Patent Signaling

Deepak Hegde
NYU & NBER

Boyan Jovanovic
NYU & NBER

Jin Liu
UPenn

March 2026

Abstract

Patent policy assumes disclosure facilitates imitation. Yet inventors often disclose even when legal protections are weak. We develop a model where, by disclosing earlier, a firm signals a greater commercialization efficiency and thereby deters entry of imitators. Using the American Inventors' Protection Act of 1999, which allowed accelerated patent disclosures, we show: (i) inventors disclose higher-quality inventions sooner; and (ii) disclosure is faster in fields where rivals learn slowly. Event-study estimates reveal that up to 50 percent of stock-market returns linked to patenting reflect informational gains from disclosure, not legal rights.

Keywords: Patent Publication, Invention Disclosure, Signaling, Innovation, Competition
JEL Codes: E02, D23, G24, L26, O34

We are grateful to Bernhard Ganglmair for data on product versus process inventions; Bronwyn Hall, Hugo Hopenhayn, Hong Luo, Virgiliu Midrigan, and Dimitris Papanikolaou for helpful comments. For any questions or comments, please reach out to us at dh99@stern.nyu.edu, bj2@nyu.edu, and jinliu52@upenn.edu. 19 West 4th Street, 6th Floor. New York, NY 10012

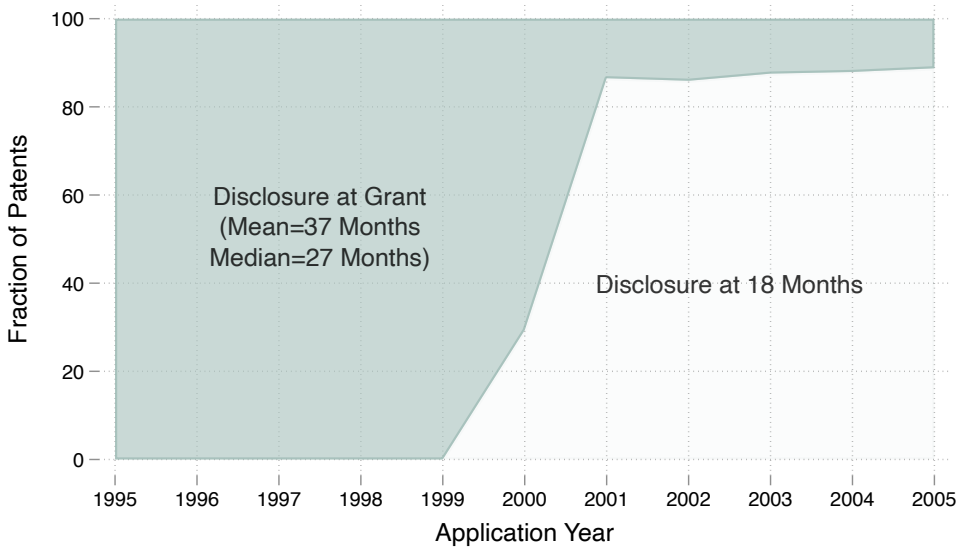
I Introduction

At least since Arrow (1962), economists have recognized that disclosing intellectual property (IP) can be costly: it enables imitation by competitors and expropriation by partners (*e.g.,* Horstmann, MacDonald and Slivinski, 1985; Modigliani et al., 1999; Anton and Yao, 2004). Yet in practice, inventors routinely disclose their inventions earlier than necessary. Following the American Inventors' Protection Act (AIPA) of 1999, inventors could choose to publish patent applications either 18 months after filing or at patent grant—typically about 37 months after filing. Despite this choice, 91 % of U.S. applicants opted for early disclosure at 18 months. Even among inventors who forgo foreign patent protection (and face no legal requirement to disclose early), 87.3% still disclosed at 18 months (Figure 1).

Why do inventors reveal their ideas early despite imitation risks? This paper proposes a simple answer: under competition, disclosure can act as a signal of invention quality. We build a model where inventors privately know the quality of their ideas. Rivals attempt to infer this quality before deciding whether to imitate or innovate around the disclosed invention. Disclosure reveals the invention's quality, but high-quality inventors choose to disclose early because they can still profit after rivals' entry; high quality inventions imply lower profits for rivals and thus stronger deterrence. Low-quality inventors, by contrast, face worse outcomes after disclosure and prefer to delay. In equilibrium, disclosure timing signals invention quality.

The model delivers two key empirical predictions: (1) Higher-quality inventions are disclosed sooner; (2) Early disclosure is more likely when competitors learn slowly from disclosed information. We test these predictions using inventors' disclosure decisions after AIPA. Our empirical strategy exploits variation both at the extensive margin (whether inventors disclose early) and at the intensive margin (the delay between application and disclosure among those who wait for grant). Because inventors can affect grant timing through prosecution tactics (e.g., by delaying responses to the patent office), even post-18-month disclosure timing tied to patent grant is informative about strategic disclosure behavior.

Figure 1: Publication Lag Since Application Date



Notes: This figure plots the fraction of USPTO patent applications that are published at grant, versus those that are published at 18 months after application date. The figure uses the population of regular utility patent applications filed at the USPTO between 1995 and 2005. Patents with foreign priorities are excluded. All U.S. patent applications were published at grant prior to the American Inventors' Protection Act (AIPA). AIPA provided the choice of 18 month disclosure to patent applications filed on or after November 29, 2000.

Consistent with the model, we find that high-quality inventions—measured by forward citations and claim counts—are disclosed sooner. Moreover, inventors' disclosure behavior appears to differ sharply across technology: in fields like pharmaceuticals and chemistry, where rivals are likely to learn gradually from disclosures (we call these "gradual learning" fields), patent disclosures tend to occur more quickly; in contrast, in fields like software and information technology, where rivals can learn relatively quickly from disclosures (we call these "Eureka learning" fields), incumbents tend to delay patent disclosures more. Overall, patent disclosure timing across fields seem to correlate with the patterns predicted by our theory. While signaling is the primary mechanism we study, alternative motivations like attracting external financing could also encourage early disclosure. However, unlike signaling, external financing motives do not predict that higher-quality inventions should systematically disclose sooner—and our regressions control for observable proxies of external finance needs (e.g., firm size, capital intensity).

To further validate the signaling interpretation, we conduct an event study which measures

stock-market response to the patent disclosures of U.S. publicly listed firms. We find that patent disclosure triggers a sharp stock-market response, suggesting that information about invention quality was private prior to disclosure. Moreover, earlier patent disclosures are associated with larger positive stock market returns, consistent with higher-quality inventions revealing themselves sooner.

Our findings advance innovation theory by integrating disclosure into a competitive signaling model, bridging insights from (Arrow, 1972) and (Spence, 1973). Unlike prior economic models that treat disclosure purely as a cost (e.g., Bhattacharya, Glazer and Sappington, 1992; Anton and Yao, 1994, 2004; Bhattacharya and Guriev, 2006), we show that disclosure can credibly separate inventors by quality. Accordingly, we formalize and test the conjectures made by (Long, 2002) . Our analysis differs from models of patent races where firms disclose early to preempt rivals (e.g., Hopenhayn and Squintani, 2011). Here, secrecy—and the timing of disclosure—serves as a strategic tool for signaling, not just preemption.

Our findings complement empirical work documenting the social benefits of disclosure (e.g., Johnson and Popp, 2003; Furman, Nagler and Watzinger, 2021; Lück et al., 2020; Gross et al., 2019; Ganglmair and Reimers, 2019; Hegde, Herkenhoff and Zhu, 2023) by demonstrating its private benefits to inventors. Specifically, our stock-market event study extends previous influential work by (Kogan et al., 2017), and teases apart market response to invention disclosures, from market response to patent grants. We demonstrate that early patent disclosures reduce information asymmetry sooner by informing rivals and investors sooner, and benefit inventors of high-quality inventions.

The rest of the paper proceeds as follows. Section II develops our theoretical model of invention disclosure under competition. Section III describes the data. Section IV presents empirical tests and estimates model parameters. Section VI offers concluding thoughts.

II Model

Consider a two-stage game between an inventor and a potential imitator, both competing for the same market. The inventor, referred to as the "incumbent," has a type α, representing the output generated by his idea, which we refer to as the idea's "quality." The incumbent's only decision is how early to disclose the idea through the patent application.

The imitator is a potential entrant, or "challenger," who has a type c, representing his cost of entering the market should he decide to do so. The parameter c is privately known by the challenger, while the incumbent only knows that it is drawn from the cumulative distribution function $F(c)$. This is similar to the setup in Milgrom and Roberts (1982). Likewise, the parameter α is privately known by the incumbent. The challenger only knows that it is drawn from the cumulative distribution function $G(\alpha)$.

We denote by δ the measure of how well the challenger can copy, develop, and possibly extend the incumbent's idea if given x time periods to do so. It is measured in units of output that the challenger can produce, and δ is increasing in x, or the amount of time that the challenger has to work on the idea. In practice, x can be thought of as the time lag between the incumbent's patent publication and product market entry or commercialization.

The patent application is filed at date 0, and $t \geq 0$ is the disclosure, or "publication date," measured as the time elapsed since the application. Production starts at date T, at which point the challenger decides whether or not to enter. If he does enter, his output will be $\delta(x)$, where

$$x = T - t. \tag{1}$$

The incumbent discloses at date t. However, it is more convenient to use $x = T - t$ as the decision variable. Let $\tau : A \to [0, T]$ be the policy function, i.e., $x = \tau(\alpha)$ is optimal for an incumbent of type α. If a higher α means that disclosure happens earlier, then $\tau(\alpha)$ is an increasing function.

The copying function δ is increasing in x. Therefore, if α were common knowledge, it

would be optimal to set $t = T$, making $x = 0$ and δ as small as possible. Since δ does not depend on α, all patents are equally helpful to the challenger. This is a simplifying assumption that helps us solve the model by hand.

