

Patents and the Performance of Voluntary Standard Setting Organizations *

Marc Rysman

Boston University, Department of Economics

Tim Simcoe

J.L. Rotman School of Management, University of Toronto

DRAFT July 2005

*Financial support for this research was provided by CITRIS, and the Intel Robert Noyce Memorial Scholarship. Useful comments were received from Kevin Lang, Josh Lerner, David Mowery, Bronwyn Hall, Avi Goldfarb and the seminar audience at the NBER Productivity Lunch. Address for correspondence: Joseph L. Rotman School of Management, 105 St. George Street, Toronto, ON M5S 3E6, Canada. E-mail: timothy.simcoe@rotman.utoronto.ca

Patents and the Performance of Voluntary Standard Setting Organizations

Abstract

This paper measures the technological significance of voluntary standard setting organizations (SSOs) by examining citations to patents disclosed in the standard setting process. We find that SSO patents are cited far more frequently than a set of control patents, and that SSO patents receive citations for a much longer period of time. Furthermore, we find a significant correlation between citation and the disclosure of a patent to an SSO, which may imply a marginal impact of disclosure. These results provide the first empirical look at patents disclosed to SSOs, and show that these organizations not only select important technologies, but may also play a role in establishing their significance.

1 Introduction

This paper studies the economic and technological impact of voluntary collaborative non-market standard setting organizations (SSOs). SSOs are an important catalyst for coordination in many industries where consumers value inter-operability (e.g. telecommunications and computing). These organizations provide a forum for collective decision-making and an alternative to standardization through market competition or through government regulation. SSOs come in a variety of shapes and sizes—from large industry associations to small consortia—and are often involved in a variety of different activities, including collaborative R&D, compatibility testing, and product certification. SSOs generally do not have formal powers to enforce their recommendations. As a result, these groups work to create a consensus around particular technologies that can serve as a focal point for industry coordination or lead to a bandwagon process among adopters.

Several authors have documented the substantial resources devoted to SSOs and the standard setting process (Farrell 1996; Cargill, 1997). However, our knowledge of the economic and technological impact of these institutions remains quite limited. Evaluating the role of SSO's is difficult because they operate in diverse markets and their effect on such standard variables as price and quantity is uncertain.¹ However, a ubiquitous problem for SSO's is the treatment of intellectual property. Participants regularly must disclose relevant patents to SSO's in the process of negotiating a standard. In this paper, we use these patents as a window into the role of SSO's in technological innovation. Patents are easily compared across time and industries, and many properties are well-known as a result of a large amount of research in economics.

Following the literature on patents in economics, we use patent citations as a measure of economic and technological innovation (Jaffe and Trajtenberg, 2002). We use patents identified in the intellectual property disclosure records of four SSOs: the European Telecommunications Standards Institute (ETSI), the Institute for Electrical and Electronic Engineers (IEEE), the Internet Engineering Task Force (IETF), and the International Telecommunications Union (ITU). We construct control samples based on technological class and application year of the patents. In our first set of results, we find that SSO patents receive far more citations than an average patent, around 3 times higher. More surprisingly, SSO patents receive citations over a much longer time period. Building on techniques of Hall, Jaffe and Trajtenberg (2001), we show that the age profile of patent citations is higher for SSO's than control patents and

¹We know of a few empirical studies of standard setting organization in the economics literature, all very recent. See Chiao, Lerner and Tirole (2005), Gandalf, Gantman and Genesove (2005), Toivanen (2004) and Blind (2005). There is a large, mostly case study literature (for example, Bolin (2004) and the new journal *International Journal of IT Standards and Standardization Research*) that has few members of academic economics departments, although this is changing (see Greenstein and Stango, 2005).

that the difference is economically and statistically significant. Interestingly, this difference is greater when we compare SSO patents to a group of highly cited control patents.

Two reasons that SSO patents differ from other patents are that the SSO *selects* patents that represent important technologies and that the SSO actually *causes* technologies to have the citation profile we observe. That is, we may wonder whether SSO patents would have had similar citation patterns if they had never been associated with an SSO. The selection effect is natural given that SSO's explicitly attempt to identify the best technology to serve a given need. Finding that the selection effect is important suggests that SSO are successful in identifying important technologies. The causal effect may arise because an SSO embeds a technology in a standard that then exhibits long-lasting economic importance because of network effects and lock-in. Another source for a causal effect may be that because an SSO disclosure represents a public announcement, it attracts attention to a patent. Finding a causal effect for SSO's suggests that over and above the stated goals of SSO's in facilitating interconnection between complementary markets, SSO's have a further role in determining the path of technological innovation into the future.

In this paper, we exploit the timing of disclosures to separate between the selection and causation effects. That is, the extent to which the citation pattern changes after a patent is disclosed to an SSO gives a measure of the causal effect of the SSO. We are cautious in this interpretation as the timing of disclosure depends on the economic environment. Below, we discuss why the endogeneity of disclosure could lead us to over or under-estimate the causation effect. However, given the lack of a truly exogenous determination of disclosure, we find this approach a logical starting place.

Our regression approach compares disclosed to undisclosed SSO patents and compares patents before and after disclosure. We find an economically and statistically significant correlation of citations with disclosure. To the extent that we measure the causal effect of an SSO, it appears that the causal effect represents between 20% and 26% of the difference in citations between SSO and non-SSO patents.

This paper contributes to a growing empirical literature that examines the impact of particular institutions on the process of technological change. Examples of this research include Furman and Stern's (2004) study of biological resource centers, and studies of the university-industry interface, including Mowery et al (1999) and Markiewicz (2004). In the next section, we describe the four SSOs that are examined in the paper and how they treat intellectual property. Section 3 describes the data set, while Section 4 takes an initial look at the difference in citation patterns between the SSO and control samples. Section 5 examines the post-disclosure increase in citation rates. Section 6 offers some conclusions.

