Learning to Live with Patents: Assessing the dynamic adaptation to the law by the scientific community

Fiona Murray, MIT Sloan School

Scott Stern, Northwestern University and NBER

April 2008

Learning to Live with Patents:
Assessing the dynamic adaptation to the law by the scientific community

Abstract

Law plays a central, but under-appreciated role in knowledge work. While the broad institutional foundations of the role of law, specifically the patent system, in shaping the production of knowledge have been articulated, much of the current literature is based on equilibrium assumptions about the role of law. However this approach has led to considerable debate over the degree to which patents stifle knowledge work, particularly among scientific communities. We argue that law and knowledge work should be analyzed within from a more dynamic perspective - examining the ways in which knowledge communities adapt to law over time. Within this adaptive framework we identify three ways in which patents impact scientific knowledge production: through an initial shock, a cumulative tax and adaptation. This allows us to provide more nuanced insights into the way scientists learn to life with patents. Further we show that different sub-communities and different members of the community learn at different rates. This work contributes to our understanding of the role of law in knowledge work, the nature of law in practice and the role of law in scientific communities.

Knowledge work has come to be recognized as a key driver of economic growth, enhancing the competitive advantage of firms, regions and nations (Romer 1990; Aghion and Howitt 1997). Scholars have responded by exploring the nature of knowledge work and the practices of knowledge-based organizations. However, the ability of knowledge workers to produce knowledge and to build on knowledge generated by others is not an inherent property either of knowledge or of knowledge communities (Mokyr 2002). Instead it relies on the provision of an institutional framework for disclosure, access and attendant rewards (Stern 2004; Murray and O'Mahony 2007). Thus a critical agenda for scholars of knowledge work is to examine its "dynamic institutional context" (North 1993, p. 6) and understand the ways in which the institutional context – as constituted by material culture, social networks and broader rules and norms - shape the ability of innovators engage effectively in knowledge work.

Much of the literature has highlighted the role of material culture (Kohler 1994), artifacts (Latour, 1987), and boundary objects (Bechky, 2003; Carlile, 2004) helping individuals recall, replicate and combine distinct pieces of knowledge, offering a tangible form of access that can help translate knowledge from one community of experts to another and as a means of become immersed in prior knowledge. There is also a growing understanding of the ways in which social structures embedding knowledge communities enable (or constrain) their knowledge work (Van de Ven, 1993; Hansen, 1999; Hargadon and Sutton, 1997; Fleming, 2001; Knorr-Cetina, 1999). A large body of work illustrates the importance of an individual's or firm's network position in providing opportunities for knowledge accumulation through brokerage and by facilitating knowledge exchange (Hansen et al., 2005, Burt, 2004; Uzzi and Spiro, 2005; Flemming 2007). These studies highlight the factors shaping knowledge flows but provide few insights into the conditions under which different knowledge workers will share their knowledge with one another. For example, Hargadon and Sutton argue that firms such as IDEO use their industry position to gain access to a variety of knowledge and later effectively recombine that knowledge as the foundations for further knowledge accumulation (1997: 723). What is not considered are the conditions under which knowledge can be shared and accumulated from one client to the next – what are the terms and conditions of access, reuse and accumulation faced by IDEO or other knowledge workers? Thus discussions of knowledge work either assume that the law is ever present in the background shaping action in an unspecified fashion, or provide a law-free perspective on norms and practices.

Scholarship in the institutional tradition does provide important insights into the role of norms and informal practices in providing rules of access to knowledge (Dasgupta and David 1994; Fauchart and Von Hippel 2007) grounded in the sociological approach to the norms of sharing and exchange in science articulated by Merton (1957). Much of this literature however is silent on the role of law as it pertains to the rules shaping knowledge work. Nonetheless, with the rise of intellectual property (IP) rights in the past several decades (Jaffe and Lerner 2004), legal institutions do provide another pervasive source of rules shaping knowledge work: scientists must concern themselves with patents on research tools, rap artists sampling music confront possible copyright infringement suits, software engineers confront complex copyright and copyleft provisions as they combine code. Organizations, communities and individuals are now "immersed in a sea of law" (Edelman & Suchman 1997). This transformation engenders heated debate: Does the expansion of intellectual property, designed to protect and encourage knowledge work in many domains, now stifling the very work it sought to encourage?

Current research addressing this question is contradictory particularly as it pertains to the impact of intellectual property (IP) on the production of scientific knowledge by the academic research community and speaks to the limitations of current approaches to law and knowledge work. On the one hand, a significant amount of research has highlighted the benefits of IP rights (Kitch, 1977; Arora, Forsfuri and Gambardella, 2001). Recent empirical research on commercial discoveries suggests that IP 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 & Nelson, 1990, 1994; Arora, Fosfuri, and Gambardella, 2001; Gans and Stern, 2000). Indeed, within the context of university research (particularly publicly-funded university research), it has been suggested that IP offers important incentives to move nascent discoveries out of the "ivory tower" and into commercial practice. In other words, from the perspective of an individual discovery, IP may enhance the ability to realize its commercial and social benefits (Kitch, 1977, Hellman, 2005). However, the "anti-commons" perspective argues that the expansion of IP (in the form of patents and/or copyrights) is "privatizing" the scientific commons, and limiting scientific progress (Heller and Eisenberg, 1998; Argyres & 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 cumulatively build on each other's discoveries (Heller & Eisenberg, 1998; David, 2003, 2000; Lessig, 2002; Etzkowitz, 1998; Krimsky, 2003). The empirical evidence is also contradictory. On the one hand, Heller and Eisenberg (1998) have provided empirical evidence that academic scientists have been stifled in their knowledge work. Their perspective has been supported by more recent quantitative analysis by Murray and Stern (2007) showing that the grant of IP rights over scientific knowledge leads to a reduction in citations to that knowledge. In contrast, Cohen, Walsh and co-authors use extensive survey-based evidence to document that while more than 20% of scientists have been involved in seeking IP over their discoveries in the two years prior to the survey, few admit to paying attention to IP rights, and only 20% report that specific projects have been delayed or diverted because their most recent request for materials had been declined or delayed while they negotiate access to knowledge and materials via Material Transfer Agreements(which may or may not be associated with patents) with other universities and academics (Walsh et al. 2002, 2003, 2005).

In this paper, we argue that this debate can be resolved by taking a dynamic approach to the study of law and knowledge work. The current approach to the law exemplified by the debate over IP rights and the scientific community exemplifies the equilibrium perspective on the law: First, it assumes that the law has a direct impact on individuals, unmediated by practices, norms and culture. Second, it assumes that actors respond rapidly and rationally to changes in the law and adapt to any legal "shocks" quickly and thoroughly (Meyer and Argyris 2004). We suggest that a dynamic approach to the role of law provides a more nuanced insights into the role of law in knowledge work, are closer to the empirical phenomenon and allow us to resolve the controversy in the debate over IP and the scientific community. This dynamic approach builds on current literature in law and society which highlights the nature of law as it is constituted in the daily practice of individuals and communities (Edelman and Suchman 1997; Silbey xxx). If the dynamic approach to law and knowledge work is a close representation of knowledge communities, then we would expect to find dramatic shifts in the role of law in knowledge work over an extended period. This view is supported by recent qualitative longitudinal analysis of a scientific community engaged in mouse genetics impacted by the patenting of one of their central research tools – the Oncomouse (Murray 2008). The study highlights how scientists learn to live with patents and the complex temporal dynamics through which the scientists responded to the Oncomouse patents. In this paper we hypothesize that this adaptive approach to the law can be generalized to other scientific knowledge communities and the effects of IP on the knowledge work of scientists identified through three empirically distinct, separable measures: An initial shock brought by a change in the law, the tax effect that increases the impact of the law on knowledge work over time, and a countervailing, adaptation effect that mediates and reduces the impact of law on knowledge work. Using a sample of life science publications also subject to IP rights – known as patent-paper pairs (Murray 2002; Murray and Stern 2007) – we develop a formal empirical test of the longitudinal impact of law on the scientific community. We measure the relative strength of these three separate effects by examining the follow-on knowledge production in the period when the knowledge was published and shared through normative (academic) schema and after the "shock" of IP rights. Rather than taking the traditional equilibrium approach and estimating only the average patent shock, we measure the initial shock of patents but also measure the adaptive response to IP over time and the changing tax imposed by any given generation of patents. Our results suggest that knowledge workers the initial imposition of IP rights over their work has a chilling effect that is increasing in the life of the patent. However there is also strong evidence for adaption to the law particularly among the academic community of scientists, who learn to live with patents and over time find mechanisms through which patents no longer impact their productivity.

These findings have some important implications for the debate over IP and knowledge work. In particular, the dynamics of knowledge work suggests that cross-sectional approaches will not capture the range of ways in which communities learn to live with and adapt to the law. Specifically, its suggests a resolution of the current empirical impasse in the anti-commons debate: It is possible that in the period studied by Heller and Eisenberg (1998), impact of patents was significant and the rapid rise of patents had caused a shock to the academic community and was of growing salience as universities sought to impose their IP rights on many community members. However by the time of the survey-based analyses, the life science community had adapted to patents, they continued to file patents but their impact on knowledge work had become curtailed as the community a combination of contractual and normative mechanisms to limit the deleterious impacts of IP rights and reinforce the traditional practices of the academic community (Murray 2008). More broadly, our findings suggest a promising research agenda examining the role of law as it shapes the daily lives of knowledge communities across a broad range of sectors and work. Rather than follow the equilibrium approach and simply document whether or not law impacts knowledge work, this agenda can encompass rich quantitative and qualitative analyses of the ways in which different communities learn to live with the law, the rate of adaptation and the extent of adaptation. Work in areas as diverse as fashion (Sprigman 2007), cuisine (Fauchart and Von Hippel, 2007) and software (O'Mahony 2003) suggest that different communities will develop distinctive adaptive strategies to IP rights.

