Title
Essays on Using the IP System to Promote Innovation and Economic Growth

Author
Sun, Zhen

Publication Date
2016

Peer reviewed|Thesis/dissertation

Essays on Using the IP System to Promote Innovation and Economic Growth

By

Zhen Sun

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Agricultrual and Resource Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Brian D. Wright, Chair
Professor Michael Anderson
Professor Jeffrey Perloff
Professor Lee Fleming

Spring 2016

Abstract

Citations as Indicators of Patent Value
— A non-linear reinterpretation of the empirical evidence

by

Zhen Sun

Doctor of Philosophy in Agricultrual and Resource Economics

University of California, Berkeley

Professor Brian D. Wright, Chair

One of the main purposes of the patent system is to facilitate the flow of knowledge by ensuring access to information related to inventions. In the United States, the patent applicant has a legal duty to disclose knowledge of relevant prior art in the form of citations, thus placing the innovation in context and enhancing knowledge of this prior art. However, the applicant might miss relevant prior art due to ignorance, or conceal it for strategical reasons, which could reduce the value of public information generated by the applications and also reduce the accuracy of patent grant decisions. Patent examiners can add citations not known to the applicant, because they supposedly have less barrier in accessing knowledge from different regions. The examiners, specialized in a given technical area, might also identify and add prior art that are deliberately omitted by applicants thus tend to render the patent obvious or non-novel. Prior literature generally agrees the examiner citation is more important than applicant citation as a signal of patent value or knowledge flow. Using US patents granted since 2001, I test the hypothesis that examiner citations contain more important information on patent value and knowledge flow than applicant citations. Surprisingly, on a variety of measures, I find evidence to support the contrary. I also show that the evidence supporting the current hypothesis was attained by a model mis-specification. The findings shed light on many innovation studies that rely on patent citations, and also have important policy implications regarding the operation of the patent office and in general the patent examination system.

1. Introduction

The patent system is considered an instrumental part of the institutions that contribute to modern economic growth (Mokyr, 2009). A patent is a property right of potential commercial use, which grants its owner a limited period of monopoly over the invention. In exchange, the proprietary knowledge embedded in a patent must be disclosed to the public. In this way, the patent system facilitates the flow of knowledge. It also means the inventors are required to disclose any relevant necessary knowledge that they use to achieve the invention. In fact, the novelty requirement for a patent is evaluated based on the existing related technology (the so-called "prior art"). In US, applicants are obliged to disclose any knowledge of the prior art (the "duty of candor" requirement), which is included in the patent documents in the form of citations. The legal purpose of these citations is to indicate which parts of the knowledge described are claimed in the patent and which parts have been claimed by previous patents or non-patent. From an economic perspective, patent citations are the "paper trails" of knowledge flows (Jaffe et al., 2000) and contain important information about the inventions. In principle, when a patent cites another patent, this indicates that the knowledge embodied in the cited patent has been useful in someway for developing the newknowledge described in the citing patent. Therefore, patent citations are important in the study of knowledge flow and value as they place the innovation in context of inventive trajectories and indicate the importance of a cited patent.

However, it has long been noted that the patent applicants could conceal or miss relevant prior art. Importantly, they could strategically withhold known information from publicizing them on patent application document, such that they may be granted low quality patents with intellectual property claims which would not otherwise have been granted (Lampe, 2012; Steensma et al., 2014). On the other hand, the citations could be missed due to the inventors' ignorance of the prior art. For example, there is evidence suggesting that knowledge flows embedded in patents are constrained by geographical barriers (Jaffe et al., 1993; Henderson et al., 2005; Peri, 2005). The seminal work done by Jaffe et al. (1993) identified a strong "localization effect" on knowledge spillovers by comparing locations of patents that appear in citation pairs with locations of carefully constructed patent pairs (control group) that are not linked by citation. The "localization effect" indicates that inventors may miss some innovation activities across the geographical boundaries. Moreover, inventors may have disincentive to conduct a thorough prior art search, as the "duty of candor" requirement only bears obligation for the inventors to disclose what they are aware of (Wagner, 2002). The "missing citations", from either strategical withholding or ignorance, can be considered as "frictions" that hinder the flow of knowledge and undermines the function of the patent system.[1]

An embedded feature of the patent system, the examination of patents, could help fill these "missing citations". In most patent offices in the world, the decision regarding which citations to include ultimately rests with the patent examiners, who are experts in their respective technology field and thus supposed to be able to identify relevant prior art that the applicant misses or conceals (Hall et al., 2005). In fact, Cotropia et al. (2013) report that patent examiners rely almost exclusively on their own prior art in narrowing patent ap-

[1] On the other hand, inventors usually do not to cite patents unnecessarily, as it may reduce their claims to novelty of the invention and therefore affect the scope of the patent rights.

plications, and their decisions are not affected by the citations added by applicants. These examiner hand-picked citations are therefore highly likely to be relevant to the patent application and sometimes considered the "citations that matter". If the examiners can fill the missing citations from the applicants, we would expect examiner citations send stronger signal to value of the cited patent (both in terms of technical importance and commercial value) than applicant citations. Because if the missing citation is strategical, it indicates significant value of, or important knowledge flow from the cited patent. Even if the missing citation is truly due to ignorance, the examiner-added citation sends signal to the inventors of the later patent about the technical relevance and importance of the cited patent. The signal comes as a "shock" to the citing firms, which could stimulate further inventions informed by the cited patent.

Because of the higher likelihood of relevance associated with the hand-picked examiner citations it can be valuable to separate them from the applicant-added citations in any citation analysis. This was nearly impossible until 2001, when the US patent office (USPTO) started to distinguish examiner from applicant citations in its patent database. The importance of examiner citations have been discussed in several studies since then. They in general find evidence in support of the reletive importance of examiner citations. For example, compared to applicant citations, examiner citations are more likely to traverse geographical boundaries (Alcacer and Gittelman, 2006; Thompson, 2006; Criscuolo and Verspagen, 2008), providing supportive evidence on "localized knowledge spillovers".[2] Hegde and Sampat (2009) find that examiner (forward) citations to a patent are stronger predictors for patent renewals than inventor citations. The authors speculate that examiner citation is perhaps a better measure of patent scope (subsequent inventions it potentially blocks), while application citation measures technology spillovers that do not necessarily generate private rents for patent owners.[3] Building on the assumption that examiner citations reveal important information links between patents, (Lampe, 2012) find inventors strategically withhold roughly one third of known information from publication of the patent application document.

However, there are still puzzles remaining. For example, the statistics suggest that roughly one third of the self citations (citations made to patents owned by the same patentee) are added by examiners. The proportion is almost the same as non-self citations. This could pose a challenging view on the idea of patent applicants strategically withholding citations as firms should have less incentive to "strategically withheld" self-citations than citations from other parties. Firms should in general have strong incentive to self cite their own inventions to, for example, build a thicket of patent.[4] On a similar note, Alcacer and Gittelman (2006) report a puzzling observation that the citations made to patents that share at least one inventor are more likely to be examiner citations than inventor citations. These "missing"

[2] Examiners, sitting at an office at Washington DC, could not presumably know the communication among the inventors, and thus the localization of examiner citations indicates a pooling or concentration effect (Marshall, 1920). The localization of applicant citations, beyond the level of examiner citations, thus could represent geographical barrier to knowledge flows.

[3] Similar findings are later reported in Czarnitzki et al. (2012), using EPO patents, that "blocking forward citations" have stronger correlation with patent value than other citations.

[4] Interestingly, the proportion of examiner-added self citations is about the same as the "withheld known citations" reported in Lampe (2012), who excludes self citations but only considers patents that were cited by the assignee previously (thus considered to be "known" to the assignee). Therefore, whether these examiner-added citations are truly "strategically withhold" citations is more uncertain.

citations cannot be due to ignorance and strategical citation does not seem to explain the behavior either. These and other questions prompt the current study to further investigate the role of examiner citations versus applicant citations in signalling the patent value.

In this paper, I test the hypothesis that examiner citations send stronger signal to the value of cited patents and contain more important information on knowledge flow than inventor citations. I approach the research question in two ways. Firstly, I re-examine the correlations between patent renewal and examiner forward citations *vis-à-vis* inventor forward citations. I find, contrary to existing belief, examiner citations are much less associated with patent value (measured by renewals) than inventor citations. I also show, using citation data and Monte Carlo simulation, that the existing inferences arise from a misspecified model that fails to take into account the difference in distribution of the two types of citations and the (naturally) marginally decreasing relation between the number of citations and the renewal decision. I confirm the findings by carefully examining the different combinations of the first few citations a patent receives. Secondly, I investigate the knowledge flow generated by a patent based on the type of the first citation it receives — whether the patent was first cited by the examiner or the inventor. I find that if the first citation to a patent comes from the inventor, the patent turns out to be more highly cited over time by other inventors than if it is first cited by the examiner, as measured by the total number of inventor forward citations. The evidence is not consistent with the hypothesis that examiner citation represents "ignorance" or "strategic withholding" on the applicant's side. Because under the hypotheses, if such a link is "revealed" by the examiner, we should expect future inventor citations to this patent to increase. Taken together, the findings suggest inventors may have systematically different views from examiners regarding the relevance or importance of the prior art.

The findings have important implications on business management, innovation research and policy debate. Citations are one of the few observable indicators of the value of knowledge flow. Knowledge embodied in patents is rarely traded on market and therefore difficult to evaluate, especially at the early stage of the development. At the same time, intangible capital is now increasingly becoming the most important asset of a firm. Patent citations provide one natural way to predict value of inventions. However, it's often used in a gross way, without considering the fact that citations come from both the applicants and patent examiners. It is important to distinguish these two types of citations, as they may contain different information regarding the value of the cited invention. If examiners are less capable at identifying valuable (from a commercial point of view) patents than inventors, the total citations including examiner citations could be a very "noisy" predictor of patent value. In innovation studies, misunderstanding the relative importance of the two types of citations could lead to misinterpreted research results. For example, a "missing citation" may not indicate that inventors are ignorant or strategically withhold known information, but rather inventors might have very different opinions from examiners regarding the relevance or importance of the prior art. For another example, the first few citations to a patent could be especially important in the study of knowledge flow, as information in a new invention tends to be tacit in early period and becomes codified over time (Griffith et al., 2011). However, examiner citations in general come much earlier than inventor citations as examiners have more convenient access to patent documents than inventors. The first few citations to a patent thus are disproportionally examiner citations. Research focusing on the influence of the first few citations, without distinguishing the source of the citation, could be misleading. On the policy side, the findings suggest the importance of the "duty of candor" requirement on inventors reporting their prior art, such as that is implemented at USPTO; and possible consequences of relying on examiners to find most of the prior art (for example, at EPO). Lastly, as a technical note, I also show the danger of using simple linear models when the true association of the independent variable to dependant variable is marginally decreasing and when the supports of the independent variables are very different. This simple fact, if overlooked, can produce misleading results.

The article is organized as follows. In Section 2 I provide a brief discussion on using citations to measure patent value. Section 3 describes the data and some summary statistics. In Section 4 I present the first set of empirical findings: starting with a simple linear model that leads to misleading results, I proceed to show, with the right regression strategy, inventor citations are much more strongly associated with patent renewals than examiner citations. Section 5 provides the second piece of evidence based on the type of the first citation, and shows that patents first cited by inventor to be more valuable. Section 6 concludes.

2. Patent Forward Citations and Patent Value

Patent values are difficult to measure as most patents are not traded in a market. Patent renewals provide arguably the most objective indicator of the patent's economic value. Many patent offices require patent owners to pay periodic and usually monotonically increasing maintenance fees in order to keep the patents in force. Patent renewal fees could constitute a significant financial cost to patent holders. For example, maintaining a patent to full-term in US costs $8,860 before 2013, and this fee has been increased to $12,600 after 2013. A risk neutral patentee will only renew its patent if the expected value is greater than maintenance fee. Pakes and Schankerman (1984) and Pakes (1986) are among the first attempts to estimate patent value based on renewal decisions, under some assumptions on the initial distribution of patent value and its evolution over time. Richer models were developed to take into account learning, possibility of litigation, and other patent characteristics that affect patent value (Lanjouw, 1998; Baudry and Dumont, 2006; Bessen, 2008; Deng, 2011). In the first part of the paper, I use patent renewals to evaluate whether examiner or inventor forward citations are better predictors of patent value.

Patent citation data are used in a growing body of research on technological diffusion and business management. Patent citation, in principle, indicates the citing patent builds upon a piece of existing knowledge in the cited patent. Therefore, it is considered a direct measure of knowledge flow, compared to the bulk of the literature that relies on some forms of knowledge production function (see Breschi and Lissoni, 2001 for a critical review). Numerous studies rely on patent citations to study knowledge flows and spillovers (Jaffe et al., 1998; Jaffe and Trajtenberg, 1999; Duguet and MacGarvie, 2005).

It is appropriate at this stage to discuss the concept of forward and backward citations, as they are often used for different purposes. The discussion of patent citations always begins with a focal patent. The citations that this focal patent references itself are called backward citations, since they are citations which proceeded the focal patent. Conversely, going forward in time from the focal patent, any more recent patents which cite the focal patent are called forward citations to the focal patent. Therefore, backward citations of a patent contain the prior art that this patent builds on; while forward citations provide us a measure of the amount of future research that uses the knowledge of the focal patent.

Patent forward citations carry important information about patent's technological or economic value, which are important in research on innovation because of the wide variance in patent value. Trajtenberg (1990) provides the first empirical evidence that citations to a patent contain quite accurate information on the commercial value of patented innovations. Albert et al. (1991) shows that forward citations strongly correlate with the purported technical value of the patent as well. There has since been a number of studies that show strong correlation between forward citations and patent value (Harhoff et al., 1999; Lanjouw and Schankerman, 2004; Gambardella et al., 2008; Bessen, 2008), as well as between forward citation weighted patent counts and firms' value (Shane and Klock, 1997; Hirschey and Richardson, 2004; Hall et al., 2005). However, it is noted that citations could be a noisy signal of patent value and especially knowledge flow, not in the least of which is the fact that citations come from two different sources: they are added by bothe the patent examiner and inventors (Jaffe et al., 2000). What role does examiner citation play is of immediate importance, given a large share ($\sim$40% for patent citations at USPTO, much higher at EPO) of citations are added by the examiners and their different characteristics (Alcacer and Gittelman, 2006; Thompson, 2006; Alcacer et al., 2009; Criscuolo and Verspagen, 2008).

3. Data

USPTO started to distinguish examiner citation from applicant citation since 2001. Therefore I can only fully tell whether the forward citations are from the examiner or not for the patents granted after 2001. I collect all patents that are granted between 2001 and 2009 collect the forward citations to these patents.[5] I exclude citations made by continuing patents (continuation, continuation-in-part and divisional patents) and restrict the forward citations only to those made by original patents for two reasons. Firstly, continuing patents usually automatically include most if not all citations of their parent patents. These citations perhaps do not carry the same weight as those made by original patents, and to some extent "double count" the actual citations to the cited patents. Secondly and perhaps more problematically, these citations inflate disproportionally the share of inventor citations as most of the examiner citations listed on the parent patents became applicant citations on continuing patents.[6] On the other hand, I still keep continuing patents in the cited patents. That is, forward citations (made by future original patents) to the continuing patents are still included in the data, as I do not find evidence that a patent that cites a continuing patent is likely to cite its parent patent as well. Please refer to Appendix C for more description on the data construction process.

I divide the forward citations into four groups based on whether the citation was added by

[5] The data collection stops at 2009 to leave at least one renewal period for the observations.

[6] Nonetheless, the results reported in this paper are all robust to including citations made by continuing patents. The difference between the effects of inventor citations and examiner citations become smaller after including these citations as expected due to the second reason.

the examiner or by the inventor;[7] and whether the citation is a self citation or not (hereafter we call them "self citations" and "other citations"). Self citation is defined as a citation made by the same assignee to its own prior patent. It is usually considered to be more important signal of knowledge flow or patent quality (Hall et al., 2005; Bessen, 2008). In Table 1, we report some summary statistics of all the citations (up until the end of 2014) made to patents granted between 2001-2009 based on the type of the citation.

Table 1. Characteristics of examiner forward citations

citation type	citation lag (in years)	same-class	same-country	total number
inventor other-cite	6.6 (0.001)	40.4% (0.020%)	61.0% (0.020%)	$6.1 \cdot 10^6$
examiner other-cite	5.0 (0.001)	52.4% (0.024%)	45.6% (0.024%)	$4.4 \cdot 10^6$
inventor self-cite	5.0 (0.003)	50.8% (0.051%)	91.7% (0.028%)	$9.6 \cdot 10^5$
examiner self-cite	3.6 (0.004)	61.2% (0.071%)	91.2% (0.071%)	$4.7 \cdot 10^5$

[1] Standard errors are in parenthesis.
[2] Citation lags are the difference between the issue dates of the citing and cited patents.
[3] Technology class of a patent is defined by the 3-digit main US classification code.
[4] Country of a patent is defined as the country location of its first inventor.

Consistent with the results reported in Alcacer and Gittelman (2006) and Thompson (2006) on backward citations, roughly 40% of the forward citations are added by examiners. I confirm that there are substantial differences between citations added by inventors and examiners, as well as self citations and other citations. Therefore it's justified to make such distinction on citation types. For example, examiner citations are much more likely to transverse the national boundaries than inventor citations. 61% of the inventor citations are from the same country as the cited patent, compared to only 46% of examiner citations.[8] The results are robust to using only domestic citations and focusing on the state or metropolitan boundaries.[9] This fact is often cited as evidence that knowledge spillover is hindered by geographical boundaries (Thompson, 2006; Criscuolo and Verspagen, 2008). On the other hand, examiner citations are on average about 25% more likely to come from the same 3-digit technology class than inventor citations. It seems to suggest that examiners, while having better awareness of a broader geographical space than inventors, are at the same time more restricted to their technology expertise.

In terms of the timing of the citations, I find that, not surprisingly, self citations arrive earlier (about 1.4 years earlier) than other party citations. Interestingly, examiner citations on average arrive around 1.5 years earlier than applicant citations. Considering that 40% of the total citations are added by examiners, the first few citations to a patent, which are perhaps of special importance from a knowledge spillover perspective (Griffith et al., 2011), are much more likely to be added by examiners. In fact, among patents in our sample that received at least one forward citation, roughly 70% of them were first cited by examiners.

[7] Here we use the loose term "inventor citation" to represent all citations that are not added by examiners. It should be noted that some of these citations could be added by the attorneys of the patentee in the preparation of patent applications.

[8] It's not surprising that most of the self citations are from the same country.

[9] Not reported here, the results are available upon request.

Intriguingly, I find that roughly one third of the self-citations are added by examiners. This could pose a challenging view on the idea of patentees strategically omitting citations: firms should have strong incentive to self cite their inventions to build a thicket of patents. At least they should have less incentive to "strategically withheld" self-citations than citations from other parties. However, the proportion of "withheld" self-citations is about the same as the "withheld known citations" reported in Lampe (2012), who excludes self citations but only considers patents that were cited by the assignee previously (thus considered to be "known" to the assignee).

4. Predicting Patent Renewal Using Forward Citations

In US, patent maintenance fees are due three times during the life of a patent. The fees can be paid without a surcharge at 3-3.5, 7-7.5 and 11-11.5 years after the issue date; or with a late fee surcharge at 3.5-4, 7.5-8 and 11.5-12 years after the date of issuance. Otherwise, the patent lapses at 4, 8, or 12 years after the issue date. We use forward citations within the first 5 years interval after the grant of the patent to predict the renewal of the patent. For example, for a patent that was issued on Feb 4th, 2001, we count the number of patents issued between Feb 4th, 2001 and Feb 4th, 2006 that cited this focal patent.[10] This number of forward citations is used to predict the renewals of the focal patent, which usually occurs at 3.5-4, 7.5-8 and 11.5-12 years after the date of issuance.

The rational for using a window that covers one more year beyond the first renewal date is that there is a lag between filing and grant of a patent (so-called "grant lag", averaging about 3 years), while citations could be added at the time of the filing or during the examination process (by examiners and by inventors through interaction with the patent office, Jaffe et al., 2000).[11] I keep the time window for forward citations to be the first 5 years of grant for two reasons. Firstly, business managers would like to get an estimate of the value of the patent as early as possible. It would be of little use, for example, even if the number of forward citations in year 11 is very predictive of patent's third renewal. Therefore I want to test whether the forward citations in the early years are good predictors of later patent renewals. Secondly, the first few citations are often considered to be the most important as the knowledge embedded in patents is tacit in the early stage. Therefore it makes natural sense to focus on these early citations. I plotted a schematic diagram of how the study is conducted in Figure 1, using the second renewal as an example.

Since we have patents that were granted until 2014, we can study the renewal behavior for patents issued between 2001 and 2009. The first, second and third renewals are all covered in the data for patents issued in 2001; the first and second renewals are included for patents issued in 2002-2005; and only the first renewals are included for patents issued in 2006-2009. Overall we have about 1.5 million patents that are issued from 2001 to 2009. Around 12% of them did not have any forward citations by the end of 2014. Summary statistics of the number of forward citations received within the first 5 years of the patent issuance and their

[10]For the easiness of programming, we consider a year as exactly 365 days. So a 5-year window would cover 5×365 days after the issue date of the patent.

[11]Since we do not know exactly when the citations are added, it's not justified to use the filing date of the citing patent as the citation date either. Moreover, using filing date introduces the complicated truncation problem.

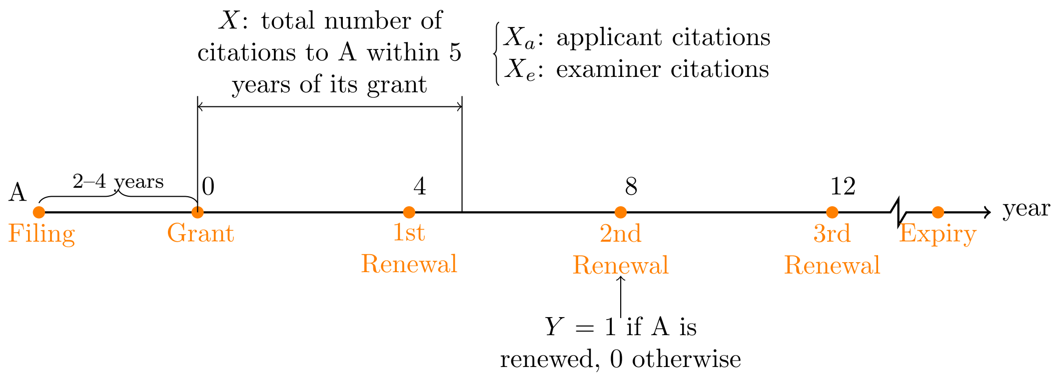

Figure 1. A schematic diagram of the correlation studies between forward citations and renewals

correlation with renewals are reported in Table 2. Because examiner citations come earlier (see citation lag in Table 1), there are in general more examiner citations within 5 years of the patent issuance than inventor citations.

Table 2. Average number of forward citations within 5 years of patent issuance

citation type	first renewal		second renewal		third renewal	
	NO	YES	NO	YES	NO	YES
applicant cite	0.77 (0.01)	1.34 (0.00)	0.63 (0.00)	1.10 (0.00)	0.70 (0.01)	1.19 (0.01)
examiner cite	1.15 (0.00)	1.64 (0.00)	1.16 (0.00)	1.70 (0.00)	1.34 (0.01)	2.00 (0.01)
self cite	0.11 (0.00)	0.36 (0.00)	0.11 (0.00)	0.29 (0.00)	0.12 (0.00)	0.29 (0.00)
examiner self cite	0.12 (0.00)	0.23 (0.00)	0.14 (0.00)	0.23 (0.00)	0.16 (0.00)	0.26 (0.00)
N	184,265	1,281,656	251,811	558,330	83,595	84,318

[1] Standard errors of the mean are in parenthesis.
[2] Citations are average count of forward citations within 5 years of the grant date of the patent.
[3] First renewal columns include all patents granted from 2001-2009; second renewal columns include patents granted from 2001-2005; and third renewal columns only include patents granted in 2001.

As illustrated in Figure 1, I use the number of forward citations to a patent in the first 5 years to predict its renewals. We observe from the table that all four types of citations seem to have predictive power over patent renewals. Patents that are renewed on average receive higher number of forward citations (of all four types) in first 5 years of its grant. To answer the question which type of citations is more predictive of patent renewals, we need to run a regression of renewals on citations and compare the coefficients on different type of citations (the so-called "horse racing" models).

4.1 Linear probability models with categorical variables

I estimate the effects of the four types of citations on the probability of renewal by the following model:

$$renewal_i \sim D_{\text{app-cite}} + D_{\text{exa-cite}} + D_{\text{self-cite}} + D_{\text{exa-self-cite}} + X_i + u_i \tag{1}$$

The model regresses patent renewal dummies on categorical variables of each type of citations.[12] It is convenient in my case because the number of citations is discrete and highly skewed. I can easily group the number of citations over a threshold into a single category and treat each value below this threshold as a separate one. The number of non-self citations is smaller than 6 in over 90%-95% of the observations, therefore we group them into 8 categories (7 categories for number of citations equalling to $0, 1, 2, \cdots, 6$, and the last category for number of citations > 6). We choose the corresponding cut-off point as 3 for self citations. The results are nonetheless very robust to the choice of other cut-off values. The models control other characteristics X_i of the focal patent, including a domestic dummy (indicating the invention conducted in the US), the grant year fixed effects (to control for the cohort effect), technology field fixed effects (3-digit US classification, to control for inherent difference across technology field in renewal tendency),[13] number of claims in the patent (a proxy to the breadth of the patent Lanjouw and Schankerman, 2001) and the grant lag (time between filing and grant, to capture patent level heterogeneity). The standard errors are clustered at the 36 sub-categories of US patent technology classes, defined by Hall et al. (2001).[14]

Let's compare the coefficients on the inventor citation categorical variables to the corresponding examiner citation categories. The results are reported in Table 3. For each type of the citation, number equalling to 0 is the omitted category. So the coefficients should be interpreted as relative to having no citations. The self citations are cut off at 3 and non-self citations at 6.

We observe that for every renewal, inventor citations have stronger association with renewal outcomes than examiner citations at all values; self citations (added by inventors) also have stronger assoication with renewals than self citations added by examiners. This is in contrast to the existing literature, especially the findings in Hegde and Sampat (2009). I'll show in Section 4.2 that if we use a simple model with citations as linear regressors we would get the reversed finding regarding the relative association of the two types of citations with renewals.

Using categorical variables, we also observe evidence of the marginally decreasing association between the number of citations and patent renewal. The first citation in general has the largest effect on renewal while the effect of following citations decreases. For example, the first inventor-added (non-self) citation is associated with 2.9 percentage point increase in the probability of the second renewal. The marginal effect is reduced to 1.7 percentage point for the second inventor citation, 1.3 percentage point for the third one, and so on and

[12]The method is widely used in the epidemiology literature to compare the predictive ability of different markers on certain disease events (for two highly cited papers, see Wilson et al., 1998; Ridker et al., 2002).

[13]For example, pharmaceutical patents in general have a longer life than software patents.

[14]Hall et al. (2001) further aggregated technology classes into 6 categories, which is more widely used in later studies for the convenience of comparison. We choose not to cluster the standard errors at this level, because these 6 categories are too broad to assume that there is considerable correlation within each class. Moreover, the statistics obtained using the clustered covariance matrix are based on asymptotics in the number of clusters (Donald and Lang, 2007; Cameron et al., 2008), which requires the number of clusters to be large. Even though t-statistics still follow a t_{g-1} distribution when the number of observations is large within each cluster, $g = 6$ would violate the full-rank condition, which requires $g - 1 > k$, where k is the number of parameters (Hansen, 2007).

Table 3. Estimation of the correlation between forward citations and renewals, when citations are included as categorical variables

citation type	first renewal	second renewal	third renewal
inv-cite			
1	0.015 (0.001)***	0.029 (0.002)***	0.030 (0.003)***
2	0.024 (0.003)***	0.046 (0.005)***	0.050 (0.004)***
3	0.027 (0.003)***	0.059 (0.005)***	0.077 (0.005)***
4	0.032 (0.004)***	0.069 (0.006)***	0.090 (0.010)***
5	0.030 (0.004)***	0.066 (0.008)***	0.104 (0.011)***
6	0.030 (0.005)***	0.068 (0.010)***	0.097 (0.015)***
> 6	0.032 (0.003)***	0.073 (0.008)***	0.110 (0.010)***
exa-cite			
1	0.011 (0.001)***	0.020 (0.002)***	0.023 (0.003)***
2	0.016 (0.001)***	0.031 (0.003)***	0.033 (0.004)***
3	0.018 (0.002)***	0.037 (0.003)***	0.052 (0.005)***
4	0.020 (0.002)***	0.039 (0.004)***	0.057 (0.009)***
5	0.020 (0.003)***	0.041 (0.005)***	0.052 (0.009)***
6	0.022 (0.003)***	0.047 (0.007)***	0.067 (0.010)***
> 6	0.023 (0.003)***	0.057 (0.008)***	0.088 (0.012)***
self-cite			
1	0.042 (0.005)***	0.081 (0.007)***	0.097 (0.010)***
2	0.040 (0.005)***	0.093 (0.008)***	0.109 (0.013)***
3	0.042 (0.003)***	0.082 (0.007)***	0.099 (0.014)***
> 3	0.046 (0.003)***	0.087 (0.005)***	0.113 (0.014)***
exa-self-cite			
1	0.031 (0.005)***	0.055 (0.007)***	0.065 (0.009)***
2	0.033 (0.008)***	0.061 (0.010)***	0.077 (0.012)***
3	0.030 (0.010)**	0.056 (0.016)***	0.059 (0.016)***
> 3	0.027 (0.009)**	0.039 (0.013)***	0.047 (0.020)**
log(claims)	0.026 (0.002)***	0.024 (0.002)***	0.018 (0.003)***
log(grant lag)	-0.012 (0.002)***	0.003 (0.003)	-0.003 (0.005)
Constant	0.847 (0.019)***	0.527 (0.032)***	0.325 (0.041)***
N	1,465,920	810,141	167,913

[1] Standard errors clustered at 36 technology categories are in parentheses. All models include issue year fixed effects and 3-digit US classification fixed effects.

