

Crafting Intellectual Property Rights: Implications for Patent Assertion Entities, Litigation, and Innovation*

Josh Feng, Harvard University
Xavier Jaravel, London School of Economics

June 10, 2018

Abstract

We show that examiner-driven variation in patent rights leads to quantitatively large impacts on several patent outcomes (including patent value, citations, and litigation). In particular, Patent Assertion Entities (PAEs) overwhelmingly purchase patents that were granted by “lenient” examiners; these examiners ask for fewer edits, and their patents are also more likely to be legally invalid and litigated by non-PAEs. This suggests PAEs are a symptom of the broader issue of poorly crafted weak patents leading to more litigation. Our results show that there is much at stake during patent examination, contradicting the influential “rational ignorance” view of the patent office.

JEL codes: O30, O31, O34, O38, K41

*A preliminary draft of this paper was previously circulated under the title “Who Feeds the Trolls? Patent Trolls and the Patent Examination Process.” We are particularly grateful to Michael Frakes and Melissa Wasserman, the RPX Corporation and Juristat for providing some of the data used in this paper. For thoughtful discussions and comments, we thank Philippe Aghion, Pierre Azoulay, Raj Chetty, Lauren Cohen, Stephen Haber, Nathan Hendren, Larry Katz, Bill Kerr, Jay Kesan, Scott Kominers, Mark Lemley, Josh Lerner, Ross Levine, Alan Marco, Arti Rai, Ben Roin, Scott Stern, John Van Reenen, Fabian Waldinger, Michael Webb, Heidi Williams, as well as seminar participants at Boston University, Duke, Harvard, Hoover, the NBER Productivity Seminar, the NBER Summer Institute, and the United States Patent Office Visiting Speaker Series. This research was funded by the Kauffman foundation.

I Introduction

A striking feature of patent systems around the world is the enormous variation in private returns, social returns and litigation risk across patents (e.g., [Pakes \(1986\)](#) and [Kogan et al. \(2017\)](#) on firms' returns, [Toivanen and Väänänen \(2012\)](#) and [Bell et al. \(2017\)](#) on inventors' returns, [Jaffe et al. \(1993\)](#) on patent citations as a proxy for social value, and [Lanjouw and Schankerman \(2001\)](#) on exposure to litigation). What are the sources of this heterogeneity in patent outcomes? Scientific or technical factors, such as the expertise of eminent scientists (e.g., [Azoulay et al. \(2010\)](#)) or a firm's learning capacity (e.g., [Cohen and Levinthal \(1989\)](#)), are likely to be important drivers of patent outcomes. Yet, the value of a patent may not be solely determined by the quality of the idea embedded in it: a patent is not a raw idea but a carefully-worded legal document, conferring to its holder the right to sue for infringement.

Using variation in the process of writing the patent description and claims at the United States Patent Office (USPTO), this paper shows that a significant amount of heterogeneity in patent outcomes is independent of technical determinants and results from the way patent rights are crafted. This finding is of particular relevance to understand a much-debated feature of the U.S. innovation system: the activities of Patent Assertion Entities (PAEs). PAEs, which acquire patents from third parties and generate revenue by asserting them against alleged infringers, have become controversial as they account for a large share of patent licensing and lawsuits.¹ We find that PAEs disproportionately purchase and assert patents from "lenient" patent examiners, who craft patents that are more likely to be litigated and to be legally invalid.² As discussed below, this pattern cannot be accounted for by mainstream theories of PAEs (which emphasize efficient intermediation or nuisance litigation); instead, it indicates that PAEs are primarily the symptom of a specific friction in the patent system, the way patent rights are crafted by lenient examiners (which affects litigation more broadly).

In the first part of the analysis, we show that the crafting of patent rights is an important driver of a wide range patent outcomes, in particular those related to litigation. To arrive at this result,

¹For instance, [RPX Corporation \(2015\)](#) reports that the share of PAEs in overall patent lawsuits went from 35% in 2010 to 70% in 2015, while [Federal Trade Commission \(2016\)](#) documents that the share of PAE in licensing revenue was 80% in the wireless chipset sector between 2009 and 2014. Some experts conjecture that recent patent reforms (inter-partes review) and Supreme Court rulings (*eBay Inc. v. MercExchange, L.L.C.*, *Alice Corp. v. CLS Bank International*) may curb PAE activity. One piece of evidence supporting this conjecture is that the number of defendants targeted by PAEs fell by 27% from 2016 to 2017. However, PAE-targeted firms still made up 56% of all defendants in IP cases, and the number of PAE lawsuits does not account for other forms of PAE activity.

²As we discuss in Section IV.D, surprisingly these examiners do not, on average, increase patents' private value, as measured by stock market responses.

we need variation in patent rights that is orthogonal to other determinants of patent outcomes, such as technical merit. Patent examiners may provide such variation as they only affect patent rights, not the underlying idea embedded in the patent. Examiners are heavily involved in the process of writing the patent description and claims through a back-and-forth process with the applicant between patent filing and patent grant (known as the “prosecution” process). By law, all examiners must ensure that the patents they grant have clear, well-defined claims with appropriate scope. In practice, we find significant variation in the way examiners craft patent rights (using prosecution data from [Frakes and Wasserman \(2017\)](#)); we can therefore use examiner assignment as a source of variation in patent rights, holding idea quality fixed.

A growing literature (e.g., [Sampat and Williams \(2015\)](#), [Gaulé \(2015\)](#) and [Farre-Mensa et al. \(2017\)](#)) suggests that patent applications can be treated as quasi-randomly allocated to examiners conditional on some covariates like application, year and technology class.³ Prior research has used examiner assignment to estimate the causal effects of obtaining a patent, as examiners differ in their grant rates. We build on this approach but differ in two ways. First, we develop new quasi-experimental approaches to address identification concerns about examiner specialization raised in more recent work ([Righi and Simcoe \(2017\)](#)); second, we exploit variation in examiner prosecution behavior *conditional* on granting the patent, rather than variation in the propensity of examiners to grant patents. We present evidence supporting the validity of our approach after reporting a set of baseline results.

Our baseline research design estimates examiner fixed effects on the set of granted patents conditional on technology by year fixed effects. Our estimator uses an Empirical Bayes shrinkage correction to prevent “overfitting” of the fixed effects, which would misattribute some of the variation from the noise to causal variation across examiners. We apply this methodology to a range of patent outcomes related to private returns (stock market response from [Kogan et al. \(2017\)](#) and payment of maintenance fees), patent citations (total citations, self citations and external citations), patent market dynamics (patent sales, in general and specifically to PAEs) and legal disputes (patent infringement lawsuits, in general and specifically from PAEs). The estimated examiner effects are large for many outcomes, in particular for those related to PAEs and litigation. For example, a one standard deviation change in examiner effects leads stock market capitalization to increase by

³Conceptually, patent outcomes may vary because of heterogeneity in idea quality, heterogeneity in the applicant’s input into patent drafting (typically via the applicant’s lawyers), and heterogeneity in the examiner’s input into patent drafting. We use variation in patent drafting from examiners, rather than from lawyers, because examiners are quasi-randomly assigned to patents while lawyer assignment may be correlated with idea quality across applicants.

3 million dollars, total citations by 24%, patent purchases by PAEs by 63%, litigation by 64%, and litigation specifically by PAEs by 46%. These estimates imply that policies affecting examiner behavior can have a substantial impact on the U.S. innovation system.⁴

We then validate the causal interpretation and the magnitudes of our baseline estimates in three ways. First, regarding identification, [Righi and Simcoe \(2017\)](#) report strong evidence that examiners working in the same technology-based group (called “art unit”)⁵ in fact specialize in specific sub-technologies, in ways that may be difficult to control for using observables. We develop two complementary quasi-experimental approaches to address this concern: (1) we show that there is a large subset of art units within which patent applications are assigned to examiners based on the last digit of the application’s serial number, implying that examiner assignment is orthogonal to potential confounds in these art units; and (2) we show that an examiner’s “busyness” can be used as an instrument for application assignment: examiners with recently disposed applications are much more likely to be assigned the next incoming application, which provides variation in assignment even in art units with significant specialization. These two alternative sources of variation yield estimates that are similar to our baseline results. Second, we show that our results are not confounded by selection effects stemming from the decision to grant a patent. Since examiners differ in their grant rates, it could be the case that patent outcomes vary across examiners because of underlying differences across examiners’ pools of granted patents, independently of the crafting of patent rights. For instance, examiners with a low grant rate might only grant patents of high technical merit. To establish that the bias is small empirically, we introduce flexible controls for examiners’ grant rates in our baseline specification and show that there is equally large causal variation in patent outcomes across examiners with the same grant rate. Third, we validate our baseline estimates in out-of-sample tests. We find that the Empirical Bayes shrinkage correction is important to suitably account for excess variance from noise and obtain unbiased estimates of examiner effects, in particular for rare outcomes such as PAE purchase and litigation.

In the second part of the analysis, we investigate why examiner effects are an important driver of the wave of patent purchases and lawsuits by PAEs, a major and controversial feature of the U.S.

⁴As a point of comparison, the teacher value-added literature has documented sizable but much smaller effects of teachers on students’ outcomes. [Chetty et al. \(2014a\)](#) and [Chetty et al. \(2014b\)](#) estimate that a one standard deviation improvement in teacher effects in one grade raises students’ earnings by about 1% at age 28.

⁵Examiners at the USPTO are divided into more than 600 working groups called “art units”, each composed of about twenty examiners who handle patent applications on relatively homogeneous technologies. Following qualitative evidence on assignment of applications to examiners reported in [Cockburn et al. \(2003\)](#), [Lemley and Sampat \(2010\)](#) and [Lemley and Sampat \(2012\)](#), the recent literature treats assignment of patents applications to examiners within the same art unit as “as good as random” (e.g., [Sampat and Williams \(2015\)](#), [Gaulé \(2015\)](#) and [Farre-Mensa et al. \(2017\)](#)).

innovation system. We focus on outcomes related to PAEs because they rank among the outcomes that are most sensitive to examiner effects, and because PAEs have generated substantial academic and policy debate.⁶ There are two main hypotheses about PAEs’ behaviors: (1) PAEs may be useful intermediaries in the patent market, fostering greater incentives to innovate by lowering the cost of matching patent holders to patent buyers (e.g., [Hagi and Yoffie \(2013\)](#) and [Abrams et al. \(2016\)](#)) and by helping enforce the patents of small inventors who lack the financial resources or legal expertise to defend themselves against large infringing companies (e.g., [Lu \(2012\)](#) and [Galetovic et al. \(2015\)](#)); or (2) PAEs may exploit imperfections in the legal system by acquiring patents with unclear claim boundaries and by asking innovative firms for licensing fees, whether or not the asserted patent is valid or infringed, in the hope that targeted firms will settle instead of risking a costly and uncertain trial (e.g., [Miller \(2013\)](#), [Council of Economic Advisers \(2013\)](#), [Cohen et al. \(2016\)](#) and [Federal Trade Commission \(2016\)](#)). Any plausible theory of PAEs should account for the new fact, documented in the first part of this paper, regarding the large sensitivity of PAEs to the way examiners craft patent rights. By analyzing which examiners drive patent acquisition and litigation by PAEs, we can assess which PAE theories are plausible.

We start by studying the characteristics of examiners who issue patents that are purchased and asserted by PAEs or by practicing firms. We correlate the causal examiner effects from the first part of the paper with measures of examiners’ prosecution behaviors based on the correspondence between examiners and applicants (from [Frakes and Wasserman \(2017\)](#)). We find that, within the same technology category, PAEs and practicing firms target patents issued by examiners with different characteristics. PAEs disproportionately purchase and assert patents that were granted by “lenient” examiners, who require applicants to make fewer changes to the text of the patent, such as clarifying a claim or withdrawing a claim deemed to be obvious or to bear on a non-patentable subject matter. Examiner leniency has a negligible impact on purchases by practicing firms, a sizable effect on litigation by practicing firms and on purchases by PAEs, and a much larger effect on litigation by PAEs. These patterns are first-order: for instance, a one standard deviation increase in a simple proxy for examiner leniency, the change in the number of words per claim between patent filing and grant, leads to an increase in litigation of 40.5% for PAEs and of 13.9% for practicing

⁶PAEs, also known as “non-practicing entities”, “patent monetization entities” or “patent trolls”, are defined as entities that generate revenue exclusively from patent licensing and litigation, without producing or selling products ([Federal Trade Commission \(2016\)](#)). Since there is no official list of PAEs, we follow the literature (e.g., [Bessen and Meurer \(2014\)](#)) and rely on a list provided by the RPX Corporation, a firm that helps companies manage risks from exposure to patent litigation. Universities, individual inventors and failed companies are excluded from the set of PAEs we consider and we show that the results are similar with alternative PAE lists from [Cotropia et al. \(2014\)](#).

firms. These results cannot be accounted for by theories of PAEs based on a generic friction in the patent market, such as matching costs or the lack of financial resources for some inventors. They are consistent with the view that PAEs have a comparative advantage in patent litigation and therefore handle patents that are subject to a higher litigation risk, induced by the way patent rights were crafted during the patent prosecution. The fact that examiner leniency is an important driver of litigation for *both* PAEs and practicing firms, although the effect is not as large for the latter, is in line with a nuanced view of PAEs (e.g., [Lemley and Melamed \(2013\)](#) and [Schwartz and Kesan \(2013\)](#)). According to this view, PAEs do not exploit imperfections of the legal system in an idiosyncratic way, but behave as litigation experts. In sum, our results show that PAEs' activities are the symptom of the way patent rights are crafted by lenient examiners, who affect litigation more broadly.

Given the evidence that patent litigation by PAEs is strongly correlated with examiner leniency, we study whether lenient examiners tend to issue patents that are more likely to be invalid according to the standards set by current patent law. Several observers have hypothesized that PAEs assert invalid patents (e.g., [Federal Trade Commission \(2016\)](#)); approaching this question in terms of examiner effects has the potential to be informative about PAEs but also about patent litigation by practicing firms, who also selectively assert patents that were crafted by lenient examiners. Patent invalidity is notoriously difficult to measure because of selection effects. For instance, court rulings on patent validity are observed only for a strongly selected set of patents, as there were only a few hundred rulings over the past decade. To address this issue, we introduce a proxy for patent invalidity available in the full sample of granted patents: patent re-issuance requests, which can be filed by the applicant when a patent is deemed wholly or partly “inoperative or invalid” through an error in the document. Using this proxy as well as two common proxies for invalidity (decisions from court rulings and trials at the patent office), we document robust and quantitatively important evidence that lenient examiners issue patents that are more likely to be invalid. The evidence is therefore consistent with the view that PAEs are willing to purchase and assert patents whose validity is questionable, but PAEs are not the only entities to assert such patents: practicing firms do so as well.⁷

These results are surprising for several reasons. First, the very large impact of patent exam-

⁷This finding does not speak conclusively to the welfare effects of PAEs, because litigation of patents issued by lenient examiners could conceivably be socially valuable, even when these patents are deemed invalid by the courts, the USPTO, or applicants themselves. For instance, [Galetovic et al. \(2015\)](#) suggest that the process of litigation might be the socially-efficient dynamic process through which the patent system defines the contours of what should be patentable in highly-innovative, rapidly changing industries.

iners on patent outcomes (conditional on grant) is unexpected. While previous work documented variation in examiners’ propensities to grant patents (Sampat and Williams (2015), Gaulé (2015) and Farre-Mensa et al. (2017)), we uncover the importance of the “intensive margin” of examiner effects (the crafting of patent rights, conditional on patent grant). This margin was not previously thought of as being of paramount importance for patent outcomes, as evidenced by (1) the fact that patent examiners are not very well paid;⁸ (2) the influential “rational ignorance” view of the patent office (Lemley (2000)), which states that there is little at stake in the patent examination process and that low-quality patent can be rationally ignored without significant consequences; (3) the focus of innovation research on the “macro-determinants” of the patent system, such as laws that establish a patent system or change the set of patentable subject matters;⁹ in contrast, we show the importance of the “micro-determinants” of patents by establishing that the specific way in which patent rights are crafted (by examiners who are all subject to the same patent law) has a substantial impact on a range of patent outcomes and is of first-order importance to understand certain features of the U.S. innovation system such as litigation (by PAEs in particular, but also by practicing firms). Additionally, our results on PAEs and litigation are also unexpected: (1) our finding that PAEs overwhelmingly purchase and assert patents from lenient examiners is not in line with mainstream theories of PAEs;¹⁰ (2) our findings imply that policies affecting examiner behavior could have a large impact on PAEs’ activities and litigation (in contrast with prior work that does not use quasi-experimental variation, e.g. Marco and Miller (2017));¹¹ (3) the examiners who drive PAEs’ activities do not tend to create greater private value in general, as discussed in Section IV.¹²

The remainder of the paper is organized as follows. Section II presents the data and descriptive statistics. Section III estimates examiner effects on a range of patent outcomes. Section IV studies the implications for PAEs’ activities. Section V concludes.

⁸In 2017, most patent examiners started with annual salaries between \$54,099 and \$82,094; salary can reach around \$130,000 for senior patent examiners, which is much lower than the typical salary of patent lawyers.

⁹For theoretical contributions, see e.g., Nordhaus (1969), Klemperer (1990) and Gilbert and Shapiro (1990); for empirical studies, see e.g., Sakakibara and Branstetter (2001), Moser (2005), Lerner (2009) and Williams (2013).

¹⁰For recent work on PAEs, see e.g. Golden (2006), McDonough III (2006), Chien (2013), Tucker (2014), Allison et al. (2016) and Haber and Werfel (2016).

¹¹Our findings are also related to the pioneering study of Cockburn et al. (2003), who document relationships between some examiner characteristics and patent invalidity rulings. We show how to recover the full magnitude of examiner effects on patent outcomes using a fixed effects estimator with a Bayesian shrinkage correction.

¹²Therefore it is not the case that examiners with high PAE effects create more valuable patents for all agents in the patent system, as would be the case if they were just allowing greater scope.

II Data

In this section, we describe the data sources, define the samples and key variables we used in the analysis, and present summary statistics.

II.A Data Sources, Samples and Variable Definitions

Patent Records. We use two types of patent data to achieve two purposes. First, we rely on data on granted patents to measure a series of post-grant patent outcomes. Specifically, we build proxies for the private returns to patents, identify high-impact patents using citations, and document transactions in the patent market. Second, we use data on both granted and ungranted patent applications to identify examiners and measure their behavior during patent prosecution.

The granted patent dataset is obtained from USPTO and extends from 1975 to 2016.¹³ We rely on several proxies for the private returns to patents. Following the literature (e.g., [Pakes \(1986\)](#)), we use the payment of patent maintenance fees as a lower bound on the private valuation of the patent by the assignee. These fees are due 4 years, 8 years and 12 years after patent grant and are increasing over time.¹⁴ We also use the estimates of firm-level returns to patents from [Kogan et al. \(2017\)](#), who run event studies to estimate the excess stock market return realized on the grant date of patents assigned to publicly-traded firms; these estimates are available for patents granted before 2010. Moreover, we use data on patent citations to identify high-impact patents. We consider alternatively total citations, self citations (i.e. the assignee of the focal patent cites it in future patents) and citations by assignees that were not listed on the focal patent. We build these measures using the disambiguated assignee names from [Balsmeier et al. \(2015\)](#). To address censoring, we focus on citations that occurred in the three years following patent grant and we document in robustness checks that the results are similar when considering all citations. Finally, we measure changes in ownership of patents by merging in data on patent re-assignments from [Graham et al. \(2015b\)](#).¹⁵

The data covering both granted and ungranted patent applications ranges from 2001 to 2015

¹³This data is obtained through the Reed Tech USPTO page: <http://patents.reedtech.com/patent-products.php>.

¹⁴For entities that do not benefit from reduced rates, the fees are \$1,600 after 4 years, \$3,600 after years and \$7,400 after 12 years. The complete fee schedule is available from the USPTO at <https://www.uspto.gov/learning-and-resources/fees-and-payment/uspto-fee-schedule#Patent Maintenance Fee>.

