

# Patent Examiner Specialization\*

Cesare Righi<sup>1</sup> and Timothy Simcoe<sup>2</sup>

<sup>1</sup>Boston University

<sup>2</sup>Boston University and NBER

March 1, 2018

## Abstract

We study the matching of patent applications to examiners at the U.S. Patent and Trademark Office. The distribution of technology classes is more concentrated than would occur under random matching. Using the application text, we show that specialization persists even after conditioning on observed technology classes and sub-classes. Examiner specialization is less pronounced in computers and software than in other technological fields. More specialized examiners have a lower grant rate and produce a larger narrowing of claim-scope during the examination process. These findings undermine the idea that random matching justifies instrumental variables based on examiner behaviors or characteristics.

---

\*We thank Sam Arts, Jim Bessen, Bruno Cassiman, Christian Catalini, Iain Cockburn, Dennis Crouch, Stuart Graham, Deepak Hegde, Mark Lemley, Alan Marco, Michael Meurer, Arti Rai, Bhaven Sampat, Mark Schankerman, 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, the NBER productivity lunch and the USPTO for useful comments and suggestions. We thank Evgeny Klochikhin for access to the PatentsView patent applications database. Contact: crighi@bu.edu; 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 proposed by Sampat & Williams (2017), used examiner characteristics as instrumental variables for examination outcomes.<sup>1</sup>

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, Chemistry, Mechanical Engineering and Semiconductor applications, and less so in the computer-related technology centers. We use a measure of textual similarity between pairs of patent applications to show that technological specialization persists even within U.S. Patent Classification System (USPC) subclasses. We also find that applications filed by the same organization are more likely to be assigned to the same examiner, even after controlling for subclasses and observable similarities between technologies. We find no evidence, however, 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 variable strategies based on examiner behaviors and characteristics. Random assignment would suffice to make leave-one-out grant rates (or any other examiner characteristic) uncorrelated with application characteristics. But when examiners specialize,

---

<sup>1</sup>Papers adopting variants on this identification strategy include Farre-Mensa, Hegde & Ljungqvist (2017), Feng & Jaravel (2017), Gaulé (2015), Kuhn (2016), Kuhn & Thompson (2017), and Sampat & Williams (2017).

examiner-based instruments can be correlated with technology, suggesting an alternative path through which the instrument could influence outcomes. While this does not necessarily invalidate examiner-based instruments, it implies that the identification strategy rests on a stronger assumption than is typically acknowledged: potential outcomes must be uncorrelated with (unobserved) technological heterogeneity.

This is the first paper to systematically test the random matching hypothesis across all of the technology areas examined by the USPTO. We use two different methods to provide evidence of specialization. First, we examine a pair of test statistics from the literature on industry agglomeration (Mori, Nishikimi & Smith 2005) that ask whether application characteristics (e.g. technology subclass) are more or less dispersed across examiners than would occur under random assignment.<sup>2</sup> These statistics are computed at the art-unit-filing-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). Our second method exploits a new measure proposed by Arts, Cassiman & Gomez (2018), and uses OLS regression to show that the probability of two applications being assigned to the same examiner increases with the textual similarity of their titles and abstracts.

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 2017, Kuhn & Thompson 2017), and experienced examiners are more lenient on average, partly because of time constraints (Lemley & Sampat 2012, Frakes & Wasserman 2017). Random matching provides another path to fairness, but forgoes the potential efficiency benefits of specialization. Our analysis shows that the amount of specialization varies across art units, leading some applicants to get tougher

---

<sup>2</sup>These methods focus specifically on the null hypothesis of random assignment, unlike instrumental variable falsification tests that ask the slightly different question of whether examiner and application characteristics are correlated.

examiners on average. But we find very little evidence that particularly important applications (with large families) or broad applications (with short first independent claims) are assigned to specific examiners.

We consider two possible mechanisms that would generate more observed specialization in Chemistry, Biotechnology, Mechanical Engineering and Semiconductors 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 other areas. Another possibility is that the USPC technology classification system works better outside computers and software, so we simply fail to observe much of the specialization that takes place within computer-related art units. The data provide support for both of these explanations. Application-pairs sharing the same examiner have less textual overlap in Computers and Communications, suggesting that examiners in the computer-related tech centers are more generalist. At the same time, application-pairs in Computers and Communications that share the same primary USPC subclass are less textually similar than in other fields, suggesting differences in the quality of the classification system.

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.

## 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.<sup>3</sup> Each application has at

---

<sup>3</sup>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

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 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.<sup>4</sup>

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.<sup>5</sup> 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. 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

---

classification for the entire period of our study.

<sup>4</sup>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>.

<sup>5</sup>The Manual of Patent Examining Procedure sec. 903.08 describes the rules governing assignment and transfer of applications between art units.

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 2017), how quickly a decision is reached (Farre-Mensa et al. 2017), and the scope and strength of an issued patent (Kuhn & Thompson 2017). This variation in standards led Cockburn et al. (2002) to conclude that “there may be as many patent offices as patent examiners.”

### 3 Methods and Data

We use two methods to measure specialization. The first approach compares the actual distribution of application characteristics to the predicted distribution under random assignment. The second approach uses regression at the application-pair level to ask whether textual similarity predicts having the same examiner.

#### 3.1 Agglomeration Test Statistics

Our first approach to measuring specialization uses two statistical tests originally developed to study industry agglomeration: the Divergence Index (D-index) and the Multinomial Test of Agglomeration and Dispersion (MTAD). In our setting, patent examiners are analogous to cities, and technology subclasses (or other application characteristics) are analogous to industries. We briefly describe the two tests here, and provide details in Appendix A.

