EXAMINER CHARACTERISTICS AND PATENT OFFICE OUTCOMES

Mark A. Lemley and Bhaven Sampat*

Abstract—In this paper, we show that there are important differences across patent examiners at the U.S. Patent and Trademark Office. We show that more experienced examiners cite less prior art, are more likely to grant patents, and are more likely to grant patents without any rejections. These results suggest that the most important decisions made by the patent office are significantly affected by the happenstance of which examiner gets an application. They also point to human resource policies as potentially important levers, hitherto neglected, in patent system reform.

I.   Introduction

UNDER the standard economic rationale, patents provide incentives to conduct research and 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 & Shapiro, 2005). In the United States, examiners at the Patent and Trademark Office (PTO) are charged with serving as gatekeepers to ensure that only novel and nonobvious inventions are granted patents.[1]

In this paper, using information from several novel data sets, we show that there are important differences across patent examiners and that these relate to the most important decisions made by the PTO. In particular, we find that more experienced examiners are less active in searching for prior art and have a much higher grant rate.

The results could be problematic 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 PTO could affect patent grant rates, an important finding amid growing concern that it grants too many low-quality patents and is in need of significant

reform (Jaffe & Lerner, 2004; Cohen & Merrill, 2003; FTC, 2003).

Several scholars have previously studied examiner heterogeneity and its effects on patenting. Cockburn, Kortum, and Stern (2003) studied the effect of particular examiners on the characteristics of issued patents and their survival in litigation. 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 (2004) 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 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 just 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 assess how examiner characteristics, and examiner experience in particular, affect the outcomes at the patent office for each of the nearly 10,000 applications filed in January 2001.[2] We also investigate how examiner characteristics affect the number of patent citations and publication citations they reference in their granted patents. These measures are proxies for how much scrutiny examiners are giving to applications and, in particular, how thoroughly they are searching the prior art.

In the next section, we describe the patent prosecution process to provide context for our analysis and motivate construction of the variables we examine. In section III, we describe how we collected the examiner data and provide descriptive statistics. Section IV presents results from linear probability models relating examiner experience and application characteristics to prior art citations. Section V relates examiner experience to the grant rate and 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 confound our results. Section VI examines this possi-

[1] As Thomas Jefferson, the first commissioner of the Patent Office, 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."

[2] As a result of changes in the law, January 2001 is the first month for which the relevant data were available.

bility and suggests that our results are not driven by selection bias. Section VII explores some possible causal mechanisms driving our results, and section VIII concludes with a discussion of potential implications for policy and caveats.

II.   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 that the applicant hopes to obtain. In the United States, applicants have a duty of candor to disclose any previous patents and publications, or "prior art," that are material to the patentability of the claimed invention. Accordingly, most (but not all) patent applicants also disclose patent and nonpatent 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 (2010) and Alcacer et al. (2009) argue that incentives for applicants to do so vary across inventions and industries.

Once applications arrive at the PTO, they are divided 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 (Manual of Patent Examining Procedure sec. 903.08(b)). One issue relevant to our empirical analyses—or any other attempt to relate examiner characteristics to PTO outcomes— is whether there is sorting, that is, 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 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 nonpatent literature (such as scientific and technical journals) to find prior art that might compromise patentability. Searching the nonpatent prior art in particular may be more difficult: Thomas (2001) argues that "in comparison to much of the secondary literature [nonpatent 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" (p. 318).

Examiners then assess the novelty and nonobviousness of the claims in the application relative to what is disclosed in the complete list of prior art—the prior art references from the applicant, plus any discovered through the examiner's own search. Examiners generally issue an initial rejec-

tion of the application, setting out the problems they find in one or more of the claims (Lemley & Sampat, 2009). 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 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 patent lists a range of bibliographic data, including the final claims and information on the applicant and examiner involved in the examination. Issued patents also list each 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 & Gittleman, 2006; Sampat, 2010; Alcacer et al., 2009).

The large number of applications facing the PTO means that examiners are subject to sharp time constraints; the 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 three to four years on average (Allison & Lemley, 2000), but the examiner spends an average of only eighteen 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 & 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 toward granting rather than rejecting applications (Jaffe & Lerner, 2004).

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.[3]

Another difference between more junior and senior examiners is the scrutiny they receive. A junior examiner's

---

[3] 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.