Demand and Production. Industry output, Q, is homogeneous and given by

$$Q = \begin{cases} \alpha & \text{if the challenger stays out,} \\ \alpha + \delta(x) & \text{if he comes in.} \end{cases}$$

The demand curve is specified as,

$$p = \begin{cases} D(\alpha) & \text{if the challenger stays out,} \\ D(\alpha + \delta(x)) & \text{if he comes in.} \end{cases}$$

Production costs are zero and the incumbent's payoff consists of sales revenues alone:

$$\text{Incumbent's payoff} = \begin{cases} \alpha D(\alpha) & \text{if the challenger does not come in,} \\ \alpha D(\alpha + \delta(x)) & \text{if he does.} \end{cases}$$

Challenger's Entry Decision. We denote by $\hat{\alpha}$ the challenger's belief about α, which is a real number in the separating equilibrium. The challenger's payoff is given by:

$$\text{Challenger's payoff} = \begin{cases} \delta(x) D[\hat{\alpha} + \delta(x)] - c & \text{if he enters,} \\ 0 & \text{if he stays out.} \end{cases} . \tag{2}$$

The challenger enters if the first line of (2) is positive. The probability of entry, ω, then is given by:

$$\omega(x, \hat{\alpha}) = F(\delta(x) D[\hat{\alpha} + \delta(x)]) . \tag{3}$$

Incumbent's policy. Given $(\alpha, \hat{\alpha}, x)$, the incumbent's payoff is

$$U(\alpha, \hat{\alpha}, x) = \omega(x, \hat{\alpha})\, \alpha D(\alpha + \delta(x)) + (1 - \omega(x, \hat{\alpha}))\, \alpha D(\alpha). \tag{4}$$

However, in a separating equilibrium, beliefs $\hat{\alpha}$ depend on the observed decision x. Thus, if the optimal decision $\tau(\alpha)$ is unique, strictly monotone, and hence invertible, then $\hat{\alpha} = \tau^{-1}(x)$. Equilibrium then requires that

$$\tau(\alpha) = \arg\max_x U\left(\alpha, \tau^{-1}(x), x\right) \quad \forall \alpha, \tag{5}$$

and the associated first-order condition is

$$U_2(\alpha, \hat{\alpha}, x)\frac{d\hat{\alpha}}{dx} + U_3(\alpha, \hat{\alpha}, x) = 0 \tag{6}$$

II.A Isoelastic demand example

We now describe the example that we shall solve and estimate. Suppose that (1) incumbent types fall within the interval $\alpha \in [\alpha_{\min}, \alpha_{\max}]$; (2) the product demand is isoelastic, i.e. $D(Q) = Q^{-\beta}$; (3) the entry cost is uniformly distributed, such that $F(c) = c$ for $c \in [0, 1]$; and (4) the challenger learns linearly as a function of time x before production starts, expressed as $\delta(x) = a + bx$, where $a \geq 0$ measures instant learning or "Eureka" learning, and $b \geq 0$ measures incremental learning.

The Appendix proves the following result:

Proposition 1. *If* $\beta = 1$ *and*

$$\frac{\alpha_{\max}}{\alpha_{\min}} < \left(1 + \frac{b}{a}T\right)^2, \tag{7}$$

then the incumbent's posting decision has the solution

$$\alpha(x) = \alpha_{\min}\left(1 + \frac{b}{a}x\right)^2 \tag{8}$$

for $x \in \left[0, \frac{a}{b}\left(\sqrt{\frac{\alpha_{\max}}{\alpha_{\min}}} - 1\right)\right]$. *Inverting (8) yields*

$$x = \frac{a}{b}\sqrt{\frac{\alpha}{\alpha_{\min}} - 1} \tag{9}$$

We derive some properties of the equilibrium signal in (9):

1. *Entry deterrence.*—By observing x, the challenger infers that his post-entry revenue would be

$$R_E = (a + bx)\, Q^{-1} = \frac{a + bx}{\alpha + a + bx} = \frac{1}{1 + \frac{\alpha_{\min}}{a^2}(a + bx)}. \tag{10}$$

The last equality follows after plugging in for α from equation (8), and dividing top and bottom of the $\alpha_{\min}\left(1 + \frac{b}{a}x\right)^2$. Thus a higher x lowers an entrant's post-entry revenue, and deters entry for a larger fraction of realizations of the entry cost.

2. *Information externality.*—If all we have data on are measures of patent quality, α, and information about the disclosure lead time x, then we cannot a and b separately, only their ratio

$$\frac{b}{a} = \frac{\text{Incremental learning}}{\text{Eureka learning}}.$$

Section II.D discusses how policy should react to the knowledge spillover.

3. *Strength of competition.*—While keeping the distribution of the entry cost c uniform, suppose now that $c \in \left[0, \frac{1}{\lambda}\right]$ so that its CDF is $F(c) = \lambda c$ for $c \in \left[0, \frac{1}{\lambda}\right]$. The higher is λ, the lower are the competitor's entry costs so λ indexes the threat of entry by the sole potential challenger. Then the Appendix shows that when $\beta = 1$, x is still given by (9) but that the entry probability is proportional to λ:

$$\omega(x, \alpha) = \min\left(1, \lambda\delta(x)\, D\left[\alpha + \delta(x)\right]\right). \tag{11}$$

II.B Public information and the challenger's beliefs

The challenger's prior beliefs before observing x are described by the single parameter $\alpha_{\min}$, which represents the smallest possible incumbent output. A large firm would be expected to produce more with any fixed quality of patent, and that would be reflected in a larger $\alpha_{\min}$. But that does not necessarily mean that the parameters a and b are larger—as Arora et al. (2023) argue, "average invention value increases with firm size even though average invention quality declines because larger firms have a greater capacity to capture value from their inventions." If true, this statement suggests that when comparing large and small firms' inventions, α relates negatively or not at all to the parameters a and b.

II.C Stock-price response to the publication gap t

The publication gap is the waiting time, t, between application date and publication. Publishing at t gives the challenger a lead time of $x = T - t$.

A type-α incumbent has an equilibrium expected utility $U(\alpha, \alpha, \tau(\alpha))$. We assume that the shareholders of a publicly traded firm do not know α any better than the challenger does. Therefore if the market knows only that the firm has applied for a patent, the value of the firm should be its expected value

$$V = \frac{1}{\alpha_{\max} - \alpha_{\min}} \int_{\alpha_{\min}}^{\alpha_{\max}} U(\alpha, \alpha, \tau(\alpha))\, d\alpha. \tag{12}$$

Then after the firm has chosen its publication gap t, i.e., its lead-time $x = T - t$, and its choice becomes public knowledge, the stock price should change by the amount

$$\Delta(x) = U\left(\tau^{-1}(x), \tau^{-1}(x), x\right) - V, \tag{13}$$

where $\tau^{-1}(x) = \alpha(x)$ in Eq. (8) so that $\alpha'(x) = \frac{2b\alpha_{\min}}{a}\left(1 + \frac{b}{a}x\right) > 0$

Because x solves the first-order condition (6), the term U_2 and U_3 drop out and the

stock-price response to a unit rise in t (i.e., its response to a unit fall in x) should be

$$-\Delta'(x) = -U_1(\alpha(x), \alpha(x), x)\,\alpha'(x) < 0. \tag{14}$$

Therefore the stock prices also will identify only the ratio b/a.

Thus, three points are relevant for what the data will show. Firstly, disclosing later leads to a smaller change in V, possibly a decline. Or if we plot x on the horizontal axis, the gain should be larger. Secondly, the jump of size Δ should occur at t if the market learns about the patent *and its application date* at the time the disclosure occurs.

Thirdly, if instead the market learns about the patent application when it occurs, then the value of the firm should immediately jump up by the amount V. I.e., the upward jump should occur at the application date. After that, value should fall continuously until the disclosure date and it then should jump up and stabilize. To explain better: suppose that a is uniformly distributed on $[\alpha_{\min}, \alpha_{\max}]$. Upon application (if that's publicly known), the price jumps up to V as given in equation (12), which we can write as

$$V = E(U \mid \alpha \in [\alpha_{\min}, \alpha_{\max}])$$

Then as time passes and there is no disclosure yet, and x shrinks, the price shrinks to

$$V_x = E\left(U \mid \alpha \in \left[\alpha_{\min}, \tau^{-1}(x)\right]\right) \tag{15}$$

and as x shrinks, $\tau^{-1}(x)$ shrinks and V falls. And then at the actual disclosure date, there is an upward jump of size

$$\Delta_x = U_1\left(\tau^{-1}(x), \tau^{-1}(x), x\right) - V_x$$

where V_x is the prior belief after t periods of nondisclosure. I.e., disclosure tells the market that α is at the upper bound of the posterior range that the market had an instant before the disclosure. But the later the disclosure (i.e., the smaller is x), the smaller the upward

jump is and when x reaches $x_{\max}$,disclosure has no effect.

In words, at application (if the market sees there was an application) V jumps by the amount equal to the value of an *average* additional patent. Then if disclosure is early, the market infers the patent was above average and the price jumps up further at disclosure. If the disclosure is later, the market keeps revising downward its belief over the patent's quality, and V gradually declines, equal to the conditional expectation of the patent's value at each date. And finally at disclosure there is an upward jump because the market then infers that the quality is at the upper bound of its posterior.

II.D Policy issues

External effect. The incumbent bestows a positive external effect: Industry output is $\alpha + \delta(x)$ in the event that the challenger enters, and this is the likeliest to happen if $x = T$. So, since δ is increasing in x, if the quality of the idea α was generated exogenously (as our paper has assumed), it would be socially optimal to require *immediate* publication so that x (and hence $\delta(x)$) would be maximized at $\delta(T)$.

Disclosure should be immediate and as detailed as possible. But the incumbents payoff conditional on α would be smaller than its equilibrium payoff, i.e.,

$$U(\alpha, \alpha, T) = \omega(T, \alpha)\,\alpha D(\alpha + \delta(T)) + (1 - \omega(T, \alpha))\,\alpha D(\alpha) \tag{16}$$

$$\leq U(\alpha, \alpha, \tau(\alpha)). \tag{17}$$

The positive external effect means that if the generation of α does require costly research, then that research should be subsidized, and the subsidy would need to be even larger if the innovator is forced to disclose his idea immediately. I.e., if the results of research are automatically publicized and perhaps imitated, the innovator would then need to be compensated more for the costs of invention.

Delaying publication amounts to a lengthening of the period in which the idea is kept

secret; Levin et al. (1987) discuss some of the issues surrounding secrecy of ideas. And Klette, Møen and Griliches (2000) survey the evidence on the effectiveness of research subsidies.

II.E Heterogeneous Capital

Let k denote a vector of incumbent characteristics that are common knowledge and, in particular, that are known by the challenger. For example, k could denote the incumbent's capital stock or mre generally "commercialization capability."

Let us then define patent quality to be θ so that

$$\text{incumbent's output } = q(\theta, k). \tag{18}$$

The production function $q(\theta, k)$ must be strictly increasing in θ and in k because they both are inputs in production.

It then turns out, again for the case of unit elastic product demand, $\beta = 1$, that one can simply relabel α by $q(\theta, k)$ and that the same equilibrium then holds with the incumbent signaling q instead of α. Formally, Appendix C shows the following result:

Proposition 2. *If* $\beta = 1$,

$$x = \frac{a}{b}\left(\sqrt{\frac{q(\theta, k)}{q(\theta_{\min}, k)}} - 1\right). \tag{19}$$

Table 5 shows that he publication lag decreases with measures of θ and k.

