The Unintended Consequences of Crowd-sourced Patent Examination

Jin-Hyuk Kim* Benjamin Mitra-Kahn†

Abstract

There have been repeated calls for improving the quality of patent examination at the United States Patent and Trademark Office. One program aimed to address this, piloted between 2007 and 2009, sought to crowd-source information for patent examiners to help them examine emerging technologies. Since then, calls have been made to re-instate or expand this Peer-to-Patent program, but there has not been a systematic study of the causal links between program participation and the final outcome of the patent applications. By comparing the participating applications to a control group, we show that the program marginally reduced the incidence of patent allowance, but it resulted in a doubling of the examiner's search effort afterwards. Contrary to the aims of the patent office, participation in the pilot also led applicants to file more requests for continued examination, and this extended prosecution did not lead to a change in

*University of Colorado at Boulder, 256 UCB, Boulder, CO 80309, USA. Tel: 1 303 492 3827. Email: JinHyuk.Kim@colorado.edu.

†IP Australia, 47 Bowes Street, Woden, ACT 2606, Australia. Tel: 61 2 6283 2291. Email: Benjamin.Mitra-Kahn@ipaustralia.gov.au

the average outcome of the patent application, meaning both applicants and the office incurred higher cost at no additional benefit.

Keywords: patent examination; patent office; open government; crowdsourcing

JEL codes: H83, L17, O34, O38

Highlights

- We examine the causal effects of Peer-to-Patent pilot program using administrative data.
- There is insufficient evidence that the pilot program provided a net benefit to the examination process.
- The pilot program appears to have imposed a substantial search cost on the examiners after the first office action.

1 Introduction

Over the last 15 years there has been considerable debate over the patent system and how it might be improved (Jaffe and Lerner, 2004). The issue is that the US Patent and Trademark Office (PTO) is seen to be granting patents that should have been narrowed or rejected (Merril et al., 2004; Bessen and Meurer, 2008). The reason for this phenomenon is quite well known. That is, according to the Federal Trade Commission (2003), the PTO has lacked enough resources and information to make thorough evaluations of patent applications in a reasonable amount of time.[1] PTO examiners can identify an application's patentability and devote more search effort to weaker patents (Lei and Wright, 2017); however, with a backlog of over half a million applications, examiner knowledge may well be under-utilised.[2]

The issue of weak patents has led economists to propose various policy mechanisms to disincentivise low quality patent applications. These include a post-grant review system along the lines of what was introduced through the America Invents Act with the Patent Trial and Appeal Board (Hall et al., 2004), a pre-licensing review to determine patent validity (Farrell and Shapiro, 2008), a penalty system for rejected applications (Caillaud and Duchêne, 2011), and tiered levels of patent examinations (Picard and de la Potterie, 2013). These proposals involve substantial changes in legal standards or political support from Congress. Thus, given the potentially high cost of reform, it would be prudent to carefully examine and document whether the initial screening can be made more effective especially without deploying much

[1] Patent examiners have between eight and 25 hours "to read and understand the application, search for prior art, evaluate patentability, communicate with the applicant, work out necessary revisions, and reach and write up conclusions" (Federal Trade Commission, 2003).

[2] Jaffe and Lerner (2004) argues that the incentive systems in the USPTO are not geared towards the goal of maintaining high quality. For instance, examiners are rewarded based on the counts of first office action on the merits, while searching for prior art is not explicitly rewarded.

resources.[3]

One method of expanding resources and gathering information is to leverage public knowledge. Noveck (2006) criticised that the current patent system places undue reliance on centralised structures of procedural expertise and decision-making, and argued that the public or the scientific community could provide useful information relevant to the claims of patent applications. Specifically, patent examiners rely on their own knowledge and that of the established literature—the "prior art"—which is mainly drawn from previous patent examination reports. Thus, it was argued that examiners at the PTO would be at a disadvantage in examining technologies that were in emerging technologies, or in areas where there was little prior art to be had from previous exam reports (Allen et al., 2008).

These issues were seen to be particularly acute in areas relating to computer software and business methods. Responding to this, the PTO launched a pilot program with the New York Law School in June 2007 to crowd-source relevant prior art information on live patent applications.[4] This platform allowed volunteer citizen-experts to submit and discuss prior art from both patent and non-patent literatures relevant to the claims of published patent applications, which was then formally delivered to the PTO examiners before the patent office took its first action on the application. The over-riding policy aim was to provide examiners with better prior art from a broader variety of sources to allow them to reject patents which were not novel.

The aim of this paper is to empirically evaluate the effects of this pilot program. The

[3] Lemley (2001) argues that the PTO exercises "rational ignorance" when it grants undeserving patents. That is, because most patents are dormant after they are granted, resources spent on rigorous up-front screening processes would be wasted if the PTO's criteria for grant remained the same.

[4] All citizen activities were conducted on the 'Peer-to-Patent' website (http://peertopatent.org/) managed by Professor Beth Noveck and students at the New York Law School's Center for Patent Innovations.

pilot's annual progress reports mainly relied on surveying the participants on both the public and patent office side, and the summary statistics reported by these reports were limited to the first decision issued by the examiners (Allen et al., 2008, 2009, 2012; Loiselle et al., 2010).[5] Patent applications however take many years to complete their examination, so investigating the efficacy of the pilot, relative to a control group, is not possible in a short time span.[6] Therefore, no formal analysis has been undertaken of this crowd-sourcing experiment. We aim to fill this gap by using a rich set of administrative data that took several years to emerge.

The first substantive examination report undertaken by the PTO is only the first step in what can be a long process to receive a final decision on a patent application. The first office action is important, and sets out the detailed view that the examiner has taken of the application, but only 11.4 percent of applications filed from 1996 to 2005 were allowed by an examiner at this stage (Carley et al., 2015). The remaining applications go through to a secondary examination stage, which is where the PTO expends the majority of its examination resources, and applicant's costs and additional delays occur. Figure 1 sets out a simplified version of the PTO examination process illustrating that the first report is only the first step.

In the literature, the use of administrative data, such as the examiner's multiple search reports and the applicant's requests for continued examination during the examination process have been often neglected. Further, the eventual outcomes of the patent applications, such

[5] For instance, Allen et al. (2012) is limited to comment that "38 applications were rejected based on 41 references submitted through Peer To Patent" when setting out the case for the pilot's merits.

[6] In fact, applications filed in 2007 would on average not have reached an examiner for a first report until more than two years later as the average pendency to first report was 25.3 months (PTO, 2008).

as allowance and citations, as well as the behaviors of the applicant and the examiner have not been evaluated in previous studies of the Peer-to-Patent pilot. Rather, despite calls for the expansion (Bestor and Hamp, 2010), embrace of the US pilot (Weatherall, 2009), and the extension of crowd-sourcing in other countries (Ghafaele and Gibert, 2011) the PTO discontinued the pilot in 2011, neither side of the argument having evidence of the efficacy of the program.

Following Jaffe et al. (1993), this paper empirically investigates what happened to the patent applications that went through the Peer-to-Patent pilot, by comparing them to a control group of applications at the three-digit patent technology class that were not in the pilot. The literature has seen some debate on whether three-digit-level controls are too broad as a sample selection criterion (Thompson and Fox-Kean, 2005); however, using controls based on the patent subclasses has its own problems and in particular it can lead to unreliable point estimates (Henderson et al., 2005). Our own observation with the data was that the initial subclass assignment of both treatment and control applications often changed throughout the examination process. Hence, our preferred classification system for constructing a control group is the three-digit level.

The challenge in estimating treatment effects is that participants chose to be in the program. Note that some authors use the quasi-random assignment of cases to examiners, and the heterogenous grant rate ("leniency") of examiners to construct an instrumental variable for which applications are granted patents (e.g., Sampat and Williams, 2015; Farre-Mensa et al., 2017). These approaches rely on the homogeneity of the examiners apart from the said leniency conditional on other observable characteristics, but also on the valid random assignment of patent applications to examiners within certain examination groups at the PTO. Unfortunately, this instrumental approach does not work with the Peer-to-Patent pilot as applicants self-selected into the pilot, and as we will show below the assignment to examiners does not appear to be random for the pilot program.

Our approach is to instrument program participation based on the examination process related to the initial filing of a patent. Between an application being received in the office, and the application being ready for the examiner there can be a delay. There are administrative actions which the office has to undertake, there may be missing documentation or formatting, and sometimes revisions have to be made. The main incentive for applicants to participate in the pilot was that their application was advanced out of turn, so examiners would undertake a first action early. Importantly, the eligibility criterion for the pilot program was based on the date of publication, not the filing of application. We thus use the initial delay between the filing and the publication of application as our instrument because the initial delay is often due to requests for missing formalities and revising documents.

