



Bundling of RAND-committed patents



Anne Layne-Farrar^a, Michael A. Salinger^{b,*}

^a Charles River Associates, One South Wacker Drive, Chicago, IL 60606, USA

^b Boston University, Questrom School of Business, 595 Commonwealth Ave., Boston, MA 02215, USA

ARTICLE INFO

Article history:

Received 31 July 2015

Received in revised form 25 February 2016

Accepted 2 March 2016

JEL classification:

O3

K21

Keywords:

Patents

Licensing

Bundling

Tying

RAND

FRAND

Hold-up

Intellectual property

ABSTRACT

Gilbert and Katz (2006) (GK) show that allowing (pure) patent bundling increases the incentives for patent owners to enter into “long-term” patent licensing that commits them not to expropriate licensees’ sunk costs in complementary assets with opportunistic licensing terms. We interpret RAND commitments as a form of long-term contracting, and extend their framework to analyze the tying of non-RAND-committed patents to RAND-committed patents. Pure patent bundling/tying is common and often has sound efficiency justifications, so we caution against prohibiting the pure bundling of RAND-committed and non-RAND-committed patents. Whether such a license honors a RAND commitment turns, however, on the licensing terms. We argue that including a non-RAND-committed patent (patent 2) in a bundle with a RAND-committed patent (patent 1) does not increase the license fee that honors the RAND commitment. If, however, the patent owner offers patent 1 separately at a RAND rate, its RAND commitment does not restrict what it charges for a bundle of patent 1 and patent 2.

© 2016 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

In this paper, we address the relationship between two common practices in patent licensing: RAND commitments and patent bundling. A RAND commitment is a commitment to license technology on “Reasonable and Non-Discriminatory” terms¹. Patent bundling (or, more precisely, pure patent bundling) is licensing patents only in bundles (or portfolios) rather than offering licenses to individual patents on an à la carte basis. We ask whether, once a patent owner makes a general commitment to license its patents on RAND terms, it is permissible to offer those patents solely in a bundle with other patents. If so, what does the RAND commitment imply for the royalties charged for the bundle? Alternatively, does a RAND commitment necessarily oblige a patent holder to license its patent on a stand-alone basis?

To address the general issue of bundling RAND-committed patents, we must understand why patent bundling is such a

common phenomenon, what effect patent bundling has on licensing terms, and why patent holders make RAND commitments. A substantial and growing economics literature has emerged on bundling in general, and some of it specifically focuses on bundling intellectual property (like computer software, music, and video entertainment), but the formal literature on patent bundling is remarkably thin². A notable exception is the *Gilbert and Katz (2006)* (henceforth “GK”) model of patent bundling.

The GK model is a natural starting point for our analysis. A key idea in the GK model is the distinction between “long-term” and “short-term” patent licensing, which turns on whether the parties enter into contracts before or after the patent user incurs sunk costs. The patent owner can impose terms that expropriate the patent user’s sunk costs with short-term contracts, but not with long-term licensing.

GK do not address how long-term licensing occurs in practice. One possibility, of course, is that the patent user seeks a license before starting to commercialize the technology. An alternative

* Corresponding author. Tel.: +1 617 353 4408; fax: +1 617 353 6667.

E-mail addresses: alayne-farrar@crai.com (A. Layne-Farrar), salinger@bu.edu (M.A. Salinger).

¹ As we understand the nomenclature, RAND is synonymous with “FRAND,” which is an acronym for “Fair, Reasonable, And Non-Discriminatory.” Licensing terms that are both “reasonable” and “non-discriminatory” are necessarily “fair.”

² Patent bundling is related to but different from patent pooling, which *Shapiro (2001)* and *Lerner and Tirole (2004)* have analyzed. Patent pooling entails licensing patents of different patent owners in a single package. Patent bundling refers to licensing multiple patents of a single patent-owner as a package.

mechanism, however, is to make a RAND commitment. As a starting point³, we interpret a RAND commitment to mean a commitment on the part of the patent owner not to charge higher licensing fees or impose other terms that are more onerous than it would have sought to impose before the patent user made irreversible decisions (such as incurring sunk costs in complementary assets) to use the patent⁴. Under this interpretation, the GK model provides a general framework for analyzing RAND commitments, and the patent license bundling in their model is of two RAND-committed patents for technologies that are both needed to implement a single standard.

