

DISCUSSION PAPER SERIES

No. 9849

SSOS: CURRENT POLICY ISSUES AND EMPIRICAL EVIDENCE

Neil Gandal and Pierre Régibeau

INDUSTRIAL ORGANIZATION



Centre for Economic Policy Research

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP9849.php

SSOS: CURRENT POLICY ISSUES AND EMPIRICAL EVIDENCE

Neil Gandal, Tel Aviv University and CEPR
Pierre Régibeau, Charles River Associates (CRA) and Imperial College

Discussion Paper No. 9849
March 2014

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INDUSTRIAL ORGANIZATION**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Neil Gandal and Pierre Régibeau

CEPR Discussion Paper No. 9849

March 2014

ABSTRACT

SSOs: Current Policy Issues and Empirical Evidence

Despite the increasing importance of and huge growth in SSOs and the flurry of research on the economics of SSOs in the last decade, no comprehensive review of the literature exists. Such an undertaking is beyond the scope of this paper. Here, we address a more modest goal: a selective review focusing on current policy issues regarding SSOs and empirical evidence on the performance of SSOs. A key part of the analysis here focuses on what light the empirical work on SSOs can shed on SSO performance.

JEL Classification: L4

Keywords: SSO performance and standard setting organizations (SSOs)

Neil Gandal
Berglas School of Economics
Tel Aviv University
Tel Aviv 46766
ISRAEL

Pierre Régibeau
Charles River Associates
99 Bishopsgate
London EC2M 3XD
UK

Email: gandal@post.tau.ac.il

Email: pregibeau@crai.com

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=144974

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=116416

Submitted 23 January 2014

SSOs: Current Policy Issues and Empirical Evidence

Neil Gandal¹ and Pierre Régibeau²

December 2013

Abstract

Despite the increasing importance of and huge growth in SSOs and the flurry of research on the economics of SSOs in the last decade, no comprehensive review of the literature exists. Such an undertaking is beyond the scope of this paper. Here, we address a more modest goal: a selective review focusing on current policy issues regarding SSOs and empirical evidence on the performance of SSOs. A key part of the analysis here focuses on what light the empirical work on SSOs can shed on SSO performance.

¹ Berglas School of Economics, Tel Aviv University and CEPR, gandal@post.tau.ac.il

² Charles River Associates and Imperial College, PRegibeau@crai.com

1. INTRODUCTION

Given the dramatic growth of the Internet and information technology industries in general, and the importance of interconnection in these networks, the economics of compatibility and standardization has become mainstream economics. Most Internet and information technology products and services exhibit network effects. In such industries questions of compatibility and standardization are important. Because of the network effects that are inherent in such industries, successful diffusion of these products is often contingent on the emergence of a single standard. For example, Postrel (1990) blames (in part) the failure of quadraphonic sound in the 1970s to competing standards. Standardization can occur through two main mechanisms: the de facto market dominance of a specific technology or the explicit coordination of product designs around commonly agreed technological dimensions.

This second mechanism is by far the most prevalent and operates through standard-setting organizations. One website lists more than 1000 SSOs. Fully 260 SSOs are classified as dealing with “interoperability” standards. Additionally, a large proportion of SSOs classified as dealing with the internet, “wireless and mobile”, “software” and “Multimedia” would likely address compatibility issues. Hence, it seems therefore reasonable to assume that compatibility standards are the main activity of the large majority of SSOs on the list. The importance of SSOs is also illustrated by an estimate that IBM spent \$500 million on standards related activities in 2005 and Hewlett Packard and Sun Microsystems each belonged to more than 150 SSOs in 2003.

Despite the increasing importance of and huge growth in SSOs and the flurry of research on the economics of SSOs in the last decade, no comprehensive review of the literature exists. Such an undertaking is beyond the scope of this paper. Here, we address a more modest goal: a selective review focusing on current policy issues regarding SSOs and empirical evidence on the performance of SSOs. A key part of the analysis here focuses on what light the empirical work on SSOs can shed on SSO performance. We should also point out that we focus our attention on “traditional” SSOs where participants retain the right to charge royalties for licensing their patents reading on the agreed upon standard. We do not therefore consider “open” standard organisations.

We first provide a short review of the early empirical work on network effects (section 2.1) and empirical work quantifying network benefits (section 2.2). In section 3, we provide a brief primer on SSOs.

In section 4, we examine policy concerns. The policy concerns fall into two broad categories: (i) concerns about the operating performance of SSOs – i.e. about how good SSOs are at “doing their job” and (ii) concerns regarding the impact of SSOs in the markets for products, technologies and standards. We first propose several criteria that might be used to distinguish more successful SSOs from less impressive ones. We then examine (in section 4.1) what light the existing empirical literature on SSOs can shed on how SSOs are meeting or not meeting these criteria. In section 4.2, we examine where the existing empirical work has to say about key competition policy concerns.

If standardization is desirable and SSOs often are the only realistic route to a standard, then whether SSOs perform well or not may be a secondary consideration. In this view, what matters is how differences in the internal rules of SSOs and in public policy towards SSOs affect the performance of these organisations. We address this issue in section 5. In section 6, we provide a very brief conclusion and suggest topics for future empirical research.

2. BENEFITS OF STANDARDISATION: EMPIRICAL EVIDENCE

SSOs are associations that include potential rivals and make decisions that affect the conditions of competition in the relevant downstream industries. As such, as we will discuss further in section 4, SSOs are immediately suspect of possibly facilitating anti-competitive behaviour. The reason why SSO activities are nevertheless encouraged is that the anti-competitive risks – minimised by an adequate competition policy – seem worth taking in light of the potential benefits from standardisation. It does therefore make sense to begin with a review of what we know about the size of such benefits.

2.1. REVIEW OF EARLY EMPIRICAL WORK ON NETWORK EFFECTS

A growing empirical literature has developed to examine technological adoption of products with network effects. The primary issue addressed by the early literature is whether network effects exist; this work typically employed reduced form models. Greenstein (1993), Gandal (1994, 1995), and Saloner and Shepard (1995) provide early evidence that the value of the “hardware” good depends on the variety of compatible complementary software.

Software for the IBM 1400 mainframe could not run on succeeding generations of IBM mainframes while software for the IBM 360 could run on succeeding models. Greenstein (1993) finds that a firm with an IBM 1400 was no more likely than any other firm to purchase an IBM mainframe when making a future purchase. On the other hand, a firm with an IBM 360 was more likely to purchase an IBM mainframe than a firm that did not own an IBM 360. This result can be interpreted as a demand for (backward) compatible software.

Saloner and Shepard (1995) test for network effects in the ATM industry by examining whether banks with a larger expected number of ATM locations will adopt the ATM technology sooner. Since expected network size is not an observable variable, they use the number of branches as a proxy. Their results suggest that banks with more branches will adopt earlier, which is consistent with virtual network effects.

Gandal (1994) estimates hedonic (quality-adjusted) price equations for spreadsheets to examine whether spreadsheet programs that were compatible with Lotus -- the de facto standard -- command a premium. The results, that consumers place a positive value on compatibility, suggest (I) direct network effects because people want to share files and (II) indirect network effects because compatible software enables the of transfer data among a variety of software programs.³

Empirical evidence for the existence of network effects can also be obtained indirectly, for example Liu, Kemerer, and Slaughter, (2007) argue that, if network effects are indeed important, then the existence of a technology that helps mitigate those effects should have a noticeable impact on observed behaviour. In particular, the existence of a good-quality conversion technology should favour the emergence of a number of competing standards and decrease the price premium enjoyed by the sponsors of the leading technological standard. The authors test this hypothesis on data from the flash memory card market. This market is characterized by the existence of a number of different standards that were made significantly more compatible by the emergence of a digital conversion technology. The authors find that the emergence of this good quality “adapter” did indeed lead to a substantial decrease in the premium earned by the leading formats, implying that, in the absence of such convertibility, network effects must have been significant.

2.2. QUANTIFYING NETWORK BENEFITS

The early papers discussed above were important for empirically establishing the presence of network effects across many industries. Recent papers in the literature have employed structural estimation which enables researchers to better quantify network benefits and conduct counterfactuals.

Gandal, Kende, & Rob (GKR, 2000) develop a dynamic structural model of consumer adoption and software entry and use the model to estimate the feedback from hardware to software and vice versa in the CD industry. GKR (2000) show that a 5 percent reduction in price would have had the same effect as a 10% increase in CD variety in terms of increasing sales of CD players. This illustrates the importance of network effects (manifested in increases in software variety) relative to changes in the price of the hardware.

³ Gandal (1995) extends the analysis to Database Management Software (DMS) and multiple standards and finds that only the Lotus file compatibility standard is significant in explaining price variations, suggesting that indirect network effects are important in the DMS market.

Two subsequent papers employ the GKR (2000) methodology in the video games industry. Clements and Ohashi (2005) use a logit model to test for indirect network effects in the U.S. video game market. Following the methodology of Gandal, Kende, and Rob (2000), they find that a 1% increase in game titles is equivalent to a 2.3% price cut in the market. Using the same market for a slightly later time period, Prieger and Hu (2006), find that a 1% increase in software variety is equivalent to a 0.4% price reduction.

Comparing these results, GKR (2000) and Prieger and Hu (2006) have similar estimates: a 2.0-2.5 percent increase in software variety is equivalent to a 1% fall in the price of the hardware. Clements and Ohashi find a larger variety effect, namely roughly a 0.5% increase in variety has the same effect as a 1% fall in price. Even if GKR (2000) and Prieger and Hu's (2006) lower estimates are the general rule, the importance of network effects (as manifested in the importance of changes in the software variety relative to changes in the hardware price) is quite strong.

