American Economic Review 2008, 98:4, 1347–1369
http://www.aeaweb.org/articles.php?doi=10.1257/aer.98.4.1347

How Strong Are Weak Patents?

By Joseph Farrell and Carl Shapiro*

We study the welfare economics of probabilistic patents that are licensed without a full determination of validity. We examine the social value of instead determining patent validity before licensing to downstream technology users, in terms of deadweight loss (ex post) and innovation incentives (ex ante). We relate the value of such pre-licensing review to the patent's strength, i.e., the probability it would hold up in court, and to the per-unit royalty at which it would be licensed. We then apply these results using a game-theoretic model of licensing to downstream oligopolists, in which we show that determining patent validity prior to licensing is socially beneficial. (*JEL* D82, K11, L24, O34)

Roughly 15,000 patents a month are issued by the US Patent and Trademark Office (PTO).[1] By law, these are supposed to cover only "novel" and "nonobvious" inventions, but an average application gets only about 15–20 hours of patent examiner time,[2] and a substantial proportion of the few patents later fully evaluated in court are held invalid. The Federal Trade Commission (FTC) and the National Academies of Science have joined consumer and industry (especially information technology) groups in expressing grave concern about the issuance of many questionable patents.[3] Responding to such concerns, the Supreme Court recently made it easier for patent examiners and courts to reject an application or overturn a patent as "obvious."[4] Congress is considering broader patent system reform.

Is it very bad if issued patents are invalid, i.e., not justified by the applicant's novel, nonobvious invention? Surely the government should not be handing out legal monopolies unless well justified by the need to reward such invention. But a patent holder cannot exclude its rivals or extract royalties without at least threatening to go to court, which would subject the patent to a thorough review for validity. Thus, a blatantly invalid patent, which clearly would be overturned in court, may never be asserted and may thus cause no harm. The bigger issue, we suggest, concerns patents that are not *clearly* invalid, but are weak—they may well be invalid, but nobody knows for sure without conclusive litigation.[5]

* Farrell: Department of Economics, University of California, Berkeley, CA 94720 (e-mail: farrell@econ.berkeley.edu); Shapiro: Haas School of Business, University of California, Berkeley, CA 94720 (e-mail: shapiro@haas.berkeley.edu). We thank seminar participants at UC Berkeley, Stanford University, Northwestern University, University of Chicago, the FTC, NBER, and the Conference in Tribute to Jean-Jacques Laffont for helpful suggestions. We also thank Bharat Anand, Richard Gilbert, Michael Katz, Mark Lemley, Josh Lerner, Robert Merges, Michael Meurer, Michael Schwartz, and Scott Stern for valuable comments on earlier drafts.

[1] In fiscal 2006, the PTO received 444,000 patent applications and issued 183,000 patents; in the past ten years, it has issued 1.7 million (US PTO 2006, tables 2 and 6). In 2006 the European Patent Office issued nearly 63,000 patents, a record.

[2] Federal Trade Commission (2003, chap. 5, p. 5).

[3] See, for instance, FTC (2003), National Academies of Science (2004), Adam Jaffe and Josh Lerner (2004), Mark Lemley and Shapiro (2005), and James Bessen and Michael Meurer (2008).

[4] *KSR International Co. v. Teleflex, Inc. et. al.*, 127 S. Ct. 1929 (2007); opinion available at http://www.supremecourtus.gov/opinions/06pdf/04-1350.pdf.

[5] This notion of a "weak patent" contrasts with "weak enforcement" of clearly valid patents, e.g., because infringement is difficult to detect.

But conclusive litigation on patent validity is rare. Even when litigation is initiated, most cases settle; and far more patents are licensed without litigation.[6] Thus, we ask: what is the social value of determining up front whether a patent is valid, versus releasing it into the commercial world of patent licensing and litigation?[7] We evaluate these social benefits both ex post (given that the patented technology exists) and ex ante (exploring incentives for invention and/or patenting).

Suppose that an issued patent's market impact is proportional to *patent strength*, i.e., the probability θ that it would be found valid and infringed if tested in court. Then, if a random group of 1,000 probabilistic patents were thoroughly examined for validity, many would be overturned, but the rest would be strengthened, with no change in their aggregate market impact.[8] The proportionality hypothesis thus suggests that improved prelicensing review would be socially desirable if, and only if, the administrative costs of examining these 1,000 patents more thoroughly were lower than the resulting savings in litigation and license negotiation costs. Lemley (2001) argues that the costs of litigating a few commercially important patents may well be smaller than the costs of more thoroughly examining a great many patent applications. Concerns about the issuance of weak patents would indeed be significantly defused if a patent's market impact were proportional to its strength.

To explore the market impact and welfare economics of probabilistic patents, we use a simple model of licensing in the shadow of patent litigation, although some of our results are much broader. In our model, the patent holder offers licenses to downstream firms, each of which can accept the license, avoid using the patented technology, or infringe, prompting litigation.[9] This naturally generalizes existing models of the licensing of ironclad (certainly valid) patents, in which a downstream firm can either accept the offered license or avoid using the patented technology.[10] Our model assumes that litigation costs are zero, but nevertheless predicts licensing without litigation. It also assumes that no interim social costs arise during the pendency of litigation, and that the patent holder cannot behave opportunistically toward those who make investments specific to the patented technology.[11] In this context, we identify special circumstances in which the market impact of a probabilistic patent is proportional to patent strength, and other circumstances in which such proportionality dramatically fails.

Ex post, suppose that the patented technology is already available to society and ostensibly covered by a patent that is valid with probability θ. If licenses can be negotiated under uncertainty, they will be (in our model). Writing $r(\theta)$ for the per-unit royalty that will be agreed, we show quite generally that the gross social benefit of determining the validity of an issued patent

[6] Lemley (2001) estimates that about 5 percent of all patents are either licensed without litigation or are litigated, and that only 0.1 percent of all patents are litigated to trial, so roughly 50 times as many patents are licensed (without litigation or to settle litigation) as are litigated to trial. Kimberly Moore (2000) reports that the percentage of patent cases going to trial has declined over time, to 3.3 percent by 1999 (table 1). Jay Kesan and Gwendolyn Ball (2006) conclude that patent litigation is largely a settlement mechanism; about 10 percent of patent cases filed in 2000 led to rulings and verdicts (table 6).

[7] The most important prelicensing review in practice, and the obvious policy variable, is additional review by the PTO, either before a patent is issued or through reexamination. We thus refer, for short, to additional "PTO review." Importantly, additional PTO review can lead to a patent never being issued, or invalidation of an issued patent. Improved prelicensing review can take other forms, however, including litigation, peer evaluation of patent applications, and the collection of prior art by public-interest groups concerned about improperly issued patents.

[8] Unlike popular commentators, we do not assume that improved PTO review would simply eliminate some "bad" patents without affecting others. Such a view is inconsistent with Bayesian statistics. Suppose, for example, that 20 percent of these 1,000 patents are expected to be valid. Thorough review would, on average, yield 200 valid patents and 800 invalid patents. The 200 surviving patents would emerge from the review far stronger: what does not kill a patent makes it stronger.

[9] To focus on issues of patent validity, we assume that downstream firms plainly infringe the claims in the patent; the uncertainty concerns the validity of those claims.

[10] Morton Kamien (1992) reviews this literature; more recently, see Debapriya Sen and Yair Tauman (2007).

[11] Shapiro (2008) studies such patent hold-up.

before it is licensed is roughly proportional to $r(\theta) - \theta r(1)$, the gap between the negotiated royalty and the expected royalty that would result if validity were determined prior to licensing. This establishes an ex post normative benchmark of $\theta r(1)$ for the per-unit royalty on a probabilistic patent.

If a patent is licensed to downstream firms that do not compete against each other or against the patent holder, our licensing model supports the optimistic or proportional perspective: the per-unit royalty rate $r(\theta)$ is approximately $\theta r(1)$, and the profit accruing to the owner of a patent of strength θ, $P(\theta)$, is approximately $\theta P(1)$.

But when downstream firms use the patented technology in competing against each other or against the patent holder, two mechanisms captured in our licensing model boost royalties and patentee profits well above those benchmarks, especially for weak (small θ) patents. First, agreeing to higher per-unit royalties raises the joint profits of the patent holder and licensees by bringing the downstream price closer to the monopoly price. We show that this joint profit motive for high per-unit royalty rates prevails for weak patents if licenses can use unrestricted two-part tariffs. If they cannot, we assume that licenses are linear or use two-part tariffs with fixed fees that are restricted to be nonnegative. Even here, a downstream firm's decision to litigate, and perhaps invalidate the patent, benefits other downstream firms as well as consumers. Because of this positive externality on rivals, incentives to challenge patents are suboptimal, and downstream firms will accept surprisingly large per-unit royalties. In our licensing model, $r(\theta)$ is a large multiple of $\theta r(1)$, and the patent holder's profits $P(\theta)$ are a large multiple of $\theta P(1)$.

This has strong implications for the ex post social benefits of improved PTO review. For the ex post analysis, θ need only be a commonly known probability; the relationship between patent strength and actual innovation does not matter.[12] In the limiting case, "ideal" PTO review that replicates judicial review would lead either to an ironclad patent (with probability θ) or to no patent (with probability $1 - \theta$). The expected ex post gross benefit of such review is $B(\theta) \equiv [\theta W(1) + (1 - \theta)W(0)] - W(\theta)$, where $W(\theta)$ is the welfare resulting when a patent of strength θ is licensed. We show that $B(\theta)$ is small (and sometimes negative) for patents licensed to downstream firms that do not compete, since proportionality holds there, but large and positive for weak patents licensed to downstream rivals, since proportionality fails.

Ex ante, how does (the prospect of) enhanced review prior to licensing affect innovation incentives? We assume that courts truly assess patent validity, so that θ is the probability that the patent holder actually contributed to society the technology covered by the patent, rather than getting a wrongly issued or overly broad patent on prior art or obvious technology; as we did ex post, we first develop general principles and then turn to our particular licensing model. We analyze the expected social contribution $K(\theta)$ made by the owner of a patent of strength θ. If the downstream firms compete and licenses can use unrestricted two-part tariffs, we find that $K(\theta) < 0$ for weak patents, distorting ex ante incentives for research and patenting. Even when $K(\theta) > 0$, ex ante incentives are distorted if $P(\theta) > K(\theta)$, and, more generally, if the ratio of private to social returns, $P(\theta)/K(\theta)$, varies substantially with θ, as we show is typical for patents licensed to downstream rivals in our licensing model. For such patents, enhanced PTO review not only yields ex post benefits $B(\theta) > 0$, but also eliminates some distortions in ex ante incentives to engage in research and apply for patents. We also establish an ex ante benchmark for $r(\theta)$ equal to θv, where v is the per-unit value of the technology.

Section I presents our licensing model for probabilistic patents. Section II establishes some general results for evaluating the ex post and ex ante welfare effects of enhanced prelicensing review. For a patent that is licensed using a two-part tariff to downstream firms that do not

[12] We assume that θ is common knowledge. We discuss the role of this assumption in Section VI.

compete against each other or against the patent holder, Section III shows that $r(\theta) = 0$ for all θ in our licensing model, so (in expectation) enhanced review generates no ex ante or ex post social benefits. In contrast, for patents licensed to multiple downstream rivals, Section IV shows, surprisingly, that $r(\theta) \geq r(1)$ in our licensing model, so $r(\theta)$ must exceed the benchmark $\theta r(1)$, especially for weak patents (small θ), and ideal PTO review generates ex ante and ex post social benefits. When θ is small, downstream firms accept this high running royalty because it increases joint profits and the patent owner shares the increase with downstream firms through a *negative* fixed fee $F(\theta) < 0$.