GK's results are that the patent owner would offer its two complementary patents as a bundle rather than à la carte. Thus, the pure bundling of the two patents does not violate the patent-owner's commitment to limit the license fee to what it would have sought before the patent user incurred any sunk costs. Moreover, they show that banning bundled patent licensing (and thereby forcing the patent owner to license its patents separately) might prevent the patent holder from being willing to enter into long-term licenses or, under our interpretation, to make RAND commitments. As commitments to avoid expropriation can be necessary to induce others to invest in complementary assets, GK urge caution about restrictions on patent bundling.

In this article, we adapt and extend the GK framework to address a different question about patent license bundling. Our focus is not the pure bundling of multiple RAND-committed patent licenses but, rather, the pure bundling of two patent licenses, one of which is RAND-encumbered and one of which is not. Doing so requires analyzing the meaning of a RAND commitment on one patent when a patent owner has not made a similar commitment on another patent that a patent user might need to use with the RAND-committed patent. Note that, while most commonly made in the context of standard setting within cooperative standard development organizations (SDOs), RAND pledges are not limited to SDOs, in large part because SDOs are not the only way that standards emerge⁵.

We analyze expropriation through tying as the outcome to what we call the "Some Opportunism Game" in which the owner of two patents can behave opportunistically in the licensing of one of its patents but not the other. As a foundation for that game, we analyze a "No Opportunism Game" that has the basic structure of the GK model but with only a single patent to license. The result of this model brings out a key distinction between our model and the GK model related to the meaning of a RAND commitment. In particular, we show that there is no universal formula for RAND. Instead, for some parameters, the RAND royalty is "value-based," meaning that it reflects the *ex ante* upper bound value created by the patented technology. For other parameters, the RAND royalty is "cost-based," meaning one that is low enough to eliminate a patent user's incentive to try to invent around the patent.

³ This interpretation is consistent with Swanson and Baumol (2005). As we show in Section 4, it has more complicated implications than one might expect, and these complications might cause some to argue for a refinement on this interpretation of RAND.

⁴ In addition to sunk investments in complementary assets, irreversible decisions include the decision not to try to invent around the patent and, as we discuss in Section 5, the decision to adopt one technology instead of an available alternative. In what follows, when we refer to a time before the patent user incurs sunk costs in complementary assets, we are referring more generally to a time before the patent user makes any irreversible decision with respect to its use of the patent.

⁵ The Program for Information Justice and Intellectual Property at Washington College of Law maintains a database of more than 150 such public non-SDO patent pledges: <http://www.pjip.org/non-sdo-patent-commitments/>. Contreras (2015) and Elhaug (2015) discuss the legal basis for enforcing RAND commitments made outside an SDO setting. For other discussions of non-SDO patent pledges, see Layne-Farrar (2014) and Harkrider (2013).

The remainder of this paper is organized as follows. Section 2 discusses the economics of bundling in general, highlighting the aspects that we consider essential for understanding patent bundling. Section 3 then explains the GK model and the role of bundling in it. Section 4 presents the No Opportunism Game, bringing out the distinction between value-based and cost-based RAND royalties. Section 5 presents a modified version of the No Opportunism Game to show that our results apply in a wider set of circumstances than might initially appear to be the case⁶. In Section 6, which contains our main results, we use the GK framework to analyze licensing two patents where the patent holder has made a RAND commitment on just one.

Section 7 contains our conclusions. Briefly, the pure bundling of patent licenses raises a basic policy dilemma. On the one hand, the practice could be a way to renege on the RAND commitment. On the other hand, as we explain in Section 2, the pure bundling of patents is prevalent because it can be an efficient form of contracting that minimizes transaction and litigation costs. We argue that because of these efficiencies, the owner of multiple patents should be allowed to engage in pure bundling of patents with and without RAND commitments, but that RAND commitments limit the license fees that it can charge. Indeed, we show that the license fees that honor a RAND commitment on a bundle of RAND-committed and non RAND-committed licenses may be less than the RAND license fee if the patent owner had made RAND commitments on both patents.

2. Bundling and tying in general

While a complete review of the extensive economics literature on bundling and tying is unnecessary for what follows, a few key points are essential.

