



WORKING PAPERS

RESEARCH DEPARTMENT

**WORKING PAPER NO. 07-6
ECONOMICS AND THE DESIGN OF PATENT SYSTEMS**

Robert M. Hunt*

Federal Reserve Bank of Philadelphia

First Draft: September 2006

This Draft: March 2007

RESEARCH DEPARTMENT, FEDERAL RESERVE BANK OF PHILADELPHIA

Ten Independence Mall, Philadelphia, PA 19106-1574 • www.philadelphiafed.org/econ/index.html

WORKING PAPER NO. 07-6
ECONOMICS AND THE DESIGN OF PATENT SYSTEMS

Robert M. Hunt*

Federal Reserve Bank of Philadelphia

First Draft: September, 2006

This Draft: March, 2007

* Research Department, Federal Reserve Bank of Philadelphia, 10 Independence Mall, Philadelphia, PA 19106. e-mail: bob.hunt@phil.frb.org

The ideas presented here have benefited from discussions with and the research of Robert Barr, John Barton, James Bessen, Wesley Cohen, Dominique Guellec, Bronwyn Hall, Brian Kahin, Mark Lemley, Robert Merges, Leonard Nakamura, Cecil Quillen, Arti Rai, Suzanne Scotchmer, and Rosemarie Ziedonis. Annette Fratantaro provided excellent research assistance. The opinions expressed, as well as any errors, are my own, and not those of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

This paper is available free of charge at www.Philadelphia-fed.org/econ/wps/index.html.

ECONOMICS AND THE DESIGN OF PATENT SYSTEMS

Robert M. Hunt

Federal Reserve Bank of Philadelphia

Abstract:

I use intuition derived from several of my research papers to make three points. First, in the absence of a common law balancing test, application of uniform patentability criteria favors some industries over others. Policymakers must decide the optimal tradeoff across industries. Second, if patent rights are not closely related to the underlying inventions, more patenting may reduce R&D in industries that are both R&D and patent intensive. Third, for reasons largely unrelated to intellectual property, the U.S. private innovation system has become far more decentralized than it was a generation ago. It is reasonable to inquire whether a patent system that worked well in an era of more centralized innovation functions as well for the more decentralized environment of today.

JEL Codes: L1, O31, O34

Keywords: Patents, Nonobviousness, Patent thickets, R&D, Entry barriers

Introduction

Economists resort to theory as a way of sorting out the effects of different phenomena on economic activity we actually observe. In the case of the topic before us, we observe firms making substantial investments in research and development, and we have in place very complicated, and sometimes costly, systems of intellectual property rights to encourage such investment. The patterns and institutions we observe are so complicated that we have a difficult time identifying cause and effect, and so we turn to economic models as a way of seeing these things more clearly.

One hopes to use the right models, just like the right pair of eyeglasses, to improve one's vision, not impair it. Of course, there is plenty of room to argue about the right ways to model this phenomenon, and there is no shortage of economists to take up that space. I will set forth, in words only, a few such models and describe their implications and leave it to the reader to decide whether these clarify or muddy our understanding.

I wish to make three main points. First, a uniform set of patentability standards will not be optimal for most industries in the sense that different standards might well lead to more innovation in some industries and less in others. In that case, policymakers face a tradeoff. Patentability standards should be set to maximize innovations in those industries that generate the greatest gains for the economy. If it is optimal to maximize the rate of innovation in industries where R&D is the most productive, patentability standards should be set relatively high.

Second, patents are sometimes more than the literal right to exclude others from producing the underlying invention. Once we recognize that there can be some slippage between inventions and the property rights granted by the government, it is possible that *more* patents can result in *less* R&D, particularly in industries that are both R&D and patent intensive. To mitigate such a perverse outcome, patent procedures should be refined to tighten the relationship between patents and the underlying inventions, or patent claims should be narrowed. If these are not possible, the cost of obtaining patents, pecuniary or otherwise, should be increased.

Third, over the last quarter century the private U.S. innovation system has become far more decentralized and for reasons that may have little to do with intellectual property rights. In that sense, we have more diversity in R&D than in the past, but we also have much more fragmented property rights. In addition, increased turnover within industries and more asymmetry between firms may complicate the ability of private agents to contract efficiently for access to state-of-the-art technologies. Is a patent system that functioned well in an era with a much more concentrated innovation system equally effective for conditions as they are today? This is, I think, an open question for researchers.