2 SSOs and Intellectual Property

Before using patent data to study the role of SSO’s in the innovation process, it is important to understand the role of patents and intellectual property in the standard setting process. This section describes the four organizations studied below and describes how each of them deals with intellectual property.

We use data collected from four major SSO’s. These groups are the European Telecommunication Standards Institute (ETSI), the Institute of Electrical and Electronics Engineers (IEEE), the Internet Engineering Task Force (IETF), and the International Telecommunications Union - Telecommunication Standardization Sector (ITU-T, or often, ITU). Both ETSI and the ITU are international institutions that focus primarily on telecommunications standards. While international in scope, the IEEE and IETF draw the majority of their participants from North America, and are usually associated with the computer hardware and software industries although some of their most significant standards are communications protocols. Table 1 provides some indication of the relative scope of these four organizations based on the technology class assigned to each disclosed patent.

Table 1: Technology Classification of SSO Patents[†]

	ETSI	IEEE	IETF	ITU-T	Totals
Computers & Communications	532	109	23	62	726
Computers Hardware & Software	94	83	53	59	289
Information Storage	4	7	2	0	13
Electronic Business Methods	0	2	2	0	4
Electrical Devices	7	7	0	1	15
Measuring and Testing	0	2	0	1	3
Power Systems	0	2	0	0	2
Semiconductor Devices	0	9	0	0	9
Misc. Electrical	2	1	0	39	42
Optics	0	1	0	9	10
Total Patents	639	223	80	171	1,113

[†]Based on subcategory classifications in the NBER US patent database.

Of the four SSOs that we examine, the ITU is the oldest, with origins dating back to around 1865. Its original mission was to promote international coordination among the various rapidly expanding domestic telephone networks. The ITU is based in Switzerland and is associated with the United Nations. Its membership consists of delegates from member nations along with representatives of the larger firms or network operators in each of these countries. The organization’s standard setting activities continue to emphasize the protocols used to operate

the international telephone network, with work areas that include numbering and addressing, network services, physical interconnection, monitoring and accounting, traffic management, and quality of service.

The IEEE is only slightly younger than the ITU. It was founded in 1884 by several pioneers in the field of electrical engineering. Although the IEEE is a professional society whose members are individual engineers, it is possible to become a corporate member when participating in its standard setting activities. The IEEE's standard setting efforts cover a wide range of subjects, from electrical safety, to cryptography, to standards for semiconductor testing equipment. In recent years, the IEEE's most commercially significant standards work has revolved around the 802.11 specifications for wireless computer networking.

ETSI was formed in 1988 to provide a less formal and more industry-driven forum than the ITU for European telecom standardization. The organization is located in southern France and participants are typically firms—as opposed to the member-state representatives of the ITU. ETSI has played a prominent role in creating several generations of mobile telephony standards that are in use throughout Europe and much of the rest of the world. In particular, it is the forum where a variety of network operators, electronics suppliers, and OEM handset manufacturers reached key agreements on the GSM and 3G wireless protocols.

Finally, the IETF is the least formal of the four SSOs studied in this paper. This organization grew out the ARPANET engineering community that emerged during the 1970s, and did not resemble a formal SSO until the late 1980s or early 1990s (Mowery and Simcoe, 2002). The IETF creates a host of protocols used to run the internet. Prominent examples include the internet's core transport protocols (TCP/IP and Ethernet), standards used to allocate network addresses (DHCP), and specifications used by popular applications such as e-mail or file transfer. From its inception, membership in the IETF and its various working groups has been open to any interested individual. Much of the IETF's work takes place in online forums sponsored by individual committees and is visible to the general public.

Because all four of the SSOs examined in this paper are more or less “open” each of them must deal with the increasing tension between open standards and intellectual property protection. The goal of most SSOs is to promote widespread implementation and adoption of the specifications they produce. However, these goals often conflict with those of individual participants who may hold intellectual property rights in a proposed standard. Patent owners frequently seek royalty payments for the use of their technology—even (or, perhaps, especially) when it is essential to the implementation of an industry standard. Moreover, many firms realize that owning intellectual property rights in an industry standard can result in substantial licensing revenues. This creates strong incentives to push for one's own technology within the SSO, and may lead to long-delays or breakdowns in the standard setting process (Simcoe 2004).

While most SSOs would like to avoid the distributional conflicts and obstacles to implementation that patents can produce, they often have no choice other than to evaluate a variety of proposals that are subject to some type of intellectual property protection. In part, this is because of the well-documented surge in patenting that began in the mid-1980s. This increase reflects a growing awareness of patents' strategic significance, as well as the actions of courts, policy-makers, and patent offices. Awareness of the tension between SSOs and their patent-holding members has also increased because of a number of high-profile legal conflicts.²

Given the increasing importance of the intellectual property issue, many SSOs have been debating their own policies for dealing with patents. Lemley (2002) presents a survey of the various policies that SSOs have adopted. All four of the SSOs examined in this paper use variations on the relatively common policy of "reasonable and non-discriminatory" licensing (RAND). Under this policy, SSO members agree to disclose any known property rights as soon as they become aware of them. (They are not, however, obliged to carry out a search.) When a patent or other piece of intellectual property is discovered, the SSO seeks assurances from the owner that they will license the technology to any interested standards implementor on reasonable and non-discriminatory terms.³ While SSO's and their individual committees are generally inclined to search for technologies that are unprotected or available on a royalty-free basis, their job is to evaluate the potential tradeoff between technical quality and openness.⁴

Figure 1 illustrates the growth in intellectual property disclosures at the four SSOs that we study. (We define a disclosure is an announcement by a single firm on a given date that it potentially owns one or more pieces of intellectual property.) While the number of intellectual property disclosures was initially quite small, it began to grow during the early 1990's. By the late 1990s, all four SSOs were experiencing significant growth.

