Innovation, Intellectual Property, and the Role of External Forces

Jonathan F. Lee*

May 2017

Abstract

Innovators must decide on an innovation strategy and how to protect their innovations from potential competitors. While the two decisions are related, they are not identical, and there is no guarantee that changes in the propensity to innovate must be reflected by changes in the propensity to protect. I separate a firm's innovation and intellectual property (IP) protection decisions and enrich her choices by allowing for multiple possible methods of IP protection. I find that the potential for competition in a firm's industry reduces the incentives to create a market through innovation, but it increases the use of IP should innovation be successful. These two channels combine so that industry competition has a broadly inverse–U shaped relationship with observed patenting, but a normal–U relationship with actual innovation. Here, the average quality of a firm's patent portfolio is a better proxy for innovation: when incentives to patent are low and innovative effort is high, firms only other to file high–quality patents.

JEL–Classification: K11, L24, O31, O34
Keywords: innovation, intellectual property, patents, trade secrets, market structure, competition

*Queen's University, Dept. of Economics. Contact: leejf@econ.queensu.ca.

I am grateful for helpful comments from Ján Zábojník, Marie–Louise Vierø, Veikko Theile, Jean de Bettignies, Olena Ivus, Corinne Langinier, and Yossi Spiegel, as well as seminar participants at Queen's University, the Canadian Economics Association 2016 Annual Meeting, and the International Industrial Organization 2016 Conference .

Innovation is generally understood as the key to economic growth and prosperity, but measuring the innovative activity in an economy is difficult. Spending on research and development, one possible measure, is only sometimes recorded by firms, and it is unclear to what extent that spending increases innovative output, as opposed to the creation of basic scientific knowledge (Nagaoka et al., 2010). The Solow residual, the classic measure of aggregate technological change, comprises the empirical leftovers after accounting for whatever we *can* measure in our data (Solow, 1957). The most common measure of innovation in wide usage is quality–adjusted patenting: a researcher counts the flow of new patents granted to an entity, weights them by some quality metric (usually citations from future patents), and uses that weighted sum as measured innovation. These patenting measures are far from exact: Griliches (1990) writes, "Not all inventions are patentable, not all inventions are patented...", but nevertheless "patent statistics remain a unique resource for the analysis of the process of technical change."

Adjusting patent counts by quality can alleviate some of the problems with using patenting as a measure of innovation, but other issues remain. Quality adjustments intend to adjust for the variation in the amount of innovation embodied by a patent, but one must assume that a marginal increase in observed patenting activity is equivalent to at least *some* increase in innovative activity. In this paper, I show that this assumption does not hold, at least when comparing firms and industries that face different levels of competitive pressure. In fact, changing competitive forces decrease a firm's incentives to innovate while simultaneously increasing her incentives to protect innovations that do occur: the propensity to invent decreases while the propensity to patent increases. Taken together, these two forces imply that observed patenting only matches underlying innovative output as a special case. In general, observed patenting is not a valid measure of how innovative effort responds to changing competitive pressure in an industry.

Below, I present a two–firm model of innovation, intellectual property (IP) protection, and competition. One firm, the Leader, exerts innovative effort to create a new product with multiple components. If she is successful, she then decides which of those components to protect with IP and what forms of IP to use. I consider patenting and trade secrecy here, though these methods can stand in respectively for any formal or informal method of IP protection more broadly. Trade secrecy determines the likelihood that a potential competitor, the Follower, is able to compete with the Leader in the market. Patenting determines the expected value of patent litigation and damages awarded if the Follower enter the market and infringe upon the Leader's patents. If these expected damages are large, the Follower declines to compete and the Leader secures a monopoly. If the expected damages are small, however, the Follower competes and the patents facilitate a transfer of profit from the Follower to the Leader. The model is agnostic about the actual form of competition that ensues; I simply specify how much producer surplus is competed away through entry and how the remaining surplus is divided between the two firms. I define "competitive pressure" as the decrease in producer surplus from competition, and I analyze its effects on the Leader's incentives to innovate and, separately, to patent, as well as on the Follower's incentives to compete in the market.

The model yields three key insights into the patent–innovation relationship. First, competitive pressure *decreases* the Leader's incentives to innovate but *increases* her incentives to protect that innovation with IP such as patents. Second, these two competing incentives combine to imply that innovation has a V–shaped relationship with competitive pressure, but that observed patenting has a broadly inverse–U shaped relationship: observed patenting is negatively correlated with underlying innovation for changes in competitive pressure. Third, this discrepancy can be addressed by using a product's average patent quality as a measure of innovation instead. When the incentive to innovate increases, the incentive to protect decreases and the Leader uses fewer patents. Since the quality of the marginal patent is lower than the average patent the Leader obtains, a decrease in patenting implies an increase in average patent quality.

These results follow from the fact that patenting here is a strategic, continuous decision: there are costs and alternative methods of protection; and therefore the Leader does not decide *whether* to patent but instead *how much*. The Leader can use her patents to create a monopoly outcome, but she can also use them to extract revenues from a competitor instead. By enriching the strategic space that patenting occurs in, and distinguishing the incentives to patent from the incentives to innovate, the model makes other predictions about how the relationships between innovation, IP usage, and observed market outcomes change as external factors vary. For instance, stronger patents elicit substitution away from trade secrets, making it easier for others to create a competing product and decreasing the likelihood of monopolization. Similarly, increased competitive pressure in the industry should be associated with increasingly monopolized markets as firms increase their IP usage across both methods.

The model can also inform innovation policy. The results capture a trade–off between increased welfare from innovation and the decreased efficiency of a patent–enforced monopoly, the classic patenting problem. Here, stronger patents can increase market entry by competitors, and thus consumer welfare, when they are used to extract licensing revenue from those competitors. But patents that are strong enough to guarantee a monopoly are welfare–reducing unless the innovation requires extensive effort. Ideally, patent policy should maximize the licensing revenue transferred from the competitor to the incumbent, as long as the competitor remains willing to enter the market and pay those transfers as a cost of business.

Relevant Literature

The paper contributes to three related bodies of research: the measurement of innovative activity, the relationship between innovation and competition, and the choice between various forms of IP. I will briefly outline each here, but the interested reader should consult the more exhaustive review articles cited below.

First and foremost, the paper refines our understanding of how patent statistics relate to the underlying rate of innovation that we wish to measure. Seminal work includes Pakes and Griliches (1980), Griliches (1981), and Pakes (1985), which relate a firm's patenting to its research and development (R&D) spending and its market value. Researchers took stock of the contemporary knowledge in Griliches et al. (1987) and Griliches (1990). The latter guided research in the following decades, noting that citation–weighted patent data were a promising avenue of future work and identifying Trajtenberg (1990) as a precursor. Citation–weighted patent counts became a standard way of weighting a patent by the amount or quality of innovation inherent within the patent, and further work refined the concept. Notably, Hall et al. (2005) validated the measure by relating citation–weighted patents to a firm's market value. The authors also developed a method of correcting for "truncation bias", to correct for the likelihood of unobserved citations that have not yet occurred. During this time, the National Bureau of Economic Research undertook a project to make these data more accessible to researchers at large and to link them with external data on firm financial performance, such as those available through CompuStat (Hall et al., 2001). Much of the work done since has relied on these data. For a more complete treatment of this topic, see Hall et al. (2010) and Nagaoka et al. (2010).

The measurement literature begins with the assertion that patents are a fundamentally sound proxy for innovative activity, and the bulk of work has attempted to refine and improve this measure. In this paper, I urge caution by documenting a case in which that fundamental assertion is not valid: patenting moves in the opposite direction from innovation and is thus not a useful proxy here. I hope that the alternative I propose, average patent quality, will prove to be a useful substitute in the researcher's toolbox.

A second broader literature explores the relationship between innovation and competition. One major strand of this research, usually attributed to Arrow (1962), argues that a monopolist would have no incentive to improve her product, since she lacks any market rivals. The other strand, beginning with Schumpeter (1942), argues that competitive firms would not make enough profit to recoup their innovation costs, so only firms with sufficient market power would innovate. Schumpeterian growth models arise from this line of thinking (Aghion and Howitt, 1992; Aghion et al., 2013). Recent work joins these opposing strands into one cohesive theory and predicts that competition and innovation have an inverse–U relationship (Aghion et al., 2005). Others have observed that Arrow's argument focuses on competition ex ante, while Schumpeter's concerns competition ex post, so the apparent contradiction is a bit of a misunderstanding (Whinston, 2011). For thorough reviews of this large body of work, see Gilbert (2006), Shapiro (2011), and Peneder (2012).

By separating the decision to innovate from the decision to use IP, the model below presents a more nuanced view of the innovation–competition relationship. Here, competitive forces in the industry influence a firm's innovative output and, separately, its IP usage. Success in innovation then determines whether a new market is created, while IP usage determines the competitiveness of that new market. In other words, observed market competition is not directly causal for innovation, and competitive industries can generate noncompetitive markets in IP–intensive sectors. It directly follows that measures of market power (such as the Lerner index) are outcomes, not determinants, of the innovation–IP process, and would likely only predict future innovative output through financial channels (*e.g.*, profits can be invested in future R&D) or through secular trends in industry–level competition.[1] Researchers in this area should consider more aggregated measures

[1] Instrumental variables are sometimes used to purge this endogeneity; see Aghion et al. (2005) and de Bettignies and Gainullin (2016) as examples.

of competition if available.

A third literature examines the patent system as just one of many methods to encourage innovation and to appropriate its returns. In most cases, patents are seen as substitutes for other means of appropriation. Alternatives include trade secrecy (Friedman et al., 1991; Arora, 1997; Denicolò and Franzoni, 2004; Hussinger, 2006; Moser, 2012; Zhang, 2012; Png, 2015), product complexity (Henry et al., 2016), and lead time (Cohen et al., 2000). However, survey evidence suggests that patents may be used in tandem with these and other methods (Levin et al., 1987; Cohen et al., 2000), and some work investigates the simultaneous use of multiple protection methods (Anton and Yao, 2004; Graham, 2004; Ottoz and Cugno, 2008; Belleflamme and Bloch, 2013; Graham and Hegde, 2014). The distinction between IP and innovation, crucial to the results of this paper, features in many of the papers cited above: *e.g.* Moser (2012) finds that many inventors who cannot patent effectively will nevertheless invent and rely on informal protection methods. Insights from the IP choice literature shape policy recommendations; see Hall (2007) and Shapiro (2011). For an extensive discussion of the choice between formal and informal intellectual property, see Hall et al. (2014).

One of the model's core assumptions is that products are "complex"; a product comprises multiple components, each of which may be covered by its own IP. Thus a single product can be protected by many patents and many secrets, and firms must decide how many of which kinds of IP to use. Holding innovation fixed, firms will increase IP usage or substitute between IP methods in response to external influences. These incentives are a core mechanism of the model, and they are independent of the direct incentives to innovate. These complex products exist in a wide variety of industries: Arora et al. (2003) estimate the average number of patents per innovation ranges from 2 in pharmaceuticals to 6.6 in industrial chemicals to 8.8 in rubber products, with an overall average of 5.6 patents per innovation. Premarin, an estrogen medication, is particularly illustrative. This pharmaceutical has been on the market since 1942, and the drug is still actively marketed and sold today. Even after patent expiration in the mid–twentieth century, no generic competitor has been cleared by the US FDA. The process used to extract the active ingredient in Premarin is protected by trade secret, even though the chemical makeup of the ingredient is commonly known and the relevant patents have expired. This secrecy has led to the extended absence of competition. The chemical compound, extraction process, and manufacturing methods can be viewed as separate components of a single product

(Premarin), and each of these products could have been patented or kept secret at the time of invention. Thus Premarin's manufacturer has reaped profits due to its mixed portfolio of IP. (Noonan, 2011; Lobel, 2013)

This paper uses patents and trade secrets to study the trade–offs and complementarities of formal and informal IP rights more broadly, and predicts when a firm will prefer a mixed IP basket. By including alternative protection methods in the firm's choice set, the model differentiates between forces that encourage protection generally and those that encourage patenting specifically. Thus firms will use both methods simultaneously because incentives to protect are high, not because of any inherent complementarity between the two methods. In the model, simultaneous usage will be most likely when incentives to protect generally are moderate, but the benefits of lead time (referred to below as "incumbency advantage") are low.

The patent system is not perfect, and clever use of patents allows for more than just a straightforward guarantee of monopoly. Often the mere threat of a patent lawsuit is enough to deter competitor entry, even if the validity of the patent is suspect (Lerner, 1995; Anton and Yao, 2004; Anton et al., 2006). Other strategic complications abound, such as defensive patenting, patent pools, or patent thickets, each with its own body of work. These interactions, while important, are not the focus of the current paper. Instead, I focus on how strategic patenting for entry deterrence and revenue extraction influences how we should measure underlying innovative activity, improves our understanding of the innovation–competition relationship, and clarifies the relationships between various forms of IP rights.

The paper proceeds as follows. Section 1 outlines the theoretical model, Section 2 details the analysis, and Section 3 presents core results. Section 4 analyzes the implications of the model for welfare and policymaking, and Section 5 discusses implications for measuring innovative activity. Section 6 concludes. Appendix A presents a technical foundation for the main model, and proofs for all results can be found in Appendix B.

1 Theoretical Environment

Consider a model in which a firm (the "Leader", feminine pronouns) attempts to create a new product. If successful, she then decides how to protect it from competition. Once protection is established, a competitor (the "Follower", masculine pronouns) may be able to enter the market. To enter, the Follower must gain complete knowledge of the product; only then can he enter the market if he wishes. Competition (or monopoly) ensues, and both the Leader and Follower receive payoffs.

The game occurs in three stages: Creation, Protection, and Entry. Figure 1 depicts each stage in sequence, with market profits given at each terminus. In the Creation stage, the Leader attempts to invent a new product that consists of a continuum of measure N of components.[2] She exerts innovation effort $e \geq 0$ at cost $c{\cdot}e$ and successfully creates a new product with probability $\phi(e)$. I assume $\phi(e)$ is an increasing, differentiable, and concave function with $\phi(0) = 0$ and $\lim_{e\to\infty} \phi(e) \leq 1$. If the Leader employs more resources in the innovative process, she is more likely to succeed at creating a new product. Note that her effort only improves the likelihood of creation; the quality and characteristics of the created product are taken as exogenous features of the product's technology.

