

Disclosure as a Strategy in the Patent Race*

Scott Baker
University of North Carolina
School of Law
sbaker@email.unc.edu

Claudio Mezzetti
University of North Carolina
Department of Economics
mezzetti@email.unc.edu

* We thank two anonymous referees and two editors for their insightful comments. This paper owes a lot to Doug Lichtman. Doug was a co-author of an earlier version, but was unable to devote time to this revision. For helpful comments and conversations, we thank Paul Edelman, Rebecca Eisenberg and workshop participants at the American Law & Economics Association Annual Meeting (May 2001).

Abstract

Research firms disclose a surprisingly large amount of information to the patent office through “targeted” disclosures; that is, disclosures intended to make the patent office aware that potentially patentable information is already in the public domain. Conventional wisdom holds that these disclosures are made for defensive purposes; the disclosing firm does not itself plan to pursue patents related to the disclosed information, so the firm discloses as a way of creating prior art that might stop rivals from patenting. But firms have an incentive to disclose even if they themselves intend to pursue patent protection. The reason is that, by making it more difficult to patent, disclosure in essence extends the patent race. If an invention of a certain quality would have been sufficient to qualify for patent protection before the disclosure, after the disclosure any invention must be that much better before it will represent a sufficient advance over the now-expanded prior art. This paper models disclosure strategies and explains how research firms can use disclosure to gain a competitive edge. It also presents empirical evidence that IBM, at least, engages in disclosures consistent with the model developed here.

I. INTRODUCTION

Research firms occasionally allow their employees to make presentations at conferences, contribute articles to peer-reviewed journals, or both. Such activities unavoidably reveal some otherwise-proprietary research information, but they are nevertheless easy to understand. The disclosures involved are typically small-scale; they reward employees for their achievements and also lead to favorable publicity for the firm, perhaps making it easier for the firm to raise capital to fund future projects. Harder to understand are large-scale disclosures targeted toward the patent office. We have three examples of “targeted” disclosures in mind. First, there are the disclosures in the publications produced by companies such as IP.com and Research Disclosures, Inc. For a fee, these companies will place non-patented technology in the public domain and make it readily available to the United States Patent and Trademark Office (PTO) and other patent offices worldwide. IP.com and Research Disclosures, Inc. provide disclosure services for a large and diverse group of companies.¹ To show that these publications reveal valuable information, note that the publication of Research Disclosure, Inc. is often referenced in later-granted United States patents.²

Second, IBM published a technical journal from 1958 until 1998.³ Like the publication of Research Disclosure, Inc., the IBM journal revealed valuable information - its articles have been cited over 48,000 times in United States patents.⁴ IBM distributed its journal to patent offices at no fee.

Third, from February 1, 1976 through April 7, 1997, the Xerox Corporation published a bimonthly technical journal in which firm employees described in detail their

¹ IP.com provides disclosure services for, among others, Abbott Laboratories, Armstrong World Industries, Avery Dennison Carrier, ChevronTexaco, Eastman Kodak, General Electric (GE), IBM, International Fuel Cells, Kimberly-Clark, Motorola Corporation, NCR Financial Solutions Group Ltd., Nippon Soda Co., Ltd., ON Semiconductor, Otis Elevator, Plantronics, Inc., Polaroid Corporation, PPG Industries, Inc., Pratt & Whitney, Pratt & Whitney – Canada, Square D Company, United Technologies Research Center. See <http://ip.com>, visited September 29, 2004. Research Disclosure, Inc. list of clients is even more extensive; it claims that “[o]ver the years an estimated 90% of the world’s leading companies and corporations have published disclosures in RD.” See <http://www.researchdisclosure.com/overview/company.html>, visited September 29, 2004.

² Our search revealed that this publication has been cited 5,894 times in granted patents. We came by this information by doing the appropriate search in the Delphion database. See <http://delphion.com>. IP.com is a relatively recent player in the disclosure market. Hence, its publication has not yet been cited in later-filed patent applications.

³ IBM’s journal was called the IBM Technical Disclosure Bulletin. Although IBM no longer publishes this journal, IBM does continue to publish a variety of periodic research reports.

⁴ See http://www.delphion.com/search-prior_art, visited September 29, 2004. According to IBM spokesmen, the Bulletin is the third most frequently cited non-patent reference in United States patent applications. See Pryor Garnett, The Case for Defensive Disclosure (remarks given at the Software Patent Institute organizational conference, July 22-23, 1991) (excerpted at <http://www.spi.org/defdis.html>, visited September 24, 2004). Our own search at the Patent & Trademark Office website revealed over 7,000 separate patents citing the Bulletin in just the last five years.

on-going research.⁵ The journal was then distributed to libraries and patent offices worldwide at no fee.

At the outset, it is important to focus on the type of disclosures considered. Although interesting, we do not set out to explain why firms disclose potentially patentable information to other competitors, the public generally, or the capital markets.⁶ Instead, we focus our attention on why firms reveal information to the PTO itself.⁷ In addition to the publicity and reward rationales, there are a host of other reasons why a firm might place potentially patentable information in the public domain. For example, a firm might reveal information to actual or potential competitors in order to discourage competition in the disclosing firm's line of research -- by, for instance, showing that the disclosing firm is very far down the research path or that its research is of a high quality. This explanation resonates; it cannot explain, however, why firms try so hard to make the PTO aware of their disclosures. After all, if discouraging competition were the motivation, then a firm would target its disclosures at competitors and potential competitors -- not the PTO.

Alternatively, a firm might disclose to a broad technological audience, hoping that someone in that audience might build on the disclosure and break down a technological barrier that the disclosing firm could not solve. This justification also makes sense, and it explains, perhaps, why firms sometimes engage in large public disclosures. Again, however, this rationale fails to explain disclosures directed at the PTO.⁸

⁵ The publication was called the Xerox Disclosure Journal. A spokeswoman for Xerox provided the dates of publication; further information about this publication is available online at <http://gort.ucsd.edu/newjour/x/msg00749.html>, visited September 29, 2004. Xerox now makes its disclosures through IP.com. According to the spokeswoman, Xerox made 50 publications in 2003 and plans to make 50 publications in 2004. Interestingly, Xerox does not list its name on the disclosures through IP.com; they are anonymous. Xerox spokeswoman, telephone conversation with author, September 29, 2004.

⁶ On using disclosure to attract capital, see Sudipto Bhattacharya and Jay R. Ritter, *Innovation and Communication: Signalling with Partial Disclosure*, 50 *REV. OF ECON. STUD.* 331 (1983).

⁷ On its website, Research Disclosure, Inc. advertises its service as facilitating disclosures to patent examiners. See <http://www.researchdisclosure.com/overview/index.html>, visited September 29, 2004 ("You just send us the invention details you wish to disclose, and we publish them in both our paper and electronic RD formats. We charge you a modest one off fee and that is it. You can then rest assured that the *patent examiners* will see it for decades to come while you concentrate on developing your invention.") (emphasis added).