For our purposes, the rise in intellectual property disclosures means that we have access to a publicly available list of patents associated with standard setting. Many features of patents,

²The most well-known is the Rambus case. The documents for the case can be found at the FTC web site <http://www.ftc.gov/os/adjpro/d9302/>

³In practice, the "reasonable and non-discriminatory" requirement in a RAND licensing policy seems to imply very few obligations on the part of prospective licensors. The reasonableness requirement is rarely taken to mean that the technology must be offered at a uniform price. When the intellectual-property holder has not made an *ex ante* commitment to some set of licensing terms, each potential implementor of the standards will negotiate their own terms. While licensors are expected to negotiate in good faith with any potential developer, the individual terms offered may vary widely. SSOs have been very hesitant to get further involved in the negotiating process. In part this reflects their own concerns about the antitrust implications associated with any type of collective pricing agreement. At the same time, it also likely reflects their fear of alienating particular members.

⁴Each of the SSOs considered below uses some variation on a RAND IPR policy. All of them have produced specifications that contain proprietary technology at some point. The IETF's policy is the least centralized, in that it leaves most of the decision-making to the discretion of its individual working groups. IEEE follows the guidelines established by ANSI. ETSI's policy is the most explicit of the four. In particular, it specifies a set of rules for dealing with a situation where some intellectual property is determined to be essential to a standards development effort.




Figure 1: Intellectual Property Disclosures

such as the number of citations they receive, are easily compared across different industries and time periods. We utilize the information contained in intellectual property disclosures to identify standards-related patents whose citation rates may provide a window onto the potential impact of SSOs.

A note of caution: patents may be disclosed for proposals that never become standards, and proposals may become standards but not require licensing of every patent that was disclosed in relation to the proposal. We observe only intellectual property disclosures, not whether they were included in the final standard, whether the proposal became a standard or often even what proposal they were disclosed in relation to. Making these distinctions might be useful for a number of questions and such data are the subject of current search.

3 Data

This section describes the sample of SSO patents that we use to study the standard setting process. All of these data were collected from the publicly available records of ETSI, the IEEE, IETF, and ITU. We begin by describing the complete sample of intellectual property disclosures. We then examine the 1,113 US patents contained in one or more of these disclosures. After discussing some of the issues associated with these patents, we conclude by describing the creation of our initial control sample.

Although the four SSOs in this study have similar intellectual property policies, the scope

and specificity of individual disclosures varies dramatically across organizations. These differences reflect variation in the participants, policies, and objectives of the four institutions. In order to provide some intuition for the type of disclosure information provided by these SSOs, we group the data using a particular definition of “disclosure.” We define a disclosure as an announcement by a single firm on a given date that it owns (or may own) one or more pieces of intellectual property related to a single standard setting initiative. When a firm claims that a single patent covers two or more standards, each one counts as a separate disclosure. When a single announcement lists more than one patent or patent application, we will refer to each piece of intellectual property in the disclosure as a claim. (Since we do not work with the claims data from individual patents, this should not lead to much confusion.)

Table 2: Intellectual Property Disclosures

	Disclosures		Claims per Disclosure			Patents	
	Earliest	Total (Count)	Mean	Median	Max	All Patents	US Patents
ETSI	April, 1990	262	36.5	4	1582	847	672
IEEE	January, 1988	125	3.4	1	37	313	252
IETF	June, 1995	314	1.5	1	27	193	97
ITU-T	October, 1983	821	1.0	1	2	339	188

Table 2 illustrates some of the variation in how intellectual property is disclosed across the four SSOs in this study. First, the data for each organization begins at a different point in time. While the ITU disclosures begin in 1983, intellectual property did not become an issue at the IETF until 1995. Second, there are substantial differences in the number of claims per disclosure. While the ITU has the largest number of disclosures, almost all of them contain a single claim. At ETSI, on the other hand, the median disclosure makes four claims, and one contains more than 1500. Finally, individual claims vary in their level of specificity. For example, it was a common practice at the IETF for several years to “disclose” the existence of an unpublished patent application without providing any information that could be used to verify its existence. This variation in claim-specificity can be seen by comparing the number of patents disclosed to the total number of claims at an SSO. Figure 2 shows the percentage of disclosures that list one or more US patent numbers at each SSO.

The final column in Table 2 shows the number of US patents contained in the data set. This figure is smaller than the number of patents claimed at each SSOs for two reasons. First, many disclosures list non-US patents. This is particularly true at ETSI, where the large number of claims per disclosure often reflects the disclosure of patent families which cover the same invention in several legal jurisdictions. Second, there are several patents that get disclosed more




Figure 2: Disclosures Listing US Patents

than once (both within and between SSO's). For example, there are a number of cryptography patents that seem to be disclosed on a regular basis when SSOs deal with issues of computer security.

After removing all of the foreign patents, patent applications, and duplicate observations, the intellectual property disclosures made at ETSI, the IEEE, IETF and ITU yield a sample of 1,113 unique US patents. We do not claim that these patents are broadly representative of the technology evaluated by these four SSOs. More likely, they are concentrated within several of the most commercially significant standard setting efforts. Nevertheless, these patents provide a unique opportunity to study the role of SSOs in the innovation process.

We obtained citation data for these patents by linking the SSO sample to the NBER US patent data file (Hall, Jaffe, Trajtenberg 2001).⁵ These data also contains several important patent characteristics, such as application and grant dates, and the name of assignees. Figure 3 shows the distribution of grant dates for the patents in the SSO sample.

It is clear from Figure 3 that the majority of patents listed in SSO disclosures were not granted by the USPTO until the mid-1990s. This is not surprising, given the surge in patenting and the timing of the disclosures in Figure 1. However, because these are relatively new patents, it is important to consider the issue of sample truncation. In particular, many of the SSO patents were granted near the end of our sample (our citation data extends to 2004). While

⁵These data have been updated through 2002 and are available on Bronwyn Hall's web site <http://emlab.berkeley.edu/users/bhhall/bhdata.html>




Figure 3: Grant Dates of SSO Patents

we would like to study the long run impacts of SSO affiliation, the data are not sufficient to consider what happens to SSO patents after about 15 years. This issue becomes even more severe when we focus on comparing the pre- and post-disclosure periods—which in many cases may only last one or two years.