If the Leader fails to innovate, the game ends. If she is successful, the Protection stage follows. The Leader then chooses n_p of the product's N components to patent, n_s to keep secret, and leaves the remaining $N - n_p - n_s$ unprotected and fully disclosed. Patents enable litigation but cost f each in filing fees, attorney retainers, etc. A larger collection of patents increases the value of legal action: in expectation, an entrant will pay the Leader damages (or negotiated licensing revenues) equal to $g(n_p)$ times the Leader's profits that are lost due to entry (a "lost profits" approach).[3] Note that $g(n_p)$ is the expected outcome of an uncertain legal process: patenting can be an imperfect instrument. I assume that $g(n_p)$ is an increasing, continuous, weakly concave, and twice differentiable function defined on $[0, N]$, with $g(0) = 0$.

Trade secrets reduce the likelihood that the Follower will learn enough to enter the market later in the game. The Leader pays s per component that she keeps secret, which represents the cost of non–disclosure agreements, the loss of efficiency from maintaining a "need–to–know basis" policy within a firm, or any costs associated with security.[4] The Follower learns all patented

[2] I use a continuum of components for tractability. Discretizing the number of components leads to similar results.

[3] Using a "reasonable royalties" approach, in which the court sets the royalty rate when lost profits cannot be determined, yields qualitatively similar, but less clean, results.

[4] By law under 18 U.S.C. § 1839 (3)(A), a firm must have "taken reasonable measures" to keep the information secret in order for it to meet the legal definition of a trade secret, and these measures certainly have a cost. If the definition is met, the owner is afforded legal protection against corporate espionage, but not independent invention or reverse

and unprotected components with certainty but may not learn all the secrets, which he must do in order to enter the market. The probability that the Follower learns all knowledge in a portfolio of secrets of size n_s (and can therefore enter the market if desired) equals $\ell(n_s)$. I assume that $\ell(n_s)$ is a decreasing, continuous, weakly convex, and twice differentiable function defined on $[0, N]$, with $\ell(0) = 1$.

The Protection stage begins as the Leader chooses an IP portfolio (n_p, n_s) and ends with the Follower's learning process. If the Follower does not learn enough to enter the market, the Leader earns monopoly profits π_M and the game ends. If the Follower does learn enough to enter, the game moves to the Entry stage and the Follower decides whether or not to enter the market. If he chooses not to enter, the game ends, the Leader earns monopoly profits π_M and the Follower's payoff is zero (the normalized value of his outside option). If the Follower does enter, however, then the two firms compete for shares of the available producer surplus. I define Δ as the loss in aggregate profits due to entry and $\lambda \in [0, 1]$ as the Leader's share of aggregate profits earned through competition.[5] Thus the Leader's market profits equal $\lambda(1 - \Delta)\pi_M$ and those of the Follower equal $(1 - \lambda)(1 - \Delta)\pi_M$.[6] Patent litigation or licensing then yields transfers in expectation from the Follower to the Leader, and expected payoffs from entry equal

$$\begin{aligned} \Pi_L &= \lambda(1 - \Delta)\pi_M + g(n_p)(1 - \lambda(1 - \Delta))\pi_M \\ \Pi_F &= (1 - \lambda)(1 - \Delta)\pi_M - g(n_p)(1 - \lambda(1 - \Delta))\pi_M \end{aligned} \tag{1}$$

since $(1 - \lambda(1 - \Delta))\pi_M$ is the size of the Leader's lost profits.[7]

The above framework is concise, but it makes an implicit assumption about the nature of the product and its components. Specifically, I assume *component independence*: the patent or secrecy status of one component does not influence the protection that could be afforded to another component via patent or secret.

engineering. Similar laws exist elsewhere, *e.g.* the Uniform Trade Secrets Act in Canada.

[5] Note that I do not assume $\Delta > 0$. Though it seems likely that entry will reduce aggregate profits (*e.g.* Cournot), it is not inevitable (*e.g.* differentiated goods markets).

[6] This parameterization is flexible and can accommodate many different assumptions about the market's structure. For instance, Cournot duopoly with linear demand and constant marginal costs is captured by $\Delta = \frac{1}{9}$ and $\lambda = \frac{1}{2}$.

[7] I assume risk–neutrality of both firms, so that the expected litigation outcome $g(n_p)$ serves as a credible threat point for the Leader in licensing negotiations. The model thus captures both licensing and litigation outcomes.

Assumption 1. *For a given n_s, $\ell(n_s)$ is constant for all n_p. For a given n_p, $g(n_p)$ is constant for all n_s.*

This assumption could be violated in two ways. First, there could be direct knowledge spillovers between components: disclosing knowledge of one component could make other components easier to discover. Second, a secret component could have higher patent potential than some patented components, and removing this secret would increase $g(n_p)$ because the Leader would swap it with a less effective patent. Assumption 1 asserts that protection is additive, addressing the first violation; and that patent potential is inversely related to secrecy potential across components, addressing the second. For the interested reader, I formally micro–found these two notions in Appendix A and show that they imply Assumption 1 and the shapes of $g(n_p)$ and $\ell(n_s)$ (concavity, etc.). For the body of the paper, though, I simply take $g(n_p)$ and $\ell(n_s)$ to be primitives of the model.

I assume that $g(N) < 1$, which implies that the Leader cannot be awarded aggregate damages higher than her lost profits, and that $|g''(n_p)| \geq \frac{g'(n_p)^2}{1-g(n_p)}$ and $\ell''(n_s) \geq \frac{\ell'(n_s)^2}{\ell(n_s)}$, which state that components are sufficiently heterogeneous in patent and secrecy potential and also act as sufficient–curvature assumptions, ensuring the Leader's problem is globally concave. If these assumptions do not hold, corner solutions can occur, but the main predictions of the model are unaffected.

2 Analysis

I solve the model by backwards induction. I first consider the entry decision of the Follower, conditional on full learning and a given patent portfolio n_p. I then consider the portfolio of IP chosen by the Leader and the consequences of her choices for incentives to innovate.

Entry

Conditional on learning all components, the Follower will enter the market if his expected profit is strictly greater than zero, which is true when

$$\Pi_F = (1-\lambda)(1-\Delta)\pi_M - g(n_p)(1-\lambda(1-\Delta))\pi_M > 0 \rightarrow g(n_p) < \frac{(1-\lambda)(1-\Delta)}{1-\lambda(1-\Delta)}$$

So the Follower enters when the number of patents is smaller than some threshold $\overline{n_p}$.[8] If such a threshold exists, it is defined by $g(\overline{n_p}) = \frac{(1-\lambda)(1-\Delta)}{1-\lambda(1-\Delta)}$. Otherwise, $\overline{n_p}$ is taken to be infinite, or at least strictly greater than N.[9]

Lemma 1. *The Follower only enters the market if $n_p < \overline{n_p}$. The Leader enjoys a monopoly if $n_p \geq \overline{n_p}$.*

Note that if $\frac{(1-\lambda)(1-\Delta)}{1-\lambda(1-\Delta)} > g(N)$, then exclusion is impossible. Since $g(N) < 1$, a necessary condition for the ability to exclude is $\Delta > 0$; *i.e.*, if total profits shrink due to entry:

Lemma 2. *Exclusion of the Follower is possible only if the firms' total surplus shrinks due to entry.*

This result only follows because of the assumption that $g(N) < 1$: the Leader can never recoup all of her lost profits through the legal system. While this is probably a reasonable assumption, there could be legal environments in which the Leader is able to recover more than what was lost. Further, environments exist in which $\Delta < 0$ and entry increases surplus, such as markets with sufficiently differentiated products. In these markets, if $g(N) \geq 1$ the Leader would never exclude, regardless of whether Δ is positive or negative.[10] The $g(N) < 1$ case is more strategically rich for the Leader, so I retain that assumption below.

Incorporating the learning mechanism (and accounting for protection costs and Lemma 1), the Leader's expected payoff from an IP portfolio (n_p, n_s) equals

$$\Pi_L = \begin{cases} \begin{array}{l} \ell(n_s)\,(\lambda(1-\Delta)\pi_M + g(n_p)(1-\lambda(1-\Delta))\pi_M) \\ +(1-\ell(n_s))\pi_M - n_p f - n_s s, \end{array} & n_p < \overline{n_p} \\ \pi_M - n_p f - n_s s, & n_p \geq \overline{n_p} \end{cases}$$

In the Protection stage, the Leader chooses n_p and n_s to maximize Π_L.

[8] This threshold creates a discontinuity in the resulting comparative statics. If, instead, Δ and λ were sufficiently uncertain to the Leader, then $\overline{n_p}$ would be as well, and the comparative statics would smooth out.

[9] One could envision a model where the Follower must pay some fixed entry cost F, say to innovate or build production capacity. In this case, $g(\overline{n_p}) = \frac{(1-\lambda)(1-\Delta)}{1-\lambda(1-\Delta)} - (1-\lambda(1-\Delta))\frac{F}{\pi_M}$ and the rest of the analysis is unchanged.

[10] Using $n_p = \overline{n_p}$, a Leader will retain π_M, but using $n_p = \overline{n_p} - \epsilon$ will yield almost the whole surplus, which is larger than π_M here by assumption.

Protection

Note that for any choice of $n_p > \overline{n_p}$, the optimal n_s is 0 and a choice of $\overline{n_p}$ is preferred. Thus if the Leader chooses to exclude, she does so optimally by choosing $(n_p, n_s) = (\overline{n_p}, 0)$. If the Leader instead decides to allow the possibility of entry, she chooses (n_p, n_s) to solve

$$\max_{n_p, n_s} \pi_M - \ell(n_s)(1 - g(n_p))(1 - \lambda(1 - \Delta))\pi_M - n_p f - n_s s$$

$$\text{subject to} \quad \begin{array}{ll} n_p \geq 0 & n_p < \overline{n_p} \\ n_s \geq 0 & n_p + n_s \leq N \end{array}$$

Note that here, the Leader is effectively choosing a fraction $1 - \ell(n_s)(1 - g(n_p))(1 - \lambda(1 - \Delta))\pi_M$ of profits to protect; call this the protection level. Given $g(N) < 1$ and the sufficient curvature assumptions, the problem is globally concave and the first–order conditions characterize an optimum. Rearranged, the first–order conditions yield

$$\frac{s}{f} = \frac{-\ell'(n_s^*)}{\ell(n_s^*)} \frac{1 - g(n_p^*)}{g'(n_p^*)} = \frac{\mathrm{MP}_s}{\mathrm{MP}_p} \tag{Ratio}$$

$$\begin{aligned} \frac{s + f}{\pi_M} &= \left[\ell(n_s^*) g'(n_p^*) - \ell'(n_s^*)(1 - g(n_p^*))\right](1 - \lambda(1 - \Delta)) \\ &= \mathrm{MP}_p + \mathrm{MP}_s \end{aligned} \tag{Sum}$$

where $\mathrm{MP}_{p,s}$ is the marginal increase in the protection level from one additional patent or secret:

$$\begin{aligned} \mathrm{MP}_p &= \ell(n_s) g'(n_p)(1 - \lambda(1 - \Delta))\pi_M \\ \mathrm{MP}_s &= \ell'(n_s)(1 - g(n_p))(1 - \lambda(1 - \Delta))\pi_M \end{aligned}$$

Equations (Ratio) and (Sum) define equilibrium, assuming entry is not forbidden. The Leader compares the payoff from this optimal strategy to $\pi_M - \overline{n_p} f$, the payoff from Exclusion, and then chooses the profit–maximizing strategy. Incorporating equations (Ratio) and (Sum) into this comparison yields the set of equilibrium portfolios with entry (n_p^*, n_s^*) that are preferred to Exclusion, which are the ones that satisfy

$$\overline{n_p} f \geq \ell(n_s^*)(1 - g(n_p^*))(1 - \lambda(1 - \Delta))\pi_M + n_p^* f + n_s^* s \tag{Entry}$$

Thus expressions (Ratio), (Sum), and (Entry) fully characterize the Leader's problem, which is depicted graphically in Figure 2. Under entry, the optimal portfolio is given by the intersection of (Ratio) and (Sum), the dashed lines. If this point lies to the left of the solid curve, then (Entry) is satisfied and entry occurs, as it does in Figure 2. Otherwise, the Leader would exclude the Follower. In either case, conditions (Ratio), (Sum), and (Entry) characterize the unique equilibrium portfolio chosen by the Leader.

The condition that $n_p + n_s \leq N$ can be depicted in Figure 2 by drawing a line from $(N, 0)$ to $(0, N)$. If the solution to equations (Ratio) and (Sum) lies beyond this boundary, then the Leader will choose some portfolio along $n_p + n_s = N$, but it is straightforward to show that the qualitative conclusions of the model will be the same.[11] For simplicity, I make the following simplifying assumption:

Assumption 2. *The parameters of the model are such that $n_p + n_s \leq N$ never binds. In other words, N is sufficiently large.*

Creation

Substituting the Leader's optimal portfolio choice into her expected payoff function Π_L yields her ex ante expected profit from an innovation:

$$\Pi_L^* = \max \begin{Bmatrix} \pi_M - \ell(n_s^*)(1 - g(n_p^*))(1 - \lambda(1 - \Delta))\pi_M - n_p^* f - n_s^* s \\ \pi_M - \overline{n_p} f \end{Bmatrix} \tag{2}$$

which the Leader takes as constant in the Creation stage. She chooses optimal innovation effort e^* to equate her marginal expected profit from effort with its marginal cost:

$$\phi'(e^*)\Pi_L^* = c \tag{3}$$

which immediately implies that innovation is increasing in Π_L^*.