Patent Standards as Industrial Policy

Most economic models of patents suggest there is an "optimal" instrument in the sense that it balances two or more tradeoffs. The most typical tradeoff is between the need to provide positive profits to an innovator to recover the cost of R&D and the distortions

associated with pricing the good or service above the marginal social cost of production.¹ We can write down more and more complicated versions of the story, but so long as extending ever larger benefits to the innovator eventually results in diminishing returns in terms of inventive output, the story is essentially unchanged.

The design of the “optimal” patent, then, will depend on a number of parameters that describe consumers’ demand for the new goods and services, characteristics of the R&D technology, market structure, and other factors. The design might be characterized in terms of patent length, the breadth of its claims, the ability to claim infringement by modest incremental improvements (what economists call leading breadth), or even the standard that determines which inventions can be patented (e.g., non-obviousness, or what I will call the inventive step).

The crucial point is that if industries can be characterized in terms of differences in these parameters, the optimal system of intellectual property will vary systematically across industries.

In my 2004 article, I consider an environment in which firms engage in R&D competition over many periods—in fact, forever. Thus, firms engage in cumulative innovation.² Assume that patents are essential—that is, without protection, innovations are imitated at no cost at all. To simplify things, we can characterize intellectual property simply in terms of the minimum inventive step required to qualify for protection. To be concrete, innovations come in all sizes, and the size of the inventive step determines which inventions will be protected and which will fall immediately into the public domain. The higher the inventive step, the smaller is the proportion of inventions that will qualify for protection. Suppose that a patentable invention earns rents only until the next patentable invention comes along. Thus, the *economic* life of a patent depends on the rate of innovation in an industry and the standard of patentability.³ Finally, suppose there is free entry into R&D competition. In other words, if firms anticipate making enough money, they will sink resources into an R&D facility and compete for the expected rewards.⁴

This model teaches three lessons—one is obvious; the others less so. The first is that in industries where R&D is more productive, there is more innovation and more firms are willing to enter into R&D competition. Second, for any given industry described by a number of characteristics, including the productivity of R&D, there is a unique inventive step that maximizes the rate of innovation by maximizing the entry of firms. In this context, this standard is an “optimal” patent. Third, the *more* productive an industry is in terms of its R&D, the *higher* is the optimal inventive step. Thus, in order to

¹ See, for example, Nordhaus (1969).

² See the seminal papers on cumulative innovation by Green and Scotchmer (1995) and Scotchmer (1996).

³ This is a simple way to abstract from issues of infringement and licensing. O’Donoghue (1998) considers both factors in a model similar to the one presented here. From his analysis, it is clear that incorporating them only reinforces the intuition presented here.

⁴ My 1999 working paper considers the opposite case, where the number of firms is fixed, and finds many similar results. Moschini and Yerokhin (2006) adapt this model to conduct a very interesting study of prior user rights.

maximize the rate of innovation in industries where R&D is highly productive, a relatively small share of innovations should qualify for protection.

The key insight from the model is that the rate of cumulative innovation depends on the degree of business stealing across different generations of innovation. This margin is affected by patent standards. Granting patents for too many innovations can dissipate the economic value conveyed by them, reducing the incentive to innovate. This is especially pronounced in industries where R&D is more productive because the profits dissipated would have been earned in the very near future rather than far into the future.

Now suppose we can specify only a single inventive step for the entire economy. If we characterize industries by differences in the productivity of their R&D, it is clear that any given standard will be closer to the optimal standard for some industries than for others. As a result, there will be more entry into those industries and more entry than in the other industries.

It is in exactly this sense that I argue that a common patent standard applied to all industries is a form of dynamic industrial policy. If one wishes to favor industries where R&D is not very productive, the inventive step should be set relatively low. The price of this policy is much less innovation in industries where R&D is more productive, and this is unlikely to be the best policy for an economy whose competitiveness is based on rapid innovation. On the other hand, if one wishes to favor industries where R&D is most productive, the inventive step should be set relatively high. Even though a relatively small share of innovations would qualify for protection, innovation in those industries would be more rapid, and this would likely contribute to a more rapid rate of growth.

Now an alternative interpretation to consider is that the optimal inventive step involves balancing two competing factors: the instantaneous benefit of patenting the marginal innovation versus the dynamic losses that result from increased business stealing. In that case, one can think in terms of applying this balancing of factors to every industry, implicitly allowing for a different inventive step in every industry. This might be interpreted as a common law standard of patentability. Under this interpretation, the role of the courts is much different than under the previous interpretation.

