Peer Reviewed Patent Applications: Evidence from a Pilot Program

Jin-Hyuk Kim[*] Benjamin Mitra-Kahn[†]

Abstract

In June 2007, the United States Patent and Trademark Office launched a pilot program designed to solicit public participation in the prior art search process. Specifically, citizen-experts could identify and evaluate prior art relevant to the claims of a published patent application before an examiner reviews it. This paper assesses whether the peer review pilot achieved the goal of improving examination quality using data generated from the first pilot. In agreement with the available qualitative evaluations, we find the expected first order effect, that more prior art does indeed lower the probability of patent allowances; and does so by 5% for every contributed piece of relevant prior art. On the other hand, the second order effects are significant and, rather surprisingly, work in the opposite direction to the first order effect. Further, while reviewer-submitted prior art has no impact on forward citations, participation in the pilot substantially increased the number of these citation counts. Finally, using

[*]University of Colorado at Boulder. Email: JinHyuk.Kim@colorado.edu
[†]IP Australia, Email: Benjamin.Mitra-Kahn@ipaustralia.gov.au

references cited in the examiner's search reports, we interpret these findings as indicating that participation in open peer review had a positive signaling effect on patent prosecution.

Keywords: peer review; patent applications; open government; patent office

JEL classification: D04, O34, O38

Disclaimer: The views and opinions expressed in this paper are those of the authors and do not reflect those of any private or government agencies, including IP Australia.

1 Introduction

The United States Patent and Trademark Office (USPTO)'s most important function is to examine whether a patent application is novel and non-obvious in the sense that no prior art disqualifies the application for a patent. By law the patent examiners should make a thorough investigation of the available prior art relating to the subject matter of the claimed invention, but the USPTO has been criticized in the past for granting patents that should have been narrowed or rejected due to existing prior art. This issue has driven a number of changes at the USPTO whose 2012 strategic goal was to "optimize patent quality and timeliness," and it has been part of the rationale for the America Invents Act and was the reason for the original Peer-to-Patent pilot program.[1]

Part of the problem for examiners is the time and resources allocated to undertake prior art searches. A 2003 report by the Federal Trade Commission says that 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" (FTC 2003: 10), while reports for the USPTO suggest examiners have 20 hours on average (Loiselle et al., 2010; Allen et al., 2012: 4). Given that the pending application backlog reached a historic proportion in 2006 (USPTO, 2006), the patent office and examiners came under increasing pressure and criticism, part of which revolved around issuing weak patents.[2] Further, the incentives facing examiners are not always aligned with the goal of maintaining high quality. For instance, examiners

[1] See http://www.uspto.gov/about/stratplan/ar/2011/mda_02_03.html.

[2] Lemley (2001) argues that this concern is lessened if the costs of litigating a far fewer number of all patents are smaller than the increase in costs of more thoroughly examining many patent applications. For a formal analysis, see Farrell and Shapiro (2008).

are rewarded based on the counts of first office actions, while searching for prior art is not explicitly rewarded.[3]

It was under these circumstances that the USPTO launched a pilot program in June 2007, where it hoped that the open review by the public would provide useful prior art in front of examiners before they make the first office action on a patent application. On June 7, 2007, the USPTO released details of the peer review pilot, which would give citizen-experts, for the first time, the opportunity to submit annotated technical references relevant to the claims of a published patent application before an examiner reviews it.[4] The scope of the first pilot was limited to computer technologies and business methods applications where much of the nonpatent literature is not readily accessible to patent examiners.

The effect of open peer review procedure in general seems to be almost unknown presently in the economic discipline. Pöschl (2012) looks at this issue from a hard science perspective, but the economics, and the impact, of open peer review remain relatively unstudied. For different reasons than the USPTO, in academic publishing many science journals have moved towards open peer review process. For instance, in 2006, *Nature* launched a trial of open peer review, where authors of newly submitted papers that survived the initial editorial assessment were invited to post their papers on the Internet for public comment. Just as the USPTO's Peer-to-Patent pilot was not intended to replace the patent examiner's search for prior art, *Nature*'s open-review trial was intended to supplement their existing reviews,

[3]See Jaffe and Lerner (2004) for a detailed study. Lemley and Sampat (2012) empirically show that the increasing counts requirement for satisfactory performance leads senior examiners to cite less art and grant more patents.

[4]The pilot program was conducted on the Peer-to-Patent website developed and managed by Beth Noveck (2006) and law students at New York Law School's Center for Patent Innovations. In May 2009, the peer review pilot was showcased by the White House as an example of the administration's Open Government Initiative.

and so articles were simultaneously subjected to standard peer review.[5]

Despite the significant interest in the *Nature* trial, only a small proportion (5%) of authors opted to participate. Further, only 54% of those participants received any comments despite significant web traffic. Editors did not think the open peer review contributed significantly to their decisions, and discontinued the trial program. In a sense, the USPTO's Peer-to-Patent pilot seems to have followed a similar path. The first pilot had its limit of 400 participating applications; however, in total there were only 226 applications posted on the Peer-to-Patent website.[6] Despite some positive survey responses from the participating examiners about reviewer submitted prior art, the second pilot was discontinued in 2011.[7]

The various evaluations and annual progress reports of the US pilot, and indeed similar evaluations of similar pilot projects in other countries, have mainly been based on surveying the participants on both the public and patent office side.[8] The use of administrative data and information on the eventual patent outcome has been restricted, and when undertaken, has limited itself to the impact of Peer-to-Patent references on first office actions, as in Allen et al. (2008).[9] This paper therefore goes beyond the existing evaluations to empirically

[5]See http://www.nature.com/nature/peerreview/debate/index.html for an overview of the trial and perspectives from leading scientists. However, after the trial *Nature* removed the manuscripts and comments posted on the preprint server from public access, so no data could be obtained.

[6]Low numbers for volunteered applications were observed in similar pilot projects run in other countries: In Australia the first round included less than 40 patent applications (Fitzgerald et al., 2010), the Japanese pilot had 39 volunteers (Allen et al., 2012: 33), while the UK pilot offered 172 applications (IPO, 2012).

[7]The final US evaluation stressed that the Peer-to-Patent program had been beneficial and recommended the USPTO to expand on this work. The report noted that on average the number of prior art references forwarded to examiners was 2.66 per application, and that "38 applications were rejected based on 41 references submitted through Peer To Patent" (Allen et al., 2012: 27). Evaluations of similar programs in other countries suggest there are benefits, but response rates were much lower, as in the UK case where 172 applications elicited observations on 11 applications, six of which were deemed "useful" by examiners (IPO, 2012: 5).

[8]See Allen et al. (2008, 2009, 2012) and Loiselle et al. (2010) for evaluations of the first USPTO pilot; IPO (2012) for the UK pilot; and Fitzgerald et al. (2010) for the Australian Pilot's first year.

[9]The final USPTO evaluation (Allen et al., 2012: 27) says that "for the five applications with Peer To Patent prior art used in a later office action, these references were all initially cited as pertinent," which

investigate what happened to the patent applications that went through peer-to-patent and how they compare to similar patents that were not in the USPTO pilot.