3 Results

With the model solved, I now determine how the economically relevant parameters of model determine the Leader's choice of IP portfolio, her ex ante

[11] Solving the Leader's constrained optimization problem yields two first–order conditions. When $n_p + n_s \leq N$ is slack, the conditions are identical to (Ratio) and (Sum). When the constraint does bind, one condition is a more general form of (Ratio) and the other is simply $n_s = N - n_p$, which slopes downward just as (Sum) does. Condition (Entry) is unchanged, so the model makes the same qualitative predictions as the unconstrained version considered in the body of the paper.

profit, and her innovative effort. I then use the results to explore the relationships between IP usage, actual innovation, and changes in industry characteristics or economic policy.

I will principally consider the effects of two industry characteristics (competitive pressure and incumbency advantage) and one policy instrument (patent strength, to be defined later). Since Δ equals the loss in profits due to entry, I interpret it as a measure of the competitive pressure within an industry. If Δ is close to zero, for example, then market entry will leave firms' profits mostly intact. If instead Δ is close to one, then market entry will eliminate most of the firms' profits and market participants face increasing pressure to deter competitor entry. Similarly, λ measures the importance of incumbency advantage in the industry, whether it be due to brand loyalty, a lead time in production, etc. If λ is low, then potential entrants can destroy an incumbent's profits easily. If λ is high, then competitors stand to gain very little from entering a market and the incumbent has less to fear. Thus, λ also measures the competitive nature of an industry.

3.1 IP Portfolio Choice

Equations (Ratio), (Sum), and (Entry) map industry parameters into IP portfolio choices. These portfolios can be broadly classified as Empty (no protection is used), Licensing (only patents are used, but fewer than $\overline{n_p}$), Concealment (only secrets are used), Mixed (both patents and secrets are used), and Exclusion ($\overline{n_p}$ patents are used). In other words, the Leader chooses whether to exclude a potential entrant or to allow entry, and then how to protect her innovation should entry occur. Proposition 1 details how the industry parameters Δ and λ shape the Leader's optimal IP portfolio.

Proposition 1. *There exist thresholds $\lambda_1(\Delta) \leq \lambda_2(\Delta) \leq \lambda_3(\Delta)$ and $\Delta_2 \leq \Delta_1 \leq \Delta_3$ such that:*

- *For $\lambda < \lambda_1(\Delta)$, the Leader chooses $n_p \in (0, \overline{n_p})$ and $n_s > 0$ (Mixed);*
- *For $\lambda \in [\lambda_1(\Delta), \lambda_2(\Delta))$,*
  - *the Leader chooses $n_p = 0$ and $n_s > 0$ (Concealment) if $\frac{s}{f} < \frac{-\ell'(0)}{g'(0)}$;*
  - *the Leader chooses $n_p \in (0, \overline{n_p})$ and $n_s = 0$ (Licensing) if $\frac{s}{f} > \frac{-\ell'(0)}{g'(0)}$;*

- *For $\lambda \in [\lambda_2(\Delta), \lambda_3(\Delta))$, the Leader chooses $n_p = 0$ and $n_s = 0$ (Empty);*
- *For $\lambda > \lambda_3(\Delta)$, the Leader chooses $n_p = \overline{n_p}$ (Exclusion).*

Further,

- *$\lambda_3(\Delta)$ is decreasing, $\lambda_3(0) = 1$, and $\lambda_3(\Delta_3) = 0$;*
- *$\lambda_2(\Delta)$ is increasing until Δ_2, and equals $\lambda_3(\Delta)$ afterwards; and*
- *$\lambda_1(\Delta)$ is increasing until Δ_1, and equals $\lambda_3(\Delta)$ afterwards.*

Figure 3 depicts the regions described by Proposition 1. The $\lambda_3(\Delta)$ line divides the Δ–λ plane into Entry and Exclusion regions. Entry occurs for lower Δ or λ, and Exclusion occurs for higher values. Under Entry, the $\lambda_1(\Delta)$ and $\lambda_2(\Delta)$ lines distinguish between different kinds of portfolios: below $\lambda_1(\Delta)$, the Leader uses both patents and secrets (Mixed); above $\lambda_2(\Delta)$, the Leader does not protect at all (Empty), and in between, the Leader uses only one kind of IP (Licensing or Concealment).

The competitive nature of the industry determines the Leader's portfolio choice. For industries with relatively little surplus loss from entry, protection is relatively weak, especially if there is a large incumbency advantage. But if entry destroys more surplus, protection increases as well. Eventually, the loss from entry is so great that the Leader will always exclude the Follower if able.

Note that, *ceteris paribus*, a Licensing portfolio and an Exclusion portfolio will look markedly different even though both are comprised solely of patents. Obviously, $n_p < \overline{n_p}$ in any non–Exclusion portfolio since a portfolio with $\overline{n_p}$ patents will exclude the Follower (recall Lemma 1). But condition (Entry) implies a stronger statement: under entry, the magnitude of $\overline{n_p} - n_p$ *cannot* be made arbitrarily small by adjusting the parameters s, f, Δ, or λ, so IP portfolios will not smoothly transition to and from Exclusion:

Corollary 1. *For a given set of functions $g(\cdot)$ and $\ell(\cdot)$, there exists $\epsilon > 0$ such that the difference $\overline{n_p} - n_p > \epsilon$ for any Licensing or Mixed portfolio. In particular, the difference cannot be made arbitrarily small by varying the parameters s, f, Δ, and λ.*

Proposition 1 describes the Leader's choice in the Protection stage, conditional on having successfully innovated. These choices will determine the expected profit from innovation and the incentives to do so in the Creation stage of the game. Industry characteristics will therefore influence market outcomes through two channels: their effects on IP portfolio choice should innovation occur, and their effects on the incentive to innovate in the first place. I explore these comparative statics and their interaction below.

3.2 Competitive Pressure: The Effects of Δ

Conditional on successful innovation, competitive pressure on the Leader is an incentive to protect that innovation. But that pressure also discourages Follower entry, since entry shrinks the potential surplus. Therefore, higher Δ leads to more use of both kinds of IP until Exclusion becomes the more viable option, and further increases in Δ simply make Exclusion feasible with a smaller portfolio.

Proposition 2. *There is a Δ_{Bd} such that:*

- *For $\Delta < \Delta_{Bd}$, the Leader uses an Entry portfolio and n_p and n_s are increasing in Δ; and*
- *For $\Delta \geq \Delta_{Bd}$, the Leader uses an Exclusion portfolio and $n_p = \overline{n_p}$ is decreasing in Δ.*

Already, the competitive nature of the industry has a nonlinear effect on IP usage (conditional on innovation). Leaders increase their patenting and secrecy as the threat of Follower entry increases. Eventually, the threat is large enough to the Leader and the benefit is small enough to the Follower that Exclusion occurs, patenting jumps upward, and secrecy drops to zero. Beyond this point, it becomes feasible to accomplish that Exclusion with fewer and fewer patents.

Proposition 2 also implies a counterintuitive relationship between competitive pressure at the industry level and competitive outcomes at the market level. Since Entry is less likely when n_s is higher and impossible when $n_p = \overline{n_p}$, industries with high competitive pressure are more likely to be made of monopolized markets when IP protection is possible.

Corollary 2. *The probability of a monopoly outcome is increasing in Δ.*

Next we arrive at our first discrepancy between the propensity to use IP and the incentives to innovate: the effects of competitive pressure Δ

on the Leader's profit Π_L^*, innovative effort e^*, and innovative output $\phi(e^*)$ are the opposite of its effects on ex–post patenting n_p. Since entry destroys more profit, innovation is less lucrative, but protecting successful innovations with IP becomes more important. Once Exclusion occurs, more competitive pressure makes that Exclusion portfolio cheaper, increasing profits and the incentives to innovate accordingly.

Proposition 3. *For $\Delta < \Delta_{Bd}$, Π_L^*, e^*, and $\phi(e^*)$ are decreasing in Δ. For $\Delta \geq \Delta_{Bd}$, Π_L^*, e^*, and $\phi(e^*)$ are increasing in Δ.*

Again, industry characteristics have a nonlinear effect on innovative activity. Competitive pressure decreases aggregate profit in the market, so less effort is exerted to innovate and create that market. However, if that pressure means Exclusion is guaranteed, the Leader's monopoly will also be ensured and incentives to innovate rise.

Figure 4 depicts Propositions 2 and 3 together, graphing both innovation $\phi(e^*)$ and ex–post patenting n_p as functions of Δ. The curvature of $\phi(\cdot)$ gives innovation a V–shape, while ex–post patenting rises, jumps up, and then falls. The break point occurs at Δ_{Bd}, the value of Δ for which (Entry) is exactly satisfied.

Propositions 2 and 3 do not necessarily imply that IP usage and innovation are inversely related in the data, however. Since innovation must be successful for patents to exist in the first place, increased innovation incentives could lead to an overall increase in patenting, at least in an expected sense. In other words, what we observe in practice is not n_p or $\phi(e^*)$, but their product $\phi(e^*)n_p$: the probability of invention times the number of patents used to protect that invention. I call this *observed patenting*, as it is what we actually observe in the data. When researchers use patenting to proxy for innovation, observed patenting $\phi(e^*)n_p$ is what they are measuring.[12] It is crucial, then, that observed patenting is correlated with actual innovation $\phi(e^*)$ if it is to be a useful proxy.

I denote observed patenting as $n_p^o \equiv \phi(e^*)n_p$. Consider how this measure changes with respect to Δ:

$$\frac{\partial n_p^o}{\partial \Delta} = \underbrace{\phi(e^*)\frac{\partial n_p}{\partial \Delta}n_p}_{\text{Protection Effect}} + \underbrace{\frac{\partial \phi(e^*)}{\partial \Delta}}_{\text{Innovation Effect}} \tag{4}$$

[12]Note that the same results also hold for ex post patent protection $g(n_p^*)$ and expected patent protection $\phi(e^*) \cdot g(n_p^*)$, which I discuss below.

For this measure to accurately capture changes in innovation, it must be that the first term is sufficiently small relative to the second term: changes in ex–post IP usage cannot outweigh changes in the underlying incentives to innovate. Proposition 4 describes this relationship between n_p^o and Δ.

Proposition 4. *There are $\Delta_{En} \leq \Delta_{Ex}$ such that:*

1. *For $\Delta < \Delta_{Bd}$ and $\Delta < \Delta_{En}$, n_p^o is increasing in Δ;*
2. *For $\Delta < \Delta_{Bd}$ and $\Delta \geq \Delta_{En}$, n_p^o is decreasing in Δ;*
3. *For $\Delta \geq \Delta_{Bd}$ and $\Delta < \Delta_{Ex}$, n_p^o is increasing in Δ; and*
4. *For $\Delta \geq \Delta_{Bd}$ and $\Delta \geq \Delta_{En}$, n_p^o is decreasing in Δ.*

2 and 3 are possible only if $|\phi''|\frac{\phi}{\phi'^2}$ is sufficiently low at Δ_{Bd}.

Proposition 4 is not encouraging: the only cases where observed patenting n_p^o and actual innovation $\phi(e^*)$ could be positively correlated are near the Exclusion threshold; *i.e.*, between Δ_{En} and Δ_{Bd} or between Δ_{Bd} and Δ_{Ex}. Worse, these cases will not exist if $\phi(e)$ is sufficiently concave; *i.e.*, when diminishing returns to innovative effort are severe. And since the Exclusion threshold imparts a discontinuity in observed patenting but not innovation (see Corollary 1), this relationship is likely to be a weak one.

Figure 5 shows the results of Proposition 4. In this particular parameterization, regions do exist where observed patenting n_p^o moves with accordance with underlying innovation. Observed patenting attains local maxima at Δ_{En} under Entry and at Δ_{Ex} under Exclusion. Innovation, meanwhile, has a V–shaped relationship with competitive pressure. Here, patenting is only a useful measure of underlying innovation for Δ between Δ_{En} and Δ_{Ex}. Elsewhere, the two move in opposite directions.

3.3 Incumbency Advantage: The Effects of λ

Whereas competitive pressure increases the need for protection and decreases incentives to innovate, a first–mover or incumbency advantage lessens the need for protection and makes innovation more profitable. With a high λ, the Leader is less threatened by entry and will therefore rely less on IP to guard against losses, and fewer patents are needed to dissuade the Follower from entry.

Proposition 5. *There is a λ_{Bd} such that:*

- *For $\lambda < \lambda_{Bd}$, the Leader uses an Entry portfolio and n_p and n_s are decreasing in λ; and*
- *For $\lambda \geq \lambda_{Bd}$, the Leader uses an Exclusion portfolio and $n_p = \overline{n_p}$ is decreasing in λ.*

Proposition 5 implies a relationship between the typical incumbency advantage in an industry and outcomes in the market. A higher incumbency advantage decreases protection under Entry, so monopoly is less likely. But eventually, the Follower is too easily dissuaded from entering the market, Exclusion occurs, and monopoly ensues.

Corollary 3. *For $\lambda < \lambda_{Bd}$, the probability of a monopoly outcome is decreasing in λ. For $\lambda \geq \lambda_{Bd}$, monopoly always occurs.*

Proposition 5 contrasts with the effect of incumbency advantage on the incentives to innovate. Whether or not entry occurs, incumbency advantage increases the expected profit of an innovation, so innovative effort increases as well. In all cases, then, IP and innovation move in opposite directions.

Proposition 6. Π_L^*, *e^*, and $\phi(e^*)$ are continuous and increasing in λ.*

Just as with competitive pressure, however, we need to determine the effects of incumbency advantage on observed IP usage in order to quantify the overall effect.

Proposition 7. *There are $\lambda_{En} \leq \lambda_{Ex}$ such that:*

1. *For $\lambda < \lambda_{Bd}$ and $\lambda < \lambda_{En}$, n_p^o is increasing in λ;*
2. *For $\lambda < \lambda_{Bd}$ and $\lambda \geq \lambda_{En}$, n_p^o is decreasing in λ;*
3. *For $\lambda \geq \lambda_{Bd}$ and $\lambda < \lambda_{Ex}$, n_p^o is increasing in λ; and*
4. *For $\lambda \geq \lambda_{Bd}$ and $\lambda \geq \lambda_{Ex}$, n_p^o is decreasing in λ.*

2 is possible only if $|\phi''|\frac{\phi}{\phi'^2}$ is sufficiently high at λ_{Bd}, and 3 is possible only if $|\phi''|\frac{\phi}{\phi'^2}$ is sufficiently low at λ_{Bd}.

