

**Do Formal Intellectual Property Rights Hinder  
the Free Flow of Scientific Knowledge?  
An Empirical Test of the Anti-Commons Hypothesis\***

**Fiona Murray  
MIT Sloan School**

**Scott Stern  
Northwestern University and NBER**

**September 2004**

\* We thank Jeff Furman, both for extremely insightful comments and suggestions as well as for generous assistance in obtaining the citation data. We also thank seminar participants at the SPRU Conference in Honor of Keith Pavitt, the 2004 International Industrial Organization Conference, and the NBER Academic Science and Entrepreneurship Conference for extremely useful comments. All errors are our own.

**Do Formal Intellectual Property Rights  
Hinder the Free Flow of Scientific Knowledge?  
A Test of the Anti-Commons Hypothesis**

Fiona Murray  
Scott Stern

September 2004

ABSTRACT: While the potential for intellectual property rights to inhibit the diffusion of scientific knowledge is at the heart of several contemporary policy debates, evidence for the “anticommons” hypothesis has been anecdotal. This paper develops and implements an empirical test for the anticommons effect exploiting two aspects of the process of disclosing and protecting scientific knowledge: (a) scientific knowledge receiving formal IP often appears also in the form of scientific research articles (a phenomena we refer to as a “patent-paper pair”) and (b) patents are granted with a substantial lag, often many years after the knowledge is initially disclosed through paper publication. The knowledge associated with a patent-paper pair therefore diffuses within two distinct intellectual property environments – one associated with the pre-grant period and another after formal IP rights are granted. Relative to the expected citation pattern for publications with a given quality level, the anticommons theory predicts that the citation rate to a scientific publication should fall after formal IP rights associated with that publication are granted. Employing a differences-in-differences estimator for 169 patent-paper pairs (and including a control group of other publications from the same journal for which no patent is granted), we find evidence for a modest anticommons effect. While publications linked to patent grants are associated with a higher overall citation rate, the citation rate after the patent grant declines by between 9 and 17%. This decline becomes more pronounced with the number of years elapsed since the date of the patent grant, and is particularly salient for articles authored by researchers with public sector affiliations.

*Journal of Economic Literature* Classification Numbers: O300, O330, O340, L330  
*Keywords:* anti-commons, intellectual property, academic science, Pasteur’s Quadrant.

## I. Introduction

Knowledge produced by investments in research often possesses the property of *duality*, simultaneously making a contribution to both basic and applied research (Rosenberg, 1974; Stokes, 1997). Moreover, over the past two decades, knowledge that was traditionally maintained in the public domain has been increasingly associated with the application for and receipt of formal intellectual property rights (IPR) such as patents. This has resulted in a vigorous contemporary academic and policy debate about the impact of IPR on the creation and diffusion of useful scientific knowledge. On the one hand, when formal IPR (in the form of patents and/or copyrights) is granted over basic knowledge that would traditionally have been maintained in the public domain, the IPR system may create an “anticommons” effect, inhibiting the free flow of scientific knowledge and the ability of researchers to cumulatively build on each other’s discoveries (Heller & Eisenberg, 1998; David, 2003, 2000; Lessig, 2002; Etzkowitz, 1998; Krinsky, 2003). Simply put, the expansion of IPR into traditional scientific fields may be “privatizing” the intellectual commons to the detriment of scientific progress (Argyres & Liebskind, 1998; David, 2001b). Alternatively, by establishing specific rights such that “trade” can occur over the applied aspects of knowledge, IPR may facilitate the creation of a market for ideas and so mitigate the disincentives to disclose and exchange knowledge which might otherwise remain secret (Merges & Nelson, 1990, 1994; Arora, 1995; Gans and Stern, 2000, Levin et al., 1987). Indeed, the classical public policy rationale for the patent system is to simultaneously provide incentives for innovation while maximizing the disclosure of new discoveries and inventions. Nevertheless, there is currently no systematic empirical evidence to adjudicate the debate one way or another.

This paper develops and implements a novel empirical methodology addressing two key questions: What is the overall scientific impact of published knowledge which is also patented? How does the *grant* of formal IPR over such knowledge impinge the diffusion and impact of scientific research findings?

This paper addresses these questions by exploiting the duality of scientific knowledge and the distinctive processes through which such dual scientific knowledge is diffused and protected. In particular, reflecting its duality, scientific knowledge that appears in the form of scientific research articles will often be subject to patent applications and receive formal IPR in

the form of patents. We refer to the simultaneous instantiation of individual research findings as both a scientific article and a patent application as a “patent-paper pair” (Murray, 2002; Ducor 2000). Second, though publication lags are usually modest (on the order of a few months), patent grant delays are substantial. In most cases, formal IP rights are granted at least two (and often three or more) years after initial application.

The scientific knowledge associated with a patent-paper pair can therefore be observed under two distinct intellectual property environments associated with the pre-grant and post-grant period. More than simply a “timing” effect, it is useful to recall that (a) until 2001, U.S. patent applications have historically been kept secret until the patent grant date, (b) patent applicants face substantial uncertainty over the scope and precise nature of the rights they ultimately receive and (c) the legal rights associated with a patent commence with the patent grant date. In other words, during the period before the patent grant date, researchers intent upon building on publicly disclosed (published) research findings are unrestricted by (and often will be unaware of) pending patent applications, and patent applicants cannot assert their rights until after the uncertainty associated with the patent grant process has been resolved.

By observing individual pieces of scientific knowledge in different intellectual property environments, we have constructed a systematic test of the Anticommons hypothesis. Specifically, if the grant of intellectual property hinders the ability of researchers to build on a given piece of knowledge, then the citation rate to the scientific publication disclosing that knowledge should fall after formal IP rights over that knowledge is granted. Of course, such an analysis must control for the fact that citation patterns vary with the underlying qualities of the article, and with the time elapsed since publication. Our empirical test, therefore, accounts for individual publication quality, the effects of publication age and scientific maturity (through the introduction of article, article age, and calendar year fixed effects, respectively). In other words, we test for the anticommons effect by calculating how the citation rate for a scientific publication *changes* after patent rights are granted, accounting for fixed differences in citation rate across articles and relative to the trend in citation rates for articles with similar characteristics.

We implement this test using the entire population of peer-reviewed scientific articles appearing between 1997 and 1999 in *Nature Biotechnology*, perhaps the leading publication for research exhibiting the duality of a basic science contribution and the potential for commercial application in the life sciences. Of all articles in this journal during this period, just fewer than

50% receive a USPTO patent grant. We examine differences in the annual forward citation patterns for publications between those with and without a patent pair. While the average cumulative citations between the two groups is relatively similar, articles linked to a patent have a higher *initial* citation rate which then *converges* with the non-patented article citation rate.

Using a differences-in-differences estimator of the change in the citation rate after a patent grant occurs, we establish three key findings. First, we find robust evidence for a quantitatively modest but statistically significant anticommons effect. Across different specifications, the article citation rate after a patent grant declines by 9 to 17%. This decline becomes more pronounced with the number of years elapsed since the date of the patent grant. Moreover, consistent with what theory would predict, the evidence for the anticommons effect is particularly salient for those articles with authors with public sector affiliations (such as a university or government laboratory). Indeed, if we restrict the sample to US authors with public sector affiliations, we are able to identify an anti-commons effect, even if we only examine articles which ultimately receive a US patent. While we are cautious in our interpretation, the evidence suggests that the anticommons effects may have a sound empirical basis, but the size of the effect may be modest.

The remainder of the paper is as follows. The next section reviews the economic foundations of the anti-commons hypothesis, including the role of knowledge duality. Section III develops the empirical test to assess the anticommons hypothesis. After a review of the data in Section IV, Section V presents our empirical findings. A final section concludes.

## **II. The Impact of Intellectual Property on the Diffusion of Scientific Knowledge**

Economists, law and technology scholars, and policymakers have become increasingly concerned about the impact of intellectual property rights on the growth and diffusion of new scientific knowledge (Heller and Eisenberg, 1998; David, 2001a, Campbell et al. 2002; Strauss et al. 2002; Walsh et al. 2002). By its very nature, scientific knowledge is non-rivalrous, so that the diffusion of that knowledge can serve repeatedly (and with little additional cost) as an input into future knowledge production. Because IP can serve to exclude follow-on researchers from exploiting scientific discoveries, the use of IP can undermine the process of cumulative scientific discovery. More precisely, in the absence of an efficient mechanism for gaining access to knowledge (e.g., through efficient licensing), IPR can be used to erect barriers that hinder the

effective exploitation of the scientific commons. It has been argued that these restrictions in the use of non-rivalrous knowledge can lower overall research productivity, leading to the so-called “anti-commons” effect (Heller, 1998; Heller and Eisenberg, 1998; David 2003). Given the role of the Federal government as the single largest source of funding for basic research under the rationale of a strong historical linkage between basic research spillovers and long-term economic growth (Bush, 1945; Adams, 1990), the anti-commons effect has become a central concern for policymakers and analysts of the impact of science and basic research on growth.

The economic logic of the anti-commons effect can only be understood by re-conceptualizing the relationship between basic and applied research. In the traditional “linear” model, the norms and institutions supporting the production and use of basic versus applied research are separable and distinct. Under this model, applied research *exploits* publicly available basic research as an input, transforming that knowledge into innovations with valuable application. Though the linear model has been sharply criticized (Kline and Rosenberg, 1986), most formal theoretical and empirical economic research remains premised on the linear model, from assessment of the impact of university research (Jensen and Thursby, 2001; Zucker et al. 1998; Mowery et al, 2001; Narin and Olivastro, 1992) to the impact of Science and basic research on economic growth (Romer, 1990; Adams, 1990).

Most research in the “new” economics of Science focuses on comparisons of alternative institutional arrangements and incentive regimes associated with different stages of knowledge production. Building on the classical distinction between “free access” and “private property rights” emphasized by Weitzman (1974) and research in the sociology of science (Merton, 1973), Dasgupta and David (1994) insightfully demonstrate the sharp differences in economic incentives between an Open Rights regime (for basic research) and a Private Property Rights regime (for applied research). On the one hand, the Open Rights regime (“Science”) depends on the equilibrium adoption of norms that allow for complete and full disclosure and diffusion of knowledge. This system includes the recognition of scientific priority by future scientific generations, the importance of demonstrating experimental replicability, and a system of public (or coordinated) expenditures to reward those who contribute to cumulative knowledge production over the long term (Merton, 1973; Dasgupta and David, 1994). By premising career rewards (such as tenure) on disclosure through publication, Open Science enables the public

goods nature of basic research and serves as a foundation for applied research.<sup>1</sup> In contrast, the incentives that govern the private property rights regime are quite different. Instead of rewarding researchers on the basis of their ultimate impact on follow-on research, private research incentives depend on the degree to which a researcher can *exclude* others and so appropriate some of the value created by their knowledge through the commercialization of new technology (Nelson 1959; Arrow 1962; Levin et al. 1987; Kremer 1997). The rationale for IPR in this regime is clear, as they provide incentives for innovation while simultaneously achieving the disclosure of technological knowledge for long-term knowledge accumulation.