¹⁵Records of the assignments (transactions) affecting US patents are maintained by the US Patent & Trademark Office and available up from 1970 to 2014. There is no express legal requirement for parties to disclose assignments to the USPO, but patent laws provide incentives for recording. For instance, failure to record an assignment renders it void against any subsequent purchaser or mortgagee (35 USC 261). See [Graham et al. \(2015b\)](#) for more details.

and is obtained from the USPTO’s Patent Examination Dataset ([Graham et al. \(2015a\)](#)). We use this dataset to obtain unique numeric identifiers for each examiner during their tenure at the patent office, which are the critical inputs needed to estimate examiner effects. We then merge in data from [Frakes and Wasserman \(2017\)](#) on the correspondence between the examiner and the applicant. When asking applicants to amend patent documents, examiners need to ground their demands in specific sections of patent law, which we describe in Section II.B.¹⁶ To characterize an examiner’s behavior during prosecution, we count the number of references made to the various sections of patent law. We also measure the examiner’s grant rate and, for granted patents, we directly measure the extent to which the text of the patent changes between application and grant by computing changes in the number of words per claim and in the number of claims.¹⁷

Our main analysis sample is the Patent Examination Dataset merged to the patent outcomes of the granted patent dataset. We implement one important sample restriction: we exclude the so-called “continuation applications”, applications that follow an earlier-filed patent application. Those applications are assigned to the same examiner as the patent they follow and, therefore, quasi-random assignment of examiner does not hold. Our main analysis sample covers each non-continuation granted patents between 2001 and 2015, for which we observe the patent outcomes of interest as well as examiners’ identity and prosecution behaviors. For robustness, we estimate examiner effects on the full sample of (non-continuation) granted patents going back to 1975 by disambiguating examiner names (given the lack of numeric identifiers in this sample), but we lose information on examiners’ prosecution behaviors.

Patent Litigation. We combine three data sources to obtain a comprehensive picture of patent litigation. Specifically, we combine data from LexMachina, Darts IP and RPX, which have been tracking intellectual property lawsuits since 2000 and thus offer full coverage for our main analysis sample. Although the datasets have significant overlap, it is sometimes challenging to identify all the patents involved in a given lawsuit, which creates differences in the lists.¹⁸

Patent Assertion Entities. Following standard practice (e.g., [Bessen and Meurer \(2014\)](#)), we rely on a list provided by the RPX Corporation, a firm that helps companies manage litigation

¹⁶When a patent is assigned to two examiners, a “primary” examiner with signatory authority and a “secondary” examiner who carries out most of the work, we treat the data as if the patent had been assigned to the secondary examiner only, following the example of [Lemley and Sampat \(2012\)](#).

¹⁷The USPTO’s Patent Examination Dataset only covers published patent applications. For ungranted patents, applicants are free to opt out of publications, which occurs in about 5% of cases during the period we consider ([Graham et al. \(2015a\)](#)). The potential selectivity issues that could arise from the omission of “nonpublic” applications are largely orthogonal to our analysis, as we only rely on ungranted applications to measure an examiner’s allowance rate.

¹⁸We manually checked a few of the differences and verified that the patents were actually involved in litigation.

risk, and exclude from the list any individual inventor, university or failed company.¹⁹ We then build the patent portfolio of PAEs by merging the PAE list to the patent re-assignment dataset of [Graham et al. \(2015b\)](#) by assignee name. We only consider patents that were purchased by PAEs (a few large PAEs, such as Intellectual Ventures, also invent their own patents). To establish that our results are robust to the choice of PAE list, we repeat the analysis using alternative PAE lists from [Cotropia et al. \(2014\)](#) and considering only the patent portfolio of Intellectual Ventures.²⁰

II.B Summary Statistics

Table 1 presents the summary statistics for the variables of interest, documenting heterogeneity in patent outcomes (Panel A), the extent to which patent documents change between application and grant (Panel B) and heterogeneity in examiner behavior (Panel C).

Statistics on private returns, citations, patent sales and patent litigation are shown in Panel A of Table 1. Private returns feature high variance: the standard deviation of the firm-level patent value estimates from [Kogan et al. \(2017\)](#) is equal to almost three times the mean. The rates of maintenance fee payments are very high in early years but are substantially lower for the more expensive 12th-year maintenance fee payment, which also indicates heterogeneity in private valuations. Citations also feature high variance, indicating that patents greatly vary in their level of impact, regardless of whether we consider total citations, self-citations or citations by other assignees. The panel also shows that about 20% of all granted patents are sold to practicing (i.e. non-PAEs) firms and 1.01% to PAEs. Only 0.65% of all granted patents are litigated. Patent litigation by PAEs involves 0.04% of patents: this fraction is very small but it indicates that PAEs’ litigation rate is over six times higher than average, given that they own only about 1% of the patent stock.²¹ The purpose of Section III is to estimate the extent to which this heterogeneity in patent outcomes results from the way patent rights are crafted by examiners.

Panel B of Table 1 shows how the patent document changes between application and grant. In most cases, the examiner issues a so-called “rejection” as her first decision on the application ([Williams \(2017\)](#)), which is effectively an invitation for the patent applicant (or their representative, typically a patent attorney) to revise the text of the patent. Panel B shows that these changes are substantial. Through the back-and-forth with the examiner, the number of words in each claim

¹⁹Excluded entities are based on classifications from RPX and [Cotropia et al. \(2014\)](#).

²⁰Intellectual Ventures holds an estimated 25-30k US patents and released a list of around 20,000 patents on their website in November of 2013, which is available at <http://patents.intven.com/data/ivpatents.csv>.

²¹In addition, PAE patents are involved in about 7 cases per litigated patent versus about 2 cases for non-PAE litigated patents, based on a simple count of District Court cases per patent in the LexMachina data.

increases by 57% on average.²² The lengthening of the claims can be interpreted as limiting the scope and clarifying the claims by making them more precise (Marco et al. (2016)). In addition, examiners tend to ask applicants to reduce the number of claims to limit the scope of the patent: while the average change is limited (-3.64%), the standard deviation across patents is high (46.14%).²³ We also observe that the examiner asks the applicant to add citations to prior patents. The changes to the patent document during the back-and-forth between the applicant and the examiner show that the examiner is engaged in an iterative process and does not simply make a one-time accept-or-reject decision. During this process, the examiner must substantiate her demands by referring to specific sections of patent law corresponding to various standards of patentability, namely that the invention is useful and its subject matter is eligible for a patent (35 U.S.C. §101), it is novel relative to the prior art (35 U.S.C. §102(a)), it is non-obvious (35 U.S.C. §103(a)), and the claims are sufficiently clear to satisfy the disclosure requirement (35 U.S.C. §112(b)). Panel B of Table 1 shows that on average non-obviousness is used significantly more frequently than other sections.

Panel C of Table 1 presents statistics at the level of examiners. We observe 10,018 examiners in our main analysis sample, who work at the USPTO for 6.35 years on average. The median number of technology areas in which an examiner works (called “art units”) is two. The average examiner processes close to 200 patents over the course of our sample. The panel shows that some examiners have a much higher grant rate than others, or have a stronger tendency to invoke specific sections of patent law during the back-and-forth with the applicant. We also observe large variation across examiners in the shares of their granted patents that is purchased by a PAE: the standard deviation across examiners is twice the average PAE purchase rate. This observed heterogeneity across examiners could merely reflect noise or the fact that different examiners are working on different technologies, or it could be driven by systematic (causal) differences in examiner behavior, which we investigate in the remainder of the paper.

II.C Illustration of Main Findings

Some of our main results in Sections III and IV can be previewed in a simple, graphical way. The various panels of Figure 1 document the relationship between patent acquisition or litigation and a simple measure of examiners’ prosecution behavior.

²²Given that the effective IP protection provided by a patent depends entirely on the content of the claims, and given that examiners affect to a great extent the words in the claims during prosecution, it is plausible that examiners may have a large impact on the legal force of the patent.

²³Following the literature, we report statistics for independent claims, leaving dependent claims aside as in Marco et al. (2016).

For each patent, we compute the average change in the number of words per claim between application and grant for all other granted patents processed by the same examiner, leaving out the focal patent. This leave-one-out examiner measure is exogenous to the focal patent. To ensure that we compare similar examiners, we include art unit by patent filing year fixed effects in all specifications. To ensure that potential extensive-margin selection effects are not confounding the results, we control for the (leave-one-out) grant rate of the examiner. Conceptually, these specifications compare patent outcomes for examiners who have the same grant rate, work in the same art unit in the same year, but differ in the way they craft property rights, as measured by the change in the number of words per claim between application and grant.

Panel (a) of Figure 1 shows that the probability that a patent is purchased by a PAE is a strongly negative function of the examiner’s propensity to ask applicants to add words to the patent claims (for instance to clarify them). Each dot in the binned scatter plot represents 5% of the data. The PAE purchase rate falls by about 25% of the baseline rate as we move from the left to the right along the x-axis, which shows very directly that the way examiners craft property rights is first-order for certain patent outcomes. Similarly large effects are found for litigation by PAEs and litigation by practicing firms, but not for purchases by practicing firms. The comparison of the various panels shows that PAEs and practicing firms respond in a similar way to examiners for the purpose of patent litigation (Panels (c) and (d)) but not for patent acquisition (Panels (a) and (b)).

This simple regression approach has the benefit that its robustness can immediately be assessed graphically. But the choice of the variable on the x-axis is arbitrary: this variable may capture only a small fraction of the relevant examiner behaviors and it may be correlated with examiner traits that would suggest different interpretations. To address this limitation, we turn to a research design that can recover the full impact of examiners on patent outcomes (Section III), and we then correlate the examiner-level causal estimates with a range of examiner characteristics (Section IV).

III Estimating Examiner Effects on Patent Outcomes

In this section, we estimate the impact of examiners on a range of patent outcomes. We assess the validity of the identifying assumptions in our baseline design using additional sources of variations and alternative specifications.

III.A Research Design

To estimate the extent to which the heterogeneity in patent outcomes results from the way patent rights are crafted, we need variation in patent rights that is orthogonal to other determinants of patent outcomes, such as technical merit. Through their back-and-forth with the applicant between initial filing and grant, examiners may provide such variation. By definition, examiners only affect patent rights, not the underlying idea embedded in the patent. Moreover, a growing literature suggests that patent applications can be treated as quasi-randomly allocated to examiners working in the same art unit in the same year (Sampat and Williams (2015), Gaulé (2015) and Farre-Mensa et al. (2017)).

Using quasi-random allocation of patent applications to examiners raises three empirical concerns, which were previewed in the introduction. First, since we are interested in recovering the full magnitude of examiner effects, conceptually we need to estimate fixed effects for all examiners, instead of projecting the data onto a specific examiner trait as in Figure 1.²⁴ Given that we have a large number of examiners and work with rare outcomes such as litigation, it is likely that we may be “overfitting” the fixed effects: we may misattribute some of the variation from the noise to causal variation across examiners. This “excess variance” problem is well-known and we address it using a standard Bayesian shrinkage methodology (e.g., Kane and Staiger (2008), Chetty et al. (2014a) and Chetty and Hendren (2016)). Our baseline research design focuses on addressing this issue. Second, recent evidence from Righi and Simcoe (2017) challenges the notion that the allocation of patent applications to examiners can be treated as “as good as random”. Third, our examiner effects could in principle be confounded by selection effects related to grant decisions. Using alternative sources of variation and specifications, we find that the last two potential threats turn out to leave our baseline estimates unaffected. We therefore proceed by presenting our baseline design and its results, before turning to validation tests addressing the other potential threats.

Our baseline research design estimates examiner fixed effects on the set of granted patents with an Empirical Bayes shrinkage correction, conditional on art unit by year fixed effects. The

²⁴Running a specification using examiner characteristics as regressors can only recover a lower bound for the overall effect of examiners, because the observed characteristics only capture a fraction of examiner behavior. A fixed effects estimator can recover the full effect, but it must be adequately adjusted to avoid excess variance due to overfitting of the fixed effects. In addition, the regression coefficients for the various examiner characteristics included in the specification should not be interpreted as causal, because random assignment occurs at the level of examiners and the observed examiner characteristics are likely to be correlated with other, unobserved examiner characteristics. For instance, in contemporaneous work, Kuhn (2016) and Kuhn and Thompson (2017) create an instrumental variable for patent scope based on an examiner characteristic they label “scope toughness”, but this characteristic could be correlated with other examiner traits that may affect the patent through channels other than scope.

identification assumption is that the allocation of (non-continuation) patents to examiner working in the same art unit in the same year is as good as random, i.e. it is not correlated with other determinants of patent outcomes. Given this assumption, we estimate examiner effects using the following statistical model:

$$\begin{aligned} Y_i &= a_{ut(i)} + v_{ij}, \\ v_{ij} &= \mu_j + \epsilon_i, \end{aligned} \tag{1}$$

where i indexes the patent, j the examiner, u the art unit and t the year. Y_i is the patent outcome of interest, $a_{ut(i)}$ denotes art unit by year fixed effects, μ_j is the causal examiner effect of interest and ϵ_i is an idiosyncratic patent-level shock. Our goal is to recover $\sigma_\mu \equiv \sqrt{Var(\mu_j)}$.

We estimate the standard deviation of the underlying distribution of examiner effects in three simple steps. We first obtain estimates of residuals $\{\widehat{v}_{ij}\}$ for each patent by estimating art unit by year fixed effects in (2) by OLS. We then compute the average estimated residual per examiner in each year:

$$\bar{v}_{jt} \equiv \frac{1}{n_{jt}} \sum_{i=1}^{n_{jt}} \widehat{v}_{ij} = \mu_j + \frac{1}{n_{jt}} \sum_{i=1}^{n_{jt}} \epsilon_i, \tag{2}$$

where n_{jt} is the number of patents processed by examiner j in year t .

Finally, we compute the covariance between an examiner’s average residuals across consecutive years:

$$\widehat{\sigma}_\mu = \sqrt{Cov(\bar{v}_{jt}, \bar{v}_{j(t+1)})}, \tag{3}$$

which yields a consistent and unbiased estimate of σ_μ , as can be seen immediately from the second equality in (2). Excess variance in the average residual is handled by isolating the “systematic” component of the variation in average residuals that persists over time. If the examiner causal effects $\{\mu_j\}$ are close to zero, we may still observe variation in the average residuals $\{\bar{v}_{jt}\}$ across examiners in any given year because of idiosyncratic shocks, but there will be no covariance between examiners’ average residuals across years because the idiosyncratic shocks are uncorrelated. We call σ_μ the “signal” standard deviation of examiner effects to contrast it with the “raw” standard deviation of residuals, which is contaminated by noise. The covariance calculation in (3) uses the counts of patents granted by each examiner $\{n_{jt}\}$ as weights to increase precision.

The signal standard deviation is our primary focus because it is informative about the overall variation from examiners, but we also compute individual estimates of causal effects for each examiner. We compute an average of the residuals \widehat{v}_{ij} over all years for each examiner, which we

denote \bar{v}_j .²⁵ We then construct the empirical Bayes posterior estimate of each examiner effect by multiplying \bar{v}_j by a shrinkage factor:

$$\widehat{\mu}_j = \frac{\widehat{\sigma}_\mu^2}{\text{Var}(\bar{v}_j)} \cdot \bar{v}_j. \quad (4)$$

The shrinkage factor is the ratio of signal variance to total variance.²⁶ We validate this research design by documenting in Section III.C that this approach yields unbiased estimates of examiner effects in out-of-sample tests, while ignoring excess variance delivers misleading results.

III.B Baseline Estimates of Examiner Effects

Table 3 reports the estimates of examiner causal effects for a range of patent outcomes. We find substantial examiner effects for private value and for outcomes related to patent litigation.

Private value is strongly affected by examiner effects. The first row of Table 2 shows that the signal standard deviation of examiner effects corresponds to a 3.32 million dollar change in patent value, using the estimates from Kogan et al. (2017). In percentage terms, one signal standard deviation in examiner effects explains 40.8% of the average patent value for publicly-traded firms. The process of creation of patent rights therefore has a first-order impact on a patent’s private value to its assigned firm. We confirm this result in rows two to four of the table by considering other proxies. The rates of payment of patent maintenance fees at the various horizons are all responsive to examiner effects. Consistent with the notion that fee payments can only give a lower bound on private valuations, especially in earlier years when the fees are smaller, the examiner effects are smaller than with the Kogan et al. (2017) estimates; the signal standard deviations are under 10% of the average payment rate.

Citations also respond to examiner effects. Considering in turn the signal standard deviations for total patent citations, self citations and citations by other assignees, we consistently find significant effects. The impact is strongest for self-citations, with a signal standard deviation of 46.06%, while the signal standard deviation for citations by other assignees is only 24.47%. This finding points to

²⁵To increase precision, \bar{v}_j is computed using weights that make \bar{v}_j a minimum variance unbiased estimate of μ_j for each examiner. This step requires estimating the variances of other shocks in the statistical model. Specifically, we allow for an examiner-by-year shock θ_{jt} and compute $\widehat{\sigma}_\epsilon^2 = \text{Var}(v_{ij} - \bar{v}_{jt})$ and $\widehat{\sigma}_\theta^2 = \text{Var}(v_{ij}) - \widehat{\sigma}_\mu^2 - \widehat{\sigma}_\epsilon^2$. We obtain $\bar{v}_j = \sum_t w_{jt} \bar{v}_{jt}$, with $w_{jt} = \frac{h_{jt}}{\sum h_{jt}}$ and $h_{jt} = \frac{1}{\widehat{\sigma}_\theta^2 + \frac{\widehat{\sigma}_\epsilon^2}{n_{jt}}}$. See Online Appendix A for a complete discussion.

²⁶Online Appendix A discusses the computation of $\text{Var}(\bar{v}_j)$. Because of the precision weights in \bar{v}_j , the shrinkage factor is lower for examiners for which more patents are observed. The estimated examiner effects $\{\widehat{\mu}_j\}$ have an empirical Bayes interpretation as the Bayesian posterior estimates of the examiner effects, starting from a normal prior distribution centered around zero with signal variance σ_μ . There is also a frequentist interpretation: the shrinkage factor is the OLS coefficient in a hypothetical regression of the true (unobserved) μ_j on the (observed) \bar{v}_j .

the role of cumulative innovation by the assignee.²⁷

We find particularly strong examiner effects for litigation and PAEs’ activities. The signal standard deviation of examiner effects accounts for over 60% of the baseline rate of patent purchases by PAEs. In contrast, the impact of examiners on the probability that a patent is sold to a practicing firm is much smaller: the signal standard deviation is 14.6% of the baseline rate. The impact of examiners on the probability that a patent is litigated is very large: the signal standard deviation is about 65% of the baseline rate. Considering the raw standard deviation of examiner effects would be very misleading: for rare outcomes like patent litigation or PAE purchase, the raw standard deviation is implausibly high, over four times larger than the signal standard deviation (Online Appendix Table A1).

We use a bootstrapping procedure for inference. We re-draw samples from the application-level dataset with replacement and repeat the estimation of the signal standard deviations.²⁸ The 95% confidence intervals are reported in Column (1) of Table 2 . The signal standard deviations are all precisely estimated, except for one extremely rare outcome, patent litigation by PAEs.

The standard deviation of shrunk examiner effects obtained from equation (4) are also substantial. Column (3) of Table 2 reports these results. For instance, the standard deviation of shrunk examiners effects accounts for 29.48% of the average patent value from [Kogan et al. \(2017\)](#), 31.11% of the baseline rate of PAE patent purchases, and 27.43% of the average rate of patent litigation.

The large signal standard deviations indicate that examiners have a first-order impact on patent outcomes. Consequently, policies affecting examiners have the potential to greatly affect the U.S. innovation system, for instance regarding litigation rates or the activities of PAEs. The large standard deviations of shrunk examiner effects indicate that, based on historical data, one can identify examiners who have a particularly large or low impact on specific outcomes.²⁹ Our analysis so far is silent on the characteristics of these examiners, which we turn to in Section IV. Before doing so, we establish the validity of our identification assumptions with a series of tests and robustness checks.

²⁷Although this finding may also reflect strategic self citations, the literature on strategic self citations has emphasized the importance of strategic continuation filings, while we focus on non-continuation patents.

²⁸We also bootstrapped by re-sampling within examiner or within examiner by filing year and obtained similar results (not reported).

²⁹We found that examiner effects do not tend to “average out” across outcomes; for instance, there is a large share of examiners who produce patents with systematically lower value, fewer citations and higher probabilities of litigation or of PAE purchase (not reported).

III.C Validation of Baseline Design: Addressing Non-Random Assignment and Selection

In this subsection, we use alternative research designs and specifications to investigate potential limitations of the baseline research design.

Alternative source of variation #1: allocation of applications to examiners using the last digit of the application’s serial number. A potential concern with our baseline research design is that there is specialization even across examiners working in the same art unit at the same time (Righi and Simcoe (2017)). If specialization patterns are correlated with other factors that affect patent outcomes, then the examiner effects document in Table 2 may reflect omitted variable bias.

To address this potential concern, we identify art units where application assignment to examiners is determined by the last digit of the serial number of the patent application. The last digit of an application’s serial number, ranging from 0 to 9, is determined by the order of submission of applications and is therefore orthogonal to potential confounding variables such as technical factors.³⁰ Anecdotal evidence suggests that some art units assign applications to examiners based on the last digit of the serial number (Lemley and Sampat (2012)). To determine which art units do so at different points in time, we compute an index of “concentration” of last digits across examiners working in the same art unit in the same year. If some examiners systematically get specific last digits, we will find a high degree of concentration. We use the concentration index initially developed by Mori et al. (2005) to study industry agglomeration, which was recently applied by Righi and Simcoe (2017) to the context of patents to study examiner specialization.³¹ Applied to our purposes, the test delivers a Chi-square statistic asking whether applications’ last digits are less dispersed across examiners than one would expect if last digits were not used for application assignment.³² We carry out the test in each year and in each art unit.