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 four and six years after hire, depending on what level they were hired at), they receive signatory authority, or the right to sign off on an application independently.

The main empirical question we explore in this paper is whether more experienced examiners behave differently in how they process applications and in their propensity to grant applications. This provides a window into whether the human resource dimensions of the PTO matter. Examiner experience is also interesting since there is a widespread belief that despite the incentives discussed above, more experienced examiners are better, and that retaining experienced examiners, or intensifying their mentorship of junior examiners, is the key to solving the patent quality crisis (Jaffe & Lerner, 2004).[4]

III. Data and Descriptive Statistics

We used data from Delphion, a patent data vendor, to identify every original utility patent application filed in January 2001 and published before April 2006. Evaluating applications rather than just issued patents allows us to obtain richer data on the behavior of examiners, data that were not available before 2001. After eliminating plant, design, and reissue patents, applications directed at foreign filing, and follow-on applications (continuations and divisionals), we were left with 9,960 original patent applications. We collected information on the status of these applications as of June 2008 and other information about the prosecution process as of April 2006 from the Patent Office's Patent Application Information Retrieval (PAIR) database.[5]

As reported in Lemley and Sampat (2009), 72% of the applications had resulted in patents by June 2008. About a quarter of the applications, 27%, 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 the applications could also be abandoned for other reasons (perhaps the applicant's firm went out of business).[6] The remaining 1.2% of the applications were pending as of June 2008.

In this paper, we focus on the set of these applications that received a final disposition—the 99% of the applications that are patented or abandoned. In our main analyses, we relate whether these 9,846 applications were granted by June 2008 to examiner characteristics.

Since the PTO employs its examiners, and examiner names are listed on issued patents as well as in application status databases, it might seem straightforward to identify the examiners associated with each application. However, official PTO examiner data are reported in a haphazard format and are rife with errors, as Cockburn et al. (2003) point out. To take just one example, we identified one examiner whose name was spelled no fewer than twenty different ways on the front page of issued patents.[7]

To solve this problem, we acquired the PTO Employee Directories from 1992 to the present. These directories list the examiner name in standardized format. We then used a mix of programming and hand-matching to link these clean examiner names to examiner names in the applications filed in January 2001. In our empirical analyses, we treat the examiner for each application or patent as the person 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.

Access to the employee directories also allowed us to determine exactly when the examiner began working at the PTO, which we used to construct experience measures. These measures are right-censored, since we lack access to employee directories before 1992. Though the applications were filed in 2001, there can be lags before they are docketed to an examiner (90% of the applications in our sample were docketed by the end of 2001 and 98% 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]

In view of the right-censoring, in our empirical analyses, we examine the effects of examiner experience nonparametrically, including dummy variables for zero or one year experience, two to four years, five to seven years, or eight or more years, though the main results are similar if we use different cohort groupings or treat "experience" as a continuous variable.

---

[4] For example, a recent *New York Times* opinion piece, "Inventing a Better Patent System," notes "the quality of American patents has been deteriorating for years" and that "one root of the problem is that patent examiners, many of whom are young or lack practical experience, are not qualified to evaluate whether complex claims in biotech or physics meet the most critical tests: whether the claim is novel relative to prior art, and whether this would be obvious to a person skilled in the art." http://www.nytimes.com/2009/11/17/opinion/17pozen.html.
[5] Unfortunately, changes to the PAIR interface make it difficult to update the prosecution data beyond April 2006.
[6] There is no actual way for the PTO to finally reject a patent (Lemley & 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.

[7] *Ponnathapura Achutamurthy's* name was listed as Achutamurthy Donnathapu, Achutamurthy P., Achutamurthy Ponnathapau, Achutamurthy 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.
[8] For a nontrivial share of the applications, 16.95%, 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 that all of our results are unchanged if we instead use 2001 minus year first in directory as our experience measure.

