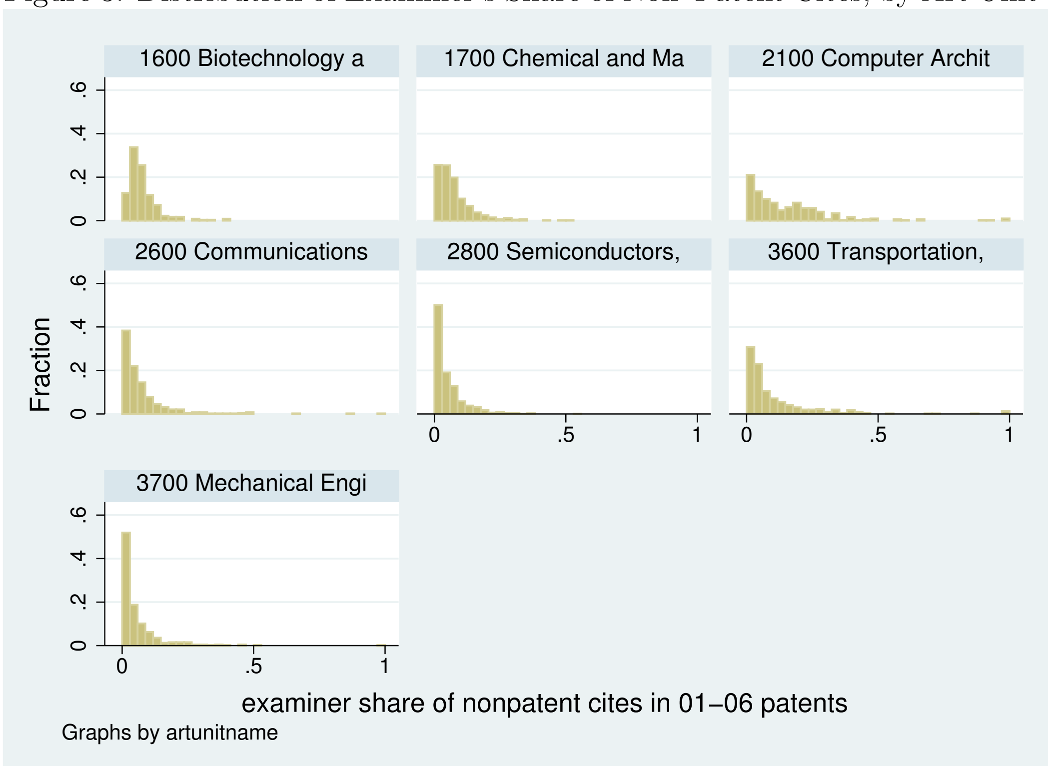

Figure 4: Examiner Propensity to Cite Prior Art, by Experience

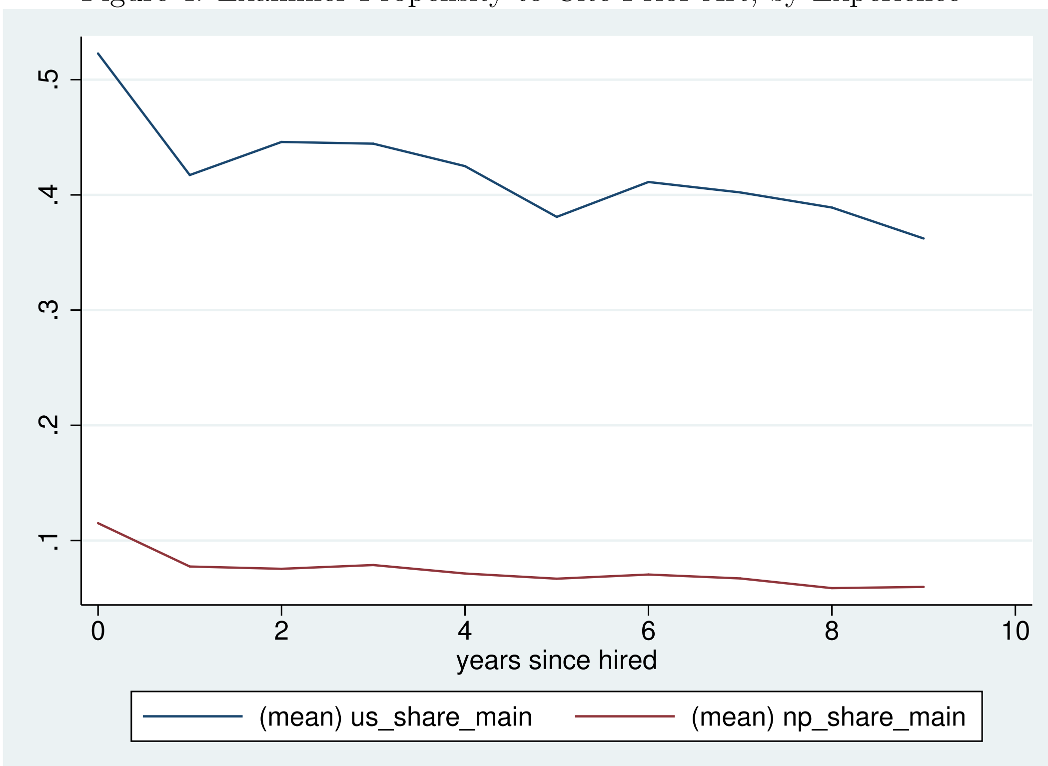

Figure 5: Claim Changes and Examiner Rejections

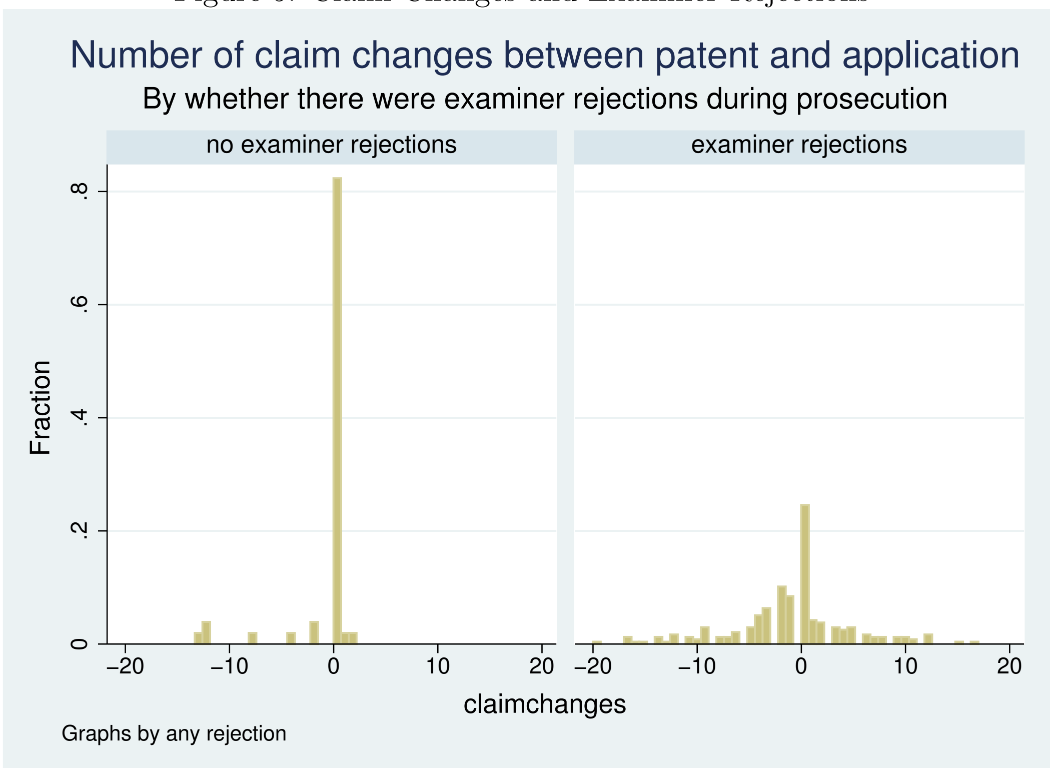

Table 1: Summary Statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
years since hired	3.807	3.574	0	9	2797
count of 01-06 patents	218.679	155.471	1	842	2708
examiner share of patent cites in 01-06 patents	0.436	0.188	0.029	1	2708
examiner share of nonpatent cites in 01-06 patents	0.081	0.112	0	1	2670

