NBER WORKING PAPER SERIES

PATENT EXAMINER SPECIALIZATION

Cesare Righi
Timothy Simcoe

Working Paper 23913
http://www.nber.org/papers/w23913

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2017

We thank Jim Bessen, Christian Catalini, Iain Cockburn, Stuart Graham, Deepak Hegde, Evgeny Klochikhin, Alan Marco, Michael Meurer, Arti Rai, Bhaven Sampat, Andrew Toole, Joel Waldfogel, Heidi Williams along with seminar participants at the BU Technology Policy Research Initiative (IP Day), the BU Strategy & Innovation department, the Copenhagen Business School, the Max Planck Institute for Innovation and Competition and the NBER productivity lunch for useful comments and suggestions. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Patent Examiner Specialization
Cesare Righi and Timothy Simcoe
NBER Working Paper No. 23913
October 2017
JEL No. H83,K11,L98,O3,O34

ABSTRACT

We study the matching of patent applications to examiners at the U.S. Patent and Trademark Office. Using test statistics originally developed to identify industry agglomeration, we find strong evidence that examiners specialize in particular technologies, even within relatively homogeneous art units. Examiner specialization is more pronounced in the biotechnology and chemistry fields, and less in computers and software. Evidence of specialization becomes weaker, but does not completely disappear, if we condition on technology sub-classes. There is no evidence that certain examiners specialize in applications that have greater importance or broader claims. More specialized examiners have a lower grant rate and produce a larger narrowing of claim-scope during the examination process. We discuss implications for instrumental variables based on examiner characteristics.

Cesare Righi
Boston University
School of Law
Technology & Policy Research Initiative
765 Commonwealth Avenue
Boston, MA 02215
crighi@bu.edu

Timothy Simcoe
Boston University
School of Management
595 Commonwealth Avenue
Boston, MA 02215
and NBER
tsimcoe@bu.edu

1 Introduction

In 2015, the U.S. Patent and Trademark Office (USPTO) received 589,410 utility patent applications. Matching each application to a qualified examiner is a fundamental part of the examination process. This matching proceeds in two steps. First, each application is assigned to an "art unit" comprised of several examiners who specialize in a particular technology. Then the application is assigned to an individual examiner within that art unit. Motivated by the accounts in Cockburn, Kortum & Stern (2002) and Lemley & Sampat (2012), several studies have suggested that the second step in this process is more-or-less random, and then, building on an idea first proposed by Sampat & Williams (2015), used examiner characteristics as an instrumental variable for examination outcomes.[1]

We re-examine the random matching assumption, and find strong evidence of technological specialization by patent examiners within art units. Examiner specialization is more pronounced in the art units that examine Biotechnology and Chemistry patent applications, and less so in the computer-related technology centers. Evidence of specialization becomes weaker, but does not completely disappear, when we condition on U.S. Patent Classification System (USPC) sub-classes. However, we find no evidence that certain examiners specialize in applications that have greater importance or broader claims. Finally, we show that more specialized examiners have a lower grant rate and produce a larger narrowing of claim-scope during the examination process.

These findings have implications for instrumental variables strategies based on examiner characteristics. Random assignment would suffice to make leave-one-out grant rates (or any other examiner characteristic) uncorrelated with potential outcomes. But when examiners specialize, their individual characteristics are likely to be correlated with technology, suggesting an alternative path through which the instrument could influence outcomes. To see whether this matters in practice, we estimate "first stage" OLS regressions for two

[1] Papers adopting variants on this identification strategy include Farre-Mensa, Hegde & Ljungqvist (2015), Feng & Jaravel (2016), Gaulé (2015), Kuhn (2016), Kuhn, Roin & Thompson (2016), and Sampat & Williams (2015).

examination outcomes (patent issuance and first-claim scope) and find that adding USPC subclass fixed-effects modestly changes our estimates. This does not necessarily invalidate examiner-based instruments. However, it implies that the IV strategy rests on a stronger assumption than is typically acknowledged: potential outcomes must be uncorrelated with (unobserved) technological heterogeneity, despite observed technological sorting.

This is the first paper to systematically test the random matching hypothesis across all of the technology areas examined by the USPTO. Our methods for detecting specialization are borrowed from the literature on industry agglomeration (Mori, Nishikimi & Smith 2005). Specifically, we compute a pair of test statistics that ask whether application characteristics (e.g. technology subclass) are less dispersed across examiners than we would expect under random assignment.[2] Our main tests are performed at the art-unit-year level, and we examine the entire distribution of p-values for various application characteristics, including technology subclass, assignee, and indicators of patent value (family size) and scope (first independent claim length).

At a substantive level, our findings illustrate how the USPTO manages a tension between efficiency and fairness (Merges 1999). One way to promote fairness is through uniform application of patentability criteria, but prior research suggests that this is difficult. Some examiners are simply tougher than others (Sampat & Williams 2015, Kuhn et al. 2016), and experienced examiners are more lenient on average, partly because of time constraints (Lemley & Sampat 2012, Frakes & Wasserman 2014). Random matching provides another path to fairness, but forgoes the efficiency benefits of further technological specialization. Our analysis shows that the amount of specialization varies across art units, leading some applicants to get tougher examiners on average. But we find no evidence that particularly important applications (with large families) or broad applications (with short first independent claims) are assigned to specific examiners.

We discuss two plausible explanations for our finding that examiners are

[2]These methods focus specifically on the null hypothesis of random assignment, unlike IV falsification tests that ask the slightly different question of whether examiner and application characteristics are correlated.

more specialized in Chemistry and Biotechnology than in the computer-related art units. One possibility is that "generalist examiners" are able to evaluate computing inventions, while more specialized skills and knowledge are required in chemistry and life sciences. Another possibility is that the USPC technology classification system works better in chemistry and biotech, so we fail to observe much of the specialization that takes place within computer-related art units. Distinguishing between these hypotheses is a good topic for future research.

Finally, we find a positive correlation between specialization and a more stringent examination process, suggesting that it is easier for examiners who specialize to find relevant prior art. Under random matching, these estimates have a causal interpretation. Alternatively, they remain important for showing how non-random matching is related to examination outcomes.

The paper proceeds as follows. Section 2 describes how the USPTO assigns applications to examiners. Section 3 explains our methods and data. Section 4 discusses results and implications, and Section 5 concludes.

2 Patent Examiner Assignment at the USPTO

When a patent application is filed, the Office of Patent Application Processing reviews the formality requirements of the application and assigns it a serial number. A contractor defines the technological classification of the application using USPC class and subclass codes.[3] Each application has at least one mandatory classification, which is defined as a unique combination of class and subclass identifiers. The current version of the USPC has roughly 450 classes and more than 150,000 subclasses.

The USPTO has eight Technology Centers (TCs) responsible for examination of utility patent applications in broad technological areas. Each TC

[3] The two main purposes of the USPC are to facilitate the retrieval of technical documents and to ease the allocation of applications to the examining personnel specialized in a particular technology. For details, see `http://www.uspto.gov/sites/default/files/patents/resources/classification/overview.pdf`. Although it was replaced by the Cooperative Patent Classification (CPC) on January 1, 2013, the USPC is the relevant classification for the entire period of our study.

is comprised of several art units, or teams of patent examiners who specialize in a particular technology. Technological classifications are used to assign each new patent application to a specific art unit.[4]

Within each art unit the initial assignment of a new application is handled by a Supervisory Patent Examiner (SPE). The SPE can refine the technological classification of a new application if it is incorrect, or request that an application be transferred to another art unit.[5] But in most cases, the SPE will assign the application to an examiner within her art unit. This is the step we analyze below.

Previous research documents that SPEs have substantial discretion in examiner assignment. Some SPEs interviewed by Lemley & Sampat (2012) mention assigning applications to examiners essentially randomly within subclasses. Other SPEs give the oldest unassigned application to an examiner when she finishes the examination of another application. Although these practices suggest random matching, some SPEs may encourage technological specialization of examiners within their art-unit, and specialization could also emerge through trading among individual examiners. Cockburn et al. (2002) suggest that the degree of technological specialization varies across art units – in some art units an individual examiner is responsible for almost all applications in a specific technology class, and in others the examiners are less specialized.

Although the USPTO constantly monitors the performance of art units and examiners to ensure a certain level of quality of the examination process, the assignment to a particular art unit and to a specific examiner can have important consequences for an application. Different practices across art units and the personal approach of each examiner can affect whether an application is eventually granted (Sampat & Williams 2015), how quickly a decision is reached (Farre-Mensa et al. 2015), and the scope and strength of an issued patent (Kuhn et al. 2016). This variation in standards led Cockburn

[4]For the current list of classes and subclasses examined by each art unit, see `http://www.uspto.gov/patents-application-process/patent-search/understanding-patent-classifications/patent-classification`.

[5]The Manual of Patent Examining Procedure sec. 903.08 describes the rules governing assignment and transfer of applications between art units.

et al. (2002) to conclude that “there may be as many patent offices as patent examiners.”

3 Methods and Data

We use two statistical tests originally developed to analyze industry agglomeration. In our application, patent examiners are analogous to cities, and technology subclasses (or other application characteristics) are analogous to industries. Each test compares observed frequencies to the distribution under a null of random assignment.

3.1 Agglomeration Test Statistics

3.1.1 Divergence Index

The D-index was developed by Mori et al. (2005), building on Kullback & Leibler (1951), and is based on the concept of relative entropy.[6] Suppose we have a set of applications characterized by category $i \in \mathbf{I} = \{1, ..., I\}$, assigned to a set of examiners denoted by $r \in \mathbf{R} = \{1, ..., R\}$. In our application, the categories i may correspond to USPC subclasses, assignees or any other predetermined observable characteristic of a patent application. Under random allocation, examiner r's share of all applications from category i should equal her share of the overall population.