There are three layers to our analysis. First, we find that program participation had a marginally negative effect on allowance (i.e., a notice that an invention qualifies for a patent); and applications which participated in the pilot had more prior art cited in the examiner's search report before the first office action, which was what the pilot program was intended to do. Program participation had no effect on the number of forward citations (i.e., the number of subsequent patents which cite that application), meaning that the open public review did not have material effects on the importance or strength of the patents or its applications.

We then focus on the treatment group to assess and compare the direct effects of proposed peer review activities (such as the quantity of prior art submitted by the public, prior art which is not patents and discussions on the Peer-to-Patent website) on the same set of outcome measures. Because the effect of prior art is likely to be endogenous, we employ the size of the Peer-to-Patent community as an instrument to identify a causal effect of those activities (which we elaborate in Section 4). Here, we find that within the treatment group additional prior art submitted by the volunteers did not affect any of the outcome measures.

However, allowance was negatively affected by participation in the pilot; that is, examiners were more likely to reject applications that had crowd-sourced prior art. Thus, we can untangle the direct effects of receiving more prior art references, and the overall impact that the pilot had. We conjecture that the reason why there was a marginally negative effect from the pilot participation was that the pilot may have influenced the examiner's behavior. When we analyse the subsequent work by examiners, we find that they had to work more on applications that were treated, and in fact undertook an additional search report on these applications.

More troubling, from a PTO perspective, is the finding that participation in the pilot is causally associated with an increase in the number of request for continued examinations (RCEs). RCEs are a way for applicants to request additional negotiation with the examiner after an applicant has received a "final" rejection and prosecution is closed. Patent prosecution then continues until the application is allowed or reaches another "final" rejection. Thus, a high number of RCEs that are eventually abandoned represent a substantial waste of resources and a problem for reducing the backlog of patent applications.

This is an issue for the PTO, because the office has been actively seeking to reduce the number of RCEs which it considers a resource drain away from examining new applications.[7]

[7] In 2012 the PTO posted a request through the Federal Register setting out its activities to reduce the demand for RCEs and commencing a public consultation seeking ideas on how to address RCE demand. The PTO has repeatedly sought to raise the price of RCEs, which it did again for financial year 2017.

Given that the crowd-sourced prior art was intended to help examiners weed out low quality patent applications and help reduce the backlog of pending applications, the above results appear to be quite the opposite to the goals of the pilot program. Further, the positive treatment effect on the number of RCEs is a concern because it would likely to lead the applicant to additional resources and time spent in prosecuting RCEs.

To test this hypothesis, we estimate another set of treatment effects using the number of examiner search reports as well as the number of references cited therein. The result is that the program participation is causally associated with an increase in the number of examiner search reports and also references cited therein, in particular, after the first office action. This evidence corroborates our conjecture that the increased number of RCEs by program participants led to additional burden on the examiners and the PTO. Given the null effect of crowd-sourced prior art on allowance, the decision to discontinue the pilot seems warranted.

This paper contributes to a few streams of literatures. One is the extensive literature on patent and innovation (see Jaffe and Trajtenberg (2002) for an overview). In particular, a number of authors construct a control group to study such topics as the pattern of knowledge spillovers (Jaffe et al., 1993), the determinants of patent suits (Lanjouw and Schankerman, 2001), and the effects of institutions on innovation (Furman and Stern, 2011). Our methodology is similar in this regard, but the difference is that we study the effect of public knowledge provision on the patent examination process which was supposed to reduce the quantity of weak patents.

There is a growing interest in the efficacy of the patent examination process and examiners in particular. For instance, Alcácer and Gittelman (2006) and Sampat (2010) find that a large share of citations on the granted patents is provided by examiners rather than by applicants. Frakes and Wasserman (2016) shows that insufficient examination time after examiners are promoted may hamper examiner search. Carley et al. (2015) examine the prosecution histories of 2.15 million US patent applications, such as allowance rate and continuation procedures. Our paper contributes to this literature by examining the applicants' and examiners' behavior in the pilot.

Another related literature is the literature on open sourcing. Johnson (2002) and Lerner and Tirole (2002, 2005) explain that volunteers may work for an open source software project when they have impact on outcomes and/or their efforts may be rewarded in the labor market, for which some empirical evidence exists (Belenzon and Schankerman, 2015). While the motivation for public participation in government crowd-sourcing projects may be similar, one difference is that in the public administration, the crowd may not completely substitute for the government agencies but it is likely to aid the agencies' decision.

The Peer-to-Patent pilot was often touted as an example of the Open Government initiative by the Obama Administration. The existing literature in the field of public administration on this topic, however, comprises mostly qualitative studies focusing on the program characteristics and/or opinion surveys (e.g., Grimmelikhuijsen and Feeney, 2016; Kornberger et al., 2017; Piotrowski, 2017) rather than quantitative studies using the administrative data. Our findings illustrate that the public participation in government administration need to be carefully designed and its outcomes be evaluated by using objective data.

2 Data Description

Our dataset comprises the 225 patent applications posted to the Peer-to-Patent website (the treatment group) and a random sample of 580 regular patent applications filed at the PTO (the control group), which are all utility patent applications. The control group was sampled 3:1 ratio in order to assure enough sample size and sufficient statistical power. To compare similar technologies, the control samples were randomly drawn from each primary, three-digit classification that was eligible to participate in the pilot program. Specifically, only applications assigned to three technologies defined at the PTO as Business Methods, Computer Architecture and Computer Networks & Cryptography and Security were eligible for participation in the Peer-to-Patent pilot.

Specifically, within each Technology Center (TC) only a sub-set of technology classes were eligible for participation: Business Methods (TC 3600): class 705; Computer Architecture (TC 2100) classes: 380, 700, 703, 706, 707, 708, 710, 711, 712, 714, 715, 717, 718, 719; Computer Networks & Cryptography and Security (TC 2400) classes: 380, 709, 713, 726. These eligible three-digit classes again have a number of subclasses, some of which are assigned to Group Art Units outside the three TCs (2100, 2400 and 3600). We eliminated such ineligible draws from our control sample. Further, to be eligible, the timing of the publication of patent application had to occur during the pilot period, which ran from June 15, 2007, to June 15, 2009.[8]

We first collect information on the status of each application in the treatment and the

[8] To be precise, the pilot was initially scheduled to last for one year, but in July 2008 the PTO extended the pilot for another year and also expanded the eligibility to include the automated business data processing technologies ('business methods' patents—class 705). Hence, our sampling period for class 705 only runs from July 16, 2008, to June 15, 2009.

control groups from the PTO's public Patent Application Information Retrieval (PAIR) database. Because a patent prosecution can take several years until the final disposition, data collection could only be done several years after the pilot had finished. There are still 4 out of the 225 applications (1.8%) in the treatment group and 13 out of the 580 applications (2.2%) in the control group that have not yet finished prosecution by April 2017. While abnormally long examination periods are interesting cases, these would not change our results given the small number of cases. We thus exclude them from our analysis in order to focus on the effect of the pilot on the average application.

For the 788 applications that have been disposed (i.e., either allowed by the PTO or abandoned by the applicant), we assigne 1 to an indicator variable, *Allowance*, if the application was allowed (document code 'NOA' from the PAIR database) and 0 if abandoned (i.e., when the notice of abandonment is mailed). A few applications were abandoned after allowance, for failure to pay the issue fee. Because we are interested in the effect on examination outcomes, we use an indicator for allowance rather than grant. Similarly, we count the *Number of requests for continued examination* (document code 'RCEX' from the PAIR database) for each application, before the application was disposed.

We then construct a measure of forward citations (i.e., the number of times a patent or its application is referenced by later patent applications or patents) for each application. Notice that citations can be made to both applications and patents, because a patent application is prior art and searchable, granted or not.[9] In fact, patent applications contain a nontrivial

[9]Patent applications are searched under 35 U.S.C. 102(e)(1); and issued U.S. patents and foreign patent documents are searched only after published patent applications are searched under 102(e)(2). While the literature often focuses on citations to patents only, citations to applications are also important for our purpose, given the long average pendancy.

amount of citations that do not re-appear in the issued patents because it takes years to be issued a patent. Using Google Patent, we counted the *Number of forward citations* for each application and patent, if granted, netting out any duplicate citations.[10]

There is considerable evidence that forward citations of patents are a proxy for the importance or the economic value of the patents (e.g., Trajtenberg, 1990; Harhoff et al., 1999; Hall et al., 2005; and Belenzon, 2012). Our use of forward citation is a little different in that what we want to know is the treatment effect of the pilot program on forward citations. Forward citations have a right-censoring problem; however, our data spans eight to ten years since the date of application publication. Further, our treatment and control samples are drawn from the same time period, so the truncation is unlikely to cause a systematic bias in our estimates.