The D-index was proposed by Mori et al. (2005), building on Kullback & Leibler (1951), and is based on the concept of relative entropy.<sup>6</sup> For a given technology subclass (or other discrete characteristic) indexed by  $i$ , the D-index is given by

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

---

<sup>6</sup>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 setting.

where  $r$  indexes individual examiners in set  $\mathbf{R}$ ;  $p_{0r}$  is the share of all applications assigned to examiner  $r$ ; and  $\hat{p}_{ir}$  is the share of all applications in subclass  $i$  assigned to examiner  $r$ . If there are  $N_i$  applications in subclass  $i$ , then  $2N_i D(\hat{p}_i | p_0)$  has a chi-square distribution with  $R - 1$  degrees of freedom under the null of random matching. In our analysis, the number of D-index 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.  $N_i > 20$ ).

MTAD was developed by Rysman & Greenstein (2005), and computes multinomial likelihood functions for an allocation of agents to a set of discrete locations. In our setting, the agents are patent applications and locations correspond to examiners. In particular, suppose we have  $R$  examiners, each receiving  $n_r$  applications, and  $C$  subclasses, each having unconditional probability  $p_c$ . The observed number of applications of type  $c$  assigned to examiner  $r$  is  $x_r^c$ . Under random matching of applications to examiners, the likelihood of observing outcome  $\mathbf{x}_r$  for examiner  $r$  is the multinomial pdf

$$\mathcal{L}(\mathbf{x}_r, n_r, \mathbf{p}) = \binom{n_r}{x_r^1, \dots, x_r^C} p_1^{x_r^1} \dots p_C^{x_r^C}. \quad (2)$$

The intuition behind MTAD is to replace  $p_c$  with the observed share of subclass  $c$ , and then compare the sample average log-likelihood to the expected log-likelihood under simulated random assignment. If the likelihood of the observed data is lower (higher) than the likelihood under random choice, MTAD indicates that the technology classes are agglomerated (dispersed).<sup>7</sup> 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.

---

<sup>7</sup>The difference between the observed and simulated likelihood is distributed asymptotically normal and we use simulation to generate its confidence intervals. 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>

## 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>), PAT-STAT, 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 1<sup>st</sup> 2013, whose examiner is affiliated with one of the eight TCs 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. The USPTO Patent Examination Research Dataset also provides the art unit of the examiner of record at the time of the last office action recorded for a given application. Under the AIPA, regular utility patent applications are generally published eighteen months after filing.<sup>8</sup>

The data have several limitations. First, applications will not appear in our data if they are abandoned or granted 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.<sup>9</sup> A second limitation is that we do not observe whether applications are transferred from one examiner to another.

Our primary analysis sample contains 2,717,032 applications examined

---

<sup>8</sup>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.

<sup>9</sup>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. Moreover, only 7% of the applications that meet all other criteria for inclusion in our main analysis sample are not published.



by 12,338 unique examiners affiliated with 590 art units. Table 1 shows the number of art units, examiners, classes, subclasses and applications in each USPTO Technology Center (TC). 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.

Table 1 also reports sample sizes for the application-pair-level analyses. The full sample contains roughly 1.1 billion application-pairs filed in the same calendar year and processed by the same art-unit. Our regressions utilize a sub-sample of 11.7 million application-pairs that have the same filing year, art unit, primary USPC subclass and have an assignee in our data.

### 3.3 Variables

We focus on several application characteristics that may influence the assignment of applications to individual patent examiners within an art-unit-filing-year.<sup>10</sup> 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. While Lemley & Sampat (2012) provide qualitative evidence that this may happen at least in a subset of art units, our analysis tests this statistically on the entire population of art units (conditional on sample size) and identifies in what TCs technological specialization is more relevant.

We use technology classification data from published applications, rather than granted patents, to avoid measuring any agglomeration created by the examination process. In particular, because the USPC classification of an

---

<sup>10</sup>We typically compute our test statistics within a art-unit-filing-year cell to account for possible changes in assignment practices over time and turnover in the pool of examiners.

application is based on its claims, which are usually amended during examination, the subclass of many applications changes over time.<sup>11,12</sup> This could lead to spurious agglomeration if certain examiners are more likely to reject claims in particular classes.

The last two columns in Table 1 show 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 TCs, with patents in Biotechnology and Chemicals changing classification more often than those in other areas.<sup>13</sup>

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 and keep only applications assigned to a single entity, 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.<sup>14</sup> After completing this process, we have missing assignee data for 584,313 applications (about 20% of our main analysis sample).<sup>15</sup> To check

---

<sup>11</sup>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.

<sup>12</sup>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.

<sup>13</sup>Many papers utilize USPC (sub)classes as a control variable, and future research might consider whether it is better to measure this variation at the time of application publication or grant.

<sup>14</sup>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/patentdatapoint/Home/posts/namestandardizationroutinesuploaded>. Details are available upon request.

