

IS THE TIME ALLOCATED TO REVIEW PATENT APPLICATIONS INDUCING EXAMINERS TO GRANT INVALID PATENTS?: EVIDENCE FROM MICRO-LEVEL APPLICATION DATA

Michael D. Frakes^{} and Melissa F. Wasserman^{**}*

This paper explores how examiner behavior is altered by the time allocated for reviewing a patent application. Insufficient examination time may crowd out examiner search effort, impeding the ability to form time-intensive prior-art-based rejections (especially, obviousness rejections) and thus leaving examiners more inclined to grant otherwise invalid applications. To test this prediction, we trace the behavior of individual examiners over the course of a series of certain promotions that carry with them a substantial reduction in expected examination time. For these purposes, we use novel micro-level application data spanning a ten year period and estimate examiner fixed-effects specifications that allow us to control flexibly for examiner heterogeneity. We find evidence demonstrating that search efforts and time-intensive rejections indeed fall, while granting tendencies rise, upon the promotions of interest. Assuming that patent examiners will tend to make the correct patentability determinations when provided sufficient examination time, our results suggest that the present schedule of time allotments may be inducing patent examiners to grant patents excessively.

^{*} Associate Professor of Law, Northwestern Law School; Faculty Research Fellow, National Bureau of Economic Research; Faculty Fellow, Institute for Policy Research, Northwestern University. J.D., Ph.D.

^{**} Associate Professor of Law; Richard and Anne Stockton Faculty Scholar; & Richard W. and Marie L. Corman Scholar, University of Illinois College of Law. J.D., Ph.D.

I. INTRODUCTION

Evidence suggests that patents play an important role in both promoting innovative activity and shaping the direction of technological growth (Moser, 2004). Yet in recent years the patent system has come under trenchant criticism (Burk & Lemley, 2004). Many believe the root cause of the system's dysfunction is that the U.S. Patent and Trademark Office (Patent Office or Agency) is issuing too many invalid patents that unnecessarily drain consumer welfare, stunt productive research, and unreasonably extract rents from innovators (Jaffe & Lerner, 2004). Although commentators have suggested a plethora of reasons as to why the Agency may be biased towards allowing patents, there exists little compelling empirical evidence that any particular feature of the Patent Office actually induces the Agency to over-grant patents.¹

This paper begins to rectify this deficiency by addressing one feature of the Patent Office that we contend may influence an examiner's decision to grant a patent: the time allotted to review a patent application. Because patent applications are presumed to comply with the statutory patentability requirements when filed, the burden of proving unpatentability rests with the Patent Office. That is, a patent examiner who fails to set forth reasons as to why the application fails to meet the patentability standards must grant the patent. To the extent that examiners are given insufficient examination time, one might expect them to conduct limited reviews of the applications and therefore grant patents at elevated rates. Much anecdotal evidence has been put forth to suggest that patent examiners indeed face binding examination time constraints, implicating such concerns.

To more comprehensively test this simple hypothesis and challenge this anecdotal sentiment, we rely upon the fact that examination times decrease upon certain types of examiner

¹ See, however, Frakes and Wasserman (2013, 2015), which explore how the Patent Office's fee schedule, along with the Office's inability to finally reject a patent application, creates an incentive for a financially constrained agency to allow additional patents.

promotion. That is, as examiners ascend in rank, they are expected to process applications more expeditiously. Our basic empirical strategy is to follow individual examiners throughout the course of their careers and to track the evolution of their examination behavior—including their granting rates—as they experience promotions that diminish the amount of examination time at their disposal. Bolstering our ability to separate the effect of examination time from other factors that may change generally upon promotion is the fact that examiner promotions and pay raises come in several varieties, some of which bear on allocated examination times and some of which do not. There is even variation among those types of promotions that come with reductions in examination times, a fact which we likewise exploit to identify the relationship of interest. Our identification strategy is further strengthened by the fact that the promotions of interest do not transpire lock-step with increases in years of experience, allowing us to decouple an experience effect from a promotion-of-interest effect.

To execute this empirical strategy, we estimate examiner fixed-effects specifications using novel, micro-level data on 1.4 million patent applications disposed of between 2002 and 2012, merged with rich, examiner roster data received from the Patent Office pursuant to a series of Freedom of Information Act Requests (FOIA). Our results suggest that as an examiner is given less time to review an application, the less active she becomes in searching for prior art, the less likely she becomes to make prior-art-based rejections (in particular obviousness rejections, which are especially time-intensive exercises),² and the more likely she becomes to grant the patent. Under the assumption that patent examiners who are allocated sufficient time to review applications will, on average, make the correct patentability determinations, our results suggest that the time allotments are inducing patent examiners to grant patents excessively.

² To the best of our knowledge, our analysis is the first to report comprehensive application level rejection data.

Our findings have several important implications for patent policy. First, they provide policymakers with guidance as to what institutional features of the Patent Office are actually inducing the Agency to over-grant patents and hence, concomitantly, insight on how to begin to solve the patent quality crisis. That is, they suggest that the hours allocated to review an application or, in particular, how these allocations decrease upon promotion, may be leading to the granting of invalid patents. Second, our results cast doubt on a widely held belief that decreasing patent examiner attrition or intensifying the monitoring of newly hired examiners is vital to increasing patent quality (Jaffe & Lerner, 2004). In fact our findings suggest that the process of promoting examiners, which is meant to reward admirable behavior on the part of examiners, may, in part, be responsible for the agency issuing patents of marginal quality.

Several scholars have previously studied examiner heterogeneity and its effect on the patenting process (Cockburn, Kortum, & Stern, 2003; Lichtman, 2004; Mann, 2014). Demonstrating that an applicant's experience with the application process is largely a function of the examiner that she randomly receives, these studies raise concerns of an inefficient Patent Office. However, these groundbreaking studies fail to explore arguably the most important outcome of this process—that is, whether the examiner granted the patent—while also failing to examine whether a particular feature of the Patent Office influenced the examiner's behavior.

Lemley and Sampat (2012) arguably come the closest in filling these gaps in the literature, estimating a monotonically increasing relationship between years of examiner experience and examiner grant rates. Given the natural connection between experience and promotion, their analysis undoubtedly captures some aspects of the impact of allotted examination time on grant rates; though, absent data on examiner promotions, they are unable to decouple an experience effect from an examination-time-allotment effect. Moreover, their

analysis is largely cross-sectional in nature (observing 10,000 patent applications filed in January 2001) and cannot fully rule out that the observed relationship is attributable to a story of selective retention—i.e., that senior examiners represent those that have elected to stay and may thus be of a distinct disposition. By tracking individual examiners over the course of a ten-year period, our fixed-effects specifications are able to overcome these concerns. While our focus is on understanding the impact of reductions in allocated examination time and not necessarily on the independent impacts of examiner experience, we note that the imposition of examiner fixed effects produces an inverse-U shape in the relationship between grant rates and experience, as opposed to the monotonic relationship documented in Lemley and Sampat (2012). Some specifications even suggest a strictly negative influence of experience (in years) on grant rates.

In the next section, we describe the patent prosecution process and the manner in which examiners are promoted, motivating a theoretical prediction that reductions in examination time will lead to less prior-art searching, less prior-art based rejections (especially obviousness) and overall higher grant rates. In section III, we describe how we collected our data and provide descriptive statistics. Section IV describes our empirical methodology, while Section V presents results from our examiner fixed-effects analysis. Finally, Section VI concludes.

II. BACKGROUND AND THEORY

A. Description of Examination Process

Each year approximately 500,000 new patent applications are filed at the Patent Office. Every patent application contains a specification, which describes the invention, and a set of claims that defines the metes and bounds of the legal rights the applicant is seeking. Patent applicants in the United States have a duty of candor to disclose material information to the Patent Office regarding the patentability of the invention. To fulfill this duty, applicants

typically disclose to the Agency “prior art,” that is previous patents, patent applications, or other publications, that are material to the patentability of their invention.

Before a patent application enters examination, the Patent Office assigns the application a classification based on the technology disclosed within the application. This classification is utilized to route the patent application to an Art Unit, an administrative unit comprising eight to fifteen patent examiners who examine applications in the same technological field. Upon arrival, the Supervisory Patent Examiner (SPE) of that Art Unit assigns the application to a specific examiner. Examination assignments are not based on the characteristics of patent applications, such as the application quality or patent worthiness. Instead, applications are largely assigned randomly. However, even when SPEs make non-random assignments, they do so not based on any characteristic that would affect the patentability of the application but instead, for instance, on an examiner’s backlog of patent applications (Lemley & Sampat, 2012). We conducted a series of telephone interviews with former SPEs to confirm these details of patent examination assignment. Our interviews further substantiated that SPEs do not make any substantive evaluation of an application before assigning it to a particular examiner.

