

**Are All Patent Examiners Equal?  
The Impact of Examiners on Patent Characteristics  
and Litigation Outcomes\***

Iain Cockburn  
Boston University and NBER

Samuel Kortum  
University of Minnesota and NBER

Scott Stern  
Northwestern, Brookings, and NBER

February 2002

\* This paper was prepared for the National Academy of Sciences STEP Board Conference on *The Operation of the Patent System: Insights from New Research*. We thank USPTO personnel for offering their time and insight, and members of the STEP Committee on the Intellectual Property Rights in a Knowledge-Based Economy, including Wes Cohen and George Elliott, for their comments and suggestions. Tariq Ashrati provided excellent research assistance. All errors, however, remain our own. Financial support for this research was provided by the National Academy of Sciences. Contact author: Scott Stern, Kellogg Graduate School of Management, 2001 Sheridan Road, Evanston, IL 60208 [s-stern2@northwestern.edu](mailto:s-stern2@northwestern.edu).

# **Are All Patent Examiners Equal? The Impact of Examiner Characteristics on Patent Statistics and Litigation Outcomes**

## **ABSTRACT**

We conduct an empirical investigation, both qualitative and quantitative, on the role of patent examiner characteristics in the allocation of intellectual property rights. Building on insights gained from interviewing administrators and patent examiners at the United States Patent and Trademark Office (USPTO), we collect and analyze a novel dataset of patent examiners and patent outcomes. This dataset is based on 182 patents for which the Court of Appeals for the Federal Circuit (CAFC) ruled on validity between 1997 and 2000. For each patent, we identify a USPTO primary examiner, and collect historical statistics derived from their entire patent examination history. These data are used to explore a number of hypotheses about the connection between the patent examination process and the strength of ensuing patent rights. Our main findings are as follows. (i) Patent examiners and the patent examination process are not homogeneous. There is substantial variation in observable characteristics of patent examiners, such as their tenure at the USPTO, the number of patents they have examined and the degree to which the patents that they examine are later cited by other patents. (ii) There is no evidence that examiner experience or workload at the time a patent is issued affects the probability that the CAFC finds a patent invalid. (iii) Examiners whose patents tend to be more frequently cited tend to have a higher probability of a CAFC invalidity ruling. The results suggest that all patent examiners are not equal, and that one of the roles of the CAFC is to review the exercise of discretion in the patent examination process.

## I. Introduction

Recent years have seen a worldwide surge in interest in intellectual property rights, particularly patents, in academia, in policy circles, and in the business community. This heightened level of interest has produced a substantial body of research in economics ranging from analyses of decisions to use patents rather than alternative means of protecting intellectual property (Cohen, et al, 1999) to studies of the ways in which patents are used and enforced once granted (see, for example, Hall and Ziedonis, 2001; Launjouw and Lerner, 1997). There has, however, been little systematic attention paid to the *process* of how patent rights are created.

Indeed, only recently have researchers begun to develop a systematic understanding of the differences in intellectual property regimes across countries and over time (Lerner, 2000). Moreover, except for some preliminary aggregate statistics (Griliches, 1984; 1990), there are no published studies of the empirical determinants of patent examiner productivity, or of linkages between characteristics of patent examiners and the subsequent performance of the patent rights that they issue.<sup>1</sup> Here, we begin to address these questions by evaluating the role that variation in examination procedures and in the exercise of discretion by individual patent examiners may play in determining the allocation of patent rights, and their subsequent testing by the courts.

Filling in this gap in our knowledge may yield a number of benefits. First, and perhaps most importantly, it is difficult to assess the likely impact of changes in the funding or operation of the United States Patent and Trademark Office (USPTO) without some understanding of the “USPTO production function.” For example, at various points in the past there have been shifts in the resources available to the USPTO as well as in the incentives and objectives provided to examiners, recently focused on reducing the time taken between initial filing of a patent application and final issuance. At the same time, court rulings and revisions in USPTO practice have broadened intellectual property protection into new areas, such as genomics and business methods, where the novelty and obviousness of inventions and the scope of claims awarded may be difficult, at least initially, to assess. Developments such as these raise a variety of difficult

---

<sup>1</sup> King (this volume) offers a complementary examination to the one conducted here, where he undertakes a detailed analysis of the impact of resource allocation per se on “art unit” performance, while our quantitative research focuses on how individual examiner characteristics might impact litigation outcomes. The literature on the use of patent statistics and the impact of patents on innovation is far too large to be summarized here, but see Levin et al (1986),

questions for policymakers. How might changes in the incentives for timeliness and patent breadth impact the overall quality of the patent system? Is there a historical tradeoff between quality and quantity? To what extent does quality reflect systematic differences across examiners, above and beyond the effects of training, experience or tenure?

Our analysis include both a qualitative and quantitative component. In the first part of the paper, we review our qualitative investigation, in which we develop an informal understanding of the process of patent examination and investigate potential areas for differences among patent examiners to impact policy-relevant measures of the performance of the patent system. The key insight from our qualitative analysis is that “there may be as many patent offices as patent examiners.” There is reason to suspect substantial heterogeneity among examiners, and to suspect that this heterogeneity affects the outcome of the examination process. This insight motivates the development of the key hypotheses in Section III.

To test these hypotheses, we construct a novel dataset that links information from USPTO “front page” information with data based on the US Court of Appeals for the Federal Circuit (CAFC) record between 1997 and 2000. We consider a sample of 182 patents on which the CAFC issued a ruling on validity during this period. For each patent, we identify the primary and secondary examiners associated with the patent, and collect the complete set of issued patents issued by that examiner during their tenure at the USPTO. We then construct measures based on this examiner-specific patent collection, including their experience with examination, workload, and measures based on the citation patterns associated with issued patents..

We present our key findings in several steps. First, we show that patent examiners differ on a number of observable characteristics, including their overall experience at the USPTO (both in terms of years as well at total number of issued patents), their degree of technological specialization, their propensity to cite their own patents, and their propensity to issue patents that are highly cited. Indeed, a significant portion of the overall variance among patents in measures such as the number and pattern of citations received, the number and pattern of citations made, and the approval time can be explained by the identity of the examiner — in the language of econometrics, “examiner fixed effects.” These examiner effects are significant even after controlling for the patent’s technology field and its cohort (i.e., the year the patent was issued).

---

Griliches (1990), Cohen et al (2000) and Hall et al (2001) for an introduction.

We then turn to an examination of whether observable characteristics of our sample of patents tested by the CAFC such as their citation rate or approval time can be tied to observable examiner characteristics such as their experience or the rate at which “their” patents receive citations. Here we find strong evidence for the impact of examiners. For example, there is a significant positive relationship between the citations received by a subsequently litigated patent and the “propensity” of its examiner to issue patents that attract a large numbers of citations. This result leads to the core of the analysis — tying these relationships to patent validity rulings. The results of our econometric analysis are striking. Although the outcome of a test of validity by the CAFC is unrelated to the number of citations received by that particular patent, validity is strongly related to the portion of the citation rate explained by the examiner’s idiosyncratic propensity to issue patents which receive a high level of citations. To the extent that the examiner-specific citation rate reflects the “generosity” of examiners in allowing claims, our empirical findings suggest that one of the roles of the CAFC is to review the discretion exercised in the patent examination process.

The remainder of the paper is organized as follows. In the next section, we review our qualitative data gathering, and motivate the evidence for our key testable hypotheses, which we state in Section III. Section IV describes the novel dataset we have constructed, and Section V reviews the results. A final section offers a discussion of our findings as well as identifying areas for future empirical research in this area.

## **II. The Patent Examination Process**

### *Methodology*

This section reviews the initial stage of our research, a qualitative investigative phase where we sought to understand the process of patent examination and the potential role of patent examiner characteristics on that process. This type of investigation is precisely what has been lacking from much academic and policy discussion of the impact of patent office practices, procedures, and personnel on the performance of the intellectual property rights system. While practitioners and USPTO personnel are intimately acquainted with these procedures, there has been little attempt to identify which aspects of the examination process can be linked through

rigorous empirical analysis to the key policy challenges facing USPTO.

Our qualitative research involved three distinct stages. First, we informally interviewed a small number of former patent examiners and patent attorneys, outside the USPTO, in order to develop a basic grounding in the process and procedures of the USPTO. This allowed us to evaluate and understand some of our initial hypotheses about the impact of patent examiner characteristics and USPTO practice on the ultimate allocation of intellectual property rights. We used this working knowledge to develop a proposal to undertake systematic interviews within the USPTO; and with the assistance of the NAS STEP Board, we engaged in several meetings with senior USPTO managers to evaluate the possibility of administering a survey which would allow us to link detailed information about examiner history (e.g., educational and employment background) with information that could be gleaned from patent statistics about differences among patent examiners (e.g., differences in approval times or citation statistics). This approach was not entirely successful as we were unable to obtain approval to distribute a systematic survey of our own design to a broad cross-section of current and former examiners. However, USPTO management did generously allow us to undertake several visits in which we were allowed to conduct informal interviews and question-and-answer sessions with a small number of examiners, mostly limited to those in a supervisory role. These conversations were very helpful in developing a number of more subtle, precise, and econometrically testable hypotheses about how the practice and procedures of the USPTO might impact the allocation of intellectual property rights. This leads to the third stage of qualitative research, where we confirmed the viability of our hypotheses with individuals external to the USPTO. Overall, our qualitative research phase included interviews with approximately 20 current or former patent examiners as well as an equal number of patent attorneys with considerable experience in patent prosecution.

### *The Examination Process*

Here we describe the patent examination process in general terms, focusing on the aspects where we identified potential sources of heterogeneity in examination practice. The USPTO is one of the earliest and among the most visible agencies of the Federal government, receiving more certified mail per day than any other single organization in the world. Located in a single campus of connected buildings, the USPTO is staffed by over 3000 patent examiners,

and has more than 6,000 total full-time equivalent employees. In recent years the examiner corps has been responsible for over 160,000 patent approvals per year. The Federal government raises nearly \$1 billion in revenue per year from the fees and other revenue streams associated with the USPTO.

