ANNUAL REVIEWS

Annual Review of Economics

How Do Patents Affect Research Investments?

Heidi L. Williams[1,2]

[1]Department of Economics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139; email: heidiw@mit.edu

[2]National Bureau of Economic Research, Cambridge, Massachusetts 02138

Annu. Rev. Econ. 2017. 9:441–69

The *Annual Review of Economics* is online at economics.annualreviews.org

https://doi.org/10.1146/annurev-economics-110216-100959

JEL codes: O3, O31, O34

Keywords

patents, innovation, research

Abstract

Although patent systems have been widely used both historically and internationally, there is nonetheless a tremendous amount of controversy over whether patent systems, in practice, improve the alignment between private returns and social contributions. In this article, I describe three parameters—how the disclosure function affects research investments, how patent strength affects research investments in new technologies, and how patents on existing technologies affect follow-on innovation—needed to inform the question of how patents affect research investments, and review the available evidence that has attempted to empirically estimate these parameters.

1. INTRODUCTION

Technological change is widely perceived to be a key driver of improved standards of living. As an example, consider the gains in longevity that have been realized in the United States over the past half century.[1] There is a widespread consensus among health economists—albeit based on somewhat indirect evidence—that the realized improvements in longevity over this period largely reflect benefits reaped from technological change in medicine. For example, David Cutler's (2004) book *Your Money or Your Life* summarizes evidence on this point from a number of case studies suggesting that medical innovation has dramatically improved the value of treatments for at-risk infants, mental health patients, and heart attack patients over this time period.

Although many factors—such as breakthroughs in basic science—influence the pace and direction of technological change, economists have focused, at least since the seminal work of Schmookler (1966), on the important role of market incentives in shaping investments in new technologies. Because it is often difficult for inventors to capture the full social value of their inventions, it has long been recognized that, in the absence of government intervention, competitive private markets may provide less innovation than is socially desirable (Nelson 1959, Arrow 1962).[2] This concern has motivated long-standing academic and policy interest in the design of public policies to increase incentives for innovation.

The patent system is one of a suite of policy levers that has been used to attempt to bring the private returns captured by inventors closer to the social value of their inventions. By providing inventors with a temporary period of market power, patents aim to allow inventors to recoup the fixed costs of their research investments. The goal is for this temporary period of supracompetitive pricing to increase research investments relative to what would be realized in a competitive market with no government intervention. In addition, the patent system requires disclosure of the invention in exchange for the patent right. Although the disclosure requirement may decrease the private benefit of patenting in at least some markets, the motivation of the disclosure requirement is to increase the social value generated by the patent system.

Although patent systems have been quite widely used both historically and internationally, there is nonetheless a tremendous amount of controversy over whether the patent system is, in practice, improving the alignment between private returns and social contributions. Edith Penrose (1951, p. 40) famously wrote, "[I]f national patent laws did not exist, it would be difficult to make a conclusive case for introducing them; but the fact that they do exist shifts the burden of proof and it is equally difficult to make a really conclusive case for abolishing them."[3] Over a half century later, the question of whether patents are effective in spurring the development of new socially valuable technologies is still at the center of an active debate. On the one hand, in a recent *Journal of Economic Perspectives* article, Boldrin & Levine (2013, p. 3) argue that the case against patents can be summarized concisely: "[T]here is no empirical evidence that [patents] serve to

[1] Nordhaus (2003) estimates that growth in longevity in the United States since 1950 has been as valuable as growth in all other forms of consumption combined.

[2] Of course, there could also be too much innovation due to business stealing effects. Recent empirical work by Bloom et al. (2013) suggests that, on net (given current policy interventions), technological spillovers outweigh such business stealing effects.

[3] Machlup (1958, pp. 79–80) made a similar argument: "No economist, on the basis of present knowledge, could possibly state with certainty that the patent system, as it now operates, confers a net benefit or a net loss upon society. The best he can do is to state assumptions and make guesses about the extent to which reality corresponds to these assumptions . . . . If we did not have a patent system, it would be irresponsible, on the basis of our present knowledge of its economic consequences, to recommend instituting one. But since we have had a patent system for a long time, it would be irresponsible, on the basis of our present knowledge, to recommend abolishing it."

increase innovation."[4] On the other hand, in a recent *George Mason Law Review* article, Haber (2016, p. 814) argues that "[T]he weight of the evidence supports the claim of a positive causal relationship between the strength of patent rights and innovation."

In this review, I argue that estimates of three key parameters are needed to inform the question of how patents affect research investments.[5] First, how does the disclosure function of the patent system affect research investments? Second, to what extent is stronger patent protection (that is, longer patent terms or broader patent scope) effective in inducing additional research investments? For example, if we had longer patent terms or broader patent rights, would that spur the development of more pharmaceutical treatments? And third, do patents on existing technologies affect subsequent research investments? For example, do patent rights covering human genes affect the likelihood that gene-based pharmaceutical treatments and diagnostic tests based on those genes are developed? Although answers to these three questions are not sufficient for informing optimal patent policy design, estimates of these three parameters would characterize how patents affect research investments, which is one important input into optimal policy design. For example, if patent-granted exclusivity rights induce socially valuable research investments, then there is a stronger case for longer or broader patents. However, if patents hinder subsequent innovation, then there is a weaker case.

Large literatures in economics, finance, management, strategy, law, and related fields have developed over the past few decades to investigate various aspects of the patent system. To name but a few examples, researchers have investigated what role patents play in encouraging the international diffusion of new pharmaceutical drugs (Kyle & Qian 2014, Cockburn et al. 2016, Duggan et al. 2016), what role patents play in enabling technology licensing (Gans et al. 2008), how to set royalties for standard-essential patents (Lemley & Shapiro 2013), and whether the establishment of the US Court of Appeals of the Federal Circuit caused growth in patenting (Kortum & Lerner 1999). The well-known article by Hall et al. (2001) documenting a linkage between the Compustat data and granted US patents has almost 3,000 citations as of August 6, 2016 (Google Scholar, **https://scholar.google.com/scholar?cites=3323304452134060568&as_sdt=40000005&sciodt=0,22&hl=en**), one summary indicator of the volume of patent-related research using that data set alone. However, somewhat surprisingly, there has been remarkably little empirical research tackling the three questions described above, which I argue are the key empirical parameters that we need to estimate to confidently make statements about whether patents are effective in encouraging the development of new technologies.

The difficulty inherent in providing answers to the three questions outlined above is that answering these questions requires constructing counterfactuals within which we can measure the set of "potential" innovations—that is, innovations that could have been developed in the sense of being scientifically feasible but that were never brought to market. If we measure the set of potential inventions and compare it with the observed set of commercialized inventions, then we can quantify the existence of—and, ideally, the social value of—the missing technologies that could have been developed were there a better-designed patent system.

Section 2 provides a brief primer on patents. I then describe in more detail these three parameters—how the disclosure function affects research investments, how patent strength affects research investments in new technologies, and how patents on existing technologies affect

[4]Plant (1934, p. 51) famously made an earlier, similar argument: "[T]he science of economics as it stands to-day furnishes no basis of justification for this enormous experiment in the encouragement of a particular activity by enabling monopolistic price control."

[5]Of course, other parameter estimates would be needed to inform broader questions, such as the questions of optimal patent length and breadth.

follow-on innovation—and review the available evidence that has attempted to empirically estimate these parameters (Sections 3, 4, and 5). Section 6 concludes.

2. PATENTS: A BRIEF PRIMER

In the United States, inventors wishing to obtain a patent submit an application to the US Patent and Trademark Office (USPTO).[6] Patent applications are composed of two key sections. First, the specification is a written description of the invention that includes references to so-called prior art, which are citations to previously filed patent applications, previously granted patents, prior scientific publications, and other sources that are known to the inventor and relevant to the patentability of the invention. Second, the claims of the patent are a specific list of what the applicant wishes to claim intellectual property over. Filing and other fees must be paid by applicants (for more information on the current USPTO fee schedule, see **https://www.uspto.gov/learning-and-resources/fees-and-payment/uspto-fee-schedule**).

The USPTO is responsible for determining whether inventions claimed in patent applications qualify for patentability. The standard for patentability is that inventions are patent-eligible (Title 35, US Code section 101), novel (Title 35, US Code section 102), nonobvious (Title 35, US Code section 103), and useful (Title 35, US Code section 101) and that the text of the application satisfies the disclosure requirements and the requirement of claim definiteness (Title 35, US Code section 112). An initial decision on whether a given patent application meets this standard for patentability is made by a patent examiner with expertise in the scientific or technical area of the application. If the examiner issues an initial allowance of the application, the inventor can then be granted a patent. In most cases, the examiner instead issues a so-called rejection as her first decision on the application. However, in practice, patent applications cannot be rejected by the USPTO, only abandoned by applicants (Lemley & Sampat 2008). What this means is that a rejection is essentially an invitation for the applicant to submit a revised patent application that, for example, eliminates one or more claims or changes the text of some claims. The patent prosecution process—a phrase used to refer to the interaction between the USPTO and the patent applicant or her representative (such as a lawyer)—can involve several rounds of rejection and revision and, in this sense, can best be conceptualized as an iterative process between the applicant and the examiner, rather than as a one-time decision by the examiner. Applicants presumably choose between revising and resubmitting or abandoning a rejected patent application by weighing the relevant costs and benefits: If a revision that accommodates the examiner's criticisms would result in a patent that, even if granted, would be too narrow to provide much economic value to the applicant, we would expect the application to be abandoned. Thus, although much attention has focused on estimating what share of patent applications is allowed by the USPTO, attention on that metric is somewhat misplaced: Presumably, most patent applications would be allowed if the applicant were willing to accept a sufficiently narrowed set of claims. However, claims that are very narrow would provide very little economic value, and thus applicants will generally not find it in their private interest to prosecute such applications through to the allowance stage.[7]

[6] Patent protection is jurisdiction-specific: If an inventor wants to hold patent protection in the United States, France, and Japan, she must file her patent application in each jurisdiction. As described by Putnam (1996), the timing of international filings relative to an initial filing is constrained under the Paris Convention.

[7] USPTO rejection rates are also difficult to measure for other reasons, including the fact that US patent applications can spawn closely related "new" applications (such as continuations or divisionals). Carley et al. (2015) use internal USPTO data to calculate simple and "family" (including patents granted to continuations or divisionals) grant rates in the universe of new (not related to any previously filed applications) utility patent applications filed at the USPTO from 1996 to 2005. In

The institutional structure of the patent system could affect incentives for innovation in several ways, but theoretical models of optimal patent design have largely focused on two parameters: patent term length and patent breadth. The US patent term length is currently 20 years from the filing date of the patent.[8] Although some cross-country variation exists in the patent term, this 20-year term exists in most other countries as well (largely due to international harmonization efforts such as the World Trade Organization's Agreement on Trade-Related Aspects of Intellectual Property Rights). As discussed below, a large literature, starting with Nordhaus (1969), has investigated how patent term length might impact incentives for innovation.

