

Do formal intellectual property rights hinder the free flow of scientific knowledge? An empirical test of the anti-commons hypothesis

Fiona Murray^{a,*}, Scott Stern^b

^a MIT Sloan School, 50 Memorial Drive E52-567, Cambridge, MA 02142, United States

^b Northwestern University and NBER, United States

Received 16 June 2005; accepted 19 May 2006

Available online 16 January 2007

Abstract

Although many scholars suggest that IPR has a positive effect on cumulative innovation, a growing “anti-commons” perspective highlights the negative role of IPR over scientific knowledge. At its core, this debate is centered on how intellectual property rights over a given piece of knowledge affect the propensity of future researchers to build upon that knowledge in their own scientific research activities. This article frames this issue around the concept of dual knowledge, in which a single discovery may contribute to both scientific research and useful commercial applications, and finds evidence for a modest anti-commons effect. A key implication of dual knowledge is that it may be simultaneously instantiated as a scientific research article and as a patent. Such patent-paper pairs are at the heart of our empirical strategy. We exploit the fact that 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 anti-commons perspective suggests that the citation rate for 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 anti-commons effect (the citation rate after the patent grant declines by approximately 10 to 20 percent). 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. © 2007 Elsevier B.V. All rights reserved.

JEL classification: O300; O330; O340; L330

Keywords: Anti-commons; Intellectual property; Academic science; Pasteur’s Quadrant

* Corresponding author. Tel.: +1 617 258 0628; fax: +1 617 253 2660.

E-mail address: fmurray@mit.edu (F. Murray).

1. Introduction

In the early 1980s, Professor Philip Leder, recently recruited to head the new Genetics department at the Harvard Medical School, developed the first genetically engineered mouse with cancer, dubbed the Oncomouse. Leder and his post-doc Tim Stewart had used novel transgenic techniques to insert an oncogene into a mouse embryo; the result was a mouse that was highly susceptible to cancer. In his pursuit of a deeper understanding of cancer, Leder came to recognize that “it could serve a variety of different purposes, some purely scientific others highly practical” (Kelves, 2002, p. 83). This research was published in *Cell* in 1984, and in 1988 a broad patent for the Oncomouse was granted by the US Patent Office (USPTO). The Oncomouse patent was more controversial than most; not only was the Oncomouse the first living mammal to be patented, but Du Pont, as the patent’s exclusive licensee, aggressively enforced the property rights. Du Pont’s strategy included demands for “reach-through” rights and review of publications that used the Oncomouse in further scientific research (Murray, 2006). Ideas that are simultaneously of value as a scientific discovery and inventive construct, such as the Oncomouse, are frequently generated in the disciplines that underpin modern biotechnology (Stokes, 1997).

Dual-purpose ideas provide their originators with multiple disclosure choices, including an option to publish research in the scientific literature *and* obtain intellectual property rights (IPR) over that knowledge. In academia, an increasing number of scientists have chosen this path of dual knowledge disclosure, which we describe as patent-paper pairs (Murray, 2002; Ducor, 2000). Patent-paper pairs are scientific articles and individual patents that disclose (and serve as a property right over) the same underlying “piece” of knowledge. Patent-paper pairs are thus more than simply a reflection of the rise in patenting by academics of knowledge unrelated (or only tangentially related) to their scientific research. Rather, by embedding the same piece of knowledge in two distinct institutional regimes, patent-paper pairs instantiate the expansion of formal intellectual property rights over knowledge that was traditionally disseminated solely through the norms of scientific publication.

The increased use of intellectual property rights (IPR) in scientific research has sparked a vigorous academic and policy debate. On the one hand, a significant amount of research has highlighted the benefits of IPR (Kitch, 1977; Arora et al., 2001). Recent empirical research on commercial discoveries suggests that IPR may facilitate the creation of a market for ideas, encourage further investment in ideas with commercial potential, and mitigate disincentives to disclose and exchange knowledge which might otherwise remain secret (Merges and Nelson, 1990, 1994; Arora et al., 2001; Gans and Stern, 2000). Indeed, within the context of university research (particularly publicly funded university research), it has been suggested that IPR offers important incentives to move nascent discoveries out of the “ivory tower” and into commercial practice. In other words, IPR may enhance the ability of society to realize the commercial and social benefits of a given discovery (Kitch, 1977; Hellman, 2007). However, a more recent “anti-commons” perspective argues that the expansion of IPR (in the form of patents and/or copyrights) is “privatizing” the scientific commons and limiting scientific progress (Heller and Eisenberg, 1998; Argyres and Liebskind, 1998; David, 2001b). Specifically, the anti-commons hypothesis states that IPR may inhibit the free flow and diffusion of scientific knowledge and the ability of researchers to build cumulatively on each other’s discoveries (Heller and Eisenberg, 1998; David, 2003, 2000; Lessig, 2002; Etzkowitz, 1998; Krinsky, 2003). Taken together, the impact of IPR on future progress in the broader scientific community remains open to debate.

Building on several key case examples in the literature, this article frames this debate around the concept of dual knowledge as instantiated in patent-paper pairs. Within this framework we then report a novel empirical strategy to evaluate the salience of IPR on the cumulative impact of scientific knowledge. Our strategy exploits the existence of patent-paper pairs as both the instantiation of the expansion of intellectual property rights over dual knowledge and as a concrete empirical starting point from which to identify the impact of such rights on the rate of scientific knowledge diffusion.

Our approach is to compare patterns of scientific citations to scientific articles that are part of patent-paper pairs, relative to citation patterns for articles that are *not* part of a patent-paper pair (but are similar along other dimensions). This allows us to evaluate several key hypotheses at the center of the anti-commons debate. First, we evaluate whether citation patterns are different for scientific research that is ultimately patented. In other words, to what extent does published scientific knowledge disclosed as a patent-paper pair differ in its future cumulative impact on public domain research (as measured by forward citations to the publication) from papers that are similar in topic, published in the same journal in the same time period, but never receive IPR? Second, we take advantage of *patent grant delay*. While publication lags are usually modest (on the order of a few months), patent grant delays are substantial (in most cases IPR is granted 2–4 years after initial application). Consequently, scientific knowledge associated with a patent-paper pair diffuses under two distinctive institutional environments, a pre-grant period where no IP rights are present and a post-grant period in which specific property rights have been granted. To the extent that a patent grant comes as a “surprise” to at least some potential follow-on researchers, this difference allows us to ask how does the *grant* of formal patent rights over such knowledge influence the trajectory of forward citations and therefore the impact of the scientific research findings in the public domain.

The “experiment” afforded by the combination of patent-paper pairs and patent grant delay allows for a set of precise tests motivated by the anti-commons perspective: if the grant of intellectual property hinders the ability of researchers to build (in the public domain) on a given piece of knowledge, and the patent grant itself is “news” to the broader scientific community, then the citation rate to the scientific publication disclosing that knowledge should be lower than for scientific publications with no IP and should fall after formal property rights are granted. Of course, such an analysis must control for the fact that citation patterns vary with the underlying quality of the article and with the time elapsed since publication. Our use of patent grant delay allows us to do so. Specifically, by observing a *given* piece of knowledge in two different institutional environments, we are able to evaluate how differences in the institutional environment affect the diffusion of a given piece of knowledge, including a fixed effect for each article in our sample. To evaluate the anti-commons hypothesis, we examine how the grant of IPR *changes* the citation rate to scientific articles, accounting for fixed differences in citation rates across articles and relative to the trend in citation rates for articles with similar characteristics.

The analysis employs two distinct (and complementary) approaches to the identification of the impact of patent grant on scientific citation. In the bulk of the analysis, we evaluate how the citation rate changes after patent grant, controlling for the trend in citation identified by articles that do not receive IPR. Moreover, to address the potential for selection of articles into patenting, we also explore a more nuanced empirical strategy that exploits the variation in patent grant delay among patented articles. Specifically, we examine the impact of patent grant on scientific citation relying exclusively on differences across patented articles in the time it takes to receive a patent. Overall, our approach employs a differences-in-differences estimator to evaluate the impact of IPR on the diffusion of scientific knowledge.

Our sample is composed of 340 peer-reviewed scientific articles appearing between 1997 and 1999 in *Nature Biotechnology*, a high-quality scientific publication and perhaps the leading publication for research exhibiting knowledge duality in the life sciences. The incidence of patent-paper pairs is quite high within this sample: for just under 50 percent of the scientific articles in our sample, a US patent has been granted over the knowledge covered in that publication. For those articles that ultimately receive a patent, there is a significant lag between scientific publication and patent grant (on average, more than 3 years). We exploit these data to establish three core findings. First, published articles also associated with formal IP are more highly cited than those whose authors choose not to file for patents; however, most of this boost is accounted for by observed characteristics such as author location and number of authors on the article. Second, there is robust evidence for a quantitatively modest but statistically significant anti-commons effect; across different specifications, the article citation rate declines by approximately 10 to 20 percent after a patent grant. Third, the anti-commons effect is particularly salient for articles with public sector co-authors.

We would like to be cautious in our interpretation. On the one hand, though the size of the effect is modest, the approach and results do seem to provide empirical evidence consistent with the anti-commons effect. With that said, the use of citation data is only a noisy indicator of the impact of any given piece of research, and our approach does not separately identify any potentially positive impact of IPR on research incentives (from the perspective of the original inventor). Moreover, we have not identified the specific institutional mechanism by which patent grant both surprises and influences researcher behavior. For example, as suggested by qualitative research, it is possible that the reduction in citation after patent grant does not arise from IPR *per se* but from how IP rights are enforced in specific circumstances. Example of this scenario include onerous licensing terms and delays associated with bargaining and negotiation over material transfer agreements among research organizations (Walsh et al., 2003, 2005; Murray, 2006).

The remainder of the paper is as follows. The next section reviews the economic foundations of our understanding of the impact of IPR on scientific knowledge and lays out the specific anti-commons hypothesis, framed around the nature and institutional foundations of dual knowledge. Section 3 develops the empirical test and predictions of the anti-commons hypothesis. After a review of the data in Section 4, Section 5 presents our empirical findings. We conclude with implications for future empirical work and innovation policy in Pasteur's Quadrant.

2. The impact of formal intellectual property rights on scientific knowledge

The relationship between formal intellectual property rights such as patents and the disclosure and diffusion of scientific knowledge has been a long standing policy concern.¹ The majority of the literature in this area makes a clear distinction between basic and applied research, with basic research being focused on questions of fundamental scientific interest and applied research focused on questions of usefulness and applications (Stokes, 1997). Scholarship in the “new” economics of science compares alternative institutional arrangements associated with these distinctive types of knowledge and primarily focuses on the cases where both the institutional regime and disclosure choices are exogenous (Dasgupta and David, 1994). Basic research is undertaken in

¹ In contrast, much knowledge produced in industry is often maintained as a trade secret or disclosed only through the patent system. However, an increasing amount of industry-produced knowledge is also disclosed through scientific publication, often to serve specific strategic purposes.

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

Fig. 1. The Stokes (1997) model.

the institutions of Open Science and disclosed through publication; applied research takes place in a Private Property regime and is disclosed through patents (to the extent that disclosure takes place at all).

Knowledge production in the Open regime (“Science”) is characterized (in the stylized case) by a distinctive set of economic incentives for cumulative knowledge production, including the adoption of norms that facilitate full disclosure and diffusion of knowledge. This system includes the recognition of scientific priority 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 leverages the public goods nature of basic research and therefore promotes cumulative innovation and “standing on the shoulders of giants.”² In contrast, the incentives that govern the private property rights regime pay little attention to the basis of a researcher’s ultimate impact on follow-on research. Instead, these incentives depend on the degree to which a researcher can *exclude* others and thus 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; Scotchmer, 1996).