A key concept in the literature on bundling is the "single monopoly profit"⁷ (or, more generally, "single rent"⁸) theorem, and this theorem plays an important role in the GK results. As is well understood, if a company has a monopoly over two goods—call them "A" and "B"—that consumers use together in fixed proportions, there is a combined price for the two goods that maximizes the company's profits. Any combination of prices that sums to this optimal bundled price will generate the same level of profits (and consumer surplus). Indeed, the company can earn this level of profits on just one of the two goods as long as the price of the other good equals marginal cost.

The single rent principle was one of the key underpinnings of the so-called Chicago critique of a wide range of antitrust policies, including tying. Arguably the dominant interpretation of the

⁶ In the GK model and in the model in Section 4, the opportunity for opportunistic licensing arises because the patent user has to invest in sunk costs to earn any commercial value from the technology. Another form of opportunism that is sometimes alleged (particularly in formal standard setting proceedings) arises when there are multiple patents that compete for inclusion in a standard (with the one chosen being essential). Prior to adoption of the standard, the licensing terms a patent owner could negotiate are constrained by the available alternatives. Once a patent becomes essential, however, the alternatives cease to be a constraint and the patent owner has an incentive to extract the entire value of the standard rather than its contribution to the value over the value that could have been realized with an alternative patent. The model in Section 5 shows that we can interpret the opportunity cost of not adopting an alternative patent as a sunk cost that has the same economic effect as the sunk costs in the GK and Section 4 models.

⁷ As evidenced by the title to Elhaug (2009) the term is widely recognized. However, we have not been able to document the source of the term. Whinston (1990) attributes the arguments to a Chicago oral tradition. Bowman (1957) recognizes the strong assumptions underlying the principle and a set of exceptions to it when those assumptions do not apply.

⁸ While economists (as well as lawyers and courts) use the phrase "single monopoly profit," we prefer the term "single rent" because the argument applies to rents of any kind, including patent royalties.

principle is that it casts doubt on claims that the incentive to tie goods together is “leveraging” or “foreclosure.” An alternative interpretation, however, is that the assumptions underlying the single rent principle are so strong that relaxing them could undo the principle.

Understanding the assumptions provides a framework for understanding why a single source of rents may not be sufficient to capture all the rents available. Farrell and Weiser (2003) and Elhauge (2009) have articulated this perspective at length, as did Krattenmaker and Salop (1986). One of the stronger assumptions underlying the single rent principle presumes that the monopoly or property right on at least one of the goods is iron-clad. The threat of entry into both goods can violate this assumption, though a threat of entry into just one of the goods does not. Suppose entry into A is not possible but entry into B is. According to the single rent principle, the seller should welcome entry into B by a company offering a lower-cost, higher quality, or differentiated version since the improvement in B would increase the profits the seller could earn from A. But, the threat of entry in A as well can make it impossible for the firm to raise the price of A to take advantage of cost reductions or improvements in B. With the potential for entry into both A and B, tying can raise entry barriers by forcing two-stage entry. This exception to the single rent principle underlies the analyses of Choi and Stefanadis (2001) and Nalebuff (2004).

Failure of the single rent principle does not, however, necessarily imply that modularized competition results in better outcomes. Because of double marginalization, the profits a company can earn by controlling the rights to both A and B exceed the total profits that separate owners of A and B would earn. Eliminating double marginalization is the reason that patent pools can be in the public interest. As Lerner and Tirole (2004) show, whether patents are complements or substitutes is more subtle than in the case of consumer goods⁹. However, to the extent that bundles of separately-owned complementary patents can be combined in the public interest, one must be cautious about condemning the tying of commonly-owned patents¹⁰.

While the formal economics literature has focused on price discrimination and foreclosure as possible motives for tying, the practice is far too common for the specialized models in the price discrimination or foreclosure literature to explain them all. Instead, transactional and organizational efficiency is likely to explain much more of the tying that occurs in practice. Virtually no one wants every section of a newspaper. Evans and Salinger (2005, 2008) argue that one cannot understand the diverse instances of tying without

recognizing the cost of product offering complexity. Before a firm decides on how much to produce (or try to sell) and how much to charge, it first has to decide exactly what it sells. Even after deciding on its general line of business, the products a firm offers are typically a small subset of the products and product combinations it could conceivably offer. Dell revolutionized the personal computer industry by putting in place systems to customize orders to consumer specifications to a degree that had previously been unimaginable. But even Dell's highly customized offerings did not include every conceivable configuration and, more importantly, Dell and the personal computer industry are the exception. In general, companies do not and indeed cannot customize their offerings to the precise desires of every (or, for that matter, any) customer.