<sup>15</sup>This is a relatively large percentage of applications with missing values. However, we are not using assignments by the inventors to organizations that are not their employers in order to measure more precisely the identity of the original applicant, and those as-

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 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.<sup>16</sup>

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<sup>th</sup> percentile in the family size distribution (within an art-unit-filing-year). We call this variable “DOCDB Family Size.”<sup>17</sup>

Finally, we consider whether some examiners are assigned applications seeking greater scope of protection. Kuhn & Thompson (2017) 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<sup>th</sup> 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.<sup>18</sup> We

---

signments that happen after applications are docketed to an examiner in order to utilize only information that is not affected by examination. Moreover, some patents are never assigned by their inventors to an organization, and so they cannot have an assignee in our data. As a comparison, we compute the percentage of granted utility patents filed during our sample period that are unassigned at the grant date. These are about 8%. While we exclude applications assigned to more than one assignee before they are docketed to an examiner, these are less than 3% of the applications assigned by the inventors to their employers before they are docketed to an examiner.

<sup>16</sup>Results of the customer number analysis are similar to those for the assignee and are available upon request.

<sup>17</sup>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.

<sup>18</sup>Kuhn & Thompson (2017) note that this measure of scope is not suitable for the

call this variable “Words in 1<sup>st</sup> Claim.”<sup>19</sup>

In our application-pair-level analyses, the outcome variable is an indicator variable equal to one for pairs that share the same examiner. The main explanatory variables are a Jaccard similarity index based on keywords in the abstract and title of each application, and an indicator variable equal to one for pairs filed by the same assignee.

To compute the Jaccard similarity, we pre-process the abstracts and titles of all published patent applications from the PatentsView patent application database following the procedure described in Arts et al. (2018). Then, for every pair where both applications have at least 10 keywords and share the same art unit and filing year, we compute the Jaccard similarity index as 100 times the number of unique keywords in the intersection of the two sets of keywords, divided by the number of unique keywords in the union of the two sets. Intuitively, this variable measures the percentage overlap of unique keywords. Arts et al. (2018) validate the Jaccard similarity index for the population of granted patents using a panel of experts and information on patent families, inventors, assignees and patent citations. In particular, they show that matching on the Jaccard index outperforms matching on the primary USPC subclass, based on the mean scores from an expert evaluation.

## 4 Results

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

### 4.1 Examiner Specialization

Figure 1 shows that patent examiners handle more applications from a given subclass or assignee than we would expect under random allocation. Specif-

---

analysis of patent scope in biotechnology. So we exclude the Biotechnology TC 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.

<sup>19</sup>Summary statistics for all variables used in the analysis are in Table B1.

ically, 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-filing-year by subclass or assignee cell containing more than 20 applications. For MTAD (bottom row) we run a separate test for each art-unit-filing-year cell containing more than 50 applications.<sup>20</sup>

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. 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).

Table 2 examines the degree of examiner specialization in different TCs, 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.<sup>21</sup> Panel A shows that there is evidence of examiner specialization in every TC, and all of the MTAD results indicate agglomeration rather than over-dispersion. The share of tests that reject random assignment is much lower, however, in the “Computer Architecture” and “Computer Networking” areas than for Biotechnology, Chemistry, Semiconductors and Mechanical Engineering. The results in Panel B, using assignees rather than subclasses, are very similar.

The lower half of Table 2 examines agglomeration for a pair of dichoto-

---

<sup>20</sup>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 B1 shows the distributions of p-values of D-index and MTAD for subclass and assignee with thresholds equal to, respectively, 10 and 25.

<sup>21</sup>Table B3 in the appendix reports analogous figures with a cutoff at the 5-percent threshold for statistical significance.

mous 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 TC we labeled as “Miscellaneous.” There is also evidence of specialization in broader patents (as measured by length of the first claim) in the Chemical and Materials Engineering and Semiconductors TCs. In general, the MTAD tests detect more evidence of specialization than the D-index. But overall, these effects are relatively small, and might easily be caused by the technological specialization observed in Panel A.

The findings in Table 2 raise two questions: (1) What explains heterogeneity in specialization across TCs?, and (2) Is examiner-application matching effectively random within narrower technology subclasses?

With respect to the the first question, there are at least two reasons why we might find less specialization in the computer-related art-units. First, examiners in the less agglomerated TCs 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. Alternatively, patent examiners in the Computers and Communications TCs might be just as specialized as their counterparts in other TCs, but this is not apparent to us because the USPC classification system is less representative of actual technological differences in these fields.

To see how well the USPC classification system works across TCs, we can use the Jaccard similarity index. The first column of Table 3 reports the mean Jaccard similarity for application-pairs within art-unit-filing-years by TC, and column (2) reports the mean Jaccard similarity conditional on sharing the same primary subclass.<sup>22</sup> Not surprisingly, pairs sharing the same primary subclass are more similar across all TCs.<sup>23</sup> To compare the

---

<sup>22</sup>Given the relatively small percentage of application-pairs sharing the same subclass, the unconditional mean similarity is almost identical to the mean similarity conditional on having two different primary subclasses, and we omit the latter statistic for brevity.

<sup>23</sup>We run a t-test for the difference in means for each TC and all tests are statistically

magnitude of these differences, column (3) reports the percentage change in mean Jaccard similarity when conditioning on sharing the same primary subclass. The percentage increase in keyword similarity is large for all TCs, indicating that the USPC system does capture a substantial amount of the variation reflected in the Jaccard measure. However, the percentage increase is much smaller in the Computer and Communication TCs than in other areas, suggesting that the USPC system works less well for those technologies.