While concise in its formulation of the relationship between the nature of knowledge and incentives for its production and distribution, this analytical framework fails when knowledge has *both* basic and applied value. By highlighting its *duality*, Stokes reformulated the traditional distinction between basic and applied research; a single discovery could simultaneously possess both applied and basic characteristics (Fig. 1). Stokes highlighted the potential for dual-use knowledge by proposing two dimensions along which research might be motivated (rather than the traditional approach which places knowledge on a linear dimension from basic to applied). One dimension pertained to whether knowledge was produced for fundamental scientific interest. The second was whether knowledge was produced for commercial gain (or in response to practical problems). Two “traditional” quadrants were then identified: the first was knowledge produced only for scientific interest, known as “Bohr’s Quadrant”, and the second with knowledge produced primarily for commercial gain (referred to as “Edison’s Quadrant”). Moreover, Stokes suggested that a significant share of all scientific research combined the two motives, resulting in “use-inspired basic research” being produced in “Pasteur’s Quadrant.” Like Leder’s Oncomouse research in the 1980s, Pasteur’s fundamental insights into microbiology served both as a foundation for the germ theory of disease and had practical application for cholera and rabies (Geison, 1995; Stokes, 1997).

² 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, 2004; Murray, 2002).

For knowledge generated in Pasteur's Quadrant, disclosure becomes endogenous: scientists have a well-defined set of disclosure choices including publication in scientific journals and application for protection through formal intellectual property rights. Moreover, these choices are *not* mutually exclusive—a given piece of knowledge can both be disclosed through scientific publication and be protected by intellectual property rights. The phenomena of “dual knowledge – dual disclosure” is perhaps most clearly instantiated in the context of “patent-paper pairs” – linked scientific articles and individual patents which disclose (and serve as a property right over) the same underlying “piece” of knowledge.³ Patent-paper pairs are thus more than simply a reflection of the rise in patenting in academia over knowledge unrelated (or only tangentially related) to scientific research. Rather, they embed the same piece of knowledge in two distinct institutional regimes, each of which is governed by specific (and potentially opposing) rules for access and attribution.

2.1. Debating the role of formal IPR

In recent years, a considerable theoretical and qualitative literature has been developed highlighting the impact of formal IPR on research, disclosure, and diffusion in Pasteur's Quadrant. The literature particularly focuses on the impact of IPR over knowledge generated in academia and traditionally placed in the public domain through scientific publication.

On the one hand, theoretical and empirical research highlights the salutary effects arising from the introduction of IPR, even as it protects scientific discoveries. First, patents provide incentives for research investment, and while many discoveries in Pasteur's Quadrant may have been pursued in the absence of IPR, it is possible that the enhanced incentives from IPR attract the entry of high-quality scientific researchers into specific research fields. Second, even if there is no impact on the incentive to produce knowledge *per se*, patents may usefully facilitate the commercialization of that knowledge and help to bridge the university-industry divide. Patents may contribute to the effective functioning of the market for ideas (Merges and Nelson, 1990, 1994; Arora et al., 2001; Gans and Stern, 2000), as well as enhance the incentives and efficiency of the process by which academic researchers search and match with potential downstream partners (Kitch, 1977; Jensen and Thursby, 2001; Hellman, 2007).⁴

In contrast to this literature, others have highlighted the potential tax imposed by intellectual property rights over scientific knowledge traditionally disclosed only through publication. In particular, a body of scholarship has emerged around the “anti-commons” effect. Scholars who take this view argue that the imposition of IP rights over areas traditionally maintained in the public commons undermines the process of cumulative scientific discovery (Heller, 1998; Heller and Eisenberg, 1998; David, 2003). 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 to future knowledge production and hence cumulative innovation. Because IP can serve to exclude follow-on researchers from exploiting scientific discoveries, the anti-commons hypothesis posits

³ There is also evidence for increased secrecy and delay in publication as a result of increased economic incentives in academia, the disclosure requirements of patent applications in the US and Europe, and the strictures placed on academics when their research is funded by private industry (Blumenthal et al., 1996; Campbell et al., 2000). The growing privatization of basic research may also be associated with refocusing of research agendas (Thursby and Thursby, 2003; David, 2003) and an increased potential for bias in research results (Nelkin, 1984; Krinsky, 2003).

⁴ It is also possible that the introduction of patents over “upstream” research can have subtle effects on research incentives and research direction for both academic and corporate researchers (Aghion et al., 2005).

that the privatization of the scientific commons will impose a “tax” on the use of prior scientific knowledge.⁵ If protected by IPR, the impact of an individual piece of knowledge on follow-on research by others is diminished, potentially resulting in a lower equilibrium level of on-going cumulative research productivity. Overall, the anti-commons theory therefore suggests that while individual researchers have strong incentives to take advantage of the protections afforded by formal IP rights, the scientific community as a whole benefits from the free dissemination and diffusion of knowledge.

2.2. *The anti-commons debate in molecular biology*

While not limited to life sciences, many of the issues that currently animate the IPR discussion surround the interaction between public and private knowledge exploitation in areas related to biotechnology (Kenney, 1986; Orsenigo, 1989; Powell et al., 1996; Gambardella, 1995; Galambos and Sewell, 1995). In this setting, emphasis on dual knowledge as embodied in dual institutions reflects at least three related forces: the expanding promise of biotechnology, reductions in the costs of academic patenting, and increases in the scope of IPR over knowledge produced in the life sciences. The rise in dual knowledge in this field can be traced back to the early 1970s; most key milestones in the field, with noted exceptions, have been disclosed as patent-paper pairs⁶:

- The techniques of recombinant DNA provided insights into the machinery of the cell but also laid the foundation for the production of recombinant therapeutic proteins (Cohen et al., 1973).
- The discovery of the HIV retrovirus both established the cause of AIDS (and scientific routes for investigating how to treat AIDS) while providing the foundation for a commercial blood test for HIV infection. The controversy over the scientific credit for the discovery of AIDS was resolved, at least in part, through the settlement of the patent dispute between the French and American governments over the blood test patent (Murray and Stern, 2006).
- The Oncomouse simultaneously provided great insight into the sources of cancer while becoming a model for investigating cancer therapies (Stewart et al., 1984).
- The discovery of RNA interference represented a further step towards explaining DNA replication and the foundation of a potentially new therapeutic category (Zamore et al., 2000).
- The development of embryonic stem cells show us how cells develop but also have the potential to serve as novel therapeutics or the foundations of organ replacement (Thomson et al., 1998).

At the same time, policy shifts encouraged academics to claim IPR over their dual knowledge. Prior to this time, patent applications filed by universities on behalf of investigators required case-by-case negotiation of the assignment of patent rights and their subsequent licensing. The 1980 Bayh-Dole Act assigned IPR (generated using Federal funds) to universities along with a duty to

⁵ Furthermore, to the extent that IPR is narrow in scope and highly dispersed across individuals and institutions, fragmentation can impose a further tax in the form of significant transaction costs (Eisenberg, 1996; Shapiro, 2001; Hall and Ziedonis, 2001; Ziedonis, 2004).

⁶ Perhaps the most interesting exception to this pattern concerns the Nobel Prize-winning work on the development of hybridomas that allowed understanding of immune systems and the creation of monoclonal antibodies (Köhler and Milstein, 1975). As Köhler and Milstein submitted their ground breaking findings to *Nature*, they also submitted the manuscript to their funding agency (the Medical Research Council) with a proposal to file for a patent. However, the request was refused on the basis that “[i]t is certainly difficult for us to identify any immediate practical applications which could be pursued as a commercial venture, even assuming that publication had not already occurred” (<http://www.path.cam.ac.uk/~mrc7/mab25yrs/index.html> last accessed March 14, 2005).

license the patents and facilitate their translation and commercialization (Mowery et al., 2001).⁷ This clarification in institutional practice likely increased the propensity for dual knowledge to be disclosed through both scientific publications and patents. Nevertheless, individual researchers and institutions varied widely in their response to this new environment (Azoulay et al., 2007; Ding et al., 2006).

Third, there was a significant expansion in the *scope* of IPR available in the life sciences. After the 1980 *Diamond versus Chakrabarty* decision and the granting of the Oncomouse patent in 1988, IPR comprehensively covered the domain of genetically modified living organisms from bacteria to mammals (Kelves, 2002). In combination with the developments in the biotech industry, “universities were literally propelled into an awareness of the potential economic value of the technology that was being generated in their research programs” (Bremmer, 2001).

Thus the ground work was laid for the debate regarding the impact of formal IPR over scientific knowledge: increasingly, the duality of knowledge and the widening scope of patents meant that many scientific discoveries now encompassed patentable inventions. Together with more streamlined institutional rules for patenting and changes in university culture, this shifted the likely set of disclosure decisions of faculty to more frequently include patenting. In the period between 1989 and 1999, US Research One universities received over 6000 life science patents (Owen-Smith and Powell, 2003), and patent-paper pairs became an important disclosure phenomenon (Murray, 2002). In short, the life science commons were increasingly covered by intellectual property rights.

Economists, law and technology scholars, and policymakers have focused on key cases to highlight their concerns over the impact of this expansion in IPR (Heller and Eisenberg, 1998; David, 2001a; Campbell et al., 2002; Straus et al., 2002; Walsh et al., 2003, 2005). Consider the Oncomouse (Murray, 2006): the knowledge described by Leder’s *Cell* publication was also covered by a “paired” Oncomouse patent (4,736,866). This patent was licensed exclusively to Du Pont Corporation who, as the donor of a broad research gift to support Leder’s lab, actually participated in the decision to patent. Soon after the patents were granted, Du Pont imposed its property rights (on commercial and academic scientists) through a series of stringent licensing terms: “Reach-through” clauses that specified royalties for Du Pont for any discoveries or products developed using the Oncomouse, oversight and control of scientific publications, and limits on informal breeding and sharing of these mice among scientific colleagues.

These limits were seen as a severe impediment to scientific research and a contravention of the norms of the scientific community. Scientists were vocal in their opposition to Du Pont; led by Harold Varmus, senior scientists in mouse genetics met at Cold Spring Harbor and began talking of revolution. In response, the NIH stepped in and sanctioned a non-profit central facility (The Jackson Laboratory) to serve as a repository for genetically altered mouse strains. However, while the Jackson Lab could distribute the mice, the Oncomouse was still constrained by restrictive licensing terms. For example, Du Pont continued to contact scientists and demand review of their articles before publication, and “nobody was able to exchange materials” (Marshall, 2000). When Varmus became the Director of the NIH, he furthered his campaign to limit Du Pont’s actions

⁷ In this context, the traditional justification of IP rights over Open Science relates not to scientific knowledge accumulation but rather to cumulative commercial innovation. IP provides incentives for further commercial investment. IPR may also facilitate a “market for ideas” by encouraging the exchange and trade of knowledge particularly with private sector researchers (Merges and Nelson, 1994; Arora et al., 2001; Gans and Stern, 2000; Gambardella, 1995). For knowledge generated in academia the third leg of the traditional argument for IP (as an inducement for disclosure over secrecy) does not apply since there is a presumption of and evidence for disclosure in publications.

on the basis that they “seriously impede further basic research,” and in 1999 an agreement was reached between NIH and Du Pont that outlined the terms under which the technology could be used in research supported by the NIH. Du Pont commented that they “are committed to making this tool available to the academic community. . .while retaining our commercial rights” (NIH, 2000).