The workflow and procedures associated with patent approval are quite systematic and well-determined.<sup>2</sup> After being received at a central receiving office, and passing basic checks to qualify for a filing date, patent applications are sorted by a specialized classification branch<sup>3</sup> which allocates them to one of 235 “Art Units” — a group of examiners who examine closely related technology, and constitute an administrative unit. Within the Art Unit, a “Supervisory Patent Examiner” (a senior examiner with administrative responsibilities) looks at the technology claimed in the application and assigns it to a specific examiner. Once the patent is allocated to a given examiner, that examiner will, in most cases, have continuing responsibility for examination of the case until it is disposed of – either through rejection, allowance, or discontinuation. The examination process therefore typically involves an interaction between a single examiner and the attorneys of the inventor or assignee. While the stages associated with this process are relatively structured (and exhaustively documented in the Manual of Patent Examining and Procedure) they leave substantial discretion to the examiner in how to deal with a particular application.

The examination of an application begins with a review of legal formalities and requirements, and an analysis of claims to determine what the claimed invention actually is. The examiner also reads the description of the invention (part of the “specification”) to ensure that disclosure requirements are met. The next step is a search of prior art to determine whether the claimed invention is anticipated by prior patents or non-patent references, and whether the claimed invention is obvious in view of the prior art. There is considerable scope for heterogeneity in this search procedure. The prior art search typically begins with a review of existing US patents in relevant technology classes and subclasses, either through computerized tools, or by hand examination of hard-copy stacks of issued patents, and may then proceed to a

---

<sup>2</sup> In this short discussion, we do not cover the legal requirements for patentability, since these are covered in great detail elsewhere. Indeed, the departure point between our analysis and more of the prior literature in this area is that we are principally concerned with the actual process of examination rather than the standards as defined by the patent law.

word search of foreign patent documents, scientific and technical journals, or other databases and indexes. USPTO's Scientific and Technical Information Center maintains extensive collections of reference materials. Word searches typically require significant skill and time to conduct effectively. The applicant may also include significant amounts of materials documenting prior art with their application. The extent to which examiners review this non-patent material may be a function of the nature of the technology, maturity of the field, and the ease with which it can be searched. For example, in science-intensive fields like biotechnology where much of the relevant prior art is in the form of research articles published in the scientific literature, and indexed by services such as Medline, examiners may rely extensively on non-patent materials. In very young technologies, or in areas where the USPTO has just begun to grant patents, there may be very limited patent prior art. In more mature technologies examiners may have only a moderate interest in non-patent materials, and limited ability to easily or effectively search them.

Once relevant prior art has been identified, the examiner obtains and reads relevant documents. Again, different examiners and different Art Units may use substantially different examination technologies. For example, while many of the mechanical Art Units have historically relied on the "shoes" (the storage bins for hard-copy patent documents), and may search for prior art primarily by viewing drawings, a typical search in the life sciences can involve detailed algorithmic searches by computer to evaluate long genetic sequences, and review of tens or hundreds of research articles and other references. Some examiners may develop and keep close to hand their own specialized collections of prior art to facilitate searching. Indeed, patent examiners identify and frequently refer to "favorite" examples of prior art that usefully describe ("teach") the technology area and the bounds of prior art in a way which facilitates the examination of a wide range of subsequent applications.<sup>4</sup>

After reviewing the specification to ensure that it provides an adequate "enabling disclosure" and an appropriate wording of claims, the initial examination is complete. The examiner then arrives at a determination of whether or not the claimed invention is patentable,

---

<sup>3</sup> This sorting function identifies and appropriately treats applications with national security implications.

<sup>4</sup> Many of these "favorites" are university or public sector patents, which may be written less strategically than those for private firms. In part, this may yield an alternative hypothesis to the finding that university patents are more highly cited than control patents by private firms (Jaffe, Henderson, and Trajtenberg, 1998).

and composes a “first action” letter to the applicant (or, normally, their attorney) that accepts (“allows”), or rejects, the claims. Some applications may be allowed in their entirety upon first examination. More commonly, some or all of the claims are rejected as being anticipated by the prior art, obvious, not adequately enabled, or lacking in utility, and the examiner will write a detailed analysis of the basis for rejection. The applicant then has a fixed length of time to respond by amending the claims and/or supplying additional evidence or argument. After receiving and evaluating this response, the examiner can then “allow” the application if it is satisfactory (the most common stage in the process at which an application proceeds on to final issuance of a patent), negotiate minor changes with the attorney, or write a “second action” letter, which maintains some or all of the initial rejections. In this letter the examiner is encouraged to point out what might be done to overcome these rejections. Though at this stage the applicant’s ability to further amend the application is formally somewhat restricted, in effect additional rounds of negotiation between the examiner and applicant may ensue. The applicant also has the opportunity to appeal decisions for re-examination or evaluation within an internal USPTO administrative proceeding. However, such actions are quite rare; most applications are allowed (or not) on the second or third action letter.

USPTO operates various internal systems to ensure “quality control” through auditing, reviewing, and checking examiner’s work. This includes the collection and analysis of detailed statistics about various measures of examiner work product flow. For example, Supervisory Patent Examiners, as well as their supervisors, routinely evaluate data relating to the distribution of times to action and the number of actions required prior to “disposal” of an application through allowance, abandonment, or appeal. These measurements are one of the many tools that USPTO uses to refine the internal management of the examination process.

It is also useful to note that examiners are allocated fixed amounts of time for completing the initial examination of the application, and for disposal of the application. However, examiners are free to average these time allotments over their caseload. Moreover, there are differences in these time allocations across technology groups, and there also have been changes over time. Though we do not explore this variation in the current study, exploiting these changes in USPTO practice across technology groups and over time could give some leverage for understanding the relationship between time constraints and patent quality.

### *Examiner Training and Specialization*

Variation among examiners in their conduct of the examination process may arise from several sources, but two in particular are worth pointing out. First, at a given point in time, or for a particular patent cohort, examiners necessarily vary substantially in their experience. Experience may affect the quality of patent examination, and this has been a source of concern in recent years as the rate of hiring into the USPTO has increased, particularly into art areas with little in-house expertise. On the other hand, our qualitative research greatly emphasized the role of the systematic apprenticeship process within the USPTO, which is likely to reduce errors made by junior examiners. For the first several years of their career, examiners are denoted as Secondary Examiners, and their work is routinely reviewed by a more senior Primary Examiner. Over time, the Secondary Examiner takes greater control over their case load, and the Primary Examiner focuses on teaching more subtle lessons about the practice of dealing with applicants and their attorneys, and instilling the delicate “not too much, not too little” balance that the USPTO is trying to achieve in the patent examination process.

Second, as alluded to above, Art Units may vary substantially in their organization and functioning. In the most traditional group structure, the allocation of work promotes a maximal amount of specialization by individual examiners. For example, in many of the mechanical Art Units, an individual examiner may be responsible for nearly all of the applications within specific patent classes or subclasses. In other Art Units, however, the approach is more team-oriented. In these groups, there is less technological specialization (multiple subclasses are shared by multiple examiners) and there is likely a higher degree of discussion and knowledge sharing among examiners. In the more specialized organization, there are far fewer checks and balances on the practices of a given examiner. When the examiner has all of the relevant technological information; the cost for an auditor to effectively review their work becomes very high. By contrast, in less specialized environments, there are likely greater opportunities for monitoring, though the loss of specialization may induce a greater degree of arbitrariness into the examination process.

### *Qualitative Findings*

Our qualitative investigation of the patent examination process both generated a number

of insights central to our hypothesis development and raised some flags about potential hazards for empirical research in this area.

The first key finding from the qualitative evaluation of patent examination can be summarized in the phrase of one of our informants: “there may be as many patent offices as there are patent examiners.” In other words, though the examination process is relatively structured, and USPTO devotes considerable resources to quality control, substantial discretion is provided to examiners in how they deal with applications, and the extent to which they exercise this discretion can potentially vary substantially across examiners. Several features contribute to this potential for heterogeneity, including the formal emphasis on specialization, variation among Art Units and individual examiners in their approach to searching prior art, the fact that much learning is through an apprenticeship system with only a small number of mentors, and the existence of differences across groups and examiners in the time allocated to specific tasks and examination procedures.

This heterogeneity might manifest itself in several ways. First, there may be substantial variation across examiners in the breadth of patent grants — some examiners may have a propensity to systematically allow a more restrictive or more expansive set of claims. One potential consequence of this use of discretion may be that patents issued by examiners who tend to allow broader claims will impinge on a greater number of follow-on inventions, and so therefore receive more citations over time. Though prior research has emphasized the degree to which the number of citations received by a patent as an indicator of its underlying inventive significance, it is important to recognize that a given patent’s propensity to receive future citations may also be related to the generosity of the examiner in allowing a broad patent, relative to an “average” examiner’s practice.

Second, examiners may vary substantially in their propensity for self-citation. Self-citation is the practice by which examiners tend to include citations to patents for which they were the examiner. A high degree of self-citation is particularly likely for examiners working in technology areas which are highly specialized, with little communication across examiners, and which are highly reliant on hard-copy technologies for the prior art search process.

Third, examiners may vary substantially in their effective average “approval time,” the length of time between initial application and the date at which the patent issues. Though a large

fraction of the lag between application and approval will, of course, be driven by external forces — the speed at which applicants respond to office actions, for example — differences across Art Units and across examiners in their workload and the type of applications they receive will likely lead to differences in average approval time. It is an interesting question whether this involves a tradeoff with other dimensions of quality, specifically the ability to withstand judicial scrutiny.

At the same time that this qualitative analysis formed the basis for our hypotheses concerning how examiners might influence the allocation of patent rights, it also suggested several limitations to any empirical work, and some challenges that must be overcome before drawing policy conclusions from it. First, and perhaps most importantly, it is important to take account of variation across technologies and patent cohorts in any empirical analysis. Our investigation suggests that there are large differences across Art Units in examination practice, and these technology effects must be controlled for. In addition, examination practice, resources, and management processes have changed over time, so it is also necessary to control in a detailed way for the cohort in which a particular patent was granted.

