Who Feeds the Trolls?
Patent Trolls and the Patent Examination Process*

Josh Feng, Harvard University
Xavier Jaravel, Stanford University

July 9, 2016

Abstract

Leveraging random examiner assignment and detailed patent prosecution data, we find that non-practicing entities (NPEs) purchase patents granted by examiners that tend to issue incremental patents with vaguely-worded claims. In comparison, practicing entities purchase a very different set of patents, but assert patents similar to those purchased by NPEs. These results show that on average NPEs purchase and assert patents productive for litigation but lacking technological merit, thus adding to overall litigation fees without providing incentives for high-quality innovations. Their activities are in part the symptom of a broader problem with the issuance of ill-defined intellectual property rights, which leads to (likely inefficient) litigation even among practicing entities. A cost-benefit calibration shows that investments in improving examination quality at the United States Patent Office would have large social returns.

*We thank Philippe Aghion, Joseph Bailey, Gary Chamberlain, Raj Chetty, Lauren Cohen, Ken Felter, Lee Fleming, Ed Glaeser, Stu Graham, Stephen Haber, Nathan Hendren, Bill Kerr, Robert Kimble, Scott Kominers, Mark Lemley, Josh Lerner, Guan-Cheng Li, Alan Marco, Filippo Mezzanotti, Richard Miller, Markus Mobius, Amanda Myers, Alexandra Roulet, Asrat Tesfayesus and seminar participants at the NBER Productivity Lunch Seminar, USPTO Visiting Speaker Series, Harvard Labor Lunch, and Harvard IO Lunch for thoughtful discussions and comments. We are particularly grateful to Michael Frakes and Melissa Wasserman for sharing their data on examiners' blocking actions. In addition, we would like to thank RPX Corporation for providing data on IP litigations and non-practicing entities. Part of this research was conducted while Jaravel was visiting the United States Patent Office. Jaravel gratefully acknowledges financial support from the HBS doctoral office and the Kauffman foundation.

1 Introduction

This paper sheds new light on the patent acquisition and assertion behavior of non-practicing entities (NPEs). NPEs, also known as patent assertion entities or patent trolls, garner profits exclusively from IP litigation and licensing without producing or selling goods. They have grown in prominence in the IP system over the past decade, attracting a large amount of scrutiny and debate along the way. In the words of President Obama, "[NPEs] don't actually produce anything themselves. They're just trying to hijack somebody else's idea and see if they can extort some money out of them."[1] The general unpopularity of NPEs within Congress has led to several attempts to pass legislation to curb their activities, including the Strong Patent Act[2] and the Innovation Act[3], but arguments over the definition of NPEs and the nature of their activities have prevented them from passing. In defense of their business model, NPEs argue that they serve as intermediaries that improve the efficiency of the market for ideas, by helping resource-constrained inventors and firms enforce their patents against infringing entities.[4] By law, all issued patents are supposed to cover only "novel" and "non-obvious" inventions, but an average application gets less than 20 hours of patent examiner time, and a large proportion of the few patents later fully evaluated in court are held invalid. A blatantly invalid patent, which would almost surely be overturned in court, would never be asserted and would thus cause no harm. The criticism of NPEs is based on the concern that they purchase and assert patents that are not clearly invalid, but are "weak",[5] lack technological merit and therefore do not provide incentives for high-quality innovations. We find strong empirical support for this view, and we show that NPEs' activities are part of a broader problem with the issuance of ill-defined patents, which could likely be solved by improving the quality of patent examination at the United States Patent Office (USPTO).

Before turning to the specifics of our research design, we preview our various findings. On average, do NPEs tend to purchase and assert "weak" patents? We give a positive answer to this question by documenting that, controlling for art unit[6] and year fixed effects, NPEs purchase and assert patents that were granted by a specific set of examiners, who tend to allow incremental patents with vaguely-worded claims. In contrast,

[1] Google Hangout Session, February 14th 2015.

[2] https://www.congress.gov/bill/114th-congress/senate-bill/632

[3] https://www.congress.gov/bill/113th-congress/house-bill/3309

[4] As explained on Intellectual Ventures' website, "We purchase patents from individual inventors, start-ups, large corporations, research institutions, and everything in between. By acquiring these inventions, we provide capital to inventors and give their ideas a better chance of getting into the marketplace."

[5] Weak patents are defined as patents that may well be invalid, but nobody knows for sure without conclusive litigation (Farrel and Shapiro, 2008)

[6] Art units are small working groups, typically composed of about twenty examiners processing patent applications in similar technology classes.

the purchasing behavior of practicing entities is not skewed toward any particular set of examiners. But practicing entities are more likely to assert patents granted by examiners similar to those who grant the patents purchased and asserted by NPEs. This finding is in line with the idea that disagreement between parties over ill-defined intellectual property rights leads to litigation.[7] We present additional evidence from European Patent Office (EPO) grant outcomes (for patents filed jointly at EPO and USPTO) and from litigation trial outcomes (for litigated patents) showing that NPE patents tend to be of particularly low technological merit, even relative to patents asserted by practicing entities. We thus conclude that on average NPEs do not provide incentives for high-quality innovations. Moreover, their activities are in part the symptom of a broader problem with the issuance of ill-defined intellectual property rights, which leads to (likely inefficient) litigation even among practicing entities.

In addition to being the first large-scale empirical study of the patent acquisition behavior of patent trolls, this paper speaks to another important debate about the patent system: do patent examiners have a significant influence on the eventual outcomes and uses of the patents they grant, or do existing rules at the USPTO constrain them such that they have little ability to reject of modify patent applications[8]? Answering this question is key to establish whether policy changes related to examination practices have an important role to play in reforming the Intellectual Property (IP) system. Recently, the USPTO began implementing the so-called *Enhanced Patent Quality Initiative* (EPQI), which aims at improving patent correctness[9] and clarity[10] by training examiners and running pilots to detemine which examination practices can best achieve these two goals. In a recent blog post[11], Michelle Lee, Director of the USPTO, suggested that the renewed focus on clarity in the EPQI was necessary because of "the evolving patent landscape" and that improving patent clarity could reduce "needless high-cost court proceedings." Given the examiner-focused policies included in the EPQI and its intended goals, whether examiners actually have a significant influence on eventual outcomes and uses of patents is an important policy question.[12] We show empirically

[7] Intuitively, if property rights are well defined and everyone knows which patent is valid or not, it will be more efficient for the parties involved to settle. In leading litigation models, litigation occurs because of asymmetric information, when there is disagreement over patent validity. In practice, it is well-known that most litigation cases involving NPEs are held at the Eastern District Court of Texas, which is widely believed to have a pro-plaintiff bias. Therefore the outcome of patent litigation may not accurately reflect the underlying quality of the intellectual property, at least in the case of NPEs.

[8] For instance, it may be very difficult for examiners to reject weak patents because the rules of the USPTO are such that the burden of proof of non-patentability rests on the examiners. This view is articulated in Lei and Wright (2009), who emphasize that other patent offices like the European Patent Office place much fewer constraints on the examiners.

[9] Correctness refers to correctly judging whether an application meets all patentability criteria.

[10] Clarity refers to whether the granted patent clearly defines the technology covered by it.

[11] http://www.uspto.gov/blog/director/entry/enhanced_patent_quality_initiative_moving

[12] This paper is thus related to a growing literature in public economics that studies how to improve the management and retention of employees in the public sector, in particular in the context of teachers (e.g. Chetty et al., 2014).

that patent examiners at the USPTO indeed have a large causal impact on how their granted patents are used within the intellectual property system. Their impact is evident both through the "extensive margin" of accepting or rejecting patent applications and through the "intensive margin" of forcing modifications to patent applications during the examination process, which includes a "back-and-forth" between the applicant and the examiner. More broadly, our evidence shows the importance of the micro-determinants of intellectual property: important patent outcomes like litigation largely depend on examiner behavior and not simply on macro-determinants of the IP system such as the statutes in Title 35 of the US code.

Our research design starts from the fact that patent applications are conditionally randomly assigned to examiners.[13] Specifically, within an art unit applications that are not continuation applications are randomly assigned to examiners. Leveraging this fact, we adapt the methodology used in the teacher value-added literature (e.g. Kane and Staiger, 2008) to estimate the causal impact examiners have on the probability that their granted patents are litigated or purchased by NPEs. The intuition underpinning this research design is as follows: start from a patent outcome such as NPE purchase, calculate the share of an examiner's granted patents that feature this outcome, and compare this share across examiners working in the same year and art unit. If there is large heterogeneity in this share across examiners for a particular outcome, and if this pattern is stable over time (i.e. an examiner's granted patent consistently feature a given outcome across years), then our methodology will yield a large causal effect from examiners for this outcome. We find that there is significant variation in the examiner causal effects for both NPE litigation and NPE purchase outcomes. The signal standard deviation of the examiner causal effect distribution for patent purchase by NPEs corresponds to about 51% of the baseline NPE patent assertion rate. The corresponding number for patent litigation by a practicing entity is 62%. In contrast, it is much lower for patent purchase by practicing entities, around 14%. This provides initial evidence that examiners do have a significant influence on important IP market outcomes, and that the patent acquisition behavior of NPEs differs from that of regular firms.

As a second step, we distinguish between mechanisms that could explain the examiner variation for each of the patent outcomes of interest. To do so, we calculate the causal effects of each examiner for various behaviors observed during prosecution, such as their tendency to edit the text of the application and to cite various patentability criteria when rejecting a patent application. Using detailed prosecution data in

[13] This fact has been discussed in Lemley and Sampat (2012). We have conducted our own interviews with patent examiners and the novelty of our approach is to exclude continuation applications (including continuation, continuation-in-part, and divisional applications). This is an important adjustment for NPEs' portfolios, which have a high share of continuation patents, as documented in Section 3.

this way allows us to measure whether a given patent examiner typically clarifies, narrows or rejects patent applications, which in turn allows us to infer the whether the patents granted by this examiner tend to be "weak." Then, we use leave-one-out versions of these examiner causal effect to predict the actual outcome of a given patent, such as NPE purchase and assertion in litigation. Another way to view this exercise is from a patent statistics perspective: we are projecting examiner characteristics onto the patents that they grant, and then comparing these randomly assigned characteristics across patents that are or are not purchased by NPEs. We discuss in the main text the selection effects inherent in this research design (some outcomes are observed only conditional on the patent application being granted).

We find that patents purchased by NPEs are, on average, granted by examiners who allow more incremental patents and patents with vaguer language. This result also holds for the subset of patents that are asserted by NPEs in litigation. In contrast, practicing entity purchasing behavior exhibits very little dependence on the behavior of examiners. Importantly, patents litigated by practicing entities also appear to be more incremental and vaguer relative to others in their technological cohort, but the effect sizes are much smaller than those exhibited by NPE litigated patents.

Taken together, our results show that NPEs purchase and assert different types of patents than regular firms, and appear to engage in a form of rent seeking by purchasing patents for reasons orthogonal to or even negatively correlated with technological merit and social value. A significant share of patents litigated by practicing entities also appear to be of low technological value, suggesting that the issuance of incremental and ill-defined patents by some examiners results in a broader problem for the IP system, beyond the activities of NPEs.

We provide additional evidence that NPEs tend to strategically buy and assert "weak" patents in two ways. First, we find that conditional on examiner and assignee fixed effects, NPEs are much more likely to purchase patents that were rejected at the European Patent Office. In other words, within the portfolio of a given examiner, NPEs are more likely to purchase more incremental patents. We find no such effect for non-NPE purchases or litigation. Second, we find that the patents of examiners with high NPE-purchase causal effects have a higher rate of invalidation in litigation cases that go to trial. This also suggests that NPE patents have low technological quality and are perhaps invalid, but it is only suggestive given that any analysis of litigation trial outcomes is beset by important selection effects.

In terms of policy, our results suggest that examiner-focused reforms at the USPTO, such as examiner

training or hiring more examiners, may have high social returns in terms of limiting the supply of inputs desired by NPEs. The lower-bound cost estimates we obtain in the paper, albeit in a partial equilibrium framework, suggest that these policies may represent excellent public investments, with expected returns close to 80%.

There are a few caveats and limitations to our study, but we believe they do not alter the general interpretation of our results. First, we only analyze observable outcomes in the IP system such as patent transfer and patent litigation, which do not encompass other important outcomes such as demand letters and licensing. Second, we may not have found all patents belonging to NPEs, given their tendencies to obfuscate their portfolios using shell companies, but this would only weaken our quantitative results - besides, our results are extremely robust across all available lists of NPEs we know of. Finally, the nature of our methodology means that we are unable to make judgments on the behaviors of specific NPEs if they have small patent portfolios. This means that there may be NPEs that act more like IP intermediaries rather than rent seekers. However, given our findings, reforms to the patent examination process would actually serve as a sharp policy instrument, affecting only NPEs that rely on vague and obvious patents.[14]

The remainder of this paper is structured as follows. Section 2 provides background on the patent examination process, non-practicing entities, and patent and USPTO policies. Section 3 presents summary statistics and preliminary evidence based on the data we collect on NPE patent purchases and USPTO examiners. Section 4 describes the empirical framework and presents the baseline results on the size of the examiner causal effects. Section 5 conducts heterogeneity analysis to distinguish betwen various mechanisms. Section 6 discusses policy counterfactuals using our previous results. Section 7 concludes and discusses directions for future research.

2 Background

2.1 USPTO Examination Process

This paper relies extensively on the data from and the features of the patent examination system at the USPTO, which we describe briefly here.[15] When the USPTO receives a patent application, it first sends it to a vendor to classify the application based on technological area. Based on this initial classification, the

[14] As we will discuss later, such reforms would be beneficial even if NPEs served as efficient screening intermediaries that sift through patents with vaguely worded claims for ones with technological value. Reforms would deal with the vague claim problem at the source rather than having NPEs fund their screening through costly litigation.

[15] See Appendix B of the Sampat and Williams working paper for a more detailed explanation. See Lemley and Sampat (2012) for a detailed exploration of the assignment system.

application is then assigned to one of around 600 art units, which are comprised of groups of examiners with similar technological expertise. Once assigned to an art unit, the patent is assigned to examiners either based on examiners' current workload or based on the last digit of the application (Lemley and Sampat (2012)). As discussed in greater detail later in the paper, this assignment process is as good as random within art units.

Once the application has been assigned to an examiner, it is then evaluated by the examiner to ascertain whether it meets all the criteria for grant. If it does, then the examiner grants the patent immediately. According to USPTO figures, the percentage of patents granted right away ("first action") is around 10% only (Carley et al (2014)). Typically, the examiner issues either a final or a non-final rejection, also known as blocking actions, citing all of the eligibility criteria that are not met by the patent application. After rejections of either type, the applicant can abandon the application or send a response to the patent office, often editing the text of the application in order to address the criteria that the patent examiner deemed unmet. This then starts a back-and-forth process between the examiner and the applicant until the application is either approved or abandoned. In our analysis, we focus on applications with either of these outcomes, also known as disposed applications. The nature of the examination process described above shows that there is a possibility for examiners to differ not only in terms of the rate at which they grant patent application (as has been extensively discussed in the existing literature using "examiner instruments"), but also in terms of how they affect the nature of the eventually granted patent (breadth, clarity, etc.).

Examiners can affect the nature of the patents they grant through the aforementioned blocking actions. Based on the criteria cited in a rejection, an applicant will accordingly edit the text of the patent application before sending it back to the examiner. The main patentability criteria that examiners cite when blocking an application are the following sections from Chapter 35 of US Code:

1. Section 101 (Patentability and Utility): The patent needs to satisfy eligibility requirements and have the potential to be useful.

2. Section 102 (Novelty): The patent's claims are not covered in the prior art or in older academic literature.

3. Section 103 (Non-obviousness): The difference between the invention and prior art is significant enough that it would not have been obvious to a "person having ordinary skill in the art."

4. Section 112, paragraph one (Sufficient disclosure/"Enablement"): The application "shall contain a writ-

ten description of the invention... in such full, clear, concise, and exact terms as to enable any person skilled in the art ... to make and use [the invention]."

5. Section 112, paragraph two (Claim clarity/"Definiteness"): "The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the inventor or a joint inventor regards as the invention."

We show later that there is consistent and substantial variation in the usage of these provisions across examiners within the same art unit in any given year. The main focus of this paper will be on Section 103 and Section 112, paragraph two (also known as 112(b)), because these two blocking actions rely on the subjective judgement of the examiner, as reflected by the adjectives used to define them ("obvious" and "clear"), and because they can be used to test for the rent seeking behavior of NPEs.

One last element relevant for our discussion is the aforementioned claims of a particular patent. As noted in Section 112(b), claims in an application attempt to summarize and delimit the boundaries of a particular invention. The interpretation of text in a patent's claims is often the main focus of patent litigation. One concern with unclear claims text is that there will be many interpretations, and therefore a higher likelihood of confusion and disagreement over whether a product infringes on the patent. Examiners who are more likely to issue 112(b) blocking actions will on average grant patents with clearer claims text.

2.2 Relationship to Existing Literature

This paper builds on and contributes to several literatures. First, there is a growing economic and legal literature on NPEs. Cohen, Gurun, and Kominers (2014) investigate the characteristics of the defendants in NPE litigation. They find that NPEs tend to target firms with positive cash flow shocks, even if the cash flow shock hits a subdivision of a conglomerate that is different from the one accused of infringement. They also document other characteristics of firms targeted by NPEs such as number of lawyers, which provide evidence that NPE behavior is unrelated to actual infringement. Our paper complements their analysis by focusing on the characteristics of the key inputs in the NPE production function, namely patents. Despite their differences in focus and methodology, in line with our results they find that the behavior of NPEs is driven by factors unrelated to conventional use of the IP system, which can be thought of as purchasing patents based on technological merit and litigating based only on patent infringement.

Another strand of the NPE literature that we rely on is the classification of NPEs, which has generated

widespread disagreement within academia and policy-making circles.[16] One simple approach is to apply an NPE label to any entity that makes all or most of its revenues from licensing and litigation. However, this broad definition would apply to technology development companies, some university-based IP entities, and failed start-ups, entities that many consider to be important to the innovation system. To address this issue, a recent paper by legal scholars Cotropia, Kesan, and Schwartz (2014) manually classifies all plaintiffs of IP lawsuits in the years 2010 and 2012, assigning each plaintiff to one of eight categories.[17] Their categories attempt to distinguish between possibly rent-seeking patent holding companies and large patent aggregators on the one hand and the aforementioned "benevolent entity types" on the other hand. We make use of multiple sources in generating our NPE portfolios, and ascertain that our core results are robust to various portfolio construction methods.

In addition to Cotropia et al. (2014), there are several papers in the legal literature that provide descriptive evidence of NPE behavior. Fischer and Henkel (2012) show that patents acquired by NPEs are classified under more technology classes and received more citations, traditional measures of patent value. An earlier study, Allison, Lemley, and Walker (2009), considered patents that were litigated multiple times, and also found that they received more citations and were most likely to be owned by NPEs. These broad patterns also appear in the data we have collected, but our approach addresses possible endogeneity problems in citation statistics, such as the "publicity effect" of litigation on citations discussed by Lanjouw and Schankerman (2001). Our analysis suggests that in our specific setting, citations and independent claims may not be a valid indicator of social or technological value.

On the USPTO side, there is a literature focusing on the patent examination process. As mentioned earlier, we use the random assignment of applications for our identification strategy, a feature also exploited by Sampat and Williams (2015) in their investigation of follow-on innovation in human genome patents. Their and our work build on Lemley and Sampat (2012), which studies the issue of random assignment in detail by interviewing USPTO examiners. They found that random assignment within art units is plausible. Our contribution here is twofold. First, we show the importance of removing continuation applications, otherwise IV estimators are biased because continuation applications are not randomly assigned. Second, we show that examiners have an effect beyond the decision of granting the patent: they affect the nature of the granted

[16] Including the recent debate in Congress over how to define the entities which the proposed fee-shifting provisions in the Innovation Act (H.R. 3309) would apply to.

[17] University/College, Individual/family trust, Large aggregator, Failed operating company/failed start-up, Patent holding company, Operating company, IP Holding company of operating company, and Technology development company.

patent (breadth, clarity, etc.), which is important to consider when gauging the validity of the exclusion restriction in any IV framework based on examiner assignment.

Beyond the issue of random assignment, the USPTO literature also discusses the validity of issued patents. There is a growing consensus, reflected in recent legislation, that US Patent Office examiners issue many patents of dubious validity. There is a debate as to whether this is primarily due to examiners' inability to distinguish these from other valid applications or to institutional constraints that make it very difficult for examiners to reject patents. Lei and Wright (2009) document that examiners distinguish weak patents from others that are stronger and, bearing the burden of proof of non-patentability, search more intensively for prior art that might bolster a case for rejecting weaker patents. They conclude that USPTO rules and procedures induce informed examiners to grant many of these weak patents. In contrast, we find that examiners have a large effect on the nature of the patents they grant and that their behavior does not appear to be significantly constrained by the existing USPTO rules and procedure.

Previous research has established that although patent examiners are charged with a uniform mandate, in practice examiners have a fair amount of discretion, and this discretion appears to translate into substantial variation in the decisions different examiners make on otherwise similar patent applications (Cockburn et al., 2003; Lichtman, 2004; Lemley and Sampat, 2010, 2012). Our research builds on the basic strategy used in Cockburn et al (2003) - we compute examiner effects with newer methodology to generate reliable magnitudes, analyze a different setting (NPEs and patent assertion), and add the extra step of correlating examiner effects with detailed data on the application process that was not available to earlier researchers.

Our methodological approach is adapted from the teacher value-added literature (Kane and Staiger (2008), Chetty et al. (2014)). Our paper is the first, to our knowledge, to apply this methdology in the setting of intellectual property. In addition, we discuss methodological issues in the case of rare and binary outcomes. Our methodology is also related to the examiner or judge leniency instrumental variables approach, used for instance by Sampat and Williams (2015) in the context of patent examination to analyze follow-on innovation. We discuss this link in greater detail in Appendix D. An additional methodological contribution here is to show that the validity of the instrument requires the exclusion of continuation applications.Our focus is on the "intensive margin", namely the effect an examiner has on a patent conditional on grant. But we also show results in the appendix using the examiner leniency methodology, which encompasses both intensive and extensive margins. The motivation for focusing on the intensive margin is that it allows us to move

beyond just the effect of approval on outcomes and analyze changes in the nature of the patent.

