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

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

December 4, 2017

Please click here for the latest version

Abstract

We show that much of the heterogeneity in patent outcomes, such as patent value, citations and litigation, is caused by the way patent rights are crafted, rather than by heterogeneity in idea quality. We obtain variation in patent rights from the quasi-random allocation of patents to examiners: a one standard deviation change in examiner effects leads stock market capitalization to increase by 3 million dollars, citations by 24%, and litigation by 64%. Patent Assertion Entities, which are very active in litigation and licensing, overwhelmingly purchase and assert patents granted by "lenient" examiners, who craft patents with higher litigation and invalidity risks.

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, 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 across 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 patents citations as a proxy for social value, and Lanjouw and Schankerman (2001) on exposure to litigation). The determinants of this large heterogeneity in patent outcomes are not well understood, although they are a key input into the effect of the patent system on innovation.[1] Scientific 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)), may appear likely to be the primary 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.

In this paper, we use variation in the process of writing the patent description and claims at the United States Patent Office (USPTO) to establish that much of the heterogeneity in patent outcomes is independent of scientific determinants and results from the way patent rights are crafted. We then show that this finding is key to understand the activities of Patent Assertion Entities (PAEs), a central and much-debated feature of the U.S. innovation system. 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.[2] We find that they disproportionately purchase and assert patents from "lenient" patent examiners, who craft patents that are more likely to be litigated and to be invalid.

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, we need variation in patent rights that is orthogonal to other determinants of patent outcomes, such as scientific 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 with the applicant between

[1] For instance, according to theoretical analyses of investment under uncertainty (e.g., Dixit and Pindyck (1994) and Bloom et al. (2007)), if a high share of the variance in the private returns to patenting results from factors outside of the control of the inventor (such as the way patent rights are crafted by examiners), then the responsiveness of innovation to demand shocks will be low.

[2] For 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.

initial patent filing and patent grant (known as the "prosecution" process). By law, examiners must all 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 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), payment of maintenance fees, future patents by the applicant), 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 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 magnitudes of our baseline estimates in three

[3]Conceptually, 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.

[4]As 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.

ways. First, regarding identification, Righi and Simcoe (2017) report strong evidence that examiners working in the same part of the technology space (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 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 two digits of the application, implying that examiner assignment is orthogonal to potential confounds; 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. 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 scientific 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. 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

[5] Examiners 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)).

[6] PAEs, 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).

useful intermediaries in the patent market, fostering greater incentives to innovate by lowering the cost of matching patent holders to patent buyers (e.g., Hagiu 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 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 reduce the range of plausible theories of PAEs.

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, a larger effect on purchases by PAEs, and a much larger effect on litigation by PAEs. These patterns 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 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 other words, our results suggest that their activities are the symptom of the way patent rights are crafted by lenient examiners, which affects 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, over the past decade there were only a few hundred of court rulings on patent validity. 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]

Our results build on and contribute to several literatures. An extensive line of research has examined how the patent system affects innovation, either theoretically (e.g., Nordhaus (1969), Klemperer (1990) and Gilbert and Shapiro (1990)) or empirically (e.g., Sakakibara and Branstetter (2001), Moser (2005), Lerner (2009) and Williams (2013)). Conceptually, this literature primarily investigates how innovation incentives are shaped by what could be called the "macro-determinants" of the patent system, such as laws that establish a patent system or change the set of patentable subject matters. 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). Compared with recent work using variation in grant rates across examiners as an instrument for patent grant (Sampat and Williams (2015), Gaulé (2015)

[7]This 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, and 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.

and Farre-Mensa et al. (2017)), we differ by uncovering the importance of the "intensive margin" of examiner effects (the crafting of patent rights, conditional on patent grant) and by providing ways of addressing the identification concerns raised by Righi and Simcoe (2017). Building on 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 using a fixed effects estimator with a Bayesian shrinkage correction.[8] Finally, we contribute to the growing literature on PAEs (e.g., Golden (2006), McDonough III (2006), Chien (2013), Tucker (2014), Allison et al. (2016) and Haber and Werfel (2016)) by uncovering the importance of examiners for patent acquisition and assertion by PAEs. Our finding that PAEs selectively purchase and assert patents from lenient examiners, which have a higher risk of litigation and invalidity, helps discipline theories of PAEs and implies that policies affecting examiner behavior could have a large and targeted impact on PAEs' activities.

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.

II Data

In this section, we describe our data sources, define the samples and key variables we use in our 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, 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.

[8] Running 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.

The granted patent dataset is obtained from USPTO[9] and extends from 1975 to 2016. We use three proxies for the private returns to patents. First, 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.[10] Second, we use the estimates of firm-level returns to patents from Kogan et al. (2017), who use event studies to estimate the excess stock market return realized on the grant date of patents assigned to publicly-traded firms.[11] Third, as a measure of cumulative innovation by the firm, for each patent we compute the number of patents granted to the same assignee within five years of the focal patent's grant date. We do so using disambiguated assignee names from Balsmeier et al. (2015). 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. 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).[12]

The data covering both granted and ungranted patent applications ranges from 2001 to 2015 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 recover 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, examiner needs to ground their demands in specific sections of patent law, which we describe in Section II.B.[13] 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

[9]This data is through the Reed Tech USPTO page: `http://patents.reedtech.com/patent-products.php`.