Rysman (2004) developed a structural model to examine the importance of network effects in the market for Yellow pages. The model includes a consumer adoption equation, advertiser demand for space, and a firm's profit maximizing behavior. He finds that consumers value advertising and advertiser's value consumer adoption, suggesting virtual network effects. He measures the importance of the network effect by calculating the surplus lost by the market's failure to internalize the network effect. This calculation is then compared to the (classical) deadweight loss from imperfect competition. He finds that the ratio of network deadweight loss to classical deadweight loss is 1.26, which suggests that the network effect is quite large, even in an industry with significant market

Finally, there are a few results that measure the magnitude of the effect of compatibility (vs. incompatibility) in the speed of technology adoption: GKR (2000) show that if it had been possible to make CD players compatible with LPs, compatibility could have accelerated the adoption process by about 1.5 years, reducing the time of adoption by more than 20 percent.

Dranove and Gandal (2003) empirically test for network effects and preannouncement effects in the DVD market by measuring the effect of potential (incompatible) competition on a network undergoing growth. They find that there are network effects and the DIVX preannouncement reduced DVD sales by approximately 20 percent. In the end, DVD was adopted as the de-facto standard. Nevertheless, the non-trivial effect of the preannouncement by such a weak entrant (see Dranove and Gandal (2003) for details) suggests that network effects were strong in this market as well.

3. A PRIMER ON SSOS

Standardisation can occur through two main mechanisms: the de facto market dominance of a specific technology or the explicit coordination of product designs around commonly agreed technological dimensions. This second mechanism is by far the most prevalent and operates through standard-setting organisations. Not surprisingly the prevalence of SSOs varies both across sectors and over time.

In terms of sectors, we would naturally expect SSOs to arise more often where the benefits of compatibility are higher. This includes so-called “complex” industries, where products need to aggregate a variety of technologies in order to be of any use to consumers. We would also expect efforts to achieve compatibility to be more intense when the usefulness of the products involved relies on some form of communication between users or calls for the use of complementary goods.

A very rough way of getting at the importance of SSOs across technologies and sectors is to simply count the number of existing organisations. Even such a simple count runs into two main difficulties. Firstly, there is no such thing as an official, exhaustive list of SSOs. Secondly, SSOs deal not only with the type of “compatibility” standards that we are interested in, they also create quality/safety standards.

The tables below rely on the list of SSOs provided at <http://www.consortiuminfo.org/links/>. As of October 14, 2013, this list included 885 SSOs. While, we are not able to distinguish between SSOs involved in the setting of “compatibility” standards as opposed to those involved in quality or safety standards, we note that there are 260 SSOs classified as dealing with “interoperability” standards. Moreover, one would think that a large proportion of SSOs dealing with the internet (141), “wireless and mobile” (135), “software” (174) and “Multimedia” (87) would also mostly address compatibility issues; it seems therefore reasonable to assume that compatibility standards are the main activity of the majority of SSOs on the list.⁴

The following tables show the number of SSOs by technical category. In the last column of the tables, we have tried to match the technology classification with the classification of technologies obtained by von Graevenitz et al (2013). These authors use the pattern of patent citations to distinguish between “complex” technologies, where many components controlled by different parties go into designing a successful product and “discrete” technologies where products are more uni-dimensional and the IPRs relating to a given product tends to be held by one or very few firms. As expected, the vast majority of the technologies or industries in our tables are classified as complex, confirming our expectation that compatibility issues are much more likely to arise in such environments. Notice however that the categories marked NA (not

⁴ Table 1 has more than 885 SSOs, because some SSOs appear in more than one category. Table 2 has less than 885 SSOs because not all SSOs can be assigned to specific industries. This is not important, since the goal is simply to give a sense that most of the technologies or industries in the tables are classified as complex and that compatibility issues are much likely to arise in such environments.

available) either have very few SSOs or are categories that one would intuitively think of as “complex” anyway. (Unfortunately, we were not able to relate every one of our technologies or industries to von Graevenitz and et al (2013.)

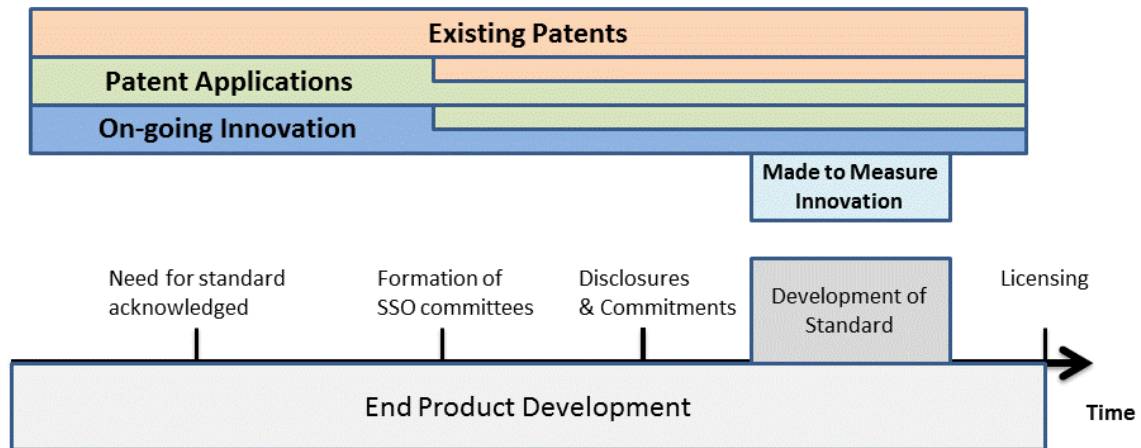
Table I: SSOs per Technical Category

Technical Category	Number of SSOs	Complexity
Audio/Video/Multimedia	87	Complex
Cloud Computing	16	NA
Credit/Debit/Smart Cards	24	NA
Electronic Media	25	Complex
Hardware	97	Complex
Imaging	30	Complex
Internet	141	Complex
Interoperability	260	NA
Languages/Protocols	27	NA
Network and Network Centric Computing	101	Complex
Operating Systems	62	Complex
Security and Cyber-security	91	Complex
Semiconductors	31	Complex
Software	174	NA
Web Services	56	Complex
Wireless and Mobile	135	Complex

Table 2: SSOs per Industry

Industry	Number of SSOs	Level of Complexity
Aeronautics	8	NA
Automotive	23	Complex
Bio IT and Life Sciences	19	Discrete
Clean Tech and Renewable Energy	40	Complex
Construction	14	Complex
Consumer Electronics and Content	27	Complex
Defence	10	NA
Digital and Distance Learning	22	NA
Electronics	100	Complex
Health and Medical	52	Med. Techno: Complex
Manufacturing	35	NA
Power and Smart Grid	30	Complex
Real Estate	5	Discrete
Telecom	89	Complex
Multi-Industries	84	NA

As for any decision-making institution, the performance of SSOs ultimately depends on its membership and on how the various interests of its members get aggregated into a final decision. The following figure (Figure 1) gives us a typical timeline for SSO activities.



STANDARD-SETTING BY SSOs: TIME LINE

The standard-setting process begins when the need for a technological standard is widely recognised in the industry. At that point, technologies that will prove relevant to that standard already exist, be it as granted patents, patent applications or even on-going research programs. The next important stage is the formation of the relevant committees or work groups, i.e. the determination of SSO *participation*. This participation has two crucial dimensions. The first one concerns the *type* of agents allowed to participate. Broadly speaking, SSO members come from three groups: IPR-holders, potential users of the standards and consumers. Clearly, these groups can overlap. In particular, IPR-holders are often also potential users of the standard. For example, Samsung, Motorola Mobile, Nokia and Ericsson hold significant shares of the patents declared “essential” for the 3G standard and are all also significant manufacturers of smartphones. The second important dimension is the distinction between more and less “equal” participants. Standard setting activities are very effort-intensive institutions. Not surprisingly, the more expertise a given member has and the more human capital she is willing to invest into the process, the greater her influence on the design of the standard. Accordingly, large IP-holders and large potential users will generally have significantly more weight within the SSO than

smaller firms or than final consumers. Although SSO membership is often open to all, it is sometimes restricted. On the other hand, membership is always voluntary. As we will see below, this is important for public policy towards SSOs as any tightening of the rules to which they are subject is likely to increase the number of agents – and especially IPR-holders – who choose to remain on the side-lines.

While membership has its privileges, it also comes with some obligations. The two most important obligations are disclosure and pricing/(F)RAND commitments, where (F)RAND = (Fair), Reasonable and Non-Discriminatory terms, to anyone wishing to practice the standard. Disclosure requires members to reveal all existing IPRs, IPR applications and even on-going research that might be relevant to the standard to be developed. Since the precise technological characteristics of the standard are, per force, not known at this stage, this duty of disclosure must be understood quite broadly: any proprietary knowledge that *might* be relevant to the standard should be revealed. Moreover, the duty of disclosure remains throughout the process: a new member or an existing member whose relevant portfolio of IPRs and research projects changes are both required to notify the SSO. In this sense, disclosure is a continuous process,

SSOs can also ask participating IPR-holders to indicate – or even commit to – a royalty ceiling as well as a “worst” set of other licensing conditions. While this falls significantly short of an ex ante commitment to IPR licensing terms, it at least provides some information as to the likely relative “price” of using alternative sources of technology. However, the proportion of SSOs that impose such a requirement appears to be relatively small. Much more common is a demand that participating IPR-holders commit to license any price of intellectual property that actually “reads” on the actual standard on so called (F)RAND terms to anyone wishing to practice the standard.