Finally, our paper indirectly contributes to the vast patent statistics literature. Key papers in this area include Pakes (1986) and Hall, Jaffe, and Trajtenberg (2001), which look at patent value in the context of optimal renewal decisions and forward citations, respectively. Other papers in the literature investigate the role of scope (Lerner (1994)) and the number of independent claims as markers of patent value. Our methodology introduces a new set of randomly assigned patent statistics based on the projection of examiner tendencies onto the patents that they grant, which avoids endogeneity issues in some of the aforementioned widely-used statistics.

3 Data and Descriptive Statistics

3.1 Data Construction

We combine several data sources for our analysis. The first is patent data on both granted and ungranted applications. The American Inventors Protection Act of 1999 stipulated that patent applications filed on or after November 29, 2000 would be published eighteen months after the filing date. Prior to the legislation, the USPTO only published application information for granted patents. We make use of the Patent Examination Research Dataset compiled by the Chief Economist Office at the USPTO, and discussed in Graham, Marco, and Miller (2015). This covers applications in the period from November 2000 through December 2014. In our data, we see about 12,000 unique examiner names, and 650 unique art unit codes. We also make use of the patent data collected by Lee Fleming and co-authors [18], which contains organized extracts of the patent information made available by Google, and process the data using their name matching algorithm to assign inventor and assignee identifiers to each patent. Finally, we also process the aforementioned organized extracts from Google to obtain additional information on non-patent citations and IPC/CPC patent classifications.

In addition to basic patent application and grant information, we also make use of application-level blocking action data collected by Frakes and Wasserman (2015), which they have generously shared with us. As discussed earlier, data on blocking actions will give us added insight into the nature of patents granted by each examiner, and has, to our knowledge, not been used before Frakes and Wasserman (2015).[19] The Frakes and Wasserman data covers all applications filed after January 2001 and that are disposed by July 2012,

[18] The data is from the Patent Database Search Tool available at http://rosencrantz.berkeley.edu/

[19] This data is also collected by IP service websites such as Juristat, an indication of the important role of examiners in the patent application process.

which is only part of the sample available through Patent Examination Research Dataset. The applications in this set, which will serve as our core analysis sample, covers about 1.9 million disposed applications and 1.27 million granted patents (versus 2.68 million disposed applications and 1.82 million granted patents in the Patent Examination Research Dataset). We perform various robustness checks on the unrestricted sample.

Second, we collect data on NPE patent portfolios. The starting point of our exploration centered around the patent portfolio of Intellectual Ventures, a prominent NPE. Intellectual Ventures holds an estimated 25-30k US patents, and released a list of around 20k on their website in November of 2013. As we note later, some of these patents are applied for by Intellectual Ventures, and we exclude them in our formal analysis. To augment this data, we constructed NPE lists using two data sources on NPE names. Our main source is a set of entity names shared with us by RPX, a defensive patent aggregator that monitors NPE activity. The list of entities identified by RPX is similar to the one used in Cohen et al. (2015), because RPX recently purchased PatentFreedom, the source of their data. However, an important point to note is that we exclude individuals and universities from the list used by RPX in their reports.[20] A second source for NPE names comes from Cotropia et al (2014).[21] The limitation to this approach is that their paper only classifies plaintiffs of IP lawsuits in the years 2010 and 2012, so it would miss entities that are only active in other years. Both lists are then matched to assignee name in the Patent Assignment Dataset recently constructed by Alan Marco, Amanda Myers, and their colleagues at the USPTO and discussed in Marco et al. (2015).

Third, we use a combination of sources to look at lawsuits involving patents. This includes data from LexMachina, Darts IP, and RPX, three organizations that track intellectual property lawsuits and NPE activity. Our NPE lawsuit indicator is also derived from RPX classifications. This follows the approach taken by Cohen et al. (2015). Each of these sources tracks intellectual property lawsuits since 2000, which is ideal for our analysis because we start observing abandoned applications in November of 2000. In addition, LexMachina tracks patents that are eventually appealed through inter partes review, which was instituted as part of the America Invents Act.

Fourth, we collect additional characteristics on each published application and granted patent. This includes the text in the claims section of both applications and granted patents, backwards citations by who added it [22], assignment information by type of event, and maintenance fee payment information. For

[20] We additionally exclude Wisconsin Alumni Research Corporation and Children's Medical Center Corporation, which RPX classify as companies.

[21] Publicly available at npedata.com

[22] Either the examiner or the applicant. This data is available back to 2001.

performing robustness checks, we also collect additional patent-level indicators for special types of patents, including ones purchased by regular firms, assigned to universities, and listed in the FDA Orange Book.

3.2 Summary Statistics

In this subsection, we provide some summary statistics based on the data sources described above, in order to offer a clearer picture of the patent examination process at the USPTO and the types of patents litigated and purchased by NPEs. Throughout the section, we will exclude continuation applications from our summary stats, as our core analysis will do the same in order to maintain random assignment. We include summary statistics based on all applications in the Appendix.

3.2.1 Art Units and Examiners

We start by describing various patterns in our USPTO examination data. We present the basic structure of art units at the USPTO and document the amount of variation across art units and examiners for various variables, particularly NPE purchases and indicators for various types of examiner behavior.

Overall, there are 670 art units in our dataset. 559 of these art units have at least one granted patent ending up in an NPE portfolio (around 83%). Art units exhibit significant variation in attributes. There is a long right tail for most statistics, and with the level of NPE activity having the highest spread. The top art units in terms of NPE Patent Rate are in the areas of Data Processing (3621, 3622, 3685, 3688) and Communications (2631, 2637). Detailed summary statistics at the art unit level are shown in Appendix A. In this section, we will focus on statistics at the art unit by application filing year level, which is the level at which there is random assignment, and patent examiners, who create the variation used in our study.

First, we look at the statistics at the art unit by year level, which is the level at which applications are randomly assigned to examiners. There are NPE purchased patents in about 50% of the 7,156 art unit by year data points. The results are displayed in Table 1. We also look at blocking actions on patent applications within each art unit and year, including appeals to requirements listed in Section 101 (patentability and utility), Section 102(a) (novelty), Section 103(a) (non-obvious), and Section 112(b) (clarity in claims text). The use of 103(a) and 112(b) appear to be somewhat frequent, although 112(b) has a higher degree of variation across art units.

Table 1: Art unit by year level statistics

Art unit by year level statistics	Median	Mean	Standard Deviation	Max
Examiners	15	15.5	15.0	155
Cases Processed	142	331	461	4510
Patents Granted	69	221	328	2538
NPE Patents	0	2.26	6.64	129
Grant Rate	0.64	0.62	0.20	1
NPE Patent Rate	0.001	0.011	0.020	0.333
Use of Section 101	0.029	0.097	0.132	0.667
Use of Section 102(a)	0.014	0.020	0.022	0.2
Use of Section 103(a)	0.47	0.45	0.20	0.96
Use of Section 112(b)	0.20	0.21	0.13	0.77

Notes: This table summarizes the statistics computed at the art unit by application year level, weighting each art unit by year equally. NPE patents are identified using the RPX list of entities, using the routine discussed in Section 3.1. Grant rate refers to the fraction of disposed applications (as of December 2014) that are granted in a given art unit by year. The NPE Patent rate refers to the fraction of granted patents in an art unit by year that have been purchased by NPEs (as of December 2014). "Use of Section" variables refer to the sections in Chapter 35 of US Code (discussed in Section 2.1) that are cited in examiner blocking actions.

Next, we show the statistics associated with patent examiners in Table 2. This gives us a basic idea of the amount of variation across examiners. There are 12,032 unique examiners, and 98.7% of the examiners are at one point assigned to an art unit with some NPE activity. There is significant variation amongst examiners in all of the computed statistics, with the NPE rate having the highest spread relative to the mean. A major part of this variation comes from across art unit differences, as discussed above. In addition, there is also noise in the sets of applications that examiners receive. In the remained of the paper, we use various approaches to eliminate the noise and identify the signal variance coming from examiners: we find that it is still substantial.

Table 2: Examiner level summary statistics

Examiner Level Statistics	Median	Mean	Standard Deviation	Max
Cases Processed	119	190	215	1600
Art Units	2	1.80	0.96	7
Years in Data	7	6.35	3.19	12
Patent Grant Rate	0.59	0.57	0.23	1
NPE Patent Rate	0	0.011	0.023	0.5
Use of Section 101	0.029	0.092	0.129	0.738
Use of Section 102(a)	0.010	0.019	0.025	0.240
Use of Section 103(a)	0.49	0.47	0.17	0.95
Use of Section 112(b)	0.19	0.21	0.13	0.76

Notes: This table summarizes the statistics at the examiner level, weighting each examiner equally. Rate computations are restricted to examiners with more than 20 cases in the data. The variables here are computed in the same way as the ones shown in Table 1.

3.2.2 NPE Purchased Patents

In this part, we explore these characteristics of patents purchased by NPEs, in order to understand NPE purchasing tendencies and the technological areas and industries in which they are the most active. Once again, we exclude continuation applications from our main analysis, and report overall numbers in the Appendix. The main findings here are that NPE-purchased patents are predominantly concentrated in computer hardware and software, and NPE patents have very different pre- and post- examination features relative to the average patent.

We start by looking at the patent classes in which NPEs are the most active. In Table 3, we display aggregate counts of all patents held by NPEs, at the NBER category and sub-category level (Hall et al (2001)). This confirms the general view that many of the patents are concentrated in IT, although there are a small number of patents in other fields. As a comparison, we have also computed similar statistics for non-NPE asserted patents in Table 22, and found a very different composition of categories. Of course, this evidence alone does not help differentiate between the two theories on NPE behavior, as it could be that the most resource-constrained firms are in IT.

Table 3: NPE Patent Holdings by NBER Technology Category

Panel A: Primary Technology Categories

NBER ID	Category Name	Patents
2	Computers & Communications	7,587
4	Electrical & Electronic	3,050
5	Mechanical	684
-	*New Classes (since 2001)*	629
6	Chemical	415
1	Others	392
3	Drugs & Medical	198

Panel B: Secondary Technology Categories

NBER ID	Subcategory Name	Patents
21	Communications	3,368
22	Computer Hardware & Software	2,868
46	Semiconductor Devices	1,147
24	Information Storage	900
-	*New Classes (since 2001)*	629
41	Electrical Devices	496
23	Computer Peripherials	451
49	Miscellaneous-Elec	411
45	Power Systems	359
42	Electrical Lighting	352
54	Optics	294
19	Miscellaneous-chemical	282
69	Miscellaneous-Others	231
59	Miscellaneous-Mechanical	211
43	Measuring & Testing	159

Notes: the NPE patents are identified based on RPX classifications (see Section 3.1). We take the primary USPTO technology class for each patent, and then aggregate up to subcategory and category levels, using the NBER patent crosswalk provided by Hall et al. (2001).

Next, we look at the sources of patents owned by NPEs. To do this, we look up the original assignee on each NPE patent in our list. The results of this exercise are shown in Table 4. The plurality of the roughly 35k NPE owned non-continuation patents are initially unassigned. In addition, some patents are initially assigned to entities associated with Intellectual Ventures, such as The Invention Science Fund and Searete. Since the focus of our paper is on patent acquisition, we will not count these patents as NPE-purchased for our formal analysis.[23] Another interesting point to note is that while some of these patents were originally assigned to companies that have gone through bankruptcy (Eastman Kodak and Polaroid Corporation), others were initially owned by firms that have not, such as GE, Lucent, and Micron. As we discover later in

[23] As detailed in the appendix, the rest of the portfolio is constructed excluding patents that are initially assigned to NPEs on our lists. On a different but related note, we verify that Intellectual Ventures patent applications appear to be randomly assigned to patent examiners, which is consistent with our identification assumption.

our analysis with initial-assignee fixed effects, firms with many granted patents also appear to selectively sell their vaguer and more obvious patents to NPEs, so our effect is only partly driven by cross-firm differences and individual inventors.

Table 4: Aggregate NPE portfolio patent counts by initial assignee

Initial Assignee	# Patents	Initial Assignee	# Patents
(Unassigned)	1,434	Harris Corporation	104
Eastman Kodak Company	806	MIPS Technologies, Inc.	93
Micron Technology, Inc.	381	NEC Corporation	84
Telefonaktiebolaget L M Ericsson	334	MOSAID Technologies	83
Nokia Corporation	322	DaimlerChrysler AG	79
Koninklijke Philips Electronics	302	NEC LCD Technologies	76
Matsushita Electric Industrial Co, Ltd.	256	Empire Technology Development	71
NXP B.V.	230	Lucent Technologies Inc.	70
Panasonic Corporation	195	Raytheon Company	69
American Express Travel Related	167	University of California	68
Global OLED Technology LLC	146	DPHI Acquisitions, Inc.	67
Hynix Semiconductor Inc.	127	Lite-On Technology Corp.	66
Industrial Technology Research Institute	125	Cypress Semiconductor Corporation	65
Virginia Tech Intellectual Properties	120	*Searete LLC*	64
The Invention Science Fund I, LLC	111	LG Electronics Inc.	60

Notes: We use assignee names from the Lee Fleming database. Italics indicate entities associated with Intellectual Ventures (patents granted directly to them). For entities other than Intellectual Ventures, we include only patents that are assigned to them, but not as part of an "employer assignment," as classified in the USPTO assignment dataset.

Finally, we compare basic statistics of NPE-purchased patents with other patents in their technological cohort. The main takeaway from this exercise is that patents purchased by NPEs have different initial and post-examination characteristics relative to other patents in the same cohort.

To construct a crude control group, we take all granted patents that had the same application year and art unit assignment as an NPE patent, and re-weight the control patents based on the number of NPE patents in their cohort. The statistics we compute include ex-ante characteristics, such as the number of independent claims in the initial application, examination characteristics, such as the frequency of blocking actions and changes in words in the primary claim, and ex-post outcomes. such as forward citations and re-assignment events.The results are shown in Table 5. The statistics presented here are robust to the NPE list we use.[24]

[24]Including removing all Intellectual Ventures patents to look at the portfolio of the rest of the NPEs and looking at the Kesan list of patent holding companies and large aggregators.

Table 5: Summary Statistics: NPE vs. Non-NPE Patents

Statistic	NPE Patents	Control Patents	Difference
Initial Application Characteristics			
Independent claims at Application	5.00	4.07	0.92***
Examination Outcomes			
Examiner leave-one-out allowance rate	0.766	0.759	0.006***
Number of Examiner Rejections	1.93	1.87	0.06***
Change in # Independent Claims	-1.58	-0.83	-0.75***
Use of Section 101 - Lack of utility or eligibility	0.085	0.090	-0.005**
Use of Section 102 - Prior art exists	0.021	0.019	0.002
Use of Section 103 - Obvious invention	0.48	0.47	0.01
Use of Section 112(a) - Improper Disclosure	0.059	0.051	0.008***
Use of Section 112(b) - Vague Claims	0.18	0.17	0.012***
Post-grant Outcomes			
Citations	8.86	6.71279	2.14***
Number of re-assignments	1.67	0.36	1.31***
Date of latest re-assignment relative to grant (days)	1245	881	364***
Involved in IP Litigation	0.035	0.005	0.030***
Inter Partes Review Instituted	0.0042	0.0002	0.0040***

Notes: For each NPE patents, identified using the RPX list as discussed in Section 3.1, we take the set of non-NPE patents in the same art unit and application year cohort, and re-weight them so that there is the equivalent of one control for each NPE patent. Then, we compute average characteristics and outcome values for each group and also compute the difference in means. Examiner leave-one-out allowance rate for each patent is computed based on Equation 1. ** p-value $<$ 0.1, ** p-value $<$ 0.05, *** p-value $<$ 0.01*

The results show that there are significant differences between NPE-purchased patents and others in the same year and technological cohort. In terms of characteristics at the start of application, the set of NPE-purchased patents have a higher number of independent claims at application. This suggests that the applications aim to claim a broad scope, but as we see in the average number of changes in independent claims, many of these are removed before final grant.

Next, the examination indicators provide a mixed message. NPE-purchased patents come from examiners with a slightly higher leave-one-out grant rate, although this is a noisy measure of true examiner approval rates. In addition, these patents tend to have more examiner rejections ("final" and "non-final") and more independent claims removed. There is also significantly more 112(b) usage on the NPE-purchased patents, although this could suggest either a vaguer application to begin with or examination by more stringent examiners.

Finally, in terms of eventual outcomes, NPE-purchased patents are more likely to be cited. Citations have been the primary measure of patent value in the literature, but as pointed out earlier, part of the differ-

ence here could be from an exposure effect.[25] NPE-purchased patents are also more frequently re-assigned, although this is partly mechanical, as we construct this group based on re-assignment data. Furthermore, conditional on re-assignment, NPE patents are more likely to be purchased later in a patent term than other purchased patents, a result similar in nature to Love (2013), which finds that NPEs assert patents late in their term. These patents are also more than seven times more likely to end up in litigation, and twenty times more likely to be instituted under the new Inter Partes Review system.[26] These post-grant outcomes are where the two groups of patents differ most, but some of this is unsurprising given that the groups are defined based on a post-grant outcome, namely purchase by an NPE.

In general, the simple comparisons are hard to interpret, as they reflect both the nature of the original application and the changes made by the patent examiner, which sometimes create offsetting differences. Of course, we could look to improve the matching criteria to create a better control group, but this is not always feasible because of the limited number of patents within each art unit cross year cohort and hard-to-measure quantities such as vagueness of the initial claims language. Instead, as we explain in our methodology, we indirectly control for differences across applications by looking for consistency of examiner differences over time, leveraging random allocation.

3.3 Examiner Variation

Next, we present evidence on random assignment of patent application to examiners and we show preliminary evidence on the heterogeneous effects examiners have on patent outcomes.

To show preliminary evidence on the heterogeneous effects examiners have on patent outcomes, we rank examiners by their allowance rate and blocking action usage, and look for variation in the fraction of litigated and NPE purchased patents across the quartiles. Specifically, for each granted patent, we compute an examiner leave-one-out allowance rate, excluding that particular case from the calculation.

$$Leave-one-out\ Allowance\ Rate_i = \frac{\sum_{k \neq i, k \in J(i)} GrantDecision_k}{N_{J(i)} - 1} \tag{1}$$

where i is the granted patent of interest, $J(i)$ is the set of (granted and ungranted) patent applications processed by the examiner who granted patent i, and $N_{J(i)}$ is the number of applications in $J(i)$. A high

[25] One particular way this could work is through the threat of "inequitable conduct" rulings. A patent holder can have his patent ruled unenforceable if he can be shown to have deliberately not cited prior art in his patent application. An applicant can avoid this risk by citing all widely known patents that are involved in lawsuits or are in the news.

[26] Instituted cases represent appeals where the petitioner has a "reasonable likelihood" of prevailing. We use this outcome as there have been very few cases with final verdicts under this new system put in place after the America Invents Act.

value for this variable indicates that patent i was granted by an examiner who grants a high share of the patent applications they process. We use a similar approach to compute leave-one-out blocking action usage.

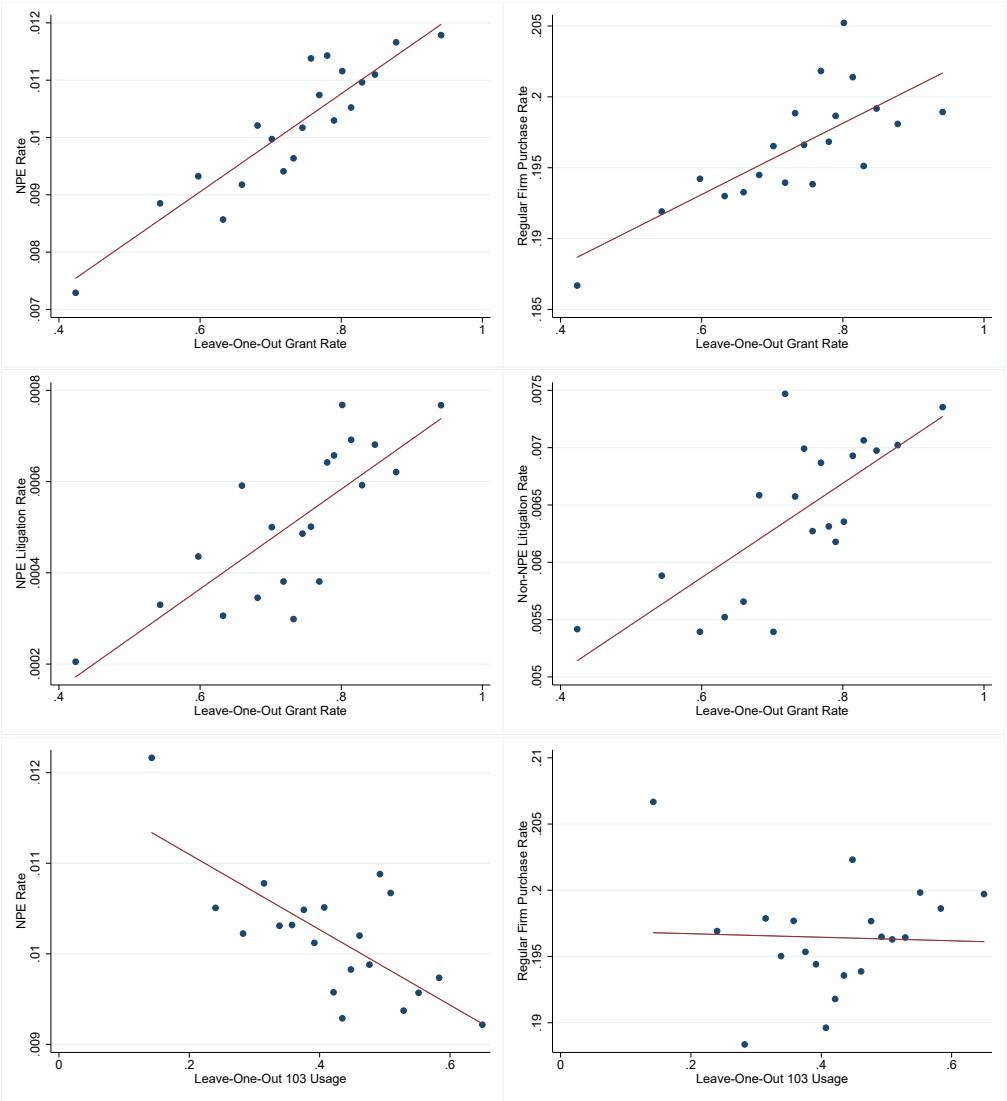
Figure 1 shows the relationship between examiner leave-one-out means over their whole career and subsequent granted patent outcomes, such as NPE purchase, non-NPE purchase, NPE-litigation and non-NPE litigation. The NPE purchase outcome shows a strong positive association with examiner leave-one-out grant rate, a relationship that is much weaker for the non-NPE purchase outcome. The relationships with examiner leav-eone-out grant rate are also strong for NPE-litigation and non-NPE litigation. In addition, NPE purchase is negatively associated with examiner usage of 103(a) blocking actions, whereas non-NPE purchase is uncorrelated.