TABLE 1.—DESCRIPTIVE STATISTICS

Variable	N	Mean	SD	Min	Max
Examiner has zero or 1 year experience	9,847	0.318	0.466	0	1
2–4 years experience	9,847	0.283	0.451	0	1
5–7 years experience	9,847	0.099	0.299	0	1
8 or more years experience	9,847	0.299	0.458	0	1
Application patented?	9,847	0.723	0.447	0	1
Patented without rejections?	7,121	0.178	0.382	0	1
Number of examiner patent cites	6,528	5.75	4.65	0	59
Number of examiner publication cites	6,528	0.158	1.17	0	46
Number of applicant patent cites	6,528	4.86	13.9	0	337
Number of applicant publication cites	6,528	1.37	5.9	0	188

One outcome measure we examine is the extent to which examiners search, as measured by the number of references in issued patents coming from examiners (rather than applicants). Since 2001, the front pages of issued patents identify whether references emanate from applicant disclosures or examiner searches. Previous research (Sampat, 2010; Alcacer et al., 2009) shows that examiners account for a surprisingly large share of citations to patents, though the examiner share varies across fields. Cotropia, Lemley, and Sampat (2010) show that the examiner-provided citations are the ones that are most likely to be used in deciding whether to grant a patent.

Some of this research uses the applicant share of references in a patent as a proxy for how intensely applicants are searching for prior art (Sampat, 2010; Alcacer et al., 2009). In our analyses, for the 6,528 patents resulting from the applications in our sample, we take an analogous approach to assessing examiner search intensity: relating the number of examiner citations to patents and the number of examiner citations to publications to examiner characteristics.[9] In some of the models, we also control for the number of applicant references to patented and nonpatent prior art.[10]

The other main outcome measures we use relates to what happens to the application, the most important decision made by the PTO. In particular, for each of the 9,846 applications, we relate examiner characteristics to whether the application was granted ("Granted") and whether it was granted on first office action ("Granted with no rejections").

Table 1 shows descriptive statistics on each of the main independent variables, calculated at the patent application or patent level.

The modal experience cohort is the first: most of our January 2001 patents have relatively new examiners. Also notable is that there are few examiners in the five-to-seven-year cohort, consistent with concerns that numerous examiners leave the PTO relatively early in their careers (GAO, 2005). But over a quarter of the examining corps are veterans, having stayed with the PTO for over eight years. Similar trends are evident if we calculate these figures across examiners: of the 2,797 unique examiners on these applications, 35% have no experience, 28% two to four years, 9% five to seven years, and 28% eight or more years.[11]

Of the applications, 72% were granted. Of those, about 18% were granted on first office action, without any rejections.

For granted patents, we also have applicant and examiner citation data. Examiners are much more active in citing patented relative to nonpatented prior art, both absolutely and as compared to applicants. These figures are consistent with arguments that examiners are better at searching patented prior art than nonpatent prior art (Thomas, 2001).[12]

IV.   Examiner Experience and Citation Patterns

We begin by relating an examiner's experience to how deeply she searches, as measured by how many citations she adds during patent prosecution. Since citations are available only for issued patents, this measure is conditional on patent issuance and calculated over all 6,528 patents resulting from the applications in our sample.

Before even turning to measured examiner characteristics, we can see the importance of examiner heterogeneity in simpler models. Table 2 summarizes information from least-squares regressions relating the number of examiner citations to patents and the number of examiner citations to

---

[9] While 7117 of the applications were granted–as measured by a notice of allowance being mailed—there is a lag between allowance and when the patent is published: only 6,528 of the resulting patents were published at the time we collected the data. The variables based on information from front pages of issued patents are collected only for these 6,528 patents.

[10] In a previous version, we looked at the share of applications supplied by an examiner. On the suggestion of a referee that the share conflates applicant and examiner side characteristics, we now use the number of examiner citations. Results using the share measures are qualitatively similar and available on request.

[11] There is considerable within- and across-field variation in experience. Unreported analyses reveal that in the computer art units, examiners are overwhelmingly newly hired. There is also a large proportion of newly hired examiners 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 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.

[12] Examiners account for 43.6% of citations to patented prior art in their issued patents, but only 9.5% of all citations to nonpatent prior art.