⁸ There are a number of other possible reasons why firms might publish valuable information. A firm, for instance, might disclose some irrelevant bit of information to mislead competitors about the direction of the disclosing firm's research agenda. Alternatively, a firm might offer a publication reward for employees who comply with the firm's confidentiality policy and submit their work for review. Similar to the reasons mentioned above, these justifications fail to explain the disclosures at issue here; that is, disclosures intended to make the patent office aware that potentially patentable information is already in the public domain. We thank a referee for bringing to our attention many of the reasons for disclosure previously discussed.

What, then, is the reason for targeted disclosures? Certainly, these disclosures have some impact on the patenting process; otherwise firms would not bother.⁹

In assessing the impact of targeted disclosures, patent attorneys typically say that these disclosures are defensive in nature. What they mean is that the disclosures are designed to preempt patents in instances where the disclosing firm does not itself plan to pursue patent protection but fears that its rivals might. The logic is as follows. In all of the world's patent systems, patent applications are evaluated in light of the prior art, and patents issue only in instances where an alleged invention is a sufficient advance over that prior art. Firms like IBM and Xerox increase the scope of the prior art when they disclose. Hence, disclosure can serve a defensive purpose: disclosure makes it more difficult for rivals to patent inventions related to the disclosed information.

When might a firm engage in this sort of defensive maneuver? One typical setting would be an instance where the firm prefers state trade secret protection over federal patent protection. In such an instance, the firm might defensively disclose, revealing enough information to thwart rival patents but, beyond that, keeping its research secret. Another typical setting would be one where the disclosing firm doubts that its research will ever lead to a profitable commercial product. In that case, defensive disclosure allows the firm to avoid the cost of applying for a patent but nevertheless keeps rivals from acquiring exclusive rights—just in case circumstances change and the disclosing firm decides to return to this line of research.

These standard articulations of defensive disclosure resonate; certainly large-scale disclosures are sometimes made in instances where the disclosing party is not seeking patent protection and the disclosure is thus simply designed to thwart rival patent applications.¹⁰ But a close look at recent patenting behavior by Xerox and IBM suggests that this explanation may be incomplete. After all, between 1996 and 2001, there were nearly 150 patents issued where Xerox had to cite *its own journal* as prior art against a patent application filed on behalf of the firm.¹¹ For IBM, during that same time span, the

⁹ Note that making the PTO aware of disclosures entails a non-trivial cost. For example, IP.com customers pay \$155 to have their first set of disclosures turned into prior art (for each additional publication IP.com offers a quantity discount). See <http://ip.com>, visited September 29, 2004. Research Disclosure, Inc. charges less, \$120, for each page of manuscript submitted for publication. See <http://www.researchdisclosure.com/publish/dasr.html>, visited September 29, 2004. It is hard to know the costs IBM and Xerox faced in making their journals freely available to the patent office.

¹⁰ Another possible reason exists for disclosing to the PTO, one not emphasized in the literature or considered in the model. A firm might disclose in cases where there is a dearth of prior art and, in effect, educate the PTO about a certain technology. In so doing, the firm could establish what the prior art is. This, in turn, might enable the disclosing firm to use the now-established prior art to prove that a future innovation is non-obvious. Again, we thank a referee for bringing this possible reason for disclosure to our attention.

¹¹ The number quoted in this sentence comes from an electronic search of the United States Patent & Trademark Office's archive of issued patents. For patents issued between January 1, 1996, and July 17, 2001, there were 362 citations to the Xerox Disclosure Journal, 139 of which

number of self-citations exceeds 2,300—or, to put it into perspective, includes nearly 1 of every 6 patents assigned to IBM in those years.¹² And while self-citations might simply be evidence of a change in the relevant firm’s research focus—deciding to abandon the project one year, then reversing that decision years later—in many of these applications the publication of the journal article is remarkably close in time to the filing of the patent application.¹³ The close dates suggest that these firms are not simply reacting to changing research priorities. Instead, the close dates of the self-citations are consistent with firms knowingly publishing information about ongoing research projects.¹⁴

In this paper we therefore set out to provide an alternative explanation for why, as part of their formal intellectual property strategies, research firms intentionally disclose sensitive information to the PTO. Our basic insight is that firms have reason to disclose even in cases where they still plan to pursue patents related to the disclosed information. The conventional wisdom does not fully capture the potential advantages of disclosure; it is not merely some spoiler strategy played by firms that expect to exit a given patent race. Disclosure can in addition be a rational strategy for firms that plan to continue racing.

The intuition is straightforward. Because patents are evaluated in light of the prior art, disclosures by one firm make it more difficult for any firm to claim a related patent. Disclosure in essence extends the race. If an invention of a certain quality would have been sufficient to qualify for patent protection before the disclosure, after the disclosure the invention must be that much better before it will represent a sufficient advance over the now-expanded prior art. Our point is that a firm actively engaged in a patent race might very well have an incentive to extend the race. For a firm trailing in a given patent race, a longer race might offer a better opportunity to catch up. For a firm

came on patents assigned back to Xerox itself. The total number of patents assigned to Xerox during this time period was 4,962.

¹² More specifically, for patents issued between January 1, 1996, and July 17, 2001, there were 9,066 citations to the IBM Technical Disclosure Bulletin, 2,316 of which came on patents assigned to IBM. The total number of patents assigned to IBM during the time period was 13,854.

¹³ Some representative examples include United States Patent No. 6,256,775 (application assigned to IBM, filed in 1997, citing 1996 article from IBM Technical Disclosure Bulletin as limiting prior art); United States Patent No. 6,253,279 (again assigned to IBM, filed in 1998, citing 1995 IBM article); United States Patent No. 6,229,114 (assigned to Xerox, filed in 1999, citing 1991 article from Xerox Disclosure Journal). A fuller case study of IBM’s citations to its own journal is presented in Part IV.

¹⁴ This point has not been stressed by patent practitioners, policy-makers, and even industry insiders. Typically, they describe disclosure’s strategic value as arising when, at the time of disclosure, the disclosing firm does not itself plan to pursue patent protection. This is the logic put forward in the materials advertising services related to disclosure. See, for example, <http://www.ip.com/about.jsp?id=protectInnovation>, visited September 29, 2004. This is also the logic typically used to explain the Statutory Invention Registration (52 U.S.C. § 157), a federal program that helps firms disseminate information about unpatented inventions. See, for example, Federal Register Document 99-30598 (Department of Commerce document explaining that Statutory Invention Registration is useful to parties who do “not want to go through the effort and expense of obtaining a patent on the invention” but “[at] the same time, . . . [want] to prevent someone else from later obtaining a patent on a like invention.”).

leading a given patent race, extending the race raises the costs of racing, a strategy that will in certain instances discourage the laggard from racing so aggressively.¹⁵

In short, this paper offers a fuller account of the reasons for disclosure. Conventional accounts stress that disclosure can benefit a firm that does not itself seek patent protection; our account shows that disclosure can in addition benefit a firm that is in fact seeking patent protection.¹⁶ We proceed as follows. Section II presents our model of disclosure in the context of a patent race. In Section III we show that disclosure can be attractive to firms engaged in a patent race, and we show how the incentive to disclose varies with changes in the legal rules and firm attributes. Section IV presents supportive empirical evidence—namely, a case study of IBM’s disclosure practices as evidenced by patents issued between January 1, 1996, and July 17, 2001. Finally, Section V concludes, providing one possible extension and showing how our work fits within the general debate over the proper scope of the prior art rules and the effectiveness of the patent prosecution process.