One issue with pooling data from an examiner's entire career is that we could be capturing secular technological variation over time. To check for this, we have also compute leave-one-out means at the art unit by year level, which is the approach taken by Sampat and Williams (2014). The results are reported in Figure 5 in Appendix A, and provide further evidence that patents granted by examiners with a higher allowance rate are more likely to be purchased by NPEs. In fact, the weak positive relationship between non-NPE purchase and grant rate disappears under this specification. Of course, neither piece of preliminary evidence addresses noise in the grant rate calculations and the consistency of examiner rankings across years, which we will address through our formal methodology in Section 4.

In order to test the validity of the random assignment of examiner, we conduct placebo tests by regressing a number of outcomes that were determined before the time of patent application on the leniency of the examiner in charge of this patent application. The outcomes we consider are either defined at the inventor level or the patent level. Specifically, we consider three outcomes that are related to the "value" of a patent: the number of patents granted to an inventor in previous years (as a proxy for inventor quality), the number of independent claims at the time of publication of the application (a widely-used proxy for the quality of a patent) and the number of words in the first independent claim (another commonly used proxy for patent quality). In addition, we consider two outcomes that speak directly to the exclusion restriction: the number of patents of the inventor that were purchased by or issue to NPEs in previous years, and the number of patents of the inventor that were litgiated in previous years. We find no statistically significant relationship between these outcomes and the leave-one-out examiner allowance rate in our preferred sample, which excludes continuation application and repeated inventor-examiner pairs and adjusts the leave-one-out

allowance rate for docketing time patterns.[27] However, if we keep the full sample of applications (which to the best of our knowledge is the current practice in the literature using patent examiner instruments), the placebo tests fails. Restricting the sample as we do is therefore key to preserve random assignment. These various results are reported in Appendix B.

Figure 1: Preliminary Evidence on the Relationship Between Examiner Tendencies and Patent Outcomes



Notes: Results are computed on the baseline sample (non-continuation granted patents covered by Frakes and Wasserman). Examiner effects are computed as a leave-one-out mean on their entire history. Outcome data is at the granted patent level (fraction conditional on grant). Results are computed absorbing art unit fixed effects, and restricting to examiners who review at least 50 applications.

Another key element of the research design is that our specifications are carried out conditional on the

[27] Continuation applications are not randomly assigned and should therefore be excluded. The rational for also excluding repeated inventor-examiner pairs (within the same artunit-year-class) is that some continuation applications may not be properly recorded. Applications that were submitted at the same time are likely to be assigned to the same examiner in a batch, therefore we compute the leave-one-out allowance rate only based on applications that were docketed in different months.

patent being granted: if an application is not granted, by definition we could not observe NPE purchase, litigation, etc... Therefore, the variation in the rate of NPE patents observed across examiners could be due to variation in the underlying quality of the pool of granted patents across examiners. If NPEs tend to purchase patents from a set of examiners with a pool of granted patents of high quality, then the interpretation that NPEs target "weak patents" would be misplaced. The underlying quality of a patent is not observed, but we find that the examiners who have a high rate of NPE purchases or litigation also have high allowance rates (as documented above) and do no narrow or clarify the claims of the patents as much as other examiners (as documented in Section5). In other words, average patent quality should be lower in the pool of patent granted by examiners with a high NPE rate. We test this hypothesis in Appendix B and indeed find support for it: inventors who get granted patents by examiner with high allowance rates tend to be less accomplished (i.e. they have been granted fewer patents in previous years). Therefore, our results across examiners are not driven by a selection effect on quality. Appendix B provides a more in-depth discussion of selection effects in our setting.

In the remainder of the paper, we first use a shrinkage methodology to estimate the causal effect of examiners on the probability that a granted patent becomes part of a NPE portfolio. We find very large effects, suggesting that the rate of patents in NPE portfolios could be reduced by over 50% by re-designing the patent examination process. Second, we systematically investigate the characteristics of examiners with a high NPE effect, studying all types of rejection during prosecution, changes to the claims and propensity to cite prior art. The results confirm the suggestive findings presented in this section: a large share of the patents that are part of NPEs' portfolio are "weak" patents, which were granted by lenient examiners and appear to be vaguely worded. The formal shrinkage and regression frameworks we use allow us to establish the statistical robustness of the preliminary findings here.

4 Estimating Examiner Causal Effects

In this section, we describe our overall research design, discuss threats to identification, and then show formal methodology and results on computing examiner causal effects.

4.1 Research Design Overview

Our research design starts from the fact that patent application are conditionally randomly assigned to examiners. Indeed, within an art unit, applications that are not continuation applications are randomly assigned to examiners.[28] The intuition of our research design is to look at the share of NPE patent (patents that are eventually purchased or asserted by NPEs) in the portfolio of granted patents of an examiner, compared with the share of NPE patents of other examiners working in the same year and art unit. If there is substantial heterogeneity across examiners, then we can conclude that the patent examination process plays an important role in the activities of NPEs and we can look at the characteristics of examiners who have a high share of NPE patents to learn about the important features of the patent examination process and infer the patent acquisition strategy of NPEs.

The first step of our analysis is to test for "excess variance" from examiners, i.e. do NPE shares differ significantly and consistently across examiners in the same art unit? To do so, we employ a shrinkage methodology filtering out the noise in the data. Intuitively, if an examiner granted only two or three patents, his estimated share of NPE patents will be extremely noisy. We address this in various ways, following the teacher value-added methodology in Kane and Staiger (2008) and then using a novel non-parametric Empirical Bayes methodology we introduce as a robustness check. This analysis is reported in the remainder of this section.

The second step, the focus of Section , is to estimate the correlation between the NPE examiner effect estimated in the first step and other examiner effects we estimate following a similar methodology (examiner allowance effect, examiner propensity to cite prior art, examiner propensity to use certain kinds of blocking actions, etc.) in order to learn about the mechanism at play.

The analogy between our setting and the teacher valued-added framework of Kane and Staiger (2008) is as follows: we treat examiners as teachers, applications within a given year by art unit as the student cohort, and we measure patent-level outcomes both at the time of grant, such as claims attributes, and eventual outcomes, such as ending up in an NPE portfolio and being asserted in IP litigation. This is analogous to the approach taken in various papers in the teacher effects literature, such as Chetty, Friedman, and Rockoff (2012), which uses test scores as the short-term indicator of teacher quality and wages at age 28 as the

[28] This fact has been discussed in Lemley and Sampat (2012). We have conducted our own interviews with patent examiners and the novelty of our approach is to exclude continuation applications. This is a very important adjustment for NPEs' portfolio, which have a high share of continuation patents, as documented in Section 3.

long-term outcome of interest. In their paper, the nature of the test scores is uncertain and higher wages are indicative of better outcomes. Here, we take the reverse approach: we use the short-term indicators such as 103(a) and 112(b) usage to better understand the nature of the long-term outcome, namely NPE acquisition.

4.2 Computing Examiner Causal Effects

The first step in our methodology is to compute an examiner causal effect. For our core results, we adapt methodology from the teacher value-added literature. This framework allows us to estimate an unbiased predictor of the effect that a given examiner has on a particular outcome for a given patent. This is done by taking the examiner fixed effect in a given year and shrinking it by a factor equal to the correlation of effects across years divided by the level of idiosyncratic noise in the outcomes, in order to extract the signal component in the data. These compute effects can then be treated as causal examiner effects, because of the random assignment structure detailed above.

More precisely, our empirical framework is as follows:

$$T_{ijt} = X_{ijt}\beta + a_{ut} + v_{ijt} \tag{2}$$

$$v_{ijt} = \mu_j + \theta_{jt} + \epsilon_{ijt}$$

where i indexes the patent, j the examiner, u the art unit and t the year. X_{ijt} represents controls for patent application characteristics, a_{ut} represents art unit by year fixed effects, to control for differences at the level of random assignment. μ_j is the examiner causal effect of interest (assumed to be constant over time), θ_{jt} represents an idiosyncratic examiner by cohort effect,[29] and ϵ_{ijt} is an idiosyncratic patent effect. We run Equation 2 using OLS to obtain the residuals. We then take a series of steps to compute examiner effect estimates using these residuals.

First, we compute the within examiner-year variance in v_{ijt} to obtain an estimate of the variance of the idiosyncratic component:

$$\hat{\sigma}^2_\epsilon = Var(v_{ijt} - \bar{v}_{jt})$$

where$\bar{v}_{jt} = \frac{1}{n_{jt}} \sum_{i=1}^{n_{jt}} v_{ijt}$, an average of residuals within a year. This step serves as an important for later adjustments for background noise. In our context, this idiosyncratic component has significant variance, as

[29]This could result from the examiner being more familiar with certain technology cohorts or just random fluctuations in examiner behavior over time. The key to using this framework is that θshould not be serially correlated, which would alter the interpretation of the calculated signal variance computed in Equation 3.

there are many features of patent applications that we are unable to observe or measure.

Second, we use the covariance between the average residual in an examiner's portfolio in year t and year $t+1$ as an estimate of the variance in the examiner component:

$$\hat{\sigma}^2_{\mu} = cov(\bar{v}_{jt}, \bar{v}_{j(t+1)}) \tag{3}$$

where the covariance calculation is weighted by the number of patents granted by each examiner (n_{jt}). This variance component - which we refer to as the "signal variance" of the examiners - is a measure of the variation in examiner effects across examiners. Intuitively, if all the true examiner μ_j's are close to zero relative to the size of the outcome, we may still get variation in the calculated annual average residuals $\bar{v}_{jt}$ from the idiosyncratic error draws, but our methodology will then pick up a very low signal covariance across years, because the idiosyncratic components are uncorrelated.

Third, the variance of the examiner-cohort idiosyncratic component is estimated as the remainder:

$$\hat{\sigma}^2_{\theta} = Var(v_{ijt}) - \hat{\sigma}^2_{\mu} - \hat{\sigma}^2_{\epsilon}$$

Next, we form a weighted average of the average residuals for each examiner in each year ($\bar{v}_{jt}$) that is a minimum variance *unbiased* estimate of μ_j for each examiner. Data from each year is weighted by its precision, with years in which the examiner granted more patents received more weight:

$$\bar{v_j} = \sum_t w_{jt} \bar{v}_{jt}$$

where

$$w_{jt} = \frac{h_{jt}}{\sum_t h_{jt}}$$

$$h_{jt} = \frac{1}{\hat{\sigma}^2_{\theta} + \frac{\hat{\sigma}^2_{\epsilon}}{n_{jt}}}$$

The last step is to construct the "empirical Bayes estimate" of each examiner's effect by multiplying the weighted average of examiner-year residuals by a shrinkage factor:

$$ExaminerEffect_j = \bar{v_j} \frac{\hat{\sigma}^2_{\mu}}{Var(\bar{v_j})}$$

where $Var(\bar{v_j}) = \hat{\sigma}^2_{\mu} + (\sum_t h_{jt})^{-1}$.

The shrinkage factor is the ratio of signal variance to total variance, and is different across examiners based

on cases examined. This final quantity, the examiner effect, has two desirable properties (see Appendix C for formal demonstrations of these properties). First, it has an empirical Bayes interpretation as the Bayesian estimate of the examiner effect, with a normal prior distribution centered around zero with variance equal to the signal variance calculated in the second step above.

Second, it also has a frequentist interpretation: the shrinkage factor represents the coefficient of a hypothetical regression of μ_j on $\bar{v}_j$. The regression coefficient is a ratio of the covariance of the two expressions, in this case $\hat{\sigma}^2_\mu$ because the other parts of $\bar{v}_j$ are uncorrelated noise, divided by the variance of the independent variable. Therefore, the estimated $ExaminerEffect_j$, although biased, has the same scale as the true examiner effect μ_j in this model, causing it to have a lower mean-squared forecast error relative to $\bar{v}_j$. We believe that the distribution of $ExaminerEffect_j$ better captures the scale of true examiner effect distribution relative to using $\bar{v}_j$, and therefore report results using the former in our main findings.[30] We later verify using a split sample approach that our examiner effect has this desired property.

One shortcoming of using this approach in our setting is that, as explained in Kane and Staiger (2008), a Bayesian interpretation of the results requires the assumption that the residuals of the regression are normally distributed, which may not be appropriate in our setting given that the outcome variable is binary. Nonetheless, we include controls, in particular art unit and year fixed effects, therefore the normality assumption is not necessarily violated. In addition, to address this this concern we also estimate the distribution of examiner effects using a binomial model, in the spirit of Ellison and Swanson (2010). The results are consistent across models, as detailed in Section 3.4 below.

4.3 Results

Here, we report the main results, computed using the Kane and Staiger framework. We report the distribution of causal examiner effects for our main outcomes of interest, and also report results on the signal variance parameter estimated as part of the methodology.

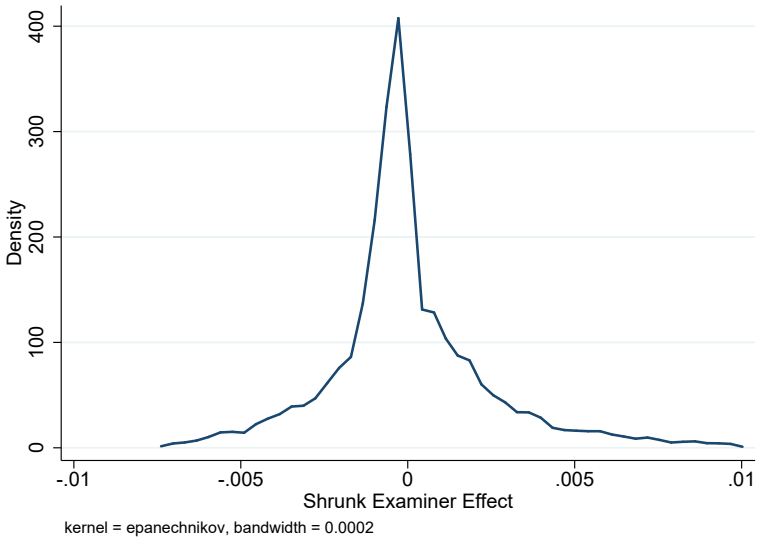
First, we plot the distribution of examiner effects on NPE purchase in Figure 2, keeping all examiners and weighting by the number of cases processed.[31] The baseline rate of NPE purchase is around 0.01, so the spread in the distribution accounts for a sizable fraction of this baseline rate.

To be more specific, we display in Table 6 provides various measures of the magnitude of these effects,

[30] Using the distribution of $\bar{v}_j$ would leave to a higher variance.
[31] The results without weighting are very similar.

expressed as a percentage of the share of NPE patents in the patent population. For each of these results, we also compute 95% confidence intervals using a bootstrap procedure, and report it in brackets in the line below the computed number. The first row reports the signal standard deviation, estimated as explained in section 4.2, which amounts to a staggering 52.95% of the NPE share. Under the Bayesian interpretation, this is the empirical Bayes estimate of the standard deviation of the distribution of true examiner effects $\hat{\sigma}_{\mu}$. Next, the standard deviation of the computed shrunk examiner effects distribution is around 24.33% of the baseline rate. This number is distinct from the signal standard deviation, as it reflects the distribution of posterior means from updating the prior for each examiner. The interquartile range of the shrunk examiner effect distribution is less than the standard deviation, still a sizable 7.25% of the NPE share. This suggests that the shrunk examiner effects distribution has excess density in the tails relative to a normal distribution. Finally, if all examiners above the 75th percentile of the distribution (i.e. examiners with an unusually high share of NPE patents in their granted patents) were replaced with examiners located exactly at the 75th percentile of the distribution, then the share of granted patents would decrease by 6.15%. Overall, these results indicate that patent examiners have a large causal effect on the probability that a granted patent becomes part of an NPE portfolio.

Figure 2: NPE Purchase Examiner Effect Distribution



Notes: We run a kernel density plot of the distribution of shrunk examiner effects across all examiners, weighted by the number of granted patents for each examiner. The shrunk examiner effects are calculated using the methodology described in Section 4.2.

Table 6: NPE Purchase Examiner Effect Distribution Characteristics

Measure of Examiner NPE Effect	Percentage of Sample NPE share
Signal standard deviation ($\hat{\sigma}_\mu$)	50.97% [33.7%,60.7%]
Standard deviation of Shrunk Examiner Effects	24.02% [12.26%,30.90%]
Difference between examiners at p25 and p75	6.86% [3.25%,8.51%]
Difference in distribution mean replacing examiners above p75 with examiners at p75	6.00% [3.02%,7.30%]

Notes: Bootstrapped 95% confidence intervals are displayed below the corresponding parameter estimate. The size of the parameters are normalized by the baseline rate of NPE-purchased patents in the sample (0.9%). The "difference in distribution mean" value is calculated by replacing shrunk examiner effects above the 75th percentile with the value at the 75th percentile, and then re-computing the average of the new distribution.

As a comparison, we also compute similar results for patents asserted by companies not on our NPE list. The baseline rate of these patents is 13.81%, and we display the results in Table 7. The major difference here is that the signal standard deviation is much lower, suggesting less heterogeneity across examiners in whether their granted patents end up being transferred on the IP market.

Table 7: Non-NPE Purchase Examiner Effects Distribution Characteristics

Measure of Examiner non-NPE Effect	Percentage of Sample NPE share
Signal standard deviation ($\hat{\sigma}_\mu$)	14.01% [10.70%,14.47%]
Standard deviation of Shrunk Examiner Effects	7.32% [5.05%,7.32%]
Difference between examiners at p25 and p75	6.97% [4.47%,6.97%]
Difference in distribution mean replacing examiners above p75 with examiners at p75	1.35% [0.95%,1.35%]

Notes: Comparable to Table 6. Bootstrapped 95% confidence intervals are displayed below the corresponding parameter estimate. The size of the parameters are normalized by the baseline rate of non-NPE purchased patents in the sample (13.81%).

Finally, we also compute similar results for patents asserted by companies not on our NPE list. The baseline rate of these patents is 0.43%, and we display the results in Table 8. The various measures are similar in magnitude to the ones reported for NPE-purchased patents.

Table 8: Non-NPE Litigated Examiner Effects Distribution Characteristics

Measure of Examiner NPE Effect	Percentage of Sample NPE share
Signal standard deviation ($\hat{\sigma}_\mu$)	62.84% [35.86%,74.45%]
Standard deviation of Shrunk Examiner Effects	25.62% [10.32%,32.00%]
Difference between examiners at p25 and p75	9.15% [3.14%,11.37%]
Difference in distribution mean replacing examiners above p75 with examiners at p75	6.96% [2.69%,8.53%]

Notes: Comparable to Table 6. Bootstrapped 95% confidence intervals are displayed below the corresponding parameter estimate. The size of the parameters are normalized by the baseline rate of non-NPE asserted patents in the sample (0.43%).

We also apply the same methodology on a variety of outcomes characterizing the examination process, such as the use of various blocking actions during prosecution, which we will make use of in the discussion on mechanisms. Examiner effect distribution graphs for some variables are included in Appendix D. In particular, we find that examiner effects are quite large for 103(a) and 112(b) blocking actions (24.0% and 46.5%, respectively), which are known to rely more on the subjective judgment of the examiner. They are also large for the practicing entity litigation outcome (signal standard deviation of 62.1%), which we will explore further in the mechanism section.

4.4 Robustness Checks

4.4.1 Robustness of Basic Methodology

In this part, we report a series of robustness checks of the basic results presented in Section 4.3. As shown on table 9, we find that the results are very similar across these different settings. Rows one to three report the

results in different samples. The sample that includes continuation applications exhibit larger NPE effects, as continuations are often assigned to the same examiner, and patents in the same family are likely to be purchased together. Rows four and five show that the results are very similar when considering other NPE lists. Finally, rows six and seven shows the robustness of the results to the inclusion of additional patent controls and assignee fixed effects. These specifications as well as additional robustness checks are described in greater detail in Section 5.3.

The fact that our core result is robust to the inclusion of assignee fixed effects is particularly interesting, because it shows that the effect is not driven by assignee-specific mechanisms. In particular, a possible explanation for our result that examiners have a large causal effect on NPE purchase could have been that more lenient examiners grant a bigger share of low-quality patents from failing firms, which then would sell their patents to NPEs. Table 9 shows that we can reject this hypothesis: the effect is not driven by such "supply" effects from assignees and, rather, results from "demand" effects from NPEs, which tend to purchase patents from a specific set of examiners even within the portfolio of a given assignee.

Table 9: Robustness checks for the NPE examiner effect distribution.

Specification	Signal SD	Shrunk Effect SD
Baseline	62.8%	25.3%
Expanded Sample (PERD)	52.9%	24.3%
Include Continuations (PERD)	77.2%	48.1%
Expanded Sample (1976-2015)	70.6%	24.8%
Excluding Intellectual Ventures Patents	71.2%	33.5%
Cortropia et al. NPE Patents	43.1%	15.5%
Additional Patent Controls	54.9%	24.8%
Assignee Fixed Effects	33.9%	13.2%

Notes: Expanded Sample (PERD) refers to the sample of all non-continuation applications in the Patent Examination Research Dataset, which covers applications filed between 2001 and 2014. Expanded Sample (1976-2015) refers to the sample of all non-continuation granted patents in that time period in the data available from Google. Excluding Intellectual Ventures Patents refers to excluding them from the analysis, and looking at patents owned by other NPEs. Additional patent controls refers to adding the number of independent claims at application and first inventor experience as controls in the basic estimating equation (Equation 2). All effects are normalized to the baseline NPE rate, which is different for the various specifications.

4.4.2 Empirical Bayes Count Model

To address the concern that the assumptions behind the Kane-Staiger (2008) methodology may not be satisfied, we use an alternative methodology. Specifically, we rely on an Empirical Bayes Beta-Binomial model, in the spirit of Ellison and Swanson (2010). This methodology allows us to flexibly model the data-generating process, using a binomial likelihood (which is non-parametric given that our outcomes are counts)

and a flexible prior on the distribution of examiner effects.