Patent breadth (or scope) is less straightforward to define. From a theoretical perspective, the economic meaning of patent breadth is clear: How different must rival products be in order to be deemed noninfringing on a given patented product (see, e.g., Gilbert & Shapiro 1990, Klemperer 1990)? But from an empirical perspective, measuring the breadth of patent applications or granted patents is quite challenging. Most definitions are based in some way on the text of the claims of the patent, which can be conceptualized as the boundaries of the intellectual property right being sought or having been granted. Marco et al. (2016) empirically investigate two of the more commonly used metrics for patent scope: the number of words in the claims and the total number of claims.[9] As shown in their analysis, the patent prosecution process tends to narrow the scope of patent claims as measured by both claim length and claim count. As motivated by the discussion above, rejections (abandonments) of patent applications can arguably be considered one extreme on the continuum of claim narrowing that tends to occur during the patent prosecution process. To the extent that inventors (or potential inventors) correctly anticipate this claim narrowing, this is a second mechanism through which the institutional structure of the patent system could affect incentives for innovation.

Once granted, in order to keep a patent in force, the owner must pay maintenance fees. For the USPTO, these fees are currently due at 3.5 years ($1,600), 7.5 years ($3,600), and 11.5 years ($7,400). Pakes (1986) and Schankerman & Pakes (1986) pioneered the idea of using renewal fees to provide lower-bound estimates on the private value of granted patents.

3. THE DISCLOSURE FUNCTION OF THE PATENT SYSTEM

Part of the so-called grand bargain struck by the design of the patent system is that, in exchange for receiving a patent, inventors must publicly disclose their invention. This contrasts with, for example, trade secrecy, in which inventions are kept secret rather than disclosed. The so-called disclosure theory of patents stresses that a major benefit of the patent system is that these disclosures will stimulate the development of new ideas. In practice, very little evidence is available on whether or how the disclosure function of the patent system affects research investments in practice.

---

that sample, 55.8% of applications were granted to patents directly, and including patents granted to children increased the allowance rate to 71.2% (see also Lemley & Sampat 2008).

[8]This is true for utility patent applications filed on or after June 8, 1995. Technically, the 20-year term runs from the filing date of the earliest US filing or the earliest patent cooperation treaty (PCT) international filing to which priority is claimed, excluding provisional applications (Title 35, US Code section 154). One exception to the 20-year term within the United States is the so-called Hatch-Waxman term extension provided to some pharmaceutical innovations (Title 35, US Code section 156), which compensates for delays accruing due to regulatory approval by other US agencies.

[9]One important limitation of these types of measures of patent breadth is that, in practice, breadth is determined not only by the text of the patent but also by courts' interpretations of these claims and the market's expectations of how the courts will interpret these claims. Because only a subset of patents is litigated, and because market expectations are difficult to measure, it is challenging to construct measures of patent breadth that incorporate those factors.

One way in which the disclosure function of the patent system could affect research investments is if, by making patenting less attractive to private inventors, the disclosure function reduces the level of incentives provided by the patent system. Firms report in some surveys that they choose not to patent some inventions because of the disclosure requirement (Levin et al. 1987; Cohen et al. 2000, 2002); whether this in turn reduces R&D incentives is, of course, an empirical question.

Setting aside this incentive effect, the disclosure theory of patents focuses on an enablement effect: That is, conditional on an invention being disclosed, this theory asserts that new research ideas will then be stimulated. Critics of the disclosure theory of patents have argued that disclosure may generate few social benefits in practice if, for example, actual patents contain little valuable technical information, or if only inventions that would be disclosed anyway are patented (Roin 2005). These are all, again, empirical questions. Ouellette (2012) develops some empirical evidence on these points by analyzing a survey of nanotechnology researchers. She documents that 64% of survey respondents report having read patents, of which 70% used them as a source of technical information.

A natural setting in which to attempt to analyze the effects of disclosure requirements on research investments would be an analysis of a recent set of policy changes in the United States and elsewhere that have expanded the scope of disclosure by requiring disclosure not just of granted patents but also of all patent applications (with some opportunities to opt out of disclosure for certain types of patent applications that are not granted patents). I am aware of two studies that have analyzed this policy variation in the United States, although neither has aimed to quantify the effects of this policy change on research investments. First, Graham & Hegde (2015) analyze the opt-out decision among eligible applicants (namely, applicants not seeking foreign patent protection) and document that 85% of these applicants opt for disclosure. They argue that these results suggest that, contrary to common belief, there may be private benefits to disclosure. Second, Hegde & Luo (2017) analyze the effect of patent disclosure on patent licensing in the biomedical industry. They document that, after the passage of the American Inventors Protection Act (AIPA), which is the policy change described above that mandated publication of patent applications 18 months after filing, US patent applications were 20% less likely to be licensed after grant and significantly more likely to be licensed between publication and grant. They argue that these results suggest that disclosure facilitates sales and transactions in the market for ideas. I hope that future researchers will develop a way to use this AIPA-type policy variation in the United States or other countries to examine the question of how disclosure affects research investments, both through the incentive channel and through the enablement effect.[10]

4. PATENTS AND RESEARCH INVESTMENTS IN A NORDHAUS FRAMEWORK

To the best of my knowledge, Nordhaus (1969) provided the first formal theoretical framework analyzing optimal patent length (see, in particular, Nordhaus 1969, chapter 5, pp. 76–86). In Section 4.1, I summarize Budish et al.'s (2016) simplified and slightly modified version of the Nordhaus model, and in Section 4.2, I discuss survey evidence as well as empirical papers that have attempted to estimate the key empirical elasticity that emerges from this framework.

[10]Both of the studies discussed in Section 5.2 on how patents affect follow-on innovation hold disclosure constant: In the work of Galasso & Schankerman (2015), invalidated patents have already been disclosed; in the work of Sampat & Williams (2015), both accepted and rejected patent applications are disclosed during the time period of their sample.

4.1. Theory

In the original Nordhaus (1969) model, a representative firm chooses its level of R&D investment as a scalar decision variable, and R&D generates both private and social benefits by lowering that firm's production costs for a single output good. To focus attention on empirical elasticities, Budish et al. (2016) instead model a representative firm's decision of which potential R&D projects to invest in and assume that R&D investment generates private and social benefits by bringing to market inventions that would not otherwise have existed.

The set of potential inventions is of unit mass; index these potential inventions by $i \in I$. The R&D investment associated with pursing a potential invention i has a cost c_i and a probability of success p_i. If successful, the invention will generate social value for T_i years, after which it will become technologically obsolete. Patents can be conceptualized in this stylized model as providing a fixed number of years t_{patent} during which a successful invention can be sold by a monopolist and after which there can be free entry. While successful inventions are nonobsolete and are priced by a monopolist, they generate annual profits of π_i and annual social value of v_i^m. While successful inventions are nonobsolete and priced competitively, they generate annual social value of v_i^c. Budish et al. (2016) ignore discounting by assuming that R&D investment is a one-time cost incurred at time 0 and assuming that the annual quantities π_i, v_i^m, and v_i^c grow at a discount rate that is the same for the firm and society.

The firm will choose to invest R&D in potential invention i if and only if the number of years of expected monopoly life (that is, the minimum of t_{patent} and the socially useful life of the technology T_i) times the risk-adjusted profits per year (p_i times π_i) exceeds the cost of the R&D investment c_i: $p_i \cdot \min(t_{\text{patent}}, T_i) \cdot \pi_i \geq c_i$. The social value of the firm bringing invention i to market is $p_i \cdot \{\min(t_{\text{patent}}, T_i) \cdot v_i^m + [T_i - \min(t_{\text{patent}}, T_i)] \cdot v_i^c\} - c_i$.

Budish et al. (2016) use this simple framework to derive expressions for the benefits and costs of a marginal increase in the patent term t_{patent}. The benefit of a marginal increase in the patent term is that more inventions will be elicited on the margin. Letting ξ denote the quantity of inventions elicited at the margin, we can write these benefits as

$$\text{Benefits} = \underbrace{\xi}_{\substack{\text{elasticity of R\&D} \\ \text{with respect to} \\ \text{patent term}}} \times \underbrace{\mathbb{E}_{EML_i \cdot \pi_i = c_i}[ETL_i \cdot v_i^c - EML_i \cdot (v_i^c - v_i^m) - c_i]}_{\text{social value of marginal inventions}}. \quad 1.$$

The cost of a marginal increase in the patent term is that inframarginal inventions—that is, inventions that would have been developed even in the absence of the patent term extension—spend a larger share of their socially useful life under monopoly pricing, which generates additional deadweight loss. We can write these costs as

$$\text{Costs} = \int_I \underbrace{\mathbf{1}_{\{EML_i \cdot \pi_i \geq c_i\}} \mathbf{1}_{\{T_i > t_{\text{patent}}\}}}_{\text{intensive margin}} \times \underbrace{(v_i^c - v_i^m)}_{\text{deadweight loss}} \, di. \quad 2.$$

Optimal policy in the Nordhaus (1969) framework equates these benefits and costs at the margin. The ξ parameter in Equation 1 quantifies the extent to which stronger patent protection (in the form of longer patent terms) is effective in inducing additional research investments. As we discuss in Section 4.2, although it is straightforward to define this parameter theoretically, this parameter is difficult to empirically estimate because it requires drawing inferences about inventions that could have been developed—in the sense of being scientifically feasible—but that

Table 1 Patents and research investments: evidence from the Mansfield (1986) survey

Industry	Percent that would not have been developed	Percent that would not have been introduced
Pharmaceuticals	65	60
Chemicals	30	38
Petroleum	18	25
Machinery	15	17
Fabricated metal products	12	12
Primary metals	8	1
Electrical equipment	4	11
Instruments	1	1
Office equipment	0	0
Motor vehicles	0	0
Rubber	0	0
Textiles	0	0

Percent of developed or commercially introduced inventions that would not have been developed or commercially introduced if patent protection could not have been obtained, 12 industries, 1981–1983. Data taken from Mansfield (1986, table 1).

were never brought to market because the prevailing patent term was too short to incentivize the necessary private research investments.

4.2. Empirics

This section provides an overview of evidence on how patent length affects innovation, including survey-based evidence (Section 4.2.1), evidence drawn from investigating changes in patent laws (Section 4.2.2), and evidence drawn from investigating variation in clinical trial length (Section 4.2.3).

4.2.1. Survey evidence. Three influential surveys have documented evidence on the importance of patents in spurring research investments: the survey by Mansfield (1986), the so-called Yale survey (Levin et al. 1987), and the so-called Carnegie Mellon survey (Cohen et al. 2000).

Mansfield (1986) conducted a survey aimed at shedding light on the impact of patents on the rate of development of and commercialization of new inventions. The survey includes a random sample of 100 firms from 12 US industries.[11] The survey asked respondents to estimate the proportion of the firm's inventions developed in 1981–1983 that would not have been developed had the firm been unable to obtain patent protection.[12] **Table 1** summarizes the key tabulation of his results across industries. In 7 of the 12 industries surveyed, patent protection is estimated to be of limited importance in both the development and commercialization of new inventions.

[11] The sampling frame was all firms on a list published in *Business Week* in 1982 of firms in those 12 industries spending over $1 million (or 1% of sales, if sales were at least $35 million) on R&D in 1981, as well as a group of smaller firms listed in *Poor's Register*. He notes (Mansfield 1986, p. 174) that "practically all" surveyed firms filled out a detailed questionnaire, with only 4 of 100 firms not responding, and the major executives of 25 firms were also interviewed.

[12] An invention was defined as a prescription for a new or improved product or process (or one or more components or facets thereof) that is prospectively useful and that is not obvious to one skilled in the relevant art at the time the idea is generated.