This paper proceeds as follows. We first discuss the equilibrium approach to the role of law and the impact of legal change on knowledge work. We then contrast this with a dynamic approach to the study of law and knowledge drawing on the law and society literature and develop a series of hypotheses on the impact of IP on the scientific community. In the following section we describe our empirical setting and empirical approach. Next we present our empirical results and then conclude with a discussion of the implications of our findings for debates over IP and scientific knowledge work and more broadly for the role of law and knowledge work.

Theoretical Perspectives

Law & Knowledge Work – an equilibrium approach

The debate over the impact of patents on the scientific knowledge community, and on creative work more broadly, is mainly grounded in an equilibrium approach to our understanding of law and knowledge work. Built upon a broad understanding of the law as part of a broader collage of institutions that provide the incentive structure of the economy and therefore shape economic growth (North 1991), the institutional approach to law and knowledge work (more widely referred to as innovation) is grounded in a study of the patent system and the broader role of intellectual property in the economy (Scotchmer 1991; Jaffe and Lerner 2004). This literature is grounded in the notion that the accumulation of knowledge – standing on the shoulders of giants – is not an inherent property of knowledge or of knowledge communities. Instead it relies on the provision of an institutional framework for disclosure, access and attendant rewards (Mokyr 2002; Stern 2004; Murray and O'Mahony 2007). That law can provide (and inhibit) such institutional support it clear. Notwithstanding recent critiques of the patent system (Lessig 2001), the design of the patent system is intended to serve as a set of legal arrangements providing incentives for innovators to invest in and then disclose novel, useful and non-obvious ideas. The incentives that govern the patent system depend on the degree to which a researcher can *exclude* others and so appropriate some of the value created by their knowledge through the commercialization of new technology (Nelson 1959; Arrow 1962; Levin et al. 1987; Kremer 1997; Scotchmer 1996).

The prevailing approach to understanding the role of patent law and knowledge work is to examine the ways in which the various elements of the system – what is considered patentable material, the length of patent protection, the scope of patents, the time to patent grant – impact the rate and direction of knowledge work. This literature makes two important assumptions about the role of law as it shapes knowledge work. First, that the law has a direct impact that can be directly linked to specific attributes of patent law. The effects of the law are modeled as unmediated by norms and practices, instead being based on a clear interpretation and direct response to the law. The second assumption implicit in most analyses of this type is to examines the equilibrium implications of the law as given or changes in the law. Central to this approach are models of how changes in the length and scope of patent protection affect the incentives for both first and second generation innovators (Katz and Shapiro, 1985; Merges and Nelson, 1990; Klemperer, 1990; Scotchmer 1991). These models consider the decisions of what to patent and whether to race with another firm to win a patent on the basis that there will be "foresight and in the knowledge that the subsequent decisions of other players will be similarly rational" (Scotchmer and Green 1990, p. 132). In other words, these equilibrium assumptions can be interpreted to imply that adaptation to any legal change is quick and thorough (Meyer and Argyris 2004). Empirical work in this tradition establishes a similar framework, exploring how a group of actors respond to a particular legal decision and examining the change in knowledge production before and after the legal "shock". Three examples illustrate this approach: Using historical patent data from over 100 US semiconductor firms Hall and Ziedonis (2001) examine whether a shift in the legal environment with the establishment of the Court of Appeals for the Federal Circuit (CAFC) which was widely interpreted as being pro-patent did in fact lead to a rise in patenting among firms in the period before and after the legal shock. At an individual level, Marx et al. (2007) examine whether the enforcement of non-compete agreements impacts the mobility patterns of inventors in knowledge-based communities. They use an inadvertent reversal of enforcement in Michigan in the mid 1980s as a "shock" in the legal institutional regime and measure mobility pre- and post- shock. In a third example, Murray and Stern (2007) examine whether the rate of follow-on academic research – as measured by citations - is impacted by the granting of more patent rights over the knowledge described in the original (cited) publication. They use the patent grant date as a "shock" to the openness and accessibility of the published knowledge. In each of these three papers, the authors use a legal shock as a means to identify the impact of law on knowledge work. In each case, they make a series of equilibrium assumptions that law will instantaneously impact the knowledge community, and that its impact will (presumably) be directly in accordance with our economic understanding of the nature of the law. While these and other studies provided significant insights into the design of the patent system, in the debate over the role of law in knowledge work this approach provides contradictory results which speak to the need for an alternative. Indeed in the anti-commons debate raging around the impact of patents on the scientific community, as noted above, both theory and empirical evidence developed within the equilibrium tradition provide countervailing predictions and findings. We argue that in fact one potential resolution is to take a more adaptive perspective on the law.

Law & Knowledge Work – an adaptive approach

The adaptive, dynamic approach to law and knowledge work is grounded in the recognition of two important characteristics of the law as it shapes the daily life of individuals. First, we argue that the opposition between the role of law and the role of norms in shaping knowledge communities is a false dichotomy. In the traditional perspective, law is regarded in opposition to more norms-based approach and is viewed as a rigid framework in which actors simply "read off" the appropriate efficient response to any legal shock. However, this ignores the possibility that law itself is embedded in norms and behaviors with individuals interpreting and mobilizing rights in response to legal structures as read through their own experiences (Fuller, Edelman and Matusik; 2000). Thus law becomes constituted largely by the meanings that are attached to IP rights and their implementation, rather than directly through the incentives they are intended to provide. This perspective is illustrated in the mouse genetics community which came to see IP rights as an alternative source of prestige (Murray forthcoming), and as a mechanism to protect the open commons of academia rather than simply as a way of protecting and profiting from their knowledge (Murray 2008). More broadly, within the overall patent system, the choice and arrangement of institutions does not follow economic efficiency criteria alone but is also influenced by cultural and social factors (DiMaggio1994). A second characteristic follows from this observation; that the transformation and interpretation of the law requires adaptive change and may not be immediate. Instead, over time, individuals come to interpret the law in the same way that they learn to operate within a wide range of contractual, moral and social orders (Meyer and Argyris 2004). This adaptation arises as law establishes normative frameworks for behaviors and for meaningfulness – with the law being interpreted through the lenses of social practices (Macaulay 1963, Hurst 1964, Mnookin and Kornhauser 1979). The literature on law and society provides further guidance in this regard, arguing that the law is as malleable – that different interpretations and different ways in which it can shape practice. The adaptive approach is not to deny the potential efficiencies associated with institutions grounded in either a legal or normative order (North 1991). Instead it suggests that adaptation to institutional change is likely to be slow and to involve a complex of individual and community-based actions that emerge through the daily life of knowledge workers or other actors (Heimer 1985). In sociology studies of law and the employment relationship have pioneered our understanding of the ways in which organizations respond to legal change (Kelly and Dobbin 1999; Dobbin and Kelly 2007). They argue that the law has had a profound impact on employer practices, for example in changing maternity leave policies. Moreover, it is not necessarily the strongest legal sanctions that lead to the greatest changes, contrary to the dominant legal scholarship (Posner 1972). Instead, legal ambiguity has provided ample opportunities for the legal profession to establish its expertise and win corporate resources (Edelman et al. 1992). It has also led to complex compliance activities subject to significant contestation, thus "legislation is just the beginning of law making" (Kelly and Dobbin 1999, p. 487) and the beginning of law in practice. This is particularly true of intellectual property. In Edelman and Suchman's terms (1997), IP rights are facilitative in that they provide their owners with a variety of legal rights that they can use in as a setting for action. They do not directly modify behavior (of the patent holder or others) but they can do so. Such adaptation arises in part because of the flexibility inherent in IP rights – they can be thought of as a right to exclude others, but they also confer prestige, can be used to control the innovation of others, and they are a source of potential economic rewards (Murray 2008). It is also grounded in the recognition that like virtually all property rights, patents have some element of uncertainty. However, the probabilistic nature of patents is "especially striking and fundamental to an understanding of the effects of patents on innovation and competition" (Lemley and Shapiro p. 4 2005). As such, members of the knowledge community may respond over time and adapt to the actions of community members and also those who may be outside the community as they understand the strength of patent rights, and the willingness on the part of IP owners to impose those rights.

Learning to Live with the Law: Shock, Tax & Adaptation

What then are the implications of the adapting to the law approach from the perspective of debates over patents in the scientific community? When placed in a equilibrium context, contrary empirical results such as those examining the impact of patents on the scientific community can only be regarded as being in opposition. However from an adaptive perspective, empirical analyses of the law generated in different time periods may instead be capturing the differential responses of the scientific community over time to the imposition of patents. Based on recent qualitative analysis of the mouse genetics community, together with current theories regarding the role of patents on knowledge communities we argue that the effect of patents on the production of scientific knowledge should be separable into three distinctive elements: shock, tax and adaptation illustrated in Figure 1.

-- Insert Figure 1 about here --

We argue that in the aftermath of patent grant, regardless of the long-run equilibrium outcome, there will be a *patent shock* - an immediate negative impact of the grant of IP rights on knowledge workers. With regard to the production and disclosure of knowledge through academic publications, IP rights will have a chilling effect on the productivity of knowledge workers. This prediction is grounded in a tradeoff between two countervailing theoretical perspectives on the role of patents in knowledge work. On the one hand, patents provide incentives for research investment, and, while many discoveries 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 (Nelson 1959; Arrow 1962). 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 and therefore spur overall knowledge production. However, in contrast to this literature, others have highlighted the potential costs imposed by IP rights over scientific knowledge traditionally disclosed only through publication. The balance of empirical evidence suggests that in the early period of patent grant, individual scientists, their Technology Transfer Office (TTO) professionals, corporate lawyers who license the patents and execute the license on others had limited expertise and often sought to impose their license not only on other commercial firms but also on academic scientists. Moreover, the ability of universities to execute university-to-university licensing and related material transfer agreements was clumsy and complex (Mowery and Ziedonis 2007). Thus we argue that early in the emergence of IP rights, the net initial patent shock will be significant and negative – in other words productivity will decline.