Throughout the analysis, we will be comparing the SSO patents to a control sample. The baseline control sample was chosen by selecting all of the patents with the same application year and primary 3-digit technology classification as any patent disclosed to one of the four SSOs.

Before turning to the analysis, it is important to consider how we should interpret the citation patterns revealed below. There are a number of papers that suggest that forward-citations (i.e. the citations received by a particular patent after it has been granted) are an indicator of economic and technological significance. For example, Hall, Jaffe and Trajtenberg (2005) show that citation weighted patent counts are more correlated with a firm’s market value than un-weighted patent counts. Other papers, such as Jaffe, Henderson and Trajtenberg (1993) interpret these citations as an indicator of knowledge flows from the cited to the citing entity.⁶ For this paper, it is not important to defend any particular interpretation of the meaning of a patent citation. As long as forward citation counts contain some information about the

⁶This interpretation raises the question of how to treat self-citations (i.e. citations to a patent owned by the same entity as the citing patent). We found that there was little difference in the results presented below when self-cites were excluded from the sample. For now, we have removed them just to be conservative.

technological or commercial significance of the cited invention, we can use them to learn about the impact of SSOs on the innovation process.

4 Citation Patterns

In this section, we examine the distribution of forward-citations to patents in the SSO and control samples. We are primarily interested in the age profile of citations – the average citation rate conditional on patent age. Hall, Jaffe, and Trajtenberg (2001) refer to this statistic as the lag distribution. (Caballero, Hall, Jaffe, and Trajtenberg—hereafter CHJT—have written a series of papers that examine the shape of this distribution for various sets of patents. See Hall and Trajtenberg, 2002) We begin with a direct comparison of the average citation rates for SSO and control patents. We then turn to an econometric model with application-year and citing-year fixed effects to account for time trends in citing propensity and differences in the “fertility” of inventions across vintage years.

Figure 4 provides a direct comparison of the age profile of citations that illustrate the two basic results that emerge from this section. First, SSO patents are cited more frequently than controls. This holds true regardless of which organization we look at. The average number of citations collected by SSO patents ranges from 26 at the ITU to 37 at the IETF. This compares with the control groups which all average between 7 and 9 citations. Second, the age profile of the SSO patent citations is different from that of the controls. In particular, SSO’s patents exhibit a later peak in citations and a longer lived citation life. This is true overall and for each SSO individually.

One of the most striking facts about Figure 4 is the particularly long citation life exhibited at the IETF. A quick search reveals that there are several notable patents appearing in the tail of the age profile for the IETF, as well as the IEEE. These patents include numbers 4,405,829 and 4,200,770, which cover the basics of public-key cryptography, as well as 4,063,220 which describes the Ethernet networking.⁷ These are exceptional patents in many respects—including the fact that they are disclosed on separate occasions in more than one SSO. So, while these patents are excellent examples of the potential impact of an SSO on the innovation process, it is hard to believe that the *average* patent from among the 400 disclosed at these SSOs in 2003 will turn out to have a similar citation trajectory. Nevertheless, it is interesting to consider whether the importance of the inventions embodied in these early patents could have enhanced the future influence of their respective SSOs.

⁷The inventors on the first patent are Ronald Rivest, Adi Shamir and Leonard Adelman, whose initials are the basis for RSA cryptography. The second patent’s inventors include Martin Hellman and Bailey Diffie. The third patent’s inventor was Robert Metcalfe, who created Ethernet while working at Xerox PARC




Figure 4: Citation Flows for Individual SSO's

While these figures turn out to be consistent with the results found below, it is important to be careful about drawing any conclusions based on the unconditional age profile. The data contains far more information about the first few years of the age profile than the later ones. The sharp increase in citation rates for later years may be caused by a combination of sample truncation and a small number of extremely significant SSO patents. Moreover, these figures do not make any adjustment for differences in the application-year or citing-year. In order to deal with some of the problems inherent in the simple comparisons of 4, we turn to an econometric model of the citation process. This approach corrects for a number of confounding factors, such as the increase in average citation rates over time or differences in the technological significance or “fertility” associated with different application-year cohorts.

Separately identifying the age profile of citations from application-year and citing-year effects can be problematic. Following CJHT, it is important to control for the application year of the citing and cited patent as the application year captures economically relevant time features. Further, CJHT define age as the difference between application year and the citing year. Therefore, the three variables are perfectly collinear and their separate effects cannot be identified without parametric restrictions. In order to proceed, CJHT make non-linear functional form restrictions on how the application year, the citing year and age affect

citations. Identifying age effects separately from cohort and time effects is a common empirical problem in economics. Hall, Mairesse and Turner (2005) review a number contexts in which it arises and review standard solutions. These solutions consist of different ways to restrict functional forms.

As an alternative, we follow an approach proposed in Mehta, Rysman and Simcoe (2005). Mehta, Rysman and Simcoe (2005) points out that the process of diffusion and obsolescence that is captured by the age process described in CJHT can be reasonably thought to start with the publication of the patent instead of the application of the patent. Under this view, we define age as the difference between the citing year (as measured by the application year of the citing patent) and the grant year of the cited patent. This approach exploits the time to grant as random exogenous noise to separately identify age effects from cohort and time effects. Mehta, Rysman and Simcoe consider the efficacy of this approach. They find support for it in the data and find that it results in similar conclusions to CJHT.⁸ Intuitively, we compare across patents with the same application year to see how many citations they receive in a given citing year. Differences in the time to grant among the cited patents allows us to identify the effect of age. We adopt the approach here because it allows for non-parametric identification of the levels of citations of both SSO and control patents.

It is natural under this definition for patents to receive citations at negative ages. That occurs whenever a patent is cited by a patent that applies before the cited patent is granted. For our assumption that age begins at grant date to be exactly correct, it must be that these citations are added by the patent examiner or turned up in a patent search as opposed to indicating an actual intellectual debt. Mehta, Rysman and Simcoe (2005) discuss this at length. In practice, we drop citations from ages below -2 from our data set.