If licenses cannot use such negative fixed fees, we assume that they will consist simply of a per-unit royalty, $s(\theta)$. Section V shows that for weak patents, $s(\theta)$ far exceeds the benchmark $\theta s(1)$ when licensees compete. In Cournot competition, the ratio $s(\theta)/\theta s(1)$ is of order N, the number of downstream firms, for very weak patents. Section VI generalizes our results to the case of a vertically integrated patent holder that competes against its licensees, and explores the effects of relaxing some of our assumptions. Section VII concludes.

I. Patent Licensing in the Shadow of Litigation

A. *Technology and Licensing Game*

An upstream patent holder P offers licenses to N symmetric downstream firms. The patented technology lowers downstream firms' unit production costs by v, the *patent size*, relative to the best alternative, or *backstop*, technology. Equivalently, the technology makes each unit of the product worth an extra v to all customers. Until Section VI, we assume that the patent holder does not compete against the downstream firms.

In general, when an upstream monopolist sells an input to downstream firms that compete, the literature on multilateral vertical contracting has shown that equilibrium depends heavily on the form of contracts allowed, on downstream firms' information, and on their beliefs about what they cannot observe.[13] If the upstream supplier can commit to arbitrary contingent contracts, it can organize a hub-and-spoke downstream cartel supporting the monopoly price downstream, even if it controls only a minor input.[14] At the other extreme, if contracts are private, the upstream supplier may be unable to charge any price above marginal cost for its input.[15] Finding those extreme outcomes unrealistic, we adopt a simple licensing model: the patent holder offers a two-part tariff $[F, r]$ to all downstream firms.[16] Nondiscriminatory offers are sometimes used in practice, are prominent in the ironclad patent licensing literature (see e.g., Kamien 1992), and are typically required for the licensing of patents incorporated into industry standards.[17]

[13] See R. Preston McAfee and Marius Schwartz (1994) and Ilya Segal (1999).

[14] The upstream firm can sell its input at a price that supports the downstream monopoly price and threaten to subsidize all other downstream firms if any one downstream firm does not buy its input.

[15] If negotiations are private and each downstream firm has "passive beliefs"—does not adjust its beliefs about other firms' contracts when offered a new contract—then the upstream firm negotiates the bilaterally efficient contract with each downstream firm, which involves a price equal to its marginal cost. See McAfee and Schwartz (1994) and Patrick Rey and Thibaud Vergé (2004).

[16] Restricting the number of licenses offered can be optimal for an ironclad patent; see Michael Katz and Shapiro (1986), Chun-Hsiung Liao and Sen (2005), and Sen and Tauman (2007). This approach does not, however, work as a licensing strategy for a probabilistic patent, since firms that do not receive licenses will infringe the patent. If P sues those firms, the equilibrium involves litigation (considered below). If P ignores infringing firms, downstream firms will be unwilling to pay for licenses.

[17] Benjamin Chiao, Lerner, and Jean Tirole (2007) study the rules of 59 standard-setting organizations; about 75 percent require essential patents to be licensed on fair, reasonable, and nondiscriminatory terms. Licensing programs in standard-setting contexts that have attracted antitrust attention include those of Rambus (computer memory) and Qualcomm (mobile telephones).

With an ironclad patent, each downstream firm accepts the offered license or uses the backstop technology. Here, a downstream firm that rejects a license has another option: infringing the patent. In that case, we assume that the patent holder sues the infringer.[18] If the patent is held invalid, all downstream firms can use the technology free of charge.[19] Alternatively, if the patent is ruled valid, we assume (until Section VI) that any licenses already signed remain in force, and that the patent holder negotiates anew with the downstream firm(s) that lack licenses. Lastly, the downstream firms compete, given the licenses they have signed and the technologies they use.

B. *Downstream Oligopoly*

The downstream oligopoly equilibrium depends on downstream firms' marginal cost curves. For ease of notation, we measure each firm's marginal cost relative to the cost that it would incur using the patented technology free of charge. With this notation, a firm that accepts a license with per-unit royalty r has marginal cost r, and a firm using the backstop technology has marginal cost v. To analyze a symmetric equilibrium, we need only consider the profits of one firm with costs a when all other firms have costs b; write $x(a,b)$ for its output and $\pi(a,b)$ for its profits, net of running royalties but gross of any fixed fee. We assume $\pi(a,b)$ satisfies three mild conditions: (1) $\pi_1(a,b) < 0$: a firm's profits are decreasing in its own costs; (2) $\pi_2(a,b) \geq 0$: a firm's profits are nondecreasing in the other firms' costs; and (3) $\pi_1(a,a) + \pi_2(a,a) < 0$: each firm's profits fall if all firms' costs rise in parallel. Writing $p(a)$ for the price charged by each downstream firm if all firms have cost a, in a wide range of simple oligopoly models $p(a)$ is linear.[20] In the text, we assume this; the Appendix shows where we actually rely on this assumption.

C. *Optimal Two-Part Tariffs*

Suppose that downstream firm D expects all its rivals to accept the offer $[F, r]$. If it, too, accepts, its payoff is $\pi(r,r) - F$. If it rejects the offer, infringes, is sued, and the patent is upheld, P would hold D down to its backstop payoff $\pi(v,r)$. Thus D's reservation payoff is $\theta\pi(v,r) + (1-\theta)\pi(0,0)$.[21] Figure 1 displays the game tree for this licensing game, simplified to focus on just one downstream firm. If P opts to avert litigation, it will set the largest F such that it is a subgame equilibrium for all downstream firms to accept $[F, r]$: thus, $F(\theta) = \pi(r,r) - \theta\pi(v,r) - (1-\theta)\pi(0,0)$. Writing total profits (divided by N) as $T(r) \equiv rx(r,r) + \pi(r,r)$, let $r(\theta)$ maximize P's payoff per downstream firm, $G(r,\theta) \equiv T(r) - \theta\pi(v,r) - (1-\theta)\pi(0,0)$, subject (we assume) to $0 \leq r \leq v$.[22]

[18] P might not want to sue an infringing firm, especially if others have signed licenses and litigating would put their royalty payments at risk: it might prefer, ex post, quietly to ignore an infringer. We discuss litigation credibility further in Section VI.

[19] The US Supreme Court has ruled that if one challenger to a patent prevails on patent invalidity, other users can rely on this result and therefore need not pay royalties, even if they had previously agreed to do so. See *Blonder-Tongue Labs, Inc. v. University of Illinois Foundation,* 402 US 313, 350 (1971). As we discuss in Section VI, we assume that the uncertainty concerns validity, or an aspect of infringement that is similarly highly correlated among the downstream firms: if downstream firm 1 proves that it does not infringe the patent, other downstream firms could readily show the same.

[20] Linear pass-through (meaning that price as a function of the royalty rate is a straight line, not necessarily through the origin) holds in Cournot oligopoly with linear demand or constant elasticity demand, in the standard Hotelling duopoly model, and in differentiated-product Bertrand oligopoly with linear demand, among others.

[21] Using the backstop technology would yield $\pi(v,r) \leq \pi(v,v) < \pi(0,0)$, so D would prefer litigating.

[22] A downstream firm might accept a running royalty rate $r > v$ combined with $F < 0$. Under patent law, a license can impose royalties only for use of the patented technology. We assume that this rule is effectively enforced. This implies $r \leq v$, since even a downstream firm that signed a license would use the backstop technology rather than pay $r > v$ to use the patented technology. If the rule is not well enforced, P can bribe each downstream firm with a negative fixed fee to accept a royalty rate $r > v$ (on all output) that supports the downstream monopoly price. In a previous

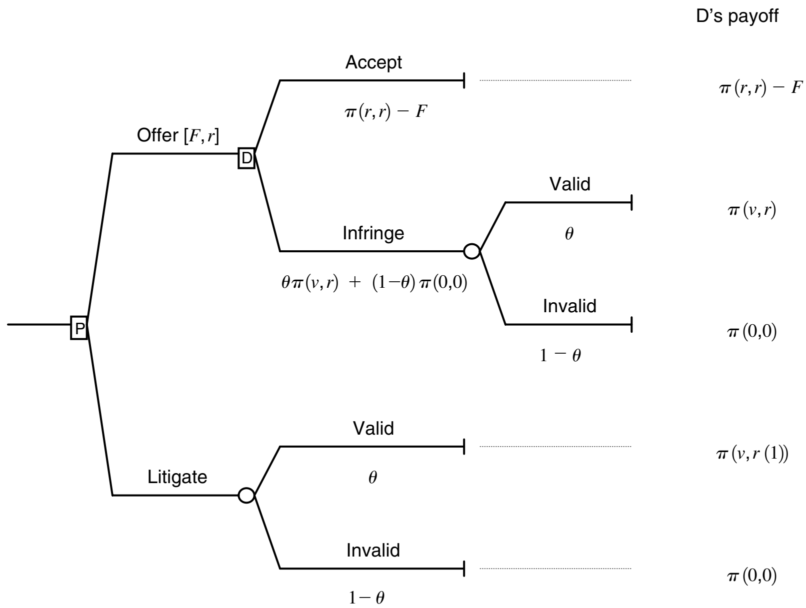

Figure 1. Licensing Game

Let $H(\theta) \equiv G(r(\theta),\theta)$ denote the patent holder's resulting no-litigation payoff. That will be the overall equilibrium if P prefers it to litigating the patent. Litigating gives P zero if the patent is declared invalid and $H(1)$ if it is upheld, for an expected payoff of $\theta H(1)$. Since P chooses between licensing and litigation, its payoff is $P(\theta) \equiv \max[H(\theta), \theta H(1)]$. Since $P(1) = H(1)$, we have $P(\theta) \geq \theta P(1)$: the patent owner can always get a fraction θ of the payoff from an ironclad patent by litigating.

II. Welfare Analysis of Probabilistic Patent Licenses

Before solving for the equilibrium two-part tariff in various settings, we develop the welfare analytics of enhanced review for patents that will then be licensed rather than litigated.[23] Section IV shows that such licensing indeed arises in our model. Figure 2 displays a simplified game tree for licensing with ideal PTO review.

A. *Ex Post Analysis*

Given that the innovation has been made and a patent of strength θ issued, ex post welfare $W(\theta)$ is the sum of the patent holder's licensing revenues, the downstream firms' profits, and the surplus enjoyed by final consumers. If all N downstream firms accept licenses with running royalty r, ex post welfare depends only on r, since fixed fees are just transfers: $W(\theta) = w(r(\theta))$,

version of this paper we showed that for sufficiently weak patents this is the equilibrium, and that $r \geq 0$ for weak patents even if $r < 0$ is allowed. We assume that both T and G are single-peaked in r on $[0, v]$.

[23] In our model, if the patent would be litigated, additional PTO review has no benefit, since the patent's validity will be determined before licensing anyway. Outside our model, the PTO might hold a comparative advantage over the courts in patent review: for instance, PTO review could be faster, reducing social costs borne in the interim before a patent validity ruling.