While the Oncomouse has become something of a touchstone for proponents of the anti-commons perspective, the salience of these concerns depends on their empirical relevance across the broader, more representative landscape of the life sciences (and beyond). Accordingly, the remainder of our analysis attempts to trace systematically the consequences of IPR on the diffusion of scientific knowledge in a large-scale sample.

3. The empirical framework

In adjudicating the positive and negative perspectives on the impact of IPR over scientific knowledge, we argue that the effects of IPR would be of particular concern if the tax that they impose on research productivity outweighs the direct incentive effect arising from the patentability of scientific research findings. While there has been little empirical work in assessing this delicate balance, these theoretical ideas have observable empirical implications: to the extent that the anti-commons effect is manifest in the scientific community, published ideas that are also covered by IPR (such as the Oncomouse) will be expected to have lower levels of cumulative scientific *impact*, relative to the impact that would be realized in the absence of IPR. Moreover, we expect the salience of this effect to be higher for certain types of researchers (e.g., for patent-pairs associated with public sector research) and certain types of discoveries (e.g., for patent-paper pairs over more complex or more important discoveries). This section first develops an empirical approach for evaluating the anti-commons effect and then links this framework to a number of specific theoretical predictions of the anti-commons theory.

The empirical challenge in documenting the anti-commons effect is straightforward. How can we compare the impact a given piece of knowledge would have *with* IPR versus the impact of that knowledge without IPR? For a given piece of knowledge *with* IPR, one cannot observe the counterfactual impact of that knowledge in the event that 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 IPR. A simple comparison between patented and non-patented knowledge may therefore be biased by unobserved heterogeneity. The key empirical challenge in adjudicating the impact of disclosure regimes is thus an identification problem; 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 compare the impact on future research for each case.

Our approach exploits the existence and nature of patent-paper pairs. A patent-paper pair is the dual instantiation of a given piece of knowledge as both a scientific research article and a patent. Consider the following example of research undertaken in the biology department at the Massachusetts Institute of Technology in the field of bacterial genetics and metabolism:

“A method has been developed for control of molecular weight and molecular weight dispersity during production of polyhydroxyalkanoates in genetically engineered organisms by control of the level and time of expression of one or more PHA synthases in the organisms. The method was demonstrated by constructing a synthetic operon for PHA production in *E. coli* . . . Modulation of the total level of PHA synthase activity in the host cell by varying

the concentration of the inducer . . . was found to effect the molecular weight of the polymer produced in the cell.” (U.S. Patent No. 5,811,272)

“A synthetic operon for polyhydroxyalkanoate (PHA) biosynthesis designed to yield high levels of PHA synthase activity in vivo was constructed . . . by positioning a genetic fragment . . . behind a modified synthase gene containing an *Escherichia coli* promoter and ribosome binding site. Plasmids containing the synthetic operon . . . were transformed into *E. coli* DH5 alpha and analyzed for polyhydroxybutyrate production. . . Comparison of the enzyme activity levels of PHA biosynthetic enzymes in a strain encoding the native operon with a strain possessing the synthetic operon indicates that the amount of polyhydroxyalkanoate synthase in a host organism plays a key role in controlling the molecular weight and the polydispersity of polymer. (Sim et al., 1997)

As outlined in these brief excerpts, the research described in both documents is based on a specific genetic modification of a bacterium (*E. coli*) designed to control the type and amount of particular chemicals (PHA) the bacteria might ordinarily produce. From the scientific perspective, the publication emphasizes that these experiments deepen our understanding of the genes that regulate particular chemical pathways in bacteria. However, as highlighted in the patent, they also provide practical techniques for the manipulation of bacteria and the optimization of their use as a source of useful biomaterials. In other words, this single discovery has been instantiated as both a publication emphasizing its scientific contribution and as a patent disclosure emphasizing its utility.

Our empirical approach exploits three key aspects of the phenomenon associated with the production of patent-paper pairs:

- (a) within the scope of published scientific research that is potentially patentable, a significant fraction of researchers choose to forego formal IPR,
- (b) among those who pursue formal IPR, there is usually a significant delay between the first publication of that knowledge within the scientific literature and the granting of formal IPR over that knowledge, and
- (c) future scientific citations are a noisy but useful measure of the impact of a scientific article on future scientific research.

Combining these factors, we exploit the existence of patent-paper pairs to evaluate the impact of IPR on the diffusion of scientific knowledge. The first step in our approach is to collect a sample of published scientific research articles that are of roughly similar “quality” (though we will account for quality variation among articles in our empirical framework) and that disclose knowledge that is potentially patentable (whether or not the researchers choose to apply for IPR). We draw a set of research articles, within a specific time frame, from a top-tier research journal that specializes in publishing dual knowledge discoveries, *Nature Biotechnology* (we discuss this choice further in Section 4). By limiting our sample to articles published within a narrow time window within a specialized journal, we ensure that all articles share some affinity in topic, and are subjected to the same peer-review process. By choosing a journal whose editorial policy is to focus on research that is simultaneously of scientific and commercial interest, we consider the publications in our dataset to be “at risk” of being associated with a USPTO patent and thus forming a patent-paper pair. As we discuss further in the next section, we are able provide qualitative support for our assumption of comparability between the articles within our sample with analysis of the article’s scientific content and review of the potential patentability of the knowledge disclosed in the non-patented

articles by an experienced intellectual property lawyer. Of course, the decision to patent a piece of dual knowledge is endogenous to the specific circumstances of individual researchers, including factors such as their institutional affiliation and their gender (Azoulay et al., 2007; Ding et al., 2006; Markiewicz and DiMinin, 2005; Agrawal and Henderson, 2002). Our procedure simply assumes that the impact of intellectual property on the use of knowledge by *follow-on* researchers is (conditionally) independent of the patent filing decision.

Within our sample of articles (only a subset of which are associated with a patent-paper pair), we then exploit the fact that for any patent-paper pair, there exists a patent grant delay—a substantial gap between the date of scientific publication and the date at which the associated patent is granted.⁸ This empirical technique exploits the insight that while publication in the scientific literature often occurs within 6 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 a 2–4-year time window). It is important to emphasize that patent grant delay is more than simply a matter of the timing of a *pro forma* administrative decision. During the time between application and grant, applicants and examiners undertake detailed negotiations about the scope and extent of the patent grant, so there is significant uncertainty about the extent of IPR prior to grant (Cockburn et al., 2003; Jaffe and Lerner, 2004). Perhaps more saliently, prior to the patent grant date, the patent applicant holds no formal IPR, and, in nearly all cases, cannot sue for infringement for activities undertaken during the pre-patent grant period. Finally, until 2001 (and thus for nearly all of the cases within our empirical work), USPTO patent applications remained *secret* until granted. In other words, for any given patent-paper pair, we observe the same “piece” of knowledge in two distinct institutional regimes: one associated with the pre-patent grant period with no formal intellectual property rights and then one in the post-patent grant period with patent rights.

Of course, the behavioral response to the patent grant date depends on the degree to which the patent grant (and associated enforcement activities) serves as “news” or as a “surprise”: to what extent were researchers aware of the impending grant, and how does patent grant (and associated post-grant enforcement) change behavior by raising the perceived “price” of building on a prior discovery? If follow-on researchers believe that a patent-paper pair is likely (as might occur, for example, for articles published by researchers from a for-profit company), the impact of patent grant on behavior is likely to be modest (since researchers will anticipate the potential for IPR in advance and incorporate these potential costs into their research decisions). On the other hand, if only a minority of university researchers engage in patenting behavior, the potential for a post-patent grant “surprise” is quite high. We argue that at least within academia, patent grant typically comes as a “surprise” to academic researchers. We ground this empirical strategy in the observation that in many cases, follow-on researchers are unaware of whether or not a particular discovery will ultimately be associated with a patent-paper pair (Walsh et al., 2003). Moreover, the nature of the “news” may not necessarily be the direct finding of granted patents but rather may come from behaviors around patent rights enforcement of the type observed in the Oncomouse case. For the purpose of the current analysis, we assume that at least some potential follow-on

⁸ The specifics of patent law regarding the precise timing of disclosure is complex, varies across countries, and has been subject to change. However, broadly speaking, under US patent law, inventors have a grace period of 12 months between public disclosure (for example in an academic publication or presentation) and filing for patenting covering that knowledge. Thus, the timing of the publication submission and patent application can vary among patent applications with some filed before publication and some after. In the Oncomouse case, a patent was filed on 22 June 1984, the article was submitted on 21 August 1984, the article was published in October 1984 and patent grant and publication occurred on 12 April 1988.

researchers may experience the patent grant as “news,” though of course the level of news may depend on factors such as the institutional affiliations or locations of the originating authors, the licensing strategies of licensees, and the amount of time elapsed since the patent grant date.⁹

We combine these elements into an empirical model by taking the annual rate of citation to an article in follow-on scientific research articles as a measure of the cumulative impact of that article through time. Though scientific citations are by no means a perfect measure of the impact of a specific article, they provide a useful (if noisy) index of the relative salience of research in follow-on research that is also disclosed in scientific publications (Cole, 2000). More precisely, if the granting of IPR raises the cost of building on a specific piece of knowledge, then the citation rate to IPR-linked scientific publications should decline post-patent grant. Of course, measuring the impact of scientific research using citations implies that we must account for its form as count data skewed to the right (and likely over-dispersed relative to a Poisson distribution). Therefore, except where noted, we employ a negative binomial model of the citations produced per year for each scientific article in our dataset. Moreover, the impact of a given piece of research, as measured by citations, will vary considerably with the underlying importance of the research discovery, with the time elapsed since initial publication, and with the year in which the citations are being considered. Finally, there are clearly delays in the process by which a patent grant would impact the rate of citation in the scientific literature since the impact of behavioral responses to patent grant (e.g., choosing another research avenue) will only be observed after these new research streams are themselves published. As such, our empirical specifications account for individual publication quality (through article fixed effects), for the effects of publication age and the overall rate of citation in a given year (through age and citation year fixed effects), and for the presence of a patent grant “window” during the year in which a patent is actually granted.¹⁰ Taken together our baseline empirical test for the anti-commons hypothesis is therefore:

$$\text{CITES}_{i,t} = f(\varepsilon_{i,t}; \gamma_i + \delta_{t-\text{pubyear}} + \beta_t + \psi_{\text{WINDOW}} \text{WINDOW}_{i,t} + \psi_{\text{POST-GRANT}} \text{POST-GRANT}_{i,t}) \quad (1)$$

where γ_i is a fixed effect for each article, $\delta_{t-\text{pubyear}}$ captures the age of the article, β_t is a fixed effect for each citation year, WINDOW is a dummy variable equal to one in the year in which a patent is granted and POST-GRANT is a dummy variable equal to one only for years after the patent grant year for an individual article.¹¹ Because we observe citations to a scientific publication both before

⁹ It is also theoretically possible that follow-on researchers will exploit the “window” between publication date and the patent grant date to take advantage of the absence of IPR over knowledge that will ultimately be protected. While we acknowledge this theoretical possibility, our fieldwork strongly mitigates against this strategy. Few laboratories are able to predict the precise timing of their research results, and this is unlikely to enter strategically and then exit a research area that will come under patent protection at an uncertain date in the future. As emphasized by Walsh et al. (2003, 2005), most academic research laboratories do not seem to undertake proactive monitoring of IPR grants in their research areas. Accordingly, it is more likely that the mechanism underlying a reduction in citation after a patent is granted is driven by the surprise arising from IPR rather than a result of a pre-grant research and publication strategy by rival research teams.