[2] * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

so forth.[15]

In order to compare the coefficients more illustratively, we plot the estimated effect of the four types of citations on the probability of patent renewal in Figure 2 (standard errors are omitted since they are generally very small). We confirm that inventor citations have

[15]This monotonically decreasing relation may not always hold in the estimated results. Moreover, the marginal effect becomes smaller and less distinguishable from 0 when the number of citations becomes large.

a larger association with renewals at every value. We also observe that at most values, especially at lower values where the majority of the observations fall and the estimates are more accurate, inventor citations have higher marginal association with patent renewal than examiner citations, conditional on having the same number of citations.

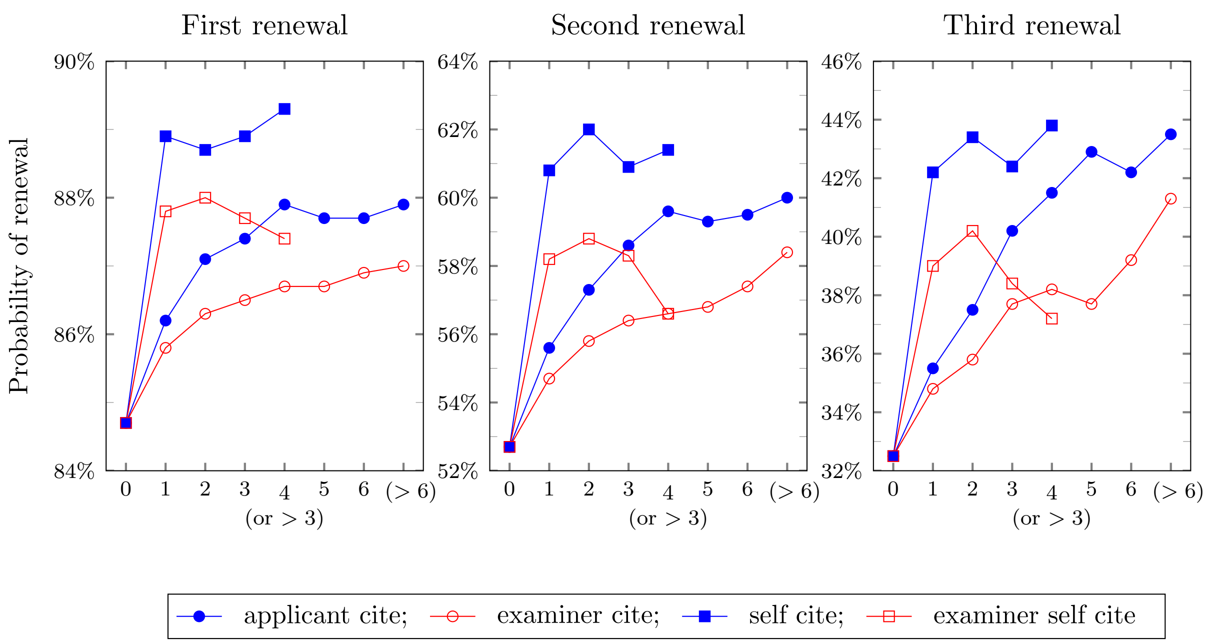

Figure 2. The relation between the probability of renewal and the number of citations

We make a few interesting observations regarding the self citations. Firstly, consistent with literature, self citations generally have stronger association with patent renewal than other citations. However, the marginal effect of self citations seems mainly fall with the first citation. In other words, whether the patent was self cited tells a lot about the probability of its renewal, while how many times it was self cited does not make a big difference. This is partly because of the relatively few number of self-citations a patent attains. Moreover, I speculate that, since self-citations indicate follow-on research, having a self-citations indicates the importance of the patent to the assignee; while having more self-citations perhaps indicates the development along the same flow of research, which may not add more value to the focal cited patent.

Therefore, the association between forward self citations with renewals reflects the overall value of a line of research, not specifically the importance of the cited patent *per se*. To test this hypothesis, I include two extra firm characteristics in Model 1: $log(size)$ and $log(patent_t)$, where $size$ is the total number of patents the owner acquired over the study period, which measures firms' overall innovation capacity (also the available pool of self citing patents);and $patent_t$ is the total number of patents the owner acquired during the year of the cited patent, which measures firms' R&D intensity in that given year. The resulting estimated marginal associations are plotted in Figure 3. The base renewal probabilities are normalized to be the same as in Figure 2 for easy comparison.

The findings are consistent with my hypothesis. After including the two characteristics that control for the patentee's overall innovation capacity and R&D intensity, the marginal

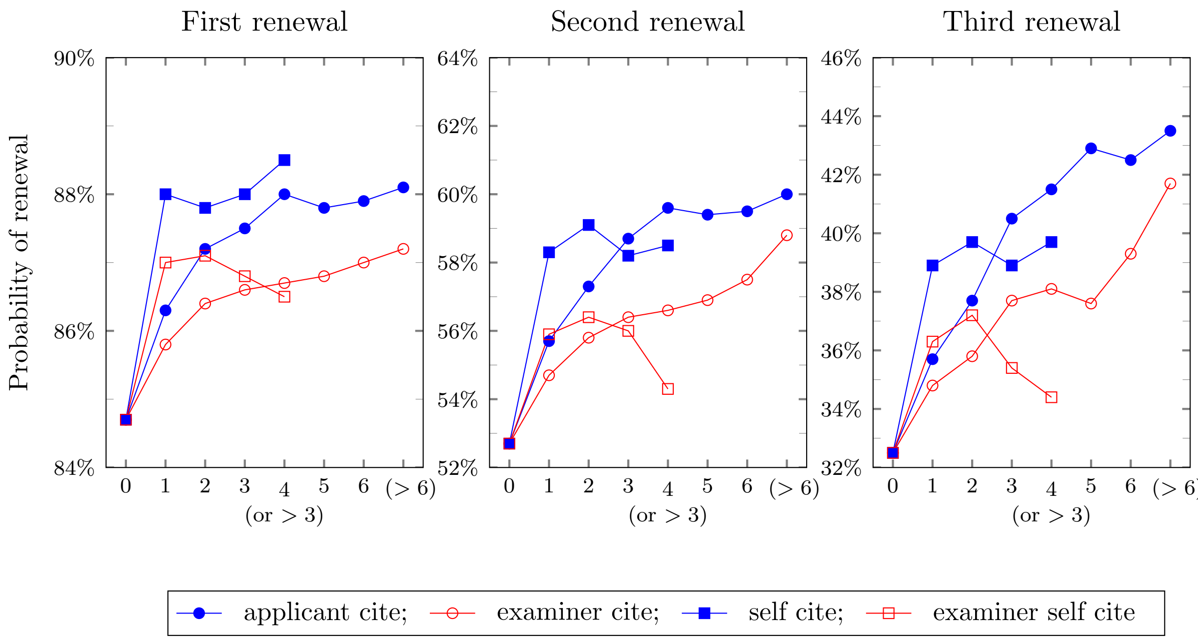

Figure 3. The relation between the probability of renewal and the number of citations

association between self citation and renewals are greatly reduced. The first 1 to 2 self citations are still more strongly associated with renewals than non-self citations, but the difference is much smaller. Since extra self citations do not add more information to the value of the cited patent, the marginal association of non-self citations to renewals quickly surpass that of self citations when there are more citations. On the other hand, the estimated coefficients for non-self citations remain almost unchanged. Therefore, we take this as evidence that self citations are strongly predictive of the renewal of the cited patent. However, self citations indicate more of a continued line of development, rather than the specific importance of the cited patent. For this reason, self citations can be used as important indicators of firms' R&D activities in different projects.

In the following section, I show that how modeling the relation between citations and renewals as linear can go wrong in the "horse-racing" models. In these simple linear models, the order of the marginal effects of inventor citations and examiner citations would be completely reversed from the findings in this section. I also explore the underlying factors that lead to this apparent reversal of relations. I show that this could occur when the relations between the independent variables and the explanatory variable is marginally decreasing, while at the same time the independent variables have very different support in their distribution. In Appendix A, I compare the results, using Monte Carlo simulation, of the model with categorical variables with other potential fixes: simple linear models cutting off the tail of the distribution of the data, logit and linear-log models. I show that, under the data structure similar to this study, models that include citation variables as categorical variables produce the most robust and accurate estimates.

4.2 Linear probability models with linear variables (the wrong way to go)

When we are given the problem of evaluating the relative association of the two types of citations with renewals, it is tempting and convenient to model the relation as linear. In this section I show the results produced using such models and argue that the results are misleading. The results from a linear probability model that regress renewal dummies on the number of the four types of citations together with other control variables are reported in Table 4.

Table 4. Estimation of the correlation between forward citations and renewals, when citations enter as linear variables

citation type	first renewal	second renewal	third renewal
inv-other	0.001 (0.000)**	0.005 (0.000)***	0.008 (0.001)***
exa-other	0.003 (0.000)***	0.007 (0.001)***	0.010 (0.002)*
inv-self	0.003 (0.001)***	0.006 (0.003)**	0.013 (0.002)**
exa-self	0.017 (0.003)***	0.031 (0.005)***	0.033 (0.005)***
log(claims)	0.028 (0.002)***	0.045 (0.002)***	0.044 (0.003)***
log(grant lag)	-0.010 (0.002)***	-0.001 (0.005)	-0.002 (0.006)
Constant	0.847 (0.020)***	0.539 (0.032)***	0.342 (0.040)***
N	1,465,920	810,141	167,913

[1] Standard errors clustered at 36 technology categories are in parentheses. All models include issue year fixed effects and 3-digit US classification fixed effects.

[2] * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The results show that all four types of citations have significant association with renewals, and self citations have stronger association with renewals than other citations. However, examiner forward citations are shown to have much stronger effect on whether a patent will be renewed than inventor forward citations. One more examiner-other citation (non-self citation) is associated with 0.3-0.5 percentage point increase in renewal probability. The magnitude is about 2-3 times that of those for inventor citations. The contrast is even more drastic for self citations: one more examiner-self citation is associated with 1.7-3.6 percentage point increase in renewal probability, while the corresponding change associated with one more inventor-self citation is only 0.3-0.4. These numbers, considered together, are broadly consistent with findings by Hegde and Sampat (2009).

The results are problematic because firstly and obviously, the true relation between citations and renewal is not linear. As the value of renewal is constrained to be a binary variable, it's natural the relation is marginally decreasing. Normally this may not be a concern as the coefficients can still be interpreted as the average marginal effect. It, however, becomes a problem if we try to compare two variables that have very different distributions (mainly very different supports). The average marginal effect for the variable that has a larger support would be made lower simply because the larger values have smaller marginal effect. It's then possible the variable that has larger marginal effect everywhere turn out to haver smaller average marginal effect, because it has much wider support. In this case, the estimated average marginal effects cannot be used to infer the relative importance of these variables.

For an illustrative example, please see Figure 4, where it's visually evident that x_1 has a larger marginal effect on y than x_2 at every value. However, due to the distribution difference and the marginally decreasing relationship between x's and y, a linear fitting of y on x's will instead produce a larger coefficient (slope) for x_2.

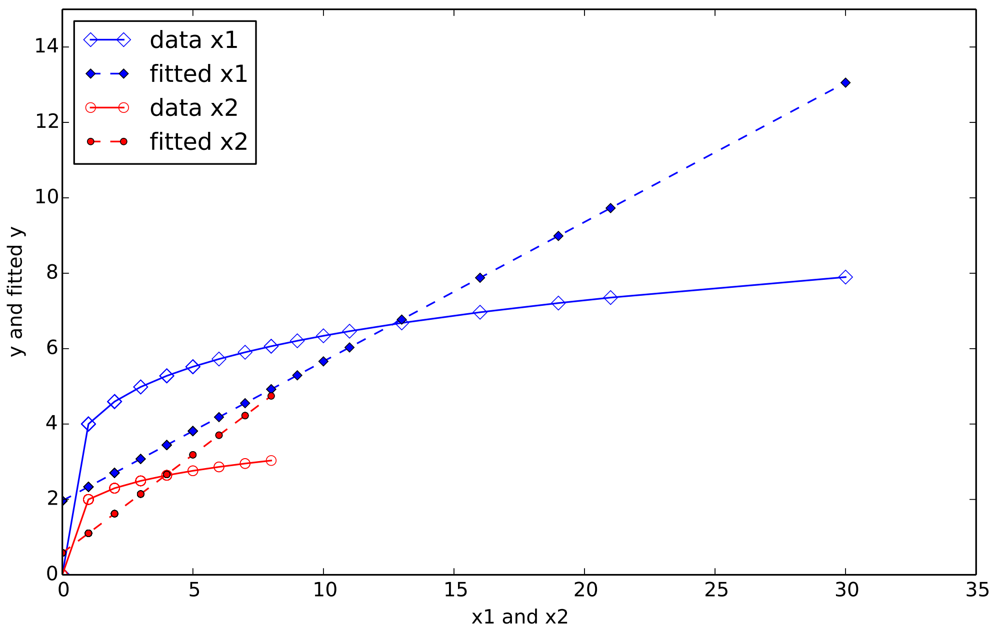

Figure 4. An illustrative example: the order of estimated average marginal effects of x_1 and x_2 is reversed from that of the the true marginal effects, when the relations between x's and y are marginally decreasing and when the distributions of x_1 and x_2 have very different support. x_1 and x_2 are negative binomial variables to imitate the skewness of citation data.

This is indeed the case for my data. Inventor citations have much larger variance and are much more skewed than that of examiner citations. Take inventor and examiner non-self citations as an example. Even though their means are close to each other (1.3 v.s. 1.6), the support of the inventor citations is $[0, 701]$ while that of examiner citations is $[0, 120]$; the standard deviation is 4.1 for inventor citations and 2.5 for examiner citations; and skewness is 31.3 for the former and 4.6 for the latter. Adding them as linear variables together would lower the coefficient of inventor citation, relative to that of the examiner citation. Because the average marginal effect for inventor citations is averaged over a much wider support, while those larger values have smaller marginal effect due to the marginally decreasing relation between citations and renewals.

4.3 Conditional marginal association of citations with renewals

We may worry the analysis in Section 4.1 obscures important interdependence of the different types of citations. For example, the marginal association of the first inventor citation with renewal if the citation is the first overall citation could be different from the association if it follows an inventor citation. I look at all the possible combinations of the first two forward citations to a patent within 5 years of its grant, and construct a fully saturated model to estimate the effect of each combination in predicting renewal. We have four types of citations: examiner citation, inventor citation (non-self citations), self citation and self citation by examiner. Therefore, altogether we have 21 different combinations for the first two citations ($21 = 4^0 + 4^1 + 4^2$).

We focus on the first two citations because the first few citations are important as over time, the information contained in a new invention gets more easily transmitted. Our analysis previously also has shown that early citations can predict later renewals quite well. We restrict the citations to within 5 years of grant for a couple reasons. We want to use forward citations to predict renewals. Since first renewal occurs at 3-4 years after grant, we want the citations to be before the first renewal. A citation lag of 5 years (issue date to issue date) is appropriate and consistent with the analysis in Section 4.1. Moreover, the construction assumes away the truncation problem. Lastly, late citations are inherently different from early citations; by restricting the attention to citations within 5 years, the comparison is less affected by the citation lags.

A break down of patents by the first two citations within 5 years of grant is shown in Table 5. Roughly 26.6% of patents do not have any forward citations within 5 years of their issuance. This group of patents is the base group we use for comparison.

We predict patent renewals based on the first few citations by fully saturating the different combinations of the first two citations (within the five year). We control the cohort effect (issuance year), technology class, and other patent characteristics X_i as in Model 1. Attention should be paid to the special type of censored data: around 38% of the patents in the sample had more than 2 citations within 5 years of their issuance. For these patents, the effect from the constructed first-two-citation combination absorbs the effect of those extra citations. Therefore, the marginal effect of the second citation would be over-estimated. Indeed, running the model without taking care of the censored data produces results where the marginal effect of the second inventor/examiner citation is greater than that of the first one, while we expect a marginally decreasing relation.

One approach would be to simply drop observations that have more than 2 citations. This is the so-called *complete case* analysis. There are two potentially serious problems with the *complete case* approach. Firstly and obviously there could be the sample selection problem, especially given that the proportion of "censored" patents is quite large in our data. Secondly and more problematically, the marginal effect of the second citation in the *complete case* regression could still be over estimated, especially when we use them to predict the second and third renewals. Because a patent that had 2 citations within the first 5 years of issuance, compared to a patent that had only 1, is also more likely to have more citations in the future, beyond these 5 years. Therefore, when we use the first two citations to predict future renewals (second and third), the marginal effect of the second citation could still absorb effects from future citations and be overestimated.[16]

Another common practice is to use a dummy variable to indicate censored data. This is

[16] I indeed find these results. The *complete case* regression produces marginally increasing relation when we try to predict second and third renewals.

Table 5. Number of patents based on the first two citation combinations within 5 years of issuance

first two citation combination	Number
no citation	390,201
inventor citation	74,757
inventor citation; inventor citation	78,981
inventor citation; examiner citation	56,743
inventor citation; self citation	6,702
inventor citation; self citation by examiner	5,302
examiner citation	181,647
examiner citation; inventor citation	112,423
examiner citation; examiner citation	314,589
examiner citation; self citation	18,707
examiner citation; self citation by examiner	31,712
self citation	16,081
self citation; inventor citation	9,537
self citation; examiner citation	14,048
self citation; self citation	25,515
self citation; self citation by examiner	5,875
self citation by examiner	29,949
self citation by examiner; inventor citation	15,483
self citation by examiner; examiner citation	44,895
self citation by examiner; self citation	10,715
self citation by examiner; self citation by examiner	22,059
N	1,465,921

the *censoring dummy* approach. However, Rigobon and Stoker (2007) point it out that the *censoring dummy* approach in general introduces bias to all estimators and is not advisable.

I realize that I actually have full information of the "censored" data: they are censored only because I want to focus on the first two citations, not because that they are not observed. Therefore I can go with the *imputation method* in dealing with censored data. My imputation is quite accurate in this case because the data are still observed: we can include a dummy variable for each extra number of citations into the model. That is, we modify the model into the following:

$$renewal_i = \alpha + \beta D_{first2citation} + \gamma_0 D_{num.\ citations=k} D_{k>2} + \gamma_1 X_i + u_i$$

The results are shown in Table 6.

The coefficients should be interpreted relative to patents that have no citations. The interpretation is straightforward and quite interesting. The marginal effect of the second citation, conditional on the first citation is of the same type, is clearly diminishing. In general, we confirm that the inventor citations have a higher marginal effect than examiner citations; self citations added by the inventor has a higher marginal effect than self citations

Table 6. Estimation of the correlation between forward citations and renewals, when citations are included as categorical variables

citation type	first renewal	second renewal	third renewal
First two-citation combinations			
inventor citation	0.021 (0.001)	0.040 (0.002)	0.043 (0.006)
inventor citation; inventor citation	0.037 (0.001)	0.075 (0.003)	0.060 (0.007)
inventor citation; examiner citation	0.034 (0.002)	0.063 (0.003)	0.072 (0.007)
inventor citation; self citation	0.082 (0.003)	0.162 (0.007)	0.180 (0.022)
inventor citation; self citation by examiner	0.076 (0.004)	0.142 (0.008)	0.154 (0.023)
examiner citation	0.016 (0.001)	0.029 (0.002)	0.027 (0.004)
examiner citation; inventor citation	0.035 (0.001)	0.061 (0.002)	0.059 (0.005)
examiner citation; examiner citation	0.027 (0.001)	0.048 (0.002)	0.048 (0.004)
examiner citation; self citation	0.076 (0.002)	0.145 (0.004)	0.159 (0.012)
examiner citation; self citation by examiner	0.061 (0.002)	0.106 (0.003)	0.109 (0.008)
self citation	0.072 (0.002)	0.123 (0.005)	0.125 (0.013)
self citation; inventor citation	0.076 (0.003)	0.139 (0.007)	0.167 (0.019)
self citation; examiner citation	0.070 (0.002)	0.128 (0.005)	0.154 (0.015)
self citation; self citation	0.072 (0.002)	0.132 (0.004)	0.138 (0.013)
self citation; self citation by examiner	0.077 (0.003)	0.146 (0.008)	0.146 (0.024)
self citation by examiner	0.058 (0.002)	0.088 (0.004)	0.092 (0.010)
self citation by examiner; inventor citation	0.060 (0.002)	0.111 (0.005)	0.113 (0.013)
self citation by examiner; examiner citation	0.049 (0.002)	0.088 (0.003)	0.092 (0.008)
self citation by examiner; self citation	0.077 (0.002)	0.136 (0.006)	0.154 (0.016)
self citation by examiner; self citation by examiner	0.062 (0.002)	0.099 (0.004)	0.101 (0.011)
total number of citations			
3	0.010 (0.001)	0.015 (0.002)	0.023 (0.005)
4	0.016 (0.001)	0.033 (0.002)	0.047 (0.005)
5	0.022 (0.001)	0.042 (0.003)	0.071 (0.006)
6	0.026 (0.001)	0.054 (0.003)	0.077 (0.007)
7	0.028 (0.002)	0.064 (0.003)	0.088 (0.008)
8	0.033 (0.002)	0.069 (0.004)	0.102 (0.009)
9	0.032 (0.002)	0.075 (0.004)	0.115 (0.010)
10+	0.044 (0.001)	0.103 (0.002)	0.164 (0.005)
domestic	0.009 (0.001)	0.019 (0.001)	0.039 (0.002)
log (number of claims)	0.029 (0.000)	0.048 (0.001)	0.048 (0.002)
log (issue lag in days)	-0.005 (0.001)	0.012 (0.001)	0.013 (0.003)
Constant	0.741 (0.004)	0.332 (0.008)	0.124 (0.019)
N	1,465,920	810,141	167,913

[1] Robust standard errors are in parentheses. All models include issue year fixed effects and six technology category effects (all insignificant).

added by the examiner.[17]

We still find that inventor self citations have the strongest predictive power in predicting renewals. Roughly consistent with our previous argument, if the inventor self citation is the first citation, any other extra citation has relatively little marginal effect. On the other hand, self citations that come later have almost additive power in predicting renewals. That is, the marginal effect (self citation | inventor/examiner citation) $\approx$ the marginal effect of the first self citation.

It's worth noting that, in general, the marginal effect of any citation is positive.[18] It suggests that, in general, follow-on patents are more likely to be complements to the cited patent, rather than substitutes. I don't find strong evidence of "creative destruction", in which the follow-on research renders the technology in the cited patents obsolete.

The overall evidence strongly suggests that examiner-added citations are much less associated with patent value, as measured by renewals, than inventor-added citations.

5. Knowledge Diffusion based on the Type of the First Citation

In this section I evaluate the relative importance of the two types of citations by investigating the knowledge diffusion process generated by different types of patents: whether a patent is first cited by an examiner or inventor. The hypothesis is, if the examiner hand-picked citations fill the missing citations from the applicants, we would expect examiner citations send stronger signal to the technical importance of the cited patent than applicant citations. The examiner-added citation sends signal to the inventors in the related industry about the technical importance of the cited patent. The signal may come as a "shock" to these firms, which could stimulate further knowledge flows from by the cited patent. Therefore, If the examiner citations represent strategical withholding or ignorance on the inventor's part, I expect a patent that is first cited by an examiner to generate more knowledge flows over its life time, compared to a patent that is first cited by an inventor.

The interest in knowledge flows traces back to the Marshallian externalities on why industries are concentrated in cities (Marshall, 1920). Study of knowledge flows is difficult because knowledge flows "leave no paper trail by which they may be measured and tracked" (Krugman, 1991). The patent citation data help solve the question because a patent citation indicates the citing patent builds upon some claims in the cited patent. In this section, I use the number of forward citations to a focal patent as a measure of the knowledge flows generated by this patent.

In order to investigate the knowledge diffusion process, I construct a model to estimate the age profile of patent citations. The model follows Mehta et al. (2010), by exploiting the relatively exogenous grant lag to identify the cohort-age-period effect. The model I use is as

[17] There are only a few exceptions. In predicting the third renewal, the marginal effect (inventor citation | inventor citation) = 0.017 < (examiner citation | inventor citation) = 0.029. Also, when the first citation is an inventor self citation, the marginal effect of an extra inventor self citation is smaller than that of an self citation added by examiner. However, in all the above mentioned exceptions, the difference is not statistically significant.

[18] This is not without exceptions. When the first citation is a self citation, the marginal effect of the next examiner citation can be 0 or even negative (though not that significant).

follows:

$$C_{it} = \alpha_c + \alpha_p + \alpha_{age} + \beta_1 D_{\text{exa-first-cite}} + \beta_2 D_{\text{exa-first-cite}} \cdot \alpha_{age} + \gamma_0 X_i + u_i t \tag{2}$$

The model regresses the yearly number of forward citations C_{it} on the cohort dummies α_c (defined as the filing year of the focal patent), period dummies α_p (defined as the grant year of the citing patent) and citation age α_{age} (defined as the difference between the grant year of the citing patent and that of the focal patent),[19] and other patent level characteristics similar to Model 1. I also include a dummy indicating the patent is first cited by an examiner and the interaction term between this dummy and citation age. The coefficients of interest are the coefficients on citation age dummies and β_1, β_2, so that I can compare the citation age profiles for the two types of patents. The estimated results are reported in Table 7 in two columns, where the second column contains the interaction terms.

Similar to Jaffe and Trajtenberg (2002), I find that inventor citations peak early (roughly 4-5 years after grant) and then started to decline. This is consistent with a technology diffusion and decay process. By including a full set of interaction terms to allow the effect to vary across age, I can estimate the citation age profiles for these two groups of patents separately. I find that a patent that is first cited by an examiner on average generate fewer forward citations (knowledge flows) compared to a patent that is first cited by an inventor. The diminishing effect related to being first cited by an examiner is mostly evident in the early years. For a better illustration, I plot the citation age profiles for the two types of patents in Figure 5. The number of citations in the first year for patents that are first cited by other inventors is set to be 0.5. I find that examiner first cited patents turn out to be less influential and generate less knowledge flows, compared to patents that are first cited by other inventors. That is, I don't find evidence to support the hypothesis that examiner-added citations send stronger signal about the importance of the cited patent to the industry, than typical inventor citations.

The knowledge flows generated by the focal patent speak more about the technical importance of the patent. On a similar note, in Appendix B, I show that inventor citations are not only more predictive to patent renewal than examiner citations, at the patent publication stage, they are also more predictive of patent grant. The finding is even surprising as examiners rely on their prior art search in the examination process Cotropia et al. (2013), it would be reasonable to expect that patent applications that are cited by examiners to have some "blocking" power and therefore more likely to be granted. The findings suggest that even when we focus on the technical quality of the patent, inventors citations have higher predictive power than examiner citations. Taken together, the findings suggest favorable consideration be made to inventor citations, rather than examiner citations, when they are used to evaluate the value (either commercial or technical) of the cited patents.

[19]One advantage of defining citation age by the difference in grant year is that I do not have to worry about negative citation ages, as very few citations have negative ages.

Table 7. Estimation of the age profile of forward citations

	citation-year	examiner-first-cited × citation-year
citation age		
1	base	
2	0.088 (0.012)***	-0.094 (0.009)***
3	0.198 (0.022)***	-0.157 (0.015)***
4	0.276 (0.030)***	-0.172 (0.020)***
5	0.320 (0.038)***	-0.169 (0.026)***
6	0.307 (0.046)***	-0.140 (0.033)***
7	0.283 (0.059)***	-0.117 (0.043)**
8	0.225 (0.070)***	-0.094 (0.048)*
9	0.138 (0.081)	-0.072 (0.050)
10	0.067 (0.089)	-0.053 (0.050)
11	-0.034 (0.091)	-0.022 (0.043)
12	-0.134 (0.089)	0.001 (0.041)
13	-0.204 (0.095)**	0.010 (0.038)
Category		
Chemical	-0.039 (0.041)	
Cmp&Cmm	0.244 (0.067)***	
Drgs&Med	0.486 (0.305)	
Elec&Elec	0.046 (0.053)	
Mechanical	-0.080 (0.044)*	
Others	base	
examiner-first-cite	-0.065 (0.039)	
constant	-0.306 (0.149)**	
N	11,380,618	

[1] Standard errors clustered at 36 technology categories are in parentheses. All models include issue year fixed effects and 3-digit US classification fixed effects.
[2] * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

6. Conclusion

By comparing the correlation between examiner citation and inventor citation with patent renewal, I show that citations added by inventors have a much stronger association to patent renewal than those added by examiners. Moreover, I find evidence, based on the first citation to a patent, that inventor citation sends a stronger signal to the value or technical importance of this focal patent than examiner citation. Patents that are first cited by inventors generate more forward citations over their life time. Therefore, contrary to existing consensus, examiner citations are less good indicators of patent value, compared to inventor citations.

I also find strong evidence that self citations are very strong predictors of patent values. This is not surprising. Self citations likely reflect more complementary research, whereas third party citations are more likely to reflect some substitute inventions. However, caution

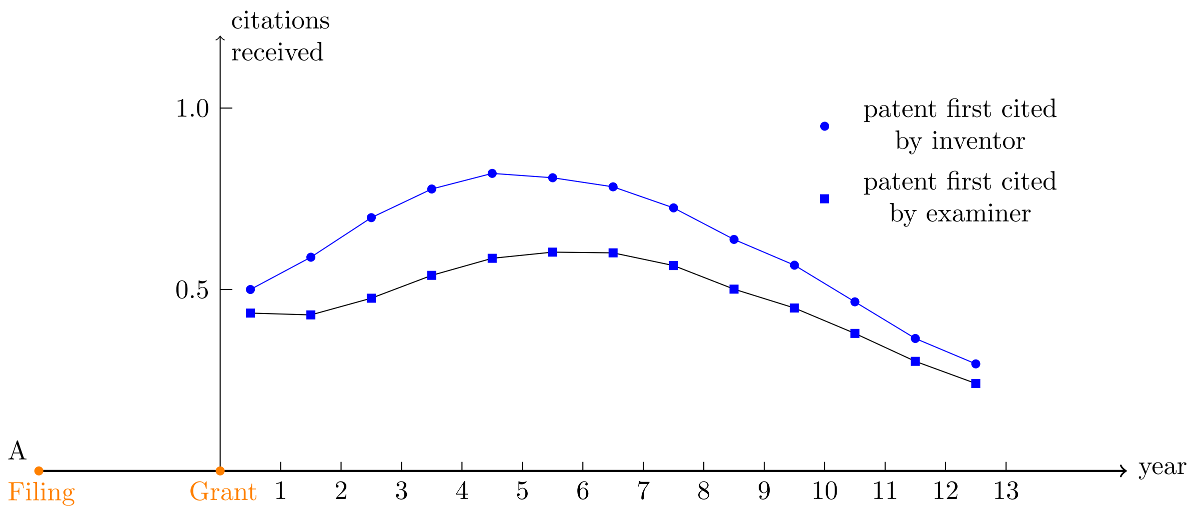

Figure 5. Patent that is first cited by an applicant generates more knowledge flow over time compared to a patent that was first cited by an examiner

should be used in interpreting the marginal association of self citations with patent value. The associated effect of self citations potentially capture the importance of a line of research, rather than the pure importance of the specific "cited patent".