What If Inventions and Patents Are Distinct Objects?

Most economic models of the patent system assume that a patent is a monopoly on the right to produce the thing that was invented. It is thus customary to think of inventions and patents as *complementary* inputs in the production of profits for the firm. In the language of economists, this means that a decrease in the cost of either input would result in an increase in both inventing and patenting. This, in turn, would increase firm profits.

In my 2006 article I consider the possibility that inventing and patenting need not be so closely related as typically assumed in the literature. I then ask a simple question: “When is it the case that inventions and patents can be *substitute* inputs in the production of firm profits?” In other words, is it possible that a decline in the cost of obtaining patents, which would clearly induce more patenting, could result in a *decline* in the invention rate? The answer is yes, under a specific set of conditions made clear by the model. When those conditions are satisfied, firms will devote too many resources to

patenting in order to ‘tax’ each other’s inventions and too few resources to the process of inventing itself.

Three key ingredients are necessary for such a perverse outcome to occur. First, firms must do a sufficient amount of R&D. In other words, the perverse outcome is not an explanation for why some firms or an industry doesn’t do R&D. But it may explain why a fairly R&D-intensive industry is doing less R&D than one might expect, given its productivity and other fundamentals, or relative to the past. Second, patents must be cheap relative to both the cost of doing R&D and in terms of the value of final output in the industry. Finally, there must be sufficient technological overlap between firms so that there is a good chance that one firm’s invention may infringe a patent obtained by another.

The first two conditions describe the characteristics of an industry and the resource cost of obtaining patents. The third requires some additional explanation and interpretation. For some industries it may simply be a question of technology. Firms may draw from similar technical fields and arrive at similar solutions even when they apply them to different problems. This is particularly true for industries that advance through cumulative innovation and where firms may rely on a largely common set of building blocks derived from previous innovations. In addition, some products incorporate several, if not dozens, of potentially patentable innovations. Two examples that come to mind are semiconductors and computer software.

The degree of overlap might also depend on the breadth of claims contained in patents themselves.⁵ If broad claims are regularly granted, it is more likely that firms will infringe each other’s patents. Under this interpretation it is possible that this degree of overlap can be regulated at the patent office, in the courts, or, in the last instance, by policymakers.

A third, and more controversial, interpretation is that firms have some facility in obtaining property rights over things they have not yet invented. While this is prohibited by U.S. patent law, it might nevertheless arise from mistakes in the examination of patent applications. This is a topic that has received considerable attention in academic and policy circles in recent years.⁶

This may be a particular problem for patents on computer programs, especially ones that implement methods of doing business, if the patent law’s disclosure requirements are not adequately enforced.⁷ In that case, one might not be certain what the applicant has invented and how far his or her claims should extend. In these areas, some

⁵ The tradeoff between patent breadth and length is studied in Gilbert and Shapiro (1990) and Klemperer (1990).

⁶ See, for example, Jaffe and Lerner (2004), Federal Trade Commission (2003), and Merrill, Levin, and Myers (2004).

⁷ See 35 USC 112. An applicant must describe his or her invention in sufficient detail that a person skilled in the relevant art can replicate it without undue experimentation.

researchers and practitioners worry that applicants can obtain relatively broad patents even though they have not really started their R&D.⁸

The model suggests that the perverse outcome is more likely to occur in high-tech industries that advance rapidly via cumulative innovation (and which also obtain many patents) than in industries that are not research intensive or which do not build up large patent portfolios. Previous empirical work has identified a number of industries with such characteristics, including electronics, computers, and semiconductors.⁹ These industries account for most of the rapid growth in U.S. patenting in recent years.¹⁰ They are also the industries where researchers identify what is sometimes called “strategic patent” behavior, including the assembly of large portfolios for wholesale cross-licensing and, in some instances, deceptive patent prosecution.¹¹

In our 2004 working paper, Jim Bessen and I present some empirical results consistent with the phenomenon modeled here in the context of patenting computer software.¹² Obtaining such patents was difficult, but not impossible, during the 1970s and early 1980s. Over time, however, courts became more receptive to such patents, and their numbers have grown rapidly, especially among firms in the industries described above (and much less so among firms in the software industry). All else equal, firms that concentrated on obtaining software patents experienced a statistically and economically significant decline in their R&D intensity relative to other firms.