We consider the following specification, where C_{it} is the number of citations received by patent i in year t , α_y are fixed effects for application year y , α_t are fixed effects for citing year t (as measured by the application year of the citing patent), α_a^{CTRL} and α_a^{SSO} are the age effects for the control patents and SSO patents at age a , and ε_{it} is a patent-year error term that is uncorrelated with the fixed effects.

$$C_{it} = f(\alpha_y, \alpha_t, \alpha_a^{CTRL}, \alpha_a^{SSO}, \varepsilon_{it}) \quad (1)$$

In practice, the function $f()$ is typically a poisson function, which generates count variables such as what we observe in our data. This specification is based on the assumption that the

⁸In particular, Mehta, Rysman and Simcoe (2005) show that patents with longer time-to-grant reach their peak in citations later (relative to application date). In addition, the distribution of citations around the grant date appears very similar for patents with differences in time-to-grant.

application-year and citing-year effects are identical for the SSO and control sample but that the age profiles can be different. This is a natural assumption as application-year and citing-year effects are meant to capture “macro” effects, such as changes in policy at the USPTO. While both the control sample and the SSO sample contribute to identifying the application-year and citing-year effects, the number of observations in the control sample dwarfs the number in the SSO sample. Conceptually, we are using the control sample to pin down application-year and citing-year effects and we are estimating age effects for each sample separately.

Note that if we defined age as $a = t - y$, no function of age would be separately identified from y and t dummies regardless of how few parameters we use. But this definition of age is the standard approach. We have also run specifications where we defined age as $a = t - y$ and dropped α_a^{CTRL} from the specification. In this sense, α_y and α_t can be thought of as controlling for application-year, citing-year *and* age effects of the control sample and α_a^{SSO} identifies the difference in age effects between the SSO and control samples. To the extent that results are comparable, we find similar results to those below. But we are interested in comparing the SSO and control sample, not just in identifying their difference. Defining $a = t - g$ where g is the grant year allows us to do so (in addition to being sensible).

We estimate Equation 1 as a poisson regression for each SSO separately.⁹ We leave a full set of regression results unreported. The most obvious result from these regressions is that the SSO age dummies are larger than the control dummies, generally implying that SSO patents receive around 3 times as many citations. This result is not surprising based on Figure 4, which reflects the fact that most patents receive very few citations. While the difference is striking, other hypothesis are difficult to evaluate directly from parameter estimates. Instead, we focus some summary statistics.

More interesting than the level of citation differences is that SSO patents receive citations over a longer time period than control patents. It is straightforward to use the regression results to generate a predicted number of citations for each age (we set the dummy variables for application year 1999 and citation year 1999 on and leave the other application and citation years off). We use these results to compute the probability distribution function of citations over age for the control and SSO patents. Then, we can compute the average age of a citation to each group of patents.

Table 3 presents these results for each SSO, both in the raw data and based on our estimation results. The estimation procedure corrects for the truncation problem inherent in observing many patents near the end of the sample period and so the average age is naturally higher in

⁹One patent disclosed to the IETF has an application year of 1977 while all the rest are applied for in 1985 or later. We drop the 1977 patent in the following analysis.

the regression results than in the raw data. The important point to see is that the average age of a citation is greater for SSO patents than control patents. The difference is positive and significant for each SSO in both the raw and regression results. Note that we have measured age as age since grant.

	Raw		From Regression Results		
	Control	SSO	Control	SSO	Difference
IETF	1.348 (0.006)	3.299 (0.064)	3.752 (0.107)	5.361 (0.126)	1.609 (0.166)
IEEE	2.021 (0.005)	5.066 (0.067)	3.498 (0.082)	4.617 (0.101)	1.119 (0.130)
ITU	2.556 (0.005)	4.250 (0.054)	3.420 (0.076)	4.133 (0.086)	0.713 (0.115)
ETSI	2.058 (0.004)	3.263 (0.023)	4.106 (0.061)	5.263 (0.063)	1.157 (0.088)

Table 3: Average age since grant of citations

Figure 5 graphs the distributions as computed from the regression results. The distributions are for ages -2 to 12. In each case, we can see that the SSO distribution is lower at low ages and higher at high ages. The IETF exhibits the most remarkably long-lived citation profile.

CJHT draw similar graphs for a number of groups of patents and always find peaks in the 4th or 5th year after the application year. That is consistent with our control groups, which show peaks 1 to 2 years after the grant year. This result contrasts with the SSO patents, which exhibit later peaks. For instance, the peak in the ETSI distribution occurs 7 years after the grant date. We know of no other group of patents for which the peak in citation rates are so late.

One concern may be that the high average age of a citation in the SSO sample stems from the fact that SSO patents represent important technologies. That is, it might be true that all highly cited patents exhibit age profiles like we see for SSO patents. In fact, the opposite is true. We show that highly cited patents have age profiles shifted towards lower ages relative to the average patent, not towards higher ages.

To show this, we break up the control patents into average patents and highly cited patents, where highly cited patents collect greater than some cut-off number of citations over their