The examination of an application will typically begin with the patent examiner conducting her own prior art search to supplement the prior art disclosed by the patent applicant. Upon completion of this search, the examiner assesses the patentability of the invention based on the criteria outlined in the Patent Act, including whether the claimed invention involves statutory subject matter (35 U.S.C. § 101) that is novel (35 U.S.C. § 102), useful (35 U.S.C. §101), and nonobvious (35 U.S.C. § 103) and whether the application satisfies the disclosure requirements (35 U.S.C. § 112). Without making any reference to prior art, a patent examiner can deny a patent based on grounds of lack of utility, lack of patentable subject matter, or failure to satisfy

the disclosure requirements. In contrast, novelty and nonobviousness assessments require the examiner to make a comparison of the claimed invention with the background art already known to the public. Because these latter rejections require this delicate comparison—along with the search underlying this comparison—they are typically viewed as being more time consuming to perform than non-art-based rejections. Obviousness-based rejections are especially time intensive in this regard, even relative to novelty rejections.³ While novelty assessments require that examiners determine whether the claimed invention is covered by a single prior publication or patent, obviousness assessments require more. That is, an obviousness determination requires an examiner to start with a prior art reference that covers only a portion of the invention and then piece together additional references or rely upon what is known to one of ordinary skill in the art. The challenge with, and thus the extra effort associated with, an obviousness rejection comes in determining whether it would be “obvious” in light of this group of multiple prior art references (and/or what is known to one of ordinary skill in the art) to modify any one of the cited prior art references to achieve the claimed invention.

After assessing the patentability of the claimed invention, an examiner composes a “first office action” letter to the applicant that accepts, or rejects, the claims. Importantly, because patent applications are presumed to meet the patentability requirements when filed, a patent examiner who fails to set forth a basis of rejection must grant the patent. Although some applications will be allowed in their entirety upon first examination, more frequently, some or all of the claims will fail to meet at least one of the patentability requirements. Thus, the first office action will typically contain a detailed analysis for the basis of rejecting the patent application. The applicant then responds by amending the patent claims or disputing the rejection. After the

³ We confirmed that obviousness rejections are, on average, more time intensive than novelty rejections through a series of interviews with former SPEs (the support for this contention was overwhelming).

response, a patent examiner may issue a final rejection in a final office action or allow the patent to issue.⁴

B. Examination-Time Allocations

Patent examiners are under considerable time constraints in assessing the patentability of claims and a number of scholars have surmised that these time constraints are partly responsible for the Patent Office allowing too many invalid patents (Jaffe & Lerner, 2004; Lemley, 2001). Although it may take several years from filing a patent application for an applicant to receive a final patentability decision from the Patent Office, on average, an examiner spends only nineteen hours reviewing an application, including reading the patent application, searching for prior art, comparing the prior art with the patent application, writing a rejection, responding to the patent applicant's arguments, and also often conducting an interview with the applicant's attorney (Frakes & Wasserman, 2014). Examiners are generally expected to meet certain workload goals, whereby they are expected to attain a certain number of credits (often referred to as "counts"). A patent examiner does not receive credit for each task associated with examining an application. Historically credits have been earned only upon the issuing of a first office action regarding the patentability of an application and at final disposal, which occurs when a patent application is allowed by the examiner or abandoned by the applicant (often after receipt of a final rejection or in anticipation of such a rejection).⁵

⁴ After receiving a final rejection, an aggrieved patent applicant can restart the examination process by filing a continuation application, appeal the denied application to Patent Board of Appeals, or abandon the application altogether.

⁵ Since 2010 examiners can also earn partial credits for final office actions and examiner-initiated interviews with the patent applicant or her attorney. Under either system, a patent examiner earns a maximum of two credits per patent application examined. While examiners are free to average these time allotments over their caseload, they are strongly encouraged to meet their credit quota on a biweekly basis.

By setting expectations regarding the number of credits examiners should attain, the Patent Office contemporaneously sets expectations regarding the amount of time examiners should spend on applications. These time allotments have largely remained unchanged since 1976.⁶ The number of expected hours allocated to review a patent application—i.e., the time allotted to earn two credits—depends on both the technological field in which the examiner is working and on her position in the general schedule (GS) pay scale. A patent examiner in a more technological complex field is allocated more hours to review an application than a patent examiner of the same grade who is working in a less technological complex field. The higher the pay grade of an examiner within a technology area the fewer number of hours the Patent Office allocates to earn two credited work units. A promotion to each subsequent pay grade is roughly equated to a ten to fifteen percent decrease in the number of hours allocated to review an application. While hour allotments generally change only with GS-level promotions, examiners often experience an hour-allotment alteration during the middle of their tenure at GS-level 13. This unique within-GS-level promotion constitutes the moment that examiners are granted “partial signatory authority,” a status which we describe below.

To demonstrate the degree to which time allocations scale with GS-level changes, we present in Table 1 the examination time expectations facing a patent examiner working in one of the most complex fields, artificial intelligence, and one of the least complex fields, compound tools. For reasons explained below, our empirical analysis focuses on the GS-7 to GS-14 range; as such, we confine this Table to those pay grades. Examiners operating at GS-level 7 are given the greatest amount of time—19.7 hours and 45.1 hours—in reviewing a patent in compound

⁶ The Patent Office has created new Art Units to review patent applications covering emerging technology. Once the Agency has set the time allotments for new technology these allocations also have largely remained unchanged.

tools and artificial intelligence respectively, whereas examiners operating at GS-level 14 are expected to review the same patent in approximately half that time.

TABLE 1: EXAMINATION HOURS ALLOCATED TO EXAMINER AS A FUNCTION OF GS-LEVEL

| | (1) | (2) |
|--------------------------|----------------|-------------------------|
| GS-level | Compound Tools | Artificial Intelligence |
| GS-7 | 19.7 | 45.1 |
| GS-9 | 17.3 | 39.5 |
| GS-11 | 15.3 | 35.1 |
| GS-12 | 13.8 | 31.6 |
| GS-13 | 12.0 | 27.5 |
| GS-13, partial signatory | 11.0 | 25.3 |
| GS-14 | 10.2 | 23.4 |

C. Promotion Process

Patent examiners are hired at different pay grades (GS-5, GS-7, GS-9 or GS-11) depending upon their educational background and prior experience. Promotions at low pay grades typically occur within a year for examiners that meet their credit quota with few errors. In contrast, promotions at the high pay grades (GS-13, 14 and 15) often require more time, as they involve the completion of additional testing or programs. It is important to note that while examiners do generally receive pay increases upon promotions to higher grades, they generally also receive pay raises as they spend more years within given GS levels, a feature of the examiner compensation structure that will be helpful to our identification strategy discussed further below.

While we contend that the most significant change associated with a promotion that bears on the examiner's decision to grant a patent application is the time allocated to earn two credited work units, there is, upon promotion within GS-13 and to GS-14, also a change in the scrutiny of

their work. Examiners at pay grades GS-13 and below must have their decisions reviewed by an examiner that has “full signatory authority.”⁷ Patent examiners at pay grades GS-13 may begin to work towards obtaining such authority, by undergoing an evaluation period, which upon successful completion will result in a promotion within GS-13 to a patent examiner with “partial signatory authority.” This latter promotion, though not associated with a change in the GS level, is associated with a decrease in the examination time allotted to the promoted examiner (as discussed above) and provides that examiner the ability to sign off independently on first office actions. A GS-13 partial signatory patent examiner can be promoted to GS-14, which is associated with full signatory authority or the right to sign off on all aspects of an application independently, upon successfully completing a second period of evaluation. The fact that variations in scrutiny of this nature does not occur upon all examination-time-reducing promotions is likewise an important component to our identification strategy, as discussed in Section V below.

Finally, we note that, to our knowledge, nothing else changes upon GS-level promotions that would affect the manner in which examiners conduct their examination. For instance, the basic structure of overtime and bonuses remains constant upon GS-level promotions as does the ways in which examiners earn work credits, in which event one would not expect examiners to face enhanced financial incentives to grant patents (to the extent that they ever face such incentives) upon promotions to higher grade levels. We confirmed that GS-level promotions are not associated with such changes through our review of examiner compensation materials made available by the Patent Office and through our interviews with former SPEs.

D. Hypothesis

⁷ Even though these “junior” examiners do most of the work on the application they are listed as secondary examiners on the application.

Our analysis effectively assumes that examiners are sufficiently qualified and that their motivations are otherwise unaffected. As such, we inherently assume that examiners will conduct their examination practices in line with proper patentability standards. Anecdotal evidence suggests, however, that examiners are in fact under substantial time pressure in making patentability determinations.⁸ Binding time constraints may force examiners of this otherwise competent disposition to decrease the degree to which they search prior art, decrease their ability to extend meaningful obviousness- and novelty-based rejections and thus increase the propensity by which they grant patents, considering that they must allow applications upon the failure to find proper bases of rejection. We surmise that examiner promotions of the variety that decrease the amount of time expected to review applications will only tighten these constraints and intensify such outcomes.

III. DATA AND DESCRIPTIVE STATISTICS

Most prior investigations into the determinants of examiner behavior have explored only issued patents (for example, Cockburn, Kortum, and Stern, 2003). Among other things, a sampling frame of this nature is insufficient to capture arguably the most important decision that an examiner must make: whether or not to grant the given patent application. Moreover, when prior studies have considered application-level data, they have done so only with respect to a subset of applications at one snap shot in time,⁹ which is insufficient to account for sources of

⁸ In an August 2010 report commissioned by the Patent Office to reassess the schedule by which they set examination-time expectations (which we obtained pursuant to a Freedom of Information Act Request), the Manhattan Strategy Group stated the following:

Examiners consistently expressed the need for additional time. This was stated mostly in concern to not being able to do a high-quality examination and to avoid taking short-cuts. As one examiner in [Technology Center] 1700 explained, “when you add it up its not enough time to do a proper job on a case.” A junior examiner expressed a similar sentiment, stating that “rather than doing what I feel is ultimately right, I’m essentially fighting for my life.”