If there are instances where patents substitute for R&D, what is appropriate policy response? It is clearly most important to modify the patent process to ensure there is indeed a close relationship between what a firm invents and the property rights it subsequently obtains from the government. As described earlier, this may involve modifications to patent law’s disclosure and enablement requirements. It might also involve new re-examination procedures or some process of objecting to recent patent grants. It could also involve a modification to the presumption of patent validity accorded in trials so that courts could engage in a more reasonable assessment of the quality of a patent’s examination.

If the breadth of claims is the issue, granting more narrow patents would likely help even if this resulted in more patents being granted.¹³ But, as noted earlier, overlapping claims may be an unavoidable outcome for some technologies or for some industries where innovation is cumulative in nature.

Finally, if these other avenues are not effective, another option is to *raise* the cost of obtaining patents. This might involve higher pecuniary charges, but it might also

⁸ See Burk (2002), Burk and Lemley (2002), Flynn (2001), and chapter 4 of FTC (2003).

⁹ See, for example, Levin et al., (1987), Cohen, Nelson, and Walsh (2000), and Hall and Ziedonis (2001).

¹⁰ This is documented in Hall (2003) and Bessen and Hunt (2004a).

¹¹ See, in particular, Grindley and Teece (1997) and Graham and Mowery (2004).

¹² A description of the findings for a more general audience is found in Bessen and Hunt (2004b).

¹³ This result is established in Proposition 4 of the working paper version of the paper..

involve raising the standards used to determine when an invention is eligible for patent protection. In other words, a higher inventive step might also mitigate the problem.

Who Is Doing All This R&D Anyway?

I have already described a relationship between patent standards, the number of firms engaged in R&D, and the overall pace of innovation. In more recent work, Leonard Nakamura and I return to the question of exactly who is doing R&D and why. We document the very significant change in the distribution of private R&D investments across the U.S. economy and attempt to sort through a number of competing explanations for these changes. One conclusion from this work is that the U.S. innovation system has become vastly more decentralized than in the past, and we think for reasons largely unrelated to the patent system. An important question, then, is whether our current system of intellectual property rights and the methods for enforcing them are well suited to this new environment.

But first, we present some facts. One can divide the U.S. economy into industries that are R&D intensive (say R&D is at least 1% of sales) and those that are not. One can apply this definition to firms in 1973 (when accounting rules required the disclosure of material amounts of R&D) and ask whether the division of industries remains accurate 25 years later. The answer, somewhat surprisingly, is yes (see Figure 1). Nearly all private R&D continues to occur among the R&D-intensive industries we observed back in 1973. If anything, R&D is more concentrated in these industries, but that is because the R&D intensity of these industries has risen significantly over time (see Figure 2).¹⁴

But one can also examine the distribution of R&D within industries, and this too changed dramatically after about 1980. Nakamura and I focus on a particular set of large, established firms, about 65 in number, which accounted for the majority of all private R&D around 1973; we call these firms the incumbents. As Figure 3 shows, these firms accounted for about the same share of private R&D until about 1980; thereafter, their share fell to about 35% of the total. Similar tabulations, based on the size distributions of public and private firms in the National Science Foundation survey of industrial R&D, tell the same story (see Figure 4). Thus the decline of particular firms in our list of incumbents does not adequately explain the trend.

It is not the case that large, established firms are doing less R&D than in the past. Indeed, on average, they are doing more R&D, and their R&D intensity has grown over time. But the R&D intensity of smaller, younger firms—which we call the entrants—has grown even more dramatically (see Figure 2).¹⁵ Indeed, it is the rise in R&D intensity

¹⁴ Both Tables 1 and 2 are derived from Standard & Poor's Compustat. We report R&D intensity as R&D divided by the operating expenses of firms because we do not wish to exclude from these calculations firms engaged in extensive R&D but not yet producing significant revenues.

¹⁵ The differences in the height among the first set of bars is statistically significant—in other words it is clear that incumbent firms used to be more R&D intensive than other firms. The differences in the height among the last set of bars is not statistically significant—in other words we cannot be sure there are significant differences in the R&D intensity of these groups of firms. Nevertheless, we can conclude that younger and smaller firms have at least caught up with the incumbent firms in terms of the research intensity.

among this group of firms that explains nearly all of the rise in the R&D intensity of the U.S. economy after 1980 (this is also evident from Figure 4).