While Evans and Salinger did not address the issue of tying patents, the framework they suggest provides a plausible—indeed, obvious—explanation for why patent tying is such a common phenomenon. The key question to ask in regard to à la carte licensing is, “What is the limiting principle?” Some large innovative companies that license their technology have thousands of patents. Yet, they might offer them in only a handful of bundles, and they might not offer any of them individually. For a company with 1000 patents, the number of possible licensing combinations of patents is on the order of 10^{301} , which is about googol cubed! The notion that the licensor is obligated to unbundle any arbitrary set and offer a discount means that it would have to set 10^{301} different prices and then monitor and enforce compliance for all those different configurations. As a practical matter, the number of different combinations that licensees might demand would likely be much smaller, but there is nonetheless a cost of having more complex product offerings. As a result, patent licensors necessarily offer a small subset of the patent bundles that they could conceivably offer; and it should come as no surprise (and should not necessarily be a public policy concern) if many licensees only use a subset of the patents they license.

It also should not be surprising if licensees want inclusive bundles. A licensee takes a license to (1) access useful technology and (2) avoid being sued for patent infringement. A company that licenses just a subset of the patents that it needs to implement a technology risks a patent-infringement suit even if it pays the royalties due on the patents it does license. As a result, licensees might demand patents in an inclusive bundle that completely protects them from the threat of a suit (referred to as “freedom to operate” licenses). To the extent that such bundles cover both RAND-committed patents and non-RAND committed patents, including all patents that a licensee might conceivably use is itself a form of commitment on the part of a patent owner not to behave opportunistically by suing for patent infringement on a non-RAND committed patent that the licensee ends up infringing. Both the licensor and the licensee can save transaction and enforcement costs by licensing at the technology portfolio level.

3. The Gilbert-Katz (GK) model

With the motives for patent bundling in mind, we turn now to patent hold-up. The Gilbert-Katz (GK) model addresses this topic in the context of patent bundling. The possibility of patent hold-up adds a layer of complexity to the implications of the single-rent theorem because there are two levels of rents to consider: the *ex ante* rents before licensees engage in sunk investment (including R&D efforts to invent around the patents) and the *ex post* rents after they do¹¹.

⁹ Observe that Lerner and Tirole (2004) make a different, though related, point to the one we focus on. Specifically, Proposition 5 in Lerner and Tirole (2004) establishes that a welfare enhancing patent pool contains complementary patents and is strongly stable, such that the presence of individual patent licensing will not unravel the pool, whereas a welfare decreasing pool contains patents that are substitutes for one another, in which case the presence of individual patent licensing will unravel the pool, as licensees can obtain lower aggregate rates through individual licensing of the subset of patents needed. In our model, the bundle (or “pooled” patent license) can be used to achieve a higher rate when the patents are complements, not substitutes, such that the presence of individual licensing (here across patents, rather than across patent holders) would unravel the bundle. The difference between Lerner and Tirole (2004) and our analysis stems from the presence of RAND commitments, something not included in Lerner and Tirole. As we argue below, combining RAND-encumbered patents with non-RAND-encumbered patents can create a mechanism to avoid the RAND commitment.

¹⁰ To be sure, participants in patent pools often offer to license their patents separately as well as part of the pool, and whether they do can be a factor in determining whether a pool is in the public interest. See the discussion in Lerner and Tirole (2004). While at least one paper, Quint (2014), develops theory demonstrating that the inclusion of patents that are not perfectly complementary need not be anti-competitive or harmful, the general consensus is that pools should be restricted to essential patents to prevent foreclosure of alternative technologies for optional features.

¹¹ We start our analysis from the point in time that the patent holder has already made its R&D investment and has a patent in hand. Thus, in our analysis here, “*ex ante*” refers to the licensee's investment, not the patent holder's.

In the formal set up of the GK model, the owner of two complementary patents (“the patent owner”) can license to a single licensee (the “manufacturer”). The technologies covered by the patents are strongly complementary. That is, the technologies create value only in combination with each other. However, the patents themselves are not strongly complementary as the manufacturer can invest in R&D to try to invent around one or both patents. The IP owner does not have manufacturing capability, so it must license its technology to the manufacturer to capture any value from it. The value that can be realized from the technology is a function of the level of complementary investment, which only the manufacturer can make. If the manufacturer does not obtain a license to the technology prior to investing in complementary assets, the IP owner might be able to expropriate the contribution of the manufacturer’s investment. In other words, if the manufacturer does not obtain an early license, the patent holder can practice hold-up.