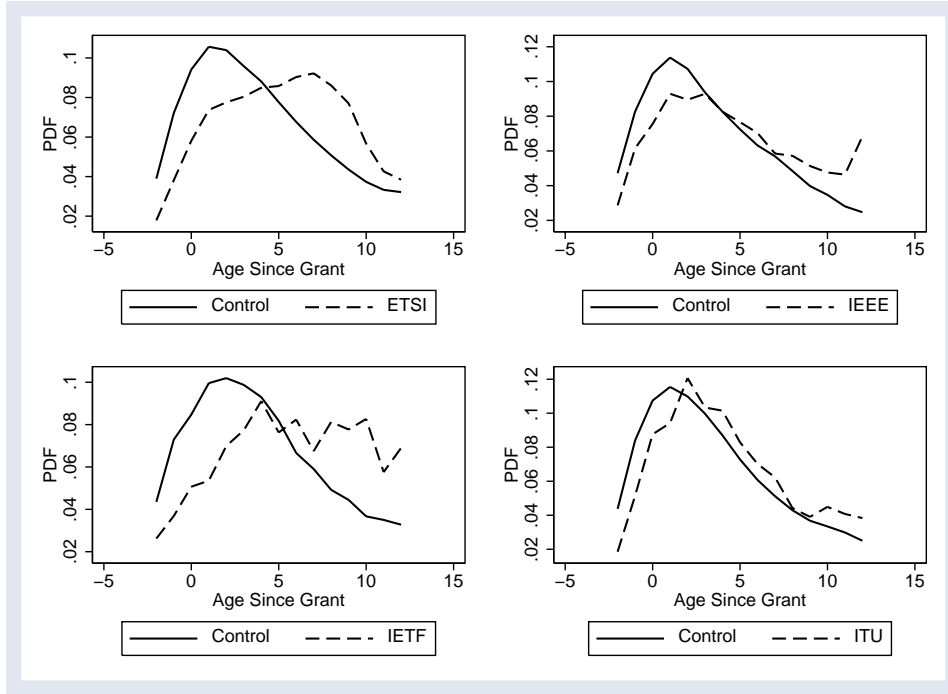


Figure 5: Age profile of citations based on regression results

lifetime.¹⁰ We re-run the poisson regression above but now we allow for separate age dummies for average and highly cited patents. In Table 4, we report the resulting average age of citations for the highly cited and SSO patents. Again, SSO patents have later average ages than control patents, both in the raw data and in the predictions from regression results.

Comparing across Tables 3 and 4, we see that the regression results predict that highly cited patents have a *lower* average citation age than the full set of control patents. This result is true for the set of controls associated with each SSO. We believe this is because the plurality of patents get no citations, which implies a flat age profile. It is the patents that actually get citations that generate the hump-shaped age profile. Removing the patents that get no citations from the sample exaggerates the hump-shaped profile. In the raw results, the highly cited control patents have a higher average age than the full set of control patents. This differs from the regression results because picking patents with a high number of absolute cites selects for patents with early application dates that get cites over a long period of time. The regression procedure corrects for this truncation issue. We have also performed these experiments defining “highly cited” patents to be those with “total citations” in the highest decile of citations for each

¹⁰The cut-offs were chosen so that the highly cited patents had slightly higher average citation rates than the SSO patents. The cut-offs were 20 for the IETF, 11 for the IEEE and the ITU and 14 for ETSI.

application year (which implies less truncation effects) and found very similar results. Note that the average age predicted by the regressions for the SSO patents changes somewhat across Tables 3 and 4 because the application and citing-year effects change. The result for SSO's in the raw data naturally does not change.

	Raw		From Regression Results		
	Control	SSO	Control	SSO	Difference
IETF	1.911 (0.009)	3.299 (0.064)	2.842 (0.222)	4.418 (0.124)	1.577 (0.255)
IEEE	2.367 (0.006)	5.066 (0.067)	3.300 (0.178)	4.527 (0.099)	1.226 (0.204)
ITU	2.952 (0.006)	4.250 (0.054)	3.208 (0.165)	4.031 (0.083)	0.823 (0.185)
ETSI	2.596 (0.005)	3.263 (0.023)	3.221 (0.143)	4.900 (0.061)	1.679 (0.155)

Table 4: Average age since grant of citations for highly cited control patents

5 Selection versus Causation

The previous section showed that patents disclosed to SSOs are cited at higher rates than the average patent. We interpret this as evidence that these patents embody significant inventions. However, this evidence is insufficient to distinguish whether SSOs select technologies that would have been important regardless or whether SSOs actually influence on the importance of these technologies. In this section, we use the timing of intellectual property disclosures to distinguish these affects. Our goal is to use the disclosure event to estimate the marginal impact of the standard setting organization on patent citation rates.

To be clear, this interpretation depends on the date of disclosure being an exogenous event. This condition is unlikely to be met in practice and the sign of the associated bias is difficult to predict. Suppose the selection effect dominates and patent holders tend to disclose important patents to SSO's, but they do not realize the importance of patents for some number of years. Then, they may choose to disclose patents at the time they can predict that citations will increase. In that case, we will observe an increase in citations around the date of disclosure but presumably, the patent would have experienced the increase without disclosure and the correlation between citations and disclosure would over-estimate the marginal impact. Con-

versely, suppose there is a large causal effect of disclosure but market participants can predict which patents will be disclosed some period in advance. In that case, patents may begin to receive citations before disclosure, which would cause the correlation between disclosure and citations to understate the impact of SSO disclosure on citations. With these concerns in mind, we interpret the correlation between disclosure and citations with caution. But as we lack truly exogenous events pushing patents into SSO negotiations, this approach seems to be the appropriate starting place for distinguishing between the selection and causation effects.

We estimate the impact of disclosure in two ways. In the first approach, we are interested in comparing the size of the SSO effect to the size of the disclosure effect. That is, we want to measure the extra effect on an SSO patent of being disclosed. We use the following regression framework:

$$C_{it} = f(\alpha_y, \alpha_t, \alpha_a, \alpha^{SSO}, \alpha_{it}^{Disc} \varepsilon_{it}) \quad (2)$$

In this equation, α_a represents age dummies, α^{SSO} represents a dummy for an SSO patent and α_{it}^{Disc} is a dummy for patent i having been disclosed by period t . In this regression, α_{it}^{Disc} is estimated entirely off of within-SSO variation. For instance, if all SSO patents were disclosed in the same year or at the same age, α_{it}^{Disc} would not be identified. Note that we assume the age process is the same for SSO and control patents and capture the SSO effect in a single dummy. Doing so makes the size of the SSO effect easily comparable to the size of the disclosure effect. Results for the disclosure dummy are very similar if we allow for a set of SSO age dummies.