¹⁰ Several subtle issues, including an incidental parameters problem, arise in incorporating multiple fixed effect vectors into a negative binomial specification. We experimented with a range of alternative approaches, including the conditional negative binomial estimator (Hausman et al., 1984) and the fixed effects estimator (Allison and Waterman, 2002). All of our qualitative findings are unchanged across these different procedures; building on recent results about the relative size and importance of the small sample versus asymptotic bias arising in count data models, we report fixed effects results using robust standard errors (Allison and Waterman, 2002; Greene, 2004).

¹¹ Of course, it is possible to constrain this specification in order to, for example, obtain a separate estimate of the impact of unobserved heterogeneity between patented versus unpatented articles versus the impact of patent grant itself. In this

and after the patent is received (and because we observe a control group of similar publications that never receive a patent), we are able to identify how the temporal pattern of citations to a scientific publication changes as the result of patent grant.¹² 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. In other words, we test for the anti-commons 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.

At its core, anti-commons theory predicts that if the grant of intellectual property hinders the ability of researchers to build cumulatively on a given piece of knowledge, then the citation rate for the scientific publication disclosing that knowledge should fall after formal IP rights are granted over that knowledge. In addition to this baseline prediction, theory and the details of our empirical setting suggest a number of more nuanced hypotheses.

First, as we discussed earlier, the ability to infer the impact of IPR from the date of patent grant (and not before) depends importantly on the fact that the patent grant represents “news” to at least some follow-on researchers. As such, we expect our proposed test to be more salient among researchers with public sector author affiliations (e.g., government or university researchers). When a purely private research team publishes a scientific result, potential follow-on researchers may (rationally) anticipate that the knowledge and/or tools described in the article will eventually be covered by IPR. Accordingly, in contrast to the “news” associated with public sector patent grant, researchers may be deterred at the outset from research areas in which private sector researchers are publishing, whether or not the patent grant has yet been issued.

Furthermore, we would expect that the “news” from a particular patent grant likely diffuses among researchers over time, particularly in the few years after the grant. Taken together, and going beyond the simple use of a 1-year patent grant window measure, we expect that the impact of patent grant changes with the time elapsed since patent grant:

$$\begin{aligned} \text{CITES}_{i,t} = & f(\varepsilon_{i,t}; \gamma_i + \delta_{t-\text{pubyear}} + \beta_t + \sum_{k=1\dots r} \psi_{\text{PRE}_k} \text{PRE-GRANT}(k)_{i,t} \\ & + \sum_{l=0\dots s} \psi_{\text{POST}_l} \text{POST-GRANT}(l)_{i,t}). \end{aligned}$$

This specification allows us to evaluate both whether the impact of patent grant was anticipated (in terms of changes of citation behavior that preceded the actual patent grant date) as well as how the impact of a patent grant changes with the years elapsed since the patent grant.

Anti-commons theory also offers several additional predictions about the type of knowledge and research that might be more closely associated with the anti-commons effect. First, it is possible

case, one could simply replace the article fixed effect with a dummy variable for those articles that ultimately receive a patent: $\text{CITES}_{i,t} = f(\varepsilon_{i,t}; \delta_{t-\text{pubyear}} + \beta_t + \lambda \text{PATENTED}_i + \psi_{\text{WINDOW}} \text{WINDOW}_{i,t} + \psi_{\text{POST}} \text{POST-GRANT}_{i,t})$.

¹² This baseline analysis does assume that the age fixed effects associated with citation do not depend on whether a paper receives a patent. In particular, a key assumption of our base model (which we later relax) is that patented articles are not simply “shooting stars” (articles that, for exogenous reasons, experience a high rate of early citation followed by a rapid decline). In part, the “shooting star” hypothesis would be counterfactual to the most well-documented pattern of scientific citation, the so-called Matthew Effect, in which articles with a high rate of early citation tend to continue to receive an ever-higher rate of citation after a favorable early record (Merton, 1973). Also, in our robustness analysis, we actually rely exclusively on a sample of *patented* articles (with varying patent grant lag times) and find a similar pattern of results.

that the impact of patent grant may be strongest for publications by US-based authors. For articles written exclusively by non-US authors, the process of applying for and having a US patent granted is likely to be a very heterogeneous process, and enforcement activities may be uneven, reducing the salience of patent grant on follow-on research activities. Second, the anti-commons effect should be greatest for patent-paper pairs with particular patent characteristics. In particular, we expect stronger effects on patents associated with more complex technologies, broad scope, more prior art, and possibly those focused on research tools. The anti-commons hypothesis is premised on the difficulty of using contracting methods to access knowledge covered by patent protection. While contractual difficulties may be relatively unimportant when there is only a single patent or little prior art, transaction costs (and bargaining break-down) are likely to be more important when there are multiple competing claims limiting access to knowledge and materials (or when the complexity of the underlying and technology make it difficult to delineate the precise scope of individual claims). In other words, broad patents in complex areas with a high degree of prior art might be expected to be associated with denser patent thickets and hence stronger anti-commons effects. Likewise, prior researchers have suggested that patents over research tools (such as the Oncomouse, cell lines, gene probes) may be associated with larger anti-commons effects for at least two reasons. First, research tools are of broad relevance to many researchers and thus may impinge on many ongoing lines of research. Second, by and large, the uses of research tools are subject to a high degree of transparency in use: materials and tools are usually covered by material transfer agreements and other institutional arrangements. While contracts such as MTAs do facilitate the dissemination of tools, it is possible that the delays over the negotiation and terms of use for MTAs discourages the use of materials and tools that come under their purview (Walsh et al., 2003, 2005). Overall, the empirical framework provides a rich setting for evaluating a number of detailed hypotheses implied by the anti-commons argument. The remainder of this paper focuses on evaluating these empirical questions.

4. The data

4.1. 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 period 1997–1999. While the journal publishes scholarly material in a variety of formats, we confine our data to research articles that are defined by the editorial policies of the journal as “a substantial novel research study” (see *Nature Biotechnology*, 2006). 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. As noted above, biotechnology is a particularly salient arena for dual knowledge. 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).¹³ 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

¹³ 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.

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 continues to play this role. It is an important journal for academics in a wide variety of disciplines that underpin modern biotechnology and has risen to gain an impact factor of over 20 (surpassed in its field only by *Nature* and *Science*). In other words, research published in *Nature Biotechnology* is both high quality and “at risk” of serving as a simultaneous foundation for future scientific studies and commercial exploitation. As such, research published in *Nature Biotechnology* lies at the heart of the anti-commons debate.

For each of the 340 articles, we determined whether a patent associated (“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 some institutions patent under distinctive institutional names; for example, patents assigned to the University of Oxford are listed under the name of its separately incorporated technology licensing organization, 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 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.¹⁴ 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. Using this procedure, 169 of the 340 articles were found to be 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 5 years of publication.

Our empirical work relies on the fact that (a) the entire sample of articles is initially “at risk” of being patented, and (b) but for being patented the articles do not vary systematically along an unobserved dimension impacting the citation time trend. We check this 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

¹⁴ 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 criterion 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.

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 et al., 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 that might be reflected in scientific research articles, such as the discovery of new scientific principles, abstract ideas, and 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 percent (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, 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.

Third, 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. By and large, matched articles were found to be qualitatively similar, both in terms of their underlying scientific content and their potential for patentability.¹⁵

Finally, we gathered a series of variables associated with the publications: number of authors, the number and type of institutional affiliations, and the author age, rank and gender. On the one hand, there are some (marginally) significant differences between the unpatented and patented articles. For example, articles with private sector affiliations tend to have a higher patenting rate than those with exclusively public sector affiliations; articles with US-based authors tend to have a higher rate of USPTO patenting than those articles with exclusively foreign authors. Perhaps more importantly, there were no significant correlations between patenting and any aspect of the *nature* of the scientific research contribution (e.g., article length, the number of authors, or the number of backward references). While we do not undertake a detailed analysis of the determinants of patenting behavior (as explored by, among others, [Azoulay et al., 2007](#); [Ding et al., 2006](#); [Markiewicz and DiMinin, 2005](#); [Agrawal and Henderson, 2002](#)), our analysis of the

¹⁵ 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 that 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 Sugan (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.

impact of patenting exploits institutional and individual variation in the idiosyncratic incentives (or disincentives) for patenting itself.

Overall, given the nature of the sampling process and the qualitative comparisons, we assume that articles within our sample are each “at risk” of being patented (in a technical sense), and both observable and unobservable institutional factors are important in determining this outcome. In most of the empirical work, we therefore assume that after controlling for the overall quality of an article, the sample population is composed of articles at risk of being patented, and, but for a patent grant, the time trend in citations for all articles follows the same stochastic process.

The dataset is drawn from three distinct sources, each noted in Table 1. Article-specific characteristics are gathered from *Nature Biotechnology*: the 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). For each article we identified the lead author (the so-called Principal Investigator) for whom we gathered individual and institutional information from institutional websites and author resumes. The annual citation counts for each article (through October 2003) are calculated using the Science Citation Index Expanded (SCI).¹⁶ For each article associated with a USPTO patent, we collected a number of patent characteristics from the public USPTO database.

4.2. 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 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 = 9.34), 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 percent of the citation-years receiving either 0 or 1 citation, but also including one publication with an annual citation count equal to 184 (Fig. 2). 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 years.

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 examined a number of article-specific characteristics. These characteristics include the number of authors (# AUTHORS, mean = 5.89) as well as their institutional affiliations. For example, US AUTHOR is a dummy variable measuring whether *at least* one of the authors lists a US address (mean = 0.59). 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

¹⁶ Maintained by the Institute for Scientific Information (ISI), SCI records reference and citation information for nearly 6000 scientific and technical journals in approximately 150 disciplines.

Table 1
Variables and definitions

Variable	Definition	Source
Citation-year characteristics		
FORWARD CITATIONS _{jt}	# of forward citations to article <i>j</i> in year <i>t</i>	Science Citation Index (SCI)
YEAR _t	Year in which FORWARD CITATIONS are received	SCI
AGE _{jt}	YEAR <i>t</i> — PUBLICATION YEAR <i>j</i>	Nature Biotech (NB)
Publication characteristics		
PUBLICATION YEAR _j	Year in which article is published	NB
# AUTHORS _j	Count of the number of authors of article <i>j</i>	NB
US AUTHOR _j	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 _j	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 _j	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 _j	# of FORWARD CITATIONS from publication date to October 2003	SCI
Patent characteristics		
PATENTED _j	Dummy variable equal to 1 if article is associated with a patent issued by the USPTO prior to October 2003	USPTO
GRANT YEAR _j	YEAR in which PATENT has been granted	USPTO
PATENT GRANT WINDOW _j	Dummy variable equal to 1 in year patent is granted, 0 in all other years	USPTO
PATENT, POST-GRANT _j	Dummy variable equal to 1 if PATENTED = 1 and YEAR > GRANT YEAR	USPTO
# INVENTORS _j	Count of the number of inventors listed in the granted patent associated with article <i>j</i> ; 0 if PATENTED = 0.	USPTO
PUBLIC SECTOR ASSIGNEE _j	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 _j	Dummy variable equal to 1 if patent reports Federal interest, indicating Federal funding of research upon which patent is based; 0 otherwise	USPTO
PATENT LAG _j	Days elapsed between patent priority application and grant dates, 0 if PATENTED = 0	USPTO
# CLAIMS _j	Count of the number of allowed claims; 0 if PATENTED = 0	USPTO
PATENT BACK CITATIONS _j	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 _j	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
TOOLS PATENT _j	Dummy variable equal to 1 if primary patent classes are associated with <i>research tools</i> (vs. composition of matter patents); 0 if PATENTED = 0	USPTO; author verification

Table 2

Means and standard deviations

Variable	<i>N</i>	Mean	S.D.	Minimum	Maximum
Citation-year characteristics					
FORWARD CITATIONS	1688	9.34	12.49	0	184
TOTAL CITATIONS	1688	54.95	60.07	2	524
CITATION YEAR	1688	1999.95	1.52	1997	2002
AGE	1688	2.05	1.52	0	5
Publication characteristics (<i>N</i> = 340 total articles)					
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
Patent characteristics (<i>N</i> = 340 total articles, 169 articles associated with USPTO patents)					
PATENTED	340	0.50	0.50	0	1
GRANT YEAR ^a	169	2000.54	1.71	1996	2003
PATENT GRANT WINDOW ^b	1688	0.08	0.28	0	1
PATENT, POST-GRANT ^b	1688	0.16	0.37	0	1
# INVENTORS ^a	169	3.02	1.59	1	8
PUBLIC SECTOR ASSIGNEE ^a	169	0.65	0.48	0	1
GOVT FUNDED ^a	169	0.29	0.45	0	1
PATENT LAG ^a	169	1126.07	480.10	238	3714
# CLAIMS ^a	169	21.12	15.00	2	94
PATENT BACK CITATIONS ^a	169	7.26	13.10	0	79
PATENT BACK REFERENCES ^a	169	28.19	37.25	0	226
TOOLS PATENT ^a	169	0.58	0.49	0	1

^a Summary statistics for these measures calculated only for those articles for which PATENTED = 1, weighted by article (i.e., *N* = 169).