From the PAIR database, the applicant's *Small entity* status claimed during the examination process was also collected.[11] The main reason to file as a small entity is that the fees are reduced by 50 percent, and we control for this variable in line with other works that use PTO data (e.g., Bessen, 2008). The common hypothesis is that fee-conscious applicants may be less likely to hire a patent attorney who can help search for prior art and write patent claims. In fact, there are some evidence that un-represented applicants file lower-quality applications on average, with a lower probability of allowance (Gaudry, 2012).

Our instrument for program participation (*Days from application to publication*) is the

[10] That is, we count the number of unique titles citing each patent or its application. This way we eliminate counting more than once the same references made by a patent application and its later-issued patent, as well as the related patent applications or patents filed at or issued by multiple offices.

[11] This information is collected from Fee Worksheet ('SB06') or Issue Fee Payment ('PTO-85B') in the PAIR database. A small entity means any small business, independent inventor, and nonprofit organization as defined under Section 3 of the Small Business Act.

number of days between the filing of the application and the earliest publication date, both of which are collected from the summary page of the PAIR database. We also collected from the PAIR database the name of the patent examiner as well as the name and location of the first inventor. In line with other works (e.g., Lemley and Sampat, 2012), we collect information on whether the examiner is a primary examiner who has the authority to dispose applications or an assistant examiner who is under supervision of a primary examiner.

For this, we look up the PTO's full-text search database for patents and patent applications. If a patent was granted, then the patent document indicates whether the examiner was a primary or an assistant examiner at the time of patent publication. If an application was abandoned, such information is not directly available. We then searched the full-text search database to find out if the examiner had issued patents as a primary or had acted only as an assistant examiner at the time when the said application was abandoned. This was relatively straightforward as we did not find examiners who were suspected of having the same name in our sample period.

In line with other works (e.g., Lanjouw and Schankerman, 2004) we also collected the *Number of independent, and dependent claims* in the patent applications from the full-text search database. While independent ("parent") claims describe the invention, dependent claims, which refer to a parent claim, tend to narrow down the scope of the invention. We also searched the database to collect the *Number of the first-named inventor's previous patent applications*, which serves as a proxy for the inventor's experience.[12] Only the patent

[12] This means that we searched the inventor name by varying the middle name convention, while limiting the search to the same state (for US applicants), or the same city and country (for foreign applicants). When an inventor name appeared in more than on location, we used other information (such as assignees) and our best judgement to construct this measure.

applications filed since 2001 are searchable, so this measure spans only the last 16 years.

Following the literature, to control for the quality of application we collected information on an application's patent family. A patent family is a set of related applications filed at multiple patent offices to protect a common invention. From the European Patent Office (EPO) database Espacenet, we measure the size of patent family as the number of distinct offices (other than the USPTO) in which related applications were published.[13] In addition, we include a dummy variable indicating whether the same or related application was filed at the EPO (*EPO filing*) as applications filed with the EPO suggest the application will be used in the European market.

Examiner's search behavior is measured by the *Number of examiner search reports* (document code '892' from the PAIR database) and the number of references cited by the examiner in each report. Specifically, the examiner has to document in the PTO form 892 a list of references relevant to the decision to be rendered. Notice that all prior art submissions from the pilot program were delivered to the examiner before the first office action. Thus, we cannot distinguish which references in the first search report were only found by the volunteers or found by both the volunteers and the examiners, before the first office action.

However, for each search report appearing after the first office action, we can identify what new references were added by the examiner in the subsequent examination period. Since the live period during which the volunteers could submit art and discuss applications on the pilot website ends before the first office action, the additional (i.e., previously unreported)

[13] This means that we searched the Espacenet database to find the set of documents having the same priority as those of our samples, which is automatically listed per item. We did not count the WIPO patents (also known as PCT applications) because the PCT application does not itself result in a grant of patent rights in any national territory.

prior art appearing in the examiner's second or later reports can be assumed to be coming from the examiner's search. Therefore, we count the *Number of new references contained in the examiner's search reports before and after the first office action*, separately.

Finally, for those applications that went through the Peer-to-Patent (P2P) pilot, we measure the *Size of P2P community* for each application, by the number of registered users in each community (i.e., a community is set up for each posted application), which include both active and inactive participants; that is, a user may register but need not post anything to the community. Anecdotal evidence suggests that the community size is influenced by the media exposure and the promotional effort. We will explain below that the community size is plausibly exogenous to the PTO examination processes and outcomes.

We also collect the *Number of patent and non-patent art references* as well as the *Number of discussions* posted on each community in the P2P website. The pilot allowed only the ten best prior art references to be forwarded to the PTO examiners for their consideration, but there are only three cases where the number of community-identified prior art exceeded ten (and relaxing this constraint does not change any of our results). While only prior art references were forwarded to the PTO, the existing reviews of the pilot indicate that the examiners often read all the comments posted on the P2P website (Allen et al., 2012).

Summary statistics are presented in Table 1. Overall, 63% of the applications were given a notice of allowance, and an average application has less than one (0.6) request for continued examination by the applicant. There is no statistically significant difference between the treatment and the control groups regarding the incidence of allowance or the number of RCEs. On the other hand, the treatment group has on average four more unique forward citations than the control group does, which is statistically significant. These summary statistics can however be misleading, and it is not wise to base a conclusion about the efficacy of the pilot on this, because they do not take into account possible correlation between the outcome and covariates as well as the endogeneity of participation.

The initial delay from application filing to publication is higher, on average, for the treatment group by 86 days, which is statistically significant at the 1% level. We have found that this is mostly due to missing formalities such as oaths or declarations and requests for new or revised application documents such as drawings. The share of applications assigned to assistant examiners is lower for the treatment group than for the control group, which is also statistically significant at the 1% level. Lemley and Sampat (2012) says that the case assignment mostly depend on familiarity with particular technologies and docket flow management. The pilot assignment may have received some special treatment.

It is also found that the treatment group is more likely to claim a small entity filing status, and it has a smaller size of related foreign patent applications as well as a lower incidence of EPO filing. This suggests that those who did not possess sufficient financial resources to apply for patents were more likely to participate in the pilot. The number of dependent claims is marginally lower for the treatment group, but there is no statistically significant difference concerning the first inventor's experience in patent applications. Finally, the number of prior art cited in the examiner's report before the first action is higher for the treatment group, which seems in line with what the pilot program was intended to do.

3 Treatment Effects on Patent Application

The challenge in estimating a program treatment effect is to overcome the omitted variable bias. That is, the patent applications in the treatment group may be selected in a way that is unobservable but correlated with the outcome variables of interest. Because we cannot observe the counterfactual outcome for the treatment sample had it not been posted to the community review, we need an instrument to identify the causal effects of the pilot program. The instrument employed in this section is the number of days between the initial filing of the application and the subsequent publication of the application once all formalities are addressed. Because such delay lengthens the examination process, the applicant may have an incentive to expedite the process by participating in the pilot.

Thus, our exclusion assumption is that initial delays are caused by relatively exogenous administrative issues that are not likely to be correlated with the error term in our regression models for patent examination and outcomes. In fact, there is a considerable variation in the days until publication, which are due to publication schedules and queues but also due to applicants being required file missing forms or drawings. Specifically, it is not uncommon for the PTO to issue pre-exam notices requesting missing oath, fees or new drawings, and the applicant has to respond to such notices in order to complete the application. Only when the application is completed is it placed in the publication queue, and a notice of projected publication date is sent to the applicant.

There are a few reasons to suspect that a link exists between the initial delay to publish and pilot participation, particularly in the computer technologies—subject matters that were eligible for the pilot. First, the 20-year patent term begins from the date when the application was filed but is not enforceable until granted by the office, which means that the time between filing and grant would be lost because the patent is not enforceable during that time period. In a field with rapidly changing technologies, a timely prosecution of the application would be valuable. Second, the initial delay taken by the applicant in responding to the pre-exam notices may reduce the patent term should it be granted later on. The longer the delay, the more likely the patent term reduction.[14]

The main benefit of the Peer-to-Patent pilot was that participating applications were advanced out of turn for the first office action at no cost to the applicant. This contrasts to the Accelerated Examination track introduced in 2006, for which the applicant needed to pay an additional $4800 fee.[15] Thus, the pilot program was an inexpensive way to accelerate the first office action. The summary statistics in Table 1 already showed that the mean delay until publication is higher in the treatment group by 86 days relative to the control group; and the no-cost acceleration benefit of the pilot could have been particularly attractive for small entity applicants who experienced initial delays. Therefore, the initial delay would be the main driving force of our two-stage least squares estimation.