⁹ For example, Lemley and Sampat (2012) consider only 10,000 applications filed in January of 2001.

examiner heterogeneity that may bias the analysis. To overcome these deficiencies and to facilitate a rich examiner-fixed-effects design, we collected data on all 1,956,493 million utility patent applications filed on or after March, 2001 that were published by July, 2012 from the Patent Office's Patent Application Information Retrieval (PAIR) database.¹⁰ By the end of 2012 49 percent of these applications had resulted in patents, 25 percent were not patented because they had been abandoned by the applicant,¹¹ and the remainder were still pending. Applicants may elect to abandon their applications when they are unable to overcome an examiner's rejection or for other reasons, such as when a business becomes insolvent or decides to change its research direction. Our study focuses on the 1.4 million utility patent applications filed from 2001 onwards that received a final disposition—those that were granted or abandoned—by July, 2012.

Though especially rich in content, the PAIR database is not readily suitable for a comprehensive analysis of granting practices considering that the data is divided into separate webpages for each individual application, with each webpage providing information via numerous tab delimited and portable document format (pdf) files. Because of the nontrivial nature of this data collection we utilized the National Center for Supercomputing Applications at the University of Illinois to aggregate information contained across the nearly 2 million different webpages. Specifically, we collected information on the status of the application as well as other

¹⁰ In November 2000, there was a change in the law that required newly filed patent applications to be published 18 months after they were filed. 35 U.S.C. § 122(b). Applicants abandoned within the first 18 months of filing, Id. § 122(b)(2)(A)(i) and applications wherein the applicant filed a special exemption to maintain confidentiality are exempted from this requirement, Id. § 122(b)(2)(B)(i). Such applications are thus absent from the PAIR database. When some or all of an applicant's claims are not allowed by the Patent Office, the aggrieved party will sometimes file a continuation application. This application is given a new serial number and may be assigned to a different examiner. Continuation applications are treated as unique applications in the PAIR database. A related and now far more commonly used device, known as a Request for a Continued Examination (RCE), does not receive a new application serial number and effectively allows an aggrieved applicant to keep the application on the examiner's docket for further prosecution. RCEs are not treated as new, unique filings in the PAIR database; rather, they are treated as a continuation in the prosecution of original applications.

¹¹ A small portion of these applications were actually abandoned after being allowed by the examiner.

information about the prosecution process, including the patent examiner charged with reviewing the application, the basis of any rejection associated with the application (e.g., obviousness), whether the application was filed by a large or small entity, the technology class the application was assigned, and the duration to disposition (i.e., the time from filing to when the application was disposed).

Critical to our analysis is determining the experience (in years) and the GS-level of the examiner tasked with reviewing each application. For these purposes, we match the examiner field in the PAIR data with the two sets of examiner rosters received pursuant to separate FOIA requests. First, to calculate the relevant examiner's experience, we take the difference between the year at the time of disposal of the application and the first year at which the examiner joined the Patent Office, as determined by observing when each examiner was first represented in annual examiner lists that we received from the Patent Office. To ensure that this approach accurately captures the experience of long-tenure examiners, we began collecting these annual rosters in 1992—that is, nearly ten years prior to the commencement of our sample period. Naturally, this cannot ensure complete precision in the experience assignment given that some examiners may have joined the Patent Office long before 1992 (making it difficult to distinguish between 10-year examiners and 20-year examiners for those applications disposed of in 2002). To alleviate these final censoring concerns, we simply focus the empirical analysis on those examiners who joined the Office in 1993 and beyond. Pursuant to a second FOIA request, we received an additional set of annual rosters from 2001 to 2012 indicating the GS level associated with each examiner on staff over those years. Furthermore, a third and final FOIA request allowed us to determine whether GS-13 examiners did or did not have partial signatory authority at that time, a distinction, as above, that bears on the hours allocated to the examiner for review.

TABLE 2: REPRESENTATION OF GS-LEVEL AND EXPERIENCE GROUPS

| | (1) | (2) |
|----------------------------------------|----------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|
| GS-level & Experience Group | Percentage of Applications Disposed of by Examiner in Indicated Group (%) | Percentage of Total Examiner Years Spent in Indicated Group Between 2002 and 2012 (%) |
| GS-7 | 1.5 | 5.0 |
| GS-9 | 4.8 | 9.9 |
| GS-11 | 8.3 | 12.1 |
| GS-12 | 11.1 | 13.4 |
| GS-13 | 16.0 | 17.1 |
| GS-13, partial signatory | 14.8 | 12.3 |
| GS-14 | 43.4 | 30.1 |
| 0-1 Years | 8.8 | 19.8 |
| 2-3 Years | 15.2 | 19.0 |
| 4-5 Years | 15.9 | 15.7 |
| 6-7 Years | 12.8 | 11.6 |
| 8-9 Years | 10.7 | 7.6 |
| 10-11 Years | 10.8 | 8.6 |
| 12-13 Years | 10.1 | 6.7 |
| 14+ Years | 15.7 | 11.0 |

The examiner field in PAIR had a number of typographical errors and variations in the spelling or formatting of names, complicating efforts to perform the above matches. To overcome this difficulty, we utilized the reclink Stata module, a “fuzzy” matching program designed to deal with variations in names over time (e.g., inclusions of full middle names versus middle initials, name changes upon marriage, etc.). Having performed this match, we then ensured the creation of a stable set of examiner field effects. All told, our analytical file contained roughly 9,000 examiners. In our analysis we treat the individual who did the majority of work on the application as the examiner charged with reviewing that application: the non-signatory examiner, when both a non-signatory and an examiner with signatory authority are

associated with an application, and the signatory examiner, when only one examiner is associated with an application.

In Columns 1 of Table 2, we set forth the percentage of applications in our sample that are disposed of by examiners in each of the relevant GS-levels and experience groups. Examiners spend considerably more time in higher GS ranges, especially GS-level 14, thus accounting for the higher percentage of applications associated with high GS-level examiners. Also contributing to the relatively weaker presence of GS-levels 7 and 9 in the data is the fact that many examiners (nearly 1/3 of new examiners) begin at GS-level 11. In Column 2, we further illuminate this breakdown by taking all of the examiner rosters over the 2002–2012 period and indicating what percentage of these total examiner years were represented by examiners in the various GS-level and experience categories. For the reasons just discussed, this representation also tends to be weighted near the higher GS ranges.

For each of the roughly 1.4 million patent applications that reached a final disposition in our sample period, we relate examiner characteristics, including their pay grade and experience level, to whether or the application was granted, our key outcome of interest. All told, 70 percent of the applications disposed of over this time period were granted.¹² To form our second set of outcome measures, we determine whether the given application had at least one claim rejected during examination based on each of the following statutory bases: 101 (lack of patentable subject matter or utility), 102 (lack of novelty), 103 (obvious), and 112 (failure to meet written description requirements). To the best of our knowledge, we are the first to report the bases of

¹² As stated previously, continuation applications, as distinct from the now more common RCEs, are counted as a rejection / abandonment of the original application and the filing of a new application within the PAIR database (RCEs, which keep the same serial number and stay with the same examiner, are not treated as new applications). Accordingly, this 70 percent rate does not necessarily capture the percentage of original applications that are ultimately allowed considering that some continuation applications may successfully issue. It is important to note that this is merely a classification question—i.e., do these events contribute or not contribute to the Patent Office’s grant rate? Our focus is largely on exploring the relationship between the grant rate, however it is defined, and certain characteristics of the examiners.

rejections for any substantial sample of patent applications. We utilized an iterative process comprising a mix of programming and hand inspection to develop a comprehensive list of phrases that examiner's utilize when making rejections. We then searched all office actions for these phrases to flag for the presence of each of the indicated rejection types.¹³ The likelihood that a given application received each of the indicated types of rejections in the sample are as follows: (1) 10 percent, lack of patentable subject matter or utility; (2) 56 percent, lack of novelty; (3) 71 percent, obvious; and (4) 36 percent, failure to meet written description requirements.

