

The Private Value of Software Patents

Bronwyn H. Hall and Megan MacGarvie¹

December 2009

Abstract

We investigate the value creation or destruction associated with the introduction of software patents in the United States in two ways. The first looks at the cumulative abnormal returns to Information and Communication Technology (ICT) firms around the time of important court decisions that impacted software patents, and the second analyzes the relationship between firms' stock market value, the sector in which they operate, and their holdings of software patents. We conclude that the market evaluated software patents as a negative development *ex ante*. *Ex post*, a greater number of firms in all ICT sectors invested in these patents, and these firms had slightly higher market values than those with no software patents. However, while we obtain clear evidence that the technological importance or quality of patented innovation mattered for the market value of hardware firms both before and after the legal changes, it is less clear that the marginal patent right *per se* was associated with increases in market value, and there are no significant valuation effects associated with patents for pure software firms after the change.

¹ Hall: 549 Evans Hall, Berkeley, CA 94720, bhall@econ.berkeley.edu. MacGarvie: 595 Commonwealth Ave., Boston, MA 02215, mmacgarv@bu.edu. This is a revision of a paper written in April 2005. We thank Anne Layne-Farrar and LECG for the use of the Corptech data, and we thank Jim Bessen for helpful suggestions and for supplying us with the list of software patents used in Bessen and Hunt (2003). We are also grateful to two anonymous referees and to Stu Graham, Bala Iyer, Michael Salinger, Lucia Silva, numerous participants in the Empirical Patent Research Conference, St. Helena, California, February 2005, the IIOC, Atlanta, April 2005, the ZEW Conference on Patents and Innovation in Mannheim, September 2005, the Seventh CEPR Conference on Applied Industrial Organization, Madeira, May 2006, the NBER Summer Institute 2006, the First Annual EPIP Conference, Munich, September 2006, and seminars at UCSD, University of Kansas, University of Michigan, University of Maastricht, LMU-Muenchen, and RIETI, Tokyo for extremely helpful comments.

The Private Value of Software Patents

Bronwyn H. Hall and Megan MacGarvie²

December 2009

1. Introduction

Ever since software became generally patentable in the United States in 1995, the wisdom of such a change has been widely debated. Proponents of the change argue that there is no statutory reason to exclude software (or computer-implemented) inventions from patentability, and also that patenting software has social benefits from disclosure and from the fact that it enables software components to be reused by others (Cohen and Lemley 2001). Critics, who are numerous, base their arguments on a series of considerations: that such patents have often been of low quality, that they discourage rather than encourage innovation, and that they have a negative effect on the growing open source model of innovation (Kasdan 1994; Barton 2000, 2002; Hunt 2001; Bakels and Hugenholtz 2002; Lunney 2001; Dreyfuss 2000, 2001; Meurer 2002). In a recent article in PC magazine, the columnist John Dvorak argued that software patenting is even bad for Microsoft, the largest patent holder among the pure software firms (Dvorak 2005). In Europe, patentability of software is still being debated, with no clear conclusion emerging (Blind et al. 2001; European Commission 2002; Spindler 2003).

The economic view of the patent system sees this debate as an inevitable one, given the nature of patents: with a patent grant we trade off short term exclusive (monopoly) rights to the

² UC Berkeley, University of Maastricht, NBER, and IFS London; Boston University and NBER.

use of an invention in return for two things: 1) an incentive to create the innovation; and 2) early publication of information about the innovation and its enablement. The argument is that without the patent system, fewer innovations would be produced, and those that were produced would be kept secret as much as possible to protect their returns from appropriation by others. In considering the economic impacts of a subject matter extension to software and business methods, the tradeoff between these benefits and the welfare cost of the grant of a monopoly right are at least as important as they are in any other technological arena.

Recent analysis also says that although *competition* may suffer when we grant a monopoly right to an inventor, it will benefit if this right facilitates entry into the industry by new and innovative firms (Hall 2003)³. Second, *innovation* will benefit from the incentive created by a patent but may suffer if patents discourage the combining and recombining of inventions to make new products and processes (Scotchmer 1991, 1996; Heller and Eisenberg 1998). There are several reasons to think that facilitating entry (the benefit of patents for competition) and discouraging recombination of elements (the cost of patents for innovation) may be particularly salient factors when considering the effects of patentability on the software industry. The software sector is characterized by low tangible assets and high intangible assets, which means that it is more important to have intellectual property rights in order to secure funding for investment. Software itself is typically cumulative and composed of a number of components which may come from a number of different sources, and where interoperability standards are

³ Cockburn and MacGarvie (2006) find evidence that firms holding patents are more likely to enter software markets. However, increases in the extent of patenting by other firms in the market reduce the rate of entry. Hall (2005) finds that firms in the ICT sector that enter the publicly traded sector after changes in patenting behavior in this sector in the mid-1980s have much higher patent valuations than those that enter earlier.

important. Both factors mean that combination and recombination of inputs is an essential part of software creation.

Evaluating the tradeoff between the benefits and costs of the introduction of software patents is a formidable, perhaps impossible, task. It is not made easier by the fact that the most important expansion of software patentability occurred in 1994-1995, exactly coincident with the internet revolution and the beginnings of its impact on the global economy's use of software for a wide range of new applications.⁴ To evaluate the impact of software patentability properly would require an examination of a counterfactual world without software patents and measurement of the differential impact on welfare and innovative activity in existing firms, new entrants, downstream buyers, and consumers. In this paper we look at one ingredient of the overall problem: the private returns to established computer hardware and software firms from the expansion of software patentability. If the patentability of software is beneficial for the economy, we expect that such firms will experience positive returns to its introduction, even if the overall economic effect is neutral or negative. That is, a finding that existing firms do not benefit financially from the introduction of software patents makes it probable that the impact on the economy as a whole is not positive either.

Existing research on software patents reveals a dramatic increase in the propensity to patent software over the last two decades, and argues that this indicates that holding patents on software has become both easier and more valuable to firms (Bessen and Hunt 2003). In this paper we explore the relationship between firm value and the patenting of software in two

⁴ It is possible that this timing is not a coincidence. The expansion of patentability has often come about as a result of the expansion of economic activity in a particular area. E.g., see Murmann (2003) on the chemical industry in Germany and Britain in the late 19th century. However the rapidity with which attempts to patent software followed on the introduction of the internet makes the causal explanation improbable. It seems more likely that the changes in software patenting in the mid-nineties were driven by the introduction of the personal computer in the mid-eighties and the consequential increased importance of stand-alone software firms.

different ways. First we conduct a series of event studies that look at the immediate market impact of changes in the patentability of software on firms in sectors where software patenting is prevalent.⁵ Second, we use the methodology of Hall, Jaffe, and Trajtenberg (2005) to estimate the market value of software patents using various patent-related measures and we explore the effect of changes in the patentability of software over time on the market value of publicly traded firms that hold such patents.⁶

In performing these analyses, we separate the software firms in the sample into those that are “upstream” and those that are “downstream,” which corresponds to the position of their products in the computing stack, from the basic operating system layer to end user application software products (Raduchel 2006). We hypothesize that firms producing software that relies on hardware or software that is upstream from their software in order to operate may face more negative consequences than other firms from the introduction of software patents (at least initially), because they are more likely to need licenses for patented software technology in order to ensure that their products will operate successfully in a particular computer environment.

The paper is organized in the following way: the next section reviews the recent history of software patenting in the United States. This is followed by a discussion of our data, which consists of a panel of publicly traded firms in the ICT sector; we also spend some time evaluating a number of definitions of “software” patents that are based on keyword searches and patent classification systems. The next two sections present our empirical results: first a series of

⁵ Based on a review of the industry sectors where firms obtain software patents described later in the paper, we identified the relevant two-digit industries as 35 (machinery including computing equipment), 36 (electrical machinery, 38 (instruments), 48 (telecommunications), and 73 (business services including software).

⁶ Noel and Schankerman (2006) perform a similar analysis of the market value effects of software patents, and obtain conclusions similar to ours. They also study the relationship between market value and measures of the cost of transacting with rival patenters, and conclude that market value is enhanced when rivals’ intellectual property is concentrated among fewer parties.

studies of the immediate stock market impact of software patentability decisions in the courts, and then the results for the relationship between firm market value and the ownership of patents. The latter results are presented for two subperiods, before and after the general acceptance of software patenting by the USPTO guidelines. The final section concludes.

2. *Background*

The law concerning the patentability of software in the United States has evolved through a series of decisions following the passage of the 1952 Patent Act (which did not exclude any subject matter statutorily), to the point where algorithms may be patented if “there is practical application for the algorithm or if it is associated with a tangible medium (Sterne and Bugaisky 2004, p. 221). Although a complete history of the use of intellectual property protection in the software industry is beyond the scope of this paper, we provide a brief summary here and refer the reader to Graham and Mowery (2003) for more detail.

In 1972, the Supreme Court’s ruling in *Gottschalk v. Benson* held that because software is essentially a collection of algorithms, it could not be patented. In 1980, Congress confirmed that the appropriate intellectual property protection for software was copyright (Samuelson 1984).⁷ However, in 1981 the court allowed for patenting of software tied to physical or mechanical processes, such as the program implemented in the method for curing rubber at issue in *Diamond v. Diehr*, a decision which seems consistent with the present day European Patent Convention.⁸

⁷ Patents and copyrights protect very different aspects of software. Copyrights are awarded to creators of “original works of authorship” and protect the specific computer code as an “original expression.” Copyright does not cover the functions performed by the code. What a software program actually *does* may be protected by a patent.

⁸ The European Patent Convention’s (EPC) treatment of patentable subject matter differs markedly from that of the United States Patent Act. Whereas in the US, the 1952 Patent Act opened the door to judicial expansions of patentable subject matter by not including any explicit limits on statutory subject matter, Article 52 of the EPC expressly excludes several categories of inventions, among them mathematical methods, schemes, rules and

However, this and subsequent decisions by the court during the next 15 years led the US Patent and Trademark Office (USPTO) to modify its position with respect to software patents, ultimately to one that allowed them even when not embodied in a physical process. As Figures 1 and 2 show, the growth of software patenting increased from a 5-10 per cent level prior to 1981 to a 15 per cent level afterwards, and then increased again in the mid-1990s, as the internet boom took off and the USPTO issued new guidelines for software patenting in response to a series of court rulings described later in the paper.

Following the *Diamond v. Diehr* decision, the software patenting standard was substantially unchanged until the early 1990s when the rise of the personal computer industry and the consequent vertical disintegration of the computing sector encouraged software and information producers to test the subject matter exclusions with respect to software again. In 1993 Compton Encyclopedia attempted to enforce a broad patent on the display of multi-media content on CD-ROMs against several firms; in response to complaints by these firms, the USPTO re-examined and revoked the patent in March 1994, which halted assertion of this particular software patent.

But then later in 1994 an important change occurred when the Court of Appeals of the Federal Circuit (CAFC) stated (*In re Alappat*) that unpatentable software was that which represented “a disembodied mathematical concept...which in essence represents nothing more than a ‘law of nature,’ ‘natural phenomenon,’ or ‘abstract idea.’” Software that could be patented

methods for performing mental acts, playing games or doing business, and programs for computers “to the extent to which a European patent application or European patent relates to such subject-matter or activities *as such*” (EPO 1989, emphasis added). In the case of software, the phrase “as such” has led to a number of software inventions actually being patented via their embodiment in hardware. See Bakels (2005) for a detailed discussion of the evolution of European policy in this area.

was “rather a specific machine to produce a useful, concrete, and tangible result.”⁹ A series of additional decisions, described in more detail in section 4 of this paper, culminated in 1995 with *In re Beauregard*, in which the CAFC ruled that the Patent Office’s rejection of a software patent application by IBM was erroneous. The Commissioner of Patents then issued new guidelines on the patenting of software, which allowed inventors to patent any software embodied in physical media (which essentially reversed the *Gottschalk* decision). (Sterne and Bugaisky 2004, p. 223) These decisions presumably lie behind the huge increases in software patenting during the 1996-1999 period shown in Figures 1 and 2.

A second important expansion of patentability took place in 1998 when Judge Rich issued the famous opinion in *State Street Bank and Trust v Signature Financial Corporation*.¹⁰ The Signature patent at issue was a “pure” number-crunching type of application, which implemented financial accounting functions. The Federal Circuit Court decision stated clearly that Section 101 (the section of the patent code that deals with subject matter for patentability) is unambiguous - “any” means ALL, and it is improper to read limitations into 101 not intended by Congress. Therefore, mathematical algorithms are non-statutory only when “disembodied” and thus lacking a useful application. The court went on to make sure that the decision was precedent-setting by stating that with regard to the business method exception, “We take this opportunity to lay this ill-conceived exception to rest.”