While the linear model allows for a concise formulation of the relationship between the nature of knowledge and the incentives provided for its production and distribution, this analytical framework fails when knowledge has *both* basic and applied value. Stokes (1997) therefore reformulated the traditional distinction between basic and applied research by highlighting the *duality* of research; rather than being empirically distinct, a single discovery could simultaneously possess both applied and basic characteristics. Figure A illustrates this essential insight. Instead of placing research on a single dimension ranging from basic to applied, research motivations can vary along two dimensions: in terms of whether they are dependent on “considerations of use” and, separately, on a “quest for fundamental understanding.” While the traditional linear model captures the distinction between pure basic research (as exemplified by the type of research conducted by Bohr) and pure applied research (as exemplified by Edison), Stokes’ formulation also allows for “use-inspired basic research,” as exemplified by the research activities of scientists such as Pasteur. Stokes highlighted the degree to which Pasteur’s fundamental insights into microbiology were developed as much for their application to problems such as cholera and rabies as for their relevance in developing and testing the germ theory of disease (Geison 1995; Stokes 1997).

The anticommons hypothesis is perhaps the key novel implication of the analytical framework suggested by Stokes. While traditional theory justifies IPR on the basis of enhancing incentives for innovation in the tradition of pure applied research, the anti-commons hypothesis posits that the equilibrium level of knowledge diffusion and subsequent research productivity may be declining in the use and restrictions imposed by IPR particularly for “use inspired basic

---

<sup>1</sup> While closely associated with university research, Open Science is also feasible (and profitably adopted) by private firms, including many within industries dependent on the life sciences (Cockburn and Henderson, 1998; Zucker et al., 1998; Stern, 1999; Murray, 2002).

research” or research with a high degree of duality (Heller and Eisenberg, 1998). Most notably, Heller and Eisenberg (1998) suggest that the assignment of IPR to basic research provides researchers with a control right to *exclude* others using that knowledge for the traditional purpose of cumulative knowledge production. In other words, when research incentives are already sufficiently high due to a high “quest for fundamental understanding,” privatizing the intellectual commons imposes a “tax” on the use of that knowledge and may restrict the diffusion of that knowledge, with few positive incentive effects. Instead of raising incentives for discovery, the use of IPR over knowledge which has been traditionally associated with Open Science can lower the equilibrium level of research productivity.

Of course, it is possible that IP facilitates a “market for ideas” by increasing incentives for disclosure (rather than secrecy) and encouraging the exchange and trade of knowledge (Merges and Nelson, 1994; Arora et al, 2001; Gans and Stern, 2000; Gambardella, 1995). The ability to use IPR to contract over knowledge can potentially ameliorate the Anticommons effect, at least with respect to its applied aspects. However, the proliferation of IPR over disparate but interrelated discoveries may limit the efficacy of bilateral contracting mechanisms. To the extent that IPR is narrow in scope and highly dispersed across individuals and institutions, fragmentation can impose significant transaction costs and limit research progress (Eisenberg, 1996; Shapiro, 2001; Hall and Ziedonis, 2001; Ziedonis, 2004).

Consider the controversy in the early 1990s over the patenting of gene fragments known as expressed sequence tags (ESTs). These tags or genetic markers are small fragments of DNA generated to identify regions of a gene using well known research methods. These fragments were subject to patent applications that claimed not only the small fragment but the entire gene discovered using the EST. This raised considerable debate among the scientific and legal communities; some argued that EST patents would distort research priorities, encourage researchers to search for and patent ESTs rather than focus on characterizing full-length genes, and preclude further research on entire genes (Kimball, 2001). Furthermore, it was anticipated that these patents would lead to complex and inefficient litigation and ownership disputes over full genes (Eisenberg, 1996). Others suggested that these patents would, like many other patents on novel ideas, spur innovation. Although no longer considered patentable by the USPTO on the basis of a lack of utility, the EST episode highlights the potential for complex transactions in the

use of genes for biomedical research and development (Merges and Nelson 1990; Lawson, 2002).

Anecdotal evidence also suggests that whole gene patents (such as those for the breast cancer genes BRCA-1 and BRCA-2) have also stifled innovation and rendered more complex knowledge accumulation around breast cancer diagnostics and therapeutics (William-Jones, 2002). In the breast cancer case, a portfolio of over a dozen US and international patents gave biotechnology firm Myriad exclusive rights to commercialize laboratory testing services, diagnostic test kits and therapeutic products that use the BRCA1/2 DNA sequences. However, because the discovery of the genes was based not only on Myriad's commercially funded research efforts but also built upon internationally generated public knowledge, private ownership of the breast cancer gene diagnostic kit was highly controversial. These gene patents like other broad patents raise the question of the degree to which patent scope shapes further innovation (Merges and Nelson, 1990; Crespi 2000).

In addition to the impact of broad and fragmented rights, the growing privatization of basic research may also be associated with a higher level of secrecy (Blumenthal et al 1996), refocused research agendas (Thursby and Thursby, 2003; David 2003) and an increased potential for bias in research results (Nelkin 1984; Krimsky 2003). Together, anecdotal evidence and the theoretical implications of the anti-commons effects have led some to sharply question the rationale for allowing IPR over basic research (David, 2001; Heller and Eisenberg, 1998). However, despite these calls for policy intervention, systematic empirical evidence for the anticommons hypothesis is sparse. In a notable exception, Blumenthal et al (1996) provide a useful survey suggesting that scientists perceive that the increased use of intellectual property rights are limiting their access to research resources. However, extant studies have yet to disentangle whether perceived limitations in knowledge accumulation are associated with the precise type of knowledge under consideration (basic or applied) or the property rights regime associated with that knowledge (non-patented versus patented).

With the rapid rise of university patenting and licensing activities over the past two decades, such concerns have become more important (Henderson et al., 1994; Mowery et al., 2001). While universities have historically served as the central institutions supporting the norms of Open Science, university technology licensing offices have been cited as increasingly limiting the use and diffusion of knowledge produced by university researchers (Campbell et al.,

2002). Given that knowledge with the property of duality lies at the heart of much of modern scientific research (particularly in the life sciences), it is important to address whether it is possible to adjudicate the impact of intellectual property on the development and diffusion of research knowledge.

The remainder of this paper proposes and executes a simple empirical test to address precisely this question. In so doing, we provide the first direct test of the anti-commons hypothesis. To achieve this objective, we implement a procedure that takes direct advantage of the “dual nature” of knowledge (by exploiting so-called “patent-paper pairs”) and the lag between the initial publication of scientific knowledge and the date of patent grant. The next section outlines this empirical approach before turning to our empirical examination.

### **III. Patent-Paper Pairs and Patent Grant Delay: An Empirical Test of the Anticommons Hypothesis**

The anticommons theory makes specific predictions about the impact of enhanced formal intellectual property for knowledge which might diffuse through the mechanisms of Open Science. Specifically, while the scientific community as a whole may benefit from the free dissemination and diffusion of knowledge, individual researchers have strong incentives to take advantage of the protections afforded by formal IP rights. If protected by IP, the impact of an individual piece of knowledge on follow-on research by others is diminished, as the expected benefits of drawing on that research are reduced for future generations. In other words, increases in the availability or strength of IPR can result in a lower equilibrium level of on-going research productivity.

Not simply an abstract theory, the anticommons hypothesis is at the heart of several contemporary policy discussions about the impact of enhanced IPR on the diffusion and impact of knowledge traditionally diffused through scientific channels. However, empirical research in this area has been hampered by a fundamental inference problem. Specifically, for a given piece of knowledge *with* IPR, one cannot observe the counterfactual impact that knowledge would have had if the IPR had been waived. For example, knowledge protected by IPR may tend to have a higher (or lower) intrinsic scientific value than knowledge that is not protected by IP. A simple comparison between patented and non-patented knowledge may therefore be biased by unobserved heterogeneity. From an experimental perspective, the econometrician would ideally

like to observe a given piece of knowledge in two distinct institutional environments (e.g., a non-patent versus patented environment), and then compare the impact of that knowledge across those two regimes.

While one cannot replicate this ideal experimental design, this paper develops and implements an econometric strategy that allows a precise estimate of the impact of IPR on the diffusion of scientific knowledge. Our approach exploits two key elements of the system by which scientific research is diffused and IPR are granted. First, we take advantage of the fact that research that is simultaneously both basic and applied (i.e., in Pasteur's Quadrant) is increasingly instantiated as both a scientific publication and a patent. Building on Murray (2002) and Ducor (2000), we refer to these instances as "patent-paper pairs." Such pairs represent instances in which a given experimental result or piece of scientific evidence results in two distinctive observables -- a patent and a publication.

Second, we exploit the fact that, for a given patent-paper pair, there can exist a substantial gap between the date of scientific publication and the date at which the associated patent is granted. While publication in the scientific literature often occurs within six months (or less) after initial submission to a journal, the delay between the initial application and receipt of a patent is often many years (in most cases, resulting in a 2-4 year time window). For any patent-paper pair, then, we can observe both a pre-patent grant period and a post-patent grant period. Moreover, for each period, it is possible to estimate the impact of that knowledge, measured as the rate of citation to the initial article by follow-on scientific research articles.

We combine the phenomena of patent-paper pairs and the existence of a patent grant delay to provide a precise test of the anticommons hypothesis. Specifically, our dataset is composed of scientific research publications "at risk" of being associated with a USPTO patent and thus forming a patent-paper pair. Indeed, as described in the next section, approximately 50% of the publications in our dataset in fact result in a patent-paper pair. Because we observe citations to a scientific publication both before and after the patent is received (and because we observe a control group of publications which never receive a patent) we are able to identify how the pattern of citations over time to a scientific publication changes as the result of the receipt of a patent. This test goes beyond the potentially biased test of whether patented publications are more or less highly cited than those that are not associated with patents.