TABLE 2.—SUMMARY INFORMATION FROM OLS REGRESSIONS RELATING THE NUMBER OF EXAMINER CITATIONS IN A PATENT TO ART UNIT EFFECTS AND EXAMINER EFFECTS

	Number of Examiner Citations to Patents		Number of Examiner Citations to Publications	
	R^2	Adjusted R^2	R^2	Adjusted R^2
Model with art-unit fixed effects only	0.219	0.182	0.221	0.185
Model with art-unit and examiner fixed effects	0.694	0.527	0.625	0.421
Joint significance of examiner fixed effects	$F(2210,4227) = 3.7$		$F(2210,4227) = 3.1$	
	$p < 0.01$		$p < 0.01$	
N	6,528	6,528	6,528	6,528

TABLE 3.—OLS REGRESSIONS RELATING THE NUMBER OF EXAMINER CITATIONS TO EXAMINER EXPERIENCE AND APPLICANT CITATIONS

	Model 3.1: Number of Examiner Cites to Patents	Model 3.2: Number of Examiner Cites to Publications	Model 3.3: Number of Examiner Cites to Patents	Model 3.4: Number of Examiner Cites to Publications
2–4 years Experience	−0.342	−0.064	−0.349	−0.064
	(0.246)	(0.039)	(0.246)	(0.039)
5–7 years Experience	−0.974***	−0.011	−0.989***	−0.011
	(0.293)	(0.059)	(0.292)	(0.057)
8+ years Experience	−1.709***	−0.078**	−1.712***	−0.081**
	(0.241)	(0.032)	(0.241)	(0.031)
Number of applicant cites to patents			−0.012***	−0.002
			(0.004)	(0.002)
Number of applicant cites to publications			−0.003	0.022
			(0.013)	(0.015)
Constant	6.55***	0.205***	6.613***	0.186***
	(0.195)	(0.028)	(0.195)	(0.029)
Art unit effects	Yes; jointly significant	Yes; jointly significant	Yes; jointly significant	Yes; jointly significant
N	6,528	6,528	6,528	6,528

Robust standard errors (clustered on examiners) reported in parentheses. Asterisks denote statistical significance: *10%, **5%, ***1%.

publications (nonpatent prior art) to art unit indicators (for over 300 art units, to control for technology class) and examiner indicators. The large increments to R^2 and adjusted R^2 in the models that include examiner effects provide strong initial evidence of examiner heterogeneity in prior art citation patterns. Note also that for both patent and publication citations, the examiner effects are jointly significant at the 99% level.

Next, we estimated models relating the number of examiner citations in a patent to examiner experience. Here again, these models are estimated over each of the granted patents resulting from the January 2001 applications and include fixed effects for each art unit. Table 3 shows results.

Model 3.1 shows that the number of examiner patent citations decreases monotonically with examiner experience, with the most experienced examiners (eight or more years) citing nearly two fewer patents than the least experienced (the left-out category). Model 3.2 shows that the magnitude of the experience effect is much smaller for publication citations, though the most experienced examiners have a significantly lower volume of such cites.

Since examiner citations in a patent could themselves be affected by the extent of prior art an applicant provides (Sampat, 2010), for each dependent variable we also estimated models where we included the number of applicant citations in the patent (to patents and publications separately) as independent variables. Interestingly, there is little

evidence of substitutability here: more applicant cites do not seem to be associated with fewer examiner cites. (Model 3.3 shows a statistically significant but small negative relationship between the number of applicant cites to patents and the number of examiner cites.) Moreover, including these variables does little to change the magnitude of the experience effects documented in models 3.1 and 3.2.

Overall, the data suggest that within art units, more experienced examiners tend to cite less prior art, with the magnitude of this effect more pronounced for citations to patented prior art than for publication citations.

V. Examiner Experience and the Patent Grant Rate

Next, using the full sample of applications (those granted and not granted), we relate examiner experience to the most important decision made by the PTO: whether to grant a patent. Table 4 shows results. Model 4.1 shows results from linear probability models relating whether an application was granted to examiner experience and art unit effects. The grant rate increases monotonically with experience, with the two most experienced groups having an 11 percentage point higher grant rate.

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

Table 4.—Linear Probability Models Relating Whether an Application Is Granted and Whether Granted with No Rejections, to Examiner Experience

	Model 4.1: Was the Application Granted?	Model 4.2: Was the Application Granted with No Rejections?
2–4 years experience	0.057***	0.060***
	(0.013)	(0.013)
5–7 years experience	0.110***	0.111***
	(0.018)	(0.020)
8+ years experience	0.112***	0.133***
	(0.013)	(0.015)
Constant	0.663***	0.105***
	(0.009)	(0.009)
Art unit effects	Yes; jointly significant	Yes; jointly significant
N	9,846	7,117

Robust standard errors (clustered on examiners) reported in parentheses. Asterisks denote statistical significance: *10%, **5%, ***1%.