Proposition 7 presents a less dire picture than Proposition 4. Since innovation is increasing in λ under Entry, observed patenting now correlates with innovation for most λ that elicit Entry, but again the relationship is imperfect.

3.4 Patent Strength: The Effects of Scaling $g(\cdot)$

The model also sheds light on how innovation policy can impact IP usage, innovative effort, and market outcomes. Here I focus on patent strength, defined as the expected legal payoff from exercising a patent right. Patent strength encompasses many possible policy considerations: How likely is a granted patent to be upheld in court? How generous are damage awards? How broad is a patent's coverage? For how long is the patent in effect? To model these changes, I suppose that policymakers can influence aggregate patent strength by costlessly scaling $g(n_p)$ by a factor γ so that the total litigation value of a patent portfolio is now $\gamma \cdot g(n_p)$.[13]

The effects of patent strength on IP usage vary. With weak patents, an Exclusion portfolio will be prohibitively expensive, the Leader will use patents to extract revenue, and any increases in patent strength elicit a substitution away from secrecy. If instead patents are fairly strong, then the Leader can Exclude the Follower cheaply and any further increase in patent strength allows the Leader to accomplish Exclusion with fewer patents. Proposition 8 details these relationships.

Proposition 8. *There is a γ_{Bd} such that:*

- *For $\gamma < \gamma_{Bd}$, the Leader uses an Entry portfolio, n_p is increasing in γ, and n_s is decreasing in γ; and*
- *For $\gamma \geq \gamma_{Bd}$, the Leader uses an Exclusion portfolio and $n_p = \overline{n_p}$ is decreasing in γ.*

It is intuitive that increased patent strength leads to increased patenting and decreased secrecy, but this also implies that stronger patents can actually decrease the likelihood of a monopoly, as long as the Leader still prefers Entry over Exclusion.

Corollary 4. *For $\gamma < \gamma_{Bd}$, the probability of a monopoly outcome is decreasing in γ. For $\gamma \geq \gamma_{Bd}$, monopoly always occurs.*

[13]The discrepancy between the effective life of a patent and that of a trade secret is key to a firm's IP decision: see Denicolò and Franzoni (2004) as an example. I do not model time here, but one can easily read the parameters of the model as the net present value of per–period flows. Explicitly modeling the finite life of a patent and the (potentially) infinite life of a secret does not change the results below.

Thus stronger patents can encourage market entry and competition, as long as they are not made overly strong; otherwise monopoly is guaranteed.

Stronger patents also increase the Leader's incentives to innovate, since the Leader finds it easier to forbid Entry but also to extract revenue when Entry occurs.

Proposition 9. *Π_L^*, e^*, and $\phi(e^*)$ are continuous and increasing in γ.*

Here again, observed patenting does not necessarily correlate with the Leader's innovative effort. If IP protection is imperfect, stronger patents allow the Leader to specialize her portfolio, which increases information disclosure, market entry, and revenue to the party that bore the cost of innovation. These forces increase patenting and innovation in tandem. However, if patents already guarantee perfect IP protection (*i.e.*, they guarantee a monopoly outcome), then stronger patents just accomplish the same goal more cheaply, and innovation increases while patenting decreases. Taken together, these effects imply that patenting has a broadly inverse–U relationship with patent strength, even while innovation is monotonic.

Proposition 10. *There is a γ_{Ex} such that:*

1. *For $\gamma < \gamma_{Bd}$, n_p^o is increasing in γ;*
2. *For $\gamma \geq \gamma_{Bd}$ and $\gamma < \gamma_{Ex}$, n_p^o is increasing in γ; and*
3. *For $\gamma \geq \gamma_{Bd}$ and $\gamma \geq \gamma_{Ex}$, n_p^o is decreasing in γ.*

2 is only possible if $|\phi''|\frac{\phi}{\phi'^2}$ is sufficiently low at γ_{Bd}.

Thus if a policy is enacted and patenting rates fall, it does not necessarily imply that innovation has fallen as well. The effects of a change in policy depend not only on how firms change their innovative effort but also on how they allocate resources to protection.

The above results are a first step in understanding the effects of policy here. To conduct a more thorough, normative analysis, the model must be augmented with information on welfare outcomes for consumers and how welfare relates to the market's existence and its competitiveness.

4 Policy and Welfare

Consumers derive welfare from two channels in the model. First, they directly benefit from the market's existence: the new product created by the Leader is purchased because it represents an improvement in the consumers' *status quo* levels of utility. Second, the extent to which consumers can participate in this newly–created market depends on its efficiency: if firms possess significant market power, consumers cannot easily access the potential utility increases from the new product. Optimal policy must then balance the gains from market existence with the deadweight loss of imperfect competition. Below, I model these channels explicitly and determine the welfare–maximizing level of patent strength. Other policy changes could be modeled and analyzed in the same way.

To measure welfare, I parameterize consumer surplus under monopoly and under competition. I assume that consumer surplus is equal to $\sigma\pi_M$ in a monopolized market and equal to $(\sigma + \psi\Delta)\pi_M$ in a market with Follower entry. The factor σ ties consumer surplus to the inherent value of the product (captured by the magnitude of π_M), and the factor ψ measures the increased surplus from entry as a share of the loss in producer surplus. I assume that $\psi > 1$ so that market entry increases efficiency: entry transfers $\Delta\pi_M$ from firms to consumers and also creates $(\psi - 1)\Delta\pi_M$ new surplus through increased market efficiency, resulting in a total gain of $\psi\Delta\pi_M$ to consumers. If innovation does not occur, consumer surplus is equal to the consumers' outside option, which I normalize to zero.

Expected consumer surplus is therefore a function of the Leader's IP portfolio choice and innovative effort. Specifically,

$$\mathrm{E[CS]} = \begin{cases} \phi(e)(\sigma + \ell(n_s)\psi\Delta)\pi_M, & n_p < \overline{n_p} \\ \phi(e)\sigma\pi_M & n_p \geq \overline{n_p} \end{cases} \tag{5}$$

When the Leader would choose an Entry portfolio, consumers gain welfare if innovation is successful (through $\phi(e)$) and more if the Follower successfully learns the Leader's secrets (through $\ell(n_s)$). When the Leader would choose Exclusion instead, consumer welfare is only derived from increased odds of market creation. The Leader's effort e and IP choices are functions of industry parameters and economic policy, as before.

A simple comparison of the two expressions then determines which parameters and policies induce the Leader to choose a welfare–maximizing IP portfolio. Parameters that induce an Entry portfolio are optimal when

$$\ell_{En} \cdot \frac{\psi\Delta}{\sigma} > \frac{\phi_{Ex} - \phi_{En}}{\phi_{En}} \tag{6}$$

where $\ell_{En} = \ell(n_s)$ under Entry and $\phi_{En,Ex} = \phi(e)$ under Entry and Exclusion, respectively. The intuition behind (6) is clear: inducing an Entry portfolio is preferred when the probability of entry ℓ_{En} times the percentage increase in consumer surplus upon entry $\frac{\psi\Delta}{\sigma}$ is greater than the percentage increase in the probability of innovation $\frac{\phi_{Ex}-\phi_{En}}{\phi_{En}}$. Whether (6) is actually satisfied depends on the equilibrium values of each line in (5).

Consider the optimal choice of patent strength γ. Additional patent strength increases consumer surplus under both Entry and Exclusion:

$$\frac{\partial \mathrm{E[CS]}}{\partial \gamma} = \begin{cases} (\sigma + \ell_{En}\psi\Delta)\pi_M \frac{\partial \phi_{En}}{\partial \gamma} + \phi_{En}\psi\Delta\pi_M \ell'_{En} \frac{\partial n_s}{\partial \gamma}, & n_p < \overline{n_p} \\ \sigma\pi_M \frac{\partial \phi_{Ex}}{\partial \gamma}, & n_p \geq \overline{n_p} \end{cases} \tag{7}$$

both lines of which are positive. Assuming an Entry portfolio, the optimal γ is just below γ_{Bd}; define ℓ_{Bd} and ϕ_{Bd} accordingly. Assuming an Exclusion portfolio implies that the optimal γ is as high as possible. Since $\phi \leq 1$, there must be some upper bound beyond which higher γ has no additional effect; define that bound as γ_{Max} and the corresponding ϕ as $\phi_{Max} = \lim_{e\to\infty} \phi(e)$. We can then use (6) to determine the optimal γ overall.

Proposition 11. *Optimal patent strength equals γ_{Bd} if $\ell_{Bd} \cdot \frac{\psi\Delta}{\sigma} > \frac{\phi_{Max}-\phi_{Bd}}{\phi_{Bd}}$ and equals γ_{Max} otherwise.*

In words, patents should be of moderate strength unless the increase in innovative effort from monopolization is substantial. Ideally, patents should not actually induce a monopoly outcome, but instead should elicit market entry and revenue extraction by the Leader.

Proposition 11 also captures the inherent tension between innovation policy and competition policy. Suppose innovation policy's mandate is to maximize innovative effort in the economy, and suppose that competition policy's mandate is to maximize the likelihood of competitive outcomes in markets.[14] In the context of the model, then, innovation policy intends to maximize $\phi(e)$

[14]These mandates are oversimplifications, to be sure, but they serve to capture the basic tension.

and competition policy intends to maximize $\mathbf{1}(n_p < \overline{n_p}) \cdot \ell(n_s)$. Using Proposition 9 and Corollary 4, choosing γ_{Max} fulfills innovation policy's mandate while γ_{Bd} fulfills the competition authority's.

Of course, no policy can be carried out without properly measuring the objective at hand, which in this case is innovative effort. Propositions 4, 7, and 10 show that raw patent counts are, in general, insufficient as a measure of innovation. Section 5 argues that even properly–weighted patent counts suffer the same flaw and proposes an alternative patent–based metric that can perform better.

5 Measuring Innovation

The model presented above separates a firm's innovation decision from its protection decision, and enriches the set of methods available to the firm for that protection. This division has important consequences for the measurement of innovative activity: the most widely used measure of innovation, quality–adjusted patent counts, can be decreasing even as innovative activity increases.

We have actually already seen the relationship between unadjusted patent counts and innovation. Propositions 3, 6, and 9 describe how the Leader's innovative effort $\phi(e)$ is affected by exogenous factors, while Propositions 4, 7, and 10 detail how observed patenting activity n_p^o varies with these factors. For patenting to be a useful measure of innovation, it should move in the same direction as innovation as exogenous factors change. The results above are not encouraging: for changes in competitive pressure Δ, patenting is positively related to innovation only near Δ_{Bd} (and only if $|\phi''|\frac{\phi}{\phi'^2})$ is low); for incumbency advantage λ, patenting is always positively related to innovation for low λ, only sometimes for moderate λ, and is never related for high λ; and for patent strength γ, patenting is positively related to innovation only for low to moderate values. In the context of this model, at least, raw patent counts are an unreliable measure of innovation at best. However, raw counts are not typically used directly to measure innovation; quality–adjusted totals are. Below, I discuss how observed quality–adjusted patent counts relate to innovation in the context of the model, show when they can be unreliable, and propose an alternative metric. I use exogenous variation in competitive pressure Δ as an instructive example here.

5.1 Quality–adjusted patenting

Since innovations vary in their quality and value, researchers usually try to weight the measure by some metric of patent quality, such as citations by other patents or the patentee firm's financial performance. In the context of this paper, these quality adjustments make the problems described in Proposition 4 less severe, but the basic issue persists.

Since $g(n_p)$ measures the value of a set of patents to the firm, and each patent within that set covers some innovative component of the product, $g(n_p)$ can be taken as a quality–adjusted measure of the innovative content of a patent portfolio.[15] Its observed counterpart will then be $g^o(n_p)$, defined as $\phi(e^*)g(n_p)$, and we can determine how $g^o(n_p)$ varies over different values of Δ, and how that variation compares to changes in innovation.

Since $g(\cdot)$ is simply a monotone transformation, the same problem will persist: increased competitive pressure increases the incentive to protect, so $g(n_p)$ will rise just as n_p does. However, since $g(\cdot)$ is concave, this increase will be smaller at higher values of n_p. For $g^o(n_p)$, the innovation effect will be emphasized and the protection effect will be muted. Thus using a quality–adjusted measure will lead to better performance.

Proposition 12. *Quality–adjusted patenting $g^o(n_p)$ is a better measure of innovation than raw patenting n_p^o: there are quality–adjusted $\Delta_{En}^Q \leq \Delta_{Ex}^Q$, with $\Delta_{En}^Q < \Delta_{En}$ and $\Delta_{Ex}^Q > \Delta_{Ex}$, such that:*

1. *For $\Delta < \Delta_{En}^Q$ and $\Delta < \Delta_{Bd}$, $g^o(n_p)$ is increasing in Δ;*
2. *For $\Delta \geq \Delta_{En}^Q$ and $\Delta < \Delta_{Bd}$, $g^o(n_p)$ is decreasing in Δ;*
3. *For $\Delta < \Delta_{Ex}^Q$ and $\Delta \geq \Delta_{Bd}$, $g^o(n_p)$ is increasing in Δ; and*
4. *For $\Delta \geq \Delta_{En}^Q$ and $\Delta \geq \Delta_{Bd}$, $g^o(n_p)$ is decreasing in Δ.*

2 and 3 are possible only if $|\phi|\frac{\phi}{\phi'^2}$ is sufficiently low at Δ_{En}^Q and Δ_{Ex}^Q, respectively.