[10]For entities that do not benefit from reduced rates, the fees are $1,600 due at 4 years, $3,600 due at 8 years and $7,400 due at 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.

[11]Their sample extends up to 2010.

[12]Records 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.

[13]When 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.

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.[14]

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.[15]

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.[16]

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 risk, and exclude from the list any individual inventor, university or failed company.[17] 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 our 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.[18]

[14]The 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.

[15]In particular, in the full sample of granted patents we cannot control for the grant rate of the examiner, which is important to ensure that the effects we document are driven by differences in the way property rights are crafted and not by potential confounding factors from the extensive margin (e.g., examiners with a low grant rate might only grant patents of high scientific merit).

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

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

[18]Intellectual Ventures holds an estimated 25-30k US patents and released a list of around 20k on their website in November of 2013. Currently available at: `http://patents.intven.com/data/ivpatents.csv`.

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 limited by useful 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: there are only 516 such cases in our main analysis sample, but for those patents information on whether the court deemed the patent invalid or found an infringement is available.

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".[19] We obtain this 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.

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

[19] Patent 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 increasing claim scope is not driving the patterns.

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. The standard deviation of future innovation by the assignee is high, more than twice the mean. 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. 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.

The last part of Panel A reports statistics on outcomes that are commonly viewed as proxies for patent invalidity. Court rulings are observed for only 516 patents, or about 0.0004% of our sample. Conditional on observing a court 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.[20] For both court rulings and IPRs, the invalidity rates appear to be high, but they are observed conditional on a very stringent form of sample selection. Finally, 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.

Panel B of Table 1 documents how patent documents change between application and grant. In most cases, the examiner issues a so-called "rejection" as her first decision on the application

[20] According 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).

(Williams (2017)), which is effectively an invitation for the patent applicant (or her representative, such as a lawyer) 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 increases by 57% on average. 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, examiner 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%).[21] 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 examiners work (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

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

simple measure of examiners' prosecution behavior.

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 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 scientific 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. 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. Our 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:

$$Y_i = a_{ut(i)} + v_{ij} \tag{1}$$
$$v_{ij} = \mu_j + \epsilon_i$$

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 number 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 $\bar{v}_j$.[22] We then construct the empirical Bayes posterior estimate of each examiner effect by

[22]To 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 $\widehat{\sigma_\epsilon}^2 = Var(v_{ij} - \bar{v_{jt}})$ and $\widehat{\sigma_\theta}^2 = Var(v_{ij}) - \widehat{\sigma_\mu}^2 - \widehat{\sigma_\epsilon}^2$. 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 the Online Appendix for a complete discussion.

multiplying $\bar{v}_j$ by a shrinkage factor:

$$\widehat{\mu_j} = \frac{\widehat{\sigma_\mu}^2}{Var(\bar{v_j})} \cdot \bar{v}_j. \tag{4}$$

The shrinkage factor is the ratio of signal variance to total variance.[23] 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 3 shows that the signal standard deviation of examiner effects corresponds to a 3.3 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% of the average patent value for publicly-traded firms. The process of creation of patent rights therefore has a first-order impact on patents' private values. We confirm this result in rows two to five 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 that with the Kogan et al. (2017) estimates; the signal standard deviations are under 10% of the average payment rate. We also find that examiners have a sizable impact on the number of future patents obtained by the assignee, which can be interpreted as a measure of cumulative innovation.

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 findings points to the role of cumulative innovation by the assignee.[24]

We find particularly strong examiner effects for litigation and PAEs' activities. The signal

[23]The Online Appendix discusses the computation of $Var(\bar{v_j})$. 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) $\bar{v}_j$.

[24]This may also reflect strategic or legal considerations, although much of the strategic discussion in the literature focuses on continuation filings, while we are focusing on non-continuation patents.

standard deviation of examiner effects account for over 60% of the baseline rate of patent purchase 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. In the Online Appendix, we show that 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 eight times larger than the signal standard deviation.

We use a bootstrapping procedure for inference. We draw sample from the application-level dataset with replacement[25] and repeat the estimation of the signal standard deviations. The 95% confidence intervals are reported in Column 2 of Table 3. 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 3 reports these results. For instance, the standard deviation of shrunk examiners effects accounts for 29.4% of the average patent value from Kogan et al. (2017), 31.11% of the baseline rate of PAE patent purchases, and 27.4% 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. 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.

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

[25] We also try re-sampling within examiner or within examiner by filing year, and find similar results.

outcomes, then the examiner effects document in Table 1 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 digits of applications' serial numbers are determined by the order of submission of applications and are therefore orthogonal to potential confounding variables such as scientific factors.[26] 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.[27] Applied to our purposes, the test delivers a Chi-square statistic asking whether applications' last digits are less dispersed across examiners than we would expect if last digits were not used for application assignment.[28] 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 5%, indicating that these art units use application last digit to assign patents. The test only rejects the null that last digits are *not* used and it can of course not guarantee that in art units with a p-value below 5% 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

[26]When 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.

[27]Righi 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.

[28]Formally, 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. See Mori et al. (2005), Righi and Simcoe (2017) and the Online Appendix for details.

applications according to the Chi-square test (p-value < 0.05), 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 > 0.05), 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 0.05 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.[29] 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.