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.

But it may also be that whether a patent is granted is too blunt a measure of the quality of examination. As discussed in section II, 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 patent. Another measure of the rigor or quality of patent prosecution is the extent to which this narrowing occurs.

To examine this, for the 7,121 applications that were allowed, we collected data from the Patent Application Information Retrieval (PAIR) database on whether there were any rejections during patent prosecution and investigated how the likelihood of grant with at least one rejection relates to experience. We were unable to locate prosecution histories for four of these, reducing our sample to 7,117. For those where we could, 82% of these granted applications received a nonfinal rejection, and 26% a final rejection, during patent prosecution. But 18% did not receive any rejections before they were issued; they were issued on the first office action, as is.

Model 4.2 in table 4 shows results from linear probability models relating whether a patent was issued with no rejections to experience (conditional on a patent having issued). The likelihood of granting without rejections increases sharply with experience, with the most experienced cohort 13 percentage points more likely to do so.

The finding that examiner experience increases the probability of having had no rejections, conditional on issuing a patent, also undermines one possible explanation for the relationship between examination and grant rates reported above: that experienced examiners are better at separating the wheat from the chaff, or identifying what is patentable in an application, getting quickly to the patentable piece of an application by negotiating with the applicant to amend

his claims appropriately. First office action allowances involve little such parsing or deliberation.

VI. Selection

Taken together, the data from the previous two sections show that not only do more experienced examiners cite less prior art, they also have a significantly higher grant rate and are significantly more likely to grant on the first office action (that is, grant without rejections).

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 patentworthy, our estimates would be biased. 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.

The PTO does not publish rules regarding the assignment of applications within art units to particular individuals. To investigate the possibility of selection effects, we conducted written interviews with roughly two dozen current and former examiners and supervisory patent examiners. Those interviews revealed no evidence of deliberate selection or assignment of applications on the basis of the characteristics of the application. Applications are assigned to art units in the first instance on the basis of their technology classifications; the head of the art unit then decides which subclass the application belongs to and assigns it accordingly. Within subclasses, SPEs for the most part assigned applications randomly, assigning applications to particular examiners on the basis of the last digit of the application serial number. Because application serial numbers are assigned sequentially in the central PTO, this assignment system, while not truly random, is not subject to manipulation by applicants or supervisory examiners. A minority of supervisory examiners assigned applications on the basis of docket management, giving the oldest unassigned application to an examiner when that examiner finished examining a prior application. This minority approach would have the effect of disproportionately assigning applications to examiners who are quickest to resolve cases but does not involve assignment on the basis of observable characteristics. Only one interviewee mentioned the possibility of selection of applications, and that person mentioned it to dismiss it— referring to supervisory examiners having to "resist cherry-picking what appeared to be the easiest cases."

Thus, our interviews reveal no evidence that supervisory examiners 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

TABLE 5.—APPLICATION-LEVEL CHARACTERISTICS VERSUS EXAMINER EXPERIENCE

	Model 5.1: Pages	Model 5.2: Family Size	Model 5.3: Applicant Patent Cites	Model 5.4: Applicant Publication Cites	Model 5.5: Applicant Patent Volume in 2000
2–4 years	0.450	0.113	−0.642	−0.070	0.284
	(0.447)	(0.071)	(0.502)	(0.227)	(31.428)
5–7 years	0.176	0.259**	−1.194**	−0.107	−13.676
	(0.535)	(0.105)	(0.605)	(0.258)	(44.379)
8+ years	0.241	0.088	−0.260	0.092	5.714
	(0.524)	(0.080)	(0.531)	(0.186)	(31.972)
Constant	14.574***	2.972***	5.278***	1.379***	525.395***
	(0.298)	(0.051)	(0.387)	(0.158)	(23.983)
Art unit effects	Yes; jointly significant	Yes; jointly significant	Yes; jointly significant	Yes; jointly significant	Yes; jointly significant
N	9,846	9,846	6,528	6,528	4,503

Robust standard errors (clustered on examiners) reported in parentheses. Asterisks denote statistical significance: *10%, **5%, ***1%.