Our empirical specification is

$$Outcome_i = \beta Treat_i + \gamma' X_i + Class_i + Center_i + \varepsilon_i,$$

where $Outcome_i$ is the incidence of allowance, the number of RCEs, or the number of forward

[14] Section 154(b) of the Patent Act adds additional time to the patent term if the PTO failed to process or examine a patent application in unreasonable time, and it reduces the patent term for extensions or delay taken by the applicant.

[15] Other petitions to obtain a fast-track examination were available but they were based on some specific areas (e.g., environmental technology) which did not overlap with the patent classes that were eligible to participate in the pilot.

citations of patent application i, as elaborated in the previous section. $Treat_i$ is an indicator for the treatment group, and X_i includes a set of controls introduced above. $Class_i$ is a set of three-digit class dummies (where 709 is the excluded class); and $Center_i$ is a set of technology center dummies (where 2100 is the excluded center). We allow the error term ε_i to be heteroscedasticity-consistent and correlated within the class.

Table 2 presents the first-stage result that projects the participation in the pilot. As in the main (second-stage) regressions that will follow, the first-stage regression includes patent class as well as technology center dummies, for which we do not report point estimates for brevity. The number of days from application until publication is indeed significantly associated with the program participation at the 1% level. To see if the excluded variable can plausibly identify the treatment effect, we report the standard and Cragg-Donald F statistics, which test the null hypothesis of under-identification, at the bottom of the table. Both statistics render support for our instrumental variable, and we also reject the null of "weak" instrument, as suggested by Stock and Yogo (2005).

The questions we seek to answer is whether the pilot participation affected the eventual outcome of allowance, the number of RCEs filed by applicants, and the number of forward citations received by participating applications and patents, if granted. Table 3 shows the estimation results for the three outcome variables, where we present the OLS and 2SLS estimation results for comparison. The OLS estimates would be inconsistent if the program participation is likely to be correlated with the error term, but if the exclusion restriction is valid, our 2SLS estimates would yield consistent point estimates. Notice that the patent classes are not perfectly nested within technology centers as some classes can be dealt with in more than one center.[16]

We find several things. The average treatment effect of the pilot on allowance is marginally significant at the 10% significance level. Importantly, the sign of point estimate changes from positive in the OLS to negative in the 2SLS estimates. If we accept that the instrumental variable indeed purged the endogeneity between the pilot and the outcome, the result implies that the pilot marginally helped reduce the number of granted patents, although the mechanism behind this finding remains to be seen. For instance, this could be because the prior art submission by volunteers indeed helped reject unworthy applications, or the treated applications incited some reaction by the examiners, or both. We will try to tease these out in the subsequent sections.

In the literature, the existing significant predictors of whether an application is granted include the examiner's job status (Lemley and Sampat, 2012) and the whether the applicant is a small and medium-sized enterprise (e.g. Bessen, 2008). Our results in Table 3 confirm both these previous findings as assistant examiners are significantly less likely to allow patent applications, and that small entities are less likely to succeed. Our 2SLS estimates suggest that the pilot program had a negative effect on allowance ($-.35$) of the same size as the effects of being assigned to an assistant examiner ($-.23$) and having the small entity status ($-.12$) combined.

The results on the number of RCEs (in the middle column of Table 3) is unexpected. The 2SLS uncovers a statistically significant positive impact of the treatment (at the 1% level) on the number of RCEs filed by the applicant. Given that RCEs are potentially a waste of

[16] Clustering standard errors by technology center as well does not change the main qualitative results as reported in this paper; however, it leads to too few observations within cluster. Hence, clustering by patent class is our preferred specification.

resources from both the applicant's and the examiner's end, unless it leads to a reversal of the final rejection, this indicates a major problem for the pilot program. The results also show that small entities are less likely to file RCEs, and there is also a positive, significant relationship between the number of dependent claims and RCEs, and between the size of patent family and the number of RCEs, which seems intuitive.

Moving on to the citation outcome (in the right-hand column), we find that the participation in the pilot had no statistically significant effect on the number of forward citations. This implies that going through the pilot program did not have a lasting impact on the inventors or the examiners, as some might have thought.[17] At the same time, this implies that our instrumental variable appropriately purged unobserved heterogeneity between treatment and control groups if indeed there should be no treatment effect of the pilot on forward citation.

4 Direct Effects of Prior Art

In this section, we examine the effects of crowd-sourced prior art on the same set of outcomes by using only the treatment sample. The literature on treatment effects often abstracts from the direct effects of treatment activities because relevant data may be lacking; however, the underlying forces may not be easily captured in a reduced-form manner between the treatment indicator and the outcome variables of interest. We thus believe that examining the direct effects can yield useful checks against the stated goals of the pilot program and also refine the interpretation of the results obtained in the previous section. For instance, we can

[17] This finding is in a sense consistent with McCabe and Snyder (2015), who shows that open-access journals need not boost forward citations of the articles published therein, once fixed effects for journals (or journal volumes) are included.

find out whether or not the negative treatment effect on allowance was due to effectiveness of the community review.

We have data on the quantity of patent and non-patent prior art as well as the number of discussions for each application posted on the Peer-to-Patent (P2P) website. Because an invention is not novel (hence, not patentable) if it is already known by the public, the submission of prior art, if useful and relevant to the application, should have a negative effect on allowance. Similarly, it was hoped that the prior art identified by the P2P community at an early stage of the patent prosecution would reduce informational asymmetry between the applicant and the examiner, so that the application would be disposed (abandoned or allowed) earlier than otherwise. The results below cast some doubt that these goals were met. That is, we will show that a larger number of crowd-sourced prior art references had no statistically significant effects on allowance, the number of RCEs or forward citations.

One issue in estimating the causal effect of prior art submission is the likely endogeneity of the quantity of prior art to the outcome variables of interests. For instance, weaker applications are more likely to receive more prior art submissions because the marginal impact of prior art is higher and there are also more references that can be found. We argue that the size of Peer-to-Patent community can be used as a plausible instrument for these activities. The various reviews of the pilot as undertaken by Allen et al. (2008, 2009) and others state that volunteering in the P2P community was driven by invitations sent out by student facilitators at the New York Law School asking potential reviewers to provide prior art submissions and discuss a particular application.

For instance, Allen et al. (2009) says "the student facilitator clearly and concisely summarizes the patent application. This process allows the students to develop refined keyword searches in order to identify potential reviewers. The students then use these summaries for recruiting individuals specializing solely in a particular application field in an attempt to increase expert participation in the review process and the contribution of quality prior art. To enlist these potential peer reviewers, individuals are contacted at academic research universities, private companies, and various journals or blogs." Since the contacts were initiated by the student facilitators who were not experts, the community size is not likely to be endogenous to the application's quality.

The first stage regression is reported in Table 4, with the columns providing estimation for the total number of prior art; patent art only; non-patent art only; and the number discussions related to the application. This is instrumented with the excluded variable, the size of Peer-to-Patent community for each application. In all four columns, the P2P community size is a significant positive predictor (at the 1% level) for each of the four types of input we are interested in. At the bottom of each column are the F statistics. In all cases, the standard F statistics stays above 10, the rule of thumb number; and in all cases the Cragg-Donald F statistics is larger than the critical values suggested by Stock and Yogo (2005), which means that any bias due to weak instrument is not likely to drive our results.

Using the first stage results, we can then use 2SLS to investigate the impact which different types of prior art and discussion have on our three outcomes variables of interest. Table 5 shows the main estimation results from the second stage regressions, where the only difference from the previous specification in Table 3 is that we now have the total number of prior art submissions from the Peer-to-Patent participants in the first row instead of the treatment status, and the data is limited to the treatment sample.

The results are rather unexpected, and need to be considered in the context of the overall treatment results already discussed. First, the number of crowd-sourced prior art is not causally associated with a lower incidence of allowance, which contrasts with the stated goal of the pilot. This result suggests that the prior art submissions by the volunteers did not help with reducing the number of patents, some of which probably should not have been issued. From the previous section, we know that the treatment effect on allowance is marginally negative. Our interpretation is that the crowd-sourced prior art itself was not likely to be helpful, but the program had some effect on the examiners, which led them to reject applications with a marginally higher probability.