Table 2 Patents and appropriability: evidence from the Yale survey

Industry	Process patents		Product patents	
	Mean	Standard error	Mean	Standard error
Pulp, paper, and paperboard	2.6	0.3	3.3	0.4
Cosmetics	2.9	0.3	4.1	0.4
Inorganic chemicals	4.6	0.4	5.2	0.3
Organic chemicals	4.1	0.3	6.1	0.2
Drugs	4.9	0.3	6.5	0.1
Plastic materials	4.6	0.3	5.4	0.3
Plastic products	3.2	0.3	4.9	0.3
Petroleum refining	4.9	0.4	4.3	0.4
Steel mill products	3.5	0.7	5.1	0.6
Pumps and pumping equipment	3.2	0.4	4.4	0.5
Motors, generators, and controls	2.7	0.3	3.5	0.5
Computers	3.3	0.4	3.4	0.4
Communications equipment	3.1	0.3	3.6	0.3
Semiconductors	3.2	0.4	4.5	0.4
Motor vehicle parts	3.7	0.4	4.5	0.4
Aircraft and parts	3.1	0.5	3.8	0.4
Measuring devices	3.6	0.3	3.9	0.3
Medical instruments	3.2	0.4	4.7	0.4
Full sample	3.5	0.06	4.3	0.07

Data taken from Levin et al. (1987, table 2).

In only two industries—chemicals and pharmaceuticals—are patents deemed to be essential to the introduction of 30% or more of inventions. Mansfield (1986, table 2) also documents that, in industries where patents are self-reported to be more important (as in **Table 1**), firms also self-report patenting a higher share of patentable inventions.

Levin et al. (1987) report tabulations from the so-called Yale survey, which aimed to shed light on the R&D appropriability conditions in different industries. The authors identified a sample of publicly traded firms with R&D expenses and sent questionnaires to employees (determined through communication with the CEO) with knowledge of the major lines of business at the firm.[13] The authors asked respondents to rate the effectiveness of alternative means of capturing and protecting the competitive advantages of new or improved processes and products on a seven-point scale: patents to prevent duplication, patents to secure royalty income, secrecy, lead time, moving quickly down the learning curve, sales efforts, and service efforts. In aggregate, patents were reported to be the least effective mechanism of appropriation for new processes, with secrecy being the most effective. Product patents were reported to be relatively more effective, although less so than several other strategies, such as lead time and sales or service efforts. **Table 2** documents tabulations of these survey responses at the industry level for the 18 industries with

[13]The sampling frame was, again, the *Business Week* annual R&D survey, which was used to identify all publicly traded firms that reported R&D expenses in excess of 1% of sales or $35 million (a total of 746 firms in 1981). Dun and Bradstreet's *Million Dollar Directory* was used to assign each firm to its major line of business. The authors received 650 responses (representing 130 lines of business) from the final sampling frame of 1,562 units.

10 or more survey responses. As in the survey of Mansfield (1986), the chemical industries stand out as reporting higher patent effectiveness for both process and product patents. Even among that group, only the pharmaceutical (drugs) industry reported patents to be the most effective mechanism of appropriation for both product and process patents. Levin et al. (1987) conjecture that patents may be more effective in the chemical industry due to relatively clearer standards of patent validity and infringement due to molecules being discrete and differentiable.

Cohen et al. (2000) report tabulations from the so-called Carnegie Mellon survey, which built closely on the Yale survey and was aimed at evaluating whether appropriability conditions changed over time with a series of reforms in the 1980s that were perceived as strengthening the patent system. The authors surveyed R&D labs and units of manufacturing firms.[14] The survey asked respondents about the effectiveness of various appropriability mechanisms (patents, secrecy, lead time, complementary sales and service, and complementary manufacturing facilities and know-how), where effectiveness was measured as the percentage of product or process innovations for which each appropriability mechanism was effective (less than 10%, 10–40%, 41–60%, 61–90%, greater than 90%). Cohen et al. (2000, tables 1 and 2) indicate that, consistent with the earlier Mansfield (1986) and Yale surveys, patents are reported to be important for product innovations in the drug and medical equipment industries, but in no industry are patents regarded as the most effective appropriability mechanisms.

Across industries, all three surveys documented evidence that, on average, firms report that patents have limited effectiveness as an appropriation mechanism relative to other levers, such as lead time and trade secrecy, but that, in a few industries—namely, chemicals and pharmaceuticals—firms report that patents are essential for spurring R&D investments. This survey evidence is useful to keep in mind, partly because several of the empirical analyses of the relationship between patents and research investments described below focus on particular industries.

4.2.2. Evidence from patent law changes. Moving beyond survey evidence to empirically estimate the relationship between patent strength and research investments requires finding variation in the patent strength that society provides to different types of potential inventions. Several researchers have relied on changes in patent law as a source of such variation.

Sakakibara & Branstetter (2001), in the first such analysis of which I am aware, provide a very careful study of a 1988 Japanese patent reform. Although this reform packaged a number of legal changes, the authors argue that perhaps the most important feature was a change from a single-claim patent system to the type of multiclaim patent system common in the United States and elsewhere, in which one patent application can include multiple claims. They note that patent experts argue that this reform broadened the scope of patent protection provided under Japanese law. Using data from the Japanese Patent Office, they document a sharp and immediate increase in the number of claims per patent application, consistent with this reform being salient to applicants. The authors then use data on the R&D investments of approximately 300 publicly traded Japanese manufacturing firms to empirically test whether broader patent rights—as induced by this reform—increase research investments. The time series evidence presented in **Figure 1** documents that, although there was an increase in R&D investments over time that predates the 1988 reform, there is little evidence for a deviation in this trend after the 1988 reform. Based on this, as well as on similar evidence from a variety of similar analyses, the authors conclude that,

[14]The sampling frame of 3,240 labs was randomly drawn from eligible labs in the Directory of American Research and Technology or belonging to firms in Standard and Poor's Compustat database stratified by a three-digit Standard Industrial Classification code. The authors received 1,478 responses and further restricted the sample to firms with at least $5 million in sales or business units with at least 20 employees for a final sample of 1,165 observations.

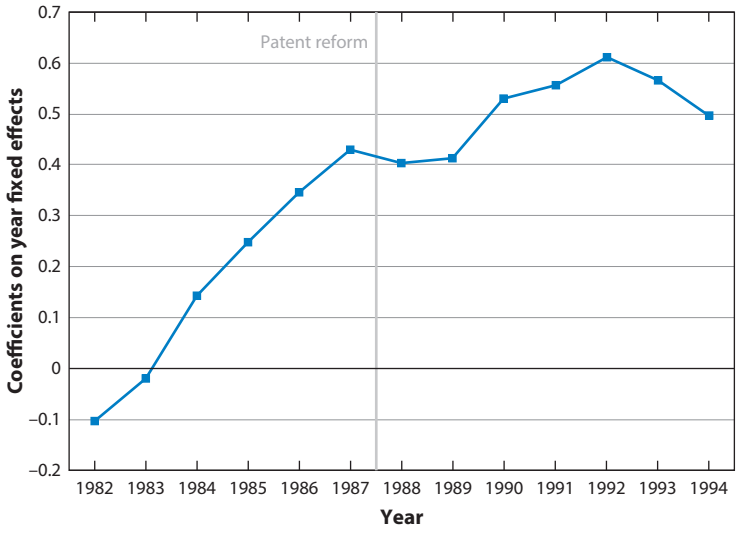

Figure 1

Patent law changes and research investments. The figure displays estimates from Sakakibara & Branstetter (2001, table 2, column 1). The dependent variable is log(real R&D spending), with data drawn from a *Kaisha Shiki Ho* R&D survey. The independent variables included are Tobin's Q, log(sales), year dummy variables (for years 1982–1994; the reference year is 1981), one-period lagged R&D, and a constant. The figure displays the coefficient estimates for the year dummy variables (year effects). Figure adapted with permission from Sakakibara & Branstetter (2001, figure 5).

although the patent reform clearly induced a behavioral response by Japanese firms (as evidenced by the change in claim behavior), there is little evidence that the broader patent protection provided by the reform resulted in increased levels of R&D investment.

Perhaps the best-known analysis of patent law changes is that of Lerner (2009). Lerner collected records of major patent policy changes in 60 countries over 150 years. Because cross-country data on R&D investments analogous to that used by Sakakibara & Branstetter (2001) is not available historically, Lerner instead constructs a novel measure of R&D investment: He chooses a country (Great Britain) that had a relatively stable patent policy over the time period of study and analyzes how foreigners' decisions to file patents in Great Britain change as a function of patent law changes in their home country. For example, if Japan lengthens its patent term, do Japanese firms increase their propensity to apply for patents in Great Britain? The advantage of this approach is that, even though Japan's patent law change may induce changes in firms' decisions over whether or not to file for patent protection in Japan on existing research investments, filings by Japanese firms in Great Britain should change only if these Japanese firms change their research investments. **Figure 2** graphically depicts the average change in patent applications around the date of patent-strengthening policy changes, net of an index that controls for changes in patent filings that are common across countries within a year. Filings by domestic entities in either their home country or in Great Britain were flat or, if anything, fell after the patent laws were strengthened. Based on this, as well as similar evidence from a variety of similar analyses, Lerner concludes that there is little evidence that stronger patent laws result in increases in R&D investments, as measured by patent filings.

As noted by Sakakibara & Branstetter (2001) and Lerner (2009), the lack of an observed relationship between stronger patent rights and research investments is a puzzling result. Lerner

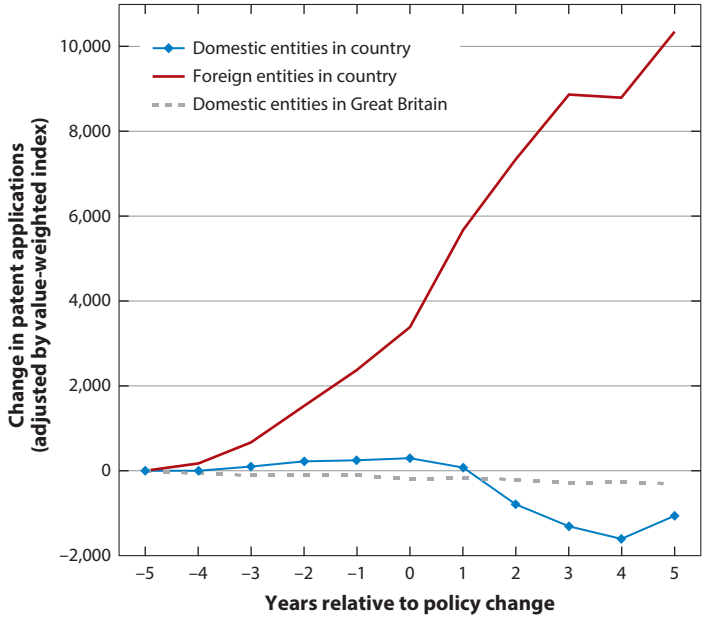

Figure 2

Patent law changes and research investments, showing domestic entities in their home country, foreign entities in their country, and domestic entities in Great Britain. Figure reproduced from Lerner (2009, figure 2).