For each generation of patents, we hypothesize that there will be a *patent tax* imposed on follow-on productivity. This patent tax will accumulate each year, meaning that the salience of a patent on the productivity of a specific research line will be increasing (i.e. more negative each year). This effect is driven by several factors. First, as argued by scholars proposing the "anti-commons" effect, 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). Because IP can serve to exclude follow-on researchers from exploiting scientific discoveries, the anti-commons hypothesis posits that the privatization of the scientific commons will impose a "tax" on the use of prior scientific knowledge through significant transaction costs (Eisenberg, 1996; Shapiro, 2001; Hall and Ziedonis, 2001). In its precise formulation, the anti-commons effect is grounded in the proliferation and fragmentation of IP rights (Ziedonis, 2004; Huang and Murray 2008), the associated transaction costs and potential for royalty stacking. Therefore, for any given patent, the tax it imposes on follow-on research who seek to build on an individual piece of knowledge will be increasing in time as related (and fragmented) IP may be generated – by the original researcher or by other competing research teams for whom there will likely be strong incentives to take advantage of the protections afforded by formal IP rights, even though the scientific community as a whole benefits from the free dissemination and diffusion of knowledge. Growing awareness of the role of IP rights may also cumulatively lead researchers to exit a research line rather than compete in an area that is fraught with complex licensing requirements – as was observed in the mouse genetics community, many of whom exited from transgenic mouse research due to the costs imposed by DuPont and the patent license. A final factor contributing to the negative patent tax is the growing reach of the TTO and patent licensees which may lead to greater enforcement of the terms of licensing requirements.

Notwithstanding the initial and increasing costs of any generation of patents on the scientific research covered by the IP rights, we hypothesize that we will observe the declining impact of all patents in every calendar year from the initial year of patent grant in our analysis which we refer to as *patent adaptation*. The notion that any scientific community adapts to patents over time is grounded in the notion that while patents provide strong property rights, over time, knowledge workers learn how likely these rights are to be imposed, over whom and for what types of infringing actions. As the mouse genetics community came to realize, through orchestrated resistance to the patent, the mobilization of powerful organizations such as the National Institutes of Health (NIH), and pressure on TTO professionals, it was possible to adapt to patents. As a result, even while patents may be proliferating, the impact of successive patent generations on the academic community may be decreasing. We also argue that adaptation through the diminution of the cost of patents may also be accompanied by a rise in the benefits of patents. Specifically, over time, patents may contribute to the effective functioning of the market for ideas (Merges & Nelson, 1990, 1994; Arora, Fosfuri, and

Gambardella, 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, 2005; Aghion et al. 2005).

Empirical Approach

Empirical Setting – the Life Science Community

Many of the issues that currently animate the debate over patents and the academic community are focused on the life science communities (Kenney, 1986; Orsenigo, 1989; Powell et al., 1996; Gambardella, 1995). In this setting, the rise in patenting over knowledge produced by knowledge workers in academia has been driven by a number of factors: The expanding promise of biotechnology, reductions in the costs of academic patenting, and increases in the scope of IP over knowledge produced in the life sciences. The rise in useful, inventive knowledge in this field can be traced back to the early 1970s. At the same time, policy shifts encouraged academics to claim IPR over their 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 IP (generated using Federal funds) to universities along with a duty to license the patents and facilitate their translation and commercialization (Mowery et al 2001). In this context, the traditional justification of IP rights relates not to scientific knowledge accumulation but rather to cumulative commercial innovation. IP provides incentives for further commercial investment. Finally, there was a significant expansion in the *scope* of patents available in the life sciences. After the 1980 *Diamond vs. Chakrabarty* decision and the granting of the Oncomouse patent in 1988, IP comprehensively covered the domain of genetically modified living organisms – from bacteria to mammals (Kevles 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 IP over scientific knowledge. The new institutional environment shifted the likely set of disclosure decisions of faculty to more frequently include patenting. Specifically, faculty more frequently disclosed their knowledge as patent-paper pairs - rather than simply publishing their ideas, faculty also instantiated the same underlying knowledge in patents. In the period between 1989 and 1999 US Research One universities received over 6,000 life science patents (Owen-Smith and Powell 2003), patent-paper pairs became an important disclosure phenomenon (Murray 2002), and by the late 1990s, over 25% of life science faculty had patents (Ding, Murray and Stuart 2007). In short, the life science commons were increasingly covered by intellectual property rights and the stage was set for the debate over law in the life science community.

Identifying the Impact of the Law on Knowledge Production

While the growing scope of the patent landscape over life science knowledge such as human genes has been documents (Jensen and Murray 2005), the ability to undertake large-scale empirical analysis of the impact of these patents on knowledge production is more challenging in part due to a lack of appropriate quantitative methods. In particular, the question of whether the law shapes the production of knowledge, and more specifically whether and how the grant of IP rights effects over knowledge effects the use of that knowledge by follow-on researchers, is difficult to measure empirically. We confront these challenges by attending to three characteristics of knowledge work:

- "Pieces of knowledge" are disclosed and described in publications and in patents making them available (under a variety of access terms) for follow-on innovators to build upon.
- Citations to knowledge disclosed (in patents or publications) are a proxy of follow-on innovation building and knowledge accumulation that also identify follow-on researchers.
- After a piece of knowledge is produced the institutional environment in which it is available to others can be subject to a variety of shifts which can potentially impact the rate and nature of follow-on innovation.

First, large-scale empirical work requires a sample of "pieces of knowledge" that are comparable to one another although their organizational and institutional environment might differ. Since the widespread availability of comprehensive patent statistics, patents have become a staple measurement tool in studies of knowledge and innovation (Hall, Jaffe and Trajtenberg, 2001). Notwithstanding all the caveats associated with patents, many studies aggregate patent data at the firm level as a starting point for the analysis of knowledge building (Almeida and Kogut 1999, Cockburn and Henderson 1998; Ahuja 2000; Hoeteker and Agrawal 2007; Ziedonis 2004; Fleming and Sorenson 2004). Another stream in the literature examines the productivity of individual knowledge workers, used patents to identify knowledge workers - in both academia and industry - and trace ongoing knowledge production (Paruchuri et al., 2006). Given our interest in the role of law on knowledge generated in the scientific community, we are interested in the way in which patent grant impacts the use of the underlying knowledge. We therefore turn to a different measure of knowledge – scientific papers. Less extensively analyzed, particularly by organization scholars, papers are "inscriptions" that provide another codified form of knowledge disclosure (Latour and Woolgar 1979). Particularly among academic scientists, they not only serve to disclose and describe knowledge but they also establish priority, provide a critical "currency" for prestige and promotion, and bring credit (Merton 1973; Biagioli 1998). Studies in the sociology of science have made widespread use of publication data to measure individual productivity (Levin and Stephan 1991). Organization studies have also attended to the production of papers by firms as measure of innovative output in addition to patents (Gittelman and Kogut 2003; Gittelman 2007).

Having defined a sample of "pieces of knowledge", any analysis of knowledge building or knowledge accumulation must track the incorporation of that knowledge into follow-on innovation. For both patents and papers follow-on citations (by the original knowledge workers – self-citations and by others) provide a "knowledge trail." In using citations we follow a long literature that has used both patent and publication citations to trace the flow of ideas and their follow-on use and accumulation by others in their ideas (de Solla Price, 1965; Jaffe & Trajtenberg, 1996). In the case of patents, the inclusion of citations in follow-on innovations is required by law when those innovations are considered to be "prior art". Thus a citation indicates that the patent has been built upon by others (see Hall, Jaffe and Trajtenberg 2001) and makes it possible to it is possible to track knowledge accumulation across people, firms, countries, regions, and time (Almeida et al. 2007). We have chosen to use publication citations rather than patent citations because of our particular phenomenon of interest; whether there are changes in the diffusion and follow-on use of scientific papers as a result of that knowledge also being claimed in patents. We rely on the seminal work of Merton (1973), Hagstrom (1965) and de Solla Price (1965) in articulating the importance of priority and publication in the system of scientific rewards and noting the importance of publication citations in tracking the rate and direction of scientific progress.[1] Like patent citations, paper citations allow us to track follow-on innovators by the nature of their affiliation (including status, location, organizational type), the type of follow-on research, and the mode of follow-on research (collaborative or single authored). Unlike patent citations, scientific citations are more informal and are not enforced by the law. They are however part of a strongly enforced community norm of scientific exchange in the scientific community (Hagstrom, 1965; Merton, 1988) and therefore provide a useful index of the degree to which an idea is incorporated into follow-on research (Cole, 2000). They also serve to contextualize and frame the contributions of a particular scientific idea (Latour, 1987) and hence have an important meaning in the scientific literature. There is a long history of empirical analysis of scientific citations to measure networks of scientists (Crane, 1969). These citations may represent only the tip of the iceberg of research that builds on a published piece of knowledge but they are one critical element needed by follow-on researchers for continued knowledge accumulation (Murray & O'Mahony, 2007).

The third element of our empirical approach tackles the problem of how to identify whether and how institutional changes impact the rate at which follow-on innovators build upon a piece of knowledge. Without a well identified "experiment" it is difficult to disentangle the multiple factors that might underlie an observed change in the pattern of follow-on knowledge accumulation. For example, if as a starting point we were to compare knowledge associated in one institutional environment (e.g. with patents) and knowledge in another (e.g. without patents) we would confront several possible explanations for any observed differences in follow-on use. First, knowledge associated with patents may simply be different in *quality* (higher or lower) compared to unpatented knowledge i.e. the citation curve might be shifted up or down. Indeed there is a growing body of evidence that scientific knowledge published in academic journals that is also patented is, on average, more highly cited over its lifetime than unpatented articles (from the same journal) (Huang and

[1] We recognize that bibliometric analysis is a noisy indicator of scientific progress (see, e.g., Garfield (1979) and Schubert and Braun (1993)): For a number of reasons, small differences in the citation rate of a single paper (particularly early in its publication history) are of limited value in distinguishing the importance of research or its use by the research community. We take care to minimize the impact of these limitations by drawing comparisons among large samples of publications, comparing across control samples, and assessing the impact of policy changes by drawing comparisons within articles across time.

Murray 2008; Lissoni and Montobbio 2007; Murray & Stern, 2007). Moreover, high quality scientists are more likely to patent than low quality scientists (Agrawal and Henderson 2002; Azoulay et al., 2007). Second, knowledge associated with patents may be *different* from non-patented knowledge and might thus exhibit different citations patterns i.e. the shape of their citation curve may be different. It is only having accounted for "quality" and "difference" in our analysis that we examine whether changes in the institutional environment in which the knowledge is embedded actually lead to differences in observed follow-on use patterns among scientists with different organizational affiliations.