The challenge in estimating a causal effect of open peer review is that such programs are not controlled experiments, but participants self-select into the program. Therefore, a qualitative analysis of survey responses and summary statistics does not prove or disprove the hypothesis that open peer review has any effect on the quality of the examination process or the issued patents. The primary contribution of this paper is the approach used to estimate the causal effect (also called the treatment effect) of the open peer review of patent applications as well as third party submitted prior art (conditional on participation) on such outcome measures as the probability of allowance and the number of citations.

In order to identify the causal effect of a variable, one or more variables are required that affect the variable of interest, but does not directly affect the main outcome equation. The instrument employed in this paper for participation in the pilot program is the applicant's country of residence. As the final USPTO evaluation (Allen et al., 2012) notes, the relatively low level of participation was a result of insufficient familiarity with the program. Although a number of published articles and online blogs increased the program's profile, many of them were US sources. Accordingly, many eligible foreign applicants either did not become aware of the pilot in a timely manner or were reluctant to participate.

On the other hand, the patent laws of the United States make no discrimination with respect to the citizenship of the inventor. Approximately half of US patents in recent years are issued to the residents of foreign countries. In fact, recent work by Dent et al. (2012) sug-

suggests the authors have looked at some of the later impacts of the program, but no formal investigation appears to have been undertaken.

gests that the USPTO examination is indeed neutral to the origin of applicants, in contrast to work by Webster et al. (2011) on the Japanese and European Patent Office which found some national biases in the examination process. Therefore, the evidence seems to support our point that applicant residency should have little direct impact on USPTO applications and hence be a good instrument for the treatment status.

Examining a treatment group of 180 applications and a control group of 462 applications that have finished the prosecution, we show that there is an economically significant (despite relatively imprecise), positive treatment effect on the probability of patent applications being allowed by the USPTO. We also evaluate the causal effect of reviewer submitted prior art using only the set of participating applications. Here, instrumenting with the size of Peer-to-Patent communities for each application to take account of the endogeneity of the number of prior art, we find a negative and statistically significant effect on the probability of allowance in a variety of specifications. On average, the increase in the probability of allowance due to positive treatment effect is more than double the size of the reduction in the probability due to submission of prior art.

Similar contrasting results are obtained for the treatment effect and the prior art effect when the number of forward citations within three years of the publication date is used as the outcome measure. To be precise, the treatment effect of the pilot is to increase the citation counts by, on average, five, and this effect is statistically significant in a number of specifications. On the other hand, as one would expect, there is virtually no statistically or economically significant effect of prior art submissions on the number of forward citations, which is a common measure for the quality of the underlying innovation.[10] Therefore, par-

[10]See Trajtenberg (1990) for the classic reference on using forward citations to indicate quality of the

ticipation in open peer review seems to have a positive signaling effect on the examination and boost quality measures of the innovation, which may been neglected in the USPTO's cost-benefit analysis.

To further examine our preferred interpretation and an alternative hypothesis, we use data from the PTO examiner's search behavior as contained in the PTO form 892. Specifically, we look at whether the higher likelihood of allowing treatment applications is due to the substitution from the reviewer submitted prior art. Both in terms of number of examiner search and number of references cited by the examiners, we do not find evidence supporting this hypothesis. That is, examiners search more, if anything, and the number of references cited before the first office action is significantly higher for the treatment applications. Further, the number of new references cited after the first office action is also significantly higher by the same order of magnitude, which would not be observed if peer review and examination were substitutes.

This paper is related to the growing literature that applies a quasi-experimental design approach to the citation analysis to account for the knowledge spillovers. For instance, Jaffe et al. (1993) construct a control sample of patents based on the classification and the grant date to account for the geographic localization of knowledge spillovers. Furman and Stern (2011) construct a control sample of the most-related articles published in the same journal and publication year to assess the impact of biomaterial deposit on follow-on research. Our paper is similar in spirit to these papers. The novel contribution of our paper is to analyze the effects of open peer review on patent applications using a number of measures. Our

innovation.

paper can also offer insights into the effect of open peer review in academic publishing.[11]

The outline of this paper is as follows. Section 2 summarizes the data set used in the analysis. Section 3 presents the estimates of the treatment effect of the pilot program, and Section 4 the estimates of the first-order effect of prior art submission. Section 5 looks at the issue of whether peer review and examiner search are substitutes. Section 6 contains concluding remarks.

2 The Data

The data used comprise a stratified random sample of 462 patent applications filed at the USPTO ("control group") and a total of 180 patent applications posted to the Peer-to-Patent website ("treatment group"). These are all utility patent applications, and the random sample was drawn from each primary classification that were eligible to participate in the first pilot program. Specifically, only applications assigned to three Technology Centers (TC) were eligible. These were Computer Architecture (TC 2100), Computer Networks & Cryptography and Security (TC 2400) and Business Methods (TC 3600). And within these, only a section of classes were in turn eligible for the pilot.[12] Furthermore, the timing of the publication of eligible applications had to occur during the pilot, which ran from June 15,

[11] In science journals, open peer review process is getting some traction. For instance, *Atmospheric Chemistry and Physics* was launched in 2001 by the European Geosciences Union, and by 2005 it has become the highest impact journal in the category "Meteorology & Atmospheric Sciences" in the annual Journal Citation Reports. In economics, the journal *Economics: The Open-Access, Open-Assessment E-Journal* was launched in 2007 following the similar open peer review process wherein interested researchers can comment on the manuscript evaluation.

[12] For TC 2100 these are classes 380, 700, 703, 706, 707, 708, 710, 711, 712, 714, 715, 717, 718, 719; for TC 2400 they are 380, 709, 713, 726 and for TC 3600 only Class 705 was eligible.

2007, to June 15, 2009.[13]

We collected information on the status of each application in the control and the treatment groups (current as of March 2013) from the USPTO's public Patent Application Information Retrieval (PAIR) database. A nontrivial fraction of these applications are still being prosecuted by the USPTO. For instance, of the 226 applications subjected to the peer review, 46 applications (20%) were still pending.[14] Further, the three-digit eligible classes have a number of subclasses, some of which are assigned to Group Art Units (GAUs) that belong to other Technology Centers than 2100, 2400, or 3600. Since such applications were not eligible for the pilot, we eliminated them from our sample.[15] For the eligible 642 applications that have been prosecuted, we assigned 1 to a dummy variable, *Allowance*, if an application was allowed, and 0 if it was abandoned.[16]

Another outcome variable of interest is the number of (forward) citations, meaning the number of times a publication is referenced by later patent applications or patents.[17] Notice that this measure is distinct from the list of references cited in a given patent document, which are usually referred to as backward citations.[18] In our context, we do not observe whether

[13] The pilot was initially restricted to patent applications in the computer-related arts and was scheduled to last for one year. However, in July 2008 the USPTO extended the pilot for another year and expanded the scope of the program to include the automated business data processing technologies (class 705). Thus, the sampling period for class 705 is from July 16, 2008, to June 15, 2009.

[14] For this reason, we do not analyze data available from the second pilot program, which ran from October 25, 2010 to September 30, 2011. The second pilot also expanded the scope of the program beyond computer technologies and business methods applications.