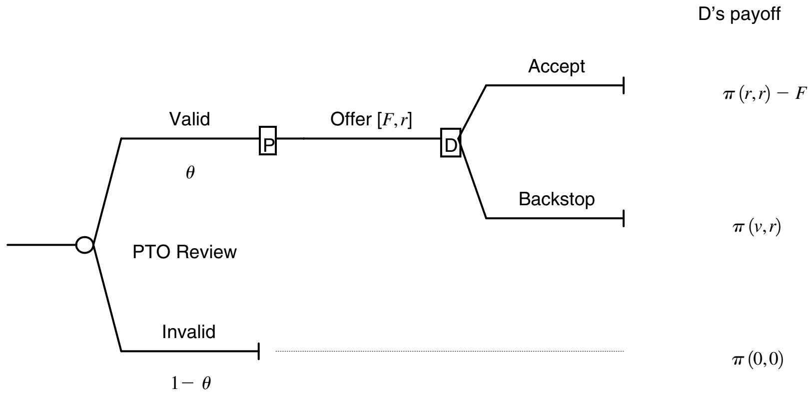

Figure 2. Licensing Game with Ideal PTO Review

where $w(r)$ denotes ex post welfare with royalty r. Increasing r reduces output, lowering welfare if the downstream price exceeds social marginal cost, as it does at $r > 0$ and even at $r = 0$ with imperfect downstream competition, so then $w'(r) < 0$ for all $r \geq 0$. Define

$$\lambda_{max} \equiv \max_{0 \leq t \leq v} |w'(t)|, \ \ \lambda_{min} \equiv \min_{0 \leq t \leq v} |w'(t)| > 0, \ \text{ and } \mu \equiv \frac{\lambda_{max}}{\lambda_{min}}.$$

Ideal PTO review (or litigation prior to licensing) gives expected welfare $\theta w(r(1)) + (1 - \theta) w(0)$, so its expected benefit is $B(\theta) = [\theta w(r(1)) + (1 - \theta) w(0)] - w(r(\theta))$. The Appendix proves:

THEOREM 1: $B(\theta) \geq [r(\theta) - \mu \theta r(1)] \lambda_{min}$.

Theorem 1 tells us that the ex post social benefit of ideal PTO review depends on the impact of review on the expected royalty rate, $r(\theta) - \theta r(1)$, and on the social benefit from a lower royalty rate, which is at least λ_{min}. For small values of v, $B(\theta) \approx [r(\theta) - \theta r(1)]|w'(0)|$.

Theorem 1 implies that $B(\theta) > 0$ if $r(\theta)/\theta r(1) > \mu$. This justifies the intuitive benchmark $\theta r(1)$, in that $B(\theta) > 0$ if $\mu \approx 1$, as is the case for small v, and $r(\theta)$ nontrivially exceeds the benchmark. For larger innovations, μ may not be close to 1, but the Appendix shows that, for instance, $\mu \leq (x(0,0)/x(v,v))[1 + v/(p(v) - v - c)]$ in Cournot oligopoly. By comparison, with $N \geq 2$, we identify below cases where $r(\theta)/\theta r(1) \approx N$, so $B(\theta) > 0$ for a wide range of patent sizes.

Theorem 1 casts the ex post analysis in terms of total welfare, but our model ignores litigation costs, which are borne by the patent holder and downstream firms. Given the PTO's review standards, those parties can choose whether or not the patent is litigated, but consumers cannot, so an externality-inspired approach would consider the effects of PTO review (or litigation) on consumers. With our assumption that $p(r)$ is linear, the Appendix proves:

THEOREM 2: *If* $r(\theta) \geq \theta r(1)$, *then consumers benefit from ideal PTO review.*

Intuitively, consumers are risk-loving in price, hence (with linear $p(r)$) in royalty rate, so if PTO review lowers the expected royalty rate, it benefits consumers in expectation. In the cases identified below where $r(\theta) > \theta r(1)$, the downstream firms have too little incentive to challenge a weak patent.[24] In such cases, consumers will value the right to trigger patent reexamination.[25]

B. *Ex Ante Analysis*

A firm's private incentive to engage in the R&D and patenting activities that lead to a patent of strength θ is $P(\theta)$. How does this compare to the social contribution $K(\theta)$?

We assume that θ is the true probability that the patent holder contributed the patented technology to society.[26] Denoting by $\bar{W}$ the welfare that would result if the patented technology were not available to society, $K(\theta) \equiv W(\theta) - [(1-\theta)W(0) + \theta\bar{W}]$. The Appendix proves:

THEOREM 3: *If downstream firms get their reservation (litigation) payoff, then if* $r(\theta) > \theta v$, *then* $P(\theta) > K(\theta)$.

That is, if the per-unit royalty exceeds the intuitive benchmark θv, namely the per-unit cost savings from the patented technology times the probability that the patent holder actually contributed this technology to society, then the patent holder's private return exceeds its social contribution. In expectation, the patent holder has inflicted a negative externality on others (it may have invented something, but certainly is charging high royalties); marginally profitable activities leading to such patents lower expected welfare.[27]

Theorem 3's benchmark, θv, differs from that in Theorems 1 and 2, $\theta r(1)$. For evaluating ex post benefits of PTO review in Theorems 1 and 2, what matters is the expected royalty level, $\theta r(1)$ versus $r(\theta)$. For comparing ex ante expected profits and contribution in Theorem 3, what matters is the patent holder's royalties versus its expected contribution.

Even if $P(\theta) < K(\theta)$, the relative incentives to pursue patents of different strengths may be biased. Consider a firm allocating its R&D and patenting budget between two activities. The first activity is a "conventional" line of research that, if it succeeds technically, will produce a useful but unsurprising technology, so there may already be prior art or a court may later deem the invention obvious. Thus, this activity generates patents of strength $\theta < 1$. A second, more creative line of research, if technically successful, will generate clearly novel and nonobvious results, leading to ironclad patents. The firm will allocate its R&D budget based on the relative reward to the two kinds of patents, $P(\theta)/P(1)$. For social efficiency, the allocation should be based on the relative contributions, $K(\theta)/K(1)$. If $P(\theta)/P(1) > K(\theta)/K(1)$, the firm will devote too much of its budget to the conventional line of research. In an extreme case, the firm may do little actual R&D and devote most of its resources to applying for patents covering technologies

[24] Jay Pil Choi (2002, 2005) argues that patent holders have weak incentives to challenge one another's patents if multiple weak patents are contributed to a patent pool. Our focus is instead on challenges by direct purchasers of the patented technology (downstream firms). Direct purchasers seem more likely to have legal standing, and although we are not aware of systematic evidence, we suspect that most patent licenses do not involve patent pools.

[25] The Electronic Frontier Foundation has a "patent-busting project" (www.eff.org/patent) that seeks to overturn some Internet and software-related patents by gathering prior art and requesting reexamination by the PTO. See Ian Austen, "Claiming a Threat to Innovation, Group Seeks to Overturn 10 Patents," *New York Times*, July 5, 2004. Joseph Miller (2004) advocates rewarding firms that successfully invalidate patents.

[26] Reiko Aoki and Jin-Li Hu (1999) analyze ex ante incentives of probabilistic patents, but focus entirely on the dilution (and other changes) of incentives for true innovation that arise because of the costs of enforcing properly issued patents; James Anton and Dennis Yao (2003) have a similar focus. Their approach assumes that when a patent is held invalid or not infringed, it is a court (or legal system) error. Alan Marco (2006) attempts to estimate the frequency of such errors if financial markets always get it right. By contrast, we assume that when a patent is held invalid, it is because the court's thorough scrutiny shows that it truly did not represent a novel, useful, nonobvious contribution of the patent holder.

[27] This can happen even with ironclad patents: in our linear Cournot example, $P(1)$ typically exceeds $K(1)$.

that very likely are already known or are obvious, if such weak applications often yield valuable weak patents rather than rejections. Under ideal PTO review, weak applications mostly yield rejections and occasionally yield ironclad patents; the expected contribution and reward to a technical success in the conventional research line thus become $\theta K(1)$ and $\theta P(1)$, respectively, eliminating the bias. The Appendix proves:

THEOREM 4: *If* $B(\theta) > 0$, *then* $P(\theta)/P(1) > K(\theta)/K(1)$.

Theorem 4 links ex post and ex ante analysis: if there are ex post benefits of ideal PTO review, then there is also an ex ante bias toward seeking weak patents, which such review eliminates.

Theorems 1–4 do not depend on a specific model of licensing: thus, for instance, they apply if royalties are elevated because of holdup concerns (Shapiro 2008), or are negotiated in the shadow of asymmetric litigation costs. We now turn to our specific model of licensing.

III. Downstream Firms That Do Not Compete

Suppose the downstream firms operate in separate markets, so each firm's profits do not depend on others' costs: $\pi_2(a,b) \equiv 0$. Then, running royalties would reduce profits through double marginalization and $T(r)$ is maximized at $r = 0$. Nor would running royalties help P extract a bigger share of joint profits, since, simplifying notation, a downstream firm's expected payoff from litigation is $\theta\pi(v) + (1-\theta)\pi(0)$, independent of r, so $G(r,\theta) = T(r) - [\theta\pi(v) + (1-\theta)\pi(0)]$. Thus, $r(\theta) = 0$ for all θ. It follows that $B(\theta) = 0$: further PTO (or judicial) review of validity does not affect total welfare or consumer surplus. Moreover, $F(\theta) = \theta F(1)$, for otherwise either D or P would prefer litigation. Hence P's payoff is proportional to θ: specifically, $P(\theta) = \theta[\pi(0) - \pi(v)]$. Indeed, nobody cares in expectation, ex post, whether the patent is licensed under uncertainty or further reviewed before licensing. Ex ante, $K(\theta) \equiv \theta[W(0) - \bar{W}] > 0$, so $P(\theta)/K(\theta) = [\pi(0) - \pi(v)]/[W(0) - \bar{W}]$, which is between 0 and 1 and is independent of θ, so incentives are not biased among R&D and patenting strategies that lead to patents of different strengths. Summarizing, we have a reassuring benchmark:

THEOREM 5: *If* $r(\theta) = 0$ *for all* θ, *additional PTO review of patents generates no ex ante or ex post benefits in our model. If downstream firms are not rivals, then* $r(\theta) = 0$ *for all* θ.

IV. Downstream Firms That Compete

When downstream firms compete against one another, their prices fall below the joint profit maximizing level. As a result, total profits (including P's) rise when all downstream firms face a small positive running royalty r. If the downstream industry is reasonably competitive, the running royalty m that supports the downstream monopoly price, maximizing total profits $T(r)$, can be large. We assume that $m \geq v$, so $r = v$ maximizes $T(r)$ in the feasible range $0 \leq r \leq v$.[28]

Since fixed fees allow P to capture joint profits $T(r)$ minus downstream firms' reservation payoffs, one might thus expect $r(\theta) = v$. That is correct for small θ, although not in general, because P can lower each downstream firm's reservation payoff $\theta\pi(v,r) + (1-\theta)\pi(0,0)$ by lowering

[28] $m \geq v$ if and only if the downstream price charged by an integrated monopolist using the new technology, $p(m)$, is no lower than the oligopoly equilibrium price with the old technology, $p(v)$. With even moderate downstream competition, this will hold for quite substantial innovations.

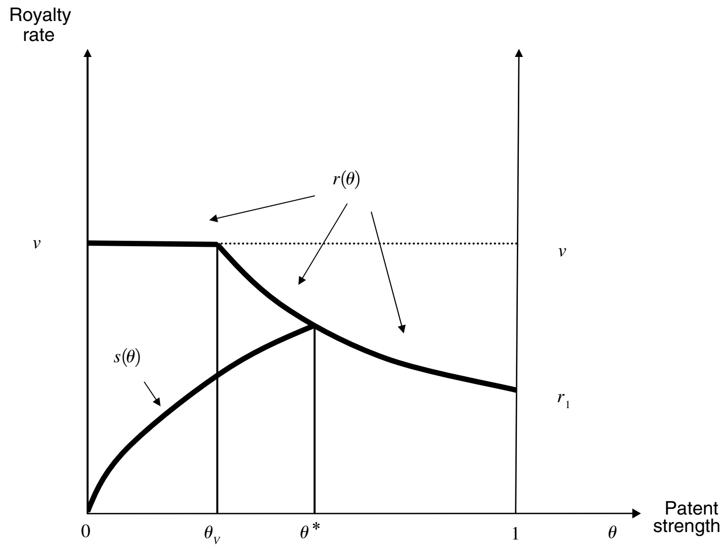

Figure 3. Equilibrium Royalty Rates

the running royalty rate r to the firm's rivals.[29] P sets r to maximize $G(r,\theta)$, not $T(r)$, and $G_r(r,\theta) = \mathrm{T}'(r) - \theta\pi_2(v,r) < \mathrm{T}'(r)$. But this rent-shifting effect is proportional to θ, since an infringer faces rivals with marginal cost r only if the patent is found valid. For weak patents, rent-shifting does not much modify joint profit maximization and sharing:[30] for $\theta \leq \theta_V \equiv \mathrm{T}'(v)/\pi_2(v,v)$, we have $G_r(v,\theta) \geq 0$ and $r(\theta) = v$.[31] A downstream firm accepts this license because $F(\theta) = -(1 - \theta)[\pi(0,0) - \pi(v,v)] < 0$.[32] As the Appendix shows, the patent holder strictly prefers such licensing to litigation. In sharp contrast with Theorem 5, we thus have:

THEOREM 6: *For weak patents* $(\theta \leq \theta_V)$ *licensed to downstream rivals using unrestricted two-part tariffs,* $r(\theta) = v$ *and* $F(\theta) = -(1 - \theta)[\pi(0,0) - \pi(v,v)] < 0$.