From an experimental perspective, scholars of knowledge work would ideally observe a given piece of knowledge in distinct institutional environments and compare the impact of that knowledge across regimes. To do so, our analytical framework relies on the fact that institutional changes may induce changes in the production of scientific articles or changes in citation behavior relative to baseline levels.[2] Moreover this natural experiment approach exploits the fact that the institutional environment changes over time in ways that do not impact the original "piece of knowledge" but which do impact the incentives and opportunities for follow-on researchers to exploit that piece of knowledge in their own research. In particular, we investigate the extent to which changes in institutional environments induce changes in the production of or citations to scientific articles in the pre- and post- policy (treatment) period, relative to a set of control articles which are not impacted by the policy treatment. To the extent that these changes are exogenous, observation of citation patterns in the pre- and post- period help us evaluate the precise impact of the change.

The organizations and economics literature has used a variety of shifts in the environment of a piece of knowledge as an empirical strategy to deepen our understanding of innovation processes (Rysman & Simcoe, 2007; Hoetker & Agarwal, 2007). Most relevant to our analysis are studies that examine the impact of shifts in access to scientific knowledge on scientific citations: For example, Furman and Stern (2006) have explored

---

[2] There are, of course, some important caveats to this approach. First, not all research is disclosed in the scientific literature; indeed, for-profit entities may decline to publish research results either to increase the costs of rivals' research or in the event that such results are disadvantageous for the firm. Second, a increases citations (relative to a baseline) may occur not because of the increased importance of a particular 'unit' of knowledge, but simply because of the ease of its availability relative to alternative pieces of knowledge or for other reasons (such as changes in author prominence or position) that do not reflect changes in the actual use of knowledge. Such problems would average out across the areas we study, unless these changes are closely correlated with the specific policy or institutional changes we study.

whether depositing enabling scientific materials linked to specific publications increases citations to the related papers. More recently, scientific citations have been used to assess the impact of policies to provide unencumbered access to patented research materials by observing changes in citations in the pre and post access period (Aghion et al., 2007).

Identifying the Impact of Patents on Follow-on Innovation- Patent-Paper Pairs

For the purposes of this paper, the specific shift in access we are interested examining is the granting of IP rights over scientific knowledge and the impact of this shock on the follow-on use of knowledge disclosed in scientific papers. The empirical phenomenon motivating our study - the expansion of patents over knowledge traditional disclosed only through publications – has an observable empirical implication which we will exploit as a source of identification. As noted above, with increasing frequency scientists generating scientific knowledge make a decision to disclose that knowledge in a scientific paper *and* to seek patent rights over that knowledge. This leads to a phenomenon referred to as patent-paper pairs – a paper and its "paired" patent that disclose the same piece of knowledge (Ducor, 2000; Murray, 2002; Murray and Stern 2007b). 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:

> "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." (Snell; Kristi D. (Belmont, MA); Hogan; Scott A. (Troy, MI); Sim; Sang Jun (Seoul, KR); Sinskey; Anthony J. (Boston, MA); Rha; Chokyun (Boston, MA) 1998, 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 SJ, Snell KD, Hogan SA, Stubbe J, Rha CK, Sinskey AJ , Nature Biotechnology 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.

The use of patent-paper pairs as an identification strategy for the impact of IP rights over published scientific knowledge has been explored in several prior papers (Murray and Stern 2007a; Huang and Murray 2008; Sampat 2005). As these studies elaborate, prior to patent grant, knowledge in papers is available to follow-on researchers in an essentially open and public environment. Post patent-grant the same knowledge is subject to patent rights which therefore shift a piece of knowledge into a private property environment. The legal shock to the environment in which follow-on researchers use the knowledge takes place with the grant of patent rights claiming the knowledge. The identification embodied in patent-paper pairs relies on the patent-grant delay – a substantial gap between the date of scientific publication and the date at which the associated patent is granted.[3] This empirical technique exploits the insight that while publication in the scientific literature often occurs within six months (or less) after initial submission to a journal, the delay between the initial application and receipt of a patent is often many years (in most cases 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

[3] The specifics of patent law regarding the timing of disclosure are complex and have been subject to change. Under US patent law, inventors have a grace period of twelve 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.

detailed negotiations about the scope and extent of the patent grant, and so there is significant uncertainty about the extent of IPR prior to grant (Cockburn, Kortum, and Stern, 2002; 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 and then a regime shift into the post-patent grant period. Recent research shows that this shift from uncertain to certain property rights is salient in terms of the timing patent licensing (Gans, Hsu & Stern, 2007).

The patent-paper pair design and the identification strategy it embodies makes a number of assumptions. We assume that the impact of intellectual property on the use of knowledge by *follow-on* researchers is (conditionally) independent of the patent filing decision. This is in spite of the fact that the decision to patent a piece of knowledge is endogenous to the specific circumstances of individual researchers, including factors such as their institutional affiliation and their gender (Azoulay, Ding and Stuart, 2007; Ding, Murray, and Stuart, 2007; Markiewicz and Diminin 2004; Agrawal and Henderson 2002). A second critical assuming is that the timing of patent grant is random and is not anticipated by those who use and cite the paper, such that the impact of patent grant on follow-on is observable only in the post-patent grant period. In other words, if patents matter, the *rate* of paper citation before and after patent grant should be different. 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, 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). For the purpose of the current analysis, we assume that at least some potential follow-on researchers may experience the patent grant as "news," although 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.[4]

We also rely on the fact that there is no strategic citing pre and post patent-grant (i.e. there are no authors who continue to build on a particular paper but deliberately choose not to cite that paper in the post-grant period). If this was the case, then a reduction in citations would not represent a reduction in the follow-on use of a piece of knowledge. Finally, we assume that citations to the paper in follow-on papers are a good proxy for follow-on use. In particular we argue that a shift to secrecy would still represent a reduction in follow-on use because the reduction in *disclosed* knowledge would have implications for future generations of scientists. As a caveat we note that the paper citation measure does not capture the positive potential benefits of patenting through citations to the paired patent in other patents.

Data & Methods

Patent-Paper Pairs Sample

The first step in our approach is to collect a sample of published scientific research articles which are of roughly similar "quality" (though we will account for quality variation among articles in our empirical framework) and which disclose knowledge that is potentially patentable (whether or not the researchers choose to apply for IPR). The sample for this study is drawn from a population of 340 scientific research papers analyzed in prior work by Murray and Stern (2007a). The papers are published in a narrow time window (1997-1999) in a top-tier research journal *Nature Biotechnology* (impact factor over 20). The journal's

---

[4] 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 which 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 so are unlikely to strategically enter 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 proactively monitor IPR grants in their research areas.

policy explicitly aims at research with potential applications to biotechnology: *"[the journal] aims to publish high-quality original research that describes the development and application of new technologies in the biological, pharmaceutical, biomedical, agricultural and environmental sciences, and which promise to find real-world applications in academia or industry. We also have a strong interest in research that describes the application of existing technologies to new problems or challenges, and basic research that reports novel findings that are directly relevant and/or of interest to those who develop biology into technology*." 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 and therefore "at risk" of being associated with a USPTO patent and thus forming a patent-paper pair.

While the journal publishes scholarly material in a variety of formats, the Murray-Stern (MS) dataset is confined to research articles - defined by the editorial policies of the journal as "a substantial novel research study" (see *Nature Biotechnology*, A Guide to Authors). For each of the 340 articles it was 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. A number of approaches to this pairing have been devised (Ducor 2000, Murray 2002, Lissoni and Montobbio 2007, Franzoni and Scellato 2007, Huang and Murray 2007). In this instance, 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. 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*. 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" i.e. a verification of whether the material described in the abstract of the article was incorporated into the description, claims and/or examples of the granted patent.[5] Using this procedure, 169 of the 340 articles were found to be associated with a paired patent as of October, 2003. In other words,

[5] One of the authors 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.

approximately half of all publications in *Nature Biotechnology* are associated with a patent-paper pair within five years of publication.

For the purposes of this analysis we sampled only those papers authored by public United States institutions. Under these criteria, the dataset consisted of 174 unique research articles of which 93 (53%) are associated with a granted United States patent. Murray and Stern (2007a) show that the grant of IP rights has the greatest impact on forward citations to papers published by authors affiliated with public-sector institutions. It is therefore the most salient context in which to examine the longer term and dynamic influence of patent rights on knowledge accumulation. Our sampling decision was driven by three additional factors. First, the theoretical debate on the role of intellectual property rights over knowledge previously in the public domain is most salient for research generated by public sector researchers. Indeed the surprise for knowledge generated in the private sector is that it is published in the peer-reviewed literature at all –patents are the traditional disclosure the mode for knowledge production in industry. Second, by focusing on one setting for knowledge production and patent enforcement we limit the range of possible mechanisms at work as follow-on researchers learn to live with the law. Third, paired patents are generated from searches of the US patent office. We believe that non-US researchers may patent in outside the US first (or only outside the US) leading to inaccuracies in our characterization of IP rights over non-US authored research.

For each of the 174 articles and 93 patents we gathered variables on observable characteristics: number of authors/inventors, number of institutional addresses/assignees, date of publication/application etc. Finally we gathered all the forward citations to the 174 research articles. Among the complete set of forward citations we selected only those that were designated as "research articles" according to ISI Web of Science. This amounted to 14,688 forward citations each of which was then coded for a series of variables: number of authors, the number and type of institutional affiliations (public versus private sector), the rank of institutional affiliation, and the country of institutional affiliation.