More precisely, if the grant of intellectual property hinders the ability of researchers to build on a given piece of knowledge, then the citation rate to the scientific publication disclosing that knowledge should fall after formal IP rights over that knowledge is granted. Of course, such an analysis must account for the fact that we are measuring the impact of scientific research using citations, which comes in the form of count data which are skewed heavily to the right. Therefore, except where noted, we employ a negative binomial model of the citations produced per year for each scientific article in our dataset. As well, the impact of a given piece of research will vary considerably with its underlying importance and with the time elapsed since initial publication. To control for these factors, our empirical specifications will account for individual publication quality (through article fixed effects) and for the effects of publication age and scientific maturity (through age and vintage effects). Taken together, our baseline empirical test for the anticommons hypothesis is therefore:

$$CITES_{i,pubyear,t} = f(\varepsilon_{i,pubyear,t}; \gamma_i + \beta_{pubyear} + \delta_{t-pubyear} + \psi POST - GRANT_{i,t}) \quad (1)$$

where  $\gamma_i$  is a fixed effect for each article,  $\beta_{pubyear}$  is a cohort effect,  $\delta_{t-pubyear}$  captures the age of the article, and POST-GRANT is a dummy variable equal to one only for years in which a patent is in force for an individual article.<sup>2</sup> In other words, we test for the anticommons effect by calculating how the citation rate for a scientific publication *changes* after patent rights are granted, accounting for fixed differences in the citation rate across articles and relative to the (non-parametric) trend in citation rates for articles with similar characteristics.<sup>3</sup>

It is important to emphasize that our exploitation of patent grant delay allows us to capture more than simply examining articles at different points in time. Indeed, simple “timing” effects are captured by the use of fixed age effects. Instead, patent grant shifts a given piece of publicly disclosed scientific knowledge into an alternative institutional regime. In particular, until 2001, patent applications were kept secret until granted. As well, when patents are pending, patent applicants face substantial uncertainty over the scope of the rights they will receive.

---

<sup>2</sup> As well, it is possible to separately identify the degree of unobserved heterogeneity versus the impact of patent grant itself. In particular, one can replace the full set of fixed effects in the specification with an overall effect capturing the difference between patented and non patented articles, (i.e.,

$$CITES_{i,pubyear,t} = f(\varepsilon_{i,pubyear,t}; \beta_{pubyear} + \delta_{t-pubyear} + \lambda PATENTED_i + \psi POST - GRANT_{i,t})).$$

<sup>3</sup> To the extent that there are differences among articles in the time lag between scientific publication and patent grant, it is possible to identify  $\psi$  even if one conditions on a sample of articles that ultimately receive a patent. By identifying the impact of patent grant from differences in the length of time until a patent is received, we account for unobserved differences between articles that do or do not receive a patent.

Finally, patent rights only commence with patent grant. In other words, during the pre-grant period, researchers building on findings that are disclosed in the scientific literature are unrestricted by (and often unaware of) pending patent applications.

In addition to estimating the overall average impact of patent grant, we can gain additional insight into the anti-commons hypothesis by exploring several variations of (1). First, it is likely that the impact of a patent grant increases over time, particularly in the first few years after the grant. As such, we can estimate whether the impact of patent grant changes with the time elapsed since patent grant. In addition, we can check that our estimate of the impact of patent grant is not simply the result of a separate trend for patented articles by estimating the pre-grant trend for those articles which are ultimately patented. Further, we can interact POST-GRANT with discoverer (researcher) characteristics. For example, because articles linked to researchers from private sector firms may be *assumed* to be protected by some form of IP, anticommons theory would suggest that the impact of patent grant itself should be higher for scientific articles associated with public sector researchers. More generally, we can consider how the effect of patent grant varies with technological and institutional factors associated with the scientific research article and the patent itself. We review each of these factors in our discussion of the dataset in the following section.

## **IV. The Data**

### *IV.A. Sample Definition*

The data for this study is based on the entire population of peer-reviewed research articles published in the journal *Nature Biotechnology* over the three year period from 1997 to 1999. While the journal publishes scholarly material in a variety of formats, we confine our data to research articles which are defined by the editorial policies of the journal as “a substantial novel research study” (see *Nature Biotechnology*, A Guide to Authors). Under these criteria, the dataset consists of 340 unique research articles.

Our sample population was chosen to focus on research exhibiting the duality emphasized by Stokes (1997). Biotechnology defines an important research arena which, since its inception, was seen as making a contribution both to scientific progress and technological progress, leading to a complex interaction between the institutions of open Science and those of Private Property Rights (Kenney, 1986; Orsenigo, 1989; Powell et al., 1996; Gambardella,

1995). In *Nature Biotechnology*'s first issue in 1996, the editorial mission of the journal was described as: "cover[ing] business, financial, and regulatory matter: not to do so would be perverse and self-defeating. But its emphasis will be unashamedly on research and development, the fuel for biotechnology's fire." (Nature Biotechnology, 1996).<sup>4</sup> The publishing policy adopted by *Nature Biotechnology* explicitly aimed at research with potential applications to biotechnology: "[the journal] aims to publish high-quality original research that describes the development and application of new technologies in the biological, pharmaceutical, biomedical, agricultural and environmental sciences, and which promise to find real-world applications in academia or industry. We also have a strong interest in research that describes the application of existing technologies to new problems or challenges, and basic research that reports novel findings that are directly relevant and/or of interest to those who develop biology into technology." Since its inception, *Nature Biotechnology* has established itself as the leading outlet (in terms of measured scientific impact) for refereed scientific research relating to biotechnology and applied microbiology. In other words, research published in *Nature Biotechnology* is both "at risk" of serving as a foundation for future scientific studies and for commercial exploitation.

For each article, we determined whether a patent associated (or paired) with the article had been granted by the USPTO. Using the USPTO search engine, we defined a series of searches for each article. The basic search included i) the first, last and corresponding authors for the article and ii) the list of institutions found in the article "address field" in the Web of Science database. For some institutions, specific name variations were used to account for the fact that institutions may not patent under the full institutional name. For example, patents assigned to the University of Oxford are listed under the name of its separately incorporated technology licensing office, ISIS Innovation. Different combinations of authors and/or institutions were used from the most to the least inclusive in order to identify all issued patents associated with the authors and institutional affiliations whose research appeared in *Nature Biotechnology*. For example, since some patents were assigned to individuals (rather than an institution), the search procedure examined whether each author for each article received a patent within the time frame in question. After establishing the set of patent grants received by

---

<sup>4</sup> The new journal was not entirely new. It "picked up the torch from Bio/Technology", a journal founded in 1983 to explore and publish leading edge science in biotechnology largely from academics but also from the newly founded biotechnology firms.

individuals and institutions represented in the articles, patent abstracts and claims were read to establish the presence of a patent-paper “pair.” To do so, we verified whether the material described in the abstract of the article was incorporated into the description, claims and/or examples of the granted patent.<sup>5</sup> By checking the precise content of patents granted to those whose research is published in *Nature Biotechnology*, our procedure provides a consistent way to identify the subset of articles within our overall sample which are also patent-paper pairs.

Of the 340 articles in the dataset and using this procedure, 169 articles are associated with a patent as of October, 2003. In other words, approximately half of all publications in *Nature Biotechnology* are associated with a patent-paper pair within five years of publication. Of course, for the differences-in-differences estimator to be consistent, patented and unpatented articles must be comparable. Specifically, the bulk of the empirical work assumes that the entire sample of articles is initially “at risk” of being patented, and that, but for being patented (and a control for overall article quality), patented and non-patented articles follow the same citation time trend. The results may be subject to bias if non-patented articles are systematically different than patented articles in terms of their potential for patentability (or along an unobserved dimension impacting the citation time trend).

We checked the comparability assumption in several ways. First, the sample design ensures that the articles are comparable insofar as all are drawn from the same (reasonably specialized) high-quality journal. All articles have undergone a similar refereeing process, and editorial decisions are presumably made with the journal’s editorial mission in mind.

Second, for a subset of 34 of the non-patented articles, we undertook a detailed evaluation of their innate patentability. The standard for patentability (defined in 35 U.S.C. Section 101) is defined as “new and useful processes, machines, manufactures, compositions of matter; and any new or useful improvements thereof...subject to the conditions of patentability” where conditions of patentability include novelty, non-obviousness and utility (see, for example, Merges, 2003, for an overview of the law and economics of the patent system). It is important to note that this standard excludes important categories of knowledge which might be reflected in scientific research articles, such as the discovery of new scientific principles, abstract ideas, and

---

<sup>5</sup> One of the authors (Murray) holds a PhD in Applied Sciences and has conducted detailed qualitative research on the scientific content of contemporary biotechnology and applied microbiological research (Murray, 2002). The criteria used to assign a patent-paper pair was conservative insofar as there had to be a direct connection between the disclosures in the article abstract and patent record. In the vast majority of cases, the presence (or not) of a patent-paper pair was unambiguous.

the identification of naturally occurring materials. However, novel “research tools” and “compositions of matter” are patentable, and constitute the bulk of patented technologies in the sample. An experienced patent attorney (graciously) undertook an examination of the publication abstracts and was asked to make a “conservative” determination of whether the research findings included a potentially patentable discovery. Of the articles submitted for review, more than 75% (27 out of 34) were considered to be obviously patentable; of the remaining, most contained at least some potential for patentability (i.e., while they failed the conservative test we requested, they likely would have passed a more lenient (but still plausible) standard for patentability). In particular, most of the articles not considered patentable under our test reported research results using standard techniques on pre-existing materials, and so the abstract did not include a description of a novel research tool or composition of matter. While these evaluations do not constitute a formal legal opinion, this check does provide support for the assumption that most (if not entirely all) articles within the sample are at risk of being patented.

Further, we directly compared the similarity of articles within the sample. The MedLine database includes a feature allowing the identification of “similar” articles, based on keyword matching. For each patented article, we identified the “most similar” non-patented article within the sample and qualitatively evaluated a subset of these for comparability. For the entire dataset, we were able to identify a non-patented match for each of the 169 patented articles. By and large, matched articles were found to be qualitatively similar, both in terms of their underlying scientific content and their potential for patentability.<sup>6</sup>

Overall, given the nature of the sampling process and the qualitative comparisons of the patentability and scientific content of patented and non-patented articles, the assumption of comparability seems plausible. As such, the bulk of the empirical work assumes that, after controlling for overall quality, the sample population is composed of articles at risk of being

---

<sup>6</sup> Consider the following matched example. One article, published in June, 1997, describes a research study (by researchers at John Hopkins University) of a novel method of using bacteriophage which express ligands on their surface to detect the interaction between key proteins, thus allowing “a powerful approach to the molecular studies of protein-protein interactions” (Li, 1997). The second article, published in December, 1999, by researchers at Sugen (a biotechnology firm) describes a novel display technique to examine the interaction among specific proteins using a library of DNA fragments that contain specific mutations used to reduce non-specific binding that will render results imprecise and difficult to analyze. The method was applied to key signal transduction pathways and the authors suggest could be “a rapid and efficient tool for elucidating protein networks” (Zozulya et al. 1999). Though both articles are concerned with extremely similar scientific issues (methods for identifying protein interactions) and both are clearly describing (patentable) research tools, only the first is associated with a USPTO patent grant.

patented and, but for a patent grant, the time trend for all articles follows the same stochastic process. Of course, it is possible to relax each of these assumptions, and we assess alternatives in our empirical evaluation.