To formalize that idea, define n_{ir} as the number of applications in category i assigned to examiner r, and $N_i = \sum_{r=1}^{R} n_{ir}$ as the total number of applications in category i. The reference distribution $p_0 = (p_{0r} : r \in \mathbf{R})$, where $p_{0r} = \frac{\sum_{i=1}^{I} n_{ir}}{\sum_{i=1}^{I} N_i}$ measures examiner r's share of all applications, is the share we expect her to be allocated from each category under the null of random assignment.

Let p_{ir} denote the true probability that a randomly sampled application in category i is assigned to examiner r, so the distribution across examiners

[6]Statisticians often refer to the D-index as a G test statistic. The main advantage of a G-test relative to a chi-squared test of independence occurs when some cells in a frequency table have very small expected counts, which is the case in our application.

for the category is $p_i = (p_{ir} : r \in \mathbf{R})$. We can measure the divergence between p_i and p_0 using the relative entropy of p_i with respect to p_0, called the D-index by Mori et al. (2005):

$$D(p_i|p_0) = \sum_{r \in \mathbf{R}} p_{ir} \ln\left(\frac{p_{ir}}{p_{0r}}\right).$$

$D(p_i|p_0)$ is nonnegative, achieves its minimum at $p_i = p_0$ and its local maxima when all applications in category i are assigned to a single examiner.

To estimate the D-index, we use the observed data to estimate the probabilities p_{ir}, with $\hat{p}_{ir} = \frac{n_{ir}}{N_i}$, thus estimating:

$$D(\hat{p}_i|p_0) = \sum_{r \in \mathbf{R}} \hat{p}_{ir} \ln\left(\frac{\hat{p}_{ir}}{p_{0r}}\right). \tag{1}$$

These probability estimates converge to the true value exponentially fast with the increase in sample size for a given category N_i.

As shown by Mori et al. (2005) the D-index can be related to the the log likelihood ratio (λ):

$$-\frac{\ln \lambda}{N_i} = \sum_{r \in \mathbf{R}} \frac{n_{ir}}{N_i} \ln\left(\frac{\hat{p}_{ir}}{p_{0r}}\right) = D(\hat{p}_i|p_0).$$

Given that $-2\ln\lambda$ is distributed asymptotically as a chi-square with $R-1$ degrees of freedom, we can use this relationship to test the null hypothesis that $p_i = p_0$ (see Mori et al. (2005) for details).[7] In our application, the number of tests will equal the number of categories (e.g. one per technology subclass) and we examine the distribution of p-values from all of these tests conditional on a given sample-size threshold (e.g. $N_i > 20$).

3.1.2 Multinomial Test for Agglomeration and Dispersion

MTAD computes multinomial likelihood functions for an allocation of agents to a set of discrete locations. In our application, the agents are patent ap-

[7] In practice, we compute $2N_iD(\hat{p}_i|p_0)$ and use it for a chi-square test with $R-1$ degrees of freedom.

plications and locations correspond to examiners. If the likelihood of the observed data is lower (higher) than the likelihood under random choice, MTAD indicates that the agents are agglomerated (dispersed). This approach differs from the D-index because the statistic is computed for an entire art unit, and because it can detect whether deviations from random assignment are due to agglomeration or over-dispersion.

To provide a brief formal description of MTAD, we adapt the notation provided in Rysman & Greenstein (2005). Suppose we have R examiners, each receiving n_r applications, with $r = 1, ..., R$. The variable n_r is bounded between $\underline{n} = 0$ and $\overline{n} = \infty$ and distributed according to the discrete distribution $f(n_r)$. Each examiner can be assigned applications of c types. The unconditional probability of being assigned type c is p_c for $c = 1, ..., C$. The observed number of applications of type c assigned to examiner r is x_r^c. Define $\mathbf{x}_r$ as the vector of elements $x_r^1, ..., x_r^C$, $\mathbf{p}$ as the vector of probabilities $p_1, ..., p_C$, $\mathbf{n}$ as the $R \times 1$ vector of applications assigned to each examiner, and $\mathbf{X}$ as the $R \times C$ matrix of allocations. If examiners are assigned applications independently, the likelihood of observing outcome $\mathbf{x}_r$ for examiner r is the multinomial pdf

$$\mathcal{L}(\mathbf{x}_r, n_r, \mathbf{p}) = \binom{n_r}{x_r^1, ..., x_r^C} p_1^{x_r^1} ... \, p_C^{x_r^C}$$

and the average log-likelihood for the data is

$$l(\mathbf{X}, \mathbf{n}, \mathbf{p}) = \frac{1}{R} \sum_{r=1}^{R} \ln \Big(\mathcal{L}(\mathbf{x}_r, n_r, \mathbf{p}) \Big).$$

We want to compare this log-likelihood with the value we would observe under independent random assignment. Let the random variable $l(f, \mathbf{p})$ be distributed according to the distribution $l(\mathbf{X}, \mathbf{n}, \mathbf{p})$ if $\mathbf{X}$ was *actually* drawn from a multinomial distribution and n_r was drawn from f. Then the expected log-likelihood under random allocation is given by

$$E[l(f, \mathbf{p})] = \sum_{n_r} \Bigg(\sum_{\mathbf{z} \in \Phi(n_r)} \ln \mathcal{L}(\mathbf{z}, n_r, \mathbf{p}) \times \mathcal{L}(\mathbf{z}, n_r, \mathbf{p}) \Bigg) f(n_r)$$

where $\Phi(n_r)$ is the set of all possible allocations of the n_r applications. To compute $E[l(f, \mathbf{p})]$ we treat $\mathbf{p}$ as known and take f to be the empirical distribution of n_r. The MTAD test-statistic is

$$t(\mathbf{X}, \mathbf{n}, \mathbf{p}) = l(\mathbf{X}, \mathbf{n}, \mathbf{p}) - E[l(f, \mathbf{p})]. \tag{2}$$

A negative (positive) value of $t(\mathbf{X}, \mathbf{n}, \mathbf{p})$ signals agglomeration (dispersion) of patent applications compared to the null of random assignment. This statistic is distributed asymptotically normal and we use simulation to generate its confidence intervals.[8]

3.2 Data

Our main data source is the USPTO Patent Examination Research Dataset (Graham, Marco & Miller 2015), which is based on information from the Public Patent Application Information Retrieval system (Public PAIR). We also use information from PatentsView (`http://www.patentsview.org`), PATSTAT, the USPTO Patent Assignment Dataset (Marco, Myers, Graham, D'Agostino & Kucab 2015) and the Patent Claims Research Dataset (Marco, Sarnoff & deGrazia 2016).

We restrict our analysis to published utility patent applications filed on or after the enactment of the American Inventor's Protection Act of 1999 (November 29, 2000) and before January 1st 2013, whose examiner is affiliated with one of the eight technology centers responsible for the examination of utility patent applications. The USPTO Patent Examination Research Dataset provides information on the examiner of record for each application as of January 24, 2015. This is the examiner as of that date for pending applications and the examiner at the time of disposal for disposed applications. We assign art units based on the examiner of record at the time of the last office action recorded for an application. Under the AIPA, regular utility

[8]See Rysman & Greenstein (2005) for details on the test. Timothy Simcoe developed a software module to easily perform this test in Stata, available at the following link: `https://ideas.repec.org/c/boc/bocode/s457205.html`

patent applications are generally published eighteen months after filing.[9]

The data have several limitations. First, applications will not appear in our data if they are abandoned before publication, or if the applicant files only in the United States and requests that the application not be published. Previous research suggests that these outcomes are relatively rare.[10] A second limitation is that we do not observe whether applications are transferred from one examiner to another.

We exclude applications filed after 2012 to avoid problems related to publication lags and a change in the USPTO technological classification scheme. We also exclude serialized continuations (continuation applications, continuations in part and divisional applications) because these applications are usually assigned to the same examiner of the original application, and would therefore lead us to overstate the extent of agglomeration. Our primary analysis sample contains 2,717,032 applications examined by 12,338 unique examiners affiliated with 590 art units. Table 1 shows the distribution of applications, art units, examiners, classes and subclasses by technology center.

3.3 Variables

We focus on several application characteristics (indexed by i or c above) that may influence the assignment of applications to individual patent examiners within an art-unit-year.[11] The first of these characteristics is the primary USPC classification of the application, which is defined by a unique combination of primary class and primary subclass codes (for brevity, subclass). If patent examiners specialize in evaluating applications related to particular technologies, we expect to see agglomeration on this variable.

We use technology classification data from published applications, rather than granted patents, to avoid measuring any specialization created by the

[9] As in Graham et al. (2015) and in the Public PAIR data, we use the term "regular utility patent application" to distinguish nonprovisional utility patent applications from provisional, PCT, reissue or re-examination applications.

[10] Graham et al. (2015) show that about 95% of the regular non-provisional utility patent applications filed between 2001 and 2012 can be found in Public PAIR.

[11] We typically compute our test statistics within a filing-year-art-unit cell to account for possible changes in assignment practices over time and turnover in the pool of examiners.

examination process. In particular, because the USPC classification of an application is based on its claims, which are usually amended during examination, the subclass of many applications changes over time.[12,13] This could lead to spurious agglomeration if certain examiners are more likely to reject claims in particular classes.

Table 1 shows that for patents granted before July 21, 2015, twenty percent of all applications change primary class during the examination process, and almost seventy percent change primary subclass. There is heterogeneity across technology centers, with patents in Biotechnology and Chemicals changing classification more often than those in other areas.[14]

The identity of the applicant is a second variable that could influence the allocation of applications — either directly or due to technological specialization. We measure this with the assignee of an application. Specifically, we retrieve information on the assignment of applications, identify the assignments made by the inventors to their employers before the application is docketed to an examiner, clean and standardize the assignee names and create clusters of names that are likely to belong to the same organization, to which we assign a unique identifier.[15] After completing this process, we have missing assignee data for 584,313 applications. To check the robustness of our assignee measurement, we utilize a second measure of the applicant identity: the customer number assigned by USPTO to each application. This number

[12] The data in Public PAIR provide only the most recent classification of an application, so we utilize the primary classification of applications at publication from PatentsView, which is more likely to reflect the classification contractor's original assignment. We thank Evgeny Klochikhin for access to the PatentsView patent applications database.