Figure 2 presents the results. Panel A shows the distribution of the p-values of the Chi-square tests across art units. There is a large number of art units with a p-value below 1%, indicating that these art units use application last digit to assign patents. The test only rejects the null that last

³⁰When a patent application is filed, the Office of Patent Application Processing assigns it a serial number. The first part of the serial number indicates the technology category while the last digits reflect the order of arrival of applications.

³¹Righi and Simcoe (2017) use this test to document specialization of examiners in the same art unit and year, specifically testing for failure of random assignment with respect to technological features of the patent. We use the same test, but for the opposite purpose: we use the test to identify art units that allocate applications based on their last digits, which implies quasi-random allocation with respect to technological features of the patent.

³²Formally, we are testing the null that applications assignment is independent of their last digit; this test can be viewed as a multivariate generalization of a t-statistic comparing observed frequencies to the distribution under random assignment. For details, see Online Appendix A as well as Mori et al. (2005) and Righi and Simcoe (2017).

digits are *not* used and it can of course not guarantee that in art units with a p-value below 1% all applications are assigned to examiners solely based on last digits. To address this limitation, we use a split-sample procedure to quantify the extent to which examiners get consistently assigned the same last digits. We split our main sample into two 50% samples at random. For each of the two subsamples, we compute the share of each last digit in an examiner’s pool of assigned applications. We then test whether the shares computed in the first subsample are predictive of those in the second subsample (comparing assigned shares for the same examiner in the same year in the two samples). Panel B of Figure 2 presents the results. For the art units that use last digits to allocate applications according to the Chi-square test (p-value < 0.01), we find a strong correlation between the last digit shares that were independently estimated in the two subsamples, with a slope close to one. This result indicates that the use of last digits for allocation of patents is quantitatively important (i.e. the Chi-square tests are not identifying statistically significant but quantitatively small rejections of the null that last digits are not used for application assignment). In contrast, in the art units for which we cannot reject that last digits are not used for application assignment (p-value > 0.01), there is no relationship between the last digit shares across the two samples. The two panels of Figure 2 thus establish that there is a large number of art units that use last digits for application assignment and that they do so in a quantitatively important way.

Panel A of Table 3 shows that the signal standard deviations estimated for art units that allocate patents using last digits are quantitatively similar to those from the baseline design. Column (1) shows the signal standard deviations for various outcomes in the sub-sample of art units with a p-value below 0.01 in the Chi-square test. Moreover, Column (2) repeats the estimation of the signal standard deviation in the subsample of art units belonging to Information Technologies.³³ The results are similar in this subsample as well, which is comforting because [Righi and Simcoe \(2017\)](#) report that they find no evidence of examiner specialization in Information Technologies.³⁴

Alternative source of variation #2: a busyness instrument. A limitation of using art units that allocate applications using last digits is that these art units account for only about a third of all art units. There is anecdotal evidence that some art units allocate applications to examiners based

³³This subsample includes the following technology centers: Computer Architecture and Software (21); Computer Networks, Multiplex, Cable and Cryptography/Security (24); Communications (26); and Business Method art units (3620s, 3680s, 3690s). We exclude technology center 2800 (Semiconductors), which [Righi and Simcoe \(2017\)](#) identify as having significant examiner specialization.

³⁴The signal standard deviation for patent value from [Kogan et al. \(2017\)](#) is smaller in the IT subsample (3) than in the full sample (Table 2). But this is due to heterogeneity in the signal standard deviation of examiner effects across technology categories, rather than to endogeneity concerns: Online Appendix Table A2 reports smaller signal SDs for patent value in IT-related technology categories.

on the timing of arrival of applications (Lemley and Sampat (2012)). When a new application arrives at the patent office, an examiner who recently finished processing another application may be particularly likely to be assigned the new application, because they happen to have more time on their hands.

To proxy for how busy an examiner is when a given new application arrives, we measure the number of cases closed by the examiner in the two preceding weeks. For each incoming application, we compute assignment probabilities across all examiners working in the relevant art unit and time period based on the number of cases closed in the previous two weeks, art unit by year fixed effects and examiner fixed effects. Within an art unit and a year, assignment probabilities vary only because of changes in (relative) busyness across examiners. We estimate assignment probabilities using a simple linear probability model, presented in Online Appendix A.

Using the estimated assignment probabilities across examiners, we instrument for the characteristics of the examiner who actually processed the application. For instance, if an application arrives in the art unit at a time when only “lenient” patent examiners (who tend to ask the applicant to make only a few changes to the patent) happen to be free, then the application should be more likely to receive a more lenient treatment. Using this source of variation, we can document the relationship between any given examiner characteristic and any patent outcomes. Specifically, we can use the estimated assignment probabilities to compute the expected examiner characteristic, which we can relate to the actual characteristic of the examiner who handled the application (the “first stage”) and to any patent outcome of interest (the “reduced form”).

Figure 3 presents the results of the busyness approach. The panels are based on the following specifications:

$$E_{j(i)} = \beta_1 \left(\sum_{j \in ut(i)} p_{ij} E_j \right) + a_{ut(i)} + \nu_i, \quad (5)$$

$$Y_i = \beta_2 \left(\sum_{j \in ut(i)} p_{ij} E_j \right) + a_{ut(i)} + \kappa_i, \quad (6)$$

where i indexes the patent, j the examiner, u the art unit and t time. p_{ij} denotes the application-specific examiner assignment probability; E_j denotes the examiner characteristic, measured using a leave-one-out procedure that does not use information on patent i ; $E_{j(i)}$ is the (leave-one-out) characteristic of the examiner who actually processed application i ; and Y_i is the patent outcome of interest. Figure 3 estimates these specifications, considering the (leave-one-out) change in the number of words per claim as the examiner characteristic and the (actual) purchase by a PAE as

the outcome of interest. This choice of variables allows for a comparison with Figure 1, which did not use the busyness instrument and was using raw variation in the examiner’s propensity to change the number of words per claim between application and grant.

Panel A of 3 reports the relationship between the actual and expected examiner characteristics, as in (5). The slope is strong and positive and the binned scatter plot is close to linear, indicating that the busyness instrument has power. Panel B of 3 shows the relationship between PAE purchase and the expected examiner propensity to increase the number of words per claims: there is a strong downward relationship. These patterns are similar to Figure 1, which used the raw variation in examiner characteristic instead of the busyness instrument. These results provide additional evidence that departures from random assignment of examiners to applications do not bias our estimates.

Accounting for potential selection effects on the extensive margin. Another potential concern with our baseline research design is that our estimates may be confounded by selection effects stemming from the decision to grant a patent. Examiners differ in their grant rates, therefore it could be the case that patent outcomes vary across examiners because of underlying differences across examiners’ pools of granted patents, independently of the crafting of patent rights. For instance, examiners with a low grant rate might only grant patents of high technical merit. To investigate this possibility, we introduce controls for the examiner’s leave-one-out grant rate in equation (1) and then repeat the estimation of the signal standard deviation using equation (3). With this specification, we are now estimating the amount of systematic variation in patent outcomes across examiners who work in the same art unit, in the same year, and have the same grant rate.

Panel B of Table 3 reports the results and shows that our baseline estimates remain virtually unaffected. Column (1) controls for the grant rate in (3).³⁵ The estimated signal standard deviations are very similar to our baseline estimates from Table 2. In principle, it may be possible for extensive margin effects to operate even across examiners with the same grant rate. For instance, an examiner may systematically grant patents with underlying technological characteristics that appeal to PAEs, while another examiner (with a similar overall grant rate) may tend to systematically reject those patents and grant others. To assess how strong this effect might be empirically, Column (2) introduces controls for a host of initial characteristics of the patent application, namely: the application’s initial number of independent claims and number of words per claim; the assignee’s

³⁵To flexibly control for the grant rate, we introduce a quartic polynomial in the grant rate. The results are similar when controlling linearly for the grant rate or introducing higher-order polynomials (not reported).

number of applications, grants and citations prior to the filing date; and the first inventor’s number of applications, grants and citations prior to the filing date. The estimates of signal standard deviations are not sensitive to these controls, indicating that extensive margin effects are unlikely to bias our estimates in any meaningful way.

Accurately accounting for excess variance. The preceding discussion indicates that our results are robust to failures of random assignment and extensive margin selection effects. A remaining potential concern is that the Empirical Bayes shrinkage correction used in our baseline research design may fail to account for noise perfectly. To address this point, we first discuss some plausible limitations of our baseline design, in particular for rare binary outcomes such as litigation; we then present an alternative approach which addresses these limitations and produces similar results. Finally, we use out-of-sample tests to directly show that our baseline design accurately accounts for excess variance.

Our baseline research design yields very large signal standard deviation estimates for rare binary outcomes, such as litigation or purchase by a PAE, but the Bayesian shrinkage correction may not be appropriate in such cases. Indeed, for binary outcomes our statistical model in equation (1) may be misspecified as it does not impose the constraint that the predicted value should lie between zero and one. Given that rare binary outcomes have a particularly high estimated signal standard deviation in Table 2, it appears important to assess whether these results are sensitive to a change in the underlying statistical model.

We repeat the analysis using an Empirical Bayes Beta-Binomial count model, a common statistical model that can fit count data in a flexible way (Ellison and Swanson (2010)). To see how this framework operates, consider the example of patent purchases by PAEs. For each examiner j , we observe data of the form (n_j, r_j) , where n_j is the examiner’s total number of granted patents and r_j is the number of patents granted by the examiner that were purchased by PAEs. We assume that the probability p of granting a patent purchased by a PAE follows a Beta distribution across examiners working in the same art unit in the same year: $p \sim \text{Beta}(\alpha, \beta)$. Given that we are examining the count of PAE purchases across examiners, the likelihood function for the data is a binomial distribution. Using the fact that the beta distribution is the conjugate prior of the binomial distribution, we show in Online Appendix A that the integrated likelihood is:

$$L(r_j | n_j, \alpha, \beta) = \binom{n_j}{r_j} \frac{\Gamma(\alpha + \beta) \Gamma(r_j + \alpha) \Gamma(n_j - r_j + \beta)}{\Gamma(\alpha) \Gamma(\beta) \Gamma(n_j + \alpha + \beta)},$$

which we estimate via maximum likelihood in each art unit by year. Having recovered estimates of

the hyperparameters, $\hat{\alpha}$ and $\hat{\beta}$, we compute the posterior mean for each examiner:³⁶

$$\widehat{\mu}_j^{BetaBinomial} = \frac{\hat{\alpha} + r_j}{\hat{\alpha} + \hat{\beta} + n_j}. \quad (7)$$

Panel C of Table 3 reports the standard deviation of the estimates: we continue to find large examiner effects. This finding indicates that our large estimates for the impact of examiner on patent litigation and purchase by PAEs is not an artifact of the statistical model used in our baseline design.

To conclude this section, we conduct out-of-sample tests of the examiner effects estimated in our baseline research design to check that we have recovered estimates of the correct magnitude. After splitting the main analysis sample into two 50% samples at random, in each subsample we compute the raw examiner effects using equation (2) and the shrunk examiner effects using equation (4). To test predictive accuracy, we regress the raw examiner effect from the first subsample on the shrunk examiner effects from the second subsample.³⁷ We also regress the raw examiner effect from the first subsample on the raw examiner effect from the second subsample to assess whether a standard regression approach would suffer from excess variance. We do so in the full sample but also in a reduced sample of examiners who granted more than fifty patents, as measurement error may no longer be a problem if sufficiently many patents are observed per examiner.

Figure 4 reports the results and shows that the Empirical Bayes shrinkage approach yields unbiased estimates of examiner effects, in contrast with standard regression analysis. A regression coefficient of one indicates unbiased prediction, while a coefficient below one indicates attenuation bias and implies that the estimates suffer from excess variance due to noise. Figure 4 shows that our baseline design delivers unbiased estimates of examiner effects even for rare outcomes such as patent purchase by PAEs or patent litigation. The point estimates are very close to one and are precisely estimated. In contrast, the specifications without shrinkage always deliver a coefficient well below one, indicating that the raw variation in examiner effects contains a lot of noise. This problem is less acute for outcomes that are more common, such as the patent value measure of [Kogan et al. \(2017\)](#) (with a regression coefficient close to 0.5 full sample), than for rare outcomes like patent litigation (with a regression coefficient close to 0.1 in the full sample). Restricting the analysis to examiners who handle a lot of patents does not solve the problem, which offers another vindication of our baseline research design.

³⁶Intuitively, this procedure shrinks an examiner’s PAE share towards the mean PAE share in the art unit. The amount of shrinkage is larger for examiners who have granted fewer patents.

³⁷We regress raw effects on shrunk effects because the shrinkage factor in the shrunk effects addresses measurement error, which poses an issue for the independent variable but not for the dependent variable.

III.D Robustness Checks

Table 4 shows the robustness of the signal standard deviations when using alternative samples and specifications. The first row repeats the analysis including continuation applications; the second row includes all granted patents from 1976 to 2015; the third row controls for the length of time between filing and grant to assess whether the results may be driven by delays rather than by the way patent rights are crafted; the fourth row includes fixed effects for examiner experience as in [Frakes and Wasserman \(2017\)](#).³⁸ The results are very similar across samples and specifications. Finally, the Online Appendix shows that the signal standard deviations are of comparable magnitudes across technology categories (Online Appendix Table A2) and reports the distributions of the shrunk examiner effects (Online Appendix Figure A2).

IV Implications for Patent Assertion Entities

Our analysis so far has established that the crafting of patent rights is an important driver of a wide range of patent outcomes, in particular those related to PAEs and litigation. In this section, motivated by the large sensitivity of PAEs to the way examiners craft patent rights, we investigate the features of examiner behavior that drive PAEs’ responses. We find that “lenient” examiners, who issue patents with higher litigation and invalidity risks, produce a much higher share of patents purchased and asserted by PAEs. We discuss how this evidence helps discipline theories of PAE behavior.

IV.A Research Design

There are two standard views of the role played by PAEs in the patent market. According to the first view, PAEs could be useful intermediaries who address standard frictions in the patent market by lowering transaction costs and solving liquidity problems ([Hagi and Yoffie \(2013\)](#), [Abrams et al. \(2016\)](#), [Lu \(2012\)](#) and [Galetovic et al. \(2015\)](#)). The second view suggests that PAEs do not help address any particular friction but, rather, exploit limitations of the legal system by asserting patents of questionable validity in the hope that targeted firms will pay them settlement fees instead of risking a costly and uncertain trial ([Miller \(2013\)](#), [Council of Economic Advisers \(2013\)](#), [Cohen et al. \(2016\)](#) and [Federal Trade Commission \(2016\)](#)).

³⁸An alternative to the inclusion of examiner experience fixed effects in our baseline specification is to look for discontinuities in patent outcomes around examiners’ promotions; we find no discontinuity (Online Appendix Figure A1), which confirms that examiner experience effects play a second-order role compared with the examiner fixed effects we focus on.

We investigate the extent to which the two standard views can account for the (quantitatively large) patterns related to examiners in the data. The way examiners craft patent rights has a first-order impact on PAEs: a one standard deviation change in examiner effects shifts the probability of patent acquisition by a PAE by over 60% of the baseline rate (Table 2). This fact may not be incompatible with the two standard views of PAE behavior. For instance, the process of creation of patent rights may create frictions affecting both PAEs and practicing firms (in line with the first view) or may lead to the issuance of questionable patents that only PAEs are willing to purchase and exploit via frivolous litigation (in line with the second view).

We examine this question using detailed data on the prosecution behaviors of examiners, drawing a contrast between the responses of PAEs and practicing firms. We start by characterizing the prosecution behaviors that are predictive of future purchase or litigation by a PAE or practicing firms (Section IV.B); we then investigate whether these prosecution behaviors are predictive of patent invalidity (Section IV.C). Specifically, we run regressions of the following form:

$$Y_i = \beta E_{j(i)} + a_{ut(i)} + \epsilon_i, \quad (8)$$

where i indexes the patent, j the examiner and ut the art unit-by-year; Y_i is the patent outcome of interest; and $E_{j(i)}$ is a (vector of) examiner behavior(s), estimated using a leave-one-out procedure that does not use information on patent i . We scale the examiner behavior measures $E_{j(i)}$ by their signal standard deviations, which are estimated using (3). This standardization gives us the proper scaling to compare the quantitative importance of various examiner traits.³⁹

We rely on a variety of proxies reflecting different aspects of examiner behavior to isolate robust correlations with the potential to inform theories of PAE behavior. The estimates from specification (8) cannot be interpreted as causal because quasi-random assignment occurs at the level of examiners working in the same art unit at the same time, and not at the level of examiners’ traits. Given that quasi-random assignment is at the level of examiners, the only causal effect that can be recovered is the effect of the examiner “as a whole” on patent outcomes (as in Section III).⁴⁰ In contrast, the relationships between specific examiner traits and patent outcomes may be biased by potential omitted variables (i.e. other traits of the examiner that are unobserved). To address this limitation, we use several proxies to control for various aspects of examiner behavior and we focus on establishing correlations which (1) are quantitatively large and robust to the inclusion of

³⁹Specification (8) is analogous to the regression underlying Figure 1, except that we are now using properly scaled regressors.

⁴⁰One would need a quasi-experiment that directly affects specific behaviors (e.g., a training program) in order to recover more granular causal impacts.

additional controls; and (2) can be interpreted as reflecting a more general trait of the examiner, such as the propensity to let the applicant keep the text of the claims relatively unchanged between application and grant (“leniency”).

IV.B PAEs and Examiner Behavior

In this subsection, we document which examiner traits correlate with patent acquisition or litigation by PAEs and practicing entities. We use specification (8) and consider seven measures that capture different aspects of examiner behavior.

We use three general proxies for the degree of “leniency” of the examiner. By examiner leniency, we refer to the extent to which the examiner makes demands on the applicant during prosecution. First, the percentage change in the number of words per claim (averaged across claims) indicates the extent to which the examiner asks the applicant to refine the claims. Second, the percentage change in the number of claims reflects the extent to which the examiners affects the overall structure and scope of the patent document. Third, the examiner’s grant rate can be interpreted as another proxy for leniency, given that examiners who are more demanding on applicants also have lower grant rates.

To characterize in greater detail the examiner behaviors that drive PAEs’ activities, we measure examiners’ propensities to cite specific sections of patent law when asking the applicant to revise the patent. As mentioned previously, the examiner must substantiate any demand by referring to specific sections of patent law corresponding to various standards of patentability. An examiner who is less lenient should tend to refer more often to any of the sections compared with other examiners working in the same art unit at the same time. The relative frequency of usage of the various sections may differ across examiners depending on their examination styles. Examiners who place more emphasis on the invention being useful and eligible for a patent should use section 101 more often; those who particularly care about prior inventions should refer section 102 frequently; section 103 should be invoked more often by examiners who are particularly sensitive to the requirement that the invention should be non-obvious to someone who knows the field; and section 112(b) should be used by examiners who focus on the requirement of claim clarity.⁴¹

Table 5 presents the results with patent acquisition as the outcome.⁴² In both panels, the

⁴¹Although all examiners are supposed to apply the same standards for patent grant, which are determined by patent law, we find large causal variation across examiners in terms of their propensity to refer to the various sections (Online Appendix Table A3)

⁴²The sample is restricted to art units that are part of Information Technologies since PAEs are primarily active in these art units (Online Appendix Table A4). All results reported in this section are similar in the full sample (Online Appendix Tables A5, A6 and A7).

first seven columns run univariate regressions, while columns (8) and (9) consider multivariate regressions. Panel A shows that all proxies of examiner leniency deliver a similar message: more lenient examiner grant substantially more patents that are eventually purchased by PAEs. The regression coefficients are standardized by the signal standard deviations of the regressors and expressed as a percentage of the outcome. Column (1) shows that a one standard deviation increase in the distribution of examiner effects for the change in number of words per claim implies a 13.9% decrease in the probability of purchase by a PAE. This fraction is relatively large, given that a one standard deviation change in the overall examiner effect accounts for about 60% of the baseline rate (Table 2). Columns (2) and (3) show that the effect goes in the same direction, with a similar magnitude, for the other broad proxies for examiner leniency: a one standard deviation increase in the change in number of claims implies a 7.3% increase in the probability of PAE purchase,⁴³ the corresponding number for grant rates is 11.4%. Columns (4) to (7) show that the same finding holds when considering the use of various sections of patent law: examiners who use sections more often tend to have a lower rate of purchase by PAEs (although some specifications are noisy). Column (8) presents the results of a specification that simultaneously includes all types of references to patent law. In this specification, the section relating to the obviousness of the invention is the most important. Finally, specification (9) includes all regressors simultaneously. The results become more noisy because of collinearity, but the coefficient on the change in the number of words per claim remains large, significant and similar in magnitude to the univariate regression in Column (1). These findings show that PAEs have a preference for purchasing patents that were issued by lenient examiners.