The dataset is drawn from three distinct sources. First, article-specific characteristics are initially gathered from *Nature Biotechnology*. These characteristics include the precise date of publication, the number of authors, and the location and institutional affiliation of each author (available from the address list provided for each article). The citation counts (through October, 2003) for each article are calculated using the Science Citation Index Expanded (SCI). Maintained by the Institute for Scientific Information (ISI), SCI records reference and citation information for nearly six thousand scientific and technical journals in approximately 150 disciplines. For each article associated with a USPTO patent, we collected a number of patent characteristics from the USPTO public database. These characteristics include the date of patent application and patent grant, the number of inventors and assignees, the location and institutional affiliation of each inventor and assignee, the primary US patent class, the number of citations to prior patented material, the number of citations to prior non-patented material, the number of allowed claims, and whether the patent is associated with a “Federal interest” (indicating Federal funding of research upon which the patent is based).

#### *IV.B. Summary Statistics*

For the variables used in our analysis, Table 1 provides variable names and definitions and Tables 2 and 3 reports summary statistics. For each article in the dataset, we track citations beginning in the year in which the article was published and continuing until the end of 2002. The total number of articles in the dataset is 340, and the total number of article-year observations is 1688.

The key dependent variable in our analysis is FORWARD CITATIONS, the number of articles that reference the focal article in a given year. Not surprisingly given the prestige and quality of *Nature Biotechnology*, the average level of annual citations received by articles in this dataset is quite high, relative to randomly selected academic articles (mean = 9.35), and, by the end of 2002, the average article had received more than 54 total cumulative citations. Consistent with prior citation analysis studies, the distribution of citation counts is quite skewed, with nearly 20% of the citation-years receiving either 0 or 1 citation, but also including one

annual citation count equal to 181 (Figure B). Because we observe article-years from 1997 through 2002 (but only observe articles published in 1998 or 1999 for a shorter set of years), the average CITATION YEAR is at the margin of the 2000 calendar year, and the average AGE observed within the sample is just a little over 2.0.

While the heart of the analysis incorporates article fixed effects to account for differences between articles in terms of their impact and overall quality, we have collected a number of article-specific characteristics. These characteristics include the number of authors (# AUTHORS (mean = 5.89) as well as information associated with the institutional affiliations of those authors. For example, US AUTHOR is a dummy variable measuring whether *at least* one of the authors lists a US address (mean = 0.59). As well, we observe the organization for which the individual works. We assign university and government researcher affiliations as “public sector” institutions and pharmaceutical and biotechnology affiliations as “private sector” organizations. We then define two dummy variables, PUBLIC SECTOR AUTHOR and PRIVATE SECTOR AUTHOR, which are equal to one if *at least* one author is associated with a public sector or private sector organization, respectively. Interestingly, nearly 90% of the articles in the sample have at least one public sector author, and more than 30% have at least one private sector author.

Finally, our data includes information about the patent associated with each patent-paper pair. Just under 50% of articles are associated with a patent, and the average date of patent grant (weighted by article) is mid-2000.<sup>7</sup> The average lag between the patent application and patent grant date is over 1100 days (i.e., just over 3 years). PATENTED, POST-GRANT is the dummy variable indicating a citation-year is linked to a patent-paper pair, and the patent has already (or just been) issued and is in force.<sup>8</sup> Conditioning on only those articles which are associated with patent-paper pairs, just under half of the citation-year observations are observed in the post-grant period.

We also measure several patent-specific characteristics. Along with the article-specific characteristics, each of these measures can be interacted with PATENTED, POST-GRANT to

---

<sup>7</sup> Some patents have been issued to articles in 2003 (mostly associated with 1999 publications). Inclusion or exclusion of these 23 articles from the analysis does not change any of the qualitative findings (primarily because the differences-in-differences strategy relies on those articles where we observe a change in the IP regime during those years where we observe a citation count).

<sup>8</sup> PATENT, POST-GRANT is equal to one from the year of the patent grant date. All of the qualitative results are robust either to including the years associated with the year of the patent grant in the “pre-grant” category or dropping these article-year observations from the analysis.

assess whether the impact of patent grant on the citation rate depends on observable characteristics of the article or associated patent.

First, we measure the number of listed inventors (# INVENTORS); interestingly, the average for # INVENTORS is just over half that associated with # AUTHORS. We also calculate two measures of the level of prior art cited by the patent, including the number of citations to prior patents (PATENT BACK CITATIONS) and the number of citations to prior non-patent references (PATENT BACK REFERENCES). As the number of prior art references increases, the potential for a “patent thicket” increases (Shapiro, 2001); in the spirit of Heller and Eisenberg (1998), the presence of a patent thicket may exacerbate the anticommons effect and result in a greater decline in the post-grant citation rate of patented articles. Interestingly, relative to the overall means for citations made to patented prior art by “biotechnology” patents reported by Allison and Lemley (2002), the averages for both PATENT BACK CITATIONS (7.26) and PATENT BACK REFERENCES (28.19) are somewhat (though not extraordinarily) high.

We also include two measures of the type and scope of the patent. # CLAIMS is simply the number of allowed claims. Though # CLAIMS is an extremely imperfect (and noisy) measure of the strength of patent rights, it does provide at least a weak measure of patent scope and so increases in # CLAIMS should intensify the impact of patent grant on article citation rates. TOOLS PATENT is a dummy variable equal to one if the primary class for the patent is within the 435 and 800 patent classes. Out of the 11 3-digit patent classes represented across the patents within the sample, these two 3-digit classes are most closely associated with processes and tools. Since research tools have been of particular concern within the anticommons debate (relative to composition of matter patents), the TOOLS PATENT dummy will allow us to assess whether the impact of patent grant on citation is greater for patents covering research tools and methods.

Finally, the dataset includes two measures related to the nature of the research funding associated with the project. First, similar to PUBLIC SECTOR AUTHOR, PUBLIC SECTOR ASSIGNEE (mean = 0.65) is a dummy variable equal to one if there is *at least* one assignee from the public sector. In addition, when the research upon which a patent is based is funded (even in part) by the Federal government, the applicant must disclose a Federal “interest” (indeed, the ability to retain patent rights in most cases despite this Federal interest is at the heart of the Bayh-

Dole Act). Of the 169 articles associated with patents in our dataset, 49 report a Federal interest. To the extent that contemporary policy proposals related to the anticommons effect would most likely impact Federally funded research, the impact of patent grant on citation is particularly interesting for those patents for which GOVT FUNDED = 1.

#### *IV.C. Patented Versus Non-Patented Articles*

Table 3 compares the means of patented and unpatented articles within the sample. A few notable differences stand out. First, the average rate of citation is relatively similar across the two groups, with the patented articles receiving, on average, an additional citation per article-year over the sample. However, this 10% average difference masks more substantial differences that manifest themselves over time. In Figure C, the average FORWARD CITATIONS are plotted by AGE (years since publication). During the year of publication and in the subsequent three years, PATENTED articles have a significant citation advantage, equivalent to an approximately 15-18% “boost” over the citations rates for non-patented articles. However, in the fourth and fifth years after disclosure in the literature, patented articles converge to the citation rate associated with non-patented articles. As we explore further in the next section, it is during these later years in which patented articles are in the post-grant phase, perhaps helping to explain the decline relative to trend experienced by patent-paper pair publications. In addition to these overall differences in citation rates, it is important to recognize that there are also differences in article characteristics. Relative to non-patented articles, patented articles have a significantly higher chance of having at least one US author, or at least one author from a private sector organization (the differences in means in # AUTHORS and PUBLIC SECTOR AUTHOR are not significant).<sup>9</sup>

These data suggest that the number of citations and article characteristics vary across the margin of whether or not an article is part of a patent-paper pair. While the drop-off in the citation advantage associated with patented articles is consistent with the presence of an anticommons effect, such an effect could result from differences in the characteristics of articles represented in the different age-cohort categories. In order to precisely identify the source of these differences, the next section undertakes a systematic econometric analysis of these patterns.

---

<sup>9</sup> Appendix A explores these conditional means in more detail by breaking them out according to whether a public sector author is associated with article. Notably, among patented articles, there is a distinct citations advantage associated with those articles with PUBLIC SECTOR AUTHOR equal to 0.

## V. Empirical Analysis

The empirical analysis proceeds in several stages. We first compare the cross-sectional differences in citation rates between patented and unpatented articles allowing for controls for other article characteristics. We then turn to the principal empirical exercise, examining how the citation rate changes with the grant of formal intellectual property under a variety of control structures and examining how the effect of IP manifests itself over time. Finally, we examine how our *diffs-in-diffs* estimates vary with article and patent characteristics. As is well known, appropriate analysis of citation data must take account of its discrete nature and the skewness of the underlying citation distribution. As such, except where noted, all specifications employ a negative binomial regression. Coefficients are reported as *incident rate ratios*, and so should be interpreted as a multiplicative effect on the expected number of citations received in a given year resulting from a one unit change in a regressor (i.e., the null hypothesis of no effect yields a coefficient of 1.0).

We begin in Table 4 with two negative binomial specifications, each of which includes YEAR and AGE fixed effects. While PATENTED articles are associated with a higher rate of citation without any additional controls, this effect is reduced in magnitude and statistically insignificant when controls for # AUTHORS, US AUTHOR, and PUBLIC SECTOR AUTHOR are included. In other words, the overall citation advantage observed in the conditional means can be (primarily) explained by differences in observable article characteristics.<sup>10</sup>

This brief cross-sectional analysis motivates our main empirical analysis of the impact of *patent grant conditional on article quality*. We begin in (5-1) with a regression, in which the dependent variable is equal to the natural log of FORWARD CITATIONS + 1; this OLS regression includes fixed effects for each calendar year, article age, and article fixed effects. Though this specification does not account for the nature of citations as skewed count data, the results provide support for nearly a 10% decline in the rate of citation after patent grant (significant at the 10% level). The remainder of this table turns back to negative binomial regression. In (5-2) and (5-3), we estimate both the overall difference between patented and non-patented articles and the marginal impact of being in the post-grant phase. For each of these specifications, we include the same article controls included in (4-2); as well, whereas (5-2)

---

<sup>10</sup> Though it is not the principal focus of the analysis, several unreported specifications confirm this qualitative finding. Though PATENTED is significant when controls are not included, the estimated effect becomes both smaller and imprecise when even a single observable article characteristic is included.

employs a parametric treatment for the impact of AGE (including a linear and quadratic term), (5-3) includes a complete set of age and year fixed effects. In both cases, patented articles enjoy approximately an 18% overall citation boost; however, the receipt of a patent effectively erases this advantage, resulting in an estimated 15-17% decline in the expected citation rate. Finally, in our differences-in-differences estimate in (5-4), we include a separate fixed effect for every article (and a complete set of year and age fixed effects); as such, these estimates are identified exclusively off the within-article contrasts between pre-grant and post-grant citation levels (and after accounting for the impact of age and year). According to this specification, the estimated post-grant decline is estimated at just over 9% with a high level of statistical significance. Moreover, these results are robust to alternative specifications and sample definitions.<sup>11</sup> At face value, these estimates provide concrete evidence for the existence of an anticommons effect. Simply put, the impact of a given piece of scientific research on scientific research declines after IP rights are granted. At an intuitive level, the results suggests that between 1 in 11 and 1 in 6 researchers building on a given paper may forego that specific research direction after IPR are granted over that research.