Theorem 6 holds because high per-unit royalty rates maximize joint profits and this dominates royalty setting for weak patents. As a result, consumers gain nothing from the new technology, since each downstream firm's private marginal cost is the same as under the backstop technology. Ideal PTO review unambiguously benefits consumers and efficiency ex post. If the patent is ruled invalid, royalties drop to zero. If it is upheld, royalties become $r(1) \leq v$ rather than $r(\theta) = v$, so consumer and total welfare will either rise or be unchanged. The Appendix shows that if downstream competition involves strategic substitutes (as in Cournot competition), then $r(1) < v$ and welfare strictly improves even if the patent is upheld. Figure 3 displays $r(\theta)$ in the case where $r(1) < v$.

[29] This rent-shifting effect is recognized in the literature on the licensing of ironclad patents; see Sen and Tauman (2007). Segal (1999) studies this effect much more generally.

[30] Rent-shifting would be a big impediment to cartelizing an industry without a patent, because each downstream firm might hope to be (very profitably) the only one outside the cartel. Inadvertently, Blonder-Tongue ensures that if such an outsider successfully challenges a weak patent for its own use, it also disrupts the cartel.

[31] Since $G_{r\theta}(r,\theta) = -\pi_2(v,r) < 0$, the optimal running royalty $r(\theta)$ is weakly decreasing in θ.

[32] There exists $\theta^* \geq \theta_V$ such that negative fixed fees are optimal for all $\theta \leq \theta^*$, as shown in Figure 3.

Many policies that benefit consumers and raise total welfare ex post also reduce patentees' payoffs, thereby worsening ex ante incentives. Here, however, the prospect of ideal PTO review, while reducing patentees' profits, strictly improves ex ante incentives in several respects, whatever the elasticity of supply of innovation. The Appendix proves:

THEOREM 7: *For weak patents* ($\theta \leq \theta_V$) *licensed to downstream rivals using unrestricted two-part tariffs,* $r(\theta) = v$, $B(\theta) > 0$, *and* $P(\theta) > K(\theta)$. *If also* $\theta < [w(0) - w(v)]/[w(0) - w(v) + vx(v,v)]$, $K(\theta) < 0$. *In the range where* $K(\theta) > 0$, $P(\theta)/K(\theta)$ *strictly decreases with* θ. *Ideal PTO review ensures that the patent holder's social contribution is positive and that the ratio of profits to social contribution does not vary with patent strength.*

By equating $P(\theta)/K(\theta)$ across patent strengths, ideal PTO review eliminates profitable opportunities to do harm ($K(\theta) < 0$) and eliminates a bias toward seeking weak patents.

For Cournot oligopoly with linear demand and constant marginal costs, one can directly calculate $r(\theta)$, $F(\theta)$, $P(\theta)$, $B(\theta)$, and $K(\theta)$ in terms of N and v/A,[33] where A is the difference between the demand intercept and the production cost using the patented technology.[34] For example, with $N = 5$ and $v/A = 0.1$, $r = v$ for $\theta \leq 0.41$, $F(\theta) < 0$ for $\theta < 0.48$, $K(\theta) < 0$ for $\theta < 0.18$, and $P(\theta)/K(\theta) > 2$ for $\theta \leq 0.48$.

V. Negative Fixed Fees Not Feasible

Section IV's results involve negative fixed fees, but we do not know how often such fees are feasible or used in practice.[35] Large negative fixed fees may induce entry, and may carry antitrust risk.[36] If a patent for which $F(\theta) < 0$, when feasible, is licensed when negative fixed fees are not feasible, it will be licensed with no fixed fee (since G is single-peaked in r).

If all downstream firms pay a pure running royalty s, P's income per downstream firm is $R(s) \equiv sx(s,s) = T(s) - \pi(s,s)$. In the range $0 \leq s \leq v$, $T(s)$ increases with s, and $\pi(s,s)$ falls with s, so $R(s)$ increases with s and, unless it prefers to litigate, P will license at the highest royalty that downstream firms will accept rather than litigate. That is, $s = s(\theta)$, defined by $\pi(s(\theta), s(\theta)) = \theta\pi(v,s) + (1-\theta)\pi(0,0)$.[37] The Appendix gives conditions under which P prefers such linear licensing to litigation in our model.[38]

How does $s(\theta)$ compare to our benchmarks? If D litigates and loses, it will be at a cost disadvantage $v - s$ relative to its licensed rivals, so its downside from litigating is proportional to $v - s$

[33] The Supplementary Materials associated with this paper (available at http://www.aeaweb.org/articles.php?doi=10.1257/aer.98.4.1347) work out all of these functions in this special case, which is often used in the oligopoly and licensing literature (e.g., Kamien 1992; Sen and Tauman 2007).

[34] Alternatively, $2(v/A)$ approximates the proportionate increase in first-best welfare from the innovation.

[35] We are unaware of any systematic empirical evidence on how often licenses contain negative fixed fees. Bharat Anand and Tarun Khanna (2000) assemble a sizeable database of licensing contracts but lack sufficient information on the use of running royalties versus fixed fees to reach reliable conclusions. The Federal Trade Commission (2002, 2005) reports on the use of negative fixed fees in certain pharmaceutical patent agreements.

[36] The Federal Trade Commission has brought several antitrust cases challenging negative fixed fees, known as "reverse payments" in antitrust circles, in agreements between vertically integrated patent holders (branded pharmaceutical suppliers) and would-be generic competitors. See Jeremy Bulow (2004), Herbert Hovenkamp, Mark Janis, and Lemley (2003), Shapiro (2003), and Robert Willig and John Bigelow (2004). The patent holder may be able to disguise negative fixed fees (for example, it might transfer know-how to the licensee, or agree to a side deal).

[37] D's payoff if it litigates and loses is $\pi(v,s)$ if P would then hold D to its backstop payoff. P will indeed do so if it can charge a positive fixed fee or, if licenses are constrained to be linear, if it would optimally charge a running royalty of v. The supplementary materials show that this is optimal in the linear Cournot example for small values of v, and we assume it below.

[38] In the linear Cournot case with $N = 5$, $v/A = 0.1$ and $\theta = 0.2$, and the patent holder's expected payoff from licensing is about twice as large as from litigating. With $N = 10$, the ratio is about three to one.

and to $|\pi_1(s,s)|$. In contrast, if D litigates and *wins* it will not gain any competitive advantage over its rivals. Rather, it will have lowered industry-wide costs from s to zero, raising its profits from $\pi(s,s)$ to $\pi(0,0)$. Its upside is thus proportional to s and to $|\pi_1(s,s) + \pi_2(s,s)|$.

We thus define the oligopoly's *relativity* coefficient $\rho \equiv \{|\pi_1(0,0)|\}/\{|\pi_1(0,0) + \pi_2(0,0)|\}$ as the relative importance to a firm of small changes in its own costs versus small changes in industry-wide costs (evaluated at $s = 0$). The extent to which $\rho > 1$ measures the strength of downstream competition. For example, if downstream firms are symmetric Cournot oligopolists with constant marginal costs, the Appendix shows that $\rho \geq N$ for linear or constant-elasticity demand (and gives a more general expression for ρ). If the downstream industry is a Bertrand duopoly with differentiated products, ρ is higher, the closer substitutes are the two downstream products. Using $\pi(s,s) = \theta\pi(v,s) + (1-\theta)\pi(0,0)$, the Appendix proves:

THEOREM 8: *For small* v, $s(\theta) \approx \theta v[\rho/(1 + (\rho - 1)\theta)]$.

If $\rho > 1$, Theorem 8's approximation for $s(\theta)$ exceeds θv for $\theta \in (0,1)$ and is in turn approximately $\rho\theta v$ for small θ. As $\rho \to 1$, $s(\theta) \to \theta v$, confirming that it is now relativity that enables the running royalty to exceed the ex ante benchmark level θv (in Theorem 3). In contrast, in Theorem 6, the mechanism was joint profit maximization (tempered for stronger patents by rent shifting). Figure 3 displays $s(\theta)$ and $r(\theta)$ in the case where $r(1) < v$.

If a patent is linearly licensed as in Theorem 8, Theorem 1 (recalling $r(1) \leq v$) implies $B(\theta) \geq \theta v\{\rho/[1 + (\rho - 1)\theta] - \mu\}\lambda_{\min}$. For weak patents, the expression in brackets is robustly positive; for instance, with Cournot oligopoly, linear demand, and constant costs, $\rho = N$ and $\mu \leq [x(0,0)/x(v,v)][1 + v/(p - c - v)]$. For $N = 5$ and $v/A = 0.1$, linear licenses are used for $\theta < 0.48$. In this range, $s(\theta)/\theta v$ declines from nearly 4 near $\theta = 0$ to 1.7. While $K(\theta) > 0$ for all $\theta > 0$, $P(\theta)/K(\theta)$ declines with θ from over 7 near $\theta = 0$ to 2.0 at $\theta = 0.48$; for an ironclad patent, $P(1)/K(1) = 1.13$. The ratio $B(\theta)/[B(\theta) + K(\theta)]$ declines with θ but for very weak patents is near 0.5, meaning that ideal PTO review contributes roughly as much to society as (in expectation) did the patent holder without such further PTO review.

VI. Variations and Extensions

A. *Vertically Integrated Patent Holder*

Our analysis extends easily to the case in which the patent holder is vertically integrated, competing downstream with N other downstream firms. Define $\pi^I(a,b)$ as the profits of a downstream firm with cost a, given that the other downstream firms have cost b and the patent holder competes using the patented technology, and write $\phi(r)$ for P's product market profits if all rivals pay royalty r. We assume that $\pi^I(a,b)$ satisfies the three conditions assumed above for $\pi(a,b)$, and that $\phi'(r) \geq 0$. The Appendix shows how our analysis and main results carry over. With unrestricted two-part tariffs, for weak patents $r(\theta) = v$ and $K(\theta) < 0$, and for all patents $B(\theta) > 0$. If the patent holder faces only one downstream rival, then $r(\theta) = v$ for all θ and $F(\theta) < 0$ for all $\theta < 1$. If no downstream firm using the backstop technology could profitably compete against the patent holder, then, again, $r(\theta) = v$ for all θ and $F(\theta) < 0$ for all $\theta < 1$: the patent holder pays each downstream firm to agree not to infringe or challenge the patent, which is tantamount to exit.[39] If negative fixed fees are not feasible, then, as in Theorem 8, $s(\theta) \approx \rho^I\theta v$ where now

[39] If P's downstream division were less efficient than downstream firms, these outcomes would not maximize joint profits. If feasible, P would prefer to commit to shut down its downstream division. Technically, T would not be increasing in r, as we are assuming.