The findings suggest that inventors and examiners may hold systematically different views regarding the relevant prior art. When the citations are used for the purpose of evaluating the patent value, inventor citations should be given more consideration. Using citations in a gross way without considering the difference in inventor and examiner citations could be noisy signals. Research that focuses on the first few citations needs to take extra care, as those citations are more likely examiner citations. The findings also have policy implications regarding the practice of the patent office. The "duty of candor" requirement in US patent office seems to be an efficient instrument to elicit information regarding patent value and knowledge flows. Moreover, I do not find much evidence that patent examiners are able to introduce important findings into the knowledge flow, as was hoped by some economists and policy makers. Lastly, in general, I find follow-on patents are complement to the cited patent, rather than substitute. Knowledge disembedded in a patent that is captured by a citing patent is in general positively correlated with the value of the cited patent. I do not find strong evidence of "creative destruction".

A. Monte Carlo Simulations of Different Regression Models

In this Appendix I simulate two variables (*inv* and *exa*) that behavior qualitatively similarly to our observed inventor and examiner citations, and I assume a marginally decreasing relationship between them and the patent value (*value*). In Appendix A.1 I assume the patent value is observed so that I can use the continuous patent value as the dependant variable directly. In Appendix A.2 I modify this assumption to align with the actual situation: only the renewal decision (*renewal*) is observed, the probability of which is positively related to patent value.

The data are generated such that *inv* has a larger marginal effect on *value* than *exa* at every possible value. We show that under the two conditions:

- The effects of both *inv* and *exa* on *quality* are marginally decreasing.
- The distributions of *inv* has much larger variance and support than that of *exa*.

It would be possible to have the estimated average marginal effect of *exa* much larger than that of *inv*, if both variables are included linearly into a linear probability model. We show that common fixes to such problems, including logit models and dropping "outlier" observations, do not produce satisfactory results. The linear-log model, which is often used when the effect of independent variable on explanatory variable is marginally decreasing, shows improvement and can serve as a quick diagnosis of problems. Linear probability models with independent variables included as categorical variables (hereafter the "categorical variables") produce the most accurate and robust results. The general conclusion does not depend on whether the dependant variable is bounded or not.

A.1 Patent value is observed

Firstly I lay out the basic set up of the model. The number of inventor citations (*inv*) and examiner citations (*exa*) are constructed by the following procedure:

$$inv = x_1 + z$$
$$exa = x_2 + z$$

where

$$x_1 \sim Negative\ Binomial(0.2, 0.1)$$
$$x_2 \sim Negative\ Binomial(0.8, 0.5)$$
$$z \sim Negative\ Binomial(0.5, 0.35)$$

i.e., they are the sum of two negative binomial variables, and they are correlated. Negative Binomial variables are chosen so that both *inv* and *exa* are highly skewed and take only non-negative integer values. We set the "number of failure" parameters (the first parameter) to be less than 1, so that they have mode of 0, which is consistent with the observed data. The "probability of success" parameters (the second parameter) are chosen so that the resulting distributions and correlation between *inv* and *exa* are qualitatively similar to the observed inventor and examiner citations, which I use non-self citations during the second window as the comparison group. We show some summary statistics of the simulated *inv* and *exa* with those of the observed data in Table 8. The statistics for *inv* and *exa* are computed by the means from $r = 1000$ simulation, with $n = 10000$ observations in each simulation.

Table 8. Comparison of summary statistics of the simulated variables with the observed inventor and examiner citations in windows 2

statistics	app-other	exa-other	*ratio*	*inv*	*exa*	*ratio*
max	580	118	4.9	64.0 (11.0)	20.1 (2.8)	3.2
median	1	1	1	1.0 (0.0)	1.0 (0.0)	1
mean	2.6	1.6	1.6	2.7 (0.0)	1.7 (0.0)	1.6
variance	42.5	6.9	6.2	20.6 (1.0)	4.2 (0.1)	4.9
skewness	10.6	5.6	1.9	3.8 (0.3)	2.0 (0.1)	1.9
kurtosis	328.1	79.8	4.1	22.7 (5.0)	5.9 (1.0)	3.8
correlation	0.4			0.3 (0.0)		

[1] Standard deviation of estimated statistics are in parenthesis.
[2] The *ratio* columns are computed as the ratio of the (mean) statistics from previous two columns.

As can be seen, the distributions of the actual citations are still more skewed than our simulated data. Nonetheless, the relative magnitude of statistics between our simulated *inv* and *exa* generally falls on a similar scale as that of the observed data.

Lastly, the patent quality is constructed as a marginally decreasing function of *inv* and *exa*, in which *inv* has larger marginal effect on *value* than *exa* at every point of support:

$$value = -4.2 + 4 \times inv^{0.1} + 2 \times exa^{0.1} + u$$

where

$$u \sim Normal(0, 1)$$

We estimate different models of *value* on *inv* and *exa*, with $n = 10000$ observations in each simulation. The models include simple linear models, linear-log models to take into account the marginally decreasing effect of citations on patent value,[20] linear models dropping the right tail of the citations (cut off at 10 and 5 respectively), and lastly, the categorical models. We repeat the process for $r = 1000$ times and report the average and standard deviation of the estimated coefficients for each model in Table 9.

We show that the relative magnitude of the estimated average marginal effects from a simple linear regression could give rise to very misleading interpretation. From the column “linear model”, the effect of *exa* is more than twice as large as that of *inv*, while quite contrarily, it’s clear from the patent value construction function that the inventor citation is twice more important than the examiner citation. Dropping observations that have more than 10 of either type of citations, which removes about 6% of the observations, still produces the wrong relative magnitude (column “linear, cutoff at 10”). Further reducing the cut-off threshold to 5 helps, at a cost of removing 17% of the observations, though the resulting relative magnitude of 1.7 is still lower than the true value of 2.

[20] We use $inv + 1$ and $exa + 1$ in the dependant variables, since many of the *inv* and *exa* take value of 0.

Table 9. The simulation results of the effects of inventor and examiner citations on patent quality

citation type	linear model	linear-log	linear (cutoff at 10)	linear (cutoff at 5)	categorical
inv	0.24 (0.01)	2.04 (0.02)	0.60 (0.01)	1.14 (0.02)	
exa	0.59 (0.02)	1.68 (0.03)	0.52 (0.01)	0.68 (0.02)	
$inv_{categorical}$					
1					4.00 (0.03)
2					4.29 (0.04)
3					4.47 (0.04)
4					4.60 (0.05)
5					4.71 (0.06)
> 5					5.05 (0.03)
$exa_{categorical}$					
1					2.00 (0.03)
2					2.14 (0.03)
3					2.23 (0.04)
4					2.30 (0.05)
5					2.34 (0.06)
> 5					2.40 (0.05)
Constant	-1.66 (0.04)	-3.08 (0.03)	-2.14 (0.03)	-2.79 (0.03)	-4.20 (0.02)

[1] Standard deviation of estimated statistics are in parenthesis. All estimates are significant at 0.1% level.

Both the linear-log and categorical models show the correct relative importance of *inv* and *exa*. We compute the estimated marginal effects from these two models and compare them with the true marginal effects calculated directly from the *value* function.[21] Table 10 report the comparison of the estimated and true marginal effects for both inventor and examiner citations over $[0, 1, 2, 3, 4, 5]$. It's clear that the categorical model produce the more accurate estimates. In fact, all of the estimated marginal effects for citations in $[0, 1, 2, 3, 4, 5]$ are very close to the theoretical values. Moreover, comparing these marginal effects to the average marginal effects estimated from linear models (even with cut-off at 5), it's clear those estimates are way off from the true marginal effects. Interpretation based on the linear models thus could be unfounded.

A.2 Only renewal decisions are observed

The basic set up of the model is the same as in Appendix A.1, except now I assume patent value (*value*) is a latent variable. We only observe the renewal decisions ($renewal = 1$ if the patent is renewed). The patent is renewed with probability p, which is an increasing function of the value of the patent. For simplicity, I model p as a logistic function of patent

[21] The marginal effects for the log-linear model can be calculated from the fitted equation $value = -3.08 + 2.04log(inv+1) + 1.68log(exa+1)$. The marginal effects for the categorical model are the differences between the adjacent estimated coefficients.

Table 10. The marginal effects of inventor and examiner citations on patent value

citation	$value_{inv=i} - value_{inv=i-1}$			$value_{exa=i} - value_{exa=i-1}$		
i	theoretical	linear-log	categorical	theoretical	linear-log	categorical
1	4.000	1.415	4.002	2.000	1.161	2.001
2	0.287	0.828	0.286	0.144	0.679	0.142
3	0.177	0.587	0.180	0.089	0.482	0.087
4	0.130	0.456	0.129	0.065	0.374	0.066
5	0.104	0.372	0.109	0.052	0.305	0.042

value:

$$p = \frac{e^{value}}{1 + e^{value}}$$

and

$$renewal \sim Bernoulli(p)$$

From Appendix A.1 we can see the linear model truncating at 10 does not perform well. I drop it in the comparison in this section. Rather, I include a logit model which is usually used in predicting 0-1 outcome variables.[22] The results are reported in Table 11.

Again the coefficients from the linear model seem quite off: the average marginal effect of *exa* is estimated to be more than twice that of *inv*. The logit model also provides the wrong inference: the coefficient from *exa* is larger than that of *inv*. The linear-log model and the linear model with cut-off at 5 provide, on the other hand, coefficients in the correct order, as well as our model of choice shown in the "categorical" column.

Under the model assumption, the marginal effect of *inv* on probability of renewal also depends on the values of *exa*, and *renewal* is not related to the citations by a closed-functional form. Therefore it's not straightforward to evaluate the performance of different models by comparing these coefficients. I compute the theoretical probability of renewal for all combinations of *inv* and *exa* in $[0, 1, 2, 3, 4, 5]$, altogether 36 combinations. The probability is computed using the assumed logistic function with the error term $u = 0$. I calculate the predicted renewal probability from different models using the estimates in Table 11. These predicted probabilities can be compared to the theoretical probability to evaluate the performance of the models. The results are reported in Table 12.

It's immediately clear the predictions from the two models with linear independent variables are quite off from the theoretical probabilities. As is well known, linear probability models can produce estimates not bounded by $[0, 1]$. The linear-log model and logit models both have their own advantages in certain ranges of dependant variables. The linear-log model predicts well at small value of *exa*, while the logit model has good prediction at large values of *inv*. They produce better prediction than the categorical model in a few cases. But

[22] One note should be made when we use logit models: the coefficient (or rather the exponential of the coefficient) gives the odds ratio, not the marginal percentage change that we are interested in (sometimes termed "relative risk"). This is especially a problem in our case since the base renewal rate is quite high, which leads to large discrepancy between odds ratio and relative risk. For a comparison of these two statistics, see Zhang and Kai (1998).

Table 11. The simulation results of the effects of inventor and examiner citations on patent renewal

citation type	linear model	linear-log	logit	linear (cutoff at 5)	categorical
inv	0.03 (0.00)	0.29 (0.01)	0.39 (0.03)	0.17 (0.00)	
exa	0.07 (0.00)	0.20 (0.01)	0.49 (0.02)	0.08 (0.00)	
$inv_{categorical}$					
1					0.60 (0.01)
2					0.64 (0.01)
3					0.66 (0.02)
4					0.67 (0.02)
5					0.68 (0.02)
> 5					0.71 (0.01)
$exa_{categorical}$					
1					0.21 (0.01)
2					0.23 (0.01)
3					0.24 (0.01)
4					0.25 (0.02)
5					0.26 (0.02)
> 5					0.26 (0.02)
Constant	0.32 (0.01)	0.13 (0.01)	-1.30 (0.04)	0.17 (0.01)	-0.02 (0.00)

[1] Standard deviation of estimated statistics are in parenthesis. All estimates are significant at 0.1% level.

over the majority of the values of *inv* and *exa*, our model of choice still generate the closest prediction.

In order to make the comparison clearer and more informative, I compare the average (estimated) marginal effects over the values of *inv* and *exa* in $[0, 1, 2, 3, 4, 5]$. I regress a linear model of the probabilities on *inv* and *exa* and compare which model gives the closest point estimates to those obtained from the model with the theoretical probability.[23] The results are reported in Table 13.

The categorical models produce the closest estimates of the average marginal effect to those estimated by the theoretical probability (I call them "true" values). They are very close to the true values in both the magnitude and ratio, and even the standard errors. The linear-log model produces the second closest estimates. The logit model still provides estimate of wrong order: *exa* has a larger average marginal effect than *inv*. The linear model with cutoff at 5 produces estimates of the correct order, but the estimates are way off from the true values.

I conclude that with the two problems laid out in Appendix A, the categorical model corrects the misleading results from a simple linear model, and also outperform other commonly used models that attempt to fix the problems. Moreover, based on our simulation, an empirically quick test of the existence of the two problems is to run a linear-log model and check whether the order of the coefficient changes dramatically. If the answer is yes, then

[23] The point estimates from the models with linear regressors are of course the same as in Table 11.

Table 12. The predicted probabilities from models for different values of inventor and examiner citations

inv	*exa*	theoretical probability	linear	linear-log	logit	linear (cutoff at 5)	categorical
0	0	0.01	0.32	0.13	0.21	0.17	-0.02
0	1	0.10	0.40	0.27	0.31	0.25	0.19
0	2	0.11	0.47	0.35	0.42	0.32	0.21
0	3	0.12	0.54	0.41	0.54	0.40	0.22
0	4	0.13	0.62	0.45	0.66	0.48	0.23
0	5	0.14	0.69	0.48	0.76	0.56	0.24
1	0	0.45	0.36	0.33	0.29	0.33	0.58
1	1	0.86	0.43	0.47	0.40	0.41	0.79
1	2	0.87	0.50	0.55	0.52	0.49	0.81
1	3	0.88	0.58	0.61	0.64	0.57	0.82
1	4	0.89	0.65	0.65	0.74	0.65	0.83
1	5	0.90	0.72	0.69	0.82	0.73	0.83
2	0	0.52	0.39	0.45	0.37	0.50	0.62
2	1	0.89	0.46	0.59	0.49	0.58	0.83
2	2	0.90	0.54	0.67	0.61	0.66	0.85
2	3	0.91	0.61	0.72	0.72	0.73	0.86
2	4	0.92	0.68	0.77	0.81	0.81	0.87
2	5	0.92	0.76	0.80	0.87	0.89	0.87
3	0	0.57	0.42	0.53	0.47	0.66	0.64
3	1	0.91	0.49	0.67	0.59	0.74	0.85
3	2	0.92	0.57	0.75	0.70	0.82	0.87
3	3	0.92	0.64	0.81	0.79	0.90	0.88
3	4	0.93	0.72	0.85	0.86	0.98	0.89
3	5	0.93	0.79	0.89	0.91	1.06	0.89
4	0	0.60	0.45	0.60	0.57	0.83	0.65
4	1	0.92	0.53	0.73	0.68	0.91	0.86
4	2	0.93	0.60	0.81	0.78	0.99	0.88
4	3	0.93	0.67	0.87	0.85	1.07	0.89
4	4	0.94	0.75	0.91	0.90	1.14	0.90
4	5	0.94	0.82	0.95	0.94	1.22	0.91
5	0	0.62	0.49	0.65	0.66	0.99	0.66
5	1	0.92	0.56	0.79	0.76	1.07	0.87
5	2	0.93	0.63	0.87	0.84	1.15	0.89
5	3	0.94	0.71	0.92	0.89	1.23	0.90
5	4	0.94	0.78	0.97	0.93	1.31	0.91
5	5	0.95	0.85	1.00	0.96	1.39	0.91

care should be taken in interpreting the results obtained from linear probability models. In this case, I suggest using categorical independent variables to get the most accurate and robust estimates of the effects.

Table 13. The average marginal effect of inventor and examiner citations on predicted renewal probabilities from different models

citation type	theoretical probability	linear	linear-log	logit	linear (cutoff at 5)	categorical
inv	0.12 (0.02)	0.032	0.10 (0.01)	0.07 (0.00)	0.166	0.10 (0.02)
exa	0.05 (0.02)	0.074	0.07 (0.01)	0.09 (0.00)	0.079	0.04 (0.02)
Constant	0.31 (0.09)	0.324	0.25 (0.02)	0.27 (0.02)	0.166	0.36 (0.07)

[1] Standard errors are in parenthesis. Standard errors are 0 for models with linear independent variables so they are omitted.

B. Supportive Evidence from Patent Grant

In this section, I show that not only inventor citations are more predictive to patent renewal than examiner citations, they are also more predictive of patent grant. Here a forward citation is defined as a patent that cited a focal patent application publication and was filed before the grant/abandonment/withdrawal of the focal patent application.[24] I investigate the relative association between the number of the two types of forward citations during the application publication and the final status of the patent. The data include all published patent applications from 2001 to 2014 that have a final status (grant and abandonment are the two main categories).

I control publication lag (i.e., the number of days between publication and application dates) as there is evidence that early publication indicates more valuable patents. I also control for four variables that document the interaction between examiners and inventors during the patent examination process:

- Number of final rejections: the total number of final rejections during the patent office action.
- Number of non-final rejections: the total number of non-final rejection during the patent office action, including "Non-Final Rejection" and "Informal or Non-Responsive Amendment after Examiner Action", which is often sent out on the same day when the patent office received a response (perhaps the response is deemed informal).
- Number of restrictions: the total number of "Restriction/Election Requirement" office actions, which usually restricts the claims of the applications.
- Number of interviews: the total number of "Examiner Interview" during the examination. We believe the interview indicates a lot of effort input from the examiner's side for this patent.

The last control variable I include is the days between patent application publication and its final status date. I call this variable the "exposure", i.e., the time window that the patent application is exposed for citation. The results are reported in Table 14.

[24] Note since we do not know the exact time this citation was made, it's still possible the citation was made after the grant of the focal patent application but still cited the publication number instead of the patent number.

Table 14. Estimation of the correlation between forward citations and patent grant decisions

citation type	grant
inventor citation	
1	0.034 (0.005)***
2	0.047 (0.007)***
3	0.054 (0.007)***
4	0.062 (0.008)***
5	0.064 (0.010)***
6	0.068 (0.011)***
7	0.068 (0.009)***
8	0.064 (0.011)***
9	0.066 (0.010)***
10	0.071 (0.010)***
> 10	0.086 (0.008)***
examiner citation	
1	0.020 (0.004)***
2	0.032 (0.006)***
3	0.037 (0.006)***
4	0.037 (0.008)***
5	0.036 (0.008)***
6	0.031 (0.011)**
7	0.032 (0.010)**
8	0.027 (0.009)**
9	0.040 (0.008)***
10	0.030 (0.005)***
> 10	0.022 (0.008)*
log(exposure)	-0.011 (0.005)
log(publication lag)	-0.028 (0.003)***
No.(final rejection)	-0.059 (0.004)***
No.(interview)	0.108 (0.005)***
No.(non-final rejection)	0.014 (0.005)**
No.(restriction)	-0.020 (0.004)***
Constant	0.941 (0.029)***
N	1513890

[1] Standard errors clustered at 36 technology categories are in parentheses. All models include issue year fixed effects and 3-digit US classification fixed effects.
[2] * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

I find that the association between citations at the publication stage and patent grant is much stronger and more significant for inventor citations than for examiner citations. Since examiners rely on their own prior art in the examination process Cotropia et al. (2013), it would be reasonable to expect that patent applications that are cited by examiners to have some "blocking" power and therefore more likely to be granted. Thus the finding that inventor citations are still more associated with grant than examiner citations is truly surprising. The findings suggest that even when we focus on the technical quality of the patent, inventors citations have higher predictive power than examiner citations.

I also have a few other interesting findings. The publication lag is indeed negatively related to patent grant, but we don't think the effect is economically significant. A 10% reduction in publication lag, which is roughly equivalent to 1 month earlier publication based on the data, is only associated with 0.28 percentage point increase in grant probability. Number of interviews is positively related to patent grant: having an interview is associated with 10 percentage point in crease in the probability of grant. Number of restrictions, not surprisingly, reduces the probability of patent grant. The number of final rejections works in the opposite direction as the number of non-final rejections. Non-final rejections are usually followed by response and amendment from inventors, therefore its number perhaps indicate the effort made by inventors to obtain the patent.[25]

C. Brief Description of the Data Construction

The data contain bibliographic information and legal events on granted utility patents from 1976 to 2014. The data combine information from mainly three sources: the USPTO Bulk Download Project, Google Patent webpages, and Patent Application Information Retrieval (PAIR), together with examiner roster data from 1992 to 2014 from the patent office pursuant to a Freedom of Information Act (FOIA) request.

The bibliographic information is parsed from patent grant full text (USPTO Bulk Downloads Project), which is hosted by *Reed Technology and Information Services Inc.*[26] The text information of patents is stored in structured ASCII files (pre-2001) or concatenated XML documents (post-2001). The concatenated XML file is not immediately readable by an ordinary XML parser. A python program is developed to split the file into individual parse-able XML documents. I extracted the following information from the USPTO Bulk Download Project: the patent number, filing date and issue date, type of the patent (divisional, continuation, continuation-in-part, or original), the number of total claims and exemplary claims, the full list of cited patents, inventors' names and addresses, assignees' names and address, names of the primary and assistant examiners. This dataset (hereafter the USPTO dataset) is used as the basis for the construction of the full dataset.

The patent renewal information (fee payments at year 4, 8 and 12 and whether the patent is lapsed) is obtained from Google Patent, which stores the information of each patent as a separate webpage. The intersection of the Google patents dataset (hereafter

[25]Care should be taken to interpret this coefficient, as there is very high positive correlation between the number of final and non-final rejections (about 0.65). Therefore the coefficient is only interpreted as conditional on the number of final rejections.

[26]The data can be accessed at `www.reedtech.com`.

Google data) and USPTO dataset are constructed.[27] Missing information from USPTO data is supplemented with Google data information (mainly inventor and assignee names).

Patent application information are downloaded and parsed from the Public Patent Application Information Retrieval (PAIR) data, which is not currently available in bulk form. USPTO entered agreements with *Reed Technology and Information Services Inc* to provide the data to the public and update them on a daily basis. Each application file is compressed into a *zip* file. Since applications made after 2001 include image file wrappers, the size of the *zip* file ranges from several Megabytes to over 2 Gigabytes, which would take tremendous effort to download and read the more than 6 million patent applications. To circumvent the problem, I write a python program to dynamically read only the portion of the files that are needed.

Patent examiners are known to exhibit significant heterogeneity (Lemley and Sampat, 2012), but the recording of the examiner names in the USPTO patent full text document is notoriously erroneous. To this end, I acquired the USPTO employee directories from 1992 to 2014 via a FOIA request to help disambiguate the names. Some substantial disambiguation work on assignee names, examiner names, locations was conducted based on Damerau-Levenshtein distance and alphabetical clusters. A detailed description of the disambiguation procedure is available upon request.

I construct a concordance between patents and its cited patents, including information indicating whether the citation comes from examiner. I only include domestic utility patent citations. Forward citations are often used as indicator of the importance of a patent and are usually more useful statistics than backward citation. However, construction of forward citations are naturally more difficult since it involves the inversion of citation data. Because I am mainly interested in questions related to examiner citation, I restrict the analysis to patents granted after 2001. I can only fully tell whether the forward citations are from the examiner or not for those patents, since USPTO only started to record examiner citation to patents granted after 2001.

One complication is introduced by another reform in patent policy at 2001, which requires a patent to be published within 18 months of its application, if the patentee has filed or intend to file the patent abroad. As a result, the majority of patents filed after 2001 were published before their issuance. Published patents become prior art and therefore could be cited by later filed patents. In fact, a substantial percentage ($\sim 24\%$) of patent citations are made to published patent applications, instead of issued patents. About 40% of these citations to published patents were added even when the cited patents had already been issued. Therefore it seems inventors/examiners often choose to search the published patent applications instead of issued patents when they look for prior art. Even if some of these applications later get granted, the citation information on the citing patent would not be updated (i.e., the record on the citing patent still cites the publication number, instead of the patent number of the cited patent). Omitting these citations would lead to a serious undercount of forward citations. Therefore, in constructing the citing-cited patent pairs, I

[27] Those that are in USPTO dataset but not in the Google dataset are "withdrawn patents numbers". USPTO maintains a full list of such numbers (see `http://www.uspto.gov/patents-application-process/patent-search/withdrawn-patent-numbers`). Those that are in Google Patent but not in USPTO dataset are patents that miss from the full-text database (only around 10k utility patents, see `http://patft.uspto.gov/netahtml/PTO/help/contents.htm`).

need to update all publication numbers to patent numbers for the cited patents if they were granted.

Abstract

What Happens When Politics Intervenes in the Patent System?

by

Professor Zhen Lei
Department of Energy and Mineral Engineering
Penn State University

and

Zhen Sun
Doctor of Philosophy in Agricultrual and Resource Economics
University of California, Berkeley

and

Professor Brian D. Wright
Agricultrual and Resource Economics
University of California, Berkeley

We investigate what happens when a centralized political power intervenes in the intrinsically decentralized patent system by studying the seasonal characteristics of patent filings at China's State Intellectual Property Office (SIPO). We find a strong monthly variation in domestic filings which peak in December every year compared to foreign filings. The peak gets more significant after 2000, when China started to emphasize innovation and IP strategy in its so-called "10^{th} five-year plan". The findings suggest some of the domestic patent filings may be driven by the yearly quotas set by the local governments, rather than commercially driven. These quotas are important assessment criteria for some local officials. We conjecture that under political pressure, firms file more low quality applications or break up their inventions to get more applications. Firms may also shift patent filings from the following year to previous November and December to meet the quota. We provide evidence on these hypotheses based on three patent characteristics: the grant rate, the proportion of shared co-inventors and the number of forward citations to the granted patents. The findings suggest the quota incentive may be counter-effective.

1. Introduction

Patents have long been used as an indicator of innovative activity and technological growth (Griliches, 1990; Kortum, 1997). In 2011 China became the world leader in the number of published invention patent applications, outpacing the United States, Europe and Japan. Patent applications in China increased from 63450 in 2001 to 391177 in 2010, an annual rate of 22.6%;[28] and domestic applications have grown even faster, from 30038 to 293066, by 28.8% per year during the period.[29] The driving forces behind this China patent boom have been debated. Some observers consider the boom to be an indicator of significant and genuine strides in China's innovative capacity, resulting from China's focused efforts to promote its indigenous innovation and technological development and transform its economy from "made-in-China" to "invented/designed in China". Other observers, however, believe that the patent boom is largely due to various government incentives for patenting. They argue that these incentives induce applicants to file opportunistic applications for inventions of low patentability or low value that would otherwise have not been filed. The government incentives can be in the form of patent subsidies, patent rewards, or patent quota. Thus these observers claim that most filings in this China patent boom are so-called "junk inventions".[30]

In this paper we investigate the patent quota incentive and its effect. Firms in China seem to produce more innovative output during the latter part of the year and file patent applications before the calendar year ends.[31] We study the seasonal pattern of patent filings counts in China from 1994 to 2007. By comparing domestic filings with foreign filings at China's State Intellectual Property Office (SIPO), we find a much stronger monthly pattern of domestic filings which peaks in December every year. The peak gets more significant after 2000, when China started to emphasize innovation and IP strategy in its so-called "10$^{\text{th}}$ five-year plan". The findings suggest some of the domestic patent filings may be politically driven, rather than commercially driven. We conjecture that under political pressure, firms and individual file more low quality applications to meet the quota. Political pressure may also cause an inter-temporal shift in patent filings from the following year to previous November and December. Firms also are also likely to split and repackage their existing innovation to generate more applications. We provide evidence on the hypotheses based on two patent application characteristics: the grant rate and the proportion of shared co-inventors. We find a significantly lower grant rate for domestic filings in late year compared to other months, using foreign filings as controls. The decrease in grant rate for late year filings by domestic firms seems to get stronger after around 2000. We also find evidence that domestic patent filings in December have higher proportion of shared co-inventors from the applicant firm, which may be indicative of "patent split"—patentees break up their patents to get more ap-

[28] Growth was much slower elsewhere: 4.6% in America, 5.6% in South Korea, 3.6% in Europe and -2.7% in Japan during the same period.

[29] *Patent applications by patent office (1983-2010)*, source: WIPO Statistics Database, available at `http://www.wipo.int/ipstats/en/statistics/patents/`

[30] For example, see Economist article "Patents, yes; ideas, maybe–Innovation in China" at `http://www.economist.com/node/17257940`, and a more recent article "Valuing patents–How innovative is China" at `http://www.economist.com/news/business/21569062-valuing-patents`.

[31] "China is more innovative in the Fall", as commented by Mark Cohen, former Senior Intellectual Property Attaché at the U.S. Embassy in Beijing and the Attorney-Advisor in the Office of International Relations at USPTO.

plications (Lei et al., 2013), and we find a significant increase in this phenomenon after 2000. Pulling together the evidence, one plausible explanation of the surge in domestic patenting in December is that applications are made under political pressure to meet yearly quotas set by the local governments. These quotas are important assessment criteria for some local officials. Firms under political pressure may file some applications they consider worthless, break up inventions to get more applications, or shift applications from later months to meet temporary pressure from government. The findings have important implications to China's patent strategies.

The findings shed light on the debate of China's phenomenal performance in patent applications. To our knowledge, this is the first empirical paper to look at the quota's effect on a patent system. Previous studies investigate the growth of China patenting from other perspectives. Tests of spill-over effect from foreign direct investment (FDI) were inconclusive (Liu, 2002; Cheung and Lin, 2004; Hu and Jefferson, 2009; Girma et al., 2009). Hu and Jefferson (2009) also suggest that other factors, including an intensification of research and development (R&D), entry into WTO, and more importantly the 2000–2001 amendments to the patent law that offered stronger protection to patent holders, all of which contributed to the rise in patenting. Li (2012) and Lei et al. (2013), on the other hand, believe the patent subsidy policies at all levels of government are an important factor in the patent growth in China. Our results suggest, besides the other government incentives, the quota requirements set by local governments may have played an important role in this patent application surge.

