DISCLOSURE AND INVESTMENT AS STRATEGIES IN THE PATENT RACE*

Scott Baker
University of North Carolina
Department of Economics
sbaker@email.unc.edu

Doug Lichtman
University of Chicago
School of Law
dgl@uchicago.edu

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

abstract

Research firms disclose a surprisingly large amount of information to the public. 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. Extending the patent race can be an attractive strategy for a firm trailing in a given race since a longer race might offer that firm a better opportunity to catch up. Extending the race can similarly be attractive to a leading firm, since making the race longer raises the costs of racing, a strategy that will in certain instances discourage trailing firms from pursuing aggressively. This paper models disclosure strategies and explains how research firms use disclosure to gain a competitive edge. It also studies certain interactions between the decision to invest in accelerated research and the decision to disclose. Finally, it presents empirical evidence that IBM, at least, does engage in disclosures of the sort described here.

* For helpful comments and conversations, we thank Marty Adelman, Douglas Baird, Rebecca Eisenberg, Jack Goldsmith, Scott Kief, Bill Landes, Mark Lemley, Saul Levmore, Michael Meurer, Rob Merges, Randy Picker, Eric Posner, David Weisbach, and workshop participants at the Boston University Law School (April 2001), the Harvard University Law & Economics Workshop (May 2001), and the American Law & Economics Association Annual Meeting (May 2001).

Research firms occasionally allow their employees to make presentations at conferences and/or contribute articles to peer-reviewed journals. These activities unavoidably reveal some otherwise-proprietary research information, but they are nevertheless easy to understand. Small-scale disclosures like these not only reward employees for their achievements, but also lead to favorable publicity for the firm. Harder to understand are large-scale disclosures. The Xerox Corporation, for instance, publishes a bimonthly technical journal in which firm employees describe in detail their on-going research.[1] The journal is distributed to libraries and patent offices worldwide. IBM published a similar journal from 1958 until 1998;[2] and, as evidence that the IBM journal revealed valuable information, note that its articles have been cited over 48,000 times in later-filed United States patent applications.[3] If small-scale disclosures reward employees and generate favorable publicity for the firms, what purpose do these larger, more systematic disclosures serve?

Patent attorneys typically respond that large-scale 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 keep 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. But a close look at recent patenting behavior by Xerox and IBM suggests that this explanation must be

---

[1] The publication is called the *Xerox Disclosure Journal*. Further information is available online at http://www.xerox.com/xdj.

[2] 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.

[3] *See* http://www.patents.ibm.com/tdb. 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.htm). Our own search at the Patent & Trademark Office website revealed over 10,000 separate patent applications citing the *Bulletin* in just the last five years.

incomplete. After all, in just the last five years, there have been 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.[4] For IBM, that number exceeds 2,300—or, to put it into perspective, includes nearly 1 of every 6 patents assigned to IBM over the last five years.[5] And while some of these 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.[6] The close dates suggest that these firms are not simply reacting to changing research priorities, but are, instead, knowingly publishing information about active patent races.[7]

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 public. 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 is wrong; disclosure 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 here 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

---

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

[5] 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.

[6] 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 V.

[7] This is the point that patent practitioners, policy-makers, and even industry insiders all seem to miss. Conventional wisdom is that disclosure's strategic value comes only from settings where, at the time of disclosure, the disclosing firm does not itself plan to pursue patent protection. This is the logic law firms put forward in the materials they use to advertise services related to disclosure. *See, e.g.*, http://www.ip.com (advertising services of, among others, the law firm of McDermott, Will & Emery). 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, e.g.*, 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 "[a]t the same time, . . . [want] to prevent someone else from later obtaining a patent on a like invention."). Ironically, this is even how Xerox explains its own *Xerox Disclosure Journal*; *see* http://www.xerox.com/research/xdj ("Xerox, like several other large R&D corporations, publishes on the subject matter of inventions . . . for which patent protection is not warranted."). But these statements all overlook the fact that disclosure can benefit a firm actively pursuing related patents. These statements also provide no basis by which one might explain the empirical evidence sketched in the preceding footnotes and presented more fully in Part V.

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 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.

In short, this paper offers a fuller account of defensive disclosure. Conventional accounts point out 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.[8] We proceed as follows. Part I presents our basic model of disclosure in the context of a patent race. 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. Parts II and III then refine the model by considering interactions between the decision to invest in research and the decision to disclose. For example, we examine the extent to which a firm might increase its research investments so as to increase its ability to disclose research information. Part IV discusses some limitations to the model. Part V 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, Part VI concludes with two possible extensions.

I. The Baseline 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; Reinganum (1989) 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 Fraja (1993) and Anton & Yao (1999) consider the possibility of intentional spillovers, a modeling problem very similar to our own. Third and finally, several recent papers on strategic disclosure—including a paper by two of the three current authors—helped us to understand the strengths and limitations inherent in a variety of possible modeling approaches. In specific, while not containing formal models, the text discussions in Parchomovsky (2000) and Eisenberg (2000) outlined several ways of articulating pure spoiler strategies; and the incomplete information model in Lichtman, Baker & Kraus (2000) 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 capture the richer interaction between disclosure and investment that we have ultimately been able to model here.

---

[8] 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 to not 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, e.g.,* S. Salop & D. Scheffman, Raising Rival's Costs, 73 American Econ. Rev. 267 (1983).

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.[9] Naturally, $\min\{m, n\} \geq p$; in words, each firm knows at least as much as is in the public domain.