(2009, p. 347), for example, notes, "It runs not only against our intuition as economists that incentives affect behavior, but also counter to the findings in the 'law and finance' literature that stronger property rights . . . encourage economic growth." In joint work with Eric Budish and Ben Roin, Budish et al. (2016) argue that these results are less puzzling than they may seem on first glance. The authors of these studies note that, as implemented in practice—often as part of trade negotiations—patent law changes may affect domestic R&D through other mechanisms, such as changes in foreign competition. More generally, these types of studies, by construction, investigate how R&D investments by domestic firms respond to domestic changes in patent laws, but it is not clear that this thought experiment is the right one. First, if a large economy (like that of the United States) were to lengthen its patent term, that would affect the investment incentives not only of US firms but also of non-US firms that aim to sell their goods to US consumers. Second, because technologies are developed for a global market, patent law changes in small economies likely generate relatively small sources of variation in global R&D incentives. For both reasons, we might expect estimates of the elasticity of research investments with respect to patent terms that are based on country-specific law changes to be biased toward zero.

4.2.3. Evidence from variation in clinical trial length. Given the limitations inherent in using patent law changes as variation, a natural question is whether we can isolate other alternative forms of variation in effective patent terms. Budish et al. (2015), as described below, provide one such attempt.

On paper, the US patent system provides an identical 20-year patent term to all potential inventions. However, in practice, some inventions receive fewer than 20 years of effective patent protection. One example of such uneven effective patent protection arises in the case of drug

development. Before a drug can be sold to consumers, firms are required by regulators such as the US Food and Drug Administration (FDA) to document evidence from randomized clinical trials that their drug is safe and effective. However, private firms face strong economic and legal incentives to file for patent protection on candidate drug compounds very early in the drug discovery process, prior to starting clinical trials. What that means is that, in practice, society has designed its patent system to provide shorter effective patent terms to drugs developed to treat patient groups who require longer clinical trials: With a fixed patent term, every additional day in clinical trials is one fewer day that a firm can charge supracompetitive prices on a drug that is being sold to consumers.

Because evidence of effectiveness is traditionally defined by the FDA as evidence that a drug improves survival outcomes of patients, the length of a given clinical trial is in large part determined by the time it will take to show improved survival for the treated patient population. The simple statistics of power calculations imply that clinical trials will need to be longer if the drug targets a patient population with a lower mortality rate: For example, it will take less time to document a statistically significant improvement in survival among late-stage breast cancer patients than it will among early-stage breast cancer patients. Of course, having to conduct longer clinical trials may well matter for reasons other than receiving a shorter patent term: Longer clinical trials likely cost more, both in direct financial costs and in the sense of the time value of money. Such alternative factors complicate use of this empirical context to directly estimate an elasticity of research investments with respect to the patent term, as discussed below.

Budish et al. (2015) analyze this distortion in the context of cancer drug development. The treatment of cancer tends to be organized by the organ in which the cancer originates (e.g., breast, lung, prostate) and the stage of the cancer at the time of diagnosis (local, regional, metastatic). This medical categorization of cancer types provides a natural framework within which we attempt to categorize observed and potential research investments into cancer treatments. That is, even if we never observe private firms attempting to develop drugs to treat patients with a particular type of cancer, we know from a medical perspective the full set of cancer patient groups that exist and can use this categorization to test for the existence of missing inventions due to a lack of economic incentives to make R&D investments in certain patient groups. The cancer-stage classification outlined above is natural for this purpose in part because cancer drugs are approved by the FDA for treatment of a patient group at the same cancer stage level. For example, Genentech' s drug bevacizumab was approved by the FDA in 2004 for the treatment of patients with metastatic carcinoma of the colon and rectum. Budish et al. (2015) construct measures of research investments for patients diagnosed with different cancer stages by parsing listings of which patient groups are eligible to enroll in different clinical trials, as these enrollment lists provide relatively precise information about which patient groups the sponsor is attempting to develop a new drug to treat.

For each cancer stage, we also need to observe a proxy for how long it would likely take to develop a drug to treat that patient group. Because many cancer-stage patient groups have never had a drug developed, using observed clinical trial lengths is not possible.[15] Budish et al. (2015) focus on survival time as a proxy for clinical trial length, given that survival time predicts potential clinical trial length in the statistical sense described above and that, for cancer patients, relatively

[15] In addition to this nonobservability constraint, using observed clinical trial lengths as a proxy for firms' expectations would likely be biased because the set of trials that have been conducted is—as posited by Budish et al.'s (2015) theoretical model—not a random subset of the universe of potential trials. In particular, we would expect that clinical trials to treat patient groups with long survival times may be more likely to occur in practice if, for some reason, they are allowed to conduct shorter trials via, e.g., showing evidence that the drug improves a surrogate end point rather than improving survival.

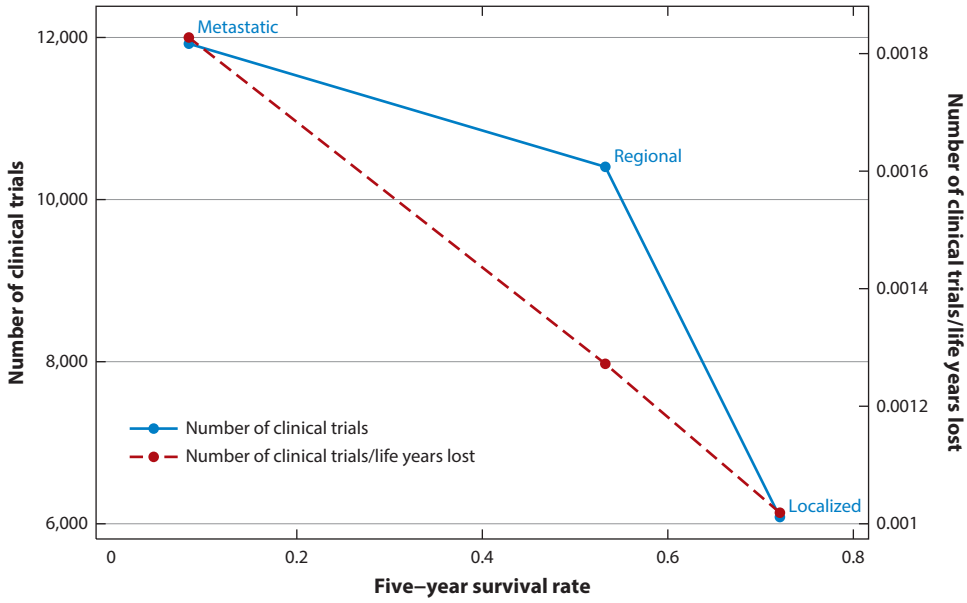

Figure 3

Research investments and clinical trial length, showing the number of clinical trials and the number of clinical trials divided by the number of life years lost by cancer stages, arranged by the average 5-year survival rate for each cancer stage. Life years lost are measured as age-, gender-, and year-specific life expectancy in the absence of cancer in the year of diagnosis, less observed survival time in years, averaged over patients diagnosed with that cancer stage in the SEER cancer registry data between 1973 and 1983 (to minimize censoring), multiplied by the number of patients diagnosed with that cancer stage. Figure adapted with permission from Budish et al. (2015, figure 1, panel *a*).

rich data are available tracking the survival time of individuals diagnosed with different cancer stages (for example, from the SEER cancer registry).

Budish et al. (2015) use these data to document that, consistent with their conjectured distortion, patient groups requiring longer clinical trials (as proxied by having higher survival rates) tend to be less frequently the target of drug development efforts. **Figure 3** provides one tabulation of this result. Metastatic cancer patients have a much lower 5-year survival rate (on the order of 10%) than regional (on the order of 50%) or localized (on the order of 70%) cancer patients, and metastatic patients are much more frequently targeted by clinical trials, being allowed to enroll in approximately 12,000 clinical trials, compared to approximately 10,000 for regional patients and approximately 6,000 for localized patients.

Although consistent with their conjectured distortion, the pattern documented in **Figure 3** is difficult to interpret for two reasons. First, a correlation between commercialization lags and research investments need not reflect a causal relationship. For example, the science behind the development of drugs to treat localized cancers could simply be more difficult than the science behind the development of drugs to treat metastatic cancers, so that, even if clinical trials evaluating treatments for localized cancers were shorter, we would not observe more of those trials. Second, even if this correlation did reflect a causal relationship, it may not be evidence of a distortion. As we clarify in a simple theoretical framework, both private firms and the social planner are more likely to pursue projects with shorter commercialization lags, which can be completed more

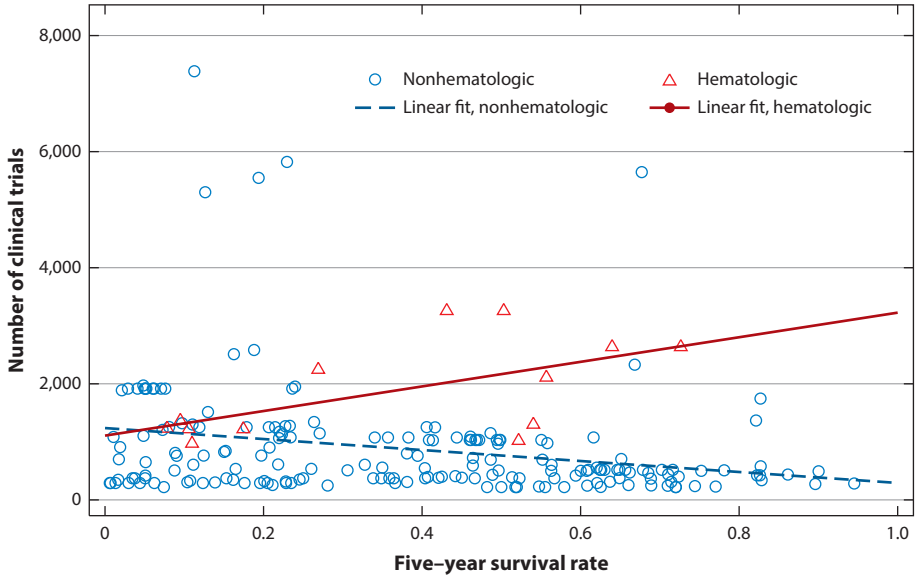

Figure 4

Research investments and surrogate end points for cancer drugs for nonhematologic and hematologic cancers. Figure adapted with permission from Budish et al. (2015, figure 4).

quickly. Budish et al. (2015) present the results of two empirical tests which together address these concerns.

First, we document evidence consistent with the idea that shortening clinical trial lengths can increase research investments. **Figure 4** provides one tabulation of this result. Our key idea is to take advantage of the fact that, for some cancers, firms are not required to show that drug treatments improve survival but instead are allowed to show that drug treatments improve an intermediate or surrogate end point that is observed more quickly. As an example, clinical trial end points for leukemia treatments are generally based on blood cell counts and related bone marrow measures, rather than on observed survival outcomes. For cancers for which such surrogate end points are consistently and predictably allowed, we expect to see less of a negative relationship between commercialization lags and research investments. **Figure 4** contrasts hematologic cancers, which are essentially always approved on the basis of surrogate end points, with other cancers and documents that, among hematologic cancers, there is not a negative relationship between 5-year survival rates (our proxy for commercialization lags) and research investments.[16] This fact is consistent with there being a causal relationship where, if firms are allowed to conduct shorter clinical trials for a given cancer, more trials will be conducted.