More generally, the paper brings a new insight to the long list of discussion of using policy intervention to promote R&D growth. Financial incentives and quotas are like the two sides of the coin, it is puzzlingly that we have seen relatively few quota requirement on R&D output. Though rare, it is not unheard of. Korea's patent application data also exhibit a strong "December effect", which is likely a result of corporative quota incentive. When HP's former CEO Carly Fiorina encouraged the company to double its patenting, they eventually "tripled the rate of innovation to 11 patents a day". Many other goal based R&D policies have a similar taste to the quota system. For example, many governments (the EU, Portugal, China, Canada, to name a few) project the R&D intensity to a certain percentage of the national GDP (Carvalho 2011, which has a nice review on R&D tax incentives). This paper, to the best of our knowledge, conducts the first study of the effect of quota system on patenting.

The remainder of the paper is organized as follows. In section 2 we provide a brief review of China's patent systems and the quota requirement set by local government. Section 3 explains our data in more detail. Section 5 shows our methodology and the econometric results. Section 6 concludes.

2. Background

China started to envision its first modern patent law in July 1978, shortly after its open-door policy. After a heated debate and 25 revisions, the patent law, modeled after those of Japan and Germany, became effective on April 1^{st}, 1985. The patent law has since been amended three times, in 1992, 2001 and 2008, respectively. The first two amendments were made during the process of China's negotiation of entry into the WTO and largely aimed at harmonization of the Chinese Patent Law with those in other countries and the WTO/TRIPS principles.[32] The Patent Law was amended again in 2008, as part of China's effort to promote indigenous innovation.

There are three types of patents in China: invention patents, utility model patents and design patents. Invention patents are granted after a substantive examination of utility, novelty and non-obviousness; the other two need only a preliminary examination. Invention patents, utility models and design patents have life terms of 20 years, 10 years and 10 years, respectively. After a patent is issued, the patentee needs to pay annual renewal fees to maintain the patent right; otherwise, the patent is deemed to be abandoned. Invention patents are of the main interest to researchers due to their similarity to the most prevalent types of patents elsewhere with respect to the standard of patentability and terms of protection. The analysis in this paper will focus on the invention patents.

Since it opened to the world economy three decades ago, China's economic development has been one of "made in China", relying on its low-cost manufacturing of existing products. However, promoting innovation and technological development has long been an important theme for the Chinese government. Since around 2000 China has been making systematic efforts to strengthen its R&D and plan its indigenous innovation, in its ambitious effort to transform the manufacturing-based economy to be innovation-based. In 1999 the central government revealed the "Decisions on Strengthening Technological Innovation, High Tech Development and Industrialization by the Central Committee of Chinese Communist Party and State Council", emphasizing the importance and urgency for China to become an "innovation-oriented" society and a global leader in science and technology. The effort to strengthen China's indigenous innovation culminated in 2006, with the announcement of the "Medium to Long Term Plan for the Development of Science and Technology" strategy (MLP). The MLP aims at building up domestic R&D capability, advancing Chinese firms' innovation, and reducing China's reliance on foreign know-how and technologies. Following the MLP strategy, a web of policies have been implemented at both national and local levels.

Creation of its own intellectual property, in particular patents, has been considered by China as one major endeavor in its push for indigenous innovation push. The central government has announced multiple national plans on patenting that set ambitious and quantitative goals on patenting by domestic inventors, some of which have been labelled "mind-blowing" in nature.[33] For instance, the "National 10th Five-year Plan on Patents (2001-2005)", drafted in around 2000 and announced in 2001, specified a target of an annual growth rate of 14% for domestic patent applications between 2001 and 2006, and another target of more than

[32] The first amendment was made when China started its effort to become a member of the WTO and the second right before China entered the WTO.

[33] In a 2011 interview, David J. Kappos, director of the U.S. Patent and Trademark Office at the time, called the Chinese patent targets "mind-blowing numbers".

300,000 annual domestic application by 2005 (including all the three types of patents allowed in China). The MLP in 2006 set the goal of becoming one of the top five countries in terms of invention patents granted to its citizens by 2020. The "National Patent Development Strategy (2011-2020)", revealed in 2010, set an astounding goal of 2 million domestic patent applications annually (all the three types included) by 2015 and another goal of becoming one of the top two countries in terms of granted invention patents.

The targets set by these national plans on patenting are claimed by the central government to function as guidance and not to be binding. However, under such mentality of innovation-by-number and the legacy of top down planning, governments at lower levels (provincial, municipal and county) have accordingly, from top downward, established their own strategic plans on patenting, specifying ambitious quantitative targets on patenting from their jurisdiction and establishing a variety of incentives, including patenting subsidies to achieve this goal. These patenting targets are then included as important factors in assessing the performance of government officials and government-controlled enterprises. See Appendix A for some illustrative anecdotes of patenting targets imposed by governments from top down.

3. Data

The dataset from Chinas State Intellectual Property Office (SIPO) covers patent applications (including all the three types of patents: invention patents, utility models and design patents) that were published by SIPO from April 1985 (when the patent law was enacted in China and the SIPO was established) through December 2011.[34] We supplement the legal outcomes with information from Google Patents, which have the updated legal outcomes for patent applications at SIPO until December 2013. The dataset, for each published patent application, contains a rich set of bibliographic information including backward and forward citations, number of claims, IPC classes, inventor names, origins (and detail addresses if domestic) of applicants, legal status with dates of and information about events during patent prosecution (examination request, patent grant, etc).

Our study focuses on invention patent applications filed at SIPO between January 1994 and December 2007. Year 1994 is the start of our study period because China joined the Patent Cooperation Treaty (PCT) in 1994 and foreign applicants have since filed patent applications at SIPO either directly or through PCT filings at WIPO (World Intellectual Property Office) first. We choose year 2007 as the last year of the study period to: (1) avoid potential differences between domestic filings and foreign filings at SIPO due to the financial crisis that started in 2008; and (2) to minimize the problem of data truncation. There are two sources of data truncation. Firstly the data may be truncated due to the lag between filing and publication of an invention patent application at SIPO. In principle, the maximum time allowed between filing and publication (the "publication lag") of invention patent applications is 18 months. But in practice the lag could be longer. In fact, using a cut-off time of 19 months (considering the requirement should be based on calendar month, instead of exact date), we find roughly 8% of the applications in the data have a publication lag exceeding the required time (not including PCT filings). Nonetheless, considering the

[34]Data source: CD-ROM on published patent applications that was released by SIPO in January 2012.

our patent data are collected 6 years after the filing of the last application in our study, the truncation should be ignorable. Secondly, we want to minimize the truncation of SIPO examination outcome to the collected patent application data, due to the lag between filing and grant (the "grant lag"). This truncation would cause biases to our analysis on patent quality, because patents filed later are more likely to have missing examination outcome, and patent grant usually takes longer time than patent rejection. Figure 6 provides scatter plots of the ratio of patent applications filed in each month that have examination outcome. By the end of 2013, the truncation due to grant lag is very small, except for a few outlier months.[35]

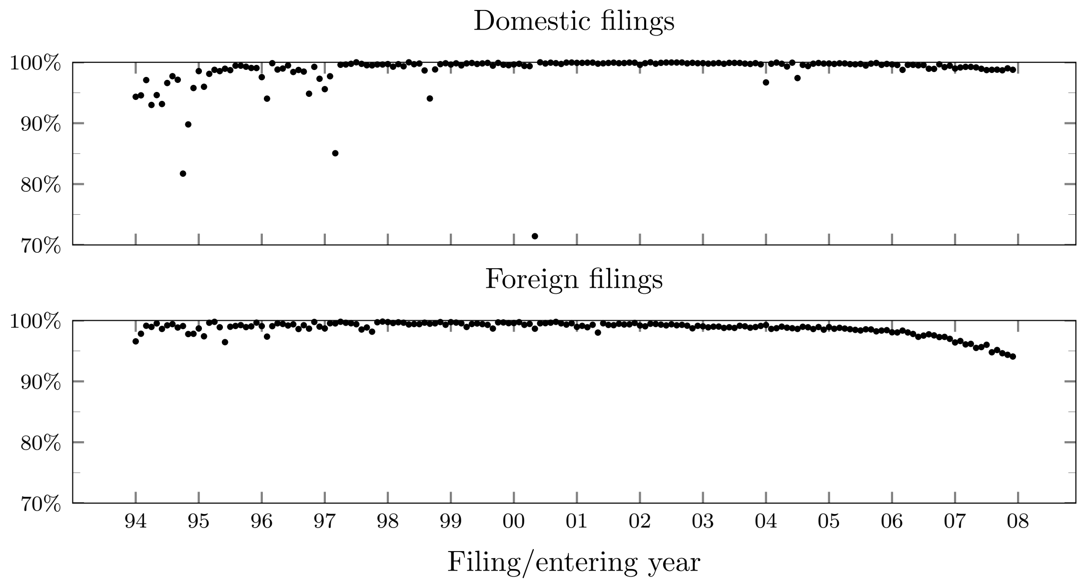

Figure 6. Ratio of patents applications that have received the examination decision from SIPO by December 2013

We separate invention patent applications into domestic filings (by domestic applicants) and foreign filings (by foreign applicants), based on the country of the first applicant.[36] We further distinguish foreign applications by the way they enter China: conventional filings (directly filed at SIPO) or PCT filings (filed at WIPO first and enter SIPO through PCT, the national phase). For a conventional foreign application, we focus on its "application date" recorded in the SIPO database, which is indeed the filing date at SIPO; while for a PCT foreign application, we focus on the date it enters SIPO, as its "application date" is the date of filing at WIPO (more accurately, at a WIPO designated Receiving Office, often the patent office in the applicants home country), which is the priority date of the PCT applications. It may take up to 30 months for a PCT application to enter China, from its WIPO filing date. Hereafter we pool the two types of foreign applications to make comparison with domestic applications.

[35]For PCT filings, we use the enter date (the date the application enters China, see discussion below) instead of the filing date at the receiving office.

[36]For the purpose of this study, we classify applications from Hong Kong, Macao and Taiwan as foreign.

When we make the comparison between domestic and foreign patent filings, we focus on invention patent applications filed by firms (including academic institutions), excluding those filed by individuals.[37] The reason is three-fold: (1) the number of foreign filings by individuals is much smaller than domestic individual flings at SIPO; (2) domestic individual filings usually have much lower quality compared to firm filings; and (3) in China, individuals are less likely to be faced with political and administrative pressure on patenting than firms. Table 15 shows a breakdown of the number of patent applications by applicant origin and types.

Table 15. Break-down by applicant type and origin of the invention patent applications (1994-2007)

	Applicant type		
Origin	Firm	Individual	Total by origin
Domestic	340,696	204,308	545,004
Foreign (conventional)	331,898	34,825	366,723
Foreign (PCT)	282,577	14,124	296,701
Total by applicant type	955,171	253,257	1,208,428

From 1994 to 2007, domestic firms filed altogether 340,696 invention patent applications at SIPO, which is comparable to the number of foreign conventional filings by firms (331,898) and PCT filings by firms (367,468). On the other hand, a much greater number of domestic applications are filed by individuals, compared to the foreign filings.

4. A Seasonal-trend Decomposition Illustration

In this section we illustrate the seasonal characteristics of log monthly application data using a seasonal-trend decomposition procedure based on loess (STL). STL is a filtering procedure for decomposing a time series into three components: trend, seasonal, and remainder, using the loess smoother (Cleveland et al., 1990). Some of the key parameters to choose for the STL operations are: n_p, the number of observations in each cycle of the seasonal component, which is 12 in our study; n_l, the smoothing parameter for the low-pass filter, which is usually taken to be the smallest odd integer greater than or equal to n_p; n_t, the smoothing parameter for the trend component, the recommended value of which is given in Cleveland et al. (1990); and most importantly, n_s, the smoothing parameter for the seasonal component. Cleveland et al. (1990) recommends n_s to be an odd integer greater than or equal to 7, in order for it to smooth the data and remove low frequencies from seasonal frequencies. The larger n_s is, the more smoothing we have for the seasonal component (pooling the values for the same month across more years). An example of STL decomposition is shown in Figure 7, using

[37]We separate the two groups by whether the applicant can be matched with one of the inventors. This proves to be remarkable way of separating the two types of applicants as most of the individual applicants are inventors themselves. We cross check applicant names by their length (Chinese people's names usually have less than or equal to 3 characters) to take account for the individuals that filed applications invented by other inventors.

$n_s = 9$ for domestic monthly invention patent applications. The results are robust at other values.[38]

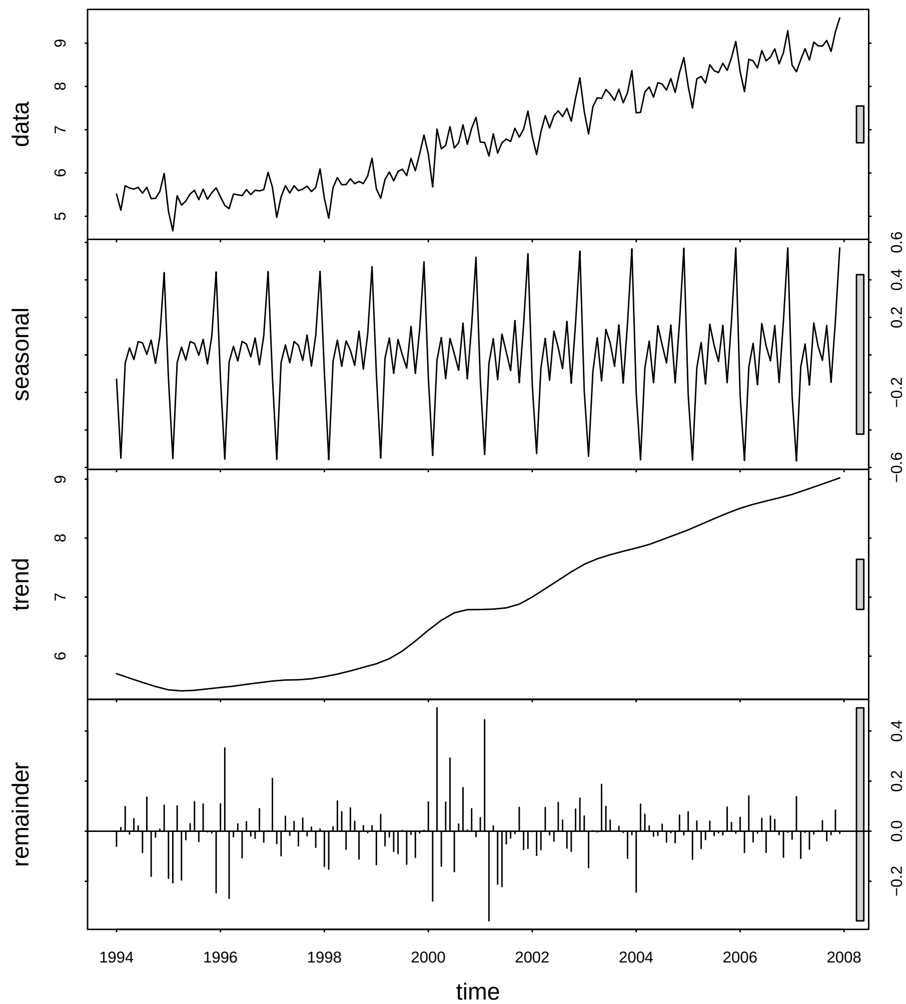

Figure 7. STL decomposition of domestic application data

The figure confirms that for applications by domestic firms, there has been a clear peaking at the end of a year (the "seasonal" panel), in particular since 2000. Though it is reasonable to expect that firms may file more patent applications when the fiscal/calendar year approaches to the end, to clear up the stock of innovations or use up some (earmarked) funds, the ever-growing end-of-year peaking since 2000 seems to suggest that Chinese firms rush to file for more patent applications at the end of a year. In order to better understand this "December effect", we compare the seasonal component of domestic firm filings to the foreign filings.

[38]The STL procedure has been implemented as a package in R.

4.1 Illustration of the "December effect" for domestic filings

We look at the seasonality of foreign applications at SIPO to compare to the seasonality of domestic applications. Figure 8 plots the standardized seasonal components of domestic firm applications compared to that of foreign firm applications at SIPO. The standardized seasonal component equals to the seasonal component (in Figure 7) divided by the average number of monthly applications (similar to the coefficient of variation), so that both serials are unitless and can be compared under the same scale.

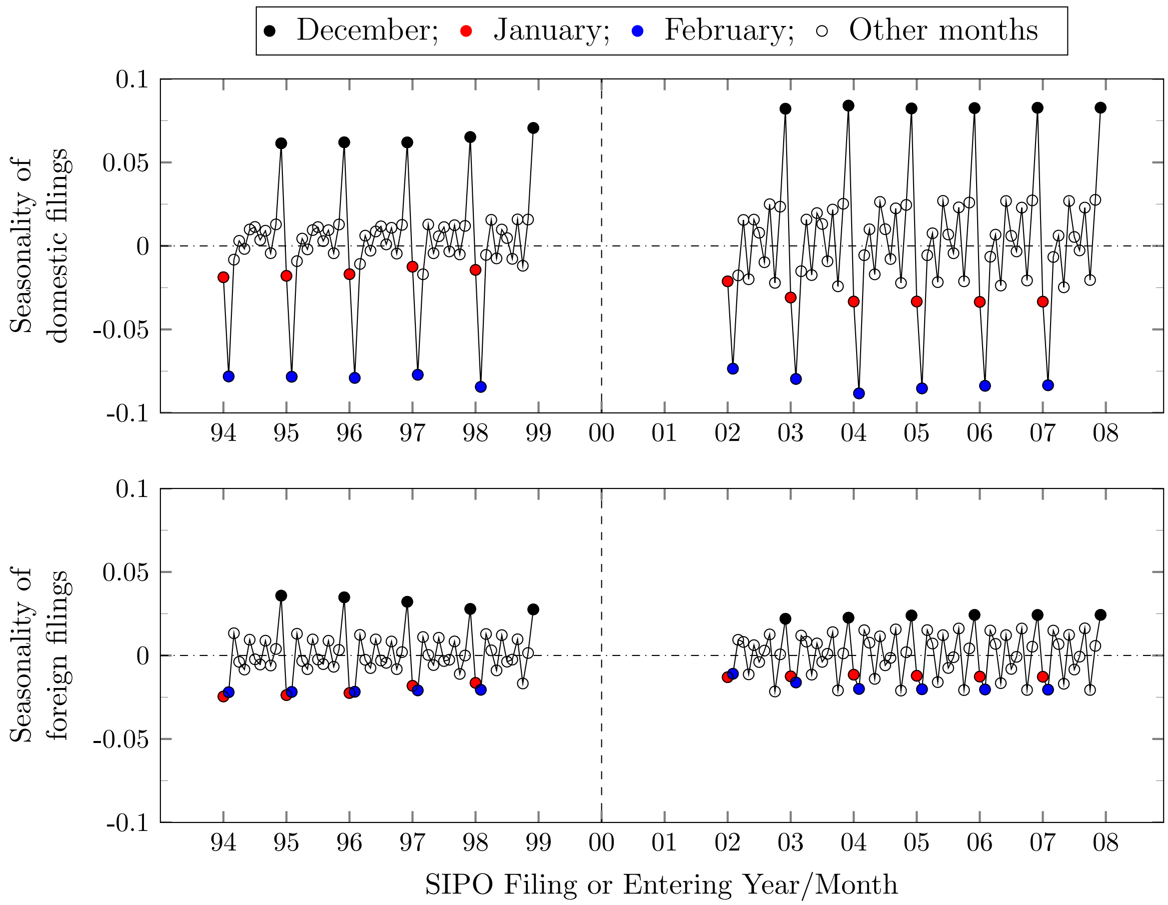

Figure 8. Comparison of seasonality between domestic and foreign firm patent filings

Both series of domestic and foreign applications show some seasonality. For example, firms file fewer patents in January and February due to holidays: the lowest number for domestic applications usually occurs in February, during which the Chinese New Year falls in; while the lowest number for foreign applications tends to occur in January, when employees leave for the new year break. However, from the comparison of the two figures, some distinctions stand out: domestic applications have a much stronger peak in December than the domestic applications after 2000. For foreign applications, though patent filings in December are usually higher than other months, which is expected due to the end of fiscal year and "holiday effect", the gap between December and other months is much smaller compared to domestic applications.

A very plausible explanation is that, after 2000, domestic patent applications near the end of the year may be politically driven, rather than commercially driven. Since 2000, when China's "10^{th} five-year plan" kicked off, China has been making systematic efforts to strengthen its R&D and indigenous innovation, in its ambitious effort to transform the manufacturing-based economy to be innovation-based. Consequently there has been an increasing emphasis on patent applications as a criterion to assess local government's achievement after 2000. To meet the target on patenting, many local governments have set goals for patent applications in each year, i.e., they need to file a certain number of patent applications in a given year. These numbers can be important for assessment of the local officials. They may check the progress periodically. If as December approaches it appears that the goals will not be met, local governments may use political pressure or financial incentives to motivate firms in their locality to increase their applications.

Since foreign firms may observe different holidays from domestic firms (January is not a holiday month for domestic firms), it's possible that foreign firms work less during December compared to domestic firms. The "holiday effect" for foreign firms may be offset by the less working days. Moreover, it's possible the "December effect" was caused by some policies that only target domestic applicants (for example, financial incentives to R&D and patenting). To make a stronger case of our story, we compare the seasonal patterns of domestic firm applications to domestic individual applications. The rational is that individuals follow the same holidays as the firms, and they should be less likely to respond to political/administrative pressure than firms; while they should both be responsive to other policies that target domestic applicants. Moreover, the total number of applications from domestic firms and individuals are on a similar scale, which makes their comparison more informative.

We plot the standardized seasonality component of applications filed by firms and individuals in Figure 9. The plots roughly agree with our explanation. Before 2000, the December application peak was not very evident for both types of applications. After around 2000, the individual applications show some December peaks, which may be due to the increasing financial support for patenting which was implemented since around 2000.[39] The financial incentives can also have a "December effect" to patentees. The budget for patent subsidy and award could be rolled out on a yearly basis such that there is incentive for innovators to claim the benefit before the new fiscal year. Moreover, filing late in the year means less discounting, as the financial reward can usually be claimed during the next year.[40] However, when we compare the firm filings with individual filings, it's clear the former has a much more evident end-of-year hook, which apparently cannot be accounted for by patent subsidy and reward policies. Figure 9 provides us with more support that the "December effect" could be caused by political and administrative pressure.

We also observe that as the domestic firm application peak in December grows more

[39] Shanghai implemented China's first patent subsidy policy in 1999. Almost all provinces have had some subsidy policies in place since around 2003, and in addition many cities have their own subsidies for patent applications. The majority of the subsidies compensate the applicants for the application fees; while others also have a prize awarded to the granted patents.

[40] Another possible explanation for the "December effect" for individual innovators is that the fact that many individuals work in firms but file patent applications by themselves. This explanation does not change our interpretation.

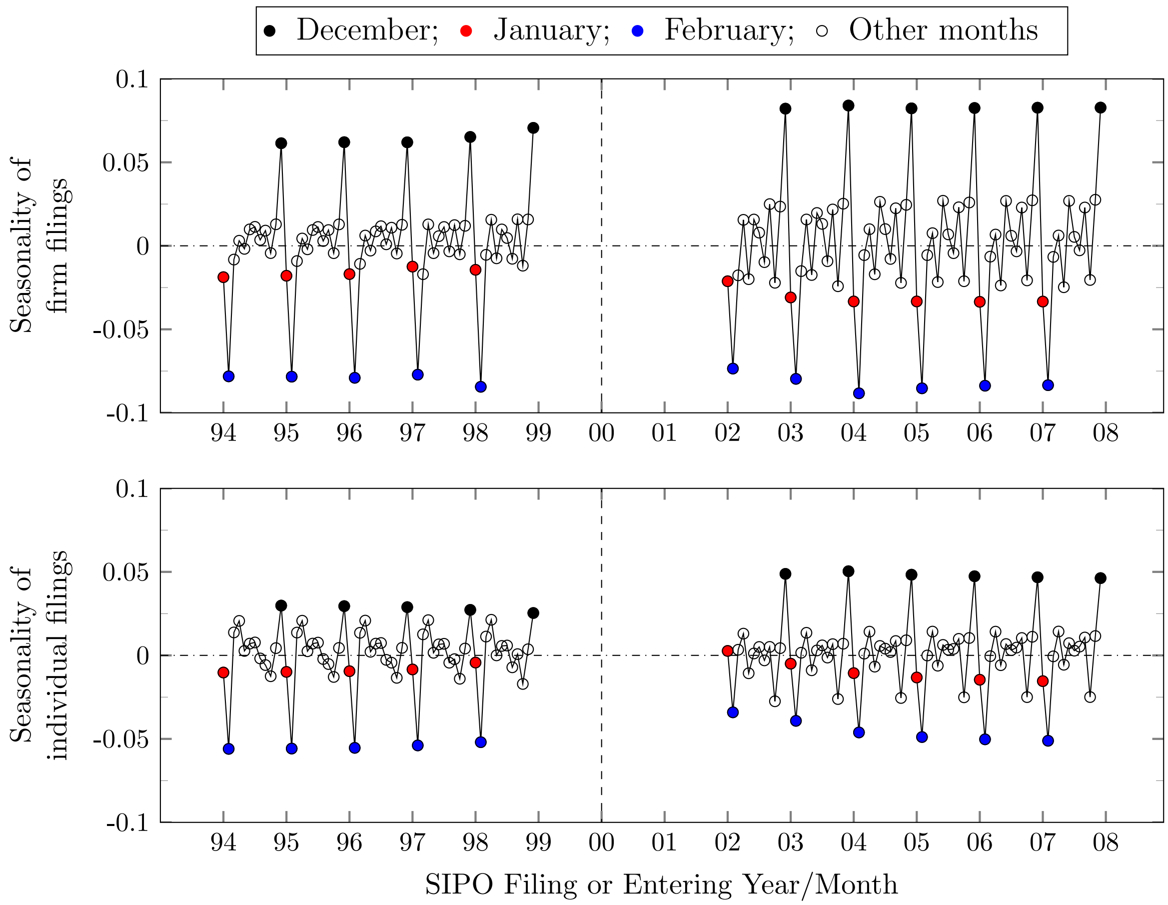

Figure 9. Comparison of seasonality between domestic firm and individual patent filings

evident in recent years, there is a clear trend of a decrease in filings in the early months in the following year, i.e., domestic filings in January, February and even March seem to be falling after 2000. A plausible explanation is that in order to meet the quota by the end of the year, firms shift some of their patent applications from the early months of the following year to previous year. This inter-temporal shift of patent applications may affect the quality of the applications and consequently the patenting strategy of the firms. In the next section, we similarly use the STL decomposition to investigate the change in quality of applications before and after 2000.

4.2 Illustration of the change in patent quality

What does the "December effect" have on patent quality? In this section we look at the standardized seasonal components of grant rate, number of claims and number of forward citations for domestic and foreign applications.

One obvious measure of the applications quality is the grant rate. Figure 10, again employing the STL procedure, shows that since around 2000 there have been "end-of-year" troughs in grant rates of invention patent applications at SIPO by both Chinese institutions, suggesting a worsening in the patentability of domestic applications that are filed at the end

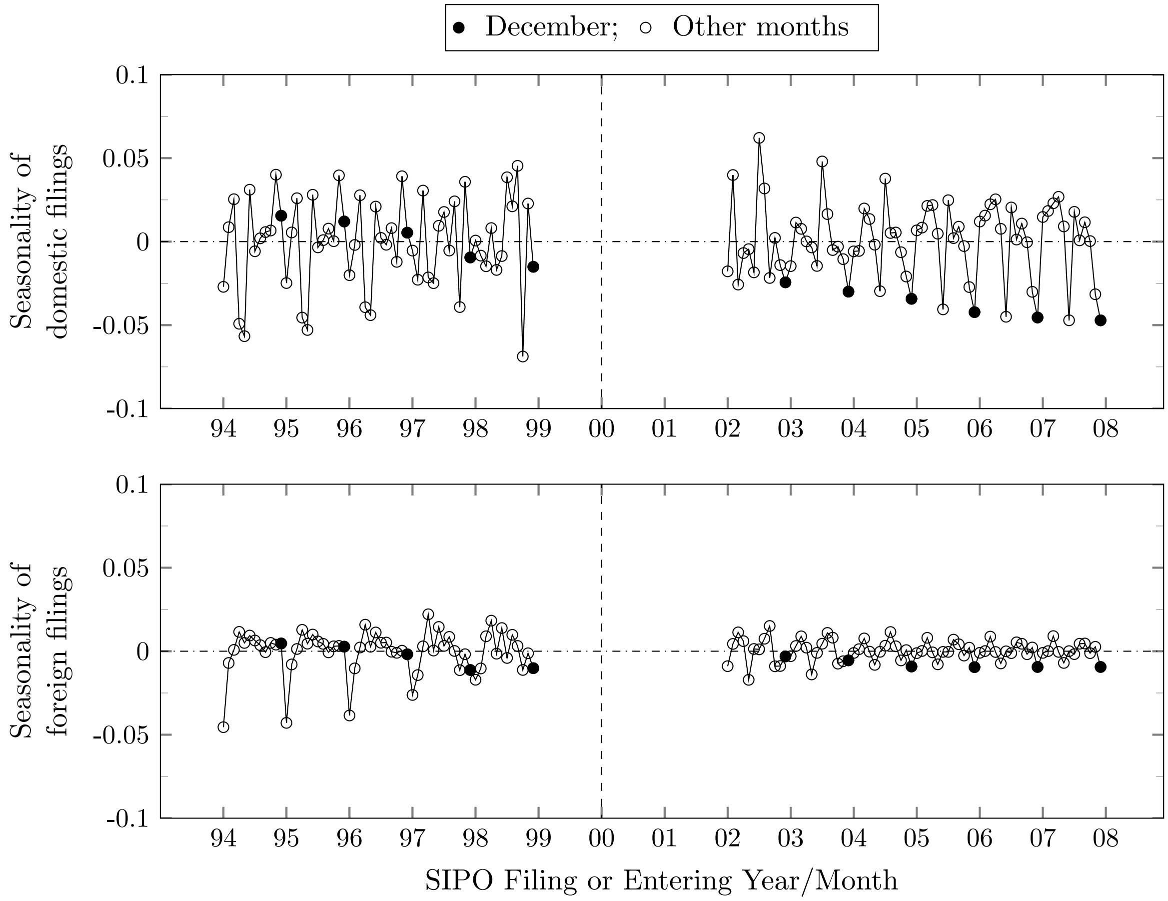

Figure 10. Comparison of seasonality in grant rate between domestic and foreign firm applications

of year. By contrast, no similar seasonality is apparent in grant rates for foreign applications at SIPO.