A patent issues in this model whenever a given firm's knowledge exceeds the prior art by some sufficient measure (determined by law).[10] Label the necessary margin Δ. As figure 1 shows, then, variables *m* and *n* mark the firms' relative positions at the start of the race, while variables *p* and Δ combine to determine the threshold of patentability. Note that only one firm can win any given patent, so in this setting the patent will go to the first firm to achieve total knowledge $p + \Delta$. 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.

```
Firm M     ----------------- m ---------------------->
Firm N     ---------------------------- n ------------>
Prior Art  ------ p ---------------------------------->
                  :-------------- Δ --------------:
```

Figure 1. Basic race structure.

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 unique about this model is that, here, we consider changes to *p* as well. Firms in our model are allowed to disclose information, 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, introduces needless complexity in that later periods basically repeat the interactions represented by the three-period game.

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

[9] We discuss the likely implications of different assumptions in Part IV.

[10] We do not delve here into the nuances of the prior art rules, mainly because two of the current authors have already written at length on the relevant statutory provisions. Interested readers are encourage to consult Lichtman, Baker & Kraus (2000), at pp. 2180-89 & 2197-98, for a full discussion.

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.

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.[11] 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.

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 $F(\bullet)$, where $F(k) = \text{prob}\ (k_m \leq k)$. In words, $F(\bullet)$ is the probability that *M*'s draw in this period will be less than or equal to some number, k. Obviously, if that knowledge puts *M* over the patentability threshold, the game will immediately end. 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 $G(\bullet)$, where $G(\bullet)$, like $F(\bullet)$, is simply a distribution of the form $G(k) = \text{prob}\ (k_n \leq k)$. Note that there is no specific relationship between $F(\bullet)$ and $G(\bullet)$. So, for example, if firm *M* has a better research staff than does firm *N*, $F(\bullet)$ will on average yield higher knowledge draws than $G(\bullet)$. Conversely, if *N* has the better scientists, $G(\bullet)$ will tend to yield better knowledge draws than $F(\bullet)$.[12]

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

---

[11] 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. For this reason, we will only consider the case in which firms do not disclose in periods two and three.

[12] For technical reasons that will become clear later, we assume that the distributions *F* and *G* have log concave densities *f* and *g*.

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 patent. Since *N* will win the patent if its knowledge satisfies the inequality

$$n + w_n \geq p + d + \Delta ,$$

we know that

$$w_n = p + d + \Delta - n . \quad (1)$$

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 the patent in period three (firm *M*'s "knowledge threshold"). *M* will win if

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

and thus

$$w_m = p + d + \Delta - m - k_m . \quad (2)$$

If firm *N* fails to win in the second period, then, the probability that firm *M* will win in the third period is simply $(1 - F(w_m))$.

If firm *M* fails to win 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 probability of *M* then winning in the third period. In the language of the model, *M* will attempt to choose *d* so as to maximize

$$G(w_n)(1 - F(w_m)) \quad (3)$$

subject to the constraints $d \geq 0$ and $d \leq m + k_m - p$. Note that disclosure is a double-edged sword in equation (3): it increases the odds that firm *N* will lose in the second period (good for *M*), but it also decreases the odds that firm *M* will win in the third period (bad for *M*).

We can now state and prove the following three propositions regarding strategic disclosure. The propositions are designed to capture the basic relationships between disclosure and the three model parameters, *m*, *n*, and Δ. We will explain the intuitions behind these propositions and also work through a specific example immediately 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 Δ, the distance by which patentable inventions must exceed the prior art. This is true no matter whether the disclosing firm leads or trails in the race.

A. Intuitions

The intuition for proposition 3 is easiest to understand. Proposition 3 asserts that increases in Δ lead to decreases in disclosure. This follows because, all else held equal, an increase in 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 Δ already make *N* less likely to win, *M*'s incentive to disclose is diminished.

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 is more complicated. 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 specific example, but, as the proposition states, we know across all examples 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.

B. Proofs

Define d^* to be the optimal level of disclosure from firm *M*'s perspective. Our three propositions require us to show, in turn, that:

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

(5) $$\frac{\partial d^*}{\partial n} \geq 0 \text{ and}$$

(6) $$\frac{\partial d^*}{\partial \Delta} \leq 0$$

for all parameter values $\{m, n, \Delta\}$. The first step in doing so is to find the value of d^* that maximizes the expression in equation (3). At an interior solution ($m+k_m-p > d > 0$) we have the following first order condition:

(7) $$\frac{\partial w_n}{\partial d}\frac{\partial G(w_n)}{\partial w_n}(1-F(w_m)) - \frac{\partial w_m}{\partial d}\frac{\partial F(w_m)}{\partial w_m}G(w_n) = 0$$

We can rewrite $\frac{\partial G(w_n)}{\partial w_n}$ as $g(\bullet)$, where $g(\bullet)$ is the density associated with the distribution $G(\bullet)$. Similarly, we can rewrite the partial derivative $\frac{\partial F(w_m)}{\partial w_m}$ as $f(\bullet)$. This eliminates some of the notational complexity in (7).

Under the assumption that $(1-F(\bullet))$ and $G(\bullet)$ are log concave, the solution to the first order condition (7) gives a maximum.[13] Log concavity of $(1-F(\bullet))$ and $G(\bullet)$ is guaranteed by the assumption that the densities $f(\bullet)$ and $g(\bullet)$ are log concave. This is a technical requirement about the shape of these distributions, but it is not a particularly limiting restriction. Most familiar distributions—for example, the uniform distribution, the normal distribution, and the exponential distribution—satisfy it.[14]

Now, it follows from equation (2) 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 (7) 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 (7) as:

---

[13] This is a sufficient but not a necessary condition.