Table 2: Correlations Between Examiner Characteristics

Variables	years since hired	share patcites	share npcites
years since hired	1.000		
share patcites	-0.296	1.000	
	(0.000)		
share npcites	-0.164	0.286	1.000
	(0.000)	(0.000)	

Table 3: Linear probability model: whether application granted vs. examiner experience

	(1)	(2)
years since hired		.013***
		(.001)
1 year experience	.011	
	(.021)	
2 years experience	.055***	
	(.017)	
3 years experience	.068***	
	(.018)	
4 year experience	.071***	
	(.023)	
5 years experience	.102***	
	(.029)	
6 years experience	.106***	
	(.027)	
7 years experience	.149***	
	(.027)	
8 years experience	.098***	
	(.029)	
9+ years experience	.121***	
	(.015)	
Const.	.655***	.668***
	(.010)	(.007)
Obs.	9690	9690
R^2	.146	.145

*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 application granted vs. experience, patent citation propensity

	Full Sample	Restricted Sample
	(1)	(2)
years since hired	.006***	.011***
	(.002)	(.002)
count of 01-06 issued patents	.0003***	
	(.00004)	
share of patent citations in 01-06 patents	-.098***	-.123***
	(.035)	(.039)
Const.	.671***	.744***
	(.021)	(.021)
Obs.	9505	9029
R^2	.147	.129

*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: Linear probability model: whether application granted vs. experience, non-patent citation propensity

	Full Sample	Restricted Sample
	(1)	(2)
years since hired	.007***	.012***
	(.002)	(.001)
count of 01-06 issued patents	.0003***	
	(.00004)	
share of nonpat citations in 01-06 patents	-.259***	-.309***
	(.054)	(.074)
Const.	.648***	.709***
	(.012)	(.009)
Obs.	9383	9029
R^2	.148	.13

*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: whether application granted vs. experience, citation propensity

	Full Sample	Restricted Sample
	(1)	(2)
years since hired	.006***	.011***
	(.002)	(.002)
count of 01-06 issued patents	.0003***	
	(.00004)	
share of patent citations in 01-06 patents	-.069*	-.075*
	(.039)	(.040)
share of nonpat citations in 01-06 patents	-.231***	-.272***
	(.057)	(.077)
Const.	.681***	.742***
	(.022)	(.020)
Obs.	9383	9029
R^2	.149	.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 7: Patented applications with final or non–final rejections

Variable	Mean	Std. Dev.	Min.	Max.
any nonfinal rejection	0.805	0.396	0	1
any final rejection	0.256	0.437	0	1
N		6459		

Table 8: Linear probability model: whether patented application had rejections during prosecution, vs. examiner experience

	Nonfinal	Final	Nonfinal	Final
	(1)	(2)	(3)	(4)
years since hired	-.014***	-.014***		
	(.002)	(.002)		
1 year experience			-.021	-.045
			(.024)	(.028)
1 year experience			-.021	-.045
			(.024)	(.028)
2 years experience			-.044**	-.036
			(.020)	(.023)
3 years experience			-.086***	-.111***
			(.022)	(.025)
4 year experience			-.057	-.068*
			(.035)	(.037)
5 years experience			-.103***	-.142***
			(.029)	(.034)
6 years experience			-.098***	-.107***
			(.032)	(.037)
7 years experience			-.185***	-.136***
			(.045)	(.037)
8 years experience			-.179***	-.183***
			(.045)	(.041)
9+ years experience			-.132***	-.138***
			(.018)	(.021)
Const.	.872***	.322***	.886***	.343***
	(.009)	(.011)	(.013)	(.016)
Obs.	6423	6423	6423	6423
R^2	.1	.12	.102	.122

*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 patented application had rejections during prosecution, vs. examiner patent citation propensity

	Nonfinal	Final	Nonfinal	Final
	(1)	(2)	(3)	(4)
years since hired	-.013***	-.009***	-.015***	-.013***
	(.002)	(.002)	(.002)	(.002)
count of 01-06 issued patents	-.00008	-.0003***		
	(.00006)	(.00005)		
share patcites	.041	.068	.111**	.133**
	(.051)	(.053)	(.052)	(.056)
Const.	.870***	.344***	.827***	.260***
	(.030)	(.031)	(.027)	(.029)
Obs.	6329	6329	6189	6189
R^2	.103	.126	.109	.122