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. For example, recall our finding that more experienced examiners are much more likely to grant applications without amendment: if supervisory examiners 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 patentworthy an application is by looking at whether the same applications were granted in Europe (at the European Patent Office, EPO).

A. Selection on Observables

There are two difficulties in examining potential selection on observable variables. The first is that for patent applications, much of the front-page data available for issued patents (including citations and assignee information) are 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 patentworthiness 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 nonobvious in light of the prior art.[13]

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, that is, 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 results.

Models 5.1 and 5.2 relate each of these variables to the nonparametric experience measures, and art unit effects. In model 5.1 there is no clear pattern in any of the coefficients, and none are statistically significant. There is thus little evidence that more experienced examiners get systematically longer or shorter applications. For the "family size" measure, the coefficients also vary haphazardly with experience, and only one of the estimates (for five to seven years) is statistically significant at the 5% level. Thus, here too there is no evidence of any systematic selection.

As another check for selection, we looked at application-level characteristics for the subset of granted patents. (Granted patents include more information, including the variables below, than applications.) This analysis requires characteristics of the application that could plausibly influence patent-worthiness but would not 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 the patent prosecution process itself (such as 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 nonpatented prior art, and the number of patents the applicant was issued in the previous year (2000), a measure of the applicant's experience.[14] Models 5.3 to 5.5 show results from regressions of these variables on examiner experience. These

---

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

[14] Data on the application patent volume variable are available for only the 4,503 patents that were assigned at issue and thus have assignee information listed.

TABLE 6.—U.S. VERSUS EPO OUTCOMES

	Not Patented by PTO	Patented by PTO	Total
Rejected by EPO	344	525	869
Patented by EPO	128	1143	1271
Pending at EPO	79	512	591
Total	551	2,180	2,731

models too show no evidence that more experienced examiners get different types of applications.

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.[15] In the next section, we examine a proxy variable, whether the same application was granted by the EPO.

B. Selection: Evidence from the EPO

We also collected information from Delphion on whether the applications in our sample were also filed at the 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 or withdrawn or were currently pending.[16] Following Webster et al. (2007), we limited our sample to the 8,905 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, 2,761 were filed at the EPO, which had granted patents on 46% of them as of June 2008. Table 6 shows a cross-tabulation of EPO and PTO outcomes as of June 2008. Of the applications that were granted in the United States, slightly more than half (52.1%) have been granted by the EPO.[17] By contrast, of those granted by the EPO, the vast majority (88%) are also granted by the PTO.

Table 7 shows the effects of the experience and search intensity measures on probability of U.S. grant, conditional on EPO status. Model 7.1 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 United States. But controlling for EPO status, more experienced examiners have a statistically and qualitatively higher likelihood of grant than less experienced examiners.

We also assessed the effects of experience on likelihood of U.S. patent grant separately for applications that were patented at the EPO (model 7.2) and rejected at the EPO (model 7.3). Model 7.2 shows that examiner experience does not affect the likelihood that applications patented at the EPO are patented in the United States: 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 7.3 shows that more experienced examiners are more likely to grant applications that the EPO rejected.

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 that the experience effect does not reflect younger examiners being more likely to reject applications patented at the EPO; instead, older examiners are more likely to patent applications rejected by the EPO. That the experience effect persists when we control for EPO decisions is further reason to believe there is not a selection effect at work.

VIII. 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, the PTO faces significant employee attrition, particularly among examiners who have been with the agency less than five years (GAO, 2007). If examiners who were more diligent, more thorough, more technically sophisticated, or more highly educated were more likely to leave the PTO 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 PTO. Since we lack any additional demographic information about our examiners, this is difficult to assess directly. Moreover, since we have only 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 PTO employee directories do give 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% of the examiners in our sample had left by the end of 2006. Attrition is more pronounced for the newest examiners (consistent with

---

[15] For the same reason, it is unlikely that supervisory patent examiners (SPEs) are judging patent-worthiness in their initial assignment decisions either.

[16] http://www.epoline.org/.

[17] This number differs somewhat from that in Jensen et al. (2006), who show that 63% of granted applications in the United States 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.