[15] Essenatially, by imposing two eligibility criteria (i.e., three-digit classes and Technology Centers), the application process made only a subset of those three-digit classes eligible for the pilot. Notice that eligible subclasses cannot be used as an instrument for the pilot participation because an ineligible subclass (i.e., an eligible class that was not assigned to TC2100, TC2400, or TC3600) is a perfect predictor of nonparticipation.

[16] A few applications were abandoned after allowance for failure to pay the issue fee. Because we are interested in the examination, we focus on the notification of patent allowance rather than patent grant.

[17] There is considerable evidence supporting the usefulness of patent citations as a measure of the importance or value of innovations. See, e.g., Trajtenberg (1990), Harhoff et al. (1999), and Hall et al. (2005).

[18] Recently, Alcácer and Gittelman (2006) and Sampat (2010) use the share of references provided by the applicant and the examiner, respectively, in U.S. patent documents to assess how intensely applicants or

any references provided by the Peer-to-Patent community were also independently found by the examiner. Thus, we cannot readily use backward citation in patent documents to isolate the effect of peer reviewed prior art. From searching on Google Patent, we measured the *Number of Citations* for each application as the number of references made by other patent applications (plus those of patents, if granted, and net of any duplicate counts) within the first three years of the publication date.[19]

We also hand-collected data on the examiner status, the first named inventor's characteristics, the filing status of application, the number of claims, and the patent family from such sources as Google Patent and the EPO's Espacenet as well as the USPTO database. As Lemley and Sampat (2012) show, the examiner's tenure is an important control variable that could affect the probability of allowance.[20] If a patent is granted, then the publication of the patent indicates whether the examiner is a primary or an assistant examiner (the *Examiner Assistant* dummy variable), who does and does not have the authority to dispose applications assigned to them, respectively. If an application is abandoned, we searched the USPTO Full-Text database to find if the examiner had issued any patent as a primary or had acted only as an assistant at the time when the said application was abandoned.

Similarly, from the PAIR database we collected data on the number of the first named

examiners search for prior art. They show that the examiners are responsible for 40-60% of citations on the average patent.

[19] Published patent applications are the predominant prior art documents. 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). By restricting attention to the first three years, we focus on the relatively short-term impact of patent applications.

[20] Potential selection bias due to the assignment process to junior and senior examiners is unlikely to be a concern here because every nonprovisional application must be classified and assigned to an examiner at least several weeks before the publication of the application, which was the timing criterion for the pilot. The initial assignment may be transferred and disputed, but Lemley and Sampat (2012) plausibly argue that the assignment process is not subject to strategic manipulation.

inventor's previous patent applications, which serves as a proxy for the inventor's experience.[21] From the PAIR database, the inventor's country of residence is readily available as well as whether the *Small Entity* status was claimed during the examination process.[22] The main benefit to small entity status is that the fees are reduced by 50 percent, and we control for their presence in line with similar work on USPTO data (e.g., Bessen, 2008). The main reason for this is that fee-conscious applicants may be less likely to hire a patent attorney who can help with searching for prior art and writing patent claims, and we know that un-represented applicants are more likely to abandon applications; and they file lower quality applications on average, with a lower probability of allowance (Gaudry, 2012).

It is straightforward to collect data on the number of independent and dependent claims from the published applications, which many researchers use as a standard control in the patent literature (e.g., Lanjouw and Schankerman 2004).[23] To control for the patent-worthiness of an application, we also use information on patent family. A patent family is a set of related applications filed in multiple countries to protect a common invention. In this paper, we measure the *Size of Patent Family* as the number of foreign countries in which related applications are filed.[24] However, one concern is that not every national patent office applies the same patentability standard. Following Lemley and Sampat (2012),

[21] As a result of the 1999 American Inventors Protection Act, U.S. patent applications are routinely published 18 months after they were filed. Thus, only those applications filed after 2001 are available for electronic searching on the USPTO website.

[22] This is found in Fee Worksheet (SB06) or Issue Fee Payment (PTO-85B) under 'Image File Wrapper' in the PAIR database. A small entity means any small business, independent inventor, and nonprofit organization as defined by relevant regulations.

[23] While including too many claims may increase the chance of multiple rejections, dependent claims tend to narrow down the scope of the independent claims. Given the heterogeneity of claimed inventions, it is a priori uncertain whether there exists an empirical correlation between the number of claims and the quality of applications.

[24] Specifically, we searched the Espacenet database to find the set of documents having the same priority as those in our sample. 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.

we add a dummy variable indicating whether the same or related applications were filed at the European Patent Office (EPO).

For those applications in the treatment group, we collected additional information on the size of reviewing community (i.e., the number of registered reviewers), the number of prior art references posted by citizen-experts as well as the number of discussion postings for each application. Additionally, total number of prior art references are broken into patent references and non-patent prior art. To be clear, the pilot allowed at most ten best prior art references and only comments concerning the relevance of the submitted document to be forwarded to the examiners for their consideration.[25] However, the examiner survey indicates that the examiners were free to read all comments posted on the Peer-to-Patent website to see how the community reviewers interpret submitted references (Allen et al., 2012), so we include these measures.

Finally, examiner's search behavior is measured by two variables. One is how often the examiner searches for prior art, and the other is how may references are found in each reported search. Both measures can be obtained from looking at PTO form 892 ("list of references cited by examiner") documented in the PAIR database for each application. We further refine the number of references appearing on the examiner's search reports by the type of prior art (patent and non-patent art) and the timing of 892 posting (before and after the first office action). Notice that all prior art submission from the Peer-to-Patent website has to occur before the first office action. Examiner's search reports typically contain duplicate references. We only count nonduplicate references contained in each search report

[25] Thus, when the number of reviewer identified prior art exceeded 10, we recorded the number of submissions as 10; however, this constraint binds only for three cases in our data set. Relaxing this constraint does not affect our findings.

to capture the extent of incremental search efforts.

Summary statistics are presented in Table 1. Overall, 66% of the applications were allowed, and the remaining 34% were abandoned. There is a marginally significant difference in these proportions between the treatment and the control group. In general, there exists sizable variation in the variables used to control for observed heterogeneity. Test of mean difference shows that the share of applications assigned to assistant examiners is higher for the control group (59%) than for the treatment group (44%). Inventors in the control group are marginally less experienced and tend to be foreign residents. The control group also has a higher proportion of related applications filed at foreign patent offices compared to the treatment group. The average number of forward citations is 1.71, but there is a large variation across applications. In terms of both patent and non-patent literature, the treatment group has a larger number of references cited by examiners during the examination process.

3 Estimating Treatment Effect

The challenge in estimating a program treatment effect is to overcome the omitted variable bias. That is, the patent applications that went through the open review process are self-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 a suitable instrument to plausibly estimate the treatment effect. That is, the simple comparison of means in Table 1, where the pilot appears to increase the probability of allowance, could be a misleading estimate of causal effects and should not drive any conclusions as to the effects of the program.