[13] In a previous version of the paper we utilized the class/subclass codes provided by Public PAIR for our agglomeration analysis. The results were similar to those reported in the current version of the paper, but showed a greater degree of agglomeration. We thank Deepak Hegde, Bhaven Sampat, Andrew Toole and Heidi Williams for helpful conversations that improved our understanding of the classification process.

[14] Many papers utilize USPC (sub)classes as a control variable, and future research might usefully consider whether it is better to measure this variation at the time of application publication or grant.

[15] We employ an assignee name cleaning and standardization routine that builds upon Thoma, Torrisi, Gambardella, Guellec, Hall & Harhoff (2010) and the name standardization routines developed for the NBER Patent Data Project available at `https://sites.google.com/site/patentdataproject/Home/posts/namestandardizationroutinesuploaded`. Details are available upon request.

identifies the correspondent for application-related matters and is usually either the law firm representing the applicant or the legal department of the firm filing the application.[16]

We would like to examine whether some examiners are assigned a larger share of "high value" applications. The size of a patent family is often used as a proxy for economic value of the invention because increased value leads patentees to file in more countries (Harhoff, Scherer & Vopel 2003, Putnam 1996). We count the number of applications in the same DOCDB patent family, with filing dates on or before the focal application date, to construct an indicator variable that equals one if a focal application is above the 95$^{\text{th}}$ percentile in the family size distribution (within an art-unit and filing-year). We call this variable "DOCDB Family Size."[17]

Finally, we consider whether some examiners are assigned applications seeking greater scope of protection. Kuhn et al. (2016) show that the length of the first independent claim in a patent is a good measure of patent scope. The idea behind this measure is that shorter claims provide broader scope of patent protection because every word added to the text of the claims can potentially introduce additional elements or characteristics that must be present to establish infringement. We create an indicator variable that equals one if and only if a patent application falls below the 5$^{\text{th}}$ percentile of the word count distribution for the first independent claim in the subsample of applications with the same filing year examined by the same art unit.[18] We call this variable "Words in 1$^{\text{st}}$ Claim."[19]

[16]Results of the customer-number analysis are similar to those for the assignee and are available upon request.

[17]We test the robustness of these results using the INPADOC patent families. The results are similar to those for DOCDB patent families and are available upon request.

[18]Kuhn et al. (2016) note that this measure of scope is not suitable for the analysis of patent scope in biotechnology. So we exclude the Biotechnology technology center from the analysis of this variable. We also check the robustness of results based on the length of the first independent claim utilizing measures built upon the number of claims and independent claims. The results are similar and available upon request.

[19]Summary statistics for all variables used in the analysis are in Table A1.

4 Results

This section presents evidence of patent examiner specialization, and then regression results linking specialization to examination outcomes.

4.1 Examiner Specialization

Figure 1 shows that patent examiners handle more applications from a given USPC subclass or assignee than we would expect under random allocation. Specifically, each panel shows a histogram of p-values from a sample of hypothesis tests. For the D-index (top row), we run a separate test for each art-unit-year by subclass or assignee cell containing more than 20 applications. For MTAD (bottom row) we run a separate test for each art-unit-year cell containing more than 50 applications.[20]

Under the null of random assignment, the p-values in Figure 1 should be uniformly distributed between zero and one. However, in each panel a large share of the test-statistics fall below the 1 percent statistical significance threshold, providing strong evidence of specialization. The two histograms in the left column indicate that about 25 percent of the D-index and MTAD tests for random USPC assignment have a p-value below 0.01. The two histograms in the right column show somewhat weaker evidence of specialization by assignee, with about 10 to 20 percent of the p-values falling below the 1-percent threshold. The agglomeration by assignee becomes much weaker if tests are conducted within USPC subclasses (see below), suggesting that these findings are primarily a result of technological specialization of examiners and applicants. Overall, Figure 1 shows that the allocation of applications within art units is often far from random, and that SPEs take into account the technological classification when assigning applications to an examiner, as described in Lemley & Sampat (2012).

[20] All of our results are robust to varying the within-cell sample size cutoffs, but going much below these thresholds leads to large numbers of uninformative tests. Figure A1 shows the distributions of p-values of D-index and MTAD for subclass and assignee with thresholds equal to, respectively, 10 and 25. Figure A2 uses these lower thresholds to check the robustness of the analysis in Figure 2.

Table 2 examines the degree of examiner specialization in different Technology Centers, and for an additional pair of application characteristics. Specifically, the table reports the share of D-index or MTAD tests that reject the null hypothesis of random allocation at a 1-percent significance level.[21] Panel A shows that there is evidence of examiner specialization in every technology center. However, the "Computer Architecture" and "Computer Networking" areas are less agglomerated than Biotechnology, Chemistry, Semiconductors and Mechanical Engineering. The results in Panel B are similar.

Although our data do not speak to the underlying causes of variation in examiner specialization across technology centers, there are several possible explanations for this pattern. First, examiners in the less agglomerated technology centers may be "generalists" who are capable of evaluating most applications within their art-unit. This would naturally lead SPEs to adopt a more random allocation process, and might also influence examiners' application-trading practices. Alternatively, patent examiners in the Computers and Communications technology centers might be just as specialized as their counterparts, but this is not apparent to us because the USPC classification system is less representative of actual technological differences in these fields.

The lower half of Table 2 examines agglomeration for a pair of dichotomous variables: "DOCDB Family Size" and "Words in 1st Claim." Both of these variables focus on extreme outcomes because we are interested in whether SPEs assign unusual applications to a specific set of examiners. The data suggest that, for the most part, they do not. There is some evidence that applications from very large families are concentrated among a smaller set of examiners for Chemicals, Communications, Semiconductors, Mechanical Engineering and the technology center we labeled as "Miscellaneous". And there is some evidence that certain examiners specialize in broader patents (as measured by length of the first claim) in the Chemical and Materials Engineering and Semiconductors technology centers. But these effects are not

[21] Table A3 in the appendix reports analogous figures with a cutoff at the 5-percent threshold for statistical significance.

large, and might easily be caused by the technological specialization observed in Panel A.

The results presented thus far beg the question of whether examiner specialization is purely technological. To explore that idea, we test for agglomeration *within* art-unit-year-USPC-subclasses to see whether conditioning on observed technological heterogeneity changes our results. There are two caveats to keep in mind. First, we cannot condition on *unobserved* technological heterogeneity. And second, many USPC subclasses receive only a few applications per year, so these tests exclude a large amount of data. However, if examiners seem to be randomly assigned within large sub-classes, we might be more comfortable that most of the specialization we observe within art-unit-years is based on technology rather than other patent characteristics.

Figure 2 examines agglomeration by assignee, within art-unit-years, both within and without conditioning on USPC subclass. Each panel presents a quantile-quantile plot that compares the distribution of the D-index (top row) or MTAD statistic (bottom row) for the observed data to the distribution under simulated random assignment. For the D-index, agglomeration leads observed values of D_i to exceed simulated values of D_i, so that scatter points fall below the 45-degree line. For MTAD, a negative value of $t(X, n, p)$ indicates agglomeration, and a positive test-statistic indicates over-dispersion. So the scatter points will fall above the 45-degree for $t < 0$ when there is agglomeration, and below the 45-degree for $t > 0$ when there is over-dispersion.

The left column in Figure 2 shows that the observed quantiles of the D-index are higher, and the observed quantiles of MTAD are lower (for $t < 0$), than the simulated quantiles under random allocation. In other words, there is strong evidence of specialization, as we saw above. The righthand column shows that the evidence for agglomeration is much weaker once we condition on the USPC subclass, although the MTAD test does appear to detect some specialization by assignee.[22] Note how the sample size falls dramatically as we move from the left to the right column in this figure.[23]

[22]In both columns, we observe a similar distribution of $t(X, n, p)$ when the statistic is positive, suggesting that any over-dispersion is in fact random.

[23]The D-index discards any assignee that does not submit more than 20 applications to a given art-unit in a particular year (and, for the analysis conditional on USPC subclasses,

Table 3 examines agglomeration within art-unit-application-year-class-subclass bins, and reports results by technology center.[24] For the analysis in Panel A we retain all the applications with at least one secondary subclass (1,311,532 applications) and generate a data set with observations at the application-secondary-subclass level. The D-index analysis rejects random assignment for about 20% of the tests in Biotechnology, suggesting agglomeration. However, about 13% of the MTAD tests in Biotechnology and Chemicals reject random assignment in favor of dispersion.

Panel B focuses on the allocation of assignees. For the D-index, we often have a small number of tests, because it is unusual for a single assignee to file many applications in a single subclass-year. Nevertheless, the D-index tests reject random assignment more than 15 percent of the time in the Biotechnology technology center and almost 10 percent of the time in the Chemistry technology center. MTAD rejects random assignment less often, but also finds more agglomeration in biotechnology and, to a lower extent, chemistry. Because several technology centers have only a few subclasses large enough to produce reliable inference, we re-ran this analysis after pooling all years in our sample, and found very similar results (see Appendix Tables A5 and A6).

Panels C and D in Table 3 find no evidence that SPEs in any technology center allocate "outlier applications" (in terms of family size or first independent claim scope) to a specific set of examiners after conditioning on observed technological differences.

Overall, these results show that patent examiners specialize in particular technologies, even within relatively homogeneous art units. We find no evidence that certain examiners specialize in "outlier" patent applications. Moreover, much of the agglomeration by assignee disappears if we condition on primary USPC subclasses. However, we do find evidence of agglomeration or over-dispersion by secondary USPC subclass and assignee, even within primary USPC subclasses, for the Biotechnology and Chemistry technology

in a particular subclass). This excludes the large majority of applicants. MTAD retains more data because it uses all applications filed to an art unit-year-(subclass).