Panel B of Table 5 shows the results for patent purchase by practicing firms, which stand in sharp contrast with the patterns for PAEs. First, the effects are all much smaller in magnitude than in Panel A. In the first seven columns of the table, the effects are almost all insignificant and are never larger than 2%. Second, the relationship with examiner leniency does not appear to be robust: it switches signs across proxies or specifications. For instance, in the univariate regression in Column (1) we obtain a precisely estimated zero for the correlation with the change in the number of words per claim. But the regression coefficient becomes positive in specification (9), suggesting that practicing firms may have a preference for less lenient examiners, although the coefficient is relatively small (3.49%). Overall, there appears to be no quantitatively large or statistically robust

⁴³More lenient examiners tend to reduce the number of claims by less, which means that a higher change in the number of claims (in absolute value) reflects higher leniency. In contrast, a more lenient examiner increases the number of words per claim by less, i.e. a higher change in the number of words per claim reflects lower leniency.

relationship between purchases by practicing firms and examiner leniency.

The fact that only PAEs selectively purchase patents issued by lenient examiners is not consistent with the view that PAEs solve a generic friction in the patent market. If PAEs were primarily lowering transaction costs or solving liquidity problems, there would be no reason for them to selectively purchase patents from lenient examiners, which in contrast do not affect patent acquisitions by practicing firms. To examine whether PAEs may rather be addressing a patent-specific friction related to the patent examination process itself, we now investigate the correlates of patent litigation.

Table 6 presents the results with patent litigation as the outcome. Panel A reports the results for patent litigation by PAEs. The patterns are similar to those found in Table 5 for PAEs, except that the magnitudes are much larger. Column (1) shows that a one standard deviation increase in the examiner effect for the change in the number of words per claim implies a 40.5% increase in the rate of litigation by PAEs. This effect is very large in itself but also relative of the overall examiner effects documented in Table 2, according to which the signal standard deviation of examiner effects for PAE litigation is 46% (although it is imprecisely estimated). This result suggests that a simple proxy for examiner leniency can account for most of the relationship between examiner effects and PAE litigation. Moreover, the other columns of Table 5 indicate that this pattern is very robust. The other general proxies for examiner leniency, the change in the number of claims and the grant rate, go in the same direction and are larger in magnitude than when considering patent purchases. Considering the use of the various sections of patent law, as for patent purchase by PAEs the section relating to the obviousness of the invention is the most important, but the magnitude of the effect is now substantially larger. In the multivariate regression including all examiner effects simultaneously in Column (9), the patterns still point to the role of leniency as the predictive power loads on the grant rate, with a coefficient indicating that a one standard deviation increase in the grant rate implies an increase in the rate of PAE litigation close to 50%.

Panel B of Table 6 reports the results for patent litigation by practicing firms, which are qualitatively similar to the patterns for PAEs but are smaller in magnitude. Across all proxies and specifications in this panel, we consistently find that lenient patent examiners — who increase the number of words per claim by less, have a higher grant rate and reference patent law less often — issue patents with a higher litigation risk. The magnitude of the effects is less strong than for litigation by PAEs but is comparable to the magnitude of the effects for purchases by PAEs (Panel A of Table 5). For instance, a one standard deviation fall in the examiner effect for the change

in the number of words per claim implies a 13.8% increase in the rate of litigation and a 13.9% increase in the rate of PAE purchase.

The finding that patent litigation by both practicing firms and PAEs is driven by examiner leniency challenges the view that PAEs engage in idiosyncratic frivolous lawsuits. The merit of the lawsuits involving patents issued by lenient patent examiners may be questionable, but PAEs are not the only entities to selectively assert patents from lenient examiners: practicing firms do so as well. PAEs purchase patents that are different from those handled by practicing firms in the market for patents (Table 5) but their propensity to assert patents issued by lenient examiners is merely a more extreme version of the litigation behavior of practicing firms (Table 6).

The patterns in the data are therefore difficult to reconcile with the mainstream views of PAEs, either as intermediaries solving a generic friction in the patent market or as perpetrators of frivolous lawsuits. Rather, it appears that much of the activities of PAEs is driven by a specific friction in the patent market, which is caused by the way examiner craft patent rights and which strongly correlates with examiner leniency. Our findings are therefore in line with a nuanced view of PAEs, suggesting that PAEs' activities are the symptom of features of the patent system that affect litigation more generally (e.g., [Lemley and Melamed \(2013\)](#) and [Schwartz and Kesan \(2013\)](#)). PAEs behave as litigation experts and much of their activities stem from the way patent rights are crafted by lenient examiners, who affect litigation more generally. Although we can only document correlations with examiner traits, we emphasize that the underlying causal examiner effects are quantitatively large and should therefore be accounted for by any convincing theory of PAEs' activities.⁴⁴

IV.C PAEs and Patent Invalidity

In this subsection we study whether lenient examiners, who play an important role for litigation in general and for litigation by PAEs in particular, tend to issue patents that are more likely to be invalid. Various observers (e.g., [Federal Trade Commission \(2016\)](#)) have hypothesized that PAEs may be asserting patents that are “invalid”, in the sense that these patents should not have been issued in the first place because they do not comply with the standards set by U.S. patent law. Given the evidence that patent litigation by PAEs is very strongly correlated with examiner leniency, we can re-cast this question in terms of examiner effects: do lenient examiners tend to issue patent that are more likely to be invalid? Approaching this question in terms of examiner effects has the

⁴⁴Of course, even though the causal examiner effects from Table 2 are large, they do not account for the entirety of PAEs' patent acquisition and assertion behaviors. We only speak to the (substantial) part of PAEs' activities which is caused by examiner effects and point out that the two standard views of PAEs cannot account for these patterns.

potential to be informative about PAEs but also about patent litigation by practicing firms, since they also selectively assert patents that were crafted by lenient examiners.

Proxies for Patent Invalidity. Patent invalidity is notoriously difficult to measure because of selection effects (e.g., [Miller \(2013\)](#)). To assess whether a robust relationship exists between examiner leniency and patent invalidity, we rely on three complementary proxies for patent invalidity. We consider two restricted samples to study two common proxies for patent invalidity, which are subject to substantial sample selection but are standard in the literature. We also introduce a third proxy available in the full sample of granted patents.

First, for a small number of cases, patent litigation does not result in a settlement and a court trial closes the case (see [Allison et al. \(2013\)](#) for a review). We obtain this data from Lex Machina. The sample of cases for which trial outcomes are available is very selected: in our main analysis sample, there are only 516 cases with information on whether the court deemed the patent invalid or found an infringement.

The second common proxy for patent invalidity is a procedure for challenging the validity of a patent at the USPTO, known as an “inter partes review” (IPR). IPRs were introduced in 2012 as a defensive tool for those seeking to defeat meritless infringement claims (see [Chien and Helmers \(2015\)](#) for a review). The procedure can be initiated by any party other than the patent owner and requires the patent office to review the validity of the patent based on specific sections of patent law. This sample is also very selected: there are 989 IPR cases in our main analysis sample.

Third, we use patent re-issuance requests as another proxy for patent invalidity. A re-issue application can be filed by the applicant “whenever any patent is, through error, deemed wholly or partly inoperative or invalid”.⁴⁵ We obtain this information from the continuation data in the Patent Examination Dataset. Re-issue applications are a useful metric for our purposes as they are available for all granted patents and provide a direct measure of examiner mistakes from the perspective of the patent applicant.

Table 7 reports summary statistics on our proxies for patent invalidity. Court rulings are observed for only 516 patents, or about 0.0004% of our sample. Conditional on observing a court

⁴⁵Patent law states that “Whenever any patent is, through error, deemed wholly or partly inoperative or invalid, by reason of a defective specification or drawing, or by reason of the patentee claiming more or less than he had a right to claim in the patent, the Director shall, on the surrender of such patent and the payment of the fee required by law, reissue the patent for the invention disclosed in the original patent, and in accordance with a new and amended application, for the unexpired part of the term of the original patent.” (35 USC 251(a)). Re-issue applications can petition for an increase in the scope of claims only if they are filed within two years from grant of the original patent (35 USC 251(d)). We repeat our analysis considering only re-issues applications beyond this threshold to establish that attempts to increase claim scope are not driving the patterns.

ruling, the rate of invalidity is close to 19%. In 31.9% of cases, the court declares that the patent is infringed, which indirectly attests to its validity. The panel also indicates that an IPR procedure is filed for 0.0003% of patents. Conditional on filing, 78.5% of IPRs are “instituted”, meaning that the patent office deems it likely that the patent is at least in part invalid.⁴⁶ Because the “institution” rate of IPRs is very high, close to 80%, either the occurrence of an IPR or the institution of an IPR can be used as proxies for patent invalidity. For both court rulings and IPRs, the invalidity rates appear to be high, but they are observed conditional on a stringent form of sample selection.

Finally, Table 7 shows that re-issue applications are submitted for about 0.002% of patents. According to patent law, a re-issue application indicates that the applicant believes that the patent is wholly or in part invalid because of a mistake in the document. To address the potential concern that some applicants may violate patent law and strategically exploit re-issue applications to obtain greater scope, instead of correcting a mistake, we consider re-issue applications that are submitted more than two years after grant. After the two-year delay, re-issue applications cannot petition for an increase in scope; they account for about 0.0004% of all granted patents. This fraction is very small but it is comparable in magnitude to the number of observations for court rulings and IPRs and has the advantage of being available for the full sample of granted patents.

Results. We run specification (8) with our patent invalidity proxies as outcomes. The regressors are examiner effects for the change in the number of words per claim and the grant rate, which were the most powerful univariate predictors of patent acquisition and assertion by PAEs in Tables 5 and 6. We also consider the best linear predictor for patent purchase by PAEs using the specification in Column (9) of Table 5. The results are reported in Table 8.

We find a very strong and robust relationship between examiner leniency and our preferred proxy for patent invalidity, the reissuance of granted patents. Panel A of Table 8 reports this finding. The various rows of this panel correspond to separate univariate regressions. The first row of Column (1) indicates that, conditional on year fixed effects, a one standard deviation increase in the examiner effect for the change in the number of words per claim (i.e. less leniency) leads to a 26% decline in the probability of reissuance. Columns (2) and (3) show that the coefficient is very stable as art unit by year fixed effects and art unit by year by technology class fixed effects are introduced. Similarly strong and robust patterns are documented in the other rows of the tables for the grant rate and the linear predictor for PAE acquisition. Column (4) to (6) show that the

⁴⁶According to patent law, “An *inter partes* review may be instituted upon a showing that there is a reasonable likelihood that the petitioner would prevail with respect to at least one claim challenged” (35 USC Ch. 31, §311 - §319).

patterns are even stronger when we consider the reissuance rate two years or more after grant, the delay beyond which a reissuance request cannot petition for an increase in the scope of the claims. For instance, the coefficient for the change in the number of words per claim hovers between 55% and 61% across specifications. Since PAEs selectively assert patent granted by lenient examiner (more so than practicing firms), they are more likely to assert patents that are likely to contain mistakes, as reflected by the reissuance rates.

Panel B of Table 8 shows that common proxies for patent invalidity based on court rulings cannot deliver conclusive results due to data limitations. For a small sub-sample of litigated patent, we observe rulings in which the courts may indicate that the patent is invalid (Columns (1) to (3)) or that an infringement is found (Columns (4) to (6)). The various regression coefficients reported in this panel show that with such proxies the research design is under-powered, regardless of the set of fixed effects. The points estimates switch signs across specifications and are very imprecisely estimated.

Panel C of Table 8 uses IPR occurrence and IPR institution as proxies for patent invalidity from the perspective of the Patent Office. Columns (1) to (3) of Panel C of Table 8 document that examiner leniency is a very strong predictor of the occurrence of an IPR. For instance, the first row of Column (2) indicates that a one standard deviation increase in examiner effects for the change in the number of words per claims (lower leniency) implies a 41% fall in the probability of an IPR. The regression coefficients are all large and very stable across specifications that include different sets of fixed effects. In contrast, Columns (4) to (6) do not deliver conclusive results regarding IPR institution, because the selected sample of patents that go through an IPR is too small to provide adequate power.

In sum, Table 8 indicates that, when using suitable proxies for patent invalidity that do not suffer from small sample issues, there is strong and robust evidence that lenient examiners issue patents that are more likely to be invalid. These examiners account for a disproportionate share of patent litigation, in particular by PAEs. This finding indicates that examiner behavior during patent prosecution is a quantitatively important determinant of patent invalidity, suggesting that PAEs specialize in purchasing and asserting patent that should not have been issued as such in light of the standards set by current patent law.⁴⁷

⁴⁷This finding does not speak conclusively to the welfare effects of PAEs, because litigation of patents issued by lenient examiners could conceivably be socially valuable, even when these patents are deemed invalid by current patent law. The standards set by current patent law may not be social optimal and are dynamically evolving. For instance, [Galetovic et al. \(2015\)](#) point out that the process of litigation helps defines the contours of patent law in highly-innovative, rapidly changing industries.

IV.D Robustness Checks and Additional Results

In the final part of this section, we discuss the robustness of our PAE results across samples, specifications, and PAE types. In addition, we use data on patent value, auction prices, and European Patent Office decisions to shed further light on PAE behavior.

Robustness across samples and specifications. Table 9 documents the robustness of the signal standard deviations of examiner effects for PAE purchases across alternative specifications and subsamples. Row (A) reports the baseline estimate in our main analysis sample, as in Table 2. Row (B) shows that the signal standard deviation remains similar when introducing assignee fixed effects in equation (1): PAEs selectively purchase patents coming from specific examiners even within the portfolio of a given assignee. Row (C) to (E) show that the signal standard deviation is very similar across PAE lists. Row (C) reports similar estimates when excluding from the sample the patents purchased by the largest PAE, Intellectual Ventures. Conversely, Row (D) shows that the results are comparable when considering only patents purchased by Intellectual Ventures.⁴⁸ The estimates also remain stable when using the list of PAEs defined by [Cotropia et al. \(2014\)](#), as shown in Row (E).

Results by PAE type. Existing research has hypothesized that large and small PAEs may behave differently ([Cotropia et al. \(2014\)](#)). Row (F) of Table 9 shows that PAEs with a small portfolio of patents, as defined by [Cotropia et al. \(2014\)](#), are as responsive to examiner effects as the average PAE. Another plausible hypothesis is that PAEs that primarily work with small firms or individual inventors may have a different behavior with respect to examiners, for instance because they may be focused on addressing frictions that specifically affect these firms and inventors. Row (G) considers a subset PAEs which bought over 50% of their patents from small entities.⁴⁹ In this subsample as well, the signal standard deviation is very similar to the baseline.

Furthermore, Figure 5 investigates whether different types of PAEs react differently to examiner leniency. Using specification (8), this figure reports the correlation between PAE purchase rates

⁴⁸The estimates reported in rows (C) and (D) do not average out to the estimate in (A), implying that there is not as much covariance between the two outcomes (purchase by Intellectual Ventures and purchase by a PAE other than Intellectual Ventures) as there is within outcomes. This result indicates that there is some segmentation of the market between PAEs, and that examiners effects are strong everywhere.

⁴⁹We define patents from small entities as patents that either were unassigned (i.e., the inventor is the owner) or that were assigned to a firm that the USPTO classifies as a “small entity” (if there is an assignee, each patent reports whether it was initially assigned to a small entity, i.e. a small firm). On average, PAEs purchase only 15.7% of their patents from small firms (19% when excluding continuation applications). Likewise, the share of unassigned patents in PAEs’ purchases is low, ranging from 6.2% when including continuations to 10.7% without continuations. These low shares are difficult to reconcile with the view that the typical PAEs is addressing frictions that specifically affect small firms or individual inventors.

and the main proxy for examiner leniency from Section IV.B, the change in the number of words per claim between application and grant. We consider in turn Intellectual Ventures, all PAEs but Intellectual Ventures, small PAEs, and PAEs which purchased over 50% of their patents from small entities. We find that they all selectively purchase patents from more lenient examiners, with relationships of very similar magnitudes across PAE groups.⁵⁰ The leniency-bias of PAEs is therefore a very stable feature.

Additional results. Although we have documented that examiners with high PAE effects are lenient, we also find that these examiners do not create greater private value for patent holders, suggesting a distinction between breadth and vagueness of claims. Online Appendix Table A8 shows a small and statistically insignificant relationship between examiner PAE effect and examiner private value effect, as measured by stock market response. This result suggests that these examiners are not simply granting patents with greater scope, which should create higher private value. An alternative explanation is that the patents they grant contain less well-defined or vaguer language, which is consistent with the negative relationship between Section 112(b) blocking action usage and non-PAE litigation shown in Panel B of Table A6.⁵¹

Finally, the Online Appendix reports additional results shedding light on the mechanisms leading PAEs to selectively purchase patents from more lenient examiners. First, in a subsample of patents for which auction prices are available, we find that patents issued by more lenient examiners tend to sell at a lower price (Online Appendix Figure A3). Second, considering all patents that were jointly filed at the USPTO and at the European Patent Office (EPO), we find that PAEs are much more likely to purchase patents that were rejected by the EPO (panel A of Online Appendix Table ??). Interpreting EPO grant decisions as a measure of a patent's inventive step size (Picard and Van Pottelsberghe de la Potterie (2011)), this result suggests that PAEs target patents that bear on more incremental, less innovative technology. Furthermore, we find that PAEs selectively purchase patents that were rejected by the EPO only when these patents were issued by specific examiners, with a large causal impact on PAE purchases (panel B of Online Appendix Table ??). It is therefore plausible that these patents are particularly productive for litigation, as they are closer to existing intellectual property than average (given the small step size revealed by EPO rejections) and their claims may be less well-defined and harder to interpret than average (given the examiners who

⁵⁰The results are similar with other proxies for examiner leniency, such as the grant rate, as well as when considering the full sample of patents instead of IT patents only (not reported).

⁵¹Section 112(b) is typically used to clarify indefinite claims language. Under a simple model, there would only be litigation in equilibrium if there is disagreement between parties, which would not happen if claims were broad but clear.

granted them).

V Conclusion

In this paper, we have shown that significant heterogeneity in patent outcomes results from the process of creation of patent rights and is independent of technical merit. We established this result by using the allocation of patent applications to examiners as a source of quasi-random variation in patent rights. To address identification concerns, we accounted for potential examiner specialization within detailed technology categories by developing new sources of quasi-experimental variation, based on assignment mechanisms at the patent office related to patent application serial numbers and examiner busyness. These techniques could be used to investigate a host of issues related to the crafting of patent rights in future research.

We have also shown that the process of creation of patent rights is of first-order importance to understand a central and much-debated feature of the U.S. innovation system, the activities of PAEs. We found that PAEs selectively purchase and litigate patents issued by “lenient” examiners; these examiners tend to issue patents that are more likely to be litigated, but not purchased, by practicing firms. These patterns are quantitatively important and cannot be accounted for by standard PAE theories, which describe PAEs either as intermediaries solving a generic friction in the patent market (such as transaction costs and illiquidity) or as perpetrators of frivolous lawsuits. Instead, we found that the activities of PAEs are best characterized as a response to a specific friction in the patent system, which is caused by the way lenient examiners craft patent rights and which affects litigation more broadly. These findings imply that policies affecting the behaviors of patent examiners, and specifically of lenient examiners, have the potential to greatly affect PAEs and litigation. In contrast, the current policy debate has focused on a possible reform of patent law to reduce PAEs’ activities and litigation, which observers have noted may be difficult to implement ([Schwartz and Kesan \(2013\)](#)).

More broadly, our results call for a greater focus on understanding the impact of the crafting of patent rights on innovation dynamics. This paper provided a set of tools to conduct such investigation and showed the explanatory power and potential policy relevance of this line of inquiry in the context of the debate over PAEs.

References

Abrams, David, Ufuk Akcigit, and Gokhan Oz, “The Patent Troll: Benign Middleman or

Stick-Up Artist?,” *Working Paper*, 2016.

Allison, John R, Mark A Lemley, and David L Schwartz, “Understanding the realities of modern patent litigation,” *Tex. L. Rev.*, 2013, *92*, 1769.

– , – , and **David Schwartz**, “How Often Do PAEs Win Patent Suits?,” *BERKELEY TECH. LJ**, 2016, *31*, 3–4.

Azoulay, Pierre, Joshua S Graff Zivin, and Jialan Wang, “Superstar extinction,” *The Quarterly Journal of Economics*, 2010, *125* (2), 549–589.

Balsmeier, Benjamin, Alireza Chavosh, Guan-Cheng Li, Gabe Fierro, Kevin Johnson, Aditya Kaulagi, Doug O Reagan, Bill Yeh, and Lee Fleming, “Automated disambiguation of us patent grants and applications,” *Working Paper*, 2015.

Bell, Alex, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen, “Who Becomes an Inventor in America? The Importance of Exposure to Innovation,” *NBER Working Paper No. 24062*, 2017.