Once those pre-conditions are set, the actual work of building a mutually acceptable standard begins in earnest. As pointed out above, this process is very intensive in human capital... and very costly. Farrell and Simcoe (2012) cite estimates that a medium to large size firm will spend about \$50,000 a year to just to participate in a single SSO in one year. This includes *salary*, travel, and membership fees, but not any costs for contributing to the technology. In other words, it does not include the salaries of those who develop firm strategy regarding SSOs, those who are involved in technical work on behalf of the standard, or those who work on the ‘legal’ aspects. Chiao, Lerner, and Tirole (2007) cite an estimate that IBM spent \$500 million on standards related activities in 2005. According to Andrew Updegrave (2003,) Consortium Standards Bulletin) Hewlett Packard and Sun Microsystems each belonged to more than 150 SSOs in 2003.

Moreover, in many cases, developing a standard involves much more than sifting through available technologies; it also requires the development of technology from scratch in order to bridge the gaps between what is currently available. One should therefore think of the set of technologies that might be relevant to the standard as being in a perpetual state of flux as the proposed standard itself evolves, patent applications are granted or rejected, on-going research bears or does not bear fruit and standard-specific research becomes necessary.

Most SSOs claim that the standards that actually emerge from their activities are reached “through consensus” between participants. This is of course hard to take as literally true as conflicts are bound to arise between participants. Still, since most SSOs do not have explicit “voting rules”, we actually know very little as to how such conflicts are resolved.

Once the standard - or at least the first generation of the standard – is determined, IP-holders can declare the intellectual property that they own and that they feel to be “essential” to the standard. This process of declaration is often misunderstood. Firstly, while SSO members are obliged to declare as essential all patents or other IPRs for which they will seek royalties from those who practice the standard, IP-holders who chose to sit on the side-lines and not take part in the SSO’s deliberations have no such obligation. Secondly, in the vast majority of cases, there is no third-party verification of the truly “essential” characteristics of the IPRs declared as such. Quite naturally then, since failing to declare a truly essential patent would be a breach of the commitments made to the SSO, there is a natural tendency for participating IPR-holders to over-declare essential IPRs.

The standard-setting process does not actually stop when an-agreed upon standard is obtained. For the standard to be successful it must be made accessible to users on reasonable terms and it must be updated as the needs of the industry change and the set of available technologies widens. Importantly, it is not possible for users to purchase “access to the standard” from a single agent, for two main reasons. Firstly very few SSOs mandate or even recommend that all standard essential IPRs held by members be grouped into a corresponding patent pool, where access to all intellectual assets can be obtained for a single fee (or variable royalty). So potential users of the standards must usually obtain separate licensing agreements from all SSO members holding relevant patents. Secondly, even if all SSO members agreed to include their relevant IPRs into a standard-specific pool, some standard essential IPRs are also bound to be held by agents who did not take part in the standard-setting process, making separate negotiations with these agents unavoidable.

Standards are not static. For example, the 3G standard has many versions. This need to update the standard raises a number of issues. Under what conditions are new members admitted to the table? Is there an efficient mechanism to ensure that existing members cannot push for continued use of their technology even though a superior technology might now be available? How does the advent of a new version of the standard affect the terms of existing licensing contracts? These are all aspects of SSOs about which we do not have any systematic, reliable information.

To get a more complete picture of the standard-setting process, it is important to realise that SSOs do not exist in splendid isolation. In particular, the last ten years have witnessed a sharp increase in standards *consortia*.⁵ A standard consortium is an organisation that is both less formal and less inclusive than traditional SSOs. Although some consortia go on to develop their own standards – in which case we would, for the purpose of this paper, treat them as SSOs – most of them intervene ahead of participation in more formal SSOs and involve coordination of research strategies significantly more “upstream” than SSOs. They can therefore be seen as organisations that help reduce the duplication of R&D efforts and smooth conflict of interests ahead of formal SSO proceedings. Of course such “coordination” is not necessarily innocuous. In this respect, it is interesting to note that Baron and Pohlman (2013) find that consortia tend to

⁵ See Baron and Pohlman (2013)

arise more frequently when the relevant technologies are more fragmented but that they tend to be made of companies with technologies that are close competitors.

4. POLICY CONCERNS

Policy concerns regarding SSOs can be divided into two broad families: concerns about the operating performance of SSOs – i.e. about how good SSOs are at “doing their job” – and concerns regarding the impact of SSOs in the markets for products, technologies and standards. In order to assess the performance of SSOs, we must first define their task, i.e. the criteria that might be used to distinguish successful SSOs from less impressive ones. We propose the following list:

- SSOs should coordinate on the best, i.e., most technologically efficient standard, unless there are compelling cost-reasons (e.g. licensing terms commitments) that justify the use of inferior substitute technologies. (see section 4.11)
- SSOs should develop their standards in a timely manner. (section 4.12)
- The resulting standards should be accessible (and affordable) for efficient competitors in the relevant downstream product markets. SSOs should not distort competition in either the market for technologies or the market for standards. (section 4.13)
- The standard should be periodically updated to include more efficient technologies that have become available and SSOs should not have a chilling effect on innovation. (section 4.14)

The first two criteria pertain to the operating performance of SSOs, while the last two *mostly* relate to competition policy concerns. The links between these criteria and the existing empirical literature are illustrated in Figure 2 in the appendix:

4.1 OPERATING PERFORMANCE OF SSOs

The first two criteria have been analysed theoretically by Farrell and Saloner (1988) who compare the performance of a SSO to market-based or “de facto” standardisation. In the market-based mechanisms, market participants can commit to a given technology in the hope that others will follow. Whether or not effective coordination on a standard is achieved therefore depends on how strong the “bandwagon” effect. Not surprisingly then, markets do not always lead to standardisation as equilibria with incompatible technology choices can emerge. On the other hand, when markets lead to de facto standardisation, they tend to do this rather quickly (precisely because standardisation only arises when bandwagon effects are large). By contrast, standardisation committees work around the principle of “unanimity”, which means that participants cannot make unilateral commitments to technologies and must instead reach agreement with others. Because SSO members each have their own vested interests, agreement requires some form of negotiation. Farrell and Saloner model this negotiation process as a war of attrition. In such a game, actual delays arise in equilibrium as each player tries to convince the others that its vested interests are strong. Still, in the end agreement is reached. Overall then, SSOs are better than the market mechanism at ensuring standardisation but they take more time to get there. Still Farrell and Saloner show that expected welfare is higher under the SSO process...although the best performance is obtained where SSO participants can also “play the market game” by making unilateral commitments in the relevant markets. Intuitively, the threat that the market-based process might converge on a standard that they do not like spurs SSO members to settle their differences more quickly.

While it is difficult, if not impossible to directly examine the prediction of the Farrell-Saloner (1988) analysis empirically, there are a few empirical papers that shed some light on the role played by standard setting organizations.

4.1.1. Choice of technology

We are only aware of one “direct” test of the quality of the technologies that SSOs assemble into a standard. A traditional – if imperfect – manner of measuring the importance of patented technologies is to use the number of backward citations that those patents receive. There is an obvious difficulty however. In order to assess whether the SSOs end up relying on good technologies, one would want a measure of the intrinsic value of these technologies, untainted by the fact that they were selected as part of the standard. To address this issue, Rysman and Simcoe (2008) examines patent citations over time, distinguishing between the citations received before the start of the standard-setting process and those obtained afterwards. They conduct their study for a sample of US patents which were disclosed during the standard-setting process at four SSOs. The citation pattern for these patents is then compared to the pattern for a control group of patents that were not used or disclosed within the SSO. The authors find that the baseline citation rate for SSO patents is about double that of non-SSO patents. Thus there is a selection effect, since SSOs indeed chose technologies that are intrinsically important. They

also find that disclosure itself leads to a 20%-40% increase in citation rates. Hence, SSOs also “help establish” the technologies that they endorse.

While this is the best empirical evidence that we currently have on the equality of SSO’s choices, the approach has two main drawbacks. Firstly, technologies that are disclosed to a SSO are not necessarily actually *used* by a SSO. In this sense, the Rysman and Simcoe result tells us that SSOs attract agents who hold important patents, but it does not necessarily tell us that the SSO does a good job of selecting the best patents from this set. Secondly, there is a difference between the quality of the patents that read on the standard and the quality of the standard itself since that quality also depends on how well the technologies used for different parts of the standard actually fit together.

In the absence of more direct evidence, one can also try to infer the likely quality of the standards chosen by SSOs based on the *composition* of the relevant SSO committees. The general idea is that the better the active participants in the SSOs reflect the composition of the industry (including both users and IPR holders) and the less the internal SSO process is biased in favour of some types of participants, the more likely it is that the standard will be chosen mostly on grounds of technological merit and cost.

The earliest econometric study of standardization committees was conducted by Weiss and Sirbu (1990). In their paper, they considered eleven cases from the computer communications hardware. They focused on settings in which the standard was reached in a standards committee, two technologies were competing for inclusion into the standard and only one of these technologies was eventually chosen by the committee. Their results suggest that the size of the firms in the coalition supporting a technology and the efforts they made through written contributions were significant determinants of which technology was chosen. In addition, the technologies whose sponsors weighed market factors more highly than technical factors were more likely to be adopted in the standards decision studied.