[14] The condition also holds if the distribution is logistic, chi-squared, Laplace, and for some parameters Weibull, gamma, and Beta. *See* Drew Fudenburg & Jean Tirole, Game Theory (MIT Press 1996) at 267.

(8) $$\frac{g(w_n)}{G(w_n)} - \frac{f(w_m)}{1-F(w_m)} = 0$$

The second term in (8) is the hazard rate of $F(\bullet)$. The hazard rate is a familiar notion in the statistics literature; one important property of log concave density functions is that the hazard rate of their distributions is a strictly increasing function.[15] Thus, in equation (8), an increase in w_m will lead to an increase in the value of the second term.

An increase in w_n, by contrast, will lead to a decrease in the first term in equation (8). This follows from the log concavity of $G(\bullet)$. If $G(\bullet)$ is log concave, then the second derivative of log $G(\bullet)$ is negative. The first derivative of log $G(\bullet)$ is $g(\bullet)/G(\bullet)$. For $G(\bullet)$ to be log concave, this first derivative must be decreasing.

The left-hand side of equation (8) defines a function Φ, which depends on the choice variable d^* and (through w_m and w_n) the parameters m, n, k_m, and Δ. Because the first term in equation (8) decreases in w_m and the second term increases in w_n, we know by using (1) and (2) that

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

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

(11) $$\frac{\partial \Phi}{\partial d} \leq 0 ,$$ and

(12) $$\frac{\partial \Phi}{\partial \Delta} \leq 0 .$$

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

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

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

[15] *See* Fudenburg and Tirole at 267.

(15) $$\frac{\partial d^*}{\partial \Delta} = -\frac{\partial \Phi / \partial \Delta}{\partial \Phi / \partial d} \leq 0 .$$

C. A Specific Example

To examine a specific case, suppose that firms *M* and *N* both draw knowledge *k* from the uniform distribution on [0,1]. Mathematically, $F(k) = G(k) = k$. Equation (8) in this example becomes

(16) $$\frac{1}{w_n} - \frac{1}{1-w_m} = 0$$

which, when we substitute in equation (2) for w_m and equation (1) for w_n, simplifies to

(17) $$d^* = \frac{1}{2}(1 - 2p - 2\Delta + m + n + k_m)$$

Figure 2 is a table that shows the optimal level of disclosure (d^*) for various values of parameters *m*, *n*, and Δ. For simplicity, in the table we normalize *p* to zero and set k_m equal to zero as well. The trends suggested by propositions 1 through 3 are confirmed in this simple case, as is the claim that disclosure occurs regardless of the firms' relative positions in the race.

m	***n***	***Δ***	***d****	***Comments***
0.30	0.10	0.40	0.30	Disclosure, with m>n
0.10	0.40	0.70	0.05	Disclosure, with m<n
0.30	0.10	0.50	0.20	Proposition 1 (d^* increases with m)
0.35	0.10	0.50	0.23	
0.40	0.10	0.50	0.25	
0.40	0.10	0.70	0.05	Proposition 2 (d^* increases with n)
0.40	0.20	0.70	0.10	
0.40	0.30	0.70	0.15	
0.40	0.20	0.50	0.30	Proposition 3 (d^* decreases with Δ)
0.40	0.20	0.60	0.20	
0.40	0.20	0.70	0.10	

Figure 2. Numeric example of basic race.

II. Investment in Response to Disclosure

In the model thus far, we have assumed that firm *N* cannot increase the intensity of its research efforts in response to firm *M*'s disclosures. Thus, in the second period, firm *N* gets the same research increment (k_n) no matter what happens in period one. While that framework nicely represents settings where it is difficult to quickly adjust research intensity—for example, settings where it is difficult to identify additional qualified researchers or isolate additional research paths—in this section we refine the model to consider cases where firm *N* can invest in accelerated research. Our goal is to show that disclosure remains a viable strategy even for a firm whose rival is capable of ramping up its research program in response to any disclosure.

To that end, suppose that firm *N* can make an investment, call it *i*, that will enable its researchers to generate better research results on expectation. The extent of the improvement varies with the investment. So, to be specific, if we define the distribution $H(\bullet)$ to be a distribution such that $H(k) \leq G(k)$ for all k,[16] for $i \in [0,1]$ the firm's resulting research distribution will be the weighted average

(18) $$iH(\cdot) + (1-i)G(\cdot).$$

In words, the higher the value of *i*, the more firm *N*'s researchers will behave as if they produce the results in distribution $H(\bullet)$ and the less they will behave as if they produce the results in distribution $G(\bullet)$. If *N* chooses not to invest (i=0), it is left with its original distribution, $G(\bullet)$. Let the cost of this improved research effort be $i^2/2$, a term that increases at an increasing rate and thus captures the intuition that there are diminishing marginal returns to investment.

In period two, *N* now must choose an investment level. Let *V* represent the expected value of the patent. Taking w_n as given, *N* will choose the investment level, i^*, that maximizes the expression

(19) $$\left(1-\left(i(H(w_n))+(1-i)(G(w_n))\right)\right)V-\frac{i^2}{2}$$

where the first term represents the expected gains from research (odds of winning multiplied by the value of the patent) and the second term represents the costs of research. The first-order condition that establishes *N*'s optimal investment is therefore

$$-VH(w_n)+VG(w_n)-i=0$$

which gives

(20) $$i^* = V\left[G(w_n)-H(w_n)\right].$$

[16] Note that $H(\bullet)$ thus represents a superior set of expected research results as compared to $G(\bullet)$; the probability that $k_n \leq k$ is lower for all values of k in distribution $H(\bullet)$ than it is for those same values of k in distribution $G(\bullet)$. This condition is called first order stochastic dominance.