Second, the quantity of prior art that was crowd-sourced had no effect on the number of RCEs the applicant filed during the examination process. This implies that the additional art found by the public at the outset was not effective in facilitating the negotiation between the applicant and the PTO examiner, as applicants continued to file RCEs as before. This problem is more serious because at the overall treatment level, program participation increased the number of RCEs by an average of 1.5 RCEs per treated application. One hypothesis is that participating in the open peer review gave the applicants a false ground to pursue negotiation after "final" rejection on the belief that the invention is worthwhile given the community review.

Third, we find no evidence that the peer review had a statistically significant direct effect on forward citations, which is consistent with the insignificant treatment effect on forward citations. One possible reason is that the pilot did not guarantee the novelty of patent claims, so neither the fact that an application was subjected to the community review nor the quantity of prior art or discussion received by an application received any special attention from the inventor community or the PTO examiners.

Table 6 breaks down the effect by the type of activities, which suggests that the baseline results in Table 5 are robust. For instance, the number of non-patent art did not have any effect on allowance, while the expectation at the outset was that the public review would be particularly helpful by pointing to relevant non-patent art in subject matters (computer technologies) where such references may not be readily available to examiners. Similarly, the number of discussions provided by each community did not have any effect on allowance, suggesting that they were likely to be irrelevant to the patent prosecution. These results are quite unexpected. Thus, we now turn to the behavior of the examiners, which might provide an answer to this puzzle.

5 Treatment Effects on Examiner Burden

A significant driver of the Peer-to-Patent initiative was the potential to lighten the burden of examiners and the PTO. The aim was to provide examiners with more relevant prior art and then to resolve examination faster as there would be less need for original searches by examiners for prior art, and fewer additional examinations requested by the applicant. In this section, we investigate the effects of the pilot on patent examiners. That is, we return to the treatment effects framework of Section 3, where we use both treatment and control groups to identify the causal effects of the pilot.

As mentioned previously, the existing reviews of the Peer-to-Patent program focused on whether the crowd-sourced prior art helped examiners make rejections in the first office action, and did not look beyond the first office action. Given the results in this paper that the pilot program increased the number of RCEs, the pilot might have imposed additional burdens on the examiner's workload. Notice that the crowd-sourced prior art might well substitute for the examiner's search effort; that is, the PTO examiners could have found the same relevant prior art had the application not been subjected to the peer review. Thus, this is a benefit of the program because it reduces the examiner's workload up to the first office action. This initial benefit, however, need to be traded off against potential cost at a later stage.

In the data available it is not possible to tell apart which prior art was identified by the examiner only and which was found by both the crowd and the examiner, because the examiner's search report does not contain such information. We suspect that there is at least some degree of substitution of the crowd-sourced prior art for the examiner's search efforts, but we cannot precisely show the extent of this initial benefit.[18] Hence, our focus is on the examiner's search effort beyond the first office action. If the pilot program was indeed helpful for reducing the examiner's workload per application, the treatment effect should be negative or at least neutral with respect to the examiner's search effort, as proxied by the number of searches conducted and the number of new references cited in those additional search reports.

To be more precise, our first measure of examiner effort is the number of PTO form 892 associated with each application, on which the examiner must list all prior art (patent and non-patent) references that have been applied in making a decision. The examiner's first search report almost always appears before the first office action and new reports are added in subsequent rounds of continued examination. The total number of search reports

[18] The crowding-out of public provision by private provision has been documented in alternative contexts. For example, Peltzman (1973) and Becker and Lindsay (1994) show that private donations to higher education institutions resulted in significant reductions in government funding.

can thus proxy the examiner's additional search effort beyond the first office action. Our second and third measure of examiner effort are the number of additional references cited by the examiner before the first office action (for which there is typically one 892 report), and during the rest of the prosecution period.

If the pilot caused the examiners to search more, before or after the first office action, then we will see a positive treatment effect on the number of references cited in form 892, before or after the first office action. We already know from Section 3 that the pilot caused the number of RCEs to increase, so we can expect that the number of examiner search reports, which precedes final or non-final rejections, will likely increase. If we view the number of form 892's as the extensive margin of the examiner's search effort, then the number of newly identified prior art (i.e., net of any duplicates) contained in the examiner's subsequent search reports can be viewed as the intensive margin of the examiner's search effort.

The first stage regression is the same as that reported in Table 2. The second-stage estimation results using these three outcomes are presented in Table 7. The 2SLS estimate of the treatment effect on the number of examiner search reports shows that examiners had to search for prior art on average 2.4 times more than they would, had the application not participated in the pilot (which is statistically significant at the 1% level). Combined with the fact that applications in the pilot had an average of 1.5 more RCEs, this implies that the treatment increased the number of examiner searches by almost one during the normal course of office action, that is, before reaching the point at which an applicant would file an RCE, and this should be considered an extra burden of the pilot.

Looking at the number of prior art references cited before the first office action, the OLS estimate for the treatment effect is significantly positive while the 2SLS estimate is marginally significant at the 10% level. Thus, more references were cited by the examiners prior to the first action, which implies that the crowd-sourced prior art may not have fully substituted for the examiner's search. This is because either the examiner may not have found some of the crowd-sourced prior art that was relevant to the decision, or the examiner, being conscious of the community review, in fact searched harder for more prior art. In the former, there is a less than full substitution between the public and the examiner search while a full substitution is still a possibility with the latter.[19]

Finally, we look at the number of references to prior art that examiners had to find after the first office action, that is, after the crowd-sourced prior art has been accounted for. The OLS estimate for the treatment effect is less than one and negative (and statistically significant at the 5% level), but the 2SLS estimate is almost six and it is positive and statistically significant at the 5% level. This seems consistent with the previous findings and interpretations in that the pilot led to a higher burden to search on the part of examiners overall, and the large point estimate for additional art indicates that the pilot did not effectively identify as much relevant prior art as desired. Overall, the pilot increased the examiner's search effort both in terms of extensive and intensive margins.

6 Conclusion

We have investigated the causal effects of Peer-to-Patent pilot program using a rich set of administrative data. Unlike the hope that crowd-sourcing prior art would assist patent

[19] As previously alluded, we cannot easily distinguish these two cases, but at least the treatment effect is not negative, which suggests that the crowd-sourced prior art did not more than substitute for the examiner's search prior to the first action.

examiners to improve patent quality and reduce the backlog, our findings render insufficient evidence that the pilot provided a net benefit to the PTO or indeed the applicant. The pilot seems to have imposed a substantial cost after the first office action, requiring more searching by examiners and adding an additional burden onto the PTO in the form of additional RCEs from applicants who were no more successful than other applicants—suggesting the pilot did not weed out low quality patents any more than the normal examination process. Thus, our analysis reveals a number of unintended and costly consequences of the pilot.

The idea behind a direct public contribution to patent examination and the opportunities presented by crowd-sourcing deserves, and is receiving, more attention. The main obstacle to these programs appear to be the ability to scale up and encourage a sufficient amount of qualified experts to participate. Similar crowd-sourced patent pilots from Australia (Fitzgerald et al., 2010), Japan (Allen et al., 2012), and the UK (IPO, 2012) have seen low numbers of both volunteering patent applicants and expert participation. While the incentivisation of participation is left for future research, we note that the results of our analysis could be different if the peer review program could elicit a large number of qualified reviewers, although the additional prior art found through the pilot did not help examiners after the first office action.

Our results do not imply that the idea of public input to patent administration is not desirable. What the results show is that a well-considered pilot, with resources to identify experts, and an incentive for participation by applicants, delivered few short term benefits and several unexpected long term costs. We believe this result is important for patent offices around the world considering the value of crowd-sourced patent examination to complement the traditional expert review. For offices that have the scale of the USPTO, with more than 8,000 examiners specialising in subsets of technology, crowd-sourcing prior art does not appear to be of assistance to the office or the applicant.

Crowd-sourcing to improve government administration holds a lot of promise for gathering more information, or additional expertise, but the question is whether decision makers have the most efficient amount of information. In the case of examining patents, the additional information elicited by the crowd led to more noise for applicants who were more likely to pursue their applications which they would not have otherwise, and the examiner who had to undertake additional searches and reviews. As such, it appears to us that in this instance, the pilot was not successful and the PTO were wise to stop it, even if it was for budgetary reasons, and not perhaps the evidence on the pilot's performance.