To form a final outcome measure that captures how intensively examiners are searching for prior art, we focus on the sample of patents issued over the above-specified time period (as distinct from the sample of applications over this time period) and likewise tasked the National Center for Supercomputing Applications at the University of Illinois with collecting information on the number of prior art references listed in each issued patent that emanate from the examiner rather than the applicant, along with the share of all prior art references attributable to the examiner. Since 2001, the front page of issued patents identifies prior art references by their

¹³ The following phrases were searched for rejected under; rejected are under; rejected as unpatentable; as being unpatentable; rejected as failing to define; objected to; election of species; fails to define a statutory; antecedent basis; new title is required; title of the application will; notice to comply; part of paper; prior art made of record and; rejected as being based; rejected as being directed; rejected on the ground; restriction to one of the fol; restriction is required under; status identifiers; fail to meet; fail to comply; fail to contain; fail to provide; fail to identify; fail to include; do not comply with; not in accordance with; cannot be patented; defective because; non-compliant because; renders the claim indefinite; not of sufficient quality; filed after the issue fee was; filed after publication; drawings in compliance with; declaration is missing; are not consistent; is not a proper submission; not include a support; claim rejections; this is a provisional obvious; because it is unsigned; not filed a certified copy; is non-responsive; required to furnish; introduce new matter; not contain a concise explan; the following omission; request for information; requirement for information; abstract exceeds 150 words; elect a single disclosed spec; elect disclosed species; not properly annotated; not signed by an; not authorized to sign; not been submitted; not appear to be relevant to; non-elected subject matter in; terminal disclaimer needs to; associate poa are no longer; include common ownership as; other copending United States; application conflict with cla; contain every element of cla; believed to interfere; has not been met; not indicated the appropria; contain(s) every element of; claimed invention is differe; contains every element of cl; declaration in compliance wi; does not have publication da; do not have corresponding pa; filed well after the applica; list of all patents. Publica; notice of non-compliant amen; reference relevant to the ex; required information for the; requires that the summary of; restriction is hereby require; the appropriate statement ac; Website is not considered a.

origin: examiner found or applicant disclosed. Previous investigations have reported that examiners are more likely to rely upon prior art they discovered during their own search, rather than art disclosed by an applicant, to reject a patent application (Cotropia, Lemley, and Sampat, 2010). Additionally, several studies have used the share of references listed in an issued patent originating from the applicant or examiner as a proxy for the extent to which the party in question (examiner or applicant) searched the prior art (Lemley and Sampat, 2012; Alacer et al., 2009).

IV. METHODOLOGY

To explore how patent examination practices change upon promotions that leave examiners with less examination time, we estimate the following specification:

$$GRANT_{aikt} = \alpha + \gamma_i + \lambda_t + \theta_k + \beta_1 (GS_{it}) + \beta_2 EXPER_{it} + \beta_3 X_{aikt} + \varepsilon_{aikt} \quad 1)$$

where a indexes the individual application, i indexes the individual examiner, k indexes the technology associated with the application and t indexes the year in which the application is disposed of by the examiner. $GRANT_{aikt}$ indicates whether or not the given application was allowed by the examiner. GS_{it} represents a set of dummy variables capturing the incidence of the examiner assigned to the underlying application falling into each of the general schedule (GS) pay-grade levels. In most specifications, as discussed further in Section V, this variable also includes separate categories for GS-13 without partial signatory authority and GS-13 with partial signatory authority, considering that this unique within-GS-level promotion likewise carries with it reductions in examination-time expectations. The existence of this within-GS-

level variation may allow us to rule out concerns that the results are merely driven by things that change with GS-level promotions other than examination-time allotments.

Furthermore, \mathbf{EXPER}_{it} captures a set of dummy variables for the incidence of the relevant examiner falling into a range of experience-level categories, where experience is signified by the number of years at the time of the application's disposition that the relevant examiner has been with the Patent Office. In other specifications, as discussed in detail in Section V, we nest experience within GS-level categories and thus create dummy variables capturing a series of experience categories within each GS-level. Including experience levels are not just important for our analysis in order to decouple an experience effect from an hours-reducing promotion effect. Even if no GS-level promotions occur, increases in years of experience are often generally associated with pay-raises of their own (known as "step" increases). By observing how the grant-rate response to changes in experience differs from the grant-rate response to changes in GS-levels, we may be able to achieve greater confidence in interpreting the results as stemming from the reductions in examination hours associated with GS-level promotions as opposed to the increases in pay likewise associated with such promotions.

In certain robustness checks, we include a set of technology fixed effects, $\mathbf{\theta}_k$, to alleviate concerns that examiners may be reassigned over time to different technologies as they ascend to higher pay-grades, along with certain individual characteristics of the applications, \mathbf{X}_{aikt} , including the entity size status of the applicant (large versus small) and the length of time being the filing and the disposition of the application (and its square).

Importantly, a set of examiner fixed effects are captured by $\mathbf{\gamma}_i$. Such fixed effects help address concerns that more experienced examiners and higher GS-level examiners are

fundamentally different from their more junior counterparts, for reasons beyond mere differences in seniority and promotion levels—e.g., concerns that examiners who have reached higher grade levels and thus who have been successful in attaining promotions may be those with a stronger inherent disposition towards granting in the first place, along with concerns that more experienced examiners may also differ from less experienced examiners simply because they elected to stay at the Patent Office.

On a final methodological note, we exclude applications examined by individuals at GS-level 15 from the analytical sample, though we note that this choice is not of much significance for our results, as demonstrated by the Online Appendix. Only a small percentage of the original PAIR sample was examined by GS-15 examiners (roughly 3 percent, as compared with the nearly 40 percent of the sample that are examined by GS-14 examiners). This is, in part, due to the fact that GS-15 examiners spend little of their time examining patents, having generally reached a point of supervisory status. The possibility that their predominantly supervisory duties may alter the manner in which time constraints impact the examination decisions that they do make may complicate our modeling of examiner behavior in the face of career progressions that modify examination-time expectations.¹⁴

V. RESULTS

¹⁴ Another key concern driving this exclusion relates to our uncertainty regarding the examination time expected of GS-15 examiners in our sample. According to the examination-time scaling schedule provided by the Patent Office, GS-15 examiners may either receive 67 percent or 71 percent of the time extended to their GS-12 counterparts (the reference group). Hoping that this would provide another within-GS-level degree of variation, we filed an additional FOIA request with the PTO asking for the specific examination-time allotments associated with each GS-15 examiner, for each year from 2002-2012. The provided data, however, list the majority of such GS-15 examiners as having the same amount of time allocated to GS-12 and GS-13 examiners, as opposed to the anticipated 67 percent and 71 percent values. In the Online Appendix, we estimate specifications that treat all GS-15 examiners alike under the assumption, as expected from their published schedule, that GS-15 examiners would receive even less time than their GS-14 counterparts. However, the uncertainties in the data received pursuant to this final FOIA request leaves us inclined to treat this as a supplementary exercise only. On a final note, only 0.2 percent of the original sample were examined by GS-5 examiners. Given such a small level of representation, we exclude these applications from the analysis, though again we note that this exclusion is of little significance to our findings.

A. *Grant-Rate Analysis*

1. Primary Results

We begin our exploration into the effects of allotted examination time—as identified by changes in examiner pay-grade level—on the practices of patent examiners by plotting the evolution of grant rates observed over the course of a given examiner’s career as they rise in the ranks. More specifically, in Figure 1, we plot results from a regression of the incidence of an application being granted on a set of dummy variables capturing each of the relevant examiner pay grades (7, 9, 11, 12, 13, and 14), in addition to a set of year fixed effects and examiner fixed effects.¹⁵ This figure suggests that as an examiner moves from GS-level 7 to GS-level 9, they increase their grant rates by 2.9 percentage points (or by roughly 4 percent). As the examiner ascends even higher in ranks and thus as the examiner receives less and less time to review her applications, this increase in grant rates continues monotonically until the point at which her grant rate at GS-level 14 is nearly 17.9 percentage points (or nearly 26 percent) higher than it was when she was at GS level 7.

Of course, one cannot simply rely on the relationship between grant rates and GS-level dummies alone in order to infer that the reductions in time allocations stemming from promotions are inducing higher granting proclivities. After all, promotions will naturally correlate with additional years of experience within the Agency. As such, one might be concerned that the above figures are merely capturing an experience effect—for example, a story in which patent examiners become better over time in identifying patentable subject matter. We begin to confront these concerns and to separate the influence of promotion from the influence of

¹⁵ Standard errors are clustered at the examiner level to account for autocorrelation over time in examiner-specific residuals. Given computational considerations in light of the over-1-million observations and nearly 9,000 examiner fixed effects, we elect to estimate linear probability models throughout. We note, however, that the pattern of results we present are virtually unchanged when we instead take a 10-percent random sub-sample of examiners and estimate conditional logit specifications.

experience by including alongside the GS-level dummies a set of dummy variables capturing the duration of experience of the examiner associated with the given application.

Fortunately for these purposes, examiners do not always receive promotions lockstep with experience. This is especially true at higher grade levels within the Agency. Over 75 percent of examiners who have reached the stage of GS-level 14 stay at that grade level over a year. In fact, at least 20 percent of such examiners stay at GS-level for at least 8 years. On average, GS-14 examiners stay at that grade for over 4.5 years. On the other hand, only 16 percent of examiners who have been at GS-level 7 stay at that grade beyond 1 year. It is perhaps not until GS-level 12 when examiners begin to routinely spend longer than 1 year at the respective grade. As such, the ability to reliably isolate the independent impacts of promotion is greater as we focus on the upper end of the promotion schedule.

In Figures 2 and 3, we present results from this separation specification. Though the specification simultaneously includes dummy variables capturing both promotion and experience, we divide the results into two figures for ease of presentation. The trend presented in Figure 2 presents the coefficients of the GS-level dummies. The corresponding trend in Figure 3 presents coefficients of the experience group dummies, capturing the cumulative effect on grant rates (relative to the baseline experience category) of a given examiner increasing her experience duration. Importantly, this exercise demonstrates that the pattern of GS-level dummy coefficients presented in Figure 1 is robust to the inclusion of flexible, non-parametric controls for the years of experience of the examiner, though the magnitude of the increases in grant rates upon promotion are less pronounced once including experience controls. We find that once examiners ascend to GS-level 14, their grant rates are nearly 8 percentage points higher than they were at GS-level 7.