To compare “examiner generality” across TCs, we repeat the same exercise, conditioning on application pairs with a common examiner instead of a shared primary subclass. Columns (4) and (5) of Table 3 show the results.<sup>24</sup> Not surprisingly, having the same examiner is less related to technological similarity of applications than having the same subclass: the number of subclasses is much bigger than the number of examiners, and many examiners receive applications from a large number of subclasses. Nevertheless, in each TC, application-pairs assigned to the same examiner are more similar than those with different examiners.<sup>25</sup> The percentage changes are much smaller in the Computers and Communications TCs, suggesting that examiners in that area are assigned more heterogeneous applications. The increase in technological similarity is especially high in Chemistry, where having the same examiner is associated with a 55% increase in the Jaccard similarity.

The results in Table 3 provide evidence supporting *both* of our potential explanations for differences in examiner specialization across TCs. In particular, we find that applications sharing a primary USPC subclass exhibit more similarity in Biotechnology and Chemistry and less in computing. On the other hand, examiners in Computers and Communications examine more diverse applications, which is consistent with a more random allocation of applications.

We turn now to our second question — is examiner-application matching effectively random within subclass? To provide an answer, we construct a data set of all application-pairs sharing the same art unit, filing year and

---

significant at 1%.

<sup>24</sup>Once again, means conditional on having a different examiner are very close to the overall means by TC, so we do not report them in Table 3.

<sup>25</sup>Also all these differences in means are statistically significant at 1% level.

primary subclass, and regress a dummy for having the same examiner on two main explanatory variables: a dummy variable for having the same assignee and the Jaccard similarity index (multiplied by 100 and standardized).<sup>26</sup> We use the Jaccard similarity as a measure of the residual technological specialization that is not captured by the subclasses. Conditional on the subclasses and the Jaccard similarity, the same assignee dummy should capture the agglomeration based on the identity of the assignee that is not explained by observable similarities in technology.<sup>27</sup> Under the null hypothesis of random assignment within primary subclass, the coefficients of the Jaccard similarity index and of the same assignee dummy should be close to zero (and statistically insignificant).<sup>28</sup>

Table 4 presents the results. For the full sample, the estimates in column (1) imply that having the same assignee is correlated with a 5.3 percentage point increase in the probability of having the same examiner. This is a large increase, given that 11% of the application pairs in this sample have the same examiner. The coefficient on the same assignee dummy drops to 2.5 with the addition of art-unit-filing-year-subclass fixed effects, but remains statistically significant and economically large. The Jaccard similarity index is also positively associated with the probability of having the same examiner. A one standard deviation increase in the Jaccard similarity of an application-pair is correlated with a 1 percentage point increase in the probability of having the same examiner in column (1), and a 2.6 percentage

---

<sup>26</sup>In a previous draft, we computed D-index and MTAD statistics for groups of applications sharing the same application-year, art-unit and primary USPC subclass. However, that approach has two drawbacks. First, it is not clear how to measure technology — some patents have secondary classes, but there is no clear hierarchy among them. Second, and more importantly, many USPC subclasses receive only a few applications per year, leading to subclass-year cell sizes that are too small to perform the tests.

<sup>27</sup>We limit the sample to application pairs where both applications have an assignee and art-unit-filing-year-subclasses with at least two pairs. In this sample, the mean Jaccard similarity is 0.079 (standard deviation 0.1). About 11% of the application pairs in this sample have the same examiner, and about 7% have the same assignee. The application-pairs with the same assignee are technologically very similar: the mean Jaccard similarity for application-pairs assigned to the same organization in this sample is 0.25. Also, this mean is particularly high in Biotech, where it is 0.74, and relatively lower in the computing-related TCs.

<sup>28</sup>We cluster at the art-unit-filing-year level in order to avoid the issue of within-dyad correlations, as discussed in Cameron & Miller (2014).



point increase in column (2).

The lower half of Table 4 shows the results by TC.<sup>29</sup> The main message of this set of estimates is that we still observe technological specialization within subclasses. The same assignee dummy is statistically significant at 1% and relatively large in all TCs except Biotechnology. Similarly, the coefficient of the Jaccard similarity is positive, statistically significant at least at 5% and large in all TCs.

Overall, the results in this sub-section show that patent examiners specialize in particular technologies, even within fine-grained USPC technology subclasses. There is more specialization in Chemistry, Biotech, Mechanical Engineering and Semiconductors, and less in computer-related technology. Technological specialization produces agglomeration by assignee, and this agglomeration is robust to controlling for observed differences in technology. But we find very little evidence that certain examiners specialize in “outlier” patent applications, as measured by family size or the scope of the first claim. Finally, we find that the USPC classification is relatively more effective at capturing actual technological similarities in Biotech and Chemistry, and less in computing. However, the lower degree of technological specialization is not entirely an artifact of the lower quality of the technological classification in this area, as examiners in Computers and Communications do examine less similar applications, consistent with the idea that they are more “generalist” than other examiners.

## 4.2 Implications for Examiner-based Instruments

Random matching guarantees that examiner characteristics are uncorrelated with application characteristics. When examiners specialize, however, their attributes and behaviors might be correlated with (or even caused by) the technologies they examine, which could lead to a violation of the exclusion restriction in instrumental variable estimates.

One response to this problem is to condition on fine-grained measures of technology. While this is clearly a good idea, the results in Table 4 (based on

---

<sup>29</sup>For these estimates, we standardize the Jaccard similarity index within each TC.

the Jaccard similarity index) show that even within USPC subclasses, some unobserved specialization will remain. The sheer number of subclasses creates an additional difficulty. Adding subclass fixed effects to an IV regression will not purge all specialization-induced technological variation from a leave-one-out grant rate or similar IV constructed at the examiner-year level.<sup>30</sup> But there are typically too few observations within examiner-subclass-year cells to construct the instrument at this more disaggregated level.