^b Summary statistics for PATENT, POST-GRANT and PATENT GRANT WINDOW are calculated over all articles, weighted by citation year.

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, 90 percent of the articles in the sample have at least one public sector author, and more than 30 percent have at least one private sector author.

Finally, our data includes information regarding the 169 patents associated with each patent-paper pair. The average date of patent grant (weighted by article) is mid-2000 with the average

Table 3

Means conditional on patent status

	Unpatented articles	Patented articles
# Publications	171	169
FORWARD CITATIONS	8.86	10.16
# 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




Fig. 2. Annual citation distribution.

lag between the patent application and patent grant date just over 3 years.¹⁷ Interestingly, while the patent grant lag is correlated with some traditional factors (such as the number of backward references in the patent), the overall patent grant delay is not robustly correlated with institutional affiliations (PUBLIC SECTOR AUTHOR, PRIVATE SECTOR AUTHOR, or US AUTHOR).

PATENTED (mean = 0.50) is a dummy variable equal to one for all citation-years associated with an article that receives a patent at any point during the sample period. We then define two more nuanced measures relating to the patent grant status of articles that are ultimately patented. PATENT WINDOW (mean = 0.08) is a dummy variable equal to one during the year that the patent is granted, while PATENT, POST-GRANT (mean = 0.16) is a dummy variable equal to one for all years *after* the patent is granted. The patent grant period is divided into the WINDOW and POST-GRANT measures for two reasons: (a) while we observe annual data, patent grant takes place at any time during the calendar and (b) the impact of patent grant on published scientific research likely operates with a reasonable lag. We investigate this second issue more fully by evaluating how the impact of a patent grant depends on the number of years since the patent grant occurred.

We also measure key patent characteristics. The patents in our sample are associated with a reasonably high number of listed inventors (# INVENTORS, mean = 3.04); interestingly, though a reasonably high number relative to other samples of patented inventions (see, for example, Jones (2005)), the average for # INVENTORS is just over half of the mean of # AUTHORS (similar in flavor to Ducor). We then define two measures that relate directly to the policy issues raised by the anti-commons debate. Similar to PUBLIC SECTOR AUTHOR, PUBLIC SECTOR ASSIGNEE (mean = 0.65) is 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 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 (GOVT FUNDED = 1). We are particularly interested in whether the impact of patent grant is salient for those patents associated with public sector authors (or assignees) and those for which Federal funding is disclosed.

In addition, to the extent that the anti-commons effect is more salient for more contentious areas of research, the impact of patent grant may be higher for patents associated with more complex or more important areas of technology. We investigate this possibility by examining

¹⁷ Some patents have been issued since the end of 2002. Inclusion or exclusion of these 23 articles from the analysis does not change any of the qualitative findings (primarily because our strategy relies on articles where we observe a change in the IP regime during those years where we observe a citation count).

several measures related to the complexity and importance of the technology. We first define two measures related to the outcomes of the patent process. PATENT LAG is the elapsed length of time (in days) between the date of the initial patent application and the patent grant date, and # CLAIMS is simply the number of allowed claims. Though both measures are noisy, each may be associated with the degree of complexity or the underlying importance of the technology.¹⁸ 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 anti-commons 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 Lemley and Allison (2002), the averages for both PATENT BACK CITATIONS (mean = 7.26) and PATENT BACK REFERENCES (mean = 28.19) are somewhat high. Finally, TOOLS PATENT is a dummy variable equal to one if the primary class for the patent is within the 435 and 800 patent classes (mean = 0.58). Out of the 11 three-digit patent classes represented across the patents within the sample, these two three-digit classes are most closely associated with processes and tools. Since research tools have been of particular concern within the anti-commons debate (relative to composition of matter patents), the TOOLS PATENT dummy (mean = 0.58) allows an assessment of whether the impact of patent grant on citation is greater for patents covering research tools and methods. The remaining patents are largely associated with composition of matter patents.

4.3. *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, just over an additional citation per article-year over the sample. However, this 10 percent average difference masks more substantial differences that manifest themselves over time. In Fig. 3, the average FORWARD CITATIONS are plotted by AGE (years since publication). During the year of publication and in the subsequent 3 years, PATENTED articles have a significant citation advantage, equivalent to nearly a 20 percent “boost” over the citation 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. 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 (although the differences in means in # AUTHORS and PUBLIC SECTOR AUTHOR are not significant).¹⁹

¹⁸ An emerging literature has focused on the determinants of patent grant lag itself, including Popp et al. (2004), Regibeau and Rockett (2003), and Harhoff and Wanger (2005). While there is some evidence in these studies that longer grant lags may be associated with more highly cited patents, each of these studies also highlights the high level of unexplained variation in patent grant lags.

¹⁹ Appendix B 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.

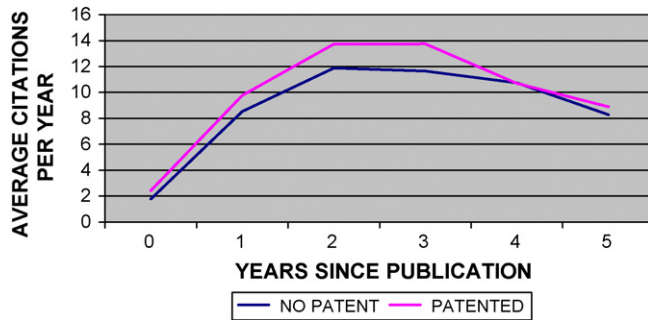


Fig. 3. Citations by type by age.

These data suggest that the number of citations and the 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 anti-commons effect, such an effect could result from differences in the characteristics of articles represented in the different age-cohort categories, which we explore in our more detailed empirical analysis.

5. 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 differences-in-differences estimates vary with article and patent characteristics. As discussed above, except where noted, all specifications employ a negative binomial regression. The tables report both the estimated coefficient and the coefficient in the form of an *incident rate ratio (IRR)*; since the IRR can be readily 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 focus on the IRR as we discuss the results.²⁰

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

This brief cross-sectional analysis motivates our main empirical analysis of the impact of *patent grant conditional on article quality*. We begin in second column of Table 5 with a regression in which the dependent variable is equal to the natural log of FORWARD CITATIONS + 1; this OLS regression includes a simple treatment for the evolution of citations over time (AGE and AGE²), as well as a complete set of article fixed effects. Though this specification does not account for the nature of citations as skewed count data, the results are suggestive. There is a negative

²⁰ Robust standard errors for the coefficients are reported in the third line of each row of each negative binomial specification.

Table 4
Cross-sectional results

NEGATIVE BINOMIAL SPECIFICATIONS				
Dependent Variable = Count of Forward Citations ^a				
	Baseline model		With publication controls	
PATENTED	[1.179] 0.164 (0.056)		[1.110] 0.104 (0.053)	
# AUTHORS			[1.033] 0.032 (0.007)	
US AUTHOR			[1.264] 0.234 (0.052)	
PUBLIC SECTOR AUTHOR			[0.934] −0.068 (0.118)	
Parametric restrictions				
Age FEs = 0	# Restrict	5	# Restrict	5
	χ^2	214.10	χ^2	227.98
	<i>p</i> -value	0.000	<i>p</i> -value	0.000
Year FEs = 0 ^b	# Restrict	5	# Restrict	5
	χ^2	21.24	χ^2	17.47
	<i>p</i> -value	0.001	<i>p</i> -value	0.004
Log-likelihood	−5263.99		−5238.52	
# of observations	1688		1688	

^a Values are [incident rate ratios] non-exponentiated coefficients (Robust coefficient standard errors).

^b Year FEs included for 1998–2002 (1997 is excluded).

though insignificant coefficient on PATENT GRANT WINDOW, and PATENT, POST-GRANT is associated with more than a 13 percent decline in the rate of citation (significant at the 5 percent level). The remainder of this table returns to negative binomial regression. We begin in third column of Table 5 with a straightforward specification that allows us to estimate both the overall difference between patented and non-patented articles and the marginal impact of being in the post-grant phase; similar to second column of Table 4, our controls include several publication characteristics, as well as controls for the age and publication year of the article. Three key findings stand out: while patented articles enjoy approximately a 20 percent overall citation boost, a patent grant is associated with an insignificant (though negative) impact on citation in the year in which the patent is granted. However the post-grant effect is associated with a statistically significant 19 percent decline in the expected citation rate. In other words, the initially higher citation rate for patent-paper pairs is erased in the years after a patent is granted. Finally, in fourth column of Table 5, we report a differences-in-differences estimate, including a separate fixed effect for every article, as well as a complete set of fixed effects for age and citation year; as such, these estimates are identified exclusively from the within-article contrasts between pre-grant and post-grant citation levels (after accounting for the impact of article age and year). According to this specification, the estimated post-grant decline is over 10 percent (and is significant at the 5 percent level). Moreover, these baseline results are robust to alternative specifications and sample definitions.²¹

²¹ 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.