*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: whether patented application had rejections during prosecution, vs. examiner non-patent citation propensity

	Nonfinal	Final	Nonfinal	Final
	(1)	(2)	(3)	(4)
years since hired	-.014***	-.008***	-.016***	-.014***
	(.002)	(.002)	(.002)	(.002)
count of 01-06 issued patents	-.00009*	-.0003***		
	(.00006)	(.00005)		
share nonpat cites	-.055	.193**	-.027	.232**
	(.090)	(.093)	(.094)	(.097)
Const.	.899***	.368***	.882***	.308***
	(.016)	(.017)	(.011)	(.013)
Obs.	6281	6281	6189	6189
R^2	.105	.127	.108	.122

*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 11: OLS Models: Examiner Characteristics versus application characteristics

	Exp	ShrPatCit	ShrNpCit	Exp	ShrPatCit	ShrNpCit
	(1)	(2)	(3)	(4)	(5)	(6)
pages	.0002		-.00002		.0001	
	(.003)		(.0001)		(.0001)	
family size		.015		-.0008		.0002
		(.015)		(.0006)		(.0004)
Const.	4.052***	4.009***	.450***	.452***	.073***	.074***
	(.082)	(.085)	(.003)	(.004)	(.003)	(.002)
Obs.	9695	9695	9510	9510	9388	9388
R^2	.259	.259	.453	.453	.311	.31

*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 12: Linear probability model: Examiner Characteristics versus application characteristics, patented applications only

	Exp	ShrPatCit	ShrNpCit	Exp	ShrPatCit	ShrNpCit	Exp	ShrPatCit	ShrNpCit
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
applicant pat cites	-.001			-.0003**			-1.62e-07		
	(.003)			(.0001)			(.00006)		
applicant nonpat cites		.004			-.0005*			.0003	
		(.007)			(.0003)			(.0002)	
applicant patent volume			7.81e-06			-2.53e-07			1.77e-07
			(6.52e-06)			(2.34e-07)			(1.40e-07)
Const.	4.581***	4.570***	4.556***	.435***	.435***	.435***	.061***	.061***	.061***
	(.081)	(.081)	(.083)	(.003)	(.003)	(.003)	(.002)	(.002)	(.002)
Obs.	6493	6493	6497	6399	6399	6402	6349	6349	6352
R^2	.25	.25	.251	.465	.465	.465	.275	.276	.276

*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 13: **patented at USPTO by patented at EPO**

	patented at EPO		
patented at USPTO	No	Yes	Total
No	427	122	549
Yes	1,095	1,050	2,145
Total	1,522	1,172	2,694

Source:

Table 14: OLS Model: Examiner Characteristics vs. Whether Filed, Granted in EPO

	Exp	ShrPatCit	ShrNpCit	Exp	ShrPatCit	ShrNpCit
	(1)	(2)	(3)	(4)	(5)	(6)
application at EPO	.039	-.004	.003	.116	.006	.002
	(.079)	(.003)	(.002)	(.142)	(.005)	(.003)
patented at EPO				.116	.006	.002
				(.142)	(.005)	(.003)
Const.	4.013***	.453***	.074***	4.183***	.421***	.071***
	(.041)	(.002)	(.001)	(.088)	(.003)	(.002)
Obs.	8859	8688	8578	2681	2644	2623
R^2	.259	.446	.311	.284	.524	.369

*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 15: Linear probability model: whether application granted vs. experience, propensity to cite prior art, whether patented in EPO

	(1)	(2)	(3)	(4)
years since hired	.006**	.007***	.005**	.007***
	(.002)	(.002)	(.002)	(.002)
share of patent citations in 01-06 patents	-.162**		-.173***	
	(.068)		(.067)	
share of nonpat citations in 01-06 patents		-.396***		-.414***
		(.138)		(.134)
patented at EPO			.153***	.153***
			(.016)	(.016)
Const.	.845***	.799***	.785***	.735***
	(.033)	(.016)	(.033)	(.018)
Obs.	2544	2544	2544	2544
R^2	.202	.204	.232	.234

*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*