Second, we test for evidence of a distortion in private research investments by contrasting public and private research investments. Although, consistent with our theoretical model, both private and public research investment levels are higher for clinical trials with shorter commercialization lags, the correlation between commercialization lags and research investments is economically

[16]The linear fit line for hematologic cancers in fact displays an upward slope in **Figure 4**. Our theory gives no prediction on what the slope of this relationship should be, but rather only predicts the comparative static that the commercialization lag distortion should make the slope of this line more negative.

and statistically significantly more negative for private research investments, which is consistent with our conjectured distortion.

Taken together, these results suggest that private research investments are distorted away from long-term research projects. As mentioned above, the patent system is one factor that could be generating this distortion, but our analysis is not able to isolate the importance of patents per se as opposed to other factors. A back-of-the-envelope calculation based on the comparison between hematologic and nonhematologic cancers described above suggests a semi-elasticity of R&D investment with respect to a 1-year change in commercialization lags between 7% and 23%. That estimate would correspond to the semi-elasticity of R&D investment with respect to the patent term only under the very strong assumption that commercialization lags affect R&D investments only through changing the patent term.[17]

To summarize, analysis of data on patent law changes has provided little evidence that stronger patent rights encourage research investments, but such law changes are likely underpowered to detect such effects. Although the variation in clinical trial lengths analyzed by Budish et al. (2015) provides evidence consistent with patent length having a quantitatively important impact on research investments, the authors of that study are not able to rule out several other plausible mechanisms for the results. Identifying new sources of variation and empirical contexts in which this relationship can be more convincingly estimated is a very important avenue for future research.

5. PATENTS AND RESEARCH INVESTMENTS IN A SCOTCHMER FRAMEWORK

A relatively well-developed theoretical literature has explored the question of how patents on existing technologies affect follow-on innovation (see, for example, Scotchmer 1991). Consider, as an example, a patent on a human gene sequence, such as the gene patents granted to the firm Incyte Pharmaceuticals. Sequenced genetic data may be a useful input into follow-on scientific discoveries, such as the documentation of evidence that variations of a given gene are linked to various diseases, and into follow-on commercial applications of those scientific discoveries, such as the development of gene-based diagnostic tests or pharmaceutical treatments. The question in this case is whether patents on human genes will encourage or discourage the development of follow-on innovations such as gene-based diagnostic tests.

If Incyte had knowledge of all follow-on innovations related to their patented gene, private and social incentives would be aligned: Incyte has the private incentive to develop all socially valuable inventions. However, if ideas are scarce in the sense of Scotchmer (2004), some individuals or firms other than Incyte may have ideas for follow-on innovations that are not known to Incyte. In this case, follow-on inventions will require a license from Incyte. Green & Scotchmer (1995) clarify that, if the correct licensing agreements are signed—namely, ex ante agreements that are signed prior to follow-on inventors sinking their research investments—then all socially valuable follow-on inventions will still be developed. However, if there are impediments to ex ante licensing agreements, such as asymmetric information, then follow-on innovation may be discouraged (e.g., Bessen 2004).

[17] Interpreting this estimate economically also requires making an assumption on whether factors like credit constraints or other input constraints are relevant, in the sense that having more research on certain types of projects could imply that we have less research on other types of projects. The pharmaceutical industry, the focus of Budish et al. (2015), is arguably not an industry where we think that credit constraints are relevant, but this issue is worth keeping in mind in interpreting the results of studies on this topic.

Work dating back at least to that of Kitch (1977) has argued that patents may facilitate the coordination of investments in follow-on innovations, providing an incentive for maximizing the value of the original patent in a way that might not otherwise be realized if the original technology were instead in the public domain. This so-called prospect theory of patents provides one reason why we might expect a positive relationship between patents and follow-on innovation.

Taken together, this theoretical literature provides ambiguous predictions for how patent rights will impact follow-on innovation. In Section 5.1, I outline a simple theoretical framework aimed at clarifying the main ideas that have been discussed in the theoretical literature, with the goal of clarifying the key empirical elasticity that emerges from this literature.[18] In Section 5.2, I discuss survey evidence as well as empirical papers that have attempted to estimate this elasticity.

5.1. Theory

Assume all actors are risk neutral.[19] One firm (Incyte) holds a set of upstream technologies (genes). There exists a supply of downstream product developers, each of whom has a probability π of discovering an innovation that yields the ability to provide a downstream commercial application (gene-based diagnostic tests) at zero marginal cost. I assume that ideas are scarce and that the set of potential innovations is large enough to abstract away from cases in which two downstream product developers discover the same innovation (apart from generic entry, discussed below). The downstream market has demand $D(p)$, which I assume is differentiable and decreasing with respect to price [i.e., $D'(p) < 0$].

Ex post—that is, conditional on entry and successful product development—downstream product developers choose the optimal price to maximize profits:

$$\max_p \left[\Pi = p \cdot D(p)\right]. \qquad 3.$$

Differentiating Equation 3 with respect to price p implies that $p^m \cdot D'(p^m) + D(p^m) = 0$, which can be rewritten to solve for the optimal monopoly price p^m as

$$p^m = -\frac{D(p^m)}{D'(p^m)}. \qquad 4.$$

Let $q^m \equiv D(p^m)$ denote monopoly output. Then, profits for the monopolist are

$$\Pi^m = -\frac{D(p^m)^2}{D'(p^m)}. \qquad 5.$$

Each downstream product developer has an idiosyncratic opportunity cost of time, c. Let the number of downstream product developers with costs below any given level x be denoted by the differentiable, strictly increasing function $N(x)$.

I consider two cases: In the first, the upstream technologies are in the public domain, and in the second, the upstream technologies are granted patents.

First, consider the case in which the upstream technologies are in the public domain, in which case any firm is free to independently develop downstream products.[20] In this case, assume there is a

[18]Note that this simple framework fails to capture some important theoretical models in the prior literature, such as the contribution of Bessen & Maskin (2009), and also sets aside the disclosure function discussed in Section 3.

[19]This framework is drawn from unpublished work contained in my Harvard Ph.D. dissertation; I am very grateful to Ed Glaeser for comments while I was a graduate student. This framework is very similar in spirit to the analytical framework presented by Galasso & Schankerman (2015).

[20]I assume the upstream technologies (genes) do not compete in the same product market as the downstream technologies (diagnostic tests).

probability $1-\delta$ that the downstream patent cannot be protected because other firms can costlessly create a generic product; in this case, the price of the downstream product is pushed to marginal cost (assumed to be zero) and the downstream product developer realizes zero profits. However, with probability δ, the downstream patent can be protected, in which case the downstream firm earns monopoly profits. Firms will enter as long as their cost c is less than expected monopoly profits—adjusted by the probability of successful innovation π and the probability the downstream patent can be protected δ. Thus, firms will enter as long as c is less than $\pi \cdot \delta \cdot \Pi^m$, and the total amount of entry will equal $N(\pi \cdot \delta \cdot \Pi^m)$.

Social welfare in this case will be the product of three terms: first, the fraction π of firms who enter that will successfully discover an innovation; second, the total amount of entry $N(\pi \cdot \delta \cdot \Pi^m)$; and third, a weighted sum of the level of social welfare generated under monopoly pricing (weighted by δ) and the level of social welfare generated under marginal cost pricing (weighted by $1-\delta$). Denoting the choke-off price (that is, the lowest price at which there is no demand) by $\bar{p}$, we find that social welfare under monopoly pricing is equal to the sum of consumer surplus and producer surplus:

$$SW^m = \int_{p^m}^{\bar{p}} D(p)\mathrm{d}p + \Pi^m. \tag{6.}$$

Social welfare under marginal cost pricing is

$$SW^c = \int_0^{\bar{p}} D(p)\mathrm{d}p. \tag{7.}$$

We can thus write social welfare in the case in which the upstream technologies are in the public domain as

$$SW^{\text{public}} = \pi \cdot N(\pi \cdot \delta \cdot \Pi^m) \cdot [\delta \cdot SW^m + (1-\delta) \cdot SW^c]. \tag{8.}$$

Next, consider the second case, in which the upstream technologies are granted patents. I assume all downstream products are cumulative in the sense of infringing on the upstream firm's patents, so that any firms developing downstream products must obtain a license from the upstream firm. I assume that patents on the upstream technology ensure that there will be no generic entry in the downstream market (because generics would also infringe on the upstream firm's patents). However, because licensing negotiations are assumed to occur ex post (that is, after the downstream product developer has already sunk her R&D costs), there is a hold-up problem where the upstream firm can extract rents after a downstream product is discovered. I capture the reduced-form outcome of these ex post negotiations by assuming the downstream firm captures a share λ of the profits from the downstream product. Firms will enter as long as their cost c is less than the share of expected monopoly profits they would capture, $\pi \cdot \lambda \cdot \Pi^m$, so the total amount of entry will equal $N(\pi \cdot \lambda \cdot \Pi^m)$.

Social welfare in this case will again be the product of three terms: first, the fraction π of firms that enter that will successfully discover an innovation; second, the total amount of entry $N(\pi \cdot \lambda \cdot \Pi^m)$; and third, the level of social welfare generated under monopoly pricing. We can thus write social welfare in the case in which the upstream technologies are granted patents as

$$SW^{ip} = \pi \cdot N(\pi \cdot \lambda \cdot \Pi^m) \cdot SW^m. \tag{9.}$$

Comparing the amount of entry in the two cases is straightforward: There is more entry when the upstream technologies are in the public domain, relative to when they are granted patents, if and only if $\delta > \lambda$ (for a formal derivation of this result, see the appendix in Section 7). One can also straightforwardly derive expressions for the level of social welfare in both cases: There is a cutoff value of λ, denoted λ^*, such that social welfare from the upstream technologies being

in the public domain is the same as social welfare from the upstream technologies being granted patents. For values $\lambda < \lambda^*$, social welfare is strictly higher when the upstream technologies are in the public domain relative to when they are granted patents. For values $\lambda > \lambda^*$, social welfare is strictly higher when the upstream technologies are granted patents relative to when they are in the public domain. The value of λ^* is greater than δ.

5.2. Empirics

This section provides an overview of evidence on how patents affect follow-on research investments, including survey-based evidence (Section 5.2.1) and evidence drawn from investigating variation in patent protection induced by heterogeneity across judges and patent examiners (Section 5.2.2).

5.2.1. Survey evidence. In a series of surveys, Cohen, Walsh, and their colleagues have documented evidence of how patents affect follow-on research investments. Walsh et al. (2003) interviewed 70 individuals in various positions related to research and intellectual property in biomedical and pharmaceutical research.[21] The goal of the survey was to elicit responses on whether increasingly broad patents or increasing complexity in navigating patents on research tools has impeded follow-on research. Survey respondents confirmed (uniformly) that the patent landscape had indeed become more complex in the years preceding the survey, with more patents per innovation (including patents on research tools) and increased patenting of upstream technologies (such as targets for drug discovery). However, almost no respondents reported abandoning worthwhile projects because of issues of access to intellectual property. Rather, both university and industrial researchers reported adopting various working solutions that allowed their research to proceed, including licensing, inventing around patents, going offshore, developing and using public databases and research tools, initiating court challenges, and simply using the technology without a license (i.e., infringement). Industrial intellectual property holders who addressed the issue reported tolerating academic research infringing on their intellectual property on research tools (with the exception of patents on diagnostic tests used in clinical research), partly because academic research can increase the value of the patented technology. Industrial respondents also noted that the small prospective gains from lawsuits would not be worth the legal fees, the risk of the patent being narrowed or invalidated, and the bad publicity from suing a university.