[24]Table A4 presents the same analysis using a 5% statistical significance threshold.

centers. This last result suggests that there may be examiner specialization based on unobserved technological differences in some art units even after conditioning on the observed USPC subclasses.[25]

4.2 Implications for Examiner-based Instruments

Examiner specialization undermines a common justification for using examiner characteristics as instrumental variables. Under random matching, examiner characteristics are uncorrelated with potential outcomes by construction. But specialization implies that these characteristics can be correlated with technology, which might lead to a violation of the exclusion restriction.

It is not possible to test the exclusion restriction. However, it is possible to examine whether observed technologies are correlated with examiner characteristics. We do this by estimating a "first stage" OLS regression with and without subclass fixed-effects, and testing whether the coefficient on a proposed instrument changes. This exercise can be viewed as an additional test of the random assignment hypothesis, since random matching implies that observed (and unobserved) technology is uncorrelated with examiner characteristics. Alternatively, one can take it as a test of the weaker assumption that the proposed instrument is uncorrelated with subclass effects.

Table 4 presents first stage OLS estimates for two instruments used in the literature: (1) an examiner's leave-one-out grant rate on patent issuance, and (2) an examiner's leave-one-out "scope change" (i.e. the mean number of words added to the first independent claim of other applications) on the scope of a focal patent.[26] The sample for this analysis contains all applications in our primary sample that were either granted or abandoned by the end of

[25] To complement the analysis describe in this section, we also run a set of Kolmogorov-Smirnov tests of the equality of distributions of the statistics produced by the D-index and MTAD analysis and their p-values for the tests on the real allocations and the simulations of random assignment. The results are consistent with those reported in the paper and are available upon request.

[26] Since examiners can change art unit within a given year, we compute both leave-one-out variables within art-unit-filing-year-examiner. Given the results of our agglomeration analysis, we would like to compute the two instruments within subclasses. However, this would exclude more than 90 percent of our data because many subclasses have only a very small number of applications in a given year.

the sample period; whose leave-one-out IV is computed with at least 10 applications; and whose art-unit-filing-year-subclass cell in the estimation sample contains at least two applications. The estimates for the patent scope models use only granted applications. Each coefficient reported in the table is based on a separate OLS regression. To ease the interpretation of the results, we standardize the two instruments, and report the ratio of the two estimates (with and without subclass effects) in a third column.

The full sample estimate in the first row of Table 4 implies that a one standard deviation increase in the leave-one-out grant rate is associated with a 16.3% increase in the probability that a focal application is granted. Adding subclass fixed effects causes this coefficient to fall by 10 percent. The remainder of columns (1) and (2) show that adding subclass effects produces a decline in the first-stage IV coefficient within every technology center. The changes are larger in Biotech and Chemistry, where evidence of specialization is stronger, and weaker in the computer-related technology centers. All of these differences are statistically significant at the 5 percent level.[27] Columns (3) and (4) in Table 4 present similar findings for the leave-one-out scope-change instrument. For the latter variable, we find a particularly large change in coefficient estimates for applications assigned to art-units in the Chemicals technology center, and no change for applications in computing technology center 2100.

Based on these results, we can (again) reject the random matching hypothesis, and also the conjecture that observed technology is uncorrelated with a specific pair of examiner-based instruments. On the other hand, controlling for subclass produces only modest changes in the first-stage estimates, particularly for several of the computing-related technology centers.

We propose the following implications for those who still wish to use ex-

[27] To test whether differences between the coefficients in columns (1) and (2) are statistically significant, we demean the leave-one-out variables within art-unit-filing-year and within art-unit-filing-year-subclass, re-run the models (without the fixed effects, as demeaning within groups at the level of the fixed effects produces the same coefficients) and test the statistical significance of the differences in the coefficients. We also run a battery of likelihood ratio tests to compare models analogous to those in columns (1)-(2) and (3)-(4) without clustering the standard errors. All tests are statistically significant at the 1 percent level.

aminer characteristics as instruments, perhaps because (like us) they see it as a clear step forward in terms of measuring the causal impacts of intellectual property. First, it is important to carefully control for any observable differences in technology. Subclass fixed effects are not a panacea, since there is almost certainly some residual unobserved technological specialization, but they are a step in the right direction. Second, instead of claiming that applications are randomly matched to examiners, authors should clearly explain the key identification assumption: conditional on observables, examiner characteristics must be uncorrelated with potential outcomes, regardless of any technological sorting. Finally, the assumption that technology is uncorrelated with potential outcomes appears most plausible for information technology art-units, a bit less so for mechanical and miscellaneous technologies, and potentially problematic for art-units examining chemical and biotechnology applications.

4.3 Specialization and Examination Outcomes

As a final step in our empirical analysis, we explore the relationship between examiner specialization and patent prosecution outcomes. We focus on three outcomes: (i) whether an application is granted, (ii) the change in the number of words in the first independent claim between the published application and the granted patent, and (iii) the number of days required to process the application (the difference between the date an application is docketed to an examiner for the first time and its disposal date).[28] Our sample consists of all applications belonging to an art-unit by examiner by filing year cell containing more than 10 applications. To account for truncation, we exclude pending applications and those filed after year 2009.

Our unit of analysis is the application, and we adopt a measure of specialization that varies across both examiners and applications. Specifically, our main explanatory variable is the share of an examiner's applications (within

[28]An application is never ultimately rejected by the USPTO. If an applicant is not granted a patent, she can file a Request for Continued Examination (RCE), a continuation application or a continuation-in-part. We do not study the implications for RCE or continuation filings in this paper.

an art unit-filing year cell) having the same USPC subclass as a focal application. To be more precise, define the set $k_{it}(j)$ of all patents (except for patent i) assigned to examiner j in year t.[29] Let n_{jt} represent the total number of patents reviewed by examiner j in year t, and define an indicator 1_{mn} that equals one if and only if two patents (m and n) have the same USPC subclass. Our main explanatory variable can be written as:

$$Share_{ijt} = \frac{\sum_{m \in k_{it}(j)} 1_{mi}}{n_{jt} - 1}.$$

Intuitively, $Share_{ijt}$ equals the probability that a random draw from the pool of applications assigned to examiner j in year t has the same USPC subclass as the focal application.

Table 5 presents estimates from a series of OLS panel-data regressions that examine the correlation between $Share_{ijt}$ and prosecution outcomes. To ease interpretation, we standardize $Share_{ijt}$ and the outcome variables except the dummy for granted patents.[30] Standard errors are clustered at art-unit-filing-year level in all models.

Columns (1) through (3) report coefficient estimates from a within-examiner regression with art-unit-examiner-filing-year fixed effects. The coefficients of $Share_{ijt}$ are all positive but very close to zero. Columns (4) through (6) report the results from a between-examiner analysis, where we regress the mean outcome for each art-unit-examiner-filing-year on the mean of $Share_{ijt}$ (i.e. the probability that two random draws from the pool of patents assigned to that examiner will belong to the same USPC subclass).

The coefficient in Column (4) indicates that a one standard deviation increase in $Share_{ijt}$ leads to a 4 percentage point drop in the grant rate. This suggests that specialized examiners are also more stringent. The coefficient in column (5) also suggests that specialization leads to more stringent examination. However, the economic magnitude of this result is rather small: a one standard deviation change in $Share_{ijt}$ produces a 0.08 standard deviation change in the number of words added to the first claim. Finally, in

[29]For this analysis we consider an examiner affiliated with two (or more) different art units in the same year as two (or more) examiners.

[30]Table A2 displays summary statistics for all variables used in this part of the analysis.

column (6) we find a small but statistically significant positive association between specialization and the time required to process a patent examination.[31]

The overall message of this part of the analysis is that examiner specialization is related to more stringent examination, although the economic magnitudes are not dramatic. This relationship is driven by differences across examiners, as showed by the "between" estimators in Table 5, while we do not find important differences in the relationship between specialization and examination outcomes "within" examiners. One plausible explanation for the finding is that it is easier for examiners that are more specialized to find relevant prior art because they are more familiar with certain fields of technology, leading to narrower claims and an increased probability of application abandonment. Under random assignment, these estimates are causal. We prefer a descriptive interpretation. Nevertheless, these results confirm the importance of differences across examiners for examination outcomes.

5 Conclusions

We study a key stage of patent prosecution: the assignment of applications to examiners. The first half of our empirical analysis focuses on characterizing the degree of examiner specialization. Using two statistical tests designed to study industry agglomeration, we find strong evidence that examiners specialize in particular technologies, even within relatively homogeneous art units. The degree of specialization varies across fields. The USPTO technology centers associated with Computers and Communications exhibit relatively little specialization, while examiners in the "Biotechnology and Organic Chemistry" and "Chemical and Materials Engineering" technology are relatively more specialized. In the latter technology centers, we find assignee agglomeration even after conditioning on USPC subclasses.

The second part of our analysis shows that observed technological classifications are correlated with potential instruments based on examiner char-

[31] Results of the analysis in Table 5 by technology center are available on request.

acteristics. And our last set of results shows that more specialized examiners are more stringent on average — they have a lower grant rate, and produce a larger reduction in the scope of issued patents' first independent claim.

It may not seem surprising that we can reject the hypothesis of random matching between applications and examiners. After all, one reason for having a patent classification system is to help route applications to appropriate examiners. However, several studies have argued that more-or-less random matching (within art-units) provides a justification for using examiner characteristics as an instrument for examination outcomes. While our findings do not invalidate this identification strategy – patent examiner characteristics might still satisfy the relevant exclusion restrictions – they do imply that we cannot rely on random assignment to justify the approach. Our findings suggest that examiner-based instruments are more plausible in studies that include subclass fixed effects and focus on computer-related art units.

On a more positive note, our results suggest that the USPTO's patent examination process strikes a reasonable balance between efficiency and fairness. Technological specialization is efficient. Fairness can be achieved by enforcing uniform examination standards, which is difficult, or through random assignment, which guarantees all applicants an equal shot at the more friendly examiners. Conditional on technology, examiner assignment appears relatively random in the computer-related technology centers. And even without controlling for technology, there is no evidence that certain examiners within a given art unit handle more patents with large families or broad claims. We leave to future researchers the question of whether procedural fairness to applicants is also the best policy in terms of social welfare.