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 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 by year fixed effects. Within an art unit and a year, assignment probabilities vary only because of changes in (relative) busyness across examiners. We estimate assignment

[29] This 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.

probabilities using a simple linear probability model, presented in the Online Appendix, to reduce the computational burden.

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:

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

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

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 scientific 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 linearly in (3). The estimated signal standard deviations are very similar to our baseline estimates from Table 2. The results remain similar when introducing higher-order controls for the examiner grant rate. 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) also controls for initial characteristics of the patent application, namely the number of previous patents of the assignee and of the inventor and the initial number of independent claims. 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 works, 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: $p \sim Beta(\alpha, \beta)$. 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 the Online Appendix that the integrated likelihood is:

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

which we estimate via maximum likelihood in each art unit. Having recovered estimates of the hyperparameters, $\hat{\alpha}$ and $\hat{\beta}$, we compute the posterior mean for each examiner:[30]

$$\widehat{\mu_j}^{BetaBinomial} = \frac{\hat{\alpha} + r_j}{\hat{\alpha} + \hat{\beta} + n_j}.$$

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

[30] Intuitively, 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.

test predictive accuracy, we regress the raw examiner effect from the first subsample on the shrunk examiner effects from the second subsample.[31] 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 by 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.

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 discriminate between various theories of PAE behavior.

[31] We 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.

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 (Hagiu 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)).

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:

$$Y_i = \beta E_{j(i)} + a_{ut(i)} + \epsilon_i, \tag{7}$$

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.[32]

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 specifica-

[32] Specification (7) is analogous to the regression underlying Figure 1, except that we are now using properly scaled regressors.

tion (7) 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).[33] 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 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 (7) 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 demands 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

[33] One would need a quasi-experiment that directly affected specific behaviors (e.g. a training program) in order to recover more granular causal impacts.

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.[34]

Table 4 presents the results with patent acquisition as the outcome.[35] In both panels, the 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 high, 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 7.4% increase in the probability of PAE purchase;[36] 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.

[34] Although all examiners are supposed to apply the same standards for patent grant, which are determined by patent law, the Online Appendix shows that there is wide causal variation across examiners in terms of their propensity to refer to the various sections.

[35] The sample is restricted to art units that are part of Information Technologies since PAEs are primarily active in these art units. The results are similar in the full sample, as shown in the Online Appendix.

[36] More 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.

Panel B of Table 4 shows the results of patent purchase by practicing firms, which offer an interesting 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 (3.5%). Overall, there appears to be no quantitatively large or statistically robust relationship between purchases by practicing firms and examiner leniency.

The fact that only PAEs selectively purchase patents issued by lenient examiners challenges the view that PAEs solve a generic friction in the patent market. If PAEs were primarily lowering transaction costs or solving liquidity problems, we would not expect their purchases to be driven by examiner effects related to examiner leniency, which in contrast does not affect patent acquisitions by practicing firms. To examine whether PAEs may be addressing a more patent-specific friction related to the patent examination process itself, we now investigate the correlates of patent litigation.

Table 5 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 4 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 suggest 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 4 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 section 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 mulitvariate regression including all examiner effects simulatenously 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 5 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 4). For instance, a one standard deviation increase 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 patent 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 patent (Table 4) but their propensity to assert patents issued by lenient examiners is merely a more extreme version of the litigation behavior of practicing firms (Table 5).

The patterns in the data are therefore difficult to reconcile with both 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 is the symptom of the way patent rights are crafted by lenient examiners, which affects litigation more broadly. 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.[37]

[37] Of 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.

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 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 (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 run specification (7) 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 4 and 5. We also consider the best linear predictor for patent purchase by PAEs using the specification in Column (9) of Table 4. The results are reported in Table 6.

Our three proxies for patent invalidity were described in detail in Section II.A. As a brief reminder, we introduce a proxy for patent invalidity which does not suffer from selection effects: 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 (Panel A of Table 6). We also use two common proxies for patent invalidity in selected sub-samples: court rulings (Panel B of Table 6) and trials at the patent office called "Inter Partes Reviews" (Panel C of Table 6)

We find a very strong and robust relationship between examiner leniency and our first proxy for patent invalidity, the reissuance of granted patents. Panel A of Table 6 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 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. The relationship between examiner leniency and reissuance rates is therefore very robust and strong. 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 6 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 6 uses a proxy for patent invalidity from the perspective of the Patent Office. Since 2012, it has been possible for any party other than the patent owner to initiate an "Inter Partes Review" (IPR) procedure, which requires the patent office to review the validity of the patent based on specific sections of patent law. If the patent office deems that it is likely that the patent contains at least one mistake, the procedure escalates to the next stage and the IPR is said to be "instituted." Because the "institution" rate of IPRs is very high, close to 80%, either the occurrence of and IPR or the institution of the IPR can be used as proxies for patent invalidity. Columns (1) to (3) of Panel C of Table 6 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 because the selected sample of patents that go through an IPR is too small to provide adequate power.

In sum, Table 6 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.[38]

IV.D Additional Results and Robustness Checks

In this subsection, we report a series of additional results and robustness checks.