Walsh et al. (2007) conducted a similar survey of academic biomedical researchers, investigating the impact of patents on access to the knowledge and material inputs used in follow-on research.[22] **Table 3** documents tabulations of these survey responses, showing the reasons academics give for choosing not to pursue research projects. In general, terms demanded for access to needed research inputs (i.e., unreasonable terms; 10%) and too many patents (3%) are relatively rarely cited as reasons for project abandonment. Unreasonable terms are somewhat more frequently cited as a determinant of research project choice for research related to drug discovery (21%) relative to basic research (9%). Other survey results (not reported) suggest that only 5% of survey respondents report regularly checking for patents on knowledge or tangible inputs related to

[21] Interviewees included intellectual property attorneys, scientists, and managers at 10 pharmaceutical and 15 biotechnology firms, academic researchers and technology transfer officers at six universities, and other intellectual property attorneys and personnel working in government and trade associations.

[22] The authors conducted a survey of 1,125 biomedical researchers in universities, government labs, and nonprofits; they received 414 responses. They also surveyed a second sample of 270 researchers working on signal proteins, from which they received 93 responses.

Table 3 Patents and follow-on research: survey evidence from Walsh et al. (2007)

Reason for not pursuing projects		Random sample	Research goal			Signal proteins samples		
			Drug discovery	Basic research	Other	CTLA-4	EGF	NF-κB
No funding	%High	62	86	60	58	63	54	82
Too busy	%High	60	55	60	59	53	58	48
Not feasible	%High	46	41	46	47	33	55	53
Not scientifically important	%High	40	24	41	45	40	36	50
Not interesting	%High	35	24	36	33	20	30	29
Too much competition	%High	29	21	32	21	27	29	29
Little social benefit	%High	15	21	14	15	13	5	22
Unreasonable terms	%High	10	21	9	6	7	9	19
Not help w/ promotion/job	%High	10	21	7	15	0	13	5
Too many patents	%High	3	3	2	3	0	4	0
New firm unlikely	%High	3	3	2	3	0	4	0
Little commercial potential	%High	2	3	2	3	0	4	0
Little income potential	%High	1	3	1	3	0	4	0
Not patentable	%High	1	3	1	3	0	4	0
Respondents	*N*	274	28	213	33	16	24	22

%High is the percent answering 4 or 5 on a 5-point scale ranging from "1: not at all important" to "5: very important." Data taken from Walsh et al. (2007, table 3).

their research, suggesting that—consistent with the findings of Walsh et al. (2003)—many of the surveyed academics may simply be ignoring patents.

5.2.2. Evidence from judges and patent examiners. When we move beyond survey evidence, two empirical challenges must be overcome to estimate how patents on existing technologies affect follow-on innovation. First, because we expect that the process through which some technologies are patented and others are not is nonrandom, identifying a causal effect of patents on follow-on innovation requires isolating a quasi-experiment in which patent protection is as good as randomly assigned to some technologies and not others. Second, even though there is widespread agreement that follow-on innovation is quite common—in the sense that many technologies build on existing technologies—it is challenging to construct data that can trace these links between technologies and follow-on innovations in practice. Two recent studies by Galasso & Schankerman (2015) and Sampat & Williams (2015) attempt to circumvent each of these empirical challenges.

Galasso & Schankerman (2015) focus on the empirical context of patent validity cases reviewed by the US Court of Appeals for the Federal Circuit. This setting is empirically useful because they argue that judges are assigned to patent cases through a computer program that randomly generates three-judge panels, with decisions governed by majority rule. This random assignment, together with the fact that judges vary in their propensity to invalidate patents, allows for a novel instrumental variables (IV) approach in which the decision of the court on a given case can be

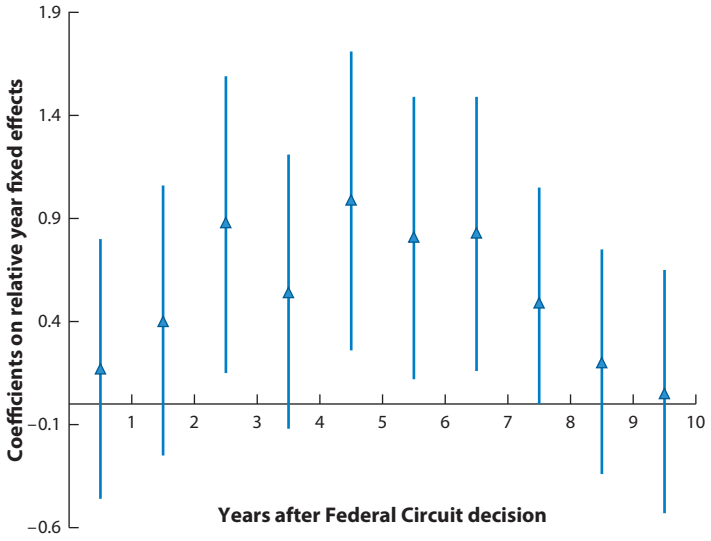

Figure 5

Patents and cumulative innovation. The lines depict the instrumental variable estimates and 90% confidence intervals of the effect of patent invalidation in each of the 10 years following a Federal Circuit decision rejecting the patent application. Figure adapted with permission from Galasso & Schankerman (2015, figure II).

instrumented with the predicted propensity to invalidate of the assigned three-judge panel. Because the judge panels are randomly assigned, this instrument should only affect follow-on innovation by affecting the validity ruling on the case.

Of the patents that reach the Federal Circuit, some are upheld and others are invalidated. In order to measure the (instrumented) effect of patent invalidation on follow-on innovation, Galasso & Schankerman measure subsequent citations to the litigated patent. They motivate this measure by noting that, because patents constitute prior art, subsequent innovators are still required to cite patents when relevant even if those patents have been declared invalid and the underlying technologies have moved into the public domain.[23] Their baseline estimates suggest that (instrumented) patent invalidation leads to an approximately 50% increase in subsequent citations to the focal patent. **Figure 5** documents the IV estimates of the effect of patent invalidation in each of the 10 years following the Federal Circuit decision, along with the associated 90% confidence intervals. This graph suggests that the effects of invalidation on subsequent citations persists for 7 years after the decision.

By construction, this empirical strategy can only be applied to patents challenged in cases that reach the Federal Circuit. Although this is certainly a selected set of valuable patents, the sample covers patents in a relatively diverse set of technology fields, allowing Galasso & Schankerman to explore the heterogeneity of their results across different types of technologies. These results suggest that patent invalidation has an impact on follow-on innovation in some fields, such as computers and communications, electronics, and medical instruments, but not in other fields, such as drugs, chemicals, and mechanical technologies. They also present results suggesting that

[23] For two of the technology fields in their sample, the authors also construct alternative non-citation-based measures of follow-on innovation. The findings they document for these alternative outcome measures are consistent with those uncovered using their citation-based outcomes, so I do not focus on discussing these measures in this review.

the impact of patent rights on follow-on innovation is driven completely by the invalidation of patents owned by large firms and by increases in the number of small innovators subsequently citing the focal patent.

Sampat & Williams (2015) focus on a case study of one particular technology: patents on human genes. They start with the sample of published USPTO patent applications claiming intellectual property over human genes. Since the early 1990s, USPTO patent applications claiming human genes have been required to list the DNA sequence of the relevant gene(s) in the text of the patent. By applying bioinformatics methods like those of Jensen & Murray (2005), it is possible to link the DNA sequences claimed in patents with a standard set of gene identifiers, which can in turn be linked to measures of follow-on innovation on those genes. Specifically, Sampat & Williams construct data measuring the scientific publications related to each gene (an indicator of scientific research investments) and measure the use of genes in pharmaceutical clinical trials and medical diagnostic tests (two indicators of commercial research investments). Because gene patents have been widely perceived as sufficiently broad that these follow-on innovations would require gene patent licenses (see, e.g., Heller & Eisenberg 1998), this data construction corresponds closely with the type of follow-on innovation described theoretically in the literature.

Using these data, we can compare the amount of follow-on innovation on human genes that are patented and not patented. However, as in the study of Galasso & Schankerman (2015), we are concerned that gene patents are likely not randomly assigned across genes. Indeed, a naive comparison of follow-on innovation across patented and nonpatented genes in our data suggests that we observe higher levels of follow-on innovation on patented genes, but that is true both before and after these genes are patented, highlighting a selection bias in the way in which genes are patented.

To address this selection bias concern, we construct two quasi-experimental approaches. First, we present a simple comparison of follow-on innovation across genes claimed in accepted and rejected patent applications. As described in Section 2, patents are not actually rejected by the USPTO, but we can compare genes claimed in patent applications that were granted patents with genes claimed in patent applications but never granted a patent—which is useful if genes that were included in unsuccessful patent applications can serve as a valid comparison group for genes that were granted patents. Second, we develop a novel IV approach in which the existence of a patent on a given gene is instrumented with the leniency of the assigned patent examiner.[24] Previous qualitative research has suggested that the assignment of patent applications to patent examiners at the USPTO is plausibly random conditional on some covariates, such as technology type and application year (Cockburn et al. 2003, Lemley & Sampat 2010), and we confirm empirically that this assumption looks plausible in our sample of gene patent applications, suggesting that examiner leniency can serve as a valid instrument.

Both of our quasi-experimental approaches suggest that gene patents have not had quantitatively important effects on follow-on innovation. **Figure 6** graphically documents the main results from the first quasi-experimental approach. Among human genes included in at least one patent application, approximately 30% are included in a granted patent. **Figure 6*a*** documents the time path of when genes in the patented group receive their (first) patent grant over time: Approximately half have received a patent grant by 2005, giving some sense of when we might expect any differences in follow-on innovation between the two groups to be observed. **Figure 6*b*,*c*** documents trends in follow-on innovation by year. Comfortingly, the number of scientific publications

[24]Note that this IV strategy has subsequently been reapplied in a number of other studies, including those of Gaule (2015) and Farre-Mensa et al. (2016).