The lessons from the Weiss and Sirbu study are not completely clear. On the one hand, their results indicate that SSOs suffer from the same drawbacks as most political/voting process: because each party that favours a given technology has an incentive to free-ride on the efforts of those who share the same preference, highly concentrated coalitions are more likely to prevail. This would suggest that SSOs do not necessarily choose the best technological standards. On the other hand, the implications of the importance the efforts exerted by the sponsors of the each technology are ambiguous. It might just be another consequence of the “free rider” argument just outlined, but it might also reflect the fact that sponsors of better technologies have greater incentives to push for them, which would imply that SSOs would indeed tend to converge on the correct standard. There are two mechanisms that would imply that greater effort is spent on promoting the better technology. Firstly, a better standard would be more likely to be widely adopted and the products designed around it would be more likely to generate large sales. This translates into both greater benefits for users and higher potential revenues for IPR holders. Secondly, SSO participants with interests on the user side of the market would naturally prefer the better of the two rival technologies. In this sense, the Weiss and Sirbu study is incomplete as it does not systematically analyse *who* was pushing for what. What was the composition of the SSOs in terms of users, strict IPR owners or vertically integrated entities? How did the composition of the winning coalitions compared to the composition of the SSO itself? Which type of actor exerted the most effort? Without answers to

these questions, it seems hard to assess the implications of the Weiss and Sirbu analysis for the accuracy of SSO's technological choices.

Gandal, and Genesove (2006) empirically examine the interaction between intellectual property and participation in standardization committee meetings in the modem industry. They combine patent data on modems (modem in the title) and SSO "meeting" data from the Telecommunications Industry Association (TIA,) the SSO responsible for developing voluntary (consensus) standards in the analog modem market in the U.S. While many firms obtained "modem" patents and many firms participated in standardization meetings, only a small subset of 45 firms both obtained patents and participated in the standardization meetings. These firms accounted for a significant percentage of the patents received and the total number of meeting attendees. Their main finding is that, while participation in standards meetings predicts future intellectual property (both un-weighted and citation weighted patents), the reverse is not true: patents and citations are not good predictors of future meeting attendance. The authors propose several explanations that are consistent with this phenomenon: (i) Firms with pending, but not yet granted, patents attend the committee to have the standard incorporate their intellectual property. (ii) The information garnered at these meetings help advance firms' intellectual property portfolio. (iii) Knowledge diffusion at the meetings may lead to firms citing patents of other firms attending standard meetings.

Of these three explanations, only the third one seems innocuous in terms of the SSO's ability to converge on the best possible standard. By contrast, participation in committees in order to further the cause of pending patents suggests that the final standard might be biased towards the (new) IPRs of participants. Finally, the gathering of information at SSO meetings lends itself to at least two rather different interpretations. On the one hand, one might see it as a sort of "spillover" that improves the technological knowledge of participants and hence improve their future innovative performance. On the other hand, one might fear that participants use the knowledge gathered at SSO meeting to either file new patent applications covering unprotected aspects that seem likely to become part of the standard or – more insidiously – in order to modify current patent applications to ensure that the wording of their claims actually read on the standard under development. In this respect, it is interesting that Berger, Blind, and Thumm (2012) find that SEPs typically have more claims and amendments than non-essential patents.

Using a small sample of European companies, Blind and Thumm (Research Policy, 2004) find that the greater the patent intensity of the firms, the less likely they are to join an SSO standardization process. This is clearly a concern as it indicates that SSOs might be unlikely to draw on some of the state of the art technologies. This concern is reinforced by the work of Bekkers, Bongard, and Nuvolari (Research Policy 2011.) They consider the case of W-CDMA and examine the factors that influence firms to claim essential patents. They find that the involvement in the standardization process is a stronger factor than the technical merit of the patent in determining the likelihood that a firm will declare a patent to be essential.

4.1.2. Timeliness

The main difficulty in assessing the timeliness of SSO decisions is to find some form of counterfactual: timely compared to what? Since what matters is the time spent *to design a standard*, the general performance of SSOs could only be compared to that of a different process that also achieves standardisation. This essentially leaves de facto market standardisation (or possibly open standard organisations) as possible benchmarks. However, the paucity of such benchmarks (there are few de facto standards) and the fact that different standard-setting mechanisms seem likely to emerge endogenously in industries with very different characteristics make such an exercise essentially hopeless.

However, it is still possible to get some insights on how the timeliness of SSO decisions might be affected by changes in their environment or changes in their internal rules. This is the approach used by Simcoe (2012) who looks at the time taken to develop standards for SSOs in charge of producing many of the standards used to run the internet. Simcoe's hypothesis is that, because these SSOs rely on consensus – or at least on super-majority rules, a sharpening of vested interests across members would be likely to lead to longer delays. This hypothesis is derived from a random bargaining model. In such models some negotiation delay is required to ensure that SSO participant converge on the best possible standard. However, when each party has vested interests in a particular technology, then bargaining delays become excessive. In this sense the predictions of the random bargaining model used by Simcoe have very much the same flavour as those obtained by Farrell and Saloner (1988).

To test the prediction of the model empirically, Simcoe needs some measure of vested interests. One approach would be to compute measures of the distribution of vested interests for each SSO, presumably relying on patent-ownership data. However, such an exercise would not be particularly reliable as such measures would be heavily influenced by one's interpretation of the technological strengths of each party and would also be vulnerable to strategic patenting behaviour. Simcoe avoids such drawback by relying on the "natural experiment" created by the increasing commercialisation of the internet. When the internet was not a source of significant wealth, the vested interests of SSO participants must also have been rather muted. As the commercial value of the internet increased, one would naturally expect divergent preferences to become more prominent. This empirical strategy is implemented by looking at internet-related SSOs over the period 1993-2003. For each SSO, an index of commercial significance is developed based on the proportion of private sector participants (the "suit to beard" ratio). A special feature of the IETF rules makes it possible to adopt a difference in difference approach: the IETF publishes both standards and "non-standard" ideas that are developed within the same type of committees (and indeed often the same committees) as standards. Since such non-standard ideas should not raise conflict of interests linked to the commercial nature of the internet they can be used to estimate the time that it should take to reach a decision on standards in the absence of distributional conflict stemming from vested interests. Simcoe finds that a one-percent increase in private sector participation leads to an additional delay of almost 8 days. Overall, he estimates that the total increase in private sector participation observed over the period contributed an additional eight months to the average SSO process. There is therefore reliable evidence that distributional concerns contribute significantly to the cost of

standardisation through SSOs. As we discuss in section 5, this of course raises the question of how such distributive conflicts might be minimised.

4.1.3. Access on reasonable terms

There are two main facets to the issue of accessibility on reasonable terms. The first aspect relates to the FRAND commitments that the vast majority of SSOs require from participating IPR-holders. As we will see below, FRAND commitments are designed to address the adverse consequences of the ex post market power wielded by the owners of IPRs that read on the selected standard. As such this issue is better discussed in the section on competition policy concerns.

The second aspect of accessibility – and hence pricing – is the so-called “royalty-stacking” issue. This is familiar territory. Once a standard has been defined, IPRs that read on this standard are complements from the point of view of the users who want to implement the standard. As is well-known since the seminal work of Cournot (1838,) independent pricing of complements leads to a higher total price for the “bundle” than if all prices had been chosen by a monopolist. Notice that this “stacking” problem has very little to do with market power: while stacking only arises if components are priced above variable costs, it can lead to very large differences between the sum of independently set prices and the sum of monopoly prices even if the market power of the bundle is not very strong.

The issues of hold-up and royalty-stacking have been conflated in much of the recent literature on the antitrust treatment of Standard Essential patents.⁶ This is unfortunate as royalty stacking is not fundamentally a competition policy issue. Competition policy deals with market power. While royalty stacking can only arise if the owner of a given component has *some* market power, very little market power can lead to very high “stacks” if the number of complementary components is large. Hence royalty-stacking is essentially a problem arising from the sub-optimality of independent price setting decisions in the presence of strong complementarity links, which is exactly the opposite of traditional competition policy concerns about coordination of decisions across independent entities. There is another fundamental difference between the hold up and stacking issues. While an ex ante commitment to royalty terms and condition would solve the first issue, it would not resolve royalty-stacking since the complementary goods pricing effect is the same ex ante as well as ex post (when the standard has been chosen.)

From an empirical point of view, the obvious questions about stacking are “is there stacking?” and “how large are the consequences of stacking on the total price of accessing a bundle of technologies”. Given how much ink has recently been spilled over the “stacking” problem, it is remarkable that the existing empirical literature is so scarce.

⁶ See for example the literature review in Gerardin, Layne-Farrar and Padilla (2008), where the authors note that early disclosure of patenting terms within SSOs has been proposed as a solution to the royalty-stacking problem. While such disclosure would allow for a more informed choice of the economically efficient standard, it does nothing to address the issue of complementarity.

Since the theoretical basis for the royalty-stacking problem is the inefficient pricing of complements by independent patent owners, it seems reasonable, as a first step to study the prevalence of such “patent thickets”. Two necessary conditions must be satisfied for the existence of such thickets: a given product (e.g. a smartphone) must involve the use of a large number of patent rights and the ownership of the relevant patents must be sufficiently dispersed.

Von Graevenitz et al (2013) use a specificity of the European Patent Office (EPO) patent application process to build a patent-specific measure of “thickets”. At the EPO, patent examiners rate each patent cited in a given patent application as potentially blocking on its own, potentially blocking in conjunction with others or not potentially blocking. The authors can then determine whether triples of patents held by different agents ‘block’ each other in the sense just defined. Relying on the incidence of such mutually blocking triplets, the authors obtain what they refer to as a measure of “technology complexity” for the patent portfolio of each firms. These measures can then be used to rank technology areas from the most to the least “thicket prone”. They find significant thicket presence in nine out of thirty technology areas. Notice however that the Von Graevenitz et al. measure mostly addresses the first of the two necessary conditions for the existence of “thickets”, i.e. the need to clear a number of IPRs in order to produce a product. By focusing on triplets, it does not fully consider the extent of the dispersion of the ownership of such rights.