Thus adjusting for patent quality improves the range of Δ for which the measure correlates with innovation; the range's lower bound decreases

[15] Strictly speaking, $g(n_p)$ measures the share of the product's market value attributable to the patented components. Market value is not necessarily equivalent to technological value, but it is the closest proxy available in the model.

from Δ_{En} to Δ_{En}^{Q} and the upper bound increases from Δ_{Ex} to Δ_{Ex}^{Q}. Figure 6 depicts the results of Proposition 12. The figure adds quality–adjusted patenting to Figure 5 as a dashed line, and the thin dotted verticals show the increased region of effectiveness. Here, the improvement is small, but it does exist. Another benefit is that the discontinuity at Δ_{Bd} is smaller, and the difference in the measure under Entry and Exclusion is smaller more generally. This partially mitigates the increased variance in measured innovation around the Exclusion threshold.[16]

Still, we see that quality–adjusted patenting is a poor measure of innovation for many possible values of Δ. Away from the Exclusion threshold, the protection incentive outweighs the effect on root innovation, and the measure is compromised. Next, I propose an alternative measure that solves this problem without the need for additional data.

5.2 Average patent quality

Because firms have scarce resources, they will choose to patent the most valuable components first, and they will only patent lower–quality innovations afterwards if resources allow. Thus the quality of the average patent decreases as a firm's patent propensity rises. But since changes in competitive pressure increase patent propensity precisely when innovation falls, the average quality of a product's patents can be used as a proxy for innovative activity. In the context of the model, I define observed average patent quality as $q^o = \phi(e^*) \cdot \frac{g(n_p)}{n_p}$.

Proposition 13. *For $\Delta < \Delta_{Bd}$, average patent quality q^o is decreasing in Δ. For $\Delta > \Delta_{Bd}$, average patent quality q^o is increasing in Δ.*

Figure 7 compares innovation to observed values of patenting, quality–adjusted patenting, and average patent quality across values of Δ. The figure adds average patent quality to Figure 6, and we see that its shape is much closer to the changes in underlying innovation than either of the other two patent measures. A discontinuity still exists at Δ_{Bd}, but here it is much smaller.

[16] Note that in these graphics, the functions have been scaled by their highest value for easier comparison. The level of the function is not important for measurement; how the function changes is more relevant.

6 Conclusions

Innovators must solve a two–faceted problem. How does one innovate? How should one protect that innovation from imitators and competitors? External forces and incentives shape the innovator's answers to these two questions in related but distinct ways. Traditionally, however, these two facets are paired: the existence of an IP protection automatically implies an innovation exists, so changes in IP usage correlate with changes in innovative output.

The model presented above disentangles these two channels. Increased incentives to innovate will necessarily increase IP usage as well (the "innovation effect" from equation (4)), but this increase could be mitigated or offset completely if the post–innovation need for IP protection decreases (the "protection effect" from (4)). Distinguishing these two effects leads to a better understanding of the relationships between innovation, IP usage, and external forces such as industry competitiveness or economic policy.

I also highlight a concern with patent–based measures of innovation: patenting rates can only be a good proxy for innovation if the protection effect is small or coincides with the innovation effect. This caveat is particularly apparent when studying how innovative effort and IP usage relate to the competitive pressure in an industry. Here, observed patenting will only correlate with innovation if competitive pressure is moderate and the returns to R&D do not diminish significantly at high levels of investment.

Fortunately, the model also suggests an alternative patent statistic that correlates better with innovation here. The protection effect increases IP usage, and the marginal patent will be of lower–than–average quality. Thus a firm's average patent quality correlates with their innovative effort when the protection and innovation effects counteract each other.

By separating and enriching a firm's IP portfolio choice problem, the model also highlights the different roles a patent can play. A granted patent does not systematically establish a monopoly, it only grants the holder a right to seek damages in the legal system. In fact, it is precisely when litigation damages are awarded or licensing agreements are reached that patent policy is working properly, and optimal patent policy will maximize the value of those inter–firm transfers. A legally established monopoly will only be welfare–maximizing if the product is highly unlikely to exist otherwise.

Researchers and policymakers should be cautious when using patent statistics as a measure of innovation. One should be sure that the incentives to protect profits align with the incentive to innovate in the setting under study, and should also consider the suitability of alternative measures such as average patent quality. When the nuances of the innovation–IP relationship are properly acknowledged, we can be more confident in patent–based measurements of innovation and the conclusions reached by their use.

A Foundations of $g(n_p)$ and $\ell(n_s)$

In this section I derive conditions on a product's component characteristics which guarantee the existence of the well–behaved $g(n_p)$ and $\ell(n_s)$ functions used in the body of the paper.

The Leader's product still consists of a measure N of components, but now suppose that a lawsuit concerning component i is won by the Leader with probability p_i and awards a share α_i of lost profits due to entry. The expected compensation from a lawsuit is thus $p_i \alpha_i \equiv g_i$ times lost profits. The expected litigation value of a portfolio of patents $\mathcal{P}$ is therefore $\int_{i \in \mathcal{P}} g_i \, \mathrm{d}i$ multiplied by lost profits. I assume that g_i is a continuous and differentiable function of i defined on $[0, N]$.

Also suppose that components differ in their secrecy potential ℓ_i, which are additive: the probability that the Follower fails to learn all of the secrets in a secrecy portfolio $\mathcal{S}$ equals $\int_{i \in \mathcal{S}} \ell_i \, \mathrm{d}i \leq 1$.[17] As with patent potential, I assume that ℓ_i is continuous and differentiable over $[0, N]$.

In principle, the values of g_i and ℓ_i need not have any particular relationship for a given component i. But to obtain useful $g(n_p)$ and $\ell(n_s)$, the following assumption on their relationship is needed.

Assumption 3. *g_i is weakly decreasing and ℓ_i is weakly increasing over $[0, N]$.*

This assumption implies that components with high patent potential necessarily have low secrecy potential and vice–versa. Assumption 3 reflects the idea that less complicated components are both easier to reverse–engineer (low ℓ_i) and to demonstrate equivalence to a component in a competitor's product (high g_i).

Imposing Assumption 3 reduces the dimensionality of the Leader's portfolio choice problem considerably:

Lemma 3. *If Assumption 3 holds, then $\int_{i \in \mathcal{P}} g_i \, di = \int_{i=0}^{|\mathcal{P}|} g_i \, di$ and $\int_{i \in \mathcal{S}} \ell_i \, di = \int_{i=N-|\mathcal{S}|}^{N} \ell_i \, di$ for any optimally–chosen portfolios $\mathcal{P}$ and $\mathcal{S}$: the protection afforded by an optimal portfolio is fully described by its size.*

[17]It is tempting to instead define ℓ_i as the independent probability that the Follower learns component i. Then the probability of learning all secrets would be the product of all the ℓ_i, which would be zero for any secrecy portfolio of nonzero measure unless one resorts to product integrals or other exotic beasts. In any case, the temptation is not likely to lead to a tractable model.

Without Assumption 3, changing the number of secret components could change the cost of a given level of total patent protection $\int_{i\in\mathcal{P}} g_i \, \mathrm{d}i$: if the patent potential of a secret component is higher than that of any patented components, reducing the number of secrets will allow the Leader to achieve the same amount of aggregate litigation value with fewer total patents (a violation of the "component independence" assumption discussed in Section 1). With this assumption, however, aggregate patent potential and secret potential can be represented by separate functions. I therefore define

$$g(n_p) \equiv \int_{i=0}^{n_p} g_i \, \mathrm{d}i \text{ and } \ell(n_s) = 1 - \int_{i=N-n_s}^{N} \ell_i \, \mathrm{d}i$$

The assumptions on g_i and l_i imply that $g(n_p)$ and $\ell(n_s)$ are continuous and twice differentiable. It also follows that $g(n_p)$ is weakly concave with $g(0) = 0$ and $\ell(n_s)$ is weakly convex with $\ell(0) = 1$. Note, however, that the assumptions of $g(N) < 1$ and sufficient concavity of $g(n_p)$ and $\ell(n_s)$, which were made in Section 1, are independent of the results of this appendix; they represent additional conditions under which the main results of the paper hold.

B Proofs

Proof of Lemma 1. Given in the body of the paper. □

Proof of Lemma 2. Given in the body of the paper. □

Proof of Proposition 1. I first suppose (Entry) is satisfied so that Exclusion does not occur and determine when the Leader chooses a Mixed, Licensing, Concealment, or Empty portfolio. Then I compare each of those outcomes to Exclusion to see when the Leader chooses it instead of allowing entry.

Consider a Mixed portfolio and increase $s + f$, holding $\frac{s}{f}$ fixed. Doing so decreases n_p and n_s. If n_p decreases to zero before n_s does, then (Ratio) becomes $\frac{s}{f} = \frac{-\ell'(n_s)}{\ell(n_s)} \frac{1}{g'(0)}$. If n_s goes to zero first, then (Ratio) becomes $\frac{s}{f} = \frac{-\ell'(0)}{1} \frac{1-g(n_p)}{g'(n_p)}$. Since $\frac{-\ell'(n_s)}{\ell(n_s)} \frac{1}{g'(0)} \leq \frac{-\ell'(0)}{g'(0)} \leq -\ell'(0) \frac{1-g(n_p)}{g'(n_p)}$ by the sufficient concavity assumptions, if the Leader only uses one method of protection, it will be secrecy if $\frac{s}{f} < \frac{-\ell'(0)}{g'(0)}$ and patenting if $\frac{s}{f} > \frac{-\ell'(0)}{g'(0)}$. For the remainder of the proof, assume $\frac{s}{f} > \frac{-\ell'(0)}{g'(0)}$. Identical arguments will hold for the other case.

Suppose the Leader is indifferent between Empty and Licensing portfolios; that is, the first–order necessary condition of her maximization problem for n_p holds with equality at $n_p = 0$, and that of n_s is slack. Then $g'(0)(1-\lambda(1-\Delta))\pi_M = f$, and she would prefer Empty if f increased slightly. Rearranging yields $\lambda > \left(1 - \frac{f}{g'(0)\pi_M}\right)\frac{1}{1-\Delta}$ as the condition for when Empty is preferred over Licensing. Then $\lambda_2(\Delta) \equiv \left(1 - \frac{f}{g'(0)\pi_M}\right)\frac{1}{1-\Delta}$, which is increasing in Δ.

Now suppose the Leader is indifferent between Licensing and Mixed: both first–order conditions hold exactly and the one for n_s holds at $n_s = 0$. Increasing s slightly would give $-\ell'(0)(1-g(n_p))(1-\lambda(1-\Delta))\pi_M < s$, which when combined with the first–order condition for n_p and rearranged gives $\lambda > \left(1 - \frac{s+f}{\pi_M(g'(n_p)-\ell'(0)(1-g(n_p))}\right)\frac{1}{1-\Delta}$ as the condition for when Licensing is preferred to Mixed. Then $\lambda_1(\Delta) \equiv \left(1 - \frac{s+f}{\pi_M(g'(n_p)-\ell'(0)(1-g(n_p))}\right)\frac{1}{1-\Delta}$, which is increasing in Δ.

Now it remains to consider when Exclusion is preferred to each of Empty, Licensing, and Mixed. This occurs precisely when (Entry) is satisfied. which can be expressed as

$$\overline{n_p} \geq \frac{1-g(n_p)}{g'(n_p)} + +n_p + n_s\frac{-\ell'(n_s)}{\ell(n_s)}\frac{1-g(n_p)}{g'(n_p)}. \tag{8}$$

To compare Empty to Exclusion, (8) becomes $\overline{n_p} \geq \frac{1}{g'(0)}$. Substituting $\overline{n_p} = g^{-1}\left(1 - \frac{\Delta}{1-\lambda(1-\Delta)}\right)$ (from the definition of $\overline{n_p}$) yields $\lambda \leq \frac{1-g\left(\frac{1}{g'(0)}\right)\frac{1}{1-\Delta}}{1-g\left(\frac{1}{g'(0)}\right)} \equiv \lambda_3^0(\Delta)$, where I define $\lambda_3^0(\Delta)$ as the portion of $\lambda_3(\Delta)$ directly below which Empty is the preferred portfolio under entry.

To compare Licensing to Exclusion, (8) becomes $\overline{n_p} \geq \frac{1-g(n_p)}{g'(n_p)+n_p}$. Using the same substitutions and rearrangements as the previous case, (8) becomes $\lambda \leq \frac{1-g\left(\frac{1-g(n_p)}{g'(n_p)+n_p}\right)\frac{1}{1-\Delta}}{1-g\left(\frac{1-g(n_p)}{g'(n_p)+n_p}\right)} \equiv \lambda_3^1(\Delta)$ where I define $\lambda_3^1(\Delta)$ as the portion of $\lambda_3(\Delta)$ directly below which Licensing is the preferred portfolio under entry.

Finally, to compare Mixed to Exclusion, (8) becomes $\overline{n_p} \geq \frac{1-g(n_p)}{g'(n_p)+n_p+n_s\frac{s}{f}}$. Using the same substitutions and rearrangements as the previous cases, (8) becomes $\lambda \leq \frac{1-g\left(\frac{1-g(n_p)}{g'(n_p)+n_p+n_s\frac{s}{f}}\right)\frac{1}{1-\Delta}}{1-g\left(\frac{1-g(n_p)}{g'(n_p)+n_p+n_s\frac{s}{f}}\right)} \equiv \lambda_3^2(\Delta)$ where I define $\lambda_3^2(\Delta)$ as the portion of $\lambda_3(\Delta)$ directly below which Licensing is the preferred portfolio under entry.

Note that all of $\lambda_3^0(\Delta)$, $\lambda_3^1(\Delta)$, and $\lambda_3^2(\Delta)$ are decreasing functions of Δ, and join at the Empty/Licensing and Licensing/Mixed boundaries to form a decreasing, continuous, piecewise function. Define this function as $\lambda_3(\Delta)$. The remaining claims are immediate. $\square$

Proof of Corollary 1. This is seen directly from (8). One can choose parameters (s, f, Δ, λ) such that (8) binds and $n_s = 0$. Varying these parameters to increase n_p will further decrease the difference.

Suppose, by way of contradiction, there exists a sequence of tuples $(s_i, f_i, \Delta_i, \lambda_i)$ that elicit Licensing or Mixed portfolios such that $\overline{n_p} - n_p(s_i, f_i, \Delta_i, \lambda_i) \to 0$. Further, the RHS of (8) approaches $\frac{1-g(\overline{n_p})}{g'(\overline{n_p})}$. But since this is positive, there must be some i in the sequence for which (8) did not hold.