A simple numerical example illustrates the GK model, the intuition behind it, and the role of patent tying in the GK framework. Suppose that without any investment by the manufacturer, the patented technology yields a value of 20 (unrealizable by the patent holder, by assumption). With efficient investment in complementary assets by the manufacturer, the technology embodied in an end product yields gross benefits (*i.e.*, before taking account of the cost of the manufacturer’s investments) of 100 (unrealizable by the manufacturer absent the initial contribution by the patent holder). The commercialization investments needed to generate the additional value (the end product) have a cost of 30, so that the potential net value of the combined technology (patent plus manufacture) embodied in the end product is 70.

Absent any frictions, the manufacturer would not invest 30 in complementary assets without obtaining a license for the patented technology. If it did so, the IP owner could insist on a license fee of 100 (or just below it) and the manufacturer would rationally accept this offer. The IP owner has an incentive to commit up front to a license fee of no more than 70 for the bundle (100–30), which allows the manufacturer to recover its investment (and also to choose the most efficient level of complementary investment)¹². GK refer to licenses entered into before the manufacturer invests as “long-term” licenses and licenses signed after the manufacturer invests as “short-term” licenses. As we have argued, we can interpret their long-term licenses as RAND-committed licenses.

Up to this point, whether the technology is based on one patent or two is irrelevant. The manufacturer needs both. Thus, whether the IP owner offers the patents separately or in a bundle is also irrelevant. It can offer the patents as a bundle at a license fee of 70 or à la carte with individual prices that add to a cumulative license of 70. Profits and consumer welfare (and, therefore, total surplus) are all the same under the various options for charging a total of 70. The irrelevance of bundling under these conditions is an application of the single rent principle. One patent that is essential for a product gives the patent holder the ability to extract the same total license fee as it can with two such patents.

The feature of the GK model that makes bundling potentially relevant is the ability of the manufacturer to invest in R&D to invent around one or both patents. The outcome of the manufacturer’s R&D is random. For a given level of investment, the manufacturer

¹² A simplifying assumption in GK is that the IP owner makes take-it-or-leave-it offers to the manufacturer, which implies that the IP owner is able to capture the entire available surplus. As they explain, the model would be more complicated if the IP owner and manufacturer bargained over the available surplus, but the risk would remain that the IP owner would expropriate part of the manufacturer’s investment if the manufacturer invested prior to obtaining a license.

decides to invent around “patent 1,” “patent 2,” both, or neither¹³. The manufacturer’s incentive to invest in R&D depends on whether it has licensed the patents on a bundled or an à la carte basis. If it licenses the bundle, it only earns a return on its R&D if it succeeds in inventing around both patents. If it licenses on an à la carte basis, then successfully inventing around just one of the patents will lower the license fees it owes the IP owner. Thus, holding the total license fee constant, the manufacturer has more of an incentive to invest in R&D under à la carte licensing than under bundled licensing. The IP owner can limit the manufacturer’s incentive to invest in R&D to invent around its patents by licensing them solely as a bundle. In the GK model, the advantage the IP holder gets from patent bundling is that it discourages attempts to invent around its patents, and that benefit is the IP owner’s incentive to enter into long-term patent licenses¹⁴.

4. RAND commitments: The no opportunism game

In this section, we present a model of the No Opportunism Game to clarify the meaning of a RAND commitment on a single patent. The model closely mirrors the GK model except that there is only one patent to be licensed instead of two.

Consider a patent that yields net benefits of B but requires a commercialization expenditure of sunk cost S ¹⁵. With expenditure R , the licensee can invent around the patent with probability p ¹⁶. Table 1 lays out the timing of the “No Opportunism Game.” At Time 1, the patent owner sets a royalty, L . At Time 2, the manufacturer decides whether to invest in R&D to circumvent the patent and in complementary inputs. If the manufacturer invests in R&D, the outcome is revealed at Time 3. At Time 4, the manufacturer decides whether to produce. The key feature of the game is that the patent owner sets the royalty before the manufacturer invests in complementary assets or R&D.