One could theorize numerous reasons as to why examiners might alter their granting practices over time alone. For instance, informational deficiencies may cause them to cautiously over-scrutinize early in their careers, only to find themselves better equipped with time to determine the validity of the applications under review. On the other hand, it could be the case that examiners simply lessen their scrutiny as time goes by in the Patent Office due to fatigue or due to an increased tendency to shirk. To the extent that any such stories are even present in the first place—which we address more directly below—the findings demonstrate a distinct jump in grant rates that occurs upon GS-level promotion independent of any flexible pattern of grant rates that examiners exhibit over time itself. Considering that the key channel by which the *act* of promotion may theoretically impact subsequent examination behavior stems from its effect on the time allotted to examination (as discussed above), these results provide greater confidence that (1) time constraints may be binding on examiners and (2) that tightening such constraints may leave examiners with less time to adequately challenge the patentability of applications, thus leading to more permissible granting behaviors.¹⁶ We further support this contention below with even richer methods of decoupling experience from promotions and with investigations into examiner search efforts and rejection patterns. Beforehand, however, we briefly discuss the pattern of results in Figure 3 bearing on the independent relationship between an increase in an examiner’s experience in years and her grant rate.

¹⁶ We acknowledge that some examiners may attempt to increase their chances of promotion by granting more permissively as a general matter of course, either because such behavior may facilitate the processing of a greater number of applications or in light of the financial interests of the Agency in over-granting patent (Frakes and Wasserman 2013). We stress that our results are not merely picking up such behavior given that each specification includes examiner fixed effects. In other words, our methodology is designed to place inherent granting tendencies aside—including those stemming from promotion-seeking behavior—and instead focus on within-examiner changes in behaviors over the course of a career. For a story of this nature to explain the results, it would have to be the case that promotion-seeking behavior elevates in intensities upon each promotion.

2. Experience Effects

This relationship between examiner experience and granting patterns was recently examined with care by Lemley and Sampat (2012), who found that grant rates increase monotonically with seniority. While they acknowledged the possibility that their findings could be attributable to changes in time allotments upon promotions—i.e., the focus of the present study—they did not have the ability to decouple experience from other factors to allow them to draw any such inference. In addition to this methodological limitation, we further note that Lemley and Sampat’s experience effect itself does not remain monotonic once one more flexibly accounts for examiner heterogeneity. Their analysis was largely cross-sectional and could not fully account for the possibility that the results are driven by selective retention—i.e., that those who stayed with the Agency longer and thus formed the senior group were of a different disposition.¹⁷ As presented in the Online Appendix, we do replicate this monotonic relationship when we likewise take a cross-sectional approach that includes only year and experience dummies.¹⁸ However, as demonstrated by Figure 3 and by [Figure A__] in the Online Appendix (where we include only experience dummies and thus exclude GS-level dummies), the relationship between examiner experience (in years) and grant rate instead follows an inverse-U pattern once one includes examiner fixed effects. Controlling for grade level dummies, grant rates do increase by close to 7 percentage points as an examiner moves from 0-1 years of experience to 2-3 years of experience. The grant rate effectively stays at this level through 5

¹⁷ To partially address selection concerns, Lemley and Sampat (2012) did, however, include as a control a dummy variable for whether or not the examiner associated with the given application ultimately stayed with the Agency for at least five years. This approach cannot account for as many sources of examiner heterogeneity as can be provided by an examiner fixed effects specification.

¹⁸ Moreover, we note that the monotonically increasing pattern of results from this cross-sectional specification remains nearly unchanged with the inclusion of a control variable for whether or not the examiner associated with the application ultimately stays at least five years, following Lemley and Sampat (2012).

years of experience and thereafter begins to fall, until the point at which the grant rate at 14+ years of experience is identical to the 0-1 year experience level.

3. Within-Grade-Level Change in Time Allotments

The primary data source employed above does not distinguish among those examiners at GS-level 13. As explained in Section II, however, many examiners are granted partial signatory authority during the course of their tenure at GS-level 13, a unique within-grade promotion that likewise comes with a reduction in the examination time allotted to the affected examiners. To identify these examiners, we filed an additional FOIA request with the PTO. In this subsection, we build upon the above analysis and include separate dummy variables for whether the given application was disposed of by a GS-level 13 examiner without partial signatory authority and with such authority. There were slight inconsistencies in the treatment of within-year promotions between the examiner roster data employed above and the secondary roster of GS-13 examiners we received pursuant to our second FOIA request. That is, many of those examiners receiving this within-GS-13 promotion were registered as doing so during the year in which our primary roster indicated that they ascended to GS-level 14. In this alternative approach, we give priority to the timings of promotions set forth in this secondary data source. Given the inconsistencies in the timing of promotions across the separate data sources, we elect to maintain transparency by presenting these results separately from the results set forth above.

We present the results of this alternative exercise in Figure 4. Encouragingly, we find that grant rates increase with each promotion from GS-level 11 onwards, including increases as given examiners initially ascend to GS-level 13 and subsequently ascend to GS-level 13 with partial signatory authority. In this alternative specification, grant rates do fall slightly upon promotions from GS-level 7 to 9 and from GS-9 to GS-11. As discussed above, however, the

ability to reliably isolate the independent effect of promotion (decoupled from experience) is stronger in the later GS ranges.

The ability to draw upon a within-GS-level change in the time allotment extended to examiners provides us with a welcome opportunity to challenge the argument that the results may be purely driven by factors changing with GS-level promotions other than examination time allocations. If that were the case, one would not expect to observe an increase in grant rates upon a promotion of a non-GS-level variety that also corresponds to a change in examination workload expectations.

As discussed in Section II, once an examiner reaches the second GS-13 classification and GS-level 14, she attains more authority of her own to sign off on decisions, thus representing a decline in the level of scrutiny placed on her by her superiors. One may be concerned that the increase in grant rates observed upon promotion are merely a reflection of this lightening of scrutiny. However, the fact that this pattern of increasing granting tendencies occurs over earlier promotions, which do not come with the extension of greater authority and less oversight, lends support to the idea that the documented pattern of results may stem from the allotted-hours reduction associated with these promotions and not simply from changes in the degree of oversight.

4. Within-Grade Experience Effects

In this sub-section, we take an alternative approach to separating the effects of grade level (and thus allotted examination time) on the one hand and experience in years on the other hand. Instead of simply estimating the overall impacts of being at the PTO for a given number of years, we nest experience years within grade level changes. In other words, we estimate specifications that include a series of dummy variables capturing the presence of specific years

within specific grade levels—e.g., 0-1 years in GS-13, 2-3 years in GS-13, 0-1 years in GS-14, 2-3 years in GS-14 etc.¹⁹ A specification of this nature allows us to more comprehensively follow the course of a hypothetical examiner over the various stages of a career and thus better visualize the independent impacts of grade-level promotions (in addition to the within-grade promotion of interest at GS-level 13 that likewise leads to an examination workload adjustment). For this analysis, we focus only on those examiners in GS-12 and above considering that the majority of those within lower grade levels achieve promotions within their first year at those grades, providing little ability to reliably track the evolution of grant rates over years while at GS-7, 9 or 11.

Figure 5 plots the results of this exercise, presenting the coefficients of each of these separate dummy variables, with the 0-1 year period at GS-12 serving as the omitted reference group. The results of this exercise only further solidify the contention that examination practices change upon the occurrence of career events with respect to which the time allocated to examiners is reduced. Upon each such promotion, the observed grant rate jumps. Importantly, these promotion-level increases do not appear to be mere reflections of continuing trends in grant rates over the duration of an examiner's tenure at the PTO, which might otherwise suggest a simple experience-level story. This is perhaps most convincing upon viewing the experience of GS-14 examiners considering the fact that examiners spend the longest at this grade level. While the grant rate jumps distinctly once one enters this GS level (to a level that is 8 percentage points higher than the reference period), the grant rate actually begins to fall thereafter. By the time a GS-14 examiner reaches her 5-6th year at that level—a milestone that occurs for about 45 percent

¹⁹ Given that the roster data we obtained is only at the year level, it is difficult for us to identify when during the year an examiner experiences a promotion. For this reason, we hesitate to make much of any effects at 12-month increments and instead group time into 24-month increments. Nonetheless, this choice is of little difference, as the corresponding figure when focusing on annual changes is nearly identical.

of examiners who have reached GS-14—her grant rate has fallen by 2 percentage points below the rate she applied in her first year at GS-level 14. In the period represented by her 9th year and beyond—a point reached by nearly 20 percent of those examiners who have reached GS-14 at all—her grant rate is 8 percentage points below the initial GS-14 grant rate. If the grant rate had incrementally continued to rise over such years, especially at levels commensurate with those experienced upon grade level changes, it would instill less confidence in an interpretation of the results as emanating from reductions in the amount of time at the disposal of examiners.