This decision was followed by an increase in applications for “business method” patents, most of which are arguably also software patents, because they describe the implementation of a

⁹ *In re Alappat*, 33 F.3d 1526, 1544 (Fed. Cir. 1994), cited in Sterne and Bugaisky (2004), p. 222.

¹⁰ *State Street Bank and Trust Co., Inc. v. Signature Financial Group, Inc.*, 149 F.3d 1368 (Fed. Cir. 1998). See Lerner (2002) for a fuller discussion of this case.

business method on a computer or the internet. However, they are still a small fraction of all software patents (Hall 2003), at least when we use the relevant patent class (705) to define them.¹¹ Most of the analysis in this paper uses data that ends in 1999 or earlier, so these patents will play a relatively small role.

2.1 The critique of software patentability

Much of the critique of software patents centers on the perception that such patents are often of low quality, rather than the principle that software should never be patentable. The “quality” of a patent is a somewhat ill-defined catchall term for all the characteristics that the particular analyst would like patents to have. Hall (2003) and Hall et al (2003) review these criteria and argue why the failure of issued patents to satisfy them may be costly for firms and society. These criteria include satisfying the statutory definition of a patentable invention (novelty, non-obviousness, and utility),¹² sufficient disclosure to enable those “skilled in the art” to understand the invention, and relatively little uncertainty over the validity of the patent. A number of legal and technical scholars have critiqued the software patents issued after the series of court decisions in 1994-1995 on all of these grounds.

First, because software development took place over a long period before software was patentable, when the USPTO began to handle large numbers of patents in this area in 1994-1995, they did not have examiners with the relevant training and lacked adequate data bases with software prior art. For example, on February 7, 1995, the following exchange between a patent

¹¹ While the majority of business method patents are found in this class, it should be noted that they can be found in other classes as well, such as classes as 235, 380, 707, and 709.

¹² See Lunney (2001) for an argument that the non-obviousness test has been weakened since the creation of the Federal Circuit Court of Appeals in 1982.

examiner and the editor concerning access to non-patent prior art took place in Aharonian's Patent Newsletter:¹³

The examiner: "The problem with obviousness is evidence. When an examiner uses common sense, attorneys scream hindsight. Also, a problem with ordering non-patent publications or translations of foreign documents is the time it takes to process these requests. An examiner cannot simply call a company whose making, or is believed to have made, a product which appears to infringe on a claim. At best, the examiner could ask a librarian at our library to call a company to request info, but again that takes time. With ten hours to do a case, movement is paramount."

Aharonian: "Additionally for subjects like software, the cost of purchasing copies of technical papers would exceed the application fee, so I doubt many examiners would get the authority to spend such sums. Since for most software patent applications, the most relevant prior art is non-patent materials, between the statistics I cited on citing prior non-patent prior art (an average of two out of about 30) plus the above and below comments, one could make a good case that it is impossible for the PTO to conduct adequate novelty analyses."

The result of lack of access to adequate prior art was that many poor quality software patents were issued (Barton 2001; Kasdan 1994, 1999; Lunney 2001).

Second, this sector is an extreme case of cumulative innovation, where one person's invention depends on those of many other people and the transactions costs associated with licensing a large number of patents for any particular software product may exceed the profits

¹³ Aharonian (1995). <http://www.bustpatents.com/>

attainable from its creation and therefore discourage innovation altogether. Third, some claim that the disclosure function of patents is particularly badly served by software patents, which rarely include the source code implementation and are often quite vaguely and broadly worded. Even in the case of patents in general, both Cohen et al. (2002) (using US survey data) and Arundel et al. (2002) (using European survey data) report that firms rank a number of other means of acquiring information such as customers, exhibits, conferences, journals, suppliers, competitors, and nonprofit institutions ahead of patent disclosures as technical information sources. Furthermore, in contrast to other industries, it has been argued that patents on software do not generally serve to diffuse information about the patented technology. Mann (2004) states that:

“It is clear that [disclosure] is an important benefit of patents in some industries, although the software industry in its current form is probably not one of them...As others have noted, with software cases the Federal Circuit has interpreted the disclosure requirement in such a way as to minimize the likelihood that disclosures will require information that is directly useful to competing inventors. Moreover, given the rapid pace of innovation, it will not often be the case that information disclosure in a patent application filed years earlier will be of immediate value to competitors in the industry.”

14

Finally, the growing open source movement has been extremely critical of such patents, because of the obvious fact that coexistence of open source with patented software is problematical. When independent invention is not a defense against infringement (as in the case of patents), a developer that issues a GPL or copyleft license to others on code that he has

¹⁴ Mann (2004), p. 49-50

developed cannot be sure that he has the right to do so. Thus the presence of patents on some innovations implemented via software may foreclose many avenues of open source development, unless low cost methods of licensing are developed.¹⁵

3. *Data*

To perform our study, we combine data on publicly traded firms and their market value for the 1975-2002 period with a version of the NBER patent and citations dataset that has been updated to 2002.¹⁶ The patent data are matched to Compustat data on firms in the following SIC categories: 35 (machinery including computing equipment), 36 (electrical machinery), 38 (instruments), 48 (telecommunications), and 73 (business services including software). This match was done using a version of Hall's Compustat-assignee match updated to include firms that were not included in the previous (1995) version.¹⁷ Using this updated match, we are able to identify patenting entities associated with 1,290 of the 2,379 Compustat firms in the SIC classes listed above. Our dataset is an unbalanced panel, and Appendix A3 lists the number of firms and the number of software firms in the dataset in each year of the sample. Both series have increased during the period, but software firms have clearly become more important as a share of the ICT sector.

Table 2 contains summary statistics for the dataset. In order to identify the particular software area in which our firms operate, we have also incorporated information from the

¹⁵ For more on the debate on software patents and open source, see Evans and Layne-Farrar (2004).

¹⁶ The update is available at <http://emlab.berkeley.edu/users/bhhall>

¹⁷ Details on the matching algorithm can be found at <http://emlab.berkeley.edu/users/bhhall/pat/namematch/namematch.html>

Corptech directory of technology companies on the type of software they produce (systems software, middleware, applications, or software-related services).¹⁸

3.1 Defining a software patent

One difficulty that all researchers in this area encounter is that the definition of a software patent is rather unclear (Layne-Farrar 2005). Many scholars in the area may feel that "we know one when we see one" but this is not a practical way to choose a particular set of patents out of the 3 million or so in our datasets. Although all patents are classified into a number of technology classes and simply choosing those patents that are associated with software classes might seem the desirable way to go, it is an unfortunate truth that many of the relevant classes are broad enough to contain both software and hardware patents, and some software patents end up classified in classes that do not appear to have anything to do with software at first glance.¹⁹ For this reason, different researchers have taken a number of approaches to identify software patents. In early versions of this paper, we explored the use of several different definitions and finally settled on a combination of them.²⁰

The results reported in the body of the paper were obtained using a definition of a software patent that combined three ways of defining software patents. The three definitions are that used by Graham and Mowery (2003), that used by Bessen and Hunt (2003), and one that we

¹⁸ We gather this information from Corptech's SOF product codes and the classification developed in Gao (2005), which is listed in the appendix to this paper.

¹⁹ The four US patent classes with the largest number of patents assigned to our software firms are 382 (Image Analysis), 345 (Selective Visual Display Systems), 341 (Coded Data Generation or Conversion), and 700 (Data Processing: Generic Control Systems or Specific Applications), all of which can hold hardware as well as software applications. On the other hand, there are over 1000 such patents in 435 (Molecular Biology) and over 500 in 84 (Music).

²⁰ For comparison, results using each of the definitions are found in an appendix.

constructed based on the patent classes and subclasses that contain patents assigned to fifteen of the largest software firms (which we call Hall-MacGarvie). Our combined definition was the union of the set of patents in all relevant IPC and US patent classes (the union of Graham-Mowery and Hall-MacGarvie) intersected with the set of patents found using a keyword search of title and abstract (Bessen-Hunt).

As a check, we compare the results of these automated classification systems with the sample of software and internet business method patents identified manually by John Allison via a reading of the claims and description in the patents (Allison and Lemley 2000; Allison and Tiller 2003). In the remainder of this section of the paper, we describe the definitions more fully and present some comparative statistics based on them.

Graham and Mowery (henceforth GM) identify as software patents those that fall in certain International Patent Classification (IPC) class/subclass/groups. Broadly defined, the class/subclasses are “Electric Digital Data Processing” (G06F), “Recognition of Data; Presentation of Data; Record Carriers; Handling Record Carriers” (G06K), and “Electric Communication Technique” (H04L).²¹ Graham and Mowery selected these classes after examining the patents of the six largest producers of software in the U.S. (based on 1995 revenues) between 1984 and 1995. Patents in these classes account for 57% of the patents assigned to the hundred largest firms in the software industry.²²

Bessen and Hunt (henceforth BH) define software patents as those that include the word “software,” or the words “computer” and “program,” in the description and/or specification. Patents that meet these criteria and also contain the words “semiconductor,” “chip,” “circuit,”

²¹ The groups included are G06F: 3,5,7,9,11,12,13,15; G06K: 9,15; H04L: 9.

²² Graham and Mowery, p. 232. The firms are Microsoft, Adobe, Novell, Autodesk, Intuit, and Symantec.

“circuitry” or “bus” in the title are excluded, as they are believed to refer to the technology used to execute software rather than the software itself.²³ Patents containing “antigen”, “antigenic”, or “chromatography” in the description/specification are also excluded.

Our third algorithm for defining software patents (which we label HM) is similar to that used by Graham and Mowery. We identified all the U.S. patent class-subclass combinations in which fifteen software firms (Microsoft, Adobe, Novell, Autodesk, Symantec, Macromedia, Borland, Wall Data, Phoenix, Informix, Starfish, Oracle, Veritas, RSA Security, and Peoplesoft) patented and then categorized patents falling in these class-subclass combinations as “software.”

²⁴ We refer to this definition of software patents as the Hall-MacGarvie definition.

As a check of the accuracy of these various definitions, we used a sample of manually identified software and business method patents compiled by John Allison.²⁵ Using a sample of 1000 patents issued between 1996 and 1998, Allison and Lemley (2000) identified 100 software patents by reading the claims and descriptions. Allison and Tiller (2003) augmented this sample with 230 internet business method patents, most of which are arguably members of the class of software patents. Table 1 shows the results of comparing this list with the patents selected by the three definitions, Bessen-Hunt, Graham-Mowery, and Hall-MacGarvie. From this procedure we can learn whether our samples include patents identified by Allison, although we are unable to ascertain how many patents in our sample would *not* have been labeled software by his

²³ Bessen and Hunt, p. 4.

²⁴ One complication in using the U.S. patent classification system is that patents are continually reclassified as new classes are opened up. Software has been particularly subject to this reclassification, because there were no specific software classes until the “700” classes were created after the issuance of the USPTO guidelines in 1996. We are using the 2002 classification of all the patents in our dataset, which is late enough so that software should be correctly classified.

²⁵ We are grateful to James Bessen for making these patents available to us in machine-readable form, and for updating the data to 2002.

procedure. That is, we can measure Type I error (missing a software patent we should have identified) but not Type II error (calling a patent a software patent when it is not).

The results are fairly clear: the keyword method of Bessen and Hunt is much more accurate in the sense of avoiding Type I error than either the International Patent Classification (IPC)-based classification or the US patent class-based classification, either of which identify only about half of these patents. Interestingly, Graham-Mowery is better at identifying software patents, whereas Hall-MacGarvie is much better at internet business method patents. This is probably because the US patent class system explicitly admits the existence of such patents, while the IPC does not, so they may end up classified in a wide range of IPC classes.

Layne-Farrar (2005) describes an attempt to determine the Type II error in these different definitions of software patents. She reports that when software experts read a random sample of patents from each dataset to see if they were truly for software, they classified about 5 per cent of the patents held by packaged software firms in the Graham-Mowery and Allison-Tiller sets as not software.²⁶ Layne-Farrar finds that a larger proportion of the Bessen-Hunt sample were classified as not relating to software, which would suggest that using only this definition is problematic. However, since the latter sample (BH) includes patents held by firms across all industries, the results of the comparison between the three samples are somewhat difficult to interpret.

However, because the GM definition does well on Type II error, and because GM and HM seem to select different groups of patents, whereas BH seems to be more comprehensive

²⁶ Layne-Farrar did not have access to the sample selected by our US patent class rule, because her paper was written in 2005. It should also be noted that the Allison-Tiller sample consists of quite a specific set of internet-related business method and software patents in classes 705, 707 and 709, and is not representative of software patents more generally.

(better at Type I error), our preferred definition of software patents combines these definitions in the following way: we take the union of the GM and HM samples, and then intersect this with the BH sample. On the assumption that the HM US class/subclass approach does something similar to the GM IPC class/subclass approach, this intersects a sample that should have approximately 5% type II error (that is, it rarely misidentifies patents as software that are not software) with an inclusive sample that covers almost all software patents plus many other computing patents. The result is shown in the final row of Table 1: our new definition captures over 80 per cent of the patents identified by Allison as software and internet patents, but not at the cost of being too inclusive. We therefore chose to use this combined definition as our preferred sample of software patents, and to present results using the other definitions in an appendix.