At Time 1, there are three levels of royalties to consider: $L_E = B$, $L_V = B - S$, and $L_C = R/p$. (The subscripts denote “expropriation,” “value,” and “cost,” respectively¹⁷.) Since our interest in the No Opportunism Game is when the patent holder has an incentive to commit not to expropriate the value of sunk investments, we restrict attention to parameter values that create such an incentive¹⁸. Thus, to capture the effects we are interested in,

¹³ The assumptions and the phenomenon that the GK model examines are similar to those explored by Choi and Stefanadis (2001), discussed above.

¹⁴ GK assume that R&D to invent around the IP-owner’s patents is entirely duplicative and that it offers efficient licensing terms that do not distort product prices. (That is, they assume lump-sum license fees). Given these assumptions, the R&D to invent around the IP-owner’s patents is socially wasteful and yields no consumer benefits. In practice, the R&D to avoid the need to license one or more of the IP-owner’s patents might yield consumer benefits. Even if that is so, however, the GK results clarify why allowing pure bundling in long-term licensing can make an IP holder more willing to engage in the long-term licensing (perhaps by making RAND commitments) needed to generate complementary investment.

¹⁵ GK allow for different levels of S and for B to be a function of S . The levels of sunk cost that are optimal for both the licensee and society are then part of the model solution. We assume a single possible value of S and B , which can be thought of as the optimized values.

¹⁶ GK allow for a range of R and let p be an increasing function of R . Just as we assumed single possible values for S and B , we simplify the GK framework by assuming single possible values for R and p . GK also assume that the patent holder’s patents have some probability of not being valid or that the technology does not work. They explicitly consider the possibility that the patent owner would offer a long-term license that would obligate the patent user to pay a license fee regardless of whether it uses the technology. They then derive the result that the patent owner would not choose such a license if there is asymmetric information about whether the technology is good or the patents are valid. We assume that the technology is good and the patents are valid and we restrict licenses to those that require payment only when the manufacturer actually uses the patents.

¹⁷ We explain why R/p is cost-based below.

¹⁸ An alternative modeling approach would be to assume that the manufacturer chooses S only after it observes the outcome of its R&D. Under that assumption, the

Table 1
No opportunism game.

Time	1	2	3	4
Decision maker	Patent owner	Manufacturer	Random	Manufacturer
Decision	Set royalty L	R&D, complementary investment	R&D succeeds or fails	Production

assume that $pB < R + S$. When $pB > R + S$, the manufacturer would invest in both R&D and complementary assets even if the IP holder chooses L_E at Time 1, and the IP holder would, for some parameters, choose L_E to charge B with probability $(1 - p)$. The solution to the No Opportunism Game is the RAND rate; depending on the parameter values, that rate can be either L_V or L_C ¹⁹.

Given the above parameter restriction (which allows us to focus on the choice between L_V and L_C at Time 1), the sub-game perfect solution of the game is as follows. At Time 4, the manufacturer produces if it has invested in invent-around R&D and the R&D succeeded. Alternatively, if either it did not invest in R&D or if it invested and the R&D failed, it produces if $L \leq L_V = B - S$. Since the patent holder only gets a positive payoff if the manufacturer produces, it will choose $L = L_V$ at Time 1. As a result, the manufacturer will produce at Time 4 and invest in complementary inputs at Time 2.

If the manufacturer does not invest in R&D, its payoff is $B - S - L$. If it does invest, its expected payoff is $p(B - S - R) + (1 - p)(B - S - R - L)$. The value for L that makes the manufacturer indifferent between investing and not investing is $L_C = R/p$ ²⁰.

If $L_C > L_V$, then the IP owner charges L_V . Investing in R&D to invent around the patent is not profitable in this case, so the IP owner can charge a value-based royalty and not face a risk of having its patent invented around. If $L_C < L_V$, then the IP owner gets an expected payoff of $(1 - p)L_V$ if it chooses value licensing and L_C if it chooses cost licensing. It chooses cost-based licensing if:

$$(B - S) \geq \frac{R}{p} \geq (1 - p)(B - S), \tag{1}$$

and $L = L_V = B - S$ otherwise. If p^* is the probability that makes the patent owner indifferent between the value and cost-based license fee, then the condition for p^* is quadratic with roots:

$$p^* = \frac{1 \pm \sqrt{1 - 4(R/(B - S))}}{2} \tag{2}$$

Based on the above discussion, we can establish:

Proposition 1. At Time 1 in the No Opportunism Game, the patent owner charges $L_V = B - S$ when $B - S \leq R/p$. It also charges $L_V = B - S$ when $B - S > R/p$ and p is within the range given by Eq. (2). Otherwise, it charges $L_C = R/p$.