Second, as we proceed to test hypotheses about observable measures of the examination process, such as citations, we should be careful to recognize how noisy the underlying data generation process is likely to be. Much of the variation in any observable patent characteristic is likely to reflect the nature of the invention, the behavior of the applicant, and other unobserved factors. It is important that we develop a test that will be able to discern rather subtle relationships in the data — the impact of examiner effects is likely to be in light of the overall noisiness of the data generating process. Finally, our hypothesis development must recognize and incorporate the fact that the USPTO has multiple objectives and that there is no single “silver bullet” measure of performance, particularly among easily available statistics. While, all else equal, shorter approval times are socially beneficial (particularly in the era when disclosure did not occur until the patent was issued), speed is not a virtue in and of itself if achieving shorter approval times involves trade-offs with other important objectives. With these caveats in mind, we now turn to a fuller development of key testable hypotheses associated with examiner characteristics.

### **III. Hypothesis Development**

Our empirical analysis is organized around two sets of hypotheses, those reflecting the relationship between patent characteristics and examiner characteristics, and those reflecting the relationship of patent litigation outcomes to patent and/or examiner characteristics.

#### *The Impact of Examiners on Patent Characteristics*

One of the key insights from our qualitative data gathering is the potential for heterogeneity across examiners in their prosecution of patent applications. Perhaps the single most important potential consequence of examiner heterogeneity would be systematic differences across examiners in the average scope of the claims in patents issued under their review. Inventors who receive patent rights with substantial scope will, on average, have been allowed more valuable rights. Identifying the impact of this variation in examiner “generosity” is subtle. The key comes from recognizing that patents with broader claims more likely constrain the claims granted to future inventors. As a result, beyond their innate inventive importance, patents with broader allowed claims will tend to be more highly cited. Conversely, if examiners use discretion similarly, and examiners receive applications with a similar distribution of inventive importance, then the average level of citation should not vary by patent examiner.

While a “generous” grant is a boon to the inventor associated with the application, such treatment may reduce incentives for future inventors, as the hurdle associated with achieving a significant inventive step increases in the breadth and scope offered to inventors from the past. From the perspective of these follow-on inventors, one mechanism to earn a higher return on their own inventions is to seek to invalidate the broad scope associated with a given patent, resulting in specific instances of litigation among the population of an examiner’s patents.

Of course, a number of additional factors determine the number of citations received by a patent or even the average level of citations received by patents associated with a patent examiner, including the particular type of technology and the amount of time that has passed since the application. However, after controlling for technology and cohort effects, variation in the exercise of discretion may still lead examiners to induce different citation levels, yielding our

first hypothesis:<sup>5</sup>

H1: Patent examiners will vary in terms of the average level of citations received by the patents they examine. This variation will exceed what can be attributed to the technological area of the patents they examine.

In addition to this variation among examiners in their level of “generosity,” there is likely variation among examiners in their ability to use search technologies that identify the broadest range of possible prior art. As well, differences in the organization of different Art Units will likely result in different levels of communication and monitoring among examiners, and among examiners and their supervisors. As discussed earlier, one of the consequences of this heterogeneity among examiners is that some examiners may tend towards a more autarkic approach to examination, principally relying on their past experience examining in a particular technological field, while others will draw on a wider range of resources. This discussion motivates our second set of hypotheses:

H2: Examiners will vary in their level of self-citation, and not simply due to the technological area of the patents they examine. Self-citation should be decreasing in the adoption of more advanced prior art search procedures and increasing in the technological specialization of the examiner.

Finally, examiners will vary in the workload they are given, and in the allocations of time for particular tasks associated with the examination process. As several examiners related to us, however, this variation may be in place in order to allow examiners to more effectively achieve other objectives of the examination process, such as precision or effective communication with the patent bar community in their technological specialty. Thus, we offer a third hypothesis about the role of the approval time:

H3: Examiners will vary in their average approval time, above and beyond what can be

---

<sup>5</sup> It is possible that variation in citations received by an examiner reflects selectivity in the assignment of applications to examiners (e.g., SPEs tend to allocate particularly important inventions to particularly able examiners). We discuss this hypothesis further when considering the impact of average citations received on litigation outcomes.

attributed to the technology of the patents examined. This variation will be negatively correlated with other dimensions of performance.

### *The Impact of Examiners on Patent Litigation Outcomes*

Ultimately, we are interested in tying examiner characteristics to more objective measures of the performance of the examination process. We organize this portion of the analysis around patent litigation outcomes. Specifically, we are interested in the possibility that the type of heterogeneity implicit in H1, H2, and H3 (as well as other examiner characteristics) will manifest itself in imperfections in the scope of patent rights that are allowed by examiners. As a preliminary foray into this area, we focus on findings of invalidity by the Court of Appeals for the First Circuit.<sup>6</sup> By focusing on *invalidity*, we develop hypotheses relating to the role that heterogeneity among examiners might play in leading to the excess allocation of patent rights; however, in future work, we hope to explore the converse possibility that this same heterogeneity may also occasionally manifest itself as under-provision.

Perhaps the most obvious potential source of variation among examiners is their overall level of examination experience. In recent years, various commentators have hypothesized that the rapid growth in patent applications and the concomitant rise in the number of examiners has reduced the experience of the average examiner, particularly in technology areas such as business methods, which have only recently begun to receive patent rights. Implicit in this argument is the proposition that less experienced examiners are more likely to inappropriately allow patent rights that should not be granted. While it is likely true that experience is helpful in the examination process, the procedures of the USPTO explicitly recognize the value of experience through practices such as the division of responsibilities between primary and secondary examiners, and the strong culture of internal promotion. There may therefore be competing effects that mitigate the impact of experience on litigation outcomes. However, in order to be precise about the specific theory that has been put forth, we offer a testable hypothesis about the impact of examiner experience.

H4: The probability of a litigated patent being ruled valid will be increasing in the experience

---

<sup>6</sup> We discuss how this particular sampling choice may impact our results in Section IV, where we present the data and our sampling scheme in more detail.

of the examiner.

In addition, H1, H2, and H3 offer at least three potential sources of heterogeneity which may be associated with excess allocation of patent rights, and therefore with invalidity findings. First, H1 states that some examiners are more “generous” than others, and that this should be associated with a higher level of citations received by their patents. To the extent that the claims allowed by over-generous examiners will tend to be more likely to be found invalid by the CAFC, the probability of validity should be declining in examiners’ average level of citations received. Similarly, to the extent that it may be easier to overturn the validity of patents based on less thorough searchers of the prior art, the probability of a ruling of validity may be declining in the self-citation of the examiner. Finally, if there is a tradeoff between the speed of approval and the quality of the examination, then the probability of validity will likely be declining in the approval time of the examiner. This discussion motivates the following hypotheses:

H5: The probability of a litigated patent being ruled valid will be declining in the examiner’s average citations received per patent, and in the self-citation rate of the examiner, but should be increasing in the examiner’s average approval time.

Of course, as suggested by our earlier discussion, a test of H5 requires the ability to control for the technological specialization of the examiner, as well as controls for the cohort of the litigated patent. Our empirical analysis includes detailed controls for each of these effects. As well, to the extent that the assignment of patent applications to examiners is subject to *selectivity* (i.e., particularly important technologies, associated with higher citation rates, are assigned to more able examiners), our ability to find evidence for the impact of generosity become more difficult. Since selectivity holds the opposite prediction about the impact of examiner average citations on validity findings, the simple test suggested by H5 results in a *lower bound* on the role of generosity in judicial review.

Finally, the richness of patent data allows us to explore the impact of examiner heterogeneity more precisely. When determining validity, the CAFC will, of course, only consider the merits of the patent under review rather than the historical record of a particular examiner. To determine the impact of examiner characteristics on the probability of validity, only that portion of the citations received by the patent that are due to the examiner’s generosity

is relevant. If H1 is true, i.e. the number of citations received by the litigated patent is a function of the examiner's average number of citations per patent, then this relationship allows us to estimate this portion econometrically.<sup>7</sup> This motivates the following hypothesis:

H6: The probability of validity should be declining in the predicted number of citations received by a patent, where the prediction is based on the examiner's generosity.

Together, these hypotheses provide several potential observable consequences of examiner heterogeneity, with clear implications for patent policy. Consider, for example, the perennial policy issue of patent "disposal" times. By linking approval to other outcomes (such as validity rulings) these hypotheses offer potential insight into the potential for tradeoffs associated with speeding of the examination process. To empirically test these propositions, we must tie these hypotheses to a specific set of data, to the description of which we now turn.

#### **IV. The Data**

##### *The Sample*

Data for this study were derived from the USPTO's public access patent databases, and from the Lexis-Nexis database of decisions of the Court of Appeals for the Federal Circuit.

We began by searching for CAFC decisions in cases where the validity of a patent was contested. In the years 1997-2000, there were 216 such cases, of which 34 were excluded from further consideration because they involved plant patents, re-examined patents, or other complicating factors were present. For each of the remaining 182 "CAFC-tested" patents, we determined whether or not the CAFC found the patent to be valid or invalid, and on what grounds: novelty, subject matter, obviousness, procedural errors etc. Note that in many instances the CAFC found the patent invalid for more than one reason. In just over 50% of these 182 cases, the patent was found to be invalid. Of these, the CAFC found problems with novelty

---

<sup>7</sup> Specifically, H6 can be tested using an "instrumental variables" estimator where the validity ruling is regressed on the predicted level of citations associated with the litigated patent, with the excluded exogenous variable in the validity equation is the examiner's generosity, as measured by the average level of citations received. Intuitively, this procedure is equivalent to a two stage estimation procedure. In the first stage, total citations received by each patent is regressed on the examiner's average citations per patent, and other controls. Predicted values of the total citations variable, i.e. the portion of citations attributable to the examiner's generosity, are then used in the second stage validity regression.

(Section 101) in 37% of cases, with obviousness (Section 102) in 47% of cases, and with the specification of the patent (Section 112) in 15% of cases.