Now what explains these patterns? In the paper, we develop a simple model that permits us to sort among a number of competing explanations by examining differences in their predictions about trends in the observable data. Setting aside all the details, the crucial point is that the patterns in the data are best explained by a decline in fixed costs firms incur *after* the success of their R&D projects. These are the costs firms pay to reach a national or international market for their goods and services. In the paper, we focus on the decline in the costs of administration, distribution, communication, and other costs of coordination that resulted as computing power became both scalable and ubiquitous. In the paper, we describe these as investments in *marketing capital*. A decline in the cost of marketing capital leads to an increase in R&D competition among both smaller and large, established firms. All else equal, this increased competition reduced the market value of the incumbent firms and increased the value of the entrant firms.

Why is this digression relevant to our discussions about patent policy? There are two important points to make. First, the trends we observe in R&D and other observable measures could be explained by dynamics other than patent policy. For example, in some circles the overall rise in R&D intensity of the U.S. economy is taken as evidence of the success of Congress and the Court of Appeals for the Federal Circuit in improving our patent system. But without controlling for these other dynamics, it is extremely difficult to establish a causal argument about the effect of these earlier changes in patent policy and aggregate measures of investment in R&D.

The second point is that, for reasons I think are largely independent of patent policy, the private U.S. innovation system has undergone a dramatic structural change. It is entirely reasonable to inquire whether our patent system serves this environment as well as it did in the era of much more concentrated R&D activity some 40 years ago. We also have more independent research programs pursuing solutions to similar technical problems. So in that sense, we have more diversity in R&D programs than in the past. But it is almost certainly true that some of these programs overlap and so will their results.

But one could also argue the democratization of R&D is an outgrowth of the “strengthening” of intellectual property rights in the U.S. after 1982. The argument is that increased certainty about enforcement and enhanced penalties for infringement facilitates entry because firms are better able to protect their intangible investments, obtain financing, and license their technologies.¹⁶ This is an appealing argument, but one for which we have relatively little empirical evidence.¹⁷ But even if this is an important channel, it is complementary to our story, since declines in the cost of marketing capital will improve the terms of trade for young firms in their licensing negotiations with established firms. Nor is there a consensus among economists that these changes in the

¹⁶ I thank Wesley Cohen for making this point. Our empirical analysis attempts to control for these effects by, among other things, taking into account firms’ patenting activity over time.

¹⁷ One exception is the work of Hall and Ziedonis (2001) on the semiconductor industry. See also Arora, Fosfori, and Gambardella (2001).

U.S. patent system have favored younger, smaller firms over larger, more established firms.

Regardless of how it came about, this decentralization of R&D poses a number of potential challenges to our patent system. What is the best way to coordinate a patchwork of property rights that are owned by a much more numerous and diverse population of inventors? Can we rely on Coasian bargaining to resolve these conflicts without additional government intervention? One reason to think not is that, with more rapid turnover of firms, combined with the increased asymmetry between firms, the mutual forbearance we expect from repeated interactions within an industry may be undermined. In such environments, activities such as licensing, patent pools, and standard setting become all the more important.¹⁸ While there is an extensive theoretical literature that studies these activities, the existing empirical literature is rather thin. We don't know as much as we should about how they work in practice.

I won't pretend to summarize the conclusions from this literature except to make two points. First, one should not assume that the associated transactions costs are minimal or, for that matter, impossibly high. Transactions costs will matter, and it is up to us to find out how large they are.¹⁹ The second point is that the effects of each of these mechanisms (licensing, patent pools, and standard setting organizations) depend on the quality of the property rights being granted. Facilitating the licensing or pooling of bad patents is probably bad policy.²⁰ A largely unexplored question is whether the optimal quality of patent examinations should depend on the market structure of the private innovation system. If so, it is reasonable to consider whether our examination process should be revised to reflect the current environment.

Two additional factors may be relevant. First, what should be the criteria to determine when an injunction should be imposed, especially if the questions of validity and infringement are not yet decided? Second, what is the appropriate method of calculating damages for infringement in an environment where one firm has a patent and nothing else and another has made substantial investments in order to bring a product to market under conditions of significant uncertainty? Should the rules depend on whether there are many active firms or only a few?