Empirical Specification

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 Poisson). Therefore, except where noted, we employ a negative binomial model of the annual citations 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 for which the citations are being considered. 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.[6] As an overall measure of the impact of patent grant, the baseline model that is specified by Murray and Stern (2007a) incorporates a dummy variable POST-GRANT equal to one only for years after the patent grant year for an individual article. By observing citations to a scientific publication both before and after the patent is received (and because we observe a control group of similar publications which 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.[7] Specifically, this baseline estimator is:

$$CITES_{i,j,pubyear(j),t} = f(\varepsilon_{i,j,t}; \gamma_i + \beta_t + \delta_{t-pubyear} + \psi POST-TREATMENT_{i,t}) \qquad (1)$$

where (γ_i) is a fixed effect for each article, β_t is a year effect, $\delta_{t\text{ - pubyear}}$ captures the age of the article, and POST-TREATMENT is a dummy variable equal to one only for years after the knowledge linked to the article is affected by the institutional or policy change. The coefficient on POST-GRANT (□) indicates the

---

[6] 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, Griliches, and Hall, 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).

[7] 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.

marginal impact of the intervention on the set of treated articles. Thus, we test for the impact of patenting by calculating how the citation rate for a scientific publication *changes* following such interventions, 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.

While this specification provides an aggregate assessment of the impact of the IP rights on forward citations in the years following patent grant it does not provide any insight into the dynamic nature of intellectual property rights as they shape follow-on researchers with different affiliations. In order to tease out these effects we provide a more nuanced baseline specification with three variables capturing the varying forces that contour the impact of patent grant on follow-on scientific research. We identify three distinctive parameters to account for the changing impact of patents on forward citations over time. The first is the "baseline" impact of patent with is the initial "shock" of a piece of knowledge moving into the post-grant institutional regime. The next is a "patent tax" with identifies the impact of the patent in each year following patent grant and provides an estimate of the trend associated with a given patent on citations to the paired publication in the years after the initial "baseline" impact. Finally we identify an "adaptation" or learning variable which pertains to the impact of any patent in a given calendar year after 1999 – the initial year in which patents are granted in our sample. This variable is intended to capture the waning impact of all patents enforce in a given year.

Taken together, our more dynamic empirical test therefore:

(2)
$$\begin{aligned} CITES_{i,t} = f(\varepsilon_{i,t}; \gamma_i + \beta_t + \delta_{t-pubyear} + \psi_0 POST-GRANT_{i,t} \\ + \psi_{PatentTax}(t - grantyear_i) * POST-GRANT_{i,t} \\ + \psi_{Adaptarion}(t-1999) * POST-GRANT_{i,t}) \end{aligned}$$

While the preceding analysis focuses on the impact of institutional changes on the overall count of citations to a given piece of knowledge, as we have outlined, the grant of IP rights is likely to have quite different implications for different subpopulations. To do so, we take advantage of the citation-level data that facilitates detailed coding of the types of citations that are received by *Nature Biotechnology* and break down forward citations into subpopulations: academic vs. industry; high status vs. low status; collaborative vs. non-collaborative. To estimate the dynamic impact of patent grant on each of these subpopulations, we can aggregate these individual citations into counts of the number of citations received by a given article in a given year by a given subpopulation of citers:

$$(3)\quad CITES_{i,l,t} = f(\varepsilon_{i,j,t}; \gamma_i + \lambda_l + \beta_t + \delta_{t-pubyear} + \sum_{l=1,\ldots L} \psi_l \iota_l POST-TREATMENT_{i,t})$$

In other words, ψ_l is the average impact of the treatment on sub-population *l*, conditional on a fixed effect for each article, and age and citation-year fixed effects.

Results

Our results proceed in several stages. In Table 3, we begin by first replicating the earlier results from the MS dataset. The dataset employed here is different in two ways. First, and most importantly, our sample is focused on the set of papers with *at least* one US author, a strategy we employed to further enhance the matching and similarity between patented and unpatented articles. As well, the precise number of citations for individual articles in a given year may differ because our current sample is drawn from a "micro" dataset composed of individual citations (as opposed to a manual count by hand of citations per article per year. Despite these modest changes, the basic patterns of results for the sample that is *most similar* to that used in our early analysis accords with our prior findings. Specifically, when we examine citations through 2002, either focusing on the full sample or only those where a patent is actually received, we observe a significant decline in the rate of citation after patent grant (similar to our earlier findings, we find 13% decline for citations through 2002 using both patented and unpatented papers). However, when we extend the analysis to the sample through 2005, we observe a very different "aggregate" result – patents are associated with a modest and marginally significant positive increase in the citation rate). In other words, as we increase the period of time of our sample, our evidence for a modest "anti-commons" effect is reduced, and indeed we find some evidence consistent with the development of a "market for ideas." The remainder of our empirical analysis explores this pattern in a more structured way.

Specifically, in Table 4, we present our main evidence for the role of communitywide learning in the impact of patents on cumulative scientific research. As discussed in the empirical framework section, our specification accounts for three separate effects:

- The "baseline" impact of a patent granted in 2000 for citations during the year 2000
- The "patent tax" which captures the increasing impact of a patent as the years since the patent is granted increases
- The "patent learning" trend which captures how the "baseline" effect is changing over time, relative to the year 2000.

The results are quite striking. First, and most importantly, while the baseline effect is quite negative (a 25% decline in the citation rate) and the patent tax is also significant (the negative impact of patents increases at a rate of 13% per year since the date of grant), the rate of learning is also impressive – with a 19% increase in the "baseline" impact of patents each calendar year. This result implies that, by 2003, the "net" impact of a patent in its first year of grant was actually positive, and that, by 2005, the impact of the patent system was a net "positive" for essentially all patent vintages. This effect is documented even more strikingly in the second column where we estimate a "baseline" effect for each year, from 2000 to 2005. The predicted "baseline" impact of patents becomes more favorable in each and every year, going from a 30% reduction in 2000 to more than a 70% predicted increase in citations as of 2005.

Learning in Different Research Communities

While the results in Table 4 provide useful evidence for our core hypotheses related to the impact of patents over time on the aggregate research community, our detailed micro-data allows us to evaluate these issues more precisely by comparing the impact of patents across different communities. We begin in Table 5-I examining the difference in the impact of patents for public sector and private sector authors. The results accord well with our hypotheses. Specifically, while patent grant has very *little* impact on private sector behavior, our core results of a negative baseline effect and then learning between 2000 and 2005 are well captured in the behavior of public sector citation behavior. In other words, while most companies likely have procedures and experience in conducting innovation, public sector researchers seemed to have faced significant costs in managing intellectual property at the beginning of our sample and have become more adept at that over time.

A similar pattern can be observed across the remaining "cuts" of the data. In particular, there seems to be significant differences between "high-quality" journal publication versus other journals, and, most striking, our results are particularly salient for research teams located in a single institution. Overall, our pattern of results accords with a model where the impact of patents has changed significant in life sciences research over time, and the "anti-commons" environment of the late 1990s seems to have been mostly replaced by a more productive "market for ideas."

Conclusions

Our analysis provides us with an opportunity to explore whether and to what extent research scientists have been able to incorporate patenting into their daily scientific activities, without any concomitant loss of productivity. Recent qualitative analysis has suggested that while patents, particularly those incorporated into aggressive licensing strategies, have been the source of much dissent and resistance among scientists, that over time scientists learn to accommodate patents. They have developed strategies to either work alongside the intellectual property rights, work around them, or use patents to further their scientific work – for example by co-opting industry into developing more standardized methods and materials. However there is little or no quantitative evidence to support this perspective; current evidence has been mixed with some finding evidence that patents do decrease productivity (as measured by forward citations) and other research suggesting that patents have no impact. Using longitudinal analysis of the citation patterns to a large sample of publications which are also subject to intellectual property rights, and incorporating the identification strategy that is afforded by the patent grant delay, we are able to disentangle these two competing claims. We find that in fact both claims are correct. But more importantly we find strong empirical support for the precise way in which they are both correct: in the early period, patents were detrimental to follow-on research, slowing progress with every year that elapsed since patent grant. However for every new generation of patents, the impact of the patent was dampened, presumably as the licensing offices, corporations and scientists learned how to live the patents. Thus patents have become commonplace in laboratory life and as a result, scientists seem to be able to proceed with their research, either allowing patents to serve only as a mechanism to facilitate the transfer of knowledge into commercialization or (as other research suggests) to maintain an open commons for ongoing scientific exploration.

References

Aghion, P., M. Dewatripont, J. Kolev, F. Murray, S. Stern. 2007. Of Mice and Growth: The effects of openness on follow-on research. Working Paper.

Aghion, P., P Howitt. 1997. Competition, imitation and growth with step-by-step innovation. *Review of Economic Studies*. **68** (3) 467-92.

Agrawal, A., R. Henderson. 2002. Putting patents in context: Exploring knowledge transfer from MIT. *Management Science*. **48**(1) 44-60.

Ahuja, G. 2000. Collaboration networks, structural holes, and innovation: A longitudinal study. *Administrative Science Quarterly*. **45**(3) 425-455.

Almeida, P., B. Kogut. 1999. Localization of knowledge and the mobility of engineers in regional Networks. *Management Science*. **45**(7) 905-917.

Arora, A., A. Fosfuri, A. Gambardella. 2001. *Markets for Technology: Economics of Innovation and Corporate Strategy*. MIT Press, Cambridge, MA.

Arrow, Kenneth. 1962. Economic Welfare and the Allocation of Resources for Invention. Richard R. Nelson, ed. *The Rate and Direction of Inventive Activity*, Princeton University Press, Princeton, NJ 609-25.

Azoulay, P., W. Ding, T.E. Stuart. 2007. The determinants of faculty patenting behavior: Demographics or opportunities. *Journal of Economic Behavior and Organization*. **64**(4) 599-623.

Bar-Gill, O., G. Parchomovsky. 2003. The Value of Giving Away Secrets. *Virginia Law Review* **89** 1857-1895.

Bechky, Beth. 2003. Sharing meaning across occupational communities: The transformation of knowledge on a production floor. *Organization Science*. **14** 312-330.

Biagioli, Mario, P. Galison, eds. 2003. *Scientific Authorship: Credit and Intellectual Property in Science*. Routledge, New York.

Biagioli, Mario. 2000. Replication or monopoly? The economies of invention and discovery in Galileo's observations of 1610. *Science in Context*. **13**(3-4) 547-590.