Figures 1 and 2 show the trends in software patent growth using the four definitions described above plus an aggregate software patenting series supplied by Aharonian (2001) in his email newsletter.²⁷ Figure 1 shows the absolute numbers and Figure 2 shows the share of software patents among all granted patents. The Aharonian definition appears to track that from the Bessen-Hunt keyword definition more closely than the others, whereas our combined definition is conservative and lies somewhere between the two patent class definitions (those using IPC and US patent classes). Using any of the definitions, there is substantial growth since 1976 that accelerates in around 1995-1996. The apparent jump from 1997 to 1998 using any of the definitions results from a jump in applications three years earlier, that is presumably due to the series of court decisions that are described in the next section of the paper.

²⁷ Greg Aharonian is a patent agent specializing in Information Technology sector patenting who has been tracking software patents using his own subjective evaluation of subject matter for some time. We include his figures to give some idea of the view of an observer familiar with the sector and its patenting.

4. *Event studies*

As discussed earlier, in 1994 and 1995, the Court of Appeals for the Federal Circuit (CAFC) handed down a number of decisions that affected the scope of software patents. These decisions include *In re Alappat*, 33 F.3d 1526 (1994, en banc); *In re Warmerdam*, 33 F.3d 1354 (1994); *In re Lowry*, 32 F.3d 1579 (1994), *In re Trovato*, 60 F.3d 805 (1995, en banc), and *In re Beauregard*, 53 F.3d 1583 (1995). In response to the confusion created by the court's determination that much of software was patentable, the USPTO proposed new guidelines for software patentability on May 12, 1995, and published these guidelines on March 29, 1996. In this section of the paper, we describe the decisions and then present the results of several studies of the effects of the decisions on the value of firms in the software industry or holding software patents. Table 3 presents a timeline for the decisions we consider, with the abnormal returns experienced by our firms around the date of each decision (these results are discussed in section 4.1)

In re Alappat

In 1988, Alappat filed an application for a patent on a means of smoothing the appearance of a waveform displayed by a digital oscilloscope. The PTO rejected the application's claims as non-patentable subject matter. Alappat appealed the decision to the Board of Patent Appeals, which reversed the PTO's decision on June 26, 1991. The patent examiner then requested a reconsideration of the decision by an expanded panel of the Board of Appeals, which on April 22, 1992 reversed the decision of the original panel. On July 29, 1994, the CAFC handed down a decision stating that the invention was not an "abstract idea", but rather, "a

specific machine to produce a useful, concrete, and tangible result.”²⁸ This decision was interpreted as a clear expansion of the patentability of software.²⁹

In re Warmerdam

The *Alappat* decision was followed on August 11, 1994 by *In re Warmerdam*, partly affirming and partly reversing the decisions of the Board of Patent Appeals, which had upheld the PTO’s rejection of an invention as non-patentable subject matter. The CAFC ruled that the invention, a mechanism for generating data structures for collision avoidance systems, was “nothing more than the manipulation of basic mathematical constructs, the paradigmatic ‘abstract idea’”, and sustained the rejection of five of the patent’s six claims.³⁰ However, the court affirmed one claim, implying that a machine containing in memory the data structure created by the mechanism would indeed be patentable (Huttner and Strobert 1995).

In re Lowry

Warmerdam was followed on August 26, 1994 by *In re Lowry*, which reversed the PTO’s rejection of an application on the grounds that a data structure held in memory was unpatentable under the “printed matter” doctrine. The “printed matter” doctrine provides that “an article of manufacture which consists of printed matter on a substrate cannot be statutory if it differs from other substrates only by the informational content of the printed matter except where there is

²⁸ *In re Alappat* 33 F.3d 1526, 31 USPQ2d 1545

²⁹ The decision was first mentioned in the Wall St. Journal on August 8, 1994. We also tried computing abnormal returns around this date, but the CARs were not statistically significant.

³⁰ *In re Warmerdam*, 3 F.3d 1354, 31 USPQ2d 1754

some functional interaction between the printed matter and the substrate."³¹ In this instance, the PTO had held that Lowry had not shown that the data structures were functionally related to the memory in which they were stored, but the CAFC ruled that the printed matter doctrine did not hold when the information in question was processed by a machine rather than by the human mind.

In re Trovato and In re Beauregard

In *In re Trovato* (December 19, 1994), the CAFC upheld the PTO's rejection of a software patent for containing non-statutory subject matter. The patent covered a new technique for calculating the most efficient path of an object between states in physical space. The court ruled that "[p]utting Trovato's claims in their most favorable light, the most they provide is a systemic way in which to compute a number representing the shortest path. A new way to calculate a number cannot be recognized as statutory subject matter."³² On May 12, 1995, the CAFC ruled in *In re Beauregard* that the PTO's rejection of a patent's claims under the printed matter doctrine was incorrect given the precedent set in *In re Lowry*. The patent office admitted its mistake in the case.³³

New USPTO Guidelines

In response to the above decisions, the USPTO proposed new guidelines for software patentability on May 12, 1995, and published these guidelines on March 29, 1996. The guidelines expanded on *Warmerdam*'s finding with respect to data held in the memory of a

³¹ USPTO 1995. See also *In re Gulack*, 703 F.2d 1381, 217 USPQ 401.

³² *In re Trovato*, 60 F.3d 805.

³³ *The Computer Lawyer*, Vol. 12, No. 5; Pg. 30, May 1995

machine, stating that the PTO would begin to allow claims related to software embedded in physical media. Claims should be considered processes unless they relate to some type of machine or physical apparatus. The guidelines clarified the definition of the following criteria to be used by examiners to distinguish between

- “a) a computer or other programmable apparatus controlled by software as a statutory ‘machine’;
- b) a computer-readable memory used to direct a computer, such as a memory device, a compact disc or a floppy disk as a statutory ‘article of manufacture’; and
- c) a series of steps to be performed on or with the aid of a computer as a statutory ‘process.’”³⁴

The guidelines stated that inventions that were to be considered non-statutory included data compilations or structures independent of physical elements, encoded information representing creative or artistic expression, and processes that do "nothing more than manipulate abstract ideas or concepts.”³⁵

4.1 Null hypotheses on the effects of the patentability decisions

As stated previously, the expansion of the patentability of software may have had both positive and negative effects. Firms already holding patents at the time of the decision may be positively affected due to an increased ability to exclude competitors from a market or collect licensing revenues. We will thus compare the effects of patentability on patent holders and non-patent holders as of the time of the event, with the expectation that the market reaction should be

³⁴ USPTO guidelines quoted in Huttner and Strobert (1995).

³⁵ *op. cit.*

larger for patent holders. Note that another set of firms that may benefit from software patents are start-ups for whom patents may help secure financing. Because we are focusing on the market value of publicly traded firms, we are unable to study this effect.³⁶

We also consider the potential negative effects of the decisions. For all firms, the decisions may have increased costs by creating the need to engage in licensing negotiations and by increasing the potential for hold-up. We hypothesize that the latter effect is especially relevant for firms in “downstream” market segments, i.e., firms producing software that must interact with or operate on top of other software or platforms. We thus test whether the market reacted negatively for firms specializing in applications software and software-related services relative to firms producing for “upstream” segments (middleware, systems software, and hardware).³⁷

4.2 Results of the event studies

Table 3 contains the results of the event studies that identify the market’s response to each of the above-mentioned decisions, for the entire sample of firms and for software firms only. We calculate the cumulative abnormal returns (CARs) associated with the events, and their standard errors, using the standard market model approach as described in Appendix B. To obtain the market model for each stock, we regress the stock returns for firms in our sample on the market return (using the CRSP value-weighted market index) using data from the period (-90, -30) relative to each event.³⁸ We then estimate the excess returns for the period (-1,+3)

³⁶ Cockburn and MacGarvie (2006) study the effects of patent “thickets” on entry in narrowly-defined software product areas.

³⁷ See the appendix for our classification of firms into these segments.

³⁸ Our study includes some events that are less than one month apart. We choose the 30 day lag for the window to minimize any contamination arising from overlapping events.

around the event by summing the forecasted residuals from the market model and computing their standard errors, taking account of any correlation that might be present.

For each event we show the average CAR for all firms in our sample and for four subgroups of firms: 1) firms that only produce software (firms whose primary SIC code is 73XX and who do not also produce hardware), 2) firms that specialize in applications software or software-related services, as defined by the Corptech directory or by our investigation using web searches,³⁹ 3) firms that had not applied for patents in the year prior to the event, and 4) firms that had not applied for software patents in the year prior to the event. We also performed rank tests of the difference in the CAR distribution between the each group and other firms using the methodology described in Appendix B.⁴⁰ We group the CARs in this way because we expect the effect of the decisions to differ depending on whether firms can expect to take advantage of the increased strength of software patents. We expect that firms holding software patents at the time of decisions that broaden patentable subject matter to include some types of software will be positively affected relative to firms without patents. We also expect that firms in “downstream” sectors – like applications software and software-related services – will be negatively affected relative to upstream firms in hardware and middleware or systems software, because they may be forced to license in patented technology in order to ensure that their products will function on top of middleware or the operating systems.

³⁹ As of the date of our data (2002), Compustat classified IBM into software services. Because of the importance of this firm in software patenting in general and because for most of our time period IBM was primarily a hardware firm, we have reclassified it into computing machinery. However, throughout this paper, we found that estimates with and without including observations on IBM were substantially the same.

⁴⁰ We found no significant differences for software firms versus hardware firms or application/services firms versus the others, so the results of these tests are not shown in the table.

The first event we consider is the Supreme Court's decision in *Diamond v. Diehr*. Surprisingly, this decision has a positive overall return of almost one per cent in our sample, and is also positive for all of the subsamples, although the number of software firms in the sample is too small in 1981 to yield significant results. One possible interpretation is that this development, which was a rather narrow extension of patentability to software embodied in hardware, was viewed as positive for the majority of firms then in the sample, almost 90 per cent of whom were not pure software firms.

The next two events were the assertion of the Compton patent and its subsequent revocation.⁴¹ The assertion was announced on November 16, 1993 at Comdex.⁴² Clearly the assertion was viewed as highly negative news for all ICT firms, and especially for software firms, although the difference is not significant. There was little significant reaction to the revocation of the patent, suggesting that the event was anticipated. What reaction there was is slightly negative.

The *In re Alappat* decision is widely viewed as the groundbreaking one on software patentability and this is reflected in the abnormal returns at the time of the decision.⁴³ All firms experienced negative abnormal returns, statistically significant at the 5 per cent level using the

⁴¹ Unfortunately, it is not possible to measure the reaction of the patent owner value to this event. Compton Encyclopedia had been purchased by Encyclopedia Britannica, a privately-held firm, in 1961, so no market value data is available. The patent was owned by Compton's NewMedia, which was owned by the Tribune Company at the time of its assertion.

⁴² John Eckhouse, writing in the San Francisco Chronicle the day before, reported that the patent had been kept secret since it was issued in August, but that it would be publicly announced at Comdex the next day. He also reported that "Attorneys who have seen some of the supporting documents say the patent will startle the industry because of its unusual breadth."

⁴³In reference to *In re Alappat*, Evans and Layne-Farrar (2004) state that "This ruling cemented the statutory standing of software patents." (p. 6). Mann (2004) says Alappat "cleared the way for software patents" Cohen and Lemley (2001) state that "In 1994, the en banc Federal Circuit decided *In re Alappat*, opening a new era in software patent protection"(p. 10).

rank test. The negative returns were primarily due to those firms that had not applied for software patents by 1993 (the year preceding). However, the magnitude of the effects is small, less than one half of a percent.

After *Alappat*, the next three decisions, *In re Warmerdam*, *In re Lowry*, and *In re Trovato* had little impact. The first two followed closely on *Alappat* and all three took the form of clarifying, partially affirming, and partially reversing that decision. *In re Warmerdam* was viewed as having ambiguous implications, and was seen to counteract previous extensions of software patentability while simultaneously affirming the patentability of a machine storing “data structures” in memory.⁴⁴ While this latter aspect of the decision is today viewed (with the benefit of hindsight) as a step towards patentability, contemporary press reports suggest it may have been primarily interpreted as a reversal. The headline in *The Computer Lawyer* in September 1994 was “Federal Circuit Holds Collision Avoidance System Unpatentable.”⁴⁵ *In re Lowry*, a clear expansion of patentability but one that essentially confirmed *Warmerdam* and *In re Trovato* had little impact.