A second response to potential endogeneity is to check whether a particular instrument is correlated with observed differences in technology. Of course, failing to reject the null hypothesis within a particular sample does not validate the instrument — more data or better measures could lead to a different result. For example, Sampat & Williams (2017) regress patent grant on their measure of examiner leniency with and without subclass fixed effects, and obtain very similar point estimates (not statistically distinguishable) for the two specifications. We performed a similar exercise using our full sample of applications, along with new data on the primary USPC subclass at the time of publication, and find that adding subclass effects produces modest changes in the size of first-stage estimates — between 5 and 10 percent for the full sample.<sup>31</sup> In our models, the subclass effects are always jointly significant, and the change in the first-stage coefficient is generally larger for TCs that exhibit more examiner specialization in our prior results.

We propose the following implications for those who still wish to use examiner characteristics as instruments, perhaps because (like us) they do not take a Manichean view of causal inference, and see the approach as a clear step forward in terms of measuring the causal impacts of intellectual prop-

---

<sup>30</sup>Formally, suppose that the probability of issuance for patent  $i$ , in class  $c$ , assigned to examiner  $j$  is given by  $Pr(G_{ijc} = 1) = \theta_j + \eta_c$  (i.e. an examiner-specific leniency and a subclass effect). The examiner-year level IV would then be  $Z_{ijc} = E_{k \neq i}[G_{kjc}] = \theta_j + E[\eta|j]$ . The last term depends on the classes assigned to examiner  $j$ , and will vary across examiners within a subclass-year.

<sup>31</sup>Our examiner leniency measure differs from that in Sampat & Williams (2017). Table B4 presents the results of this exercise for two different IV-endogenous variable combinations: (1) an examiner’s leave-one-out grant rate as an instrument for 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) as an instrument for scope of a focal patent (i.e. number of words in the first independent claim).

erty. First, it is important to carefully control for any observable differences in technology. Subclass fixed effects are not a panacea, but they are nevertheless a step in the right direction. Controlling for identity of the applicant, for example with assignee fixed effects, may also help to capture technological heterogeneity. 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. This assumption is not testable, and its reasonableness will vary from one application to the next. Finally, our results suggest it is more plausible to assume that technology is uncorrelated with potential outcomes for information technology art-units, while the same assumption is potentially problematic for art-units in other TCs.

### 4.3 Specialization and Examination Outcomes

As a final step in our empirical analysis, we explore the relationship between examiner specialization and patent examination 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).<sup>32</sup> 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 an art-unit-filing-year cell) having the same subclass as a focal application.

---

<sup>32</sup>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.

To be more precise, define the set  $k_{it}(j)$  of all patents (except for patent  $i$ ) assigned to examiner  $j$  and filed in year  $t$ .<sup>33</sup> Let  $n_{jt}$  represent the total number of patents reviewed by examiner  $j$  filed in year  $t$ , and define an indicator  $1_{mn}$  that equals one if and only if two patents ( $m$  and  $n$ ) have the same subclass. Our main explanatory variable can be written as:

$$Share_{ijt} = \frac{\sum_{m \in k_{it}(j)} 1_{mi}}{n_{jt} - 1}. \quad (3)$$

Intuitively,  $Share_{ijt}$  equals the probability that a random draw from the pool of applications filed in year  $t$  assigned to examiner  $j$  has the same 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 examination outcomes. To ease interpretation, we standardize  $Share_{ijt}$  and the outcome variables except the dummy for granted patents.<sup>34</sup> 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 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

---

<sup>33</sup>For this analysis we consider an examiner affiliated with two (or more) different art units in the same year as two (or more) examiners.

<sup>34</sup>Table B2 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.

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. 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 examination: the assignment of applications to examiners. We first focus 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 subclasses, even within relatively homogeneous art units. The degree of specialization varies across fields, with examiners in the Computers and Communications area exhibiting relatively little specialization compared to those in other TCs. We also use a measure of textual similarity (and the identity of the assignee) to illustrate residual technological specialization within subclasses. Finally, we show 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. Also, it is reasonable that SPEs take into account the different skills of their examiners when they assign applications. However, several studies have argued that more-or-less random matching (within art-units) provides a justification for using examiner behaviors and characteristics as instruments for examination outcomes. Our findings do not invalidate this identification strategy – patent examiner characteristics might still satisfy the relevant exclusion restrictions – but they do imply that we cannot rely on purely random assignment to justify the approach.

On a more positive note, our results provide some evidence on how the USPTO balances the demands of efficiency and fairness. Technological specialization is likely to be 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. Examiner assignment appears relatively more (but not entirely) random in the computer-related TCs. And even without controlling for technology, there is little 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 the current level of procedural fairness to applicants is also the best policy in terms of social welfare.

Finally, our analysis opens up several avenues for research into the causes and consequences of how work is organized at the USPTO. For example, it might be interesting to study how specialization varies over time for individual examiners, or within a given art-unit. Future research could also study additional outcomes, such as re-examinations, re-issues or invalidation rates of patents issued by examiners with different levels of specialization to better understand the consequences of specialization for patent quality.