II. THE MODEL

Our model draws on three related literatures. First and most obviously, this is a patent race model, and in structuring it we clearly benefited from the existing literature on patent races; Reinganum¹⁷ provides a helpful survey. Second, the literature on knowledge spillovers suggested to us several ways to model transfers of knowledge between firms. This literature originally focused on inadvertent transfers of knowledge, but recent contributions like De Fraja¹⁸ and Anton & Yao¹⁹ consider the possibility of intentional spillovers, a modeling problem very similar to our own. Third, several papers on strategic disclosure helped us to understand the strengths and limitations inherent in a variety of possible modeling approaches. In specific, while not containing formal

¹⁵ In formalizing this point, we incorporate within our more general framework the idea that leaders in a patent race can use disclosure to reduce the expected value of the patent for the laggard in the race. See Douglas Lichtman, Scott Baker & Kate Kraus, *Strategic Disclosure in the Patent System*, 53 *VAND. L. REV.* 2175, 2205-07 (2000) (discussing how a leading firm might use disclosure to reduce the expected value of the patent to the laggard).

¹⁶ Interestingly, both explanations fit within the more general framework of strategies where a firm attempts to raise its rival’s costs. In the traditional explanation, this has only a negligible effect on the disclosing firm’s own costs since the firm is assumed not to be pursuing patent protection. In our explanation, by contrast, disclosure does increase the disclosing firm’s costs but nevertheless benefits the disclosing firm in certain cases. There is a large industrial organization literature on strategies of this sort. See, for example, Steven C. Salop & David T. Scheffman, *Raising Rival’s Costs*, 73 *AM. ECON. REV.* 267 (1983).

¹⁷ Jennifer Reinganum, *The Timing of Innovation: Research, Development, and Diffusion*, 1 *Handbook of Industrial Organization* (Richard Schmalensee & Robert Willig, eds., 1989).

¹⁸ Giovanni De Fraja, *Strategic Spillovers in Patent Races*, 11 *INT’L J. INDUS. ORG.* 139-146 (1992).

¹⁹ James J. Anton & Dennis Yao, *Patents, Invalidity, and the Strategic Transmission of Information*, 12 *J. ECON. & MGMT.* (2003).

models, the text discussions in Parchomovsky²⁰ and Eisenberg²¹ outlined several ways of articulating pure spoiler strategies. Along related lines, Parchomovsky and Bar-Gill's recent paper²² shows that firms might disclose to weaken protection for their own patents. A weak patent, they argue, is less likely to block other firms from recouping investment in complementary innovations. As a result, complementary investment and innovation increase, rendering the weak patent more valuable to the disclosing firm. Finally, from the literature, the incomplete information model in Lichtman, Baker & Kraus²³ confirmed for us that the incomplete information framework, while useful to show the various signaling effects that were central to that paper, would not adequately explain why firms try so hard to make the PTO aware of their disclosures.

Now, we turn to our model setup. Our model considers strategic disclosure in the context of a two-firm patent race. One firm, M , begins the race with knowledge m . The other firm, N , begins the race with knowledge n . Prior art relevant to the race and already known to the public is represented by the variable p . For ease of exposition, knowledge in this race is ordered in a linear fashion. Thus, $m > n$ means that firm M knows everything firm N knows, plus more.²⁴ Naturally, $\min\{m, n\} \geq p$; in words, each firm knows at least as much as is in the public domain.²⁵

A patent issues in this model whenever a given firm's knowledge exceeds the prior art by some sufficient measure (determined by law).²⁶ Patents vary in their reach. Some only advance prior art by a small amount, other by a larger extent. Patents that have

²⁰ Gideon Parchomovsky, *Publish or Perish*, 98 MICH. L. REV. 926 (2000).

²¹ Rebecca Eisenberg, *The Promise and Perils of Strategic Prior Art Creation Through Publication: A Response to Professor Parchomovsky*, 98 MICH. L. REV. (2000).

²² Owen Bar-Gill & Gideon Parchomovsky, *The Value of Giving Away Secrets*, 89 VA. L. REV. 1857 (2003).

²³ Lichtman, Baker & Kraus, *supra* note 15.

²⁴ This is of course a simplification. In the real world, firms can pursue different research paths, and one firm might very well have less total information than its rival but nevertheless know some information that its rival does not. These sorts of complexities are likely to make disclosure more difficult. For instance, they introduce the possibility that the disclosing firm could accidentally reveal information that its rival needs, and the possibility that the disclosing firm could reveal information only to find out later that it was completely irrelevant to (and thus did not in any way delay) the rival's patent application.

²⁵ There is no reason to believe that having only two firms in any way distorts the results. Additional firms increase the need for disclosure in many cases, but they also increase the odds that a given disclosure will inadvertently assist a rival.

²⁶ We do not delve here into all the nuances of the prior art rules. Interested readers are encouraged to consult Lichtman, Baker & Kraus, *supra* note 15, at 2180-89 & 2197-98, for a full discussion. One important thing to note, however, is the statutory bar for prior publications. See 35 U.S.C.A. § 102(b). Under this provision, a firm that publishes a result may still pursue patent protection if it files its application within a year of the publication. If it doesn't file within a year, the publication prevents patentability. The model considers publications that occur inside the statutory bar (more than one year in advance of the filing). That is, we assume that if a firm chooses to disclose it cannot, then, seek patent protection for the disclosed knowledge.

a broader scope also have more value to the firm. To capture this in the simplest way, we consider two types of patents: narrow and broad. We denote by Δ the knowledge margin necessary to obtain a narrow patent and by $\Delta+\beta$ the margin necessary to obtain a broad patent, $\beta > 0$. We let V be the value to the firm of a narrow patent and $(1+\alpha)V$ the value of a broad patent, $\alpha > 0$. As figure 1 shows, then, variables m and n mark the firms' relative positions at the start of the race, while variables p , Δ and β combine to determine the threshold of patentability. Note that only one firm can win any given patent. In our setting the patent will go to the first firm to achieve total knowledge of at least $p+\Delta$. The interesting case to consider is the one in which $\max\{m,n\} < p+\Delta$, so that at the beginning of the race no firm is in a position to win without acquiring some additional knowledge.