For this reason, the authors also consider a measure of “fragmentation” first introduced by Ziedonis (2004) who measures the existence of thickets by constructing an index of the dispersion of patent-ownership for a number of firms. For each firm, the index is based on the dispersion of backward citations found in the firm’s own patent portfolio, capturing the idea that the owners of the cited patents are likely to be companies from which the firm might need a license. Ziedonis herself only uses this index to study the semi-conductor industry. In principle, however, the proportion of firms facing significant thickets could be used as an alternative manner of identifying “thicket-prone” sectors of activities. Interestingly though Von-Graevenitz et al (2013) find that the correlation between their measure of “complexity” and the Ziedonis measure of “fragmentation” is extremely weak. Since both of the dimensions captured by these two indices ought to be present for a thicket problem to truly arise, their results can actually be seen as casting doubt as to the prevalence of fully fledged “thickets”.

Of course the mere existence of patent thickets is not in itself of great interest to the analysis of SSOs since, by their very operation, SSOs do inevitably create thickets of patents around the standards that they develop. On the other hand, the effects of patent thicket on firm behaviour are relevant, even if the thickets involved do not arise as part of a standard-setting process.

Even the presence of thickets does not automatically make royalty-stacking a significant issue as mechanisms such as cross-licensing, patent pools and reputation might suffice to minimise the problem. It is therefore important to also have credible empirical evidence on the *effect* of thickets. Unfortunately, estimating the direct effect of thickets on the total royalty paid by downstream developers runs into a number of fearful obstacles. Not the least of these obstacles is the fact that there is very little systematic information on the royalties paid by licensee, not only because such data is just not collected but because licensing terms are often covered by confidentiality clauses. Moreover, even if reliable royalty data were available *and* one could

identify the set of licenses required to produce a given product, simply regressing total royalties onto the number of patents and their ownership concentration (ideally with an interaction term) would only provide useful estimates if one could also control for the different “value” of the downstream product and if one believed that the effect of other factors such as hold-up were truly independent of the pure “complementarity effect”. Accordingly, the empirical literature has looked for indirect ways of gauging the effect of thickets on licensing behaviour.

The first approach consists at asking whether – and how – the presence of patent thickets has affected patenting behaviour. In that vein, Ziedonis (2004) tests the hypothesis that the rate of patenting increases when thickets are present. The link between thickets and patenting behaviour rests on the idea is that, when transaction costs are high and ownership of required patents is dispersed, firms that want to use the technology are likely to face adverse licensing terms unless they have a significant patent portfolio of their own. This creates an incentive to invest more in licensing *for defensive* motives. Ziedonis further reasons that the desire to be well armed for ex post negotiation should be higher in industries where users are intensive in (sector-specific) capital, since this intensity makes them more vulnerable to ex post-licensing bargaining in the first place. She finds that, the effect of increased patent dispersion on patenting is five times higher for capital-intensive firms, confirming the initial hypothesis.

While the Ziedonis results do suggest that dispersion of ownership matters, it might not necessarily tell us much about the existence of “stacking”, for two main reasons. Firstly, patenting occurs for both defensive reasons and more traditional “innovation” reasons. While it is reasonable to assume that thickets increase the incentives for defensive patenting, one should also factor in the fact that the presence of thickets – precisely because they lead to suboptimal pricing that decreases the welfare of both licensors and licensees – should also decrease incentives to patent for “innovative” reasons. Whether one believes that Ziedonis’ approach adequately controls for this depends on whether or not one thinks that this “offensive” effect is also correlated with the capital-intensity of the sector. Secondly, and more importantly, it is important to be precise about the economic reasoning that underlies the tested hypothesis. When developing her hypothesis, Ziedonis’ argues that concentration of patent ownership is more likely to lend itself to *ex ante* licensing, while dispersion would tend to lead to ex post negotiations, i.e. negotiations that arise when the potential licensor has already sunk a significant part of its technology-specific investment. The reason for such a difference is not dispersion of ownership itself but the fact that *because of transaction costs* dispersion implies that it is harder for potential users to identify ex ante the set of patents that they would likely infringe.⁷ Once licensing occurs ex post, it follows immediately that potential licensors with greater amounts of technology-specific sunk capital are more vulnerable and hence would want to protect themselves more by developing their own defensive portfolio. Putting the two parts of the argument together insures that there will be a correlation between the dispersion of ownership of patents likely to “read” on a given firm’s products and that firm’s own patent investment. However, it should by now be clear that this correlation has absolutely nothing to do with the complementary pricing issue which is at the basis of royalty “stacking”. At its core, the

⁷ Transaction costs linked to actual negotiations is not a factor since it would affect ex ante negotiations as well as ex post negotiations.

Ziedonis paper is a joint test of the existence of the type of transaction costs just described and the importance of hold up effects. From the point of view of stacking, whether negotiations occur ex ante or ex post should essentially not matter.

Noel and Schankerman (2013) also examine the effect of two types of strategic patenting behaviour in the computer software industry. The first strategic dimension is portfolio size which, as in Ziedonis (2004), is seen as a proxy for the strength of a firm's bargaining position. In this respect, the authors find that having a large patent portfolio (compared to rivals) has a positive effect on the firm's market value, its patenting and its R&D. However, they find little evidence that portfolio expanding strategies are a sort of prisoner's dilemma where all firms try to improve their bargaining position at the eventual expense of the profitability of all involved: a proportional increase in all portfolios is still found to have a positive impact on the value of the firms. The second aspect investigated is the fragmentation of patent rights, which reduces market value but increases R&D and patenting.

So again, while the Noel and Schankerman paper, does not provide us with any evidence as to the size of the "complementary pricing" effects which is the essence of royalty stacking, it shows that the bargaining issues (including hold-up) that arise with ex post licensing and the fragmentation of ownership do affect the behaviour and performance of innovative firms. Interestingly though, their evidence also suggests that the net effect of such factors on R&D activity and the value of innovative firms is not necessarily negative.

As discussed in Gerardin et al. (2008), the price complement mechanism has received a lot of attention in both the biotech and the telecommunication industries. While the discussion of patent thickets in telecoms is too recent to have generated significant empirical research, the debate on the importance of the "complements" effects has been raging for a while in biotech. The growing consensus is that there is very little evidence that royalty stacking (or the piling up of transaction costs) has been significant enough to cause noticeable harm. Along the same lines, since the complement pricing problem also leads to lower profits for licensees, it should be associated with lower levels of R&D and patenting, patterns that are simply not found.

Patent pools

If one believes that independent pricing of complementary goods is indeed a sizeable issue, then a natural solution to the stacking problem would be to let the owners of standard essential IPRs set their licensing terms *jointly*. This could be achieved by mandating that SEP-holders all join a standard-specific patent-pool, where users of the standard could clear all relevant IPRs for a single fee determined jointly by pool members. Alternatively, voluntary participation in standard-specific patent pools could at least partially alleviate the royalty stacking problem. In this sense, patent pools and their properties can be seen as an important part of the evaluation of SSOs.

The effectiveness of patent pools as a way of resolving the royalty stacking problem hinges on three factors: pool participation, the effect of patent pool on innovation and the antitrust issues that patent pools themselves might raise. We will deal with innovation in section 4.2. We begin with antitrust issues.

The most obvious concern related with patent pools is that they could be used to actually raise prices (royalty). Lerner and Tirole (2004) have shown that, if we focus on patent pools that would form spontaneously – i.e. without policy intervention – then a clause that ensures that pool members can still license their own IP separately outside of the pool suffices to ensure that the patent pools that do emerge would be socially desirable. This is true whether or not those pools include patents that are strict complements or whether their membership is more diverse. This seems to imply that patent pools would be an ideal solution to the SSO-related stacking problem...but this would be misleading, for two reasons. Firstly it is not correct to think of the set of patents declared essential to a standard as strict complements. In practice, many of these patents are actually not essential to the standard. Moreover, for legal certainty reasons, it makes sense to also license some substitutes: it is not always clear whether a feature of a standard is only covered by patent A or whether patent B might also have a claim on it. Secondly, as soon as we depart from the case of perfect complements, the fact that the patent pools analysed by Lerner and Tirole *emerge spontaneously* is crucial to their finding. One cannot therefore assume that *mandated* pools would have the same desirable properties.

It is also important to realise that Lerner and Tirole do not directly address the issue of patent-participation. Their patent pools are industry-wide. Indeed, Brenner (2009) has shown that, if the extent of patent pool participation is endogenous, the “independent licensing clause” emphasised by Lerner and Tirole is no longer sufficient to ensure the socially optimal formation of patent pool.....unless pool membership is determined by unanimity rule, a feature that Competition Authorities might find problematic (see section 4.2). The fact that, in practice patent pools include only about 50% of industry actors and that even SSO-related pools usually do not include all SEPs suggest that the issue of participation is crucial and deserves to be investigated further.

4.1.4. Updates

The impact of SSOs on the rate of innovation in a given sector of activity involves the SSO's impact on technological markets at large and is therefore a competition policy concern. By contrast, the ability of a SSO to keep updating its own standard to incorporate new developments and meet the changing needs of the sector is an important aspect of its operating performance. Unfortunately we are not aware of any systematic empirical research investigating this aspect of SSO behaviour. Using data on 3,500 ICT standards, Baron, Blind, and Pohlmann (2011) find that the number of SEPs that (claim to) read on an standard has a negative effect on drastic innovation in standards, i.e., there is a bias towards preserving and improving the current standard, rather than adopting a new standard.