Throughout this paper, we have used a control group of patents to identify application and citing-year effects under the assumption that are supposed to capture “macro” issues that affect SSO and control patents equally. Doing so is not necessary for results in which we do not want to compare to a control group. In our second approach, we estimate Equation 2 using SSO patents only, dropping the parameter α^{SSO} which is no longer necessary. Therefore, application-year effects, citation-year effects, the age profile and of course the disclosure effect are identified purely from SSO data.

One issue we face is that disclosure is a relatively recent phenomenon. Figure 1 shows that there were almost no disclosures before 1988. Patents granted well before this time presumably faced a different process than ones that came later. To fix ideas, compare the cohort of patents applied for in 2000 to that applied for in 1990 and consider which members of these two groups are disclosed to the IETF, which receives its first disclosure in 1995. Suppose 10 patents from the 2000 group are disclosed. Even if 10 patents from 1990 have the potential to be disclosed to the IETF, most of them are discarded by the time disclosure becomes a normal action. We can imagine that only 1 or 2 from the 1990 group would be disclosed, and those would be the

ones that were very established and received many citations before disclosure. Because SSO patents that were granted before disclosure was an option tend to collect many citations before disclosure, these patents lead to a low estimate of the disclosure parameter.

In order to eliminate this problem, we drop patents with application years before disclosure to a given SSO was a reasonable event. The optimal cut-off year would be the one in which norms guiding disclosure stop changing, although arguably these norms are still changing today. Choosing earlier dates provides a conservative estimate of the disclosure effect. For the IEEE and the ITU we keep patents with application years in 1989 or later. For the IETF, we use 1994 and for ETSI, we use 1995. For the IETF and ETSI, these are the first years in which we observe disclosure. In general, making these cut-offs earlier leads to lower (and possibly insignificant) estimates of the disclosure parameter and making them later leads to a higher parameter. The counter-example is ETSI. In fact, only 3 patents were disclosed in 1995 and after that, no patents were disclosed until 2000. We interpret the ETSI results with caution given the short post-disclosure data that we have.

Table 5 presents results. We leave application-year, citing-year and age parameters unreported. The first column of each panel presents the results using the control sample. Not surprisingly, the SSO dummy is positive and precisely estimated for all four SSO's. In addition, the disclosure dummy is positive and significant as well for all but ETSI. For the IETF, IEEE and ITU, the disclosure parameter ranges between 25% and 35% of the SSO dummy. If we interpret the dummy on disclosure as representing the marginal impact of the SSO on the citation count, we can say that between 20% and 26% of the high citations counts for SSO patents are due to being disclosed to an SSO, and the rest is a selection effect. This result strikes us as very reasonable, although we do not have strong priors over this statistic.

	IETF		IEEE		ITU		ETSI	
	SSO only		SSO only		SSO only		SSO only	
SSO	0.937		0.841		0.843		0.685	
	(0.028)		(0.017)		(0.026)		(0.015)	
Disclosure	0.284	0.475	0.215	0.150	0.297	0.205	0.082	0.291
	(0.062)	(0.075)	(0.035)	(0.043)	(0.035)	(0.043)	(0.046)	(0.061)
obs.	142175	499	400858	1915	415165	1455	301050	2934
cut-off year	1994		1989		1989		1995	

Notes: A poisson regression of citations in a year on application year and citation year dummies, and a single set of age dummies, as well as dummies for SSO patents and for periods in or after disclosure to the SSO. Column II of each panel uses only SSO patents. Regressions use only patents with application years in or after the cut-off year, chosen for when disclosures appear at that SSO.

Table 5: Average age of citations for highly cited control patents

The second column of each panel uses only SSO data to estimate the disclosure effect. In this case, the parameter is significant for all four SSO's, although note that the parameter for

ETSI would be insignificant if we excluded the 3 patents disclosed in 1995. The magnitudes are similar to the case with the control sample, except for the IETF which is about twice the size. Differences are due to the different application, citing and age effects, but the similarity in coefficients suggests that the application and citing-year processes appear to be similar in the SSO and control samples. Visual inspection of the parameter estimates (not reported) also suggest they are similar.

These results suggest an economically significant disclosure effect. In addition, we find similar results if we define disclosure to occur one year after the reported disclosure year in order to account for some sort of lag. However, it would be interesting to do more complex analysis. For instance, we might be interested in how citations vary in the years just before and after disclosure, or how age profiles change when we control for disclosure. We are limited in our ability to answer these questions because of the scarcity of data when using only within-SSO variation. For these purposes, we are currently exploring the use a control sample comparison. While the process of matching patents to a control sample brings up well-known problems with unobserved heterogeneity, the larger sample size allows us to pose new and interesting questions.

An alternative approach would be to consider using patent fixed effects. That is, instead of making within-SSO comparisons, we could make comparison purely within patents. As pointed out in Hall, Mairesse and Turner (2005), patent fixed effects (which capture application year) would be collinear with the combination of age and citing year effects. In this case, the technique of Mehta, Rysman and Simcoe (2005) would not be sufficient to obtain separate identification because the patent fixed effects would control for the time-to-grant as well as the application year. One could estimate just patent fixed effects and age effects, but keeping in mind that these variables are collinear with citing-year effects means the “age effects” would be difficult to interpret as such. Therefore, using patent fixed effects would not be desirable in for identifying age profiles as above. However, one could estimate in this fashion and still estimate a disclosure effect, which would be identified and interpreted as we have done so far. Such an approach is the subject of current research.

6 Conclusions

While the importance of SSOs has been widely remarked by academics and practitioners, there have been few attempts to systematically measure their role in economic performance or technological change. Moreover, since much of the evidence for SSOs’ importance is based on specific examples of technologies they have endorsed, there continues to be some debate over whether they actually influence the process of cumulative development or merely choose

to select and evaluate important technologies. This is the first paper to address these issues using patent citations as a measure of SSO performance.