Bessen, James and Michael J Meurer, “The Direct Costs from NPE Disputes,” *Cornell Law Review*, 2014.

Chetty, Raj and Nathaniel Hendren, “The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates,” 2016, (May).

– , **John N Friedman, and Jonah E Rockoff**, “Measuring the impacts of teachers I: Evaluating bias in teacher value-added estimates,” *The American Economic Review*, 2014, *104* (9), 2593–2632.

– , – , and – , “Measuring the impacts of teachers II: Teacher value-added and student outcomes in adulthood,” *The American Economic Review*, 2014, *104* (9), 2633–2679.

Chien, Colleen, “Startups and patent trolls,” *Stan. Tech. L. Rev.*, 2013, *17*, 461.

Chien, Colleen V and Christian Helmers, “Inter partes review and the design of post-grant patent reviews,” *Working Paper*, 2015.

Cockburn, Iain M, Samuel Kortum, and Scott Stern, “Are All Patent Examiners Equal? Examiners, Patent Characteristics, and Litigation Outcomes,” *Patents in the knowledge-based economy*, 2003.

- Cohen, Lauren, Umit G Gurun, and Scott Duke Kominers**, “The growing problem of patent trolling,” *Science*, 2016, *352* (6285), 521–522.
- Cohen, Wesley M and Daniel A Levinthal**, “Innovation and learning: the two faces of R & D,” *The economic journal*, 1989, *99* (397), 569–596.
- Commission, The Federal Trade**, “Patent Assertion Entity Activity: an FTC Study,” *Report*, 2016.
- Corporation, The RPX**, “NPE Litigation, Patent Marketplace, and NPE Cost,” *Report*, 2015.
- Cotropia, Christopher A, Jay P Kesan, and David L Schwartz**, “Unpacking Patent Assertion Entities (PAEs),” *Minnesota Law Review*, 2014.
- Ellison, Glenn and Ashley Swanson**, “The Gender Gap in Secondary School Mathematics at High Achievement Levels: Evidence from the American Mathematics Competitions,” *Journal of Economic Perspectives*, 2010, *24* (2), 109–128.
- Farre-Mensa, Joan, Deepak Hegde, and Alexander Ljungqvist**, “What is a Patent Worth? Evidence from the U.S. Patent Lottery,” *Working Paper*, 2017.
- Frakes, Michael D. and Melissa F. Wasserman**, “Is the Time Allocated to Review Patent Applications Inducing Examiners to Grant Invalid Patents? Evidence from Micro-Level Application Data,” *The Review of Economics and Statistics*, 2017.
- Galetovic, Alexander, Stephen Haber, and Ross Levine**, “An Empirical Examination of Patent Holdup,” *Journal of Competition Law & Economics*, 2015, *11* (3), 549–578.
- Gaulé, Patrick**, “Patents and the success of venture-capital backed startups: Using examiner assignment to estimate causal effects,” 2015.
- Gilbert, Richard and Carl Shapiro**, “Optimal patent length and breadth,” *The RAND Journal of Economics*, 1990, pp. 106–112.
- Golden, John M**, “Patent trolls and patent remedies,” *Tex. L. Rev.*, 2006, *85*, 2111.
- Graham, Stuart, Alan Marco, and Richard Miller**, “The USPTO Patent Examination Research Dataset : A Window on the Process of Patent Examination,” *Working Paper*, 2015, p. 117.

- Graham, Stuart JH, Alan C Marco, and Amanda F Myers**, “Patent transactions in the marketplace: Lessons from the uspto patent assignment dataset,” *Working Paper*, 2015.
- Haber, Stephen H and Seth H Werfel**, “Patent trolls as financial intermediaries? Experimental evidence,” *Economics Letters*, 2016, *149*, 64–66.
- Hagiu, Andrei and David B Yoffie**, “The new patent intermediaries: platforms, defensive aggregators, and super-aggregators,” *The Journal of Economic Perspectives*, 2013, *27* (1), 45–65.
- Hall, Bronwyn H, Adam B Jaffe, and Manuel Trajtenberg**, “The NBER Patent Citations Data File: Lessons, Insights, and Methodological Tools,” 2001.
- III, James F McDonough**, “The myth of the patent troll: an alternative view of the function of patent dealers in an idea economy,” *Emory LJ*, 2006, *56*, 189.
- Jaffe, Adam B, Manuel Trajtenberg, and Rebecca Jaffe** 1993 **geographic Henderson**, “Geographic localization of knowledge spillovers as evidenced by patent citations,” *The Quarterly journal of Economics*, 1993, *108* (3), 577–598.
- Kane, Thomas and Douglas Staiger**, “Estimating Teacher Impacts on Student Achievement: An Experimental Evaluation,” 2008.
- Klemperer, Paul**, “How broad should the scope of patent protection be?,” *The RAND Journal of Economics*, 1990, pp. 113–130.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman**, “Technological innovation, resource allocation, and growth,” *The Quarterly Journal of Economics*, 2017, *132* (2), 665–712.
- Kuhn, Jeffrey M**, “Property Rights and Frictions in the Sale of Patents,” *Working Paper*, 2016.
- **and Neil Thompson**, “The Ways We’ve Been Measuring Patent Scope are Wrong: How to Measure and Draw Causal Inferences with Patent Scope,” *Working Paper*, 2017.
- Lanjouw, Jean O and Mark Schankerman**, “Characteristics of patent litigation: a window on competition,” *RAND journal of economics*, 2001, pp. 129–151.
- Lemley, Ma and B Sampat**, “Examining patent examination,” *Stanford Technology Law Review*, 2010, *181* (2008).

- Lemley, Mark A**, “Rational ignorance at the patent office,” *Nw. UL Rev.*, 2000, *95*, 1495.
- **and A Douglas Melamed**, “Missing the Forest for the Trolls,” *Columbia Law Review*, 2013, pp. 2117–2189.
- Lemley, Mark A. and Bhaven Sampat**, “Examiner Characteristics and Patent Office Outcomes,” *Review of Economics and Statistics*, 2012, *94* (3), 817–827.
- Lerner, Josh**, “The empirical impact of intellectual property rights on innovation: Puzzles and clues,” *The American Economic Review*, 2009, *99* (2), 343–348.
- Lu, Jiaqing**, “The Economics And Controversies Of Nonpracticing Entities (NPEs): How NPEs And Defensive Patent Aggregators Will Change The License Market Part II,” *Journal of the Licensing Executives Society*, 2012, *47* (2), 147.
- Marco, Alan C, Amanda F. Myers, Stuart Graham, Paul D Agostino, and Kirsten Apple**, “The USPTO Patent Assignment Dataset: Descriptions and Analysis,” *USPTO Economic Working Paper*, 2015, *2015* (2).
- **and Richard Miller**, “Patent Examination Quality and Litigation: Is There a Link?,” *Working Paper*, 2017.
- **, Joshua D Sarnoff, and Charles deGrazia**, “Patent Claims and Patent Scope,” *Working Paper*, 2016.
- Miller, Shawn P**, “Where’s the Innovation: An Analysis of the Quantity and Qualities of Anticipated and Obvious Patents,” *Va. JL & Tech.*, 2013, *18*, 1.
- Mori, Tomoya, Koji Nishikimi, and Tony E Smith**, “A divergence statistic for industrial localization,” *The Review of Economics and Statistics*, 2005, *87* (4), 635–651.
- Moser, Petra**, “How do patent laws influence innovation? Evidence from nineteenth-century world’s fairs,” *The American Economic Review*, 2005, *95* (4), 1214–1236.
- Nordhaus, William D**, “An economic theory of technological change,” *The American Economic Review*, 1969, pp. 18–28.
- of Economic Advisers, The Council**, “Patent Assertion and U.S. Innovation,” *Executive Office of the President*, 2013.

- Pakes, Ariel**, “Patents as Options : Some Estimates of the Value of Holding European Patent Stocks,” *Econometrica*, 1986, *54* (4), 755–784.
- Picard, Pierre M and Bruno Van Pottelsberghe de la Potterie**, “Patent office governance and patent system quality,” *Working Paper*, 2011.
- Righi, Cesare and Timothy Simcoe**, “Patent Examiner Specialization,” *SSRN Working Paper*, 2017.
- Sakakibara, Mariko and Lee Branstetter**, “Do stronger patents induce more innovation? Evidence from the 1988 Japanese patent law reforms,” *The Rand Journal of Economics*, 2001, *32* (1), 77.
- Sampat, Bhaven and Heidi L Williams**, “How Do Patents Affect Follow-On Innovation? Evidence from the Human Genome,” 2015.
- Schwartz, David L and Jay P Kesan**, “Analyzing the role of non-practicing entities in the patent system,” *Cornell L. Rev.*, 2013, *99*, 425.
- Toivanen, Otto and Lotta Väänänen**, “Returns to inventors,” *Review of Economics and Statistics*, 2012, *94* (4), 1173–1190.
- Tucker, Catherine E**, “Patent trolls and technology diffusion: The case of medical imaging,” *Working Paper*, 2014.
- Williams, Heidi L**, “Intellectual property rights and innovation: Evidence from the human genome,” *Journal of Political Economy*, 2013, *121* (1), 1–27.
- , “How Do Patents Affect Research Investments?,” *Annual Review of Economics*, 2017, *9*, 441–469.

Table 1: Summary Statistics

Panel A: Heterogeneity in Patent Outcomes				
Outcomes	Mean	Median	S.D.	Sample Size
Patent value from Kogan et al. (2017), \$M	9.0188	2.56	25.39	356,375
4th-year fee payment rate	0.8708	1	0.3354	1,247,958
8th-year fee payment rate	0.6098	1	0.4877	697,918
12th-year fee payment rate	0.2089	0	0.4065	373,207
Total patent citation within 3 years of grant	0.5256	0	1.461	988,585
Patent citations by same assignee within 3 years of grant	0.1134	0	0.7257	988,585
Patent citations by other assignees within 3 years of grant	0.4122	0	1.1992	988,585
Rate of patent acquisition by non-PAEs	0.1965	0	0.3974	1,270,082
Rate of patent acquisition by PAEs	0.0102	0	0.10045	1,270,082
Rate of patent litigation by non-PAEs	0.0065	0	0.0804	1,270,082
Rate of patent litigation by PAEs	0.0004	0	0.0202	1,270,082

Panel B: Changes to Patent Document between Application and Grant

Outcomes	Mean	Median	S.D.	Sample Size
Change in number of words per claim, % (average over all claims)	57.32	25.24	84.58	1,110,272
Change in number of claims, %	-3.64	0	46.14	1,110,912
Use of Section 101 - Lack of utility or eligibility	0.0541	0	0.226	1,270,210
Use of Section 102(a) - Prior art exists	0.0174	0	0.130	1,270,210
Use of Section 103(a) - Obvious invention	0.419	0	0.493	1,270,210
Use of Section 112(b) - Vague claims	0.056	0	0.231	1,270,210
Patent citations added by examiner	0.185	0	0.388	1,270,210

Panel C: Heterogeneity in Examiner Behavior

	Mean	Median	SD	Sample Size
Number of years at the U.S. Patent Office	6.35	7	3.19	10,018
Number of art units active in	1.80	2	0.96	10,018
Total patent applications processed	190	119	215	10,018
Patent grant rate	0.55	0.57	0.27	10,018
Use of Section 101 - Lack of utility or eligibility	0.09	0.02	0.14	10,018
Use of Section 102(a) - Prior art exists	0.02	0.006	0.03	10,018
Use of Section 103(a) - Obvious invention	0.45	0.48	0.21	10,018
Use of Section 112(b) - Vague Claims	0.19	0.17	0.15	10,018
Rate of patent acquisition by PAEs	0.011	0	0.032	10,018

Notes: In Panels A and B, patents are the unit of observation. In Panel C, patent examiners are the unit of observation. All statistics are unweighted. See Section II.A for details on the sample and variable definitions.

Table 2: Signal Standard Deviations of Examiner Causal Effects on Patent Outcomes

	Signal S.D.		S.D. of Shrunk Effects, % of Average (3)	Sample Size, Patents/Examiners (3)
	% of Average (1)	Level (2)		
Patent value from Kogan et al. (2017), \$M	40.80 (38.94—41.95)	3.32	29.48	356,375/7937
4th-year fee payment rate	3.76 (3.64—3.91)	0.0328	2.18	1,247,958/9,543
8th-year fee payment rate	10.79 (10.40—10.82)	0.0658	6.32	697,918/8,580
12th-year fee payment rate	22.62 (21.44—23.37)	0.0472	11.50	373,207/8,289
Log total patent citation	23.79 (23.27—24.15)	0.0610	14.04	988,585/8,620
Log patent citations by same assignee	46.06 (43.62—48.63)	0.0278	25.65	988,585/8,620
Log patent citations by other assignees	24.47 (23.88—24.80)	0.0512	14.10	988,585/8,620
Rate of patent acquisition by non-PAEs	14.61 (13.60—15.41)	0.0287	7.66	1,270,082/9,564
Rate of patent acquisition by PAEs	62.96 (52.95—70.93)	0.0064	31.11	1,270,082/9,564
Rate of patent litigation by non-PAEs	64.25 (52.79—72.73)	0.0042	27.43	1,270,082/9,564
Rate of patent litigation by PAEs	46.04 (0—147.76)	0.0002	4.84	1,270,082/9,564

Notes: This table reports the signal standard deviations of examiner effects as a percentage of the mean (Column 1) and in level (Column 2), as well as the standard deviations of shrunk examiner effects (Column 3). The Bayesian shrinkage methodology used to obtain these estimates is presented in Section III.A. In Column 2, 95% confidence intervals are obtained by bootstrapping. The log patent citation variables refer to the log of one plus the number of citations within three years of grant. The patent value variable is right-winsorized at the 99th percentile. See Section II.A for details on the sample and variable definitions.

Table 3: Validation of Baseline Estimates of Examiner Effects

Panel A: Accounting for Violations of Random Assignment

	Signal S.D., % of Average	
	(1)	(2)
Patent value from Kogan et al. (2017)	44.32	15.89
Log total patent citation	21.10	22.78
Rate of patent acquisition by non-PAEs	15.38	10.01
Rate of patent acquisition by PAEs	40.01	41.25
Rate of patent litigation by non-PAEs	55.65	64.36
Sample	Art units allocating patents by last digits, according to Chi-square test	Art units in Information Technology
Number of art units	243	254

Panel B: Accounting for Extensive Margin Selection Effects

	Signal S.D., % of Average	
	(1)	(2)
Patent value from Kogan et al. (2017)	40.46	41.48
Log total patent citation	18.66	22.17
Rate of patent acquisition by non-PAEs	14.31	17.03
Rate of patent acquisition by PAEs	62.64	76.20
Rate of patent litigation by non-PAEs	63.06	89.93
Controls	Examiner grant rate	Examiner grant rate and application characteristics

Panel C: Accounting for Excess Variance with Empirical Bayes Beta-Binomial Count Model

	S.D. of Shrunk Examiner Effects, % of Average
Rate of patent acquisition by PAEs	46.72
Rate of patent acquisition by non-PAEs	7.99
Rate of patent litigation by non-PAEs	48.95

Notes: Panel A reports the signal standard deviations of several examiner effects using the Bayesian shrinkage methodology in two subsamples in which there is no examiner specialization within art unit. Panel B repeats the calculation of the signal standard deviations of examiner effects in the same sample as Table 2, but adding controls to address potential selection effects. Panel C reports the standard deviation of average shrunk examiner effects using the Empirical Bayes Beta-Binomial Count model. See Section III.C for a description of the methodologies underlying each panel.

Table 4: Robustness Checks on Examiner Causal Effects on Patent Outcomes

	Signal S.D., % of Average				
	Patent value from Kogan et al. (2017)	Log total patent citations	Purchase by practicing firm	Purchase by PAE	Litigation by practicing firm
(A) Including continuations	41.8%	24.9%	16.8%	78.9%	90.7%
(B) Granted patent from 1976 to 2015	36.7%	22.4%	15.8%	72.3%	62.6%
(C) Including Review Time Controls	40.8%	23.2%	14.6%	62.9%	64.3%
(D) Including Examiner Career Controls	41.5%	24.6%	13.5%	55.8%	55.4%

Notes: Row (A) adds continuation applications which were filed and granted between 2001 and 2012 to our baseline analysis sample, which covers the same period. Row (B) uses the sample of all non-continuation granted patents from 1976 to 2015. Row (C) controls for “time under review” in equation (1) with a quadratic polynomial in the number of years between filing and grant. Row (D) controls for examiner career effects in equation (1) with experience fixed effects, as defined in [Frakes and Wasserman \(2017\)](#) (namely, the examiner’s GS level by bins corresponding to 0-1, 2-3, 4-5, 6-7, 8+ years experience at that level). All reported values are normalized by the average in the relevant sample.

Table 5: Patent Acquisition and Examiner Behavior

Panel A: Patent Acquisition by PAEs

Leave-one-out Examiner Effects	Patent Purchase by PAE								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
% Change in Number of Word per Claim from Application to Grant	-0.139*** (0.030)								-0.115*** (0.0490)
% Change in Number of Claims from Application to Grant		0.073** (0.034)							0.0519 (0.0345)
Grant Rate			0.114*** (0.028)						-0.0298 (0.0637)
Use of Section 101 - Lack of utility or eligibility				-0.061* (0.036)				-0.0468 (0.035)	-0.0225 (0.0366)
Use of Section 102(a) - - Prior art exists					0.007 (0.021)			0.0171 (0.021)	0.00835 (0.0216)
Use of Section 103(a) - Obvious invention						-0.0602*** (0.024)		-0.050** (0.026)	-0.0223 (0.0285)
Use of Section 112(b) - Vague claims							-0.037 (0.027)	-0.004 (0.026)	-0.00392 (0.0291)
Fixed Effects				Year by Art Unit					
<i>N</i>	274,464	274,537	311,615	311,470	311,470	311,470	311,470	311,470	274,464

Panel B: Patent Acquisition by Practicing Firms

Leave-one-out Examiner Effects	Patent Purchase by Practicing Firm								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
% Change in Number of Word per Claim from Application to Grant	0.0071 (0.0081)								0.0349*** (0.011)
% Change in Number of Claims from Application to Grant		-0.0003 (0.006)							-0.0001 (0.006)
Grant Rate			0.022*** (0.0082)						0.062*** (0.012)
Use of Section 101 - Lack of utility or eligibility				0.0147** (0.0065)				0.0154** (0.00711)	0.0174** (0.0073)
Use of Section 102(a) - -Prior art exists					-0.0037 (0.005)			-0.00498 (0.00556)	-0.007 (0.006)
Use of Section 103(a) -Obvious invention						0.0065 (0.005)		0.00539 (0.00633)	0.007 (0.006)
Use of Section 112(b) -Vague claims							0.002 (0.005)	-0.00419 (0.00619)	0.003 (0.006)
Fixed Effects				Year by Art Unit					
<i>N</i>	274,464	274,537	311,615	311,470	311,470	311,470	311,470	311,470	274,464

Notes: The sample is retracted to IT patents. Regressors are standardized by their standard deviations and coefficients are expressed as a fraction of the mean of the outcome. Standard errors are clustered by examiners. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Patent Litigation and Examiner Behavior

Panel A: Patent Litigation by PAEs

Leave-one-out Examiner Effects	Patent Litigation by PAE								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
% Change in Number of Word per Claim from Application to Grant	-0.405*** (0.083)								-0.097 (0.12)
% Change in Number of Claims from Application to Grant		0.127*** (0.067)							0.05 (0.07)
Grant Rate			0.567*** (0.099)						0.48*** (0.14)
Use of Section 101 - Lack of utility or eligibility				-0.105 (0.077)				-0.09 (0.08)	0.05 (0.08)
Use of Section 102(a) - - Prior art exists					0.0178 (0.089)			0.019 (0.08)	0.023 (0.082)
Use of Section 103(a) - Obvious invention						-0.156*** (0.075)		-0.176** (0.083)	-0.039 (0.08)
Use of Section 112(b) - Vague claims							-0.0003 (0.079)	0.102 (0.08)	0.085 (0.086)
Fixed Effects				Year by Art Unit					
<i>N</i>	274,464	274,537	311,615	311,470	311,470	311,470	311,470	311,470	274,464

Panel B: Patent Litigation by Practicing Firms

Leave-one-out Examiner Effects	Patent Litigation by Practicing Firm								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
% Change in Number of Word per Claim from Application to Grant	-0.138*** (0.043)								0.017 (0.071)
% Change in Number of Claims from Application to Grant		0.022 (0.031)							-0.015 (0.034)
Grant Rate			0.24*** (0.045)						0.23*** (0.067)
Use of Section 101 - Lack of utility or eligibility				-0.068* (0.037)				-0.0205 (0.0397)	0.005 (0.04)
Use of Section 102(a) - - Prior art exists					-0.008 (0.04)			0.0150 (0.0406)	0.026 (0.04)
Use of Section 103(a) - Obvious invention						-0.075** (0.034)		-0.0387 (0.0370)	0.008 (0.04)
Use of Section 112(b) - Vague claims							-0.118*** (0.032)	-0.0978** (0.0366)	-0.065 (0.04)
Fixed Effects				Year by Art Unit					
<i>N</i>	274,464	274,537	311,615	311,470	311,470	311,470	311,470	311,470	274,464

Notes: The sample is retracted to IT patents. Regressors are standardized by their standard deviations and coefficients are expressed as a fraction of the mean of the outcome. Standard errors are clustered by examiners. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Summary Statistics on Proxies for Patent Invalidity

	Mean	Median	S.D.	Sample Size
Rate of invalidity decision by court, conditional on court ruling	0.1880	0	0.3911	516
Rate of infringement decision by court, conditional on court ruling	0.3198	0	0.4668	516
Rate of IPR filing	0.0003	0	0.0164	1,833,464
Rate of IPR institution, conditional on IPR filing	0.7858	1	0.4105	719
Rate of re-issuance	0.0020	0	0.0458	1,833,464
Rate of re-issuance more than two years after grant	0.0004	0	0.0206	1,833,464

Notes: This table reports summary statistics on several proxies for patent invalidity. See Section IV.C for variable definitions and Section II.A for information on the sample.