TABLE 7.—LINEAR PROBABILITY MODELS RELATING WHETHER AN APPLICATION IS GRANTED TO EPO OUTCOMES, AND EXPERIENCE

	Model 7.1: Patented in United States?	Model 7.2: Patented in United States and Patented by EPO	Model 7.3: Patented in United States and Rejected by EPO
2–4 years experience	0.0369*	–0.004	0.051
	(0.021)	(0.026)	(0.053)
5–7 years experience	0.080***	0.011	0.174**
	(0.028)	(0.032)	(0.071)
8+ years experience	0.073***	0.028	0.118**
	(0.021)	(0.026)	(0.054)
Rejected by EPO	–0.269***		
	(0.019)		
Pending by EPO	–0.009		
	(0.018)		
Constant	0.844***	0.890***	0.536***
	(0.016)	(0.018)	(0.036)
Art unit effects	Yes; jointly significant	Yes; jointly significant	Yes; jointly significant
N	2,731	1,271	869

Robust standard errors (clustered on examiners) reported in parentheses. Asterisks denote statistical significance: *10%, **5%, ***1%.

TABLE 8.—OLS MODELS RELATING PROSECUTION CHARACTERISTICS AND OUTCOMES TO WHETHER EXAMINER WILL LEAVE IN FIVE YEARS

	Model 8.1: Examiner Patent Cites	Model 8.2: Examiner Publication Cites	Model 8.3: Was the Application Patented?	Model 8.4: Was the Application Patented with No Rejections?
Examiner will leave in 5 years	–0.174	0.069	–0.006	.051***
	(0.212)	(0.048)	(0.013)	(.014)
2–4 years experience	–0.366	–0.058	0.056***	.064***
	(0.249)	(0.040)	(0.013)	(.013)
5–7 years experience	–1.014***	–0.001	0.109***	.117***
	(0.296)	(0.059)	(0.018)	(.020)
8+ years experience	–1.725***	–0.076**	0.112***	.136***
	(0.243)	(0.032)	(0.013)	(.015)
Applicant patent cites	–0.011***	–0.002		
	(0.004)	(0.002)		
Applicant publication cites	–0.003	–0.023		
	(0.013)	(0.015)		
Constant	6.657***	0.169***	0.664***	.094***
	(0.208)	(0.033)	(0.009)	(.009)
Art unit effects	Yes; jointly significant	Yes; jointly significant	Yes; jointly significant	Yes; jointly significant
N	6,528	6,528	9,846	7,117

Robust standard errors (clustered on examiners) reported in parentheses. Asterisks denote statistical significance: *10%, **5%, ***1%.

the figures in GAO, 2007) with 27.1% 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 PTO by 2006. Table 8 shows the results.

Models 8.1 and 8.2 respectively relate the number of patent and publication references in a patent to whether the examiner would exit by 2006, conditional on experience. Conditional on experience, examiners who would leave within five years have a lower number of patent citations and a higher number of publication citations, though neither effect is statistically significant. For patented applications, we also examined if whether an application was granted (model 8.3) and the probability of granting with no rejections (model 8.4) related to whether an examiner would leave. The relationship between whether an examiner will

leave and the grant rate is essentially zero and statistically insignificant. And examiners who would exit have a higher likelihood of granting without rejections, the opposite of what we would expect were selective retention the explanation for why more experienced examiners tend to be more likely to grant without rejections. Collectively, 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 using data from issued patents, we can determine which examiners were definitely tenured by 2001. 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% of examiners with two to four years of experience tenured by 2001, 41.6% of examiners with five to seven years of experience, and 90.7% of examiners with eight or more years experience. The strong relationship between tenure and experience also makes it difficult to separately identify a tenure effect separate from an experience effect. More important, tenure is nonrandom: for a given level of experience, those who are tenured are likely to be systematically different from those who are not.

While we cannot rule out a tenure effect, the result that there are strong differences in patent application outcomes even between the two- to four-year cohort (the bulk of whom are not tenured) and brand-new examiners suggests 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 PTO. 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 PTO could affect outcomes.

IX.   Implications

Consistent with the qualitative and quantitative results of Cockburn et al. (2003), our data emphasize that examiners matter. In particular, more experienced examiners cite less art, issue fewer initial rejections, and are more likely to grant patents.