Having obtained this list of CAFC-tested patents, we then used it to construct a sample of “CAFC-tested” examiners.<sup>8</sup> To do so we identify the 196 individuals listed as either the primary or secondary examiner for each of the 182 CAFC-tested patents.<sup>9</sup>

For each CAFC-tested examiner, we search for all patents granted in the period 1976-2000 on which the individual was listed as a primary or secondary examiner. This search was conducted using a fairly generous “wild card” procedure to allow for typographical errors in the source data and variations in the spelling or formatting of names. Results were then carefully screened by hand to ensure that individuals were correctly identified. For example, our procedure would recognize “Merrill, Stephen A.”, “Merril, Stephen”, “Merrill, S.A.” and “Merrill, Steve” as being the same person, but would exclude “Meril, S.” or “Merrill, Stavros A.” If anything, this process erred on the side of caution, so that we may be slightly under-counting examiners’ output. The initial search returned just over 316,000 candidate patents, from which we excluded about 6% miss-identified patents to arrive a base dataset of 298,441 patents attributable to the 196 CAFC-tested examiners.<sup>10</sup>

Using the dataset of 298,441 patents we constructed complete histories of each CAFC-tested examiner’s patent output during the sample period, as well as various measures of their productivity, experience with examination, workload, and examining practice. Each of these patents was matched to the NBER Patent Citation Data File (Hall, Jaffe, and Trajtenberg, 2001) to obtain data on each patent's technology classes, citations made, and citations received, as well as variables computed from these data which measure the breadth of citations.<sup>11</sup>

In our empirical analysis that follows, we focus on the primary examiner for each of the 182

---

<sup>8</sup> In using the phrase “CAFC-tested” we do not mean to imply that a ruling of invalidity by the CAFC necessarily implies any shortcoming on the part of the examiner.

<sup>9</sup> In future work it would be possible to conduct parts of our analysis of patent examiners on a much wider sample of individuals who performed this function at the USPTO. A useful feature of our small sample, however, is that each examiner in the sample has examined at least one patent that was “tested” for validity by the CAFC.

<sup>10</sup> Since we have not been able to obtain a definitive matching of examiner ID numbers with issued patents, and have had to work from published data sources this search misses a small number of patents. We are confident, however, that missing observations are missing at random, and therefore do not bias our results.

<sup>11</sup> Jaffe, Henderson and Trajtenberg (1998) computed two measures of the breath of citations across technology classes: “Originality” which captures the extent to which citations made by a patent are spread across technology classes, and “Generality” which captures the extent to which citations received by a patent are spread across

CAFC-tested patents. Since the same primary examiner may show up several times in the sample, we actually have 136 CAFC-tested primary examiners. In computing statistics across examiners, we weight examiner characteristics of by the number of times that each examiner shows up in our data. Table 1 gives variable definitions, and Table 2A presents descriptive statistics for our linked datasets on CAFC-tested patents, CAFC-tested primary examiners, and patent histories of these examiners.

The set of 182 CAFC-tested patents is a highly selective sample; these patents are not at all representative of the population of all granted patents. Table 2B compares mean values of some key variables for the 182 CAFC-tested patents with those typical of a utility patent applied for in the decade of the 1980s.<sup>12</sup> On average, the CAFC-tested patents contain more claims, make more citations, receive more citations, and take longer to issue. This is not surprising, since litigants who pursue CAFC review likely perceive a high value for intellectual property over a given technology. As well, given that the litigants have not settled, these patents are likely associated with a higher level of ambiguity than an average patent (perhaps an additional reason for the longer time to approval).

Though the CAFC-tested patents are quite selective, there is little reason to believe that CAFC-tested primary examiners are very different from the population of all examiners in a way which undermines our empirical strategy.<sup>13</sup> On the one hand, due to the way we have constructed our sample, the probability of an examiner being in our dataset is likely proportional to the examiner's experience (measured in terms of total patents examined) at the USPTO. Thus relative to the set of examiners working at the USPTO on any given day, we are under-sampling inexperienced examiners. As well, our sample may under-represent the degree of variation in the degree of generosity; patents associated with the least generous examiners are less likely to be subject to an appellate validity claim. As such, our empirical design is providing a lower bound of the impact of examiner experience or generosity on patent litigation outcomes. Because examiner patent histories begin in our dataset in 1976, the measures of experience are slightly downward biased. About 30% of the examiners in our sample first appear in the dataset

---

technology classes. See Table 1 for definitions.

<sup>12</sup> The statistics for a typical patent are based on the tables and figures in Hall, Jaffe, and Trajtenberg (2001).

<sup>13</sup> We intend to test this proposition more carefully by randomly sampling examiners. Initial comparisons with the small set of examiners caught in the wildcard search, but rejected as poor matches, find no substantive differences

in 1976, some fraction of which must be assumed to have begun their careers somewhat earlier. Similarly, citations to patents granted before 1976 cannot be evaluated as self-citations (or not) since we do not have information on who the examiner was, and information based on citations received by patents granted in recent years is limited by the truncation of the dataset in 2001.

## V. Results

We present our results in several steps. First, we review evidence of the existence of heterogeneity among examiners, and show that an important component of the overall variation in commonly used patent statistics can be explained by examiner “fixed effects.” Having established the existence of observable examiner heterogeneity, we then examine the sensitivity of various characteristics of CAFC-tested patents to observable examiner characteristics. We then turn to a discussion of the determinants of patent validity. Consistent with our discussion in Section III, we evaluate a reduced-form model of the sensitivity of validity to examiner characteristics as well as a more nuanced instrumental variables estimation which only allows examiner characteristics to impact validity through their impact on characteristics unique to the CAFC-tested patent.

### *The Nature of Examiner Heterogeneity*

Our analysis begins with a set of figures that display the heterogeneity among examiners along four distinct dimensions: experience, the level of citations received per patent, the degree of self-citation, and the degree of technological specialization in the patents examined. Figure 1 plots EXPERIENCE (# of patents) across examiners. We see that while the average examiner in our sample has a lifetime experience of over 2000 patents, a large number are associated with over 4000 patents, with a few outliers over 7000. This distribution is consistent with the substantial variation we see in the examiners’ length of tenure at the USPTO. For example, nearly a third of the CAFC-tested examiners have over 24 years experience at the USPTO.<sup>14</sup>

We next turn to an evaluation of the extent to which examiners specialize in particular

---

between them and the sample of CAFC-tested primary examiners in terms of experience and other characteristics.

<sup>14</sup> This may be somewhat biased upward since patent examiners may not “exit” in the current way we have computed this particular statistic

technology classes over the course of their career. One simple way to measure this specialization is to compute the number of distinct technology classes appearing among the patents examined by a particular examiner. Using six broad technology classes, this measure (EXAMINER TECH. EXPERIENCE) is displayed in Figure 2. We see that it is most common to have examined patents in nearly all of the six classes. Yet even if an examiner has dealt with all types of patents, he may still be highly specialized within a single technology category with only an occasional patent elsewhere. A more sophisticated approach to deal with this issue is to compute a Herfindahl type index of the dispersion of an examiner's patents over technology classes.<sup>15</sup> This measure (EXAMINER SPECIALIZATION) is plotted in Figure 3. While some noise is inherent in this measure due to the nature of the technology classification system, its mean level across examiners (0.75) indicates a high average degree of specialization. As Figure 3 indicates, however, there is also considerable variation: while the modal examiner is highly specialized, with a specialization index near 1, there are still a significant number with a much greater degree of dispersion of patents across technology classes.

Perhaps more interestingly, there is also substantial variation among examiners in the characteristics of "their" patents. Figure 4 shows the distribution of the average number of citations received overall all patents issued by each examiner (EXAMINER CITES PER PATENT). The distribution is highly skewed. The coefficient of variation associated with examiner cites received per patent issued is over 0.5; and over 10% of examiners have citation rates more than double the average citation rate. Similarly, as shown in Figure 5, while the average self-citation rate (SELF-CITE) is relatively low, particularly given the technological specialization of examiners, some examiners have self-citation rates more than three times the sample mean. Another method for understanding the importance of heterogeneity across examiners is to use ANOVA analysis to formally test for the presence of examiner effects in several key statistics associated with the examination process. An advantage of this statistical approach is that we can condition on other variables that might explain the observed differences across examiners, such as the technological areas of the patents they examine. Recall that in H1, H2, and H3 we hypothesized that the differences across examiners were not simply a reflection

---

<sup>15</sup> A Herfindahl index is a commonly used measure of concentration, based the sum of the squares of the share of a variable across categories.

of the technological area of the patents they examined.

In Table 3A, we present a simple ANOVA analysis based on our complete sample of 298,441 patents attributed to the 196 CAFC-tested examiners. The results indicate that examiners matter: a significant share of the variance in this sample in the four variables capturing the volume and pattern of citations by and to a particular patent (CITATIONS MADE, CITATIONS RECEIVED, ORIGINALITY, and GENERALITY) is accounted for by fixed examiner effects, with a particularly strong effect in the ANOVA of CITATIONS RECEIVED. A similar result is obtained for the length of time between application and grant: about 8% of the variance in this measure can be attributed to differences among examiners. A much smaller share of variance is explained for the number of claims on each patent. These results are robust to controlling for differences across technology classes. As Table 3B shows, there are visible differences across technology classes in the fraction of variance explained by examiner effects. There appears to be much more homogeneity across examiners in examination of Mechanical patents, with significantly less homogeneity in CITATIONS MADE for Chemical patents, and in the approval time for Electrical/Electronic patents. Overall, these results confirm the intuition we developed in our qualitative investigation: there is substantial heterogeneity across examiners, even after controlling for the important technology and cohort effects.