Indeed, if anything, this picture depicts a story in which experience (in years) alone ultimately corresponds to a *reduction* in granting tendencies, standing in stark contrast with the positive relationship documented in Lemley and Sampat (2012). With respect to each of the four given promotion categories considered in Figure 5, the grant rate ultimately begins to fall over time as one stays within the respective category long enough.²⁰ These drops in grant rates with experience are periodically corrected by successive promotions of the sort that leave examiners with diminished examination time. If anything, the declines in grant rates observed over the temporal dimension of Figure 5—that is, over the increases in years within the various grade levels—perhaps suggest a story in which examiners in general learn over time how to form more effective bases of rejection (thus contributing to falling grant rates), only to have this learning

²⁰ To be sure, our identification of these declining grant rates over time within lower GS-level groups is drawn from those examiners that happen to stay many years at such grade levels. Since only a minority of examiners at GS-levels 12 and 13 will stay at those levels over the course of many years, one may be concerned that their experiences do not generalize to that of everyone. As such, perhaps the most conservative way to interpret our results is that with respect to at least some examiners—that is, those that happen to achieve promotions relatively more slowly—the effects of increased temporal experience on grant rates appears to generally be negative (with occasional spikes upwards upon promotion). For those other examiners that experience early promotions more rapidly, they will at least stay for a long time at GS-14, at which point their grant rates do indeed fall with more years of experience. What is difficult to say with confidence is what role experience plays as distinct from GS-level changes for these quick risers during the earliest years of their career. For those examiners, it could indeed still be that experience effects follow an inverse U pattern as opposed to a universally negative pattern.

process interrupted by occasional promotions that diminish the amount of time they have to derive such rejections (thus re-elevating grant rates).

In discussing Figure 5, it also bears mentioning that examiners may continue to receive salary increases throughout their tenure at each GS-level. Such increases occur as they are promoted to different “steps”—e.g., Step 1 at GS-12, Step 2 at GS-12, etc. “Step” increases generally transpire with increases in experience over time, as distinct from merit based promotions, and generally entail a meaningful increase in salary level. For instance, a GS-level 14 at “Step 5” is paid \$128,941 while a GS-level 14 at Step 10 is paid \$147,900. The presence of such alternative types of promotions—that is, within-GS-level increases in salary that are tied only to experience—are further helpful for our analysis in providing support against an argument that the primary findings set forth above are attributable merely to a story in which the higher incomes associated with GS-level promotions are somehow inducing examiners to adjust their granting behaviors. If such a story were driving the results, one would further expect to observe increases in grant rates over the course of years while in specific GS-levels, especially GS-level 14, where examiners stay many years on average. As such, these findings supplement those of the previous sub-section, which separately analyzes the two types of GS-13 examiners, in supporting the view that it is the timing expectations associated with the promotions of interest in this analysis that is likely attributable to the increase in observed grant rates as distinct from other factors (e.g., salary) that may be changing upon promotion.

5. Robustness Checks

In the Online Appendix, we challenge the above grant-rate results through a range of robustness exercises. For instance, we demonstrate that the above results remain virtually unchanged when we include variables in the underlying specifications for the entity size status of

the applicant (large or small entity) and for the duration (in days) of the period between filing and final disposition of the application, along with the square of this duration. The results are also not affected by the inclusion of technology class dummies to rule out concerns that the results may be a reflection of examiners switching to different technologies (i.e., those with inherently high application qualities) as they are promoted.

The primary specifications estimated above use data from all applications disposed of between 2002 and 2012. We exclude 2001 since the PAIR only began collecting application data for filings commencing in March of 2001, leaving very few final dispositions of such applications in 2001. Given that we only observe applications filed after this date, one may be concerned with some level of imbalance in the sample. The essence of this concern is that applications reaching a final disposition in the early years in the sample will be disproportionately comprised of quicker moving applications, whereas those observed in the later years in the sample represent a richer mix of quick- and slow-moving applications. This may be of consequence considering that prosecution durations may impact grant rates due to the higher likelihood of applicants abandoning their applications during long durations.

However, an increased incidence of longer-duration prosecution periods later in the sample does not necessarily confound the above analysis considering (1) the imposition of year fixed effects to capture any general trends in granting practices, (2) that we observe overlapping cohorts of examiners, in which event examiners are moving among each of the various grade levels during every year of the sample (as opposed to a situation in which we only track a new cohort of examiners over time) and (3) that controls are available for the time between filing and disposition of each application. Nonetheless, to more comprehensively address any inconsistency in the set of applications under investigation, we also estimate an alternative

specification that imposes a more balanced sample over time. In this alternative approach, we begin the period of observation in 2004 and confine the sample of applications to those that are disposed of within a three-year period. By focusing only on applications of limited prosecution duration, we ensure consistency in the relative mix of application durations observed over the full sample period. Of course, imposing this duration limitation forces us to start our analysis in a year in which we at least begin to see application durations close to this cutoff—i.e., 2004. In other words, we exclude 2002 and 2003 in this alternative approach as there will be few applications disposed of in these years that fall near the 3-year duration mark, despite the fact that we would observe more of such applications in the later sample years, which could otherwise undercut the balancing impulse of this exercise. In the Online Appendix, we demonstrate that the various patterns of results presented in Figures 1–5 are likewise robust to this alternative sampling approach.

B. Analysis of Rejection Patterns

1. Obviousness

Remember that the decision to grant a patent can be viewed as the default decision one must make upon failure to find a basis for rejection. The ability to reach such a rejection, we contend, is likely a function of the amount of time extended to examiners. As discussed in Section II, some forms of rejection may be more sensitive to timing concerns than others—e.g., rejections contending that the proposed claims are obvious in light of existing knowledge. Accordingly, we now attempt to illuminate the above grant-rate findings by exploring associated effects of promotions on the incidence of obviousness and related rejections among the underlying applications. A key prediction set forth above is that examiners will begin to perform fewer and fewer rejections based on the argument that the proposed claims are obvious—an

especially time intensive analysis—upon the occurrence of promotions that leave them with less and less allocated examination time.

To first explore this hypothesis, we estimate specifications similar to that employed above but using the incidence of an obviousness rejection as the relevant dependent variable. Figures 6 and 7 present results of this exercise, replicating Figures 2 and 4 above, respectively—that is, each controlling for experience-level dummies, but taking separate approaches to the treatment of GS-13 examiners (with Figure 7 exploring the separate effects of moving to GS-13 without signatory authority and subsequently to GS-13 with such authority). Though the incidence of an obviousness rejection appears to rise from the promotion of GS-level 7 to GS-levels 9 and 11, it falls with each examination-hours-reducing promotion thereafter. Bearing in mind again that the reliability of disentangling the effects of such promotions from that of experience is far weaker in the lower GS ranges, where promotions are more typically lock-step with time, these findings are generally consistent with the expectations of a reduction in obviousness rejections upon the promotions of interest.

One limitation of the data that we collected for this analysis, highly novel though it is, is that we simply capture the incidence of any obviousness rejection without knowing the full force of such rejection. Does it simply cover one claim or many claims? Is it easy to overcome or difficult? Such questions cannot be adequately resolved with the data collected. With this limitation in mind, we take an alternative approach where we do not view obviousness rejections in an absolute sense, but instead specify the dependent variable as the ratio of obviousness rejections to total rejections, more specifically the incidence of an obviousness rejection divided by the sum of the incidence of the following types of rejections: obviousness, novelty, patentable subject matter, and the disclosure requirements. Though each of the variables underlying this

ratio suffer from the above limitation, this measure at least provides us with a sense of the relative effort spent on obviousness rejections. In Figures 8 and 9, we replicate the previous two figures using this alternative dependent variable. The results paint an even stronger picture that examiners begin to spend less and less of their efforts on time-intensive obviousness analyses upon promotions that leave them with less and less time at their disposal to conduct examinations, in this instance with less pronounced evidence of an increase in obviousness activity in connection with the early promotions.

In Figure 10, we follow the approach of Figure 5 and track how the incidence of an obviousness rejection evolves as an examiner increases in tenure over time within each of GS-level 12, GS-13 without partial signatory authority, GS-13 with partial signatory authority and GS-14. This more comprehensive specification presents arguably the most compelling evidence of a promotion-level impact on obviousness patterns. Complementing Figure 5, which demonstrated a general trend to decrease grant rates over time within given grade levels, Figure 10 demonstrates a corresponding tendency over the time dimension to increase rates of obviousness rejections. This may even be consistent with a learning story in which examiners get better and better at forming obviousness determinations over time. Periodically, however, examiners will experience promotions that cut short the time they have to make such rejections. Figure 10 suggests that upon such instances, the rates at which they are able to form obviousness rejections fall back down. As such, Figure 10 collectively implicates a positive relationship between experience and obviousness patterns and a negative relationship between examination-time-reducing promotions and the incidence of obviousness rejections.