$\rho^I = |\pi_1^I(0,0)|/(|\pi'^I_1(0,0) + \pi_2^I(0,0)|)$. In the linear Cournot example, $\rho^I = (N + 1)/2$, so again $B(\theta) > 0$ for all θ.

B. *Short-Term versus Long-Term Licenses*

We assumed above that a downstream firm's license remains in force if the patent is upheld after P litigates it with another downstream firm. Outside our model, licensees often make specific investments to use the patented technology, which provides an efficiency reason to design licenses that way. The Appendix shows that our results grow stronger if the patent holder can offer "short-term" licenses that do not survive a finding of validity. With unrestricted two-part tariffs, $r(\theta) = v$ for all θ, so the ex post welfare analysis is the same as it was above if $\theta_V = 1$. Ex ante, for all $\theta < 1$, $P(\theta)$ is higher than we derived above, and $K(\theta)$ is unchanged or lower, so the bias resulting from $P(\theta)/K(\theta) > P(1)/K(1)$ is stronger than above.

C. *Linear Licenses*

As Kamien (1992) noted, running royalties appear to be common. It is our impression that this is true (contrary to Theorem 5) even if licensees do not compete; the reasons are presumably outside our model, such as risk aversion, asymmetric information, and moral hazard.[40]

When licensees do not compete and pass-through is linear, the Appendix shows that consumers' risk preferences between a certain royalty $s(\theta)$ and the uncertain result of PTO review are reflected in each downstream firm's similar preferences. Since a downstream firm will infringe if $s(\theta)$ is too high, replicating through litigation the validity gamble of PTO review, consumers are protected by downstream firms as their agents against royalties that hurt them relative to first determining patent validity. Since consumers are risk-loving in price, this implies that $s(\theta) < \theta s(1) = \theta v$; however, to first order, $s(\theta) = \theta v$.

When licensees compete, the equilibrium running royalty is the much higher $s(\theta)$ calculated in Theorem 8 for all θ, not just for $\theta \leq \theta_V$. Since $s(\theta) > \theta v$, $s(\theta)x(s(\theta),\ s(\theta)) > \theta v x(v,v)$ so the patent holder prefers licensing to litigation. The Appendix shows that $B(\theta)/[B(\theta) + K(\theta)] \geq [s(\theta)/\theta v - \mu]p'(s)$. For small θ, $s(\theta)/\theta v \approx \rho$ so $B(\theta)/[B(\theta) + K(\theta)] \geq [\rho - \mu]p'(s)$. This can easily exceed unity, in which case $K(\theta) < 0$.

D. *Enhanced Review of Patents*

Short of "ideal" review, more realistic "enhanced" patent review prior to licensing uncovers some additional information about patent validity, inducing a mean-preserving spread on patent strength. Enhanced review thus decreases (increases) the expected value of any concave (convex) function of θ. Enhanced review reduces the expected royalty rate if $r(\theta)$ is concave in θ. By Theorem 2, this is a sufficient condition for consumers to benefit from enhanced review. By Theorem 1, enhanced review generates ex post benefits if the concavity of $r(\theta)$ is sufficient relative to μ, a mild condition for small values of v. Theorem 5 directly applies to any enhanced review. If $\theta_V = 1$, Theorem 6 shows that $r(\theta)$ is concave (not only weakly, because $r(0) = 0$), so enhanced review is ex post beneficial if and only if it has positive probability of actually invalidating the patent.[41] If $\theta_V < 1$, $r(\theta)$ is globally concave if it is concave on $[\theta_V, 1]$.

[40] For example, Sugato Bhattacharyya and Francine Lafontaine (1995) construct a model in which linear sharing rules are optimal due to two-sided moral hazard. Because our model does not capture these reasons for variable royalties, we do not pursue the observation that a rule of lump-sum royalties might be appealing only in the model.

[41] The litigation process induces a series of mean-preserving spreads on patent strength, with $\theta = 0$ never arising until final judgment. For patents licensed to downstream rivals, private parties may not pursue litigation to final

If negative fixed fees are not feasible, Theorem 8 implies that $s(\theta)$ is concave in θ for small v.[42] Therefore, enhanced review reduces the expected royalty rate, so long as the approximation in Theorem 8 is good. Theorem 1 then tells us that such review generates ex post social benefits, at least for small v, and Theorem 2 implies that it benefits consumers.

We also can evaluate the impact of a marginal enhancement in PTO review that with small probability invalidates the patent and otherwise strengthens it slightly. Such enhancement reduces the expected per-unit royalty if and only if $r(\theta)/\theta > r'(\theta)$.[43] Applying Theorem 6, with unrestricted two-part tariffs we know that $r'(\theta) \leq 0$ for all θ, so this condition is satisfied for all θ, and a marginal enhancement in PTO review lowers the expected royalty rate, raising ex post welfare and benefiting consumers. Another notion of a marginal enhancement in PTO review is one that leads to a distribution of posteriors, all near the prior θ, with mean θ. The change in expected welfare resulting from this type of enhancement depends purely on local curvature of the composite function $w(r(\theta))$.

E. *Patent Validity and Patent Scope*

We have cast our analysis so far in terms of uncertainty about the validity of relevant claims in the patent. There is often also (or instead) uncertainty about whether a downstream firm's product actually infringes those claims.[44] The two kinds of uncertainty are equivalent if there is just one licensee. Our analysis extends to cases where patent scope or infringement rather than patent validity is the key issue, if a finding of noninfringement against one downstream firm implies that other downstream firms also are not infringing and that these firms can stop paying running royalties.

F. *Litigation Costs, Bargaining, and Private Information*

Litigation costs make licensing even more attractive relative to litigation than our model suggests. How do they affect the terms on which a probabilistic patent is licensed in the shadow of litigation? If, as above, P makes take-it-or-leave-it offers, of course, it can demand more; if downstream firms had commitment power, they could offer less. Extending the model to include litigation costs would thus seem most natural if we also extended it to more general bargaining, which becomes complex when competing downstream firms bargain with P. If litigation costs are unrelated to (θ, v) they may dominate the bargaining for small, weak patents, but if bargaining skill and litigation costs are symmetric, their effect will tend to be neutral, restoring our results.[45]

judgment. After a verbal ruling dismissing Rambus's patent infringement case against Infineon, but before a written opinion that could have set a precedent for other infringement cases brought by Rambus, Rambus and Infineon settled. See Don Clark, "Rambus, Infineon Reach Settlement," *Wall Street Journal,* March 22, 2005.

[42] More precisely, it is approximated by a strictly concave function, so it is effectively concave for mean-preserving spreads that are not too small relative to the quality of the approximation.

[43] Writing $\gamma \leq 1$ for the probability the patent survives the review, the strength of the surviving patent is θ/γ. The expected per-unit royalty resulting from the review is thus $\gamma r(\theta/\gamma)$. The review reduces the expected per-unit royalty if and only if this expression is increasing in γ near $\gamma = 1$. Differentiating with respect to γ and requiring the resulting expression to be positive gives the inequality in the text.

[44] Michael Waterson (1990) studies how uncertainty about infringement (patent scope) affects rivals' design decisions; Bessen and Meurer (2008) view this lack of clarity about property rights as the key problem in patent policy.

[45] Farrell and Robert Merges (2004) explore the role of relativity in determining parties' effort (expenditure) in patent litigation, noting that this makes both litigation costs and the resulting probability θ endogenous. Shapiro (2008) provides a model in which royalties for a probabilistic patent are determined by bargaining; his model involves a single licensee and focuses on patent hold-up.

In general, both the patent holder and downstream firms may have initially private information concerning the patent's strength; θ will be common knowledge, as we assumed, if each side finds it easy enough, perhaps aided by the pre-enhancement PTO review process and the resulting published patent, to acquire all the information initially held by the other, so each will do so before negotiating a license. While this seems a natural first-cut simplifying assumption, it does rule out one concern about weak patents: that a patent holder who knows its patent is weak might be able to fool less well-informed downstream firms into paying substantial royalties.[46]

G. *Litigation Credibility*

As we noted above, if $N - 1$ firms sign lucrative licenses and one infringes, the patent holder might be reluctant to litigate and put its licensing revenues at risk. Choi (1998) models this question of litigation credibility, but does not analyze licensing terms. In his model, either the vertically integrated patent holder excludes rivals completely, or they enter with no royalties. Yet, clearly, P has a strong incentive to ensure that it has a credible threat to sue an infringer. Several mechanisms may help it. First, if infringers divert substantial sales from licensees, as they might (especially if the downstream industry is highly competitive) due to their cost advantage from not paying royalties, litigation may well be credible. Second, reputation effects can make litigation credible. Third, licenses may contractually commit P to sue, for instance, by allowing licensees to stop paying royalties if P fails to challenge an infringer.

If none of these (or other) mechanisms establishes litigation credibility, P may have to adjust its licensing terms to do so. Generalizing our model in this direction exposes the inherent relationship between litigation credibility and relativity. Litigation becomes more credible if licensees must continue paying royalties even if the patent is overturned. *Blonder-Tongue* limits what a license can do in this respect, but licenses may be able to bundle trade secrets (or other patents) with a weak patent. However, the more effectively the license ensures that running royalties continue even if another downstream firm successfully challenges, the greater is the upside to a challenge, and the less the patent holder can exploit relativity. These issues will be a fertile area for future work.

VII. Conclusion

Since far more patents are licensed than litigated, the economic impact of questionable patents depends largely on how they are licensed. We modeled how licensing terms vary with patent strength. In our licensing model, weak patents licensed to downstream firms that are not rivals (to each other or to the patent holder) command correspondingly low royalties, so there are no (expected) social benefits of examining these patents more closely. In sharp contrast, weak patents on technology used by downstream firms that are rivals (to each other or to the patent holder) command surprisingly large running royalties, especially if licenses can use unrestricted two-part tariffs. There are large social benefits, ex post and, perhaps more importantly, ex ante, of better examining commercially significant patents that will be licensed to downstream rivals.

Closely scrutinizing the hundreds of thousands of patent applications filed each year, many of which end up having no commercial significance, would be very costly. Our analysis suggests a more targeted approach: reexamination of issued patents covering valuable technology that is useful to multiple downstream firms that compete against each other or against the patent holder.

[46] A patent applicant is required to disclose to the PTO relevant prior art that it knows about, but is not required to search, and it is difficult to enforce this requirement. Meurer (1989) considers signaling issues when a vertically integrated holder of a probabilistic patent can litigate with a single downstream rival.

Our analysis thus supports current proposals to expand post-grant review of commercially significant patents, but also identifies downstream competitive conditions as a key indicator of the value of further review.

We have analyzed two mechanisms by which weak patents can have strong (and adverse) economic effects. With unrestricted two-part tariffs, weak patents can be used to impose high per-unit royalties (along with negative fixed fees), raising downstream marginal costs and thus moving the downstream price closer to the monopoly price. With linear licenses, weak patents still command surprisingly high per-unit royalties because challenging the patent is a public good for the downstream firms. We stressed that these effects do not merely worsen the ex post deadweight loss from patents. Perhaps worse, they distort the innovation incentives that patents are meant to provide. Further empirical work is needed to determine the importance of each of these mechanisms in practice. Weak patents may be socially costly in other ways as well: they can lead to costly litigation, they can create a danger of patent hold-up (both of users and of subsequent innovators), and they can induce defensive patenting, which can itself lead to yet more weak patents in a vicious cycle. Sound patent policy calls for the targeted application of resources to review patent applications, and issued patents, where these ex post social costs are greatest and where ex ante incentives are most skewed.

Appendix

PROOF OF THEOREM 1:

$B(\theta) \geq [r(\theta) - \mu\theta r(1)]\lambda_{\min}$.