Table 5
The impact of patent grant

	OLS: Dep Var = ln(FORWARD CITATIONS + 1)		NEGATIVE BINOMIAL SPECIFICATIONS Dependent Variable = Count of Forward Citations ^b	
	Marginal impact, with article FEs ^a		Selection and marginal effects w/controls	Marginal effects, with article FEs
Patent characteristic				
PATENTED			[1.195] 0.178 (0.068)	
PATENT GRANT WINDOW	−0.052 (0.059)		[0.893] −0.113 (0.108)	[0.943] −0.058 (0.044)
PATENT, POST-GRANT	−0.136 (0.068)		[0.817] −0.202 (0.099)	[0.893] −0.112 (0.056)
Control variables				
# AUTHORS			[1.034] 0.034 (0.007)	
US AUTHOR			[1.286] 0.251 (0.053)	
PUBLIC SECTOR AUTHOR			[0.912] −0.092 (0.121)	
AGE	0.960 (0.033)		[3.118] 1.137 (0.065)	
AGE * AGE	−0.162 (0.006)		[0.834] −0.181 (0.014)	
PUBLICATION YEAR = 1998			[1.310] 0.270 (0.060)	

Table 5 (Continued)

OLS: Dep Var = ln(FORWARD CITATIONS + 1)		NEGATIVE BINOMIAL SPECIFICATIONS	
Marginal impact, with article FEs ^a		Dependent Variable = Count of Forward Citations ^b	
		Selection and marginal effects w/controls	Marginal effects, with article FEs
PUBLICATION YEAR = 1999		[1.281] 0.248 (0.071)	
Parametric restriction			
Article FEs = 0	# Restrict χ^2 <i>p</i> -value	340 10882.97 0.00	# Restrict χ^2 <i>p</i> -value
Age FEs = 0			5 500.44 0.000
Year FEs = 0 ^c			# Restrict χ^2 <i>p</i> -value
			5 2.63 0.756
Regression statistics			
log-likelihood		−5270.42	−4021.44
R^2	0.75		
# of observations	1688	1688	1688

Differences-in-differences estimates.

^a OLS standard errors are robust.
^b Values are [incident rate ratios] non-exponentiated coefficients (Robust coefficient standard errors).
^c Year FEs included for 1998–2002 (1997 is excluded).




Fig. 4. Impact of patent grant on forward citations, by year before and after patent grant (negative binomial with article FEs).

These estimates provide evidence that is consistent with the existence of an anti-commons effect. Simply put, the impact of a given piece of scientific research (as measured by its citation rate) on subsequent scientific research declines after IP rights are granted. However, we are cautious in attempting to interpret this empirical finding. On the one hand, taken at face value, the results suggests that after IPR are granted, between 1 in 11 and 1 in 6 researchers (or publications) who might otherwise build on a given paper may forego a specific research project (or a particular research approach) that would necessitate citation to the article in that patent-paper pair. However, it is important to emphasize that, by themselves, these findings do not preclude alternative interpretations. For example, follow-on researchers may not shift their actual research agenda but may instead be engaged in a more “strategic” pattern of citation, avoiding those references that directly link to intellectual property rights. While the econometric analysis does not rule this out, detailed survey-based evidence and more qualitative evaluations of this phenomenon suggest that it is unlikely that strategic citing behavior is driving these results.²² Perhaps more importantly, these findings do not provide direct evidence about what activities follow-on researchers substitute into after patent grant. To the extent that researchers have access to reasonable substitutes, the impact on citation rates that we identify may significantly overstate the impact in terms of overall research productivity.

With these caveats in mind, it is useful to dig more deeply into the core drivers of our findings. For example, the results would be far less persuasive if the erosion of the initial citation advantage associated with patented articles actually began prior to the patent grant date, and indeed it is likely that the anti-commons would likely take some time to manifest itself after patent grant itself. To explore these hypotheses, we estimate a fixed effects negative binomial regression including specific dummy variables for each year preceding and following the patent grant date. In Fig. 4, we display the coefficients from that specification. Overall, patented articles tend to have a higher citation rate (relative to the baseline of 0, which is the value of the “average” fixed effect in the entire sample). While articles receive an up-tick in citations in the year prior to patent grant, there is an immediate (?) decline in the year of patent grant which continues steadily through 4 years following the date of patent grant. The difference between the pre-grant average and the average

²² In particular, Walsh et al. (2003) suggest that academic researchers have only limited awareness of the precise rights granted to individuals, so it is unlikely that citations in the scientific literature are systematically strategically manipulated to avoid citation to individual patent-paper pairs.

4 years after patent grant is more than 25 percent, suggesting that the impact of IPR accumulates over time; taking these patterns at face value, a patent grant is associated with a sizable decline in citation after IP rights have been in force for several years.

5.1. *The role of institutional affiliations, location and funding*

As noted above, the focus of the anti-commons debate is on scientific research conducted with US government funds that involves university or other public sector researchers. In Table 6, 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 fourth column of Table 5). In second and third column of Table 6, we examine how the impact of patent grant varies according to 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 is at least one public author (and is even more salient when there are *no* private authors). It is possible that articles associated with private sector researchers are *assumed* to have intellectual property associated with them, 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 dampen the incentives for follow-on research activities in that particular research stream or approach. Also, though the results are noisy (and may be driven by a smaller number of observations with no US authors), the impact of patent grant is significant for those articles associated with exclusively non-US authors; since we had hypothesized that US-authored articles may facilitate enforcement of IPR, this finding is something of a puzzle, which we hope to address in future research.

We extend our analysis of the impact of institutional affiliation on the anti-commons effect by evaluating the role of patent assignee affiliations and the source of research funding. In second column of Table 7, 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 third column of Table 7, 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 reported research funding in the patent (the individual coefficients are noisy and not different from each other). Though policy interventions associated with reducing the patentability of research in Pasteur's Quadrant might have a direct impact on Federally funded academic research, it is useful to recall that a significant fraction of research in these areas may be funded by private funds and/or may involve the private assignment of IPR (particularly outside the US).

5.2. *Does the impact of patent grant depend on patent type?*

At its heart, the anti-commons hypothesis is premised on the difficulty of using contracting methods to overcome the limitations on access to knowledge associated with patent protection. As noted above, we would expect contractual difficulties or the imposition of patent rights to be strongest for patents with multiple complex claims, extensive prior art, and potentially in areas such as research tools. Table 8 explores this hypothesis by examining the interaction between the POST-GRANT dummy and the patent characteristic measures. In second column

Table 6
Diffs-in-diffs results by institutional or national affiliation

	NEGATIVE BINOMIAL SPECIFICATIONS		
	Dependent Variable = Count of Forward Citations ^a		
	No public author vs. public author	No private author vs. private author	No US author vs. US author
Article characteristics			
PATENT GRANT WINDOW	[0.943] −0.059 (0.045)	[0.943] −0.059 (0.045)	[0.945] −0.056 (0.044)
NO PUBLIC AUTHOR	[0.929] −0.074 (0.105)		
PATENT, POST-GRANT* PUBLIC AUTHOR	[0.886] −0.121 (0.058)		
PATENT, POST-GRANT* NO PRIVATE AUTHOR		[0.873] −0.136 (0.066)	
PATENT, POST-GRANT* PRIVATE AUTHOR		[0.920] −0.083 (0.069)	
PATENT, POST-GRANT* NO US AUTHOR			[0.837] −0.178 (0.075)
PATENT, POST-GRANT* US AUTHOR			[0.921] −0.082 (0.062)

Table 6 (Continued)

NEGATIVE BINOMIAL SPECIFICATIONS Dependent Variable = Count of Forward Citations ^a						
No public author vs. public author			No private author vs. private author		No US author vs. US author	
Control variables						
Article pair FEs = 0	# Restrict	338	# Restrict	338	# Restrict	338
	χ^2	11308.14	χ^2	11370.21	χ^2	11616.61
	<i>p</i> -value	0.00	<i>p</i> -value	0.00	<i>p</i> -value	0.00
Age FEs = 0	# Restrict	5	# Restrict	5	# Restrict	5
	χ^2	544.17	χ^2	502.79	χ^2	499.77
	<i>p</i> -value	0.00	<i>p</i> -value	0.00	<i>p</i> -value	0.00
Year FEs = 0 ^b	# Restrict	5	# Restrict	5	# Restrict	5
	χ^2	41.01	χ^2	4.03	χ^2	9.91
	<i>p</i> -value	0.00	<i>p</i> -value	0.55	<i>p</i> -value	0.08
Regression statistics						
log-likelihood	−4021.33		−4021.21		−4020.79	
# of observations	1688		1688		1688	

^a Values are [incident rate ratios] non-exponentiated coefficients (Robust coefficient standard errors).

^b Year FEs included for 1998–2002 (1997 is excluded).

Table 7
Diffs-in-diffs results by nature of research funding

NEGATIVE BINOMIAL SPECIFICATIONS				
Dependent Variable = Count of Forward Citations ^a				
	No public sector assignee vs. public sector assignee		Non-government funded vs. government funded	
PATENT GRANT WINDOW	[0.944] −0.058 (0.044)		[0.943] −0.058 (0.045)	
PATENT, POST-GRANT* NO PUBLIC SECTOR ASSIGNEE	[0.880] −0.128 (0.067)			
PATENT, POST-GRANT* PUBLIC SECTOR ASSIGNEE	[0.906] −0.098 (0.067)			
PATENT, POST-GRANT* NON-GOVT. FUNDED			[0.903] −0.102 (0.059)	
PATENT, POST-GRANT* GOVT. FUNDED			[0.874] −0.135 (0.082)	
Parametric restrictions				
Article pair FEs = 0	# Restrict	338	# Restrict	338
	χ^2	11454.23	χ^2	11287.99
	<i>p</i> -value	0.00	<i>p</i> -value	0.00
Age FEs = 0	# Restrict	5	# Restrict	5
	χ^2	548.28	χ^2	502.40
	<i>p</i> -value	0.00	<i>p</i> -value	0.00
Year FEs = 0 ^b	# Restrict	5	# Restrict	5
	χ^2	52.20	χ^2	12.91
	<i>p</i> -value	0.00	<i>p</i> -value	0.02
Regression statistics				
log-likelihood	−4021.36		−4021.35	
# of observations	1688		1688	

^a Values are [incident rate ratios] non-exponentiated coefficients (Robust coefficient standard errors).

^b Year FEs included for 1998–2002 (1997 is excluded).

Table 8
Diffs-in-diffs results with patent characteristic interactions

	NEGATIVE BINOMIAL SPECIFICATIONS		
	Dependent Variable = Count of Forward Citations ^a		
	Patent outcome interactions	Patent reference interactions	Technology type interactions
“Direct” effect of patent grant			
PATENT GRANT WINDOW	[0.950] −0.053 (0.044)	[0.943] −0.059 (0.044)	[0.941] −0.061 (0.044)
PATENT, POST-GRANT	[0.888] −0.119 (0.057)	[0.891] −0.115 (0.058)	
Interaction effects			
PATENT, POST-GRANT* PATENT LAG	[0.9997] −0.0002 (0.0001)		
PATENT, POST-GRANT* PATENT CLAIMS	[1.000] 0.000 (0.002)		
PATENT, POST-GRANT* PATENT BACKWARD CITES		[1.000] 0.0002 (0.003)	
PATENT, POST-GRANT* PATENT BACKWARD REFS		[1.000] 0.0001 (0.001)	
PATENT, POST-GRANT* NON-RESEARCH TOOL			[0.821] −0.197 (0.069)
PATENT, POST-GRANT* RESEARCH TOOL			[0.943] −0.058 (0.064)

Control variables						
Article pair FEs	# Restrict	338	# Restrict	338	# Restrict	338
	χ^2	11384.48	χ^2	11131.89	χ^2	11563.14
	<i>p</i> -value	0.00	<i>p</i> -value	0.00	<i>p</i> -value	0.00
Age FEs	# Restrict	5	# Restrict	5	# Restrict	5
	χ^2	557.31	χ^2	512.26	χ^2	501.59
	<i>p</i> -value	0.00	<i>p</i> -value	0.00	<i>p</i> -value	0.00
Year FEs ^b	# Restrict	5	# Restrict	5	# Restrict	5
	χ^2	78.49	χ^2	19.16	χ^2	3.72
	<i>p</i> -value	0.00	<i>p</i> -value	0.002	<i>p</i> -value	0.59
Regression statistics						
log-likelihood		−4019.25		−4021.42		−4019.86
# of observations		1688		1688		1688

^a Values are [incident rate ratios] non-exponentiated coefficients (Robust coefficient standard errors).