Conclusions

This paper reviews a number of economic analyses and offers several tentative suggestions about the design of our patent system. First, unless some form of balancing test that reflects variations in industry/technological characteristics can be implemented, a uniform set of patentability criteria will favor innovation in some industries over others. It is up to policymakers to determine where innovations should be maximized—presumably where R&D is the most productive. Second, for industries that are both R&D and patent intensive, it is vital that the property rights granted conform closely to the

¹⁸ In addition, Ziedonis (2004) points out that firms may respond to fragmented property rights by accumulating even more patents—bargaining chips for use in subsequent licensing disputes.

¹⁹ For a discussion, see chapter 3 of FTC (2003).

²⁰ See, for example, the papers by Choi (2003), Farrell and Shapiro (2005), and Shapiro (2001, 2003).

innovations they seek to protect. This may require changes in the patent examination process, the presumption of patent validity, post-grant opposition, interpretation of patent claims, and possibly even raising the pecuniary and other costs associated with obtaining patents. Third, given that the private innovation system in the U.S. is far less centralized than it was a generation ago, considerably more research is required to determine how efficiently the private sector can coordinate and price access to technologies developed by an extremely heterogeneous pool of inventors. While more theoretical research is needed, at present the primary impediment to our further understanding of the efficacy of these mechanisms is the limited amount of empirical work. This needs to change.

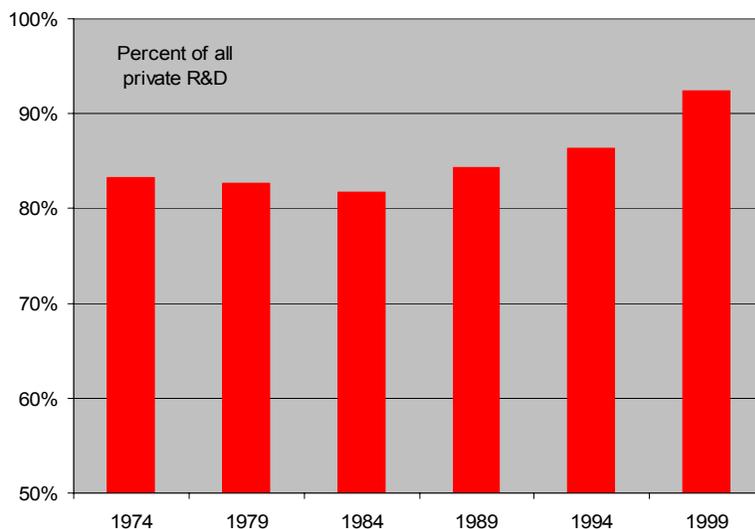
References

- Arora, Ashish, Andrea Fosfori, and Alfonso Gambardella. 2001. *Markets for Technology*. Boston: MIT Press. 2001
- Bessen, James, and Robert M. Hunt. 2004a. "An Empirical Look at Software Patents," Working Paper 03-17/R, Federal Reserve Bank of Philadelphia.
- Bessen, James, and Robert M. Hunt. 2004b. "The Software Patent Experiment," Federal Reserve Bank of Philadelphia *Business Review*, 3rd Quarter, 2004.
- Burk, Dan L. 2002. "Patent Disclosure Doctrines: Enablement and Written Description," Public Hearings on Competition and Intellectual Property Law and Policy in the Knowledge-Based Economy. Washington: Federal Trade Commission.
- Burk, Dan L., and Mark A. Lemley. 2002. "Is Patent Law Technology Specific?" *Berkeley Technology Law Journal*, Vol. 17, pp. 1155-1206.
- Choi, Jay Pil. 2003. "Patent Pools and Cross-Licensing in the Shadow of Patent Litigation," mimeo, Michigan State University.
- Cohen, Wesley M., Richard R. Nelson, and John P. Walsh. 2000. "Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not)," NBER Working Paper No. 7552.
- Farrell, Joseph, and Carl Shapiro. 2005. "How Strong Are Weak Patents?" mimeo, University of California at Berkeley.
- Federal Trade Commission. 2003. *To Promote Innovation: The Proper Balance of Competition and Patent Law and Policy*. Washington, DC: Federal Trade Commission.
- Flynn, John D. 2001. "Comments on the International Effort to Harmonize the Substantive Requirements of Patent Laws," IBM submission at <http://www.uspto.gov/web/offices/dcom/olia/harmonization/TAB42.pdf>
- Gilbert, Richard, and Carl Shapiro. 1990. "Optimal Patent Length and Breadth," *RAND Journal of Economics*, Vol. 21, pp. 106-12.
- Graham, Stuart H., and David C. Mowery. 2004. "Submarines in Software? Continuities in U.S. Software Patenting in the 1980s and 1990s," *Economics of Innovation and New Technology*, Vol. 13, pp. 443-456.