This establishes the relationship between *M*'s disclosure and *N*'s investment. When *M* discloses, w_n increases; w_n, in turn, affects i^* in the manner stipulated by equation (20).

We can now solve for firm *M*'s optimal disclosure strategy given the possibility of investment by firm *N*. In the original model, *M* chose *d* to maximize equation (3):

$$G(w_n)(1-F(w_m))\,.$$

The first term represented the probability that firm *N* would fail to draw the requisite amount of knowledge, w_n, in the second period. The second term represented the probability that firm *M* would successfully draw its requisite amount of knowledge, w_m, in the third period.

This time, by contrast, the first term must account for the fact that firm *N* is not working with the distribution $G(\cdot)$, but is instead working with the weighted average of $G(\cdot)$ and $H(\cdot)$, evaluated at the optimal investment point, i^*. Making the appropriate substitutions, we see that *M* is now choosing *d* to maximize

$$[i^* H(w_n)+(1-i^*)G(w_n)][1-F(w_m)] \qquad (21)$$

subject to the constraints (a) $d \geq 0$; (b) $d \leq m + k_m - p$; and (c) equation (20).

A firm can benefit from disclosure even in settings where its rival is capable of responding to disclosure by investing in accelerated research. In particular, we can now state and prove the following proposition. As before, we explain the intuition and also work through a simple example immediately below.

> **Proposition 4.** In cases where the rival's research falls sufficiently short of the patentability threshold (that is, $p+\Delta-n$ sufficiently large), the possibility of an investment response will always lead to more disclosure than would occur were no response possible.

A. Intuitions

The ability to invest makes the rival a more formidable research opponent. Thus, just as an increase in the rival's knowledge would lead to increased disclosure, one rightly expects that the possibility of an investment response will lead to increased disclosure as well. The additional complexity here, however, is that the rival might react to that increased disclosure by further increasing its research investment. There are thus two competing factors at play. On the one hand, the disclosing firm wants to disclose more information to compensate for the fact that it is playing against a more capable rival. On the other hand, the disclosing firm needs to be careful not to change the payoff structure too much, else the rival might invest even more and in that way begin to undermine the benefits of disclosure.

The proposition can be explained by reference to these two competing effects. To begin, note that disclosure remains a viable strategy even in the face of a possible investment response. This follows directly from the competing factors outlined above. In cases where the "better player" effect dominates, the possibility of an investment response will make disclosure even more attractive than it would otherwise be. Moreover, even in cases where the rival's redoubled efforts are significant, disclosure can still sometimes be attractive; the factors only tell us that, in such cases, disclosure will be less attractive than it would be in the absence of an investment response.

Proposition 4 predicts increased disclosure in cases where $p+\Delta-n$ exceeds some case-specific lower bound. The logic here is as follows. The first effect—the "better player" effect—always pushes toward more disclosure; so the only thing that needs to be explained is why, in this case, the second factor is weak. This follows easily. After all, as $p+\Delta-n$ grows, the rival increasingly knows that its odds of winning the patent will remain low even if it marginally increases its research investment. Thus, in cases where $p+\Delta-n$ is relatively large, the rival has little reason to answer disclosure with further investment; the expected benefit is small relative to the additional expense. In these cases, then, the first factor dominates the second, and we can safely predict increased disclosure.

B. Proofs

Let us define d^{**} to be the optimal level of disclosure from firm *M*'s perspective given the possibility of an investment response by firm *N*. Proposition 4 can be established by showing that, for fixed *m* and k_m, there exists some lower bound, call it μ, such that $d^{**}>d^*$ for $p+\Delta-n > \mu$.

Using (20), equation (21) can be rewritten as

$$\text{(22)} \qquad [G(w_n)-\frac{i^{*2}}{V}][1-F(w_m)]\,.$$

While *d* is not explicitly shown in this equation, every term here is a function of *d*. Specifically: the knowledge threshold w_n is a function of *d* as is established by equation (1); the knowledge threshold w_m is a function of *d* according to equation (2); and the optimal investment i^* is a function of *d*, as defined most clearly in equation (20).

We can now solve for the d^{**} that maximizes equation (18) subject to the constraints (a) $d \geq 0$; (b) $d \leq m+k_m-p$; and (c) equation (20). For an interior solution, the first-order condition necessary for maximization is

$$\text{(23)} \qquad g(w_n)[1-F(w_m)]-f(w_m)G(w_n)+f(w_m)\frac{i^{*2}}{V}-\frac{2i^*}{V}[1-F(w_m)]\frac{\partial i^*}{\partial d}=0\,.$$

The left-hand side of equation (23) defines a function, Ψ, which depends on the choice variable d^{**} and (through w_m, w_n, and i^*) model parameters p, Δ, m, n and k_m. The second-order condition for maximization at an interior solution requires

$$\text{(24)} \qquad \frac{\partial \Psi}{\partial d} \leq 0\,.$$

Now return to equation (20), the equation that establishes the relationship between *N*'s optimal research investment and *M*'s optimal disclosure. Taking the partial derivative of (20) with respect to w_n yields