First, we document the robustness of the signal standard deviations of examiner effects for PAE purchases and litigation. Table 7 reports the results. Row (A) includes in the sample granted patents from continuation applications. Row (B) uses the sample of all non-continuation granted patents from 1976 to 2015, available from Google. Row (C) excludes from the main sample the patents purchased by Intellectual Ventures and only considers patents purchased by other PAEs. Row (D) use an alternative list of PAEs from Cotropia et al. (2014). The specification used in row (E) controls for the number of independent claims at application and for the first inventor's number of previous patents. Row (F) repeats the estimation including assignee fixed effects. Across all sub-samples and specification, we find consistent and large examiner effects.

Second, we investigate the amount of heterogeneity across different types of PAEs. Existing research has hypothesized that large and small PAEs may behave differently (Cotropia et al. (2014)). We test this hypothesis in the context of examiner effects. Figure 5 reports the correlations between PAE purchase rates and the two main proxies for examiner leniency from Section IV.B, namely the grant rate and the change in the number of words per claim. We find no noticeable differences across PAE types. Large PAEs, small PAEs, and Intellectual Ventures considered separately all selectively purchase patents from more lenient examiners.

Additional results and robustness checks are reported in the Online Appendix. First, we document that examiner effects are stable over time and that career effects do not have a significant impact on patent acquisition or litigation by PAEs. Second, we show that the results in Tables 4, 5 and 6 are robust to the inclusion of additional controls and to the use of alternative measures for the use of specific sections of patent law by examiners. Finally, we report evidence that the patent

[38]This 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.

from lenient examiners purchased by PAEs are cheaper (using patent auction data) and close to existing technologies (using EPO grant decisions as a measure of step size, as in Picard and Van Pottelsberghe de la Potterie (2011)), which is consistent with the notion that these patents are productive for litigation.

V Conclusion

In this paper, we have shown that much of the heterogeneity in patent outcomes results from the process of creation of patent rights and is independent of scientific determinants. 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 policy debate has focused on a possible reform of patent law to reduce PAEs' activities and litigation, which 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.

Bloom, Nick, Stephen Bond, and John Van Reenen, "Uncertainty and investment dynamics," *The review of economic studies*, 2007, *74* (2), 391–415.

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.

Dixit, Avinash K and Robert S Pindyck, *Investment under uncertainty*, Princeton university press, 1994.

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"," 2017.

Frakes, Michael D. and Melissa F. Wasserman, "Is the Time Allocated to Review Pattent 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.

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 Rebeccajaffe1993geographic 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 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, 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
Future patents filed 1-5 years (first-time patentees)	0.6173	0	12.447	63,385
Total patent citation w/n 3 years of grant	0.5256	0	1.461	988,585
Patent citations by same assignee w/n 3 years of grant	0.1134	0	0.7257	988,585
Patent citations by other assignees w/n 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
Rate of invalidity, conditional on court ruling	0.1880	0	0.3911	516
Rate of infringement, 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

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.

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

	Signal S.D.		S.D. of Shrunk	Sample Size,
	% of Average	Level	Effects, % of Average	Patents/Examiners
	(1)	(2)	(3)	(3)
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 effets 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. 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)
Rate of patent acquisition by PAEs	38.16	41.25
Rate of patent acquisition by non-PAEs	14.52	10.01
Rate of patent litigation by non-PAEs	41.85	64.36
Sample	Art units allocating patents by last digits, according to Chi-square test	Art units in Information Technology
Number of patents in sample	557,562	311,686
Art Units	243	254

Panel B: Accounting for Extensive Margin Selection Effects

	Signal S.D., % of Average (1)	(2)
Rate of patent acquisition by PAEs	62.64	65.02
Rate of patent acquisition by non-PAEs	14.31	13.51
Rate of patent litigation by non-PAEs	55.45	60.95
Controls	Examiner grant rate	Examiner grant rate and application characteristics
Fixed Effects	Year by art unit	

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, described in Section III.C. 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, described in Section III.C.

Table 4: 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	-0.139***								-0.115***
Claim from Application to Grant	(0.030)								(0.0490)
% Change in Number of Claims		0.073**							0.0519
from Application to Grant		(0.034)							(0.0345)
Grant Rate			0.114***						-0.0298
			(0.028)						(0.0637)
Use of Section 101				-0.061*				-0.094	-0.0225
- Lack of utility or eligibility				(0.036)				(0.082)	(0.0366)
Use of Section 102(a) -					0.007			0.0196	0.00835
- Prior art exists					(0.021)			(0.089)	(0.0216)
Use of Section 103(a)						-0.0602***		-0.176***	-0.0223
- Obvious invention						(0.024)		(0.083)	(0.0285)
Use of Section 112(b)							-0.037	0.102	-0.00392
- Vague claims							(0.027)	(0.087)	(0.0291)
Fixed Effects	Year by Art Unit								
N	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	0.0071								0.0349***
Claim from Application to Grant	(0.0081)								(0.011)
% Change in Number of Claims		-0.0003							-0.0001
from Application to Grant		(0.006)							(0.006)
Grant Rate			0.022***						0.062***
			(0.0082)						(0.012)
Use of Section 101				0.0147**				0.0154**	0.0174**
- Lack of utility or eligibility				(0.0065)				(0.00711)	(0.0073)
Use of Section 102(a) -					-0.0037			-0.00498	-0.007
-Prior art exists					(0.005)			(0.00556)	(0.006)
Use of Section 103(a)						0.0065		0.00539	0.007
-Obvious invention						(0.005)		(0.00633)	(0.006)
Use of Section 112(b)							0.002	-0.00419	0.003
-Vague claims							(0.005)	(0.00619)	(0.006)
Fixed Effects	Year by Art Unit								
N	274,464								