Clearly, there is room for further research. Of particular interest would be the link between the pattern of ownership of SEPs and the rate at which standards are updated. For example, does greater concentration hinder updating (because more parties have to agree to the changes) or does it favour it (because each party loses less from adding some additional patent holders to the mix)?

4.2. COMPETITION POLICY CONCERNS

To understand the various competition policy concerns related to SSOs one must first distinguish between three types of markets that might be affected by the activities of SSOs. The first and most obvious of these markets are the markets for the products that would make use of the standard. There are two main concerns relating to behaviour in the technology market. The first concern is that SEP holders that are vertically integrated might seek to exclude rivals from the relevant downstream markets by either refusing to license the technologies required to practice the standard or by imposing discriminatory terms⁸. The second concern is the fear the SEP holders would exploit the additional market power gained through the inclusion of their IPRs into the standard to charge excessive prices. Two additional issues arise because of the vertical relationship between the technology markets and the markets for standards. Firstly, competition authorities might suspect that SSO membership rules might be devised to exclude some holders of technology from the standard-setting process. This not only would lead to the choice of inferior standard but it might eventually reduce the innovation incentives of firms that find themselves systematically locked out. Secondly, when rivalry between standards is a realistic scenario, SSOs might be tempted to forbid their members from making their technology available to rival standard-setting bodies. Again, the resulting lack of “standard competition” would not bode well for either the quality or the timeliness of the chosen standards.

4.2.1. Exclusion

As discussed, one of the main sources of antitrust concerns with respect to SSOs (and patent pools) is the fear that these organisations might be used to exclude rivals. Some types of exclusions are easily guarded against. For example, Competition Authorities usually require that patent-holders who take part in the activities of one SSO should still be able to make their patents available to a rival SSO.⁹ This should be sufficient protection against the loss of inter-standard competition. Other potentially exclusive tactics are trickier to handle. Let us first consider the question of participation. Can SSO membership be restricted by a group of firms in order to exclude others? The simple answer is that, in that absence of appropriate competition policy, it definitely could...and indeed definitely would. Lampe and Moser(2012a) provide convincing evidence on this issue. Although their study concerns patent pools rather than SSOs, it is not hard to see why it is also relevant for SSOs which are after all patent pools aimed at coordinating technologies (i.e. at forming a technological bundle) rather than aimed at

⁸ Notice that, contrary to what is often argued, “non-practicing entities” that hold SEP are not an issue in this respect as they have no incentive to foreclose downstream market competitors. See Schamlensee (2009) for a thorough discussion.

⁹ See the European Commission’s *Horizontal Guidelines*.

coordinating sales of technologies (i.e. forming a “commercial” bundle). Moreover, patent pools often arise in the shadow of standard setting processes.

The main attraction of the Lampe and Moser analysis is that the authors look at a period over which competition policy was deliberately “relaxed” (or even suspended) in order not to impede economic recovery in the post-depression era. Without any “fear of the gendarme”, patent pool members were found to practice exclusion with some delight. Out of the twenty patent pools reviewed by the authors, eleven did not license a single outside firm, three pools licensed up to three companies not belonging to the pool and only 6 pools licensed extensively outside of the pool itself. This pattern was not only explained by a reluctance to license to downstream rivals, it was also linked to “inter-standard” competition, as the authors conclude that “archival evidence suggests that the pools studied.....used licensing as a means to limit competition with substitute technologies”.

4.2.2. Access: Hold up

In economic terms, there is a potential hold up problem if at least one of two parties to an agreement must make relationship-specific investments at a time where the terms of the agreement are not set fully and irrevocably. In other words, “hold up” problems only emerge if two necessary conditions are fulfilled: the relationship involves specific investments and there is room for opportunism.

Thanks to a few high-profile antitrust cases in the smartphone industry, SSOs have recently found themselves in the middle of a “hold up” storm. Unfortunately, the discussion of hold up in this context has been far from limpid. Let us therefore re-state a few simple points. Firstly, as stated above, hold up requires that two conditions be met. The first condition is the presence of investments whose value is higher within the relationship between the relevant parties than outside of this relationship. In other words, one needs (partially) sunk investment. The second necessary condition is that there is room for opportunism by the party that did not make the investment. This arises if fool-proof complete contracts cannot be either written or enforced. We can then ask where and when in the standard-setting process these two conditions are met. To do this, we must first identify the relevant “bilateral relationship”. Clearly, in this case, this is the relationship between SEP-holders and potential users of the standard. This relationship involves the licensing of SEPs to the users. There are three types of investments that are linked to this relationship: the R&D investments made by the SEP holder before the start of the SSO process, the product development investments of the potential users and the effort and additional standard-specific R&D that goes into the SSO process itself.

The next step is to identify which of these investments is in fact “relationship specific”, as required for hold up effects to arise. This of course depends on when the terms of the licensing agreement are irrevocably set. As we have seen, although potential SEP holders who participate in SSOs often have to agree to license on “FRAND” terms, the actual licensing terms are usually determined only after the standard is agreed upon and they are usually determined through bilateral negotiations between each SEP-holder and each standard user. Even though

some SSOs require the disclosure of “worst” terms before the standard is chosen, it seems fair to say that such commitment still leave a lot to be determined once the standard is actually set. So, let us proceed under the assumption that meaningful bilateral negotiations only arise after the standard is agreed upon.

At that stage, the R&D costs incurred to obtain the discoveries covered by SEPs has clearly been sunk. This is true of both SEPs that existed before the start of the SSO process and of those that were developed in to “fill in the gaps” as the standard was being developed. In terms of specificity though there is a potential difference between these two types of SEPs. SEPs corresponding to “fill the gap” technologies can safely be assumed to have little value outside of the chosen standard. In that sense the corresponding investment is both sunk and relationship-specific, setting the stage for a potential hold up *by the potential users*. Pre-existing patents that become SEPs might on the other hand have commercially profitable applications outside of the SSO context. To that extent then, it would not be appropriate to consider the full value of the sunk investment as relevant to the SSO-related hold up problem. Still, it should also be remembered that some types of technologies – such as those covering communication protocols – have by their very essence very little value if they cannot be used to communicate...which usually requires some form of standard. For such technologies, then, almost the entirety of R&D costs should be seen as relationship-specific sunk costs. It should therefore be clear that, although hold up of users by SEP-holders has gotten most of the attention, the conditions for significant hold up of SEP holders by users are also fulfilled. As Williamson (1983) pointed out, bilateral hold-up can be a solution to the hold-up problem itself as it amounts to an “exchange of hostages”. Hence, even if we have the right conditions for hold-up on both side of the relationship, it is not at all a priori clear that hold up would have a significant effect on the equilibrium terms of licensing transactions.

Other costs that have been sunk – and are undeniably relationship specific – are all other costs of participating in the standard-setting decision. Indeed, if there were no such costs, then hold-up of licensees would not be an issue at all as any excessive royalty demand would lead to the (costless) redesign of the standard.

On the side of the potential licensee, the level of relationship-specific sunk investment will depend on the industry. In fast-moving industries, where first mover advantages are crucial, one should not expect product developers to wait until a standard has been finalised before starting to work on their products. Indeed, in these industries, standard-setting efforts are often a response to new needs that become apparent because potential users are working on new products. In this case, a substantial proportion of product development costs might well qualify as relevant specific sunk costs in the post-standard bilateral negotiations, exposing the users to hold up on the part of SEP-owners.

It is remarkable that, given the undeniable existence of a potential hold up issue on both sides of the SEP-holder-users relationship, competition authorities have focussed their attention uniquely on the plight of the licensee. One incarnation of this concern is the series of on-going cases regarding the use of “injunctions” by SEP-holders. The case against widely allowing SEP-owners to use injunctions against potential licensees is based on the following steps. Firstly, SEP holders benefit from the “unearned” increased market power that having their technology included into a standard confers. There is no denying this claim as it is at the core of the hold-up issue: ex post market power is higher than ex ante market power. In order to avoid this issue,

SSO participants usually commit to licensing on FRAND terms. FRAND terms are then usually seen as the terms that would have prevailed if negotiations had occurred ex ante rather than ex post. The Competition Authorities' fear is then that injunctions – which undeniably increase the relative ex post bargaining power of SEP-owners compared to potential licensees – can only be used to extract terms that are “not FRAND”. Furthermore, it is often pointed out that the problem of hold up by SEP-owners is especially worrisome because it arises from the coordination of efforts between technological rivals to design a standard.

Although broadly accepted, this approach has serious flaws. Firstly, in the absence of any objective measure of “FRAND” terms or even any attempt to determine what ex ante terms would have been, FRAND is simply a red herring. From an economic point of view, we find ourselves in a situation of bilateral hold-up. Given this, the question of whether allowing for broad injunctions is or is not desirable simply hinges on the relative strength of the two “hold up” issues. If the hold up of licensees is more severe than the hold up of SEP-owners then injunctions are not desirable since they further increase the bargaining power of SEP-owners. On the contrary, if the hold up of patent-holders is more severe than the hold-up of users, then allowing for injunctions might in fact be a way of reducing the adverse consequences of the bilateral hold up problem. It is therefore unfortunate that the whole debate on injunction has proceeded without *any* evidence as to the strength of the hold-up issues on either side of the licensing relationship. It is even more unfortunate that we are not aware of any systematic evidence on this issue. Of course, the correct policy on injunctions also depends on how much of a change in bargaining power is produced by allowing the SEP-owner to use injunctions fairly readily. Injunctions are a rather blunt tool that can shift the balance of power in litigation quite significantly. Even if, on balance, the hold-up issue experienced by SEP holders were more serious than the hold-up issue faced by potential licensee, it might be that a relatively free use of injunctions would switch bargaining power too drastically in favour of the SEP-owners. This possibility seems most realistic in industries where the market “moves fast”: by being absent from the market even for a relatively short period of time, a licensee under an injunction could lose not only large volumes of sales (think Christmas season) but could also be left behind in the race to establish its product. By contrast, the absence of an injunction might not have such a quick and destructive impact on SEP-holders, at least if they can be assured that an infringing standard users will eventually be made to pay appropriate royalties.