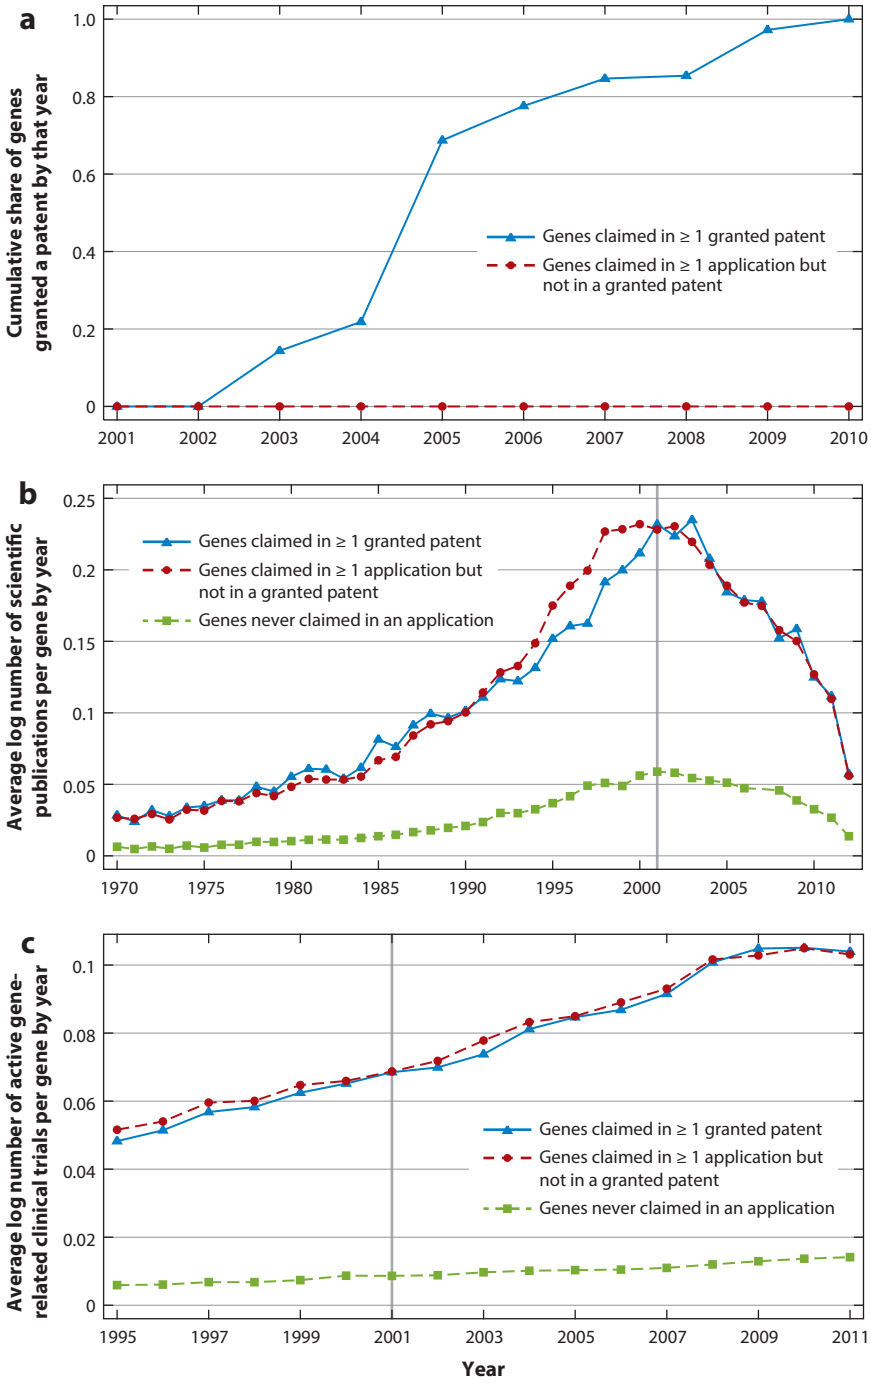

Figure 6

Patents and cumulative innovation in gene research. (*a*) Share of genes receiving a patent grant. (*b*) Gene-level scientific publications. (*c*) Gene-level clinical trials. The vertical line at 2001 in panels *b* and *c* marks the year in which patent applications in our sample started to be filed; all years prior to 2001 can be considered a preperiod and used to estimate the selection of genes into patenting based on prefiling measures of scientific research (publications) and commercialization (clinical trials). Figure adapted with permission from Sampat & Williams (2015, figure 2).

(**Figure 6*b***) and the number of clinical trials (**Figure 6*c***) look quite similar across genes claimed in accepted and rejected patent applications prior to 2001, the year when our data on gene patent applications start.[25] But the two data series also look quite similar post-2001, providing little evidence that patent grants have had any effect on either follow-on innovation outcome.

The estimates from our first approach—comparing genes claimed in accepted and rejected patent applications—are reasonably precise and are able to reject the types of declines in follow-on innovation that have been observed in quasi-experimental studies of nonpatent forms of intellectual property (Williams 2013, Murray et al. 2016). Although the estimates from the second instrumental variables approach are less precise, the fact that both approaches generate similar conclusions is comforting. Most directly, these estimates speak against the arguments made by the US Supreme Court in the case of *Association for Molecular Pathology v. Myriad Genetics*, in which the Court unanimously ruled to invalidate patents on human genes, arguing that such patents "would 'tie up'. . . [genes] and . . . inhibit future innovation" (US Supreme Court 2012, p. 11). More broadly, our analysis is complementary to that of Galasso & Schankerman (2015) in providing evidence on when and how patents will affect follow-on innovation.[26]

6. CONCLUSION

The patent system is a widely used policy lever attempting to better align the private returns to developing new technologies with the social value of those inventions. The past few decades have seen the development of large academic literatures in a variety of fields—including economics, law, and strategy—investigating various aspects of the patent system. However, surprisingly little research has focused on empirically estimating the key parameters needed to evaluate the social costs and social benefits of the patent system. A half century ago, Penrose (1951) and Machlup (1958) argued that insufficient empirical evidence existed to make a conclusive case either for or against patents. Today, I would argue that, given the limitations of the existing literature, we still have essentially no credible empirical evidence on the seemingly simple question of whether stronger patent rights—either longer patent terms or broader patent rights—encourage research investments into developing new technologies. Although researchers have recently begun to make progress on the more limited question of how patents on existing technologies affect follow-on innovation (Galasso & Schankerman 2015, Sampat & Williams 2015), evidence on the overall effects of patents on research investments is needed as one input into optimal patent policy design.

Some of the evidence discussed in this article suggests that there may be important heterogeneity across industries or across technologies in how patents affect research investments. Although legal constraints limit the ability to tailor current patent law away from the current one-size-fits-all reward approach, researchers such as Roin (2014) have argued that such tailoring may be feasible, at least in principle, given observable economic determinants of optimal patent strength. Development of additional empirical evidence informing the questions of what optimal patent terms are and how they vary across industries is an important direction for future work, which could quantify the potential benefits of such tailored patent awards.

[25] In contrast, the level of follow-on innovation on genes never included in a patent application is notably lower throughout the sample. Although these genes are not part of our empirical analysis, that pattern is consistent with strong selection of genes into patent filings.

[26] Consistent with our results, Galasso & Schankerman (2015) find no evidence that patent invalidations increase patent citations in the National Bureau of Economic Research technology field that includes most of our gene patents (drugs and medical).

7. APPENDIX

7.1. Proposition 1

Social welfare is higher when the upstream technologies are in the public domain, relative to when they are granted patents, whenever SW^{public} is larger than SW^{IP} (where IP denotes intellectual property)—that is,

$$\pi \cdot N(\pi \cdot \delta \cdot \Pi^m) \cdot [\delta \cdot SW^m + (1-\delta) \cdot SW^c] > \pi \cdot N(\pi \cdot \lambda \cdot \Pi^m) \cdot SW^m. \quad 10.$$

There exists a value of λ, denoted λ^*, such that social welfare from the upstream technologies being in the public domain is the same as social welfare from the upstream technologies being granted patents. For values $\lambda < \lambda^*$, social welfare is strictly higher when the upstream technologies are in the public domain relative to when they are granted patents. For values $\lambda > \lambda^*$, social welfare is strictly higher when the upstream technologies are granted patents relative to when they are in the public domain. The value of λ^* is greater than δ and is increasing in δ if and only if the elasticity ε of the supply of innovators with respect to the level of expected returns [that is, $\frac{x \cdot N'(x)}{N(x)}$] is such that

$$\varepsilon > \frac{\frac{SW^c}{SW^m} - 1}{1 + \frac{1-\delta}{\delta} \cdot \frac{SW^c}{SW^m}}. \quad 11.$$

7.2. Proof of Proposition 1

The function $N(\cdot)$ is strictly increasing, which implies that social welfare when the upstream technologies have been granted patents is strictly increasing in λ.

At $\lambda = 0$, social welfare with patents is $\pi \cdot N(0) \cdot SW^m$. Comparing this level of social welfare with SW^{public}, we see that, because $N(0)$ is strictly less than $N(\pi\delta\Pi^m)$ and SW^m is weakly less than $\delta \cdot SW^m + (1-\delta) \cdot SW^c$, at $\lambda = 0$ granting patents on the upstream technologies is strictly worse than the upstream technologies being in the public domain.

At $\lambda = 1$, social welfare with patents is $\pi \cdot N(\pi \cdot \Pi^m) \cdot SW^m$. Comparing this level of social welfare with SW^{public}, we see that, in this case, social welfare with patents is greater than SW^{public} if and only if

$$\frac{N(\pi\Pi^m)}{N(\pi\delta\Pi^m)} > \delta + (1-\delta) \cdot \frac{SW^c}{SW^m}. \quad 12.$$

If Equation 12 holds at $\lambda = 1$, then granting patents on the upstream technologies is preferable to the upstream technologies being in the public domain.

At $\lambda = 1$, Equation 12 implies that granting patents on the upstream technologies is strictly worse than the upstream technologies being in the public domain as long as $\frac{N(\pi\Pi^m)}{N(\pi\delta\Pi^m)} < \delta + (1-\delta) \cdot \frac{SW^c}{SW^m}$. Because social welfare when the upstream technologies have been granted IP is strictly increasing in λ, this implies that, for all values of λ, granting patents on the upstream technologies is strictly worse than the upstream technologies being in the public domain as long as $\frac{N(\pi\Pi^m)}{N(\pi\delta\Pi^m)} < \delta + (1-\delta) \cdot \frac{SW^c}{SW^m}$.

Because social welfare when the upstream technologies have been granted patents is strictly lower than when the upstream technologies are in the public domain at $\lambda = 0$, strictly higher at $\lambda = 1$ when Equation 12 holds, and monotonically increasing in λ, the continuity of $N(\cdot)$ implies that, when Equation 12 holds, there exists a value of λ at which social welfare is the same when the upstream technologies have patents as when they are in the public domain.

The value of λ at which social welfare is equal across these two cases, denoted λ^*, is implicitly defined by

$$\pi \cdot N(\pi \cdot \delta \cdot \Pi^m) \cdot [\delta \cdot SW^m + (1-\delta) \cdot SW^c] = \pi \cdot N(\pi \cdot \lambda^* \cdot \Pi^m) \cdot SW^m, \quad 13.$$

which we can rewrite as

$$N(\pi \cdot \lambda^* \cdot \Pi^m) = N(\pi \cdot \delta \cdot \Pi^m) \left[\delta + (1-\delta) \cdot \frac{SW^c}{SW^m}\right]. \quad 14.$$

Note that the bracketed term, $\delta + (1-\delta) \cdot \frac{SW^c}{SW^m}$, is greater than 1, implying that $N(\pi \cdot \delta \cdot \Pi^m) < N(\pi \cdot \lambda^* \cdot \Pi^m)$, which in turn implies that $\delta < \lambda^*$.