(25) $$\frac{\partial i^*}{\partial w_n} = \frac{\partial i^*}{\partial d} = V\left[g(w_n) - h(w_n)\right]$$

which allows us to rewrite (23) as follows

(26) $$\frac{g(w_n)}{G(w_n)} - \frac{f(w_m)}{1-F(w_m)} = -\frac{V\left[G(w_n)-H(w_n)\right]^2}{G(w_n)}\left(\frac{f(w_m)}{1-F(w_m)} - 2\frac{g(w_n)-h(w_n)}{G(w_n)-H(w_n)}\right).$$

Using that, we can now establish proposition 4. The left-hand side of (26) coincides with the left-hand side of equation (8). At $d=d^*$, then, the left-hand side is equal to zero. From equation (11), we know that if the right-hand side of (26) is less than zero, the value of d that satisfies (26) will be $d>d^*$ and, thus, $d^{**}>d^*$. Conversely, we know that if the right-hand side of (26) is greater than zero, the value of d that satisfies (26) will be $d<d^*$ and, thus, $d^{**}<d^*$. Signing the various terms, the right-hand side of (26) will certainly be less than zero if $g(w_n)–h(w_n)<0$, so to prove the proposition all we need to show is that there exists a $p+\Delta–n = \mu$ such that, for $p+\Delta–n > \mu$, $g(w_n)–h(w_n)<0$.

That, it turns out, is easy to do—even without assuming any specific distributional form for $G(\bullet)$ and $H(\bullet)$. After all, given stochastic dominance, it is always true that $g(w_n)–h(w_n)>0$ for low values of w_n and $g(w_n)–h(w_n)<0$ for high values of w_n. The size of these ranges will vary depending on the exact distributions at issue. For example, if $G(\bullet)$ is the uniform distribution on [0,1] and $H(\bullet)$ is the quadratic distribution, also on [0,1], $g(k)–h(k)<0$ for all $k>0.5$. For all distributions, however, there must exist some μ above which we are in the appropriately high end of the range and, thus, $g(w_n)–h(w_n)<0$.

C. A Specific Example

To examine a specific case, let us work out the example referenced just above: without investment, firms *M* and *N* draw knowledge *k* from the uniform distribution on [0,1], but, with investment, firm *N* can make its knowledge draw more closely resemble the quadratic distribution, also on [0,1]. Mathematically, $F(k) = G(k) = k$, and $H(k) = k^2$.

If we normalize *V* to 1, equations (20), (25), and (26) become, respectively:

(27) $$i^* = w_n - {w_n}^2$$

(28) $$\frac{\partial i^*}{\partial d} = 1 - 2w_n,$$ and

(29) $$\frac{1}{w_n} - \frac{1}{1-w_m} = -\frac{\left(w_n - {w_n}^2\right)^2}{w_n}\left(\frac{1}{1-w_m} - 2\frac{1-2w_n}{w_n - {w_n}^2}\right).$$

Solving (29), we find the results shown in Figure 3. Note that, to facilitate comparisons with the earlier numeric example, here again we set p=0 and k_m=0.

m	***n***	***Δ***	d^*	d^{**}	***Comments***
0.40	0.10	0.60	0.15	0.22	It can be that $d^{**}>d^*$
0.10	0.70	0.80	0.10	0.09	It can be that $d^{**}<d^*$
0.10	0.70	0.90	0.00	0.00	It can be that $d^{**}=0$
0.40	0.10	0.70	0.05	0.12	Proposition 4 (large p+Δ–n means $d^{**}>d^*$)
0.50	0.30	0.70	0.20	0.26	Proposition 4 (large p+Δ–n means $d^{**}>d^*$)

Figure 3. Numeric example of basic race with possibility of rival investment response.

III. Investment by the Disclosing Firm

In the previous section, we considered the possibility of an investment response by the rival firm. Our purpose there was to show that disclosure remains a viable strategy even in settings where the rival can accelerate its research. In this section, we consider instead investment by the disclosing firm. Our purpose now is to show how the disclosing firm can use investment to offset some of the costs of disclosure. That is, we set out to consider cases where a firm discloses in order to reduce its rival's chance of winning the race and then makes up for that disclosure by increasing its own research intensity.

To begin, return to the baseline disclosure model and allow firm *M* to invest in the third period. Assume that firm *M*'s third period investment has all the same characteristics as firm *N*'s investment in Part II. In particular, *M* makes an investment, *i*, in period three that allows *M* to draw from the weighted average of two distributions: its original distribution, *F*(•), and a significantly improved distribution, *H*(•). The resulting research distribution is

(30) $$iH(\cdot)+(1-i)F(\cdot)\,.$$

As before, a unit of investment costs $i^2/2$. By the same logic as in the previous proof, then,

(31) $$i^*=V\left[F(w_m)-H(w_m)\right] \text{ and}$$

(32) $$\frac{\partial i^*}{\partial w_m}=\frac{\partial i^*}{\partial d}=V\left[f(w_m)-h(w_m)\right].$$

The disclosure decision in this version of the model, however, is different. After all, now *M* must take into account the availability of its own investment response when deciding how much to disclose. *M* thus chooses *d* to maximize

$$[G(w_n)][1-(i^* H(w_m)+(1-i^*)F(w_m))] \tag{33}$$

subject to the standard constraints: (a) $d \geq 0$; (b) $d \leq m+k_m-p$; and (c) equation (31).

We can now state two propositions regarding how a firm's disclosure decision changes given the ability to make up for any disclosures through additional research investments. As in the previous sections, we explain intuitions and also work through a simple example immediately below.

> **Proposition 5.** In cases where the rival's research is sufficiently close to the patentability threshold (that is, $p+\Delta-n$ sufficiently small), the disclosing firm will always disclose more and then use increased investment to compensate.