2. Novelty

As stated above, obviousness rejections require an additional analytical step beyond that required for a rejection analysis based on a lack-of-novelty theory, thus representing a potentially more burdensome exercise. Nonetheless, novelty analyses do require prior art search efforts of their own, which are likely to be more time intensive exercises than non-prior-art based rejections such as those contending that the proposed claims do not cover patentable subject matter. Leaving aside obviousness rejections, one may thus wonder whether we observe a similar pattern of declining novelty rejections upon the promotions of interest, relative to the number of non-prior-art-based rejections. Building off of Figure 8, we present the results of this relative-novelty specification in Figure 11, focusing, for the purposes of brevity, only on the specification in which we include both types of GS-13 examiners. The results are largely in line with expectations that time constraints may also be crowding out the relative time spent on novelty analyses, with Figure 11 documenting a decline in rates of novelty rejections relative to non-prior-art rejections upon the relevant promotions (though showing a slight increase in connection with the first promotion and a leveling out at the end of the promotion scale).²¹

As such, our findings demonstrate a decline in efforts spent on both obviousness and novelty analyses. Considering that the vast majority of patents that are ultimately invalidated in court are done so on the basis of lack-of-novelty or obviousness, with the production of new prior art, these findings lend further support to a story in which the additional patent grants stemming from the promotions of interest documented in Figures 1–5 are of a marginally less valid nature.

²¹ As demonstrated by the Online Appendix, however, the overall incidence of a novelty rejection remains essentially flat over the series of relevant promotions. In the Online Appendix, we likewise present results focusing on the incidence of the non-prior bases of rejections—i.e., patentable subject matter / utility and the disclosure requirements. Such incidences either stay flat or actually increase with GS-level changes, thus explaining the relative decline in novelty rejections.

C. Examination of Prior Art Citations

To further illuminate whether the above results are indeed a reflection of reduced examination effort stemming from binding examination-time constraints, we next estimate the relationship between GS-level promotions and the share of total prior-art citations listed in the final patent that are provided by the examiner (as opposed to the applicant), a proxy for the search effort of the examiner. In Figure 12, we estimate this relationship over the full sample of patents issued between 2002 and 2012. This specification is, of course, somewhat compromised by the fact that it relies only on issued patents, the incidence of which we already know (Figures 1 to 5 above) is likely to increase upon the promotions of interest, leading to possible selection concerns. With this caveat in mind, we note that the findings parallel those of the obviousness-rejection analysis above. Except for the first move from GS-level 7 to GS-level 9, we document a decline in this share upon each GS-level promotion thereafter. In the Online Appendix, we estimate this specification using the number, as opposed to the share, of examiner-drive prior art citations as the dependent variable (omitted now for purposes of brevity), documenting an identical pattern. As an examiner moves from GS-level 11 to GS-level 14, we estimate that she will offer on average 0.7 fewer citations, representing a roughly 13 percent reduction.²²

These findings support the obviousness analysis above in demonstrating that time pressures associated with promotions may be leading examiners to reduce their scrutiny. Moreover, considering that examiners must grant patents to applicants upon the failure to formulate suitable bases of rejections, these examiner-search findings support the inference above that time pressures may be leading examiners to grant more and more patents. In perhaps

²² We further note that no discernable pattern exists when we estimate the relationship between examiner grade-level promotions and the number of prior-art citations made by applicants. We likewise present such results in the Online Appendix. Such findings are encouraging and supportive of the contention of a random application assignment process.

broader terms, these findings suggest that time constraints may indeed be binding during the examination process.

VI. DISCUSSION AND CONCLUSION

Our data finds that as examiners are given less time to review applications, the less prior art they cite, the less likely they are to make time-consuming prior art rejections, and the more likely they are to grant patents. We identify the effects of reduced time allocation by observing the behavior of examiners over the course of up to six different promotions, each of which is associated with an iterative 10 to 15 percent decrease in the time allocated to review an application. While we contend the most important change upon such promotions that bears on how an examiner approaches the examination of an application is the time allotted to review applications, we also rule out whether other changes upon promotion, including salary raises and a decrease in the scrutiny associated with obtaining signatory authority, may be driving our results. Starting from the premise that examiners who have sufficient time to review applications will tend to make the correct patentability decisions, our findings suggest that the time currently allocated to patent examiners to review application is inducing examiners to allow arguably invalid patents on the margin.

Setting the time allocated to review patent applications is undoubtedly complicated and at a minimum involves a trade-off between patent quality and examination capacity. A Patent Office whose sole objective is to maximize patent quality would set the hour allotments much higher in order to ensure that examiner error was minimized. The Agency, however, also seeks to provide timely review of applications. As a result, the Patent Office must set allotments so that it also can process a sufficient number of applications in a given time period. With this tradeoff in mind, the Patent Office decision to decrease hour allotments upon promotion appears


prudent. That is, examiners who have repeatedly demonstrated their ability to provide high quality patent examination, and are rewarded for their admirable behavior by promotion, are likely to be able to complete a review of an application faster than an examiner who has yet to demonstrate this competency. Nevertheless, our results suggest that the current scaling of the time allotments upon promotion—a scaling that leaves GS-14 examiners with nearly half the time to review applications relative to GS-level 7 examiners—may be rather aggressive, especially considering that time allotments may be tight even for those at the lowest pay grades. Whether this scaling is too aggressive is perhaps beyond the scope of the present paper. A full welfare analysis would entail balancing the costs of expecting the Patent Office to raise additional resources to provide additional examination with the costs associated with an overly permissive granting stance on the part of the Office’s examiner corps. This open question aside, our results nonetheless suggest that time constraints may nonetheless be binding as they are and, concomitantly, demonstrate that such time constraints may be one source behind the much discussed over-granting tendencies of the Agency.

Additionally, our analysis highlights the inequitable nature of patent outcomes, building on the prior work of Cockburn, Kortum, and Stern (1993) and Lemley and Sampat (2012), each of which had raised concerns over the equity of the examination process. The decision to grant or reject a patent is intimately dependent, in part, upon the examiner that is by and large randomly assigned to the application. To the extent we value a patent system that rewards equally meritorious applications alike, our findings are perhaps worrisome.

REFERENCES


- Juan Alcacer, Michelle Gittelman, and Bhaven Sampat, “Applicant and Examiner Citations to Patents: An Exploration and Overview.” *Research Policy*, 2009, pp. 415–427.
- James Bessen & Michael J. Meurer, *Patent Failure: How Judges, Bureaucrats, and Lawyers Put Innovators at Risk* (2008).
- Dan L. Burk and Mark A. Lemley, *The Patent Crisis and How the Courts Can Solve It* (2009).
- Chris Cotropia, Mark Lemley, and Bhaven Sampat, “Do Applicant Citations Matter?”, working paper (2012).
- Michael D. Frakes and Melissa F. Wasserman, “Does Agency Funding Affect Decisionmaking?: An Empirical Assessment of the PTO’s Granting Patterns.” *Vanderbilt Law Review*, 2013, pp. 67–147.
- Michael D. Frakes and Melissa F. Wasserman, “Does the U.S. Patent and Trademark Office Grant Too Many Bad Patents?: Evidence from a Quasi-Experiment.” *Stanford Law Review*, forthcoming 2015.
- A.B. Jaffee and Josh 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)
- Mark A. Lemley and Bhaven Sampat, “Examiner Characteristics and Patent Office Outcomes.” *The Review of Economics and Statistics*, 2012, pp. 817–827.
- Mark A. Lemley, “Rational ignorance at the Patent Office.” *Northwestern University Law Review*, 2001, pp. 1495–1532.
- Douglas Lichtman, “Rethinking Prosecution History Estoppel.” *University of Chicago Law Review*, 2004, pp. 151–82.
- Ronald Mann, “The Idiosyncrasy of Patent Examiners: Effects of Experience and Attrition.” *Texas Law Review*, 2014, pp. 2149–76.
- Petra Moser, “How Do Patent Laws Influence Innovation? Evidence from Nineteenth-Century World’s Fair.” *American Economics Review*, 2004, pp. 1214–1236.

Figure 1: Relationship between Examiner GS Levels and Grant Rates



Notes: this figure presents results from a regression of the incidence of a granted application on dummy variables representing each General Schedule level between 7 and 14. The dummy variable for GS-level 7 is omitted, representing the reference group. The vertical bars represent 95% confidence intervals for the estimated coefficients. Regressions include examiner and year fixed effects. Standard errors are clustered at the examiner level. Examiner roster data and patent application data are from the PTO.

Figure 2: Relationship between Examiner GS Levels and Grant Rate Controlling for Experience Length Groups



Notes: this figure replicates that of Figure 1, except that it includes as controls dummy variables representing the incidence of 8 different experience (in years) groups.