INSERT FIGURE ONE HERE

Most patent race models focus on the process through which firms gain knowledge and move closer to earning the patent. Conventional models, in other words, focus on changes to variables like m and n . What is innovative about this model is that, here, we consider changes to p as well. Firms in our model are allowed to disclose information to the PTO, thereby increasing p and raising the threshold of patentability, $p+\Delta$. This gives the firms what can fairly be described as a unilateral right to extend the race. The only constraint is that a firm can extend the race only if it has relevant knowledge to disclose.

We will focus on a sequence of moves in three periods. As will become clear, it takes three periods to capture the core disclosure interaction. A smaller number of periods offers no real incentive for disclosure by either firm. A larger number of periods, by contrast, makes the model more complicated, without adding any important new insight.

In the first period, firm M conducts research, earns an increment of knowledge, and then either wins the patent or makes a disclosure. Firm M can choose to disclose any amount of information, ranging from nothing to its full knowledge. In the second period, it is firm N 's turn to conduct research, increase its knowledge, and then either win the patent or make a disclosure. Finally, in the third period, M conducts research once more, acquiring knowledge and again either winning the patent or making a disclosure. It is important to note that any knowledge that is strategically disclosed by a firm cannot be patented in the future.²⁷ This is because disclosed knowledge becomes prior art and a patent is won only if a firm produces a minimum knowledge increment over prior art (Δ for a narrow patent and $\Delta+\beta$ for a broad patent).

²⁷ This is subject, of course, to the qualification about the statutory bar discussed earlier. See *infra* note 26.

The allure of this three-period model is that disclosure, if it happens at all, will happen only in the first period. Think, for example, about the third period. Firm M has no reason to disclose in the third period since the game ends there regardless. In the second period, firm N similarly has no reason to disclose. After all, if N has not won the patent by the end of the second period, N has no chance of winning later in the game (it has no moves left) and thus has nothing to gain from disclosure.²⁸ So the first period is the heart of the three-period disclosure model. If firm M does not itself win the patent in the first period, M might disclose in the hope that a disclosure will decrease firm N 's chance of winning in the second period and thereby increase the chance that firm M will have the opportunity to try again in the third period.

III. RESULTS

We can solve for firm M 's optimal disclosure strategy in the first period. Suppose that, in the first period, firm M draws a random increment of knowledge (call it k_m) from some distribution $F(\cdot)$, where $F(k) = \text{prob}(k_m \leq k)$. In words, $F(k)$ is the probability that M 's draw in this period will be less than or equal to the number k . If that knowledge puts M over the patentability threshold, the game will immediately end.²⁹ If not, however, firm M has to make a decision about disclosure.

In making its decision about disclosure, firm M must consider the fact that, in the next period, firm N will get a knowledge draw and will itself have a chance to win the patent. In particular, firm N will draw a random increment of knowledge (call it k_n) from a distribution $G(\cdot)$, where $G(\cdot)$, like $F(\cdot)$, is simply a distribution of the form $G(k) = \text{prob}(k_n \leq k)$. Note that there is no specific relationship between $F(\cdot)$ and $G(\cdot)$. So, for example, if firm M has a better research staff than does firm N , $F(\cdot)$ will on average yield higher knowledge draws than $G(\cdot)$. Conversely, if N has the better scientists, $G(\cdot)$ will tend to yield better knowledge draws than $F(\cdot)$.³⁰

Using these variables, we can now calculate the specific probabilities affecting M 's disclosure decision. Define d to represent the size of firm M 's disclosure in the first period. A disclosure of size d increases the level of the prior art from p to $p + d$. Naturally, d cannot be less than zero since M cannot make a negative disclosure, and d

²⁸ To be precise, it is also true that firms M and N do not have anything to lose from disclosing in the second and third periods; however, it would be sufficient to add a small cost to disclosure—a reasonable assumption—to make “no disclosure” strictly dominant. See the cost data listed *supra* note 9. For this reason, we will only consider the case in which firms do not disclose in periods two and three.

²⁹ Here we focus on the case in which firm M prefers to obtain a narrow patent in period 1, if it can, as opposed to waiting period 3 in the hope of obtaining a broad patent but risking that firm N wins the patent in period 2. For sufficiently small Δ or α this will be the case. More precisely, a sufficient condition is $1 > G(\Delta)(1 + \alpha)$. See *infra* note 31 for a derivation of this condition.

³⁰ For technical reasons that will become clear later, we assume that the distributions F and G have logconcave densities f and g .

cannot be more than $m + k_m - p$ since, at most, M can disclose everything it knows that is not already publicly known.

Define w_n to be firm N 's "knowledge threshold"; that is, w_n is the minimum knowledge draw that firm N needs to win the (narrow) patent. Since N will win the patent if its knowledge satisfies the inequality

$$\max\{n, p + d\} + w_n \geq p + d + \Delta,$$

we know that

$$w_n = \min\{\Delta, p + d + \Delta - n\}.$$

Note that the probability that firm N will fail to win the patent in period two is $G(w_n)$, and the probability that firm N will win the patent is $(1 - G(w_n))$.

Similarly, define w_m to be the minimum knowledge draw that M will need to win a narrow patent in period three (firm M 's "knowledge threshold"). M will win a narrow patent if

$$p + d + \Delta + \beta \geq m + k_m + w_m \geq p + d + \Delta,$$

and thus

$$(1) \quad w_m = p + d + \Delta - m - k_m.$$

Clearly, $w_m + \beta$ is the minimum knowledge draw that M will need to win a broad patent in period three. If firm N fails to win in the second period, then, the probability that firm M will win a narrow patent in the third period is simply $(F(w_m + \beta) - F(w_m))$, while the probability that it will win a broad patent is $(1 - F(w_m + \beta))$.

If firm M fails to win the narrow patent in the first period (that is, if $k_m + m < p + \Delta$), M will choose a disclosure to maximize the product of two terms: the probability of firm N losing in the second period, and the expected value of winning the patent in the third period. In the language of the model, M will attempt to choose d so as to maximize

$$(2) \quad G(w_n)[V(1 - F(w_m)) + \alpha V(1 - F(w_m + \beta))]$$

subject to the constraints $d \geq 0$ and $d \leq m+k_m-p$.³¹ Note that, as long as $p+d < n$ or $d < n-p$, disclosure is a double-edged sword in equation (2): it increases the odds that firm N will lose in the second period (good for M), but it also decreases both the odds that firm M will win and the odds that it will obtain a broad patent in the third period (bad for M). Note, too, that the optimal level of disclosure d^* will never be greater than $n-p$ because firm M has no incentive to raise the level of the prior art above firm N 's knowledge level. Doing so, after all, would hurt M without in any way harming N ; that is, it would reduce $1-F(\cdot)$ without increasing $G(\cdot)$.³² Thus, we can rewrite w_n as follows

$$(3) \quad w_n = p + d + \Delta - n .$$

We can now state and prove five propositions that capture the basic relationships between strategic disclosure and the six model parameters, m , n , p , Δ , α and β . We will explain the intuitions behind these propositions below.