Notes: The sample is retricted 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 5: 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	-0.405***								-0.097
Claim from Application to Grant	(0.083)								(0.12)
% Change in Number of Claims		0.127***							0.05
from Application to Grant		(0.067)							(0.07)
Grant Rate			0.567***						0.48***
			(0.099)						(0.14)
Use of Section 101				-0.105				-0.09	0.05
- Lack of utility or eligibility				(0.077)				(0.08)	(0.08)
Use of Section 102(a) -					0.0178			0.019	0.023
- Prior art exists					(0.089)			(0.08)	(0.082)
Use of Section 103(a)						-0.156***		-0.176**	-0.039
- Obvious invention						(0.075)		(0.083)	(0.08)
Use of Section 112(b)							-0.0003	0.1	0.085
- Vague claims							(0.079)	(0.08)	(0.086)
Fixed Effects	Year by Art Unit								
N	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	-0.138***								0.017
Claim from Application to Grant	(0.043)								(0.071)
% Change in Number of Claims		0.022							-0.015
from Application to Grant		(0.031)							(0.034)
Grant Rate			0.24***						0.23***
			(0.045)						(0.067)
Use of Section 101				-0.068*				-0.0205	0.005
- Lack of utility or eligibility				(0.037)				(0.0397)	(0.04)
Use of Section 102(a) -					-0.008			0.0150	0.026
- Prior art exists					(0.04)			(0.0406)	(0.04)
Use of Section 103(a)						-0.075**		-0.0387	0.008
- Obvious invention						(0.034)		(0.0370)	(0.04)
Use of Section 112(b)							-0.118***	-0.0978**	-0.065
- Vague claims							(0.032)	(0.0366)	(0.04)
Fixed Effects	Year by Art Unit								
N	274,464								

Notes: The sample is retricted 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: Examiner Behavior and Likelihood of Patent Invalidity

Panel A: Reissuance of Granted Patents

	Leave-one-out Examiner Effects	Reissuance Rate			Reissuance Rate Two Years or More after Grant		
	(separate regressions)	(1)	(2)	(3)	(4)	(5)	(6)
(A)	% Change in Number of Word per Claim from Application to Grant	-0.26***	-0.24***	-0.25***	-0.55***	-0.57***	-0.61***
		(0.07)	(0.06)	(0.068)	(0.15)	(0.14)	(0.15)
(B)	Grant Rate	0.29***	0.27***	0.28***	0.54***	0.53***	0.54***
		(0.06)	(0.06)	(0.061)	(0.13)	(0.13)	(0.13)
(C)	Linear Predictor for PAE Acquisition	0.094***	0.099***	0.10***	0.17**	0.19***	0.22***
		(0.035)	(0.034)	(0.035)	(0.076)	(0.07)	(0.07)
	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	Invalidity Rate			Infringement Rate		
	(separate regressions)	(1)	(2)	(3)	(4)	(5)	(6)
(A)	% Change in Number of Word per Claim from Application to Grant	0.02	0.068	0.11	-0.01	-0.0001	-0.0002
		(0.06)	(0.29)	(0.32)	(0.06)	(0.24)	(0.28)
(B)	Grant Rate	0.02	-0.039	-0.019	-0.03	0.10	-0.12
		(0.03)	(0.26)	(0.54)	(0.03)	(0.54)	(0.22)
(C)	Linear Predictor for PAE Acquisition	-.0748*	-0.014	-0.051	0.01	0.017	0.018
		(0.0414)	(0.186)	(0.212)	(0.03)	(0.10)	(0.11)
	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	IPR Rate			Institution Rate of IPR		
	(separate regressions)	(1)	(2)	(3)	(4)	(5)	(6)
(A)	% Change in Number of Word per Claim from Application to Grant	-0.36***	-0.41***	-0.40***	-0.05	-0.05	-0.20
		(0.08)	(0.08)	(0.07)	(0.052)	(0.26)	(0.27)
(B)	Grant Rate	0.41***	0.44***	0.44***	0.05	-0.03	-0.034
		(0.08)	(0.08)	(0.08)	(0.047)	(0.11)	(0.12)
(C)	Linear Predictor for PAE Acquisition	0.26***	0.30***	0.29***	0.036	0.03	0.213
		(0.09)	(0.08)	(0.08)	(0.035)	(0.28)	(0.37)
	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 retricted 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 4. Standard errors are clustered by examiners. *$p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Robustness Checks on the Effect of Examiners on Patent Acquisition by PAEs

	Signal S.D. for PAE Purchases, % of Average (1)	S.D. of Shrunk Examiner Effects for PAE Purchases, % of Average (2)
(A) Including continuations	77.2%	48.1%
(B) Granted patent from 1976 to 2015	70.6%	24.8%
(C) Excluding Intellectual Ventures Patents	71.2%	33.5%
(D) Cotropia et al. (2014) PAE Patents	43.1%	15.5%
(E) Including Patent Controls	54.9%	24.8%
(F) Including Assignee Fixed Effects	33.9%	13.2%

Notes: This table report statistics on examiner effects for patent acquisitions by PAEs using alternative samples and specifications. Row (A) includes in the sample granted patents from continuation applications. Row (B) uses the sample of all non-continuation granted patents from 1976 to 2015, available from Google. Row (C) excludes from the main sample the patents purchased by Intellectual Ventures and only considers patents purchased by other PAEs. Row (D) use an alternative list of PAEs from Cotropia et al. (2014). The specification used in row (E) controls for the number of independent claims at application and for the first inventor's number of previous patents in equation (1). Row (F) repeats the estimation including assignee fixed effects. All reported values are normalized by the baseline rate of PAE purchases, which differs across samples.