## References


- Arts, S., Cassiman, B. & Gomez, J. C. (2018), ‘Text Matching to Measure Patent Similarity’, *Strategic Management Journal* **39**(1), 62–84.
- Cameron, C. & Miller, D. (2014), ‘Robust inference for dyadic data’, *Workign Paper* .
- 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. (2017), ‘What is a Patent Worth? Evidence from the U.S. Patent “Lottery”’, *NBER Working Paper 23268* .
- Feng, J. & Jaravel, X. (2017), ‘Crafting Intellectual Property Rights: Implications for Patent Assertion Entities, Litigation, and Innovation’, *Working Paper* .
- Frakes, M. D. & Wasserman, M. F. (2017), ‘Is the Time Allocated to Review Patent Applications Inducing Examiners to Grant Invalid Patents? Evidence from Microlevel Application Data’, *Review of Economics and Statistics* **99**(3), 550–563.
- 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. & Thompson, N. C. (2017), ‘How to Measure and Draw Causal Inferences with Patent Scope’, *Working Paper* .
- Kullback, S. & Leibler, R. A. (1951), ‘On Information and Sufficiency’, *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* **14**(2), 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. (2017), ‘How Do Patents Affect Follow-on Innovation? Evidence from the Human Genome’, *NBER Working Paper 21666* .
- 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 15851* .



# 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 (cleaned and standardized). 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)

Table 1: Summary Statistics by Technology Center

	Art Units	Examiners	Classes	Sub- Classes	Applications	App pairs*	App pairs, same subclass*	Class Changed†	Subclass Changed†
Biotechnology (1600)	57	1,013	268	11,530	221,586	51.1M	2.5M	32.3	75.3
Chemicals (1700)	75	1,377	411	34,509	367,371	166.6M	1.2M	31.3	77.7
Comp/Comm (2100)	79	1,733	303	7,314	208,102	31.1M	1.2M	23.4	70.7
Comp/Comm (2400)	77	1,289	159	4,492	157,852	23.7M	0.8M	18.5	59.7
Comp/Comm (2600)	82	2,046	310	14,393	338,088	208.0M	3.7	15.4	65.7
Electrical (2800)	80	2,161	382	31,815	637,929	321.0M	4.3M	17.5	65.2
Miscellaneous (3600)	75	1,617	410	38,760	360,691	114.5M	2.0M	16.8	68.8
Mechanical (3700)	65	1,949	411	40,240	425,413	180.5M	1.9M	18.6	68.9
Full sample	590	12,338	452	119,448	2,717,032	1,096.4M	17.6M	20.4	68.5

\* Application pairs with same art unit, filing year (and subclass) with at least 10 unique keywords for both applications. † 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 TCs. Abbreviated TC names from Graham et al. (2015). Full names of the TCs 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-Filing-Year (Share Rejecting Random Allocation at 1% Significance Level, by Technology Center)

Panel A: USPC subclass					
Technology Center	D-index		Agg.	MTAD	
	Rej.	N		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 <sup>†</sup>	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					
Technology Center	D-index		Agg.	MTAD	
	Rej.	N		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 <sup>†</sup>	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					
Technology Center	D-index		Agg.	MTAD	
	Rej.	N		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 <sup>†</sup>	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 1 <sup>st</sup> Claim					
Technology Center	D-index		Agg.	MTAD	
	Rej.	N		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 <sup>†</sup>	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-filing-year cells with more than 20 applications for the D-index and more than 50 applications for MTAD. <sup>†</sup> Miscellaneous = “Transportation, Construction, Electronic Commerce, Agriculture, National Security and License & Review.”

Table 3: Mean Jaccard Similarity by Technology Center

Technology Center	(1) Mean Jaccard Similarity	(2) Mean Jaccard Similarity Same Subclass	(3) % Change Different - Same Subclass	(4) Mean Jaccard Similarity Same Examiner	(5) % Change Different - Same Examiner
Biotechnology (1600)	4.9	11.4	149	5.8	19
Chemicals (1700)	2.9	7.3	152	4.4	55
Comp/Comm (2100)	4.4	5.5	27	4.6	6
Comp/Comm (2400)	4.6	6.1	33	4.8	5
Comp/Comm (2600)	4.5	6.7	52	4.9	10
Electrical (2800)	4.5	7.8	74	5.2	16
Miscellaneous (3600)	3.3	5.6	73	3.7	16
Mechanical (3700)	3.2	6.6	109	3.8	21

We use a sample of patent application pairs to compute the statistics reported in this table. See main text for the definition of the Jaccard similarity. We compute the Jaccard similarity between applications for all application pairs within art-unit-filing-year whose number of keywords is at least 10 for both applications (1,096,408,105 pairs), and multiply it by 100 for easier interpretation. Column (1) reports the mean Jaccard similarity for pairs in each TC. Columns (2) and (4) report the mean Jaccard similarity for pairs that have the same subclass and same examiner, respectively. Columns (3) and (5) report the percentage changes in Jaccard similarity going from different subclass (examiner) to same subclass (examiner). T-tests for the difference in means between application-pairs with the same subclass (examiner) and those with different subclass (examiner) are statistically significant at 1% level in all TCs.