In re Beauregard was a test case submitted by IBM in the wake of the *Lowry* decision, essentially looking for confirmation from the courts that the “printed matter” doctrine did not apply to software on physical media. In fact, the PTO conceded its mistake in rejecting the patent application in this case and the CAFC never rendered a decision. The result was PTO acknowledgement that a computer-readable medium qualifies as a statutory article of manufacture to the extent program code stored on the medium and executable by a processor

⁴⁴ Hollander (2003), Paragraphs 18-22.

⁴⁵ See also Stobbs (1995), who contrasts *Warmerdam* with *Lowry* by writing “*Warmerdam*'s data structure is not patentable subject matter... *Lowry*'s data structure is patentable subject matter.”

functionally defines the medium. The overall market reaction to this acknowledgement was significantly positive, but that for firms without patents and without software patents was significantly lower than for firms with patents.

On the heels of the *Beauregard* decision, the USPTO announced a set of proposed new guidelines for software patents in May 1995. The announcement was viewed by the stock market as significantly negative for all firms in ICT, and especially so for software firms. However, the final issuance of the guidelines somewhat reversed this reaction, suggesting either the resolution of uncertainty or confirmation that the expansion of patentability would not be as broad as feared.

5. *Market value and software patenting*

In this section of the paper we compare the financial market's valuation of patents identified as software-related to its valuation of patents in general, both before and after the changes in software patentability rules in 1995. Hall, Jaffe, and Trajtenberg 2005 (henceforth, HJT 2005) propose a methodology for estimating the contribution of patents and citations-weighted patents to market value (Tobin's Q). We use a variation of their methodology that is designed for the particular set of questions in which we are interested. In what follows, we first describe the econometric specification, and then turn to describing results for the overall sample followed by results for sub-samples of hardware and software, or applications/services and middleware/systems software firms.

As noted in the HJT 2005 paper, economic theory is of limited help in determining the specification of an equation that relates market value to patents and citations, and the interpretation of the resulting coefficients can be difficult. There are two distinct problems: first, the market value relationship is a (moving) equilibrium relationship rather than a causal

relationship and its coefficients measure the marginal shadow value of a particular asset at a particular point in time. Second, the asset for which patents are a proxy contributes to value in at least two ways: patents are both an indicator of the private value of innovative output and they are also intrinsically valuable, since they represent the right to exclude others from practicing the invention(s) which the innovation output incorporates. In the absence of patentability, the second source of value would not be present. However, the two sources are inextricably linked, because the private value of any particular innovation is inevitably related to the existence of the patent right. For this reason, it is not possible to construct or indeed even to imagine a pure *ceteris paribus* experiment using only market value data that would separate the two sources of value.⁴⁶

What we do here is attempt to provide some indication of the ways in which the relationship between a firm's patent holdings and its financial market valuation changed as a result of changes in the patentability of software. To do this, we make use of the fact that we have two patent measures: patent counts and a patent "quality" measure, the number of citations subsequently received by the patent.⁴⁷ By including both of these in our regression, we are able to make at least a partial distinction between the value of the patented inventions, and the value of the patent rights associated with them.

However, the separation into these two value components is not perfect for at least two reasons. First, and related to the discussion above, we might expect that the pure value of a patent right is correlated with the value of the underlying innovation. Second, measurement error

⁴⁶ The best way to measure the value of the patent right is probably to use patent renewal data, which measures the firm's willingness to pay for the patent. However, such data is unable to say much about the value of the more valuable patents, given the low level of renewal fees, and most of the value of the portfolio is in patents renewed to full term. See Schankerman and Pakes () for an example of this methodology.

⁴⁷ Harhoff et al. (1999) provide evidence from a survey of inventors that forward citations are significantly positively correlated with the value of patents. Also see Nagaoka et al. (2010) for a recent survey of the use of patent data.

in any of these variables will contaminate the coefficients of all the variables that are correlated with that variable. In the case considered here, measurement error is inevitable, given that the variables are integer counts that merely proxy for the underlying concepts. For this reason our specifications will rely on variable constructions that are approximately orthogonal, in order to minimize the biases from measurement error and correlations.⁴⁸ The details of the specification are discussed in the next section after we present the basic estimation model.

The HJT 2005 paper specifies a firm-level market-value equation that is linear and additively separable [following Griliches (1981)]. In that paper, the market value of firm i in year t was modeled as:

$$V_{it} = q_t (A_{it} + \gamma K_{it})^{\sigma} \quad (1)$$

where A_{it} stands for physical assets, and K_{it} the firm's knowledge assets. Some manipulation of this equation yields a nonlinear expression for the log of Tobin's q as a function of the knowledge asset intensity.

In this paper we chose to use a simpler first order logarithmic expansion of the log of market value for several reasons. The first was that our sample was very heterogeneous and skew in its patenting behavior, especially in its software patenting behavior, and we were anxious to mitigate the effects of outliers by using a robust functional form. The second was the presence of the software patenting variables and our desire to model them in a transparent way, one that made the results somewhat easier to interpret. Thus we started with the following simple equation for the market value of a firm:

$$\log V_{it} = f(A_{it}, K_{it}) \quad (2)$$

⁴⁸ See Hall (2004) "A Note on Measurement Error and Proxy Biases" for more on these biases.

Because the knowledge assets K are not observed directly, we model them as function of the R&D, patents, and citation-weighted patents:

$$K_{it} = g \left[R_{it}, \frac{C_{it}}{R_{it}}, \frac{SC_{it}}{R_{it}}, P_{it}, SP_{it} \right] \quad (3)$$

where R_{it} is firm i 's R&D capital stock in year t , C_{it} is a measure of citation stock in year t , P_{it} is a measure of patent stock in year t , SP_{it} is a measure of the software patent stock in year t , and SC_{it} is a measure of the software citation stock in year t .

The description of these variables is the following:

R&D stock (R) – the stock of past R&D spending

Citation stock/R&D stock (C/R) – a proxy for the success rate of the firm's R&D program, this measures innovative output relative to innovative input. The citation stock is the discounted sum of citations received by granted patents that have been applied for in the current year and prior years.

Software patent citation stock/R&D stock (SC/R) – a proxy for the quality or success of the firm's software efforts. Adding this variable to an equation that also includes C/R allows software output quality to be measured as in increment to that of other innovations.

Patent stock (P) – the stock of granted patents by application date

Software patent stock (SP) – similarly for software patents only; note that these patents are also included in P .

Substituting equation (3) into (2) and using a first order logarithmic expansion yields the following estimating equation:

$$\log(V_{it} / A_{it}) = \log Q_{it} = \log q_t + \gamma_0 \log A_{it} + \gamma_1 \log(RD_{it} / A_{it}) + \gamma_2 \log(C_{it} / RD_{it}) + \gamma_3 \log(SC_{it} / RD_{it}) + \gamma_4 \log(P_{it} / A_{it}) + \gamma_5 \log(SP_{it} / A_{it}) + \varepsilon_{it} \quad (4)$$

In this equation, the coefficient γ_2 measures the impact of (patented) innovation success controlling for the R&D that produces it and γ_3 measures the premium or discount associated

with patented software innovation success when compared with ordinary innovation. Assuming that C and SC control for the success and quality of innovation output, the coefficient γ_4 measures the additional value of having patents on this output, and γ_5 measures the premium or discount assigned by the market to software patents. Note that success in innovation that is not associated with patents will not be captured by this model and remains in the residual; this is inevitable given our lack of measures of such success.

Hall, Jaffe, and Trajtenberg (2005) provides evidence that citations to a firm's own patents (self-cites) are about twice as valuable as other citations. This result suggests that self-cites proxy for stronger appropriability of a particular technology as well as for the "quality" of that technology. That is, if a firm cites its own patents heavily, those patents have presumably been important in opening up a particular technological area to exploitation by that firm. In order to control for this possibility, we include the log of the ratio of self-cites to R&D in the regression as well.⁴⁹

When estimating the model in equation (4), we included a complete set of year dummies for each of the 5 two-digit industries in our sample (machinery, electrical machinery, instruments, telecommunications, and software).⁵⁰ The R&D, citation, and patent stock measures are constructed using the usual declining balance method with a depreciation rate of 15 per cent:

$$K_{it} = (1 - \delta)K_{i,t-1} + R_{it} \quad (5)$$

⁴⁹ Because self-cites to all patents are highly correlated with self-cites to software patents in our sample, we did not include a separate variable for software patent self cites.

⁵⁰ We include the interaction between year and industry to deal with the large increase in valuation for software firms at the end of our sample period, which was presumably due to the dotcom "bubble".

where R_{it} is R&D spending, granted patents applied for in year t , or citations to granted patents applied for in year t .⁵¹

Our data include observations with no R&D, observations with no patents, observations with no software patents, observations with no self-citations, and various combinations of these. A simple way to deal with this problem without losing observations is to include a set of dummies for each of the various combinations; unfortunately several of the cells thus generated are either zero or very small, so that we need to drop some of them. In Appendix C, we present the distribution of missing or zero observations and explain how we chose a slightly reduced set of dummies for the analysis. The Appendix also presents a set of regression tables that are similar to Tables 4 and 5, but that include only these dummies. The six dummies we ended up using are the following (the left out category is nonzero R&D, patents, and software patents): 1) has R&D but no patents; 2) has patents but no R&D; 3) neither patents nor R&D; 4) has no self-citations; 5) has R&D and patents, but no software patents; 6) has patents but no R&D and no software patents. The latter two dummies are only included in those regressions that also include the continuous software patent stock variables.

5.1 Results for all sectors

Equation (4) was estimated using ordinary least squares, and the results are reported in Table 4. Eicker-White standard errors clustered by firm were computed to ensure robustness to heteroskedasticity. There are three columns in each of the two panels in the table, one set

⁵¹ Our patent data end in 2006 and our regressions in 2002. This means that we have four years to observe the issuance of the youngest patents in the sample, which is enough time to observe most of them. However, it is not sufficient time to observe forward citations for the most recent patents, so we have adjusted the number of citations received by each patent by the ratio of the total number expected to be received by a patent in that technology class to the average number such patents have received by the time of the relevant citation lag, using a methodology described in Hall, Jaffe, and Trajtenberg (2001).

corresponding to the whole period, one set for the pre-guidelines period (1980-1994) and one for the post-1994 period (1995-2002), in order to focus on the question of whether financial markets placed a higher value on software inventions and software patents following the *Beauregard* decision and the Commissioner of Patents' issuance of new guidelines on software patentability in 1995.

The test for slope equality of the patent coefficients between the two periods rejects strongly for the specification without software patents, but not at all when we include the software patent variables. Thus the relationship between patenting and market value in the ICT sector appears to be stable across the two periods, provided we allow software patents to be valued differently (in fact, slightly more highly) from other patents. The coefficient on the log of firm assets is also significantly less negative in the second period than in the first, while the coefficient on the log of the ratio of R&D stock to assets increases slightly in both specifications.

The impact of the citation-R&D ratio (a measure of innovation output relative to innovation input) is fairly stable across the two specifications, with a decline in the coefficient after 1995 (from around 0.08 to around 0.03 in both specifications). For firms that have patents, a one standard deviation increase in log cites per R&D increases market value by approximately 12 per cent in the earlier period and 5 per cent in the later. Self-citations earn a substantial premium; however, it needs to be kept in mind that self-citations are only about 5 per cent of total citations on average, so that the net effect of increasing them is roughly the same as that given only by their own coefficient. Citations to software patents are valued the same as those to other patents, suggesting that the correlation between market value and invention quality is similar for software and ordinary patents.

Once citations (that is, the quality of patent output) are controlled for, the ratio of patent stocks to assets is negative and significant in the first period and insignificant in the second. This implies that, holding constant the quality of the innovation, filing additional patents hurt market value before the expansion of software patentability, and had no effect on value after the legal changes. Presumably, this is because some of the patents were filed with uncertain benefits prior to the legal change. In contrast, the ratio of the software patent stock to assets enters positively and significantly in the pooled sample and the second period, but is insignificant in the first period. For firms that have software patents, the net impact of a one standard deviation increase in the log of software patent stock per assets is approximately zero in the first period and 3 per cent in the second.⁵³ Thus, the negative effect of patents on market value (for a given level of invention quality) is mitigated if the patent is a software patent, and the value of both software and non-software patents rose after 1994, but neither the change in the software patent coefficient nor the coefficient itself are statistically significant.

Our F-test for differences in valuation between ordinary patents and software patents is a test of whether the coefficients of the software patent citation stock-R&D ratio and the software patent stock-assets ratio along with the two software patent dummies are nonzero. In the first period this test marginally rejects (significance level 0.09) and in the second period it rejects strongly. Although the number of firms with software patents and the number of software patents do increase between the two periods, the increase is not large enough to account for the change in significance via increased power of the test. It is more likely that the patents we have identified as software were to some extent different (and more valuable) after the legal changes

⁵³ The net effect of an increase in the log of software patents per asset is given by $\gamma_4 (SP/P) + \gamma_5$. The average share of software patents for firms that have them is 27% pre-1995 and 35% afterwards, which yields $-0.10 \cdot 0.27 + 0.03 = 0$ and $-0.06 \cdot 0.35 + 0.05 = 0.03$.

in 1995. Note that the associated inventions do not appear to be more valuable, suggesting that there may have simply been slightly improved appropriability.