References

[1] Alcácer, J., and M. Gittelman. 2006. "Patent Citations as a Measure of Knowledge Flows: The Influence of Examiner Citations." *Review of Economics and Statistics* 88: 774–779.

[2] Allen, N., J. Ingham, B. Johnson, J. Merante, B. Noveck, W. Stock, Y. Tham, M. Webbink, and C. Wong. 2008. "*Peer to Patent: First Anniversary Report.*" New York: The Center for Patent Innovations, New York Law School, http://www.peertopatent.org/wp-content/uploads/sites/2/2013/11/P2Panniversaryreport.pdf [accessed January 2017].

[3] ———, A. Casillas, J. Deveau-Rosen, J. Kreps, T. Lemmo, J. Mer-

ante, M. Murphy, K. Osowski, M. Webbink, and C. Wong. 2009. "*Peer to Patent: Second Anniversary Report.*" New York: The Center for Patent Innovations, New York Law School, http://www.peertopatent.org/wp-content/uploads/sites/2/2013/11/CPI_P2P_YearTwo_lo.pdf [accessed January 2017].

[4] ———, A. Casillas, S. Chichetti, M. DeFrances, T. Kabir, C. Segro, and M. Webbink. 2012. "*Peer to Patent: First Pilot Final Results.*" New York: The Center for Patent Innovations, New York Law School, http://dl.dropbox.com/u/2541719/First%20Pilot%20Final%20Results.pdf [January 2017].

[5] Becker, E., and C. Lindsay. 1994. "Does the Government Free Ride?" *Journal of Law and Economics* 37: 277–296.

[6] Belenzon, S. 2012. "Cumulative Innovation and Market Value: Evidence from Patent Citations." *Economic Journal* 122: 265–285.

[7] ———, and M. Schankerman. 2015. "Motivation and Sorting of Human Capital in Open Innovation." *Strategic Management Journal* 36: 795–820.

[8] Bessen, J. 2008. "The Value of U.S. Patents by Owner and Patent Characteristics." *Research Policy* 37: 932–945.

[9] Bestor, D., and E. Hamp. 2010. "Peer to Patent: A Cure for Our Ailing Patent Examination System." *Northwestern Journal of Technology and Intellectual Property* 9: 16–28.

[10] Caillaud, B., and A. Duchêne. 2011. "Patent Office in Innovation Policy: Nobody's Perfect." *International Journal of Industrial Organization* 29: 242–252.

[11] Carley, M., D. Hegde, and A. Marco. 2015. "What is the Probability of Receiving a U.S. Patent?" *Yale Journal of Law and Technology* 17: 203–223.

[12] Farrell, J., and C. Shapiro. 2008. "How Strong Are Weak Patents?" *American Economic Review* 98:1347–1369.

[13] Farre-Mensa, J., D. Hegde, and A. Ljungqvist. 2017. "What is a patent worth? Evidence from the U.S. Patent 'Lottery'." NBER Working Paper 23268.

[14] Federal Trade Commission. 2003. "*To Promote Innovation: The Proper Balance of Competition and Patent Law and Policy.*" Washington, DC.

[15] Fitzgerald, B., B. McEniery, and J. Ti. 2010. "*Peer-to-Patent Australia: First Anniversary Report.*" Brisbane: Queensland University of Technology, Law Department, http://eprints.qut.edu.au/39350/1/39350_P2PAU_1st_Anniversary_Report.pdf [accessed January 2017].

[16] Frakes, M., and M. Wasserman. 2016. "Is the Time Allocated to Review Patent Applications Inducing Examiners to Grant Invalid Patents? Evidence from Micro-Level Application Data." *Review of Economics and Statistics*, forthcoming.

[17] Furman, J., and S. Stern. 2011. "Climbing atop the Shoulders of Giants: The Impact of Institutions on Cumulative Research." *American Economic Review* 101: 1933–1963.

[18] Gaudry, K. 2012. "The Lone Inventor: Low Success Rates and Common Errors Associated with Pro-Se Patent Applications." *PLoS ONE* 7: e33141.

[19] Ghafaele, R., and B. Gibert. 2011. "Crowdsourcing Patent Application Review: Leveraging New Opportunities to Capitalise on Innovation?" *Intellectual Property Quarterly* 3: 23–33.

[20] Grimmelikhuijsen, S., and M. Feeney. 2016. "Developing and Testing an Integrative Framework for Open Government Adoption in Local Governments." *Public Administration Review*, forthcoming.

[21] Hall, B., S. Graham, D. Harhoff, and D. Mowery. 2004. "Prospects for Improving U.S. Patent Quality via Postgrant Opposition." *Innovation Policy and the Economy* 4: 115–143.

[22] ———, A. Jaffe, and M. Trajtenberg. 2005. "Market Value and Patent Citations." *RAND Journal of Economics* 36: 16–38.

[23] Harhoff, D., F. Narin, F. Scherer, and K. Vopel. 1999. "Citation Frequency and the Value of Patented Inventions." *Review of Economics and Statistics* 81: 511–515.

[24] Henderson, R., A. Jaffe, and M. Trajtenberg. 2005. "Patent Citations and the Geography of Knowledge Spillovers: A Reassessment: Comment" *American Economic Review* 95: 461–464.

[25] Intellectual Property Office. 2012. "*Peer to Patent – Pilot.*" London, England.

[26] Jaffe, A., M. Trajtenberg, and R. Henderson. 1993. "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations." *Quarterly Journal of Economics* 108: 577–598.

[27] ———, and M. Trajtenberg. 2002. *Patents, Citations, and Innovations: A Window on the Knowledge Economy*. MA: MIT Press.

[28] ———, and J. Lerner. 2004. *Innovation and Its Discontents: How Our Broken Patent System Is Endangering Innovation and Progress, and What to Do About It*. Princeton, NJ: Princeton University Press.

[29] Johnson, J. 2002. "Open Source Software: Private Provision of a Public Good." *Journal of Economics & Management Strategy* 11: 637–662.

[30] Kornberger, M., R. Meyer, C. Brandtner, and M. Höllerer. 2017. "When Bureaucracy Meets the Crowd: Studying "Open Government" in the Vienna City Administration." *Organization Studies* 38: 179–200.

[31] Lanjouw, J., and M. Schankerman. 2001. "Characteristics of Patent Litigation: A Window on Competition." *RAND Journal of Economics* 32: 129–151.

[32] ———, and ———. 2004. "Patent Quality and Research Productivity: Measuring Innovation with Multiple Indicators." *Economic Journal* 114: 441–465.

[33] Lei, Z., and B. Wright. 2017. "Why Weak Patents? Testing the Examiner Ignorance Hypothesis." *Journal of Public Economics* 148: 43–56.

[34] Lemley, M. 2001. "Rational Ignorance at the Patent Office." *Northwestern University Law Review* 95: 1497–1532.

[35] ———, and B. Sampat. 2012. "Examiner Characteristics and Patent Office Outcomes." *Review of Economics and Statistics* 94: 817–827.

[36] Lerner, J., and J. Tirole. 2002. "Some Simple Economics of Open Source." *Journal of Industrial Economics* 52: 197–234.

[37] ———, and ———. 2005. "The Economics of Technology Sharing: Open Source and Beyond." *Journal of Economic Perspectives* 19: 99–120.

[38] Loiselle, J., M. Lynch, and M. Sherrerd. 2010. "*Evaluation of the Peer to Patent Pilot Program.*" Report Sponsored by the USPTO, Worcester Polytechnic Institute, http://www.wpi.edu/Pubs/E-project/Available/E-project-122109-150816/unrestricted/usptofinalreport.pdf [accessed January 2017].

[39] McCabe, M., and C. Snyder. 2015. "Does Online Availability Increase Citations? Theory and Evidence from a Panel of Economics and Business Journals." *Review of Economics and Statistics* 97: 144–165.

[40] Mitra-Kahn, B., A. Marco, P. Evans, C. Frey, N. Sultan, M. Carley, and P. D'Agostino. 2013. "*Patent Backlogs, Inventories, and Pendency: An International Framework.*" UK-IPO-USPTO Joint Report, http://www.ipo.gov.uk/pro-ipresearch.htm [accessed January 2017].

[41] Noveck, B. 2006. "Peer to Patent: Collective Intelligence, Open Review and Patent Reform." *Harvard Journal of Law & Technology* 20: 123–162.

[42] Peltzman, S. 1973. "The Effect of Government Subsidies-in-Kind on Private Expenditures: The Case of Higher Education." *Journal of Political Economy* 81: 1–27.

[43] Picard, P., and B. de la Potterie. 2013. "Patent Office Governance and Patent Examination Quality." *Journal of Public Economics* 104: 14–25.