- Green, Jerry, and Suzanne Scotchmer. 1995. "On the Division of Profit in Sequential Innovation," *RAND Journal of Economics*. Vol. 26, pp. 20-33.
- Grindley, Peter C., and David J. Teece. 1997. "Managing Intellectual Capital: Licensing and Cross-Licensing in Semiconductors and Electronics," *California Management Review*, Vol. 39, pp. 8-41.
- Hall, Bronwyn H., and Rosemarie H. Ziedonis. 2001. "The Patent Paradox Revisited: Determinants of Patenting in the US Semiconductor Industry, 1979-95," *RAND Journal of Economics*, Vol. 32: 101-28.
- Hall, Bronwyn H. 2003. "Exploring the Patent Explosion," invited lecture at the ZEW Workshop on Empirical Economics of Innovation and Patenting, Mannheim, Germany, March 14/15, 2003.
- Hunt, Robert M. 1999. "Nonobviousness and the Incentive to Innovate: An Economic Analysis of Intellectual Property Reform," Working Paper No. 99-3, Federal Reserve Bank of Philadelphia.
- Hunt, Robert M. 2004. "Patentability, Industry Structure, and Innovation," *Journal of Industrial Economics*, Vol. 52, pp. 401-25.
- Hunt, Robert M. 2006. "When Do More Patents Reduce R&D?" *American Economic Review*, Papers & Proceedings, Vol. 96 (2006): 87-91.
- Hunt, Robert M., and Leonard Nakamura. 2007. "The Democratization of U.S. Research and Development after 1980," mimeo, Federal Reserve Bank of Philadelphia.
- Jaffe, Adam B., and Josh Lerner. 2004. *Innovation and Its Discontents: How Our Broken Patent System Is Endangering Innovation and Progress, and What to Do About It*. Princeton, NJ: Princeton University Press.
- Klemperer, Paul. 1990. "How Broad Should the Scope of Patent Protection Be?" *RAND Journal of Economics*, Vol. 21, pp. 113-30.
- Levin, Richard C., Alvin K. Klevorick, Richard R. Nelson, and Sidney G. Winter. 1987. "Appropriating the Returns from Industrial Research and Development," *Brookings Papers on Economic Activity*, Vol. 3, pp. 783-820.
- Merrill, Stephen A., Richard C. Levin, and Mark B. Myers, Eds. 2004. *A Patent System for the 21st Century*. Washington: National Academy Press.
- Moschini, GianCarlo and Oleg Yerokhin. 2006. "Patents, Research Exemptions, and the Incentive for Sequential Innovation," Working Paper No. 06019, Department of Economics, Iowa State University.
- Nordhaus, William. 1969. *Invention, Growth and Welfare: A Theoretical Treatment of Technological Change*. Cambridge, MA: MIT Press.
- O'Donoghue, Ted. 1998. "A Patentability Requirement for Sequential Innovation," *RAND Journal of Economics*. Vol. 29, pp. 654-679.
- Scotchmer, Suzanne. 1996. "Protecting Early Innovators: Should Second Generation Products Be Patentable?" *RAND Journal of Economics*. Vol. 27, pp. 322-31.

- Shapiro, Carl. 2001. "Navigating the Patent Thicket: Cross Licenses, Patent Pools, and Standard-Setting," in Adam B. Jaffe, Josh Lerner, and Scott Stern, eds., *Innovation Policy and the Economy*, Vol. I. Cambridge, MA: MIT Press.
- Shapiro, Carl. 2003. "Antitrust Limits to Patent Settlements," *RAND Journal of Economics*, Vol. 34, pp. 391-411.
- Ziedonis, Rosemarie. 2004. "Don't Fence Me In: Fragmented Markets for Technology and the Patent Acquisition Strategies of Firms," *Management Science*, Vol. 50, pp. 804-20.