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.[18] The fact that seniority is correlated with more first-action allowances—allowances where there was no rejection—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 PTO outcomes. From this perspective, the tenure system, the count system, and examiner recruitment and retention policies should 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 & Sampat, 2009). 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 preparing this study, we determined that the computer industry had by far the highest percentage of new examiners: more than 60% of examiners in that art unit had less than a year of experience compared with less than 20% 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.

Finally, our data suggest that whether the patent office grants or rejects a patent is significantly related to the happenstance of which examiner is assigned the application. That may not be an encouraging result if our goal is a system that rewards deserving patent applicants while denying patents to the rest.[19] Nor is it clear that in this environment, all issued patents should enjoy the same presumption of validity.

---

[18] 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.

[19] As a referee usefully pointed out to us, to the extent that the PTO's objectives are broader than quality—instead including speed and quantity of applications processed, for example—heterogeneity across experienced and less experienced examiners would be less of a problem. Thus, even if some examiners search less and are more likely to grant quality patents than others, they may compensate by handling more applications per unit of time. There is some evidence that more experienced examiners at least receive more applications: examiners with zero to one year of experience received an average of 3.2 new applications in January 2001, compared to 3.6 for examiners with two to four years of experience, 3.6 for examiners with five to seven years of experience, and 3.9 for examiners with more than eight years of experience.

REFERENCES

Alcacer, J., and M. Gittleman, "Patent Citations as a Measure of Knowledge Flows: The Influence of Examiner Citations," this REVIEW 88:4 (2006), 774–779.

Alcacer, J., M. Gittleman, and B. Sampat, "Applicant and Examiner Citations to Patents: An Exploration and Overview," *Research Policy* 38 (2009), 415–427.

Allison, J. R., and M. A. Lemley, "Who's Patenting What? An Empirical Exploration of Patent Prosecution," *Vanderbilt Law Review* 56 (2000), 2099–2174.

Cockburn, I., S. Kortum, and S. Stern, "Are All Patent Examiners Equal? Examiners, Patent Characteristics, and Litigation Outcomes," In W. M. Cohen and S. A. Merrill (Eds.), *Patents in the Knowledge-Based Economy* (Washington, DC: National Academies Press, 2003).

Cohen, W. M., and S. Merrill, *Patents in the Knowledge-Based Economy* (Washington, DC: National Academies Press, 2003).

Cotropia, C., M. Lemley, and B. Sampat, "Do Applicant Citations Matter?" working paper (2012).

Federal Trade Commission (FTC), "To Promote Innovation: The Proper Balance of Competition and Patent Law and Policy," Federal Trade Commission technical report (2003).

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

——— "Hiring Efforts Are Not Sufficient to Reduce the Patent Application Backlog" (2007).

Jaffe, A. B., and J. Lerner, *Innovation and Its Discontents: How Our Broken Patent System Is Endangering Innovation and Progress, and What to Do about It* (Princeton, NJ: Princeton University Press, 2004).

Jensen, P. H., A. Palangkaraya, and E. Webster, "Disharmony in International Patent Office Decisions," *Federal Circuit Bar Journal* 15 (2006), 679–704.

Lemley, M. A., "Rational Ignorance at the Patent Office," *Northwestern University Law Review* 95 (2001), 1495–1532.

Lemley, Mark A., and C. Shapiro, "Probabilistic Patents," *Journal of Economic Perspectives* 19 (2005), 75–98.

Lemley, M. A., and K. A. Moore, "Ending Abuse of Patent Continuations," *Boston University Law Review* 82 (2004), 63–118.

Lemley, M. A., and B. Sampat, "Is the Patent Office a Rubber Stamp?" *Emory Law Journal* 58 (2009).

Lichtman, D., "Rethinking Prosecution History Estoppel," *University of Chicago Law Review* 71 (2004), 181–208.

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

Norhaus, W., *Invention, Growth, and Welfare* (Cambridge, MA: MIT Press, 1969).

Sampat, B., "When Do Applicants Search for Prior Art?" *Journal of Law and Economics* 53 (2010), 399–416.

Thomas, J., "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, and P. H. Jenson, "Characteristics of International Patent Application Outcomes," *Economic Letters* 95 (2007), 362–368.