First, we aggregate data for each examiner j in year t and art unit a into the form (n_{jat}, r_{jat}), where n denotes the total number of granted patents for a given examiner and r the total number of NPE patents (or some other outcome) for this examiner. We then model the data generating process directly with a binomial likelihood on each data point: each examiner has some true probability p of having an NPE purchase a granted patent.

Next, we set up an empirical Bayes estimation of the prior of p. For each art unit in each year, we start with a flexible prior distribution:

$$p \sim Beta(\alpha, \beta)$$

We can then form an integrated likelihood in order to estimate the hyperparameters α and β of the prior distribution:

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

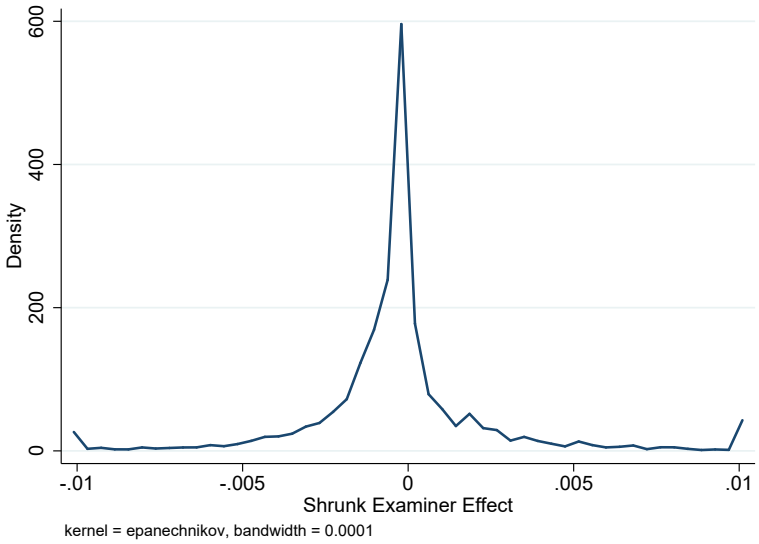
where the second step just conjugates the inside to integrate to one based on the probability density function of the Beta distribution. We then pool the data points of examiners in the same year and art unit, and estimate the hyperparameters via maximum likelihood. These parameters allow us to calculate the equivalent of the signal standard deviation in the basic framework.

To compute the posterior mean for each examiner, we take the estimated parameters $(\hat{\alpha}, \hat{\beta})$ and updated them using the data: the posterior distribution is given by $Beta(\hat{\alpha}+r, \hat{\beta}+n-r)$. The mean of this distribution, $\frac{\hat{\alpha}+r}{\hat{\alpha}+\hat{\beta}+n}$, gives us the posterior mean for the examiner. Intuitively, this procedure shrinks an examiner's NPE share towards the mean NPE share in the art unit, more so when the examiner has granted few patents.

Finally, we compute deviations from art unit prior mean for each examiner in a given year. We average these shrunk examiner effects across years, in order to make them comparable to the Kane Staiger shrunk examiner effects. The distribution of shrunk examiner effects computed using this procedure are shown on Figure 3. This distribution looks very similar to Figure 2, with a standard deviation around 44.4% of the baseline and an inter-quartile range of 8.8%, as opposed to the 24.0% and 6.9% reported in Table 6. The

main difference is that the distribution computed using the Empirical Bayes Count Model has thicker tails. The fact that the two distributions have similar characteristics is a re-assuring fact, since the Empirical Bayes Beta-Binomial count model imposes much weaker parametric restrictions on the prior distribution.[32] This shows that we can treat the results from the more tractable Kane-Staiger methodology as an approximation to a more flexible model of examiner effects.[33]

Figure 3: Distribution of NPE Examiner Effects (Empirical Bayes Count Model)



Notes: A kernel density plot of shrunk examiner effects computed using the Empirical Bayes Count Model described in Section 4.4.2. The figure is re-scaled to match Figure 2, and examiner effects above 0.01 and below -0.01 are moved to 0.01 and -0.01, respectively. The examiner effects are weighted by cases examined.

5 Which Examiner Characteristics Predict NPE Purchase and Litigation?

In this section, we present a series of results showing that the large NPE purchase and non-NPE litigation examiner causal effects, which we documented in Section 3, can be linked to some examiner behaviors observable in our data. Specifically, we show that the typical behavior of an examiner during prosecution is

[32] In fact, we can compare the consistency of the examiner rankings across the two approaches. We show evidence of this in Appendix D.

[33] The results presented in Figure 3 still rely on an untestable parametric assumption about the prior, which we have assumed to be a Beta distribution. Beta distributions are known to be flexible. An even more flexible prior would involve a mixture of Beta distributions, but the downside of such an approach is having to estimate additional hyperparameters with only twenty or so data points on examiners in a given art unit and year.

predictive of NPE purchase and non-NPE litigation. More specifically, we find that an examiner's tendency to use 103(a) and 112(b) blocking actions is *negatively* predictive of NPE outcomes. In other words, on average NPEs tend to purchase patents from "lenient" examiners who tend to allow incremental patents with vaguely-worded claims. Similar effects exist for non-NPE litigation, but they are less strong. In addition, we find evidence that this operates through changes to patent claims language between the initial application and the final grant. We do not find such results for non-NPE purchases, which confirms that NPE behave differently from practicing entities.

5.1 Predictive Regressions Methodology

Our core methodology in this section is to predict outcomes such as NPE purchase using causal examiner effects, which are computed using the methodology from the previous section. In other words, we predict patent outcomes using variation in the characteristics of the randomly assigned examiner, in order to shed light on the mechanism at play.

Formally, we compute causal examiner effects leaving out data from the patent application whose outcome we are trying to predict. This purges the regression of any mechanical correlations based on the fact that the nature of the patent itself will in general affect the behavior of the examiner. We then set up the following regression at the patent level:

$$NPE_{ij} = \beta \hat{E}_j + \epsilon_{ij}$$

where i indexes the patent and j indexes the examiner. $\hat{E}$ represents some vector of computed leave-one-out examiner effects, such as propensity to use 103(a) blocking actions or to edit the claims. This is a simple way of documenting which examiner behaviors correlate with NPE purchase.

We first run predictive regressions with a single examiner effect, in order to analyze each effect in isolation. Next, we also run some "horse race" regressions, including multiple examiner effect measures. [34]

From a methodological perspective, the leave-one-out examiner effect should be uncorrelated with the error term ϵ_{ij}, which captures unobserved characteristics of the patent. Note that selection effects might introduce a bias: the pool of granted patents allowed by each examiner varies by examiner characteristics. For example,

[34] As discussed in the results section, we find that once put in a horse-race with other variables such as allowance rates and changes in patent claims text, the blocking action variables tend to lose importance. Our interpretation of this result is that various blocking actions force different types of changes to the patent text, but our text-based measures are too crude to distinguish between these types of changes. Therefore, it makes more sense to trace the changes back to the source, and focus on the blocking actions for interpretability. We also run horse-race regressions using only blocking action examiner effects.

if examiners with a high rate of 103(a) blocking actions have higher quality patents on average, and NPEs purchase lower quality patents, then ϵ will be negatively correlated with $\hat{E}_{103a}$ because of this selection effect. We check for this by including additional controls based on inventor experience and application characteristics (similar to the "Additional Patent Control" specification in Table 3), and we find a negligible impact on the estimated coefficients.

As a robustness check, we use another strategy: we start from the full set of causal effects for each examiner, and correlate them with each other at the examiner level (e.g. correlate examiner NPE effect with examiner 103(a) effect). We can thus estimate the signal correlation between the various examiner effects. We carry out this procedure by splitting the sample in order to avoid picking up mechanical correlations induced by the measurement of the examiner causal effects. We obtain broadly similar results to those reported below.

5.2 Results

We run pairwise regressions predicting NPE purchase using various causal examiner effects, including propensity to change the number of words per claim, propensity to remove independent claims, and propensity to use the various blocking actions discussed in Section 2.1. In the results reported in the main text, we use blocking actions on eventually granted patents in order to better capture intensive margin effects - but we also report results for blocking action usage at the application level in the Appendix.

We compute how much a one-standard-deviation change in the examiner causal effect would change the outcome in question (e.g. NPE purchase), as a fraction of the baseline outcome rate. These results are reported in the tables in this section. In addition, we present summary graphs showing the pairwise predictive power of various examiner effects. In the graphs, we normalize the computed effects by the "own variable" predictive effect. This own variable effect (e.g. the examiner NPE effect) will generally have the strongest predictive power, since it is constructed to be predictive of the outcome, but this result is not mechanical and does not have to hold. Finally, in the graphs we color the bars blue for variables that are positively predictive of the outcome and red for the variables that are negatively predictive. The p-value for the coefficients are displayed in white.

5.2.1 Purchased Patents

As we see in the first column of Panel A of Table 10, the correlation between NPE outcome and NPE examiner effect is positive and the strongest predictor, with a one standard deviation change in examiner effect translating into a 44.7% change relative to the NPE purchasing baseline. As shown on Panels A and B of Table 10, several examiner effects are *negatively* predictive of NPE purchase, namely an examiner's propensity to add more words to claims, to issue blocking actions based on failure to satisfy subject patentability (101), non-obviousness (103(a)) and claim clarity (112(b)). Each of these have a sizable amount of predictive power relative to the NPE examiner effect, ranging from 5-15% of the baseline rate. This is more easily seen in Figure 4, which reports the same results in a graphical way, ordering the effects by their relative magnitude.

Let's now turn to the interpretation of these results and what they suggest about the strategy of NPEs. Panel A of Table 10 shows that NPEs purchase patents that have fewer words added and fewer independent claims struck down during examination than the average patent. The interpretation of this finding in terms of the strategy pursued by NPEs is ambiguous: patents with fewer words per claim and more independent claims could be less specific, less clear, but also broader than the average patent. On the basis of this evidence alone, it is difficult to conclude whether NPEs target vague or broad patents.

The examiner usage of blocking actions allows us to refine our interpretation of NPEs' patent acquisition behavior and distinguish between the clarity and breadth channels. We find that examiners who use more 101, 103(a) and 112(b) blocking actions are less likely to have their patents purchased by NPEs (with a 6%, 10% and 5% effect per examiner effect standard deviation, respectively). This suggests that NPEs purchase incremental patents (given 103(a) effect) that have not been properly specified (given 112(b) effect) and whose subject may not be eligible for a patent (given 101 effect). It appears unlikely that these various effects would be significant if the breadth channel was driving the purchase decision of NPEs.

As a comparison, we run the same routine for patents purchased by practicing firms . This exercise confirms that practicing firms active in the markets for technology behave very differently from NPEs. The results are reported in columns 2 of Panels A and B of Table 10. In general, examiner effects have much less predictive power for non-NPE purchase than for NPE purchase, and in particular the specific effects for 101, 103(a) and 112(b) usage are much smaller in magnitude (0%, 1.3% and 0%, respectively).[35] Later in the

[35] As an aside, the 112(a) examiner effect appears to be unrelated to non-NPE purchase, suggesting that addition clarity in technological disclosure does not facilitate more transfer through purchases. It would be interesting to see if this holds in the pool of licensed patents - unfortunately, licensing data is not available.

section, we formally confirm that these are indeed statistically significant differences by running regressions with interaction effects explictly testing for such differences.

Table 10: Pairwise Examiner Effects Results - Purchase

Panel A: General Examiner Effects

	NPE Purchase	Non-NPE Purchase
NPE Purchase	0.447***	
	(0.076)	
Non-NPE Purchase		0.110***
		(0.006)
Allowance	0.098***	0.020***
	(0.017)	(0.004)
Words Per Claim Change	-0.148***	-0.019***
	(0.021)	(0.004)
Independent Claims Change	0.078**	0.005
	(0.030)	(0.006)
Patent-Level Controls	x	x
N	1,269,623	1,269,623

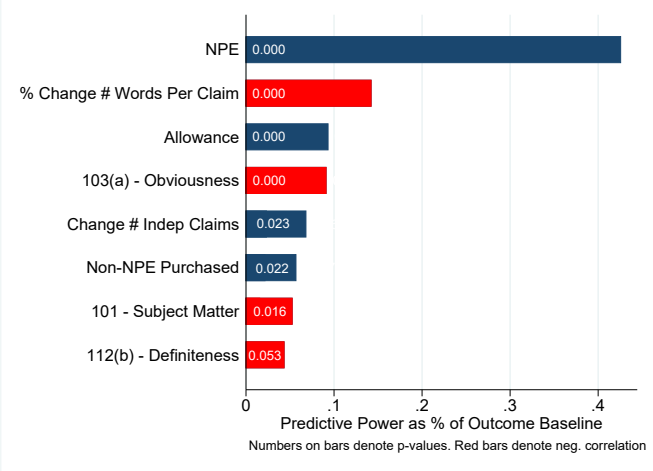
Panel B: Examiner Blocking Action Effects

	NPE Purchase	Non-NPE Purchase
101	-0.064**	0.002
	(0.024)	(0.003)
102(a)	0.012	0.006
	(0.025)	(0.005)
103(a)	-0.099***	-0.013**
	(0.023)	(0.005)
112(a)	-0.021	0.003
	(0.018)	(0.004)
112(b)	-0.047**	0.001
	(0.023)	(0.004)
Patent-Level Controls	x	x
N	1,269,623	1,269,623

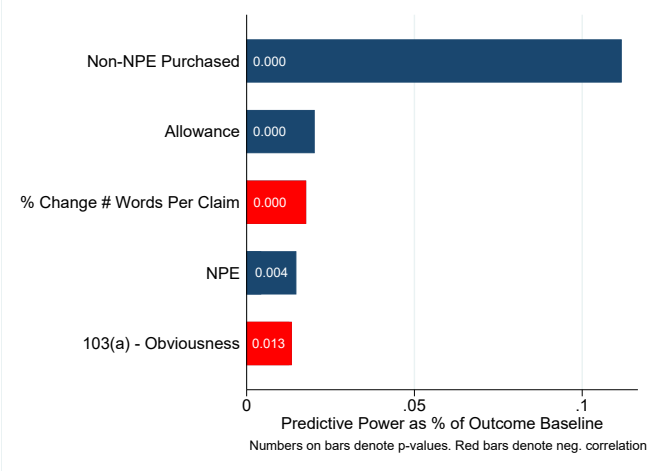
Notes: All results are produced using the restricted Frakes and Wasserman data. All coefficients are normalized to represent the impact of a one standard deviation in examiner effect relative to the baseline rate of the outcome in question (1.0% for NPE patents, 19.7% for non-NPE purchased patents). Words per claim variable refers to the percentage change in word count between application and grant. Patent-Level Controls include the history of the first inventor and the log of the number of independent claims at application. Standard errors are clustered at the examiner level. ** p-value <0.10, ** p-value < 0.05, *** p-value < 0.01*

Figure 4: Graphical Summary of Pairwise Examiner Effects Results - Purchase

(a) Predicting NPE Patent Purchase



(b) Predicting Non-NPE Patent Purchase



Notes: Results presented are based on the pairwise predictive regressions reported in Table 10. Effect sizes are again calculated in terms of fraction relative to the baseline rate of NPE purchase per one standard deviation change in the examiner effect, and then normalized to the size of the examiner effect calculated from the outcome variable in question (NPE purchase, non-NPE purchase). Blue bars indicate a positive predictive coefficient and red bars indicated a negative predictive coefficient.

5.2.2 Litigated Patents

Next, we examine litigation behavior. We repeat the same exercise as before, but with NPE litigation and non-NPE litigated as the outcomes. We report the corresponding results in Table 11. One challenge to note is that the baseline rate of non-continuation, NPE-litigated patents is around 0.033% of the patent population (about one-tenth the rate of non-NPE litigated patents).[36] Our methodology for computing examiner effects may not work well with very rare outcomes, as it relies on consistency to extract the underlying signal. Therefore, while we do compute an examiner "NPE Litigated" effect, it is quite noisy and not a reliable predictor of outcomes. However, we can use the other examiner causal effects to predict the occurence of NPE litigation.

The results show that NPEs litigate patents that have some similarity with patents litigated by practicing firms. However, NPEs once again appear to have a much stronger sensitivity to other examiner effects, such as 103(a) and 112(b), although non-NPE litigated patents do show a relatively strong sensitivity to the same two examiner effects. We interpret the results for non-NPE litigated patent as in line with standard litigation

[36] One important issue to note here is that NPEs often use the same set of patents against many plaintiffs, much more so than regular firms. Therefore, despite having about a tenth of the number of asserted patents, NPEs are still a major presence in terms of the number of firms they sue for infringement. In the Cotropia et al classifications of 2010 and 2012 IP lawsuits, around 35% of plaintiffs are NPEs. These numbers may also be a little misleading, as pre-AIA joinder rules meant NPEs could target multiple defendants in the same case, which we observe in the litigation data provided by RPX.

models (e.g. Klerman and Lee, 2014), which predict that plaintiffs and defendants only go to court if there is disagreement over the chances of winning. Vague claims language and uncertainty over obviousness may create such disagreements.

Table 11: Pairwise Examiner Effects Results - Litigation

Panel A: General Examiner Effects

	NPE Litigated	Non-NPE Litigated
NPE Litigated	-0.119	
	(0.132)	
Non-NPE Litigated		0.267***
		(0.040)
Allowance	0.321***	0.086***
	(0.052)	(0.018)
Words Per Claim Change	-0.346***	-0.061***
	(0.046)	(0.016)
Independent Claims Change	0.130*	0.049**
	(0.054)	(0.019)
Patent-Level Controls	x	x
N	1,269,623	1,269,623

Panel B: Examiner Blocking Action Effects

	NPE Litigated	Non-NPE Litigated
101	-0.123*	-0.032**
	(0.066)	(0.015)
102(a)	0.011	-0.011
	(0.054)	(0.016)
103(a)	-0.240***	-0.039**
	(0.066)	(0.017)
112(a)	-0.140***	-0.036**
	(0.044)	(0.017)
112(b)	-0.139**	-0.040**
	(0.055)	(0.018)
Patent-Level Controls	x	x
N	1,269,623	1,269,623

Notes: Results are run on the restricted Frakes and Wasserman sample. All coefficients are normalized to represent the impact of a one standard deviation in examiner effect relative to the baseline rate of the outcome in question (0.5% for non-NPE litigated and 0.033% for NPE litigated). ** p-value <0.10, ** p-value < 0.05, *** p-value < 0.01*

5.2.3 Testing for Differences

Finally, we look for examiner effects that distinguish between NPE purchase vs. non-NPE purchase, NPE litigated vs. NPE purchase, and NPE vs. non-NPE litigated. As an additional exercise, we test for a distinction between Intellectual Ventures and other NPEs. So far, we have compared groups of patents with

interesting outcomes to the rest of the patent pool. Here, we restrict the analysis sample to patents with at least one of a pair of outcomes, and then looking for examiner effects that are conditionally predictive of the first effect in each pair mentioned. This is a more rigorous way to compare types of patent outcomes.

The results are shown in Table 12. NPEs purchase patents from examiners who use 103(a) and 112(b) blocking actions less often than their peers, and strongly prefer patents with fewer added words during examination. NPEs selectively litigate more obvious patents, but there is no significant additional 112(b) effect. Another interesting point to note is that the patents purchased by Intellectual Ventures do not exhibit significant differences relative to other NPE purchased patents, as shown by the several precisely estimated zero results in column 4. This provides some evidence that the purchasing activities of Intellectual Ventures, which ignores their in-house inventions, are not particularly different from other NPEs, contrary to narratives that attempts to distinguish between various types of NPEs.

Table 12: Pairwise Differential Outcome Regressions

	NPE vs. non-NPE		Litigated vs. Purchased	Int. Vent. vs. Other
	Purchased	Litigated	NPE	Purchased
101	-0.042***	-0.073	-0.046	-0.012
	(0.015)	(0.050)	(0.033)	(0.010)
102(a)	0.011	0.018	-0.012	-0.005
	(0.020)	(0.056)	(0.044)	(0.016)
103(a)	-0.071***	-0.190**	-0.123**	-0.004
	(0.020)	(0.067)	(0.049)	(0.013)
112(a)	-0.019	-0.074^{+}	-0.133***	0.026**
	(0.014)	(0.042)	(0.045)	(0.012)
112(b)	-0.039**	-0.074	-0.066	-0.000
	(0.019)	(0.053)	(0.046)	(0.014)
Words Per Claim Change	-0.114***	-0.298***	-0.178***	-0.016
	(0.017)	(0.060)	(0.043)	(0.012)
Independent Claims Change	0.066**	0.080	0.056	-0.003
	(0.028)	(0.065)	(0.055)	(0.028)
N	262,511	7,156	12,953	12,953

Notes: The results are is run on the restricted Frakes and Wasserman data. Int. Vent. refers to Intellectual Ventures. All coefficients are normalized to represent the impact of a one standard deviation in examiner effect relative to the baseline rate of the outcome in question. ** p-value <0.10, ** p-value < 0.05, *** p-value < 0.01*

5.2.4 Discussion

The general pattern in the results presented in this Section show that patents with fewer words added during examination and fewer independent claims removed are more likely to show up in all of the patent groups we have analyzed. Digging further down, NPE purchased and NPE litigated patents have much more in common with non-NPE litigated patents than with non-NPE purchased patents. NPEs tend to purchase and litigate

patents granted by examiners who are less likely to judge a patent application to have vague claims and to be obvious. This holds up in comparisons to both the general patent pool and to corresponding non-NPE purchased and litigated patents. We dig further into the mechanisms that could explain these patterns in Section 6

Similar patterns, although less strong, hold for litigated patents. These patterns are interesting in and of themselves, independent of the debate about the nature of NPEs' activities. Our results show that many patents are litigated because of the examiners the patent application was assigned to, and the examiners with a high litigation causal effects tend to be "lenient" examiners with high allowance rate who tend to allow obvious and unclear patents. These results are in line with standard litigation models where litigation occurs because of disagreement over ill-defined and incremental patents overlapping with existing patents. This suggests that a large fraction of litigation is "inefficient" because it is caused by the prosecution process at USPTO, results in high litigation fees, and could likely be avoided by increasing the quality of the examination process.

As an additional test of the idea that certain examiners cause patent litigations because they issue ill-defined patent rights, we test whether the non-NPE litigated examiner effects are still a reliable predictor of litigation in a sample of patents linked to drugs that were approved by the FDA (in other words, we restrict attention to the patents listed in the FDA Orange Book). If the key issue driving the examiner effects for litigation is lack of clarity, then examiners should no longer matter when we consider a set of patents whose associated products (drugs) benefit from a separate and unambiguous source of market power (via the FDA). Indeed, we find that in this subsample examiner effects are no longer statistically significant. Appendix Figure 11 reports the binned scatter plot showing this.