The evidence for an anticommons effect is even stronger when we examine the relative citation rate for an article in the years preceding and following the patent grant. To do so, we estimate a fixed effects negative binomial regression including specific dummy variables for each year preceding and following the patent grant date. In Figure D, we display the coefficients from that specification. The pattern is encouraging (though individual coefficients are estimated imprecisely). Though articles receive an up-tick in citations in the year prior to patent grant, there is an intermediate decline in the year of patent grant which continues steadily through four years after the date of patent grant. The difference between the pre-grant average and the average four years after patent grant is more than 25%, suggesting that the impact of IPR accumulates over time; taking the pattern in Figure D at face value, the “size” of the anticommons effect becomes a sizable deterrent to future research efforts once patent protection is in force for several years.

Though encouraging, it is important to emphasize that the results so far have relied on the presence of the control group of non-patented articles. Table 6 explores the impact of excluding

---

<sup>11</sup> For example, the results are similar (though a bit more noisy) if broken out by individual years of publication, or if (the incomplete record for) 2003 citations are included as an additional year of data.

the control group and relying exclusively on a sample composed of articles which are ultimately associated with a patent-paper pair. In the absence of a control group, it is difficult to disentangle the impact of patent grant from the impact of age on the citation rate, particularly if one simultaneously controls for fixed quality differences across articles. Indeed, if the amount of time elapsed between publication and patent grant were constant across articles, the impact of post-grant would not be separately identified from a set of age fixed effects. In other words, after controlling for age effects (and excluding a control group), the estimate of the impact of patent grant relies on differences across articles in the amount of time between publication and patent grant. The first column of Table 6 suggests that the anticommons effect is indeed identified by this variation. This specification controls for article quality by including #AUTHORS, US AUTHOR, and PUBLIC SECTOR AUTHOR, as well as a full set of age and year fixed effects. The coefficient on POST-GRANT is quantitatively and statistically significant, with a magnitude similar to that found in the second and third columns of Table 5 (patent grant is associated with a 14-17% decline in the citation rate).

The last two columns of Table 6 address these issues a bit more directly by including article fixed effects. The difference between the two specifications is that, whereas (6-2) includes the entire sample, (6-3) only focuses on articles which have at least one US author and at least one author with a public sector affiliation. There are two reasons for drawing this contrast. First, the timing of patent grants associated with articles with only foreign authors tends to be much longer, and to have a more muted effect than the other samples. As well, from a policy perspective, the focus of the anti-commons debate is focused on scientific research conducted with US government funds that involves university or other public sector researchers. The contrast is instructive. While there is no significant effect using the full sample, there is a statistically and quantitatively significant effect in the restricted sample, similar to the magnitude of the effect found in the earlier specifications.<sup>12</sup> Overall, these results both provide useful evidence for the anticommons effect while highlighting the difficulty of distinguishing the anticommons effect from the impact of article age.

---

<sup>12</sup> Of course, if we assume that the control group of non-patented articles is valid (as in Table 5), this concern does not apply since the age and year effects are separately identified by those articles.

### *The Role of Institutional Affiliations and Location*

In Table 7, we move beyond the baseline analysis to examine how the effect of patent grant on citations varies with the affiliations of article authors. Each specification consists of a negative binomial with fixed effects for articles, article age, and calendar year (as in (5-4)). In (7-1) and (7-2), we examine how the impact of patent grant varies with whether an article has any public authors and whether it has any private authors. The results are striking. While there is an insignificant effect when there is at least one private author (or all private authors), the decline in citations is most closely associated with the case where there are *no* private authors. It is possible that articles associated with private sector researchers are *assumed* to have intellectual property associated with them, and so the actual receipt of a patent has little impact on behavior by those who might build on that research. In contrast, until the patent grant date, researchers put a lower probability on purely public sector authored papers having a patent associated with them; the effect of IP in this case is to chill follow-on research endeavors. Interestingly, there is much less difference in the citation decline depending on whether the article is associated with at least one US author; moreover, (7-4) demonstrates that the anticommons effect is identified even if one *excludes* all articles without at least one US author. In other words, though the impact of patent grant is muted for those articles without at least some public sector involvement (perhaps due to expectations about IPR by follow-on researchers), these results do not provide evidence for the hypothesis that the sensitivity of citation to patent grant depends on author location.

### *Does the Impact of Patent Grant Depend on Patent Type?*

At its heart, the anticommons hypothesis is premised on the difficulty of using contracting methods to overcome the limitations on access to knowledge covered under patent protection. While contractual difficulties may be relatively unimportant when there is only a single patent or area of prior art at issue, transaction costs (and bargaining break-down) are likely to be more important when there are multiple competing claims limiting access to knowledge and materials. Table 8 explores this hypothesis by examining the interaction between the POST-GRANT dummy and our two measures of the degree of prior art cited in the patent (PATENT BACKWARD CITATIONS and PATENT BACKWARD REFERENCES). In both cases, there is a quantitatively large and statistically significant interaction. A one standard deviation

increase in the level of prior patent art citation made by a patent-paper pair is associated with an additional 6% decline in the expected citation rate after the patent is issued.

Table 9 shifts the focus from the extent of the patent thicket to the scope and nature of the patent grant. Though # CLAIMS is an extremely imperfect (and noisy) measure of the strength of patent rights, it does provide at least a weak measure of patent scope; evidence for the anticommons effect is strengthened by the fact that the predicted decline of citations as the result of patent grant is larger as the number of patent claims increases. In contrast, though a key tenet of the anticommons theory is that the effects are particularly salient for research tool patents, there is no evidence that the impact of patent grant depends on the nature of the invention.

#### *Does the Source of Research Funding Matter?*

Finally, we assess whether the source of funding for research mediates the sensitivity of the citation rate to patent grant. In (10-1), we compare the coefficients associated with patents for which at least one of the assignees is a public sector organization versus those that are exclusively assigned to private firms. In (10-2), we divide out the impact of patent grant according to whether the patent applicant acknowledges Federal funding for the research upon which the patent is based. In both cases, one cannot distinguish a separate impact according to the source of research funding. However, since policy interventions aimed at the anticommons effect would be primarily focused on research funded by public sector organizations or the Federal government itself, it is important to emphasize the anti-commons effect may exist in organizations beyond the direct province of the most likely policy interventions.

## **VI. Conclusions**

This paper provides the first systematic empirical test of the anti-commons hypothesis. Our empirical approach exploits the fact that the duality of knowledge is captured in the phenomena of patent-paper pairs (and that the granting of patents occurs only after a substantial lag). Our evidence suggests that knowledge duality is a quite important phenomenon (nearly 50% of articles published in *Nature Biotechnology* are associated with a pair) and that for scientific knowledge exhibiting this duality and being subject to both Open and Private Property regimes, the granting of IPR is associated with a significant but modest decline in knowledge accumulation as measured by forward citations. Moreover, the decline associated with patent

grant becomes more pronounced with the number of years elapsed since the grant, specifically for papers more closely associated with university or public research, and for patents that cite a high level of prior art. Overall, we are able to reject the null hypothesis that IPR have no impact on the diffusion and use of associated scientific knowledge. However, we also find evidence that articles which are associated with patenting do continue a high level of citation, similar to the level associated with articles that never receive patent protection.

These patterns provide a novel perspective on the economic consequences of increasing the scope of IPR to cover knowledge which had traditionally been produced under the norms of Open Science. Rather than simply serving to facilitate a “market for ideas,” IP may indeed restrict the diffusion of scientific research. However, erecting a barrier to the costless use of knowledge does not eliminate all use of that knowledge. The demand for specific research findings by follow-on researchers is relatively inelastic, limiting the empirical impact of the anti-commons effect.

Moreover, it should be emphasized that our evidence for the anti-commons effect captures only one aspect of the impact of IP on dual use knowledge. IPR may enhance incentives for (unobserved) research (particularly by private sector organizations) or lead to more effective commercialization (which is far more costly than the basic research component itself). Without a detailed accounting of the size of these effects versus the effect identified here, it is impossible to infer the elements of an optimal policy.

The importance of adjudicating this debate should no be understated. Rather than a fringe activity, dual-use knowledge is increasingly the norm in scientific research. Biotechnology and the life sciences have grown enormously as a share of overall research activities, and similar questions arise in the context of areas such as nanotechnology and open source software. Because biotechnology and related disciplines simultaneously offer the potential for fundamental scientific discoveries *and* commercial breakthroughs, traditional justifications for IPR and for the norms of Open Science are open to question. By providing a window into the impact of IP rights on the diffusion of scientific knowledge, this paper offers some insight into the policy tradeoffs associated with the interaction between public science and private commercialization.

## References

- Adams, J. (1990) "Fundamental Stocks of Knowledge and Productivity Growth," *Journal of Political Economy*, 98, 673-702.
- Argyres, N. and J. Liebskind (1998). "Privatizing the Intellectual Commons: Universities and the Commercialization of Biotechnology Research," *Journal of Economic Behavior and Organization*..
- Arora, A. 1995. Licensing Tacit Knowledge: Intellectual Property Rights And The Market For Know-How", *Economics of Innovation and New Technology*, 1995, Vol. 4, 41-49.
- Arora, A., Andrea Fosfuri, Alfonso Gambardella (2001) "Markets for Technology: Economics of Innovation and Corporate Strategy", Cambridge, MA: MIT Press, December 2001
- Arrow, Kenneth. "Economic Welfare and the Allocation of Resources for Invention." In *The Rate and Direction of Inventive Activity*, edited by Richard R. Nelson, 609-25. Princeton, NJ: Princeton University Press, 1962.
- David Blumenthal, Nancyanne Causino, Eric Campbell, and Karen Seashore Louis (1996). "Relationships between Academic Institutions and Industry in the Life Sciences — An Industry Survey" *New England Journal of Medicine*, Volume 334:368-374.
- Bush, V. (1945). *Science: The Endless Frontier, A Report to the President*. Washington, D.C., U.S. Government Printing Office.
- Campbell, Eric G., Brian R. Clarridge, Manjusha Gokhale, Lauren Birenbaum, Stephen Hilgartner, Neil A. Holtzman, and David Blumenthal. 2002. Data Withholding in Academic Genetics: Evidence from a National Survey. *Journal of the American Medical Association* 287 (4):473 - 479.
- Cockburn, Iain, and Rebecca Henderson "Absorptive Capacity, Coauthoring Behavior, and the Organization of Research in Drug Discovery." *Journal of Industrial Economics*, 1998, 46(2), pp. 157-182.
- Crespi, R. S. 'Patents on Genes: Can the Issues be clarified?' (1999/2000) 5 *BioScience Law Review*, 199 – 204.
- Dasgupta, Partha. and David, Paul. 1994. "Towards a New Economics of Science" *Research Policy* (23), 487-521.
- David, Paul. 2003. "Can 'Open Science' be Protected from the Evolving Regime of IPR Protections?" Stanford Working Paper #03-011.
- David, Paul. 2001a. "From Keeping 'Nature's Secrets' to the Institutionalization of 'Open Science'" Stanford Working Paper #01-006.
- David, Paul. 2001b. "Will Building Good Fences really make Good Neighbors in Science" Stanford Working Paper #01-005.
- David, Paul. 2000. "The Digital Technology Boomerang: New Intellectual Property Rights Threaten Global "Open Science" Stanford Working Paper #00-016.
- Ducor P., 2000. Intellectual property. Coauthorship and coinventorship. *Science* 289(5481), 873-875.
- Eisenberg, R. 1996. "Public Research and Private Development: Patents and Technology Transfer in Government-Sponsored Research." *Virginia Law Review* :1663-1727.
- Etzkowitz, H., 1998. The norms of entrepreneurial science: cognitive effects of the new university-industry linkages. *Research Policy* 27, 823-833.
- Galambos L. and J. Sewell. 1995. *Networks of Innovation: Vaccine Development at Merck, Sharp & Dohme, and Mulford, 1895-1995*. (New York: Cambridge University Press.
- Gambardella, A. 1995. Science and Innovation. Cambridge University Press, Cambridge UK.