We also examine seasonality in the quality of granted Chinese patents, using the new forward citations data as our measure. In Figure 11, we plot the number of forward citations per year received by granted patents based on their patent filing months at SIPO. For patents granted to foreign institutions, no seasonality in forward citations is apparent. For patents granted to Chinese institutions, there have since 2000 been clear "end-of-year" troughs.

It seems that conditional on patent grant, patents granted to Chinese firms/institutions filed at years end are of lower quality than those filed in other months. Institutions, likely subjected to patenting targets in addition to financial incentives, seem to submit more, perhaps narrower, applications that are patentable, but of less interest to subsequent patentees.

5. Empirical Strategy and Econometric Analysis

Based on the visual analysis in Section 4, we conduct a more rigorous econometric study in this section to analyze the "December effect". In each of the tests, we investigate whether the quantity or quality of domestic applications in December are significantly different from

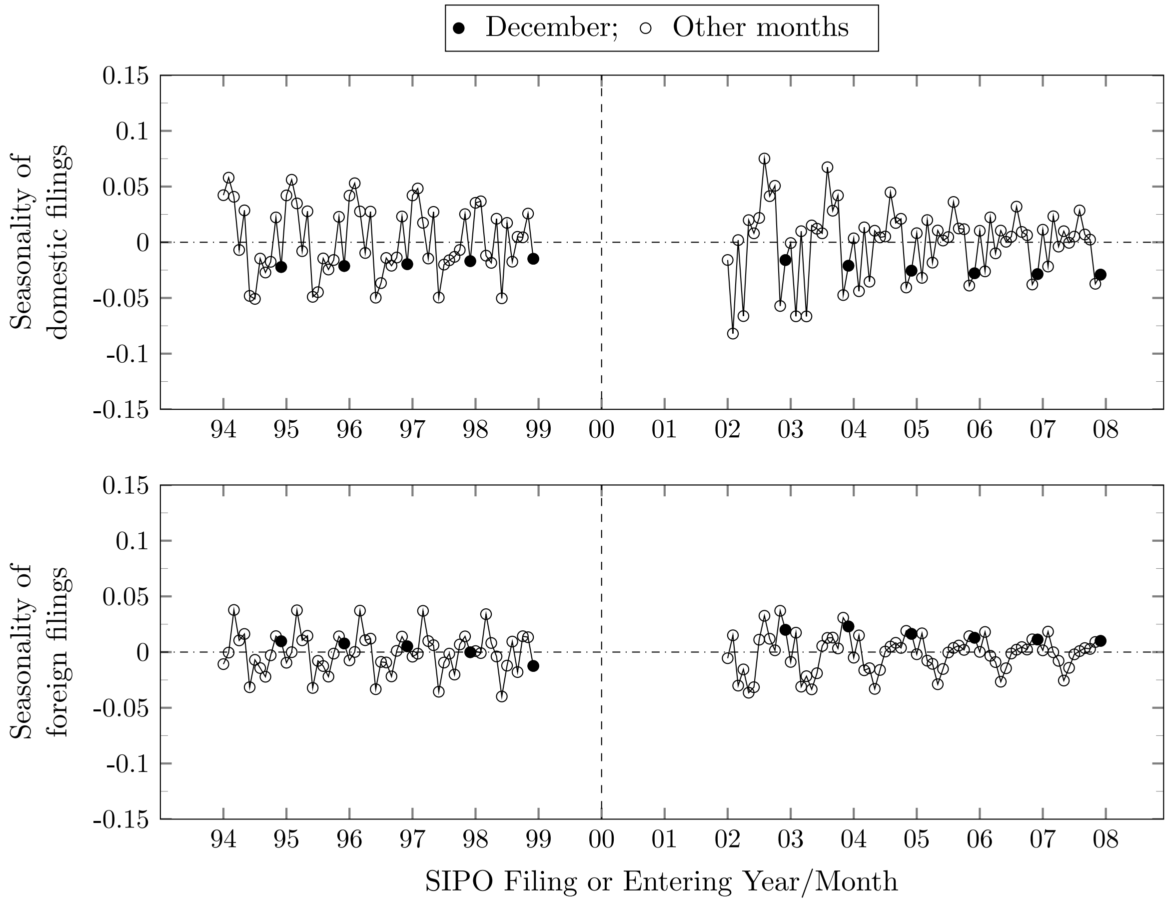

Figure 11. Comparison of seasonality in the number of yearly average forward citations between domestic and foreign firm patent applications

other months after year 2000. We conduct a difference-in-differences study to compare the filings in December to those in other months, before and after 2000. Furthermore, using foreign firm filings as a further control, we conduct a triple-differences analysis to confirm the results.

5.1 Turbulent Years

Firstly we provide evidence that there was some discontinuous change for domestic patent applications around 2000, in anticipation of or in response to the "10^{th} five-year plan", in which the Chinese government started its heightened efforts in planning on innovation and patenting. We plot the monthly number of applications and average grant rate in Figure 16. It seems that around year 2000 domestic firms discontinuously filed more applications. More evidently, the grant rate of applications by domestic firms in those months are much lower compared to other months. By contrast, foreign patent filings at SIPO were smooth during the same period. It seems during this time interval, there was considerable uncertainty or confusion over what to expect and how to respond to the new policies.[41] In light of these behaviors, we remove the year 1999, 2000 and 2001 from our empirical analysis to make sure the analysis is not affected by these "outlier" years.

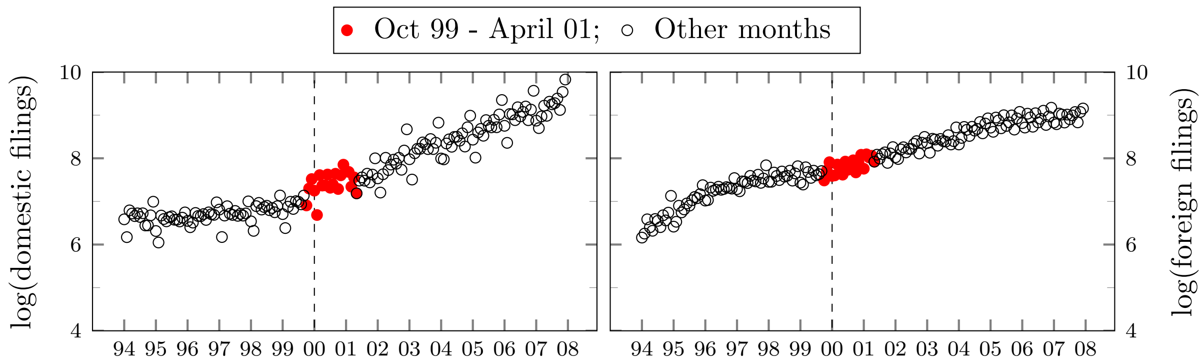

(a) monthly number of patent filings

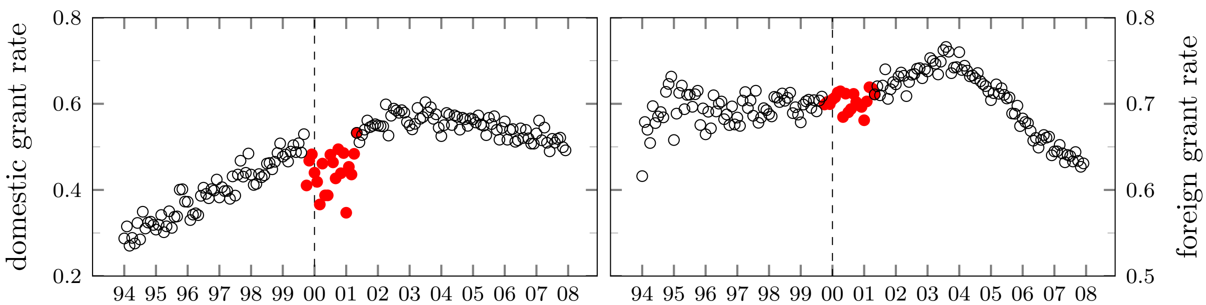

(b) monthly average grant rate

Figure 12. Turbulent patent filings by Chinese firms around 2000

5.2 Quantity

We look at the quantity of monthly applications and test whether the "December effect" becomes more evident after 2000. The comparison using the total domestic and foreign applications is not statistically informative. The results could be driven by firms from some special provinces that have different policies.[42] Therefore we generate a balanced panel of monthly applications by region (provinces in China and foreign countries, altogether 129 regions) and test whether domestic applications in December after 2000 are higher than in other months, using domestic applications before as controls (note we have already dropped observations from 1999, 2000 and 2001). Moreover, we can use foreign applications as an extra control to test the "December effect". A few regions that did not have invention

[41]In Appendix B, we also plot the graph for the monthly average number of claims and number of subclasses. Both of them show the discontinuous jump for domestic firm applications around 2000. Moreover, as can be seen from Appendix B, the "turbulence" occurs only with firm patents, not much with individual patents. Therefore, the phenomenon cannot be due to data input error or data truncation. However, we do not have a good explanation to the phenomenon.

[42]The major effective policies are usually at the provincial level.

applications in every month are dropped from the regression analysis automatically. The model we use is a difference-in-differences model as follows:

$$log(y_{rtd}+1) = \alpha + D_{after\ 2000,\ month,\ region} + \sum_{k=1}^{12} \beta_k \cdot D_{after\ 2000 \times month_k} + \varepsilon_{rtd} \tag{3}$$

where y_{rtd} is the domestic ($d = 1$)/foreign ($d = 0$) firm applications for region r at the t_{th} month.[43] We use month dummies and region dummies to control for any inherent seasonal differences and differences across regions. A linear time trend is also included. We use a mid-year month (July) as the base month, since applications in these months are relatively stable over the years (see Figure 8). The coefficients of interest are for the interaction terms between different months ($k = 1, \cdots, 6, 8, \cdots, 12$) and domestic firms. We cluster the standard errors at region level since applications within the same region are strongly correlated by nature.

The results are reported in "DID" column of Table 16. The coefficients for the months ("after 2000" interacted with months) indicate that compared to domestic applications in July, domestic applications in November and December are significantly higher after 2000. However, at the same time domestic applications in January and February become significantly lower. The slow down in applications in early year may be related to the surge of applications in the end of previous years, i.e., there could be intertemporal shift in applications in order to fulfil the quota.

In the second model we study whether the pattern (higher applications by the end of the year and less applications in the early months of the following year) still exists when we use foreign applications to control for macro-economic situation. Essentially, this is a triple-differences set up. The following model is estimated:

$$\begin{aligned} log(y_{rtd}) =& \alpha + D_{domestic,year \times month,region,after\ 2000} + D_{domestic \times month,domestic \times period,month \times period} + \\ & \sum_{k=1}^{12} \beta_k \cdot D_{domestic \times after\ 2000 \times month_k} + u_{rtd} \end{aligned} \tag{4}$$

where $after$ 2000 dummy denotes whether the observation falls before or after 2000. Other variables are defined in the same way as in model 3. In the month dummies, we are now able to assign a different dummy variable for each month (instead of a seasonal control for 12 months only). It allows for a better control of the unobserved characteristics over time. The results are shown in "Triple-difference" column of Table 16.

Compared to foreign applications, domestic applications have significantly increasing trends in late year filings (November and December), followed by, unfortunately, significantly decreasing trends in the first quarter (January through March). The results confirm our hypothesis. The late year surge has been growing because of the increasing importance of using patent applications as assessment for local government performance since 2000.

[43] We add 1 to the dependant variable to take into account cases of zero patent applications. The results are robust to dropping these months (use an unbalanced panel).

Table 16. Estimates of the number of monthly applications from different regions

Month	DID	Triple-difference
Jan	-0.276 (0.0652)***	-0.307 (0.0695)***
Feb	-0.228 (0.0648)***	-0.178 (0.0701)**
Mar	0.00208 (0.0516)	0.00499 (0.0557)
Apr	-0.0375 (0.0612)	-0.0939 (0.0658)
May	-0.124 (0.0679)*	-0.121 (0.0706)*
Jun	0.0365 (0.0680)	0.00915 (0.0719)
Aug	-0.0167 (0.0542)	-0.0447 (0.0598)
Sep	0.0816 (0.0671)	0.0595 (0.0712)
Oct	-0.0742 (0.0656)	-0.0270 (0.0692)
Nov	0.138 (0.0627)**	0.115 (0.0686)*
Dec	0.262 (0.0578)***	0.223 (0.0628)***
N	4092	18348

Standard errors clustered at region level are in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

However, this does not occur without consequences. The monthly applications in the first quarter of the year tend to be lower and this lower trend has been growing more significantly during the study period. The the surge in applications in December due to political pressure seems to have a "starving effect" to the applications in the first quarter of the following year.

5.3 How about the quality?

A following question we want to ask is whether the domestic patents filed in late year are of lower quality, since they are more likely to be filed to meet application quota, without caring about their quality. There are two dimensions of patent quality. The first one is technical quality (i.e., patentability): whether the patent meet the requirement of novelty and non-obviousness. The other one is economic quality (i.e., value): whether the patent is commercially important to the firm.

The most reliable patentability criterion is whether the patent get granted. If firms file more "junk" patents in December to meet the requirement, they are likely to be rejected by examiners. It's also possible that firms withdraw or let the application lapse without requesting substantive examination, which asks for an extra fee. In either case, we expect applications filed in late year to have on average lower grant rates. We look at the grant rates of patent applications in different months and compare them to the base month. The grant rate is defined as the granted patent applications over the total number of applications.

Similarly we estimate DID and triple difference models for grant rate. The first model looks at whether post-2000 domestic applications filed near the end of the year have a lower grant rate than those filed in other months, compared to pre-2000 applications. The second model uses foreign applications as an extra control to take into account the possibility of change in grant criteria at SIPO over time, as well as the slight truncation in application outcomes due to grant lag, which gives a lower grant rate for more recent applications (see Figure 6)[44]. The first model we use is as follows:

$$g_i = \alpha + \sum_{k=1}^{12} \beta_k \cdot D_{after\ 2000 \times month_k} + D_{month,year,region,IPC} + \varepsilon_i \tag{5}$$

where g_i, is a dummy variable indicating whether the patent application is granted. Since the patent grant rate is likely to depend on technology field, we control for 31 main International Patent Classification (IPC) classes. The other variables are defined in the same way as in Model 3.

The second model includes foreign applications as an extra control to estimate a triple-difference model:

$$\begin{aligned} g_i =& \alpha + D_{domestic,year \times month,region,IPC} + D_{domestic \times month,\ domestic \times after\ 2000,\ month \times after\ 2000} + \\ & \sum_{k=1}^{12} \beta_k \cdot D_{domestic \times after\ 2000 \times month_k} + u_i \end{aligned} \tag{6}$$

where we include a dummy indicating whether the patent was filed by a foreign firm. We similarly control for region and IPC dummies. The other control variables are defined similarly as in Model 4 and 5.

The results from Model 5 are reported in the first column of Table 17. The grant rates of domestic applications near the end of the year are significantly lower than those in July. The grant rates of domestic applications filed in November and December after 2000 are roughly 7 percentage point lower than those filed in other months. In model "Triple Difference" of panel "Grant rate", we report the estimation results using foreign applications as an extra control. The effect on grant rate of domestic applications in November and December is robust and consistently at around negative 7 percentage point. Considering that the average grant rate for domestic applications after 2000 is about 61%, the decrease in grant rate in Nov and December is sizeable (about 12% lower) and economically significant. It seems that the stronger political pressure after 2000 indeed induced more end-of-year filings of lower quality. *We do not find a change in quality for patent applications in other months. We do not find strong evidence of a change in quality of filings at the beginning of the year. Therefore the inter-temporal shift of applications did not seem to affect the quality of patents filed in the early months of the following year. (need more discussion)*

An important indicator for economic value of patent is whether the patent is cited by future patents (i.e., forward citation). We can use forward citation as a further argument, which is very meaningful as all published applications, in principle, can be cited and thus can have forward citations. We run the same models as in Model 5 and 6, using average

[44]For more recent applications, there is slightly more truncation for foreign applications than for domestic applications. Because patent grant on average takes longer than patent rejection, more truncation leads to larger downward bias in grant rate for foreign applications. The difference in truncation, therefore, biases against us when we evaluate the grant rate for domestic applications filed near the end of the year.

Table 17. Estimates of grant rates and forward citations

	Grant rate		Forward citations	
month	DID	Triple Difference	DID	Triple Difference
Jan	0.0427 (0.0172)**	0.00892 (0.0233)	0.00470 (0.00410)	-0.00109 (0.00424)
Feb	0.0212 (0.0223)	0.00949 (0.0237)	0.00358 (0.00682)	-0.000917 (0.00654)
Mar	0.00945 (0.0187)	-0.000585 (0.0223)	0.00811 (0.00701)	0.00531 (0.00671)
Apr	0.0414 (0.0226)*	0.0354 (0.0256)	0.00784 (0.00533)	0.00571 (0.00513)
May	0.0357 (0.0179)*	0.0354 (0.0227)	0.00387 (0.00518)	0.00528 (0.00616)
Jun	-0.0278 (0.0240)	-0.0261 (0.0243)	0.00254 (0.00420)	0.00251 (0.00437)
Aug	-0.0162 (0.0186)	-0.0140 (0.0214)	-0.00196 (0.00528)	0.00182 (0.00567)
Sep	-0.0144 (0.0205)	-0.0135 (0.0221)	-0.00319 (0.00420)	0.000141 (0.00458)
Oct	-0.0287 (0.0192)	-0.0301 (0.0208)	-0.00839 (0.00494)	-0.00355 (0.00576)
Nov	-0.0693 (0.0217)***	-0.0680 (0.0250)**	-0.0189 (0.00324)***	-0.0154 (0.00344)***
Dec	-0.0714 (0.0226)***	-0.0647 (0.0239)**	-0.0181 (0.00512)***	-0.0127 (0.00521)**
N	313189	835683	313189	835683

Standard errors clustered at IPC level are in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

number of forward citations per year as dependant variables. We choose the average number of forward citations per year since the publication of the patent as the dependant variable to further reduce possible bias caused by truncation (forward truncation is always right-censored as earlier published naturally accrues more forward citations by the end of our study period).[45] The results are reported in the last two columns of Table 17. It's clear that the economic value of patents filed in the end of the year after 2000 are significantly lower compared to other months. The results are robust to using foreign patents as controls. On average, domestic patents accrue 0.1 forward citation per year. Therefore, the number of forward citations to patents filed in the late year is about 15% lower compared to those filed in other months.

5.4 Intertemporal shift?

In this section we analyze the possibility of firms shifting their planned patents in early next year to the year before, in order to meet the political pressure. If the pure effect of the political intervention is to cause an intertemporal shift, without decreasing the quality of the patents, then perhaps we do not need worry too much.

One straightforward way to considering the intertemporal shift is to pool the patent applications made in the late year and early next year together. Therefore we pool patent

[45] The results are robust to using full number of forward citations as dependant variable. By using average number of forward citations, we rely on another assumption that citations grow linearly. This assumption usually does not hold. However, we checked for Chinese patents within 9 years of publication, the average number of forward citations grows roughly 0.1 per year starting from year 3. The results are also robust to using only granted patents or only considering Chinese patent citations. Please see Appendix C

applications made in November and December together with those in January and February of the following year, and estimate the quantity and quality effect. The results are reported in Table 19. It seems there is evidence of intertemporal shift in patenting, as the total number of patents filed do not increase after we pool the neighboring months across the new year boundary.[46] However, it's clear the quality of the patents still dropped, which suggest interesting consequence of the intertemporal shift.

Table 18. Estimates of the number of monthly applications from different regions (pooled months for Jan-Feb of following year to Nov-Dec of previous year)

Month	DID	Triple-difference
Mar	0.0316 (0.0543)	0.0231 (0.0596)
Apr	-0.0385 (0.0657)	-0.105 (0.0705)
May	-0.115 (0.0760)	-0.124 (0.0789)
Jun	0.0362 (0.0713)	-0.00390 (0.0761)
Aug	-0.0194 (0.0583)	-0.0659 (0.0643)
Sep	0.103 (0.0702)	0.0686 (0.0748)
Oct	-0.0777 (0.0721)	-0.0524 (0.0762)
Pooled	0.0981 (0.0484)*	0.0748 (0.0544)
N	3720	16560

Standard errors clustered at region level are in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.5 Signs of patent split?

Lei et al. (2013) find with Chinese patent application data that when the patent subsidy increases, there is evidence that applicants break up their inventions to generate more applications thus to claim more subsidy and reward. In our case, it's also possible that firms break up their applications to meet the yearly quota, especially in the late months of the year. Splitting patents does not make sense under normal condition since it leads to higher application cost and longer examination time for the firms. However, under political or administrative pressure, splitting innovation output to come up with more applications may be an easy way to meet the quota.

Assume that a group of R&D personnels have worked on an R&D project before they decide to file, then naturally, these employees, or at least some of these employees, will become the co-inventors for the filed applications. If two (or more) applications are filed from the project, they will share some or all of their co-inventors. Therefore, if the firm decides to split their innovation to meet the quota, more patent applications from this firm

[46]Note in pooling the months, we lost one year of observations (Jan and Feb of 1994 and Mar and later of 2007 are dropped). The results are robust to including 2007 as an extra year.

Table 19. Estimates of grant rates and forward citations (pooled months for Jan-Feb of following year to Nov-Dec of previous year)

	Grant rate		Forward citations	
month	DID	Triple Difference	DID	Triple Difference
Mar	0.00678 (0.0187)	-0.00317 (0.0217)	0.00376 (0.00700)	0.00139 (0.00650)
Apr	0.0308 (0.0220)	0.0246 (0.0253)	0.00364 (0.00538)	0.00223 (0.00545)
May	0.0296 (0.0189)	0.0283 (0.0232)	0.00338 (0.00581)	0.00575 (0.00667)
Jun	-0.0232 (0.0241)	-0.0218 (0.0240)	0.00170 (0.00456)	0.00151 (0.00470)
Aug	-0.0202 (0.0188)	-0.0198 (0.0206)	-0.00244 (0.00579)	0.00211 (0.00622)
Sep	-0.0231 (0.0214)	-0.0207 (0.0227)	-0.00661 (0.00418)	-0.00276 (0.00459)
Oct	-0.0408 (0.0199)**	-0.0402 (0.0213)*	-0.0113 (0.00498)**	-0.00671 (0.00587)
Pooled	-0.0664 (0.0185)***	-0.0612 (0.0185)***	-0.0168 (0.00332)***	-0.0115 (0.00334)***
N	229700	667720	229700	667720

Standard errors clustered at IPC level are in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

will have shared co-inventors, *ceteris paribus*. Since we have the names of the co-inventors for each application, we can construct the proportion of shared co-inventors for each application as a metric to indicate the likelihood of this patent application being a "split patent".

Since we focus on monthly application pattern, we define a co-inventor as "shared" if he/she is also a co-inventor for another patent application filed in the same month within the same firm.[47] Each patent will be given a share metric s_i which equals to the proportion of its co-inventors that are "shared". For example, if one patent application has three co-inventors, and two of them also filed other applications in the same month, then the share metric for this patent application is $s_i = 2/3$. Using this metric, we study the seasonal component of the shared co-inventor across different months, again by the STL decomposition. The results are plotted in Figure 13. The figure seems to suggest some significant and rising co-inventor shares at the end of year after 2000.

We run a regression model using the proportion of shared co-inventors (s_i) for each patent application as the dependant variable, and test whether domestic applications filed in December after 2000 have higher proportion of shared co-inventors compared to base month and foreign applications. The regression we run are basically the same as Equation 5 and 6, except that we use the constructed proportion of shared co-inventors s_i as the dependent variable:

$$s_i = \alpha + \sum_{k=1}^{12} \beta_k \cdot D_{after\ 2000 \times month_k} + D_{month,year,region,IPC} + \varepsilon_i \tag{7}$$

[47]The situation that one inventor works for multiple firms during one month is rather rare.

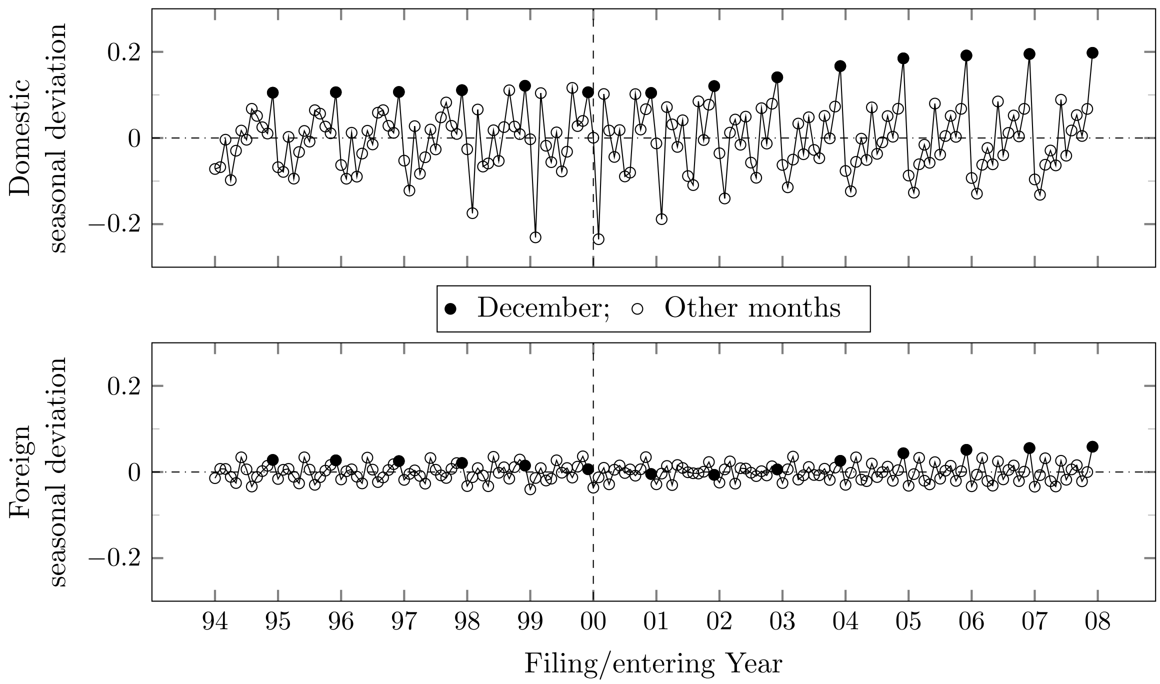

Figure 13. Comparison of seasonality in the proportion of shared inventors between domestic and foreign firm patent applications

$$s_i = \alpha + D_{domestic,year \times month,region,IPC} + D_{domestic \times month,\ domestic \times after\ 2000,\ month \times after\ 2000} + \sum_{k=1}^{12} \beta_k \cdot D_{domestic \times after\ 2000 \times month_k} + u_i \tag{8}$$

The estimates from Equation 7 and 8 are reported as "DID" and "Triple Difference" in Table 20. The proportion of shared co-inventors for domestic applications in December after 2000 is significantly larger compared to applications in other months. The results are robust to using foreign applications as controls. This finding has important implications to the current patent policies in China. In order to meet the growing political/administrative pressure, firms choose to break up patents to come up with more applications. This filing behavior is economically inefficient and leads to a waste of administrative resource for both the patentees and the SIPO.

6. Discussion and Conclusion

By looking at the seasonality characteristics of domestic invention patent applications and comparing them with foreign applications, we find an intriguing surge in patenting behavior in Decembers for domestic application. One plausible explanation of this phenomenon is that these applications are made under administrative pressure to meet yearly quotas set by the local governments. These quotas are important assessment criteria for the performance of some officials. We then compare the grant rate of patent applications by month and find the grant rate for domestic applications in December to be lower than the base month after

Table 20. Estimates of the proportion of shared co-inventors

month	DID	Triple Difference
Jan	0.00492 (0.0190)	-0.0000274 (0.0218)
Feb	-0.0173 (0.0200)	-0.0112 (0.0305)
Mar	0.0125 (0.0239)	0.00223 (0.0284)
Apr	0.0338 (0.0152)**	0.0277 (0.0190)
May	0.00774 (0.0200)	-0.000491 (0.0224)
Jun	0.0477 (0.0207)**	0.0451 (0.0252)*
Aug	0.00521 (0.0176)	-0.0146 (0.0216)
Sep	0.0203 (0.0235)	0.00740 (0.0260)
Oct	0.0128 (0.0177)	0.0157 (0.0289)
Nov	0.0449 (0.0159)***	0.0479 (0.0179)**
Dec	0.0572 (0.0208)***	0.0434 (0.0248)*
N	313189	835679

Standard errors clustered at IPC level are in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

2000. A further look at the shared co-inventors for each application reveals that domestic applications in December have significantly more shared co-inventors compared to other months, using foreign applications as controls. This suggests that firms under political pressure may file lower-quality patents or split their innovation output to come up with more applications, such that these applications share a higher number of co-inventors compared to the normal filings. If this is the case, it put doubt to the effectiveness of China's political promotion of innovation using the quota instrument.

The ineffectiveness of the quota instrument is not really surprising. It touches a broader literature on using price or quantity as instrument to achieve an optimal level of output, when the market clearance outcome is not the first best due to externality (Weitzman, 1974). When there is incomplete information with random disturbance in the associated cost and benefit, the choice of price or quantity instrument depends on the relative curvature of the cost and benefit functions. The problem can be formulated in the patent system in a similar way. Assume in a certain market of an end product, there is a plethora of small patents (incremental improvements, as are the common cases) that can be made. An extra patent helps a firm to improve slightly its end product and thus to capture an extra slice of the market. Innovation has positive externality to the whole society and the we need policy instrument to achieve the optimal level of innovation from a social point of view. That is essentially how the patent system comes into being, i.e., the monopoly power endowed with the patent is used to at least partly correct the market failure associated with innovation.

The patent system is more like a price instrument. It relies on the market to determine the value of the innovation and consequently the amount of resource firms put into the R&D process. The patent system endows the innovator a monopoly power to facilitate and encourage disclosure of innovations into the public domain for the common good. Assume that the monopoly power is insufficient to correct the market failure associated with innovation (otherwise no regulation is needed), the output of the patents may be too low from a social point of view. For this reason, sometimes we try to stimulate innovation using other policy instruments. Many of these policies also work as price instruments, such as patent subsidy/reward or R&D subsidy/tax credit.[48] On the other hand, the quota system is very much like a quantity instrument to induce the "optimal" amount of innovation.