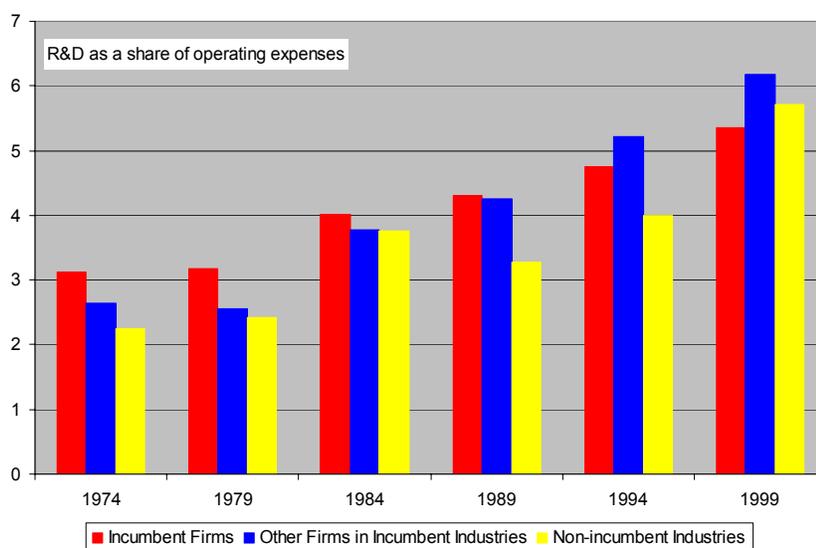
Figure 1: Private R&D Remains in the “Old” R&D Industries*



Source: Hunt and Nakamura (2007)

*: 3-4 digit SIC industries with R&D/Sales $\geq 1\%$ in 1973

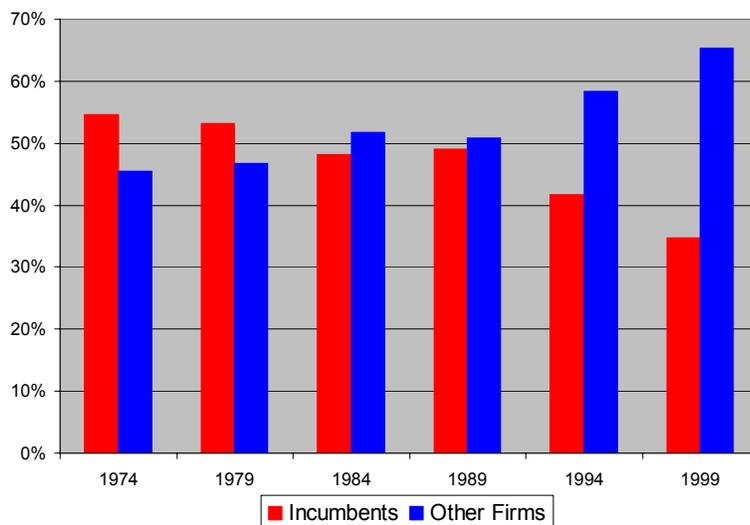
Figure 2: Private R&D Intensity Rises over Time



Source: Hunt and Nakamura (2007).

*: Incumbents are firms with 25,000+ employees in 1965

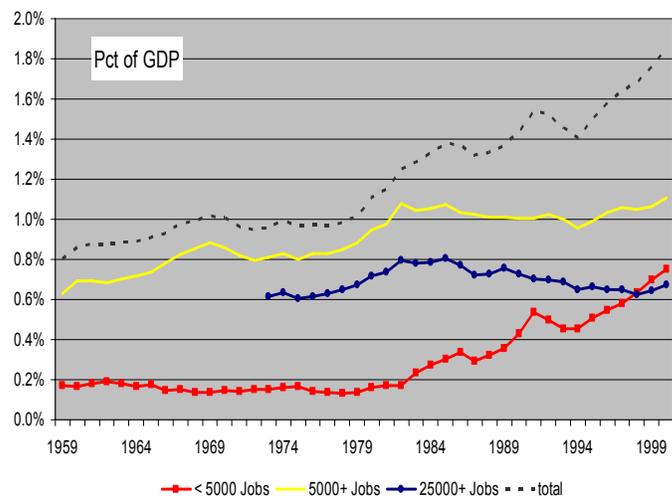
Figure 3: Share of Private R&D from Incumbent Firms Declines*



Source: Hunt and Nakamura (2007)

*: Incumbents are firms with 25,000+ employees in 1965

Figure 4: Contribution to U.S. Private R&D Intensity by Firm Size



Source: NSF Survey of Industrial R&D and author's calculations