- Gans, J. and S. Stern 2000. Incumbency and R&D Incentives: Licensing the Gale of Creative Destruction," (with. J. Gans), *Journal of Economics and Management Strategy*, 2000
- Geison, G. L. 1995. The Private Science of Louis Pasteur. Princeton, NJ: Princeton University Press.
- Hall, B.H. and R.H. Ziedonis (2001). "The Patent Paradox Revisited: An Empirical Study of Patenting in the US Semiconductor Industry, 1979-95." *RAND Journal of Economics*, 32(1): 101-128.
- Hardin, G. 1968. "The Tragedy of the Commons," *Science*, 162(1968):1243-1248.
- Heller, M. 1998. "The Tragedy of the Anticommons: Property in the Transition from Marx to Markets," 111 *Harvard Law Review* 621 (1998).
- Heller, M. and R. Eisenberg. 1998. "Can patents deter innovation? The anticommons in biomedical research", *Science*, vol. 280.
- Henderson, R., A. Jaffe, M. Trajtenberg. 1994. Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations. *Quarterly Journal of Economics*, August 1993, Vol. 434, pp. 578-598
- Jensen, Richard and Marie C. Thursby. 2001. "Proofs and prototypes for sale: The licensing of university inventions." *American Economic Review*, 91:1, pp. 240-59.
- Kenney M., 1986. Biotechnology: The University-Industrial Complex. New Haven: Yale University Press.
- Kimball, D. 2001. "Patenting Genes: Risks and Rewards versus Politics and Policies", Christie Parker & Hale LLP Publications.
- Kline S, and Nathan Rosenberg, 1986, "An Overview of Innovation," in *The Positive Sum Strategy: Harnessing Technology for Economic Growth*, eds. R. Landau and N. Rosenberg, National Academy Press, Washington (DC).
- Kremer, M. 1997. "Patent Buy-Outs: A Mechanism for Encouraging Innovation," NBER, Working Paper no. 6304.
- Krimsky, Sheldon. 2003. *Science and the Private Interest*. Rowman-Littlefield Publishing Co., August 2003.
- Lawson, Charles. 2002. Patenting genes and gene sequences and competition: patenting at the expense of competition. *Federal Law Review* 30:97-133.
- Lemley, M.A. and Allison, J.R. (2002). "The Growing Complexity of the United States Patent System," *Boston University Law Review*, 82, p. 77-144.
- Lessig, L. 2002. *The Future of Ideas: The Fate of the Commons in a Connected World*. Random House.
- Levin, R., A. Klevorick, R. Nelson, S. Winter. 1987. Appropriating the returns from industrial R&D. *Brookings Papers Econom. Activity* 3 783-831.
- Li, M. 1997. Use of a modified bacteriophage to probe the interactions between peptides and ion channel receptors in mammalian cells. *Nature Biotechnology* 15:6 (559-563)
- Merges, R., Mennell, P.S. and M. Lemley (2003), *Intellectual Property in the New Technological Age*, Aspen Publishers.
- Merges, R. and Nelson, R. R. 1990. "On the Complex Economics of Patent Scope," *Columbia Law Review*, Vol. 90.
- Merges, R.P. and Nelson, R.R. (1994). On Limiting or Encouraging Rivalry in Technical Progress: The Effect of Patent Scope Decisions. *Journal of Economic Behavior and Organization*, 25 (1), 1-24.
- Merton, R. *The Sociology of Science: Theoretical and Empirical Investigations*. Edited by Norman Storer. Chicago: University of Chicago Press, 1973.

- Mowery, David, Richard Nelson, Bhaven Sampat, and Arvids Ziedonis. 2001. "The Growth of Patenting and Licensing by U.S. Universities: An Assessment of the Effects of the Bayh-Dole Act of 1980." *Research Policy* 30:99-119.
- Murray, F. 2002. "Innovation as overlapping Scientific and Technological trajectories: exploring tissue engineering", *Research Policy*, Vol. 31.
- Narin, F., Olivastro, D., 1992. Status Report: Linkage Between Technology and Science. *Research Policy* Vol 21(3).
- Nature Biotechnology 1996. Editorial: A Broad and Inclusive Enterprise. March 1996 Volume 14 Number 3 p235.
- Nelkin, D. 1984. *Science as Intellectual Property: Who Controls Scientific Research?* New York: MacMillan Publishing Co.
- Nelson, R.R. (1959) "The Simple Economics of Basic Scientific Research," *Journal of Political Economy* 67(3), 297-306.
- Orsenigo, L. 1989. *The Emergence of Biotechnology: Institutions and Markets in Industrial Innovation.*, Pinter Publishers, London.
- Powell, W., Koput, K., Smith-Doerr, L. 1996. Inter-organizational Collaboration and the locus of Innovation: Networks of learning in Biotechnology. *Administrative Science Quarterly* 41, 116-145.
- Romer, Paul M, 1990. "Endogenous Technological Change," *Journal of Political Economy*, Vol. 98 (5) pp. S71-102.
- Rosenberg, N. (1974). "Science, Invention, and Economic Growth." *The Economic Journal* 84(333): 90-108
- Shapiro, C. (2001), "Navigating the Patent Thicket: Cross Licenses, Patent Pools, and Standard Setting," in A.B. Jaffe, J. Lerner and S. Stern (eds.), *Innovation Policy and the Economy*, Vol. 1.
- Stern, Scott. 1999. "Do Scientists Pay to be Scientists". NBER Working Paper w7410.
- Stokes, Donald. 1997. *Pasteur's Quadrant: basic science and technological innovation.* Washington D.C.: The Brookings Institution.
- Straus, Joseph, Henrik Holzapfel, and Matthias Lindenmeir. 2002. Empirical Survey on "Genetic Inventions and Patent Law". Munich.
- Thursby, G., Thursby M. 2003. University Licensing Under Bayh-Dole: What are the Issues and Evidence? Working Paper accessed from [opensource.mit.edu/papers/Thursby.pdf](http://opensource.mit.edu/papers/Thursby.pdf).
- Walsh, John P., Ashish Arora, and Wesley M. Cohen. 2002. The Patenting and Licensing of Research Tools and Biomedical Innovation: US National Academies' Science, Technology and Economic Policy Board.
- Weitzman, M. "Free Access vs. Private Ownership as Alternative Systems for Managing Common Property." *Journal of Economic Theory* 8, no. 2 (June 1974): 225-234.
- Williams-Jones, B. (2002) "History of a Gene Patent: Tracing the Development, Marketing, and Application of Commercial BRCA Testing in Canada" *Health Law Journal*, Vol. 10, p. 121-144.
- Ziedonis, R. (2004). "Don't Fence Me In: Fragmented Markets for Technology and the Patent Acquisition Strategies of Firms," *Management Science*, June.
- Zucker, L., Darby, M., Brewer, M., 1998. Intellectual Human Capital and the Birth of U.S. Biotechnology Enterprises *American Economic Review* 88, 290-306.
- Zozulya, S., Lioubin, M., Hill, R.J., Abram, C., Gishizky, M. 1999. Mapping Signal transduction pathways by phage display. *Nature Biotechnology*, 17:2 (1193-1198).

**TABLE 1**  
**VARIABLES & DEFINITIONS**

VARIABLE	DEFINITION	SOURCE
<b>CITATION-YEAR CHARACTERISTICS</b>		
FORWARD CITATIONS <sub>jt</sub>	# of Forward Citations to Article <i>j</i> in Year <i>t</i>	Science Citation Index (SCI)
YEAR <sub>t</sub>	Year in which FORWARD CITATIONS are received	SCI
AGE <sub>jt</sub>	YEAR – PUBLICATION YEAR	Nature Biotechnology (NB)
<b>PUBLICATION CHARACTERISTICS</b>		
PUBLICATION YEAR <sub>j</sub>	Year in which article is published	NB
# AUTHORS <sub>j</sub>	Count of the number of authors of Article <i>j</i>	NB
US AUTHOR <sub>j</sub>	Dummy variable equal to 1 if <i>at least</i> one of the institutional affiliations associated with Article <i>j</i> is in the US; 0 otherwise	NB
PUBLIC AUTHOR <sub>j</sub>	Dummy variable equal to 1 if <i>at least</i> one of the institutional affiliations associated with Article <i>j</i> is a university or government organization; 0 otherwise	NB
PRIVATE AUTHOR <sub>j</sub>	Dummy variable equal to 1 if <i>at least</i> one of the institutional affiliations associated with Article <i>j</i> is a biotechnology or pharmaceutical firm; 0 otherwise	NB
TOTAL CITATIONS <sub>j</sub>	# of FORWARD CITATIONS from publication date to 2003	SCI
<b>PATENT CHARACTERISTICS</b>		
PATENTED <sub>j</sub>	Dummy variable equal to 1 if Article is associated with a patent issued by the USPTO prior to October, 2003	USPTO
GRANT YEAR <sub>j</sub>	YEAR in which PATENT has been granted	USPTO
PATENT, POST-GRANT <sub>jy</sub>	Dummy variable equal to 1 if PATENTED = 1 and YEAR > GRANT YEAR	USPTO
# INVENTORS	Count of the number of inventors listed in the granted patent associated with Article <i>j</i> ; 0 if PATENTED = 0.	USPTO
PATENT LAG <sub>j</sub>	Days elapsed between patent priority application and grant dates, 0 if PATENTED = 0	USPTO
PATENT BACK CITATIONS <sub>j</sub>	Count of the number of citations to <i>patented prior art</i> included in the granted patent associated with Article <i>j</i> ; 0 if PATENTED = 0	USPTO
PATENT BACK REFERENCES <sub>j</sub>	Count of the number of citations to <i>non-patent prior art</i> included in the granted patent associated with Article <i>j</i> ; 0 if PATENTED = 0	USPTO
# CLAIMS	Count of the number of allowed claims; 0 if PATENTED = 0	USPTO
TOOLS PATENT <sub>j</sub>	Dummy variable equal to 1 if primary patent classes are associated with <i>research tools</i> (as opposed to composition of matter patents); 0 if PATENTED = 0	USPTO; author verification
PUBLIC SECTOR ASSIGNEE <sub>j</sub>	Dummy variable equal to 1 if <i>at least</i> one of the assignees on the patent associated with Article <i>j</i> is a university or government organization; 0 otherwise	USPTO; author verification
GOVT. FUNDED <sub>j</sub>	Dummy variable equal to 1 if patent disclosure reports a Federal interest, indicating Federal funding of research upon which patent is based; 0 otherwise	USPTO