Let us briefly examine potential sources of selection bias. The main benefits of the pilot are that the application is advanced out of turn for the first office action and additional prior art references will be made of record on any patent issuing from the application. Because the reviewer-submitted references might be material to patentability, those who file weak applications may have little incentive to participate in open review process. Further, higher quality patents (proxied by forward citations, renewals and other metrics) tend to have a higher value to firms, so those who have a higher chance of getting a patent are more likely to benefit from the advancement out of turn.[26] While this is plausible, a careful collection of the relevant facts suggests that this selection hypothesis should not be a major driving force for participation in the program.

To be precise, participating applications were not granted Accelerated Examination status. That is, beyond the first office action the examination process for the treatment applications was performed in a normal manner following established methods. For those who wanted to have their applications placed on an accelerated track throughout the entire prosecution, such an option was available at a moderate fee. On the other hand, Table 1 shows that the treatment group in fact comprises applications that have on average a lower patent family size as well as a lower incidence of EPO filing compared to the control group. Since these two variables are proxies for patent-worthiness, this is not necessarily consistent with the above hypothesis that high-quality applications self-select into open peer review

[26] See Pakes and Shankerman (1984) for the relation between renewals and value or Griliches et al. (1987) for the skewedness of value. Hall et al. (2005) provides empirical estimations of the relationship between forward citation and firm valuation.

process.[27]

Instead, we argue that controlling for other factors, it seems plausible that the program participation was related to the applicant's awareness of the pilot. According to the final USPTO evaluation (Allen et al., 2012) of the first pilot, the main reason why the pilot did not attain its limit of 400 participating applications is simply that many eligible applicants did not become aware of the pilot. Although coverage in several published sources and online blogs helped increase the profile, they were predominantly US-based or only available in English. Other patent offices in Australia, Japan, and the UK announced similar pilot programs later, but their announcements came only after the completion of the first US pilot.

Therefore, if an inventor was a resident of a foreign country, then he or she is less likely to be informed about the pilot. On the other hand, by law an inventor's foreign residency should be unrelated to the probability of allowance, which is supported by Dent et al.'s (2012) results. Similarly, there should be little reason why a foreign resident's US patent application is less (or more) valuable as a prior art reference than a comparable invention by a US resident. Therefore, in the following we identify the causal effect of the pilot using a dummy variable that is set equal to 1 if the first named inventor's country of residence is not the U.S. There are other related instruments (e.g., Asian country dummy) that can be used simultaneously and pass the overidentification test, but these do not change the estimation results.[28]

[27] Another problem with this selection hypothesis is that it presumes that applicants know ex ante whether their applications are more likely to be patentable or not, while in most cases applicants would (and should) have reasons to believe that their inventions are patentable.

[28] Specifically, the overidentification test using the foreign country and the Asian country dummies yields a p-value of 0.92. Because this still assumes that the foreign country instrument is valid in a just-identified model, we present results using a smaller number of instruments.

The first column of Table 2 presents the first-stage estimates for the determination of the program participation. It is shown that the inventor's foreign residency is the most significant predictor variable for nonparticipation. The treatment group also differs from the control group in such observed characteristics as the seniority of the examiner and the size of patent family. In addition, a full set of class fixed effects as well as Technology Center dummies is included to absorb cross-sectional differences in the subject matter, and also a flexible (quadratic) monthly time trend is included to account for office-level nonlinear time trends such as changes in workload and the number of applications.[29] Both estimation using ordinary and two-stage least squares allow for arbitrary heteroskedasticity in errors.

Two sets of estimation results are presented in Table 2. Results are shown separately for the probability of allowance (columns (2)-(3)) and the number of forward citations (columns (4)-(5)). Columns (2) and (4) present OLS estimates of the respective outcome equation. In both cases, the coefficients on the treatment status are positive, but they are not statistically different from zero. This seems to suggest that the pilot had negligible effect on the choice to allow a patent and also the quality of patent applications. The other significant variables are consistent with our expectations. As in Lemley and Sampat (2012), assistant (primary) examiners are less (more) likely to grant patents. The probability of allowance decreases if the inventor claimed the small entity status. Both the number of claims and the inventor's experience are positively associated with future citation counts.

2SLS estimates are presented in columns (3) and (5). For both outcome measures, the point estimates of the treatment status increase significantly. To be precise, the treatment

[29] To be precise, using June 2009 as the reference month, a patent application published in the month of June of 2008, for instance, was assigned a value of 12 for a linear term and $(12)^2 = 144$ for a quadratic term.

effect is now estimated to raise the probability of allowance by more than 25 percent, not 2 percent as in the OLS estimates, and the number of forward citations go up by more than five. Only in the latter case is the estimate statistically significant although both 2SLS estimates are relatively more precise than their OLS coefficients. The IV results also suggest that the size of patent family is positively associated with forward citation, but the inventor's experience is no longer a significant predictor. Therefore, overall the role of the open peer review appears one of facilitating diffusion of prior art knowledge similar to the one described by Furman and Stern (2011), or it raises the perceived quality of applications when examiners come to examine this particular set of patent applications.

Table 3 presents a range of sensitivity analyses. Only the treatment effect estimates are reported in the table, although the full set of covariates including class and technology center dummies as well as monthly time trend are included in all specifications. The top row of the table presents the baseline results, which correspond to those reported in Table 2. The 2SLS estimates are not sensitive to exclusion of the dummy variables or time trend. However, the results using subsample analysis suggest that there are substantial heterogeneity across the Technology Centers and the timeframe. That is, the effect of the treatment on allowance is a positive 0.69 for those applications assigned to TC 2400, but a negative 1.6 for those assigned to TC 3600. Similarly, the treatment effect does not boost the forward citation for those assigned to TC 2400 in contrast to a significant boost for those assigned to either TC 2100 or TC 3600, which implies that one should be very careful about drawing general conclusions.[30]

[30] Substantial heterogeneity also exists in online journal platforms. Using panel data on citations, Snyder and McCabe (2013) show that while in the aggregate the online-access effect is not significantly different from zero, JSTOR had a significant effect on boosting citations at around 10%.

On the other hand, restricting the sample to the first year of the pilot leads to a large increase in the estimate of the treatment effect on forward citations, which is over 13, and also a concurrent increase in the probability of allowance to 0.47. However, estimation using the second-year sample reduces the treatment effect on citation as well as probability of allowance to less than a quarter of the magnitude of the first-year sample's treatment effect. This suggests that the effect of peer review levelled off after the first year perhaps because the pilot's performance did not meet the initial expectation of examiners and the public. The bottom two rows of Table 3 show that the impact of the pilot on citation of applications that were subsequently abandoned is just as large as that of applications that were allowed. Since citation counts would be typically lower for abandoned applications, this supports the role of open peer review in facilitating diffusion of prior art knowledge.

4 Estimating Direct Effect

To our knowledge the existing literature on treatment effectiveness abstracts from the underlying mechanism that gives the key results. The implicit assumption is that the details of the underlying mechanism can be captured in a reduced-form relationship between the treatment dummy and the outcome variable of interest. The Peer-to-Patent pilot gave citizen-reviewers the opportunity to review patent applications and submit technical references and comments on what they believe to be the best prior art to consider during the examination. Since a patent is not granted on subject matter already known by others, the prior art submission and discussion posted to the Peer-to-Patent website should have a first-order effect. In this section, we investigate whether this direct, first-order effect is consistent with the above treatment effect.