Figure 3: Relationship between Examiner Experience Groups and Grant Rate, Controlling for GS-Levels

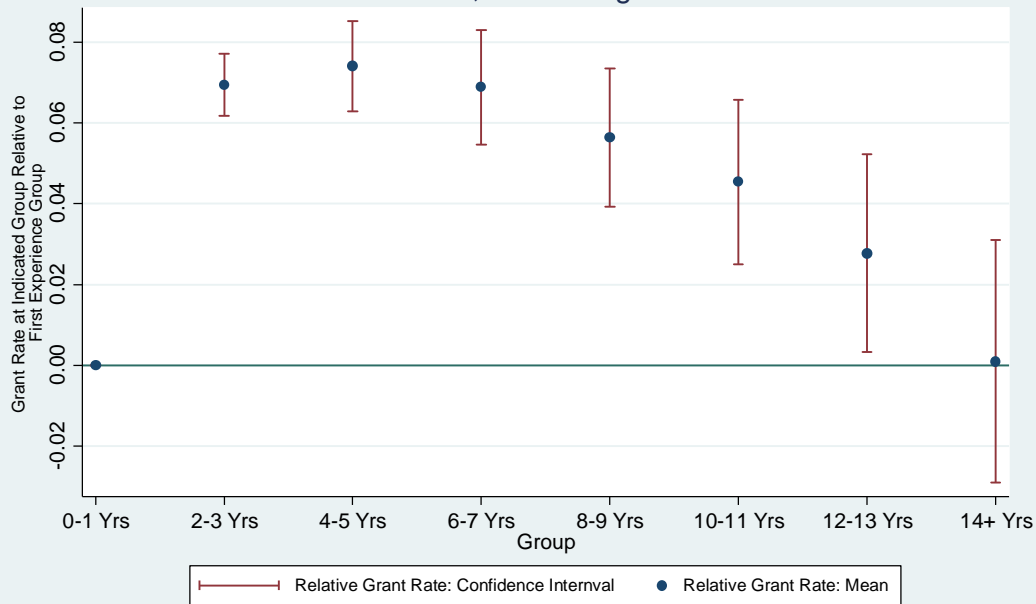


Figure 4: Relationship between Examiner GS Levels and Grant Rate Controlling for Experience Length Groups Including Both Types of GS-13 Examiners

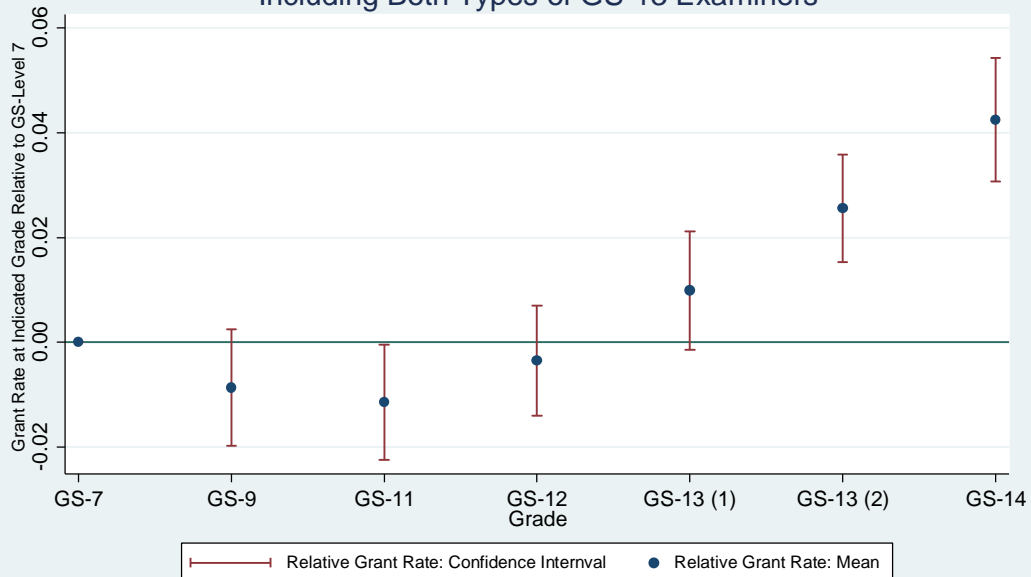

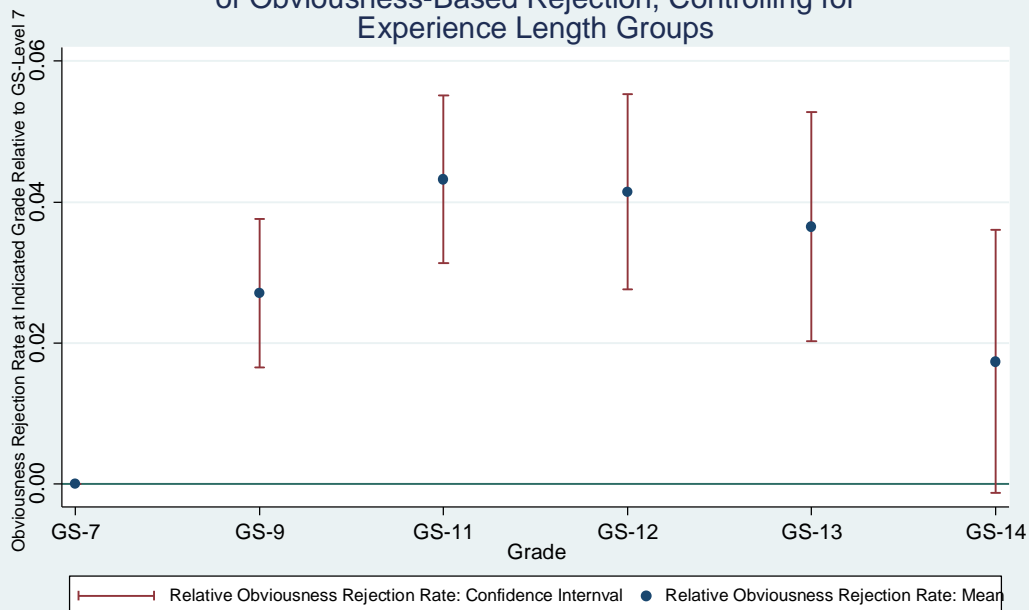


Figure 5: Relationship between Grant Rate and Increases in Experience Years within Distinct Grade Levels



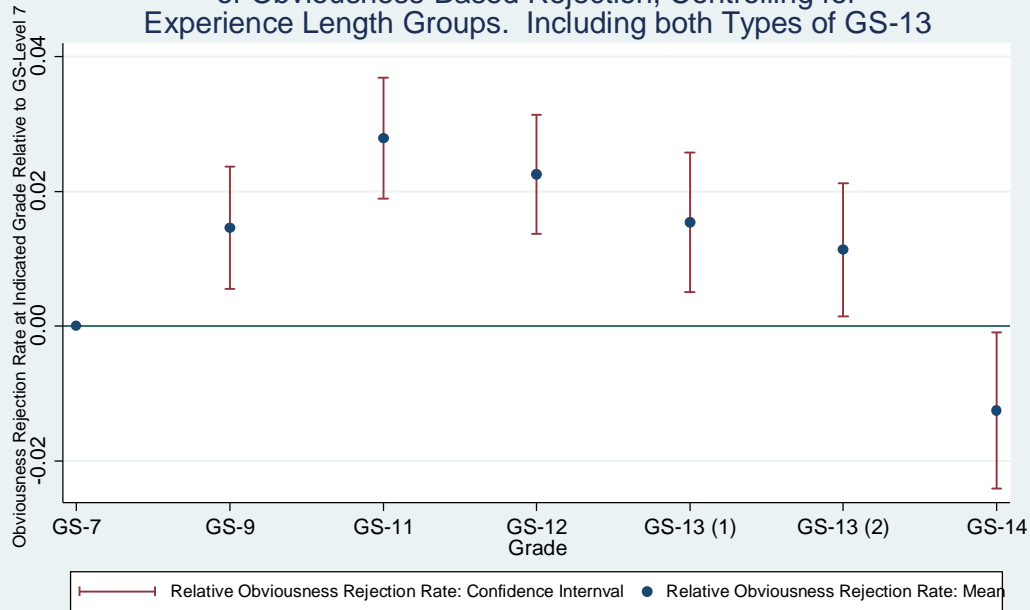
Notes: In the specification underlying this figure, we regress the incidence of the application being granted on a series of dummy variables capturing specific experience years within each grade level, beginning at GS-level 12. We track examiners for 1-2, 3-4, 5-6, 7-8 and 9+ years within GS level 12 and then the same within each of GS-level 13 without signatory authority, GS-level 13 with signatory authority and, finally, GS-level 14. Specifications include both examiner and year fixed effects.

Figure 6: Relationship between Examiner GS Levels and Incidence of Obviousness-Based Rejection, Controlling for Experience Length Groups




Notes: this figure replicates that of Figure 2, except that it replaces the dependent variable with the incidence of any obviousness-based rejection.

Figure 7: Relationship between Examiner GS Levels and Incidence of Obviousness-Based Rejection, Controlling for Experience Length Groups. Including both Types of GS-13



Notes: this figure replicates that of Figure 6, except that it includes separate groups for GS-13 examiners with and without partial signatory authority.

Figure 8: Relationship between Examiner GS Levels and Share of Rejections based on Obviousness, Controlling for Experience Length Groups



Notes: this figure replicates that of Figure 8, except that it replaces the dependent variable with the share of rejections for the given application that are based on obviousness.







Figure 11: Relationship between Examiner GS Levels and Share of Rejections (excluding Obviousness) based on Lack of Novelty



Notes: this figure replicates that of Figure 9, except that it specifies the dependent variable as the incidence of a novelty rejection divided by the sum of the incidences of the various non-prior-art based rejections, capturing the relative effort spent on novelty analyses.

Figure 12: Relationship between Examiner GS Levels and Share of Prior Art Citations Originating from Examiner



Notes: this figure replicates that of Figure 4, except that (1) it is based on the universe of issued patents between 2002 and 2012 and (2) it formulates the dependent variable as the share of the prior art citations associated with the issued patent that originate from the examiner as opposed to the applicant.