Applying the intermediate value theorem to $w(r)$, and using $0 \leq r(\theta), r(1) \leq v$, there exist $t_1 \in [0,v]$ and $t_2 \in [0,v]$ such that $w(r(1)) = w(0) + r(1)w'(t_1)$ and $w(r(\theta)) = w(0) + r(\theta)w'(t_2)$. Substituting into $B(\theta) = [\theta w(r(1)) + (1-\theta)w(0)] - w(r(\theta))$ and simplifying gives $B(\theta) = \theta r(1)w'(t_1) - r(\theta)w'(t_2)$. Since $w'(t_1) < 0$ and $w'(t_2) < 0$, we have $B(\theta) = r(\theta)|w'(t_2)| - \theta r(1)|w'(t_1)| \geq r(\theta)\lambda_{\min} - \theta r(1)\lambda_{\max} = [r(\theta) - \mu\theta r(1)]\lambda_{\min}$.

We next discuss upper bounds on μ, precisely in the Cournot case and heuristically more generally. We have $w'(r) = [p(r) - c](d/dr)Nx(r,r)$, where c is the marginal social cost of production using the patented technology. Differentiating again, and using $Nx(r,r) \equiv X(p(r))$, yields $w''(r) = [p'(r)]^2[X'(p) + (p-c)X''(p)] + p''(r)(p-c)X'(p)$. Assuming $p''(r) = 0$, this implies that $w''(r)$ has the sign of $X'(p) + (p-c)X''(p)$, which is negative if demand is linear (or concave) in the range $p(0) \leq p \leq p(v)$, whatever the oligopoly behavior. That implies that $\lambda_{\min}$ occurs at $r = 0$, $\lambda_{\max}$ occurs at $r = v$, $\mu = |w'(v)|/|w'(0)|$, and $B(\theta) \geq r(\theta)|w'(0)| - \theta r(1)|w'(v)|$.

For Cournot oligopoly, our formula for $w'(r)$ yields $|w'(r)| = [p'(r)][\varepsilon N(p(r) - c)/p(r)][x(r,r)]$ where ε is the absolute value of the elasticity of demand. Since μ is a ratio of $|w'(t)|$'s, it is the product of the ratios of the three factors in brackets making up $|w'(r)|$. If $p'(r)$ is a constant, the first ratio will be unity. In Cournot oligopoly, $(p(r) - r - c)/p(r) = 1/(\varepsilon N)$, so $N\varepsilon\{(p(r) - c)/p(r)\} = \{(p(r) - c)/(p(r) - r - c)\}$, which equals unity at $r = 0$ and is increasing in r. Therefore, the second ratio is bounded above by $(p(v) - c)/(p(v) - v - c) = 1 + v/(p(v) - v - c)$. The third ratio is bounded above by $x(0,0)/x(v,v)$, which reflects the proportionate increase in output resulting from the innovation, if it is available royalty-free. Therefore, in Cournot oligopoly $\mu \leq [x(0,0)/x(v,v)][1 + v/(p(v) - c - v)]$. Below, we will be comparing μ to numbers typically above two.

PROOF OF THEOREM 2:

If $r(\theta) \geq \theta r(1)$, then consumers benefit from ideal PTO review.

We prove this for weakly concave p, $p''(r) \leq 0$, not just for linear $p(r)$. With $p''(r) \leq 0$, for any t_1 and t_2 and $\lambda \in [0,1]$, $p(\lambda t_1 + (1-\lambda)t_2) \geq \lambda p(t_1) + (1-\lambda)p(t_2)$. Write $CS(p)$ for consumer surplus as a function of the downstream price. Since $CS(p)$ is decreasing in price, this in turn implies that $CS(p(\lambda t_1 + (1-\lambda)t_2)) \leq CS(\lambda p(t_1) + (1-\lambda)p(t_2))$. Since $CS(p)$ is convex in price, $CS(\lambda p(t_1) + (1-\lambda)p(t_2)) \leq \lambda CS(p(t_1)) + (1-\lambda)CS(p(t_2))$. Combining these two inequalities gives $CS(p(\lambda t_1 + (1-\lambda)t_2)) \leq \lambda CS(p(t_1)) + (1-\lambda)CS(p(t_2))$, so the composite function $S(t) \equiv CS(p(t))$ is convex in t. This proves that consumers are risk-loving in the royalty rate (a fact useful beyond this theorem), so they benefit from mean-preserving spreads in r. Since they also prefer lower royalty rates, they welcome ideal PTO review, which is a combination of a mean-preserving spread and a possible reduction in the expected royalty rate when $r(\theta) \geq \theta r(1)$.

PROOF OF THEOREM 3:

If $r(\theta) > \theta v$, then $P(\theta) > K(\theta)$.

Each downstream firm's reservation (litigation) payoff is $\theta\pi(v,r) + (1-\theta)\pi(0,0)$. Without the patent holder's activities, there is a probability θ that the patented technology would be unavailable, resulting in per-firm profits of $\pi(v,v)$; with probability $1-\theta$, the technology would be available without royalties, resulting in per-firm profits of $\pi(0,0)$. The difference between equilibrium payoff and expected payoff without the patent holder is thus $\theta[\pi(v,r) - \pi(v,v)]$, which is (weakly) negative when $r \leq v$ since profits are increasing in rivals' cost levels. When downstream firms actually compete, $(\pi_2 > 0)$ and $r < v$, downstream firms are strictly hurt.

Without the patent holder's activities, with probability θ the patented technology would be unavailable, which for consumers is the same as there being a royalty rate of v. With probability $1-\theta$ the patented technology would be available without royalties. The proof of Theorem 2 showed that consumers are weakly risk-loving in the royalty rate. Since they also prefer lower royalties, if $r(\theta) \geq \theta v$, they are harmed by the patent holder's activities, and strictly so if $r > \theta v$.

Since the patent holder's activities harm both downstream firms and consumers, the patent holder's profits must exceed its social contribution; the proof shows that if $\theta < 1$ and $\pi_2 > 0$, at least one group is strictly harmed, so the comparison is strict. Like Theorem 2, Theorem 3 holds for $p''(r) \leq 0$, not just for linear $p(r)$.

PROOF OF THEOREM 4:

If $B(\theta) > 0$, then $P(\theta)/P(1) > K(\theta)/K(1)$.

By definition, $B(\theta) \equiv [\theta W(1) + (1-\theta)W(0)] - W(\theta)$ and $K(\theta) \equiv W(\theta) - [(1-\theta)W(0) + \theta\overline{W}]$. Adding these together gives $B(\theta) + K(\theta) = \theta[W(1) - \overline{W}]$. Since $K(\theta)$ is the patent holder's contribution if the patent is not examined more carefully, and $B(\theta)$ is the additional benefit arising from ideal PTO review, their sum is the patent holder's expected contribution under ideal PTO review, which is precisely the contribution from an ironclad patent times the probability that the patent will indeed be found valid under ideal PTO review.

Since $K(1) = W(1) - \overline{W}$, this implies that $B(\theta) + K(\theta) = \theta K(1)$. Writing this as $K(\theta) = \theta K(1) - B(\theta)$, when $B(\theta) > 0$ we must have $K(\theta) < \theta K(1)$. Since $P(\theta) \geq \theta P(1)$, $P(\theta)/K(\theta) \geq \theta P(1)/K(\theta)$. Using $K(\theta) < \theta K(1)$, we get $P(\theta)/K(\theta) \geq \theta P(1)/K(\theta) > \theta P(1)/\theta K(1) = P(1)/K(1)$.

PROOF OF THEOREM 6:

The patent holder strictly prefers licensing to litigation.

Since $H(\theta) = \max_r G(r;\theta)$ is the upper envelope of linear functions of θ, it is convex in general and linear where r does not vary with θ.

For $\theta \leq \theta_V$, $r(\theta) = v$ and $H(\theta) = T(v) - [\theta\pi(v,v) + (1-\theta)\pi(0,0)]$. Since $\lim_{\theta\to 0} H(\theta) = T(v) - \pi(0,0) = T(v) - T(0) > 0$ (recall that $T(r)$ increases with r for $0 \leq r \leq v$), for sufficiently weak patents we must have $H(\theta) > \theta H(1)$.

If $r(1) = v$, then $r(\theta) = v$ for all $\theta > 0$, making H linear in θ. Therefore, the licensing payoff $H(\theta)$ is a straight line that starts above 0 and ends up at $H(1)$. The litigation payoff is a straight line starting at 0 and also ending up at $H(1)$. So the payoff from licensing is strictly greater than the payoff from litigation for all $\theta < 1$.

Alternatively, if $r(1) < v$ then in the range $0 < \theta \leq \theta_V$, $H(\theta)$is as just discussed. For $\theta \geq \theta_V$, r varies, so $H(\theta)$is a convex function of θ on $(\theta_V, 1]$. Therefore, if the $H(\theta)$ curve lies above the $\theta H(1)$ line as $\theta \to 1$, where the two meet, then $H(\theta) > \theta H(1)$ for all θ. But, since it is convex and begins above the line, the $H(\theta)$curve lies above the $\theta H(1)$ line near $\theta = 1$ if and only if $H'(1) \leq H(1)$. Now $H'(1) = \pi(0,0) - \pi(v, r(1))$ and $H(1) = T(r(1)) - \pi(v, r(1))$. So $H'(1) \leq H(1)$ if and only if $\pi(0,0) = T(0) \leq T(r(1))$, a condition that must be satisfied since $T(r)$ is increasing in r for $r \leq v$.

If the downstream oligopoly game involves strategic substitutes, then $r(1) < v$.

For $\theta = 1$, the patent holder maximizes $G(r;1) = rx(r,r) + \pi(r,r) - \pi(v,r)$. Since $G(r;1)$ is single-peaked, $r(1) < v$ if and only if $G(r;1)$ is declining in r at $r = v$. Differentiating $G(r;1)$ with respect to r and evaluating at v gives:

$$G_r(v;1) = x(v,v) + \pi_1(v,v) + v[x_1(v,v) + x_2(v,v)].$$

Since output declines when costs rise uniformly, the term in square brackets is negative. To sign the sum of the first two terms, note that the second term is the effect on profits of marginally higher own unit costs, π_1. We can decompose π_1 into a "direct" effect of higher costs on given output, which is just $-x$, canceling the first term, and an "indirect" effect on the firm's profits that arises through rivals' response to learning that the firm has higher costs. The sum of the first two terms is thus just that indirect effect. With strategic substitutes, including Cournot oligopoly, the indirect effect is negative: when rivals learn a firm has higher costs, they expect it to produce less output; as a result, rivals raise their own output, which reduces firm 1's profits. Therefore $G_r(v;1) < 0$ and $r(1) < v$.

PROOF OF THEOREM 7:

For $\theta \leq \theta_V$ and $\theta < [w(0) - w(v)]/[w(0) - w(v) + vx(v,v)]$, $K(\theta) < 0$. For $\theta \leq \theta_V$, if $K(\theta) > 0$, then $P(\theta)/K(\theta)$ strictly decreases with θ.

When $r = v$, $P(\theta) = vx(v,v) - (1-\theta)[\pi(0,0) - \pi(v,v)]$ and $K(\theta) = \theta vx(v,v) - (1-\theta)[w(0) - w(v)]$. Rearranging this last equation shows that $K(\theta) < 0$ for $\theta < [w(0) - w(v)]/[w(0) - w(v) + vx(v,v)]$.