References

Cockburn, I. M., Kortum, S. & Stern, S. (2002), 'Are all patent examiners equal? The impact of examiner characteristics', *NBER Working Paper 8980* .

Farre-Mensa, J., Hegde, D. & Ljungqvist, A. (2015), 'The bright side of patents', *Working Paper* .

Feng, J. & Jaravel, X. (2016), 'Who feeds the trolls? patent trolls and the patent examination process', *Working Paper* .

Frakes, M. D. & Wasserman, M. 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 20337).

Gaulé, P. (2015), 'Patents and the Success of Venture-Capital Backed Startups: Using Examiner Assignment to Estimate Causal Effects', *Working Paper* .

Graham, S. J. H., Marco, A. C. & Miller, R. (2015), 'The USPTO Patent Examination Research Dataset: A Window on the Process of Patent Examination', *Working paper* .

Harhoff, D., Scherer, F. M. & Vopel, K. (2003), 'Citations, family size, opposition and the value of patent rights', *Research policy* **32**(8), 1343–1363.

Kuhn, J. M. (2016), 'The right to exclude: Property rights and frictions in the sale of patents', *Working Paper* .

Kuhn, J. M., Roin, B. N. & Thompson, N. C. (2016), 'Causal inference on patent protection', *Working Paper* .

Kullback, S. & Leibler, R. A. (1951), 'On information and sufficiency', *The annals of mathematical statistics* **22**(1), 79–86.

Lemley, M. A. & Sampat, B. (2012), 'Examiner characteristics and patent office outcomes', *Review of Economics and Statistics* **94**(3), 817–827.

Marco, A. C., Myers, A. F., Graham, S., D'Agostino, P. & Kucab, J. (2015), 'The USPTO Patent Assignment Dataset: Descriptions, Lessons, and Insights', *Working paper* .

Marco, A. C., Sarnoff, J. D. & deGrazia, C. (2016), 'Patent Claims and Patent Scope', *Working paper* .

Merges, R. P. (1999), 'As many as six impossible patents before breakfast: Property rights for business concepts and patent system reform', *Berkeley Technology Law Journal* pp. 577–615.

Mori, T., Nishikimi, K. & Smith, T. E. (2005), 'A divergence statistic for industrial localization', *Review of Economics and Statistics* **87**(4), 635–651.

Putnam, J. (1996), 'The value of international patent protection', *Unpublished Ph. D. Dissertation, Yale University* .

Rysman, M. & Greenstein, S. (2005), 'Testing for agglomeration and dispersion', *Economics Letters* **86**(3), 405–411.

Sampat, B. & Williams, H. L. (2015), 'How do patents affect follow-on innovation? Evidence from the human genome', *NBER Working Paper 15851* .

Thoma, G., Torrisi, S., Gambardella, A., Guellec, D., Hall, B. H. & Harhoff, D. (2010), 'Harmonizing and combining large datasets–an application to firm-level patent and accounting data', *NBER Working Paper 21666* .

Tables and Figures

Figure 1: Distribution of P-values from D-index (top) and MTAD (bottom) for USPC subclass and Assignee

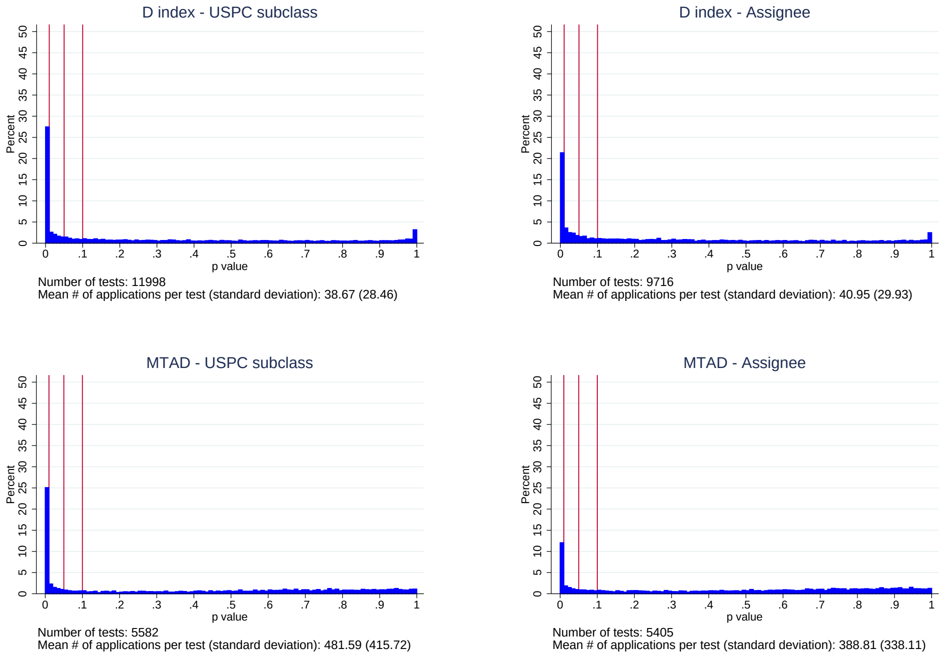

Distribution of p-values of D-index and MTAD analysis for USPC subclass and Assignee codes. Tests on subsamples with more than 20 applications for D-index and 50 applications for MTAD. Vertical red lines are standard thresholds for statistical significance (0.01, 0.05 and 0.10)

Figure 2: Quantile-Quantile Plots of D-index (top) and MTAD (bottom) by Art-Unit-Year (left) and Art-Unit-Year-USPC-Subclass (right) for Assignee

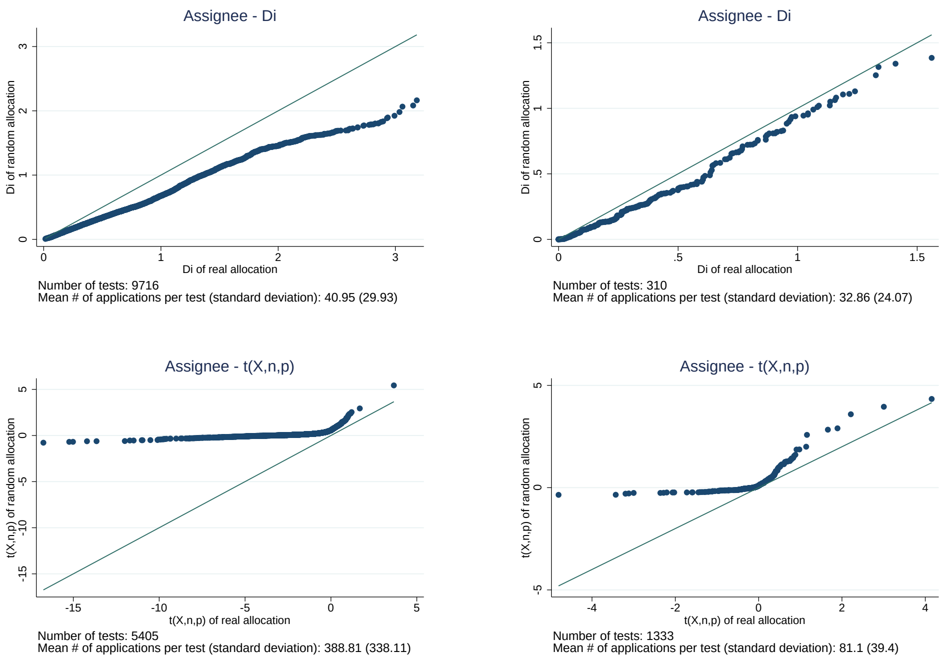

Each panel plots the quantiles of the D-index (top row) or MTAD statistic (bottom row) for the observed distribution (X-axis) against a simulated distribution under random assignment (Y-axis). Tests on subsamples with more than 20 applications for D-index and 50 applications for MTAD. If the observed distribution is random, the quantiles should be the same and the scatter points will fall along the 45-degree line. We observe large deviations from random assignment at the art-unit-year level, but much less evidence within art-unit-year-USPC-subclasses.

Table 1: Summary Statistics by Technology Center

	Applications	Art Units	Examiners	Classes	Sub-Classes	Class Changed†	Subclass Changed†
Biotechnology (1600)	221,586	57	1,013	268	11,530	32.3	75.3
Chemicals (1700)	367,371	75	1,377	411	34,509	31.3	77.7
Comp/Comm (2100)	208,102	79	1,733	303	7,314	23.4	70.7
Comp/Comm (2400)	157,852	77	1,289	159	4,492	18.5	59.7
Comp/Comm (2600)	338,088	82	2,046	310	14,393	15.4	65.7
Electrical (2800)	637,929	80	2,161	382	31,815	17.5	65.2
Miscellaneous (3600)	360,691	75	1,617	410	38,760	16.8	68.8
Mechanical (3700)	425,413	65	1,949	411	40,240	18.6	68.9
Full sample	2,717,032	590	12,338	452	119,448	20.4	68.5

†Percent of applications with change in (sub)class conditional on grant before July 21, 2015. Classification data for published applications and granted patents from PatentsView. Full sample counts remove duplicates across technology centers. Abbreviated technology center names from Graham et al. (2015). Full names of the technology centers currently responsible for examination of utility patent applications are:

- 1600 - Biotechnology and Organic Chemistry
- 1700 - Chemical and Materials Engineering
- 2100 - Computer Architecture, Software, and Information Security
- 2400 - Computer Networks, Multiplex communication, Video Distribution, and Security
- 2600 - Communications
- 2800 - Semiconductors, Electrical and Optical Systems and Components
- 3600 - Transportation,Construction, Electronic Commerce, Agriculture, National Security and License & Review
- 3700 - Mechanical Engineering, Manufacturing, Products

Table 2: D-index and MTAD Tests within Art-Unit-Application-Year (share rejecting random allocation at 1% significance level, by technology center)