According to Proposition 1, the owner of a single patent would, in seeking to negotiate license terms before the manufacturer incurs sunk costs or decides whether to invest in R&D, choose cost-based licenses in some cases and value-based licenses in others. It would never seek to charge above the value-based royalty even when the cost-based is higher because the manufacturer would

never agree to pay more than the value-based royalty. One might have guessed that long-term licensing would imply cost-based royalties whenever the expected cost of inventing around a patent is less than the value the patent creates (when the patent user in fact practices it). According to Proposition 1, however, the patent owner's choice between a value-based and cost-based license is more subtle.

Figs. 1 and 2 illustrate this point. Fig. 1 shows a case when Eq. (1) has no real roots. In that case, the patent owner chooses the minimum of the cost-based and value-based fees (as one might expect).

In Fig. 1, the horizontal line is the value-based royalty, the curve is the cost-based royalty, and the downward-sloping line is the patent owner's expected royalty if it charges a value-based royalty. For all values of p , the expected cost of inventing around the patent exceeds the patent owner's expected royalties if it charges a value-based royalty. However, when the expected cost of inventing around the patent exceeds $B - S$, the most the patent-holder can charge is $L_V = B - S$. Given the parameter values underlying Fig. 1, the patent owner maximizes its expected royalties with value-based licensing when p is between 0 and 0.3 and with cost-based licensing for values of p above 0.3.

Matters are more complicated in Fig. 2, however, where Eq. (1) does have real roots. As in Fig. 1, the patent owner charges L_V whenever $L_C > L_V$ (again, because the patent user would never accept a license fee above L_V). However, in the range between the two real roots in Fig. 2, the patent owner chooses L_V even though $L_C < L_V$. When it does so, it sacrifices getting the cost-based royalty with certainty in order to get the value-based royalty with some probability less than 1.

One can debate what this result implies about RAND royalties. On the one hand, Proposition 1 formalizes the principle of setting a non-opportunistic license fee as any royalty that does not exceed what the patent holder would have sought under long-term contracting. On the other hand, some might argue for a stricter view of RAND as a commitment not to charge a rate that leaves an incentive to invent around the patent. Under this stricter view, RAND would be the lesser of the value-based and cost-based royalties. In the context of standard setting, the patent holder is setting a rate that would allow its patented technology to be chosen for the standard. If it sets a rate too high, another SDO member (licensee in our model) will attempt to develop its own technology to define the standard or the SDO members will select an already available alternative, in which case the patent holder would not be chosen for the standard. (We discuss this interpretation more formally in the next section.)

5. Ex ante competition to be the standard—An alternative interpretation of B and S

An important context for the debate over patent tying is standard setting. In particular, technologies typically compete to define a standard and RAND commitments can play an important role in that competition. Many formal standard development organizations request RAND commitments, and even outside of such formal efforts unilateral RAND commitments can be crucial in securing industry support for informal standard competitions within a market (see footnote 4). This section addresses this context, interpreting the model as defining technologies competing to

IP owner would never successfully charge a license fee of B . In Section 5, however, we show that we can interpret S as the opportunity cost of not pursuing an alternative standard, and that interpretation is particularly relevant for cases when IP holders make RAND commitments. Under that interpretation, it would not be reasonable to assume that it is possible to delay the investment decision.

¹⁹ The condition for the value-based royalty to be greater than the cost-based royalty is $B - S > R/p$. Note that the condition $pB < R + S$ does not rule out this possibility.

²⁰ Given that the patent owner is not going to charge above L_V , the manufacturer produces regardless of whether its R&D succeeds. Thus, the expected value of the R&D is pL .