Table 8: Examiner Behavior and Likelihood of Patent Invalidity

Panel A: Reissuance of Granted Patents						
Leave-one-out Examiner Effects (separate regressions)	Reissuance Rate			Reissuance Rate Two Years or More after Grant		
	(1)	(2)	(3)	(4)	(5)	(6)
(A) % Change in Number of Word per Claim from Application to Grant	-0.26*** (0.07)	-0.24*** (0.06)	-0.25*** (0.068)	-0.55*** (0.15)	-0.57*** (0.14)	-0.61*** (0.15)
(B) Grant Rate	0.29*** (0.06)	0.27*** (0.06)	0.28*** (0.061)	0.54*** (0.13)	0.53*** (0.13)	0.54*** (0.13)
(C) Linear Predictor for PAE Acquisition	0.139*** (0.038)	0.136*** (0.035)	0.142*** (0.036)	0.24*** (0.079)	0.26*** (0.075)	0.27*** (0.078)
Fixed Effects	Year	Year by Art Unit	Year by Art Unit by Class	Year	Year by Art Unit	Year by Art Unit by Class
N		274,464			273,839	

Panel B: Court Rulings						
Leave-one-out Examiner Effects (separate regressions)	Invalidity Rate			Infringement Rate		
	(1)	(2)	(3)	(4)	(5)	(6)
(A) % Change in Number of Word per Claim from Application to Grant	0.02 (0.06)	0.068 (0.29)	0.11 (0.32)	-0.01 (0.06)	-0.0001 (0.24)	-0.0002 (0.28)
(B) Grant Rate	0.02 (0.03)	-0.039 (0.26)	-0.019 (0.54)	-0.03 (0.03)	0.06 (0.10)	-0.12 (0.22)
(C) Linear Predictor for PAE Acquisition	-0.057 (0.044)	-0.031 (0.19)	-0.069 (0.19)	0.01 (0.03)	-0.007 (0.14)	0.007 (0.15)
Fixed Effects	Year	Year by Art Unit	Year by Art Unit by Class	Year	Year by Art Unit	Year by Art Unit by Class
N		111			111	

Panel C: Trials at the Patent Office (“Inter Partes Reviews”)						
Leave-one-out Examiner Effects (separate regressions)	IPR Rate			Institution Rate of IPR		
	(1)	(2)	(3)	(4)	(5)	(6)
(A) % Change in Number of Word per Claim from Application to Grant	-0.38*** (0.098)	-0.43*** (0.094)	-0.42*** (0.087)	-0.03 (0.057)	-0.05 (0.27)	-0.23 (0.29)
(B) Grant Rate	0.41*** (0.085)	0.44*** (0.081)	0.44*** (0.082)	0.05 (0.047)	-0.03 (0.11)	-0.034 (0.13)
(C) Linear Predictor for PAE Acquisition	0.28*** (0.088)	0.34*** (0.086)	0.33*** (0.077)	0.024 (0.04)	0.089 (0.21)	0.21 (0.26)
Fixed Effects	Year	Year by Art Unit	Year by Art Unit by Class	Year	Year by Art Unit	Year by Art Unit by Class
N		274,537			180	

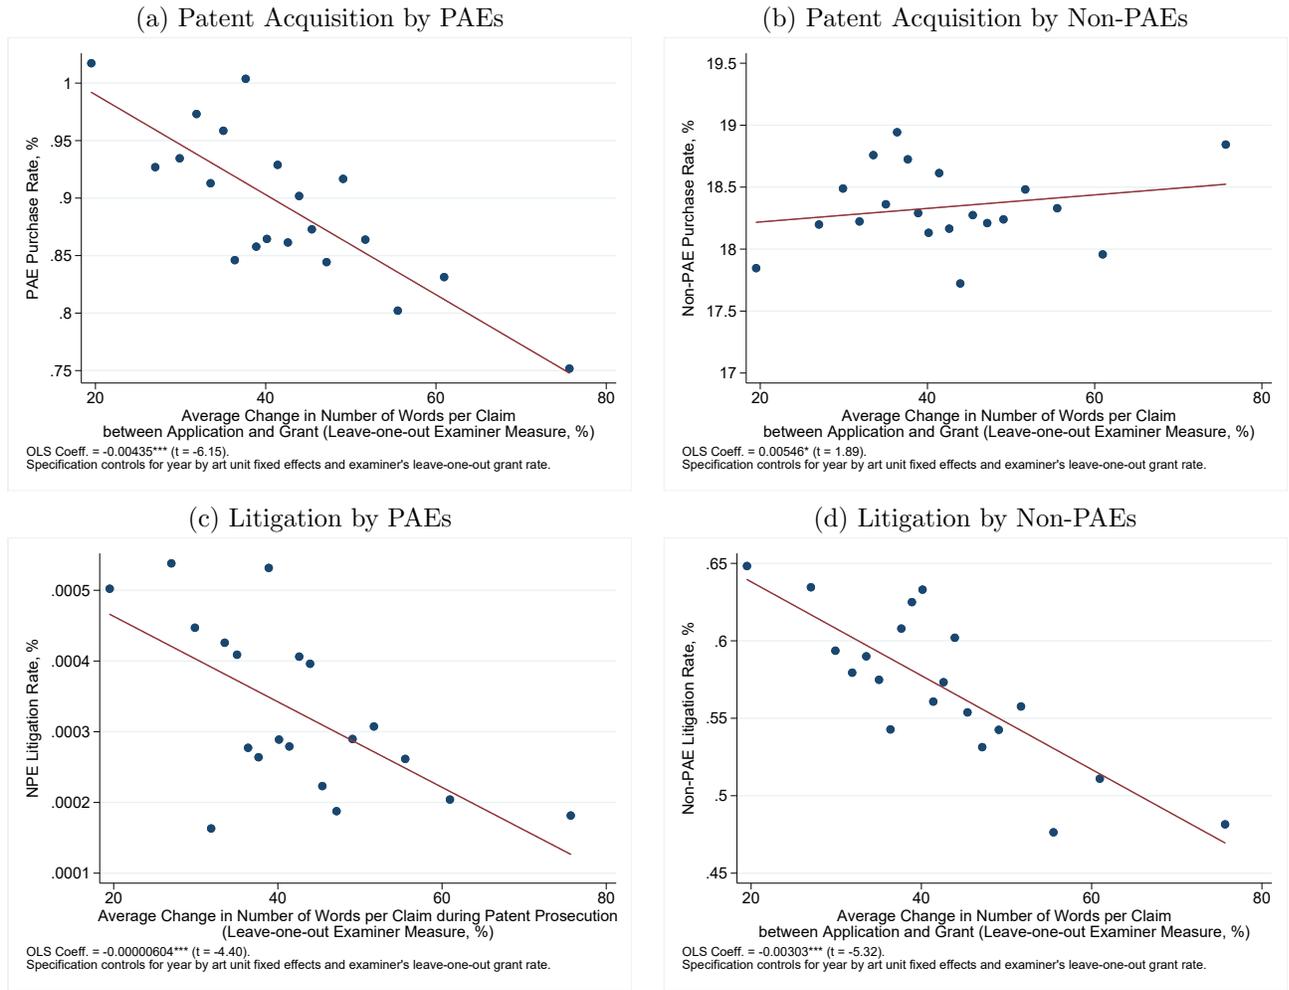
Notes: The sample is restricted to IT patents. Regressors are standardized by their standard deviations and coefficients are expressed as a fraction of the mean of the outcome. The linear predictor for PAE acquisition is given by specification (9) in Table 5. Standard errors are clustered by examiners. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: Heterogeneity in Examiner Causal Effects on Patent Acquisition by PAEs

	Signal SD, % of Average	Average Purchase Rate
(A) Baseline	63.0%	1.02%
(B) Including Assignee Fixed Effects	44.5%	1.02%
(C) Excluding Intellectual Ventures	82.7%	0.55%
(D) Intellectual Ventures Only	82.0%	0.47%
(E) PAE list from Cotropia et al. (2014)	67.8%	0.60%
(F) Small PAEs	40.7%	0.07%
(G) PAEs purchasing from small entities/unassigned patents	70.8%	0.11%

Notes: This table reports the signal standard deviation of examiner effects using different specifications and different PAE outcomes, using the full sample. The Bayesian shrinkage methodology used to obtain these estimates is presented in Section III.A. Row (A) reports the baseline estimate from Table 2. In row (B), the specification includes assignee fixed effects. Row (C) uses purchase by a PAE other than Intellectual Ventures as the outcome. Row (D) considers purchases by Intellectual Ventures as the outcome. Row (E) uses the PAEs list from Cotropia et al. (2014). Row (F) examines purchases by PAEs with a small patent portfolio, as identified by Cotropia et al. (2014). Row (G) considers purchases by PAEs whose portfolios have more than 50% of patents that either were unassigned (i.e., the inventor is the owner) or that were assigned to a firm that the USPTO classifies as a small entity.

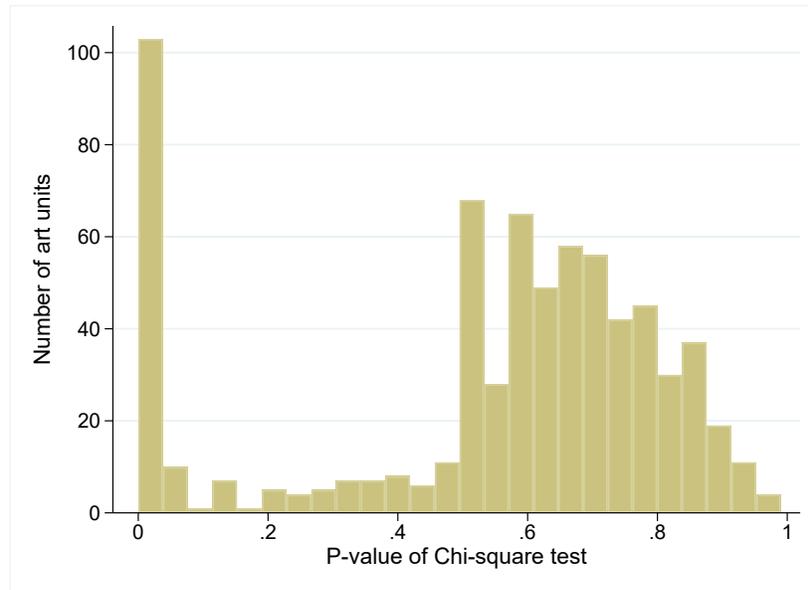
Figure 1: The Effect of Examiners on Patent Acquisition and Litigation



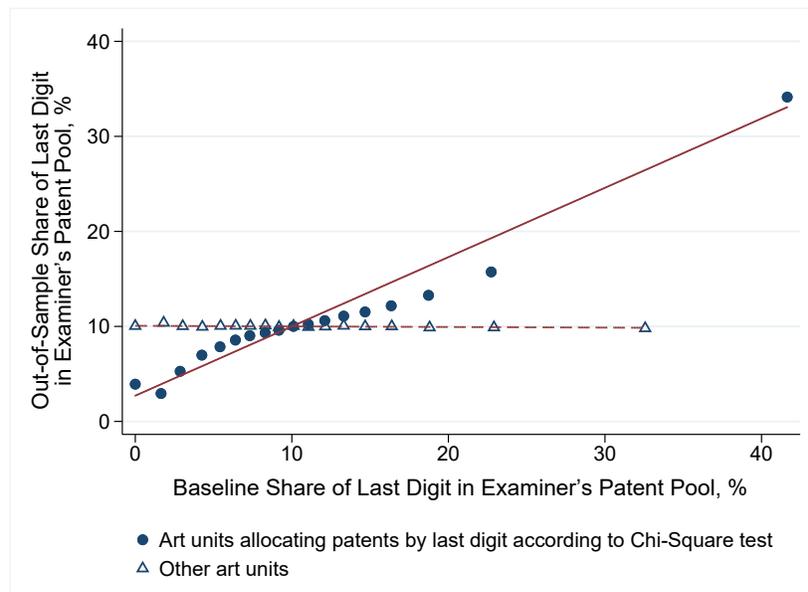
Notes: In the various panels of this figure, the level of observation is a patent. The average change in the number of words per claim is measured at the level of an examiner, leaving out the focal patent. All specifications include art unit by year fixed effects and address potential extensive margin effects by controlling for the examiner's leave-one-out patent grant rate (see text for details). The sample is the full patent grant sample described in Section II.A, excluding examiners in the top 1% of the distribution of the total number of granted patents. The total number of patents granted by the examiner is used as weights in all panels. Each dot represents 5% of the data and the OLS best-fit lines are reported. Standard errors are clustered by examiner.

Figure 2: The Allocation of Patent Applications to Examiners by Application's Last Digit

Panel A: Distribution of p-values of Chi-square tests



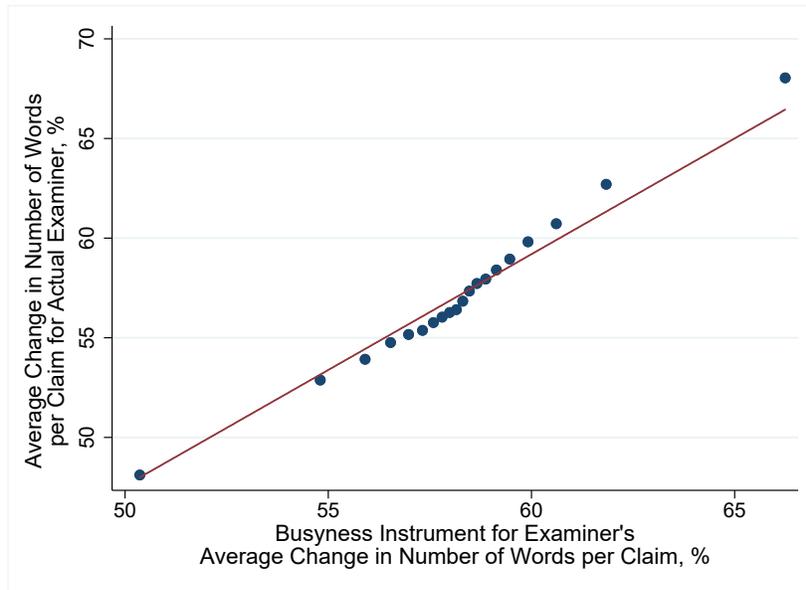
Panel B: Graphical Evidence on Allocation by Application's Last Digit



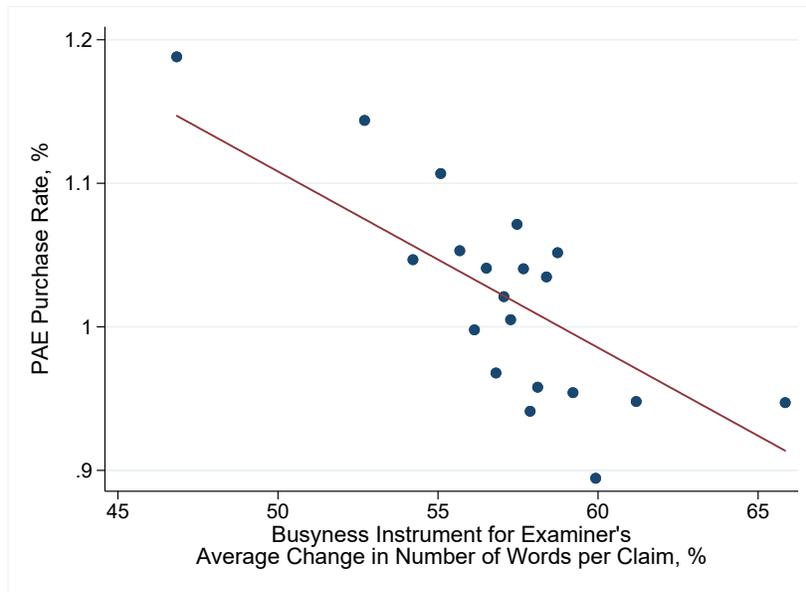
Notes: In Panel A, the level of observation is an art unit. This panel reports the distribution of the p-values of the Chi-square tests described in the main text; a p-value below 0.01 indicates excess concentration of patent applications across examiners by application's last digit. In Panel B, the level of observation is an examiner-by-application's last digit cell. Two binned scatter plots are reported with the corresponding best fit lines; each cell is weighed by the total number of applications processed by the examiner.

Figure 3: A Busyness Instrument for the Effect of Examiners on Patent Acquisition by PAEs

Panel A: First Stage

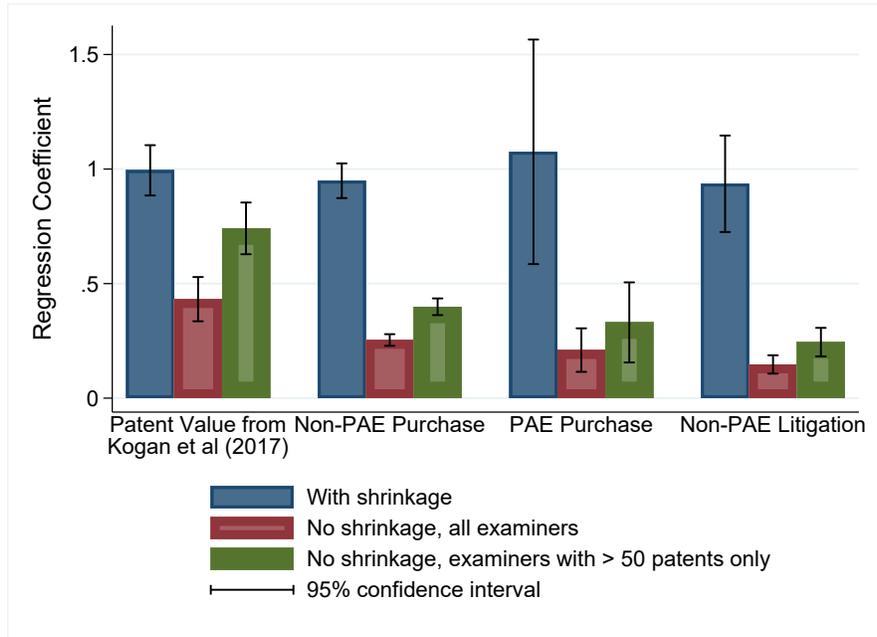


Panel B: Reduced Form



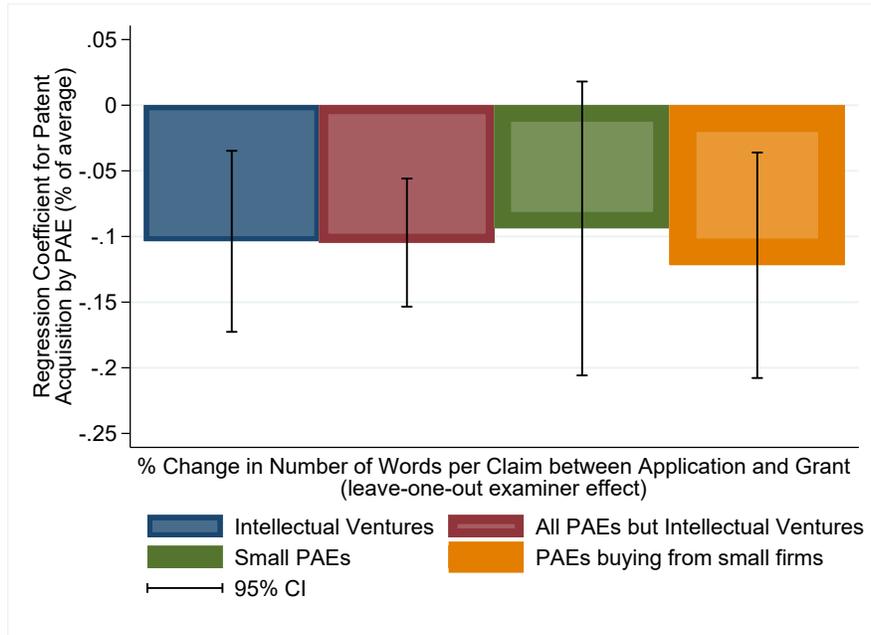
Notes: Panel A shows the relationship between the busyness instrument (described in the main text) for an examiner's propensity to change the number of words per claim during application and grant and the propensity of the examiner whom the application was actually assigned to. Panel B depicts the relationship between the busyness instrument and the purchase rate by PAEs. On both panels, each dot represents 5% of the data and OLS best-fit lines are reported.