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

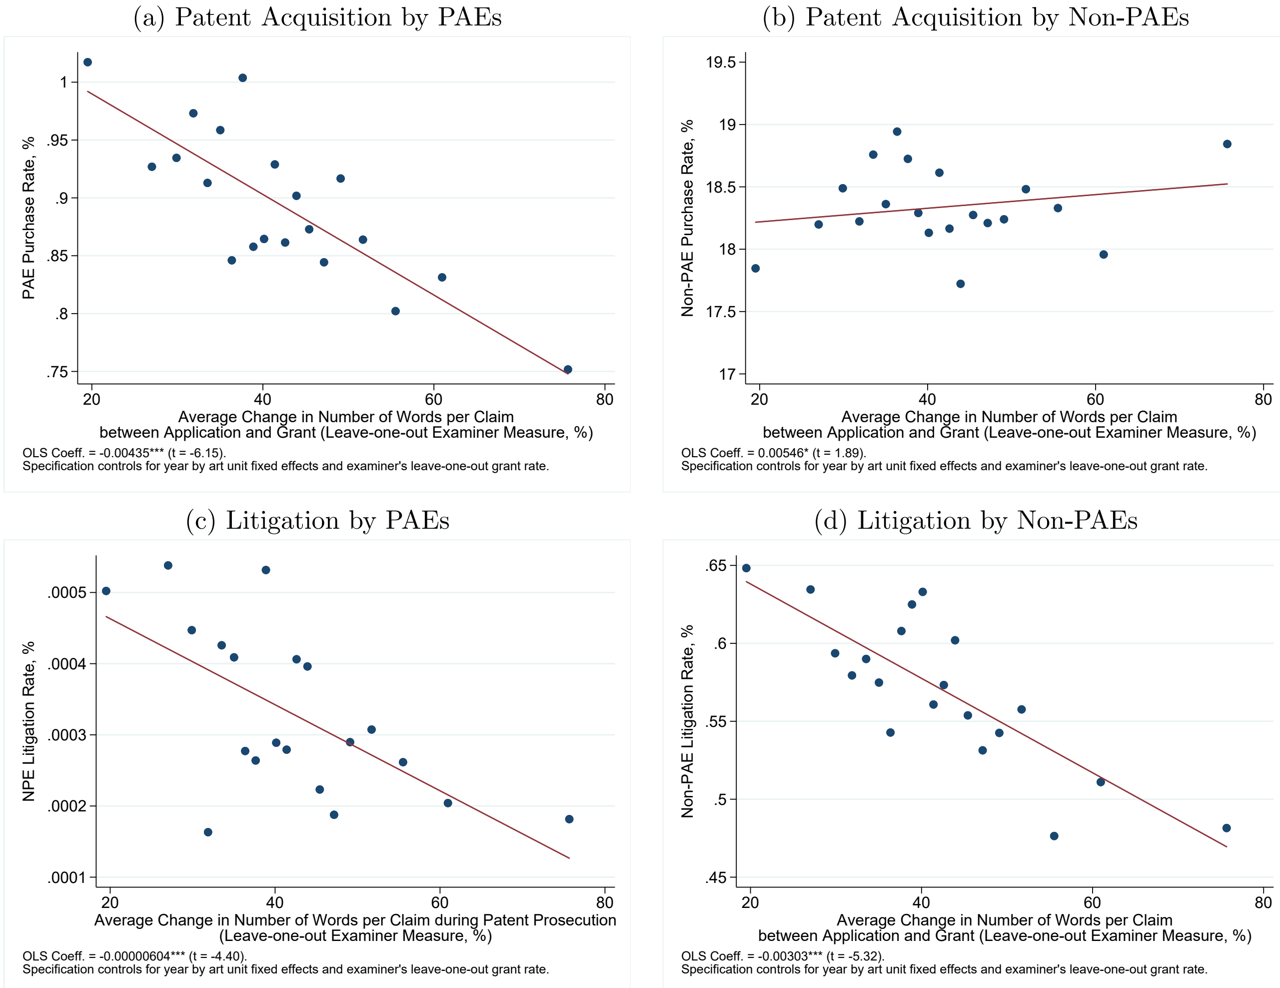

Notes: In the various panels of this figure, the level of observation is a patent. The average change in the number of words per claim is measured at the level of an examiner, leaving out the focal patent. All specifications include art unit by year fixed effects and address potential extensive margin 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