Table 4: Probability of Having the Same Examiner within USPC Subclasses

Outcome		1[Same Examiner]*100	
Technology Center	Regressors	(1)	(2)
Full sample	Same assignee	5.30 (0.37)	2.50 (0.11)
	Jaccard Similarity*100	0.96 (0.40)	2.56 (0.16)
	Art-unit-year-subclass FEs		✓
Biotechnology (1600)	Same assignee	2.27 (2.47)	1.35 (1.29)
	Jaccard Similarity*100	1.37 (0.64)	4.94 (0.45)
Chemicals (1700)	Same assignee	12.60 (2.46)	3.88 (0.47)
	Jaccard Similarity*100	3.17 (0.37)	3.23 (0.15)
Comp/Comm (2100)	Same assignee	3.01 (0.25)	2.74 (0.21)
	Jaccard Similarity*100	1.88 (0.09)	1.72 (0.08)
Comp/Comm (2400)	Same assignee	4.82 (0.34)	4.52 (0.30)
	Jaccard Similarity*100	1.84 (0.10)	1.73 (0.09)
Comp/Comm (2600)	Same assignee	3.08 (0.29)	2.39 (0.21)
	Jaccard Similarity*100	1.00 (0.14)	1.15 (0.07)
Electrical (2800)	Same assignee	3.73 (0.41)	1.66 (0.13)
	Jaccard Similarity*100	0.45 (0.16)	1.25 (0.08)
Miscellaneous (3600)	Same assignee	5.45 (0.48)	2.72 (0.28)
	Jaccard Similarity*100	1.79 (0.55)	1.67 (0.30)
Mechanical (3700)	Same assignee	2.86 (0.52)	2.76 (0.26)
	Jaccard Similarity*100	3.47 (0.23)	2.77 (0.14)

Unit of observation is a patent application pair. In all models, the outcome is an indicator variable equal to one if applications in the pair have the same examiner, multiplied by 100 for easier interpretation of the coefficients. The two main explanatory variables are the Jaccard similarity between keywords of the applications in the pair (multiplied by 100 and standardized within estimation subsample) and an indicator variable equal to 1 if the applications in the pair have the same assignee. Column (2) also includes art-unit-filing-year-subclass fixed effects. The sample includes application pairs from our agglomeration analysis sample (i) examined by the same art unit, (ii) filed in the same year, (iii) classified in the same primary subclass, (iv) whose number of keywords utilized for the computation of the Jaccard similarity is at least 10 for both applications, and (v) with an assignee for both applications. We exclude from the sample art-unit-filing-year-subclasses with just one pair. The full sample contains about 11.7 million patent application pairs. All models estimated by OLS. Robust standard errors in parenthesis, clustered at art-unit-filing-year level.

Table 5: Examiner Specialization and Examination Outcomes.

Model Outcome	Within Examiner			Between Examiner		
	Granted (1)	Words (2)	Days (3)	Granted (4)	Words (5)	Days (6)
$Share_{ijt}$	0.00*** (0.00)	0.00 (0.00)	0.01*** (0.00)	-0.04*** (0.00)	0.08*** (0.01)	0.10*** (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-filing-year-examiner for the between regressions. Sample contains all applications in our primary analysis sample 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. Outcome in column (1) is an indicator variable equal to one for granted applications. Outcome in column (2) is the change in the number of words in the first independent claim between the published application and the granted patent (standardized). Outcome in column (3) is the difference in days between the date an application is docketed to an examiner for the first time and its disposal date (standardized). See the main text for the definition of  $Share_{ijt}$ . Also this variable is standardized. The mean of the outcome of the regression in column 1 is 0.65. Between regressions in columns (4)-(6) estimated on the group means. Standard errors clustered by art-unit-filing-year in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.10

# Appendix A: D-index and MTAD Statistics

## Divergence Index

Suppose we have a set of applications characterized by category  $i \in \mathbf{I} = \{1, \dots, I\}$ , assigned to a set of examiners denoted by  $r \in \mathbf{R} = \{1, \dots, R\}$ . 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 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:

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

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 ( $\lambda$ ):

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

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

## Multinomial Test for Agglomeration and 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  $r = 1, \dots, R$ . The variable  $n_r$  is bounded between  $\underline{n} = 0$  and  $\bar{n} = \infty$  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  $c = 1, \dots, C$ . 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  $x_r^1, \dots, x_r^C$ ,  $\mathbf{p}$  as the vector of probabilities  $p_1, \dots, p_C$ ,  $\mathbf{n}$  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

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

and the average log-likelihood for the data is

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

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

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

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

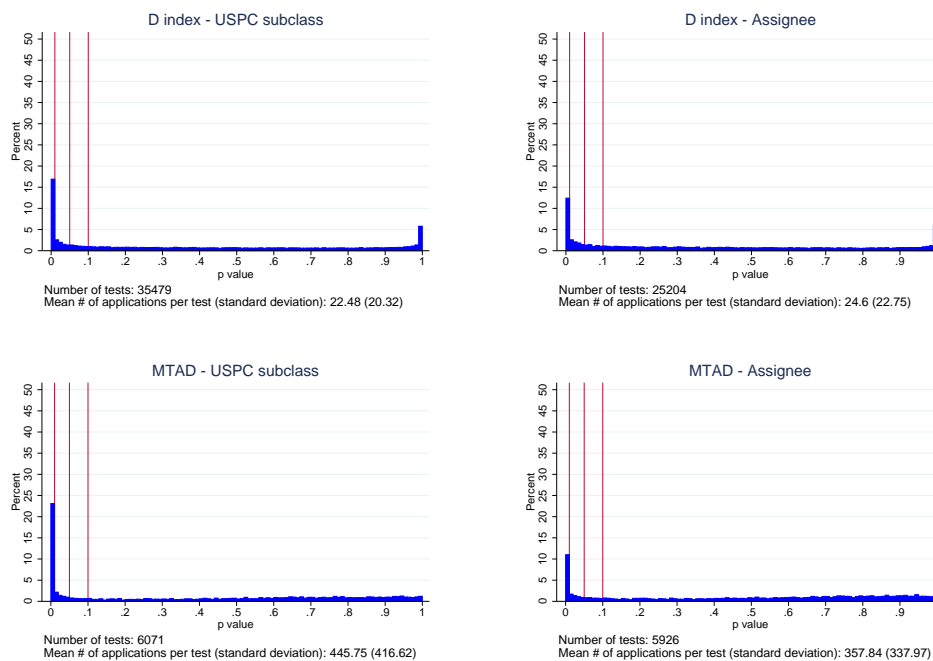
$$t(\mathbf{X}, \mathbf{n}, \mathbf{p}) = l(\mathbf{X}, \mathbf{n}, \mathbf{p}) - E[l(f, \mathbf{p})]. \quad (4)$$

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.