Biagioli, Mario. 2006. *Galileo's Instruments of Credit: Telescopes, Instruments, Secrecy*. University of Chicago Press, Chicago.

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.

Brown, J. S., P. Duguid. 2001. Knowledge and organization: A social-practice perspective. *Organization Science*. **12**(2) 198-213.

Brown, J. S., P. Duguid. 1991. Organizational learning and communities-of-practice: Toward a unified view of working, learning, and innovation. *Organization Science*. **2**(1) 40-57.

Burt, R. 2004. Structural holes and good ideas. *American J. of Sociology*. **110** 349-399.

Campbell, E. G., B. R. Clarridge, M. Gokhale, L. Birenbaum, S. Hilgartner, N.A. Holtzman, D.Blumenthal. 2002. Data withholding in academic genetics: Evidence from a national survey. *Journal of the American Medical Association*. **287**(4) 473 - 479.

Carlile, P. 2004. Transferring, translating, and transforming: An integrative framework for managing knowledge across boundaries. *Organ. Science*. **15**(5) 555-569.

Cockburn, I., R. Henderson. 1998. Absorptive capacity, coauthoring behavior, and the organization of research in drug discovery. *Journal of Industrial Economics***.** **46**(2) 157-182.

Cockburn, Iain, S. Kortum, S. Stern. 2002. Are all patent examiners created equal? Examiners, patent characteristics and litigation outcomes. Intellectual Property in the Knowledge-Based Economy. *National Academy of Sciences* 2003.

Cole, J. 2000. A short history of the use of citations as a measure of the impact of scientific and scholarly work. E. Garfield, B. Cronin & H. B. Atkins eds. *The web of knowledge: A festschrift in honor of Eugene Garfield*. Information Today, Medford, NJ, 281 -300.

Crane, D. 1969. Social structure in a group of scientists: A test of the invisible college hypothesis. *American Sociological Review*. **34**(3) 335-352.

Dasgupta, P., P. A. David. 1994. Towards a new economics of science. *Research Policy*. **23** 487-521.

David, Paul. 2001. From keeping 'Nature's Secrets' to the institutionalization of 'Open Science'. Working Paper #01-006, Stanford.

David, Paul. 2003. Can 'Open Science' be protected from the evolving regime of IPR protections? Working Paper #03-011, Stanford.

de Solla Price, D. J. 1965. Networks of scientific papers. *Science*. **149**(3683) 510-515.

Ding, W., F. Murray, T. E. Stuart. 2006. Gender difference in patenting in the academic life science. *Science* **313**(5787) 665-667.

Kelly E., Dobbin, F. 1999. Civil rights law at work: Sex discrimination and the rise of maternity leave policies. *American Journal of Sociology*. **105**(2) 455-492.

Dobbin F., E. Kelly. 2007. How to stop harassment: The professional construction of legal compliance in organizations. *American Journal Of Sociology*. **112**(4) 1203-1243

Ducor, P. 2000. Intellectual property: Coauthorship and coinventorship. ***Science,*** 289: 873-875.

Edelman, Lauren, M. Suchman. 1997. The legal environment of organizations. *Annual Review of Sociology.* **23** 479-515.

Eisenberg, R. 1996. Public research and private development: Patents and technology transfer in government-sponsored research. *Virginia Law Review* 1663-1727.

Eisenberg, R., M A Heller. 1998. Can patents deter innovation? The anticommons in biomedical research. *Science.* **280**(5364) 698-701.

Etzkowitz, H. 1998. The norms of entrepreneurial science: cognitive effects of the new university-industry linkages. *Research Policy.* **27** 823-833.

Fauchart, E. E. Von Hippel. 2007. Norms-Based Intellectual Property Systems: The Case of French Chefs. *Organization Science.*

Fleming, L. 2001. Recombinant uncertainty in technological search. *Management Science.* **47**(1) 117-132.

Fleming, Lee, O. Sorenson. 2004. Science as a map in technological search. *Strategic Management Journal.* **25(**8-9) 909-928.

Furman, J., S. Stern. 2006. Climbing atop the shoulders of giants: The impact of institutions on cumulative research. Working paper, National Bureau of Economic Research, Cambridge, MA.

Gallini, Nancy T. 2002. The Economics of patents: Lessons from recent U.S. patent reform. *Journal of Economic Perspectives.* **16**(2) 131-154.

Gambardella, A. 1995. *Science and Innovation.* Cambridge University Press, Cambridge UK.

Gans, J., S. Stern. 2000. Incumbency and R&D incentives: Licensing the gale of creative destruction. *Journal of Economics and Management Strategy*. **8**(4) 484-511.

Gans, J., D. Hsu, S. Stern. 2007. The impact of uncertain intellectual property rights on the market for ideas: Evidence from patent grant delays. *Management Science***.** Forthcoming.

Garfield, E. 1979. *Citation Indexing*. Philadelphia, PA: ISI.

Gieryn T F. 1983. Boundary-work and the demarcation of science from non-science: strains and interests in professional interests of scientists. *American Sociological Review* **48** 781-95.

Gilbert, R., C. Shapiro. 1990. Optimal patent length and breadth. *RAND Journal of Economics.* **21**(1) 106-112.

Gittelman, M., B. Kogut. 2003. Does good science lead to valuable knowledge? Biotechnology firms and the evolutionary logic of citation patterns. *Management Science*. **49**(4) 366-382.

Greene, William. 2004. Fixed Effects and the incidental parameters problem in the tobit model. *Econometric Reviews*. **23**(2) 125-147.

Hagström, W. O. 1965. *The scientific community.* Basic Books, New York.

Hall, B. H., R. H. Ziedonis. 2001. The patent paradox revisited: An empirical study of patenting in the U.S. semiconductor industry, 1979-95. *Rand Journal of Economics*. **32**(1) 101-128.

Hall, B. H., A. B. Jaffe, M. Trajtenberg,. 2001. The NBER patent citation data file: Lessons, insights and methodological tools. Working paper no. 8498, National Bureau of Economic Research, Cambridge, MA.

Hansen, M. T. 1999. The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits. *Administrative Science Quarterly.* **44** 82-111.

Hansen, M. T., M. L. Mors, B. Lovas. 2005. Knowledge sharing in organizations: Multiple networks, multiple phases. *Academy of Management Journal.* **48** 776-793.

Hargadon, A., R. I. Sutton. 1997. Technology brokering and innovation in a product development firm. *Administrative Science Quarterly.* **42** 716-749.

Hargadon, A., Y. Douglas. 2001. When innovations meet institutions: Edison and the design of the electric light. *Administrative Science Quarterly.* **46** 476-501.

Hausman, J. A., B. H. Hall, Z. Griliches. 1984. Econometric models for count data with an application to the patents-R&D relationship. *Econometrica.* **52** 909-938.

Heller, M. A., R. S. Eisenberg. 1998. Can patents deter innovation? The anti-commons in biomedical research. *Science.* **280**(5364) 698-701.

Heller, Michael. 1998. The tragedy of the anticommons: Property in the transition from Marx to markets. *Harvard Law Review* **111**(3).

Hellman, Thomas. 2005. The role of patents for bridging the science to market gap. *NBER Conference on Academic Science and Entrepreneurship.*

Henderson, Rebecca, I. Cockburn. 1994. Measuring competence? Exploring firm effects in drug discovery. *Strategic Management Journal.* **15** 63-84.

Hoetker, G. R. Agarwal. 2007. Death hurts but isn't fatal: The post-exit diffusion of knowledge created by innovative companies. *Academy of Management Journal***.** **50**(2) 446-467.

Huang, Kenneth, F. Murray. 2008. Golden helix or tangled web: Assessing the effect of gene patents on the use of genetic knowledge. Working Paper.

Hurst WJ. 1964. *Law and Economic Growth: The Legal History of the Lumber Industry in Wisconsin 1836-1915*. Harvard University Press, Cambridge, MA.

Jaffe, A. 2000. The U.S. patent system in transition: Policy innovation and the innovation process. *Research Policy*. **29**(4-5) 531-557.

Jaffe, A. B., J. Lerner. 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, NJ.

Jaffe, A. B., M. Trajtenberg. 1996 Flows of knowledge from universities and federal laboratories: Modeling the flow of patent citations over time and across institutional and geographic boundaries. *Proceedings of the National Academy of Sciences*. **93** 12671-12677.

Jensen, K., F. Murray. 2005. The intellectual property landscape of the human genome. *Science*. **310** 239-240.

Jensen, Richard, M.C. Thursby. 2001. Proofs and prototypes for sale: The licensing of university inventions. *American Economic Review*. **91**(1) 240-59.

Kelves, D. J. 2002. Of mice & money: The story of the world's first animal patent. *Daedalus*. **131**(2) 78.

Kenney, M. 1986. *Biotechnology: The University-Industrial Complex*. Yale University Press, New Haven & London.

Kitch, Edmund W. 1977. The nature and function of the patent system. *Journal of Law and Economics*. **20**(2) 265-290.

Knorr Cetina, K. 1999. *Epistemic Cultures: How Sciences Make Knowledge*. Harvard University Press, Cambridge, MA.

Latour, B. 1987. *Science in Action: How to Follow Scientists and Engineers through Society*. Harvard University Press, Cambridge, MA.

Latour, B. S. Woolgar. 1979. *Laboratory Life: The Construction of Scientific Facts*. Sage Publications, Beverly Hills.

Lemley, Mark, C. Shapiro 2005, Probabilistic patents. *Journal of Economic Perspectives*. **19**(2) 75-98.

Levin, S., P. Stephan. 1991. Research productivity over the life cycle: Evidence for academic scientists. *American Economic Review*. **81**(1) 114-132.

Linder, J., S. Jarvenpaa, T. H. Davenport 2003. Towards an innovation sourcing strategy. *Sloan Management Review*. **44**(4) 43–49.

Lissoni F., Montobbio F. 2007. Guest authors or ghost inventors? Inventorship attribution in academic patents. *Seventh Annual Roundtable For Engineering Entrepreneurship Research College Of Management Room*, Georgia Tech College Of Management.

Macaulay S. 1963. Non-contractual relations in business. *Amercian Sociological Review*. **28** 55-70.