^b Year FEs included for 1998–2002 (1997 is excluded).

of Table 8, we examine the impact of two patent process outcome measures (PATENT LAG and # CLAIMS), and in third column of Table 8, we evaluate the impact of the extent of prior art (through interactions with PATENT BACKWARD CITATIONS and PATENT BACKWARD REFERENCES).²³ Finally, in fourth column of Table 8, we compare the interaction effect for those patents associated with RESEARCH TOOLS versus not. The principal result is that there is a statistically significant (though quantitatively) modest incremental impact of longer patent lags on subsequent citations and a much more salient effect for non-research tools. In other words, though there is some evidence for interactions with patent characteristics, the findings are quite noisy; furthermore, though a key tenet of the anti-commons theory is that the effects are particularly salient for research tool patents, there is no evidence that the impact of patent grant is significant for these types of inventions.

5.3. *Exploiting variation in patent grant delay*

Though encouraging, it is important to emphasize that our results so far have relied on the presence of the control group of non-patented articles. While we have emphasized the comparable nature of our control group and included variables to account for observable differences, the potential for unobserved differences in the two groups that correlate to the pattern of forward citations remains a possibility. We examine this by exploring an empirical approach relying exclusively on a sample composed of articles that 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.

Nonetheless, in Table 9, we estimate the impact of a patent grant relying solely on differences across articles in the amount of time between publication and patent grant. We begin with a specification similar to that of third column of Table 5; we include a number of publication characteristics, as well as controls for article age and publication year. The coefficient on POST-GRANT is quantitatively and statistically significant, with a magnitude similar to that found in third column of Table 5 (patent grant is associated here with a 16 percent decline in citation rate). However, when we turn to an article fixed effects specification in the third column of Table 9, the overall impact of a patent grant is both small and insignificant. By and large, this negative result is driven by the fact that both the patent grant lag and the overall evolution of citations seem to be different for non-US authors (and, to a more limited extent, for private sector authored publications). When we condition the sample on precisely the type of research at the heart of the policy debate (patented research articles written by US public sector authors), we are able to identify a quantitatively significant (and marginally statistically significant) impact of a patent grant on subsequent citation rates. In other words, limiting ourselves to the variation associated with differences in patent grant lags and focusing on a smaller sample of articles with a US academic association, our empirical approach suggests that a patent grant is associated with a modest decline in the forward citation rates to scientific research articles over which patent rights are granted.

²³ To facilitate interpretation, each of the interaction effects are calculated relative to the mean of the distribution of that characteristic, so the coefficient on the direct effect for PATENT, POST-GRANT can continue to be interpreted as the average impact of PATENT GRANT at the sample mean of the patent characteristics under discussion.

Table 9
Patented articles only

	NEGATIVE BINOMIAL SPECIFICATIONS		
	Dependent Variable = Count of Forward Citations ^a		
	Marginal effects w/controls	Marginal effects, with article FEs (full sample)	Marginal effects, with article FEs (US public sector authors only)
Patent characteristics			
PATENT GRANT WINDOW + PATENT, POST-GRANT	[0.846 – 0.167 (0.091)	[0.988] –0.012 (0.051)	[0.894] –0.112 (0.067)
Control variables			
# AUTHORS	[1.037 0.036 (0.010)		
US AUTHOR	[1.491 0.399 (0.075)		
PUBLIC SECTOR AUTHOR	[0.887] –0.120 (0.152)		
AGE	[3.068 1.121 (0.094)		
AGE * AGE	[0.835 – 0.181 (0.020)		
PUBLICATION YEAR = 1998	[1.221 0.199 (0.084)		
PUBLICATION YEAR = 1999	[1.200 0.182 (0.103)		

Table 9 (Continued)

NEGATIVE BINOMIAL SPECIFICATIONS					
Dependent Variable = Count of Forward Citations ^a					
	Marginal effects w/controls	Marginal effects, with article FEs (full sample)		Marginal effects, with article FEs (US public sector authors only)	
Parametric restrictions					
Article FEs = 0		# Restrict	167	# Restrict	89
		χ^2	5792.48	χ^2	4540.42
		<i>p</i> -value	0.00	<i>p</i> -value	0.00
Age FEs = 0		# Restrict	5	# Restrict	5
		χ^2	341.67	χ^2	250.88
		<i>p</i> -value	0.00	<i>p</i> -value	0.00
Year FEs = 0 ^b		# Restrict	5	# Restrict	5
		χ^2	53.80	χ^2	100.57
		<i>p</i> -value	0.00	<i>p</i> -value	0.00
Regression statistics					
log-likelihood	−2698.94		−2041.96		−1103.44
# of observations	849		849		462

^a Values are [incident rate ratios] non-exponentiated coefficients (Robust coefficient standard errors).

^b Year FEs included for 1998–2002 (1997 is excluded).

6. Discussion

This paper provides the first differences-in-differences empirical test of the impact of IPR on the diffusion of scientific knowledge and the specific anti-commons hypothesis. Our empirical approach exploits the fact that the duality of knowledge is captured in the phenomenon of patent-paper pairs. Specifically, we exploit variation in both the occurrence of patent-paper pairs among a group of similar scientific publications and the patent grant delay, thus allowing us to examine pieces of knowledge in distinctive institutional settings. Our evidence suggests that knowledge duality in general and patent-paper pairs in particular are an important phenomenon in high quality research in the life sciences (nearly 50 percent of articles published in *Nature Biotechnology* are associated with a pair) and that for scientific knowledge subject to both Open Science and private property institutional regimes, the granting of IPR is associated with a statistically significant but modest decline in knowledge accumulation as measured by forward citations (in academic publications). Moreover, the decline in forward citations becomes more pronounced with the number of years elapsed since patent grant. As the anti-commons hypothesis suggests, the largest decreases occur for knowledge produced within university (or other public) institutions, where the patent grant is likely to be “news”. Interestingly, while the salience of patent grant is increasing in the patent grant lag, the decline in forward citations is actually higher for patents outside of traditional “research tools” categories and for papers with non-US authors. Overall, we are able to reject the null hypothesis that IPR have no impact on the diffusion of scientific knowledge.

These patterns provide a novel perspective on the economic consequences of the privatization of the scientific commons. Rather than simply serving to facilitate a “market for ideas,” IP may indeed restrict the diffusion of scientific research and the ability of future researchers to “stand on the shoulders of giants,” at least for research of the type published in *Nature Biotechnology*. However, erecting a (property rights) barrier to the accumulation of knowledge does not eliminate all Open Science use of that knowledge. The demand for specific published findings by follow-on researchers may be relatively inelastic; taken at face value, our estimates suggest that patent grant is associated with a reduction in the use (or attribution) of a discovery by approximately 1 in 10 to 1 in 6 follow-on research projects or at least follow-on publications (which may be part of a single research project by one researcher with many foregone publications, or 6–10 distinctive projects by different researchers).

With that said, we are cautious in the interpretation of our findings. First, we have not identified the underlying institutional mechanisms by which patent grant shifts citation behavior. We cannot disentangle whether the response to patent grant comes about through an awareness of the expansion of IPR over a particular idea, through interactions with aggressive licensees of such patents, or perhaps through the administrative burdens placed by materials transfer agreements (which are typically more onerous after patenting). Also, while patent grant may discourage follow-on research, it may simply serve to shift the specific research conducted (or cited) by follow-on researchers; in other words, it is possible that follow-on researchers have relatively good substitutes, in which case the tax imposed by patent grant (and institutional mechanisms such as MTAs) may have only a small impact on overall scientific research productivity. Second, this paper does not investigate the *source* of the reduction in citations. For example, we do not address whether the observed reduction in citation is concentrated on those who already cited the work prior to the patent grant or in a reduction in the rate of “entry” to citation to a research article. Similarly, we do not address whether the impact of patent grant is particularly salient for future public or private researchers, for collaborations or independent follow-on work, or for articles which appear in high-quality or low-quality journals.

The magnitude and distribution of the anti-commons effect among distinctive groups of researchers has important implications for policy and for the degree to which anti-commons issues should be used to animate and justify changes in patent policy over dual knowledge generated in academia. For example, the policy implications of our findings would depend considerably on whether the “missing” citations are associated with independent public-sector researchers versus being centered exclusively on follow-on research by private sector teams and whether they are associated with high or low quality publications. Adjudicating among these possibilities will provide us with much deeper insights into the mechanisms through which our observed citation decline is actually occurring. In ongoing work, we are investigating each of these hypotheses by examining how the *composition* of citation changes after IPR are introduced. Such an analysis should help deepen our understanding of the source of the patent grant effect, particularly in light of survey-based evidence that academic researchers claim to be mostly unaware of the status of IPR in their area (at least in terms of proactive monitoring) and may be engaging in pervasive (unknowing) infringement (Walsh et al., 2003, 2005).

Finally, it should be emphasized that our evidence for the anti-commons effect captures only one aspect of the impact of IP on dual knowledge. We have not observed or examined the other side of the “ledger”—whether or not the IP rights encompassing academic research articles have enhanced the incentives for (unobserved) research building on these ideas (particularly by private sector organizations), led to more effective (or rapid) commercialization (which is far more costly than the basic research component itself), or allowed for cumulative innovation through patents and future patent citations. Without a detailed accounting of the size of these (positive) effects of IPR rights on welfare, it is impossible to calculate the optimal innovation policy approach towards the establishment of IPR over knowledge that has traditionally been available in the scientific commons.

Despite these caveats, the importance of continuing to adjudicate this debate should not be understated. The production of dual-use knowledge is increasingly central in scientific research and of great interest to those who fund research. Biotechnology and the life sciences have grown enormously as a share of overall research activities, and similar questions arise in the context of other high-margin research areas such as nanotechnology and open source software. Because biotechnology and related disciplines simultaneously offer the potential for fundamental scientific discoveries *and* commercial breakthroughs, our traditional understanding of the role of private property institutions as compared to the norms of Open Science is open to question. By providing a window into the impact of IP rights on the diffusion and accumulation of scientific knowledge, this paper offers some insights into the tradeoffs associated with embedding the production and distribution of knowledge in two distinctive and potentially opposing institutional regimes.

Acknowledgements

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 Philippe Aghion, Tim Bresnahan, Iain Cockburn, Wes Cohen, Maryann Feldman, Roberto Fernandez, Joshua Gans, Shane Greenstein, Rebecca Henderson, Josh Lerner, Ed Roberts, Roberto Rigobon, Paula Stephan, seminar participants at the NBER Academic Science and Entrepreneurship Conference, as well as numerous conferences and seminars. We also thank Max Bulbin, Kyle Jensen, Jason O'Connor, Mitun Ranka and Kranthi Vistakula for excellent research assistance. All errors are our own.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.jebo.2006.05.017.