In a market with full information, price and quantity instrument should work equally well. However, since the local government does not have full information of the R&D capacity of the firms within its jurisdiction, the effect of the two types of instruments will depend, among other things, the relative curvature of the cost and benefit function of patent production. On a simple note, since patents are intermediate output, which has substitution effect with other goods of similar kinds, the demand is likely to be flatter. In this case, price instrument is generally preferred to quantity instrument. Moreover, when we consider many firms engage in the same type of innovation activities and if their marginal cost functions are correlated, the advantage of the price instrument is even larger. The results from this paper indicates that the target set by the authorities are perhaps overshoot, which could lead to significant dead weight loss to the society.

[48]There could be significant differences between the two types of subsidies, unless we assume a proportional relationship between the R&D input and patent output, which usually does not hold given the risks involved in R&D activities. Moreover, in policy intervention, each instrument has its own pitfalls: subsidies on output (patents) necessarily faces the moral hazard problem while subsidies on innovation input (R&D activities) have to deal with the adverse selection problem. We will not go into details here.

A. Illustrative anecdotes of patenting targets and mandates

Although the targets on patenting set by Chinas national plans are claimed to function only as guidance and not to be binding, governments at lower levels (provincial, municipal and county) have specified quantitative targets on patenting within their jurisdiction, which are then included as important factors in assessing performance of government officials and government-controlled enterprises. Our online search found numerous news reports on annual patenting targets, set by a province or city level governments for lower level governments or enterprises that it controls, by China's State-owned Assets Supervision and Administration Commission (SASAC) to state-owned enterprises (SOEs), or by SOEs to their subsidiary firms. Some targets specify the number of patent applications and/or patent grants, and other growth rate in patenting.

For example, the province of Liaoning has since 2002, at latest, included "meeting the patent application target" as one of its annual "Assessment criteria for city governments". Beijing included "the number of patent applications" as an important assessment criterion for managers of enterprises since at latest 2003. The provincial government of Anhui required in 2012 that cities in the province had a yearly 20%-40% increase in the number of patent applications, based on their previous performance. In September 2010, the city of Dali in Yunnan province reported that it had accomplished the 2010 quota for patent applications during the first half of the year. The quotas were set by the provincial government, for all three types of patent applications and for invention patent applications specially. On November 25th 2011, Guangxi Province issued its decision on "Assessment of patent application growth 2011" to all cities in the province, stating that an assessment would be conducted in January 2012 to evaluate the performance in patent filing by all cities during 2011 to see if the number of patent applications increased by 25% from the previous year. Even though such an evaluation might had been expected, this announcement definitely give impetus to cities in Guangxi to come up with more applications before the year ended.

Moreover, Chinas State-owned Assets Supervision and Administration Commission (SASAC) has used innovation and patenting as one important criterion for assessing the performance of state-owned enterprises (SOE). For example, in 2006, Commissioner Rongrong Li expressed concerns that half of the state-owned enterprises did not have patents on their major products and said that R&D expenditure and innovation would be used for performance assessment of the SOE management. In April 2009, the SASAC promulgated "the Guide on Strengthening Intellectual Property Right by State-owned Enterprises" and required that each SOE work out and implement its own "strategic plan on intellectual property" by the end of the year.

There are also news reports that SOEs in turn set patenting targets for its subsidiary companies. For example, *Shanxi Taiyuan Electronic Power Company* reported on June 28th, 2012, that it had accomplished its 2012 patent application quota set by the *State GridShanxi Company*.

Because most of news reports on government official websites and government online news archives were not very well maintained in China, it is more difficult to find news reported during the period of our study. Nonetheless, it seems that setting patenting targets and using them to assess performance of government officials and state-owned enterprise management has been a common practice since 2000.

B. Turbulent years

Turbulent filings from domestic firms around 2000 based on the monthly average number of claims and subclasses.

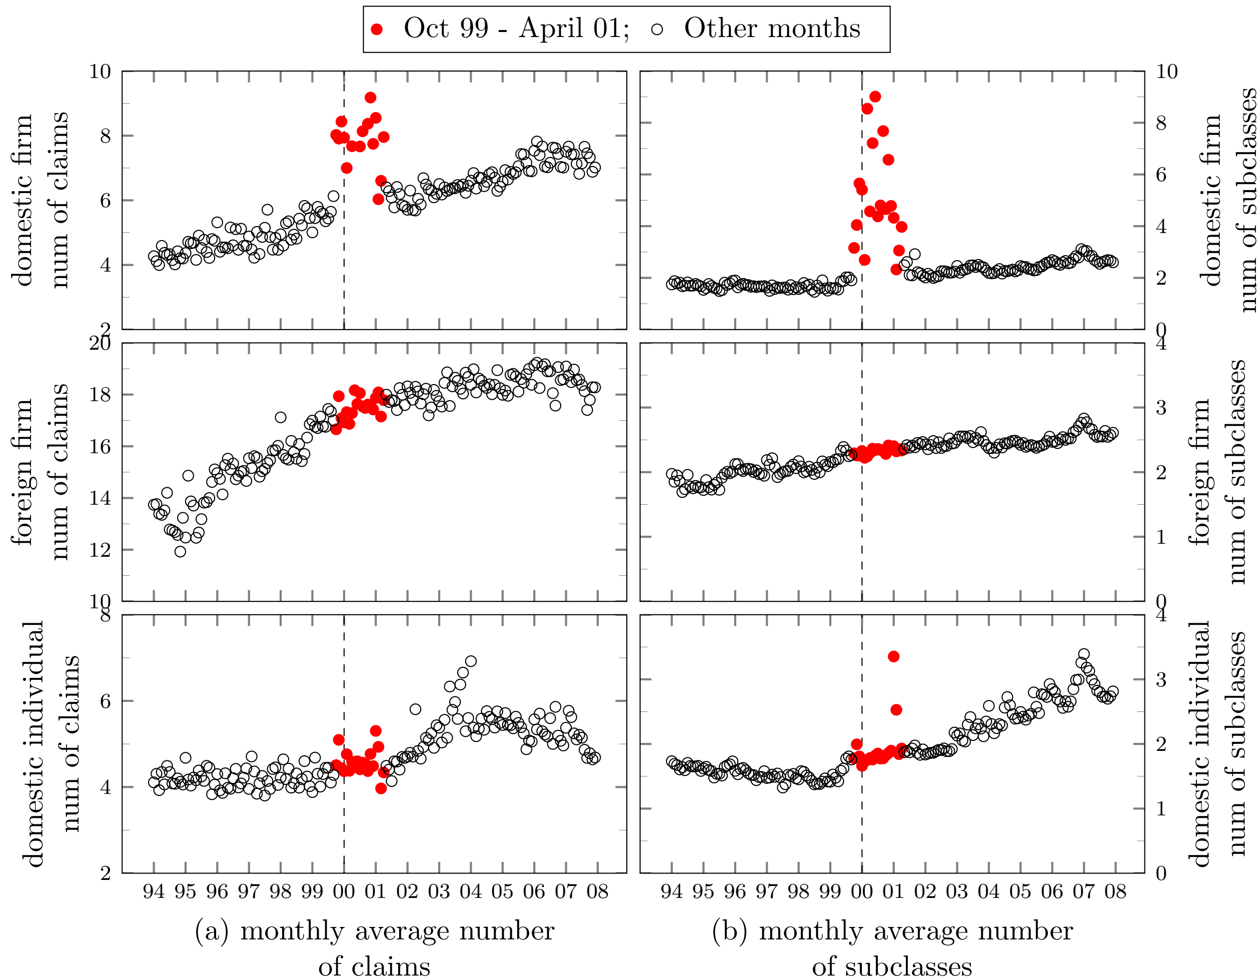

Figure 14. Turbulent patent filings by Chinese firms around 2000 in other characteristics

C. Robustness Check of Forward Citations Results

Table 21 and 22 report the results on the change in patent quality measured by forward citations. Table 21 show that the "end-of-year" drop in quality is robust to studying only the granted patents. In Table 22, we limit the citations to those that are cited by SIPO patents. The results are again consistent.

Table 21. Estimates of grant rates and forward citations

	Grant rate (applications)		Forward citations (granted patents)	
month	DID	Triple Difference	DID	Triple Difference
Jan	0.0427 (0.0172)**	0.00892 (0.0233)	-0.00296(0.00542)	-0.00941(0.00589)
Feb	0.0212 (0.0223)	0.00949 (0.0237)	-0.00744(0.0103)	-0.0146(0.00997)
Mar	0.00945 (0.0187)	-0.000585 (0.0223)	0.00174(0.00944)	-0.0000343(0.00953)
Apr	0.0414 (0.0226)*	0.0354 (0.0256)	0.00434(0.00823)	0.00107(0.00725)
May	0.0357 (0.0179)*	0.0354 (0.0227)	0.00117(0.00481)	0.00301(0.00597)
Jun	-0.0278 (0.0240)	-0.0261 (0.0243)	0.00460(0.00572)	0.00364(0.00534)
Aug	-0.0162 (0.0186)	-0.0140 (0.0214)	0.000111(0.00805)	0.00345(0.00853)
Sep	-0.0144 (0.0205)	-0.0135 (0.0221)	-0.00451(0.00618)	-0.0000755(0.00620)
Oct	-0.0287 (0.0192)	-0.0301 (0.0208)	-0.0103(0.00596)*	-0.00605(0.00654)
Nov	-0.0693 (0.0217)***	-0.0680 (0.0250)**	-0.0220(0.00477)***	-0.0173(0.00451)***
Dec	-0.0714 (0.0226)***	-0.0647 (0.0239)**	-0.0200(0.00761)**	-0.0133(0.00587)**
N	313189	835683	191186	553864

Standard errors clustered at IPC level are in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 22. Estimates of grant rates and forward citations (keep only SIPO citations)

	Grant rate (applications)		Forward citations (granted patents)	
month	DID	Triple Difference	DID	Triple Difference
Jan	0.0427 (0.0172)**	0.00892 (0.0233)	-0.00417 (0.00768)	-0.0125 (0.00794)
Feb	0.0212 (0.0223)	0.00949 (0.0237)	-0.0113 (0.0105)	-0.0191 (0.0104)*
Mar	0.00945 (0.0187)	-0.000585 (0.0223)	-0.000927 (0.00886)	-0.00397 (0.00868)
Apr	0.0414 (0.0226)*	0.0354 (0.0256)	0.00407 (0.00796)	-0.000709 (0.00721)
May	0.0357 (0.0179)*	0.0354 (0.0227)	-0.000426 (0.00553)	-0.000217 (0.00642)
Jun	-0.0278 (0.0240)	-0.0261 (0.0243)	-0.000271 (0.00668)	-0.00147 (0.00664)
Aug	-0.0162 (0.0186)	-0.0140 (0.0214)	-0.00457 (0.00665)	-0.00272 (0.00743)
Sep	-0.0144 (0.0205)	-0.0135 (0.0221)	-0.00603 (0.00628)	-0.00257 (0.00573)
Oct	-0.0287 (0.0192)	-0.0301 (0.0208)	-0.0134 (0.00814)	-0.00768 (0.00805)
Nov	-0.0693 (0.0217)***	-0.0680 (0.0250)**	-0.0222 (0.00467)***	-0.0153 (0.00460)***
Dec	-0.0714 (0.0226)***	-0.0647 (0.0239)**	-0.0174 (0.00567)***	-0.00997 (0.00485)**
N	313189	835683	191186	553864

Standard errors clustered at IPC level are in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Abstract

What Happens When Politics Intervenes in the Patent System?

by

Professor Zhen Lei

Department of Energy and Mineral Engineering

Penn State University

and

Zhen Sun

Doctor of Philosophy in Agricultrual and Resource Economics

University of California, Berkeley

and

Professor Brian D. Wright

Agricultrual and Resource Economics

University of California, Berkeley

This study examines the impacts of patent subsidy policies on patent filings in China. China had rapid growth in patenting in recent years and became the number one in patent filings in 2011. We study five neighboring cities in Jiangsu province, where in June 2006 one city, Zhangjiagang, not only significantly increased the amount of subsidy per patent application but also included a large reward for granted applications, while subsidies in the other cities remained unchanged. Using a difference-in-differences strategy we find that the number of invention patent filings from Zhangjiagang significantly increased, but the grant rate of an average patent application from Zhangjiagang did not drop after the policy change, compared to those from the other cities. Moreover, the total number of claims for each applicant in Zhangjiagang remained the same whereas the average number of claims per application from Zhangjiagang declined after June 2006. Thus, we find that the policy is ineffective. The increase in patent applications appears to be achieved by reducing the number of claims included in each application without increasing the total number of claims. Our findings indicate that applicants in Zhangjiagang were not significantly constrained by the cost of patenting before the policy change. They also suggest that applicants have significant discretion in the number of patents they can receive to protect a given number of claims. In the case studied here, the former was responsive to local financial incentives, while the latter was not.

"The generosity of China's incentives for patent-filing may make it worthwhile... to patent even worthless ideas... Patents are easy to file,... but gems are hard to find in a mountain of junk."

- *"Patents, yes; ideas, maybe?"*, The Economist, Oct 14th, 2010

1. Introduction

In 2011 China became the world leader in the number of published invention patent applications, outpacing the United States, Europe and Japan. Patent applications in China increased from 63450 in 2001 to 391177 in 2010, an annual rate of 22.6%;[49] and domestic applications have grown even faster, from 30038 to 293066, by 28.8% per year during the period.[50] The driving forces behind this China patent boom have been debated. Some observers consider the boom to be an indicator of significant and genuine strides in China's innovative capacity, resulting from China's focused efforts, iconized by its "Medium to Long Term Plan for the Development of Science and Technology (2006)" (hereafter MLP), to promote its indigenous innovation and technological development and transform its economy from "made-in-China" to "invented/designed in China". Other observers, however, believe that the patent boom is largely due to various patent subsidy policies implemented by local governments to meet and/or exceed the patenting targets specified by the government.[51] They argue that patent subsidies incentivize applicants to file opportunistic applications for inventions of low patentability or low value that would have not been filed without those subsides. Thus they claim that most filings in this China patent boom are so-called "junk inventions".

This paper, to the best of our knowledge, conducts the first applicant-level empirical study on the effects of patent subsidies on both the quantity and quality of patent filings in China. We compare five neighboring cities (Zhangjiagang, Taicang, Suzhou, Kunshan and Changshu), all within the Suzhou Municipality. In June 2006, the citiy of Zhangjiagang increased its patent subsidy for an invention patent application from YMB 1500 to YMB 3000 and added a reward of YMB 10000 if the application is granted. Around that time, patent subsidies in the other four cities remained unchanged.[52]

At first glance, one might reasonably conjecture the sizable increase in subsidy for patent filings and the hefty reward for patent grants would provide good incentives for applicants in Zhangjiagang to file applications for inventions that would have not otherwise received attention from the patent system. In particular, applicants could be incentivized to file inventions that are patentable but of little or low expected value, because of the patent application subsidy and grant reward. Applicants can always abandon these patents if they turn out not to be commercially useful.[53] If the conjecture is true, the increase in patenting

[49]Growth was much slower elsewhere: 4.6% in America, 5.6% in South Korea, 3.6% in Europe and -2.7% in Japan during the same period.

[50]*Patent applications by patent office (1983-2010)*, source: WIPO Statistics Database, available at `http://www.wipo.int/ipstats/en/statistics/patents/`

[51]The MLP and the ensuing National Intellectual Property Strategy (2008) specify overall national patenting targets, which are then allocated to local governments.

[52]Note that the level and structure of patent subsidies have been different across the five cities.

[53]Patents are renewed annually in China.

elicited by the local incentives is likely to be accompanied by sharply declining quality.

We implement a difference-in-differences strategy and study patent filings before and after June 2006 by a panel of more than 3000 applicants in Zhangjiagang and the control cities, in a time window of July 2004 through December 2007. We find a significant increase in the number of invention patent filings from Zhangjiagang after June 2006. The patentability of patent applications from Zhangjiagang did not decrease after the policy change, as indicated by their grant rates compared to those from the other cities. However, the total number of claims for each applicant in Zhangjiagang remained the same, and the average number of claims per application from Zhangjiagang actually declined after June 2006.

These results do not support the argument that the increase in patent subsidies incentivize applicants in Zhangjiagang to produce and file applications for inventions of lower quality (less patentable or less valuable) that they would not have been filed without the policy change. If this were the case, relative to those from the control cities, the total number of claims per applicant in Zhangjiagang should have increased because of more inventions being filed for patents, and/or the grant rate of patent applications from Zhangjiagang should have worsened due to less patentable inventions being filed. Given the total patenting expense of YMB 8000 from filing to grant, if an applicant files for an invention that is of low patentability and unlikely to get granted, it will, even with a patent filing subsidy of YMB 3000, incur a net loss in the end. Applicants in Zhangjiagang also did not seem to have additional inventions waiting "in the attic" which they turned into application in response to the hefty reward (otherwise the number of claims per applicant should have increased). Instead, they only seem to split their applications into multiple filings to collect more of the hefty reward for patent grants.

The finding that no additional inventions is evident in Zhangjiagang after the subsidy change suggests that, before the policy change in 2006, applicants in Zhangjiagang had little financial constraint. It is not a surprising result given that this is an economically developed region in China and that the cost of patenting in China is relatively small. Few of their inventions remain unpatented due to the low cost of patenting. What they could do was to break up existing inventive claims into multiple patent filings that were as patentable as previous filings and to collect greater reward, which they did.

Our empirical findings shed light on the effects of patent subsidy policies in China. A majority of patent filings in China are from economically developed regions such as the coastal province in which Zhangjiagang is located. Assuming that our finding that applicants have little financial constraint in patenting holds true for other developed regions, it does not seem necessary to have those local patent subsidy policies, which, depending on their design, might merely boost the number of patent filings without actually increasing the stock of patented inventions, or have little effect.[54]

Our results suggest that, with little financial constraint, applicants in Zhangjiagang likely file applications for all patentable inventions. These findings provides a novel perspective to a broad literature, both theoretical and empirical, on the optimal design of the patent system (see, e.g., Gallini, 2002; Farrell and Shapiro, 2010), and more particularly, on the use

[54]Our findings only show that the increase in subsidies did not have an effect on invention measured by patented claims. Given that the increase doubled the previous 1500 RMB subsidy and added a grant reward, it is unlikely that the previous subsidy was more effective.

of patent fees as a policy tool that often assumes an expected profit maximizing patenting strategy for applicants (Scotchmer, 1999; Cornelli and Schankerman, 1999; Gans et al., 2004; Hunt, 2006; Marco and Prieger, 2009; Caillaud and Duchêne, 2011; De Rassenfosse, 2012. Also see De Rassenfosse and Van Pottelsberghe de la Potterie, 2012 for a review).[55]

Finally, our findings might have interesting policy implications for on-going implementation or discussion of patent fee policies, in the U.S., Europe[56] and other patent offices in the world. The United States Patent Office (USPTO), for example, proposes a new fee structure to "subsidize filing, search, and exam fees to enable lower cost of entry into patent system", and includes a 50% reduction for small entities and a 75% reduction for micro entities.[57] Underlying these policies is the view that applicants, small size firms in particular, are financially constrained in filing for patents, which may or may not hold in all countries that differ in firm innovative activity and patenting fees.

The remainder of the paper is organized as follows. In section 2 we provide a brief review of China's innovation strategies and the subsidy policies, with more information for the six cities studied in the paper. Section 3 explains our data and methodology. Section 4 shows the estimated effects of the policy change on the quantity and quality of patent applications. Section 5 concludes.

2. Background

2.1 China's Recent Indigenous Innovation Strategy

Since it opened to the world economy three decades ago, China's economic development has been one of "made in China", relying on its low-cost manufacturing of existing products. However, promoting innovation and technological development has long been an important theme for the Chinese government, exemplified by its 863 Program and 973 Program.[58] Since 2000, China has made more systematic and ambitious efforts to strengthen its innovation capacity, increasing its spending on R&D (a roughly 10% increase each year since 2000) and expanding enrollment in higher education.

The "Medium to Long Term Plan for the Development of Science and Technology

[55]Our paper focuses on the effects of patent subsidies on applicants' filing behavior, as we investigagte the immediate increase in patent filings from Zhangjiagang, right after the policy change in June 2006. Patent subsidies may impact firms' R&D behavior as well, as a subsidy on patent filings is also a subsidy on R&D (Segerstrom, 1991; Almus and Czarnitzki, 2003; González et al., 2005; Özçelik and Taymaz, 2008). Because of the lag between R&D activity and patent filings, our results are unlikely related to change in firms' R&D behavior.

[56]For a recent discussion of patent fees by Nikolaus Thumm, EPO Chief Economist, see `http://is.jrc.ec.europa.eu/pages/ISG/patents/documents/NikolausThummfeesandpricing.pdf`

[57]The Executive Summary of Patent Fee Proposal, submitted to the Patent Public Advisory Committee on Feburary 7, 2012 by the USPTO in accordance with the Leahy–Smith America Invents Act, is available at `http://www.uspto.gov/aia_implementation/fee_setting_-_ppac_hearing_executive_summary_7feb12.pdf`

[58]The 863 program, established in March 1986 and also known as the State High-Tech Development Plan, is a program funded by the central government to stimulate the development in a range of key technological fields, including biotechnology, space, information technology, new materials, etc. The 973 program, initiated in 1997 and also known as National Basic Research Program, is a basic research program dedicated to areas such agriculture, health, energy, environment, etc.

(MLP)", initiated in 2006, was the culmination of the extended effort by the Chinese government to promote its "indigenous innovation", enabling China to become an "innovation-oriented society" and a global leader in science and technology. The MLP encompasses several of the Chinese government's long-term policy goals, including building domestic R&D capabilities to upgrade Chinese firms' innovative capacity and promoting domestic firms' contributions to the Chinese economy rather than relying on foreign know-how and technology. In June 2006 the State Council issued a list of rules for implementation of the supporting policies for the MLP. These policies are implemented by government ministries and agencies at all levels.[59]

2.2 China's IP Strategy and Patent Subsidy Policies

The MLP sets objectives for IP creation and commercialization by the year 2020 in areas such as patents and technical standards.[60] Further emphasis on the importance of IP for the goals of the MLP came with the announcement of the National Intellectual Property Strategy in 2008 (planned since 2005) and the recent National Patent Development Strategy (2011-2020) in 2010. Both Strategies set goals for China to become a country generating a comparatively large flow of domestic IP and urge support of market entities to create IP through the use of policies including finance, investment and government procurement. The strategies also support the inclusion of indicators of IPRs in the system for assessing the performance of local governments and state-owned enterprises.

Under these IP goals and mandates, local governments at all levels have implemented patent subsidy policies, to improve patenting awareness and boost innovation.[61] Shanghai implemented China's first patent subsidy policy in 1999, and by around 2003 almost all provinces had some subsidy policies in place. In addition many cities have their own subsidies for patent applications. Patent subsidies at province and city levels come in a variety of forms: some provide a fixed amount of reimbursement for patent applications, regardless of the actual costs or whether the application gets granted or not; some subsidize patent filings based on applicants' actual out-of-pocket spending, usually with a cap; still others compensate applicants with a portion of the application fee and award a prize (usually much larger) if applications get granted.[62] Table 23 shows some of the basic fee estimates for the three types of patent applications.

[59] There are altogether 99 supporting policies. The National Development and Reform Commission (NDRC) is responsible for the largest number of these policies (29), followed by the Ministry of Finance (MOF) with 21 policies, the Ministry of Science and Technology (MOST) with 17, and the Ministry of Education (MOE) with 9 polices. (Serger and Breidne, 2007).

[60] For example, one goal is to become one of the top five countries in terms of invention patents granted to its citizens.

[61] The central government also has a patent subsidy for international filings, which is implemented by the Ministry of Finance.

[62] The reward is only for invention patents, since utility model and design patent applications naturally have an almost 100% grant rate.

Table 23. Patent fees

Type	Application	Examination	Attorney fees	Maintenance/year
Invention	950	2500	4000+[a]	900-8000[b]
Utility Model	500	N/A	2500+	600-2000
Design	500	N/A	1500+	600-2000

[a] The exact legal fee depends on patents and lawyer offices.
[b] The maintenance fee increases incrementally roughly every 3 years.

2.3 China's Patent System and Recent Patenting Surge

China started to envision its first modern patent law in July 1978, shortly after its open-door policy. After a heated debate and 25 revisions, the patent law, modelled after those of Japan and Germany, became effective on April 1st, 1985. The Patent Law has since been amended three times, in 1992, 2001 and 2008, respectively. The first two amendments were made during the process of China's negotiation of entry into the WTO and largely aimed at harmonization of the Chinese Patent Law with those in other countries and the WTO/TRIPS principles.[63] The Patent Law was amended again in 2008, as part of China's effort to promote indigenous innovation.

There are three categories of patents in China: invention patents, utility model patents and design patents. Invention patents are granted after a substantive examination of utility, novelty and non-obviousness; the other two need only a preliminary examination (which results in an almost 100% grant rate). Currently invention patents, utility models and design patents have life terms of 20 years, 10 years and 10 years, respectively. After a patent is issued, the patentee needs to pay annual renewal fees to maintain the patent right; otherwise, the patent is deemed to be abandoned. The renewal fees for invention patents are higher than those for utility models and design patents. Invention patents are of the main interest to researchers due to their similarity to the most prevalent types of patents elsewhere with respect to the standard of patentability and terms of protection.

Patent application growth in China in recent years has drawn international attention (Hu and Mathews, 2008; Eberhardt et al., 2011). Figure 15 shows the invention patent application at China's State Intellectual Property Office (SIPO) since 1985 when the patent law was first implemented. For a long period of time, the number of domestic applications at SIPO has been on par with that of foreign filings. However, domestic filings have been growing more rapidly particularly after 2003-2004, and have since greatly surpassed foreign applications. There is no sign of a slow-down even after 2008, when the global recession reduced the growth of foreign applications at SIPO.

The recent surge in China patenting is especially intriguing,[64] as patents have long been

[63] The first amendment was made when China started its effort to become a member of the WTO and the second right before China entered the WTO.

[64] Some studies investigate the growth of China patenting before 2004. Tests of spillover effect from foreign direct investment (FDI) were inconclusive (Liu, 2002; Cheung and Lin, 2004; Hu and Jefferson, 2009; Girma et al., 2009). Hu and Jefferson (2009) also suggest that other factors, including an intensification of research and development (R&D), entry into WTO, and more importantly the 2000–2001 amendments to the patent law that offered stronger protection to patent holders, all of which contributed to the rise in patenting.

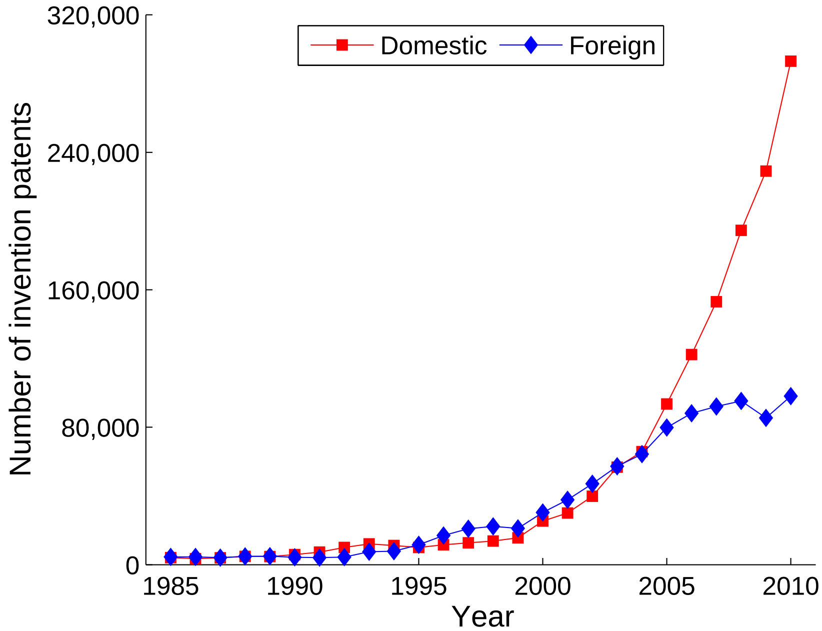

Fig. 15. Patent application growth at SIPO

used as an indicator of innovative activity and technological growth (Griliches, 1990; Kortum, 1997). Some observers argue that one major force driving this China patent boom is the patent subsidy policy implemented by local governments to meet and/or exceed patenting targets specified by the central government.[65] They argue that patent subsidies increase inventors' propensity to file applications for relatively minor or trivial inventions that would have not been filed without those subsides. Therefore, most filings in this China patent boom are "junk inventions", and do not suggest that China has made significant progress in its pursuit of "indigenous innovation".[66]

A recent paper, Li (2012), using invention patent application data at the provincial level, shows that patent subsidy programs at the provincial level are an important factor in the recent patent growth in China, and that the grant rate of patent applications after the

[65]The MLP and the ensuing National Intellectual Property Strategy (2008) specify overall national patenting targets, which are then allocated to local governments.

[66]Patent subsidies can be also viewed as one type of R&D subsidy. Both theoretical and empirical studies have shown that R&D subsidies stimulate private R&D activities and promote economic growth (Segerstrom, 1991; Almus and Czarnitzki, 2003; González et al., 2005; Özçelik and Taymaz, 2008), although some find a crowding-out effect between public and private R&D spending (for example, David et al. (2000)). The impacts of patent subsidies on firms' R&D behavior, though not the focus of the paper, is an interesting research question.

implementation of those subsidy policies has actually increased.[67] However, it is a provincial-level study and cannot take into account the fact that many cities have their own subsidy programs on top of the provincial subsidy program.

3. Research Design and Data

We conduct an applicant level empirical study on the effects of patent subsidies on patent filings, in terms of both quantity and quality of patent filings in China.

We compare six neighboring cities (Zhangjiagang, Wujiang, Taicang, Suzhou, Kunshan and Changshu), all within the Suzhou Municipality in east China's Jiangsu Province.[68] We take advantage of a policy change in patent subsidies in Zhangjiagang in June, 2006, and investigate the effects of this change on patent filings by applicants in Zhangjiagang, using as a control group applicants in other cites whose patent subsidies remained unchanged.