[44] Piotrowski, S. 2017. "The "Open Government Reform" Movement: The Case of the Open Government Partnership and U.S. Transparency Policies." *American Review of Public Administration* 47: 155–171.

[45] Sampat, B. 2010. "When Do Applicants Search for Prior Art?" *Journal of Law and Economics* 53: 399–416.

[46] Stock, J., and M. Yogo. 2005. "Testing for Weak Instruments in Linear IV Regression." J.H. Stock and D.W.K. Andrews (eds), *Identification and Inference for Econometric Models: Essays in Honor of Thomas J. Rothenberg*. New York, NY: Cambridge University Press.

[47] Thompson, P., and M. Fox-Kean. 2005. "Patent Citations and the Geography of Knowledge Spillovers: A Reassessment." *American Economic Review* 95: 450–460.

[48] Trajtenberg, M. 1990. "A Penny for Your Quotes: Patent Citations and the Value of Innovations." *RAND Journal of Economics* 21: 172–187.

[49] United States Patent and Trademark Office. 2008. "*Fiscal Year 2009 President's Budget*." Alexandria, VA.

[50] Weatherall, K. 2009. "It's Not Just Competitors: Acknowledging and Accommodating "Interfering Busybodies" and their Challenges to Patent Validity." *Journal of World Intellectual Property* 12: 500–523.

[51] Zhang, X., and F. Zhu. 2011. "Group Size and Incentives to Contribute: A Natural Experiment at Chinese Wikipedia." *American Economic Review* 101: 1601–1615.

Figure 1: The PTO Examination Process

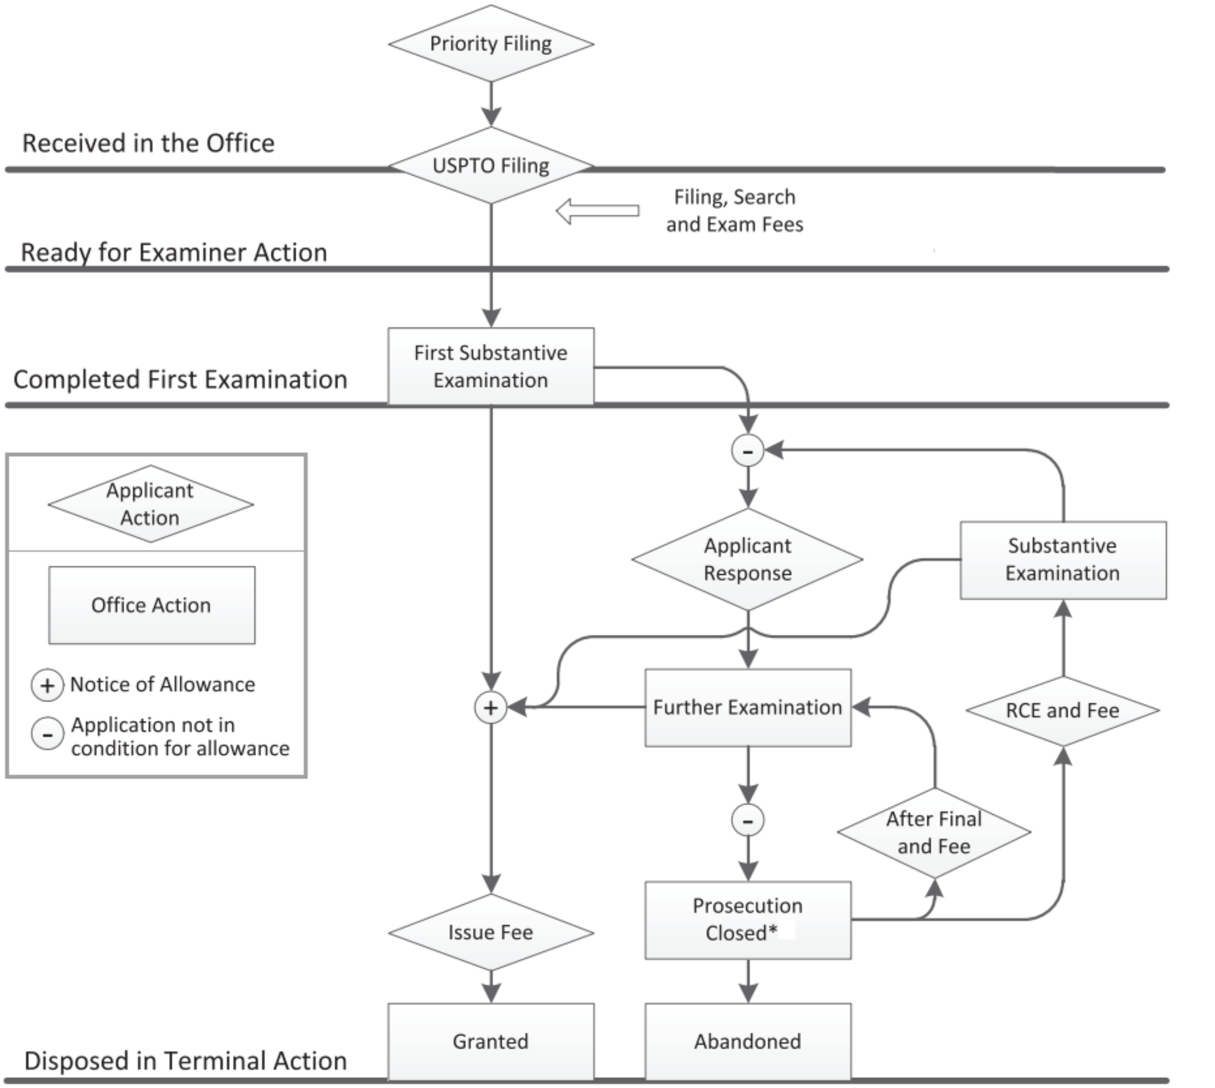

* In the US, once prosecution is closed an applicant may file an after final submission, a Request for Continued Examination (RCE), or an appeal (not pictured). The availability of after final submissions is limited by a 6 month clock which begins when prosecution is closed. The number RCEs is not limited, but, generally, must be filed within 6 months of prosecution closing.

Source: Adapted from Mitra-Kahn et al. (2013)

Table 1. Summary Statistics

Variable	Full Sample		Treatment	Control	Difference
	Mean	Std. dev	Mean	Mean	P-values†
Allowance (1=yes)	0.63	(0.48)	0.66	0.62	0.30
Number of requests for continued examination	0.60	(0.87)	0.57	0.62	0.55
Number of forward citations	14.6	(25.9)	17.6	13.4	0.04
Number of days from application to publication	403	(196)	465	379	0.00
Assistant examiner (1=yes)	0.51	(0.50)	0.39	0.56	0.00
Number of previous patent applications	11.8	(41.2)	13.1	11.3	0.59
Small entity (1=yes)	0.22	(0.41)	0.26	0.20	0.06
Number of independent claims	3.30	(1.95)	3.20	3.34	0.40
Number of dependent claims	18.0	(11.2)	16.9	18.4	0.10
Size of patent family	0.93	(1.63)	0.41	1.13	0.00
Filed with the EPO (1=yes)	0.18	(0.38)	0.08	0.22	0.00
Number of examiner's search reports	2.01	(1.41)	1.93	2.04	0.33
Number of references cited before first office action	7.09	(6.60)	10.0	5.95	0.00
Number of references added after first office action	3.27	(5.88)	2.75	3.48	0.12
Size of P2P community	5.30	(5.65)	5.30		
Number of P2P prior art submissions	2.67	(2.48)	2.67		
of which patent art	1.27	(1.74)	1.27		
of which non-patent art	1.40	(1.69)	1.40		
Number of P2P discussions	3.34	(6.39)	3.34		
Number of observations	788		221	567	

†P-values are shown based on two-tailed tests of the equality of means or proportions

Table 2. First-stage Regressions for Treatment

Variable	Participation	
	Mean	Std. dev.
Days from application to publication	.0004***	(.0001)
Assistant examiner	-.1365***	(.0336)
No. of previous applications	.0003	(.0004)
Small entity	.0735	(.0439)
No. of independent claims	.0067	(.0109)
No. of dependent claims	-.0038**	(.0016)
Size of patent family	-.0320***	(.0090)
EPO filing	-.0609	(.0382)
Patent Class dummies	Yes	
Technology Center dummies	Yes	
F statistic of excluded instrument	10.60	
Cragg-Donald Wald F statistic	23.96	
N	788	