However, note that if $\Delta \geq 0$ and $g(N) > 1$, this outcome is possible since $\frac{1-g(\overline{n_p})}{g'(\overline{n_p})}$ may take negative values. $\square$

Proof of Proposition 2. The Leader's first–order necessary conditions for maximization of Π_L are given by

$$\ell(n_s)g'(n_p)(1 - \lambda(1 - \Delta))\pi_M - f \leq 0 \tag{9}$$

$$-\ell'(n_s)(1 - g(n_p))(1 - \lambda(1 - \Delta))\pi_M - s \leq 0 \tag{10}$$

which I will assume hold with equality at lead to a Mixed portfolio; comparative statics for an Empty portfolio are trivial and those for Licensing or Concealment are simplified versions of what is to follow. The proposition makes claims about $\frac{\partial n_p}{\partial \Delta}$, among others. To proceed, totally differentiate both (9) and (10) with respect to n_p, n_s, and Δ. Stacking as a matrix equation and rearranging leads to

$$\begin{bmatrix} \mathrm{d}n_p \\ \mathrm{d}n_s \end{bmatrix} = \frac{1}{-g''(1-g)\ell''\ell - g'^2\ell'^2} \cdot \begin{bmatrix} g'(1-g)(\ell''\ell - \ell'^2) \\ -\ell'\ell(|g|''(1-g) - g'^2) \end{bmatrix} \cdot \frac{\lambda}{1 - \lambda(1-\Delta)} \cdot \mathrm{d}\Delta \tag{11}$$

Note that I have suppressed the n_p and n_s in parentheticals for clarity. Dividing by $\mathrm{d}\Delta$ yields expressions for $\frac{\partial n_p}{\partial \Delta}$ and $\frac{\partial n_s}{\partial \Delta}$, both of which are positive thanks to the sufficient concavity assumptions made (note that these assumptions also imply the entire maximization problem is concave; the first denominator above will be positive, which is equivalent to the second–order condition that the determinant of the Hessian be positive).

For the Exclusion claim of the Proposition, note that the Leader will use exactly $\overline{n_p}$ patents and zero secrets, so the quantity of interest is $\frac{\partial \overline{n_p}}{\partial \Delta}$. Differentiating the definition of $\overline{n_p}$ leads to $\frac{\partial \overline{n_p}}{\partial \Delta} = \frac{g(\overline{n_p})}{g'(\overline{n_p})} \frac{-1}{(1-\Delta)(1-\lambda(1-\Delta))}$, which is negative as claimed. $\square$

Proof of Corollary 2. Under Entry, the Follower learns enough to enter with probability $\ell(n_s)$, so monopoly occurs with probability $1 - \ell(n_s)$. Under Exclusion, entry is impossible. Since an increase in Δ implies a higher n_s and an eventual switch to Exclusion, the proposition holds. $\square$

Proof of Proposition 3. Assume Entry. First, $\frac{\partial \Pi_L^*}{\partial \Delta} = -\ell(1-g)\lambda\pi_M$ by the Envelope Theorem; this is negative. Next, recall that Π_L^*, e^*, and $\phi(e^*)$ are related through (3). The chain rule implies $\frac{\partial e^*}{\partial \Delta} = \frac{\phi'^2}{-\phi'' \cdot c^2} \cdot \frac{\partial \Pi_L^*}{\partial \Delta}$, where the e^* in parentheticals have been suppressed for clarity. Since $\phi'' < 0$, $\frac{\partial e^*}{\partial \Delta}$ is negative as well. Since ϕ is monotone, then, $\frac{\partial \phi(e^*)}{\partial \Delta}$ is also negative.

Now assume Exclusion. Now $\Pi_L^* = \pi_M - \overline{n_p} f$. Since $\frac{\partial \overline{n_p}}{\partial \Delta}$ is negative (Proposition 2), all of Π_L^*, e^*, and $\phi(e^*)$ are increasing in Δ as claimed. $\square$

Proof of Proposition 4. Assume Entry. By the product rule, $\frac{\partial n_p^o}{\partial \Delta} = \frac{\partial n_p}{\partial \Delta}\phi(e^*) + n_p \frac{\partial \phi(e^*)}{\partial \Delta}$. Substituting these derivatives, which were found in Propositions 2 and 3, and rearranging, $\frac{\partial n_p^o}{\partial \Delta} > 0$ if

$$n_p < \left(\frac{g'\left(\ell'' - \frac{\ell'^2}{\ell}\right)}{-g''(1-g)\ell''\ell - g'^2\ell'^2} \right) \left(\frac{{\Pi_L^*}^2}{(1-\lambda(1-\Delta))\pi_M} \right) \left(-\phi'' \frac{\phi}{\phi'} \right) \tag{12}$$

and vice–versa. This inequality is true for small values of n_p and false for large values. Since n_p is increasing in Δ, this yields an inverse–U relationship. The proof for n_s^o is mechanically identical.

Now assume Exclusion, so that $\frac{\partial n_p^o}{\partial \Delta} = \frac{\partial \overline{n_p}}{\partial \Delta}\phi(e^*) + \overline{n_p}\frac{\partial \phi(e^*)}{\partial \Delta}$. Using the same mechanics as the Entry case, $\frac{\partial n_p^o}{\partial \Delta} > 0$ if $\overline{n_p} > \left(\frac{{\Pi_L^*}^2}{f}\right)\left(-\phi'' \frac{\phi}{\phi'}\right)$ and vice–versa. Since $\overline{n_p}$ is decreasing in Δ, a low Δ implies a high $\overline{n_p}$ and therefore an increasing n_p^o. Similarly, a high Δ implies a decreasing n_p^o. Together, this implies an inverse–U relationship.

Given the above, two things will determine the overall shape of n_p^o: whether the Δ at the apex of n_p^o given Entry is lower than the Δ at the apex of n_p^o given Exclusion, and how those two Δ compare to the one that satisfies (Entry) with equality.

Let Δ_{En} be the Δ for which $\frac{\partial n_p^o}{\partial \Delta} = 0$ under Entry, let Δ_{Ex} be the Δ for which $\frac{\partial n_p^o}{\partial \Delta} = 0$ under Exclusion, and let Δ_{Bd} be the Δ for which (Entry) holds with equality. If $\Delta_{En} \geq \Delta_{Ex}$, then n_p^o will be inverse–U shaped with a jump discontinuity at Δ_{Bd} and apex at Δ_{Bd}. If instead $\Delta_{Bd} \leq \Delta_{En} < \Delta_{Ex}$ or $\Delta_{En} < \Delta_{Ex} \leq \Delta_{Bd}$, the shape will also be inverse–U with a jump discontinuity and an apex at Δ_{Ex} or Δ_{En}, respectively. Finally, if $\Delta_{En} < \Delta_{Bd} < \Delta_{Ex}$, then n_p^o will be M–shaped with a jump: the first apex will be at Δ_{En}, the jump at Δ_{Bd}, and the second apex at Δ_{Ex}.

Manipulating expressions for $\frac{\partial n_p}{\partial \Delta}$, $\frac{\partial \overline{n_p}}{\partial \Delta}$, and (8) shows two things:

$$\Delta_{En} \gtreqless \Delta_{Bd} \;\leftrightarrow\; -\phi'' \frac{\phi}{\phi'^2} \gtreqless \frac{-\varepsilon^{\Delta}_{\Pi_L^*}\big|_{En}}{\varepsilon^{\Delta}_{n_p}} \text{ at } \Delta = \Delta_{Bd} \tag{13}$$

$$\Delta_{Ex} \gtreqless \Delta_{Bd} \;\leftrightarrow\; -\phi'' \frac{\phi}{\phi'^2} \lesseqgtr \frac{\varepsilon^{\Delta}_{\Pi_L^*}\big|_{Ex}}{-\varepsilon^{\Delta}_{\overline{n_p}}} \text{ at } \Delta = \Delta_{Bd} \tag{14}$$

where $\varepsilon^{\Delta}_{\Pi_L^*}\big|_{En}$ is the elasticity of Π_L^* under Entry with respect to Δ, $\varepsilon^{\Delta}_{\Pi_L^*}\big|_{Ex}$ is the elasticity of Π_L^* under Exclusion, $\varepsilon^{\Delta}_{n_p}$ is the elasticity of n_p with respect to Δ, and $\varepsilon^{\Delta}_{\overline{n_p}}$ is the elasticity of $\overline{n_p}$ with respect to Δ. Computing these elasticities shows that $\frac{\varepsilon^{\Delta}_{\Pi_L^*}\big|_{En}}{\varepsilon^{\Delta}_{n_p}} > \frac{\varepsilon^{\Delta}_{\Pi_L^*}\big|_{Ex}}{\varepsilon_{\overline{n_p}}}$ iff $g(\cdot)$ is sufficiently concave relative to $\ell(\cdot)$. We can then state the following:

1. When $-\phi'' \frac{\phi}{\phi'^2}$ is low, the function is inverse–U for $\Delta < \Delta_{Bd}$ and inverse–U for $\Delta > \Delta_{Bd}$,

2. When $-\phi'' \frac{\phi}{\phi'^2}$ is moderate and $g(\cdot)$ is sufficiently concave relative to $\ell(\cdot)$, the function is inverse–U for $\Delta < \Delta_{Bd}$ and decreasing for $\Delta > \Delta_{Bd}$,

3. When $-\phi'' \frac{\phi}{\phi'^2}$ is moderate but $g(\cdot)$ is not sufficiently concave relative to $\ell(\cdot)$, the function is increasing for $\Delta < \Delta_{Bd}$ and inverse–U for $\Delta > \Delta_{Bd}$, and

4. When $-\phi'' \frac{\phi}{\phi'^2}$ is high, the function is increasing for $\Delta > \Delta_{Bd}$ and decreasing for $\Delta < \Delta_{Bd}$.

Similar claims can be proven for n_s^o.

The result for $\phi(e^*)$ follows immediately from Proposition 3: The trough of the V–shape occurs at Δ_{Bd}. □

Proof of Proposition 5. The proof is mechanically identical to the proof of Proposition 2. The relevant matrix equation is now

$$\begin{bmatrix} \mathrm{d}n_p \\ \mathrm{d}n_s \end{bmatrix} = \frac{-1}{-g''(1-g)\ell''\ell - g'^2\ell'^2} \cdot \begin{bmatrix} g'(1-g)(\ell''\ell - \ell'^2) \\ -\ell'\ell(|g|''(1-g) - g'^2) \end{bmatrix} \cdot \frac{1-\Delta}{1-\lambda(1-\Delta)} \cdot \mathrm{d}\lambda \tag{15}$$

which the observant reader will note is inversely proportional to the equations for Δ. Thus both $\frac{\partial n_p}{\partial \lambda}$ and $\frac{\partial n_s}{\partial \lambda}$ are negative, as claimed.

For Exclusion, we now have $\frac{\partial \overline{n_p}}{\partial \Delta} = \frac{g(\overline{n_p})}{g'(\overline{n_p})} \frac{-\Delta}{(1-\lambda)(1-\lambda(1-\Delta))}$, which is again negative as claimed. □

Proof of Corollary 3. As with Corollary 2, the probability of Follower entry is $\ell(n_s)$ until $n_p = \overline{n_p}$, so the probability of monopoly is $1-\ell(n_s)$ under Entry and 1 under Exclusion. Since λ decreases n_s, the result follows. □

Proof of Proposition 6. Again assume Entry. Now $\frac{\partial \Pi_L^*}{\partial \lambda} = \ell(1-g)(1-\Delta)\pi_M$ by the Envelope Theorem, which is positive. Again by the chain rule and the characteristics of $\phi(e)$, the claim follows.

Now assume Exclusion. Since $\frac{\partial \Pi_L^*}{\partial \lambda}$ is negative (Proposition 5), the claim again follows. □

Proof of Proposition 7. Assume Entry. By the same mechanics as in Proposition 4, $\frac{\partial n_p^o}{\partial \lambda} > 0$ if

$$n_p > \left(\frac{g'\left(\ell'' - \frac{\ell'^2}{\ell}\right)}{-g''(1-g)\ell''\ell - g'^2\ell'^2} \right) \left(\frac{{\Pi_L^*}^2}{(1-\lambda(1-\Delta))\pi_M} \right) \left(-\phi'' \frac{\phi}{\phi'} \right)$$

which is the same as the condition in Proposition 4 except that the inequality is reversed. But since $\frac{\partial n_p}{\partial \lambda} < 0$ now, the inverse–U relationship again appears.

Now assume Exclusion. The same steps as in Proposition 4 imply $\frac{\partial n_p^o}{\partial \lambda} > 0$ if $\overline{n_p} > \left(\frac{{\Pi_L^*}^2}{f} \right) \left(-\phi'' \frac{\phi}{\phi'} \right)$ and vice–versa. Again since $\overline{n_p}$ is decreasing in λ, it follows that n_p^o has an inverse–U relationship with λ.

Again following Proposition 4, let λ_{En} be the λ such that $\frac{\partial n_p^o}{\partial \lambda} = 0$ under Entry and let λ_{Ex} be the λ for which $\frac{\partial n_p^o}{\partial \lambda} = 0$ under Exclusion. Situation 2 of the Proposition will occur only if $\lambda_{En} < \lambda_{Bd}$, and likewise situation 3 will occur only if $\lambda_{Ex} > \lambda_{Bd}$.