This setting provides us a pseudo natural experiment to identify the effects of patent subsidies on patent filings. In the context of our study, patentees in Zhangjiagang are comparable to those in the other five cities in the same Suzhou Municipality. Also, various policies at the provincial and municipal levels that could be relevant for patent filings are controlled as they are the same for applicants in both Zhangjiagang (the treatment city) and the control cities.

The Suzhou Municipality is close to Shanghai and is one of the most economically developed regions in China. In a survey released in 2005 by The National Bureau of Statistics on economic competence of Chinese small cities, the six cities all ranked among the top 10. Moreover, the economy of Suzhou Municipality is dominated by private small or medium size enterprises.[69] Therefore, the results in the paper may be relevant to the more developed regions in China. Table 24 exhibits some summary statistics of the cities in 2008.

The six cities all implemented patent application subsidies in the year 2003, and have since raised the subsidies a few times. The subsidies came from the city budget and were administered by the Science and Technology Bureau in each city. To get the patent subsidy, an applicant needs to show a receipt for patent filing issued by SIPO that certifies that the patent application has been received and application fees have been paid.

Between July 2004 and December 2007, the subsidy policies remained unchanged in all cities but Zhangjiagang, which had an increase in subsidies for all three types of patent applications on June 12th, 2006. More specifically, Zhangjiagang increased its subsidies for

[67] Another main concern of the subsidy programs is that they may reduce quality, by encouraging applicants to file more low quality patents (Jiachun et al., 2008), with the result that a higher proportion of patent applications are rejected by the patent office, resulting in a waste of resources. An increase in applications may also increase the average workload of the patent examiners, making it more difficult for them to perform a fully comprehensive search of the prior art. This leads to more dubious applications being granted and lower criteria of the patent examination (Philipp, 2006). More application can also mean longer examination times, which serve as a hidden cost of delay. This has not happened in China, as SIPO has greatly enhanced its workforce in the past decade in terms of both the number and qualifications of patent examiners. The average examination time for an invention patent has been stable at roughly 24 months since 2005.

[68] As will be shown, Wujiang fails the parallel trend criterion for being a valid control city and is removed from the study

[69] At the end of 2008, 83.9% of the firms in Suzhou Municipality were privately owned, which are usually much smaller compared to the State owned enterprises.

Table 24. Summary statistics for the cities[a]

	Changshu	Kunshan	Suzhou	Taicang	Wujiang	Zhangjiagang
Area (km^2)	1094.0	864.9	1649.7	620.0	1092.9	772.4
GDP (billion Yuan)	115.0	150.0	271.8	52.8	75.0	125.0
Primary industry	1.9	1.2	2.2	1.8	1.8	1.5
Secondary industry	66.0	97.9	153.7	31.5	46.9	78.3
Tertiary industry[b]	47.1	50.9	115.9	19.5	26.3	45.3
Population (thousand)	1451	1241	3332	665	1096	1189
Per-capita GDP (Yuan)[c]	79263	120881	81571	79449	68434	105156
Number of firms	15435	23798	60504	8064	16346	17132
Average Size of firms	26.7	13.8	13.2	17.2	14.6	21.7

[a] All data are collected at the end of 2008.
[b] R&D is included in this sector.
[c] The national average in 2008 is 23708 Yuan.

patent filings from YMB 1500, 1000, 500 to 3000, 1500, 1000 for applications of invention patents, utility model and design patents, respectively. Moreover, it started to award a prize of YMB 10000 for grant of an invention patent application.[70] See Table 25 for an overview of subsidy policies in the six cities between July 2004 and December 2007.

Table 25. Policy overview – subsidy for the three types of patents

City	July 2004 - June 2006			July 2006 - December 2007[b]		
	type 1	type 2	type 3	type 1	type 2	type 3
Zhangjiagang	1500	1000	500	3000+10000[a]	1500	1000
Wujiang	2000	1000	800		unchanged	
Taicang	4000+5000	1000	1000		unchanged	
Suzhou	4000	1000	1000		unchanged	
Kunshan	4000	1000	500		unchanged	
Changshu	2000	1500	1000		unchanged	

[a] The subsidy after "+" is for granted patents.
[b] We cut off at December 2007 because of three facts: the subsidy policy in Changshu changed after April 2008; a new version of patent law was drafted in early 2008, and there is considerable data truncation after 2008 in our dataset.

As shown in Table 23, the estimated total expenditure, from filing to issuance including legal fees, for an invention patent application is about YMB 8000, and the estimated costs for utility models and design patents are about YMB 3000 and 2000, respectively. Thus the changes in patent subsidies in Zhangjiagang are sizable relative to the total expense for filing patent applications in China. Invention patents are of the highest quality and subject to a similar level of scrutiny as their counterparts in the developed world. They draw most

[70] This award does not apply to filings of utility model and design patents, which are almost sure to be approved due to a lack of substantive examination.

interest from researchers and therefore our study focuses on invention patents only.

3.1 The endogeneity issue

The first question we need to ask, before we move to the treatment effect of the policy change, is whether the subsidy in Zhangjiagang was a response to industry demand. Essentially, we ask whether the policy is endogenous with respect to the patenting activities in Zhangjiagang. We found on the government official website of the treated city Zhangjiagang the following information that we believe was related to the subsidy policy change: On Dec 23, 2005, the city government made some changes in their leadership, and for the first time, a vice director (Mr Yuan, Xu) was assigned to be responsible for the "patent department". Following the change, on Jan 23, 2006, the patent department made an announcement to clarify its duties, which include, among others, drafting and implementing IP policy, building the city as an IP model city, and rolling out the patent subsidy. The subsidy increase was announced on June 12, 2006. Since it's quite common in China for new leaders to bring about new policies in their favor, we believe the information can answer a big part of the endogeneity question of the subsidy policy: it is the result of a leadership reshuffle, which is not likely to be a response to the industry's need.[71]

3.2 The model

We use a difference-in-differences method to study the treatment effect of the subsidy increase in Zhangjiagang, using the other cities as control cities. The model to be estimated is:

$$y_{ict} = \beta \cdot x_{ct} + \alpha_c + \lambda_t + \varepsilon_{ict} \tag{9}$$

where y_{ict} is the number of patent applications by applicant i in city c during half-year t. The policy variable is x_{ct}, which is a dummy term, equaling to 1 for Zhangjiagang after June 2006. The city fixed effect is α_c.[72] The half-year time fixed effect is λ_t, and ε_{ict} is an idiosyncratic error term. The coefficient of interest, the average treatment effect of the policy change on applicantions in Zhangjiagang, is β.

In this paper, we address the robustness of the results in two steps. Firstly, we compute the difference-in-differences estimates of the treatment effect in Zhangjiagang with respect to each of the control cities and the pooled control cities, to make sure that the possible effect we observe is not due to some specific events in one or a few control cities. Secondly, we construct a "placebo treatment" to test the validity of the control cities. For each control city (for example, Changshu), we assume a policy change occurred in June 2006; then we compare Changshu to the remaining control cities by Equation 9 to estimate the "treatment effect"

[71] It is possible that the policy was implemented because the local government sensed a need for the city to improve its firms' patenting activities, i.e., it is a targeted policy, rather than "randomly" assigned. This is not a problem per se, since the policy is still "exogenously" assigned to local firms.

[72] Since the policy variation is at the city level and the panel is balanced, we do not control applicant fixed effects, which should not affect identification. Moreover, in this canonical difference-in-differences setting with no within-group-time-varying explanatory variables, controlling for applicant-fixed effects gives exactly the same estimates as controlling for city fixed effects (though standard errors are slightly different). A simple algebra derivation can be found in Appendix A.

of the policy change on applications in Changshu. If we indeed find a significant treatment effect, it implies that applications in Changshu may not have trend parallel to those of the other control cities. Moreover, it will put doubt on our analysis in this case, since there is no guarantee that applications in the treated city Zhangjiagang would be similar (in the absence of the policy change) to those in the control cities either.

3.3 Data description

We obtained a rich dataset from SIPO covering all three types of patents filed from these cities from 2002 to 2010. We use the period from July 2004 to December 2007 for the purpose of our study. The application data include patent application information, patent type, legal status, and applicant and inventor information. During the study period, there are 3582 applicants (firms and individuals) who applied for 42035 patents of all types (7386 invention patents, 9148 utility models and 25501 utility models). We aggregate patent applications at the applicant level (firms or individuals) and divide the time period into 7 half-years.

Since most of the applicants filed very few invention patents during the study period; while there are some very "large" applicants (for example, Foxconn) that applied for a great number of patents, we remove these "large" patentees to make sure our results are not driven or affected by these few applicants. These large applicants may also respond very differently from the small and mid-sized applicants to the subsidy policies. Figure 16 plots the histogram of total invention applications in all the cities. To make the figure more readable, we do not show applicants that made totally less than (or equal to) 5 invention applications. It appears that 100 applications is a good "cut-off" level.[73] By this criterion, we remove 13 applicants (10 firms and 3 individuals). After this step, we are left with a panel of 3569 applicants with 35414 applications (4399 invention patents, 6957 utility models and 24058 design patents) over 7 time periods.

In Table 26 we provide summary statistics for the number of invention patent applications in the six cities.[74] It can be seen that, in terms of pre-policy-change applications, the treated city (Zhangjiagang) is similar to Changshu and Taicang. A comparison of their post-policy-change applications shows that Zhangjiagang has the largest number of applications. In fact, Zhangjiagang has the largest increase in average applications among all the cities, and the change is significant compared to all the control cities except Wujiang. However, as shown in Section 3.4, the pre-treatment trend in Wujiang is significantly different from that of the treated city Zhangjiagang. Therefore Wujiang is not a valid control city in our study and we will remove it from our controls.

3.4 Test of the parallel trend assumption

In order to estimate the impact of a policy, we need the so-called "parallel trend assumption" to hold: in the absence of a policy change, the period-specific unobservables exhibit parallel

[73] The results are robust to whether we include applicants near this "cut-off" value.

[74] We do not separate firm applicants from individual applicants because the distinction is not clear-cut. Many of the individual applicants are themselves employees or employers of some firms but they file individual patent applications. Moreover, we do not expect the two groups of patentees to respond differently to the policy change.

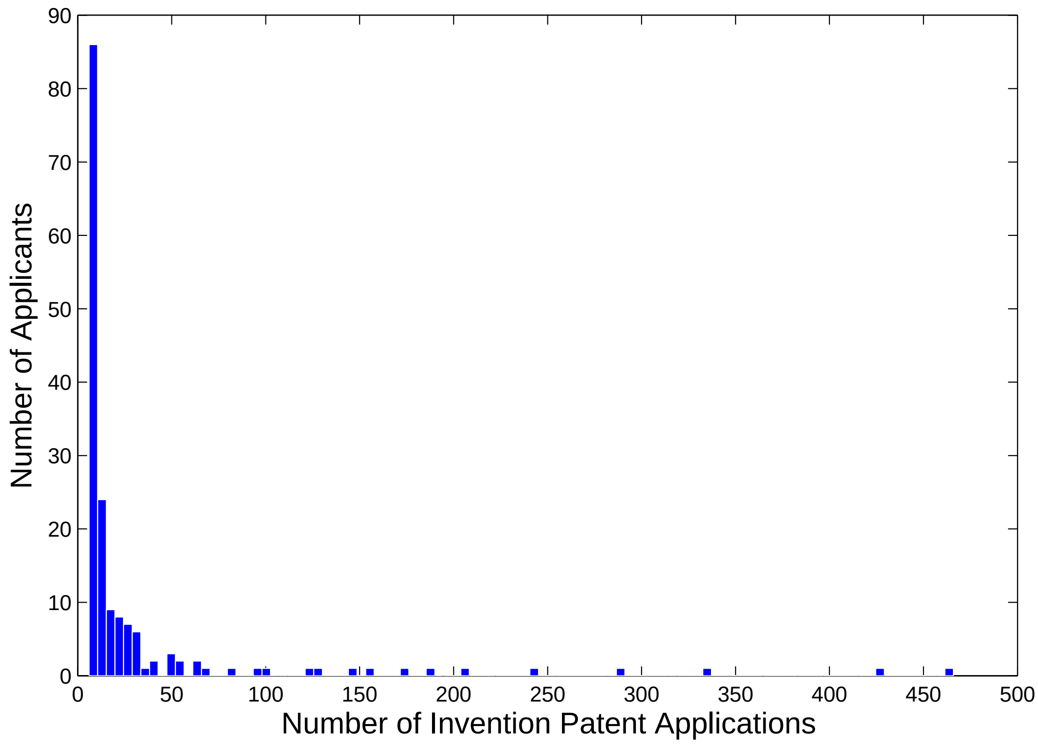

Figure 16. The histogram for applicants of total invention applications

Table 26. Comparison of average invention applications before and after the policy change in Zhangjiagang

City	Before June 2006	After June 2006	# of applicants
Changshu	0.24 (0.04)	0.72 (0.08)	484
Kunshan	0.37 (0.07)	0.84 (0.17)	547
Suzhou	0.43 (0.06)	0.78 (0.10)	1480
Taicang	0.19 (0.04)	0.76 (0.11)	279
Wujiang	0.43 (0.07)	1.37 (0.36)	314
Zhangjiagang	0.19 (0.04)	1.19 (0.13)	465

Standard errors are reported in parentheses.

trend between the treated and control units. In this section, we use the data before the policy change, and test whether a linear time trend interacted with a dummy for being in Zhangjiagang (the treated city) is significant:

$$y_{ict} = \gamma \cdot t \cdot I_{Zhangjiagang} + \alpha_c + \eta \cdot t + \varepsilon_{ict} \tag{10}$$

where α_c is the city dummy and $\eta \cdot t$ controls for the common linear trend.

The results are shown in Table 27. We do not find evidence against the parallel trend assumption for any of the cities except Wujiang, which seems to have a significantly different trend from the treated city in the pre-policy change period. Therefore, the policy effect on the filings of invention patents based on the comparison of Zhangjiagang to Wujiang using the difference-in-differences method may not be very informative. In light of this finding, we do not include the city of Wujiang in the controls.

Table 27. Test of parallel trends between Zhangjiagang and the control cities
$y_{ict} = \gamma \cdot t \cdot I_{Zhangjiagang} + \alpha_c + \eta \cdot t + \varepsilon_{ict}$

	Changshu	Kunshan	Suzhou	Taicang	Wujiang	Pooled (w/o Wujiang)
γ	-0.0018	0.0004	0.0057	-0.0125	-0.0408**	0.0015
	(0.0095)	(0.0112)	(0.0162)	(0.0122)	(0.0175)	(0.0108)

Robust standard errors clustered at applicant level in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

After removing Wujiang from our control cities, the panel consists of 3255 applicants with 27317 applications (3832 invention patents, 6173 utility models and 17132 design patents) over 7 time periods. Next we present some qualitative evidence in graphs. Figure 17 shows the average number of invention patent applications from applicants in Zhangjiagang compared to applications from the pooled control cities (the bar denotes one standard error). Keep in mind that half-years 1–4 constitute the pre-treatment period and half-years 5–7 are in the post-treatment period.[75]

The pre-treatment trends of applications growth in Zhangjiagang and in control cities seem to be parallel. It appears that a difference-in-differences comparison for invention patent applications may result in a positive effect for Zhangjiagang.

4. Results

4.1 The Effect on the Quantity of Invention Patent Applications

4.1.1 Results from Difference-in-differences Study

In this section we report the econometric estimates of the policy treatment effect on the quantity of invention patent applications. As explained in Section **??**, we conduct a pairwise difference-in-differences analysis to estimate the treatment effect, the results of which are reported in the first row of cells in Table 28. The remaining rows of cells provide the estimated "placebo treatment effect". The controls are considered "good" or more valid if all "placebo treatment effects" turn out to be insignificant. We cluster the standard errors at the applicant level to control for serial correlation of the applications for each patentee.

Following our discussion in Section **??**, we observe from Table 28 that the subsidy increase has a rather consistent and significant effect on patentees in Zhangjiagang. On average, the policy change increased the invention patent filed by patentees in Zhangjiagang by roughly

[75]The rational for comparing Zhangjiagang to the pooled control cities relies on the assumption that the control cities have similar trend to each other. We provide comparisons of invention patent applications in Zhangjiagang with each control city in Appendix B. Indeed, except for the city of Wujiang, which is removed from the analysis based on Table 27, all control cities seem to have a parallel trend to the treated city before the policy change.

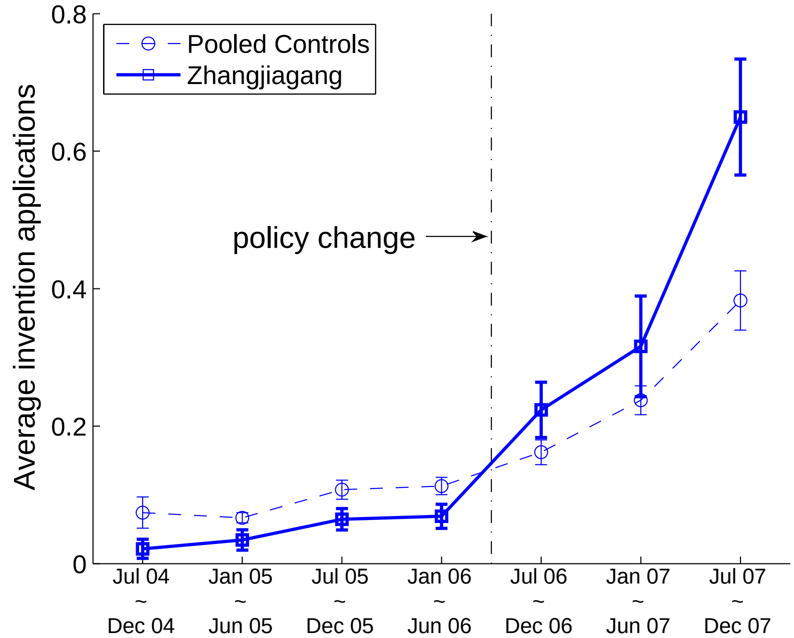

Figure 17. Comparison between Zhangjiagang and the pooled control cities

0.179 applications/half-year, compared to other cities. The average number of invention patent applications per patentee in Zhangjiagang increased by 0.35 applications/half-year, from the pre-treatment average of 0.047 (0.19 divided by 4 half-years) applications/half-year to the post-treatment average of 0.397 (1.19 divided by 3 half-years) applications/half-year. Therefore, the policy change can explain roughly 50% of the increase in Zhangjiagang.[76] Moreover, the "placebo treatment effect" results are very encouraging—none of the placebo policy changes produces a significant estimate. This finding also gives us more confidence in the parallel trend assumption for these cities.

4.1.2 Robustness checks

We conduct the estimation using alternative specifications to test the robustness of the results. Firstly we include applicant level fixed effect in the analysis. As explained in Appendix A, the coefficients will remain the same but the standard errors can be different. The results are reported in the first column of Table 29. It seems when we cluster the standard errors at the applicant level, controlling for applicant level fixed effects causes virtually no change even in the standard errors of the estimates.

One concern is that what we observe in Section 4.1.1 is simply due to the "Ashenfelter's dip" (Ashenfelter, 1978), i.e., applicants in Zhangjiagang anticipated the policy changed and strategically delayed their applications to post-treatment period in order to claim the subsidy. In our study this may be of less concern since applications made during the first half of 2006 (half year before the announcement of the policy change) were also eligible for

[76] See Section 4.1.3 for more analysis on inference.

Table 28. The effect on the quantity of invention patent applications
$y_{ict} = \beta \cdot x_{ct} + \alpha_c + \lambda_t + \varepsilon_{ict}$

Treated/Control	Taicang	Suzhou	Kunshan	Changshu	Pooled Controls
Zhangjiagang	**0.145****	**0.196*****	**0.160****	**0.167*****	**0.179*****
	(0.0567)	(0.0523)	(0.0686)	(0.0500)	(0.0471)
# of applicants	744	1945	1012	949	3255
# of observations	5208	13615	7084	6643	22785
Taicang		0.0514	0.0151	0.0224	0.0379
		(0.0480)	(0.0655)	(0.0455)	(0.0433)
# of applicants		1759	826	763	2790
# of observations		12313	5782	5341	19530
Suzhou			-0.0363	-0.0290	-0.0368
			(0.0616)	(0.0399)	(0.0396)
# of applicants			2027	1964	2790
# of observations			14189	13748	19530
Kunshan				0.00736	0.0237
				(0.0597)	(0.0578)
# of applicants				1031	2790
# of observations				7217	19530
Changshu					0.0141
					(0.0352)
# of applicants					2790
# of observations					19530

Robust standard errors clustered at firm level in parentheses
* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

the subsidy.[77] Since we do not have information of the legal history of the subsidy plan, we still want to be sure that the treatment effects are robust to omitting the period between shortly before and after the policy change. In column 2 of Table 29, we report the estimates omitting data from 2006, half year before and after the policy announcement. The estimated effect is very similar to model (1).[78]

Next we make a few refinements on our applicant selection. Since a lot of the applicants have never made an invention patent application (i.e., they file only for utility model patents and design patents), we exclude these applicants to test whether the effect is robust. Technically this specification removes from the dataset many applicants that made 0 invention patent applications both before and after the policy change. The results, reported in column 3, turn out to be significant and larger.

[77]In an announcement made in August 2006, the first batch of subsidies after the policy change was given to applications filed in the first half of 2006.

[78]Since the reward for the first half of 2006 is retrospective, we also consider the case that applicants knew this fact before 2006 and delayed their applications from late 2005 to early 2006. We therefore drop one year before the policy change (and half year after) and the results are almost the same with estimated effect of 0.215 and p-value of 0.001.

Table 29. Robustness Checks
$y_{ict} = \beta \cdot x_{ct} + \alpha_c + \lambda_t + \varepsilon_{ict}$

	(1) Firm-fixed effect	(2) Ashenfelter's Dip	(3) Invention Firms Only	(4) Unbalanced Panel
β	0.179***	0.215***	0.400***	0.212**
	(0.0471)	(0.0636)	(0.1066)	(0.0875)
Clusters	3255	3255	1237	1684
N	22785	16275	8659	9457

Robust standard errors clustered at applicant level in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

(1) Use applicant fixed-effect instead of city fixed-effect.
(2) Omit data half year before and after the policy announcement.
(3) Exclude applicants that didn't make any invention patent applications.
(2) Use only applicants that "exist".

In the last refinement, we select our applicants by a stricter criterion: since we can only infer the activity of an applicant based on patent filing date, we classify an applicant to be in "existence" if it has ever filed a patent (any type). In this way we construct an unbalanced panel that only consists of applicants that are more surely conducting innovative activities. In this model we include applicant fixed effects which can control non-parametrically for differences across cities in the distributions of different types of applicants, since we have firms enter the sample at different time periods. We report the results in the last column of Table 29. The effect is still significant and comparable to those in model (1) and (2).

4.1.3 Inference

In this section, we discuss the inference problem due to our data constraints. In all the previous analyses, we cluster the data at the applicant level, which addresses the serial correlation problem but assumes independence among firms/inventors within a city. However, applicants in the same city may have a grouped structure and thus are correlated with each other by common group errors. Failing to account for the presence of a common group error (at the city level in our case) can lead to downward biased standard errors (Moulton, 1986). Clustering at city level using a robust covariance estimator can solve the problem only if the number of groups is large (Donald and Lang, 2007). In our study we only have 5 cities, so it is not possible to consistently estimate a robust variance-covariance matrix. This is also the main critique for some influential papers like Card (1990) and Card and Krueger (1994).

In this section, we use two techniques to deal with the inference problem. The first method is the two-step estimator proposed by Donald and Lang (2007). We first average the data at city-by-treatment cells, effectively collapsing the data into 10 cells (5 cities $\times$ 2 periods). Then we calculate the change in average applications (per applicant per half-year) for each city between the two periods, i.e., the first differences in a difference-in-differences setting. We regress these differences on a dummy for being the treated city. If we assume the group errors are homoskedastic,[79] the t-statistic from this regression is distributed asymptotically as t_{c-2} when the number of observations in each group goes to infinity, where c is the number of cities in our case. This is a reasonable approximation in our study since we have hundreds of observations in each city-by-treatment cell. The estimated coefficient from this method is 0.167 with a standard error of 0.0241, which gives a t-statistic of 6.91. With 3 degrees of freedom, p-value is 0.0062, which is significant at the 1% level. The estimate is also very close to the value we get in Section 4.1.1, though the significance level is somewhat lower (as expected).

The two-step estimator still relies on some distributional assumptions that might not hold for our data. To address this issue, we propose a permutation style estimator, which is closely related to Fisher's exact test (Anderson, 2008; Bertrand et al., 2004). We test the null hypothesis that the policy change had no effect in Zhangjiagang. Suppose the policy change is assigned to a randomly chosen city at a random time. We can estimate the effect of the pseudo policy change on the treated cities as in Section 4.1.1. Since the policy is assigned at city level, the autocorrelation structure within the cities are preserved. The calculated placebo policy effects from all the policy assignments form the empirical distribution of the policy effect under the null hypothesis. Essentially, we ask "how likely is it that I observe a change of at least this level if the policy were randomly allocated?" The p-value is thus approximately the proportion of the policy effects that exceed the observed effect. The strength of this test is that it does not require any distributional assumption.[80] Since we have 5 cities over 7 time periods in our study, we can enumerate all cases of the random assignment of the policy. There are $5 \times 6 = 30$ different cases in total.[81] The observed t-statistics of 3.331 is the largest of them in absolute terms (see Figure 18). Therefore we have a p-value of $1/30 = 0.0333$, which implies that the observed t-statistic is significant at 5% level.

4.2 The effect on the grant rate of invention patent applications

One natural question following the results of Section 4.1 is whether the quality of the patent applications dropped after the subsidy increase. In other words, whether applicants took advantage of the subsidy on filing fees and filed some dubious applications that have low patentability. In this section we investigate the change in patentability of the patent applications after the policy change. We consider the grant status of patent applications and conduct a difference-in-differences analysis to take into account any possible change in examination criteria. The model we use is similar to Equation 9. The model to be estimated is:

$$g_{pct} = \beta \cdot x_{ct} + T_i + \gamma_c + \lambda_t + u_{pct} \tag{11}$$

where g_{pct} is a dummy variable indicating the grant of patent p. The policy variable is x_{ct}, which equals 1 for Zhangjiagang after June 2006. Similarly we use the half-year time fixed effect λ_t and city fixed effect γ_c. We add 31 technology fixed effects T_i to the model, to

[79] This assumption may not hold in the difference-in-differences case since the group errors have a time dimension.

[80] Its weakness, however, is that it cannot provide an estimate of the policy effect.

[81] There are 5 cities with 6 time point for each city that we can assign the placebo policy.

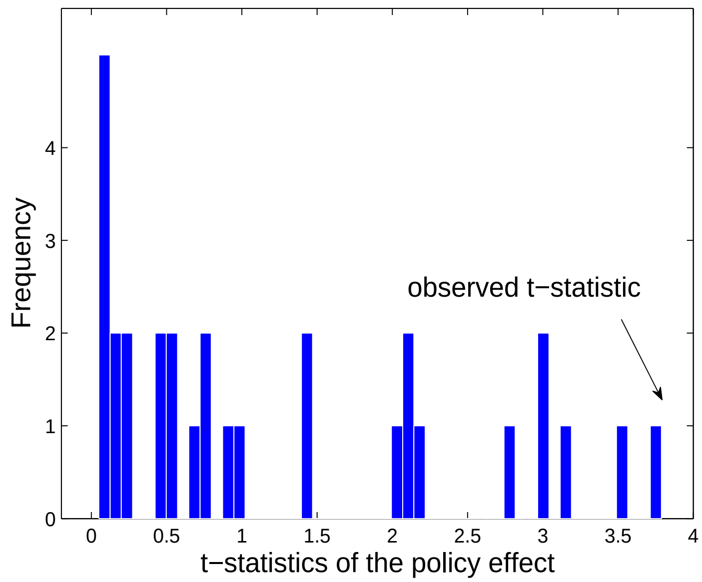

Figure 18. The histogram of the t-statistics from the permutation test

control for any differences in grant rate in technology fields even in the absence of a policy change. u_{pct} is an idiosyncratic error term. The coefficient of interest is still β.

We have the final legal status on almost all patent applications in our sample by the end of 2013 (only 11 of them have legal status missing). Among the 3832 patent applications, 1640 were granted, while the remaining were either rejected or withdrawn. In Table 30 we once more show a table with pair-wise comparison and placebo treatment effect estimation. Based on the estimates, there is no evidence of a change in quality in Zhangjiagang after the subsidy increase. In fact, the difference-in-differences estimates show Zhangjiagang experienced an increase in grant rate after the policy change compared to all cities except Kunshan (though not significant). Therefore there is especially no evidence of a drop in quality compared to the patent applications in control cities.

The results are not surprising if firms and inventors have some reasonably good understanding of the patentability of their inventions. The low subsidy on filing fees is far from enough to cover the application cost, such that applicants will still incur a considerable loss if their applications are rejected. Therefore the subsidy scheme did not encourage them to file more patents that were less patentable.

4.3 The effect on the number of claims

The applications in Zhangjiagang increased significantly after the policy change, while the quality of these applications did not decrease. This might seem to imply that applicants in economically developed regions in China still respond at the margin to financial incentives on patenting, i.e., they have R&D output that were not filed for patent applications because

Table 30. The effect on the grant rate of invention patent applications
$g_{pct} = \beta \cdot x_{ct} + T_i + \gamma_c + \lambda_t + u_{pct}$

Treated/Control	Taicang	Suzhou	Kunshan	Changshu	Pooled Controls
Zhangjiagang	0.286***	0.0399	-0.0339	0.0469	0.0539
	(0.0732)	(0.0762)	(0.0928)	(0.0824)	(0.0716)
# of tech class	28	31	30	28	31
# of observations	904	2434	1301	1105	3821
Taicang		-0.238**	-0.209	-0.211**	-0.217**
		(0.105)	(0.124)	(0.0995)	(0.103)
# of tech class		31	30	28	31
# of observations		2056	923	727	3180
Suzhou			-0.0180	0.0532	0.0493
			(0.0708)	(0.0559)	(0.0507)
# of tech class			31	31	31
# of observations			2453	2257	3180
Kunshan				0.0390	0.0328
				(0.0860)	(0.0724)
# of tech class				30	31
# of observations				1124	3180
Changshu					-0.0378
					(0.0556)
# of tech class					31
# of observations					3180

Robust standard errors clustered at technology class level in parentheses
* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

the cost of patenting is too high.[82] However, before drawing that conclusion, we need to find evidence that these "extra" applications indeed came from the stock of innovation that would not be filed without the subsidy increase. We do not observe firm's R&D behavior or patenting strategies, so we can only infer them from the application data. One important patent characteristic is the number of claims, which is usually a good indicator of patent breadth. If the "extra" applications came from the existing innovation that would otherwise not be filed for patents, we would expect to see an increase in the total number of claims over the patent applications filed by firms and individuals in the treated city.[83] We use the panel of applicants that filed for invention patents to test the hypothesis. We use exactly

[82]A further implication is that firms conducted more innovative work due to the subsidy increase and therefore filed for more applications. We would expect their value or their probability of success to reflect diminishing marginal returns. However, given the short time span of the study, we believe the effect of the subsidy on innovation may not be found yet.