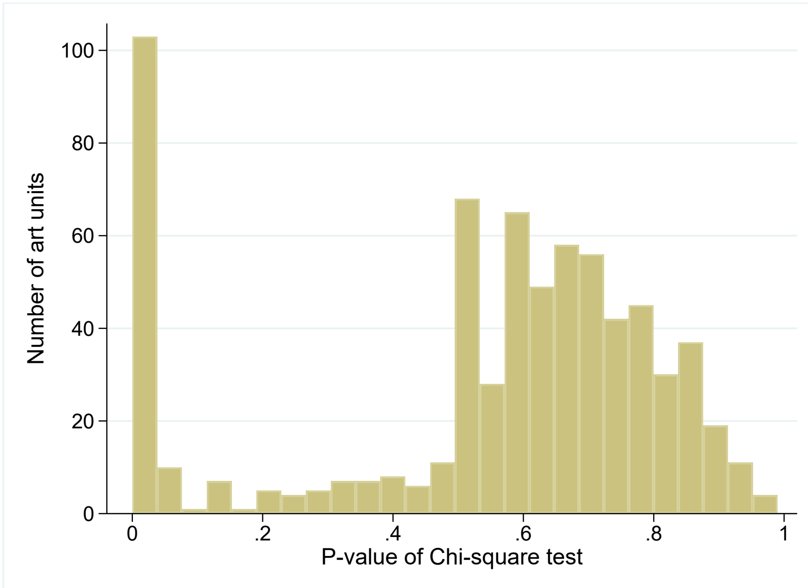

Panel B: Graphical Evidence on Allocation by Application's Last Digit

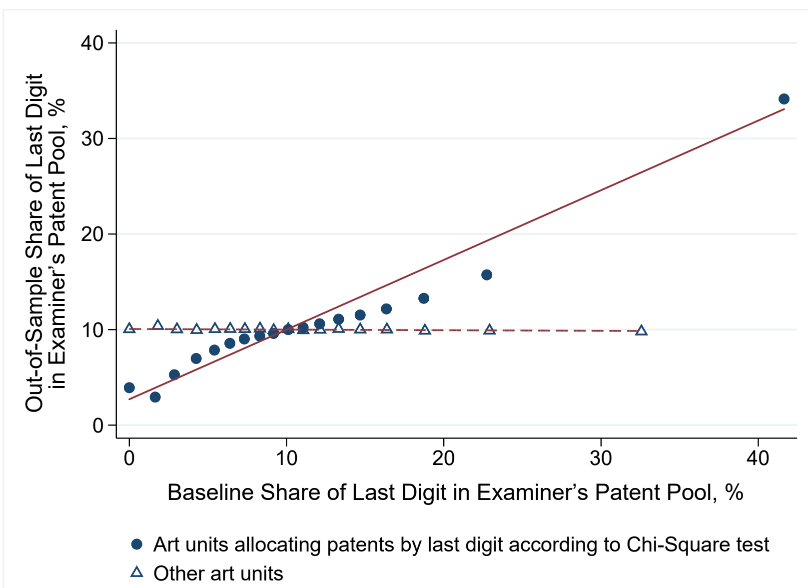

Notes: In Panel A, the level of observation is an art unit. This panel reports the distribution of the p-values of the Chi-square tests described in the main text; a p-value below 0.05 indicates excess concentration of patent applications across examiners by application's last digit. In Panel B, the 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