One probable source of bias in estimating the causal effect of peer reviewed prior art on allowance and forward citation is the likely endogeneity of the number of prior art submission with respect to the outcome variables. That is, weak applications are likely to increase the marginal impact of prior art submission by reviewers, and may therefore tend to have a large number of prior art identified by the community of voluntary reviewers. Similarly, applications with broad claims may tend to have a large number of prior art and discussions which may be material to patentability of the claimed invention. In what follows, we propose that the number of community reviewers registered on the Peer-to-Patent website for each application in the treatment group can be used as a plausible instrument.

Whether the community size can be plausibly excluded from the outcome equation is a difficult issue. Although we do not know what motivated the reviewers to subscribe for each application, the size of these communities (which averages around 5.42 per application) are deemed too small to be representative of the general interest level of inventors. Therefore, we assume that the community size does not have a direct impact on the quality of any given application. On the other hand, there is both theoretical and empirical support from the literature on open source software regarding the relationship between community size and project development. For instance, Johnson (2002) finds that it is unlikely that an increase in the developer base will lead to fewer developments, and in fact using the case of Chinese Wikipedia, Zhang and Zhu (2011) show a positive relationship between community size and contribution levels.

Hence, if the exclusion restriction is valid, we can use the number of reviewers who subscribed for each application as a plausible instrumental variable to identify a causal effect of the number of prior art submissions and also discussion postings on our outcome measures of interest. Table 4 presents the baseline results of OLS and 2SLS estimation in the same format as Table 2, where the difference is that we now have the number of prior art submission in the first row instead of the treatment status.[31] As before, we report heteroskedastic-robust variance estimators, and the first column of Table 4 shows the first-stage estimation results. The community size is indeed a significant predictor of the number of prior art submission. An F test statistic for weak identification is greater than the usual critical values compiled by Stock and Yogo (2005).

Columns (2) and (4) present OLS estimates of the effect of prior art on allowance and forward citation, respectively. In contrast to the findings in the previous section, prior art submission has a statistically significant, negative effect on the probability of allowance, lowering the probability of allowance, after controlling for other variables, by about three percent for every piece of prior art submitted. Further, its effect on the number of forward citations is not statistically different from zero, and the estimate is effectively zero. A couple of the other covariates also yield different results. An inventor's small entity status is no longer negatively related to allowance, and instead the inventor's experience (as proxied by the number of previous patent applications) is positively associated with allowance. The number of claims is only marginally correlated with citation while the size of patent family is now negatively associated with it. The latter result is likely due to the increase in the coefficient of the EPO filing.

Columns (3) and (5) provide the corresponding 2SLS estimates of the impact of prior

[31] In Table 4's specification, we did not include class dummies due to the smaller sample size; however, the qualitative results are not sensitive to this exclusion.

art submissions. For allowance, the point estimate further decreases, so for every prior art submitted, after partialling out any controls, the probability of allowance decreases by about five percent. On the other hand, instrumenting increases the impact on the number of forward citations, but the point estimate is nowhere close to the level of treatment effect shown in Table 2 or 3, and is still effectively zero. Therefore, a comparison of the results in this section and the previous section suggests that there is the expected first-order effect, namely the reduction in allowance (with little impact on citation) from additional prior art. But this appears to be exceeded by significant second-order effects that work in the opposite direction, offsetting the reduction in the likelihood of allowance and significantly boosting forward citations at the program level.

Table 5 presents the results of sensitivity analysis using different endogenous regressors one at a time, where the same instrumental variable as well as the full set of covariates are maintained in all specifications. The top row of the table presents the effect of total number of prior art, which correspond to those reported in Table 4. The 2SLS estimates are robust across reviewer activities and to the exclusion of technology center dummies and the monthly time trend. Although identifying nonpatent art was expected to prove particularly useful for making first office action, the difference between patent and nonpatent prior art submission is not significant at all.[32] Finally, reviewer's discussion postings are shown to have a significant, negative effect on allowance although the magnitude of the estimate is smaller than that of prior art submissions. As for the forward citation, most of the estimates

[32] "Studies have shown that when our patent examiners have the best data in front of them, they make the correct decision," said Jon Dudas, the then director of the USPTO. "Examiners, however, have a limited amount of time to find and properly consider the most relevant information. This is particularly true in the software-related technologies where code is not easily accessible and is often not dated or well documented" (USPTO Press Release, #07-21).

are less than one and insignificant.

5 Examiner's Search Effort

The original aim of the pilot was to address patent backlogs and improve examination quality, and one would expect that more prior art would narrow claims or reduce allowances. Hence, that the direct, first-order effect of prior art submission is several times less than the indirect, second-order impact at just the treatment level seems surprising. There are likely multiple hypotheses that would be consistent with these findings. Although we cannot tease out all the potential explanations, in this section we aim to rule out at least one alternative hypothesis (the "crowding-out" of examiner search by peer reviewed prior art) and put forth our favored interpretation (positive "signaling" through participation in the open peer review). Although reducing the examiner's workload is a part of the intention of the pilot program, we find no evidence that examiners have substituted their own search activity for the reviewer's prior art search.

To be more precise, there are at least two possible channels through which the divergence between the treatment-level and the prior art effects may have arisen, although our interpretation at this stage is inherently speculative. In the literature on private and government provision of public goods, it has been suggested that private contributions could crowd out government support. For instance, Peltzman (1973) and Becker and Lindsay (1994) show that private donations to higher education institutions resulted in significant reductions in government funding. The degree to which government and private sources are substitutes could depend on the nature of the relationship between the two. In the patent examiner-applicant relationship, their incentives are clearly not aligned with each other. Hence, in the standard examination, a patent examiner is unlikely to substitute his own search for prior art with the information disclosed by the applicant.[33]

Given that prior art search is labor intensive and examiners operate within strict timelines for delivery (FTC, 2003; Jaffe and Lerner, 2004), a system designed to support examination by providing prior art deemed important by the expert reviewers seems more aligned to the examiner's incentives. By locating prior art and generating discussion on them, the pilot can provide a basis for substitution between the examiner's and the reviewer's search effort. Hence, one possibility is that examiners allow more applications that enter the treatment group because free-riding, leading to worse exam quality. Evaluated at the mean number of prior art submission per application (2.59), our baseline 2SLS estimates suggest that the second order effects can be estimated as a 0.41 increase in the probability of allowance.[34] This ignores heterogeneous responses, but the magnitude of this difference suggests that the second-order effect does exist.

A complementary explanation of the above results, which we favor, is that participation in the peer review process acted as a signaling mechanism of the perceived quality of the treatment applications. Because participating applications were subject to volunteer scientific and technical experts' review and criticism, they may have had more exposure to the relevant inventor and scientific community. While prior art submitted by the reviewers

[33] Parties involved with prosecution of a U.S. patent application have a duty to disclose all material prior art known to them (see Title 37, Code of Federal Regulations 1.56 for the full US disclosure requirement). However, there is no duty to search for prior art; that is, if the inventor or assignee know of no prior art, then they have no duty to disclose it.