Both $P(\theta)$ and $K(\theta)$ are increasing and linear in θ. Their difference is $P(\theta) - K(\theta) = (1 - \theta)[vx(v,v) + (w(0) - w(v)) - (\pi(0,0) - \pi(v,v))]$, which is zero at $\theta = 1$. (These linear functions do not apply for $\theta > \theta_V$; we are using this fact only to demonstrate the properties of $P(\theta)/K(\theta)$ in the range $\theta \leq \theta_V$.) Therefore, showing that $P(0) > K(0)$ is sufficient to conclude that $P(\theta)/K(\theta)$ strictly decreases with θ. $P(0) > K(0)$ if $vx(v,v) + (w(0) - w(v)) - (\pi(0,0) - \pi(v,v)) > 0$. Writing $w(r) = T(r) + S(r)$, where $S(r)$ is the consumer surplus when royalties are r, this expression is equivalent to $vx(v,v) + T(0) + S(0) - T(v) - S(v) - \pi(0,0) + \pi(v,v) > 0$. Simplifying, this becomes $S(0) > S(v)$, which holds.

Licensing versus Litigation without Negative Fixed Fees

The patent holder strictly prefers licensing to litigation if $s(\theta)x(s(\theta), s(\theta)) > \theta P(1)$. If $s(\theta) = k\theta v$, this becomes $kvx(s,s) > rx(r,r) + [\pi(r,r) - \pi(v,r)]$, where $r = r(1)$. Since $vx(s,s) \geq vx(v,v) \geq rx(r,r)$, this condition is satisfied if $(k-1)vx(s,s) > \pi(r,r) - \pi(v,r)$. If $r(1) = v$ then this becomes $(k-1)vx(s,s) > 0$, which is satisfied for all $k > 1$.

If $r(1) < v$, we can use the intermediate value theorem to write $\pi(r,r) - \pi(v,r) = (r - v)\pi_1(t,r) = (v - r)|\pi_1(t,r)|$ for some $t \in [r(1), v]$. Substituting, the sufficient condition becomes $(k-1)vx(s,s) > (v-r)|\pi_1(t,r)|$.

Now $\pi_1(t,r) = -x(t,r) + IE$, where IE is the indirect effect of the higher costs on the firm's profits that arises because the firm's rivals adjust their behavior. With strategic complements, including Bertrand oligopoly, the indirect effect is positive, so $|\pi_1(t,r)| < x(t,r)$. In this case, the sufficient condition for the patent holder to prefer licensing is satisfied if $(k-1)vx(s,s) > (v-r)x(t,r)$. If $s \leq r(1)$, then $x(s,s) > x(r,r) \geq x(t,r)$, so the patent holder prefers licensing to litigation for $k \geq 2$. This condition is sufficient but far from necessary.

With strategic substitutes, including Cournot oligopoly, the indirect effect is negative. With linear demand and constant marginal costs, $|\pi_1(t,r)| = x(t,r)[2N/(N+1)] < 2x(t,r)$, so we get the sufficient condition $(k-1)vx(s,s) > 2(v-r)x(t,r)$. If $s \leq r(1)$, this condition is satisfied for $k \geq 3$. Again, this condition is sufficient but far from necessary.

PROOF OF THEOREM 8:

For small v, $s(\theta) \approx \theta v[\rho/(1+(\rho-1)\theta)]$.

Recall that $s(\theta)$ satisfies $\pi(s,s) = \theta\pi(v,s) + (1-\theta)\pi(0,0)$. By the intermediate value theorem, there exist $a,b \in (0,1)$ such that the left-hand side is equal to $\pi(0,0) + s[\pi_1(as,as) + \pi_2(as,as)]$ and the right-hand side is equal to $\pi(0,0) + v\pi_1(bv,bs) + s\pi_2(bv,bs)$, where the subscripts denote partial derivatives. Therefore, $s[\pi_1(as,as) + \pi_2(as,as) - \theta\pi_2(bv,bs)] = \theta v\pi_1(bv,bs)$, or

$$\frac{s}{\theta v} = \frac{\pi_1(bv,bs)}{\pi_1(as,as) + \pi_2(as,as) - \theta\pi_2(bv,bs)}.$$

Since $0 \leq s \leq v$, for small v, one can approximate s by $s/\theta v \approx \pi_1(0,0)/[\pi_1(0,0) + (1-\theta)\pi_2(0,0)]$. Using the definition of ρ, this is equivalent to $s(\theta) \approx \theta v[\rho/(1+(\rho-1)\theta)]$.

The result is approximate because we substituted (0, 0) for the varying arguments in the partial derivatives of π. Because we are concerned with a ratio, we need to bound the *proportional* error introduced by that substitution. Technically this requires that $\pi_1(0,0)$ and $\pi_1(0,0) + \pi_2(0,0)$ are nonzero (otherwise, continuity would bound only the *absolute* error in numerator or denominator, leaving open the possibility of large proportional errors). Since we have assumed that (see Section IB), the theorem indeed holds as a limiting statement for small enough v. But how is it likely to fare for moderate but not infinitesimal v? In the course of calculating ρ in Cournot oligopoly next, we show that the partial derivatives of π vary with output. At least in simple cases, this implies that the proportional error introduced by substituting for the varying arguments is bounded by the proportional difference in output as (a,b) varies over $[0,v] \times [0,v]$. In those cases, and (we suggest) plausibly in general in moderately competitive markets with moderate v, that error factor will not be large compared to the ratio by which the approximation exceeds the benchmark θv.

Relativity Ratio in Cournot Oligopoly: Comparison of ρ and N

With constant marginal costs and Cournot oligopoly, the first-order condition for firm i output choice is $p(X) + x_i p'(X) - c_i = 0$. Totally differentiating this, we get $[p'(X) + x_i p''(X)]dX + p'(X)dx_i - dc_i = 0$. Following the notation from Farrell and Shapiro (1990), we define $\lambda_i \equiv [-p'(X) - x_i p''(X)]/(-p'(X))$, so with $dc_i = dr_i$ we have $dx_i = -\lambda_i dX + dr_i/p'(X)$. Writing $\Lambda \equiv \sum_i \lambda_i$ and adding up across all firms gives $dX/dr_1 = 1/[1+\Lambda]p'(X)$. Substituting for dX using this expression, we get $dx_1/dr_1 = (1+\Lambda-\lambda_1)/[1+\Lambda]p'(X)$ and $dx_j/dr_1 = -\lambda_j/[1+\Lambda]p'(X)$, $j \neq 1$.

For each firm $j \neq 1$, by the envelope theorem, the profit impact of a small increase in firm 1's running royalty is given by firm j's equilibrium output x_j times the change in price resulting from the equilibrium change in output by all *other* firms, $dX - dx_j$. This price change is given by $p'(X)[dX - dx_j]$, which equals $[(1 + \lambda_j)/(1 + \Lambda)]\ dr_1$. Since this expression does not contain any parameters specific to firm 1, the effect on firm j's profits of a small increase dr in *all other firms'* running royalties is given by $(N - 1)x_j[(1 + \lambda_j)/(1 + \Lambda)]\,dr$. Returning to our main notation, we therefore have $\pi_2 = (N - 1)x_j[(1 + \lambda_j)/(1 + \Lambda)]$.

Similarly, the effect on firm 1's profits of a small increase dr_1 in its own running royalty is equal to the direct cost effect, $-x_1 dr_1$, plus the effect of the price change caused by other firms' output changes, $x_1 p'(X)[dX - dx_1] = -(\Lambda - \lambda_1)/(1 + \Lambda)dr_1$. Therefore $|\pi_1| = x_1(1 + 2\Lambda - \lambda_1)/(1 + \Lambda)$.

Putting these together, starting at a symmetric equilibrium where each $\lambda_i = \lambda$ and $x_i = x_j$, and simplifying, we get

$$\frac{|\pi_1|}{|\pi_1 + \pi_2|} = \frac{1 + (2N - 1)\lambda}{2 + N\lambda - N} = N\left[1 + \frac{N(1 - \lambda)}{2 - N(1 - \lambda)}\right].$$

In a symmetric equilibrium, we also have $\lambda_i \equiv [-p'(X) - Xp''(X)/N]/(-p'(X)) = 1 + Xp''(X)/Np'(X)$. Writing $E \equiv -Xp''(X)/p'(X)$ for the elasticity of the slope of the inverse demand curve, we have $\lambda = 1 - E/N$ or $E = N(1 - \lambda)$. Hence, we obtain $|\pi_1|/|\pi_1 + \pi_2| = N(2/(2 - E))$, or, equivalently,

$$\frac{|\pi_1|}{|\pi_1 + \pi_2|} = N\left[1 + \frac{2}{2 + Xp''(X)/p'(X)}\right].$$

Note that if demand is linear or convex, $p''(X) \geq 0$, then $E \geq 0$ and $|\pi_1|/|\pi_1 + \pi_2| \geq N$. For linear demand, $E = 0$, so $|\pi_1|/|\pi_1 + \pi_2| = N$. When demand has constant elasticity $\varepsilon > 1$ ($\varepsilon > 1$ is the regularity condition for $\pi_1 + \pi_2 < 0$), we have $E = 1 + 1/\varepsilon$, so $|\pi_1|/|\pi_1 + \pi_2| > N$.

Vertically Integrated Patent Holder

Define $\pi^I(a,b)$ as the profits of a downstream firm with cost a, given that the other downstream firms have cost b and the patent holder competes using the patented technology. This downstream firm's output is $x^I(a,b)$. We assume that $\pi^I(a,b)$ satisfies the three assumptions that Section IB assumed for $\pi(a,b)$.

Write $\phi(r)$ for P's product market profits if the rivals all pay royalty r. We make the very mild assumption that $\phi'(r) \geq 0$; P earns no less profits from its downstream operations, the higher are the royalties paid by other downstream firms.

Define $T_I(r) = \phi(r) + Nrx^I(r,r) + N\pi^I(r,r)$ as the joint profits of P and the downstream firms if all downstream firms pay royalties r. We assume that $T_I(r)$ is increasing with r in the range $0 \leq r \leq v$, now even if $N = 1$. With these definitions, the analysis proceeds just as in the nonintegrated case, using $T_I(r)$ rather than $T(r)$ and $\pi^I(r,r)$ rather than $\pi(r,r)$. So $r_I(\theta)$ maximizes $G^I(r,\theta) = T_I(r) - \theta N\pi^I(v,r) - (1 - \theta)N\pi^I(0,0)$ subject to $r \leq v$.

If $N = 1$, then $G^I(r,\theta) = T_I(r) - \theta\pi^I(v) - (1 - \theta)\pi^I(0)$, which increases with r in $r \leq v$, so $r_I(\theta) = v$ for all θ. A similar logic applies if downstream firms using the backstop technology cannot profitably compete against the patent holder: $\pi^I(v,v) = 0$. This condition implies that $\pi^I(v,r) = 0$ for all $r \leq v$, so $G^I(r,\theta) = T_I(r) - (1 - \theta)N\pi^I(0,0)$, and thus $r_I(\theta) = v$ for all θ. More generally, we have $r(\theta) = v$ for all $\theta \leq \theta_{VI}$, where $\theta_{VI} \equiv T_I{}'(v)/(N\pi_2^I(v,v))$. As in the nonintegrated case, $B(\theta) > 0$ for all θ, and $K(\theta) < 0$ for weak patents.

The analysis without negative fixed fees also closely parallels the case of the nonintegrated patent holder. P still sets the highest acceptable royalty rate, all the more so if $\phi'(r) > 0$. The same acceptance condition applies, using $\pi^I(a,b)$ instead of $\pi(a,b)$. Therefore, for small values of θ, we get $s(\theta) \approx \rho^I \theta v$, where $\rho^I = |\pi_1^I(0,0)|/|\pi_1^I(0,0) + \pi_2^I(0,0)|$. With a symmetric Cournot oligopoly and linear demand, the Supplemental Materials show that $\rho^I = (N+1)/2$. For small v we again have $B(\theta) > 0$ for all θ.