Our conclusion from the results in Table 4 is that software patents were valued the same as other patents prior to 1995, but that after the change in patentability the market valued them somewhat more highly than other patents, even though the quality of the inventions themselves was not significantly different. In the next section of the paper we look at the way these findings differ across different ICT subsectors.

5.2 Results for hardware and software

In the top section of Table 5 we split the sample into software firms (5375 observations) and hardware firms (13,192 observations on machinery, electrical machinery, instruments, and telecommunications) and show the results for the same specification as in columns 4-6 of Table 4, but only for the pre-1994 period and the 1995-2002 period separately. In this table we show only the R&D and patent coefficients for simplicity, although the regressions also include the assets variable and the industry-time dummies. The bottom of Table 5 presents results from the same regressions estimated separately for middleware/systems software (1092 observations) and applications software/software-related services/internet firms (4283 observations).

These tables reveal some interesting detail obscured by the results based on aggregated data in Table 4. We note first that the valuation of patents does not change between periods for the hardware firms, while it changes significantly for the software firms, suggesting that the role of patents did indeed change for these firms after the Alappat and related court decisions. In addition, the premium for self-citations exists only for the hardware firms, and is insignificant for software. Perhaps software firms benefit as much from inventions by other firms that build on

their inventions as from their own subsequent inventions. This is consistent with the idea that software firms depend more on recombinant invention involving other firms.

However, in general the quality of inventions in software firms as proxied by citations/R&D is not significantly correlated with market value in either period, regardless of the type of citation. The software patent and citation elasticities have larger positive point estimates than for the hardware firms, but also have larger standard errors and are always insignificant. For software firms that patent the average share of patents that are software are 0.52 and 0.65 for the two periods, which implies that the net effect of the patent-assets ratio is about zero in the first period, and 0.06 in the second.

In the lower portion of Table 5, we dig deeper by separating the software category into middleware/systems software (MS) and applications/software-related services (AS). We find some differences between these two groups of firms. First, the MS group generally holds more patents, even when we measure this relative to the size of their R&D program; 41 per cent of these firms patent at least once during the period, whereas only 31 per cent of the AS firms have a patent. Second, the negative value of holding patents relative to assets before 1995 was due only to the AS firms; for the MS firms, the elasticity of patents/assets is statistically indistinguishable from zero. Unfortunately, given the large standard errors (which are due to the relatively small sample size), we cannot really conclude that patents do not affect value for these firms, merely that we cannot detect an effect.

For AS firms prior to 1994, the coefficients suggest a surprisingly large elasticity with respect to software citations per R&D dollar. A ten percent increase in the stock of citations to software patents/R&D is associated with a 3% increase in market value. So patents are correlated with value for AS firms in the early period, but only if they are “important” (highly cited) patents

– and especially if they are important software patents. This positive premium for software citations disappears after 1994, which can be interpreted in two ways: either the marginal value of a successful software invention has declined, or invention quality is proxied better by R&D in the post-1994 period.

It is perhaps worthwhile to note here that these estimates are based on the analysis of the correlation between a firm's market value and its *own* patent stock, not factoring in the impact of changes in *rivals'* patent holdings. Noel and Schankerman (2006) perform a similar market value analysis in which they also incorporate information on patent holdings by technological rivals,⁵⁸ concluding that “a 10 percent increase in rivals' patent propensity reduces a firm's value by 1.3 percent.” While our results are not directly comparable due to differences in specifications, it is interesting to compare this number with the elasticities described above, which generally find little benefit to firms of filing additional software patents. Noel and Schankerman's estimate of the negative externalities of an across-the-board doubling of patent propensities would seem to exceed our relatively modest estimates of any positive effect on firm value of such a doubling.⁵⁹

6. *Conclusions*

We obtain two sets of conclusions based on our two approaches to measuring the private value of software patents. First, we conclude that, as measured by the stock market's reaction to legal decisions expanding the patentability of software, there is some evidence that the expansion

⁵⁸ This is the sum over all rivals' patent stocks divided by R&D stocks, weighted by the rival's technological proximity to the firm using a measure based on the Jaffe (1996) technological proximity measure and the pattern of backward citations of the firm.

⁵⁹ However, Noel and Schankerman also obtain much larger positive estimates for the value of the firm's own patents, with the result that they estimate a small net *positive* impact on firm value of an across-the-board increase in patenting.

of software patentability was initially negative for firms in the software industry. There is limited evidence of relatively negative effects on some types of software producers – applications producers and providers of software-related services, and firms with no patenting experience at the time of the decisions. The most important events from the perspective of news to the market are the original assertion of the Compton multimedia patent, which affected all types of firms negatively on average, and the announcement of proposed new guidelines for software patenting.

Our second approach examines the relationship between Tobin's Q and firms' patent and citation stocks. We attempt to distinguish the quality of a patented invention from the patent right itself using information on patent citations and patent stocks normalized so as to minimize the correlation between these right-hand-side variables. We conclude that, while patenting "important" (high-cited) inventions is clearly positively correlated with market value overall, filing additional patents that do not increase the total stock of citations (i.e. increasing the patent stock holding constant invention quality), is *negatively* associated with market value. Following the expansion of software patentability in the mid-1990s, this negative relationship between patents and market value largely disappears for software firms (but not for hardware firms), but we cannot reject the hypothesis of no effect from filing additional software patents without increasing the total quality of patented inventions (i.e. the stock of patent citations). We do find, however, that firms without any software patents have lower market values on average than firms with software patents (all else equal). This suggests that while the marginal software patent filed may not be associated with much increase in market value, the decision of whether or not to patent at all seems to have been an important one.

Combining these two sets of findings, we conclude that the market evaluated software patents as a negative development *ex ante* and expected that the expansion of software

patentability would affect firms in downstream sectors and firms without patents relatively more negatively. *Ex post*, of course, a greater number of firms in all ICT sectors invested in these patents. While we obtain clear evidence that the technological importance or quality of patented innovations mattered for market value both before and after the legal changes, it is less clear that the patent rights *per se* were associated with increases in market value.

This story paints a picture of firms adjusting to a new environment and set of rules, which were initially perceived by the market as negative for the software sector, and have resulted in uncertain benefits, despite the large increases in the number of patents filed and in the share of firms filing patents in this sector. Nevertheless we should also remind the reader that the majority of the software patents acquired during the past twenty years have been acquired by non-software firms in the ICT sector (in our sample the figure is 54 per cent for software patents and 75 per cent for all patents). In addition, two-thirds of the hardware firms have at least one patent during the sample period, while only one-third of the software firms have patented. These facts suggest that the growth in these patents has been driven to a great extent by the needs of hardware firms for large patent portfolios rather than by the needs of software firms to protect their inventions.

7. References

Allison, J. R. and M. A. Lemley (2000). "Who's Patenting What? An Empirical Exploration of Patent Prosecution," *Vanderbilt Law Review* 58, 2099-2148.

Allison, J. R. and E. H. Tiller (2003). "Internet Business Method Patents," in W. M. Cohen and S. A. Merrill (eds.) *Patents in the Knowledge-Based Economy*, National Research Council, Washington: National Academies Press, 259-284.

Aharonian, G. (2005). "Critiques of Software Patent Examination," on <http://www.bustpatents.com/>

Arundel, A., J. Cobbenhagen, and N. Schall (2002). *The Acquisition and Protection of Competencies by Enterprises*. Report for EIMS project 98/180, DG Enterprise, European Union, Luxembourg.

Bakels, R. (2005). "Current IPR Policies and Law Followed by Major Organisations in the Domain," in Ghosh et al, Report to the European Commission on The Effects of Allowing Patent Claims for Computer-Implemented Inventions, August.

Bakels, R. and P. B. Hugenholtz (2002). "The Patentability of Computer Programs," A Report to the European Parliament, Amsterdam: IViR.

Barton, J. H. (2002). "Non-Obviousness," *IDEA Journal of Law and Technology* 42.

_____. (2000). "Reforming the Patent System," *Science* 287, 1933-1934.

Bessen, J. (2009). "Estimates of patent rents from firm market value," *Research Policy* 38 (10), 1604-1616.

Bessen, J. and R. M. Hunt (2003). "An Empirical Look at Software Patents", Working Paper. <http://www.researchoninnovation.org/swpat.pdf>

Blind, K., J. Edler, and R. Nack (2001). "Micro- and Macro-economic Implications of the Patentability of Software Innovations. Intellectual Property Rights in Information Technologies between Innovation and Competition," Karlsruhe/Munich: Fraunhofer Institute and Max Planck Institute.

Cockburn, I. and M. MacGarvie (2006). "Entry, Exit and Patenting in the Software Industry," Boston University: Working Paper.

Cohen, J. E. and M. A. Lemley (2001). "Patent Scope and Innovation in the Software Industry," *California Law Review* 89(1), 1-57.

Cohen, W., A. Goto, A. Nagata, R. Nelson, and J. Walsh (2002). "R&D Spillovers, Patents and the Incentives to Innovate in Japan and the United States." *Research Policy* 31,1349-1367.

Davidson, R., and J. G. MacKinnon (1993). *Estimation and Inference in Econometrics*, Oxford: Oxford University Press.

- Dreyfuss, R. C. (2001). "Examining State Street Bank: Developments in Business Method Patenting," *Computer und Recht International* 2001(1), 1-9.
- _____ (2000). "State Street or Easy Street: Is Patenting Business Methods Good for Business?" In *Intellectual Property Law*, Volume 6, Chapter 14. London: Juris Publishing Company, Ltd.
- Dvorak, John (2005). "Software Patents: Microsoft's Fatal Error," *PC magazine*, April 18, 2005.
- Eckhouse, J. (1993). "Key Patent to Shake Multimedia Industry," *San Francisco Chronicle*, November 15.
- European Commission (2002). "Proposal for a Directive of the European Parliament and of the Council on the Patentability of Computer-implemented Inventions," Brussels, Belgium.
- Evans, D. and A. Layne-Farrar (2004). "Software Patents and Open Source: The Battle Over Intellectual Property Rights," *Virginia Journal of Law and Technology* 9(10), Summer 2004.
- Gao, L. (2005). "Essays On Investments And Financing Decisions In Complementary Networks Systems", unpublished doctoral dissertation, Boston University.
- Graham, S. J. H. and D. C. Mowery (2003). "Intellectual Property Protection in the U. S. Software Industry," in W. M. Cohen and S. A. Merrill (eds.), *Patents in the Knowledge-Based Economy*, Washington, DC: National Academies Press.
- Griliches, Z. (1981). "Market Value, R&D, and Patents," *Economic Letters* 7, 183-187.
- Hall, B. H. (2005). "Exploring the Patent Explosion," *Journal of Technology Transfer* 30, 35-48.
- Hall, B. H. (2003). "Business Method Patents, Innovation, and Policy," NBER Working Paper No. W 9717.
- Hall, B. H. and Clint Cummins. (2005). *TSP 5.0 Reference Manual*, Palo Alto, CA: TSP International.
- Hall, B. H., S. J. H. Graham, S. J. H., D. Harhoff, and D. C. Mowery (2003). "Prospects for Improving U.S. Patent Quality via Post-grant Opposition," MIT Press: *Innovation Policy and the Economy* 4, 115-143.
- Hall, B. H., A. Jaffe, and M. Trajtenberg (2005). "Market Value and Patent Citations," *Rand Journal of Economics* 36, 16-38.
- Heller, M.A. and R.S. Eisenberg (1998). "Can Patents Deter Innovation? The Anticommons in Biomedical Research," *Science* 280, 698-701.
- Hollander, Andrew J. (2003). "Patenting Computer Data Structures: The Ghost, the Machine, and the Federal Circuit", *Duke Law and Technology Review*, 0033.
- Hunt, R. M. (2001). "You Can Patent That? Are Patents on Computer Programs and Business Methods Good for the New Economy?" *Philadelphia Federal Reserve Bank Business Review* 2001(Q1), 5-15.

Huttner, C. and A. Strobert (1995). "New Proposed Guidelines on Patentability Of Computer-Implemented Inventions," *New York Law Journal*, September 25, 1995.

Kasdan, J. (1999). "Obviousness and New Technologies," New York, NY: Columbia University Center for Law and Economics Studies Working Paper No. 146.

_____. (1994). "Fascinatin' Algorithm: Patent Protection for Computer Programs," Columbia University Center for Law and Economic Studies Working Paper No. 94.