That means x should be increasing in θ and k, and the result in (19) then implies that

$$\frac{q(\theta, k)}{q(\theta_{\min}, k)} \text{ must increase with } \theta \text{ and } k.$$

The ratio must rise with θ since only the numerator depends on θ, and q increases in θ. But

the ratio does not always rise with k.[1]

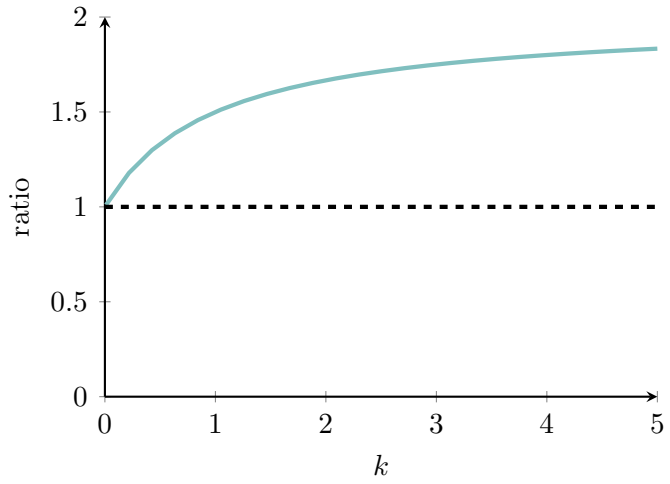


Figure 2: $\frac{q(\theta,k)}{q(\theta_{\min},k)}$ where $\theta_{\min}=1$ and $\theta=B=2$

The following production function implies that x increases (i.e., a publication lag that decreases) in both θ and k:

$$q(\theta,k)=k+\theta k^{B}.....\text{with } B>1. \tag{20}$$

Let's plot the ratio for $\theta_{\min}=1$, $\theta=2$, $B=2$ in Figure 2. It rises linearly in θ, but is strictly concave in k.

II.F Other Possible Extensions

Size of sector. If we add a multiplicative parameter to the demand curve so that

$$p=AD\left(q\right),$$

[1]For example in the Cobb-Douglas case where $q=\theta^{A}k^{B}$,

$$\frac{q\left(\theta,k\right)}{q\left(\theta_{\min},k\right)}=\left(\frac{\theta}{\theta_{\min}}\right)^{A}$$

k drops out.

then it is easy to see that this raises the probability of entry as follows: The entrant's revenue rises by the factor A, so that equation (3) reads

$$\omega\left(x,\hat{\alpha}\right)=F\left(\delta\left(x\right)AD\left[\hat{\alpha}+\delta\left(x\right)\right]\right). \tag{21}$$

With c distributed uniformly on $[0,c_{\max}]$, such that $F(c)=c/c_{\max}$, the probability of entry would rise by the factor A. To preserve the solution, $A/c_{\max}$ cannot be too large, otherwise ω would exceed the upper bound of unity, rendering the solution invalid.

Superstar inventions. They have a large α, which would require a large $\alpha_{\max}$. This could lead to a failure of condition (7) and, with the finiteness of T, would result in partial pooling—the weakest types would pool at $x=0$, i.e., at $t=T$, and enough of them would need to do so to leave sufficient time for the other types to separate on the interval $[0,t)$.

II.G Alternative Stories

Our signaling model predicts that higher-quality inventions are disclosed sooner. Two alternative mechanisms could, in principle, generate a similar correlation. We discuss each in turn and explain why signaling remains the most compelling explanation for the empirical patterns we document.

Provisional royalties. Under 35 U.S.C. §154(d), a patent holder may collect reasonable royalties from infringers for the period between the publication date of the application and the date of patent grant, provided the infringer had actual notice of the published application. Earlier disclosure therefore extends the window over which provisional royalties can accrue. If higher-quality patents are more likely to be infringed, this mechanism could generate a positive correlation between quality and early disclosure: inventors of valuable patents would disclose sooner to begin accumulating royalty claims earlier.

However, this channel is unlikely to account for the broad patterns we observe. The

exercise of provisional rights requires that the patent be eventually granted with claims substantially identical to those published, and that the patent holder pursue infringers—conditions that are rarely met in practice. Licensing revenues are concentrated among a small fraction of patents and in a narrow set of technology fields, predominantly pharmaceuticals and semiconductors (Arora et al., 2023). By contrast, the negative relationship between patent quality and publication lag that we document is pervasive: it holds across all NBER technology classes, for both small and large firms, and along both the extensive and intensive margins of disclosure timing. A mechanism that operates only for the small share of patents that generate licensing income cannot explain this universal pattern.

Preemption through the presence of a patent. Our baseline model already features entry deterrence: the challenger observes disclosure timing x, infers the incumbent's type $\hat{\alpha}$, and stays out when entry would be unprofitable. One might therefore ask whether the empirical correlation between quality and early disclosure simply reflects an earlier preemption—the incumbent disclosing early to deter rivals from even entering the product market as a potential competitor—rather than signaling of precise quality.

To clarify the distinction, consider an extension of our model to a two-stage game. In the first stage, the incumbent makes a press release about the patent, and the mere presence of the patent reveals the inventor's overall strength to potential rivals—specifically, the support $[\alpha_{\min}, \alpha_{\max}]$ of his quality distribution. A challenger must pay a fixed cost $K > 0$ to enter the second-stage game. If the challenger perceives the incumbent to be sufficiently strong—that is, if $\alpha_{\min}$ is high enough that even competing against the weakest type within the range would be unprofitable after paying K—it chooses not to enter and is preempted. This first-stage deterrence operates through the presence of the patent and the inventor's known characteristics alone. It does not require the challenger to observe disclosure timing or to learn the incumbent's precise type.

In the second stage, conditional on the challenger paying K, the signaling game proceeds

exactly as in Section II: the incumbent chooses disclosure timing x to signal his exact type α within $[\alpha_{\min}, \alpha_{\max}]$, and the challenger decides whether to produce.

The two stages thus capture distinct channels. Preemption in the first stage depends on *who* the inventor is—his aggregate strength as reflected in $[\alpha_{\min}, \alpha_{\max}]$. Signaling in the second stage depends on *what* the inventor has invented—the precise quality α of the specific patent, as revealed through the timing of disclosure. The empirical predictions differ accordingly. Preemption through the presence of a patent should matter more for large, well-known firms, whose overall strength is already visible to rivals and whose patents pose more credible competitive threats. Signaling through disclosure timing should matter more for small firms, which face greater information asymmetry and for whom the quality of a specific invention carries more weight relative to their observable characteristics. As we show in Section IV, small firms exhibit a significantly stronger relationship between patent quality and early disclosure, consistent with the signaling channel.

III Data and Descriptive Results

Sample Construction We start with USPTO's publicly available patent research dataset and restrict our sample to all regular utility patent applications filed between November 29, 2000 and December 31, 2005. We then remove observations for which important information, such as grant date or application date seemed erroneous or missing. For each patent, we calculate forward citations for the five-year and ten-year periods, excluding self-citations of the same patent owners and counting forward-citations from the disclosure date of patents (that is, 18-month publication date or grant date). For each patent, we identified the numbers of product-innovation-related and process-innovation-related claims, provided by Ganglmair, Robinson and Seeligson (2022) using text-analysis-based classification.

We also bring in Compustat financial information for publicly-listed patentees by using the crosswalk between patent assignees and Compustat firms provided by Arora, Belenzon

and Sheer (2021). Out of the 1.98 million patents in our sample, 0.55 million (27.9%) are matched to publicly listed firms.

Variable Definition. We define the disclosure date of a patent to be the earlier of its publication date and grant date date. If there is a parent application associated with the patent through continuations, the filing date of the patent is adjusted to that of the parent application (consistent with AIPA rules which require patent publication to occur 18 months from the application date of the earliest parent application). The publication gap of a patent is then defined to be the difference between its disclosure and filing dates, measured in months. The definitions of other variables are listed in Table 1.

Table 1: Variable Definitions

Variable	Definition and Source
Pub Gap	The difference between patent disclosure and filing dates, measured in months. The disclosure date equals the publication date for patents with pre-grant publications and grant date for patents without. Author's calculation from USPTO data.
X-Year Citations	The number of forward citations received within X years after disclosure. Author's calculation from USPTO data.
Allowed Claims	Total number of claims allowed at grant. From USPTO data.
Process Claims	Claims related to process innovation, identified by Ganglmair, Robinson and Seeligson (2022) using text analysis.
Product Claims	Claims related to product innovation, identified by Ganglmair, Robinson and Seeligson (2022) using text analysis.
Small Firm Dummy	Dummy variable equal to 1 if the patent-assignee qualifies as a small entity for reduced fees at USPTO. Small entities include a person, a small business (including affiliates, and having no more than 500 employees), or a nonprofit organization. From USPTO data.
Public Firm Dummy	Dummy variable equal to 1 if the patent-assigning firm is publicly traded; 0 otherwise. From Compustat data.
Past Patents	The set of patents granted to the focal patent's assignee prior to the grant of the focal patent. Author's calculation from USPTO-Compustat merged data.
Assets	The patent-assigning firm's total assets/liabilities. From Compustat variable "at."
Employment	The patent-assigning firm's number of employees. From Compustat variable "emp."
R&D Expenses	The patent-assigning firm's R&D expenditure. From Compustat variable "rdip."
Revenue	The patent-assigning firm's total revenue. From Compustat variable "revt."
Capital Intensity	The patent-assigning firm's property, plant, and equipment. From Compustat variable "ppegt."

Table 2: Summary Statistics

Panel A: Patent Characteristics

	Obs (1K)	Mean	SD	Median	p5	p95
Pub Gap (months)	1980	34.1	21.8	27.5	15.9	75.4
5-Year Citations	1980	8.9	17.4	4	0	34
10-Year Citations	1980	16.9	36.4	7	0	66
Num Allowed Claims	1980	17.6	16.4	15	3	42
Frac Product Claims	1980	64.0	48.0	100	0	100

Panel B: Patent Owner Characteristics

	Obs (1K)	Mean	SD	Median	p5	p95
Small Firm Dummy	154	0.59	0.49	1	0	1
Public Firm Dummy	154	0.05	0.22	0	0	0
Num Patents	154	11.9	227.2	1	1	19
Employment	7.4	40.7	72.7	10.2	0.1	177.1
Assets ($1B)	7.5	23.0	73.5	2.3	0.0	91.7
R&D Expenses ($1M)	7.3	0.66	0.44	1.0	0.0	1.0
Revenue ($1B)	7.5	13.1	28.8	1.9	0.0	84.3
Capital Intensity ($1B)	7.5	8.4	21.5	0.9	0.0	39.4

Notes: This table shows summary statistics for the analysis sample constructed in Section III. Panel A summarizes patent characteristics, and Panel B summarizes patent owner characteristics.