Manipulating expressions for $\frac{\partial n_p}{\partial \lambda}$, $\frac{\partial \overline{n_p}}{\partial \lambda}$, and (8) shows two things:

$$\lambda_{En} \gtreqless \lambda_{Bd} \;\leftrightarrow\; -\phi'' \frac{\phi}{\phi'^2} \lesseqgtr \frac{\varepsilon^{\lambda}_{\Pi^*_L}\big|_{En}}{-\varepsilon^{\lambda}_{n_p}} \text{ at } \lambda = \lambda_{Bd} \tag{16}$$

$$\lambda_{Ex} \gtreqless \lambda_{Bd} \;\leftrightarrow\; -\phi'' \frac{\phi}{\phi'^2} \lesseqgtr \frac{\varepsilon^{\lambda}_{\Pi^*_L}\big|_{Ex}}{-\varepsilon^{\lambda}_{\overline{n_p}}} \text{ at } \lambda = \lambda_{Bd} \tag{17}$$

where the elasticities are defined as in the proof of Proposition 4. This completes the proof. □

Proof of Proposition 8. Assume Entry. The relevant matrix equation here, analogous to those for Propositions 2 and 5, is now

$$\begin{bmatrix} \mathrm{d}n_p \\ \mathrm{d}n_s \end{bmatrix} = \frac{1}{-g''(1-g)\ell''\ell - g'^2\ell'^2} \cdot \begin{bmatrix} g'(1-\gamma g)\ell''\ell + \ell'^2\gamma g g' \\ -(-\ell'\ell\gamma g'^2 - \ell'\ell\gamma|g''|g) \end{bmatrix} \cdot \mathrm{d}\gamma \tag{18}$$

which shows that n_p is increasing with γ and n_s is decreasing with γ here.

Now assume Exclusion. Now $\overline{n_p}$ is given by $\gamma g(\overline{n_p}) = \frac{(1-\lambda)(1-\Delta)}{1-\lambda(1-\Delta)}$, which implies

$$\frac{\partial \overline{n_p}}{\partial \gamma} = \frac{-1}{\gamma} \frac{g(\overline{n_p})}{g'(\overline{n_p})} \tag{19}$$

which is negative. □

Proof of Corollary 4. As with Corollary 2, the probability of Follower entry is $\ell(n_s)$ until $n_p = \overline{n_p}$, so the probability of monopoly is $1-\ell(n_s)$ under Entry and 1 under Exclusion. Since γ decreases n_s, the result follows. □

Proof of Proposition 9. Assume Entry. The envelope theorem again implies that $\frac{\partial \Pi^*_L}{\partial \gamma}$ equals $\ell g(1-\lambda(1-\Delta))\pi_M$ under Entry and equals $\frac{f}{\gamma}\frac{g(\overline{n_p})}{g'(\overline{n_p})}$ under Exclusion, which are both positive. Therefore $\frac{\partial e^*}{\partial \gamma}$ and $\frac{\partial \phi(e^*)}{\partial \gamma}$ are positive as well. □

Proof of Proposition 10. Assume Entry. We have

$$\frac{\partial n^o_p}{\partial \gamma} = \frac{\partial \phi(e^*)}{\partial \gamma} n_p + \phi(e^*) \frac{\partial n_p}{\partial \gamma} \tag{20}$$

Each term in the expression is positive by Propositions 8 and 9, so n^o_p is increasing.

Now assume Exclusion. Now the expression implies

$$\gamma_{Ex} \gtrless \gamma_{Bd} \leftrightarrow -\phi'' \frac{\phi}{\phi'^2} \lessgtr \frac{\varepsilon^{\gamma}_{\Pi^*_L}\big|_{Ex}}{-\varepsilon^{\gamma}_{\overline{n_p}}} \text{ at } \gamma = \gamma_{Bd} \tag{21}$$

so that $\overline{n_p}$ is increasing in γ if $-\phi'' \frac{\phi}{\phi'^2}$ is low. □

Proof of Proposition 11. Given in the body of the paper. □

Proof of Proposition 12. First, assume Entry and thus $\Delta_{En} < \Delta_{Bd}$. $g^o(n_p)$ will be a better measure if it attains a local maximum for a lower Δ than n_p^o. Equation (12) defines Δ_{En}, and since $\frac{\partial g^o(n_p)}{\partial \Delta} = g'(n_p)\frac{\partial n_p}{\partial \Delta}\phi(e^*) + g(n_p)\frac{\partial \phi(e^*)}{\partial \Delta}$, the equation defining the local maximum for $g^o(n_p)$ is

$$\frac{g(n_p)}{g'(n_p)} < \left(\frac{g'\left(\ell'' - \frac{\ell'^2}{\ell}\right)}{-g''(1-g)\ell''\ell - g'^2\ell'^2} \right) \left(\frac{{\Pi^*_L}^2}{(1-\lambda(1-\Delta))\pi_M} \right) \left(-\phi'' \frac{\phi}{\phi'} \right) \tag{22}$$

This inequality will be violated at a lower n_p than (12) if

$$\frac{g(n_p)}{g'(n_p)} > n_p \rightarrow \frac{g(n_p)}{n_p} > g'(n_p) \tag{23}$$

i.e., if the average is greater than the marginal, which is always true since $g(\cdot)$ is concave, increasing, and starts at zero.

For $\Delta_{En} \geq \Delta_{Bd}$ (Exclusion), n_p is now decreasing and $\phi(e^*)$ is now increasing. Identical logic to the Entry case, and the fact that a superior measure would attain a local maximum at a higher Δ, proves the claim. □

Proof of Proposition 13. As in Propositions 4 and 12, q^o is decreasing for $\Delta < \Delta_{Bd}$ if

$$\frac{g(n_p)}{g'(n_p)\cdot n_p - g(n_p)} n_p < \left(\frac{g'\left(\ell'' - \frac{\ell'^2}{\ell}\right)}{-g''(1-g)\ell''\ell - g'^2\ell'^2} \right) \left(\frac{{\Pi^*_L}^2}{(1-\lambda(1-\Delta))\pi_M} \right) \left(-\phi'' \frac{\phi}{\phi'} \right) \tag{24}$$

But this is always true: the denominator on the left–hand side is negative since $g'(n_p) < \frac{g(n_p)}{n_p}$. Thus q^o is always decreasing for $\Delta < \Delta_{Bd}$. Identical logic applies for the $\Delta \geq \Delta_{Bd}$ case, as before. □

Proof of Lemma 3. Consider an optimally–chosen $\mathcal{P}$ and $\mathcal{S}$. $\int_{i\in\mathcal{P}} g_i \ \mathrm{d}i = \int_{i=0}^{|\mathcal{P}|} g_i \ \mathrm{d}i$ if the first $|\mathcal{P}|$ components are patented. We proceed by contradiction: suppose instead that there is at least one component $j < |\mathcal{P}|$ that is not patented, and is either therefore secret or disclosed. Let k be the index of the component with the lowest g_i in $\mathcal{P}$. If component j is disclosed, the Leader can patent it and disclose component k. This leaves her cost unchanged but increases $\int_{i\in\mathcal{P}} g_i \ \mathrm{d}i$, since g_i is nonincreasing and $j < k$ by definition. If component j is secret, the Leader can patent it and make component k secret. This again leaves her cost unchanged but increases both $\int_{i\in\mathcal{P}} g_i \ \mathrm{d}i$ and $\int_{i\in\mathcal{S}} \ell_i \ \mathrm{d}i$ since $j < k$. In either case, the supposition is contradicted since the portfolio is demonstrably sub–optimal, and the claim follows.

The proof for $\int_{i\in\mathcal{S}} \ell_i \ \mathrm{d}i = \int_{i=N-|\mathcal{S}|}^{N} \ell_i \ \mathrm{d}i$ is identical. □

References

Aghion, Philippe, Ufuk Akcigit, and Peter Howitt (2013), "What do we learn from Schumpeterian growth theory?" Technical report, National Bureau of Economic Research.

Aghion, Philippe, Nick Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt (2005), "Competition and innovation: an inverted–U relationship." *The Quarterly Journal of Economics*, 120, 701–728.

Aghion, Philippe and Peter Howitt (1992), "A Model of Growth through Creative Destruction." *Econometrica*, 60, 323–51, URL `https://ideas.repec.org/a/ecm/emetrp/v60y1992i2p323-51.html`.

Anton, James, Hillary Greene, and Dennis Yao (2006), "Policy implications of weak patent rights." In *Innovation Policy and the Economy, Volume 6*, 1–26, The MIT Press.

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

Arora, Ashish (1997), "Patents, licensing, and market structure in the chemical industry." *Research policy*, 26, 391–403.

Arora, Ashish, Marco Ceccagnoli, and Wesley M. Cohen (2003), "R&d and the patent premium." Working Paper 9431, National Bureau of Economic Research, URL `http://www.nber.org/papers/w9431`.

Arrow, Kenneth (1962), "Economic welfare and the allocation of resources for invention." In *The rate and direction of inventive activity: Economic and social factors*, 609–626, Princeton University Press.

Belleflamme, Paul and Francis Bloch (2013), "Dynamic protection of innovations through patents and trade secrets." *Available at SSRN 2537248*.

Cohen, Wesley M, Richard R Nelson, and John P Walsh (2000), "Protecting their intellectual assets: Appropriability conditions and why US manufacturing firms patent (or not)." Technical report, National Bureau of Economic Research.

de Bettignies, Jean-Etienne and Bulat Gainullin (2016), "Competition, licensing, and innovation strategy." URL https://ssrn.com/abstract=2400839.

Denicolò, Vincenzo and Luigi Alberto Franzoni (2004), "Patents, secrets, and the first-inventor defense." *Journal of Economics & Management Strategy*, 13, 517–538.

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

Gilbert, Richard (2006), "Looking for Mr. Schumpeter: Where are we in the competition–innovation debate?" In *Innovation Policy and the Economy, Volume 6*, 159–215, The MIT Press.

Graham, SJH (2004), "Hiding in the patents shadow: Firms uses of secrecy to capture value from new discoveries. georgia institute of technology ti: Ger working paper series."

Graham, Stuart JH and Deepak Hegde (2014), "Do inventors value secrecy in patenting? Evidence from the American Inventor's Protection Act of 1999." URL http://ssrn.com/abstract=2170555.

Griliches, Zvi (1981), "Market value, r&d, and patents." *Economics letters*, 7, 183–187.

Griliches, Zvi (1990), "Patent statistics as economic indicators: A survey." *Journal of Economic Literature*, 28, 1661.

Griliches, Zvi, Ariel Pakes, Bronwyn H Hall, and Zvi Criliches (1987), "The value of patents as indicators of inventive activity." In *Economic Policy and Technological Performance*, Citeseer.

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, 375–423.

Hall, Bronwyn H (2007), "Patents and patent policy." *Oxford Review of Economic Policy*, 23, 568–587.

Hall, Bronwyn H, Adam Jaffe, and Manuel Trajtenberg (2005), "Market value and patent citations." *RAND Journal of economics*, 16–38.

Hall, Bronwyn H, Adam B Jaffe, and Manuel Trajtenberg (2001), "The nber patent citation data file: Lessons, insights and methodological tools." Technical report, National Bureau of Economic Research.

Hall, Bronwyn H, Jacques Mairesse, and Pierre Mohnen (2010), "Measuring the returns to r&d." *Handbook of the Economics of Innovation*, 2, 1033–1082.

Henry, Emeric, Francisco Ruiz-Aliseda, et al. (2016), "Keeping secrets: the economics of access deterrence." *American Economic Journal: Microeconomics*, 8, 95–118.

Hussinger, Katrin (2006), "Is silence golden? Patents versus secrecy at the firm level." *Economics of Innovation and New Technology*, 15, 735–752.

Lerner, Josh (1995), "Patenting in the shadow of competitors." *Journal of law and economics*, 463–495.

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*, 783–831.

Lobel, Orly (2013), "Filing for a patent versus keeping your invention a trade secret." *Harvard Business Review*.

Moser, Petra (2012), "Innovation without patents: Evidence from World's Fairs." *Journal of Law and Economics*, 55, 43–74.

Nagaoka, Sadao, Kazuyuki Motohashi, and Akira Goto (2010), "Patent statistics as an innovation indicator." *Handbook of the Economics of Innovation*, 2, 1083–1127.

Noonan, Kevin E. (2011), "While not right for every invention, trade secret protection has its appeal." Technical report, McDonnell, Boehnen, Hulbert, & Berghoff, LLP.

Ottoz, Elisabetta and Franco Cugno (2008), "Patent–secret mix in complex product firms." *American Law and Economics Review*, 10, 142–158.

Pakes, Ariel (1985), "On patents, r & d, and the stock market rate of return." *Journal of political economy*, 93, 390–409.

Pakes, Ariel and Zvi Griliches (1980), "Patents and r&d at the firm level: A first report." *Economics letters*, 5, 377–381.

Peneder, Michael (2012), "Competition and innovation: Revisiting the inverted–U relationship." *Journal of Industry, Competition and Trade*, 12, 1–5.

Png, I.P.L. (2015), "Secrecy and patents: Evidence from the uniform trade secrets act." *Working Paper*, URL `https://ssrn.com/abstract=2617266`.

Schumpeter, Joseph Alois (1942), *Socialism, capitalism and democracy.* Harper and Brothers.

Shapiro, Carl (2011), "Competition and innovation: Did Arrow hit the bull's eye?" In *The Rate and Direction of Inventive Activity Revisited*, 361–404, University of Chicago Press.

Solow, Robert M (1957), "Technical change and the aggregate production function." *The review of Economics and Statistics*, 312–320.

Trajtenberg, Manuel (1990), "A penny for your quotes: patent citations and the value of innovations." *The Rand Journal of Economics*, 172–187.

Whinston, Michael D (2011), "Comment on "competition and innovation: Did Arrow hit the bull's eye?"." In *The Rate and Direction of Inventive Activity Revisited*, 404–410, University of Chicago Press.

Zhang, Tianle (2012), "Patenting in the shadow of independent discoveries by rivals." *International Journal of Industrial Organization*, 30, 41–49.

Figures

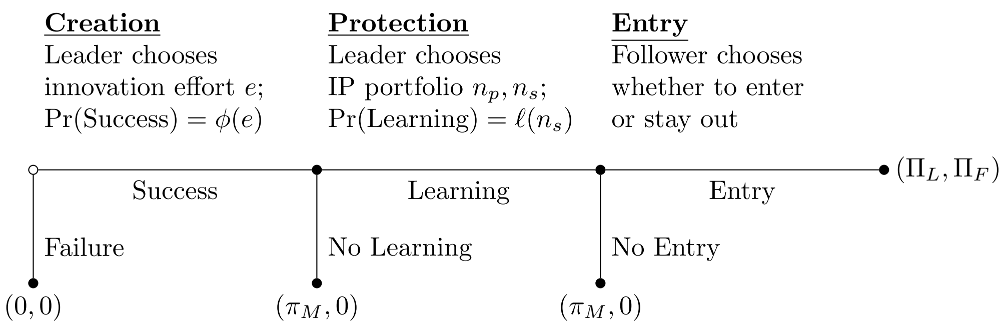

Figure 1: The model environment.