Lerner, J. (2001). "Where Does State Street Lead? A First Look at Finance Patents, 1971-2000," *Journal of Finance* 57, 901-930.

Layne-Farrar, A. (2005). "Defining Software Patents: A Research Field Guide," LECG, Chicago: manuscript.

Campbell, J. Y., A. W. Lo, and A. C. MacKinlay (1997). *The Econometrics of Financial Markets*, Princeton, NJ: Princeton University Press.

Lunney, G. S., Jr. (2001). "e-Obviousness," *Michigan Telecommunications Technology Law Review* 7, 363-422.

Mann, R. (2004). "The Myth of the Software Patent Thicket", working paper.

Meurer, M. J. (2002). "Business Method Patents and Patent Floods," *Washington University Journal of Law and Policy* 8, 309.

Murmann, J. P. (2003). *Knowledge and Competitive Advantage: The Co-evolution of Firms, Technology, and National Institutions*, Cambridge: Cambridge University Press.

Nagaoka, S., K. Motohashi, and A. Goto (2010). "Patent Statistics as an Innovation Indicator," in B. H. Hall and N. Rosenberg (eds.), *Handbook of the Economics of Innovation*, Amsterdam: Elsevier Science, forthcoming.

Noel, M. and M. Schankerman (2006). "Strategic Patenting and Software Innovation," London: LSE STICERD Working Paper No. EI43.

Raduchel, W. (2006). "The Economics of Software," in D. W. Jorgenson and C. W. Wessner (eds.), *Software, Growth, and the Future of the U.S. Economy, Report of a Workshop*, Washington DC: National Academies Press.

Salinger, M. (1992). "Standard Errors in Event Studies," *The Journal of Financial and Quantitative Analysis*, Vol. 27, No. 1 (March 1992), 39-53.

Samuelson, P. (1984). "CONTU Revisited: The Case Against Copyright Protection for Computer Programs in Machine-Readable Form," *Duke Law Journal*, 663-702.

Schankerman, M., and A. Pakes (1986). "Estimates of the Value of Patent Rights in European Countries During the Post-1950 Period," *Economic Journal* 96, 1052-1076.

Scotchmer, S. (1996). "Protecting Early Innovators: Should Second Generation Products be Patentable?," *Rand Journal of Economics* 27, 322-331.

_____. (1991). "Standing on the Shoulders of Giants: Cumulative Research and the Patent Law," *Journal of Economic Perspectives* 5, 29-41.

Spindler, G. (2003). "The European Legal Framework for Software Patents," paper presented at the EPIP conference on New Challenges to the Patent System, EPO, Munich, Germany, April 24/25, 2003.

State Street Bank and Trust Co., Inc. v. Signature Financial Group, Inc., 149 F.3d 1368 (Fed. Cir. 1998).

Sterne, R. and L. Bugaisky (2004). "The Expansion of Statutory Subject Matter Under the 1952 Patent Act", *Akron Law Review* 37, 216-229.

Stobbs, G. A. (1995). "Software Practice News: The Common Thread Unraveled", *Intellectual Property Today*, p. 6.

USPTO (1995). "Proposed Guidelines for Computer-Implemented Inventions," <http://www.uspto.gov>, 1 June 1995.

Appendix A

Additional Tables and Figures

Differences between software patent definitions

We consider several different definitions of what constitutes a software patent (see p. 12-17 for a discussion of these definitions). Figures 1 and 2 plot the total number of software patents granted and the number of software patents as a percentage of all patents granted, for each of the alternative definitions.

Table A2 presents results from the market value model for the Bessen-Hunt, Graham-Mowery, Hall-MacGarvie, and “combined” (GM or HM intersected with BH) definitions. Each column presents estimates for one of the definitions, and the differences between them highlight the importance of the choice of definition. For software firms, looking at the entire sample period (1990-1999), software patents and forward citations are positively and significantly associated with Tobin’s Q across all the definitions. The magnitudes of the elasticities are also comparable across definitions.

However, there are differences between the definitions when we look at non-software firms and when we consider the change in valuation over time. Here, the BH definition diverges from the others, showing a positive and significant valuation of citation-weighted software patents for non-software firms, even in the pre-1994 period, as well as a *decrease* over time in the elasticity of market value with respect to software patents held by software firms. This may reflect the fact that the BH set of patents appears to contain a considerably greater number of hardware patents, at least according to the sample studied by Layne-Farrar.

Table A1

Category Definitions Based on CorpTech product codes (SOF codes)

Based on Gao (2005)

SYSTEMS SOFTWARE:

- A. Utility systems software 1

MIDDLEWARE SOFTWARE:

- A. Utility systems software 2
- B. Communications systems software
- C. Internet tools
- D. Software development systems
- E. Artificial intelligence software
- F. Database/file management software
- D. Software development systems

APPLICATIONS SOFTWARE:

- A. Accounting, banking, financial, government, military, legal, real estate, insurance, health services, public utilities, transportation, sales/marketing and distribution software
- B. Technical/scientific software
- C. Construction, facilities and communications management software
- D. Manufacturing software systems
- E. Media and communications software
- F. Office automation software
- G. Educational and training software
- H. Other applications software, not elsewhere specified

SERVICES:

- A. Software related services

Table A2
Elasticities from market value equations, by software/non-software and time period

	Software firms				Hardware firms			
	Bessen-Hunt	Graham-Mowery	Hall-MacGarvie	Combined	Bessen-Hunt	Graham-Mowery	Hall-MacGarvie	Combined
1980-1999 (2173 software, 8717 non-software)								
R&D/Assets	0.331	0.322	0.317	0.321	0.287	0.289	0.289	0.289
Patents/R&D	0.000	0.000	0.000	0.000	0.009	0.016	0.017	0.017
SW Pats/R&D	0.007	0.017	0.022	0.021	0.019	0.004	0.002	0.004
R&D/Assets	0.317	0.316	0.314	0.316	0.286	0.287	0.285	0.287
Citations/Patents	0.000	0.000	0.001	0.000	0.022	0.034	0.035	0.034
SW Cites/Patents	0.022	0.022	0.021	0.023	0.016	0.003	0.004	0.003
R&D/Assets	0.318	0.316	0.323	0.323	0.286	0.287	0.287	0.287
Citations/Patents (excluding self-cites)	0.000	0.000	0.007	0.008	0.022	0.034	0.036	0.032
SW cites/Patents (excluding self-cites)	0.022	0.022	0.011	0.006	0.016	0.003	0.002	0.005
1980-1994 (924 software, 5989 non-software)								
R&D/Assets	0.365	0.348	0.335	0.350	0.260	0.265	0.266	0.265
Patents/R&D	-0.023	-0.022	-0.019	-0.020	0.010	0.016	0.017	0.017
SW Pats/R&D	0.008	0.015	0.021	0.015	0.017	0.002	-0.001	0.001
R&D/Assets	0.356	0.352	0.338	0.351	0.260	0.262	0.262	0.262
Citations/Patents	-0.017	-0.017	-0.008	-0.015	0.022	0.034	0.035	0.035
SW Cites/Patents	0.021	0.023	0.019	0.021	0.014	0.002	0.001	0.001
1995-1999 (1249 software, 2728 non-software)								
R&D/Assets	0.301	0.295	0.298	0.296	0.339	0.333	0.329	0.333
Patents/R&D	0.001	0.000	0.000	0.000	0.012	0.023	0.021	0.023
SW Pats/R&D	0.013	0.025	0.020	0.026	0.022	0.018	0.021	0.017
R&D/Assets	0.293	0.292	0.297	0.293	0.338	0.336	0.332	0.336
Citations/Patents	0.001	0.001	0.002	0.001	0.028	0.036	0.034	0.036
SW Cites/Patents	0.029	0.027	0.021	0.027	0.016	0.009	0.014	0.009
R&D/Assets	0.293	0.292	0.297	0.300	0.338	0.336	0.335	0.338
Citations/Patents (excluding self-cites)	0.001	0.001	0.003	0.010	0.029	0.037	0.038	0.037
SW cites/Patents (excluding self-cites)	0.029	0.026	0.021	0.011	0.016	0.009	0.009	0.006

All regressions include a complete set of sector-year dummies and a dummy for missing R&D.

The quantities shown are the average elasticity of the market value with respect to the variable in the first column. The significance tests are based on the corresponding coefficient and its heteroskedastic-consistent standard error.

Bold text indicates significance at the 5% level; Bold italics significance at the 1% level.

Table A3: Total number of firms and number of software firms in the sample, by year

Year	Market value regressions			Event studies		
	All firms	oftware firm	App/Services	All firms	Software	App/Services
1980	406	47	23			
1981	511	76	41	184	25	17
1982	555	89	50			
1983	664	121	71			
1984	709	134	83			
1985	722	141	88			
1986	751	156	93			
1987	758	155	96			
1988	731	161	103			
1989	719	159	99			
1990	708	153	91			
1991	730	163	95			
1992	763	181	112			
1993	833	201	118	625	152	93
1994	889	216	120	706	175	95
1995	1,003	262	145	749	194	104
1996	1,065	295	144	864	238	123
1997	1,029	292	144			
1998	879	269	122			
1999	999	382	164			
Total obs	15,424	3,653	2,002			

Appendix B

Computing the CARs and Significance Tests

There are at least two ways to compute the cumulative abnormal returns to holding a share of stock. The first is that of Salinger (1992), which simply estimates a stock market beta model for each firm that includes dummies for the event days and then combines the estimates across firms. The model is

$$\begin{aligned} r_{it} &= \alpha_i + \beta_i r_{mt} + \varepsilon_{it} & t = -90, \dots, -30 \\ r_{it} &= \alpha_i + \beta_i r_{mt} + \lambda_t + \varepsilon_{it} & t = -1, \dots, 3 \end{aligned} \quad (1)$$

where r_{it} is the rate of return to security i at day t , r_{mt} is the market return, and the λ s are the excess returns for the period $(-1,+3)$ around the event. Because the event is common across firms, estimates from this set of equations using either seemingly unrelated regression on all firms or single equation estimation firm by firm will be identical. That is, there is no efficiency gain from joint estimation across firms because the right hand side variables are the same for each firm and no cross-equation constraints are imposed (Davidson and McKinnon 1993, p. 313).

Given estimates of the α s, β s, and λ s, the CAR for firm i over the period $(-1,3)$ can be estimated as follows:

$$CAR_i = \sum_{t=-1}^3 \hat{\lambda}_{it} \quad (2)$$

The standard error of this quantity (denoted s_i) must be estimated by taking account of the full covariance matrix of the estimated parameters for each firm, but this is straightforward using the delta method.⁶⁰

⁶⁰ This method is implemented by the ANALYZ command in TSP (Hall and Cummins 2005).

Given estimates of the CARs and their standard errors, the best way to compute the mean CAR across firms is to compute a weighted average, where the weights are the square root of the precision of the estimates (the inverse of the standard errors):⁶¹

$$\overline{CAR} = \frac{\sum_{i=1}^N CAR_i / s_i}{\sum_{i=1}^N 1/s_i} \quad (3)$$

Conditional on the s_i , the variance of this estimate is straightforward to compute:

$$Var(\overline{CAR}) = \frac{N}{\left(\sum_{i=1}^N 1/s_i\right)^2} \quad (4)$$

In practice, it is easy to implement this computation of the average CAR and its standard error using weighted least squares with weights equal to the inverse of the previously computed standard errors squared.

One problem with the Salinger approach is that if the event itself changes the risk of the security, a small amount of bias may be introduced by estimating the market model over the event period. An alternative that avoids this problem is to estimate the market model for a period prior to the event window and then compute the forecasted residuals for the event window period based on the estimates. In this case, assume we have estimates for α_i and β_i . The estimate of CAR can then be computed as follows:

$$CAR_i = \sum_{t=-1}^3 (r_{it} - \hat{\alpha}_i - \hat{\beta}_i r_{mt}) \quad (5)$$

⁶¹ This method is optimal if the excess return due to the event is the same for all firms. See Campbell, Lo, and MacKinlay, p. 162, for further discussion.

Once again, the standard error of this estimate is straightforward to compute using the delta method, which takes account of the error in the estimated α and β . The mean CAR for any sample can be computed as described above.

We used both methods to estimate the CARs and found that the results were qualitatively similar and for the most part also quantitatively similar. In the case of the non-parametric rank test described below, the results were identical, which suggests that the ranking of the firm's returns was unaffected by the choice of estimation method. In the text we report estimates and tests based on those computed using the pre-event market model method.