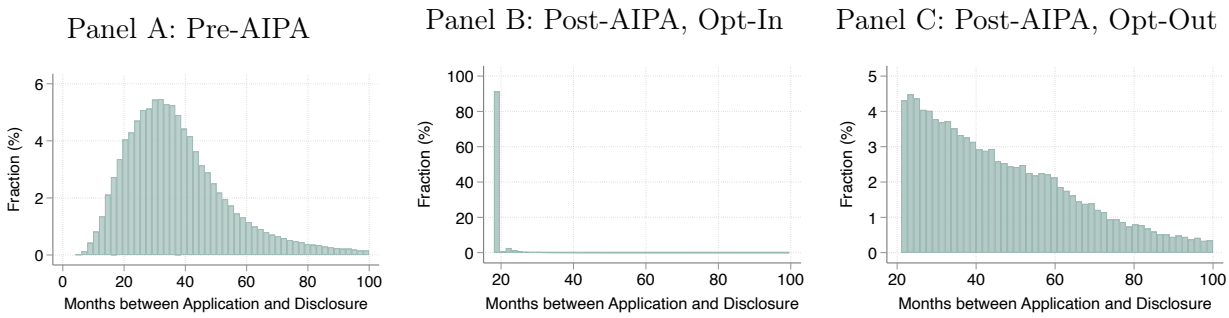
Summary Statistics Table 2 Panel A reports summary statistics for our sample of 1.98 million patents. The average publication gap is 34.1 months, with a standard deviation of 21.8 months. On average, each patent receives 8.9 citations within five years of publication and 16.9 citations within ten years. Patents are granted an average of 17.6 claims, 64% of which are related to product innovations and 36% process innovation.

Among all patent assignees, approximately 59% qualify as small entities under USPTO guidelines, while 5% are publicly traded firms found in Compustat. The average number of patents per assignee is 11.9. The average publicly listed assignee has 40.7 employees, $23 billion in total assets, $0.66 million in annual R&D expenditure, $13.1 billion in annual revenue, and $8.4 billion in capital intensity.

Descriptive Results. As shown in Figure 1, approximately 11.3% to 14.1% of patents each year opted out of pre-grant publication requirements after the enactment of AIPA. These patents are disclosed only at the time of grant, with their publication gap distribution plotted

in Figure 3 Panel C. In contrast, patents that do not opt out are mostly published around 18 months after filing, as verified in Figure 3 Panel B.

Figure 3: Distribution of Publication Lag



Notes: This graph plots the distribution of publication lags for three subsamples of regular utility patent applications filed at the USPTO between 1995 and 2005. Panel A focuses on patent applications filed before the American Inventors' Protection Act (AIPA). Panel B includes patent applications that were filed after the AIPA and disclosed at 18 months. Panel C covers patent applications filed after the AIPA with delayed disclosure. Patent applications with foreign priorities are excluded from the analysis. Before AIPA, all U.S. patent applications were published only upon grant. However, AIPA introduced the option for 18-month disclosure for applications filed on or after November 29, 2000.

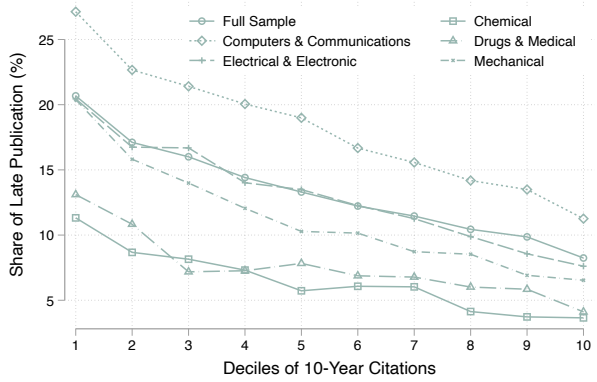
We then examine the relationship between patent quality and the timing of disclosure. There are two margins regarding when to disclose the patent content: the extensive margin, i.e. whether to opt out of the pre-grant publication requirement, and the intensive margin, i.e. after opting out, whether to logistically delay patent grant and publication.

We then examine the relationship between patent quality and the timing of disclosure. There are two relevant margins: the extensive margin—whether the applicant opts out of pre-grant publication—and the intensive margin—conditional on opting out, whether the applicant strategically delays the patent grant and publication.

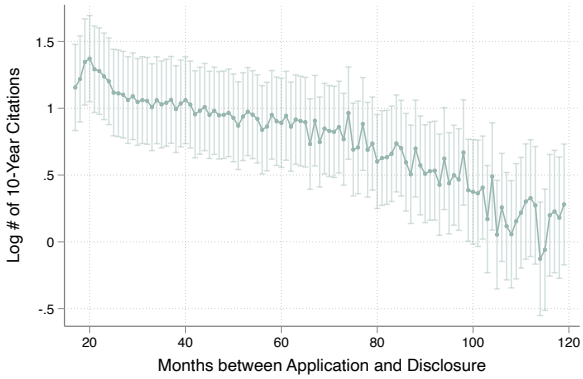
Figure 4 shows that higher-quality patents tend to be disclosed earlier along both extensive and intensive margins. Panel A plots the fraction of patents that opt out of pre-grant publication requirement across five NBER technology classes. The opt-out rate is highest in Computers and Communications (exceeding 25% for the bottom decile of quality) and lowest in Chemicals, Drugs, and Medical (below 13% for the bottom decile of quality). As patent quality increases, the opt-out rate declines significantly, and this negative relationship holds consistently across all technology classes. Panel B examines the intensive margin by plotting

Figure 4: Publication Lag and Patent Quality

Panel A: Extensive Margin



Panel B: Intensive Margin



Notes: This graph plots the relationship between patent quality, measured as 10-year forward citations, and publication lags. Panel A plots the share of late-disclosed patents across deciles of patent quality for five different NBER patent classes. Panel B shows the average patent quality for each monthly bin of publication lags, after absorbing fixed effects for patent categories (i.e. CPC codes). The sample includes all regular utility patent applications filed in the United States between November 29, 2000 and December 31, 2005, that had no equivalent foreign application and were granted before July 22, 2022.

average patent quality across monthly bins of the publication gap. Even among patents that opt out of pre-grant publication, those of lower quality tend to take longer to be disclosed. Appendix Figures A1 and A2 show that these patterns are robust when using alternative measures of patent quality such as 5-year forward citations and the number of claims.

IV Regressions and Model Estimation

IV.A Reduced-Form Analysis

To test the main prediction of our signaling model—that higher-quality inventions are disclosed sooner—we estimate the following regression specification:

$$Y_i = \beta Q_i + Z_i + Z_{f(i)} + \gamma_{\text{CPC}} + \gamma_{\text{ApplYear}} + \gamma_{\text{GrantYear}} + \varepsilon_i. \tag{22}$$

Here, Q_i denotes the quality of patent i, measured by five-year and ten-year forward citations, as well as the number of claims. The outcome variable Y_i is either a binary indicator for opting out of pre-grant publication or the log of the publication gap for patents that do opt

out. The vector Z_i includes patent-level characteristics, primarily the number of past patents assigned to the same entity. $Z_{f(i)}$ includes assignee-level controls: indicators for whether the assignee is a public firm and whether it qualifies as a small entity, and for publicly traded assignees, the log of capital intensity and R&D expenditure. All regressions include fixed effects for CPC technology class, application year, and grant year.

Regression results are presented in Table 3. Columns (1)–(3) examine the extensive margin, using a binary indicator for opting out of pre-grant publication as the dependent variable. We find a significant negative correlation between patent quality and the likelihood of opting out, controlling for patent- and assignee-level observables, as well as CPC class, application year, and grant year fixed effects. Specifically, a doubling of five-year forward citations is associated with a 2.3 percentage point reduction in the probability of opting out.

Columns (4)–(6) focus on the intensive margin, restricting to patents that opt out. The results show that higher-quality patents are disclosed sooner: a doubling of five-year citations corresponds to a 3.6% reduction in the publication gap. This negative relationship is robust—and even stronger—when restricting to patents filed by public firms, for which we further control for capital and R&D intensity.

Taken together, these results support the model's prediction: patents with higher unobserved quality (that cannot be explained by observables) are more likely to be disclosed earlier, both on the extensive and intensive margins.

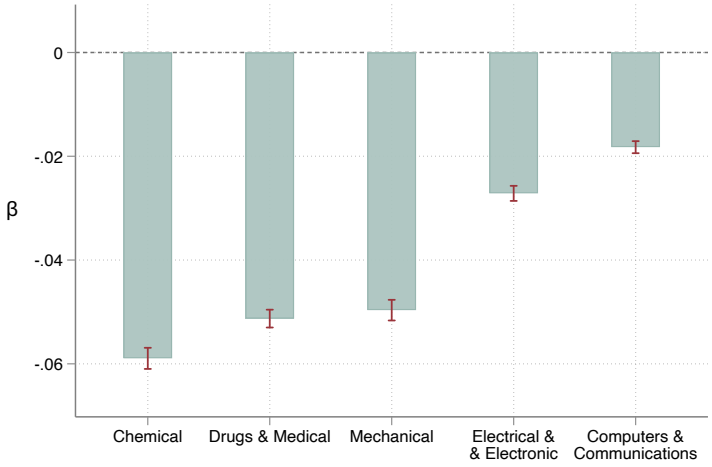
Heterogeneity across Sectors. We estimate the same regression specification separately for each of the five NBER technology classes. Figure 5 plots the main coefficient β from these regressions, using ten-year forward citations as the measure of patent quality. The negative relationship between quality and publication lag is strongest in the Chemical and Drugs & Medical sectors (approximately –0.05), and weakest in the Electrical & Electronic and Computers & Communications categories (around –0.02). Full regression results, including estimates using alternative quality measures, are reported in Appendix Table A1.

Table 3: Publication Lag and Patent Quality

	Extensive: Late Pub Dummy			Intensive | Late Publication: Log Pub Gap		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: All Patents						
Log 5-Year Citations	-0.023***			-0.036***		
	(0.000)			(0.001)		
Log 10-Year Citations		-0.031***			-0.047***	
		(0.000)			(0.001)	
Log Allowed Claims			0.000			-0.015***
			(0.001)			(0.001)
N	407044	407044	407044	132781	132781	132781
CPC FE	Yes	Yes	Yes	Yes	Yes	Yes
Appl Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Grant Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Public Firm Indicator, Small Firm Indicator, Log Number of Past Patents					
Panel B: Patents Filed by Public Firms						
Log 5-Year Citations	-0.029***			-0.040***		
	(0.001)			(0.001)		
Log 10-Year Citations		-0.037***			-0.048***	
		(0.001)			(0.001)	
Log Allowed Claims			-0.003***			-0.021***
			(0.001)			(0.002)
N	208879	208879	208879	63556	63556	63556
CPC FE	Yes	Yes	Yes	Yes	Yes	Yes
Appl Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Grant Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Small Firm Indicator, Log Number of Past Patents, Log Capital Intensity, Log R&D Expenses					

Notes: This table presents regression results that describe how patent disclosure decisions are related to patent quality and other characteristics. Panel A uses the full sample of patents filed post AIPA that do not have foreign priority, while Panel B restricts to patents filed by public firms. The outcome variable in Columns (1) to (3) is an indicator for opting out of the pre-grant publication. The outcome variable in Columns (4) to (6) is the log publication gap, conditional on opting out of pre-grant publication. Fixed effects for each application year, issue year, and CPC code are included. Standard errors are reported in parentheses. *** indicates significance at the 0.01 level, ** at 0.5, and * at 0.1.

Figure 5: Publication Lag Since Application Date



Notes: This figure plots the coefficient estimates from Equation (22), separately estimated for each of the five NBER technology classes. The 95th confidence intervals are plotted in red.