## Appendix B: Additional Tables and Figures (Online Publication Only)

Figure B1: 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 (cleaned and standardized). 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)

Table B1: Summary Statistics for Sample of Applications

Panel A: categorical variables										
Variable	# of categories	Applications per category								
		Mean	Std dev	Min	5 <sup>th</sup> percentile	1 <sup>st</sup> quartile	Median	3 <sup>rd</sup> quartile	95 <sup>th</sup> 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
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 <sup>th</sup> percentile	1 <sup>st</sup> quartile	Median	3 <sup>rd</sup> quartile	95 <sup>th</sup> percentile	Max
DOCDB family size	2,716,195	2.88	5.66	1	1	1	2	3	8	378
Words in 1 <sup>st</sup> 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 B2: 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 B3: D-index and MTAD Tests within Art-Unit-Filing-Year (Share Rejecting Random Allocation at 5% Significance Level, by Technology Center)

Panel A: USPC subclass					
Technology Center	D-index		Agg.	MTAD	
	Rej.	N		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 <sup>†</sup>	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					
Technology Center	D-index		Agg.	MTAD	
	Rej.	N		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 <sup>†</sup>	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					
Technology Center	D-index		Agg.	MTAD	
	Rej.	N		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 <sup>†</sup>	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 1 <sup>st</sup> Claim					
Technology Center	D-index		Agg.	MTAD	
	Rej.	N		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 <sup>†</sup>	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-filing-year cells with more than 20 applications for the D-index and more than 50 applications for MTAD. <sup>†</sup> Miscellaneous = “Transportation, Construction, Electronic Commerce, Agriculture, National Security and License & Review.”

Table B4: IV “First-Stage” With and Without USPC Subclass Effects

Outcome <sup>†</sup>	1[Granted]*100			log(Words-in-1 <sup>st</sup> -claim)		
Potential Instrument	Leave-one-out grant rate *100			Leave-one-out scope change		
	(1)	(2)	(1)/(2)	(3)	(4)	(3)/(4)
Full Sample	0.75 (0.00)	0.68 (0.00)	0.90	0.37 (0.01)	0.35 (0.01)	0.94
Art-unit-year FEs	✓			✓		
Art-unit-year-subclass FEs		✓			✓	
Biotechnology (1600)	0.69 (0.01)	0.61 (0.01)	0.89			
Chemicals (1700)	0.80 (0.01)	0.66 (0.01)	0.83	0.31 (0.02)	0.18 (0.02)	0.57
Comp/Comm (2100)	0.66 (0.01)	0.64 (0.01)	0.98	0.32 (0.01)	0.32 (0.01)	1.00
Comp/Comm (2400)	0.56 (0.01)	0.53 (0.02)	0.95	0.35 (0.01)	0.34 (0.02)	0.98
Comp/Comm (2600)	0.75 (0.01)	0.72 (0.01)	0.96	0.38 (0.01)	0.36 (0.01)	0.95
Electrical (2800)	0.79 (0.01)	0.74 (0.01)	0.94	0.44 (0.01)	0.43 (0.01)	0.96
Miscellaneous (3600)	0.74 (0.01)	0.63 (0.01)	0.85	0.34 (0.01)	0.30 (0.01)	0.88
Mechanical (3700)	0.79 (0.01)	0.69 (0.01)	0.88	0.34 (0.01)	0.30 (0.01)	0.88

<sup>†</sup> Outcome is the endogenous variable in an instrumental variable regression: an indicator variable equal to one if an application is granted multiplied by 100 for easier interpretation of the coefficients in models (1) and (2), and the natural logarithm of the number of words in the first independent claim of a granted patent in models (3) and (4). Coefficients and standard errors for this second set of models are multiplied by 100 for easier interpretation of the results. Sample contains all applications in our primary sample that were either granted or abandoned by the end of the sample period; whose leave-one-out potential instrumental variable is computed with more than 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. Models in columns (1) and (3) include art-unit-filing-year fixed effects, while models in columns (2) and (4) include art-unit-filing-year-subclass fixed effects. Each “first-stage” estimate in this table comes from a separate OLS regression of Outcome on Potential Instrument for applications assigned to a given TC. Robust standard errors, clustered by art-unit-filing-year, in parentheses. All estimates are statistically significant at the 1% level. We exclude Biotechnology patents (TC 1600) from the second set of estimates because Kuhn & Thompson (2017) suggest that counting words in the first claim does not yield a meaningful measure of claim-scope for those applications. To test whether differences between the estimates in columns (1) and (2) and those in columns (3) and (4) 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. All tests but those for the patent scope models in TCs 2100 and 2400 are statistically significant at 1% level. 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 likelihood ratio tests are statistically significant at the 1% level.