Nonparametric significance tests

We used the rank test method of Corrado (1989), which is also discussed in Campbell et al. (1997), Chapter 4. Using this testing approach is desirable in the case of stock market returns, which are frequently both non-normal and non-symmetric. It corresponds to asking whether the average of the within-firm ranks of the event period returns are approximately equal to the mean rank, as they would be under the null. The test is easily implemented using OLS regression once the ranks are available. First, we obtain estimates of excess returns for the entire period:

$$\hat{\lambda}_{it} = r_{it} - \hat{\alpha}_i - \hat{\beta}_i r_{mt} \quad t = -90, \dots, -30; -1, \dots, +3 \quad (6)$$

The estimated α and β are based on the pre-event period (-90, -30). For each period t , the returns are ranked and assigned numbers from 1 to N (where N is the number of firms in the sample), yielding a set of series R_{it} that correspond to the λ_{it} . Under the null hypothesis of no abnormal returns, the expected value of the rank is a constant by construction, in this case equal to $(66 + 1)/2 = 33.5$.

To obtain the test-statistic and its standard error given by Corrado (equations 8 and 9 of that paper), one simply regresses the ranks R_{it} on a constant and a dummy variable that is equal

to one if t is in the event period, and zero otherwise.⁶² It is straightforward to show that the estimated coefficient of the dummy and its standard error correspond precisely to the equations given in the Corrado paper. The advantage of the regression approach is that it is easy to implement further tests: for example, to test whether the shares of software firms responded differently than those of hardware firms to the event, we simply include the event dummy interacted with a dummy for being a software firm.

⁶² If desired, one can center the ranks before estimation by subtracting the mean and then there is no need for a constant. Centering the ranks would also be useful if each firm has slightly different numbers of observations (and therefore the expectation of the mean rank differs across firms).

Appendix C

Additional Regressions

In this appendix we describe the way in which we dealt with the fact that our data include many zeroes in the right hand side variables. Not all of our firms patent, and not all of our firms do R&D. In addition, there are many firms with no self-cites, or with no software patents. Because we use a log-log specification for our model, we had to include a set of dummy variables to control for observations that were zero for one or more of the R&D stock, patent citation stock, software patent citation stock, and self-citation stock.⁶³ A complete set of such dummies that included all possible interactions would number $16 = 2^4$, but many of the cells would be very small, and some are zero by construction (for example, zero patent citations but non-zero software patent citations). After dropping these cells, the number of nonzero cells was 8; we reduced this number to seven by treating the absence of self-citations separately (that is, by combining observations with and without R&D in that case). The left-out category was observations with both R&D and patents (whether or not these included software patents, so that we could test for their coefficients separately). The patterns of data presence and absence are shown in Table C.1.

⁶³ The correlation between self-citations to all patents and self-citations to software patents was so high (0.88) that we dropped the latter variable and considered only self-citations to all patents.

<i>Description</i>			<i>Hardware</i>		<i>Software</i>	
	<i>Obs</i>	<i>Share</i>	<i>Obs</i>	<i>Share</i>	<i>Obs</i>	<i>Share</i>
R&D, patents, SW pats, and self-cites	2,989	16.1%	2,304	17.7%	685	12.4%
R&D, patents, SW pats, but no self-cites	763	4.1%	346	2.7%	417	7.6%
R&D, patents, no SW pats, self-cites	3,189	17.2%	3,013	23.1%	176	3.2%
R&D, patents, no SW pats, no self-cites	2,332	12.6%	2,063	15.8%	269	4.9%
R&D but no patents	7,287	39.2%	4,374	33.5%	2,913	52.8%
no R&D, but patents & SW pats, no self-cites	79	0.4%	20	0.2%	59	1.1%
no R&D, but patents, no SW pats or self-cites	302	1.6%	230	1.8%	72	1.3%
no R&D, no patents	1,626	8.8%	701	5.4%	925	16.8%
Total	18,567		13,051		5,516	

About half the sample both does R&D and patents, and almost 40% do R&D but do not patent. The remainder do no R&D but a small number of them (2%) patent. Software firms are much less likely to patent (25% vs. 59%) and also somewhat less likely to do R&D (77% vs. 89%). Finally, over 90% of software firms have no self-citations, whereas only about half the hardware firms have no self-citations. However, adjusted for the fraction of firms that patent, the software firms are slightly more likely to cite their own patents.⁶⁴

Tables 4 and 5 in the body of the paper include a set of dummy variables for zeros in R&D and patenting as controls. When the continuous variables are also included in a regression, these dummy variables become difficult to interpret and we therefore did not show their estimates. However we thought it interesting to look at a regression that contained only the dummies, given the measurement error inherent in patent and citation counts. In this appendix we present the same regressions as in Tables 4 and 5 but with only the log of assets, the log of the R&D stock-assets variable and the R&D and patent dummy variables, along with the usual year-industry effects.

⁶⁴ Compare $0.62 = 15.6\%/25\%$ for software to $0.58 = 40.8\%/59\%$ for hardware.

In the first panel of Table C.2 we include dummies for doing R&D without patenting, patenting without doing R&D, doing neither, and having no self-cites to one's patents. In the second panel we add the two dummies for having software patents, with and without R&D. These regressions are of some interest in themselves. First, we confirm the Hall-Jaffe-Trajtenberg (2005) finding that citing one's own patents indicates value: there is a fairly robust and significant 17-21% discount for firms with no self-citations, controlling for the fact that they have patents. Some of this discount may be due to the fact that we have no citation count in the regression, so that this variable also proxies for lower citation intensity, but the premium for self-cites survives even in the presence of citations (Table 4 in the text).⁶⁵ The second finding in this table is that firms with no R&D stock face a discount of about 30%, whether or not they have patents. This discount corresponds to an R&D-assets ratio of approximately 0.2, which is roughly at the 15% point in the R&D-assets distribution.⁶⁶ That is, 15% of R&D-doing firms and 25% of all firms have a lower R&D-assets value. The implication is that firms that do no R&D have a valuation that is similar to firms that do very little R&D relative to their size; that is, there is no separate discount for these firms.

Note that doing R&D without patents earns a positive premium in the first period but not in the second, suggesting the increased importance of patents as a tool for appropriating value. If some of those patents are software patents, the premium doubles, even for those few observations with patents but no R&D.

⁶⁵ The alert reader will note that it is possible in principle for a firm to have patents but no citations. In fact, and partly because we use stocks, such observations are less than 1% of our sample and have an average patent stock less than one, so we chose to treat them as though the patent stock was zero, for simplicity.

⁶⁶ To compute the R&D intensity effect that corresponds to the dummy of -0.3, we set $-0.3 = \gamma_1 \log(RD/A) = 0.2 \log(RD/A)$ and solve for $RD/A = \exp(-1.5) = 0.22$.

Table C.3 corresponds to Table 5 in the paper, and repeats the same dummy variable exercise for hardware and software firms separately and by subperiod. Although some of the estimates are imprecise due to the small size of some of the cells, a few conclusions can be drawn. We first note that the self-citation premium is associated only with hardware firms; for software firms the presence of self-citations does not appear to add any value. Although the dummies for no R&D with or without patents appear higher for software firms than for hardware firms, because the R&D intensity of software firms is also higher, the “break-even” point is again in the neighborhood of 10-15% in the R&D-assets distribution. Finally, the strongest result is that during the second (post-Alappat) period, a software firm with software patents was valued at a premium of 48% over similar firms without software patents, whereas holding software patents was of no additional value for hardware firms except for those few that have patents but no R&D. This finding confirms that financial markets perceived software patents to be valuable for software firms after 1995 but not before.

Table C.2							
Market Valuation of Patents in ICT							
Dependent variable: log of Tobin's q							
	<i>1980-2002</i>	<i>1980-2002</i>	<i>1980-1994</i>	<i>1995-2002</i>	<i>1980-2002</i>	<i>1980-1994</i>	<i>1995-2002</i>
Observations (firms)	N	18567 (2062)	10605 (1433)	7962 (1552)	18567 (2062)	10605 (1433)	7962 (1552)
Log R&D stock/assets		0.222*** (0.014)	0.200*** (0.016)	0.249*** (0.018)	0.211*** (0.014)	0.192*** (0.016)	0.237*** (0.019)
R&D but no patents	7,287	0.090** (0.043)	0.141*** (0.048)	0.015 (0.060)	0.133*** (0.043)	0.168*** (0.048)	0.070 (0.060)
Patents but no R&D	381	-0.284*** (0.099)	-0.284** (0.118)	-0.265** (0.120)	-0.314*** (0.093)	-0.321*** (0.112)	-0.275** (0.116)
No R&D or patents	1,626	-0.349*** (0.060)	-0.278*** (0.076)	-0.421*** (0.076)	-0.300*** (0.060)	-0.246*** (0.077)	-0.360*** (0.077)
Patents and R&D,with SW pats	3,752				0.179*** (0.045)	0.143*** (0.052)	0.182*** (0.059)
Patents but no R&D,with SW pats	79				0.449* (0.260)	0.566* (0.315)	0.277 (0.283)
No self cites	12,389	-0.199*** (0.045)	-0.187*** (0.051)	-0.210*** (0.059)	-0.171*** (0.046)	-0.169*** (0.051)	-0.176*** (0.061)
Log assets		-0.073*** (0.009)	-0.100*** (0.010)	-0.040*** (0.011)	-0.084*** (0.009)	-0.110*** (0.011)	-0.051*** (0.012)
F-test for equality of pat coeffs across periods		2.31 (.056)*			1.28 (0.262)		
F-test for SW patent coefficients					9.2 (0.00)***	5.2 (0.00)***	5.0 (0.00)***
Standard error (R-squared)		0.857 (0.737)	0.793 (0.704)	0.928 (0.766)	0.855 (0.738)	0.791 (0.705)	0.926 (0.767)
The method of estimation is least squares based on equation (4) of the text.							
Standard errors robust to heteroskedasticity in parentheses. Significance at the 1%, 5%, and 10% levels indicated by ***, **, and *.							
All regressions include a complete set of year dummies for each 2-digit sic (machinery, elec. machinery, instruments, telecomms, software), as well as dummies for no R&D, no patents, no software patents, no self-cites, and their interactions.							

Table C.3
Market Valuation of Patents and Software Patents
By sector and time period

	N	1980-1994			N	1995-2002			N	1980-1994			N	1995-2002				
	Hardware								Software									
Log R&D stock/assets		0.170	0.018	***		0.256	0.025	***		0.246	0.035	***		0.219	0.030	***		
R&D but no patents	2946	0.205	0.050	***	1455	0.125	0.074	*	1585	-0.090	0.159		1328	0.125	0.110			
Patents but no R&D	173	-0.187	0.103	*	99	-0.310	0.118	***	66	-0.971	0.347	***	65	0.161	0.216			
No R&D or patents	407	-0.186	0.074	**	298	-0.479	0.092	***	434	-0.511	0.187	***	491	-0.198	0.129			
Patents and R&D,with SW pats	1392	0.133	0.052	**	1318	0.069	0.071		347	0.045	0.199		756	0.478	0.110	***		
Patents but no R&D,with SW pats	9	0.220	0.423		11	0.721	0.364	**	23	1.135	0.474	**	36	-0.033	0.313			
No self cites	5065	-0.199	0.051	***	2736	-0.213	0.074	***	2355	0.041	0.212		2300	-0.065	0.106			
Std. error (R-squared)		0.736 (0.664)				0.895 (0.708)				0.936 (0.767)				0.966 (0.827)				
F-test for equality of pat coeffs.		1.83 (.090)									2.73 (.012)**							
Observations (firms)		8046 (1028)				5146 (930)				2559 (405)				2816 (622)				
	Systems SW/ Middleware								Applications SW/Services									
Log R&D stock/assets		0.337	0.095	***		0.305	0.061	***		0.224	0.037	***		0.209	0.032	***		
R&D but no patents	354	-0.088	0.319		198	0.041	0.264		1212	-0.127	0.172		1122	0.149	0.119			
Patents but no R&D	19	-0.825	0.354	**	4	-0.634	0.249	**	32	-1.384	0.615	**	54	0.285	0.220			
No R&D or patents	97	-0.628	0.387		68	-0.115	0.329		333	-0.504	0.204	**	423	-0.191	0.139			
Patents and R&D,with SW pats	63	0.638	0.441		195	0.453	0.226	**	257	-0.078	0.214		527	0.483	0.125	***		
Patents but no R&D,with SW pats	2	1.810	0.252	***	1	2.205	0.128	***	21	1.485	0.768	*	35	-0.219	0.315			
No self cites	516	0.547	0.354		368	-0.192	0.223		1790	-0.061	0.235		1914	-0.057	0.121			
Std. error (R-squared)		0.890 (0.819)				1.00 (0.836)				0.939 (0.757)				0.956 (0.827)				
F-test for equality of pat coeffs.		1.99 (.070)*									2.57 (.018)**							
Observations (firms)		578 (103)				514 (146)				1981 (311)				2302 (508)				

The method of estimation is ordinary least squares based on equation (4) of the text.

Standard errors robust to heteroskedasticity in parentheses. Significance at the 1%, 5%, and 10% levels indicated by ***, **, and *.