IV.B Model Estimation

We now proceed to estimate the model parameters, both for the full sample and separately by broad technology classes. Recall the model solution

$$\alpha(x) = \alpha_{\min}\left(1 + \frac{b}{a}x\right)^2$$

where α is patent quality, x is the time left for copying, $\alpha_{\min}$ is the left bound of the distribution for α, and b/a captures the relative importance of gradual versus Eureka learning. To facilitate estimation of the model, we pick the five-year forward citation as the empirical counterpart of α and set the latest possible time for disclosure, T, to 100 months after patent application. It follows that x is the difference between T and the publication gap.

We use Nonlinear Least Squares to fit this quadratic equation and estimate two parameters, i.e. $\alpha_{\min}$ and b/a. Notice that a and b cannot be separately identified as only their *ratio* appears in the solution. Intuitively, a is the amount of Eureka knowledge that rivals can learn immediately upon the disclosure of a patent, whereas b measures how much more knowledge rivals can learn about the patent over time. The model only allows us to estimate the relative

importance of the Eureka and gradual learning.

Table 4: Estimates of Model Parameters by NBER Technology Classes

	Full Sample (1)	Chemical (2)	Computers and Communications (3)	Drugs and Medical (4)	Electrical and Electronic (5)	Mechanical (6)	Others (7)
$\alpha_{\min}$	7.066***	4.394***	11.167***	3.073***	5.210***	5.604***	4.408***
	(0.2696)	(0.6256)	(0.534)	(0.723)	(0.490)	(0.6288)	(0.4544)
b/a	0.003***	0.003***	0.002***	0.018***	0.004***	0.001	0.003***
	(0.0003)	(0.0011)	(0.0003)	(0.0036)	(0.0008)	(0.0008)	(0.0008)
N	435534	41241	138095	42665	86269	58575	63280

Notes: This table presents NLS estimates for equation 8. $\alpha_{\min}$ is the model parameter that captures the lowest quality type of patents. The ratio b/a corresponds to the relative importance of gradual learning versus Eureka learning. The first row runs NLS for the full sample, whereas the second to the last row runs NLS for subsamples that correspond to each NBER technology class. Standard errors are reported in parentheses. *** indicates significance at the 0.01 level, ** at 0.5, and * at 0.1.

Table 4 presents the estimates of model parameters for the full sample and each of the six NBER technology classes. Interestingly, b/a is estimated to be lower in Drugs and Medical where learning by competitors is gradual, and b/a is essentially zero in Computers and Communication where learning by rivals is more immediate (i.e., Eureka learning). This is consistent with our previous finding, evident from Figure 4, that the most important inventions in Chemicals or Drugs and Medicine field are about thrice as likely as their counterparts in the Computers and Communication field to be disclosed early, at 18 months.

Table 8 presents heterogenous effects based on firm ownership, size, R&D intensity, and capital intensity. The last two columns of this table also explore heterogenous effect for product-innovation versus process-innovation patents. Patents that are considered to be product-innovation patents are those in which the majority of the claims relate to invention of new products. Inventions filed by public firms, large firms, R & D intensive firms, and for process-intensive patents appear to be characterized more by gradual learning rather than instantaneous learning.

V Stock Market Analysis

When the stock market learns about the existence and content of a patent, it should respond positively if the patent is perceived to be of high quality. This logic allows us to infer the value of a patent from the stock price reaction at the time of its publication. It is therefore informative to test whether market responses align with our theoretical predictions and prior empirical findings.

More interestingly, the post-AIPA pre-grant publication policy separates the timing of patent disclosure and grant. This separation enables us to distinguish between the market's response to new information about the invention and its response to the legal conferment of monopoly rights.

In this section, we restrict the analysis to firms whose PERMNO identifiers were successfully matched to the patent database—that is, firms that filed at least one patent with the USPTO. Our primary variable of interest is the "excess return," calculated using a buy-and-hold strategy relative to a benchmark portfolio. Specifically, we define each stock's excess return as the geometrically compounded return on the stock minus the geometrically compounded return on the CRSP value-weighted market index (including dividends). The excess return is computed as follows:

$$\mathrm{ER}_j^{ab} = \Pi_{t=a}^{b} \left(1 + R_{jt}\right) - \Pi_{t=a}^{b} \left(1 + \mathrm{MR}_t\right),$$

where ER_j^{ab} is the excess return for firm j over the time period from a to b, R_{jt} is the stock's raw return on day t, and MR_t is the return on the benchmark market index. We use an event window of 61 trading days, i.e. 30 trading days prior to the patent event and 30 trading days after. That is, we set $a = -30$ and construct the excess return ER_{jt} for all $t = b \in [-30, 30]$.

We identify the following four types of events to analyze:

Pre-AIPA Grant (P). Prior to the enactment of AIPA, patent disclosure happens at the same time with grant. We identify clean events on patent grant during this period—those

without any co-occurring events within a 60-trading-day window before or after the grant date. We do, however, include instances where multiple patents were issued to the same firm on the same day; that is, we treat the issuance of one or more patents on a single day as a single event. This selection leaves us with 16,018 events.

Post-AIPA, Opt-Out Grant (L). The second type of event involves patents filed after AIPA that opted out of the pre-grant publication requirement. For these patents, disclosure occurs simultaneously with the grant. As before, we retain only clean events—those without any other patent issuance or disclosure within a 60-trading-day window before or after. We allow multiple patents granted to the same firm on the same day, treating them as a single event. This selection yields 386 events.

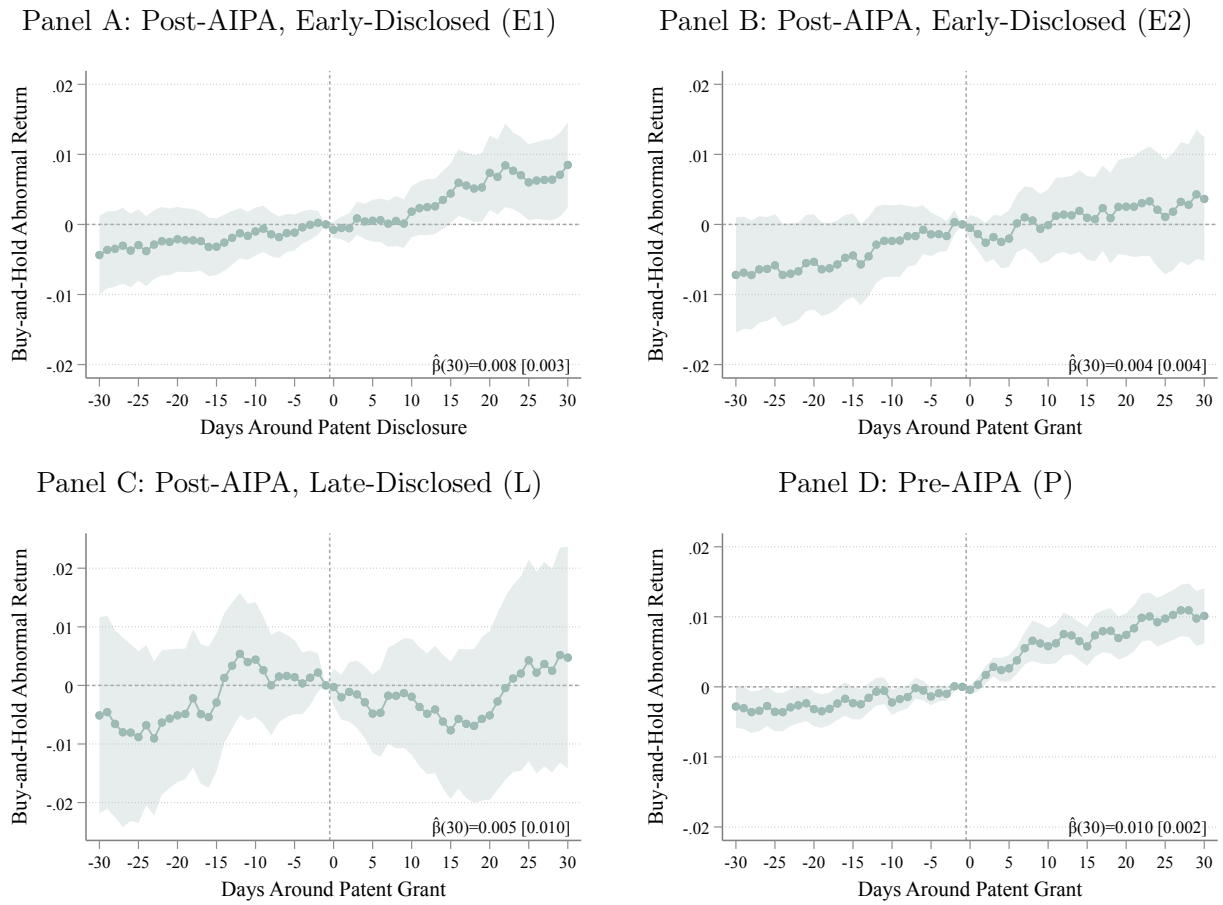
Post AIPA, Opt-In Grant (E2). In this category, we select patents filed after AIPA that were disclosed prior to issuance, i.e. opting-in for pre-grant publication. For these patents, we identify their issuance as the third type of events. We retain only clean events—those with no other patent issuance or disclosure within 60 trading days before or after. As before, we treat multiple patents issued to the same firm on the same day as a single event. This yields 2,712 events.

Post AIPA, Opt-In Publication (E1). For the last type of events, We select post-AIPA patents that opted in for the pre-grant publication requirement, and we focus on their publication event. We retain only clean events—those with no other disclosure or issuance within 60 trading days before or after. As before, multiple patents disclosed on the same day are treated as a single event. This yields 3,653 events.

Having identified the four event types, we now examine how excess returns evolve around each. Specifically, we regress ER_{jt} on relative-event-time dummies for days -30 to 30, clustering standard errors at the firm level. Results are plotted in Figure 6. We find no evidence of pre-trends, suggesting no abnormal returns prior to any event. For post-AIPA patents, we detect no significant excess return within a month of patent grant, regardless of opt-in or

opt-out status. However, for post-AIPA patents that opt in, we observe a significant positive excess return around one percentage point at the 30th trading day after disclosure. Pre-AIPA patents result in a similar gain at the 30th trading days after grant.

Figure 6: Stock Market Response to Patent Events



Notes: This figure presents coefficient estimates from event study regressions around four patent events. The shaded areas represent the 95% confidence intervals. The buy-and-hold abnormal return reflects the excess return from holding the stock relative to the CRSP Index. The event window spans 30 trading days before to 30 trading days after each event.

We report the exact estimates of excess returns on the 30th trading day after each event in Column (1) of Table 5. As a test for our theoretical prediction, recall that the model suggests that the quality of post-AIPA patents is higher for opt-in patents and lower for opt-out patents, and the quality of pre-AIPA patents will be the average so somewhere in between. From Column (1) again, we see that post-AIPA, opt-in patents lead to 1.2% excess return on the stock, opt-out patents lead to 0.5% excess return; the former is indeed higher.