> **Proposition 6.** In cases where the disclosing firm's own research falls sufficiently short of the patentability threshold (that is, $p+\Delta-m-k_m$ sufficiently large), the possibility of increased investment will not lead to increased disclosure.

A. Intuitions

The ability to invest puts the disclosing firm in a stronger position. It can disclose more information and then make up for it by increasing its own research efforts. Thus, we expect that in many instances the possibility of increased investment will lead to increased disclosure as well.

Propositions 5 and 6 examine two specific cases. In proposition 5, we isolate cases where $p+\Delta-n$ is relatively low and, hence, the rival is very likely to win the patent. These are cases where the pressure to disclose is highest, and so—while the specific cutoff value varies case-by-case—in these cases there will always be some value for $p+\Delta-n$ below which the disclosing firm will choose to increase disclosure and make up for it with increased investment. In proposition 6, by contrast, we isolate cases where $p+\Delta-m-k_m$ exceeds some case-specific upper bound and, hence, the disclosing firm is unlikely to win the patent regardless of what its rival does. Not surprisingly, in these cases the disclosing firm chooses not to aggressively invest in additional research. The costs of doing so exceed the expected benefit, namely the still unlikely event of winning the patent.[17]

---

[17] As we show in the proof that follows, in these cases we actually see a decrease in the amount disclosed. The logic here is as follows. The cost of disclosure from the disclosing firm's perspective is the decrease in the disclosing firm's likelihood of winning the patent in the third period. In cases where $p+\Delta-m-k_m$ is large and investment is not possible, this cost is virtually zero since, no matter what, the disclosing firm is not going to win. In cases where $p+\Delta-m-k_m$ is large and investment is possible, however, the cost is modestly larger; the disclosing firm is still terribly unlikely to win, but nevertheless

B. Proofs

The proofs here are analogous to the proofs presented in Part II. Specifically, solving (33) at an interior solution yields:

$$g(w_n)[1-F(w_m)]-f(w_m)G(w_n)+g(w_n)\frac{i^{*2}}{V}+\frac{2i^*}{V}G(w_n)\frac{\partial i^*}{\partial d}=0\,. \tag{34}$$

The left-hand side of (34) defines a function Θ which, again, depends on the choice variable d^{**} and (through w_m, w_n, and i^*) model parameters p, Δ, m, k_m, and n. The second-order condition for maximization at an interior solution requires

$$\frac{\partial \Theta}{\partial d}\le 0\,. \tag{35}$$

Using (31) and (32) we can rewrite (34) as follows:

$$\frac{g(w_n)}{G(w_n)}-\frac{f(w_m)}{1-F(w_m)}=-\frac{V[F(w_m)-H(w_m)]^2}{1-F(w_m)}\left(\frac{g(w_n)}{G(w_n)}+2\frac{f(w_m)-h(w_m)}{F(w_m)-H(w_m)}\right). \tag{36}$$

To prove propositions 5 and 6, note (again) that the left-hand side of (36) coincides with the left-hand side of (8) (and (26)) and, by (11), it decreases with d. Propositions 5 and 6 follow because, as before, $f(w_m)$–$h(w_m)$ is positive for low values of w_m and negative for high values of w_m. More precisely, proposition 5 holds because, for fixed m and k_m, there exists a $p+\Delta-n = \mu_1$ such that, for $p+\Delta-n < \mu_1$, the right-hand side of (36) is negative (since $g(w_n)/G(w_n)$ is large) and, hence, $d^{**}>d^*$. Proposition 6 holds because for a fixed n there exists a $p+\Delta-m-k_m = \mu_2$ such that, for $p+\Delta-m-k_m >\mu_2$, the right-hand side of (36) is positive and, hence, $d^{**}<d^*$.

C. A Specific Example

As in Part II, suppose that, without investment, firms M and N draw knowledge k from the uniform distribution on [0,1], but that, with investment, firm M can make its knowledge draw more closely resemble the quadratic distribution, also on [0,1]. That is, $F(k) = G(k) = k$, and $H(k) = k^2$.

Normalizing V to 1, equations (31), (32), and (36) become, respectively:

$$i^* = w_m - w_m{}^2\,, \tag{37}$$

$$\frac{\partial i^*}{\partial d}=1-2w_m\,,\ \text{and} \tag{38}$$

---

the disclosing firm is more likely to win than it otherwise was. Hence, disclosure becomes more costly and we therefore expect less disclosure.

(39) $$\frac{1}{w_n} - \frac{1}{1-w_m} = -\frac{\left(w_m - w_m^{\ 2}\right)^2}{1-w_m}\left(\frac{1}{w_n} + 2\frac{1-2w_m}{w_m - w_m^{\ 2}}\right).$$

Solving, we find the results shown in Figure 4. Note that, to facilitate comparisons with the two earlier numeric examples, here again we set p=0 and k_m=0.

m	***n***	***Δ***	d^*	d^{**}	***Comments***
0.60	0.60	0.70	0.40	0.42	It can be that $d^{**}>d^*$
0.10	0.70	0.80	0.10	0.09	It can be that $d^{**}<d^*$
0.10	0.70	0.90	0.00	0.00	It can be that $d^{**}=0$
0.50	0.64	0.70	0.37	0.38	Proposition 5 (small p+Δ–n means $d^{**}>d^*$)
0.70	0.80	0.90	0.35	0.36	Proposition 5 (small p+Δ–n means $d^{**}>d^*$)
0.10	0.80	0.90	0.05	0.04	Proposition 6 (large $p+\Delta-m-k_m$, thus $d^{**}<d^*$)
0.20	0.70	0.90	0.05	0.04	Proposition 6 (large $p+\Delta-m-k_m$, thus $d^{**}<d^*$)

Figure 4. Numeric example of basic race with possibility of investment by disclosing firm.