All regressions include a complete set of year dummies for each 2-digit sic (machinery, electrical machinery, instruments, telecomms, software), as well as dummies for no R&D, no patents, no software patents, no self-cites, and their interactions.

Figure 1
Software Patents Granted




Figure 2
U.S. Software patents granted as share of all patents granted




Figure 3
Share of publicly-traded US firms with granted software patents
(by date of application)




Table 1
Overlap with Allison's Samples

Definition	Original software patent list (2000)		Internet business method patent list		Union of the two lists	
	Number	Share	Number	Share	Number	Share
Allison+	100	100%	230	100%	330	100%
Bessen-Hunt (BH)	92	92%	211	92%	303	92%
Graham-Mowery (GM)	60	60%	89	39%	149	45%
Hall-Macgarvie (HM)	40	40%	150	65%	190	58%
Combined*	83	83%	183	80%	266	81%

The table shows the number of patents according to each definition that are found in the lists supplied by Allison.

+The original list is the set of patents given in Allison and Lemley (2000); the internet business method list is that from Allison and Tiller (2003).

*This set of patents is the union of GM and HM intersected with BH.

Table 2
Summary statistics for 2,062 firms 1980-2001 (18,567 observations)

Variable	Mean	Standard deviation	Median	1st Quartile	3rd Quartile
Number of employees	4.17	22.63	0.36	0.12	1.28
Tobin's Q (market-to-book ratio)	6.96	11.94	3.10	1.75	6.95
Tobin's Q corr. for R&D stock	3.12	4.26	1.76	1.06	3.28
R&D stock (millions of \$96)	182.10	1279.86	10.87	2.18	44.66
R&D Stock / tangible (book) capital	1.31	1.73	0.69	0.24	1.61
D(no R&D this year)	0.08	0.28	0.00	0.00	0.00
Patents per year	12.92	107.88	0.00	0.00	1.00
Patent stock	54.45	451.34	0.46	0.00	5.42
D(no patent stock this year)	0.47	0.50	0.00	0.00	1.00
D(no software patent stock this year)	0.79	0.41	1.00	1.00	1.00
D(never patented)	0.37	0.48	0.00	0.00	1.00
SW patent stock (combined definition)	8.41	122.52	0.00	0.00	0.00
Patent stock / real R&D stock	0.35	1.96	0.00	0.00	0.25
Cite-weighted patents per year*	273.78	2347.61	0.00	0.00	18.44
Cite-weighted patent stock*	1230.08	10709.66	2.54	0.00	106.90
Cite-wtd SW patent stock (combined)*	257.99	3291.86	0.00	0.00	0.00
Cite-wtd patent stock / real R&D stock*	6.71	31.81	0.00	0.00	4.51
Cite-weighted patents per year**	227.13	1900.34	0.00	0.00	16.61
Cite-weighted patent stock**	1033.94	8777.80	2.45	0.00	100.45
Cite-wtd SW patent stock (combined)**	219.32	2622.03	0.00	0.00	0.00
Cite-wtd patent stock / real R&D stock**	6.07	29.03	0.00	0.00	4.20
<i>Hardware firms only (11,323 obs)</i>					
R&D stock (millions of \$96)	84.18	404.17	9.51	2.21	36.10
R&D Stock / tangible (book) capital	1.01	1.31	0.59	0.24	1.21
SW patent stock (combined definition)	1.77	17.55	0.00	0.00	0.00
Patent stock / real R&D stock	0.47	2.42	0.06	0.00	0.38
Cite-wtd SW patent stock (combined)*	61.62	652.72	0.00	0.00	0.00
Cite-wtd SW patent stock (combined)**	56.53	593.46	0.00	0.00	0.00
<i>Software firms only (7,244 obs)</i>					
R&D stock (millions of \$96)	335.16	1976.11	14.32	2.11	62.83
R&D Stock / tangible (book) capital	1.78	2.15	0.95	0.25	2.49
SW patent stock (combined definition)	18.80	194.47	0.00	0.00	0.53
Patent stock / real R&D stock	0.16	0.82	0.00	0.00	0.08
Cite-wtd SW patent stock (combined)*	564.94	5191.96	0.00	0.00	7.02
Cite-wtd SW patent stock (combined)**	473.79	4118.98	0.00	0.00	6.65
<i>Applications software and software services firms only (5,741 obs)</i>					
R&D stock (millions of \$96)	313.46	1986.48	13.20	1.92	61.40
R&D Stock / tangible (book) capital	1.74	2.12	0.94	0.21	2.44
SW patent stock (combined definition)	19.35	212.73	0.00	0.00	0.27
Patent stock / real R&D stock	0.17	0.90	0.00	0.00	0.07
Cite-wtd SW patent stock (combined)*	577.24	5710.61	0.00	0.00	2.12
Cite-wtd SW patent stock (combined)**	477.61	4510.65	0.00	0.00	2.12

*Corrected for citation truncation after 2006

**Corrected for citation truncation after 2006 and self-citations excluded.

Table 3
Excess Returns [CAR(-1,+3)] to ICT Firms around Various Software Patenting Events

Event	Date	Description	CAR (-1,+3) in %					Rank test for differences in CAR	
			All firms	Software firms	App SW/ SW services only	firms without patents	firms without SW patents	without patents vs with	without SW pats vs with
Diamond v Diehr	3-Mar-81	patenting allowed for software embodied in rubber-curing process (in a physical process)	0.83***	1.26	1.40	1.40***	0.97***	0.91	1.40
Compton patent	16-Nov-93	assertion of multi-media display patent held by Compton	-1.65***	-2.37***	-2.47***	-1.79***	-1.76***	-0.37	-0.54
Compton revocation	28-Mar-94	Compton's multi-media patent re-examined by USPTO and revoked	-0.36	-0.41	-0.73	-0.30	-0.56	0.60	-0.54
In re Alappat (5440676)	29-Jul-94	33 F.3d 1526 (Fed. Cir. 1994) (en banc). Narrowing mathematical algorithm limitation on patentable subject matter and thereby expanding patentable subject matter. - oscilloscope software	-0.26**	-0.02	0.25	-0.05**	-0.32**	0.65	-1.46
In re Warmerdam (6089742)	11-Aug-94	partly affirmed, partly reversed PTO's rejection of a data structure for collision avoidance systems as non-patentable subject matter - machine containing the data structure created by the mechanism patentable.	0.08	0.34	0.32	0.05	-0.01	-0.86	-1.39
In re Lowry	26-Aug-94	data structures are patentable if held in a machine (printed matter doctrine does not apply) - confirming Warmerdam	0.16	-0.50	-0.69	0.36	0.23	1.35	0.86
In re Trovato	19-Dec-94	A new way to calculate a number cannot be recognized as statutory subject matter - countered growing trend of SW patentability	-0.12	0.25	0.19	-0.04	-0.05	0.33	1.09
In re Beauregard (5710578)	12-May-95	IBM appealed a decision to exclude using printed matter doctrine, PTO admits mistake (had contradicted <i>in re Lowry</i>)	0.76***	0.90**	-0.40	0.32	0.47***	-2.44***	-3.06***
USPTO new guidelines proposed	30-May-95	new guidelines proposed to allow patenting on software embodied in any hardware medium.	-1.00***	-1.43***	-1.04**	-0.85**	-0.86***	0.33	1.47
USPTO new guidelines final	29-Mar-96	new guidelines issued allow patenting on software embodied in any hardware medium.	0.36**	-0.61	-0.32	0.35*	0.38*	-0.56	0.01

Each cell in the first panel shows the average Cumulative Abnormal Return in per cent for a 5 day window (-1,+3) around the event date.

The cells in the second panel show the result of a rank test for differences in the distribution of the CARs during the event period.

***, **, and * denote significance at the 1, 5, and 10 per cent level respectively, using a rank test.

Table 4**Market Valuation of Patents in ICT**

	Dependent variable: log of Tobin's q					
	<i>1980-2002</i>	<i>1980-1994</i>	<i>1995-2002</i>	<i>1980-2002</i>	<i>1980-1994</i>	<i>1995-2002</i>
Observations (firms)	18567 (2062)	10605 (1433)	7962 (1552)	18567 (2062)	10605 (1433)	7962 (1552)
Log R&D stock/assets	0.250*** (0.015)	0.238*** (0.018)	0.269*** (0.020)	0.238*** (0.015)	0.230*** (0.018)	0.256*** (0.021)
Log Cite stock/R&D stock	0.049** (0.021)	0.086*** (0.024)	0.034 (0.028)	0.049** (0.022)	0.082*** (0.025)	0.029 (0.030)
Log Self Cite stock/R&D stock	0.088*** (0.015)	0.077*** (0.017)	0.090*** (0.020)	0.092*** (0.015)	0.078*** (0.018)	0.094*** (0.020)
Log SW Cite stock/R&D stock				-0.011 (0.027)	-0.004 (0.037)	0.003 (0.032)
Log Patent stock/assets	-0.058*** (0.017)	-0.097*** (0.021)	-0.034 (0.023)	-0.079*** (0.019)	-0.102*** (0.023)	-0.058** (0.025)
Log SW patent stock/assets				0.059*** (0.022)	0.030 (0.033)	0.050* (0.028)
Log assets	-0.075*** (0.009)	-0.102*** (0.010)	-0.043*** (0.011)	-0.084*** (0.010)	-0.109*** (0.011)	-0.053*** (0.012)
F-test for equality of pat coeffs across periods		2.70 (.009)***			1.30 (0.216)	
F-test for SW patent coefficients				5.8 (0.00)***	2.0 (0.09)*	3.6 (0.01)***
Standard error (R-squared)	0.850 (0.741)	0.786 (0.710)	0.921 (0.769)	0.847 (0.743)	0.784 (0.710)	0.918 (0.771)

The method of estimation is least squares based on equation (4) of the text.

Standard errors robust to heteroskedasticity in parentheses. Significance at the 1%, 5%, and 10% levels indicated by ***, **, and *.

All regressions include a complete set of year dummies for each 2-digit sic (machinery, elec. machinery, instruments, telecomms, software), as well as dummies for no R&D, no patents, no software patents, no self-cites, and their interactions.

Table 5
Market Valuation of Patents and Software Patents
By sector and time period

	1980-1994			1995-2002			1980-1994			1995-2002		
	Hardware						Software					
Log R&D stock/assets	0.213	0.021	***	0.296	0.027	***	0.256	0.036	***	0.220	0.032	***
Log Cite stock/R&D stock	0.075	0.026	***	0.053	0.034		0.095	0.084		-0.053	0.065	
Log Self Cite stock/R&D stock	0.082	0.018	***	0.104	0.022	***	-0.021	0.072		0.011	0.049	
Log Patent stock/assets	-0.083	0.024	***	-0.073	0.030	**	-0.146	0.062	**	-0.029	0.049	
Log SW Cite stock/R&D stock	-0.029	0.032		0.003	0.039		0.186	0.130		0.078	0.059	
Log SW patent stock/assets	0.028	0.032		0.025	0.032		0.052	0.102		0.080	0.053	
Std. error (R-squared)	0.728 (0.671)			0.883 (0.716)			0.929 (0.771)			0.962 (0.829)		
F-test for equality of pat coeffs.	1.40 (.166)						1.94 (.031)**					
Observations (firms)	8046 (1028)			5146 (930)			2559 (405)			2816 (622)		
	1980-1994			1995-2002			1980-1994			1995-2002		
	Systems SW/ Middleware						Applications SW/Services					
Log R&D stock/assets	0.313	0.101	***	0.338	0.064	***	0.237	0.038	***	0.206	0.034	***
Log Cite stock/R&D stock	-0.106	0.207		-0.025	0.150		0.113	0.083		-0.064	0.070	
Log Self Cite stock/R&D stock	-0.064	0.136		0.122	0.091		0.046	0.078		-0.011	0.057	
Log Patent stock/assets	-0.075	0.180		-0.108	0.142		-0.204	0.077	***	-0.019	0.053	
Log SW Cite stock/R&D stock	-0.150	0.222		0.101	0.138		0.279	0.132	**	0.080	0.065	
Log SW patent stock/assets	0.177	0.192		0.010	0.149		-0.012	0.099		0.088	0.058	
Std. error (R-squared)	0.888 (0.821)			0.999 (0.840)			0.926 (0.765)			0.952 (0.828)		
F-test for equality of pat coeffs.	2.17 (.018)**						3.58 (.000)***					
Observations (firms)	578 (103)			514 (146)			1981 (311)			2302 (508)		

The method of estimation is ordinary least squares based on equation (4) of the text.

Standard errors robust to heteroskedasticity in parentheses. Significance at the 1%, 5%, and 10% levels indicated by ***, **, and *.

All regressions include a complete set of year dummies for each 2-digit sic (machinery, elec. machinery, instruments, telecomms, software), as well as dummies for no R&D, no patents, no software patents, no self-cites, and their interactions.