The pre-AIPA patents lead to 1% excess return on average, falling in between the previous two values.

These estimates provide a test of the model's prediction that post-AIPA opt-in patents should be of higher quality than post-AIPA opt-out patents, and pre-AIPA patents have the (weighted) average quality in between. Consistent with this prediction, we find that post-AIPA opt-in patents are associated with an average excess return of 1.2%, compared to 0.5% for opt-out patents. Pre-AIPA patents yield an average excess return of 1.0%, falling between the previous two values.

Table 5: Estimated Patent Value

	(1)	(2)	(3)	(4)	(5)
	Estimated Returns at $t = 30$	In 1982 million dollars		In 2020 million dollars	
Event		Median Value	Mean Value	Median Value	Mean Value
Baseline		110	1610	295	4318
Post-AIPA, Opt-In					
Disclosure (E1)	0.8%	0.9	12.9	2.4	34.5
Grant (E2)	0.4%	0.4	6.4	1.2	17.3
Total	1.2%	1.3	19.3	3.5	51.8
Post-AIPA, Opt-Out					
Total (L)	0.5%	0.5	8.0	1.5	21.6
Pre-AIPA					
Total (P)	1.0%	1.1	16.1	2.9	43.2

Notes: This table presents the estimated cumulative stock returns for patent events, along with the corresponding monetary values expressed in 1982 and 2020 dollars.

We further use our estimates of excess returns to infer the monetary value of patents in each group. The average firm's market capitalization deflated to 1982 dollars using CPI is \$1,610 million. Multiplying this baseline value by the estimated return rates, we find that a pre-AIPA patent is valued at approximately \$16.1 million. A post-AIPA patent that opts out of pre-grant publication is valued at 8 million, while a post-AIPA patent that opts in is valued at 19.3 million. Of the latter, 12.9 million is attributable to the informational value or knowledge embedded in the patent, and 6.4 million reflects the monopoly market power conferred by patent protection. Column (3) of Table 5 reports analogous calculations based on the market value of the median firm.

VI Conclusion

We develop a theory of private information that explains the disclosure choices of inventors under competition. Our model predicts that high-quality inventors credibly signal the quality of their inventions by publishing their patents early. The model also provides a simple functional form to guide the estimation of the importance of learning over time (versus instantaneous learning) and quality thresholds across different industries and types of inventors.

Empirically, we find robust evidence for the central prediction of our signaling model that high-quality inventors publish sooner. Small inventors, those with lower investments in capital equipment, and R&D intensive firms seem to be particularly more likely to disclose their most important patents early. We also find that technology fields in which learning by competitors tends to be more gradual (e.g., Drugs and Medical) admit lower quality thresholds for disclosure, while fields in which learning is more immediate (e.g., Computers and Communication) are associated with higher quality thresholds for disclosure.

Our results inform policy discussions about the effect of patent disclosure provisions on innovation and competition. It also complements studies that document wide variation in the use of patents across different technology fields, by suggesting that at least some of the differences could be a result of how patent disclosures interact with the nature of learning in these fields.

References

Anton, James J, and Dennis A Yao. 1994. "Expropriation and inventions: Appropriable rents in the absence of property rights." *The American Economic Review*, 190–209.

Anton, James J, and Dennis A Yao. 2004. "Little patents and big secrets: managing intellectual property." *RAND Journal of Economics*, 1–22.

Arora, Ashish, Sharon Belenzon, and Lia Sheer. 2021. "Knowledge spillovers and corporate investment in scientific research." *American Economic Review*, 111(3): 871–98.

Arora, Ashish, Wesley Cohen, Honggi Lee, and Divya Sebastian. 2023. "Invention value, inventive capability and the large firm advantage." *Research Policy*, 52(1): 104650.

Arrow, Kenneth Joseph. 1972. *Economic welfare and the allocation of resources for invention.* Springer.

Bhattacharya, Sudipto, and Sergei Guriev. 2006. "Patents vs. trade secrets: Knowledge licensing and spillover." *Journal of the European Economic Association*, 4(6): 1112–1147.

Bhattacharya, Sudipto, Jacob Glazer, and David EM Sappington. 1992. "Licensing and the sharing of knowledge in research joint ventures." *Journal of Economic Theory*, 56(1): 43–69.

Friedman, David D, William M Landes, and Richard A Posner. 1991. "Some economics of trade secret law." *Journal of Economic Perspectives*, 5(1): 61–72.

Furman, Jeffrey L, Markus Nagler, and Martin Watzinger. 2021. "Disclosure and subsequent innovation: Evidence from the patent depository library program." *American Economic Journal: Economic Policy*, 13(4): 239–270.

Ganglmair, Bernhard, and Imke Reimers. 2019. "Visibility of technology and cumulative innovation: Evidence from trade secrets laws." *ZEW-Centre for European Economic Research Discussion Paper*.

Ganglmair, Bernhard, W Keith Robinson, and Michael Seeligson. 2022. "The rise of process claims: Evidence from a century of US patents." *ZEW-Centre for European Economic Research Discussion Paper*.

Gilbert, Richard J, and David MG Newbery. 1982. "Preemptive patenting and the persistence of monopoly." *The American Economic Review*, 514–526.

Graham, Stuart, and Deepak Hegde. 2015. "Disclosing patents' secrets." *Science*, 347(6219): 236–237.

Gross, Daniel P, et al. 2019. *The consequences of invention secrecy: Evidence from the USPTO Patent Secrecy Program in World War II.* National Bureau of Economic Research.

Hall, Bronwyn, Christian Helmers, Mark Rogers, and Vania Sena. 2014. "The choice between formal and informal intellectual property: a review." *Journal of Economic Literature*, 52(2): 375–423.

Hegde, Deepak, and Hong Luo. 2018. "Patent publication and the market for ideas." *Management Science*, 64(2): 652–672.

Hegde, Deepak, Kyle Herkenhoff, and Chenqi Zhu. 2023. "Patent publication and innovation." *Journal of Political Economy*, 131(7): 000–000.

Hopenhayn, Hugo A, and Francesco Squintani. 2011. "Preemption games with private information." *The Review of Economic Studies*, 78(2): 667–692.

Horstmann, Ignatius, Glenn M MacDonald, and Alan Slivinski. 1985. "Patents as information transfer mechanisms: To patent or (maybe) not to patent." *Journal of Political Economy*, 93(5): 837–858.

Johnson, Daniel KN, and David Popp. 2003. "Forced out of the Closet: The Impact of the American Inventors Protection Act on the Timing of Patent Disclosure." *RAND Journal of Economics*, 96–112.

Klette, Tor Jakob, Jarle Møen, and Zvi Griliches. 2000. "Do subsidies to commercial R&D reduce market failures? Microeconometric evaluation studies." *Research policy*, 29(4-5): 471–495.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. "Technological Innovation, Resource Allocation, and Growth." *The Quarterly Journal of Economics*, 665–712.

Levin, Richard C, Alvin K Klevorick, Richard R Nelson, Sidney G Winter, Richard Gilbert, and Zvi Griliches. 1987. "Appropriating the returns from industrial research and development." *Brookings papers on economic activity*, 1987(3): 783–831.

Long, Clarissa. 2002. "Patent signals." *The University of Chicago Law Review*, 625–679.

Lück, Sonja, Benjamin Balsmeier, Florian Seliger, and Lee Fleming. 2020. "Early disclosure of invention and reduced duplication: An empirical test." *Management Science*, 66(6): 2677–2685.

Mailath, George J. 1987. "Incentive compatibility in signaling games with a continuum of types." *Econometrica: Journal of the Econometric Society*, 1349–1365.

Milgrom, Paul, and John Roberts. 1982. "Limit pricing and entry under incomplete information: An equilibrium analysis." *Econometrica: Journal of the Econometric Society*, 443–459.

Modigliani, Franco, et al. 1999. "An open letter to the US Senate." *Retrived from http://www. eagleforum. org/patent/nobel_ letter. html.*

Nordhaus, William D. 1969. "Invention, growth and welfare: a theoretical treatment of technological change."

Spence, Michael. 1973. "Job Market Signaling." *The Quarterly Journal of Economics*, 87(3): 355–374.

Williams, Heidi L. 2017. "How do patents affect research investments?" *Annual Review of Economics*, 9: 441–469.

Table 6: Publication Lag and Patent Quality

	Extensive: Late Pub Dummy			Intensive | Late Publication: Log Pub Gap		
	(1)	(2)	(3)	(4)	(5)	(6)
Log 5-Year Citations	-0.0219***			-0.0345***		
	(0.0005)			(0.0009)		
Log 5-Year Citations × Small Firm Dummy	-0.0051***			-0.0064***		
	(0.0009)			(0.0016)		
Log 10-Year Citations		-0.0299***			-0.0447***	
		(0.0004)			(0.0009)	
Log 10-Year Citations × Small Firm Dummy		-0.0045***			-0.0068***	
		(0.0008)			(0.0015)	
Log Allowed Claims			0.0058***			-0.0080***
			(0.0008)			(0.0017)
Log Allowed Claims × Small Firm Dummy			-0.0190***			-0.0212***
			(0.0015)			(0.0029)
Small Firm Dummy	0.0144***	0.0171***	0.0678***	0.0061	0.0136**	0.0612***
	(0.0026)	(0.0029)	(0.0056)	(0.0048)	(0.0054)	(0.0105)
Log Past Patents	-0.0164***	-0.0166***	-0.0157***	-0.0081***	-0.0087***	-0.0071***
	(0.0002)	(0.0002)	(0.0002)	(0.0004)	(0.0004)	(0.0004)
Public Firm Dummy	0.0720***	0.0727***	0.0697***	0.0660***	0.0657***	0.0653***
	(0.0013)	(0.0013)	(0.0013)	(0.0025)	(0.0025)	(0.0026)
N	407044	407044	407044	132781	132781	132781
CPC FE	Yes	Yes	Yes	Yes	Yes	Yes
Appl Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Iss Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table presents regression results that describe heterogeneous effects of patent quality on disclosure decisions. Columns (1) to (3) uses a dummy variables that equals one if the patent opts for disclosure at granting as dependent variable. Columns (4) to (6) uses the log publication gap in months as dependent variable, conditional on opting for disclosure at granting. Fixed effects for each application year, issue year, and CPC code are included. Standard errors are reported in parentheses. *** indicates significance at the 0.01 level, ** at 0.5, and * at 0.1.