5.3 Robustness Checks

The main focus in this section is to test robustness to NPE classifications, adding extra controls to the basic specification, and testing for the significance of examiner career effects. We report the most important results in Tables 13 and 14 and leave the rest for the appendix.

5.3.1 Alternative Specifications

First, we verify that our results are robust across NPE lists. For this, we construct a portfolio from the Cortropia et al. classifications, focusing on entities they classify as large aggregators and small patent

holding companies. As mentioned earlier, their list only identifies entities that are plaintiffs in litigation in 2010 and 2012, and therefore may leave some true NPEs in the control group. The results using their list of NPEs are consistent with our core findings.

We also run alternative specifications to compute the residuals in our examination effects methodology, which corresponds to changing the X in Equation 2. The baseline just contains art unit by year fixed effects as controls, to reflect the level of random assignment. One major addition we try here is to add assignee fixed effects, based on the initial assignee on a given patent.

The limitation of this specification is that it eliminates many unassigned patents, which make up a sizable part of NPE portfolios, as shown in Table 4. However, this approach allows us to compare the patents sold by a firm to NPEs versus the ones they keep or sell to other entities. The results suggest that firms tend to sell the vaguer and more obvious patents within their own portfolios.[37] An additional interpretation of these results is that it rules out the story that our results are driven solely through a bankruptcy channel, namely the possibility that NPEs mostly purchase intellectual property assets during firm liquidation and firms that own the types of patents shown in our results are more likely to go bankrupt. Our results, as shown in Table 13, are generally very similar to the baseline results, but the additional noise from a reduced sample pushes the 112(b) result to become statistically insignificant.

We also run specifications adding in only observable patent application characteristics as controls. These controls include the number of independent claims at application, the applicant entity size (regular or small), and the patenting history of the firm and the inventors at the time of the application. The results are also robust to these additions.

[37] This result may be related to the concept of "patent privateering." Practicing companies face the threat of countersuit if they assert their patents directly against a competitor. Instead, they can choose to go after competitors by selling their patents to NPEs with some sort of protection provision included in the transaction.

Table 13: Robustness checks for key predictors NPE purchase.

Specification	103(a)	112(b)
Baseline	-0.091***	-0.043^{+}
	(0.023)	(0.022)
Kesan *et al* NPEs	-0.090***	-0.051*
	(0.027)	(0.023)
Non-IV NPEs	-0.088***	-0.042^{+}
	(0.027)	(0.026)
IV	-0.095**	-0.044^{+}
	(0.030)	(0.026)
Assignee Fixed Effects	-0.096***	-0.037
	(0.025)	(0.023)
Additional Patent Controls	-0.088***	-0.042^{+}
	(0.026)	(0.025)
Five Year Lag	-0.092***	-0.044^{+}
	(0.025)	(0.023)

Notes: Non-IV NPEs refers to patents purchased by NPEs other than Intellectual Ventures and IV refers to patents purchased by Intellectual Ventures. Additional Patent Controls refers to adding the same inventor and application characteristics to the basic estimating equation as the specification in Table 9. Five Year Lag refers to computing signal covariances using data five years apart instead of one, and is the same methodology used to compute the results in Table 14. *+ p-value <0.10, * p-value < 0.05, ** p-value < 0.01, *** p-value < 0.001*

5.3.2 Examiner Career Effects

One possible concern is that examiner effects are not "fixed" as in our preferred framework but rather vary substantially over the course of an examiner's career at the patent office. For instance, Frakes and Wasserman (2014) find that time constraints vary a lot during an examiner's career. Our baseline methodology does not account for this, as we assume a fixed examiner effect μ_j, along with an idiosyncratic examiner by cohort effect θ_{jt}.

We show that time-varying examiner effects are not quantitatively important in a number of ways. First, we re-compute results by computing signal covariances using different lags $\hat{\sigma}^2_\mu = cov(\bar{v}_{jt}, \bar{v}_{j(t+k)})$, for values of k beyond 1, which we use to generate our baseline results. As shown in Table 14, our core results are robust to these changes. We compute NPE effect distribution parameters for $k = 3, 5$, and compare it to the parameters estimated by using $k = 1$, but restricting the samples to examiners with at least 4 or 6 years of experience, respectively. Here, we use the non-continuation PERD sample, in order to increase the years of coverage, which is not very large relative to the lags we are using. The results suggest that our high signal covariance is not driven by examiner career effects, as we are now correlating residuals for examiners at very different points in their respective careers. The decrease in estimated signal correlations might not be driven by career effects alone, as time variation in NPE purchase rates could cause problems in our thirteen-year

sample.

As a methodological aside, relative to the non-continuation PERD sample results reported in Table 9, the modified baselines we compute here have almost identical distribution parameters, suggesting that our methodology shrinks examiners with less data in an effective manner.

Table 14: Robustness checks for the effect of Examiner Career Effects on computed NPE effect distribution.

Specification	Signal SD	Shrunk Effect SD	Shrunk Effect IQR
Three Year Lag	44.3%	19.4%	4.8%
Baseline	53.2%	26.7%	6.4%
Five Year Lag	35.8%	14.5%	3.5%
Baseline	52.8%	27.8%	6.1%

Notes: Three Year Lag and Five Year Lag Refer to computing the signal covariance $\hat{\sigma}^2_\mu = cov(\bar{v}_{jt}, \bar{v}_{j(t+k)})$ using $k = 3, 5$ respectively. The baseline in each case is re-calculated to only include examiners that have careers of at least 4 and 6 years, respectively. The results reported here are computed using the full PERD sample.

Second, we follow the methodology in Frakes and Wasserman (2015), and look for significant changes in examiner NPE grant rates over their careers, particularly before and after promotions. With minor exceptions, we generally do not find significant NPE effect differences over an examiner's career. These results are available from the authors upon request.

5.3.3 Breaking Down the Results across NPE Lists and Technological Areas

First, we break down the results by Intellectual Ventures versus other NPEs in our list. A downside to our methodology, as discussed in the context of computing NPE Litigated examiner effects, is that it has difficulty extracting signals that occur with very low frequency. Therefore, we cannot use it to describe the behavior of NPEs that hold a handful of patents. However, we are still able to check the behavior of Intellectual Ventures, which has by far the largest patent portfolio, versus the combined portfolio of other NPEs.

In addition, we also break down our results by technological areas. We do this by analyzing patents from each of the eight technology centers at the USPTO separately.[38] Results are reported in Table 32. Our results suggest that the non-obviousness result holds mainly for the four IT-related technology centers - 21 (Computer Architecture, Software, and Information Security), 24 (Computer Networks, Multiplex communication, Video Distribution, and Security), 26 (Communications), and 28 (Semiconductors, Electrical and Optical Systems and Components) - and the vague claims text results hold for art unit 26 and 28. Our core results do not hold

[38] There are a few dozen art units within each technology center.

for technology centers 16 (Biotechnology and Organic Chemistry), 17 (Chemical and Materials Engineering), and 36 (Transportation,Construction, Electronic Commerce, Agriculture, National Security and License & Review), suggesting that the IP market operates differently in those areas. The results for the non-NPE litigated patents generally follow the same patterns.

5.3.4 Horse Race Between Effects

Finally, we also run a "horse race" regression between the examiner effects to predict various outcomes. The results are shown in Table 15. The 103(a) examiner effect tends to preserve statistical significance in all of the specifications. One issue to note is that some pairs of examiner effects have low correlation, and therefore each separate effect has some predictive power, particularly the 101 effect. However, as shown in Table 20, 103(a) and 112(b) tend to be used together on applications, making it harder to separately identify the contribution of the 112(b) effect.

Table 15: Horse Race Regressions Predicting Each Outcome

Panel A: Horse Race for Outcomes

	NPE Purchase	Non-NPE Purchased	NPE Litigated	Non-NPE Litigated
101	-0.0425*	0.0041	-0.0977	-0.0314
	(0.0225)	(0.0042)	(0.0636)	(0.0185)
102(a)	0.0313	0.0053	0.0472	0.0000
	(0.0241)	(0.0051)	(0.0518)	(0.0184)
103(a)	-0.0829***	-0.0166***	-0.202***	-0.0472**
	(0.0249)	(0.0058)	(0.0635)	(0.0201)
112(b)	-0.0104	0.0014	0.0241	0.00428
	(0.0241)	(0.0054)	(0.0593)	(0.0249)
N	1,269,623	1,269,623	1,269,623	1,269,623

Panel B: Differential Horse Race Regressions

	NPE vs. non-NPE		Litigated vs. Purchased
	Purchase	Litigated	NPE
101	-0.0327**	-0.0427	-0.0226
	(0.0163)	(0.0521)	(0.0340)
102(a)	0.0223	0.0406	0.00962
	(0.0198)	(0.0567)	(0.0439)
103(a)	-0.0607***	-0.170**	-0.0989*
	(0.0212)	(0.0715)	(0.0508)
112(b)	-0.0124	0.0152	0.0449
	(0.0209)	(0.0583)	(0.0549)
N	262,511	7,156	12,953

Notes: The regressions in Panel A are run on the baseline sample. The regressions in Panel B are run on the same restricted samples as in Table 12. ** p-value <0.10, ** p-value < 0.05, *** p-value < 0.01*

6 Do NPEs Efficiently Screen Good Technology or Strategically Target Incremental Patents?

In this section, we conduct additional tests to distinguish between the remaining mechanisms that could be consistent with the evidence on NPE patent purchase and assertion behavior we have collected so far. To do so, we present additional evidence from EPO grant outcomes (for patents filed jointly at EPO and USPTO) and from litigation trial outcomes (for litigated patents) suggesting that NPE patents tend to be of particularly low technological merit, even relative to patents asserted by practicing entities. Accordingly, we conclude that on average NPEs do not provide incentives for high-quality innovations.

6.1 Hypothesis Development

Our results so far have established that on average NPEs purchase patents from a specific set of examiners, who tend to issue incremental patents (low 103(a) usage) with vaguely worded claims (low 112(b) usage). This is already informative about the role played by NPEs in the IP system: when a patent is truly innovative and well defined, there does not seem to be a big role to play for intermediaries like NPEs. Our evidence that NPEs purchase and assert patents coming from a set of examiners with similar characteristics to those with high litigation effects is in line with the idea that NPEs have a comparative advantage in litigation.

Nonetheless, two rather different views about the role played by NPEs are still consistent with these facts. The first view is that NPEs are in fact able to use their technological and legal expertise to efficiently screen patents and, in the set of vague patents, identify those patents that are truly high quality, even though they may be poorly written. In other words, NPEs might have a comparative advantage in screening (not just in litigation) and they end up purchasing primarily from "lenient" examiners precisely because it is in the pool of poorly-examined patents that screening expertise is needed to tell apart the good technology that needs to be upheld from the low-quality technology. This view predicts that *within* the portfolio of lenient examiners, NPEs in fact purchase and assert high-quality patents. In that sense, they may still play a productive role, although it is difficult to see why it would not be more efficient for the USPTO to screen out low-quality patent ex-ante.[39]

The second view is that NPEs strategically target incremental patents with vaguely-worded claims because

[39] Doing so does not appear to be beyond the abilities of the USPTO, given that at least some examiners appear to be able to appropriately reject or narrow weak patents, which results in lower litigation rates in general and lower rates of NPE purchase in particular.

these patents encroach on existing intellectual property and can be leveraged to sue practicing entities. According to this view, the average NPE's business model is broadly similar to the strategy that "patent pirates" relying on "submarine patents" used to follow.[40]

To distinguish between these two views, we examine evidence on the quality of the patents purchased by NPEs *within* the portfolio of a given examiner using grant outcome at the EPO (for patents filed jointly in the US and Europe). In line with the literature, we interpret EPO grant as a proxy for patent quality, or larger innovation step size. We find that within an examiner's portfolio (i.e. conditional on examiner's fixed effects), EPO grant negatively predicts NPE purchases, which is inconsistent with the theory that NPEs serve as effective intermediaries identifying valuable technologies in the pile of vaguely worded patents.[41] This is line with one of our previous finding from Section 5 that NPE purchased patents are unassociated with 112(a), the technological disclosure requirement, but negatively associated with 112(b) (which suggests that NPEs are better at uncovering legally useful patents rather than technologically useful patents).

Finally, to examiner whether the issue primarily comes from examiners behaving in a way that is inconsistent with existent patent law, or whether it is primarily a statutory issue (e.g. the existing statutes of U.S. Code Title 35 give examiners too much leeway during the course of the prosecution), we examine the outcome of litigation trials (for litigated patents for which there was a trial) and PTAB judgements. We find that examiners with a high NPE effects tend to be invalidated in court and reversed at PTAB, suggesting that the core issue does not lie with the existing statutes.

6.2 Research Design using European Patent Office Grant Outcomes

In order to test whether NPEs play an efficient screening role by identifying high-quality patents within the portfolio of lenient examiners, we need a proxy for patent quality. Following Lei and Wright (2009) and de Rassenfosse et al. (2016), we use grant outcomes at the EPO, for patents jointly filed in Europe and in the US, as a proxy for patent quality and step size.

We consider a sample of patent applications that were jointly filed at the USTPO and EPO. We run patent-level specifications with NPE purchase as the outcome variable, EPO grant as the independent variable, and

[40] A submarine patent is a patent whose issuance and publication are intentionally delayed by the applicant for a long time, such as several years. This strategy requires a patent system where, first, patent applications are not published, and, second, patent term is measured from grant date, not from priority/filing date. Submarine patent practice was possible previously under the United States patent law, before the U.S. signed the TRIPS agreement of the WTO in 1995. A famous example is the inventor and entrepreneur Gilbert Hyatt, who was awarded a patent claiming an invention pre-dating both TI and Intel describing a "microcontroller." The patent was eventually invalidated, but not before substantial royalties were paid out.

[41] The idea of digging up valuable patents is discussed in the book "Rembrandts in the Attic" by Kevin Rivette and David Kline.

control for art unit - year fixed effects, examiner fixed effects and assignee fixed effects. The examiner fixed effects insure that the examiner behavior is controlled for and assignee fixed effects handle the issue that some assignees are more likely to be successful with their EPO applications than others.

We run this specification in two separate subsamples, below and above median of the NPE purchase examiner effect. The "NPE as efficient screeners" view predicts that NPEs purchase higher-quality patents (i.e. patents that were granted at the EPO) in the portfolios of examiners who in general issue patents with incremental and vaguely worded claims (i.e. the subsample of examiner above median of the NPE purchase effect). In contrast, the "NPE as frivolous litigators" view predicts that NPEs purchase lower-quality patents (i.e. patents that were rejected at the EPO) in the portfolios of examiners above median of the NPE purchase effect.

6.3 Results

Table 16 presents the results, which reject the "efficient screening" view. All point estimates are standardized by the baseline rate of the outcome in the relevant sample. We find that NPEs tend to purchase patents that are *rejected* at the European office in the sample with examiner above median of the NPE examiner effect distribution, which is in line with the "NPE as frivolous litigators" view. This effect is quantitatively large: the probability of a purchase by a NPE decreased by 21% if the patent is granted by the EPO. There is no such effect for patent purchases by practicing entities (for which we obtained a very precisely estimated 0) or for non-NPE litigation. Figure 12 in Appendix D Figure reports the binned scatter plots for these various specifications.

Table 16: Testing the Efficient Screening View

Panel A: In the sample with NPE Examiner Effect above Median

	NPE Purchase	Non-NPE Purchase	Non-NPE Litigation
EPO Grant	-0.2144**	0.0037	-0.0831
	(0.1001)	(0.0133)	(0.1074)
Examiner F.E.		Yes	
Assignee F.E.		Yes	
Artunit-Year F.E.		Yes	
N		109,383	

Panel B: In the sample with NPE Examiner Effect below Median

	NPE Purchase	Non-NPE Purchase	Non-NPE Litigation
EPO Grant	0.0819	-0.0052	0.0196
	(0.0702)	(0.0126)	(0.0921)
Examiner F.E.		Yes	
Assignee F.E.		Yes	
Artunit-Year F.E.		Yes	
N		109,484	

** p-value <0.10, ** p-value < 0.05, *** p-value < 0.01*

6.4 Additional Evidence from PTAB and Litigation Trial Outcomes

To provide additional evidence that NPEs are active in the market for weak patents, we relate the patent acquisition behavior of NPEs more directly to the validity of the patents by studying whether the examiners who produce NPE patents tend to be invalidated in court (in the case of litigation by regular entities) or to be reversed during appeal procedures at the Patent Trial and Appeal Board (in the case of an appeal by a regular entity during the course of prosecution).

Our results suggest that the examiners with a high NPE effect are more likely to behave in a way that is not in accordance with the law. Figure 13, reported in Appendix B, shows that examiners with a high NPE effect tend to be reversed at the Patent Trial and Appeal Board when there is an appeal during the course of prosecution.[42]

Moreover, Table 17 shows that infringement are less likely to be found during a litigation trial if the patent was purcahsed by a NPE. Although the results of such a regression are difficult to interpret because of selection effects (a small percentage of patents are litigated, and an even smaller percentage goes to trial),

[42] More than 50% of examiners have one of the patent applications reviewed by the Patent Trial and Appeal Board during the course of their career, which mitigated concerns about the selection effects inherent in our strategy.

we view these results as suggestive evidence that the examiner with high NPE purchase causal effects are granting invalid patent and, therefore, that the fundamental issue lies with examiner behavior and not with the macro-determinants of the IP system like the various statutes in Title 35 of the U.S. Code.

Table 17: NPE Purchase and Patent Trial Outcomes

	Infringe | Trial
NPE Purchase	-0.122***
	(0.045)
Words Per Claim Change	0.057
	(0.054)
101	-0.070
	(0.046)
102(a)	-0.035
	(0.058)
103(a)	-0.010
	(0.057)
112(a)	0.026
	(0.049)
112(b)	-0.059
	(0.054)
N	595

6.5 Discussion

The fact that NPEs purchase patents from examiners who tend to grant incremental patents with vaguely worded claims, which we documented in Section 5, left open the possibility that NPEs could be playing an efficient screening role, dealing with the problems introduced into the IP system by the inconsistencies of the examination process at the USPTO.

There are various issues with this hypothesis. First, the activity of NPEs appears to be at best unrelated if not negatively related to technological attributes of the patents. Within the portfolio of examiners with a high NPE effects, we find that NPEs are significantly less likely to purchase patents that were granted at EPO, our proxy for technological quality and step size. Second, we also find that examiners with a high NPE effect tend to be reversed at the Patent Trial and Appeal Board when there is an appeal during the course of prosecution and that their patents tend to be invalidated when they reach trial during litigation. This is another clear sign that these examiners in general do not behave in accordance with the law and the patents NPEs purchase from them in particular are weak. Finally, even if NPEs play a role in screening vague

patents, the current setup involving high legal fees seems like a very costly way to deal with the problem. Our policy calculations in Section 7 suggest that it would be much more cost-effective to solve the vague patents problem at its source, namely the patent examination process.

7 Policy Implications

7.1 Calibrations

In this section, we use our estimated examiner effects in order to calibrate the returns to investment in the quality of the patent examination process. We focus here on the possible litigation fees saved from a reduction in the types of patents currently used by NPEs, and consider this to be a lower bound on the social costs of issuing such patents. Indeed, our calibration does not include other possible inefficiencies generated by NPEs, such as implicit taxes imposed on products challenged by NPEs (Tucker (2015)) and other possible distortions in the intellectual property system. More broadly, it does not include the impact a reform of the patent examination process would likely have on litigation by practicing entities.

First, we discuss here some estimates for litigation and USPTO costs that we will use to interpret our calculations. In terms of annual litigation fees associated with NPE lawsuits, there are a variety of estimates in the legal literature, each using a different methodology. The main sources for this data comes from the American Intellectual Property Law Association (AIPLA) annual reports, which surveys its members for estimates of how much they are paid for cases of different sizes, and the surveys conducted by RPX and Bessen and Meurer (2014), which surveys defendants to estimate spending on legal defense and other costs. A major difference in the figures reported by the two sources stems from the fact that the AIPLA reports median figures for each case size category, whereas Bessen and Meurer use means, which are much higher because of a long right tail. The figure arrived at by Bessen and Meurer is $29 billion in "direct costs," of which around 17% come from defensive litigation costs.[43] Therefore, this comes out to over $5 billion used in defense. Although this number may be on the high side, it also does not include resources used by NPEs, so we use this number for our calculations. In terms of USPTO spending, the patent portion of the annual USPTO budget is $3.13 billion for FY 2016.[44] Of this, around $800 million is paid out in examiner salary.[45]

The cost of policy initiatives about examiners is typically modest in comparison to the potential benefits

[43] The remainder is licensing fees, which we interpret in a static way as just a transfer, and therefore not an inefficiency.

[44] http://www.uspto.gov/sites/default/files/documents/fy16pbr.pdf

[45] Based on examiner grade-level data from Frakes and Wasserman plus data on salary by examiner grade available in the USPTO annual budget report.

mentioned above. For instance, the aforementioned *Enhanced Patent Quality Initiative* (EPQI) at USPTO will cost around $11 million for FY 2016, ramping up to about $31 million by FY 2020.[46]. Another quality initiative implemented by the USPTO was the "second pair of eyes" review,[47] which was first piloted on the set of business method patents (USPTO class 705), then partially expanded in 2003 to several other technology centers and eventually discontinued in the late 2000s.[48] The initiative introduced one hour of review per provisionally granted patent, in order to flag obvious issues. Given that examiners spend on average 19 hours reviewing each application, this suggests that such a policy, if implemented across all technology centers, would cost at most $40 million.[49]

As mentioned above, the estimated NPE litigated examiner effects are quite noisy, due to the sparsity of the outcome combined with the nature of our methodology. Instead, we can use a series of alternative approaches to come up with an estimate of the partial equilibrium cost savings from the two aforementioned policies. The results will be partly based on results reported in Table 10. We can approximate the effect of policies like the EPQI, which aim at establishing best examination practices, as moving examiners above the 75th percentile in NPE purchase effect to the 75th percentile, or moving examiners from below the 25th percentile in 103(a) and 112(b) usage to the 25th percentile. Then, we can compute the impact of such a change as a fraction of the standard deviation of the examiner effect distribution, and then scale it by the coefficient in our mechanism result.