Differentiating both sides of Equation 14 with respect to δ gives

$$N'(\pi \cdot \lambda^* \cdot \Pi^m) \cdot \frac{d\lambda^*}{d\delta} \cdot \pi \cdot \Pi^m = N'(\pi \cdot \delta \cdot \Pi^m)) \cdot \pi \cdot \Pi^m \cdot \left[\delta + (1-\delta) \cdot \frac{SW^c}{SW^m}\right] + N(\pi \cdot \delta \cdot \Pi^m) \cdot \left[1 - \frac{SW^c}{SW^m}\right], \quad 15.$$

which can be rewritten as

$$\frac{d\lambda^*}{d\delta} = \frac{N'(\pi \cdot \delta \cdot \Pi^m) \cdot \pi \cdot \Pi^m \cdot \left[\delta + (1-\delta) \cdot \frac{SW^c}{SW^m}\right] + N(\pi \cdot \delta \cdot \Pi^m) \cdot \left[1 - \frac{SW^c}{SW^m}\right]}{N'(\pi \cdot \lambda^* \cdot \Pi^m) \cdot \pi \cdot \Pi^m}. \quad 16.$$

The denominator on the right-hand side of Equation 15 is always positive [because $N(\cdot)$ is strictly increasing]. Thus, $\frac{d\lambda^*}{d\delta}$ is positive if and only if the numerator on the right-hand side of Equation 15 is positive. That is, $\frac{d\lambda^*}{d\delta}$ is positive if and only if

$$N'(\pi \cdot \delta \cdot \Pi^m) \cdot \pi \cdot \Pi^m \cdot \left[\delta + (1-\delta) \cdot \frac{SW^c}{SW^m}\right] + N(\pi \cdot \delta \cdot \Pi^m) \cdot \left[1 - \frac{SW^c}{SW^m}\right] > 0. \quad 17.$$

Define ε to be the following, representing the elasticity of the supply of innovators with respect to the level of expected returns:

$$\varepsilon = \frac{N'(\pi \cdot \delta \cdot \Pi^m) \cdot \pi \cdot \delta \cdot \Pi^m}{N(\pi \cdot \delta \cdot \Pi^m)}. \quad 18.$$

We can then rewrite Equation 17 in terms of ε as

$$\varepsilon + \varepsilon \cdot \left[\frac{1-\delta}{\delta} \cdot \frac{SW^c}{SW^m}\right] + \left[1 - \frac{SW^c}{SW^m}\right] > 0. \quad 19.$$

Equation 17 holds if and only if

$$\varepsilon > \frac{\frac{SW^c}{SW^m} - 1}{1 + \frac{1-\delta}{\delta} \cdot \frac{SW^c}{SW^m}}. \quad 20.$$

□

DISCLOSURE STATEMENT

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

I am extremely grateful to Tim Bresnahan for very thoughtful and constructive comments and to Ivan Badinski and Myles Wagner for excellent research assistance. Research reported in this

publication was supported by the National Institute on Aging and the National Institutes of Health (NIH) Common Fund, Office of the NIH Director, through grant U01-AG046708 to the National Bureau of Economic Research (NBER); the content is solely the responsibility of the author and does not necessarily represent the official views of the NIH or NBER. Financial support from National Science Foundation Grant 1151497 is also gratefully acknowledged.

LITERATURE CITED

Arrow K. 1962. Economic welfare and the allocation of resources for invention. In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, ed. R Nelson, pp. 609–26. Princeton, NJ: Princeton Univ. Press

Bessen J. 2004. Holdup and licensing of cumulative innovations with private information. *Econ. Lett.* 82(3):321–26

Bessen J, Maskin E. 2009. Sequential innovation, patents, and imitation. *RAND J. Econ.* 40(4):611–35

Bloom N, Schankerman M, Van Reenen J. 2013. Identifying technology spillovers and product market rivalry. *Econometrica* 81(4):1347–93

Boldrin M, Levine DK. 2013. The case against patents. *J. Econ. Perspect.* 27(1):3–22

Budish E, Roin B, Williams H. 2015. Do firms underinvest in long-term research? Evidence from cancer clinical trials. *Am. Econ. Rev.* 105(7):2044–85

Budish E, Roin B, Williams H. 2016. Patents and research investments: assessing the empirical evidence. *Am. Econ. Rev.* 106(5):183–87

Carley M, Hegde D, Marco A. 2015. What is the probability of receiving a US patent? *Yale J. Law Technol.* 17(1):203–33

Cockburn I, Kortum S, Stern S. 2003. Are all patent examiners equal? Examiners, patent characteristics, and litigation outcomes. In *Patents in the Knowledge-Based Economy*, ed. W Cohen, S Merrill, pp. 19–53. Washington, DC: Natl. Acad. Press

Cockburn IM, Lanjouw JO, Schankerman M. 2016. Patents and the global diffusion of new drugs. *Am. Econ. Rev.* 106(1):136–64

Cohen WM, Goto A, Nagata A, Nelson RR, Walsh JP. 2002. R&D spillovers, patents and the incentives to innovate in Japan and the United States. *Res. Policy* 31(8):1349–67

Cohen WM, Nelson RR, Walsh JP. 2000. *Protecting their intellectual assets: appropriability conditions and why US manufacturing firms patent (or not)*. NBER Work. Pap. 7552

Cutler D. 2004. *Your Money or Your Life: Strong Medicine for America's Health Care System*. Oxford, UK: Oxford Univ. Press

Duggan M, Garthwaite C, Goyal A. 2016. The market impacts of pharmaceutical product patents in developing countries: evidence from India. *Am. Econ. Rev.* 106(1):99–135

Farre-Mensa J, Ljungqvist A, Hegde D. 2016. *The bright side of patents*. Unpublished manuscript, Stern School Bus., New York Univ., New York

Galasso A, Schankerman M. 2015. Patents and cumulative innovation: causal evidence from the courts. *Q. J. Econ.* 130(1):317–69

Gans JS, Hsu DH, Stern S. 2008. The impact of uncertain intellectual property rights on the market for ideas: evidence from patent grant delays. *Manag. Sci.* 54(5):982–97

Gaule P. 2015. *Patents and the success of venture-capital backed startups: using examiner assignment to estimate causal effects*. Unpublished manuscript, Cent. Econ. Res. Grad. Educ. Econ. Inst., Prague

Gilbert R, Shapiro C. 1990. Optimal patent length and breadth. *RAND J. Econ.* 21:106–12

Graham S, Hegde D. 2015. Disclosing patents' secrets. *Science* 347(6219):236–37

Green J, Scotchmer S. 1995. On the division of profit in sequential innovation. *RAND J. Econ.* 26(1):20–33

Haber S. 2016. Patents and the wealth of nations. *George Mason Law Rev.* 23(4):811–35

Hall BH, Jaffe AB, Trajtenberg M. 2001. *The NBER patent citation data file: lessons, insights and methodological tools*. NBER Work. Pap. 8498

Hegde D, Luo H. 2017. Patent publication and the market for ideas. *Manag. Sci.* In press. **http://dx.doi.org/10.1287/mnsc.2016.2622**

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

Jensen K, Murray F. 2005. Intellectual property landscape of the human genome. *Science* 310(5746):239–40

Kitch E. 1977. The nature and function of the patent system. *J. Law Econ.* 20:265–90

Klemperer P. 1990. How broad should the scope of patent protection be? *RAND J. Econ.* 21:113–30

Kortum S, Lerner J. 1999. What is behind the recent surge in patenting? *Res. Policy* 28(1):1–22

Kyle M, Qian Y. 2014. *Intellectual property rights and access to innovation: evidence from TRIPS*. NBER Work. Pap. 20799

Lemley M, Sampat B. 2008. Is the patent office a rubber stamp? *Emory Law J.* 58:181–203

Lemley M, Sampat B. 2010. Examining patent examination. *Stanf. Technol. Law Rev.* 2010:2

Lemley M, Shapiro C. 2013. A simple approach to setting reasonable royalties for standard-essential patents. *Berkeley Technol. Law J.* 28:1135–66

Lerner J. 2009. The empirical impact of intellectual property rights on innovation: puzzles and clues. *Am. Econ. Rev.* 99(2):343–48

Levin RC, Klevorick AK, Nelson RR, Winter SG. 1987. Appropriating the returns from industrial research and development. *Brookings Pap. Econ. Act.* 1987(3):783–831

Machlup F. 1958. *An economic review of the patent system*. Stud. 15, Subcomm. Pat. Trademarks Copyr., Comm. Judic., US Senate, 85 Congr., 2nd Sess.

Mansfield E. 1986. Patents and innovation: an empirical study. *Manag. Sci.* 32(2):173–81

Marco AC, Sarnoff JD, DeGrazia CA. 2016. *Patent claims and patent scope*. Tech. Rep., Work. Group Intellect. Prop. Innov. Prosper., Hoover Inst., Stanford, CA

Murray F, Aghion P, Dewatripont M, Kolev J, Stern S. 2016. Of mice and academics: examining the effect of openness on innovation. *Am. Econ. J. Econ. Policy* 8(1):212–52

Nelson R. 1959. The simple economics of basic scientific research. *J. Polit. Econ.* 67(3):297–306

Nordhaus W. 1969. *Invention, Growth, and Welfare: A Theoretical Treatment of Technological Change*. Cambridge, MA: MIT Press

Nordhaus W. 2003. The health of nations: the contribution of improved health to living standards. In *Measuring the Gains from Medical Research: An Economic Approach*, ed. KM Murphy, RH Topel, pp. 9–40. Chicago: Univ. Chicago Press

Ouellette LL. 2012. Do patents disclose useful information? *Harv. J. Law Technol.* 532(2):531–93

Pakes A. 1986. Patents as options: some estimates of the value of holding European patent stocks. *Econometrica* 54(4):755–84

Penrose ET. 1951. *The Economics of the International Patent System.* Baltimore, MD: Johns Hopkins Press

Plant A. 1934. The economic theory concerning patents for inventions *Economica* 1(1):30–51

Putnam J. 1996. *The value of international patent protection*. PhD Diss., Yale Univ., New Haven, CT

Roin B. 2005. The disclosure function of the patent system (or lack thereof). *Harv. Law Rev.* 118(6):2007–28

Roin B. 2014. The case for tailoring patent awards based on time-to-market. *UCLA Law Rev.* 61(672):679–80

Sakakibara M, Branstetter L. 2001. Do stronger patents induce more innovation? Evidence from the 1998 Japanese patent law reforms. *RAND J. Econ.* 32(1):77–100

Sampat B, Williams H. 2015. *How do patents affect follow-on innovation? Evidence from the human genome*. Unpublished manuscript, Mass. Inst. Technol., Cambridge

Schankerman M, Pakes A. 1986. Estimates of the value of patent rights in European countries during the post-1950 period. *Econ. J.* 96(384):1052–76

Schmookler J. 1966. *Invention and Economic Growth*. Cambridge, MA: Harv. Univ. Press

Scotchmer S. 1991. Standing on the shoulders of giants: cumulative research and the patent law. *J. Econ. Perspect.* 5(1):29–41

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

US Supreme Court. 2012. *Syllabus: Association for Molecular Pathology et al. v. Myriad Genetics, Inc., et al.* Rep., US Supreme Court, Washington, DC. **https://www.supremecourt.gov/opinions/12pdf/12-398_1b7d.pdf**

Walsh J, Arora A, Cohen W. 2003. Working through the patent problem. *Science* 299(5609):1021

Walsh J, Cohen W, Cho C. 2007. Where excludability matters: material versus intellectual property in academic biomedical research. *Res. Policy* 36(8):1184–203

Williams H. 2013. Intellectual property rights and innovation: evidence from the human genome. *J. Polit. Econ.* 121(1):1–27