Table 7: Publication Lag and Patent Quality—Compustat Sample

	Extensive: Late Pub Dummy						Intensive | Late Publication: Log Pub Gap					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Log 5-Year Citations	-0.0284***			-0.0641***			-0.0384***			-0.0491***		
	(0.0006)			(0.0023)			(0.0012)			(0.0039)		
Log 5-Year Citations × Small Firm Dummy	-0.0068***						-0.0141***					
	(0.0023)						(0.0038)					
Log 5-Year Citations × Log Capital				0.0040***						0.0012**		
				(0.0002)						(0.0005)		
Log 10-Year Citations		-0.0365***			-0.0830***			-0.0467***			-0.0596***	
		(0.0006)			(0.0021)			(0.0011)			(0.0036)	
Log 10-Year Citations × Small Firm Dummy		-0.0085***						-0.0160***				
		(0.0022)						(0.0036)				
Log 10-Year Citations × Log Capital					0.0052***						0.0014***	
					(0.0002)						(0.0004)	
Log Allowed Claims			-0.0021*			0.0023			-0.0189***			-0.0191***
			(0.0011)			(0.0043)			(0.0022)			(0.0073)
Log Allowed Claims × Small Firm Dummy			-0.0211***						-0.0201***			
			(0.0045)						(0.0071)			
Log Allowed Claims × Log Capital						-0.0007						-0.0002
						(0.0005)						(0.0008)
Small Firm Dummy	-0.0812***	-0.0681***	-0.0290*	-0.0973***	-0.0927***	-0.1061***	-0.0268**	-0.0084	-0.0026	-0.0645***	-0.0607***	-0.0755***
	(0.0074)	(0.0084)	(0.0167)	(0.0037)	(0.0037)	(0.0037)	(0.0119)	(0.0133)	(0.0263)	(0.0061)	(0.0061)	(0.0062)
Log Past Patents	-0.0257***	-0.0260***	-0.0251***	-0.0258***	-0.0260***	-0.0251***	-0.0157***	-0.0164***	-0.0141***	-0.0157***	-0.0164***	-0.0141***
	(0.0005)	(0.0005)	(0.0005)	(0.0005)	(0.0005)	(0.0005)	(0.0009)	(0.0009)	(0.0009)	(0.0009)	(0.0009)	(0.0009)
Log Capital	-0.0114***	-0.0114***	-0.0111***	-0.0213***	-0.0275***	-0.0088***	-0.0008	-0.0009	-0.0009	-0.0037***	-0.0052***	-0.0002
	(0.0005)	(0.0005)	(0.0005)	(0.0008)	(0.0009)	(0.0018)	(0.0009)	(0.0009)	(0.0009)	(0.0014)	(0.0015)	(0.0031)
Log R&D Expenses	0.0102***	0.0103***	0.0103***	0.0103***	0.0104***	0.0103***	0.0019***	0.0018**	0.0027***	0.0020***	0.0019**	0.0027***
	(0.0004)	(0.0004)	(0.0004)	(0.0004)	(0.0004)	(0.0004)	(0.0007)	(0.0007)	(0.0008)	(0.0007)	(0.0007)	(0.0008)
N	208879	208879	208879	208879	208879	208879	63556	63556	63556	63556	63556	63556
CPC FE	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Appl Year FE	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Iss Year FE	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes

Notes: This table presents regression results that describe heterogenous effects of patent quality on disclosure decisions using the Compustat sample, which includes all patents filed by public firms. Columns (1) to (3) uses a dummy variables that equals one if the patent opts for disclosure at granting as dependent variable. Columns (4) to (6) uses the log publication gap in months as dependent variable, conditioinal on opting for disclosure at granting. Fixed effects for each application year, issue year, and CPC code are included. Standard errors are reported in parentheses. *** indicates significance at the 0.01 level, ** at 0.5, and * at 0.1.

Table 8: Estimates of Model Parameters — Heterogeneous Effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	All	Public		Small		R&D Intensive		Capital Intensive		Innovation Type	
		0	1	0	1	0	1	0	1	Process	Product
$\alpha_{\min}$	7.066***	8.245***	6.323***	6.605***	8.677***	6.037***	6.811***	5.619***	6.346***	6.352***	7.442***
	(0.270)	(0.376)	(0.399)	(0.313)	(0.540)	(0.561)	(0.572)	(0.498)	(0.637)	(0.409)	(0.354)
b/a	0.003***	0.001***	0.005***	0.004***	0.001**	0.006***	0.003***	0.006***	0.004***	0.005***	0.002***
	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)
N	435534	241281	194253	296629	138899	97129	97022	98287	95865	167450	268084

Notes: This table presents NLS estimates for equation 8. $\alpha_{\min}$ is the model parameter that captures the lowest quality type of patents. The ratio b/a corresponds to the relative importance of gradual learning versus Eureka learning. The first columns uses the full sample. The second to third columns use two subsamples for patents that are filed by private versus public firms. Coumns (3) to (9) are based on subsamples defined in a similar way but based on wether the patent is filed by a large versus small firms, less versus more R&D intensive firms, and less versus more capital intensive firms. The last two columns use subsamples for process-innovation patents and product-innovation patents, respectively. Standard errors are reported in parentheses. *** indicates significance at the 0.01 level, ** at 0.5, and * at 0.1.

Appendix

A Appendix: Proof of Proposition 1

The incumbent's criterion (4) simplifies to

$$U(\alpha, \hat{\alpha}, x) = \alpha D(\alpha) - \omega(x, \hat{\alpha})\, \alpha \left[D(\alpha) - D(\alpha + \delta(x))\right]$$

The multiplicative constant α cancels and the criterion is

$$\max_x \left\{D(\alpha) - \omega(x, \hat{\alpha}) \left[D(\alpha) - D(\alpha + \delta(x))\right]\right\}$$

Then the probability of entry is

$$\omega(x, \hat{\alpha}) = (a + bx)(\hat{\alpha} + a + bx)^{-\beta}.$$

Then the incumbent maximizes

$$\begin{aligned} \frac{U(\alpha, \hat{\alpha}, x)}{\alpha} &= (a + bx)(\hat{\alpha} + a + bx)^{-\beta}(\alpha + a + bx)^{-\beta} \\ &\quad + \left(1 - (a + bx)(\hat{\alpha} + a + bx)^{-\beta}\right)\alpha^{-\beta} \end{aligned} \tag{1}$$

That is,

$$\frac{U(\alpha, \hat{\alpha}, x)}{\alpha} - \alpha^{-\beta} = (a + bx)(\hat{\alpha} + a + bx)^{-\beta}\left((\alpha + a + bx)^{-\beta} - \alpha^{-\beta}\right) \equiv \hat{U}(\alpha, \hat{\alpha}, x) \tag{2}$$

A.1 The case $\beta = 1$.

If $\beta = 1$, we have

$$\omega(x, \hat{\alpha}) = \frac{a + bx}{\hat{\alpha} + a + bx} = \frac{1}{1 + \hat{\alpha}/(a + bx)} < 1$$

$$U(\alpha,\hat{\alpha},x)=\frac{1}{1+\hat{\alpha}/(a+bx)}\frac{\alpha}{\alpha+a+bx}+\left(1-\frac{1}{1+\hat{\alpha}/(a+bx)}\right)\times 1$$

The criterion simplifies to

$$\begin{aligned}\hat{U}(\alpha,\hat{\alpha},x)-1 &= \frac{1}{1+\hat{\alpha}/(a+bx)}\left(\frac{\alpha}{\alpha+a+bx}-1\right)\\ &= \frac{-1}{1+\hat{\alpha}/(a+bx)}\frac{a+bx}{\alpha+a+bx}=\frac{-(a+bx)^2}{(\hat{\alpha}+a+bx)(\alpha+a+bx)}\end{aligned}$$

or to minimize

$$\frac{(a+bx)^2}{(\hat{\alpha}+a+bx)(\alpha+a+bx)}=(a+bx)^2(\hat{\alpha}+a+bx)^{-1}(\alpha+a+bx)^{-1} \tag{3}$$

The single-crossing property is proved in the last Appendix.

A.1.1 FOC for the case $\beta=1$

The criterion is

$$\hat{U}(\alpha,\hat{\alpha},x)\equiv(a+bx)^2(\hat{\alpha}+a+bx)^{-1}(\alpha+a+bx)^{-1}$$

The FOC is

$$\hat{U}_3(\alpha,\hat{\alpha},x)+\hat{U}_2(\alpha,\hat{\alpha},x)\frac{d\hat{\alpha}}{dx}=0$$

Evaluated at $\hat{\alpha}=\alpha$ the ODE reads

$$\begin{aligned}\frac{d\alpha}{dx} &= -\frac{\hat{U}_3(\alpha,\alpha,x)}{\hat{U}_2(\alpha,\alpha,x)}=\frac{2b(\alpha+a+bx)}{a+bx}-b-\frac{b(\alpha+a+bx)}{(\alpha+a+bx)}\\ &= \frac{2b(\alpha+a+bx)}{a+bx}-2b=2b\left(\frac{\alpha+a+bx}{a+bx}-1\right)\\ &= 2b\frac{\alpha}{a+bx}\end{aligned} \tag{4}$$

The ODE in (4) has the general solution

$$\alpha(x) = C(a + bx)^2. \tag{5}$$

To get C we use the weakest type, $\alpha_{\min}$ who would set $x = 0$, so that $\alpha_{\min} = Ca^2$, which implies

$$C = \frac{\alpha_{\min}}{a^2} \tag{6}$$

Substituting from (6) into (5) we get (8)

A.1.2 Single crossing property when $\beta = 1$

The single-crossing property (Mailath, 1987) requires that the ratio of the derivatives

$$\hat{U}_3(\alpha, \hat{\alpha}, x) / \hat{U}_2(\alpha, \hat{\alpha}, x)$$

is strictly monotone in α. Using the expression in (3) which we reproduce here,

$$\hat{U}(\alpha, \hat{\alpha}, x) \equiv (a + bx)^2 (\hat{\alpha} + a + bx)^{-1} (\alpha + a + bx)^{-1}$$

$$\hat{U}_2 = -(a + bx)^2 (\hat{\alpha} + a + bx)^{-2} (\alpha + a + bx)^{-1}$$

$$\begin{aligned}\hat{U}_3 = {} & 2b(a + bx)(\hat{\alpha} + a + bx)^{-1}(\alpha + a + bx)^{-1} \\ & - b(a + bx)^2 (\hat{\alpha} + a + bx)^{-2} (\alpha + a + bx)^{-1} \\ & - b(a + bx)^2 (\hat{\alpha} + a + bx)^{-1} (\alpha + a + bx)^{-2}.\end{aligned}$$

Then the following ratio should be monotone in α:

$$\frac{\hat{U}_3}{\hat{U}_2} = \frac{2b(\hat{\alpha} + a + bx)}{-(a + bx)} + b + \frac{b(\hat{\alpha} + a + bx)}{(\alpha + a + bx)},$$

and α enters only in the denominator of the third term – so the ratio $\frac{\hat{U}_3}{\hat{U}_2}$ is indeed strictly decreasing in α for all x.