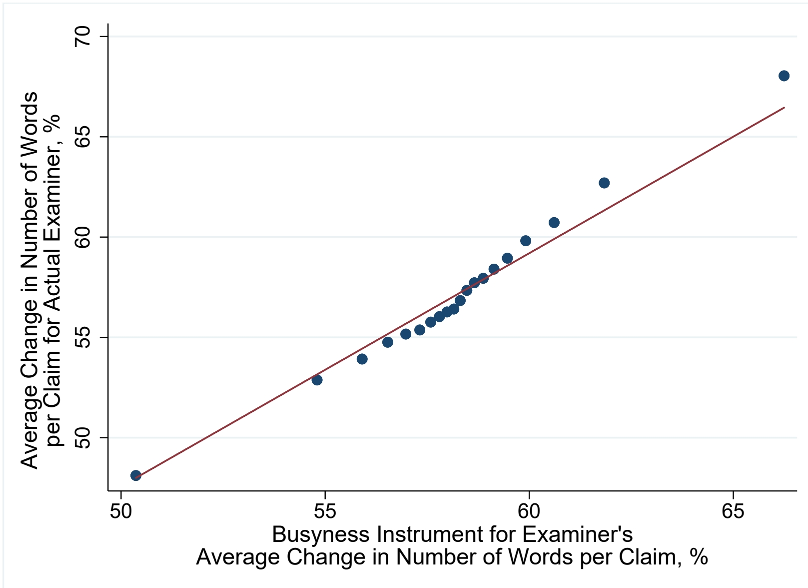

Panel B: Reduced Form

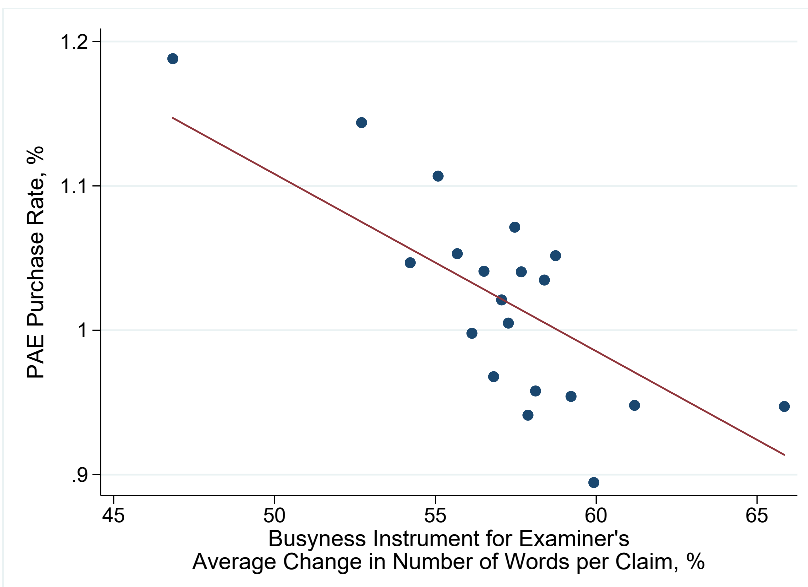

Notes: Panel A shows the relationship between the busyness instrument (described in the main text) for an examiner's propensity to change the number of words per claim during application and grant and the propensity of the examiner whom the application was actually assigned to. Panel B depicts the relationship 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

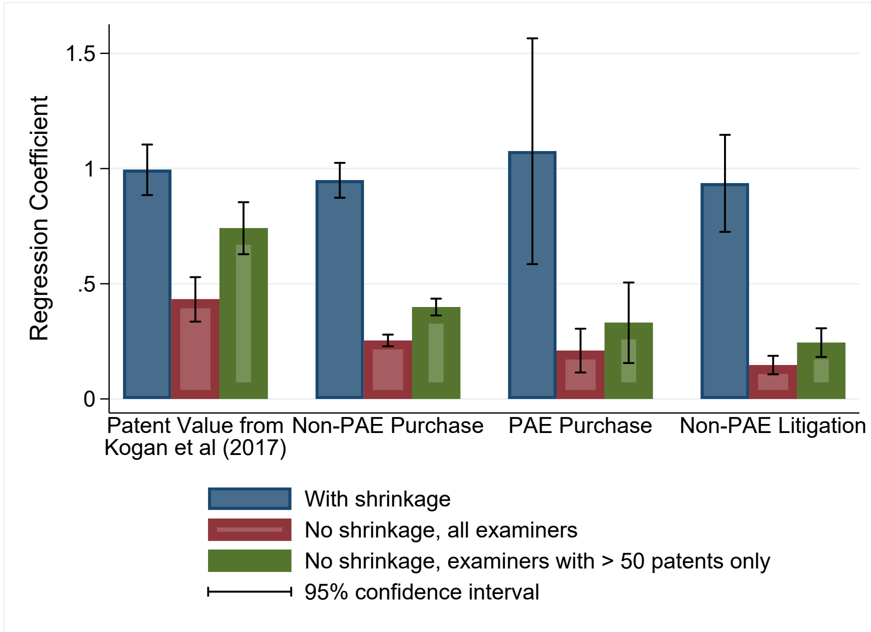

Notes: This figure reports the OLS coefficients in examiner-level out-of-sample regressions. After splitting the main analysis sample into two halves at random, we compute the raw and shrunk examiner effects on each half following the methodology outlined in Section III.C. To test predictive accuracy, 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

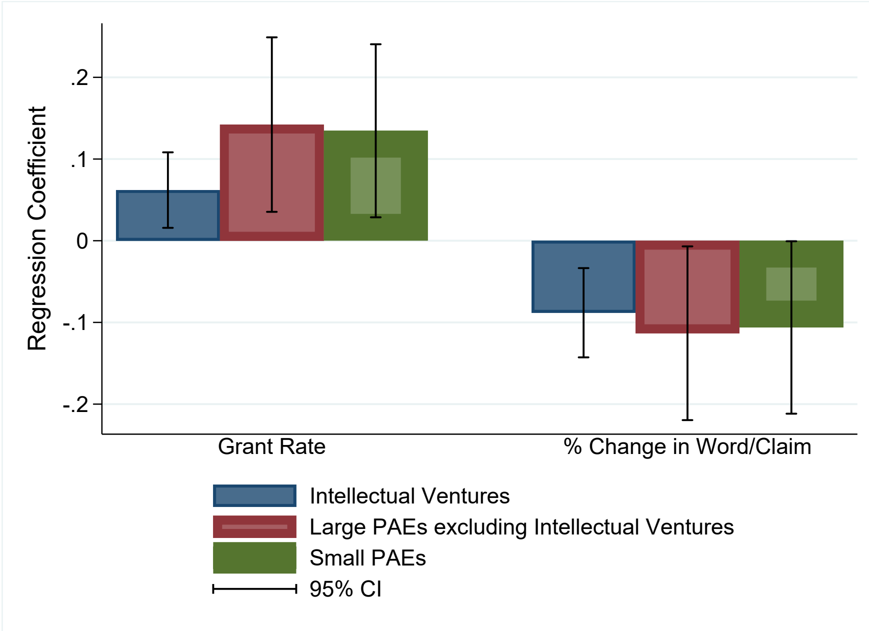

Notes: This figure documents the relationship between examiner behavior and patent acquisition for different types of PAEs. The methodology is similar to Table 4 and is described in Section IV.B. Regression coefficients are reported separately for three samples of PAEs: Intellectual Ventures, a large PAE holding many publicly-listed patents; “Large PAEs”, referring to PAEs with a patent portfolio size in the top 1% and excluding Intellectual Ventures; and “Small PAEs”, which are identified using the classification of small patent holding companies in Cotropia et al. (2014).