The calculations are reported in Table 18. The first row shows the returns if EPQI can directly target the NPE purchase examiner effects. The $297.5 million estimate suggests a very high social return on the $11 million investment. The remaining rows show the effects of EPQI if it were to target examiner usage of blocking actions. Both calculations also yield significant social returns.

[46] Based on P55 of the 2016 USPTO Budget Report, available at http://www.uspto.gov/sites/default/files/documents/fy16pbr.pdf

[47] Discussed on P28 of the 2006 USPTO Budget Report, available at http://www.uspto.gov/sites/default/files/web/offices/ac/comp/budg/fy07pbr

[48] Based on discussions with USPTO policy directors, the discontinuation resulted from examiners being uncomfortable with this review process, and possibly granting fewer patents in response, as only granted patents were scrutinized.

[49] Taking the $800 million as a base number and dividing it by the hours ratio. This will be an overestimate, as the there are more applications than granted patents.

Table 18: Simulating the Benefits of the Enhanced Patent Quality Initiative

	ΔNPE Litigation per SD	Move tail (SDs)	Total ($ millions)
NPE Purchase	0.238	6%/24% =0.25	297.5
103(a)	0.091	2.8%/19.1%=0.146	66.7
112(b)	0.043	7.7%/38.6%=0.199	42.8

Notes: Results are calculated by taking the baseline number of $5 billion per year in NPE litigation fees, multiplying it by the effect per SD, and then multiplying by the effect of the policy (represented in standard deviation units). The second column is based on results reported in Table 10. The third column, "Move tail," reflects the change by moving examiners with NPE Purchase effects above the 75th percentile to the 75th percentile and by moving examiners with 103(a) and 112(b) usage below the 25th percentile to the 25th percentile, similar to the numbers reported in Table 6, but then re-scaled to be a fraction of the standard deviation of the examiner effects distribution for that variable.

Next, we can evaluate the second-pair-of-eyes policy using a slightly more involved simulation. Here, we once again rely on translating other examiner effects into a change in NPE litigation. We simulate such a policy using our estimated primitives by randomly assigning each application i, which was assigned to examiner j, to another examiner k, and taking the minimum of the examiner NPE effects or the maximum of the blocking action effects. We winsorize examiner effects at the 1st and 99th percentile before performing the calculations:

$$\Delta NPE = \frac{1}{N}\sum_{i=1}^{N} \min(v_j, v_k) \tag{4}$$

$$\Delta 103(a) = \frac{1}{N}\sum_{i=1}^{N} \max(v_j, v_k) \tag{5}$$

In addition, we also simulate a scenario in which the second examiner has a much small effect, which is captured in Equation 6:

$$\Delta NPE = \frac{1}{N}\sum_{i=1}^{N} \min(v_j, \frac{2}{3}v_j + \frac{1}{3}v_k) \tag{6}$$

$$\Delta 103(a) = \frac{1}{N}\sum_{i=1}^{N} \max(v_j, \frac{2}{3}v_j + \frac{1}{3}v_k) \tag{7}$$

Table 19: Simulating the Benefits of the "Second-pair-of-eyes" Policy

	Effect per SD	Major Second (SDs)	Total ($ millions)	Minor Second (SDs)	Total ($ millions)
NPE Purchase	0.238	13.3%/24%=0.554	659.4	4.23%/24%=0.176	209.7
103(a)	0.091	12.5%/19.1%=0.654	297.8	4.40%/19.1%=0.230	104.8
112(b)	0.043	22.9%/38.6%=0.593	127.6	7.90%/38.6%=0.205	44.0

Notes: Results are calculated by taking the baseline number of $5 billion per year in NPE litigation fees, multiplying it by the effect per SD, and then rescaling by the impact of the policy, measured in standard deviation units. The second column is based on results reported in Table 10. "Major Second" refers to the scenario in which the second examiner has a major impact on the examination process (see Equation 4), whereas "Minor Second" refers to the scenario in which the second examiner has a minor impact (see Equation 6).

7.2 Discussion

In general, our calibration results suggest large social returns to relatively inexpensive USPTO investments. Of course, our calculations do not take into account the reactions of NPEs to a policy change. In addition, our calculations also equally weigh litigated patents, when in reality, some patents are involved in a disproportionate number of cases. The nature of our methodology somewhat limits our ability to accurately account for these weighting issues. Finally, our calculations have several moving pieces, particularly the modeling of the two policy initiatives. However, we have tended to err on the cautious side in all of our assumptions, in order to provide a lower bound for the returns to public investment. More broadly, our evidence suggests that policy reforms about patent examiners have great potential and would be a welcome addition to the current policy debate, where the discussion has focused on reforms of the statutes in Title 35 of the US code or on reforms of the court system.

Various limitations of this paper could be addressed in future work and would help refine the policy implications. First, we abstract away from the market for patents in our current analysis. A key issue when thinking about the patent market and NPEs is to determine why practicing companies do not outbid NPEs for these patents, or at least drive up the price. One simple explanation is that NPEs have greater expertise in identifying patents useful for their business model. The patent market is not particularly liquid, and some NPEs have built up extensive patent agent networks (Hagiu 2011). Another explanation is that regular companies cannot use patents in the same way as NPEs, because they are limited by the threat of a countersuit and possibly damages to firm brand, and therefore NPEs have much higher valuations for offensive patents than regular firms. Finally, practicing firms suffer from a free riding problem, because their purchasing of a patent for defensive purposes also helps their rivals in the same market. A recent trend

in the market is the emergence of defensive aggregators, such as RPX and Allied Security Trust. These defensive aggregators are funded by industry groups, and may offer a partial, market-based solution to the NPE problem by buying up problematic patents and guaranteeing that they will not assert them. In addition, NPEs may still be somewhat active even if the pool of patents possessed fewer of the attributes they currently appear to desire.

Second, we have not delved into the details of the court system. An important caveat here is that the effectiveness of patents in litigation depends not only on the claims text, but also on how judges interpret the claims. A patent with conspicuously broad or obvious claims would therefore be less useful to NPEs looking to maximize litigation revenue or the threat of litigation. An obvious channel for leveraging these patents, as frequently discussed in NPE debates, is the Eastern District of Texas, which has had a reputation for favorable treatment of plaintiffs. Another limitation in the analysis is that we do not observe legal fees and settlements, and therefore miss out on heterogeneity across lawsuits.

Finally, we did not attempt to fully evaluate the welfare effects of the activities of NPEs. Our results only address issues of static efficiency, and it may be that NPEs bolster exit value in a way that encourage entry of small firms or individual inventors, pushing against under-investment in innovation activities. However, our results suggest that NPEs would be encouraging the creation of vague and obvious patents, which seems like a blunt instrument for increasing the expected value of entry. Policies such as expanded R&D tax credits to small firms may be much more efficient socially to achieve such goals.

8 Conclusion

In this paper, we have exploited the random assignment of patent applications to examiners to show that examiners have a significant influence on a variety of patent outcomes, and to shed light on the characteristics of patents purchased and litigated by non-practicing entities. The evidence suggests that NPEs purchase more patents that are more obvious and contain vaguer claims text, attributes unrelated to or negatively associated with social value and invention quality. This research offers both evidence on the rent seeking behavior of NPEs and a natural policy lever to mitigate the impact of NPE activities. In fact, given the differing results on NPEs and practicing entities, policies that deal with vague patent language may be able to achieve their goals without disrupting the wider IP system.

In terms of policy, our quantitative estimates suggest that improving the patent examination process,

in particular by ensuring consistency in the narrowing and clarification of claims across examiners, has potentially very large returns. Indeed, a lower bound based on the estimates presented in Sections 4 and 5 is that the share of NPE patents among granted patents could be reduced by 20% by implementing a "second pair of eyes" policy. As mentioned earlier, this calculation is based on NPE litigation costs only, which are likely to underestimate the total social costs of NPEs' activities. Moreover, our evidence shows that the issuance of ill-defined patents also results in litigation among practicing entities: another likely benefit of a reform of the patent examination system would be to reduce these inefficient litigation costs, incurred due to disagreement over unclear and incremental patents.

Future research could build on our methodology in various ways. Viewed at a high level, our methodology allows us to indirectly infer rent-seeking asset purchase behavior of certain agents in a system. This is done by exploiting random assignment of the co-producer of an asset, in our case the examiner. The conditions here are to our advantage in that we have a variety of measures of co-producer behaviors that are orthogonal or negatively correlated with technological value, which we exploit to show that these co-producers are more likely to have their assets purchased by certain agents, in this case NPEs. We hope to see future work in this spirit, using asset production structure to test for rent-seeking behavior in other asset markets. Another possibly fruitful area of research is to test existing patent and innovation theories by leveraging our methodology to infer characteristics of patents. This could include looking at other ways patents are used in the IP system at large (such as licensing and commercialization of patent-protected products), and testing theories that predict the effects of patent breadth on real outcomes such as follow-on innovation.

References

[1] Allison, John R. and Lemley, Mark A. and Walker, Joshua H. (2009). Extreme Value or Trolls on Top? The Characteristics of the Most Litigated Patents. University of Pennsylvania Law Review, Vol. 158, No. 1.

[2] Bessen, James E. and Meurer, Michael J. (2014). The Direct Costs from NPE Disputes. 99 Cornell L. Rev. 387.

[3] Carley, Michael, Hedgde, Deepak, and Marco, Alan (2014). What is the Probability of Receiving a US Patent? USPTO Economic Working Paper.

[4] Chetty, Raj, John N. Friedman, and Jonah E. Rockoff. 2014. "Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood." American Economic Review, 104(9): 2633-79.

[5] Chetty, R., and Hendren, N. (2015). The impacts of neighborhoods on intergenerational mobility: Childhood exposure effects and county-level estimates. Working Paper.

[6] Cockburn, I., Kortum, S., and Stern, S. (2003). Are all patent examiners equal? Examiners, patent characteristics, and litigation outcomes. Patents in the knowledge-based economy, 35.

[7] Cohen, Lauren, Gurun, Umit G. and Kominers, Scott Duke (2014). Patent Trolls: Evidence from Targeted Firms. Harvard Business School Finance Working Paper No. 15-002.

[8] Cotropia, Christopher A., Kesan, Jay P. and Schwartz, David L. (2014). Unpacking Patent Assertion Entitites (PAEs). 99 Minnesota Law Review 649.

[9] Dobbie, W., and Song, J. (2015). Debt Relief and Debtor Outcomes: Measuring the Effects of Consumer Bankruptcy Protection. American Economic Review, 105(3), 1272-1311.

[10] Ellison, Glenn and Swanson, Ashley (2010). The Gender Gap in Secondary School Mathematics at High Achievement Levels: Evidence from the American Mathematics Competitions. Journal of Economic Perspectives, 24(2), 109-128.

[11] Fischer, Timo and Henkel, Joachim (2011). Patent Trolls on Markets for Technology - An Empirical Analysis of Trolls' Patent Acquisitions. SSRN Working Paper.

[12] Frakes, Michael D., Wasserman, Melissa F. (2014). Is the Time Allocated to Review Patent Applications Inducing Examiners to Grant Invalid Patents?: Evidence from Micro-Level Application Data. NBER Working Paper.

[13] Graham, S., Marco, A., and Miller, A. (2015). The USPTO Patent Examination Research Dataset: A Window on the Process of Patent Examination. SSRN Working Paper.

[14] Hagiu, A. (2011). Intellectual property intermediaries. Harvard Business School Strategy Unit Case, (711-486).

[15] Hall, Bronwyn H., Jaffe, Adam B., and Trajtenberg, Manuel (2001). The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools. NBER Working Paper.

[16] Kane, Thomas J. and Staiger, Douglas O. (2008). Estimating Teacher Impacts on Student Achievement: An Experimental Evaluation. NBER Working Paper.

[17] Lai, Ronald, D'Amour, Alexander, Fleming, Lee (2009). The careers and co-authorship networks of US patent-holders, since 1975. Harvard University Working Paper.

[18] Lei, Zhen and Wright, Brian D. (2009). Why Weak Patents? Rational Ignorance or Pro-'Customer' Tilt? CELS 2009 4th Annual Conference on Empirical Legal Studies Paper.

[19] Lemley, M. A., and Sampat, B. (2012). Examiner characteristics and patent office outcomes. Review of Economics and Statistics, 94(3), 817-827.

[20] Lerner, J. (1994). The Importance of Patent Scope: An Empirical Analysis. The RAND Journal of Economics, 25(2), 319-333.

[21] Lichtman, D. (2004). Rethinking prosecution history estoppel. The University of Chicago Law Review, 151-182.

[22] Love, Brian J. (2013). An Empirical Study of Patent Litigation Timing: Could a Patent Term Reduction Decimate Trolls Without Harming Innovators? University of Pennsylvania Law Review, Vol. 161, p. 1309.

[23] Maestas, N., Mullen, K. J., and Strand, A. (2013). Does disability insurance receipt discourage work? Using examiner assignment to estimate causal effects of SSDI receipt. The American Economic Review, 103(5), 1797-1829.

[24] Marco, Alan C., Graham, Stuart J.H., Myers, Amanda F., D'Agostino, Paul A and Apple, Kirsten (2015). The USPTO Patent Assignment Dataset: Descriptions and Analysis. SSRN Working Paper.

[25] Sampat, B. N., and Lemley, M. A. (2010). Examining patent examination. Stanford Technology Law Review, 2010, 2.

[26] Pakes, A. (1986). Patents as Options: Some Estimates of the Value of Holding European Patent Stocks. Econometrica, 54(4), 755-784.

[27] Sampat, B., and Williams, H. L. (2014). How do patents affect follow-on innovation? Evidence from the human genome. NBER Working Paper.

[28] Thompson, Peter (2006). Patent Citations and the Geography of Knowledge Spillovers: Evidence From Inventor- and Examiner-Added Citations. Review of Economics and Statistics, 88(2), 383-388.

[29] Tucker, Catherine (2015). Patent Trolls and Technology Diffusion. SSRN Working Paper.

Appendix A

More Descriptive Statistics

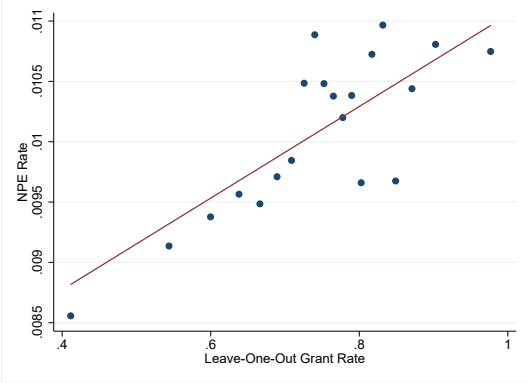
Table 20: Correlation of Blocking Actions

	101	102(a)	103(a)	112(b)
101	1.000	0.047	0.199	0.187
102(a)	-	1.000	0.104	0.086
103(a)	-	-	1.000	0.420
112(b)	-	-	-	1.000

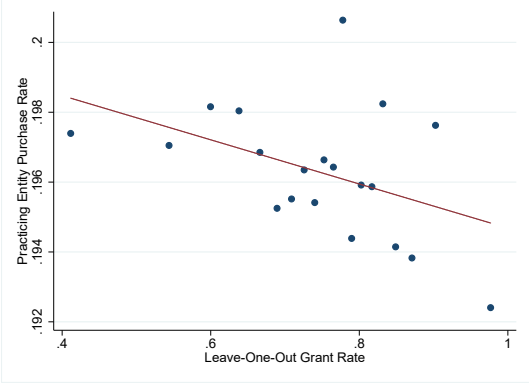
Notes: Results are computed on the baseline sample (non-continuation granted patents covered by Frakes and Wasserman).

Figure 5: Preliminary Evidence of Examiner Impact

(a) All NPE purchases



(b) Non-NPE Purchase



Notes: Similar to Figure 1, but with examiner rates computed at the year by art unit level, and with art unit by year fixed effects.

Table 21: Art unit level statistics on the 670 art units in the PAIR data. Rates are unweighted, and are computed for units with more than 50 total cases in the period

Art unit level statistics	Median	Mean	Standard Deviation	Max
Examiners	13	17.9	19.7	201
Cases Processed	2552	3961	4190	23,164
Patents Granted	1536	2669	3100	17,252
NPE Patents	7	22.2	45.9	445
Grant Rate	0.67	0.65	0.16	1
NPE Patent Rate	0.006	0.009	0.010	0.057
Use of Section 101	0.040	0.093	0.108	0.450
Use of Section 102(a)	0.015	0.018	0.013	0.093
Use of Section 103(a)	0.42	0.42	0.11	0.78
Use of Section 112(b)	0.19	0.19	0.09	0.47

Table 22: Patents Asserted in Litigation (non-NPE)

Panel A: Primary Technology Categories

NBER ID	Category Name	Patents
6	Others	4,611
2	Computers & Communications	4,175
5	Mechanical	2,859
3	Drugs & Medical	2,609
4	Electrical & Electronic	2,497
1	Chemical	1,626
-	*New Classes (since 2001)*	708

Panel B: Secondary Technology Categories

NBER ID	Subcategory Name	Patents
22	Computer Hardware & Software	2,046
69	Miscellaneous-Others	1,729
21	Communications	1,614
31	Drugs	1,009
59	Miscellaneous-Mechanical	979
19	Miscellaneous-chemical	932
32	Surgery & Med Inst.	806
-	*New Classes (since 2001)*	708
51	Mat. Proc & Handling	620
33	Biotechnology	574
65	Furniture, House Fixtures	542
61	Agriculture, Husbandry, Food	504
55	Transportation	489
62	Amusement Devices	487
68	Receptacles	471
42	Electrical Lighting	451
49	Miscellaneous-Elec	442
41	Electrical Devices	409
45	Power Systems	394
46	Semiconductor Devices	364

Table 23: NPE Patent Holdings by NBER Technology Category (RPX data)

Panel A: Primary Technology Categories

NBER ID	Category Name	Patents
2	Computers & Communications	27,156
4	Electrical & Electronic	10,660
5	Mechanical	2,709
-	*New Classes (since 2001)*	2,324
1	Chemical	1,669
6	Others	1,453
3	Drugs & Medical	1,312

Panel B: Secondary Technology Categories

NBER ID	Subcategory Name	Patents
22	Computer Hardware & Software	11,459
21	Communications	11,020
46	Semiconductor Devices	4,667
24	Information Storage	3,298
-	*New Classes (since 2001)*	2,324
49	Miscellaneous-Elec	1,626
41	Electrical Devices	1,539
23	Computer Peripherials	1,379
19	Miscellaneous-chemical	1,159
45	Power Systems	1,121
54	Optics	906
59	Miscellaneous-Mechanical	816
42	Electrical Lighting	751
69	Miscellaneous-Others	745
32	Surgery & Med Inst.	716
43	Measuring & Testing	522
44	Nuclear & X-rays	434
39	Miscellaneous-Drgs&Med	370
51	Mat. Proc & Handling	360
55	Transportation	287

Appendix B

Random Assignment and Selection Effects

Random Assignment Falsification Tests

Table 24: Random Assignment Tests in Preferred Sample

Panel A: Patent's Predetermined Outcomes

	Patent's Number of Independent Claims	Number of Words in Patent's First Independent Claim
Leave-one-out Examiner Allowance Rate	0.0020177	0.0079458
	(0.0012327)	(0.0065522)
Class-artunit-year Fixed Effects	Yes	Yes

Panel B: Inventor's Predetermined Outcomes

	Inventor's Number of Patents Granted in Previous Years	Inventor's Number of Patents Litigated in Previous Years	Inventor's Number of NPE Patents in Previous Years
Leave-one-out Examiner Allowance Rate	0.0001682	$1.50 * 10^{-6}$	0.0002451
	(0.0001975)	$(1.19 * 10^{-6})$	0.0002188
Class-artunit-year Fixed Effects	Yes	Yes	Yes

Notes: Sample excluding continuations and repeated inventor-examiner pairs. Leave-one-out examiner allowance rate adjusted for docket timing. Standard errors are clustered by examiners.

Table 25: Random Assignment Tests in Full Sample

Panel A: Patent's Predetermined Outcomes

	Patent's Number of Independent Claims	Number of Words in Patent's First Independent Claim
Leave-one-out Examiner Allowance Rate	0.0064115***	0.0057961
	(0.0012471)	(0.0066375)
Class-artunit-year Fixed Effects	Yes	Yes

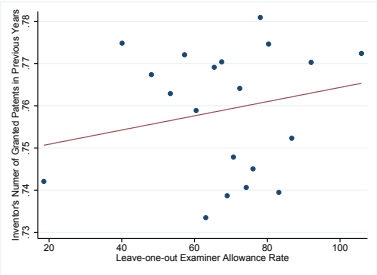
Panel B: Inventor's Predetermined Outcomes

	Inventor's Number of Patents Granted in Previous Years	Inventor's Number of Patents Litigated in Previous Years	Inventor's Number of NPE Patents in Previous Years
Leave-one-out Examiner Allowance Rate	0.0034767***	0.0000307***	0.0008802***
	(0.001084)	$(6.15 * 10^{-6})$	(0.0003454)
Class-artunit-year Fixed Effects	Yes	Yes	Yes

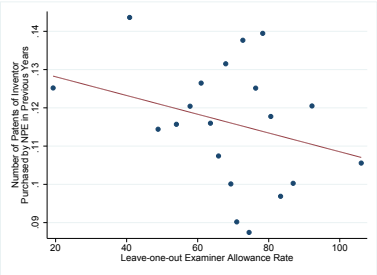
Notes: Standard errors clustered by examiners.