B Proof that entry probability is proportional to λ

Here we derive Eq. (11). If $\beta = 1$, since $F(c) = \lambda c$, we have

$$\omega(x, \hat{\alpha}) = \lambda \frac{a + bx}{\hat{\alpha} + a + bx} = \frac{\lambda}{1 + \hat{\alpha}/(a + bx)} < 1$$

so that

$$U(\alpha, \hat{\alpha}, x) = \frac{\lambda}{1 + \hat{\alpha}/(a + bx)} \frac{\alpha}{\alpha + a + bx} + \left(1 - \frac{\lambda}{1 + \hat{\alpha}/(a + bx)}\right) \underbrace{\frac{\alpha}{\alpha}}_{=1}$$

The criterion in (4) simplifies to

$$U(\alpha, \hat{\alpha}, x) - 1 = \frac{\lambda}{1 + \hat{\alpha}/(a + bx)} \left(\frac{\alpha}{\alpha + a + bx} - 1\right)$$

and λ drops out, so that the incumbent's choice of x is still given by (9).

C Proof of Proposition 2

The incumbent's payoff is $qD(q)$ if the challenger stays out, and $qD(q + \delta(x))$ if he enters. If $\hat{q}$ is the challenger's belief about q, his expected payoff to entry is $\delta(x) D[\hat{q} + \delta(x)] - c$ and with F again being the CDF of c, the probability of entry is

$$\omega(x, \hat{q}) = F(\delta(x) D[\hat{q} + \delta(x)]).$$

Given $(q, \hat{q}, x)$, the incumbent's payoff is

$$U(q, \hat{q}, x) = \omega(x, \hat{q})\, qD(q + \delta(x)) + (1 - \omega(x, \hat{q}))\, qD(q).$$

However, in a separating equilibrium, beliefs $\hat{q}$ depend on the observed decision x. Thus, if the optimal decision $\tau(q)$ is unique, strictly monotone, and hence invertible, then $\hat{q} = \tau^{-1}(x)$. Equilibrium then requires that

$$\tau(q) = \arg\max_x U\left(q, \tau^{-1}(x), x\right) \quad \forall q,$$

and the associated first-order condition is

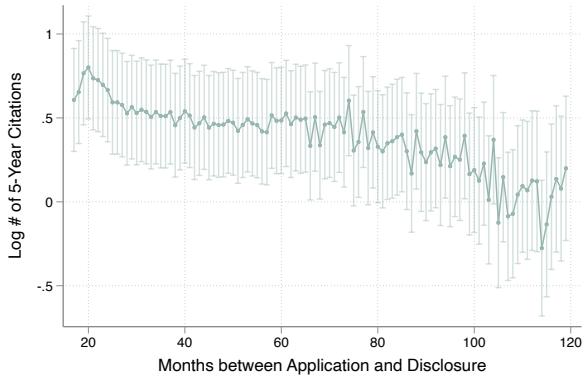
$$U_2(q, \hat{q}, x)\frac{d\hat{q}}{dx} + U_3(q, \hat{q}, x) = 0$$

Thus we have simply replaced α and $\hat{\alpha}$ with q and $\hat{q}$, the solution for x in Eq. (19) is the same function of q that it was as a function of α in Eq. (9).

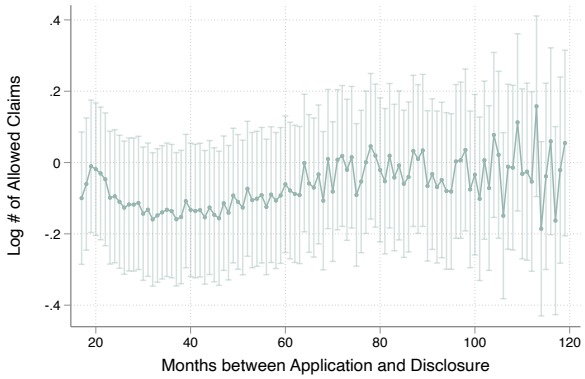
D Additional Figures and Tables

Figure A1: Publication Lag and Patent Quality

Panel A: 5-Year Citations



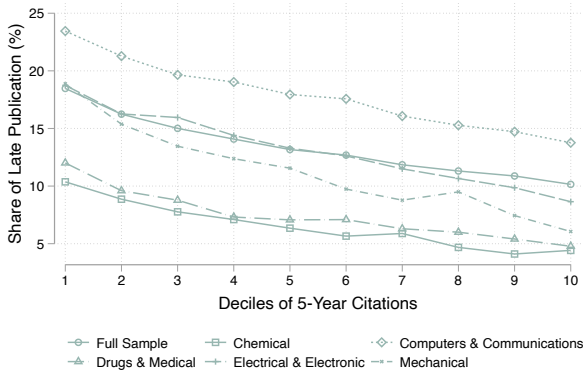
Panel B: Number of Claims



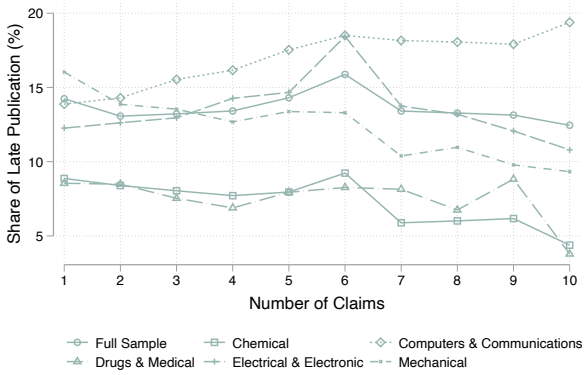
Notes: This graph plots the relationship between alternative measures of patent quality and publication lags. Specifically, it shows the average patent quality for each monthly bin of publication lags, after absorbing fixed effects for patent categories (i.e. CPC codes). Panel A measures patent quality using 5-year citations, and Panel C relies on the number of claims. The sample includes all regular utility patent applications filed in the United States between November 29, 2000 and December 31, 2005, that had no equivalent foreign application and were granted before July 22, 2022.

Figure A2: Late Publication Decision and Patent Quality

Panel A: 5-Year Citations



Panel B: Number of Claims



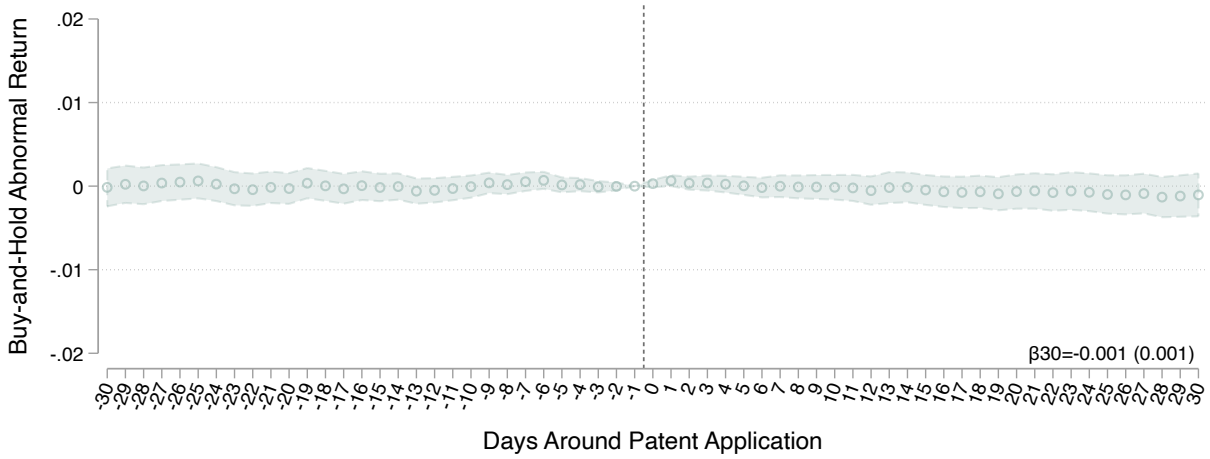
Notes: This figure plots the share of late-disclosed patents across deciles of patent quality for five different NBER patent classes. Panel A measures quality using 5-year forward citations, Panel B uses 10-year forward citations, and Panel C relies on the number of claims. Self-citations, where the citing and cited patents share common assignees, are excluded from the analysis. The sample includes all utility patent applications filed in the United States between November 29, 2000, and December 31, 2005, that were published before July 22, 2022.

Table A1: Publication Lag and Patent Quality - NBER Categories

	Extensive: Late Pub Dummy		
	(1)	(2)	(3)
Chemical × Log 5-Year Citations	-0.0531***		
	(0.0012)		
Computers&Communications × Log 5-Year Citations	-0.0059***		
	(0.0006)		
Drugs&Medical × Log 5-Year Citations	-0.0478***		
	(0.0010)		
Electrical&Electronic × Log 5-Year Citations	-0.0163***		
	(0.0008)		
Mechanical × Log 5-Year Citations	-0.0423***		
	(0.0012)		
Others × Log 5-Year Citations	-0.0547***		
	(0.0012)		
Chemical × Log 10-Year Citations		-0.0589***	
		(0.0010)	
Computers&Communications × Log 10-Year Citations		-0.0182***	
		(0.0006)	
Drugs&Medical × Log 10-Year Citations		-0.0513***	
		(0.0009)	
Electrical&Electronic × Log 10-Year Citations		-0.0272***	
		(0.0007)	
Mechanical × Log 10-Year Citations		-0.0497***	
		(0.0010)	
Others × Log 10-Year Citations		-0.0584***	
		(0.0010)	
Chemical × Log Allowed Claims			-0.0242***
			(0.0013)
Computers&Communications × Log Allowed Claims			0.0155***
			(0.0011)
Drugs&Medical × Log Allowed Claims			-0.0237***
			(0.0014)
Electrical&Electronic × Log Allowed Claims			0.0061***
			(0.0012)
Mechanical × Log Allowed Claims			-0.0153***
			(0.0013)
Others × Log Allowed Claims			-0.0282***
			(0.0013)
Small Firm Dummy	-0.1146***	-0.1114***	-0.1178***
	(0.0037)	(0.0037)	(0.0037)
Log Past Patents	-0.0192***	-0.0203***	-0.0211***
	(0.0005)	(0.0005)	(0.0005)
Log Capital	-0.0203***	-0.0197***	-0.0175***
	(0.0004)	(0.0004)	(0.0005)
Log R&D Expenses	-0.0011***	-0.0006*	0.0007**
	(0.0003)	(0.0003)	(0.0004)
N	207788	207788	207788
Appl Year FE	Yes	Yes	Yes
Iss Year FE	Yes	Yes	Yes

Notes: This table presents regression results that describe heterogeneous effects of patent quality on disclosure decisions by NBER categories. The LHS variables equals one if the patent opts for disclosure at granting as dependent variable. Fixed effects for each application year and issue year are included. Standard errors are reported in parentheses. *** indicates significance at the 0.01 level, ** at 0.5, and * at 0.1.

Figure A3: Stock Market Response to Patent Application



Notes: This figure presents coefficient estimates from event study regressions around patent application events. The shaded areas represent the 95% confidence intervals. The buy-and-hold abnormal return reflects the excess return from holding the stock relative to the CRSP Index. The event window spans 30 days before to 30 days after each event. The result indicates that there is no information leakage around patent applications.