The analysis above suggests that examiners vary, particularly in terms of the rate at which their patents tend to receive citations. But how does this variation, which we have suggested as a proxy for examiner generosity, affect our set of CAFC-tested patents? Table 4 presents regressions relating the CITATIONS RECEIVED by the CAFC-tested patents to a set of examiner characteristics, and in particular EXAMINER CITES PER PATENT. One result is particularly striking: there is a very strong relationship between EXAMINER CITES PER PATENT and CITATIONS RECEIVED by CAFC-tested patents. The effect is slightly reduced, but still quite significant, after conditioning on the patent's detailed technology sub-class, cohort, and assignee type. In each of the specifications in Table 4, increasing EXAMINER CITES PER PATENT by one patent (less than one-third of a standard deviation) increases the predicted number of citations of the CAFC-tested patent by more than one (recall that CAFC-tested patents have much higher overall citation rates). Other observable examiner characteristics have a less clear relationship with CITATIONS RECEIVED. The overall level of self-citation, experience

(both in terms of years as well as the total level of issued patents) and a measure of near-term workflow (3-MONTH VOLUME) are all insignificant in their impact on CITATIONS RECEIVED.

Many factors may affect how many citations a patent receives. Citations received are frequently thought to reflect the technological significance of the claimed invention. Pioneering inventions with broad claims and no closely related prior art will tend to be frequently cited as follow-on inventors improve on the original invention. Citations may also reflect the quality or scope of the disclosure accompanying the claims. We cannot directly measure either of these factors here, and our analysis assumes that they are randomly distributed across the patents in our sample. Nonetheless these results do indicate that a significant fraction of the variation in citations received by any particular patent is driven by a single aspect of examiner heterogeneity, the average propensity of “their” patents to attract citations. This is true even after controlling for other important attributes of the patent such as the technology class, the year when it was approved, and the type of assignee.

#### *The Impact of Examiner and Patent Characteristics on Litigation Outcomes*

We now turn to the final part of our analysis — linking examiner characteristics to litigation outcomes. While the overall probability of validity being upheld is approximately 50%, there is substantial variation in this across technological areas, year of patent approval, and even the type of assignee (see Figures 6-8). For example, while pharmaceutical and medical patents are more likely than not to be upheld, a substantial majority of computers & communications equipment patents are overturned. As well, the age of a patent seems to be an important predictor of validity — pre-1990 approvals are much more likely to be upheld by the CAFC than post-1990 approvals. As we emphasized in Section II, these findings suggest the importance of controlling for detailed technology classes and cohorts in our analysis as we seek to evaluate the sensitivity of validity findings to examiner characteristics.

We begin our analysis in Tables 5 and 6, which compare the means of examiner characteristics and patent characteristics, conditional on whether the CAFC ruled the patent valid. Several issues stand out. First, the conditional means associated with most of the patent

characteristics are roughly the same. It is useful to note that there is less than a 10% difference in the level of CITATIONS RECEIVED between the two groups. The only striking difference is in APPROVAL TIME, where the time taken to approve *invalid* patents is significantly higher than the time taken to approve those that were found to be valid. This is interesting, though certainly not dispositive, evidence against a simple tradeoff between approval times and patent quality as measured by validity rulings. Turning to the mean examiner characteristics by validity (Table 6), the striking differences are in terms of EXAMINER CITES PER PATENT and 3-MONTH VOLUME. There is no significant difference in the means according to experience level; if anything, invalid patents are associated with examiners with higher mean levels of experience, both in terms of volume and tenure. This stands in useful contrast to the most naïve interpretation of H4, which predicts that EXPERIENCE should be positively correlated with VALIDITY. In contrast, consistent with the suggestion in H5, invalid patents do seem to be associated with examiners with a higher average citation rate.

Of course, these conditional means ignore the important differences across technologies (Figure 6), cohorts (Figure 7), or assignees (Figure 8) and the potential for correlation among the examiner characteristics themselves. We therefore turn to more systematic set of regression analyses in Table 7. The dependent variable in the regressions takes the value 1 if the CAFC-tested patent is ruled valid, 0 otherwise.<sup>16</sup> The first two columns of Table 7 provide a test for H4, the sensitivity of the probability of a validity finding to the experience of the examiner. Whether detailed controls are included or not, there is no significant relationship between any measure of experience and the probability of a ruling of validity. Indeed, we have experimented with a wide variety of specifications relating to these experience measures and there is no systematic relationship with validity and these measures in these data. In the last two columns of Table 7, we turn to H5, the sensitivity of a validity ruling to other examiner characteristics. The only significant relationship is with EXAMINER CITES PER PATENT, which has a significant and large negative coefficient. Moreover, this coefficient increases in absolute value when detailed technology and cohort controls are included. According to (7-4), increasing the EXAMINER

---

<sup>16</sup> Both Tables 7 and 8 employ a linear probability model, either OLS or IV. The coefficients are therefore easily interpretable, comparable with each other, and we avoid the technical subtleties associated with an implementing instrumental variables probit in the context of a small sample. We experimented with a probit model for the reduced-form OLS results, and the results remain quantitatively and statistically significant.

CITES PER PATENT by one standard deviation (3.49), the probability of validity is predicted to decline by over 14 percentage points, from a mean of 48%. In other words, the probability of validity is strongly associated with the average rate at which that examiner's patents have received citations. One way to interpret this result is that the rulings of the CAFC serve to mitigate the value of patents allowed by particularly "generous" examiners.

This finding motivates our final set of regressions using the more nuanced instrumental variables procedure in Table 8. As discussed in Section III, we investigate the mechanism by which EXAMINER CITES PER PATENT might affect patent validity rulings by restricting its impact to the citation rate of the litigated patent. In other words, we impose the exclusion restriction that, but for its impact on CITATIONS RECEIVED, EXAMINER CITES PER PATENT is exogenous to the validity decision. The results of this IV analysis are striking. On the one hand, the OLS relationship between validity and CITATIONS RECEIVED is insignificant (8-1). However, the coefficient on CITATIONS RECEIVED in the instrumental variables equations is significant, large and negative. Though validity is unrelated to the total number of citations received by a patent, validity is strongly related to the portion of the citation rate explained by the examiner's average propensity to grant patents that attract citations. Moreover, the size of this coefficient increases substantially after the inclusion of technology, cohort, and assignee effects, as well as with the inclusion of other characteristics of CAFC-tested patents. If our results were being driven by unobserved variation across examiners in the types of technologies examined, these controls would likely condition out some of this heterogeneity; the fact that our results become stronger after the inclusion of controls makes our findings even more suggestive. In other words, even relying on a test that only allows examiner effects to matter only through their impact on the citation rate of the litigated patent, and controlling for differences in the timing and type of litigated technology, find that courts invalidate patent rights associated with "generous" examiners.

## **VI. Discussion & Conclusions**

We have conducted an empirical investigation, both qualitative and quantitative, of the role that patent examiners play in the allocation of patent rights. In addition to interviewing administrators and patent examiners at the USPTO, we have constructed and analyzed a novel

dataset on patent examiners and patent litigation outcomes. Starting with a sample of patents for which the CAFC decided on validity between 1997-2000, we collect historical data on who examined these patents at the USPTO. For each of these examiners, we also collected data on all of the other patents that they examined during their career, allowing us to compute a number of interesting examiner characteristics. The dataset obtained by matching these two sources is, of course, based on a highly selected sample, since very few patents make it to the CAFC. Nonetheless, we view it as a very useful laboratory for us to explore a number of hypotheses about the connection between the patent examination and process and the issuance of patent rights. Our results are preliminary, but they suggest a number of interesting findings.

First, patent examiners and the patent examination process are not homogeneous. There is substantial variation in observable characteristics of patent examiners, such as their tenure at the USPTO, the number of patents they have examined, the average approval time per issued patent, and the degree of specialization in technology areas. There is also systematic variation in outcomes of the examination process — such as the volume and pattern of citations made and received by patents — that can be attributed to idiosyncratic differences among examiners. Most interestingly, examiners differ in the in the number of citations made to “their” issued patents, even after controlling for technology class, issuing cohort, and other factors.

Second, we find no evidence to date favoring “naïve” hypotheses about examiner characteristics and patent quality. In particular, we find no strong statistical association between examiner experience or workload at the time a patent is issued and the probability of the CAFC finding it to be invalid if it is subsequently litigated. Thus, our work does not lead us to a policy prescription related to, for example, reducing the turnover of patent examiners.

Third, “examiners matter:” though highly structured, and carefully monitored by USPTO, patent examination is not a mechanical process. Examiners necessarily exercise discretion, and occasionally this discretion results in patent claims being allowed which are overturned by subsequent judicial review. Our core finding is that the examiners whose patents are cited most are also more likely to have their patents ruled invalid by the CAFC. Our econometric procedure distinguishes between citations received by a particular patent due to the scope of its claims, or the significance of its particular technology, and citations received due to the examiner’s “generosity” as captured by their propensity to allow patents that attract citations.

It is only the second of these that has a statistical relationship with CAFC validity rulings.

The fact that patent examination cannot be mechanistic, and that idiosyncratic aspects of examiner behavior appear to have a significant impact on the nature of the patent rights that they grant, suggests a significant role for the organization, leadership, and management of USPTO. The management literature recognizes the value of corporate culture in the form of informal rules, common values, exemplars of behavior, etc. in providing guidance on how to exercise discretion. While idiosyncratic behavior of examiners can be controlled to some extent by formal processes such as supervision, selection of examiners, training, incentives etc., the institution's cultural norms necessarily play an important role in their exercise of discretion in awarding patent rights. Policy changes which impact the organizational structure and internal culture of the USPTO should be careful to take this into account. Efforts to improve visible aspects of the examination process, such as approval times, have the potential to create long lasting and quite subtle changes on less easily measured aspects of the examination process.