**TABLE 2**  
MEANS & STANDARD DEVIATIONS

VARIABLE	N	MEAN	STANDARD DEVIATION	MIN	MAX
<b>CITATION-YEAR CHARACTERISTICS</b>					
FORWARD CITATIONS	1688	9.35	12.29	0	184
TOTAL CITATIONS	1688	54.74	58.84	2	523
CITATION YEAR	1688	1999.95	1.52	1997	2002
AGE	1688	2.05	1.52	0	5
<b>PUBLICATION CHARACTERISTICS (N=340 total articles)</b>					
PUBLICATION YEAR	340	1998.03	0.83	1997	1999
# AUTHORS	340	5.89	3.20	1	20
US AUTHOR	340	0.59	0.49	0	1
PUBLIC SECTOR AUTHOR	340	0.90	0.30	0	1
PRIVATE SECTOR AUTHOR	340	0.32	0.47	0	1
<b>PATENT CHARACTERISTICS (N = 340 total articles, 169 articles associated with USPTO patents)</b>					
PATENTED	340	0.50	0.50	0	1
GRANT YEAR*	169	2000.54	1.71	1996	2002
PATENT, POST-GRANT^	1688	0.24	0.43	0	1
# INVENTORS*	169	3.04	1.59	1	8
PATENT LAG*	169	1126.07	480.10	238	3714
PATENT BACK CITATIONS*	169	7.26	13.10	0	79
PATENT BACK REFERENCES*	169	28.19	37.25	0	226
# CLAIMS*	169	21.12	15.00	2	94
TOOLS PATENT*	169	0.58	0.49	0	1
PUBLIC SECTOR ASSIGNEE*	169	0.65	0.48	0	1
GOVT FUNDED*	169	0.29	0.45	0	1

\* Summary statistics for these measures is calculated only for those article for which PATENTED = 1 and is weighted by Article (i.e., N = 162).

^ Summary statistics for PATENT, POST-GRANT is calculated over all articles, weighted by citation year

**TABLE 3**  
**MEANS CONDITIONAL ON PATENT STATUS**

	<b>NO PATENT</b>	<b>PATENTED</b>
# Publications	171	169
FORWARD CITATIONS	8.96	10.04
# AUTHORS	5.76	6.03
US AUTHOR	0.53	0.65
PUBLIC SECTOR AUTHOR	0.93	0.86
PRIVATE SECTOR AUTHOR	0.25	0.38

**TABLE 4**

**CROSS-SECTIONAL RESULTS**

	<b>NEGATIVE BINOMIAL Dep Var = FORWARD CITATIONS (Coefs reported as incident rate ratios)</b>	
	<b>(4-1) Baseline Model*</b>	<b>(4-2) With Publication Controls*</b>
PATENTED	<b>1.156 (0.063)</b>	1.085 (0.056)
# AUTHORS		<b>1.034 (0.008)</b>
US AUTHOR		<b>1.266 (0.064)</b>
PUBLIC SECTOR AUTHOR		0.930 (0.108)
<i>Parametric Restrictions</i>		
Age FEs = 0	# Restrict 5 $\chi^2$ 219.83 p-value 0.000	# Restrict 5 $\chi^2$ 236.83 p-value 0.000
Year FEs = 0 <sup>~</sup>	# Restrict 5 $\chi^2$ 20.80 p-value 0.001	# Restrict 5 $\chi^2$ 16.76 p-value 0.005
R-squared		
Log-likelihood	-5257.86	-5230.73
P-value of Chi	0.00	0.00
# of Observations	1688	1688

\* Robust standard errors are in parentheses.

<sup>~</sup> Year FEs included for 1998-2002 (1997 is excluded).

**TABLE 5**  
**THE IMPACT OF PATENT GRANT:**  
**DIFFERENCES-IN-DIFFERENCES ESTIMATES**

	OLS: DepVar = ln(FORWARD CITATIONS)	NEGATIVE BINOMIAL Dep Var = FORWARD CITATIONS (Coefs reported as incident rate ratios)		
	(5-1) Marginal Impact, with Article FEs	(5-2) Selection and Marginal Effects w/ Controls	(5-3) Selection and Marginal Effects w/. AGE & YEAR FE	(5-4) Marginal Effects, with Article FEs
<b>PATENT CHARACTERISTICS</b>				
PATENTED		<b>1.177</b> (0.078)	<b>1.181</b> (0.081)	
PATENTED, POST-GRANT	<b>-0.097</b> (0.056)	<b>0.834</b> (0.069)	<b>0.841</b> (0.070)	<b>0.911</b> (0.039)
<b>CONTROL VARIABLES</b>				
# AUTHORS		<b>1.035</b> (0.008)	<b>1.035</b> (0.008)	
US AUTHOR		<b>1.282</b> (0.066)	<b>1.272</b> (0.065)	
PUBLIC SECTOR AUTHOR		0.906 (0.108)	0.922 (0.065)	
AGE	<b>0.959</b> (0.033)	<b>3.151</b> (0.199)		
AGE*AGE	<b>-0.160</b> (0.006)	<b>0.834</b> (0.011)		
PUBLICATION YEAR = 1998		<b>1.251</b> (0.073)		
PUBLICATION YEAR = 1998		<b>1.263</b> (0.089)		
<i>Parametric Restrictions</i>				
Article FEs = 0	# Restrict 340 $\chi^2$ 10882.97 p-value 0.00			# Restrict 338 $\chi^2$ 11038.17 p-value 0.00
Age FEs = 0			# Restrict 5 $\chi^2$ 237.92 p-value 0.00	# Restrict 5 $\chi^2$ 480.30 p-value 0.00
Year FEs = 0 <sup>~</sup>			# Restrict 5 $\chi^2$ 17.40 p-value 0.004	# Restrict 5 $\chi^2$ 8.97 p-value 0.11
<b>Regression Statistics</b>				
Log-likelihood		-5266.35	-5227.61	-4028.87
P-value of Chi		0.00	0.00	0.00
R-Squared	0.75			
# of Observations	1688	1688	1688	1688

\* Robust standard errors are in parentheses.

<sup>~</sup> Year FEs included for 1998-2002 (1997 is excluded).

**TABLE 6**  
**PATENTED ARTICLES ONLY**

	<b>NEGATIVE BINOMIAL</b> <b>Dep Var = FORWARD CITATIONS</b> <b>(Coefs reported as incident rate ratios)</b>		
	<b>(6-1)</b> Marginal Effects w/. AGE & YEAR FE	<b>(6-2)</b> Marginal Effects, with Article FEs (Full Sample)	<b>(6-3)</b> Marginal Effects, with Article FEs (US Public Sector Authors Only)
<b>PATENT CHARACTERISTICS</b>			
PATENTED, POST-GRANT	<b>0.857</b> <b>(0.079)</b>	0.978 (0.048)	<b>0.877</b> <b>(0.053)</b>
<b>CONTROL VARIABLES</b>			
# AUTHORS	<b>1.043</b> <b>(0.010)</b>		
US AUTHOR	<b>1.430</b> <b>(0.103)</b>		
PUBLIC SECTOR AUTHOR	0.893 (0.131)		
<i>Parametric Restrictions</i>			
Article FEs = 0		# Restrict 167 $\chi^2$ 5740.68 p-value 0.00	# Restrict 167 $\chi^2$ 4784.79 p-value 0.00
Age FEs = 0	# Restrict 5 $\chi^2$ 124.19 p-value 0.00	# Restrict 5 $\chi^2$ 321.26 p-value 0.00	# Restrict 5 $\chi^2$ 235.05 p-value 0.00
Year FEs = 0 <sup>~</sup>	# Restrict 5 $\chi^2$ 4.28 p-value 0.510	# Restrict 5 $\chi^2$ 51.01 p-value 0.00	# Restrict 5 $\chi^2$ 93.56 p-value 0.00
<b>Regression Statistics</b>			
Log-likelihood	-2667.07	-2039.791	-1099.151
P-value of Chi	0.00	0.00	0.00
# of Observations	849	849	462

\* Robust standard errors are in parentheses.

<sup>~</sup> Year FEs included for 1998-2002 (1997 is excluded).

**TABLE 7**  
**DIFFS-IN-DIFFS RESULTS**  
**BY INSTITUTIONAL OR NATIONAL AFFILIATION\***

	NEGATIVE BINOMIAL REGRESSIONS			
	Dep Var = FORWARD CITATIONS			
	(7-1) No Public Author v. Public Author	(7-2) No Private Author v. Private Author	(7-3) No US Author v. US Author	(7-4) US Author Articles Only
<b>ARTICLE CHARACTERISTICS</b>				
PATENTED, POST-GRANT				<b>0.900</b> <b>(0.049)</b>
PATENTED, POST-GRANT * NO PUBLIC AUTHOR	0.952 (0.106)			
PATENTED, POST-GRANT * PUBLIC AUTHOR	<b>0.903</b> <b>(0.041)</b>			
PATENTED, POST-GRANT * NO PRIVATE AUTHOR		<b>0.882</b> <b>(0.045)</b>		
PATENTED, POST-GRANT * PRIVATE AUTHOR		0.955 (0.062)		
PATENTED, POST-GRANT * NO US AUTHOR			0.913 (0.058)	
PATENTED, POST-GRANT * US AUTHOR			<i>0.910</i> <i>(0.047)</i>	
<b>CONTROL VARIABLES</b>				
Article Pair FEs = 0	# Restrict 338 $\chi^2$ 10963.05 p-value 0.00	# Restrict 338 $\chi^2$ 11205.55 p-value 0.00	# Restrict 338 $\chi^2$ 11022.10 p-value 0.00	# Restrict 198 $\chi^2$ 9398.12 p-value 0.00
Age FEs = 0	# Restrict 5 $\chi^2$ 475.44 p-value 0.00	# Restrict 5 $\chi^2$ 469.29 p-value 0.00	# Restrict 5 $\chi^2$ 469.89 p-value 0.00	# Restrict 5 $\chi^2$ 328.69 p-value 0.00
Year FEs = 0 <sup>~</sup>	# Restrict 5 $\chi^2$ 2.51 p-value 0.77	# Restrict 5 $\chi^2$ 1.00 p-value 0.96	# Restrict 5 $\chi^2$ 1.34 p-value 0.93	# Restrict 5 $\chi^2$ 3.75 p-value 0.59
<b>Regression Statistics</b>				
Log-likelihood	-4028.70	-4028.32	-4028.87	-2429.43
P-value of Chi	0.00	0.00	0.00	0.00
# of Observations	1688	1688	1688	994

\* Robust standard errors are in parentheses.