Short-Term versus Long-Term Licenses

In equilibrium, there is no litigation, so the only impact of using short-term rather than long-term licenses is on the payoff to a downstream firm that infringes rather than accepts a license. In the analysis above with unrestricted two-part tariffs, for a given θ, if the equilibrium running royalty rate is $r(\theta)$ with long-term licenses, this reservation payoff was $\theta\pi(v, r(\theta)) + (1-\theta)\pi(0,0)$. With short-term licenses, this reservation payoff becomes $\theta\pi(v, r(1)) + (1-\theta)\pi(0,0)$. The patent holder has an incentive to use short-term licenses if and only if $r(1) < r(\theta)$. If $r(1) = v$, then $r(\theta) = v$ for all θ and the patent holder is indifferent between using short-term and long-term licenses. However, if $r(1) < v$, then $r(1) < r(\theta)$ for all $\theta < 1$ and the patent holder strictly prefers to use short-term licenses. Define θ^*_{ST} such that $\pi(v,v) = \theta\pi(v, r(1)) + (1-\theta)\pi(0,0)$. For $\theta > \theta^*_{ST}$, the downstream firm's threat point is to use the backstop technology rather than infringe and engage in litigation. For all θ, the downstream firm's reservation payoff is independent of r, so the patent holder has no incentive to reduce r below v. Therefore, $r = v$ for all patent strengths. Negative fixed fees are used for all $\theta < \theta^*_{ST}$; no fixed fee is used for $\theta \geq \theta^*_{ST}$. The ex post welfare analysis is exactly the same as the case already studied in which $\theta_V = 1$. Ex ante, for all $\theta < 1$, $P(\theta)$ is higher than we had earlier and $K(\theta)$ is unchanged or lower, so the bias resulting from $P(\theta)/K(\theta) > P(1)/K(1)$ is stronger than we had earlier.

If negative fixed fees are not feasible, with long-term licenses the downstream firm's payoff from infringing was $\theta\pi(v, s(\theta)) + (1-\theta)\pi(0,0)$. With short-term licensees, for $\theta < \theta^*_{ST}$ this payoff becomes $\theta\pi(v, r(1)) + (1-\theta)\pi(0,0)$, so the patent holder has an incentive to use short-term licenses if and only if $s(\theta) > r(1)$. With $r(1) > 0$, this condition will not be met for relatively weak patents, so owners of those patents will use long-term licenses. For stronger patents, the patent holder has an incentive to use short-term licenses. For $\theta \geq \theta^*_{ST}$, the downstream firm's threat point is to use the backstop technology, so $s(\theta) = v$. Since the ability strategically to use short-term licenses raises $s(\theta)$, our welfare results are strengthened.

Consumer and Downstream Firm Risk Preferences on Linear Royalties

Each downstream firm's payoff from s is $\pi(s) = [p(x(s)) - s]x(s) = \max_x[[p(x) - s]x]$. For each x, $[p(x) - s]x$ is linear and decreasing in s, so the upper envelope $\pi(s)$ is convex and decreasing in s. Thus (as is well known), the downstream firm prefers lower s but is risk-loving in s. Since $\pi'(s) = -x(s)$ and $\pi''(s) = -x'(s)$, the downstream firm's risk preference in s is measured by the coefficient of "absolute risk aversion" $\pi''(s)/(-\pi'(s)) = (-x'(s))/x(s)$. (Because $\pi(s)$ is decreasing, this standard "risk aversion coefficient" is mathematically positive even though the downstream firm is risk-preferring.) Turning to consumers, write $V(s)$ for consumer surplus, and $p(s)$ for downstream price, as functions of s. Then $V'(s) = -p'(s)x(s)$. If pass-through is linear, $p''(s) = 0$, then $p'(s)$ is a positive constant, so $V'(s)$ is a preference-preserving transformation of $\pi'(s)$, so consumers have the same risk attitudes as the downstream firm. In more detail, we have

$$\frac{V''(s)}{-V'(s)} = \frac{-x'(s)}{x(s)} - \frac{p''(s)}{p'(s)},$$

so, when $p''(s) = 0$, consumers' risk preference coefficient equals $V''(s)/(-V'(s)) = (-x'(s))/x(s)$, which (we just saw) is also the risk preference coefficient for the downstream firm. If price is convex in s, then consumers are less risk-loving in s than the downstream firm; if price is concave, they are more risk-loving.

Benefits of Ideal PTO Review with Linear Licenses

By Theorem 1, $B(\theta) \geq [r(\theta) - \mu\theta r(1)]\lambda_{\min}$. With linear licenses, $r(1) = v$ so $B(\theta) \geq \theta v [r(\theta)/\theta v - \mu]\lambda_{\min}$ and $K(1) = vx(v,v)$. In general $B(\theta) + K(\theta) = \theta K(1)$; with $r(1) = v$ this gives $B(\theta) + K(\theta) = \theta v x(v,v)$. Therefore, $B(\theta)/(B(\theta) + K(\theta)) \geq [r(\theta)/\theta v - \mu][\lambda_{\min}/x(v,v)]$. As shown in the proof of Theorem 1, with Cournot oligopoly $|w'(r)| = p'(r)[(p(r) - c)/(p(r) - r - c)]x(r,r)$. Therefore, $\lambda_{\min} \geq p'(r)x(v,v)$, and $B(\theta)/(B(\theta) + K(\theta)) \geq [r(\theta)/\theta v - \mu]p'(r)$.

REFERENCES

Anand, Bharat N., and Tarun Khanna. 2000. "The Structure of Licensing Contracts." *Journal of Industrial Economics*, 48(1): 103–35.

Anton, James J., and Dennis A. Yao. 2003. "Patents, Invalidity, and the Strategic Transmission of Enabling Information." *Journal of Economics and Management Strategy*, 12(2): 151–78.

Aoki, Reiko, and Jin-Li Hu. 1999. "Licensing vs. Litigation: The Effect of the Legal System on Incentives to Innovate." *Journal of Economics and Management Strategy*, 8(1): 133–60.

Bessen, James, and Michael Meurer. 2008. *Patent Failure: How Judges, Bureaucrats, and Lawyers Put Innovators at Risk*. Princeton, NJ: Princeton University Press,

Bhattacharyya, Sugato, and Francine Lafontaine. 1995. "Double-Sided Moral Hazard and the Nature of Share Contracts." *RAND Journal of Economics*, 26(4): 761–81.

Bulow, Jeremy. "The Gaming of Pharmaceutical Patents." 2004. In *Innovation Policy and the Economy*. Vol 4, ed. Adam B. Jaffe, Scott Stern, and Josh Lerner, 145–87. Cambridge, MA: MIT Press.

Chiao, Benjamin, Josh Lerner, and Jean Tirole. 2007. "The Rules of Standard-Setting Organizations: An Empirical Analysis." *RAND Journal of Economics*, 38(4): 905–30.

Choi, Jay Pil. 1998. "Patent Litigation as an Information-Transmission Mechanism." *American Economic Review*, 88(5): 1249–63.

Choi, Jay Pil. 2002. "Patent Pools and Cross Licensing in the Shadow of Patent Litigation." http://www.msu.edu/%7Echoijay/Patent%20Pools-Choi.pdf.

Choi, Jay Pil. 2005. "Live and Let Live: A Tale of Weak Patents." *Journal of the European Economic Association*, 3(2–3): 724–33.

Farrell, Joseph, and Robert Merges. 2004. "Incentive to Challenge and Defend Patents: Why Litigation Won't Reliably Fix Patent Office Errors and Why Administrative Patent Review Might Help." *Berkeley Technology Law Journal,* 19(2): 943–70.

Farrell, Joseph, and Carl Shapiro. 1990. "Asset Ownership and Market Structure in Oligopoly." *RAND Journal of Economics*, 21(2): 275–92.

Federal Trade Commission. 2002. *Generic Drug Entry Prior to Patent Expiration*. http://www.ftc.gov/os/2002/07/genericdrugstudy.pdf.

Federal Trade Commission. 2003. *To Promote Innovation: The Proper Balance Between Competition and Patent Law and Policy.* http://www.ftc.gov/os/2003/10/innovationrpt.pdf.

Federal Trade Commission. 2005. *Agreements Filed with the Federal Trade Commission under the Medicare Prescription Drug, Improvement, and Modernization Act of 2003*. http://www.ftc.gov/os/2006/04/fy2005drugsettlementsrpt.pdf.

Hovenkamp, Herbert, Mark D. Janis, and Mark A. Lemley. 2003. "Anticompetitive Settlement of Intellectual Property Disputes." *Minnesota Law Review,* 87(6): 1719–66.

Jaffe, Adam B., and Josh Lerner. 2004. *Innovation and Its Discontents: How Our Broken Patent System Is Endangering Innovation and Progress, and What to Do About It*. Princeton, NJ: Princeton University Press.

Kamien, Morton. 1992. "Patent Licensing." In *Handbook of Game Theory and Economic Applications,* ed. Robert J. Aumann and Sergiu Hart, 331–54. Amsterdam: North-Holland.

Katz, Michael L., and Carl Shapiro. 1986. "How to License Intangible Property." *Quarterly Journal of Economics*, 101(3): 567–89.

Kesan, Jay P., and Gwendolyn G. Ball. 2006. "How Are Patent Cases Resolved? An Empirical Analysis of the Adjudication and Settlement of Patent Disputes." *Washington University Law Review*, 84: 237–312.

Lemley, Mark. 2001. “Rational Ignorance at the Patent Office.” *Northwestern University Law Review*, 95(4): 1497–1532.

Lemley, Mark A., and Carl Shapiro. 2005. “Probabilistic Patents.” *Journal of Economic Perspectives*, 19(2): 75–98.

Liao, Chun-Hsiung, and Debapriya Sen. 2005. “Subsidy in Licensing: Optimality and Welfare Implications.” *Manchester School*, 73(3): 281–99.

McAfee, R. Preston, and Marius Schwartz. 1994. “Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity, and Uniformity.” *American Economic Review*, 84(1): 210–30.

Marco, Alan. 2006. “Learning by Suing: Structural Estimates of Court Errors in Patent Litigation.” http://ssrn.com/abstract=913408.

Meurer, Michael J. 1989. “The Settlement of Patent Litigation.” *RAND Journal of Economics*, 20(1): 77–91.

Miller, Joseph Scott. 2004. “Building a Better Bounty: Litigation-Stage Rewards for Defeating Patents.” *Berkeley Technology Law Journal,* 19(2): 667–739.

Moore, Kimberly A. 2000. “Judges, Juries, and Patent Cases: An Empirical Peek inside the Black Box.” *Michigan Law Review*, 99(2): 365–409.

National Academies of Science. 2004. *A Patent System for the 21st Century*. http://www.nap.edu/books/0309089107/html.

Rey, Patrick, and Thibaud Vergé. 2004. “Bilateral Control with Vertical Contracts.” *RAND Journal of Economics*, 35(4): 728–46.

Segal, Ilya. 1999. “Contracting with Externalities.” *Quarterly Journal of Economics*, 114(2): 337–88.

Sen, Debapriya, and Yair Tauman. 2007. “General Licensing Schemes for a Cost-Reducing Innovation.” *Games and Economic Behavior*, 59(1): 163–86.

Shapiro, Carl. 2003. “Antitrust Limits to Patent Settlements.” *RAND Journal of Economics*, 34(2): 391–411.

Shapiro, Carl. 2006. “Injunctions, Hold-Up, and Patent Royalties.” http://faculty.haas.berkeley.edu/shapiro/royalties.pdf.

US Patent and Trademark Office. 2006. *Performance and Accountability Report, Fiscal Year 2006*. http://www.uspto.gov/web/offices/com/annual/2006/2006annualreport.pdf.

Waterson, Michael. 1990. “The Economics of Product Patents.” *American Economic Review*, 80(4): 860–69.

Willig, Robert D., and John P. Bigelow. 2004. “Antitrust Policy toward Agreements That Settle Patent Litigation.” *Antitrust Bulletin*, 49(3): 655–98.