## REFERENCES

- Cohen, W., Nelson, R.R. and J.P. Walsh (2000). "Protecting Their Intellectual Assets: Appropriability Conditions and why U.S. Manufacturing Firms Patent (or Not)." NBER Working Paper 7552.
- Griliches, Z., ed. (1984). *R&D, Patents and Productivity*. Chicago (IL): Chicago University Press.
- Griliches, Z. (1990). "Patent Statistics as Economic Indicators: A Survey," *Journal of Economic Literature*, 92, 630-653.
- Hall, B. and Ziedonis, R. (2001). "The Patent Paradox Revisited: An Empirical Study of Patenting in the US Semiconductor Industry 1979-1995," *RAND Journal of Economics*.
- Hall, B., Jaffe, A. and M. Trajtenberg, (2001). "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools." NBER Working Paper 8498.
- Jaffe, A., Henderson, R. and M. Trajtenberg (1998). "Universities as a Source of Commercial Technology: A Detailed Analysis of University Patenting, 1965-1988" *Review of Economics and Statistics*
- Jaffe, A. Trajtenberg, M. and R. Henderson (1993). "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations" *Quarterly Journal of Economics*.
- King, J. (2002). "Patent Examination Procedures as Inputs to Patent Quality," this volume.
- Lanjouw, J. and Lerner, J. (1997). "The Enforcement of Intellectual Property Rights: A Survey of the Empirical Literature," NBER WP 6296.
- Levin, R., Klevorick, A., Nelson, R. and S. Winter (1987), "Appropriating the Returns from Industrial R&D," *Brookings Papers on Economic Activity*.
- Lerner, J. (2000). "150 Years of Patent Office Practice," NBER WP 7477.

**TABLE 1**  
**VARIABLES & DEFINITIONS**

VARIABLE	DEFINITION
<b>VALIDITY</b>	
VALID	Valid = 1 if patent validity upheld by CAFC; 0 else
<b>CAFC PATENT CHARACTERISTICS</b>	
CITATIONS RECIEVED	# of Citations to CAFC Patent from grant date through 6 / 2001
CLAIMS	# of Distinct Claims for CAFC Patent
APPROVAL TIME	Patent Issue Date – Patent Application Date (Days)
GENERALITY	Jaffe-Henderson-Trajtenberg “Generality” index = $1 - \sum_j \left( \frac{Cites\ Received_j}{Total\ Cites\ Received} \right)^2$ j = technology classes
ORIGINALITY	Jaffe-Henderson-Trajtenberg “Originality” index: $1 - \sum_j \left( \frac{Cites\ Made_j}{Total\ Cites\ Made} \right)^2$ j = technology classes
<b>PRIMARY EXAMINER CHARACTERISTICS</b>	
EXPERIENCE (NO. OF PATENTS)	Cumulative Patent Production by Examiner, both primary & secondary (see Figure 1A)
EXAMINER CITATIONS	Cumulative Citations to Examiner Patents (through July, 2001)
EXAMINER CITES PER PATENT	EXAM CITATIONS divided by EXPERIENCE (NO. OF PATENTS), (see Figure 1B)
SECONDARY EXPERIENCE	Cumulative Patent Production as Secondary Examiner
SELF-CITE	Share of All Citations to Own Prior Patents (see Figure 1C)
EXAMINER TECH. EXPERIENCE	Number of broad technology classes of patents on which the examiner has experience
EXAMINER SPECIALIZATION	Herfindahl-type measure of distribution of examiner’s patents across broad technology classes
EXPERIENCE (YEARS)	Cumulative Years Observed as Issuing Examiner (both primary & secondary)
3-MONTH VOLUME	Count of Issued Patents in Three Months Immediately Prior to Issue Date
<b>CONTROL VARIABLES</b>	
TECH CLASS FIXED EFFECTS	6 Distinct Technology Categories Based on Patent Classes (see Figure 5A)
TECHNOLOGY SUB-CLASS FIXED EFFECTS	35 Distinct Technology Sub-Classes Based on Patent Sub-Classes (see Hall-Jaffe-Trajtenberg)
COHORT FIXED EFFECTS	20 Individual Year Dummies Based on CAFC Patent Issue Date
ASSIGNEE FIXED EFFECTS	4 Dummies for Type of Assignee (see, Figure 5C)

**TABLE 2A**  
*MEANS & STANDARD DEVIATIONS*

	MEAN	STANDARD DEVIATION
<b>VALIDITY</b>		
VALID	0.48	0.50
<b>CAFC PATENT CHARACTERISTICS</b>		
CITATIONS RECEIVED	16.74	21.47
CLAIMS	20.52	26.05
APPROVAL TIME	804.60	799.80
GENERALITY	0.41	0.27
ORIGINALITY	0.39	0.28
<b>EXAMINER CHARACTERISTICS</b>		
EXPERIENCE (NO. PATENTS)	2180.38	1395.65
EXAMINER CITATIONS	14201.68	12673.34
EXAMINER CITES PER PATENT	6.32	3.49
SECONDARY EXPERIENCE	207.05	137.67
SELF-CITE	0.10	0.06
EXAMINER SPECIALIZATION	0.75	0.20
EXPERIENCE (YEARS)	18.67	5.67
3-MONTH VOLUME	41.52	35.66

**TABLE 2B**  
**PATENT CHARACTERISTICS**  
**CAFC SAMPLE COMPARED TO “UNIVERSE”**

	CAFC SAMPLE (182)	TYPICAL PATENT (1980s application yr)
CLAIMS	20.5	9-14
CITATIONS RECEIVED	14.0	6-8
CITATIONS MADE	16.7	6-8
ORIGINALITY	0.36	0.3-0.4
GENERALITY	0.41	0.3-0.4
APPROVAL TIME (YEARS)	2.21	1.76-2.05

**TABLE 3A**  
*ANALYSIS OF VARIANCE OF  
 PATENT CHARACTERISTICS*

N=289,441  
 196 Examiner Effects  
 36 Technology Sub-Class Effects  
 24 Cohort Effects

Variable	Fraction of variance explained by examiner effects	F-statistic for no examiner effect	F-statistic for no examiner effect, controlling for detailed technology class and cohort
CITATIONS MADE	0.077	121.71	52.64
CITATIONS RECEIVED	0.117	193.40	51.07
APPROVAL TIME	0.083	131.77	78.92
CLAIMS	0.030	44.83	16.06
GENERALITY	0.079	105.56	38.97
ORIGINALITY	0.069	104.23	61.30

**TABLE 3B**  
*ANALYSIS OF VARIANCE OF  
 PATENT CHARACTERISTICS  
 BY TECHNOLOGY CLASS*

Variable	Fraction of Variance Explained by Examiner Effects					
	Chemical	ICT	Drug/Med	Electronic	Mechanical	Other
CITATIONS MADE	0.123	0.054	0.104	0.078	0.054	0.059
CITATIONS RECEIVED	0.058	0.099	0.110	0.066	0.076	0.072
APPROVAL TIME	0.098	0.083	0.074	0.116	0.053	0.053
CLAIMS	0.033	0.027	0.028	0.022	0.031	0.037
GENERALITY	0.084	0.112	0.078	0.086	0.055	0.081
ORIGINALITY	0.087	0.964	0.044	0.063	0.069	0.082

**TABLE 4**  
**CITATIONS-RECEIVED EQUATION**

	(4A-1)	(4A-2)	(4A-3)	(4A-4)
DEPENDENT VARIABLE	CITATIONS RECEIVED			
<b>EXAMINER CHARACTERISTICS</b>				
EXAMINER CITES PER PATENT	<b>2.68</b> <b>(0.41)</b>	<b>1.83</b> <b>(0.53)</b>	<b>1.82</b> <b>(0.54)</b>	<b>1.69</b> <b>(0.58)</b>
SELF-CITE			25.80 (26.18)	40.56 (28.03)
EXPERIENCE (YEARS)			0.13 (0.27)	0.32 (0.30)
3-MONTH VOLUME			0.01 (0.04)	0.01 (0.04)
<b>PATENT CHARACTERISTICS</b>				
GENERALITY				<b>11.28</b> <b>(6.70)</b>
ORIGINALITY				1.90 (6.33)
<b>CONTROL VARIABLES</b>				
COHORT FIXED EFFECTS		<b>SIG.</b>	<b>SIG.</b>	<b>SIG.</b>
TECHNOLOGY SUB-CLASS FIXED EFFECTS		<b>SIG.</b>	<b>SIG.</b>	<b>SIG.</b>
ASSIGNEE FIXED EFFECTS		INSIG.	INSIG.	INSIG.
<b>Regression Statistics</b>				
Adj. R-squared	0.19	0.44	0.45	0.45
# of Observations	182.00	182.00	170.00	170.00

**TABLE 5**  
**PATENT CHARACTERISTICS**  
**MEANS CONDITIONAL ON CAFC VALIDITY RULING**

	INVALID	VALID
CLAIMS	20.73	20.28
CITATIONS RECEIVED	17.38	16.04
ORIGINALITY	0.36	0.41
GENERALITY	0.41	0.41
APPROVAL TIME	845.51	760.90

**TABLE 6**  
**EXAMINER CHARACTERISTICS**  
**MEANS CONDITIONAL ON CAFC VALIDITY RULING**

	INVALID	VALID
EXPERIENCE (NO. OF PATENTS)	2276.40	2077.81
EXPERIENCE (YEARS)	18.82	18.51
EXAMINER CITES PER PATENT	6.89	5.72
SELF-CITE	0.10	0.10
3 MONTH VOLUME	45.56	37.19

**TABLE 7**  
**REDUCED-FORM OLS VALIDITY EQUATION**

	(7-1)	(7-2)	(7-3)	(7-4)
DEPENDENT VARIABLE	VALID			
<b>EXAMINER CHARACTERISTICS</b>				
EXPERIENCE (NO. OF PATENTS)	-2.97 E-05 (3.26 E-05)	-4.97 E-05 (3.52 E-05)	2.43 E-05 (4.57 E-05)	-7.60 E-05 (5.53 E-05)
EXPERIENCE (YEARS)	-0.002 (0.008)	-0.005 (0.008)	-0.0003 (0.008)	-0.002 (0.009)
SELF-CITE			-0.41 (0.67)	-0.190 (0.797)
3-MONTH VOLUME			-0.002 (0.0014)	0.001 (0.001)
EXAMINER CITES PER PATENT			<b>-0.024</b> <b>(0.011)</b>	<b>-0.041</b> <b>(0.015)</b>
<b>CONTROL VARIABLES</b>				
COHORT FIXED EFFECTS		Insig.		Insig.
TECHNOLOGY SUB-CLASS FIXED EFFECTS		<b>SIG.</b>		<b>SIG.</b>
ASSIGNEE FIXED EFFECTS		Insig.		Insig.
<b>Regression Statistics</b>				
Adj. R-squared	0.000	0.113	0.017	0.143
# of Observations	182.00	182.00	182.00	182.00