The traditional argument against injunctions is also flawed in a more subtle manner. It just is not correct to claim that the hold-up by SEP-owners is worse (or more worthy of antitrust interest) than the hold-up by licensees “because the power to hold up was obtained through concerted action”. To assess the competition law aspects of standards achieved through SSOs, the correct counterfactual cannot be a world without a standard. It must be a world where a standard would have been achieved without coordination between rivals. This could occur through de facto standardisation. Alternatively, for argument's sake, one can also think of a counterfactual where a standard would be obtained by randomly choosing a (compatible) technology for each of the standard's components. In either case, *the resulting hold-up issue is just as severe as if the standard had been chosen “jointly by rival”*. It is therefore erroneous to link the issue of hold-up, and hence the issue of FRAND, to the cooperative nature of the standard-setting process. Of course, this cooperative nature is relevant for other issues such as the incentives to exclude rival technologies from the design or the incentives to foreclose

alternative standards, but it is simply irrelevant to the main issue of hold-up and FRAND commitments.

While there is no systematic evidence regarding the relative extent of hold up on the two sides of the licensing relationship and we do not either have any evidence as to the overall effect of this two-sided hold-up problem, there is some related empirical work.

Simcoe and Graham (2009) find that SEPs have higher litigation rates than comparable non-SEP patents. This is especially true for SEPs held by small firms. The authors interpret this pattern as resulting from the fact that small firms are less likely to be “repeat players” in the standard-setting process so that they face less of a reputational cost from “violating” FRAND commitments and trying to use their enhanced market power to extract high royalties. Alternatively, they suggest that litigation might be more frequent with small firms because their limited portfolio makes traditional “cross-licensing” settlements harder to achieve. Taken at face-value, these conclusions do imply that there is indeed a significant problem of hold up by SEP-holders. If there was not, then the first explanation proposed by the authors would not make sense. While this is the implication of one of the interpretations proposed by the authors, it is not clear that this conclusion necessarily follows from the results obtained by the authors. There are two issues. Firstly litigation rates are computed as litigation per patent. However, SEPs are likely to be in higher demand than non-SEPs with the same characteristics...precisely because they are part of a standard. This means that, as SEPs will be involved in many more licensing negotiations, there are also more opportunities for disagreement and hence litigation. Ideally, then, the litigation rates should be computed with respect to the total number of bilateral licensing negotiations. As information on such negotiations is not generally available, this is cannot be done...but then it is not clear what the higher rate of litigation of SEPs actually indicates.

Turning to the differential behaviour of small SEP-owners, there are other explanations for their behaviour than the fact that they might feel less constraint not to exploit their “hold up” power. For example, small SEP holders might be more vulnerable to litigation than larger firms with more experience in defending their IP rights.

Another type of evidence is provided by Hussinger and Schwiebacher (2012, 2013). In these papers, the authors examine the benefits to firms from disclosing their IP to SSOs. Using data for the 1986 to 2005 period, they find that disclosure of patents to standard setting organizations are correlated with firms' valuations. Thus they conclude that the loss from FRAND commitment (and hence the inability to license exclusively) is outweighed by product market benefits from a standardized technology. While interesting, this result lends itself to many possible interpretations. The first one is indeed that the commitment made to have one's patented technology included in a standard are adhered to and that such commitments are worth incurring in light of the potential benefits from reading on the chosen standard. However, one could also see this as evidence that, in spite of the commitment made, SEP owners are able to extract higher royalties from potential licensees. In this sense, this could be seen as indirect evidence that the hold up by SEP owners might be more severe than the hold up by potential licensees. Yet another possibility is that the effect on firm's valuation has actually nothing to do with the specific licensing process of SEPs and hence with hold-up. Inclusion in a standard might just be seen as external validation/certification of the quality of the firm's R&D efforts.

4.2.3. Innovation

Competition policy also worries about the impact of firms' behaviour – and hence the impact of SSOs – on innovation. In what follows, we take standardisation as a *fait accompli*, i.e. we do not ask whether standardisation itself leads to more or less innovation than a world without standards. Rather we ask whether the way that SSO operate is likely to be harmful to the innovation incentives of both patent-holders and potential standard users. In particular, what is the likely effect of royalty-stacking, hold-ups and other features associate with SSO on the overall rate of innovation?

In terms of theory, standard-setting organisations can affect innovation in a large variety of ways. For example, SSOs that favour the “political strength” of firms with large portfolios at the expense of simply choosing the best combination of available technologies would clearly encourage firms to pursue “numbers” instead of quality. In this sense, any empirical evidence that suggests strong “political economy” bias within SSOs should also alert us to the possibility of a corresponding dynamic inefficiency.

Another, more subtle mechanism, is that dominant SSOs turn R&D investment into “winner takes all” races since, especially for standard-oriented patents such as telecom patents, only patents that are retained within the standard are likely to produce significant income streams. As is well-known, such “winner takes all” environments encourage excessive duplication of effort and might also discourage the pursuit of an optimal variety of research routes.

Most obviously, of course, SSOs might affect the balance between innovation in the market for standard-related innovations and the market for innovations related to the downstream products that apply the standard. The greater the hold up by patent-holders compared to the hold up by licensees, the more SSOs would change R&D incentives in favour of standard-related technologies. The same would be true if SEPs turned out to involve a truly higher rate of litigation than other similar patents.

Finally SSOs can have a more direct effect on innovation if they provide an environment that promotes knowledge sharing. Such sharing has the traditional opposite effects of potentially decreasing *ex ante* innovation while increasing the social value of any given amount of innovation. In this respect, Gandall, Gantman and Genesove's (2006) results suggest the possibility of knowledge diffusion or spillovers at standardization meetings, which would have dynamic benefits. The meeting data examined by the authors for the TIA TR-30 committee (which was responsible for setting analogue standards in data transmission systems and equipment) showed that the committee and its subcommittees met regularly, with approximately five to six meetings per year. Further, many of the participants attended consistently over a period of time. Hence, the committees created an environment for knowledge spillovers, similar to those at academic conferences.

There is very little direct empirical evidence on the non-spillover channels through which SSOs can affect innovation. There are however a few useful papers on the link between patent pools and innovations. Since SSOs have some common features with patent pools and, moreover, patent pools can be used to minimise the “royalty-stacking” issue to which SSOs might give rise,

it is interesting to briefly review this small literature. Two papers by Lampe and Moser (2010 and 2012a) are of particular interest.

Lampe and Moser (2010) look at the effect of the 19th century Sewing Machine Patent Pool. Using differences in differences, they first examine the effect of the formation of the pool on the patenting behaviour of both participants and firms that remained outside the pool. They find a significantly negative effect of the pool on the number of patents obtained by both members and non-members. It is not just that patenting decreased when the pool was formed. It also picked up as soon as the pool was dissolved. Of course there are many reasons why patenting activity might not adequately reflect innovation efforts. Note however that, for pool members at least, the traditional hypothesis is that pool participation leads to more patenting for strategic reasons. This is because pools often split the spoils based – partially at least – on the number of patents held by each member. If the strategic patenting incentive truly is to patent more once one joins a pool, then the Lampe and Moser result is only strengthened: the fact that less patenting was observed is even more likely to indicate that innovation actually slowed down. Nevertheless, the authors also construct a measure of “real” innovation that is independent of patent counts: the evolution of the number of stitches per minute that the machines were able to perform. The analysis of this indicator confirms their patent-based findings; the formation slowed down the pace of progress, which picked up again once the pool was dissolved. While the authors do not explain why innovation by pool members has decreased, they trace the decrease in the innovation of non-members to the increased litigation following the formation of the pool (pool members also pool resources for litigation).

In a similar vein, Lampe and Moser (2012a) examine patent pools in 20 industries over the period of the “New Deal” when antitrust enforcement was deliberately weak. Using a difference in difference approach, they find that on average the formation of a pool decreased patenting by 16%.

Overall then, there is rather little empirical evidence about the effect of SSOs on innovation. There seems to be reasons to believe that SSOs encourage knowledge diffusion. On the other hand, there is fairly strong evidence that (unregulated) patent pools actually decrease the rate of innovation. However, while this points to a cost of using pools to address the royalty-stacking issue stemming from SSOs, it is not clear that the patent pool results can be directly translated to the SSO themselves. In particular, SSOs do not get involved in litigation and, absent an accompanying pool, SEP owners enforce their patents separately. The “litigation” mechanism suggested by Lampe and Moser cannot therefore apply to SSOs themselves.

5. INTERNAL SSO RULES

So far we have reviewed the – relatively scarce – empirical literature that sheds some light on various aspects of SSOs performance. However, one could argue that such assessments do not help us answer the relevant question. If standardisation is desirable and SSOs often are the

only realistic route to a standard, then whether SSOs perform well or not on average is largely irrelevant: whatever their drawbacks are, this is the cost that society has to bear in order to reap the benefits of standardisation. In this view, what should be of interest is how differences in the internal rules of SSOs and in public policy towards SSOs affect the performance of these organisations. We have already discussed in section 4 how competition policy might be used to minimise the anticompetitive effects of SSOs. In this section, we turn to the SSOs' own internal rules.