Proposition 1. The disclosing firm's optimal level of disclosure increases with m , the knowledge of the disclosing firm. This is true no matter whether the disclosing firm leads or trails in the race.

Proposition 2. The disclosing firm's optimal level of disclosure increases with n , the knowledge of the rival firm. This is true no matter whether the disclosing firm leads or trails in the race.

Proposition 3. The disclosing firm's optimal level of disclosure decreases with the prior art p and with Δ , the distance by which narrow patentable inventions must exceed the prior art. This is true no matter whether the disclosing firm leads or trails in the race.

Proposition 4. The disclosing firm's optimal level of disclosure decreases with β , the distance by which broad patentable inventions must exceed narrow patentable inventions. This is true no matter whether the disclosing firm leads or trails in the race.

³¹ Suppose that firm M can obtain a narrow patent, but not a broad one, in period one; the payoff from a narrow patent is V . The payoff from waiting until period three, in the hope of obtaining a broad patent, is $VG(w_n)[1+\alpha(1-F(w_m+\beta))]$, which is bounded above by $VG(\Delta)(1+\alpha)$, since $w_n \leq \Delta$. Hence $1 > G(\Delta)(1+\alpha)$ implies that firm M will file for a narrow patent in period one instead of waiting.

³² One may suppose that research as a function of a firm's knowledge level exhibits decreasing returns. This could be modeled by conditioning the distributions G and F on, respectively, firm N and firm M knowledge; that is, we could write $G(\cdot | \max\{n, p+d\})$ and $F(\cdot | m+k_m)$. In this case, if $p+d > n$, then disclosure has the benefit of pushing firm N into a region of decreasing returns; in some extreme cases, this benefit of disclosure may outweigh the cost of increasing M 's knowledge threshold w_m , and firm M may want to disclose more than $n-p$.

Proposition 5. The disclosing firm's optimal level of disclosure decreases with α , which is a measure of the difference in scope between a broad and a narrow patent. This is true no matter whether the disclosing firm leads or trails in the race.

Proofs

[See Appendix]

It is useful to consider as an example the case in which firms M and N both draw knowledge k from the uniform distribution on $[0,1]$. Mathematically, $F(k) = G(k) = k$. In this case equation (A6) in the appendix, the relevant first order condition, becomes:

$$(4) \quad 1 - w_m + \alpha - \alpha(w_m + \beta) - w_n(1 + \alpha) = 0$$

which, when we substitute in equation (1) for w_m and equation (3) for w_n , gives us the optimal level of disclosure in the example:

$$(5) \quad d^* = \frac{1}{2} \left(1 - \frac{\alpha\beta}{1 + \alpha} - 2p - 2\Delta + m + n + k_m \right).$$

The intuition for propositions 3-5 is easiest to understand. Proposition 3 asserts that increases in p or Δ lead to decreases in disclosure. This follows because, all else held equal, an increase in p or Δ makes it less likely that firm N will win the patent. Disclosure is attractive to M only because it decreases N 's chances of winning; so, because increases in p or Δ already make N less likely to win, M 's incentive to disclose is diminished.

Propositions 4 and 5 show that an increase in either β or α also decreases disclosure. This is because disclosure has the cost of reducing the likelihood that firm M will be able to obtain a broad patent. By making the broad patent more valuable, an increase in α increases this cost – without affecting the benefit of disclosure -- and thus makes firm M less willing to disclose. Likewise, an increase in β makes it more difficult to obtain a broad patent and as such reduces the incentive for firm M to disclose.

Proposition 2 can be explained along similar lines. Proposition 2 asserts that increases in firm N 's knowledge lead to increases in disclosure by firm M . This follows because any increase in N 's knowledge increases N 's likelihood of winning the patent. Naturally, M responds with an additional disclosure in the hope of offsetting that effect.

Proposition 1 asserts that an increase in firm M 's own knowledge also leads to an increase in disclosure. The logic here is that disclosure allows firm M to trade increases in its own chance of winning (in period three) for increases in firm N 's probability of losing (in period two). Without a change in disclosure, an increase in firm M 's knowledge would affect only the first of these terms; it would increase firm M 's chance of winning in the third period but it would not at all affect firm N 's probability of losing in the second period. With an increase in disclosure, by contrast, firm M can spread the

effect across both terms. It can diminish its own (now increased) chance of winning in the third period but correspondingly increase firm N 's probability of losing in the second period. The extent to which this trade-off will benefit M depends on the specifics of each case, but, as the proposition states, we know across all cases that firm M will never react to an increase in knowledge by disclosing *less*, since that would just further skew the probabilities in favor of the first term as opposed to the second.

Note that none of the above intuitions requires any specific assumption as to whether the disclosing firm leads or lags in the race. Thus all of the propositions explicitly apply to both leader and laggard firms.

So far we have assumed that if no firm gains a patent, then when the game ends at the end of period three any partial knowledge acquired has no value. It is interesting to ask what would happen if partial progress had value. To address this issue, suppose that if no firm gains a patent then they compete in the market and each firm's profit level at the end of period three is an increasing function of its own knowledge and a decreasing function of the knowledge of the other firm.³³

If it is unable to win the race, at the end of period two firm N 's dominant strategy is to disclose all of its knowledge up to $m + k_m$, -- the present knowledge level of firm M . By raising the prior art level to $p_2 = \min(n + k_n, m + k_m)$, firm N lowers the probability that firm M will win the race in period three, without affecting its payoff in case firm M does not win the race and the two firms compete in the market. On the other hand, if firm N discloses beyond firm M 's knowledge $m_2 = m + k_m$, then it decreases its payoff in case firm M does not win the race (because it raises firm M 's knowledge), without raising the probability that firm M will not win the race (because a unit disclosure beyond the present knowledge level of firm M raises both the prior art and firm M 's knowledge by one unit, and thus does not affect the threshold knowledge increase that firm M needs to win the race in period three). This, in turn, means that it is not in firm N 's best interest to disclose more than firm M already knows. Given firm N 's dominant strategy, the minimum draw that M needs to win a narrow patent in period three is

$$(6) \quad w_m^* = p_2 + \Delta - m_2 .$$

If no firm wins the race, the levels of partial knowledge of firms M and N at the end of period three are $m_2 + k_3$ and $n + k_n$ respectively; thus, we can denote firm M 's profit level if no firms win the race as $\pi(m_2 + k_3, n + k_n)$. If firm M fails to win in the first period, then firm M chooses a disclosure d to maximize

³³ We assume that a firm's own knowledge increases its profit even if that knowledge was previously disclosed, so long as no patent has issued. This would be true if, for example, knowledge reduced the marginal cost of production, or, alternatively, if rival firms could not quickly take advantage of disclosed information.