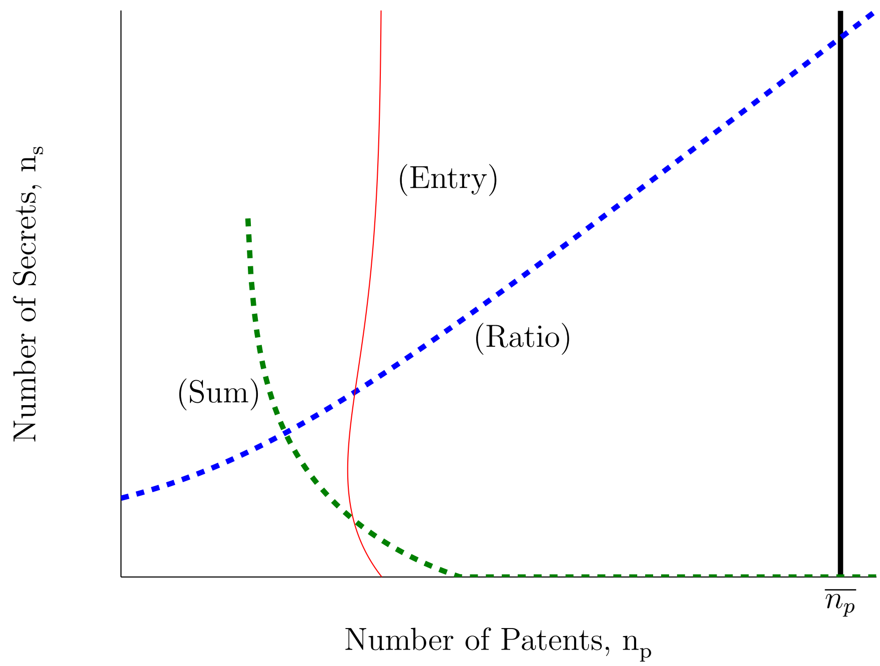

Figure 2: A graphical depiction of equilibrium as given by (Ratio), (Sum), and (Entry).

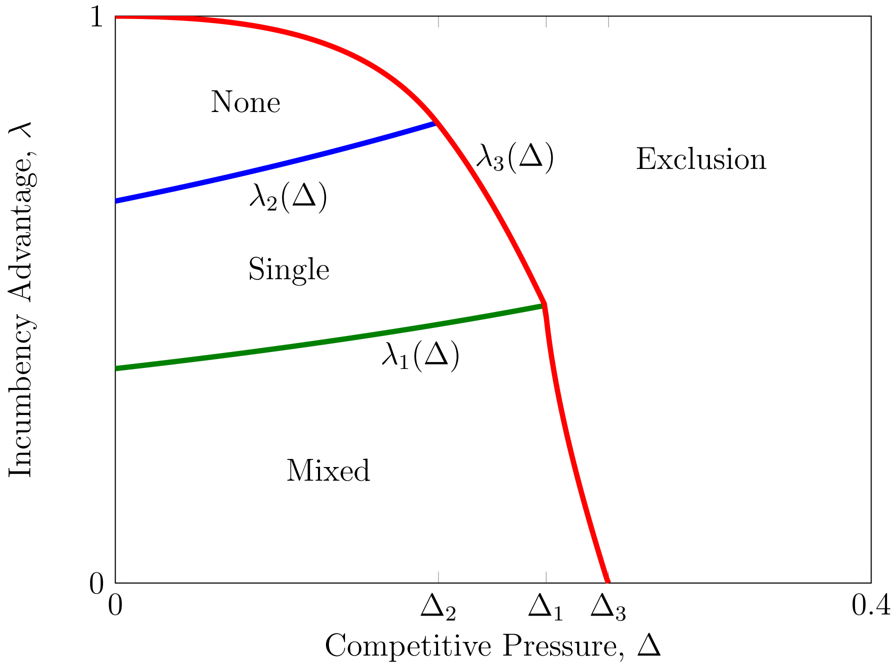

Figure 3: A graphical depiction of Proposition 1.

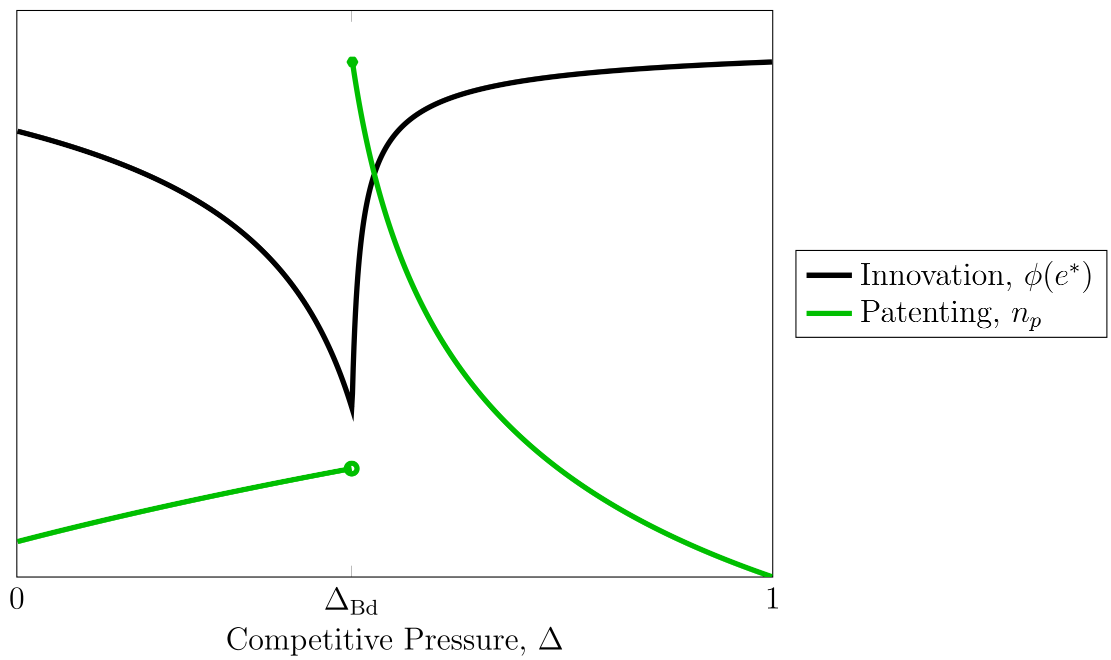

Figure 4: Ex–post patenting as a function of Δ.

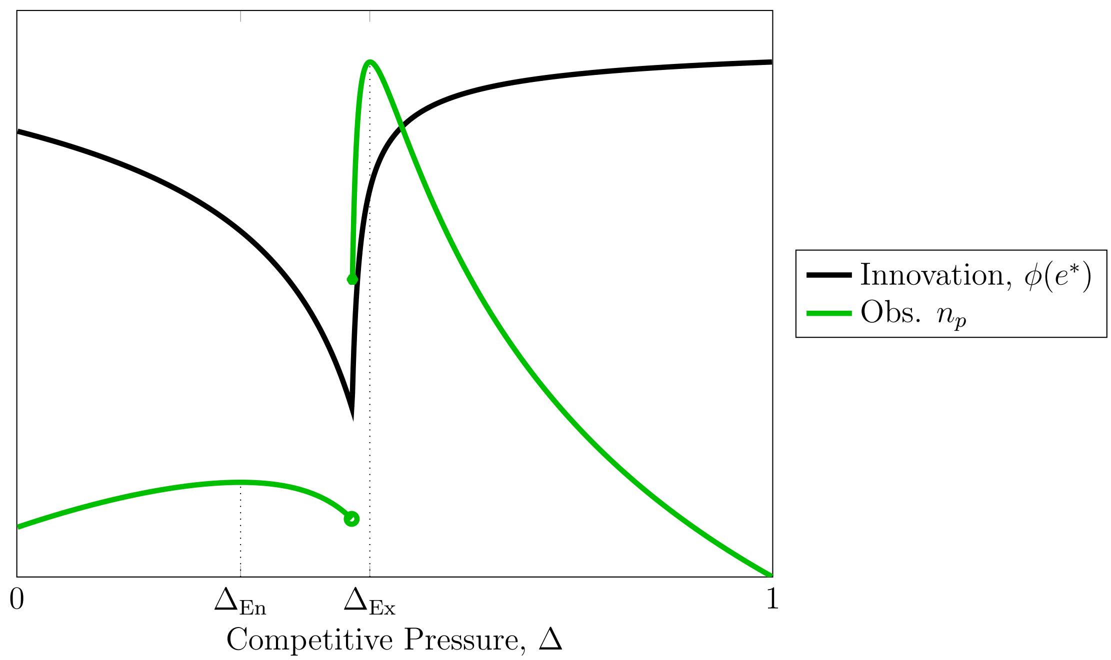

Figure 5: Observed patenting as a function of Δ.

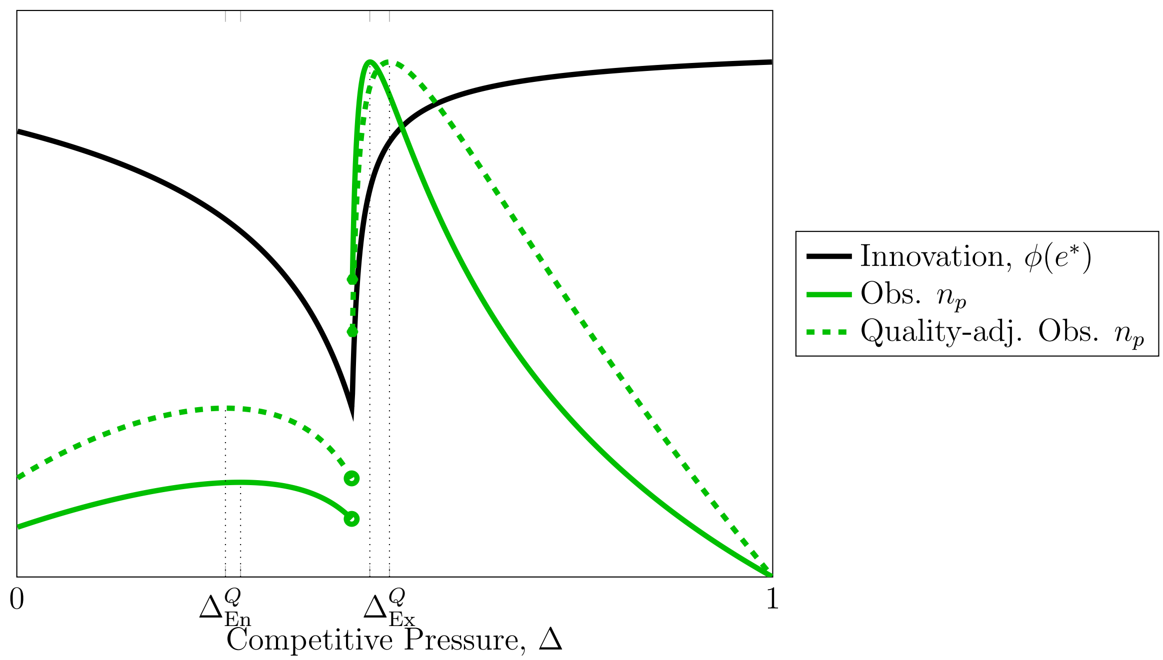

Figure 6: Patenting measures as functions of Δ.

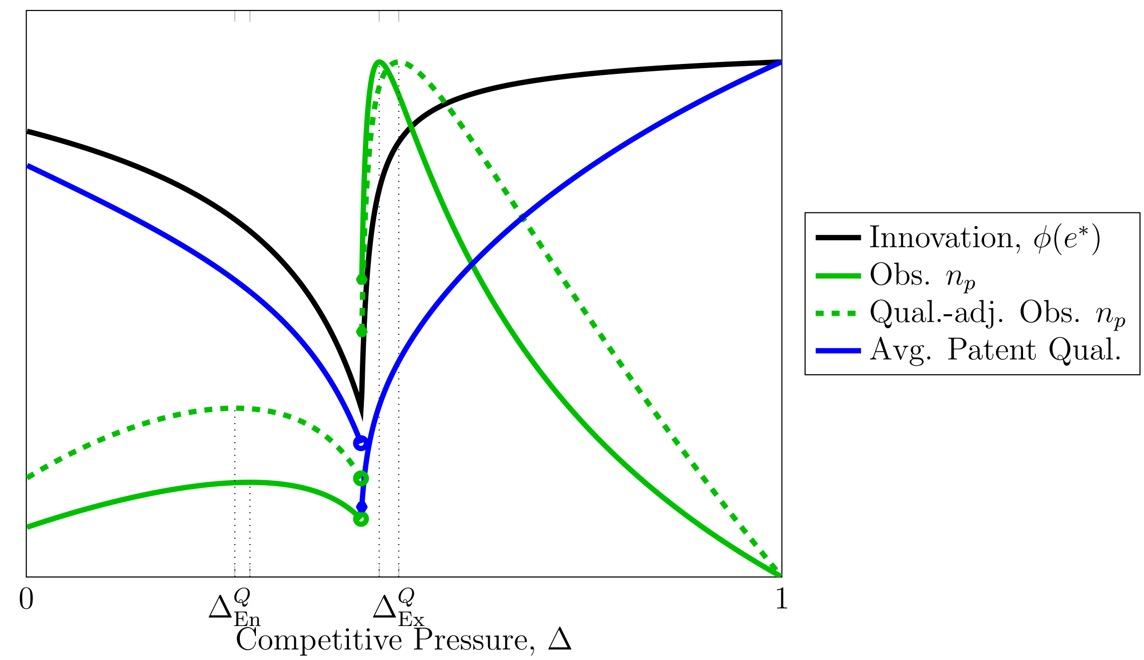

Figure 7: Patenting measures as functions of Δ.