**TABLE 8  
VALIDITY EQUATION**

	(8-1) OLS	(8-2) IV(*)	(8-3) IV(*)	(8-4) IV(*)
DEPENDENT VARIABLE	VALID			
<b>PATENT CHARACTERISTICS</b>				
CITATIONS RECEIVED	-0.0007 (0.0017)	<b>-0.0090</b> <b>(0.0042)</b>	<b>-0.0228</b> <b>(0.0106)</b>	<b>-0.0242</b> <b>(0.0111)</b>
CLAIMS				0.003 (0.002)
ORIGINALITY				0.238 (0.227)
GENERALITY				0.188 (0.268)
APPROVAL TIME				-0.000 (0.000)
<b>CONTROL VARIABLES</b>				
COHORT FIXED EFFECTS			<b>SIG.</b>	<b>SIG.</b>
TECHNOLOGY SUB-CLASS FIXED EFFECTS			<b>SIG.</b>	<b>SIG.</b>
ASSIGNEE FIXED EFFECTS			Insig.	Insig.
<b>Regression Statistics</b>				
Adj. R-squared	0.000	NA	NA	NA
# of Observations	182.00	182.00	170.00	170.00

\* IV: ENDOGENOUS = CITATIONS RECEIVED  
INST VARS = EXAMINER CITES PER PATENT

FIGURE 1: Experience of Examiners

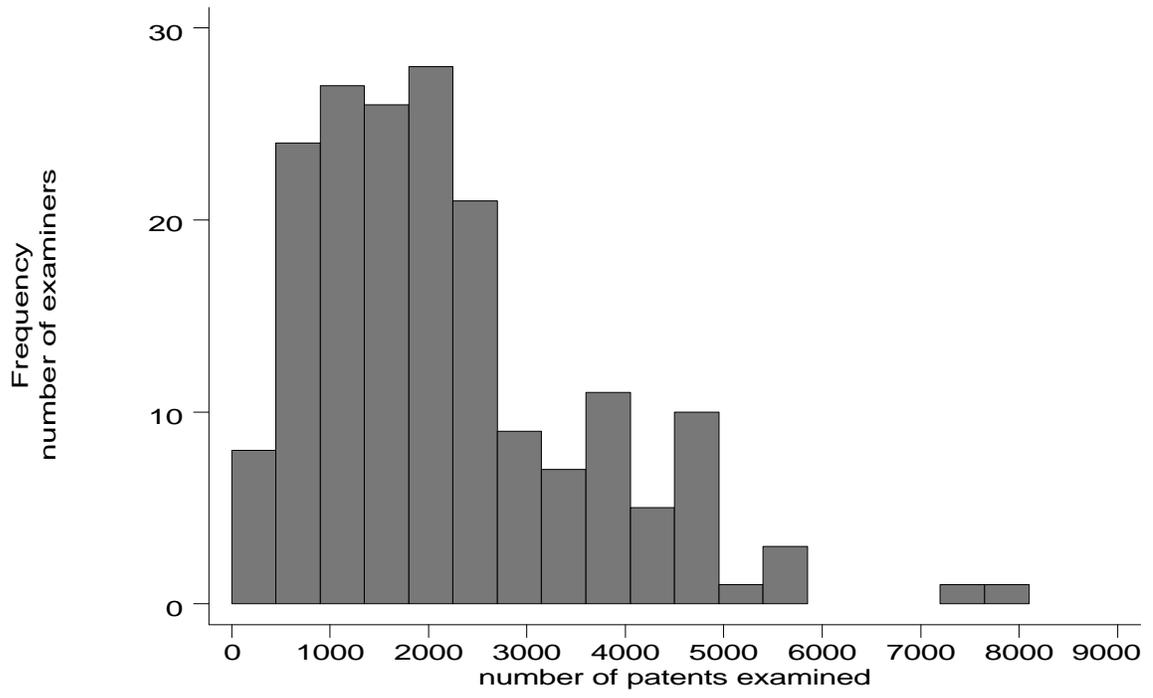


FIGURE 2: Technological Experience of Examiners

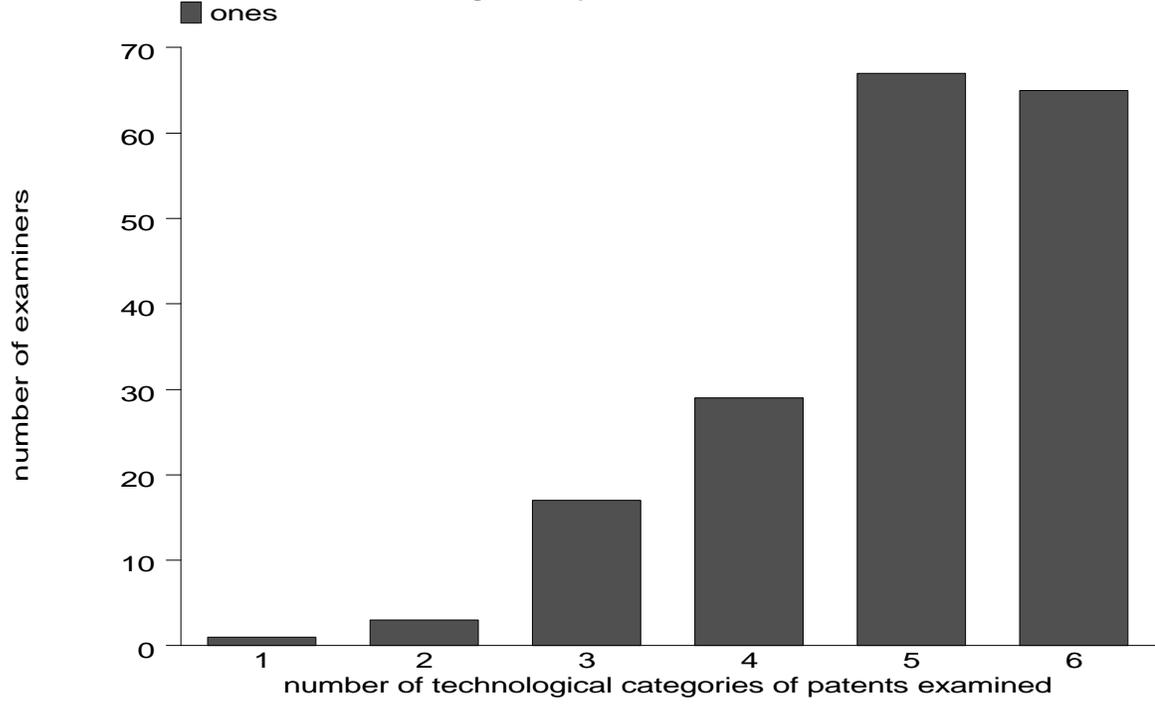


FIGURE 3: Technological Specialization of Examiners

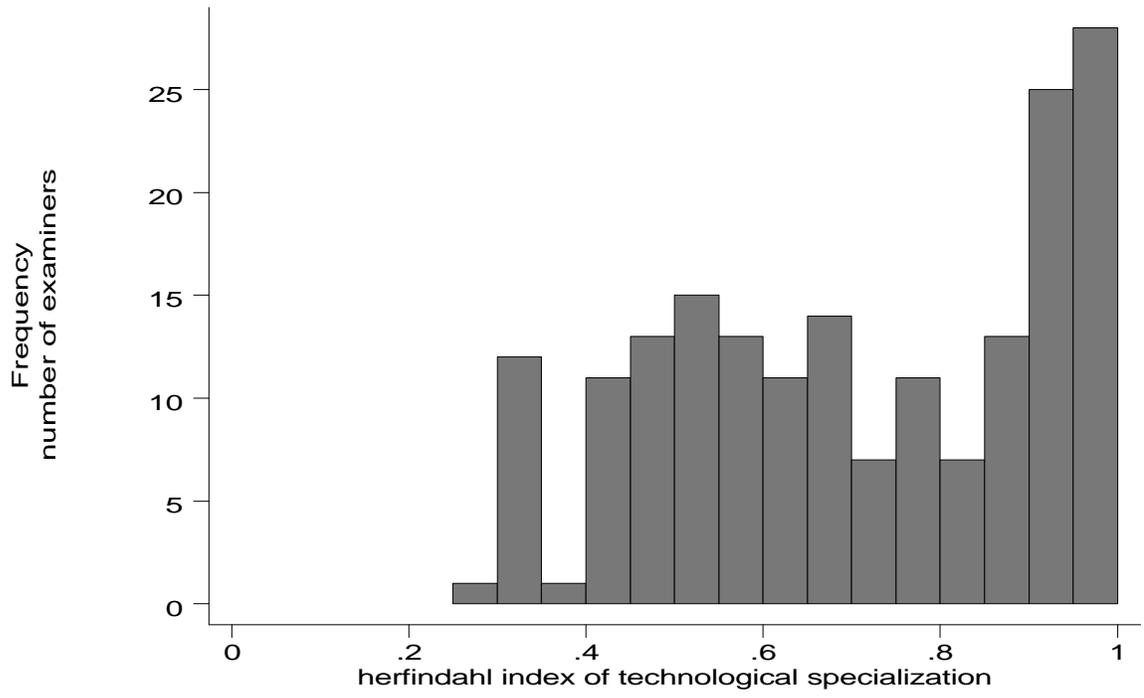


FIGURE 4: Citations Received by Examiners

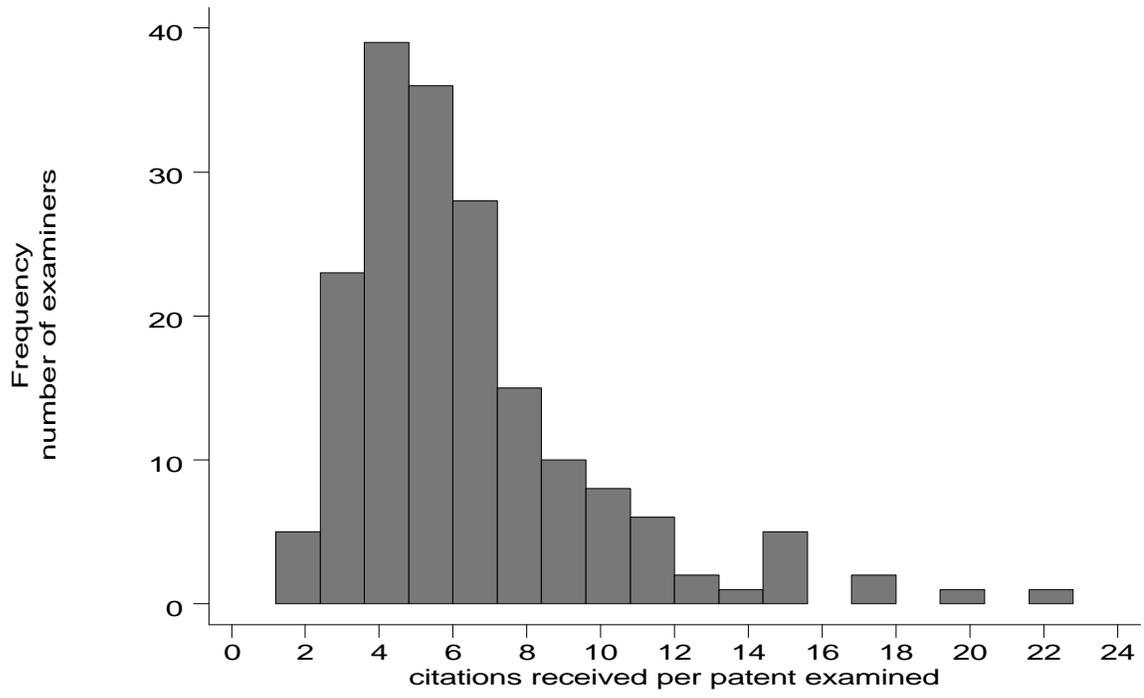


FIGURE 5: Self Citations by Examiners

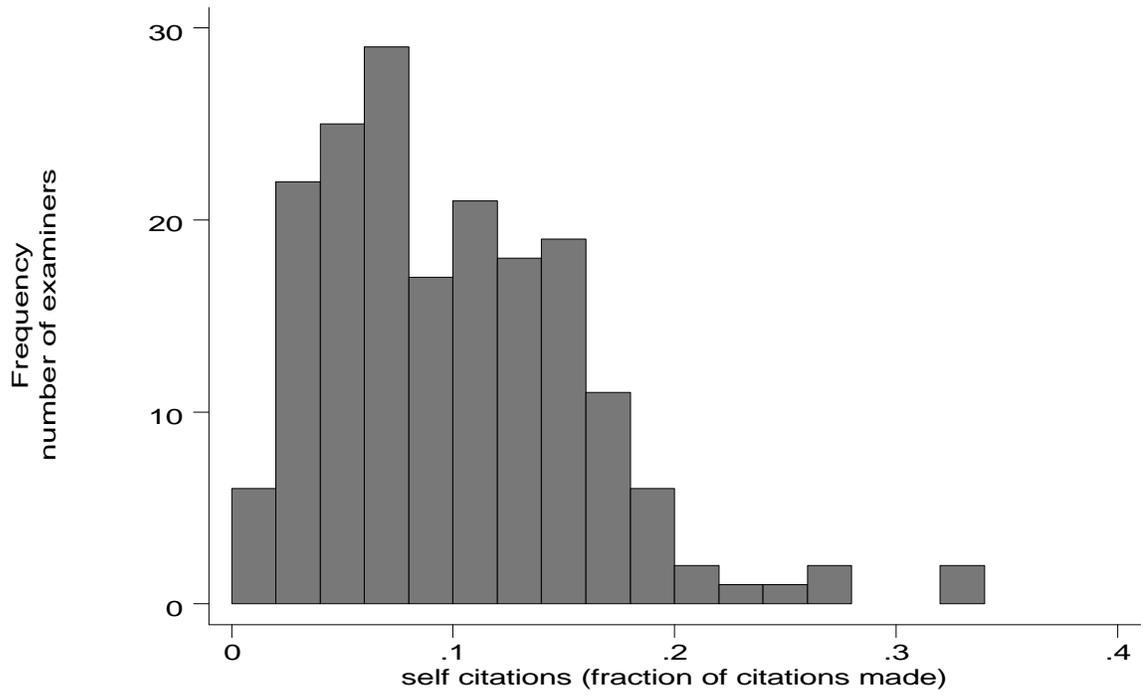


FIGURE 6: CAFC Patents by Technology

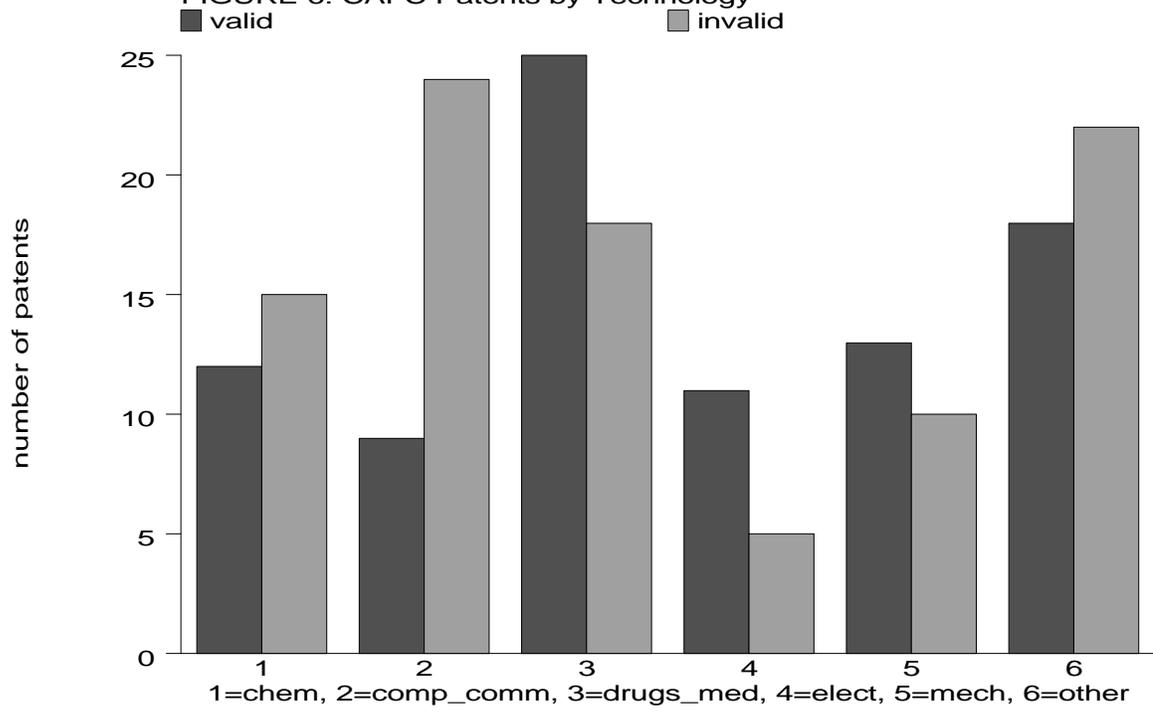


FIGURE 7: CAFC Patents by Year Issued

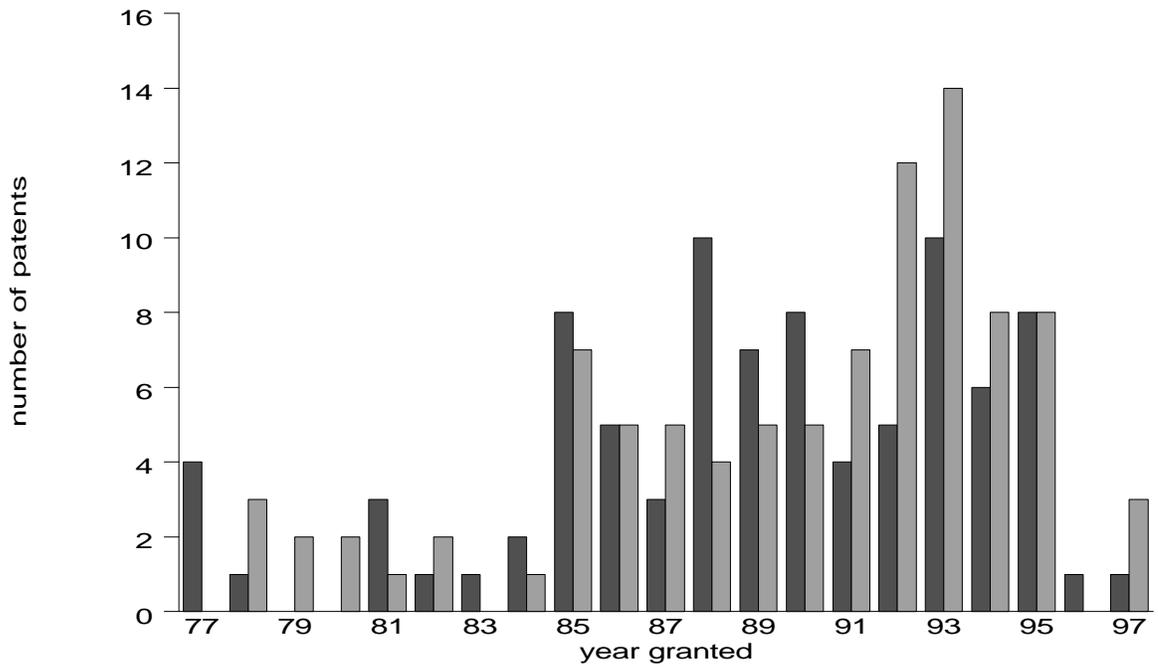


FIGURE 8: CAFC Patents by Assignee Type