Figure 4: Out-of-Sample Tests of Baseline Estimates of Examiner Effects



Notes: This figure reports the OLS coefficients in examiner-level out-of-sample regressions. After splitting the main analysis sample into two halves at random, we compute the raw and shrunk examiner effects on each half following the methodology described in Section III.C. To test predictive accuracy, we regress the raw examiner effect from the first half on examiner effects estimated in the second half, using in turn as regressors the shrunk examiner effects (“shrinkage”), the raw examiner effects (“no shrinkage, all examiners”) and the raw examiner effects for the subset of examiners who have granted more than fifty patents (“no shrinkage, examiners with > 50 patents only”). A regression coefficient of one indicates unbiased prediction. The heteroskedasticity-robust 95% confidence interval is reported.

Figure 5: Heterogeneity in Patent Acquisition Behavior across Groups of PAEs



Notes: The sample is restricted to IT patents. The regression coefficients indicate the percentage change in the probability of PAE acquisition (relative to the baseline rate) for a one standard deviation increase in the examiner effect for the change in the number of words per claim during prosecution. The methodology is described in Section IV.B (see specification (8)). Regression coefficients are reported separately for four samples of PAEs: Intellectual Ventures, PAEs other than Intellectual Ventures, PAEs with a small patent portfolio according to the classification of [Cotropia et al. \(2014\)](#), and PAEs which primarily buy patents from small entities (specifically, as described in the main text they purchase more than half of their patents from small firms or individual inventors). The 95% confidence intervals are based on standard errors clustered by examiners.

Online Appendices

A Examiner Effect Estimation Methods

In this appendix, we report additional steps involved in the research designs described in Section III to recover the causal effect of examiner on patent outcomes, namely (1) the baseline research design with the Bayesian shrinkage correction, (2) the design using applications' last digits as a source of variation, (3) the design using examiners' busyness as a source of variation, and (4) the baseline research design with the Beta-Binomial count model for binary outcomes.

1. *Bayesian shrinkage.* In what follows, we describe the two steps we take to estimate the shrunk examiner effects introduced in Section III.A. These two steps help increase the precision of the Empirical Bayes posterior estimate of each examiner effect in equation (4).

In the first step, we amend the statistical model to allow for an examiner-by-year shock θ_{jt} , i.e.

$$\begin{aligned} Y_i &= a_{ut(i)} + v_{ij}, \\ v_{ij} &= \mu_j + \theta_{jt(i)} + \epsilon_i, \end{aligned}$$

where i indexes the patent, j the examiner, u the art unit and t the year. We compute \bar{v}_{jt} using (2), $\widehat{\sigma}_\mu$ using (3), as well as $\widehat{\sigma}_\epsilon^2 = Var(v_{ij} - \bar{v}_{jt})$ and $\widehat{\sigma}_\theta^2 = Var(v_{ij}) - \widehat{\sigma}_\mu^2 - \widehat{\sigma}_\epsilon^2$.

In the second step, for each examiner we compute a weighted average of the yearly average residuals $\{\bar{v}_{jt}\}$ that has the property of being a minimum variance unbiased estimate of μ_j . This average uses weights such that years in which the examiner granted more patents are given a higher weight:

$$\bar{v}_j = \sum_t w_{jt} \bar{v}_{jt},$$

where

$$\begin{aligned} w_{jt} &= \frac{h_{jt}}{\sum h_{jt}} \\ h_{jt} &= \frac{1}{\widehat{\sigma}_\theta^2 + \frac{\widehat{\sigma}_\epsilon^2}{n_{jt}}} \end{aligned}$$

We then compute the Empirical Bayes posterior estimate of each examiner as in (4):

$$\widehat{\mu}_j = \frac{\widehat{\sigma}_\mu^2}{Var(\bar{v}_j)} \cdot \bar{v}_j,$$

with

$$\text{Var}(\bar{v}_j) = \widehat{\sigma}_\mu^2 + \left(\sum h_{jt} \right)^{-1}.$$

The shrinkage factor is the ratio of the signal variance to the total variance and varies across examiners depending on the total number of patents they granted.

The shrunk examiner effects $\{\widehat{\mu}_j\}$ have two noteworthy properties. First, under the assumption that $\mu_j \sim N(0, \sigma_\mu^2)$, $\theta_{jt} \sim N(0, \sigma_\theta^2)$ and $\epsilon_{it} \sim N(0, \sigma_\epsilon^2)$, the shrunk examiner effect is the optimal Bayesian posterior expectation of an examiner’s effect given the history of patent outcomes up to the current period. The derivation is a standard application of Bayes’ rule. Intuitively, since there is no drift in examiner effects, we can use the average patent outcome in all years prior to the current year as a sufficient statistic to form the posterior distribution of examiner effects. Second, the shrunk examiner effects also have a frequentist interpretation. The shrinkage factor is the regression coefficient in the hypothetical regression of the true (unobserved) μ_j on \bar{v}_j . The regression coefficient is naturally the ratio of the covariance of μ_j and v_j (given by $\widehat{\sigma}_\mu^2$ because the other components of \bar{v}_j are noise) to the variance of \bar{v}_j .

2. Allocation of applications to examiners using the last digit of applications’ serial numbers.

To identify art units assigning applications based on the last digit of their serial numbers, we use the USPTO patent examination database and follow three steps.

First, we prepare the data. We exclude continuation applications, since these applications are almost always assigned to the examiner who processed the parent application. For each patent application we record the “docket date”, which is the date on which the application was assigned to the relevant art unit. After an application is filed, the USPTO assigns the application to a specific art unit according to its technological features, which takes some time; therefore the docket date is typically different from the filing date and is the relevant point in time when examiner assignment occurs within the art unit.

Second, we compute the key statistics of interest. We count the number of applications falling in each last-digit-by-examiner-by-art-unit-by-docket-year cell; we denote the application count in each cell by n_{djut} , where d indexes the last digit of the application’s serial number (ranging from 0 to 9), j the examiner, u the art unit and t the docket year. In addition, for each examiner we record the total number of applications they were assigned in each art unit in each docket year, denoted by n_{jut} .

Third, for each art unit in each year, we implement a statistical test of the null hypothesis

that examiner assignment does not depend on last digits. If last digits are not used, the expected number of application with last digit d assigned to examiner j is simply one tenth of the total number of application assigned to this examiner in this art unit and docket year, which we denote by $n_{djut}^E = \frac{n_{jut}}{10}$. To assess whether the data rejects the null in a given art unit and docket year, we compute the following Pearson’s Chi-squared statistic:

$$\chi_{ut}^2 \equiv \sum_{j \in ut} \sum_{d=0}^9 \frac{(n_{djut} - n_{djut}^E)^2}{n_{djut}^E}.$$

Intuitively, we compare the *actual* number of applications with last digit d assigned to examiner j in art unit u in docket year t (n_{djut}) to the *expected* number (n_{djut}^E). If the actual assignment patterns are “too concentrated” relative to what may happen simply by chance, we reject the null. Formally, under the null that art unit u did not use last digits for examiner assignment in docket year t , χ_{ut}^2 has a Chi-squared distribution with $9 \cdot (J_{ut} - 1)$ degrees of freedom, where J_{ut} is the number of examiners in art unit u in docket year t . The degrees of freedom follow from the fact that there are ten possible last digits per examiner, minus the constraints for the total number of applications within each examiner and for the total number of applications by last digit cells. Accordingly, we compute the p-value for the null by comparing the value we obtain in the data for χ_{ut}^2 with a Chi-squared distribution with $9 \cdot (J_{ut} - 1)$ degrees of freedom.⁵²

Fourth, we draw the list of art units for which we reject the null that last digits are not used for assignment at the 1% level, i.e. with a p-value of the Chi-squared test below 0.01. We draw this list by docket years, i.e. based on statistics $\{\chi_{ut}^2\}$ that are specific to both art units and docket years, so that each art unit is allowed to change its assignment mechanism over time. To make it simple for other researchers to use this source of variation in future work, we make publicly available on our websites the list of art units by docket years for which we rejected the null at the 1% level (click [here](#) for the list).

Busyness instrument. We describe below our methodology to recover the application-specific examiner assignment probabilities p_{ij} used in equations (5) and (6) in Section III.C. Our approach delivers variation in examiner assignment solely from changes in examiner busyness over time, which is useful to validate the baseline estimates of examiner effects.

⁵²Our Chi-squared test is similar in spirit to the divergence statistics used by [Righi and Simcoe \(2017\)](#) to provide evidence of examiner specialization based on the dispersion of patent technology classes across examiners working in the same art unit. The exact formulas for their divergence tests differ from ours because they allow for technology-specific patterns of specialization within an art unit (i.e., within a given art unit, it could be that only a subset of all technology classes feature examiner specialization); we use a similar but technically different test for assignment by last digits, because it seems implausible that only a subset of last digits would be used for examiner assignment.

We start by preparing the data. We aggregate total disposals (grants plus abandonments) for examiners in each two-week period in a given year. As before, we exclude continuation applications because they tend to be assigned to the examiner who handled the parent application. For each incoming application, we create the list of all examiners that it could have been assigned to, which is given by the set of examiners who processed at least one application in that art unit and in that year. As a proxy for how an examiner’s busyness changes over time, we compute the number of patent application cases closed by the examiner in each two-week period. Intuitively, an examiner may be assigned more applications as they become less busy, i.e. in periods when they just finished working on other applications.

Next, we estimate the following linear probability model by OLS:

$$Y_{ij} = \beta D_{jt} + \delta_i + \gamma_j + \epsilon_{ijt}, \tag{A1}$$

where i indexes the application, j the examiner and t the two-week period. Y_{ij} is an indicator variable for the assignment of application i to examiner j ; D_{jt} is the number of patent application cases closed by the examiner during the relevant two-week period; δ_i is an application fixed effect which captures the fact that a larger or smaller number of examiners may be available when a given application arrives; and γ_j is an examiner fixed effect which accounts for the possibility that some examiners might be systematically assigned a large or smaller number of applications (e.g, due to seniority). The coefficient β estimates the extent to which an examiner is more likely to be assigned an application (relative to the baseline captured by the fixed effects) in a period when they just finished working on other applications.

Finally, we use the estimates from (A1) to compute the predicted assignment probabilities $p_{ij} \equiv \widehat{Y}_{ij}$, which are used in equations (5) and (6) in Section III.C. If the estimate of β were zero, there would be no variation in the application-specific examiner assignment probabilities $\{p_{ij}\}$ across applications received in the same art unit in the same year and the research design would have no power. In fact, we estimate $\beta > 0$ and obtain sufficient variation in examiner assignment probabilities to implement the busyness research design presented in Section III.C.

Beta-Binomial count model. In what follows, we derive the integrated likelihood for the Beta-Binomial count model used in Section III.C. As a reminder, we aggregate data for each examiner j in year t and art unit u into the form (n_{jau}, r_{jut}) , where n denotes the total number of granted patents for a given examiner and r the total number of patents purchased by PAEs (or some other binary outcome) for this examiner. We then model the data generating process with a binomial

likelihood on each examiner for each art unit and year: each examiner has some true probability p_{jut} of patent purchase by a PAE (conditional on grant). We drop the ut subscripts below for brevity.

For each art unit in each year, we specify a flexible Beta prior distribution on examiner effects: $p \sim \text{Beta}(\alpha, \beta)$. We then compute the integrated likelihood:

$$\begin{aligned}
L(r|n, \alpha, \beta) &= \int_{p=0}^1 \binom{n}{r} p^r (1-p)^{n-r} \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} p^{\alpha-1} (1-p)^{\beta-1} dp \\
&= \binom{n}{r} \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \int_{p=0}^1 p^r (1-p)^{n-r} p^{\alpha-1} (1-p)^{\beta-1} dp \\
&= \binom{n}{r} \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \frac{\Gamma(r+\alpha)\Gamma(n-r+\beta)}{\Gamma(n+\alpha+\beta)},
\end{aligned}$$

where the second step conjugates the inside to integrate to one based on the probability density function of the Beta distribution. Using this expression, we estimate the hyperparameters α and β by maximum likelihood.

Finally, we compute the shrunk examiner effect for each examiner. A shrunk examiner effect is simply a posterior mean: we start from the prior, which is governed by the estimates for α and β for each art unit and year, and apply Bayesian updating with the examiner's data in that art unit and year according to equation (7). We then aggregate these shrunk effects across years within each examiner by taking a weighted average (weighting by number of cases) to obtain the overall shrunk effect for each examiner.

B Building the Patent Portfolios of Patent Assertion Entities

In this appendix, we describe the procedure to we use to build the patent portfolio of Patent Assertion Entities. We proceed in four steps:

1. We start with the list of PAE names from RPX for our main sample and from [Cotropia et al. \(2014\)](#) for robustness checks. We exclude universities (e.g. Wisconsin Alumni Research Foundation) and academic hospitals (Children’s Medical Center Corporation). For the [Cotropia et al. \(2014\)](#) list, we only include entities in categories 3 (Large aggregator) and 5 (Patent holding company). This excludes failed companies and technology development compnaies.
2. We normalize entity names from both the PAE list and the USPTO Assignment Database from [Marco et al. \(2015\)](#). We do so by capitalizing all names, removing punctuation, and removing the following standard entity terms: INC, CO, COMPANY, COMPANIES, CORP, CORPORATIONS, DIV, GMBH, LLC, LC, INCORPORATED, KG, LIMITED, LIMITED PARTNERSHIP, LP, LTD, NV, PLC, SA, SARL, SNC, SPA, SRL, TRUST USA, CENTER, BV, AG, AB, GROUP, FOUNDATION, INSTITUTE, and TECHNOLOGIES.
3. We collect the identifiers of patent transactions in the USPTO Assignment Database (“Reel/Frame IDs” in the USPTO assignment data, which corresponds to one transaction) that have a normalized entity name matching the normalized name of a PAE in Step 2.
4. Using the patent transaction identifier from Step 3, we know from the USPTO Assignment Database whether the patent was assigned to the employer of the inventor(s). We only keep transactions that are non-employer assignments, to mitigate any PAE classification errors that might cause us to include patents filed by failed companies and technology development companies. We exclude transactions such as securitization, mergers, and name changes. We end up with a list of patents that were sold to PAEs on the patent market.

C Online Appendix Tables and Figures

Table A1: Raw Standard Deviations of Patent Outcomes across Examiners

	Raw S.D.	
	% of Average	Level
	(1)	(2)
Patent value from Kogan et al. (2017), \$M	106.14	8.65
4th-year fee payment rate	11.44	0.0996
8th-year fee payment rate	18.05	0.1101
12th-year fee payment rate	34.57	0.0722
Log total patent citation	42.71	0.20
Log patent citations by same assignee	92.37	0.19
Log patent citations by other assignees	46.43	0.09
Rate of patent acquisition by non-PAEs	68.41	0.1344
Rate of patent acquisition by PAEs	286.88	0.0292
Rate of patent litigation by non-PAEs	439.17	0.0285
Rate of patent litigation by PAEs	1359.17	0.0055

Notes: This table reports the raw standard deviations of examiner effects as a percentage of the mean (Column 1) and in level (Column 2). The raw standard deviations refer to the standard deviations of the average residuals (defined by equation (2)) across examiners. The raw standard deviations account for art unit by year fixed effects but not for excess variance from noise. The results in this table are directly comparable to those of Table 2, which account for excess variance. See Section II.A for details on the sample and variable definitions.

Table A2: Signal Standard Deviations of Examiner Causal Effects by Technology Categories

	Signal S.D., % of Average			
	Patent value from Kogan et al. (2017)	Log total patent citation	Non-PAE purchase	Change in number of words per claim
	(1)	(2)	(2)	(3)
(A) Biotechnology and organic chemistry	25.33	30.58	15.39	18.55
(B) Chemical and materials engineering	74.89	24.85	23.11	19.30
(C) Computer architecture, software and information security	9.00	24.69	6.43	25.60
(D) Computer networks, multiplex communication, video distribution, and security	11.80	31.80	6.10	18.06
(E) Communications	24.39	20.78	11.68	30.97
(F) Semiconductors, electrical and optical systems and components	30.83	17.59	12.08	26.47
(G) Transportation, construction, electronic commerce, agriculture, national security, and license and review	21.95	21.42	12.77	19.15
(H) Mechanical engineering, manufacturing	29.78	23.08	19.41	19.73

Notes: This table reports the signal standard deviations of examiner effects (as a percentage of the mean) for four patent outcomes across the eight technology centers of the USPTO. The means are re-computed within each technology center. The table shows that examiner effects are substantial in all technology centers. We have studied heterogeneity in signal standard deviations for the other outcomes reported in Table 2 and did not find large differences across technology centers, except for patent acquisitions by PAEs, which occur primarily in computers, software and communications (not reported). The Bayesian shrinkage methodology used to obtain these estimates is presented in Section III.

Table A3: Signal Standard Deviations of Examiner Prosecution Behaviors

	Signal S.D.		S.D. of Shrunk
	% of Average (1)	Level (2)	Effects, % of Average (3)
Change in number of words per claim, % (average over all claims)	23.37	13.39	17.47
Change in number of claims, %	136.83	4.99	82.43
Use of Section 101 - Lack of utility or eligibility	60.43	0.032	52.44
Use of Section 102(a) - Prior art exists	108.69	0.018	75.65
Use of Section 103(a) - Obvious invention	25.27	0.105	19.61
Use of Section 112(b) - Vague claims	47.72	0.088	39.09
Patent citations added by examiner, %	14.53	7.95	11.52
Citations to non-patent literature added by examiner, %	39.70	5.73	32.09

Notes: This table reports the signal standard deviations of examiner effects as a percentage of the mean (Column 1) and in level (Column 2), as well as the standard deviations of shrunk examiner effects (Column 3). The Bayesian shrinkage methodology used to obtain these estimates is presented in Section III. See Section II.A for details on the sample and variable definitions.

Table A4: PAE Patent Portfolios and Litigated Patents across Technology Categories

Panel A: Patents Owned by PAEs

Technology Category	Number of Patents
Chemical	1,669
Computers & Communications	27,156
Drugs & Medical	1,312
Electrical & Electronic	10,660
Mechanical	2,709
Others	1,453

Panel B: Patents Litigated by Non-PAEs

Technology Category	Number of Patents
Chemical	1,626
Computers & Communications	4,175
Drugs & Medical	2,609
Electrical & Electronic	2,497
Mechanical	2,859
Others	4,611

Notes: This table reports the number of patents owned by PAEs (Panel A) and the number of litigated patents by non-PAEs (Panel B) across technology categories. The technology categories are based on the primary USPTO technology class for each patent, following [Hall et al. \(2001\)](#). These panels show that PAEs tend to be most active in technology areas related to computers, communications and electronics, where patent litigation by non-PAEs is also frequent.