Using data from the patents disclosed in the standard setting process at ETSI, IEEE, IETF, and the ITU, we showed show that the SSO patents collect many more citations, typically around 3 times as many. Furthermore, they have a different age profile of citations, receiving them over a longer time span. Finally, we exploited the timing of SSO patent disclosures to show that there is a correlation between citations and the act of disclosure, representing more than 20% of the total difference between SSO and non-SSO patents.

Subject to concerns about the exogeneity of the disclosure event, the large selection effect suggests that SSOs are successful in identifying important technologies. Furthermore, the significant causal effect suggests that current SSO decisions impact the path of future technological innovation. This may occur because SSOs embed technologies in standards that are difficult to switch away from because of network effects. Alternatively, it may be simply the attention attracted to a technology by a disclosure, possibly due to the disclosure’s indication that a patent-holder is willing to license its technology.

The treatment of intellectual property at SSOs is a subject of interest to many in the technology policy-making community, and a number of recent events have increased the prominence of this issue. These events include the Rambus case, the surge in intellectual property disclosures at SSOs, the W3C’s decision to adopt a royalty-free licensing policy, and the Standards Development Organization Act of 2004—which extended certain antitrust protections originally contained in the National Cooperative R&D Act to American SSOs. While this paper emphasizes the positive question of SSOs’ role in technological change, the finding that these institutions not only select important technologies but also may influence their future significance suggests that the policy interest in these issues are justified. While we hope to address a number of these questions in future research, we should acknowledge that it is hard to draw any clear welfare implications from the current results. In particular, the impact of having patents in an industry standard will depend on the rules of the SSO, participants’ willingness to license any essential intellectual property, and whether they do so on “reasonable and non-discriminatory” terms.

Nevertheless, this paper provides some of the first large-sample statistical evidence related to the patents disclosed in the standard setting process, and should be an important starting point for future research. We will conclude by suggesting a number of possible extensions to this research. First, while several studies have tied patent citations to various measures of economic or technological significance, it would be valuable to examine whether patents disclosed to SSOs are also linked to variables such as a firm’s market value. This work would involve collecting additional data on SSO patents that provides some indication of whether a

disclosed patent was part of an important standard and/or was determined to be essential to the implementation of that specification. This data would also provide a means for identifying the underlying causes behind the SSO effects. In particular, we might be concerned that the marginal impact of an SSO on future citation flows is actually just a “publicity effect,” caused by the announcement that the patent is tied to a particular technology. If essential patents connected to important standards receive a larger post-disclosure boost, we should conclude that the citation increase actually reflects the network effects associated with standardization. Another interesting direction is to explore the relationship between patents disclosed to SSOs and other groups of “significant” patents, such as the general purpose technologies identified by Hall and Trajtenberg (2004).

References

- Blind, K. (2004) *Economics of Standards* Elgar, Northampton.
- Bolin, S. ed. (2004) *The Standards Edge*, Sheridan, Ann Arbor, MI.
- Cargill, Carl F. (1997) *Open Systems Standardization : A Business Approach*, Prentice Hall PTR, Upper Saddle River, NJ,
- Chiao, B. and Lerner, J. and Tirole, J. (2005) “The Rules of Standard-Setting Organizations: An Empirical Analysis,” mimeo, Harvard Business School.
- Farrell, J. (1996) “Choosing the Rules for Formal Standardization,” mimeo, University of California at Berkeley.
- Furman, J. and Stern, S. (2004) “Climbing Atop the Shoulders of Giants: The Impact of Institutions on Cumulative Research,” mimeo, Northwestern University.
- Gandal, N. and Gantman, N. and Genesove, D. (2005) “Intellectual Property and Standardization Committee Participation in the U.S. Modem Industry” in *Standards and Public Policy*, Greenstein, S. and Stango, V., eds., Cambridge University Press, forthcoming.
- Gandal, N. and Kende, M. and Rob, R. (2000) “The Dynamics of Technological Adoption in Hardware/Software Systems: The Case of Compact Disc Players” *RAND Journal of Economics* 31, 43-61.
- Greenstein, S. and Stango, V. eds. (2005) *Standards and Public Policy*, Cambridge Press, Cambridge, forthcoming.
- Hall, B. and Jaffe, A. and Trajtenberg, M. (2001) “The NBER Patent Citations Data File: Lessons, Insights, and Methodological Tools” in *Patents, Citations and Innovations: A Window on the Knowledge Economy*, MIT Press, Cambridge, 2002.
- Hall, B. and Jaffe, A. and Trajtenberg, M. (2005) “Market Value and Patent Citations” *RAND Journal of Economics* 36(1) 26-38.

- Hall, B. and Mairesse, J. and Turner, L. (2005) “Identifying Age, Cohort and Period Effects in Scientific Research Productivity” mimeo, University of California at Berkeley.
- Hall, B. and Trajtenberg, M. (2004) “Uncovering GPTs using Patent Data” mimeo, University of California at Berkeley.
- Hausman, J. and Hall, B. and Griliches, Z. (1984) “Econometric Models for Count Data with an Application to the Patents-R&D Relationship” *Econometrica* 52(4) 909-938.
- Henderson, R. and Jaffe, A. and Trajtenberg, M. (1993) “Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations” *Quarterly Journal of Economics* 108, 577-598.
- Jaffe, A. and Trajtenberg, M. (2002) *Patents, Citations and Innovations: A Window on the Knowledge Economy* MIT Press, Cambridge.
- Lemley, M. (2002) “Intellectual Property Rights and Standard Setting Organizations” *California Law Review* 90, 1889-1981.
- Mehta, A. and Rysman, M. and Simcoe, T. (2005) “Identifying Age Profiles of Patent Citations” mimeo, Boston University.
- Mowery, D. C. and Simcoe, T. (2002) “Is the Internet a US invention? - an economic and technological history of computer networking” *Research Policy* 31(8-9) 1369-1387.
- Toivanen, O. (2005) “Choosing Standards”, mimeo, University of Helsinki.