Standard errors are clustered at the three-digit class level and reported in the parenthesis.
**P<0.10, **P<0.05, ***P<0.01*

Table 3. Treatment Effect on Applications

Variable	Allowance		Number of RCEs		Forward citations	
	OLS	IV	OLS	IV	OLS	IV
P2P treatment	.0487*	-.3528*	-.0252	1.472***	4.573	7.939
	(.0260)	(.1857)	(.0935)	(.3447)	(3.058)	(14.14)
Assistant examiner	-.1784***	-.2331***	-.1083*	.0955	-1.724	-1.265
	(.0235)	(.0369)	(.0562)	(.0900)	(1.668)	(3.254)
No. of previous applications	-.0001	-.0001	-.0002	-.0005	.0259	.0253
	(.0004)	(.0005)	(.0004)	(.0004)	(.0434)	(.0435)
Small entity	-.1363***	-.1207***	-.2875***	-.3457***	7.825*	7.694*
	(.0424)	(.0372)	(.0380)	(.0653)	(4.414)	(4.658)
No. of independent claims	.0096	.0116	.0282	.0209	1.227	1.210
	(.0076)	(.0097)	(.0241)	(.0267)	(.7364)	(.7580)
No. of dependent claims	.0022	.0011	.0054***	.0095***	.1863***	.1955***
	(.0017)	(.0018)	(.0015)	(.0029)	(.0593)	(.0669)
Size of patent family	.0020	-.0162	.0241	.0919**	.0302	.1827
	(.0174)	(.0189)	(.0305)	(.0414)	(.6823)	(.8059)
EPO filing	.0227	.0042	-.0571	.0119	1.971	2.126
	(.0669)	(.0681)	(.1309)	(.1293)	(3.337)	(3.501)
Patent Class dummies	Yes	Yes	Yes	Yes	Yes	Yes
Technology Center dummies	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.175		0.072		0.074	
N	788	788	788	788	788	788

Standard errors are clustered at the three-digit class level and reported in the parenthesis.
**P<0.10, **P<0.05, ***P<0.01*

Table 4. First-stage Regressions for P2P Activities

Variable	Number of prior art	Number of patent art	Number of non-patent art	Number of discussions
Size of P2P community	.2546***	.1072***	.1474***	1.007***
	(.0333)	(.0331)	(.0201)	(.0751)
Assistant examiner	-.3993	-.2900	-.1092	.3222
	(.2309)	(.2074)	(.2296)	(.4270)
No. of previous applications	-.0001	-.0020	.0019	-.0017
	(.0035)	(.0022)	(.0032)	(.0037)
Small entity	.1293	.1196	.0098	.3376
	(.1441)	(.1531)	(.1591)	(.6279)
No. of independent claims	.0012	.0781*	-.0770**	.2199*
	(.0364)	(.0442)	(.0316)	(.1142)
No. of dependent claims	.0365***	.0224*	.0141	-.1078***
	(.0100)	(.0123)	(.0130)	(.0272)
Size of patent family	-.0604	-.2354	.1750	.6394**
	(.1631)	(.1976)	(.1249)	(.2479)
EPO filing	-1.159**	-.0016	-1.158**	-1.039
	(.4396)	(.4596)	(.4253)	(.7649)
Patent Class dummies	Yes	Yes	Yes	Yes
Technology Center dummies	Yes	Yes	Yes	Yes
F statistic of excluded instrument	58.47	10.50	53.96	180.1
Cragg-Donald Wald F statistic	109.60	24.97	63.69	568.7
N	221	221	221	221

Standard errors are clustered at the three-digit class level and reported in the parenthesis.
**P<0.10, **P<0.05, ***P<0.01*

Table 5. Effects of P2P Prior Art on Applications

Variable	Allowance		Number of RCEs		Forward citations	
	OLS	IV	OLS	IV	OLS	IV
Number of prior art	-.0306**	-.0288	.0242	-.0099	1.416	3.599
	(.0124)	(.0240)	(.0279)	(.0322)	(1.022)	(2.711)
Assistant examiner	-.1232	-.1228*	-.3205***	-.3293***	-7.698	-7.134
	(.0757)	(.0701)	(.0982)	(.0886)	(5.138)	(4.875)
No. of previous applications	.0017***	.0017***	.0010	.0011	.1114	.1084
	(.0005)	(.0004)	(.0015)	(.0013)	(.1519)	(.1354)
Small entity	-.0002	.0001	-.3646***	-.3703***	19.61	19.97
	(.0730)	(.0657)	(.0493)	(.0459)	(16.56)	(15.28)
No. of independent claims	.0193	.0194*	-.0217	-.0223	2.340	2.380
	(.0124)	(.0114)	(.0311)	(.0292)	(2.709)	(2.441)
No. of dependent claims	.0047	.0046	-.0042	-.0027	1.215	1.116*
	(.0074)	(.0072)	(.0121)	(.0110)	(.7232)	(.5902)
Size of patent family	.0167	.0165	.0995	.1025*	-4.476*	-4.665**
	(.0780)	(.0711)	(.0654)	(.0591)	(2.204)	(2.378)
EPO filing	-.0014	.0015	.1285	.0727	11.07	14.65
	(.2244)	(.2107)	(.2158)	(.2212)	(11.98)	(11.68)
Patent Class dummies	Yes	Yes	Yes	Yes	Yes	Yes
Technology Center dummies	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.222		0.185		0.206	
N	221	221	221	221	221	221

Standard errors are clustered at the three-digit class level and reported in the parenthesis.
**P<0.10, **P<0.05, ***P<0.01*

Table 6. Effects of P2P Activities on Applications

Specification	Allowance		Number of RCEs		Forward citations	
	OLS	IV	OLS	IV	OLS	IV
Number of prior art	-.0306**	-.0288	.0242	-.0099	1.416	3.599
	(.0124)	(.0240)	(.0279)	(.0322)	(1.022)	(2.711)
Patent art only	-.0189	-.0685	.0014	-.0235	-1.436	8.551
	(.0140)	(.0671)	(.0409)	(.0783)	(1.026)	(6.563)
Non-patent art only	-.0408**	-.0497	.0478	-.0171	4.535	6.215
	(.0145)	(.0368)	(.0447)	(.0548)	(3.086)	(4.812)
Number of discussions	-.0078*	-.0073	-.0001	-.0025	.4946	.9096
	(.0043)	(.0056)	(.0076)	(.0079)	(.5282)	(.7345)

Each row in this table represents a separate regression as in Table 5. For brevity, this table only collects the point estimate of the coefficient on P2P activity from each regression.
**P<0.10, **P<0.05, ***P<0.01*

Table 7. Treatment Effect on Examiners

Variable	Number of examiner's search reports (PTO form 892)		Number of references cited in 892 before the first office action		Number of references added in 892 after the first office action	
	OLS	IV	OLS	IV	OLS	IV
P2P treatment	-.0896	2.383***	4.096***	4.215*	-.8397**	5.865**
	(.0930)	(.6659)	(.8188)	(2.168)	(.3286)	(2.989)
Assistant examiner	-.0136	.3230**	-.4679	-.4516	-.7525	.1603
	(.0975)	(.1635)	(.4230)	(.4679)	(.4984)	(.6745)
No. of previous applications	-.0002	-.0007	-.0048*	-.0048**	-.0028	-.0040
	(.0004)	(.0012)	(.0023)	(.0021)	(.0023)	(.0047)
Small entity	-.3933***	-.4894***	.8147	.8101	-1.070***	-1.331**
	(.0624)	(.1364)	(.6582)	(.6577)	(.3626)	(.5693)
No. of independent claims	.0756*	.0637	.0270	.0264	.1950	.1625
	(.0407)	(.0423)	(.0698)	(.0626)	(.1742)	(.1730)
No. of dependent claims	.0060**	.0128***	.0356*	.0359**	.0090	.0273
	(.0023)	(.0032)	(.0182)	(.0179)	(.0202)	(.0213)
Size of patent family	-.0100	.1020*	.1181	.1236	-.0424	.2613
	(.0501)	(.0564)	(.2223)	(.1783)	(.1911)	(.2268)
EPO filing	.0917	.2057	-.7273	-.7218	.4063	.7155
	(.1795)	(.1852)	(.7785)	(.7849)	(.7710)	(.8201)
Patent Class dummies	Yes	Yes	Yes	Yes	Yes	Yes
Technology Center dummies	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.086		0.118		0.084	
N	788	788	788	788	788	788

Standard errors are clustered at the three-digit class level and reported in the parenthesis.
**P<0.10, **P<0.05, ***P<0.01*