Figure 6: Random Assignment Tests in Preferred Sample

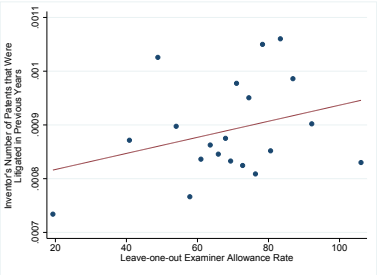
(a) Granted Patents



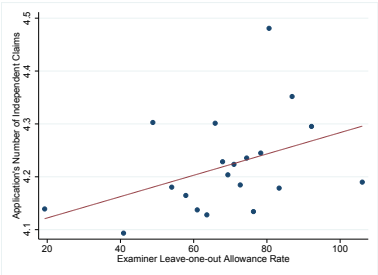
(b) NPE Patents



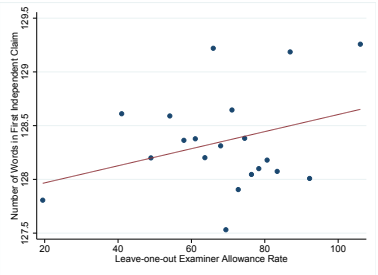
(c) Litigated Patents



(d) Independent Claims

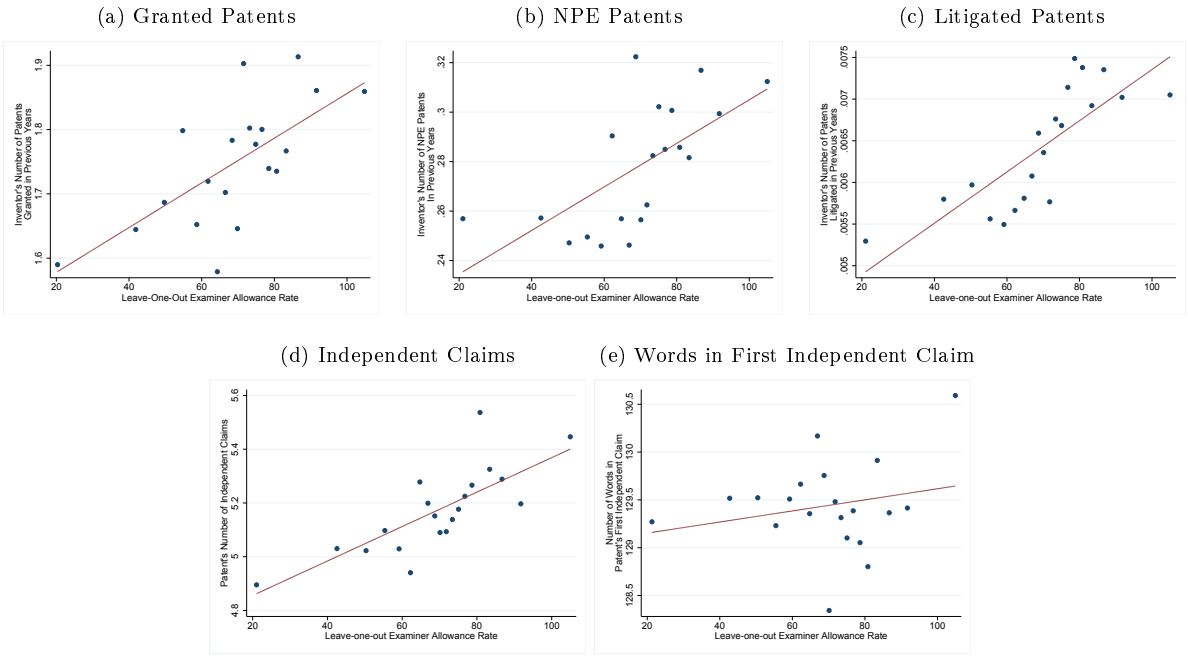


(e) Words in First Independent Claim



Notes: Sample excluding continuations and repeated inventor-examiner pairs. Leave-one-out examiner allowance rate adjusted for docket timing. Regressions include class-artunit-year fixed effects.

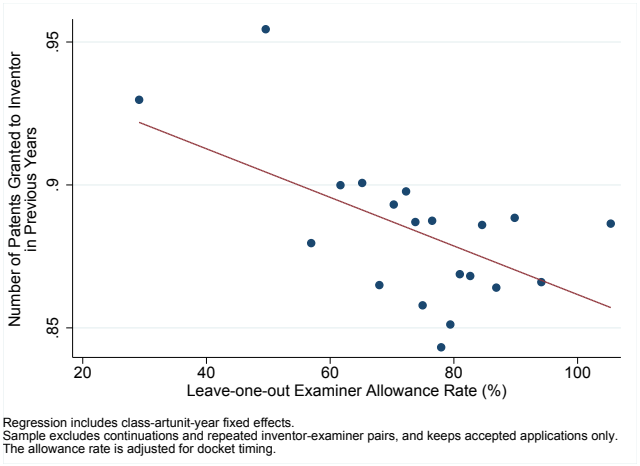
Figure 7: Random Assignment Tests in Full Sample



Notes: Regressions include class-artunit-year fixed effects.

Discussion of Selection Effects

Figure 8: Testing for Selection Effects



To understand the nature of selection effects in our setting and how we address them, consider the following framework. Within each artunit-year, each patent is defined by an examiner characteristic vector e (which includes examiner leniency, the propensity of the examiner to clarify the claims, etc...) and a patent "ex-ante"

characteristic vector p (which can be thought of as the underlying technological quality of the patent). The nature of the intellectual property is the equilibrium outcome of the interaction between the examiner and the applicant and can therefore be thought of as a flexible function of e and p. We can therefore think of e and p as the primitives that span the intellectual property space (which makes explicit the notion that the examiner effectivel co-produces the intellectual property with the applicant).

We observe the following outcomes: whether a patent is granted (G), whether a patent is purchased by a regular entity (P), whether a patent is litigated by a regular entity (L), whether a patent is purchased by an NPE (P^{NPE}), whether a patent is litigated by an NPE (L^{NPE}), and whether a patent is invalidated in court I. In a partial equilibrium framework, these outcomes can all be thought of as a (possibly stochastic) function of the primitives e and p. Most of the outcomes of interest are subject to selection effects, which can be summarized as follows:

$$G = \begin{cases} 1 \; if \; g(e,p) > 0 \\ 0 \; else \end{cases}$$

$$P = \begin{cases} 1 \; if \; p(e,p) > 0 \; \& \; G = 1 \\ 0 \; else \end{cases}$$

$$L = \begin{cases} 1 \; if \; l(e,p) > 0 \; \& \; G = 1 \\ 0 \; else \end{cases}$$

$$P^{NPE} = \begin{cases} 1 \; if \; p^{NPE}(e,p) > 0 \; \& \; G = 1 \\ 0 \; else \end{cases}$$

$$L^{NPE} = \begin{cases} 1 \; if \; l^{NPE}(e,p) > 0 \; \& \; G = 1 \\ 0 \; else \end{cases}$$

$$I = \begin{cases} 1 \; if \; i(e,p) > 0 \; \& \; L = 1 \\ 0 \; else \end{cases}$$

The definition above relate the osberved outcomes to unobserved latent variables of the form $f(e,p)$, which define mapping from the intellectual property space (spanned by e and p) to the outcomes of interest. For instance, an invalid patent is a patent for which $i(e,p) > 0$ - a well-defined mapping which can in theory be evaluated at any point of the intellectual property system. In practice, two challenges arise: first, we observe invalidation I only if the patent goes to litigation, which is itself an equilibrium outcome which depends on e and p; second, our empirical measures of e (examiner behavior) and p (underlying patent quality) are noisy.

Nonetheless, these definitions are useful to understand what we can learn from the data about the patent

acquisition behavior of NPEs and, in particular, whether they tend to purchase "weak patents." Formally, we would like to test the following:

$$H_0 : \ Cov(i(e,p), p^{NPE}(e,p)) > 0$$

where the covariance is computed over the intellectual property space. Testing whether NPEs purchase weak patents means testing whether they are active in parts of the intellectual property space that typically result in invalidations. The selection effects means that we only observe I and P^{NPE}, not the underlying mappings, and we need to find suitable measures of e and p in the data. In ongoing work, we relate our empirical findings with the theoretical covariance described above and clarify the assumptions under which the data allows us to conclude that NPEs tend to purchase weak patents.

Appendix C

Bayesian and Frequentist Interpretations of Shrinkage Methodology

Setup

Consider a setting with J examiners who are observed for T years. Assume that each examiner works in only one art unit and that each art unit has N students. The outcome of the granted patent is given by the following random effects model:

$$A_{ijt} = \mu_j + \theta_{jt} + \epsilon_{it}. \tag{8}$$

where i denotes a patent, j an examiner and t a year. A_{ijt} is the granted patent outcome (standardized by art unit-year), μ_j is an examiner effect (constant over time), θ_{jt} is an examiner-year shock, and ϵ_{it} is an idiosyncratic patent shock. We assume the following distributions: $\mu_j \sim N(0, \sigma^2_\mu)$, $\theta_{jt} \sim N(0, \sigma^2_\theta)$ and $\varepsilon_{it} \sim N(0, \sigma^2_\varepsilon)$. We are interested in estimating the vector of examiners' causal effects μ_j and using this vector for personnel policy.

Frequentist Approach

We want to forecast outcomes for patents granted by teacher j in year t using information on patent outcomes in all years prior to year t. Since all variables are demeaned, we don't need to include a constant in the set of predictors. Since there is no drift, we can use the average patent outcomes in all years prior to t, defined as $\bar{A}_{j,t-1} = \frac{1}{(t-1)N} \sum_{k=1}^{t-1} \sum_{i=1}^{N} A_{ijk}$, to predict patent outcome in year t. The optimal forecast minimizes the mean-squared error, so we must estimate α, defined in the population as:

$$\alpha = argmin_b E[(A_{ijt} - a\bar{A}_{j,t-1})^2]$$

In the random effect model above, providing an estimate $\hat{\alpha}$ of α amounts to providing estimates of examiner effects $\hat{\mu_j} = \hat{\alpha}\bar{A}_{j,t-1}$.

The first-order condition of the forecasting problem is:

$$E[2(A_{ijt}\bar{A}_{j,t-1} - \alpha\bar{A}_{j,t-1})\bar{A}_{j,t-1}] = 0$$

$$\alpha = \frac{E[A_{ijt}\bar{A}_{j,t-1}]}{E[(\bar{A}_{j,t-1})^2]} = \frac{E[(\mu_j + \theta_{jt} + \epsilon_{it}) \cdot (\mu_j + \frac{1}{t-1}\sum_{k=1}^{t-1}\theta_{jk} + \frac{1}{(t-1)N}\sum_{k=1}^{t-1}\sum_{i=1}^{N}\epsilon_{ik})]}{E[(\mu_j + \frac{1}{t-1}\sum_{k=1}^{t-1}\theta_{jk} + \frac{1}{(t-1)N}\sum_{k=1}^{t-1}\sum_{i=1}^{N}\epsilon_{ik})^2]}$$

Using independence of examiner shocks across time and patent shocks across all patents within the same art unit - year,

$$\alpha = \frac{E[\mu_j^2]}{E[\mu_j^2] + \frac{1}{t-1}E[\theta_{jk}^2] + \frac{1}{(t-1)N}E[\epsilon_{ik}^2]} = \frac{\sigma_\mu^2}{\sigma_\mu^2 + \frac{1}{t-1}\left(\sigma_\theta^2 + \frac{\sigma_\epsilon^2}{N}\right)}$$

In the sample, we can estimate α by estimating the variances σ_μ^2, σ_θ^2 and σ_ϵ^2 using the analogy principle, as described in the main text. This gives us $\hat{\alpha}$. So our estimator for μ_j is:

$$\hat{\mu_j} = \frac{\hat{\sigma_\mu^2}}{\hat{\sigma_\mu^2} + \frac{1}{t-1}\left(\hat{\sigma_\theta^2} + \frac{\hat{\sigma_\epsilon^2}}{N}\right)}\bar{A}_{j,t-1}$$

Bayesian Approach

We now derive an estimator of examiner effect in period t that is the posterior expectation of an examiner's effect given the history of patent outcomes up to period $t-1$.

Since there is no drift, we can use the average patent outcome in all years prior to t, $\bar{A}_{j,t-1}$ (defined above), as a sufficient statistic to form the posterior distribution of examiner effects. Using Bayes's rule, it is given by:

$$\pi(\mu_j|\bar{A}_{j,t-1}) = \frac{f(\bar{A}_{j,t-1}|\mu_j)\phi(\frac{\mu_j}{\sigma_\mu})}{h(\bar{A}_{j,t-1})}$$

where $f(\bar{A}_{j,t-1}|\mu_j)$ is the conditional probability density function (pdf) of $\bar{A}_{j,t-1}$ given μ_j, $\phi(.)$ is the pdf of the standard normal distribution (which we use given our prior forμ_j) and $h(\bar{A}_{j,t-1})$ is the unconditional pdf of $\bar{A}_{j,t-1}$. Since $\bar{A}_{j,t-1} = \mu_j + \frac{1}{t-1}\sum_{k=1}^{t-1}\theta_{jk} + \frac{1}{(t-1)N}\sum_{k=1}^{t-1}\sum_{i=1}^{N}\epsilon_{ik}$ and the shock are independent and normally distributed, we can write:

$$f(\bar{A}_{j,t-1}|\mu_j) = \phi\left(\frac{\bar{A}_{j,t-1} - \mu_j}{\sqrt{\frac{1}{t-1}\left(\sigma_\theta^2 + \frac{\sigma_\epsilon^2}{N}\right)}}\right)$$

So we pick the vector of teacher effects in order to maximizes the posterior probability ofμ given the data:

$$max_{\{\mu_j\}_{j=1}^J}\ \Pi_{j=1}^J \pi(\mu_j|\bar{A}_{j,t-1}) \propto \Pi_j\left[\phi\left(\frac{\bar{A}_{j,t-1} - \mu_j}{\sqrt{\frac{1}{t-1}\left(\sigma_\theta^2 + \frac{\sigma_\epsilon^2}{N}\right)}}\right)\phi(\frac{\mu_j}{\sigma_\mu})\right]$$

$$\propto \Pi_j exp(-\frac{(\bar{A}_{j,t-1} - \mu_j)^2}{\frac{1}{t-1}\left(\sigma_\theta^2 + \frac{\sigma_\epsilon^2}{N}\right)} - \frac{\mu_j^2}{\sigma_\mu^2})$$

For each j, the first-order conditions with respect to μ_j give:

$$\frac{\bar{A}_{j,t-1} - \mu_j}{\frac{1}{t-1}\left(\sigma_\theta^2 + \frac{\sigma_\epsilon^2}{N}\right)} - \frac{\mu_j}{\sigma_\mu^2} = 0$$

$$\frac{\bar{A}_{j,t-1}}{\frac{1}{t-1}\left(\sigma_\theta^2 + \frac{\sigma_\epsilon^2}{N}\right)} = \left(\frac{\frac{1}{t-1}\left(\sigma_\theta^2 + \frac{\sigma_\epsilon^2}{N}\right) + \sigma_\mu^2}{\sigma_\mu^2 \frac{1}{t-1}\left(\sigma_\theta^2 + \frac{\sigma_\epsilon^2}{N}\right)}\right)\mu_j$$

$$\mu_j = \frac{\sigma_\mu^2}{\frac{1}{t-1}\left(\sigma_\theta^2 + \frac{\sigma_\epsilon^2}{N}\right)} \bar{A}_{j,t-1}$$

We can estimate the variances σ_μ^2, σ_θ^2 and σ_ϵ^2 using the analogy principle, as described in the main text, which yields an estimator that is identical to the frequentist estimator derived earlier:

$$\hat{\mu_j} = \frac{\hat{\sigma_\mu^2}}{\hat{\sigma_\mu^2} + \frac{1}{t-1}\left(\hat{\sigma_\theta^2} + \frac{\hat{\sigma_\epsilon^2}}{N}\right)} \bar{A}_{j,t-1}$$

Intuitively, the Bayesian estimator is a precision-weighted average of the mean patent outcome for the patents granted by examiner j and the prior mean for μ_j, which is equal to 0. In other words, the mean test score observed in the data is shrunk towards 0 (empirical Bayes). As $\sigma_\theta^2 \to 0$ and $N \to \infty$, $\hat{\mu_j} \to \bar{A}_{j,t-1}$.

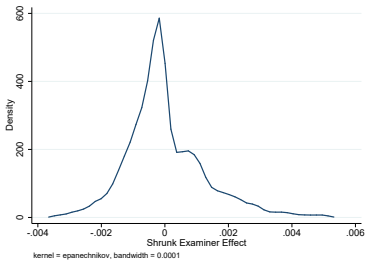
Appendix D

Additional Results on The Causal Effect of Patent Examiners

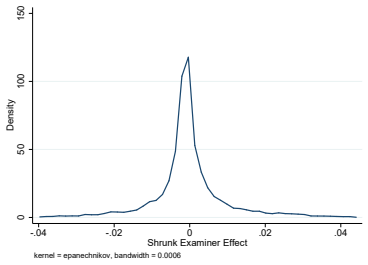
Causal Effect of Patent Examiners on Other Outcomes

Figure 9: Distributions of examiner effects and tendencies by various outcomes observed in our dataset.

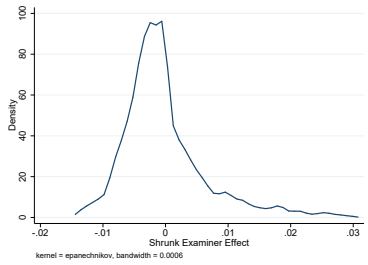
(a) Non-NPE Litigated Patents



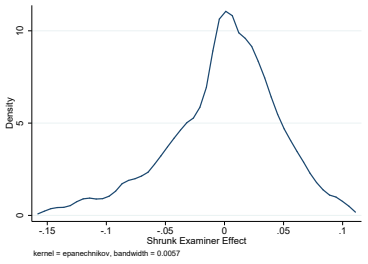
(b) 101 Rejections



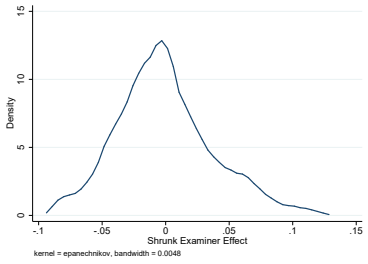
(c) 102(a) Rejections



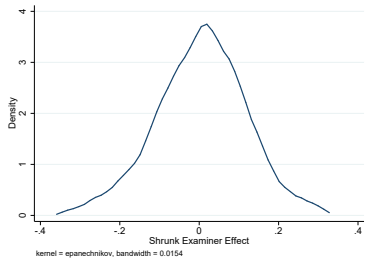
(d) 103(a) Rejections



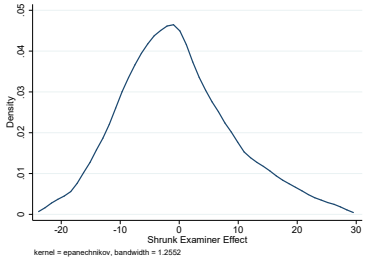
(e) 112(b)



(f) Propensity to Add Independent Claims



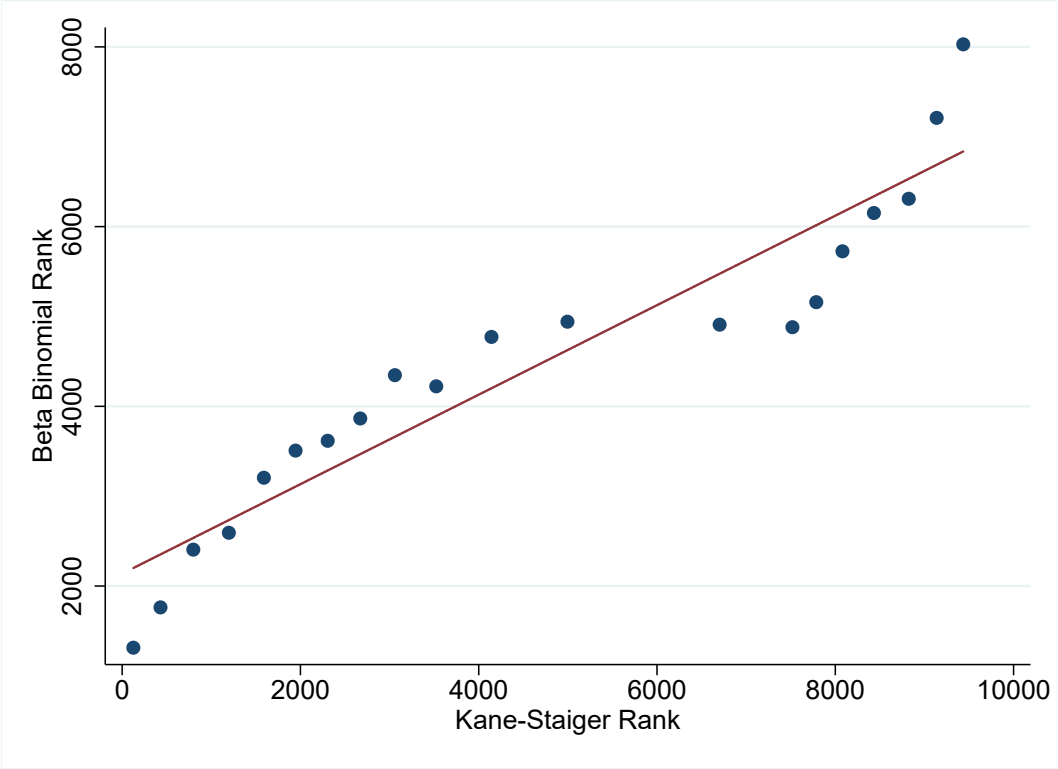
(g) Propensity to Add Words to Independent Claims



Notes: Distribution of examiner effects for other variables, computed in an analogous manner to Figure 2.

Comparison Between Kane-Staiger and Beta-Binomial Approaches

Figure 10: Comparison of Examiner Ranks



Notes: Comparison of examiner ranks, based on deviations computed through the Kane-Staiger methodology and through the Beta-Binomial methodology.

Blocking Action Robustness Check

In this section, we present additional robustness checks to the results shown in Section 5.2, by computing examiner blocking action causal effects using data from all applications they examine. This contrasts with our baseline specification, which only looks at blocking actions used on eventually granted patents. As shown in Table 26, the results are generally very similar to the ones reported in Table 10.