Table A5: Patent Acquisition and Examiner Behavior, Full Sample

Panel A: Patent Acquisition by PAEs

Leave-one-out Examiner Effects	Patent Purchase by PAE								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
% Change in Number of Word per Claim from Application to Grant	-0.076*** (0.0177)								-0.062** (0.026)
% Change in Number of Claims from Application to Grant		0.064*** (0.021)							0.057** (0.023)
Grant Rate			0.064*** (0.015)						-0.01 (0.02)
Use of Section 101 - Lack of utility or eligibility				-0.059* (0.033)				-0.057* (0.0342)	-0.052 (0.036)
Use of Section 102(a) - - Prior art exists					0.027 (0.022)			0.0343 (0.0214)	0.033 (0.022)
Use of Section 103(a) - Obvious invention						-0.033** (0.016)		-0.031** (0.015)	-0.015 (0.019)
Use of Section 112(b) - Vague claims							-0.020 (0.019)	-0.0025 (0.018)	0.01 (0.02)
Fixed Effects	Year by Art Unit								
<i>N</i>	1,109,882	1,110,522	1,270,027	1,269,751	1,269,751	1,269,751	1,269,751	1,269,751	1,109,882

Panel B: Patent Acquisition by Practicing Firms

Leave-one-out Examiner Effects	Patent Purchase by Practicing Firm								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
% Change in Number of Word per Claim from Application to Grant	0.0022 (0.005)								-0.003 (0.006)
% Change in Number of Claims from Application to Grant		0.0005 (0.005)							-0.0005 (0.0058)
Grant Rate			0.001 (0.005)						0.0014 (0.008)
Use of Section 101 - Lack of utility or eligibility				0.0136*** (0.0052)				0.012** (0.005)	0.013** (0.005)
Use of Section 102(a) - -Prior art exists					0.006 (0.004)			0.005 (0.004)	0.005 (0.004)
Use of Section 103(a) -Obvious invention						0.0045 (0.004)		0.003 (0.004)	0.004 (0.005)
Use of Section 112(b) -Vague claims							0.004 (0.004)	0.0002 (0.004)	0.001 (0.004)
Fixed Effects	Year by Art Unit								
<i>N</i>	1,109,882	1,110,522	1,270,027	1,269,751	1,269,751	1,269,751	1,269,751	1,269,751	1,109,882

Notes: Regressors are standardized by their standard deviations and coefficients are expressed as a fraction of the mean of the outcome. Results are similar with patent-level controls and assignee fixed effects (not reported). Standard errors are clustered by examiners. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A6: Patent Litigation and Examiner Behavior, Full Sample

Panel A: Patent Litigation by PAEs

Leave-one-out Examiner Effects	Patent Litigation by PAE								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
% Change in Number of Word per Claim from Application to Grant	-0.27*** (0.06)								-0.17** (0.08)
% Change in Number of Claims from Application to Grant		0.16*** (0.04)							0.11** (0.05)
Grant Rate			0.28*** (0.06)						0.09 (0.09)
Use of Section 101 - Lack of utility or eligibility				-0.17 (0.08)				-0.15 (0.091)	-0.055 (0.095)
Use of Section 102(a) - - Prior art exists					-0.014 (0.055)			0.0051 (0.055)	0.022 (0.055)
Use of Section 103(a) - Obvious invention						-0.077 (0.051)		-0.054 (0.053)	0.005 (0.062)
Use of Section 112(b) - Vague claims							-0.077 (0.048)	-0.032 (0.049)	-0.0203 (0.049)
<i>N</i>	1,109,882	1,110,522	1,270,027	1,269,751	1,269,751	1,269,751	1,269,751	1,269,751	1,109,882
Fixed Effects	Year by Art Unit								

Panel B: Patent Litigation by Practicing Firms

Leave-one-out Examiner Effects	Patent Litigation by Practicing Firm								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
% Change in Number of Word per Claim from Application to Grant	-0.066*** (0.022)								-0.07** (0.035)
% Change in Number of Claims from Application to Grant		0.029* (0.017)							0.016 (0.017)
Grant Rate			0.073** (0.028)						0.002 (0.038)
Use of Section 101 - Lack of utility or eligibility				-0.069*** (0.022)				-0.068*** (0.023)	-0.050** (0.021)
Use of Section 102(a) - - Prior art exists					-0.0016 (0.0158)			0.004 (0.016)	-0.0009 (0.016)
Use of Section 103(a) - Obvious invention						-0.021 (0.016)		-0.018 (0.017)	0.022 (0.018)
Use of Section 112(b) - Vague claims							-0.012 (0.017)	0.005 (0.019)	0.004 (0.019)
<i>N</i>	1,109,882	1,110,522	1,270,027	1,269,751	1,269,751	1,269,751	1,269,751	1,269,751	1,109,882
Fixed Effects	Year by Art Unit								

Notes: Regressors are standardized by their standard deviations and coefficients are expressed as a fraction of the mean of the outcome. Results are similar with patent-level controls and assignee fixed effects (not reported). Standard errors are clustered by examiners. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A7: Examiner Behavior and Likelihood of Patent Invalidation, Full Sample

Panel A: Reissuance of Granted Patents

Leave-one-out Examiner Effects (separate regressions)	Reissuance Rate			Reissuance Rate Two Years or More after Grant		
	(1)	(2)	(3)	(4)	(5)	(6)
(A) % Change in Number of Word per Claim from Application to Grant	-0.12*** (0.031)	-0.10*** (0.029)	-0.12*** (0.03)	-0.19*** (0.069)	-0.20*** (0.066)	-0.20*** (0.069)
(B) Grant Rate	0.13*** (0.03)	0.11*** (0.02)	0.16*** (0.031)	0.26*** (0.067)	0.25*** (0.063)	0.29*** (0.069)
Fixed Effects	Year	Year by Art Unit	Year by Art Unit by Class	Year	Year by Art Unit	Year by Art Unit by Class
N		1,109,882			1,107,565	

Panel B: Court Rulings

Leave-one-out Examiner Effects (separate regressions)	Invalidity Rate			Infringement Rate		
	(1)	(2)	(3)	(4)	(5)	(6)
(A) % Change in Number of Word per Claim from Application to Grant	0.019 (0.028)	0.11 (0.11)	0.17 (0.179)	0.046 (0.032)	0.117 (0.119)	0.082 (0.163)
(B) Grant Rate	0.005 (0.02)	-0.025 (0.09)	-0.097 (0.16)	-0.010 (0.027)	-0.027 (0.102)	-0.024 (0.165)
Fixed Effects	Year	Year by Art Unit	Year by Art Unit by Class	Year	Year by Art Unit	Year by Art Unit by Class
N		479			479	

Panel C: Trials at the Patent Office (“Inter Partes Reviews”)

Leave-one-out Examiner Effects (separate regressions)	IPR Rate			Institution Rate of IPR		
	(1)	(2)	(3)	(4)	(5)	(6)
(A) % Change in Number of Word per Claim from Application to Grant	-0.288*** (0.067)	-0.283*** (0.063)	-0.27*** (0.06)	-0.042 (0.032)	-0.027 (0.158)	-0.08 (0.209)
(B) Grant Rate	0.286*** (0.061)	0.2613*** (0.058)	0.28*** (0.06)	0.046 (0.0309)	-0.045 (0.123)	-0.109 (0.151)
Fixed Effects	Year	Year by Art Unit	Year by Art Unit by Class	Year	Year by Art Unit	Year by Art Unit by Class
N		1,109,882			523	

Notes: Regressors are standardized by their standard deviations and coefficients are expressed as a fraction of the mean of the outcome. The linear predictor for PAE acquisition is given by specification (9) in Table A6. Results are similar with patent-level controls and assignee fixed effects (not reported). Standard errors are clustered by examiners. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A8: Examiner Behavior and Other Patent Outcomes
 Panel A: Patent Value from Kogan et al. (2017)

Leave-one-out Examiner Effects	Patent Value									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	
Purchase by PAE	0.012 (0.016)									
% Change in Number of Word per Claim from Application to Grant		-0.0033 (0.008)							0.019 (0.017)	
% Change in Number of Claims from Application to Grant			-0.0026 (0.0226)						-0.005 (0.023)	
Grant Rate				0.0398*** (0.0149)					0.054** (0.024)	
Use of Section 101 - Lack of utility or eligibility					0.0308** (0.0150)				0.035** (0.015)	
Use of Section 102(a) - - Prior art exists						-0.0071 (0.013)			-0.005 (0.013)	
Use of Section 103(a) - Obvious invention							-0.011 (0.007)		-0.002 (0.013)	
Use of Section 112(b) - Vague claims								-0.008 (0.011)	-0.004 (0.013)	
Fixed Effects				Year by Art Unit						
<i>N</i>	356,250	310,264	310,332	356,318	356,250	356,250	356,250	356,250	310,332	

Panel B: Patent Citations

Leave-one-out Examiner Effects	Log Total Patent Citations									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	
Purchase by PAE	0.015** (0.006)									
% Change in Number of Word per Claim from Application to Grant		-0.065*** 0.006							0.002 (0.008)	
% Change in Number of Claims from Application to Grant			0.010 (0.010)						-0.001 (0.01)	
Grant Rate				0.111*** (0.007)					0.091*** (0.013)	
Use of Section 101 - Lack of utility or eligibility					-0.03*** (0.004)				-0.0128** (0.0049)	
Use of Section 102(a) - - Prior art exists						0.001 (0.006)			0.0099* (0.0057)	
Use of Section 103(a) - Obvious invention							-0.054*** (0.005)		-0.023*** (0.006)	
Use of Section 112(b) - Vague claims								-0.027*** (0.006)	0.0041 (0.006)	
Fixed Effects				Year by Art Unit						
<i>N</i>	988,249	848,162	848,527	988,432	988,249	988,249	988,249	988,249	848,162	

Notes: Regressors are standardized by their standard deviations and regression coefficients are expressed as a fraction of the mean of the outcome. The sample includes all technology categories. Standard errors are clustered by examiners. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

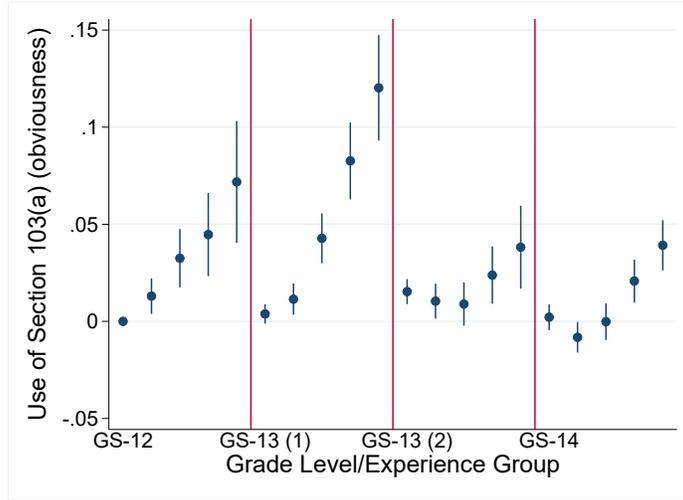
Table A9: Patent Acquisition by PAEs and Grant Decisions at the European Patent Office (EPO)
- Subsample Analysis

	Patent Acquisition by PAE		Patent Acquisition by Practicing Firm		Patent Litigation by Practicing Firm	
	(1)	(2)	(3)	(4)	(5)	(6)
EPO Grant	-0.2144** (0.1001)	0.0819 (0.0702)	0.0037 (0.0133)	-0.0052 (0.0126)	-0.0831 (0.1074)	0.0196 (0.0921)
Subsample of examiners with PAE purchase effect <u>above median</u> only	Yes		Yes		Yes	
Subsample of examiners with PAE purchase effect <u>below median</u> only		Yes		Yes		Yes
Art unit by year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Examiner Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Assignee Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
<i>N</i>	109,383	109,484	109,383	109,484	109,383	109,484

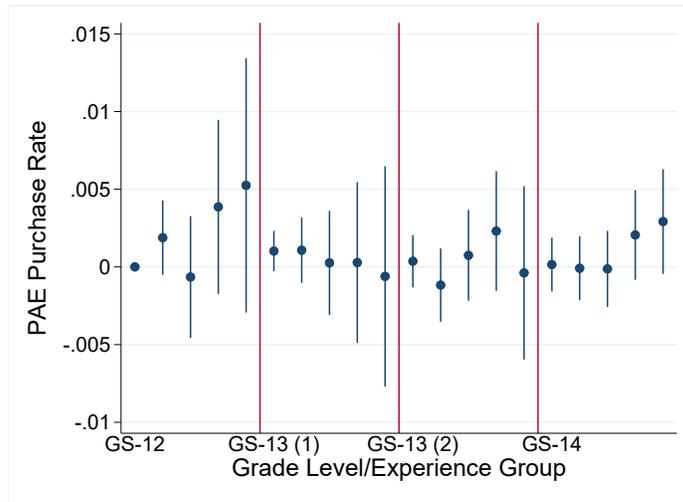
Notes: This table provides a further breakdown for our results in Table ???. Regressors are standardized by their standard deviations and regression coefficients are expressed as a fraction of the mean of the outcome. Columns (1) and (2) show that PAEs selectively purchase patents that were rejected by the EPO *only* in the patent portfolios of examiners who have a relatively large causal impact on PAE purchases (specifically, their PAE purchase effect is above median; the examiner-specific PAE purchase effects were estimated using equation (4)). These patterns suggest that PAEs selectively purchase patents with two features: (i) these patents are close to existing intellectual property because they bear on incremental technologies (hence their higher likelihood of rejection at EPO); (ii) these patents were issued by specific (lenient) examiners at the USPTO, and their claims may be less well-defined are harder to interpret than average. Given these two features, it is plausible that these patents are particularly productive for litigation, as they offer many potential litigation targets. The effect in Column (1) is quantitatively large: the probability of a purchase by a PAE decreases by 21% if the patent is granted by the EPO. As shown in Columns (3) to (6), there is no such effect for patent acquisition by practicing firms (for which we obtain precisely estimated zeros) or for non-PAE litigation. The results all account for art unit by year, examiner and assignee fixed effects; the results are similar when removing the fixed effects or changing the cutoff for the examiner PAE effects (not reported). Standard errors are clustered by examiners. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Figure A1: Examiner Career Effects

Panel A: Prosecution Behavior

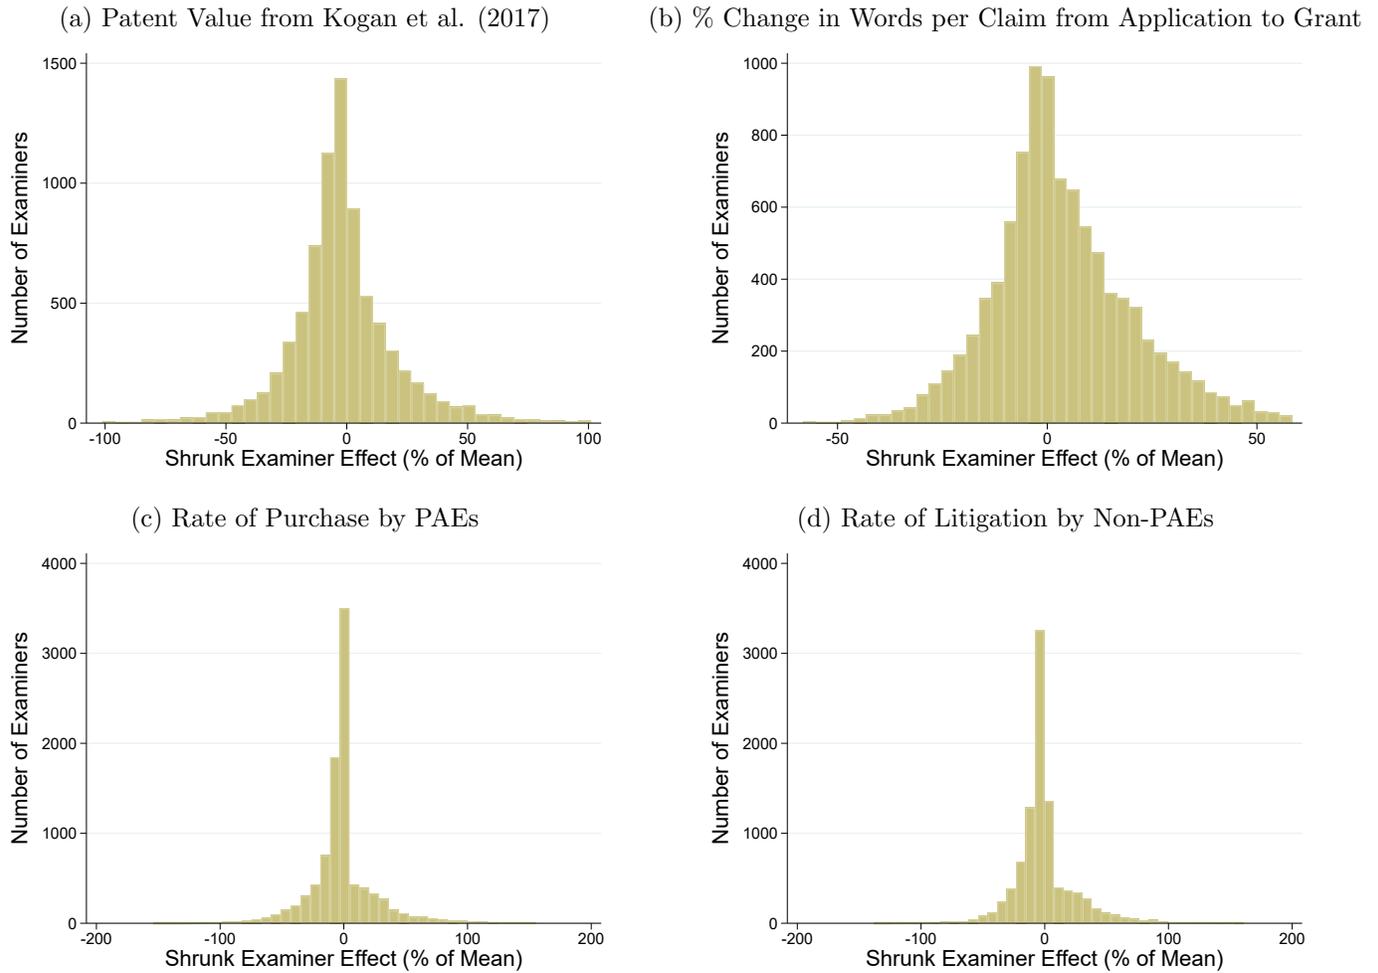


Panel B: Patent Acquisition by PAEs



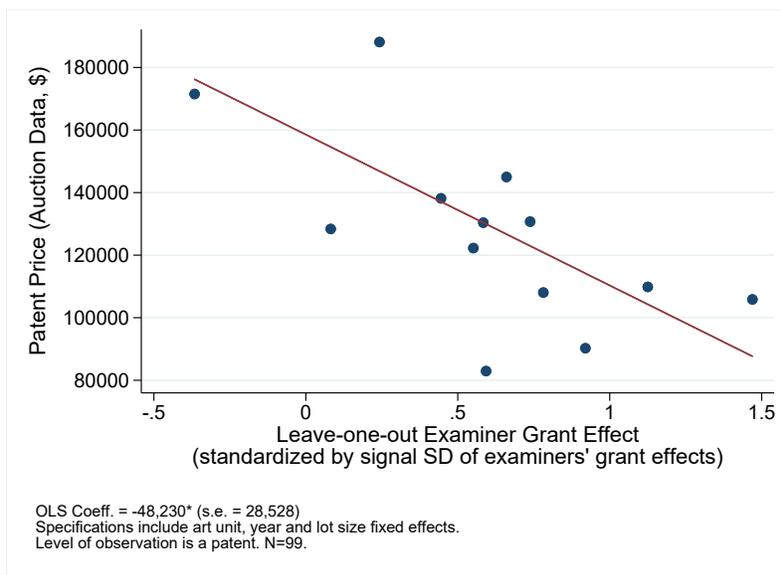
Notes: Following [Frakes and Wasserman \(2017\)](#), this figure examines the behavior of individual examiners over the course of various promotions (indicated by red bars on the figure) that carry with them reductions in examination time allocations. In each panel, we regress the outcome on a series of dummy variables reflecting examiners' experience within a grade level. For each grade level — GS-level 12, GS-level 13 without signatory authority, GS-level 13 with signatory authority, and GS-level 13 —, we track examiners for 1-2, 3-4, 5-6, 7-8, and over 9 years of experience. Specifications include examiner and year fixed effects, and standard errors are clustered by examiners. Panel A shows that after being promoted (and having less time for examination), examiners tend to make fewer demands on the applicant during the prosecution process, as evidenced by the reduction in the issuance of 103(a) blocking action (which is consistent with the findings on grant rates in [Frakes and Wasserman \(2017\)](#)). In contrast, Panel B reports that the rate of purchase by PAEs does not vary significantly around promotion events. This result indicates that examiner career effects have a second-order impact on PAE purchase, relative to the examiner fixed effects estimated in Section III.

Figure A2: Distributions of Shrunk Examiners Effects



Notes: This figure reports histograms of the shrunk examiner effects for four patent outcomes. The shrunk examiner effects are computed according to equation (4). In each panel, the shrunk examiner effects are expressed as a percentage of the mean of the outcome and the histogram is reported for shrunk effects that are within 2.5 signal standard deviations of the mean. This figure shows that there is substantial variation in shrunk examiner effects, i.e. the data delivers very different Bayesian posterior expectations across examiners. The distribution is more concentrated towards zero for rare outcomes like purchase by a PAE or litigation, because the shrinkage factors are higher for such outcomes.

Figure A3: Examiner Prosecution Behavior and Patent Prices



Notes: This figure reports the relationship between the (leave-one-out) examiner grant effect and the patent price in an auction. The examiner grant effects (on the x-axis) are computed as in Section III and are standardized by their signal standard deviation. The auction price (on the y-axis) is provided by Ocean Tomo. Each dot on the figure represents 5% of the underlying data and the OLS best-fit line is reported. Since patents are sometimes auctioned as a group rather than individually, we include fixed effects for lot size. The specification also includes art unit and year fixed effects. The negative slope shows that more lenient examiners (with a higher grant rate) issue patents that sell for less in the patent market. A patent issued by an examiner with a grant rate one signal standard deviation above the mean is sold for \$48,000 less on average.