(7)

$$G(w_n) \left[V(1 - F(w_m^*)) + \alpha V(1 - F(w_m^* + \beta)) \right] + \int_0^{w_n} \int_0^{w_m^*} \pi(m_2 + k_3, n + k_n) g(k_n) f(k_3) dk_n dk_3.$$

Equation (7) depends on d only through w_n , which is increasing in d . Hence, the optimal level of disclosure by firm M in this case is $d^* = \min(n, m + k_m) - p$, which is clearly greater than in the case in which partial knowledge has no value. The intuition for why disclosure by firm M increases when partial knowledge has value is clear. Since in period two firm N will disclose all its knowledge up to the knowledge level of firm M , in period one it is optimal for firm M also to disclose all its knowledge up to the knowledge level of firm N . By doing so, firm M decreases the probability that firm N will win the race in period two, without affecting its payoff if firm N does not win the race.

Our model does not allow firms to negotiate.³⁴ If it did, one might expect that, instead of disclosing, firms would simply threaten to disclose and then negotiate a private agreement in the shadow of that threat. These sorts of bargains would be attractive to the firms since, unlike public disclosure, a private agreement would not prolong the race. So, as long as the parties could agree on how to divide the surplus generated by an earlier patent and cheaper race, they should prefer private negotiation over public revelation. The main reason that we do not allow private negotiations is that in many instances such negotiations are extremely unlikely. For one thing, they would take place under time pressure. Delay would give the non-disclosing firm time to advance its research; perhaps enough to file for the patent and in that way nullify the disclosure threat. The disclosing firm would thus need to either strike a bargain immediately or abandon the negotiation and disclose. This would exacerbate standard impediments to bargaining—for example, disagreements as to the effect any given disclosure would have on the race, disagreements as to the size of the surplus generated, and disagreements over the fair division of that surplus.

Moreover, to whatever extent firms use disclosure as a way of undermining a rival's incentive to race, private negotiations seem particularly difficult. The difficulty comes in specifying contract terms and monitoring compliance with those terms, especially where, as here, the contract involves new and developing technology.³⁵

IV. EMPIRICAL EVIDENCE

To find empirical support for our basic argument that firms publicly disclose research information in ongoing patent races, we conducted a case study of IBM's research disclosures as reflected by patent citations to its most prominent journal, the *IBM Technical Disclosure Bulletin*. Specifically, using a database provided by the United States Patent & Trademark Office, we identified the 13,854 utility patents issued between January 1, 1996, and July 17, 2001, for which “International Business

³⁴ The possibility of negotiations between firms prior to disclosure is also considered Lichtman, Baker, & Kraus, *supra* note 15, at 2213-2216 and Parchomovsky, *supra* note 20, at 948-50.

³⁵ Lichtman, Baker, & Kraus, *supra* note 15, at 2215 outlines these difficulties in more detail.

Machines” was the patent assignee.³⁶ Examining those electronically, we found that 2,310 of them cite as prior art at least one article from the *IBM Technical Disclosure Bulletin*.³⁷ So, in approximately one out of every six patents issued to IBM during this time period, IBM’s own disclosures are cited.

Standing on its own, this data does not necessarily support our thesis. One could interpret these numbers not as evidence of the theory presented here, but instead as evidence that (often) IBM employees publish an article thinking that the firm is abandoning a given patent race but then, years later, priorities change and IBM resumes the relevant race. Stated differently, for the data to be consistent with this abandonment or defensive disclosure thesis, a significant amount of time should have elapsed between the date of the publication and the date of the patent application. After all, a long time lag would give IBM researchers a chance to abandon a certain R&D agenda and then reenter when things change. If instead the publication pertained to IBM’s active research agenda, we should observe a short time lag between publication and patent application.

Of course, it is hard to define what exactly constitutes an “active” patent race. The idea of a patent race over a single innovation is an abstraction. In reality, firms, like IBM, pursue a variety of interrelated innovations at the same time. Moreover, for each patent studied, we do not know -- and cannot know -- how many other firms were pursuing similar innovations. Furthermore, because of the secrecy surrounding R&D programs, we cannot observe whether firms dropped certain parts of their R&D agenda because of IBM publications. Our qualified claim is that the empirical results presented below are consistent with the thesis advanced in this paper and inconsistent with the most prominent explanation – defensive disclosure – set forth by patent attorneys and the academic literature.³⁸

To get at the time lapse, we analyzed each of the 2,310 patents at issue and calculated the number of months that passed between the date of the relevant publication³⁹ and the date on which the patent application was ultimately filed.⁴⁰ As

³⁶ Electronic searches were conducted using both the official patent database compiled by the United States Patent & Trademark Office (<http://www.uspto.gov>) and a database commercially available from the Delphion Intellectual Property Network (<http://www.delphion.com>). On both databases, our first search was for the phrase “International Business Machines” in the “assignee” field of any utility patent.

³⁷ Our search here was for any utility patent with the phrase “International Business Machines” in the “assignee” field and a citation to “IBM Technical Disclosure Bulletin” in its list of prior art references.

³⁸ Indeed, we want to be careful not to overstate the empirical claim. Although consistent with our thesis, self-citation is also consistent with a more mundane explanation: Inventors are familiar with the contents of their own technical journals. As a result, in providing common background information for the patent application, an inventor might cite the technical journal simply because that is the journal he is most familiar with.

³⁹ In the event of multiple citations, we used the date of the most recent one.

⁴⁰ After analyzing over 200 patents by hand, we were able to automate the process so that the computer would parse the full text of each patent and identify for us the relevant dates. For most

stated, long gaps would support the idea of changed research priorities. Short gaps, by contrast, are consistent with the disclosure strategy outlined here.

For ten of the patents, there was insufficient information to make this calculation.⁴¹ For the remaining 2,300 patents, however, we obtained the results shown below. Note that we rounded the tallies so that gaps of 1-12 months were labeled as 1 year, gaps of 13-24 months were labeled as 2 years, and so on. Gaps of over 10 years are not shown in the graph but are accounted for in the table.

INSERT FIGURE TWO HERE

patents, the necessary dates were easy to identify since filing dates are well labeled on all issued patents, and prior art cited in a patent typically also includes a well-labeled reference to the date of publication. Any patents that the computer deemed unclear were processed manually, and all results were spot-checked for accuracy.

⁴¹ These ten patents either gave incomplete or inconsistent citation information and, thus, we were unable to fairly assign an accurate publication date in these rare instances.

INSERT TABLE ONE HERE

The most important trend is simply this: in 54% of the cases, the gap between IBM's publication and IBM's patent application is less than 5 years. Since it would be surprising if IBM's research agenda were to change that quickly in so many cases, this seems to lend support to the theory underlying the paper. It seems fair to claim that the empirical evidence is consistent with IBM often disclosing information about patents that it is actively pursuing, rather than simply disclosing information relevant to races it plans to abandon.