IV. Limitations

The model developed in the previous three sections was designed to elucidate the basic incentive to disclose and, in addition, certain relationships between disclosure and investment. Like any model, however, ours is only a proxy for the relevant real-world interactions. In this section, we therefore discuss some of the wrinkles not addressed in the modeling and offer our intuitions about how those wrinkles likely affect the analysis.

First, the model abstracts away from issues related to uncertainty. In the real world, firms probably have only imperfect estimates as to the identities and accomplishments of their rivals, and so firms have to base their strategies on educated guesses as opposed to the more exact parameters used in the model. Note that this is true for all of a research firm's competitive decisions—decisions about what projects to pursue, what resources to invest, and so on—and it is true for the firm's decisions with respect to disclosure as well. Uncertainty probably explains why, in practice, one sees less disclosure than our model predicts. Uncertainty exposes the firm to the risk of giving the race away by disclosing too much, and that would tend to diminish a firm's willingness to attempt the strategy.

Introducing uncertainty into our analysis would have had another effect on our work: it would have allowed us to consider disclosure's possible use as a signal. After all, in a world of uncertainty, a firm's disclosures give its rivals some information as to the firm's relative position in the race. That might be a reason to disclose; for instance, a race leader might disclose some fraction of its information as a way of signaling its position and in that way discouraging trailing rivals from continuing to race. That might also be a reason not to disclose, for example if a race leader wanted to keep its position secret and thereby encourage rivals to waste their resources competing in a race they are likely to lose. We do not pursue these sorts of stories in this work, in part because disclosure is just one of several possible signaling mechanisms available to a research firm and thus this line of analysis seemed unlikely to provide a convincing explanation for disclosures like those made by IBM and Xerox.[18]

Second, our model does not allow firms to negotiate. 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 our model does not allow private negotiations is that we think such negotiations are extremely unlikely. For one thing, any such negotiations would take place under extraordinary 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 by contract exactly what the rival is supposed to do. Were the disclosing firm to actually disclose, its rival's behavior would be straightforward: the disclosure would change the structure of the patent race, and the rival's incentive to continue in the race would be correspondingly diminished. The firm might drop out, compete less vigorously, and so on. To simulate that result by contract, however, the firms would have to agree on exactly what the non-disclosing firm can and cannot do, all the while negotiating in a context where the technology at issue is still under development and thus poorly defined. Monitoring compliance with such a contract would likely prove both difficult and expensive, and that—taken together with the difficulties inherent in drafting the contract, the standard negotiating impediments and time pressures discussed above, and the possibility that any deal would ultimately be deemed illegal under either antitrust or patent misuse principles—in the end makes private negotiation unlikely, even though it could in theory make both disclosing and non-disclosing firms better off.

Third, in this model we have assumed that patent races proceed linearly; that is, we imagined a world where racing firms basically pursue the same research path and thus a firm with more information than its rival knows everything its rival knows plus more. 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 obviously make disclosure more difficult. For

[18] Lichtman, Baker & Krauss (2000) analyzes the possible use of disclosure as a signal.

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. Like uncertainty, then, more complex assumptions about the research process likely would lead us to expect less disclosure than the current model suggests.

Finally, like most patent race models, ours focuses on only two firms. There is no reason to believe that this in any way distorts the results. Additional firms increase the need for disclosure in many cases, but additional firms also increase the odds that a given disclosure will inadvertently assist a rival. How these various multipliers affect the overall incentive will vary case by case. The only clear implication from considering multiple firms is that having multiple firms in the race makes it even less likely that the firms will be able to negotiate a private agreement that avoids disclosure. Multiple firms, after all, increase the likelihood of holdout and free-rider problems.

V. Empirical Evidence

In the hope of finding some empirical support for our basic argument that firms publicly disclose research information even in the context of on-going 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 Machines" was the patent assignee.[19] And, examining those electronically, we found that 2,310 of them cite as prior art at least one article from the *IBM Technical Disclosure Bulletin*.[20] So, in approximately one out of every six patents issued to IBM during this time period, IBM's own disclosures seem to have come back to limit the firm's later patent claims.

Standing on its own, this data did not necessarily support our thesis. After all, one could interpret these numbers not as evidence of the theory presented in this paper, 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. To distinguish this possible explanation, we analyzed each of the 2,310 patents at issue and calculated the number of months that passed between the date of the relevant publication[21] and the date on which the patent application was ultimately filed.[22] Long gaps would support the

---

[19] 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.

[20] 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.

[21] In the event of multiple citations, we used the date of the most recent one.

[22] After analyzing over 200 patents by hand, we were able to automate the process such that the computer would parse the full text of each patent and identify for us the relevant dates. For most patents, the necessary dates were easy to identify since filing dates are well labeled on all issued patents, and prior

idea of changed research priorities. Short gaps, by contrast, would suggest that IBM was knowingly disclosing information about on-going patent races.

For ten of the patents, there was insufficient information to make this calculation.[23] For the remaining 2,300 patents, however, we obtained the results shown below. Note that we rounded the tallies such 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 had so few data points that they are not shown in the graphic but are accounted for in the table.