Panel A: USPC subclass

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Biotechnology and Organic Chemistry	32.9	906	19.6	0.2	551
Chemical and Materials Engineering	58.6	814	55.8	0.0	721
Computer Architecture, Software, and Security	2.2	1,170	0.7	0.0	723
Computer Networking and Video Distribution	6.5	753	0.8	0.0	628
Communications	17.7	2,268	16.7	0.0	694
Semiconductors, Electrical and Optical Systems	37.4	3,389	39.5	0.0	843
Miscellaneous†	15.4	1,162	21.7	0.1	742
Mechanical Engineering, Manufacturing, Products	38.9	1,536	39.4	0.0	680
All tests	27.5	11,998	25.0	0.0	5,582

Panel B: Assignee

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Biotechnology and Organic Chemistry	50.2	225	9.1	0.0	527
Chemical and Materials Engineering	46.1	866	30.9	0.0	699
Computer Architecture, Software, and Security	4.2	970	0.0	0.0	709
Computer Networking and Video Distribution	5.3	509	0.2	0.0	616
Communications	11.0	1,879	6.3	0.0	668
Semiconductors, Electrical and Optical Systems	19.6	3,360	15.0	0.1	824
Miscellaneous†	29.3	818	13.4	0.0	703
Mechanical Engineering, Manufacturing, Products	36.0	1,089	19.1	0.0	659
All tests	21.4	9,716	12.0	0.0	5,405

Panel C: DOCDB Family Size

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Biotechnology and Organic Chemistry	0.8	772	3.1	0.0	549
Chemical and Materials Engineering	3.7	1,018	5.4	0.0	716
Computer Architecture, Software, and Security	0.5	860	2.1	0.0	723
Computer Networking and Video Distribution	0.3	742	1.1	0.0	627
Communications	2.5	1,011	4.5	0.0	690
Semiconductors, Electrical and Optical Systems	3.8	1,427	6.1	0.0	841
Miscellaneous†	2.4	1,149	4.7	0.0	738
Mechanical Engineering, Manufacturing, Products	3.9	1,089	7.2	0.0	678
All tests	2.5	8,068	4.4	0.0	5,562

Panel D: Words in 1st Claim

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Chemical and Materials Engineering	3.5	1,129	5.1	0.0	721
Computer Architecture, Software, and Security	0.0	895	0.1	0.0	723
Computer Networking and Video Distribution	0.0	755	0.0	0.0	627
Communications	0.2	1,052	0.1	0.0	693
Semiconductors, Electrical and Optical Systems	2.1	1,524	5.1	0.0	843
Miscellaneous†	0.7	1,194	1.8	0.0	741
Mechanical Engineering, Manufacturing, Products	0.8	1,160	0.7	0.0	679
All tests	1.2	7,709	2.0	0.0	5,027

For D-index, columns labelled "Rej." report the share of tests that reject the null hypothesis of equality between the observed and the reference distribution at 1% level. For MTAD, columns labelled "Agg." ("Disp.") report the share of tests that reject the null hypothesis of random allocation at 1% level in favor of agglomeration (dispersion). All tests are conducted within art-unit-year cells with more than 20 applications for the D-index and more than 50 applications for MTAD. † Miscellaneous = "Transportation, Construction, Electronic Commerce, Agriculture, National Security and License & Review."

Table 3: D-index and MTAD Tests within Art-Unit-Application-Year-Subclass (1% significance level, by technology center)

Panel A: USPC secondary subclass

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Biotechnology and Organic Chemistry	21.1	720	1.5	13.8	1,613
Chemical and Materials Engineering	3.4	147	0.4	13.4	516
Computer Architecture, Software, and Security	0.0	99	0.0	0.0	143
Computer Networking and Video Distribution	0.0	33	0.0	0.0	35
Communications	0.0	120	0.0	0.0	212
Semiconductors, Electrical and Optical Systems	0.0	104	0.0	0.9	531
Miscellaneous†	3.8	26	0.0	0.7	150
Mechanical Engineering, Manufacturing, Products	0.0	55	0.0	1.2	168
All tests	12.1	1,304	0.8	8.9	3,368

Panel B: Assignee

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Biotechnology and Organic Chemistry	16.2	37	8.7	0.0	104
Chemical and Materials Engineering	9.1	33	2.5	1.2	80
Computer Architecture, Software, and Security	0.0	13	0.0	0.0	120
Computer Networking and Video Distribution	0.0	1	0.0	0.0	46
Communications	0.0	84	0.0	0.0	363
Semiconductors, Electrical and Optical Systems	0.0	82	0.3	0.0	392
Miscellaneous†	0.0	10	0.0	0.9	115
Mechanical Engineering, Manufacturing, Products	6.0	50	0.0	0.0	113
All tests	3.9	310	0.9	0.2	1,333

Panel C: DOCDB Family Size

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Biotechnology and Organic Chemistry	0.0	844	0.0	0.0	151
Chemical and Materials Engineering	0.0	768	3.2	0.0	93
Computer Architecture, Software, and Security	0.0	1,082	0.0	0.0	174
Computer Networking and Video Distribution	0.0	698	0.0	0.0	62
Communications	0.0	2,147	0.2	0.0	429
Semiconductors, Electrical and Optical Systems	0.0	3,184	0.4	0.0	458
Miscellaneous†	0.0	1,097	0.0	0.0	169
Mechanical Engineering, Manufacturing, Products	0.0	1,419	0.7	0.0	144
All tests	0.0	11,239	0.4	0.0	1,680

Panel D: Words in 1st Claim

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Chemical and Materials Engineering	0.0	747	0.0	0.0	99
Computer Architecture, Software, and Security	0.0	1,058	0.0	0.0	191
Computer Networking and Video Distribution	0.0	684	0.0	0.0	64
Communications	0.0	2,109	0.0	0.0	456
Semiconductors, Electrical and Optical Systems	0.0	3,109	0.4	0.0	520
Miscellaneous†	0.0	1,062	0.0	0.0	190
Mechanical Engineering, Manufacturing, Products	0.0	1,409	0.0	0.0	185
All tests	0.0	10,178	0.1	0.0	1,705

For D-index, columns labelled "Rej." report the share of tests that reject the null hypothesis of equality between the observed and the reference distribution at 1% level. For MTAD, columns labelled "Agg." ("Disp.") report the share of tests that reject the null hypothesis of random allocation at 1% level in favor of agglomeration (dispersion). All tests are conducted within art-unit-year-subclass cells with more than 20 applications for the D-index and more than 50 applications for MTAD. † Miscellaneous = "Transportation, Construction, Electronic Commerce, Agriculture, National Security and License & Review."

Table 4: IV "First-Stage" With and Without Subclass Effects

Outcome†	1[Granted]			Words-in-1^{st}-claim		
Potential Instrument	Leave-one-out grant rate			Leave-one-out scope change		
	(1)	(2)	(1)/(2)	(3)	(4)	(3)/(4)
Full Sample	0.163 (0.001)	0.147 (0.001)	0.90	0.224 (0.003)	0.211 (0.004)	0.94
Art-unit-year FEs	✓			✓		
Art-unit-year-subclass FEs		✓			✓	
Biotechnology (1600)	0.149 (0.002)	0.132 (0.003)	0.89			
Chemicals (1700)	0.174 (0.002)	0.144 (0.003)	0.83	0.146 (0.012)	0.060 (0.016)	0.41
Comp/Comm (2100)	0.143 (0.002)	0.140 (0.002)	0.98	0.207 (0.009)	0.208 (0.009)	1.00
Comp/Comm (2400)	0.121 (0.003)	0.115 (0.004)	0.95	0.226 (0.009)	0.217 (0.011)	0.96
Comp/Comm (2600)	0.163 (0.002)	0.156 (0.002)	0.96	0.236 (0.009)	0.226 (0.010)	0.96
Electrical (2800)	0.170 (0.001)	0.160 (0.002)	0.94	0.247 (0.005)	0.240 (0.005)	0.97
Miscellaneous (3600)	0.161 (0.002)	0.136 (0.002)	0.84	0.218 (0.009)	0.193 (0.009)	0.89
Mechanical (3700)	0.171 (0.002)	0.150 (0.002)	0.88	0.220 (0.008)	0.194 (0.008)	0.88

†Outcome is the endogenous variable in an IV regression.
Each "first-stage" estimate in this table comes from a separate OLS regression of Outcome on Potential Instrument for applications assigned to a given technology center. Robust standard errors, clustered by art-unit-filing-year, in parentheses. All estimates are statistically significant at the 1% level. See text for a discussion of the estimation sample, and variable definitions. We exclude biotechnology patents (Technology Center 1600) from the second set of estimates because Kuhn et al. (2016) suggest that counting words in the first claim does not yield a meaningful measure of claim-scope for those applications.

Table 5: Examiner Specialization and Examination Outcomes.