VI. CONCLUSION

The primary focus in this paper has been to offer a richer account of why research firms disclose information to the PTO. The traditional explanation focuses on using disclosure defensively to stop rivals from patenting in instances where the disclosing firm itself does not intend to patent. We propose here an alternative explanation: that disclosure might be a strategy through which a research firm can unilaterally extend the patent race and thereby gain advantage even in a patent race it has no intention of abandoning. Our explanation is consistent with certain elements of the evidence—for example, IBM's disclosure and patent practices—that cannot be explained solely by defensive disclosure.

Our work has other implications as well, however, and in conclusion we thought it interesting to outline two. First, in this paper we have focused almost exclusively on strategic disclosure as an uncooperative strategy, with firms disclosing for the sole purpose of harming their rivals. Note, however, that disclosure can be used in cooperative settings as well. Consider, for example, two firms both researching the same basic invention. Obviously, the firms could be made better off if they could coordinate their research agendas—say, sharing early research results, slowing the pace of invention, or agreeing to each specialize in one aspect of the invention. Enforcing this sort of coordination by contract would be difficult, however, since any such contract would have to specify behaviors related to a still-evolving technology—a tricky business to be sure.⁴²

The threat of disclosure presents a workable alternative. Indeed, just as traditional cartels work because cooperative firms can punish uncooperative firms simply by lowering the market price of a relevant good; in research settings—where the coordination relates to on-going research and hence there is no relevant price to lower—a similarly clear and effective penalty is the disclosure of research information to the public. In fact, it is hard to identify any attribute that the pricing mechanism enjoys (clear, public, effective, unilateral) that the disclosure mechanism cannot also claim.

⁴² There would also be concerns about antitrust enforcement; although, as the Department of Justice and the Federal Trade Commission both now recognize, certain types of joint research ventures likely benefit consumers and so some such contracts would not lead to antitrust challenge. See Antitrust Guidelines for Collaborations among Competitors (April 2000), available at <http://www.ftc.gov/os/2000/04/ftcdojguidelines.pdf>, visited September 29, 2004.

More broadly, then, an interesting extension to this work would be to shift the paper's analysis from the uncooperative settings that have been its focus to potentially cooperative settings where the threat of disclosure would serve as a mechanism to maintain cooperation among would-be cooperative partners.⁴³

Finally, we consider how our analysis fits into the larger debate over the scope of prior art rules and the efficacy of the patent prosecution process. Patents issued to firms like Amazon (the "one-click" patent⁴⁴) and Priceline (the reverse auction mechanism⁴⁵) have led many commentators to worry that the PTO cannot evaluate high technology patents. The problem is that most of the prior art available to patent examiners is prior art associated with preexisting patents; and so, in fields where few patents have issued, examiners simply do not have the written, archived evidence they need to conduct a thorough and accurate prior art review. Compounding the problem is the fact that patent applicants need only submit with their applications prior art that they know about.⁴⁶ There is no additional requirement that the patent applicant conducts a detailed search of prior art. The burden of searching prior art rests almost entirely with the patent examiner. Because of this fact, Mark Lemley has argued that "the PTO issues many patents that would have been rejected had the examiner possessed perfect knowledge."⁴⁷

To address this problem, scholars have advocated a variety of approaches. Jay Kesan and Marc Banik, for instance, have recently argued that patents should be given more favorable presumptions of validity in exchange for more complete prior art disclosures.⁴⁸ Jay Thomas has argued that parties should be given financial rewards for bringing forward prior art that defeats in-process patent applications.⁴⁹ Mark Lemley takes a slightly different approach. He contends that it makes sense for the PTO to be rationally ignorant about much prior art. More succinctly, Lemley contends that it is not cost effective for the PTO (or anyone for that matter) to conduct a detailed search of prior art because "ninety-five percent of patents . . . will either never be used, or will be used in circumstances that don't crucially rely on the determination of validity."⁵⁰

⁴³ The authors have begun to pursue this extension; see Scott Baker & Claudio Mezzetti, *Using the Threat of Disclosure to Enforce Knowledge Sharing in Joint Ventures Which Span Multiple Innovation Markets* (unpublished working paper, 2001) (available from authors).

⁴⁴ U.S. Patent No. 5,960,411.

⁴⁵ U.S. Patent No. 5,794,207.

⁴⁶ 37 C.F.R. § 1.56 (2002).

⁴⁷ Mark A. Lemley, *Rational Ignorance at the Patent Office*, 95 *Nw. U. L. Rev.* 1495 (2001). Notably, Lemley suggests that much relevant prior art will be hard to find because it will consist of sales and uses by third parties. *Id.* at 1500. Our analysis does not tackle the interesting strategic issues involved with creating prior art through these sorts of mechanisms.

⁴⁸ Jay Kesan & Marc Banik, *Patents as Incomplete Contracts: Aligning Incentives for R&D Investment with Incentives to Disclose Prior Art*, 2 *WASH. U. J. L. & POL'Y* 23 (2000).

⁴⁹ John R. Thomas, *Collusion and Collective Action in the Patent System: A Proposal for Patent Bounties*, 2001 *U. Ill. L. Rev.* 305.

⁵⁰ Lemley, *supra* note 47, at 1511.

Our analysis predicts that where patents are highly valuable, firms have an incentive to make the PTO aware of prior art.⁵¹ In fact, because of strategic disclosures, the PTO may make a better determination of validity in just the subset of patents where validity is especially important.

⁵¹ The marginal benefit of disclosure depends on the value of the patent, V – it is equal to V times the left hand side of equation (A7) in the appendix. We simplified our analysis by assuming that the marginal cost of disclosure was equal to zero. If, however, the marginal cost of disclosure is positive, then it is easy to show that disclosure is an increasing function of V .

Appendix

Proofs of Propositions 1 through 5:

Define d^* to be the optimal level of disclosure from firm M 's perspective. Propositions 1 through 5 require us to show, in turn, that:

$$(A1) \quad \frac{\partial d^*}{\partial m} \geq 0,$$

$$(A2) \quad \frac{\partial d^*}{\partial n} \geq 0,$$

$$(A3) \quad \frac{\partial d^*}{\partial \Delta} = \frac{\partial d^*}{\partial p} \leq 0,$$

$$(A4) \quad \frac{\partial d^*}{\partial \beta} \leq 0 \text{ and}$$

$$(A5) \quad \frac{\partial d^*}{\partial \alpha} \leq 0.$$

for all parameter values $\{m, n, \Delta, \alpha, \beta\}$. The first step in doing so is to find the value of d^* that maximizes the expression in equation (2). At an interior solution ($m+k_m-p > d > 0$) we have the following first order condition:

(A6)

$$\frac{\partial w_n}{\partial d} \frac{\partial G(w_n)}{\partial w_n} [1 - F(w_m) + \alpha - \alpha F(w_m + \beta)] - G(w_n) \frac{\partial w_m}{\partial d} \left[\frac{\partial F(w_m)}{\partial w_m} + \alpha \frac{\partial F(w_m + \beta)}{\partial w_m} \right] = 0$$