As we saw in section 2, SSOs are politico/economic institutions where influence within the SSO might matter as much as technical merit when it comes to have one's patents included in the agreed upon standard. A natural target for improving the performance of SSOs is therefore the rules that determine how power within the SSO is divided. In this spirit, Farrell and Simcoe (2002) theoretically examine which type of consensus rule is optimal for SSOs. They note that firms often have vested interests in promoting their own technology. This makes 'absolute' consensus difficult to achieve, and can significantly slow down the standardization process. Hence, they examine alternative (weaker) notions of consensus. They find that allowing neutral players to arbitrate and break deadlocks can improve on outcomes.

Another crucial determinant of SSO performance is participation. It is only if participation is broad and includes the firms with the best technology that there is any hope that the standard will be as good as technically (and economically) possible. In this respect, Layne-Farrar and Lerner (2011) find that pools that adopt proportional sharing rules based on the number of essential patents in the standard are less likely to attract potential participants. In particular, firms with high value patent portfolios are less likely to join pools with numeric proportionality.

Chiao, Lerner, and Tirole (2007) empirically examine the rules of standard setting organizations. They find that a provision requiring royalty-free licensing is negatively correlated with a provision requiring disclosure of intellectual property. This is as expected: if licensing does not involve payment, then there is no room for either patent ambush or hold up. In a similar flavour, they find that RAND licensing requirements are positively correlated with a provision requiring such disclosure. Unfortunately, this type of analysis does not tell us much about how the actual performance of SSOs might depend on the type of internal rules that they adopt. A huge empirical effort in this direction is needed if we are to further our empirical understanding of SSOs.

6. CONCLUSION

In this paper, we developed criteria to evaluate the performance of SSOs. We then examined whether the empirical work has been able to shed light on the key issues. A particularly important contribution of the paper is showing the links between these criteria and the existing empirical literature (Figure 2.) The analysis in sections 4 and 5 discusses the relevant empirical work related to each criterion. This analysis shows that while a recent flurry of empirical work

on SSOs has been conducted in the last decade or so, we still lack empirical evidence on key issues. We believe that it is especially important to examine how differences in SSO structure and internal rules affect their performance and how such differences empirically affect incentives for innovation. We encourage researchers to focus their efforts in these critical areas.

7. REFERENCES

Baron, J., K. Blind and T. Pohlman, 2011, "Essential patents and standard dynamics", *IEEE Standards and Innovation*.

Baron, J. and T. Pohlman, 2013, "Who cooperates in standards consortia – rivals or complementors?", *Mimeo, Northwestern University Law School*.

Bekkers, R., R. Bongard and A. Nuvolari, 2011, "An empirical study of the determinants of essential patent claims in compatibility standards", *Research Policy*, 40:7, pp. 1001 – 1015.

Berger, F., K. Blind and N. Thumm, 2012, "Filing behaviour regarding essential patents in industry standards", *Research Policy*: 41:1, pp. 216-225.

Blind, K. and N. Thumm, 2004, "Interrelation between patenting and standardization strategies: Empirical evidence and policy implications", *Research Policy*, 33:10, pp. 15583 – 1598.

Brenner, Steffen. 2009, "An Optimal Mechanism of Patent Pool Formation," *Economic Theory* 40(3): 373-388.

Chiao, B., J. Lerner and J. Tirole, 2007, "The rules of standard-setting organisations: An empirical analysis", *Rand Journal of Economics*, 38:4, pp. 905 – 930.

Clements, M., and H. Ohashi, (2006) INDIRECT NETWORK EFFECTS AND THE PRODUCT CYCLE: VIDEO GAMES IN THE US, 1994–2002, *The Journal of Industrial Economics* 53 (4), 515-542

Cournot, A., 1838, *Researches on the Mathematical Principles of the Theory of Wealth*,

De Lacey, B.J., K. Herman, D. Kiron and J. Lerner, 2006, "Strategic behaviour in standard-setting organizations", University of Texas, mimeo.

Dranove, D., and N. Gandal, "The DVD vs. DIVX Standard War: Empirical Evidence of Network Effects and Preannouncement Effects," 2003, (with David Dranove), *Journal of Economics and Management Strategy*, 12: 363-386

European Commission's *Horizontal Guidelines*, available at <http://ec.europa.eu/competition/antitrust/legislation/horizontal.html>

Farrell, J. and G. Saloner, 1988, "Coordination Through Committees and Markets", *Rand Journal of Economics*, 19:2, pp. 235 – 252.

Farrell and T. Simcoe, 2012, "Choosing the rules for consensus standardisation", 2012, *Rand Journal of Economics*, 43:2, pp. 235 – 252.

Gandal, N., "Hedonic Price Indexes for Spreadsheets and an Empirical Test for Network Externalities," 1994, *The RAND Journal of Economics*, 25: 160-170

- Gandal, N., "Compatibility Standards and Complementary Network Externalities in the PC Software Market," 1995, *Review of Economics and Statistics*, 77: 599-608, 1995
- Gandal, N., Gantman, N., and D. Genesove, "Intellectual Property and Standardization Committee Participation in the U.S. Modem Industry," 2006, pps. 208-230 in Standards and Public Policy, Shane Greenstein and Victor Stango editors, Cambridge University Press.
- Gandal, N., Kende, M., and R. Rob, "The Dynamics of Technological Adoption in Hardware/Software Systems: The Case of Compact Disc Players, 2000," *RAND Journal of Economics*, 31: 43-61
- Gerardin, D., and M. Rato, 2007, "Can standard setting lead to exploitative abuse: A dissonant view on patent hold-up, royalty stacking and the meaning of frand", *European Competition Journal* 101.
- Von Graevenitz, G., S. Wagner and D. Harhoff, 2013, "Incidence and Growth of Patent Thickets: The Impact of Technological Opportunities and Complexity", *Journal of Industrial Economics*, LXI:3, pp. 521 – 563.
- Greenstein, S., 1993, "Did Installed Base Give an Incumbent Any (Measurable) Advantages in Federal Computer Procurement," *RAND Journal of Economics*, 24: 19-39.
- Prieger, J., and w. Hu, (2006), An Empirical Analysis of Indirect Network Effects in the Home Video Game Market, mimeo.
- Hussinger, K. and F. Schwiebacher, 2012, "The market value of standard-setting activities for technology providers", *Druid Society*.
- Hussinger, K. and F. Schwiebacher, 2013, "The value of disclosing IPR to open standard-setting organizations", *ZEW Discussion Paper 13-060*.
- Kohler, F., K. Blind and N. Thumm, , "An analysis of the filing behaviour of patents in standards", *Idei*.
- Lampe, R and P. Moser, 2010, "Do Patent Pools Encourage Innovation? Evidence from the 19th Century Sewing Machine Industry", *Journal of Economic History*, 70:4, pp. 898 – 920.
- Lampe, R. and P. Moser, 2012a, "Do Patent Pools Encourage Innovation? Evidence from 20 US Industries under the New Deal", *NBER, August 8, 2012*.
- Lampe, R. and P. Moser, 2012b, "Patent Pools: Licensing Strategies in the Absence of Regulation", in S.J. Jalal, B.S. Silverman and M.A. Cusumano (EDs), *History of Strategy (Advances in Strategic Management, Vol. 29)*, Emerald Group Publishing Ltd, pp. 69-86.
- Layne-Farrar, A. and J. Lerner, 2011, "To join or not to join: Examining patent pool participation and rent-sharing rules", *International Journal of Industrial Organization*, 29:2, pp. 294 – 303.
- Lerner, J., and J. Tirole, 2004, "Efficient Patent Pools." *American Economic Review*, 94(3): 691-711
- Liu, C., C. Kemerer and S. Slaughter, 2007, "Standards competition in the presence of digital conversion technology: An empirical analysis of the flash memory card market", *SSRN*.
- Noel, M., and M. Schankerman, 2013, Strategic Patenting and Software Innovation," *Journal of Industrial Economics*, Volume LXI, No. 3: 481-520
- Postrel, S., 1990, "Competing Networks and Proprietary Standards: The Case of Quadraphonic Sounds," *Journal of Industrial Economics*, 39: 169-185.

- Rysman, M. and T. Simcoe, 2008, "Patents and the performance of voluntary standard setting organizations", *Management Science*, 54:11, pp. 1920 – 1934.
- Saloner, G. and A. Shepard, 1995, "Adoption of Technologies with Network Externalities: An Empirical Examination of the Adoption of Automated Teller Machines," *RAND Journal of Economics*, 26: 479-501.
- Schmalensee, R., 2009, "Standard-setting, innovation specialist and competition policy", *Journal of Industrial Economics*, 57:3, pp. 526 – 552.
- Simcoe, T.S. and S.J.H. Graham, 2009, "Competing on standards: Entrepreneurship, intellectual property and platform technologies", *Journal of Economics and Management Strategy*, 18:3, pp. 775 – 816.
- Simcoe, T.S., 2012, "Consensus governance for shared technology platforms", *American Economic Review*,
- Weiss, M.B.H. and M. Sirbu, 1990, "Technology choice in voluntary standards committee: An empirical analysis", *Innovation and New Technology*, 1:1-2, pp. 111 – 133.
- Williamson, O., 1983, "Credible Commitments: Using Hostages to Support Exchange", *American Economic Review*, 73:4, pp. 519 – 540.
- Ziedonis, R., 2004, "Don't Fence Me In: Fragmented Markets for Technology and the Patent Acquisition Strategies of Firms," *Management Science*,