[34] First-order effect + second-order effects = treatment effect at mean; so $(-0.054) * 2.59+$ second-order effects $= 0.267$; and re-arranging the equation to solve for the difference between the first-order effect and the total effect of treatment yields $0.267 + (0.054 * 2.59) = 0.41$.

reduced the likelihood of allowance, the examiner's critical assessments of the patent applications may have been affected by the fact that the peer review outcome reflects some eyeball attention of citizen-experts. Further, while prior art submission did not have any significant effect on the number of forward citations, the finding that the treatment applications received significantly more forward citations seems to appear more consistent with the signaling hypothesis than with the crowding-out hypothesis.

Although testing a signaling hypothesis is beyond the scope of this paper, we can test the crowding-out hypothesis using data from the examiner's search reports.[35] That is, the examiner must list on a form (PTO-892) all prior art patents or printed publications that have been applied in making rejections or cited as being pertinent during the examination process.[36] Since the references listed on this form were discovered by the examiner during a prior art search and do not include the references disclosed by the applicant, information contained in form 892's can serve as a proxy for the examiner's search effort. Notice that the applicant has a duty to disclose prior art known to them. Hence, we believe that it is unlikely that the applicants who participated in the pilot program behaved strategically in disclosing prior art. That is, the treatment effect on the examiner's search effort would not be confounded by the applicant's disclosure.

We measure the number of the examiner's search reports as well as the number of patent and non-patent references contained therein and estimate the treatment effect on these mea-

[35] In the literature on labor market signaling, empirically disentangling the human capital hypothesis and the signaling hypothesis has proven to be difficult. To plausibly test the signaling value of education, one would need an exogenous variation in educational degree attainment (see, e.g., Tyler et al., 2000). See also Weiss (1995) for a review of this literature.

[36] References listed on a form PTO-892 during the examination process are also indicated with an asterisk in the "References Cited" section of the front page of a patent document.

sures. We only count newly added references from each 892 reports, so that they measure the incremental search efforts. The operational hypothesis is that after controlling for other sources of variation if the peer review process encouraged the examiners to reduce their search effort, then this should be reflected as a reduction in the number of search reports as well as references cited by the examiners for those applications in the treatment group. We further divided the aggregate number of references discovered by the examiner into those cited before the first office action (for which there is typically one 892 report) and during the rest of the prosecution, each of which is used to estimate the short-run and the long-run substitution (or crowding-out) effect.

Table 6 presents the estimation results in the same manner as in Table 2. Since the first-stage estimation is the same as that in Table 2 and thus omitted. It turns out that there is no evidence that the examiners reduced their search effort for the treatment applications either in terms of the number of search reports or the number of references added. Specifically, while statistically insignificant, the treatment effect on the number of search reports has a positive sign, and in the middle columns both OLS and 2SLS estimates of the treatment effect on the number of referenced cited before the first office action (for which peer reviewed prior art was submitted) indicate that the examiners cite four to five more prior art for the first-office action. In the last column, the 2SLS estimate suggests that the examiners found more prior art for the rest of the examination process as well in the same order of magnitude as before the first office action.

Table 7 shows the sensitivity analysis of these estimates. Although there is substantial heterogeneity across Technology Centers, the overall evidence suggest that the examiners put more, not less, effort into prior art search when the applications were simultaneously subjected to open peer review. In particular, that the examiners added more references for the treatment group after the first-office action in the second year of the pilot than in the first year (that is, 7.54 more prior art versus 3.98 in the last column of Table 7) seems consistent with the previous finding in Table 3 that the treatment effect dwindled in the second year of pilot both in terms of allowance and forward citation. This may suggest that the examiners became more suspicious over time about the effectiveness of the pilot's performance and responded by increasing their own search effort, which seems to refute the crowding-out hypothesis.[37]

6 Conclusion

At the time of its announcement, the USPTO's Patent Peer Review Pilot was expected to bolster the credibility of the examination system and improve patent quality. Although the idea behind the open, public involvement in the patent examination process deserves separate attention, our analysis suggests that the evaluation of the system needs to consider first- and second-order effects, as well as use administrative data to see how the eventual patents made their way through the system, and how they end up. The pilot program was successful in adding more prior art references to applications (Allen et al., 2012), but the eventual impact of this appears to have been little impact on allowance rates and a boost in forward citations for applications that entered the pilot, both of which seem to diverge from the first-order effect.

[37]Similar conclusions hold when we estimate the direct effect of prior art submissions on the examiner's search efforts. That is, rhe direct effect on all outcome measures have a positive sign. The only difference is that the 2SLS estimate of the effect on the number of references cited after the first-office action is not statistically significant.

That the analyses at the treatment level and the activity level show a contrasting pattern in both sign and magnitude of estimated effect on patent allowance and forward citations suggests that some caution needs to be exercised in drawing conclusions on generic impacts of open peer-review process. Our approach in this paper sheds light on some basic issues in program evaluation studies, but is limited to the social costs and benefits from higher quality patent examination as we do not consider the costs to the USPTO, the examiners and applicants of managing the system. Our results may be replicated using data generated from similar pilot programs commenced in Australia, Japan, and the UK, and highlights that qualitative surveys and first-office action analysis are useful, but policy evaluation deserves more empirical scrutiny.

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://dotank.nyls.edu/communitypatent/P2Panniversaryreport.pdf [accessed December 2012].

[3] Allen, N., A. Casillas, J. Deveau-Rosen, J. Kreps, T. Lemmo, J. Merante, M. Mur-

phy, 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://dotank.nyls.edu/communitypatent/CPI_P2P_YearTwo_lo.pdf [accessed December 2012].

[4] Allen, N., 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 [accessed December 2012].

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

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

[7] Dent, C., A. Lim, and A. Christie. 2012. "A Comparative Analysis of Patent Examination: What Harmonisation." Presented at European Policy for Intellectual Property Conference, 27-28 September 2012, University of Leuven, Belgium.

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

[9] Federal Trade Commission. 2003. *To Promote Innovation: The Proper Balance of Competition and Patent Law and Policy*.

[10] 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 December 2012].

[11] 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.

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

[13] Griliches, Z., Pakes, A., and Hall, B. 1987. "The Value of Patents as Indicators of Inventive Activity." In Dasgupta, P. and P. Stoneman, eds. *Economic Policy and Technological Performance*, Cambridge, UK: Cambridge University Press.

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

[15] 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.

[16] Intellectual Property Office. 2012. *Peer to Patent – Pilot*. London: Intellectual Property Office, http://www.ipo.gov.uk/p2p-report.pdf [accessed December 2012].

[17] Jaffe, A., 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.

[18] 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.

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

[20] Lanjouw, J., and M. Schankerman. 2004. "Patent Quality and Research Productivity: Measuring Innovation with Multiple Indicators." *Economic Journal* 114: 441–465.

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

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

[23] 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 December 2012].