We can rewrite $\frac{\partial G(w_n)}{\partial w_n}$ as $g(\cdot)$, where $g(\cdot)$ is the density associated with the distribution

$G(\cdot)$. Similarly, we can rewrite the partial derivative $\frac{\partial F(\cdot)}{\partial w_m}$ as $f(\cdot)$. It follows from

equation (1) that $\frac{\partial w_m}{\partial d}$ equals 1 for all parameter values. The underlying intuition is clear:

every unit of information disclosed by firm M in the first period raises by one unit the minimum level of knowledge that firm M must uncover in the third period to win the patent. The partial derivative $\frac{\partial w_n}{\partial d}$ in equation (A6) is also equal to 1. This is true

because M 's optimal disclosure strategy would never have M disclosing information that

N does not already know; such disclosures would hurt M but not hurt N . So $d^* \leq n-p$, and $\frac{\partial w_n}{\partial d}$ equals 1 for the same reason that $\frac{\partial w_m}{\partial d}$ equals 1.

Accordingly, we can now rewrite equation (A6) as:

$$(A7) \quad g(w_n)[1 - F(w_m) + \alpha - \alpha F(w_m + \beta)] - G(w_n)[f(w_m) + \alpha f(w_m + \beta)] = 0$$

The left-hand side of equation (A7) defines a function Φ , which depends on the choice variable d^* and (through w_m and w_n) the parameters m , n , k_m , Δ , α and β . The second order condition and equations (1) and (3) imply that

$$(A8) \quad \frac{\partial \Phi}{\partial d} \leq 0, \text{ and}$$

$$(A9) \quad \frac{\partial \Phi}{\partial \Delta} = \frac{\partial \Phi}{\partial p} = \frac{\partial \Phi}{\partial d} \leq 0,$$

The next four inequalities can be derived under the assumption that the densities $f(\cdot)$ and $g(\cdot)$ are logconcave (and that β is not too large). Logconcavity is a technical requirement about the shape of these densities, but it is not a particularly limiting restriction. It guarantees that $F(\cdot)$, $1-F(\cdot)$, $G(\cdot)$ and $1-G(\cdot)$ are logconcave functions. Most familiar densities—for example, the uniform, the normal, and the exponential density—satisfy it.⁵²

$$(A10) \quad \frac{\partial \Phi}{\partial m} \geq 0,$$

$$(A11) \quad \frac{\partial \Phi}{\partial n} \geq 0,$$

$$(A12) \quad \frac{\partial \Phi}{\partial \beta} \leq 0, \text{ and}$$

$$(A13) \quad \frac{\partial \Phi}{\partial \alpha} \leq 0.$$

Applying the implicit function theorem yields our desired comparative static results:

⁵² The condition also holds if the density is logistic, chi-squared, Laplace, and for some parameters, Weibull, gamma, and Beta. See Drew Fudenberg & Jean Tirole, *Game Theory* 267 (1996).

$$(A14) \quad \frac{\partial d^*}{\partial m} = -\frac{\partial \Phi / \partial m}{\partial \Phi / \partial d} \geq 0,$$

$$(A15) \quad \frac{\partial d^*}{\partial n} = -\frac{\partial \Phi / \partial n}{\partial \Phi / \partial d} \geq 0,$$

$$(A16) \quad \frac{\partial d^*}{\partial \Delta} = -\frac{\partial \Phi / \partial \Delta}{\partial \Phi / \partial d} = -\frac{\partial \Phi / \partial p}{\partial \Phi / \partial d} = \frac{\partial d^*}{\partial p} \leq 0,$$

$$(A17) \quad \frac{\partial d^*}{\partial \beta} = -\frac{\partial \Phi / \partial \beta}{\partial \Phi / \partial d} \leq 0, \text{ and}$$

$$(A18) \quad \frac{\partial d^*}{\partial \alpha} = -\frac{\partial \Phi / \partial \alpha}{\partial \Phi / \partial d} \leq 0.$$

Bibliography

Anton, James J. and Yao, Dennis. "Patents, Invalidity, and the Strategic Transmission of Information." *Journal of Economics and Management Strategy* 12 (2) 2003; 161-178.

Baker, Scott and Mezzetti, Claudio. "Using the Threat of Disclosure to Enforce Knowledge Sharing in Joint Ventures Which Span Multiple Innovation Markets." (unpublished working paper, 2001) (available from authors).

Bar-Gill, Owen. and Parchomovsky, Gideon. "The Value of Giving Away Secrets." *Virginia Law Review* 89 (2003): 1857-1895.

Bhattacharya, Sudipto, and Ritter, Jay R. "Innovation and Communication: Signalling with Partial Disclosure." *Review of Economic Studies* 50 (1983): 331-246.

Eisenberg, Rebecca. "The Promise and Perils of Strategic Prior Art Creation Through Publication: A Response to Professor Parchomovsky." *Michigan Law Review* 98 (2000): 2358-2370.

De Fraja, Giovanni. "Strategic Spillovers in Patent Races." *International Journal of Industrial Organization* 11 (1992): 139-146.

Fudenberg, Drew and Tirole, Jean. *Game Theory*. Cambridge, MA. MIT Press 1996.

Kesan, Jay and Banik, Marc. "Patents as Incomplete Contracts: Aligning Incentives for R&D Investment with Incentives to Disclose Prior Art." *Washington University Journal of Law and Policy* 2 (2000): 23-54.

Lemley, Mark A. "Rational Ignorance at the Patent Office." *Northwestern University Law Review* 95 (2001): 1495-1532.

Lichtman, Douglas; Baker, Scott; and Kraus, Kate. "Strategic Disclosure in the Patent System," *Vanderbilt Law Review* 53 (2000): 2175-2217.

Parchomovsky, Gideon. "Publish or Perish." *Michigan Law Review* 98 (2000): 926-952.

Reinganum, Jennifer. "The Timing of Innovation: Research, Development, and Diffusion." In 1 *Handbook of Industrial Organization*, edited by Richard Schmalensee and Robert Willig. New York, NY.: Elsevier Science Publishing Co., 1989.

Salop, Steven C. and Scheffman, David T. "Raising Rival's Costs." *American Economic Review* 73 (1983): 267-271.

Thomas, John R. "Collusion and Collective Action in the Patent System: A Proposal for Patent Bounties." *University of Illinois Law Review* (2001): 305-353.

Table

Table 1

Years (Publication Application)	to	Patents in Sample (Percentage of Total)
1		7.4 %
2		12.6 %
3		12.8 %
4		11.2 %
5		10.0 %
6		5.7 %
7		5.5 %
8		3.9 %
9		3.6 %
10		3.4 %
More than 10		27.5 %

Figures

Figure 1. Basic race structure




Figure 2

IBM Self-Citation Patents