Markiewicz, Kira, A. DiMinin. 2004. Commercializing the laboratory: The relationship between faculty patenting and publishing. Working Paper.

Merges, R., R. R. Nelson. 1990. On the complex economics of patent scope. *Columbia Law Review.* **90**(4) 839-916.

Merges, R.P. R.R. Nelson. 1994. On limiting or encouraging rivalry in technical progress: The Effect of patent scope decisions. *Journal of Economic Behavior and Organization.* **25**(1) 1-24.

Merton, R. K. 1973. The normative structure of science. *The sociology of science: Theoretical and empirical investigations.* The University of Chicago Press, Chicago, IL, 267-280.

Merton, R. K. 1988. The Matthew effect in science, II: Cumulative advantage and the symbolism of intellectual property. *Isis.* **79**(4) 606-623.

Mnookin R.H., L. Kornhauser 1979. Bargaining in the shadow of the law: the case of divorce. *Yale Law Review.* **88** 950-997.

Mokyr, J. 2004. *The Gifts of Athena: Historical Origins of the Knowledge Economy.* Princeton University Press, Princeton.

Mowery, D., B. Sampat. 2001. Patenting and licensing university inventions: Lessons from the history of the research corporation. *Industrial and Corporate Change.* **10**(2) 317.

Mowery, D., R. Nelson, B. Sampat, A. Ziedonis. 2001. The growth of patenting and licensing by U.S. universities: An assessment of the effects of the Bayh-Dole act of 1980. *Research Policy.* **30** 99-119.

Mowery, David C., et al. 2004. *Ivory Tower and Industrial Innovation: University-Industry Technology Transfer Before and After the Bayh-Dole Act.* Stanford University Press, Stanford.

Murray, F., S. O'Mahony. 2007. Re-conceptualizing the institutional foundations of cumulative innovation. *Organization Science.*

Murray, F., S. Stern. 2007. Do formal intellectual property rights hinder the free flow of scientific knowledge? An empirical test of the anti-commons hypothesis. *Journal of Economic Behavior and Organization.*

Murray, Fiona, S. Scott. 2007. When ideas are not free: The impact of patents on academic science. *Innovation Policy and the Economy Meeting.* National Bureau of Economic Research.

Murray, Fiona. 2002. Innovation as co-evolution of scientific and technological networks: Exploring tissue engineering. *Research Policy.* **31**(8-9) 1389-1403.

Murray, Fiona. 2006 The Oncomouse that roared: Resistance and accommodation to patenting in academic science. Working Paper, Sloan School of Management.

Nature Biotechnology 1996. Editorial: A Broad and Inclusive Enterprise. **14**(3) 235.

Nelson, R. 1986. Institutions supporting technical advance in industry. *American Economic Review*. **76**(2) 186-89.

Nelson, R. 1959. The simple economics of basic scientific research. *Journal of Political Economy*. **67**(3) 297-306.

North, Douglass C., 1991. Institutions. *Journal of Economic Perspectives*. **5**(1), 97-112.

O'Mahony, Siobhán. 2003. Guarding the commons: How community managed software projects protect their work. *Research Policy*. **32** 1179-1198.

Owen-Smith, J. 2003. From separate systems to a hybrid order: Accumulative advantage across public and private science at research one universities. *Research Policy*. **32**(6) 1081-1104.

Owen-Smith, J., W. Powell. 2003. The expanding role of university patenting in the life sciences: assessing the importance of experience and connectivity. *Research Policy*. **32**(9) 1695-1711.

Posner R. A. 1972. *Economic Analysis of Law*. Little Brown, Boston.

Powell, W., K. Koput, L. Smith-Doerr. 1996. Interorganizational collaboration and the locus of innovation: Networks of learning in biotechnology. *Administrative Science Quarterly*. **41** 116-145.

Romer, Paul M. 1990. Endogenous technological change. *Journal of Political Economy*. **98**(5) S71-102.

Rosenberg, N. 1982. *Inside the Black box–Technology and economics*. Cambridge University Press, New York.

Rosenberg, N. 1974. Science, invention, and economic growth. *The Economic Journal*. **84**(333) 90-108.

Rysman, M., & Simcoe, T. S. 2005. Patents and the performance of voluntary standard setting organizations. Working Paper No. 05-22, NET Institute.

Sampat, Bhaven. Genomic Patenting: Bad for Science? Unpublished manuscript.

Schubert, A., T. Braun. 1993. Standards for citation based assessments. *Scientometrics*. **26** 21-35.

Scotchmer, S. 1991. Standing on the shoulders of giants: Cumulative research and the patent law. *Journal of Economic Perspectives*. **5** 29- 41.

Scotchmer, S. 1996. Protecting early innovators: Should second-generation products be patentable? *The Rand Journal of Economics*. **27** 322-331.

Scotchmer, S. 2004. *Innovation and Incentives*. MIT Press, Cambridge, MA.

Shapiro, C. 2001. Navigating the patent thicket: Cross licenses, patent pools, and standard-setting. A. B. Jaffe, J. Lerner & S. Stern, eds., *Innovation Policy and the Economy*. MIT Press, Cambridge, MA 1: 119-150.

Stern, Scott. 2004. Do scientists pay to be scientists? *Management Science*. **50**(6) 835-853.

Stokes, Donald. 1997. *Pasteur's Quadrant: basic science and technological innovation*. The Brookings Institution, Washington D.C.

Straus, Joseph, H. Holzapfel, M. Lindenmeir. 2002. *Empirical Survey on "Genetic Inventions and Patent Law"*. Max Planck Institute for Intellectual Property, Competition, and Tax Law, Munich.

Thornton, P. 2004. *Markets from Culture: Institutional Logics and Organizational Decisions in Higher Education Publishing.* Stanford University Press, Stanford, CA.

Uzzi, Brian, J. Spiro. 2005. Collaboration and creativity: Big differences from small world networks. *American Journal of Sociology***. 111** 447-504.

van de Ven, A. H. 1993. A community perspective on the emergence of innovations. *Journal of Engineering and Technology Management.* **10** 23-51.

Veggeberg, S. 1992. Controversy mounts over gene patenting policy. *The Scientist.* **6**(9) 1.

Walsh, John P., A. Arora, W. M. Cohen. 2002. The patenting and licensing of research tools and biomedical innovation. *US National Academies' Science, Technology and Economic Policy Board.*

Walsh, John, A. Arora, W. Cohen. 2003. Science and the law: working through the patent problem. *Science.* 299(5609) 1021.

Walsh, John, C. Cho, W. Cohen. 2005. View from the bench: Patents and material transfers. *Science.* **309**(5743) 2002-2003.

Ziedonis, R. H. 2004. Don't fence me in: Fragmented markets for technology and the patent acquisition strategies of firms. *Management Science.* 50(6) 804-820.

Table 1

Variables & Definitions

VARIABLE	DEFINITION	SOURCE
CITATION-YEAR CHARACTERISTICS		
ANNUAL FORWARD CITATIONS$_{jt}$	# of Forward Citations to Article *j* in Year *t*	SCI
YEAR$_t$	Year in which FORWARD CITATIONS are received	SCI
AGE$_{jt}$	YEAR – PUBLICATION YEAR	NB
CITATION CHARACTERISTICS		
CITE ARTICLE$_i$	Dummy variable equal to 1 if citation is a research article; 0 otherwise	SCI
CITE REVIEW$_i$	Dummy variable equal to 1 if citation is a review; 0 otherwise	SCI
CITER TIER 1$_i$	Dummy variable equal to 1 if citation is published in tier 1 journal; 0 otherwise	SCI / JIF
CITER TIER 2$_i$	Dummy variable equal to 1 if citation is published in a tier 2 journal; 0 otherwise	SCI / JIF
CITE PRIVATE$_i$	Dummy variable equal to 1 if *at least* one of the institutions associated with Article *j* is a biotech or pharma company; 0 otherwise	SCI
CITE PUBLIC$_i$	Dummy variable equal to 1 if *at least* one of the institutions associated with Article *j* is a public entity; 0 otherwise	SCI
CITE US$_i$	Dummy variable equal to 1 if *at least* one of the institutions associated with Article *j* is in the U.S.; 0 otherwise	SCI
CITE MULTI0$_i$	Dummy variable equal to 1 if institutional affiliations with Article *j* contain *only* 1 institution; 0 otherwise	SCI
CITE MULTI1$_i$	Dummy variable equal to 1 if institutional affiliations with Article *j* contain *more than* 1 institution; 0 otherwise	SCI
PUBLICATION CHARACTERISTICS		
PUBLICATION YEAR$_j$	Year in which article is published	NB
# AUTHORS$_j$	Count of the number of authors of Article *j*	NB
US AUTHOR$_j$	Dummy variable equal to 1 if *at least* one of institutional affiliation associated with Article *j* is in the US; 0 otherwise	NB
PUBLIC AUTHOR$_j$	Dummy variable equal to 1 if *at least* one of the institutional affiliation associated with Article *j* is a university; 0 otherwise	NB
PRIVATE AUTHOR$_j$	Dummy variable equal to 1 if *at least* one of the institutional affiliation associated with Article *j* is a biotech or pharma; 0 otherwise	NB
TOTAL CITATIONS$_j$	# of FORWARD CITATIONS from publication date to December 2005	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 AGE$_{jt}$	Age of patent defined as YEAR – GRANT YEAR	USPTO
PATENT POST-GRANT$_j$	Dummy variable equal to 1 if PATENTED = 1 and YEAR > GRANT YEAR	USPTO
PATENT TREND$_j$		USPTO
PATENT TAX$_j$		USPTO
# INVENTORS$_j$	Count of the number of inventors listed in the granted patent associated with Article *j*; 0 if PATENTED = 0.	USPTO

USPTO – United States Patent Office; NB – Nature Biotechnology; SCI – Science Citation Index; JIF – Journal Impact Factor