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

[25] Pakes, A., and M. Schankerman. 1984. "The Rate of Obsolescence of Patents, Research Gestation Lags, and the Private Rate of Return to Research Resources." In Z. Griliches, ed. *R&D, Patents, and Productivity*, Chicago: University of Chicago Press.

[26] 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.

[27] Pöschl, U. 2012. "Multi-Stage Open Peer Review: Scientific Evaluation Integrating the Strengths of Traditional Peer Review with the Virtues of Transparency and Self-Regulation." *Frontiers in Computational Neuroscience* doi: 10.3389/fncom.2012.00033.

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

[29] Snyder, C., and M. McCabe. 2013. "Does Online Availability Increase Citations? Theory and Evidence from a Panel of Economics and Business Journals." Social Science Research Network, http://ssrn.com/abstract=1746243 [accessed March 2013].

[30] 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.

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

[32] Tyler, J., R. Murnane, and J. Willett. 2000. "Estimating the Labor Market Signaling Value of the GED." *Quarterly Journal of Economics* 115: 431–468.

[33] United States Patent and Trademark Office. 2006. *Performance and Accountability Report*.

[34] Webster, B., P. Jensen, and A. Palangkaraya. 2011. "Patents and the Enforcement of the National Treatment Principle." Intellectual Property Research Institute of Australia, http://is.jrc.ec.europa.eu/pages/ISG/patents/documents/Webster_NationalTreatment_17jan2011a. [accessed December 2012].

[35] Weiss, A. 1995. "Human Capital vs. Signalling Explanations of Wages." *Journal of Economic Perspectives* 9: 133–154.

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

Variable	Sample mean	Standard deviation	Treatment group	Control group	Two-tailed test
Allowance	0.66	(0.47)	0.70	0.64	[0.12]
Examiner assistant	0.55	(0.50)	0.44	0.59	[0.00]
Inventor previous apps.	10.74	(29.19)	13.59	9.63	[0.12]
Inventor foreign	0.34	(0.47)	0.15	0.42	[0.00]
Inventor small entity	0.22	(0.41)	0.25	0.21	[0.25]
Number of claims	20.75	(11.38)	20.29	20.93	[0.52]
Number of indep. claims	3.21	(1.87)	3.22	3.20	[0.94]
Size of patent family	0.94	(1.62)	0.37	1.16	[0.00]
EPO filing	0.19	(0.39)	0.08	0.23	[0.00]
Number of forward citations	1.71	(3.45)	1.73	1.71	[0.93]
Number of examiner search	1.72	(1.09)	1.72	1.73	[0.93]
Number of examiner cited ref.	9.79	(8.98)	12.69	8.66	[0.00]
of which non-patent art	1.07	(2.13)	2.46	0.53	[0.00]
Size of P2P community	5.42	(5.92)	5.42		
Number of P2P art submissions	2.59	(2.46)	2.59		
of which non-patent art	1.31	(1.66)	1.31		
Number of P2P discussions	3.47	(6.43)	3.47		
Number of observations	642		180	462	

Table. 1: SUMMARY STATISTICS

Variable	First-stage estimation	Allowance		Citations	
		OLS	IV	OLS	IV
P2P Treatment		0.020	0.267	0.191	5.523 ***
		(0.040)	(0.213)	(0.317)	(1.513)
Examiner assistant	-0.097 ***	-0.237 ***	-0.210 ***	-0.111	0.461
	(0.035)	(0.036)	(0.043)	(0.258)	(0.369)
Inventor previous apps.	0.001	0.001	0.000	0.019 **	0.013
	(0.001)	(0.001)	(0.001)	(0.010)	(0.010)
Inventor small entity	0.007	-0.179 ***	-0.181 ***	0.576	0.529
	(0.048)	(0.049)	(0.048)	(0.461)	(0.503)
Number of claims	-0.002 *	0.002	0.002	0.032 **	0.038 ***
	(0.001)	(0.002)	(0.002)	(0.015)	(0.014)
Number of indep. claims	0.012	0.011	0.008	0.038	-0.010
	(0.009)	(0.011)	(0.010)	(0.086)	(0.094)
Size of patent family	-0.040 ***	-0.017	-0.005	0.070	0.341 **
	(0.012)	(0.016)	(0.019)	(0.142)	(0.168)
EPO filing	-0.022	0.078	0.085	0.589	0.757
	(0.052)	(0.066)	(0.065)	(0.483)	(0.552)
Inventor foreign	-0.189 ***				
	(0.034)				
Class dummies included?	yes	yes	yes	yes	yes
TC dummies included?	yes	yes	yes	yes	yes
Time trend included?	yes	yes	yes	yes	yes
F statistic of excluded instruments = 30.24; Cragg-Donald Wald F statistic = 25.03					
R^2		0.235		0.127	
N	642	642	642	642	642

Table. 2: TREATMENT EFFECT OF PILOT PROGRAM

Specification	Allowance		Citations	
	OLS	IV	OLS	IV
Baseline	0.020	0.267	0.191	5.523 ***
	(0.040)	(0.213)	(0.317)	(1.513)
Without class dummies	0.025	0.261	0.152	5.477 ***
	(0.040)	(0.215)	(0.322)	(1.523)
Without TC dummies	0.020	0.258	0.197	5.409 ***
	(0.041)	(0.210)	(0.314)	(1.490)
Without time trend	0.036	0.273	0.164	5.534 ***
	(0.040)	(0.206)	(0.292)	(1.493)
TC2100 sample	-0.002	0.339	-0.105	5.764 ***
	(0.051)	(0.237)	(0.403)	(1.780)
TC2400 sample	0.072	0.691	0.971	1.499
	(0.086)	(0.663)	(0.855)	(3.288)
TC3600 sample	0.114	-1.590	0.842	6.983 *
	(0.138)	(1.121)	(0.583)	(4.038)
First-year sample	-0.013	0.468	0.637	13.971 **
	(0.071)	(0.495)	(0.641)	(5.690)
Second-year sample	0.028	0.137	0.128	2.253 *
	(0.053)	(0.256)	(0.344)	(1.203)
Conditional on allowanace			0.254	5.277 ***
			(0.378)	(1.855)
Conditional on abandonment			0.169	5.840 **
			(0.438)	(2.650)