References

- Aghion, P., Dewatripoint, M., Stein, J., 2005. Academic freedom, private-sector focus and the process of innovation. NBER Working Paper.
- Agrawal, A., Henderson, R.M., 2002. Putting patents in context: exploring knowledge transfer from MIT. *Management Science* 48, 44–60.
- Allison, P.D., Waterman, R., 2002. Fixed effects negative binomial regression models. In: Stolzenberg, R. (Ed.), *Sociological Methodology*. Basil Blackwell, Boston.
- Argyres, N., Liebskind, J., 1998. Privatizing the intellectual commons: universities and the commercialization of biotechnology research. *Journal of Economic Behavior and Organization* 35, 427–454.
- Arora, A., Fosfuri, A., Gambardella, A., 2001. *Markets for Technology: Economics of Innovation and Corporate Strategy*. MIT Press, Cambridge, MA.
- Arrow, K., 1962. Economic welfare and the allocation of resources for invention. In: Richard, R.N. (Ed.), *The Rate and Direction of Inventive Activity*. Princeton University Press, Princeton, NJ, pp. 609–625.
- Azoulay, P., Ding, W., Stuart, T., 2007. The determinants of faculty patenting behavior: demographics or opportunities. *Journal of Economic Behavior and Organization* 63, 599–623.
- Blumenthal, D., Causino, N., Campbell, E., Seashore Louis, K., 1996. Relationships between academic institutions and industry in the life sciences—an industry survey. *New England Journal of Medicine* 334, 368–374.
- Bremmer, H., 2001. The first two decades of the Bayh–Dole Act as public policy. Address to the National Association of State Universities and Land Grant Colleges, November 11, 2001.
- Campbell, E.G., Clarridge, B.R., Gokhale, M., Birenbaum, L., Hilgartner, S., Holtzman, N.A., Blumenthal, D., 2002. Data withholding in academic genetics: evidence from a national survey. *Journal of the American Medical Association* 287, 473–479.
- Campbell, E.G., Weissman, J.S., Causino, N., Blumenthal, D., 2000. Data-withholding in academic medicine: characteristics of faculty denied access to research results and biomaterials. *Research Policy* 29, 303–312.
- Cockburn, I., Henderson, R., 1998. Absorptive capacity, coauthoring behavior, and the organization of research in drug discovery. *Journal of Industrial Economics* 46, 157–182.
- Cockburn, I., Kortum, S., Stern, S., 2003. Are all patent examiners created equal? Examiners, patent characteristics and litigation outcomes. *Intellectual property in the knowledge-based economy*. National Academy of Sciences.
- Cohen, S.C., Boyer, A.H., Helling, R., 1973. Construction of biologically functional bacterial plasmids in vitro. *Proceedings of the National Academies of Science USA* 70, 3240–3244.
- Cole, J.R., 2000. A short history of the use of citations as a measure of the impact of scientific and scholarly work. In: Cronin, B., Atkins, H.B. (Eds.), *The Web of Knowledge: A Festschrift in Honor of Eugene Garfield*. Information Today, Inc., Medford, NJ, pp. 281–300.
- Dasgupta, P., David, P., 1994. Towards a new economics of science. *Research Policy* 23, 487–521.
- David, P., 2003. Can ‘open science’ be protected from the evolving regime of IPR protections? Stanford Working Paper #03-011.
- David, P., 2001a. From keeping nature’s secrets to the institutionalization of open science. Stanford Working Paper #01-006.
- David, P., 2001b. Will building good fences really make good neighbors in science. Stanford Working Paper #01-005.
- David, P., 2000. The digital technology boomerang: new intellectual property rights threaten global open science. Stanford Working Paper #00-016.
- Ding, W., Murray, F., Stuart, T., 2006. An empirical study of gender differences in patenting in the academic life sciences. *Science* 313, 665–667.
- Ducor, P., 2000. Intellectual property coauthorship and coinventorship. *Science* 289, 873–875.
- Eisenberg, R., 1996. Public research and private development: patents and technology transfer in government-sponsored research. *Virginia Law Review* 82, 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., Sewell, J., 1995. *Networks of innovation: Vaccine development at Merck, Sharp & Dohme, and Mulford 1895–1995*. Cambridge University Press, New York.
- Gambardella, A., 1995. *Science and Innovation*. Cambridge University Press, Cambridge, UK.
- Gans, J., Stern, S., 2000. Incumbency and R&D incentives: licensing the gale of creative destruction. *Journal of Economics and Management Strategy* 9, 485–551.
- Geison, G.L., 1995. *The Private Science of Louis Pasteur*. Princeton University Press, Princeton, NJ.
- Greene, William, 2004. Fixed effects and the incidental parameters problem in the Tobit Model. *Econometric Reviews* 23, 125–147.
- Hall, B.H., Ziedonis, R.H., 2001. The patent paradox revisited: an empirical study of patenting in the US semiconductor industry, 1979–95. *RAND Journal of Economics* 32, 101–128.
- Harhoff, D., Wagner, St., 2005. Modelling the duration of patent examination at the European Patent Office. CEPR Discussion Papers 5283, C.E.P.R. Discussion Papers.
- Hausman, J.A., Griliches, Z., Hall, B., 1984. Econometric models for count data with an application to the patents-R&D relationship. *Econometrica* 52, 909–937.
- Heller, Michael, 1998. The tragedy of the anticommons: property in the transition from Marx to markets. *Harvard Law Review*, 621.
- Heller, Michael, Eisenberg, Rebecca, 1998. Can patents deter innovation? The anticommons in biomedical research. *Science* 280, 698–701.
- Hellman, T., 2007. Bridging the science to market gap: the role of patents. *Journal of Economic Behavior and Organization* 63, 624–647.
- Jaffe, A.B., Lerner, J., 2004. *Innovation and Its Discontents: How Our Broken Patent System is Endangering Innovation and Progress, and What to do About It*. Princeton University Press, Princeton.
- Jensen, R., Thursby, M.C., 2001. Proofs and prototypes for sale: the licensing of university inventions. *American Economic Review* 91, 240–259.
- Jones, B., 2005. Age and great invention. NBER Working Paper # 11359.
- Kenney, M., 1986. *Biotechnology: The University-Industrial Complex*. Yale University Press, New Haven.
- Kelves, D.J., 2002. Of mice & money: the story of the world's first animal patent. *Daedalus* 131, 78.
- Kitch, E.W., 1977. The nature and function of the patent system. *Journal of Law and Economics* 20, 265–290.
- Köhler, G., Milstein, C., 1975. Continuous cultures of fused cells secreting an antibody of predefined specificity. *Nature* 256, 495–497.
- Kremer, M., 1997. Patent buy-outs: a mechanism for encouraging innovation. NBER, Working Paper no. 6304.
- Krimsky, S., August 2003. *Science and the Private Interest*. Rowman-Littlefield Publishing Co., Lanham.
- Lemley, M.A., Allison, J.R., 2002. The Growing Complexity of the United States Patent System, vol. 82. *Boston University Law Review*, pp. 77–144.
- Lessig, L., 2002. *The Future of Ideas: The Fate of the Commons in a Connected World*. Random House, New York.
- Levin, R., Klevorick, A., Nelson, R., Winter, S., 1987. Appropriating the returns from industrial R&D. *Brookings Papers in Economic 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, 559–563.
- Markiewicz, K., DiMinin, A., 2005. Commercializing the laboratory: the relationship between faculty patenting and publishing. In: Presented at Fifth Annual Roundtable on Engineering Entrepreneurship Research (REER) Conference, Georgia Tech.
- Marshall, E., 2000. Property claims: a deluge of patents creates legal hassles for research. *Science* 288, 255.
- Merges, R., Menell, P.S., Lemley, M., 2003. *Intellectual Property in the New Technological Age*. Aspen Publishers.
- Merges, R., Nelson, R.R., 1990. On the complex economics of patent scope. *Columbia Law Review* 90, 839–916.
- Merges, R.P., 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, 1–24.
- Merton, R., 1973. In: Storer, N. (Ed.), *The Sociology of Science: Theoretical and Empirical Investigations*. University of Chicago Press, Chicago, p. 1973.
- Mowery, D., Nelson, R., Sampat, B., Ziedonis, A., 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*, 31.
- Murray, F., 2006. The mouse that roared: resistance and accommodation to patenting in academic science. Sloan School of Management Working Paper.

- Murray, F., Stern, S., 2006. When ideas are not free: the impact of patents on academic science. In: Paper Prepared for the Innovation Policy and the Economy Meeting, Washington D.C., April.
- National Institutes of Health, 2000. NIH and E.I. DuPont Sign OncoMouse Agreement. News Release January 19.
- Nature Biotechnology, 1996. Editorial: A broad and inclusive enterprise, vol. 14. March 1996, p. 235.
- Nature Biotechnology, 2006. Aims and scope of the journal. Accessed from <http://www.nature.com/nbt/about/index.html>.
- Nelkin, D., 1984. *Science as Intellectual Property: Who Controls Scientific Research?* MacMillan Publishing Co., New York.
- Nelson, R.R., 1959. The simple economics of basic scientific research. *Journal of Political Economy* 67, 297–306.
- Orsenigo, L., 1989. *The Emergence of Biotechnology: Institutions and Markets in Industrial Innovation*. Pinter Publishers, London.
- Owen-Smith, J., Powell, W.W., 2003. The expanding role of university patenting in the life sciences: assessing the importance of experience and connectivity. *Research Policy* 32, 1695–1711.
- Popp, D., Juhl, T., Johnson, D.K.N., 2004. Time in purgatory: examining the grant lag for U.S. patent applications. *Topics in Economic Analysis and Policy* 4, Article 29.
- 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.
- Regibeau, R., Rockett, K., 2003. Are more important patents approved more slowly and should they be? *Economics Discussion Papers 556*, University of Essex, Department of Economics.
- Scotchmer, S., 1996. Protecting early innovators: should second-generation products be patentable? *The RAND Journal of Economics* 27, 322–331.
- Shapiro, C., 2001. Navigating the patent thicket: cross licenses, patent pools, and standard setting. In: Jaffe, A.B., Lerner, J., Stern, S. (Eds.), *Innovation Policy and the Economy*, vol. 1.
- Sim, S.J., Snell, K.D., Hogan, S.A., Stubbe, J., Rha, C.K., Sinskey, A.J., 1997. PHA synthase activity controls the molecular weight and polydispersity of polyhydroxybutyrate in vivo. *Nature Biotechnology* 15, 63–67.
- Stokes, D., 1997. *Pasteur's Quadrant: Basic Science and Technological Innovation*. The Brookings Institution, Washington, D.C.
- Stern, Scott, 2004. Do scientists pay to be scientists. *Management Science* 50, 835–853.
- Stewart, T.A., Pattengale, P.K., Leder, P., 1984. Spontaneous mammary adenocarcinomas in transgenic mice that carry and express MTV/myc fusion genes. *Cell* 38, 627–637.
- Straus, J., Holzapfel, H., Lindenmeier, M., 2002. *Empirical Survey on Genetic Inventions and Patent Law*. Commissioned by the Federal Ministry of Education and Research, Germany.
- Thomson, J.A., Itskovitz-Eldor, J., Shapiro, Sander S., Waknitz, M.A., Swiergiel, J.J., Marshall, V.S., Jones, J.M., 1998. Embryonic stem cell lines derived from human blastocysts. *Science* 282, 1145–1147.
- Thursby, G., Thursby, M., 2003. *University licensing under Bayh-Dole: What are the issues and evidence?* Emory University Working Paper accessed from opensource.mit.edu/papers/Thursby.pdf.
- Walsh, J., Arora, A., Cohen, W., 2003. Science and the law: working through the patent problem. *Science* 299, 1021.
- Walsh, J., Cho, C., Cohen, W., 2005. View from the bench: patents and material transfers. *Science* 309, 2002–2003.
- Zamore, P.D., Tuschl, T., Sharp, P.A., Bartel, D.P., 2000. RNAi: Double-stranded RNA directs the ATP-dependent cleavage of mRNA. *Cell* 101, 25–33.
- Ziedonis, R., 2004. Don't fence me in: fragmented markets for technology and the patent acquisition strategies of firms. *Management Science*, 804–820.
- 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, 1193–1198.