[83]The increase needs not to be proportional to the increase in the number of applications, considering the extra applications may not be as valuable.

the same model as Equation 9,

$$tot_{ict} = \beta \cdot x_{ct} + \alpha_c + \lambda_t + \varepsilon_{ict} \tag{12}$$

except that the dependent variable tot_{ict} becomes the total number of claims from patent applications filed by applicant i of city c in half-year t. The policy variable is x_{ct}, $x_{ct} = 1$ for Zhangjiagang after July 2006. We control for the city and time fixed effects as in Equation 9. The results are reported in Table 31.

Table 31. The effect on the total number of claims
$tot_{ict} = \beta \cdot x_{ct} + \alpha_c + \lambda_t + \varepsilon_{ict}$

Treated/Control	Taicang	Suzhou	Kunshan	Changshu	Pooled Controls
Zhangjiagang	0.592	0.139	-1.443	-0.368	-0.149
	(0.505)	(0.680)	(1.357)	(0.610)	(0.531)
# of applicants	301	754	342	410	1237
# of observations	2107	5278	2394	2870	8659
Taicang		-0.453	-2.035	-0.959	-0.829
		(0.656)	(1.346)	(0.585)	(0.534)
# of applicants		675	263	331	1047
# of observations		4725	1841	2317	7329
Suzhou			-1.582	-0.507	-0.625
			(1.418)	(0.740)	(0.736)
# of applicants			716	784	1047
# of observations			5012	5488	7329
Kunshan				1.075	1.514
				(1.388)	(1.345)
# of applicants				372	1047
# of observations				2604	7329
Changshu					0.277
					(0.661)
# of applicants					1047
# of observations					7329

Robust standard errors clustered at firm level in parentheses
* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Surprisingly, we do not find an increase in the total number of claims for applicants in Zhangjiagang, compared to the control cities. Moreover, three of the four comparisons to the control cities turn out to be negative. Though the effects are in general not significant, it is clear there is no evidence that the total number of claims increased after the policy change. Since the number of applications increased significantly, one plausible explanation is that applicants broke up their patents to get more applications out of the same amount of R&D output. In this way, the technical quality of their applications did not decrease and thus they can claim more rewards from the subsidy program. But in this case, no more innovation output was patented due to the subsidy increase.

We divide the total number of claims by the total number of applications to get the average number of claims per patent application for each applicant (avg_{ict}). Figure 19 shows the average number of claims comparison between Zhangjiagang and the pooled control cities (including the standard error bar). It seems before the policy change, the average number of claims is lower in Zhangjiagang but it follows roughly a similar trend to that in the control cities. However, after the policy change, the average number of claims in Zhangjiagang dropped considerably, compared to that in the pooled control cities. Together with Figure 17, the inferior performance of patenting behaviors in both the quantity and number of claims in Zhangjiagang before June 2006 was likely a factor that led to the increase in subsidy and reward for patent applications. Unfortunately, even though more applications were induced by the policy change, their average number of claims actually dropped. When compared to the other cities, which experienced an increase in the average number of claims, Zhangjiagang's performance in patent quality measured by the number of claims seems to be even worse. A difference-in-differences comparison using a similar model as in Equation 12 confirms that applicants in the treated city Zhangjiagang experienced a significant decrease in the average number of claims (Table 32).

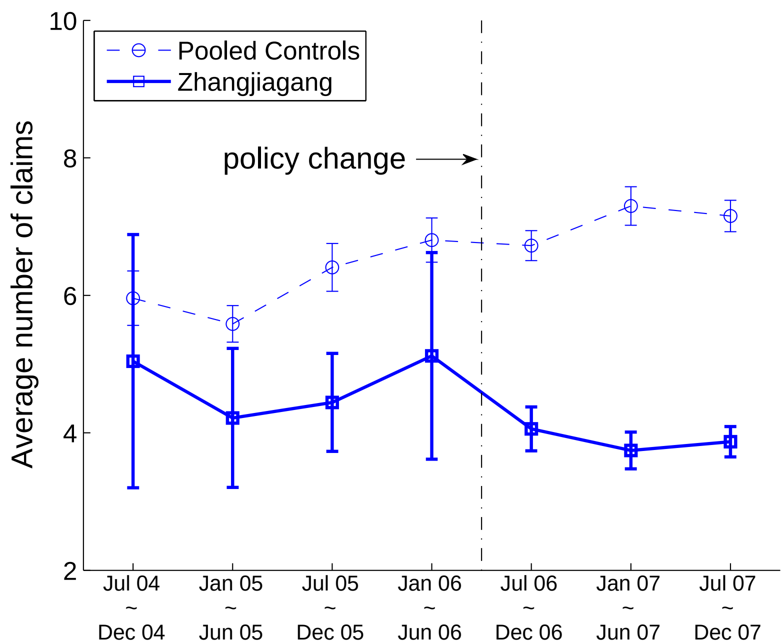

Figure 19. Comparison between Zhangjiagang and the pooled control cities

The finding has important policy implications. The subsidy scheme (which combines a low subsidy on filing fee and a high reward for granted patent applications) does not lead to a decline in patentability, but it also does not lead to an increase in the total number of claims. It seems that applicants, lured by the high reward offered to granted patents, split their patents to get more applications granted and thus more rewards. The rewards together with the subsidy on filing fees in fact give the patentees a monetary gain over their applications. Under the Chinese patent system, the patents are renewed every year. Patentees can always choose to abandon their patents if they turn out to be not valuable.

Table 32. The effect on the average number of claims
$avg_{ict} = \beta \cdot x_{ct} + \alpha_c + \lambda_t + \varepsilon_{ict}$

Treated/Control	Taicang	Suzhou	Kunshan	Changshu	Pooled Controls
Zhangjiagang	**-0.942**	**-1.493***	**-1.431**	**-2.279*****	**-1.493****
	(0.884)	(0.705)	(0.905)	(0.793)	(0.705)
# of applicants	301	754	342	410	1237
# of observations	395	1040	503	554	1712
Taicang		-0.491	-0.276	-1.231*	-0.583
		(0.659)	(0.836)	(0.653)	(0.600)
# of applicants		675	263	331	1047
# of observations		915	378	429	1452
Suzhou			-0.130	-0.825	-0.351
			(0.672)	(0.508)	(0.464)
# of applicants			716	784	1047
# of observations			1023	1074	1452
Kunshan				-0.881	-0.0252
				(0.677)	(0.621)
# of applicants				372	1047
# of observations				537	1452
Changshu					0.898*
					(0.449)
# of applicants					1047
# of observations					1452

Robust standard errors clustered at firm level in parentheses
* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Therefore, with such a subsidy scheme, it is a dominant strategy for applicants to break up their applications. This is not a desirable policy outcome.

Therefore, at least in these economically advanced cities in China, we do not find evidence that firms and individuals have financial constraint on patenting. In other words, they do not seem to have R&D output that would be filed for patent applications if the associated patent fees are reduced. Because if they do, we would expect the total number of claims to increase together with the total number of applications.

5. Conclusion

We evaluate the effectiveness of the patent subsidy policies in China by a case study in Suzhou Municipality, where the subsidy policies resemble many other regions of China. Using a panel of more than 3000 patentees between July 2004 and December 2007, we identify a significant increase in the number of invention patent applications from firms in Zhangjiagang after the city increased the patent subsidy by a considerable amount. Meanwhile, the grant rate of patent applications from Zhangjiagang did not drop. The results suggest that the effect is largely due to the award on granted patents, rather than that of the subsidy on application fees. We infer that a subsidy scheme with low subsidy on filing fees and high reward if the patent is granted may discourage applicants from filing low quality patents.

On the other hand, we find that the total number of claims for each patentee did not increase. Therefore, it's likely that the increase in the number of applications is due to patentees broke up their patents to get more applications (and thus more subsidies). We confirm by showing the average number of claims per patent indeed dropped after the subsidy increase. The fact that applicants can actually make a profit out of the subsidy program if their applications get granted provides incentive for them to split the applications.

Therefore, we find the net effect of the application subsidy and reward upon grant is to motivate patentees to file the same claims in more applications with fewer claims per application. The social welfare effect of the subsidy program is likely to be negative. The extra applications, at the least, increased the workload of both the patentees and the patent office without contributing to more effective patenting.

It seems that firms and inventors in our study region did not face financial constraints in patenting before the subsidy increase. Since the economically developed regions in China file the majority of the patent applications and many use patent subsidy strategies similar to that of Zhangjiagang, our findings put into doubt the necessity of these local patent subsidy policies. The subsidies might merely boost the number of patent filings without actually increasing the stock of patented inventions.

Moreover, our findings show that a patent subsidy that contracts on quantity (number of applications) or even patentability may not guarantee an increase in the total amount of effective patenting. Based on the findings, a better patent subsidy scheme should compensate the applicants only a very small proportion of the patent filing fees, so as to prevent the opportunistic filings; and should let the applicants bear still a small cost or at most make even over a granted patent, such that they don't have the incentive to break up their patents to get more rewards. Providing subsidies to cover part of the maintaining fees in the early years of a patent life, or providing rewards only to patents that go to the product development stage, may achieve the results.

Our finding is different from some surveys in US and EU that show cost of patenting is perceived to be one of the greatest barriers for acquisition of IP rights. In a survey of over 1000 firms in the US manufacturing sector, Cohen et al. (2000) finds that 16% of the respondents cite application cost as the most important reason not to patent. In a recent paper, Rassenfosse and Potterie (2012) find that the drop in patent fees at the EPO contributed to the observed increase in patent filings in the mid-1990s. We believe the main reason is that the cost of patenting under China's current patent system is still quite low, compared to those in the developed world.[84]

Our results raise interesting issues for further research. Do the patent splitting behavior differs across different technology fields, i.e., is it easier to manipulate the number of patents for a given number of claims in certain technology fields than others? Do firms benefit significantly from the patent subsidies, and if not, what are the other incentives for them to

[84] As shown above, the cost for filing an invention patent is around 8000 Yuan, or roughly 1300 Dollar. On the other hand, even in some early estimates, the costs of filing a patent in US and EU were much higher. The average cost of filing a patent in US was estimated to be 10000 to 30000 Dollar (Lemley, 2000). The total external pre-grant cost for a representative EPO patent application was estimated at 10000–15000 Euro in 2004 (available at `http://www.effi.org/system/files?file=cost_anaylsis_2005_study_en.pdf`).

split the patents? The dataset used in the study is small and only included six cities. To what extent can the results be extended to other cities in China, or to other countries? In other words, what are the patent fee reduction policy's effect in other developed countries (US, Korea, Singapore, etc), especially when the number of claims is at concern? Do these policies elicit other patenting behaviors, for example, double patenting?[85] It is also interesting to look at the comparison between patent subsidy policies (which targets one type of R&D output, the patent) and R&D subsidies/tax credit policies (which targets R&D activities directly. Unfortunately, our limited dataset cannot answer these questions effectively.

[85]In China, double patenting is in the form of filing two identical applications, one in invention patent category, the other in utility models. In the US, on the other hand, it is allowed to file two identical utility patents.

A. The equivalence of applicant-fixed effects to city-fixed effects

For simplicity, we assume only two cities, $c = 0$ and $c = 1$, and two time periods, $t = 0$ and $t = 1$. Then in Equation 9, we decompose the error term into two parts: $\varepsilon_{ict} = u_{ct} + \epsilon_{ict}$, such that $\frac{1}{N_{ct}} \sum_i \epsilon_{ict} = 0$. Denote $\bar{y}_{ct} = \frac{1}{N_{ct}} \sum_i y_{ict}$ for any city-time cell. The difference-in-differences estimator is

$$
\begin{aligned}
&(\bar{y}_{11} - \bar{y}_{10}) - (\bar{y}_{01} - \bar{y}_{00}) \\
=&[(\beta + \alpha_1 + \lambda_1 + u_{11}) - (\alpha_1 + \lambda_0 + u_{10})] - [(\alpha_0 + \lambda_1 + u_{01}) - (\alpha_0 + \lambda_0 + u_{00})] \\
=&[\beta + (\lambda_1 - \lambda_0) + (u_{11} - u_{10})] - [(\lambda_1 - \lambda_0) + (u_{01} - u_{00})] \\
=&\beta + [(u_{11} - u_{10}) - (u_{01} - u_{00})]
\end{aligned} \tag{13}
$$

With the common-trend assumption, $E(u_{11} - u_{10}) = E(u_{01} - u_{00})$ and we get the estimate of the policy effect β in expectation. If we use applicant-fixed effect γ_i, the city-fixed effects (α_i's) in the first differences change to $\frac{1}{N_{c1}} \sum_i \gamma_i - \frac{1}{N_{c0}} \sum_i \gamma_i$, which still is canceled out because the panel is balanced. Therefore the estimate with applicant-fixed effect is the same as in Equation 13.

B. City by city comparison of the parallel trend assumption

We compare the average applications of invention patents in Zhangjiagang with each control city in Figure 20. The figures show the average number of applications and the standard error. It seems that except for Wujiang, the parallel trend assumption holds quite well for all cities. Moreover, it seems applicants in Zhangjiagang experienced an increase in average patent applications from a difference-in-difference comparison with each of the control cities except Wujiang.

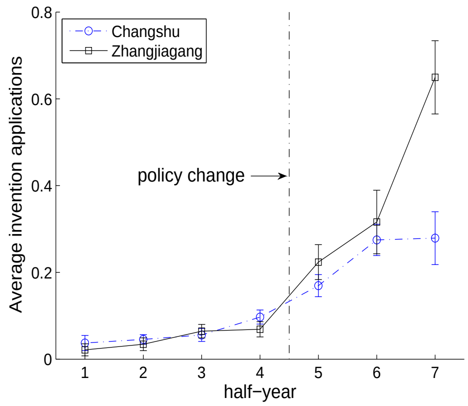

(a) Zhangjiagang vs Changshu

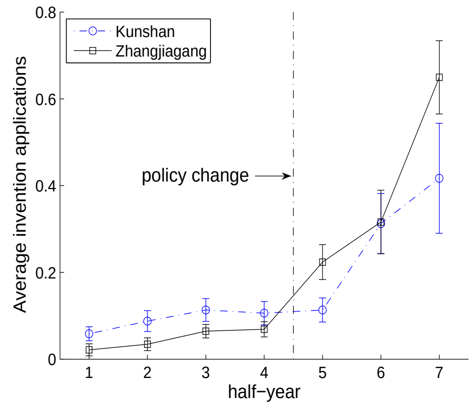

(b) Zhangjiagang vs Kunshan

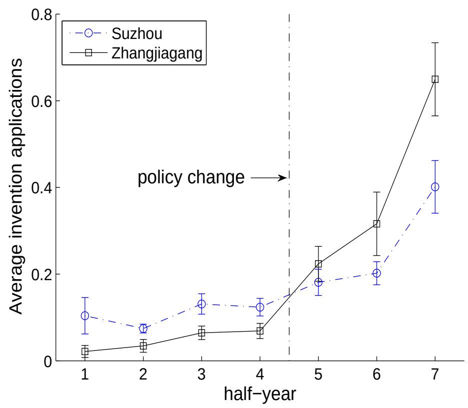

(c) Zhangjiagang vs Suzhou

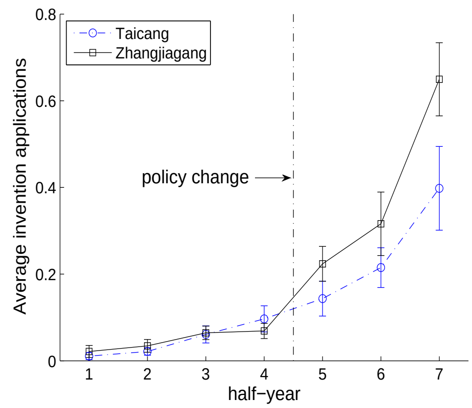

(d) Zhangjiagang vs Taicang

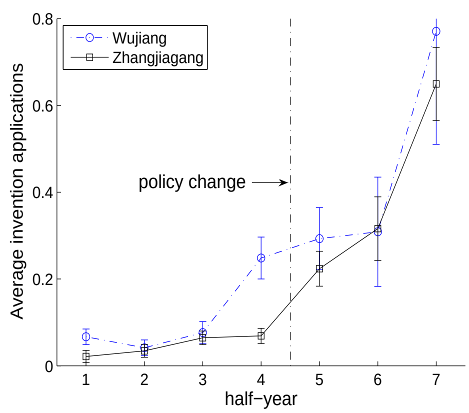

(e) Zhangjiagang vs Wujiang

Figure 20. City by city comparison for invention patent applications

References

Albert, M. B., Avery, D., Narin, F., and McAllister, P. (1991) Direct validation of citation counts as indicators of industrially important patents, *Research policy*, **20**, 251–259.

Alcacer, J. and Gittelman, M. (2006) Patent citations as a measure of knowledge flows: The influence of examiner citations, *The Review of Economics and Statistics*, **88**, 774–779.

Alcacer, J., Gittelman, M., and Sampat, B. (2009) Applicant and examiner citations in US patents: An overview and analysis, *Research Policy*, **38**, 415–427.

Almus, M. and Czarnitzki, D. (2003) The effects of public r&d subsidies on firms' innovation activities, *Journal of Business and Economic Statistics*, **21**, 226–236.

Anderson, M. (2008) Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects, *Journal of the American Statistical Association*, **103**, 1481–1495.

Ashenfelter, O. (1978) Estimating the effect of training programs on earnings, *The Review of Economics and Statistics*, **60**, 47–57.

Baudry, M. and Dumont, B. (2006) Patent renewals as options: improving the mechanism for weeding out lousy patents, *Review of Industrial Organization*, **28**, 41–62.

Bertrand, M., Duflo, E., and Mullainathan, S. (2004) How much should we trust differences-in-differences estimates?, *The Quarterly Journal of Economics*, **119**, 249–275.

Bessen, J. (2008) The value of US patents by owner and patent characteristics, *Research Policy*, **37**, 932–945.

Breschi, S. and Lissoni, F. (2001) Knowledge spillovers and local innovation systems: a critical survey, *Industrial and corporate change*, **10**, 975–1005.

Caillaud, B. and Duchêne, A. (2011) Patent Office in innovation policy: Nobody's perfect, *International Journal of Industrial Organization*, **29**, 242–252.

Cameron, A., Gelbach, J., and Miller, D. (2008) Bootstrap-based improvements for inference with clustered errors, *The Review of Economics and Statistics*, **90**, 414–427.

Card, D. (1990) The impact of the mariel boatlift on the miami labor market, *Industrial and Labor Relations Review*, **43**, 245–257.

Card, D. and Krueger, B. (1994) Minimum wages and employment: A case study of the fast-food industry in new jersey and pennsylvania, *The American Economic Review*, **84**, 772–793.

Carvalho, A. (2011) Why are tax incentives increasingly used to promote private R&D?, Technical report, University of Evora, CEFAGE-UE (Portugal).

Cheung, K. and Lin, P. (2004) Spillover effects of FDI on innovation in China: Evidence from the provincial data, *China Economic Review*, **15**, 25–44.

Cleveland, R., Cleveland, W., McRae, J., and Terpenning, I. (1990) STL: A seasonal-trend decomposition procedure based on loess, *Journal of Official Statistics*, **6**, 3–73.

Cohen, W., Nelson, R., and Walsh, J. (2000) Protecting their intellectual assets: Appropriability conditions and why US manufacturing firms patent (or not), Technical report, National Bureau of Economic Research.

Cornelli, F. and Schankerman, M. (1999) Patent renewals and R&D incentives, *The RAND Journal of Economics*, 197–213.

Cotropia, C. A., Lemley, M. A., and Sampat, B. (2013) Do applicant patent citations matter?, *Research Policy*, **42**, 844–854.

Criscuolo, P. and Verspagen, B. (2008) Does it matter where patent citations come from? Inventor vs. examiner citations in European patents, *Research Policy*, **37**, 1892–1908.

Czarnitzki, D., Hussinger, K., and Leten, B. (2012) Improving citation-based patent value measures, in *Management of Innovation and Technology (ICMIT), 2012 IEEE International Conference on*, IEEE, 589–593.

David, P., Hall, B., and Toole, A. (2000) Is public R&D a complement or substitute for private R&D? A review of the econometric evidence, *Research Policy*, **29**, 497–529.

De Rassenfosse, G. (2012) Are patent fees effective at weeding out low quality patents?, *Working Paper*.

De Rassenfosse, G. and Van Pottelsberghe de la Potterie, B. (2012) The role of fees in patent systems: Theory and evidence, *Journal of Economic Surveys*.

Deng, Y. (2011) A dynamic stochastic analysis of international patent application and renewal processes, *International Journal of Industrial Organization*, **29**, 766–777.

Donald, S. and Lang, K. (2007) Inference with difference-in-differences and other panel data, *The Review of Economics and Statistics*, **89**, 221–233.

Duguet, E. and MacGarvie, M. (2005) How well do patent citations measure flows of technology? Evidence from French innovation surveys, *Economics of Innovation and New Technology*, **14**, 375–393.

Eberhardt, M., Helmers, C., and Yu, Z. (2011) Is the dragon learning to fly? An analysis of the chinese patent explosion, *The New York Times*, 1.

Farrell, J. and Shapiro, C. (2010) How strong are weak patents?, *American Economic Review*, **98**, 1347.

Gallini, N. (2002) The economics of patents: Lessons from recent US patent reform, *The Journal of Economic Perspectives*, **16**, 131–154.

Gambardella, A., Harhoff, D., and Verspagen, B. (2008) The value of European patents, *European Management Review*, **5**, 69–84.

Gans, J., King, S., and Lampe, R. (2004) Patent renewal fees and self-funding patent offices, *Topics in Theoretical Economics*, **4**, 1–13.

Girma, S., Gong, Y., and Görg, H. (2009) What determines innovation activity in Chinese state-owned enterprises? The role of foreign direct investment, *World Development*, **37**, 866–873.

González, X., Jaumandreu, J., and Pazó, C. (2005) Barriers to innovation and subsidy effectiveness, *RAND Journal of Economics*, 930–950.

Griffith, R., Lee, S., and Van Reenen, J. (2011) Is distance dying at last? Falling home bias in fixed-effects models of patent citations, *Quantitative economics*, **2**, 211–249.

Griliches, Z. (1990) Patent statistics as economic indicators: a survey., *Journal of Economic Literature*, **28**, 1661–1707.

Hall, B. H., Jaffe, A., and Trajtenberg, M. (2005) Market value and patent citations, *RAND Journal of economics*, 16–38.

Hall, B. H., Jaffe, A. B., and Trajtenberg, M. (2001) The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools, Working Paper 8498, National Bureau of Economic Research.

Hansen, C. B. (2007) Asymptotic properties of a robust variance matrix estimator for panel data when T is large, *Journal of Econometrics*, **141**, 597–620.

Harhoff, D., Narin, F., Scherer, F. M., and Vopel, K. (1999) Citation frequency and the value of patented inventions, *Review of Economics and statistics*, **81**, 511–515.

Hegde, D. and Sampat, B. (2009) Examiner citations, applicant citations, and the private value of patents, *Economics Letters*, **105**, 287–289.

Henderson, R., Jaffe, A., and Trajtenberg, M. (2005) Patent citations and the geography of knowledge spillovers: A reassessment: Comment, *American Economic Review*, 461–464.

Hirschey, M. and Richardson, V. J. (2004) Are scientific indicators of patent quality useful to investors?, *Journal of Empirical Finance*, **11**, 91–107.

Hu, A. and Jefferson, G. (2009) A great wall of patents: What is behind China's recent patent explosion?, *Journal of Development Economics*, **90**, 57–68.

Hu, M. and Mathews, J. (2008) China's national innovative capacity, *Research Policy*, **37**, 1465–1479.

Hunt, R. (2006) When do more patents reduce R&D?, *The American Economic Review*, **96**, 87–91.

Jaffe, A. B., Fogarty, M. S., and Banks, B. A. (1998) Evidence from patents and patent citations on the impact of NASA and other federal labs on commercial innovation, *The Journal of Industrial Economics*, **46**, 183–205.

Jaffe, A. B. and Trajtenberg, M. (1999) International knowledge flows: evidence from patent citations, *Economics of Innovation and New Technology*, **8**, 105–136.

Jaffe, A. B. and Trajtenberg, M. (2002) *Patents, citations, and innovations: A window on the knowledge economy*, MIT press.

Jaffe, A. B., Trajtenberg, M., and Fogarty, M. S. (2000) Knowledge Spillovers and Patent Citations: Evidence from a Survey of Inventors, *American Economic Review*, **90**, 215–218.

Jaffe, A. B., Trajtenberg, M., and Henderson, R. (1993) Geographic localization of knowledge spillovers as evidenced by patent citations, *The Quarterly Journal of Economics*, **108**, 577–598.

Jiachun, W., Yongtao, Z., and Xuezhong, Z. (2008) Research on patent fees subsidized by local government in china, in *Information Management, Innovation Management and Industrial Engineering, 2008. ICIII'08. International Conference on*, volume 3, IEEE, 269–274.

Kortum, S. (1997) Research, patenting, and technological change, *Econometrica*, **65**, 1389–1420.

Krugman, P. R. (1991) *Geography and trade*, MIT press.

Lampe, R. (2012) Strategic citation, *Review of Economics and Statistics*, **94**, 320–333.

Lanjouw, J. and Schankerman, M. (2001) Characteristics of patent litigation: a window on competition, *RAND Journal of Economics*, 129–151.

Lanjouw, J. O. (1998) Patent protection in the shadow of infringement: Simulation estimations of patent value, *The Review of Economic Studies*, **65**, 671–710.

Lanjouw, J. O. and Schankerman, M. (2004) Patent quality and research productivity: Measuring innovation with multiple indicators, *The Economic Journal*, **114**, 441–465.

Lei, Z., Sun, Z., and Wright, B. (2013) Patent subsidy and patent filing in China, *Working Paper*.

Lemley, M. (2000) Rational ignorance at the patent office, *Nw. UL Rev.*, **95**, 1495.

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

Li, X. (2012) Behind the recent surge of Chinese patenting: An institutional view, *Research Policy*.

Liu, Z. (2002) Foreign direct investment and technology spillover: Evidence from China, *Journal of comparative Economics*, **30**, 579–602.

Marco, A. and Prieger, J. (2009) Congestion pricing for patent applications, in *Workshop on Innovation, Intellectual Property and Competition Policy*, volume 18.

Marshall, A. (1920) Principles of Economics.

Mehta, A., Rysman, M., and Simcoe, T. (2010) Identifying the age profile of patent citations: New estimates of knowledge diffusion, *Journal of Applied Econometrics*, **25**, 1179–1204.

Mokyr, J. (2009) Intellectual property rights, the industrial revolution, and the beginnings of modern economic growth, *The American Economic Review*, 349–355.

Moulton, B. (1986) Random group effects and the precision of regression estimates, *Journal of econometrics*, **32**, 385–397.

Özçelik, E. and Taymaz, E. (2008) R&D support programs in developing countries: The Turkish experience, *Research Policy*, **37**, 258–275.

Pakes, A. (1986) Patents as options: some estimates of the value of holding European patent stocks, *Econometrica*, **54**, 755–784.

Pakes, A. and Schankerman, M. (1984) The rate of obsolescence of patents, research gestation lags, and the private rate of return to research resources, in *R & D, Patents, and Productivity*, University of Chicago Press, 73–88.

Peri, G. (2005) Determinants of knowledge flows and their effect on innovation, *Review of Economics and Statistics*, **87**, 308–322.

Philipp, M. (2006) Patent filing and searching: Is deflation in quality the inevitable consequence of hyperinflation in quantity?, *World Patent Information*, **28**, 117–121.

Rassenfosse, G. and Potterie, B. (2012) On the price elasticity of demand for patents, *Oxford Bulletin of Economics and Statistics*.

Ridker, P. M., Rifai, N., Rose, L., Buring, J. E., and Cook, N. R. (2002) Comparison of C-reactive protein and low-density lipoprotein cholesterol levels in the prediction of first cardiovascular events, *New England Journal of Medicine*, **347**, 1557–1565.

Rigobon, R. and Stoker, T. M. (2007) Estimation with Censored Regressors: Basic Issues, *International Economic Review*, **48**, 1441–1467.

Scotchmer, S. (1999) On the optimality of the patent renewal system, *The RAND Journal of Economics*, 181–196.

Segerstrom, P. (1991) Innovation, imitation, and economic growth, *Journal of Political Economy*, 807–827.

Serger, S. and Breidne, M. (2007) Chinas fifteen-year plan for science and technology: An assessment, *Asia Policy*, **4**, 135–164.

Shane, H. and Klock, M. (1997) The relation between patent citations and Tobin's Q in the semiconductor industry, *Review of Quantitative Finance and Accounting*, **9**, 131–146.

Steensma, H. K., Chari, M., and Heidl, R. (2014) The quest for expansive intellectual property rights and the failure to disclose known relevant prior art, *Strategic Management Journal*.

Thompson, P. (2006) Patent citations and the geography of knowledge spillovers: evidence from inventor-and examiner-added citations, *The Review of Economics and Statistics*, **88**, 383–388.

Trajtenberg, M. (1990) A penny for your quotes: patent citations and the value of innovations, *The Rand Journal of Economics*, 172–187.

Wagner, R. P. (2002) Reconsidering estoppel: Patent administration and the failure of Festo, *University of Pennsylvania Law Review*, 159–244.

Weitzman, M. L. (1974) Prices vs. quantities, *The review of economic studies*, **41**, 477–491.

Wilson, P. W., DAgostino, R. B., Levy, D., Belanger, A. M., Silbershatz, H., and Kannel, W. B. (1998) Prediction of coronary heart disease using risk factor categories, *Circulation*, **97**, 1837–1847.

Zhang, J. and Kai, F. Y. (1998) What's the relative risk?: A method of correcting the odds ratio in cohort studies of common outcomes, *Journal of American Medical Association*, **280**, 1690–1691.