Model	Within Examiner			Between Examiner		
Outcome	Granted	Words	Days	Granted	Words	Days
	(1)	(2)	(3)	(4)	(5)	(6)
$Share_{ijt}$	0.00***	0.00	0.01***	-0.04***	0.08***	0.10***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)
Art-unit-year-examiner FEs	✓	✓	✓			
Observations	1,750,211	1,069,834	1,749,990	48,973	44,036	48,973
Art-unit-year-examiners	48,973	44,036	48,973			

All models estimated with OLS. Unit of observation is a patent application for the within regressions and an art-unit-year-examiner for the between regressions. Variables $Share_{iijt}$, Words and Days are standardized. The mean of the outcome of the regression in column 1 is 0.65. Between regressions estimated on the group means. Standard errors clustered by art unit-filing year in parentheses. *** p<0.01, ** p<0.05, * p<0.10

A Appendix: Online Publication Only

Figure A1: Distribution of P-values from D-index (top) and MTAD (bottom) for USPC subclass and Assignee (lower thresholds)

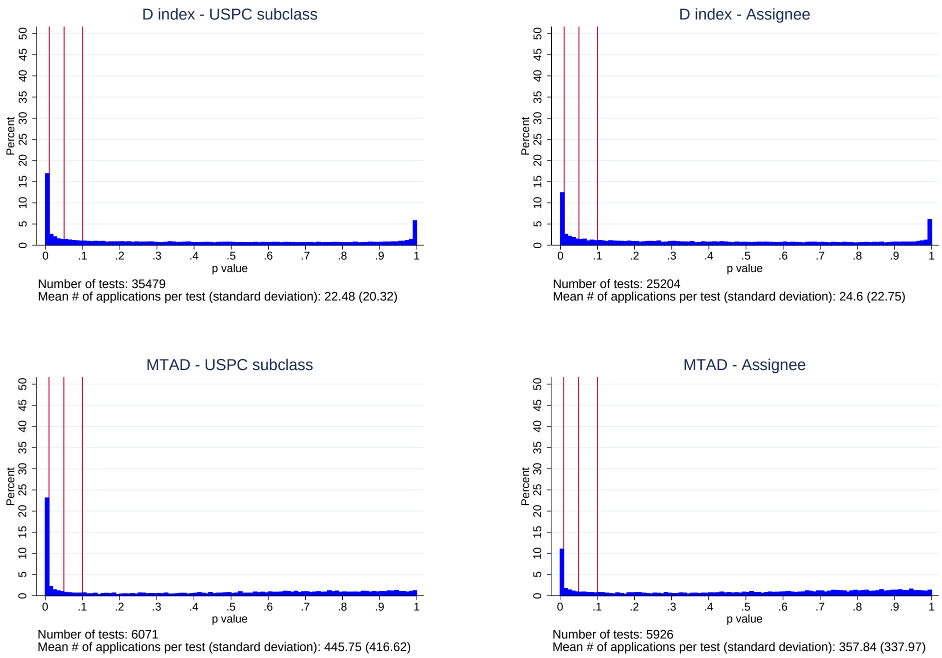

Distribution of p-values of D-index and MTAD analysis for USPC subclass and Assignee codes. Tests on subsamples with more than 10 applications for D-index and 25 applications for MTAD. Vertical red lines are standard thresholds for statistical significance (0.01, 0.05 and 0.10)

Figure A2: Quantile-Quantile Plots of D-index (top) and MTAD (bottom) by Art-Unit-Year (left) and Art-Unit-Year-USPC-Subclass (right) for Assignee (lower thresholds)

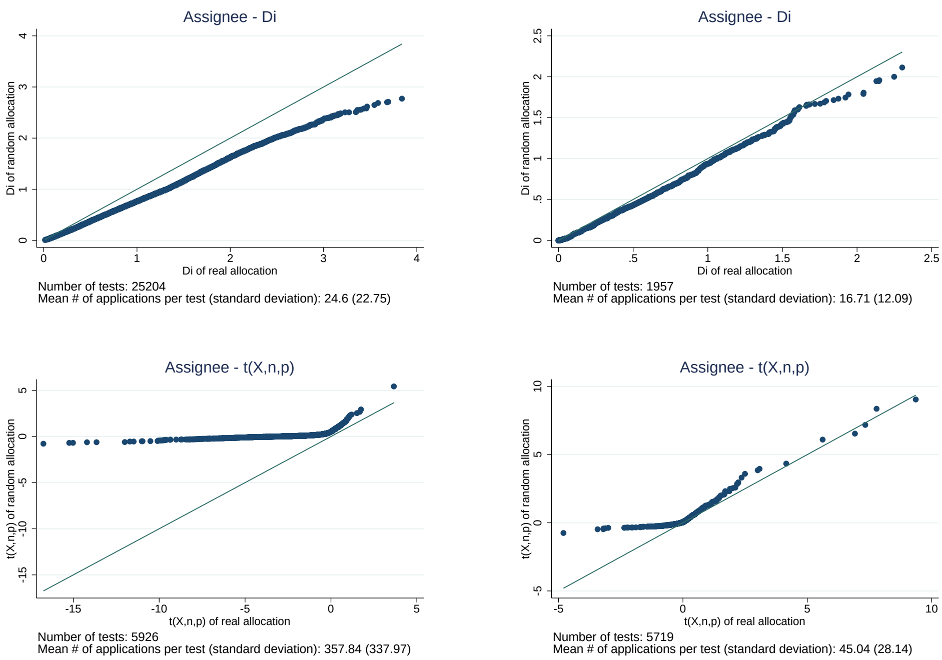

Each panel plots the quantiles of the D-index (top row) or MTAD statistic (bottom row) for the observed distribution (X-axis) against a simulated distribution under random assignment (Y-axis). Tests on subsamples with more than 10 applications for D-index and 25 applications for MTAD. If the observed distribution is random, the quantiles should be the same and the scatter points will fall along the 45-degree line. We observe large deviations from random assignment at the art-unit-year level, but much less evidence within art-unit-year-USPC-subclasses.

Table A1: Summary statistics for sample of applications

Panel A: categorical variables

Variable	# of categories	Applications per category Mean	Std dev	Min	5^{th} percentile	1^{st} quartile	Median	3^{rd} quartile	95^{th} percentile	Max
Examiners	12,338	220.22	227.54	1	2	48	156	313	714	1,655
Art units	590	4,605.14	3,942.42	3	446	2,007	3,228.50	6,157	13,459	21,905
USPC subclasses	119,448	22.75	106.65	1	1	2	5	14	85	13,836
Assignees	164,195	12.99	301.16	1	1	1	1	3	19	59,998

Panel B: quantitative variables

Variable	N	Mean	Std dev	Min	5^{th} percentile	1^{st} quartile	Median	3^{rd} quartile	95^{th} percentile	Max
DOCDB family size	2,716,195	2.88	5.66	1	1	1	2	3	8	378
Words in 1^{st} claim	2,712,367	124.95	128.00	1	35	70	103	151	269	46,194

The number of applications characterized by a big DOCDB family and a low number of words in the first independent claim are respectively 106,408 and 116,665.

Table A2: Summary statistics for examiners' specialization and examination outcomes.

Variable	N	Mean	Std dev	Min	Median	Max
$Share_{ijt}$	1,750,211	0.04	0.09	0.00	0.00	1.00
Granted	1,750,211	0.65	0.48	0.00	1.00	1.00
Days	1,749,990	918.09	510.19	0.00	826.00	17,835.00
Words	1,069,834	49.16	87.56	-10,351.00	30.00	9,248.00

Table A3: D-index and MTAD Tests within Art-Unit-Application-Year (share rejecting random allocation at 5% significance level, by technology center)

Panel A: USPC subclass

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Biotechnology and Organic Chemistry	43.8	906	28.1	0.2	551
Chemical and Materials Engineering	66.3	814	64.6	0.0	721
Computer Architecture, Software, and Security	7.1	1,170	1.4	0.1	723
Computer Networking and Video Distribution	14.9	753	2.1	0.2	628
Communications	22.5	2,268	20.6	0.0	694
Semiconductors, Electrical and Optical Systems	46.5	3,389	47.1	0.1	843
Miscellaneous†	23.2	1,162	28.8	0.3	742
Mechanical Engineering, Manufacturing, Products	49.0	1,536	49.3	0.0	680
All tests	35.3	11,998	31.0	0.1	5,582

Panel B: Assignee

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Biotechnology and Organic Chemistry	65.8	225	13.3	0.0	527
Chemical and Materials Engineering	56.0	866	40.3	0.0	699
Computer Architecture, Software, and Security	11.3	970	0.3	0.0	709
Computer Networking and Video Distribution	15.5	509	0.3	0.0	616
Communications	17.0	1,879	10.0	0.1	668
Semiconductors, Electrical and Optical Systems	31.4	3,360	22.6	0.1	824
Miscellaneous†	40.2	818	21.1	0.0	703
Mechanical Engineering, Manufacturing, Products	51.1	1,089	29.1	0.0	659
All tests	31.7	9,716	17.6	0.0	5,405

Panel C: DOCDB Family Size

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Biotechnology and Organic Chemistry	2.8	772	6.4	0.0	549
Chemical and Materials Engineering	6.9	1,018	9.4	0.0	716
Computer Architecture, Software, and Security	0.8	860	4.4	0.0	723
Computer Networking and Video Distribution	0.7	742	2.6	0.0	627
Communications	4.0	1,011	6.5	0.0	690
Semiconductors, Electrical and Optical Systems	7.6	1,427	10.5	0.0	841
Miscellaneous†	5.3	1,149	8.5	0.0	738
Mechanical Engineering, Manufacturing, Products	8.2	1,089	11.9	0.0	678
All tests	5.0	8,068	7.7	0.0	5,562

Panel D: Words in 1st Claim

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Chemical and Materials Engineering	7.2	1,129	9.7	0.0	721
Computer Architecture, Software, and Security	0.1	895	1.5	0.0	723
Computer Networking and Video Distribution	0.0	755	0.0	0.0	627
Communications	0.5	1,052	0.9	0.0	693
Semiconductors, Electrical and Optical Systems	5.0	1,524	8.8	0.1	843
Miscellaneous†	3.4	1,194	4.2	0.1	741
Mechanical Engineering, Manufacturing, Products	3.2	1,160	4.7	0.1	679
All tests	3.1	7,709	4.5	0.1	5,027

For D-index, columns labelled "Rej." report the share of tests that reject the null hypothesis of equality between the observed and the reference distribution at 5% level. For MTAD, columns labelled "Agg." ("Disp.") report the share of tests that reject the null hypothesis of random allocation at 5% level in favor of agglomeration (dispersion). All tests are conducted within art-unit-year cells with more than 20 applications for the D-index and more than 50 applications for MTAD. † Miscellaneous = "Transportation, Construction, Electronic Commerce, Agriculture, National Security and License & Review."