Table 2
Means & Standard Deviations

VARIABLE	N	MEAN	STANDARD DEVIATION	MIN	MAX
CITATION-YEAR CHARACTERISTICS					
FORWARD CITATIONS	917	11.49	19.39	0	315
CITATION YEAR	917	2001.95	2.09	1998	2005
AGE	917	4.05	2.09	1	8
CITATION CHARACTERISTICS					
CITE ARTICLE	917	8.21	13.77	0	213
CITE REVIEW	917	3.01	5.77	0	97
CITER TIER 1	917	0.37	0.83	0	6
CITER TIER 2	917	0.86	1.62	0	17
CITE NB	917	0.23	0.66	-2	8
CITE PRIVATE	917	1.92	4.70	0	61
CITE PUBLIC	917	10.59	17.29	0	288
CITE US	917	6.40	11.31	0	162
CITE MULTI0	876	5.21	9.01	0	122
CITE MULTI1	876	1.89	3.71	0	54

Table 3
Impact Of Patent Grant:
Difference-In-Difference Estimates Over Different Time Periods

Poisson Specifications	Dep Var = ANNUAL FORWARD CITATIONS [Incident rate ratios reported in square brackets] (Robust coefficient standard errors reported in parentheses)		
	3-1 Cite years 1997-2003 All articles	3-2 Cite years 1997- 2003 Patented articles only (PATENTEDj=1)	3-3 Cite years 1997-2005 All articles
PATENT POST-GRANT	*[0.877]* *(0.065)*	**[0.724]** **(0.075)**	**[1.153]** **(0.083)**
Article FE	Y	Y	Y
Age FE	Y	Y	Y
Citation-Year FE	Y	Y	Y
# Observations	524	337	917
Log-likelihood	-1314.69	-942.91	-2454.28

Table 4

Estimating Temporal Trends In Impact Of Patent Grant

Difference-In-Difference Estimates

Poisson Specifications	Dep Var = ANNUAL FORWARD CITATIONS [Incident rate ratios reported in square brackets] (Robust coefficient standard errors reported in parentheses)	
	4-1 With patent shock, patent tax & patent learning variables	4-2 With annual patent impact variables
PATENT POST-GRANT	*[0.757]* *(0.113)*	
PATENT TAX (ANNUAL)	**[0.876]** **(0.057)**	*[0.875]* *(0.062)*
PATENT LEARNING TREND (ANNUAL)	**[1.190]** **(0.083)**	
PATENT POST_GRANT IMPACT 2000		[0.716] (0.149)
PATENT POST_GRANT IMPACT 2001		[0.984] (0.099)
PATENT POST_GRANT IMPACT 2002		[1.006] (0.085)
PATENT POST_GRANT IMPACT 2003		**[1.301]** **(0.121)**
PATENT POST_GRANT IMPACT 2004		**[1.551]** **(0.288)**
PATENT POST_GRANT IMPACT 2005		**[1.797]** **(0.479)**
Article FE	Y	Y
Age FE	Y	Y
Citation-Year FE	Y	Y

Table 5

Difference-In-Difference Estimates by Institutional Affiliation

Poisson Specifications	Dep Var = FORWARD CITATIONS [Incident rate ratios reported in square brackets] (Robust coefficient standard errors reported in parentheses)	
	5-1 Citations by Public Sector Authors	5-2 Citations by Private Sector Authors
PATENT POST-GRANT IMPACT	**[0.722] (0.112)**	[1.054] (0.281)
PATENT TREND	**[1.216] (0.089)**	[1.006] (0.109)
PATENT TAX	**[0.864] (0.058)**	[1.065] (0.113)
Article FE	Y	Y
Age FE	Y	Y
Citation-Year FE	Y	Y

Table 6

Difference-In-Difference Estimates by National Affiliation

Poisson Specifications	Dep Var = FORWARD CITATIONS [Incident rate ratios reported in square brackets] (Robust coefficient standard errors reported in parentheses)	
	6-1 Citations by US Authors	6-2 Citations by Non-US Authors
PATENT POST-GRANT IMPACT	[0.788] (0.137)	*[0.707] (0.141)*
PATENT TREND	*[1.140] (0.086)*	**[1.267] (0.112)**
PATENT TAX	[0.965] (0.071)	**[0.773] (0.062)**
Article FE	Y	Y
Age FE	Y	Y
Citation-Year FE	Y	Y

Table 7

Difference-In-Difference Estimates by Institutional Status

Poisson Specifications	Dep Var = FORWARD CITATIONS [Incident rate ratios reported in square brackets] (Robust coefficient standard errors reported in parentheses)	
	7-1 Citations by Top Tier Authors	7-2 Citations by Low Tier Authors
PATENT POST-GRANT IMPACT	*[0.693]* *(0.142)*	[0.775] (0.131)
PATENT TREND	*[1.177]* *(0.101)*	**[1.202]** **(0.091)**
PATENT TAX	[0.985] (0.083)	**[0.832]** **(0.058)**
Article FE	Y	Y
Age FE	Y	Y
Citation-Year FE	Y	Y

Table 8

Difference-In-Difference Estimates by Number of Authors

Poisson Specifications	Dep Var = FORWARD CITATIONS [Incident rate ratios reported in square brackets] (Robust coefficient standard errors reported in parentheses)	
	8-1 Citations by Multiple Organization Authors	8-2 Citations by Single Organization Authors
PATENT POST-GRANT IMPACT	[0.822] (0.226)	**[0.745]** **(0.105)**
PATENT TREND	[1.043] (0.126)	**[1.219]** **(0.081)**
PATENT TAX	[1.000] (0.109)	**[0.853]** **(0.053)**
Article FE	Y	Y
Age FE	Y	Y
Citation-Year FE	Y	Y

Figure 1

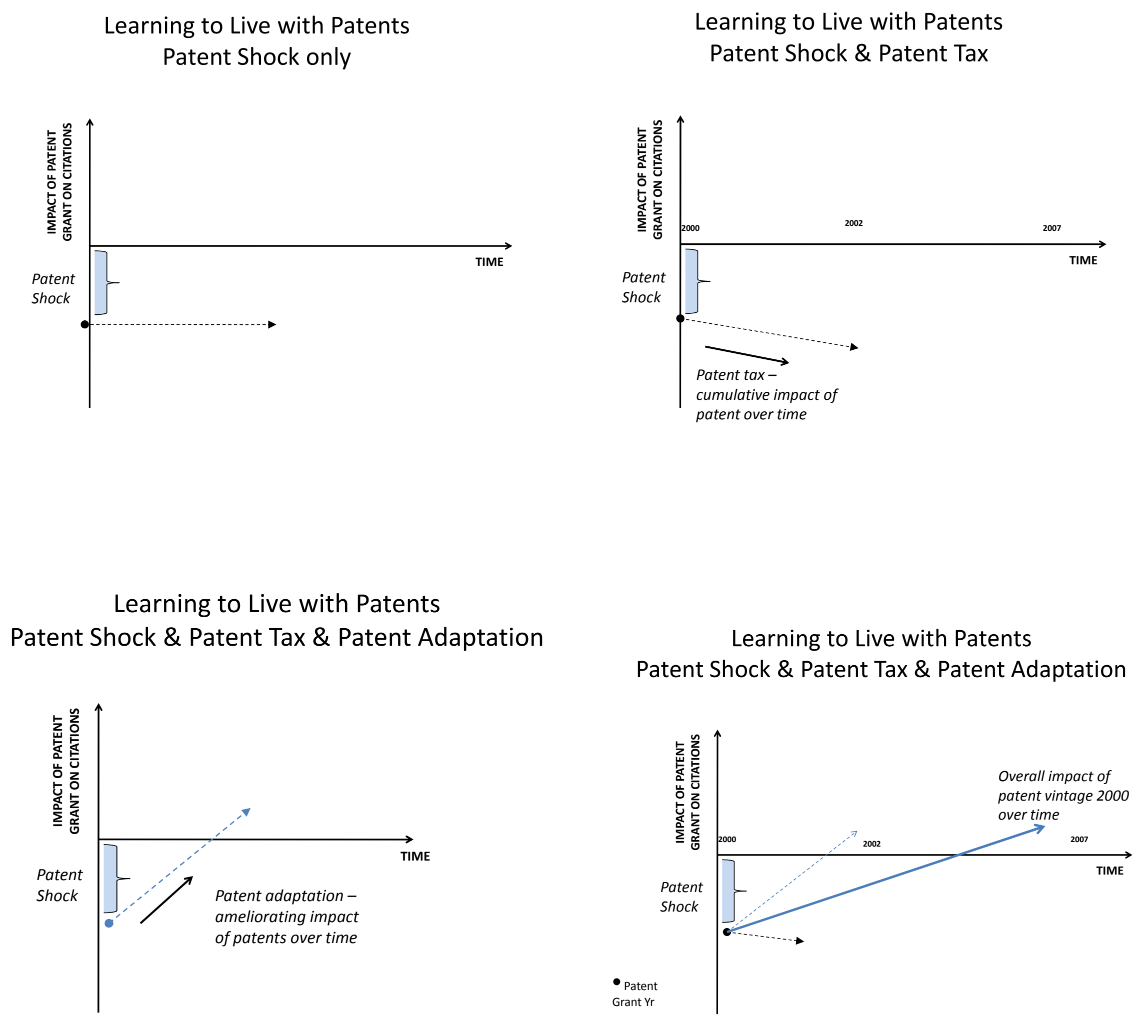
Learning to Live with Patents
Patent Shock only
IMPACT OF PATENT GRANT ON CITATIONS
TIME
Patent Shock
Learning to Live with Patents
Patent Shock & Patent Tax
IMPACT OF PATENT GRANT ON CITATIONS
2000
2002
2007
TIME
Patent Shock
Patent tax – cumulative impact of patent over time
Learning to Live with Patents
Patent Shock & Patent Tax & Patent Adaptation
IMPACT OF PATENT GRANT ON CITATIONS
TIME
Patent Shock
Patent adaptation – ameliorating impact of patents over time
Learning to Live with Patents
Patent Shock & Patent Tax & Patent Adaptation
IMPACT OF PATENT GRANT ON CITATIONS
Overall impact of patent vintage 2000 over time
2000
2002
2007
TIME
Patent Shock
Patent Grant Yr