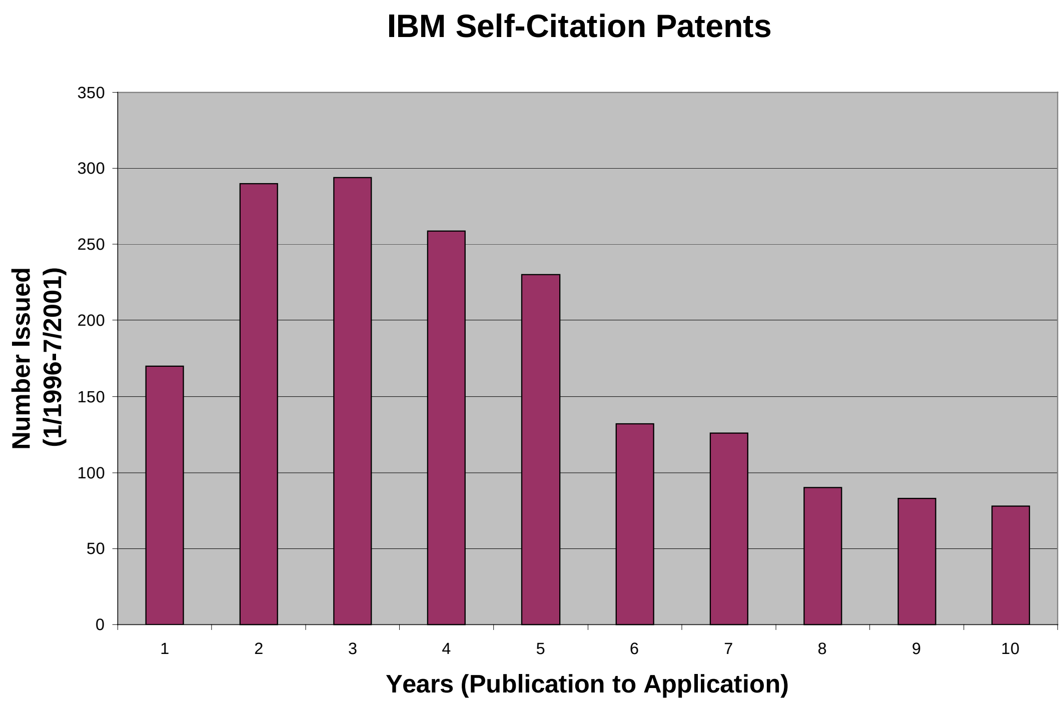

art cited in a patent typically also includes a well-labeled reference to the date of publication. Any patents that the computer deemed unclear were processed manually, and all results were spot-checked for accuracy.

[23] 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.

Years (Publication to Application)	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 %

Over 50% of the patents cite *Bulletin* articles published less than 5 years before.

Several interesting trends are revealed from the data, but the most important is simply this: in the bulk of the cases, the gap between IBM's publication and IBM's patent application is small. Since it would be surprising if IBM's research agenda were to change that quickly in so many cases, this seems to lend strong support to the theory underlying the paper. IBM is not simply disclosing information relevant to races it plans to abandon. IBM is often disclosing information about patents that it is actively pursuing.

VI. Conclusion

Our primary focus in this paper has been to offer a richer account of why research firms disclose information to the public. There are three traditional explanations: that disclosure serves to reward employees; that it generates positive publicity for the firm; and that it is a defensive action designed to stop rivals from patenting in instances where the disclosing firm itself does not intend to patent. We propose here a fourth 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 by the traditional theories.

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.[24]

The threat of public 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. More broadly, then, one 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.[25]

A second interesting extension would be to think about strategic disclosure's role in bringing prior art to the attention of patent examiners. Recent patents issued to firms like Amazon (the "one-click" patent[26]) and Priceline (the reverse auction mechanism[27]) have led many commentators to worry that the Patent and Trademark Office (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.

To address this problem, scholars and activists have advocated a variety of approaches. Greg Aharonian and Internet startup ip.com are both working to create private stores of prior art that might then be made available to PTO examiners.[28] Jay Kesan and Marc Banik have recently argued that patents should be given more favorable presumptions of validity in exchange for more complete prior art disclosures.[29] Jay Thomas has recently argued that parties should be given financial rewards for bringing forward prior art that defeats in-process or recently issued patent applications; and, in fact, Internet startup bountyquest.com has begun to do just that.[30] The arguments sketched in this paper suggest that, in addition to the three mechanisms outlined above, a possible solution to the difficulty associated with high technology patents would be to

---

[24] 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).

[25] Two of the current authors have begun to pursue this extension; *see* Baker & Mezzetti (2001).

[26] U.S. Patent No. 5,960,411.

[27] U.S. Patent No. 5,794,207.

[28] Mr. Aharonian's prior art archive is currently available at http://www.bustpatents.com.

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

[30] John Thomas, Collusion and Collective Action in the Patent System: A Proposal for Patent Bounties (2000) (on file with authors).

adjust the prior art inquiry so as to better harness the strategic incentive for prior art revelation outside the formal patent prosecution process. If incentives to disclose could be tweaked—most obviously through adjustments to the threshold of patentability—strategic disclosure could serve to make more prior art available to the PTO.

References

James J. Anton & Dennis Yao, Patents, Invalidity, and the Strategic Transmission of Information (working paper, 1999) (available from authors).

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

Rebecca Eisenberg, The Promise and Perils of Strategic Prior Art Creation Through Publication, 98 Michigan Law Review (2000).

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

Douglas Lichtman, Scott Baker & Kate Kraus, Strategic Disclosure in the Patent System, 53 Vanderbilt Law Review 2175 (2000).

Gideon Parchomovsky, Publish or Perish, 98 Michigan Law Review 926 (2000).

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