Table. 3: SENSITIVITY ANALYSIS OF THE ESTIMATED EFFECT

Variable	First-stage	Allowance		Citations	
	estimation	OLS	IV	OLS	IV
Number of prior art		-0.035 **	-0.054 **	0.005	0.151
		(0.015)	(0.026)	(0.094)	(0.229)
Examiner assistant	-0.312	-0.179 ***	-0.187 ***	-0.198	-0.137
	(0.311)	(0.068)	(0.067)	(0.601)	(0.564)
Inventor previous apps.	-0.002	0.002 ***	0.002 ***	0.014	0.014
	(0.003)	(0.000)	(0.000)	(0.012)	(0.012)
Inventor small entity	0.192	-0.059	-0.060	1.223	1.232
	(0.316)	(0.086)	(0.082)	(1.334)	(1.291)
Number of claims	0.039 **	0.004	0.004	0.086 *	0.081 *
	(0.016)	(0.006)	(0.006)	(0.047)	(0.045)
Number of indep. claims	-0.062	0.028 *	0.027 *	0.138	0.145
	(0.058)	(0.016)	(0.015)	(0.209)	(0.204)
Size of patent family	-0.054	-0.039	-0.035	-0.679 ***	-0.712 ***
	(0.206)	(0.057)	(0.057)	(0.221)	(0.225)
EPO filing	-0.992 **	0.157	0.125	1.034	1.270
	(0.494)	(0.166)	(0.170)	(1.138)	(1.163)
Size of P2P community	0.284 ***				
	(0.044)				
TC dummies included?	yes	yes	yes	yes	yes
Time trend included?	yes	yes	yes	yes	yes
F statistic of excluded instruments = 42.00; Cragg-Donald Wald F statistic = 64.58					
R^2		0.203		0.137	
N	180	180	180	180	180

Table. 4: EFFECT OF PRIOR ART SUBMISSION

Specification	Allowance		Citations	
	OLS	IV	OLS	IV
Baseline	-0.035 **	-0.054 **	0.005	0.151
	(0.015)	(0.026)	(0.094)	(0.229)
Without TC dummies	-0.040 **	-0.064 **	-0.016	0.112
	(0.015)	(0.026)	(0.092)	(0.215)
Without time trend	-0.034 **	-0.043 **	0.117	0.421 *
	(0.014)	(0.021)	(0.099)	(0.237)
Patent art	-0.037 *	-0.109 *	-0.237	0.303
	(0.021)	(0.057)	(0.181)	(0.481)
Without TC dummies	-0.040 *	-0.130 **	-0.252	0.233
	(0.021)	(0.057)	(0.194)	(0.447)
Without time trend	-0.039 *	-0.105 *	-0.165	1.035
	(0.020)	(0.057)	(0.155)	(0.688)
Non-patent art	-0.028	-0.108 *	0.338	0.299
	(0.023)	(0.059)	(0.237)	(0.447)
Without TC dummies	-0.036	-0.128 **	0.310	0.222
	(0.024)	(0.061)	(0.217)	(0.420)
Without time trend	-0.029	-0.072 *	0.458 *	0.710 *
	(0.020)	(0.038)	(0.259)	(0.392)
Discussions	-0.012 **	-0.017 **	0.008	0.047
	(0.006)	(0.007)	(0.051)	(0.072)
Without TC dummies	-0.015 **	-0.021 **	0.004	0.036
	(0.006)	(0.008)	(0.050)	(0.069)
Without time trend	-0.009 *	-0.011 **	0.077 *	0.112 *
	(0.005)	(0.005)	(0.045)	(0.066)

Table. 5: SENSITIVITY ANALYSIS OF THE ESTIMATED EFFECT

Variable	Number of examiner search reports (PTO form 892)		Number of references cited before the first office action		Number of references added after the first office action	
	OLS	IV	OLS	IV	OLS	IV
P2P Treatment	0.013	0.658	4.444 ***	5.129 *	-0.447	4.867 **
	(0.098)	(0.495)	(0.764)	(2.840)	(0.481)	(2.414)
Examiner assistant	0.079	0.148	-0.970	-0.897	-0.390	0.180
	(0.089)	(0.099)	(0.595)	(0.692)	(0.441)	(0.443)
Inventor previous apps.	0.000	-0.001	-0.003	-0.003	-0.005	-0.011 **
	(0.001)	(0.001)	(0.007)	(0.008)	(0.004)	(0.006)
Inventor small entity	-0.355 ***	-0.361 ***	0.698	0.692	-1.222 ***	-1.269 ***
	(0.090)	(0.090)	(0.852)	(0.838)	(0.389)	(0.451)
Number of claims	0.005	0.006	0.044	0.045	0.014	0.019
	(0.004)	(0.004)	(0.032)	(0.031)	(0.015)	(0.017)
Number of indep. claims	0.016	0.010	0.011	0.005	0.027	-0.021
	(0.023)	(0.024)	(0.132)	(0.128)	(0.076)	(0.091)
Size of patent family	-0.044	-0.011	0.201	0.236	-0.133	0.138
	(0.047)	(0.052)	(0.298)	(0.332)	(0.145)	(0.201)
EPO filing	0.193	0.213	-1.151	-1.130	0.576	0.744
	(0.165)	(0.162)	(1.101)	(1.058)	(0.637)	(0.698)
Class dummies included?	yes	yes	yes	yes	yes	yes
TC dummies included?	yes	yes	yes	yes	yes	yes
Time trend included?	yes	yes	yes	yes	yes	yes
F statistic of excluded instruments = 30.24; Cragg-Donald Wald F statistic = 25.03						
R^2	0.088		0.137		0.066	
N	642	642	642	642	642	642

Table. 6: EFFECTS ON EXAMINER'S SEARCH BEHAVIOR

Specification	Number of examiner search reports (PTO form 892)		Number of references cited before the first office action		Number of references added after the first office action	
	OLS	IV	OLS	IV	OLS	IV
Baseline	0.013	0.658	4.444 ***	5.129 *	-0.447	4.867 **
	(0.098)	(0.495)	(0.764)	(2.840)	(0.481)	(2.414)
Without class dummies	0.048	0.762	4.589 ***	4.610	-0.187	4.985 **
	(0.095)	(0.506)	(0.752)	(2.852)	(0.445)	(2.383)
Without TC dummies	0.007	0.666	4.447 ***	4.925 *	-0.426	4.791 **
	(0.098)	(0.491)	(0.761)	(2.793)	(0.476)	(2.374)
Without time trend	-0.004	0.650	4.136 ***	4.863 *	-0.459	4.684 **
	(0.096)	(0.477)	(0.738)	(2.745)	(0.447)	(2.289)
TC2100 sample	-0.139	0.738	4.897 ***	8.176 **	-0.583	3.814
	(0.131)	(0.580)	(1.079)	(3.269)	(0.688)	(2.340)
TC2400 sample	0.032	-1.465	3.098 **	-5.675	-0.906	4.561
	(0.200)	-1.688	(1.211)	(8.780)	(0.823)	(8.857)
TC3600 sample	0.547 **	2.196 *	4.495 **	15.513	0.640	5.438 *
	(0.220)	(1.245)	(2.187)	(9.948)	(1.200)	(3.223)
First-year sample	-0.282 *	1.194	6.890 ***	10.183	-0.273	3.975
	(0.170)	(1.222)	(1.448)	(6.499)	(0.885)	(4.387)
Second-year sample	0.291 **	0.788	2.664 ***	1.538	-0.054	7.541 **
	(0.122)	(0.568)	(0.875)	(3.145)	(0.540)	(3.535)
Patent references only			2.590 ***	4.011	-0.527	3.929 *
			(0.704)	(2.707)	(0.449)	(2.273)
Non-patent references only			1.854 ***	1.118	0.080	0.938 **
			(0.219)	(0.706)	(0.082)	(0.379)

Table. 7: SENSITIVITY ANALYSIS OF THE ESTIMATED EFFECT