Table A4: D-index and MTAD Tests within Art-Unit-Application-Year-Class-Subclass (share rejecting random allocation at 5% significance level, by technology center)

Panel A: USPC secondary subclass

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Biotechnology and Organic Chemistry	24.7	720	2.0	27.3	1,613
Chemical and Materials Engineering	7.5	147	0.8	21.7	516
Computer Architecture, Software, and Security	0.0	99	0.0	0.7	143
Computer Networking and Video Distribution	0.0	33	0.0	0.0	35
Communications	0.8	120	0.0	0.5	212
Semiconductors, Electrical and Optical Systems	0.0	104	0.4	3.0	531
Miscellaneous†	3.8	26	0.0	1.3	150
Mechanical Engineering, Manufacturing, Products	1.8	55	0.0	1.8	168
All tests	14.7	1,304	1.1	17.1	3,368

Panel B: Assignee

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Biotechnology and Organic Chemistry	27.0	37	11.5	0.0	104
Chemical and Materials Engineering	24.2	33	3.8	2.5	80
Computer Architecture, Software, and Security	0.0	13	0.0	0.0	120
Computer Networking and Video Distribution	0.0	1	0.0	0.0	46
Communications	0.0	84	0.3	0.0	363
Semiconductors, Electrical and Optical Systems	0.0	82	0.3	0.0	392
Miscellaneous†	0.0	10	0.9	0.9	115
Mechanical Engineering, Manufacturing, Products	10.0	50	1.8	0.0	113
All tests	7.4	310	1.5	0.2	1,333

Panel C: DOCDB Family Size

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Biotechnology and Organic Chemistry	0.0	844	0.0	0.0	151
Chemical and Materials Engineering	0.0	768	5.4	0.0	93
Computer Architecture, Software, and Security	0.0	1,082	0.0	0.0	174
Computer Networking and Video Distribution	0.0	698	0.0	0.0	62
Communications	0.0	2,147	0.5	0.0	429
Semiconductors, Electrical and Optical Systems	0.0	3,184	0.4	0.0	458
Miscellaneous†	0.0	1,097	3.0	0.0	169
Mechanical Engineering, Manufacturing, Products	0.1	1,419	0.7	0.0	144
All tests	0.0	11,239	0.9	0.0	1,680

Panel D: Words in 1st Claim

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Chemical and Materials Engineering	0.0	747	2.0	0.0	99
Computer Architecture, Software, and Security	0.0	1,058	0.0	0.0	191
Computer Networking and Video Distribution	0.0	684	0.0	0.0	64
Communications	0.0	2,109	0.0	0.0	456
Semiconductors, Electrical and Optical Systems	0.0	3,109	0.6	0.0	520
Miscellaneous†	0.0	1,062	0.0	0.0	190
Mechanical Engineering, Manufacturing, Products	0.0	1,409	0.5	0.0	185
All tests	0.0	10,178	0.4	0.0	1,705

For D-index, columns labelled "Rej." report the share of tests that reject the null hypothesis of equality between the observed and the reference distribution at 5% level. For MTAD, columns labelled "Agg." ("Disp.") report the share of tests that reject the null hypothesis of random allocation at 5% level in favor of agglomeration (dispersion). All tests are conducted within art-unit-year-subclass cells with more than 20 applications for the D-index and more than 50 applications for MTAD. † Miscellaneous = "Transportation, Construction, Electronic Commerce, Agriculture, National Security and License & Review."

Table A5: D-index and MTAD Tests within Art-Unit-Class-Subclass (share rejecting random allocation at 1% significance level, by technology center)

Panel A: USPC secondary subclass

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Biotechnology and Organic Chemistry	15.4	1,425	1.5	8.3	2,132
Chemical and Materials Engineering	6.0	431	0.5	10.6	1,474
Computer Architecture, Software, and Security	0.4	253	0.0	0.0	470
Computer Networking and Video Distribution	1.3	154	0.0	0.0	280
Communications	0.6	362	0.0	0.0	639
Semiconductors, Electrical and Optical Systems	1.6	741	0.1	0.2	1,847
Miscellaneous†	2.7	222	0.0	0.6	464
Mechanical Engineering, Manufacturing, Products	2.5	314	0.6	0.8	893
All tests	7.1	3,902	0.5	4.2	8,199

Panel B: Assignee

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Biotechnology and Organic Chemistry	25.9	81	3.1	0.3	355
Chemical and Materials Engineering	13.4	134	1.8	3.3	389
Computer Architecture, Software, and Security	0.4	274	0.0	0.0	569
Computer Networking and Video Distribution	1.2	85	0.0	0.0	475
Communications	0.5	430	0.0	0.0	948
Semiconductors, Electrical and Optical Systems	1.6	880	0.1	0.1	1,819
Miscellaneous†	2.3	219	0.8	0.2	525
Mechanical Engineering, Manufacturing, Products	2.8	396	0.6	0.3	797
All tests	2.9	2,499	0.5	0.3	5,877

Panel C: DOCDB Family Size

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Biotechnology and Organic Chemistry	0.1	1,590	1.0	0.0	512
Chemical and Materials Engineering	0.0	2,266	0.4	0.0	470
Computer Architecture, Software, and Security	0.0	1,879	0.3	0.0	616
Computer Networking and Video Distribution	0.0	1,605	0.0	0.0	526
Communications	0.0	2,802	0.0	0.0	981
Semiconductors, Electrical and Optical Systems	0.0	5,806	0.2	0.0	1,952
Miscellaneous†	0.1	2,789	0.4	0.0	711
Mechanical Engineering, Manufacturing, Products	0.0	3,759	0.8	0.0	1,126
All tests	0.0	22,496	0.3	0.0	6,894

Panel D: Words in 1st Claim

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Chemical and Materials Engineering	0.0	2,226	0.2	0.0	496
Computer Architecture, Software, and Security	0.0	1,880	0.0	0.0	657
Computer Networking and Video Distribution	0.0	1,588	0.0	0.0	553
Communications	0.0	2,794	0.1	0.0	1,068
Semiconductors, Electrical and Optical Systems	0.0	5,745	0.1	0.0	2,076
Miscellaneous†	0.0	2,791	0.0	0.0	748
Mechanical Engineering, Manufacturing, Products	0.0	3,723	0.0	0.0	1,187
All tests	0.0	20,747	0.1	0.0	6,785

For D-index, columns labelled "Rej." report the share of tests that reject the null hypothesis of equality between the observed and the reference distribution at 1% level. For MTAD, columns labelled "Agg." ("Disp.") report the share of tests that reject the null hypothesis of random allocation at 1% level in favor of agglomeration (dispersion). All tests are conducted within art-unit-subclass cells with more than 20 applications for the D-index and more than 50 applications for MTAD. † Miscellaneous = "Transportation, Construction, Electronic Commerce, Agriculture, National Security and License & Review."

Table A6: D-index and MTAD Tests within Art-Unit-Class-Subclass (share rejecting random allocation at 5% significance level, by technology center)

Panel A: USPC secondary subclass

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Biotechnology and Organic Chemistry	23.3	1,425	2.2	16.2	2,132
Chemical and Materials Engineering	10.4	431	0.5	19.3	1,474
Computer Architecture, Software, and Security	1.2	253	0.0	0.0	470
Computer Networking and Video Distribution	4.5	154	0.0	0.7	280
Communications	1.7	362	0.0	0.2	639
Semiconductors, Electrical and Optical Systems	3.4	741	0.1	0.7	1,847
Miscellaneous†	6.3	222	0.2	1.9	464
Mechanical Engineering, Manufacturing, Products	6.4	314	0.8	1.6	893
All tests	11.6	3,902	0.8	8.1	8,199

Panel B: Assignee

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Biotechnology and Organic Chemistry	38.3	81	4.5	0.8	355
Chemical and Materials Engineering	20.1	134	2.3	6.9	389
Computer Architecture, Software, and Security	1.5	274	0.0	0.0	569
Computer Networking and Video Distribution	7.1	85	0.0	0.0	475
Communications	1.2	430	0.1	0.0	948
Semiconductors, Electrical and Optical Systems	3.0	880	0.2	0.5	1,819
Miscellaneous†	5.5	219	1.0	0.6	525
Mechanical Engineering, Manufacturing, Products	8.6	396	1.5	0.4	797
All tests	5.8	2,499	0.8	0.8	5,877

Panel C: DOCDB Family Size

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Biotechnology and Organic Chemistry	0.1	1,590	1.4	0.0	512
Chemical and Materials Engineering	0.0	2,266	1.5	0.0	470
Computer Architecture, Software, and Security	0.1	1,879	0.3	0.0	616
Computer Networking and Video Distribution	0.0	1,605	0.0	0.0	526
Communications	0.0	2,802	0.1	0.0	981
Semiconductors, Electrical and Optical Systems	0.0	5,806	0.5	0.0	1,952
Miscellaneous†	0.2	2,789	0.7	0.0	711
Mechanical Engineering, Manufacturing, Products	0.1	3,759	1.0	0.0	1,126
All tests	0.1	22,496	0.6	0.0	6,894

Panel D: Words in 1st Claim

	D-index		MTAD		
Technology Center	Rej.	N	Agg.	Disp.	N
Chemical and Materials Engineering	0.0	2,226	1.0	0.0	496
Computer Architecture, Software, and Security	0.0	1,880	0.2	0.0	657
Computer Networking and Video Distribution	0.0	1,588	0.0	0.0	553
Communications	0.0	2,794	0.2	0.0	1,068
Semiconductors, Electrical and Optical Systems	0.1	5,745	0.5	0.0	2,076
Miscellaneous†	0.0	2,791	0.3	0.0	748
Mechanical Engineering, Manufacturing, Products	0.0	3,723	0.5	0.0	1,187
All tests	0.0	20,747	0.4	0.0	6,785

For D-index, columns labelled "Rej." report the share of tests that reject the null hypothesis of equality between the observed and the reference distribution at 5% level. For MTAD, columns labelled "Agg." ("Disp.") report the share of tests that reject the null hypothesis of random allocation at 5% level in favor of agglomeration (dispersion). All tests are conducted within art-unit-subclass cells with more than 20 applications for the D-index and more than 50 applications for MTAD. † Miscellaneous = "Transportation, Construction, Electronic Commerce, Agriculture, National Security and License & Review."