Signal Correlations

In this section, we present pairwise results using the signal correlation framework used in Chetty and Hendren (2015). The idea here is to split the sample of applications into two, and then compute the correlation of examiner effects across the samples. The signal correlation between examiner effects for variables X and Y is computed as follows:

$$\rho_{XY} = \frac{cov(E_{X1}, E_{Y2})}{\sqrt{cov(E_{X1}, E_{X2}) \cdot cov(E_{Y1}, E_{Y2})}}$$

Table 26: Pairwise Examiner Alternate Blocking Action Effects Results

	NPE Purchase	Non-NPE Purchase	NPE Litigated	non-NPE Litigated
101	-0.051*	0.003	-0.126*	-0.052**
	(0.022)	(0.004)	(0.064)	(0.018)
102(a)	0.013	0.003	-0.002	-0.011
	(0.023)	(0.005)	(0.051)	(0.018)
103(a)	-0.098***	-0.016**	-0.255***	-0.065**
	(0.024)	(0.005)	(0.064)	(0.020)
112(a)	-0.015	0.002	-0.128**	-0.054*
	(0.018)	(0.004)	(0.046)	(0.025)
112(b)	-0.047*	-0.003	-0.130*	-0.043*
	(0.021)	(0.005)	(0.057)	(0.019)
N	1,269,623	1,269,623	1,269,623	1,269,623

Notes: Similar to Tables 10 and 11, but computing examiner blocking action causal effects using their blocking actions on all applications. *+ p-value <0.10, * p-value < 0.05, ** p-value < 0.01, *** p-value < 0.001*

The idea here is to look for correlation of examiner effects on different variables across samples, normalized by the consistency of the same-variable examiner effects across the samples. The method is similar to our previous pairwise predictive regressions, but incorporates re-weighting of data points.

We report the results in Table 27, performing a bootstrap routine at the examiner level to construct confidence intervals.

Additional Legal Outcomes

In this section, we present results by predicting additional legal outcomes using examiner causal effects. The data is constructed from raw data provided by LexMachina.

Application-Level Analysis

We present analogous results to the ones in Sections 4 and 5, but counting outcomes for rejected applications as zero. The results here capture both the intensive and extensive margin effects of an examiner. As we see in Table 29, the extensive margin effects push many of the pairwise coefficients upwards in magnitude relative relative to the results reported in Table 10.

Instrumental Variable Analysis

As discussed earlier, the most commonly-used methodology in the examiner literature has been an instrumental variables approach, using the leave-one-out examiner grant rate as an instrument for allowance. This framework does not address the question we are interested in here, which is the effect of an examiner of the

Table 27: Pairwise Examiner Alternate Blocking Action Effects Results

	NPE Purchase	Non-NPE Purchase	non-NPE Litigated
Word Count	-0.095*	0.011	-0.067*
	[-0.173,-0.041]	[-0.021,0.045]	[-0.151,-0.016]
101	-0.055*	0.043*	-0.054*
	[-0.121,-0.004]	[0.009, 0.076]	[-0.124,-0.011]
102(a)	0.028	0.027	-0.011
	[-0.017,0.082]	[-0.011,0.065]	[-0.063,0.040]
103(a)	-0.049*	0.037*	-0.028
	[-0.104,-0.008]	[0.008,0.067]	[-0.086,0.014]
112(a)	-0.002	0.047*	-0.054
	[-0.038, 0.032]	[0.013,0.083]	[-0.086,0.014]
112(b)	-0.027	0.026	-0.024
	[-0.073,0.011]	[-0.004, 0.060]	[-0.080,0.019]
N	1,269,623	1,269,623	1,269,623

Notes: 95 percent confidence intervals are shown in the brackets. ** p-value $<$ 0.05*

nature of the patent granted. Instrumental variables analysis is akin to performing the analysis in Section 4, without the extra analysis that we do in Section 5. In this section, we discuss the strong connection between the approaches, and also calculate the results using an instrumental variables approach.

Framework

The basic approach here is a standard instrumental variables analysis at the application level, instrumenting for approval using the leave-out-mean of examiner leniency. Formally, the first and second stage regressions are:

$$T_{ijat} = \beta_0 + \beta_1 Patented_{ijat} + Controls + \epsilon_{ijat}$$

$$Patented_{ijat} = \gamma_0 + \gamma_1 Z_{ijat} + Controls + \eta_{ijat}$$

where i denotes application, j denotes examiner, a denotes art unit, t denotes year of application. The reduced form regression looks like:

$$T_{ijat} = \alpha_0 + \alpha_1 Z_{ijat} + Controls + \zeta_{ijat} \tag{9}$$

The examiner leave-out-mean computed at the year-art unit level, akin to accounting for cohort effects,

Table 28: Pairwise Examiner Effect Results for Other Legal Outcomes

Panel A: District Court Litigation Outcomes

	Completed Litigation	Trial | Completed Lit.	Infringe | Trial	Invalid | Trial
NPE Purchase	0.050^{+}	0.015	-0.122**	0.103
	(0.026)	(0.036)	(0.045)	(0.088)
Words Per Claim Change	-0.195***	-0.010	0.057	0.044
	(0.018)	(0.040)	(0.054)	(0.076)
101	-0.068***	-0.009	-0.070	0.040
	(0.017)	(0.037)	(0.046)	(0.074)
102(a)	-0.001	0.033	-0.035	-0.134^{+}
	(0.019)	(0.045)	(0.058)	(0.081)
103(a)	-0.151***	-0.020	-0.010	0.036
	(0.021)	(0.040)	(0.057)	(0.080)
112(a)	-0.080***	0.088*	0.026	-0.009
	(0.017)	(0.04)	(0.049)	(0.071)
112(b)	-0.083***	-0.011	-0.059	-0.043
	(0.020)	(0.040)	(0.054)	(0.078)
N	1,867,760	5,319	595	595

Panel B: Inter-Partes Review Outcomes

	IPR Filed	IPR Instituted | Filed	Final Decision | Instituted	All Claims Invalid | Final D
NPE Purchase	0.116*	-0.005	0.113*	0.015
	(0.055)	(0.019)	(0.054)	(0.042)
Words Per Claim Change	-0.246***	-0.015	-0.014	-0.019
	(0.037)	(0.018)	(0.056)	(0.043)
101	-0.112**	-0.035*	-0.071	-0.036
	(0.040)	(0.017)	(0.054)	(0.034)
102(a)	0.015	-0.022	0.032	-0.070
	(0.041)	(0.023)	(0.066)	(0.053)
103(a)	-0.167***	-0.028	-0.022	-0.011
	(0.043)	(0.022)	(0.065)	(0.042)
112(a)	-0.103**	-0.023	-0.058	0.012
	(0.034)	(0.018)	(0.052)	(0.031)
112(b)	-0.094*	-0.045*	-0.061	-0.019
	(0.039)	(0.020)	(0.057)	(0.041)
N	1,867,760	577	467	176

Notes: Similar to Table 10, but expanding the outcome dataset to incorporate all data from the PERD dataset, by assigning shrunk examiner effects to applications outside of the Frakes and Wasserman sample, while still using leave-one-out shrunk effects for applications within the Frakes and Wasserman sample. *+ p-value <0.10, * p-value < 0.05, ** p-value < 0.01, *** p-value < 0.001*

Table 29: Pairwise Examiner Blocking Action Effects Results (Application Level)

	NPE Purchase	Non-NPE Purchase	NPE Litigated	non-NPE Litigated
101	-0.071**	-0.018***	-0.143*	-0.071***
	(0.022)	(0.005)	(0.062)	(0.017)
102(a)	0.005	-0.005	-0.009	-0.018
	(0.024)	(0.006)	(0.052)	(0.017)
103(a)	-0.194***	-0.110*	-0.358***	-0.160***
	(0.026)	(0.006)	(0.068)	(0.020)
112(a)	-0.064***	-0.047***	-0.176***	-0.103***
	(0.019)	(0.006)	(0.046)	(0.025)
112(b)	-0.095***	-0.052***	-0.179**	-0.092***
	(0.023)	(0.006)	(0.057)	(0.019)
N	1,269,623	1,269,623	1,269,623	1,269,623

Notes: Similar to Tables 10 and 11, but computing examiner blocking action causal effects using their blocking actions on all applications. *+ p-value <0.10, * p-value < 0.05, ** p-value < 0.01, *** p-value < 0.001*

which in the teacher effects framework is captured by the θ_{jt} error term.

$$Z_{ijat} = \frac{n_{jat}^{grant} - Patented_{ijat}}{n_{jat}^{appl} - 1}$$

where n_{jat}^{grant} represents the number of granted patents for examiner j in art unit a in year t, and n_{jat}^{appl} represents the corresponding number for applications. More generally, we can instrument for examiner tendencies in a similar manner, replacing the *Patented* variable with an indicator for the usage of a given provision such as 103(a) and 112(b), and then instrumenting for the actual usage with a leave-one-out measure of average examiner usage of the given provision.

Comparison of Instrumental Variables to Kane-Staiger Framework

The instrumental variables (IV) framework has many similarities to the Kane-Staiger framework we applied in Sections 4 and 5.

The analysis in Section 5 is very similar in spirit to the reduced form IV specification (Equation 9), as we leave out the current data point in computing examiner effects (a measure of examiner leniency in the IV setup), and then use it to predict an outcome T, such as NPE purchase. In the context of the IV-2SLS approach, the resulting coefficient α_1 is then scaled to by $\frac{1}{\gamma_1}$ in order to normalize the effect size, although in the case of the patent examiner instrument, γ_1 is often pretty close to 1. In addition, the motivation for computing Z at the year level rather than just at the examiner level is driven by the possibility of examiner by cohort effects. This is accounted for in our methodology through the θ_{jt} term.

There are a couple of major differences between the frameworks. First, the IV analysis is performed at the application level rather than the granted patent level. Therefore the reduced form coefficient α_1 will capture both the extensive margin (stricter examiners reject more applications, which then cannot be bought buy an NPE) and the intensive margin (stricter examiners also appear to force more edits to patent applications). Second, our mechanism analysis in Section 5 uses examiner effect measures computed across all years, rather than within a single year, reducing some of the noise in the measurement.

Results using IV Methodology

Our preferred specification is to include patent applications that were examined by examiners with 10 or more cases in a given year and art unit, which covers over 91% of the applications in our sample. We also run robustness checks with 5, 15 and 25 as the cutoff. Panel A of Table 30 contains the first stage estimates. Consistent with the results reported in Sampat and Williams (2015), there is a very strong first stage relationship between the examiner leave-one-out grant rate and the decision on a given patent application. Some of the relationship is driven by year and art unit effects, but even after controlling for these, there is a precisely estimated coefficient of around 0.75.

Next, in Panel B of Table 30, we report the coefficients from the two-stage least squares instrumental variables estimation. Here, we report the coefficients in a similar manner to the teacher effects results, by normalizing the coefficients to represent the effects of a one standard deviation change in examiner grant rate on NPE outcomes, normalized to the baseline NPE rate. Note that ex-ante, the reduced form relationship between NPE purchase and approval will be at least weakly monotonically increasing, because the NPE purchase vector is component-wise weakly less than the "Decision" variable. Due to the re-scaling, the OLS coefficient essentially captures the standard deviation in examiner grant rate, because the raw coefficient absent the scaling would just be the baseline rate of NPE purchasing, as rejected applications have a zero NPE purchase rate. The 2SLS coefficients capture the reduced form relationship between NPE purchase rates and examiner leave-one-out grant rate, scaled by the first stage coefficient. The fact that the 2SLS coefficient is higher than the OLS coefficient can be interpreted as evidence that the pool of NPE purchased patents are more "marginal," in the sense that they are more sensitive to examiner tendencies. A final point to note is that the coefficients are of similar magnitude to the ones found in the examiner effect distribution analysis. A one standard deviation change in examiner grant rate leads to an NPE effect of around 30% of the baseline effect.

Table 30: Instrumental Variables Results

Panel A: First Stage Results

Decision	(1)	(2)	(3)
Examiner Grant Rate	0.885***	0.752***	0.726***
	(0.00158)	(0.00185)	(0.00190)
Art Unit by Year FE		x	x
Class FE			x
N	2,478,697	2,478,697	2,478,697
R^2	0.180	0.186	0.189

Panel B: Examiner Impact on NPE Purchase

NPE Purchase	(1) OLS	(2) OLS	(3) 2SLS	(4) 2SLS
Decision	0.284***	0.267***	0.387***	0.348***
	(0.00575)	(0.00549)	(0.0163)	(0.0184)
Art Unit by Year FE		x		x
N	2,478,697	2,478,697	2,478,697	2,478,697
R^2	0.003	0.018	0.003	0.002

Notes: Decision refers to the final disposal outcome of a given patent application. Examiner Grant Rate refers to the leave-one-out grant rate of an examiner in a given art uni and year. Analysis is restricted to examiners with more than 10 cases processed in a given year. Analysis is run on the entire Patent Examination Research Dataset. All standard errors are clustered at the examiner level. **p-value $<$ 0.05, ** p-value $<$ 0.01, *** p-value $<$ 0.001.*

We also run similar analyses with the Frakes and Wasserman blocking action variables. The results are reported in Table 31. The results are consistent with our earlier reported results. The usage of a given provision is strongly associated with usage of the provision on other applications (First Stage column). In addition, we see major discrepancies between the OLS and 2SLS results. This should not be entirely surprising, given our earlier summary statistics results (Table 5) and mechanism results. NPE purchased patents are almost equally like to receive 103(a) rejections and are more likely to receive 112(b) rejections (OLS results). However. the causal examiner contribution to this has the opposite sign, consistent with the result that examiners that tend to use more 112(b) blocking actions have fewer NPE purchased patents. As discussed earlier, the 2SLS coefficient also incorporates the direct effect of blocking actions on abandonments, so some applications mechanically cannot end up as NPE purchases.

This methodology could be useful for inferring the causal effects of specific blocking action types on outcomes, especially in cases where we can associated outcomes to ungranted patents, as in Sampat and Williams (2015) in relation to gene-based patent applications and subsequent research. However, the limitation is that if there is some correlation between usage of these provisions for a given examiner, and the instrumental variables results may pick up mechanisms associated with other provisions (it violates of the

exclusion restriction). One could then think about using a more general instrumental variables framework with multiple predictors and multiple instruments, which would solve some of these problems as long as one assumes additive effects.

Table 31: Blocking Action Instrumental Variables Results

NPE Purchase	First Stage	OLS	2SLS
101	0.482***	-0.000547	-0.188***
	(0.00282)	(0.00498)	(0.0486)
102(a)	0.434***	0.00171	0.00137
	(0.00337)	(0.00222)	(0.0322)
103(a)	0.710***	-0.00919*	-0.217***
	(0.00248)	(0.00456)	(0.0297)
112(b)	0.690***	0.0103**	-0.130***
	(0.00252)	(0.00377)	(0.0247)
Art Unit by Year FE	x	x	x
N	1,752,641	1,752,641	1,752,641

Notes: Decision refers to the final disposal outcome of a given patent application. Examiner Grant Rate refers to the leave-one-out grant rate of an examiner in a given art uni and year. Analysis is restricted to examiners with more than 10 cases processed in a given year. Analysis is run on the Frakes and Wasserman coverage range. All standard errors are clustered at the examiner level. **p-value < 0.05, ** p-value < 0.01, *** p-value < 0.001.*

Appendix E

Additional Results on Mechanisms

Table 32: Results across technology centers for key predictors of NPE purchase and non-NPE litigation.

	NPE Purchase		Non-NPE Litigation	
Specification	103(a)	112(b)	103(a)	112(b)
1600 - Biotechnology and Organic Chemistry	-0.173	0.078	-0.036	0.002
	(0.120)	(0.077)	(0.068)	(0.067)
1700 - Chemical and Materials Engineering	0.072	0.187	0.025	0.052
	(0.130)	(0.173)	(0.037)	(0.035)
2100 - Computer Architecture, Software, and Information Security	-0.064**	-0.036	-0.124**	-0.121**
	(0.023)	(0.026)	(0.005)	(0.038)
2400 - Computer Networks, Multiplex communication, Video Distribution, and Security	-0.016	0.024	-0.119^{+}	0.003
	(0.029)	(0.030)	(0.062)	(0.050)
2600 - Communications	-0.098***	-0.061**	-0.010	-0.070*
	(0.025)	(0.023)	(0.030)	(0.031)
2800 - Semiconductors, Electrical and Optical Systems and Components	-0.062**	-0.036^{+}	-0.071**	-0.013
	(0.022)	(0.021)	(0.024)	(0.024)
3600 - Transportation,Construction, Electronic Commerce, Agriculture, National Security...	-0.049	-0.062	-0.070*	-0.069*
	(0.039)	(0.039)	(0.029)	(0.027)
3700 - Mechanical Engineering, Manufacturing, Products	0.034	-0.046	-0.001	-0.002
	(0.040)	(0.043)	(0.029)	(0.028)

Notes: Same specifications as Tables 10 and 11, but restricting the sample to each technology center. Baseline NPE purchase and non-NPE litigation rates are re-computed for each technology center. $+$ *p-value <0.10,* $*$ *p-value < 0.05,* $**$ *p-value < 0.01,* $***$ *p-value < 0.001*

Figure 11: Examiner Effects Do Not Predict Litigation in Orange Book Subsample

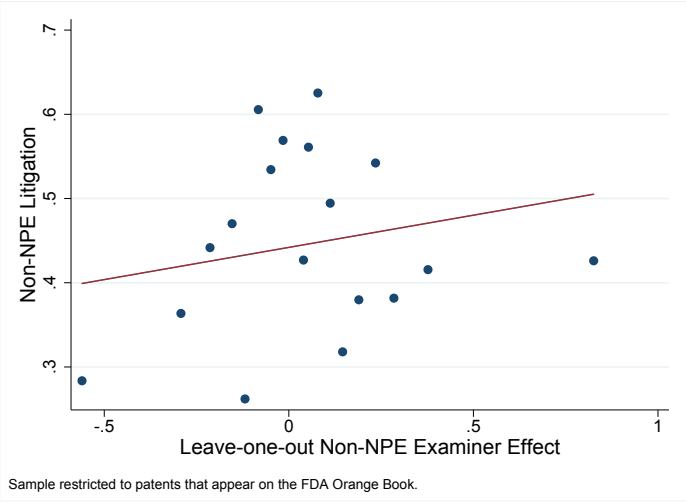
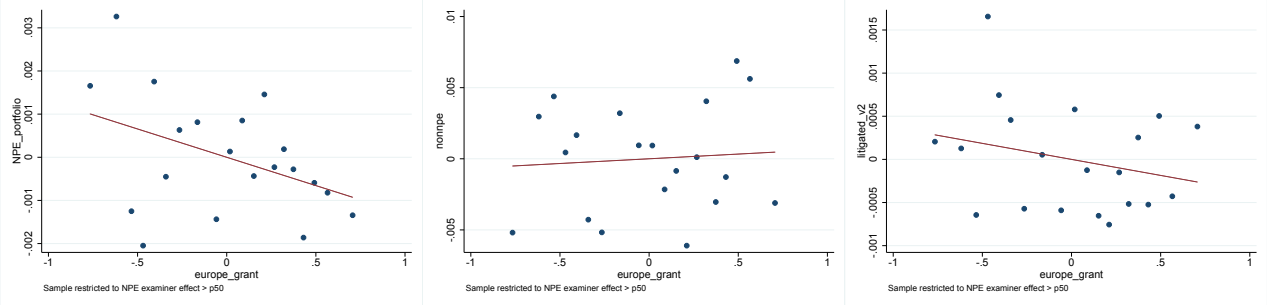


Figure 12: Results of EPO Analysis

Panel A: Sample with NPE Examiner Effects above Median



Panel B: Sample with NPE Examiner Effects below Median

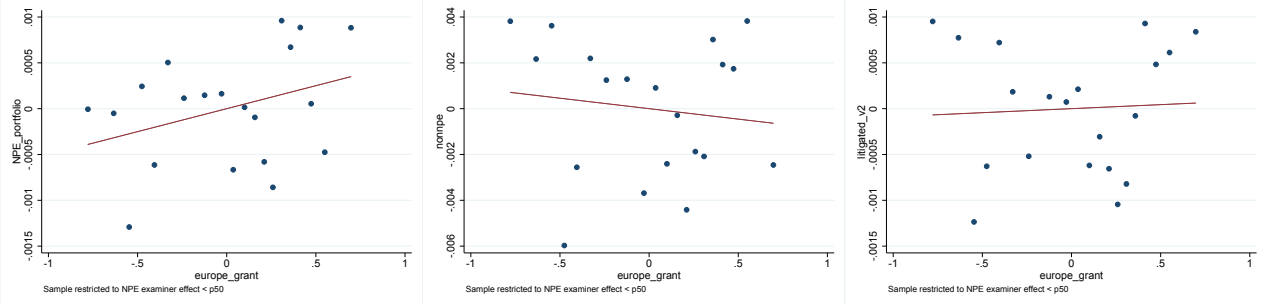
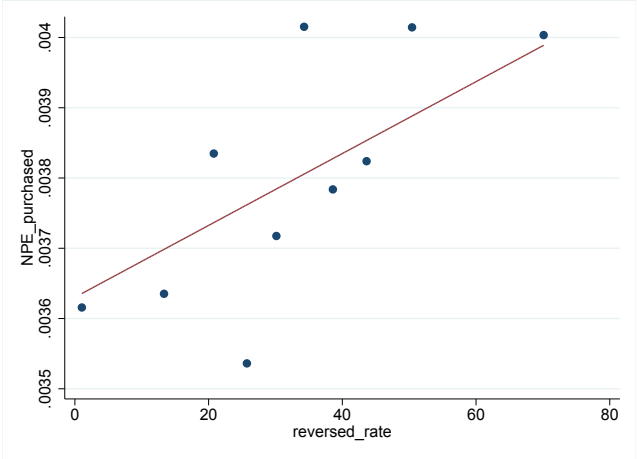


Figure 13: NPEs Tend to Purchase Patents from Examiners Who Are Reversed at PTAB



Appendix F

Data Appendix

Name Matching

In this part, we discuss our procedure for creating NPE portfolios. The algorithm proceeds as follows:

1. We start off with a list of NPE names, either from RPX or from Cotropia et al

2. We normalize entity names from both the NPE list and the USPTO Assignment Database. This is done by capitalizing all names, removing punctuation, and removing standard entity terms: INC, CO, COMPANY, COMPANIES, CORP, CORPORATIONS, DIV, GMBH, LLC, LC, INCORPORATED, KG, LIMITED, LIMITED PARTNERSHIP, LP, LTD, NV, PLC, SA, SARL, SNC, SPA, SRL, TRUST USA, CENTER, BV, AG, AB, GROUP, FOUNDATION, INSTITUTE, and TECHNOLOGIES.

3. We then collect Reel/Frame IDs of patent transactions in the USPTO Assignment Database that have a normalized entity name matching the normalized name of an NPE

4. Each Reel/Frame ID is classified in the USPTO data based on the type of transaction and whether the assignment was to an employer (essentially the first assignment). We keep transactions that are non-employer assignments. This then gives us a set of patents involved in patent purchasing, and excludes other types of transactions such as securitization, mergers, and name changes.

5. Finally, we collect the list of patents associated with each of these transactions to create our portfolios.