<sup>~</sup> Year FEs included for 1998-2002 (1997 is excluded).

**TABLE 8**  
**DIFFS-IN-DIFFS RESULTS**  
**WITH PATENT CHARACTERISTIC INTERACTIONS \***

	NEGATIVE BINOMIAL REGRESSIONS Dep Var = FORWARD CITATIONS		
	(8-1) PATENT LAG Interaction	(8-2) PATENT BACKWARD CITATIONS Interaction	(8-3) PATENT BACKWARD REFERENCE Interaction
<b>“DIRECT” EFFECT OF PATENT GRANT</b>			
PATENTED, POST-GRANT	1.043 (0.097)	0.952 (0.045)	0.987 (0.049)
<b>INTERACTION EFFECTS</b>			
PATENTED, POST-GRANT * PATENT LAG	<b>0.999</b> <b>(0.0000)</b>		
PATENTED, POST-GRANT * PATENT BACKWARD CITATIONS		<b>0.995</b> <b>(0.001)</b>	
PATENTED, POST-GRANT * PATENT BACKWARD REFERENCES			<b>0.997</b> <b>(0.0000)</b>
<b>CONTROL VARIABLES</b>			
Article Pair FEs = 0	# Restrict 340 $\chi^2$ 10965.59 p-value 0.00	# Restrict 340 $\chi^2$ 10838.08 p-value 0.00	# Restrict 340 $\chi^2$ 11070.68 p-value 0.00
Age FEs = 0	# Restrict 5 $\chi^2$ 535.25 p-value 0.00	# Restrict 5 $\chi^2$ 480.83 p-value 0.00	# Restrict 5 $\chi^2$ 471.55 p-value 0.00
Year FEs = 0 <sup>~</sup>	# Restrict 5 $\chi^2$ 76.93 p-value 0.000	# Restrict 5 $\chi^2$ 18.24 p-value 0.00	# Restrict 5 $\chi^2$ 0.89 p-value 0.97
<b>Regression Statistics</b>			
Log-likelihood	-4027.97	-4026.88	-4024.95
P-value of Chi	0.00	0.00	0.00
# of Observations	1688	1688	1688

\* Robust standard errors are in parentheses.

<sup>~</sup> Year FEs included for 1998-2002 (1997 is excluded).

**TABLE 9**  
**DIFFS-IN-DIFFS RESULTS**  
**BY TYPE OF IPR\***

	<b>NEGATIVE BINOMIAL</b>	
	<b>Dep Var = FORWARD CITATIONS</b> (Coefs reported as incident rate ratios)	
	<b>(9-1)</b> NUMBER OF CLAIMS Interaction *	<b>(9-2)</b> No Tools Versus Research Tools Patents*
PATENTED, POST-GRANT	1.017 (0.063)	
PATENTED, POST-GRANT * NUMBER OF CLAIMS	<b>0.995</b> <b>(0.002)</b>	
PATENTED, POST-GRANT * NON-RESEARCH TOOL		0.899 (0.055)
PATENTED, POST-GRANT * RESEARCH TOOL		0.918 (0.048)
<i>Parametric Restrictions</i>		
Article Pair FEs = 0	# Restrict 338 $\chi^2$ 11139.06 p-value 0.00	# Restrict 340 $\chi^2$ 11037.62 p-value 0.00
Age FEs = 0	# Restrict 5 $\chi^2$ 484.17 p-value 0.00	# Restrict 5 $\chi^2$ 481.60 p-value 0.000
Year FEs = 0 <sup>~</sup>	# Restrict 5 $\chi^2$ 10.15 p-value 0.07	# Restrict 5 $\chi^2$ 11.00 p-value 0.05
Log-likelihood	-4026.73	-4028.83
P-value of Chi	0.00	0.00
# of Observations	1688	1688

\* Robust standard errors are in parentheses.

<sup>~</sup> Year FEs included for 1998-2002 (1997 is excluded).

**TABLE 10**  
**DIFFS-IN-DIFFS RESULTS**  
**BY NATURE OF RESEARCH FUNDING \***

	<b>NEGATIVE BINOMIAL</b> <b>Dep Var = FORWARD CITATIONS</b> <b>(Coefs reported as incident rate ratios)</b>	
	<b>(10-1)</b> No Public Sector Assignee Versus Public Sector Assignee	<b>(10-2)</b> Non Govt. Funded Versus Govt. Funded
PATENTED, POST-GRANT * NO PUBLIC SECTOR ASSIGNEE	<i>0.899</i> <i>(0.056)</i>	
PATENTED, POST-GRANT * PUBLIC SECTOR ASSIGNEE	0.921 (0.049)	
PATENTED, POST-GRANT * NON-GOVT. FUNDED		<i>0.913</i> <i>(0.046)</i>
PATENTED, POST-GRANT * GOVT. FUNDED		0.907 (0.061)
<i>Parametric Restrictions</i>		
Article Pair FEs = 0	# Restrict 338 $\chi^2$ 11053.87 p-value 0.00	# Restrict 338 $\chi^2$ 10919.50 p-value 0.00
Age FEs = 0	# Restrict 5 $\chi^2$ 517.05 p-value 0.00	# Restrict 5 $\chi^2$ 471.38 p-value 0.000
Year FEs = 0 <sup>~</sup>	# Restrict 5 $\chi^2$ 48.68 p-value 0.07	# Restrict 5 $\chi^2$ 11.37 p-value 0.05
Log-likelihood	-4028.82	-4028.86
P-value of Chi	0.00	0.00
# of Observations	1688	1688

\* Robust standard errors are in parentheses.

<sup>~</sup> Year FEs included for 1998-2002 (1997 is excluded).

**APPENDIX A  
CONDITIONAL MEANS  
BY PATENTED & AUTHOR AFFILIATION \***

	No Patent		Patented	
	No Public Author	Public Author	No Public Author	Public Author
# PUBS	12	159	23	146
FORWARD CITATIONS	9.38	8.92	12.02	9.73
# AUTHORS	7.75	5.61	7.61	5.78
US AUTHOR	0.67	0.52	0.83	0.62

# FIGURE A

## THE STOKES MODEL

	Consideration of Use?		
		NO	YES
Quest for Fundamental Understanding?	NO		Pure Applied Research (Edison)
	YES	Pure Basic Research (Bohr)	Use-inspired / translational Basic research (Pasteur)

FIGURE B  
CITATION DISTRIBUTION

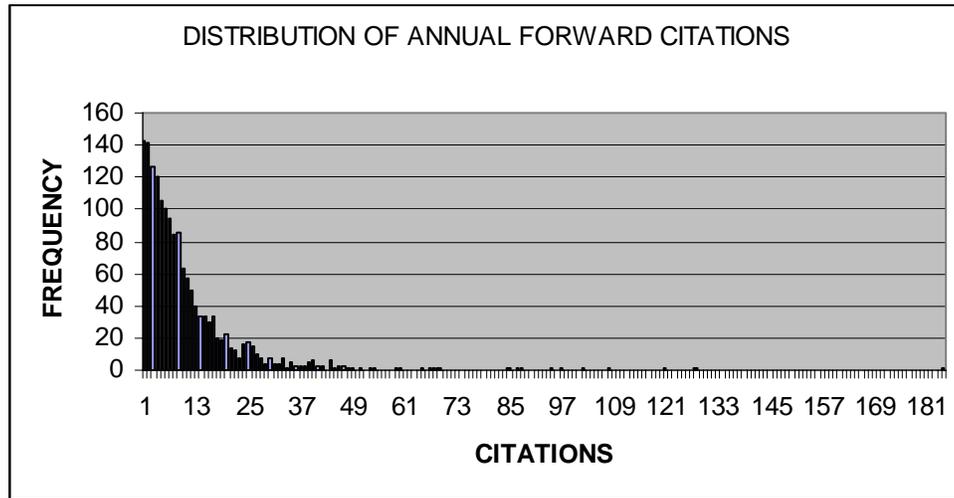
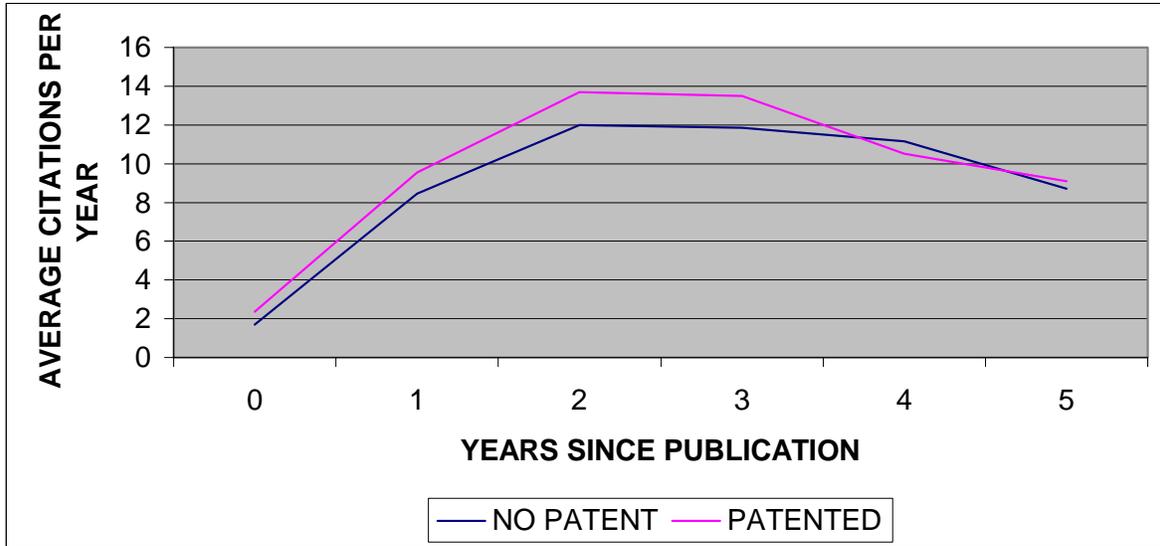


FIGURE C  
CITATIONS BY TYPE BY AGE



**FIGURE D**  
**IMPACT OF PATENT GRANT ON FORWARD CITATIONS,**  
**BY YEAR BEFORE AND AFTER PATENT GRANT**  
**(NEGATIVE BINOMIAL WITH ARTICLE FE<sub>s</sub>)**