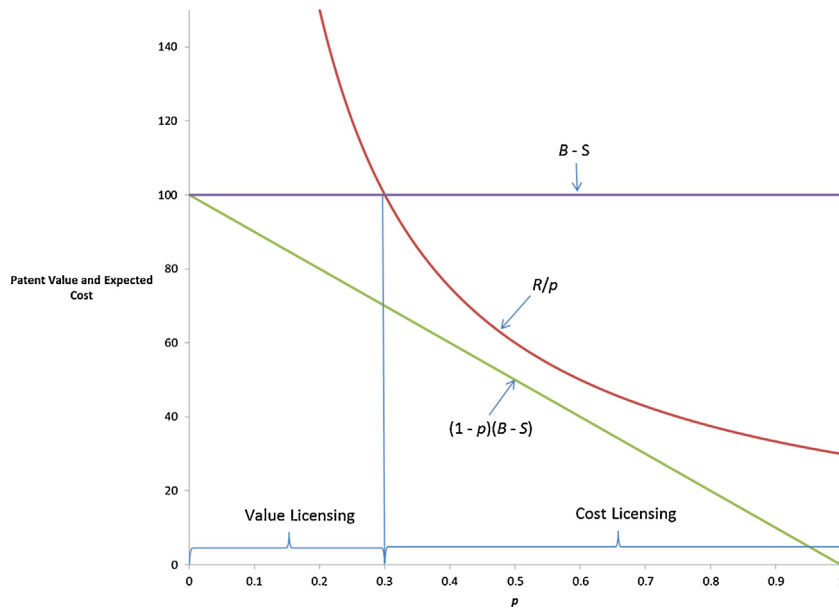


Fig. 1. Cost-based vs. value-based royalties (no real roots to Eq. (1)).

become a standard and considering how that competition impacts the resulting RAND royalty rate.

The two essential features giving rise to the need for standards are (1) competing technologies, each of which could serve as the standard, and (2) network externalities. The GK framework does not explicitly capture those elements, but can be extended to accommodate them. The “Standards Competition Game” introduced below captures both of these essential features; here the manufacturer’s sunk costs in the No Opportunism Game are reinterpreted as the opportunity cost of not adopting an alternative technology as the standard.

Suppose that there are two competing technologies, α and β , and that there are three types of customers, I , II , and III . Correspondingly, there are three types of manufacturers to cater to each type of customer. Let B_{jk}^i be the per person benefit that Group i ($i \in I, II, III$)

gets from technology j ($j \in \alpha, \beta$) given that the standard is technology k ($k \in \alpha, \beta$). Group I prefers α to β but is above all interested in standardization, so $B_{\alpha\alpha}^I > B_{\beta\beta}^I > B_{\alpha\beta}^I$. In contrast, Group II is loyal to α and will choose it even if it is not the standard, while Group III is similarly loyal to β . That is, $B_{\alpha\beta}^{II} > B_{\beta\beta}^{II}$ and $B_{\beta\alpha}^{III} > B_{\alpha\alpha}^{III}$. Let ϕ_{II} and ϕ_{III} be the proportion of customers in Groups II and III (with the proportion in Group I being $(1 - \phi_{II} - \phi_{III})$).

To keep matters simple, we assume that there is no additional cost of adopting a standard other than the opportunity cost of not adopting the alternative technology as the standard and that the probability of inventing around the patent(s) is 0. Table 2 lays out the timing of the game. At Time 1, the owners of both technologies each offer a royalty (L_α and L_β , respectively) conditional on being accepted as the standard. Further assume that only Type I manufacturers participate in the standard setting effort, so that Type II and




Fig. 2. Cost-based vs. value-based royalties (real roots to Eq. (1)).

Table 2
Standards competition game.

Time	1	2	3	4
Decision maker	Patent owners	Type I manufacturers	Owner of non-standard patent	Manufacturers
Decision	Set royalty L_α, L_β	Set std. on α or β	Set royalty $L\alpha'$ or $L\beta'$	Produce or shut down

Type III manufacturers simply take the standard setting outcome as given and have no voice in the cooperative development of the standard. Therefore, at Time 2, the Type I manufacturers choose which technology to adopt as the standard. At Time 3, the owner of the technology not selected as the standard can choose a license fee (not bounded by RAND). At Time 4, manufacturers decide whether to produce.

As with the No Opportunism Game, we assume a sub-game perfect Nash equilibrium. The solution to the game is as follows. Suppose the Type I manufacturers choose α at Time 2. Then Type III manufacturers produce at Time 4 (using the β technology, outside of the standard) if $L_{\beta'} \leq B_{\beta\alpha}^{III}$ and shut down otherwise. Type I and Type II manufacturers produce (using the standard α technology). At Time 3, the owner of the β technology chooses $L_{\beta'} = B_{\beta\alpha}^{III}$ (and earns $\phi_{III} B_{\beta\alpha}^{III}$). At time 2, Type I producers select α if $B_{\alpha\alpha}^I - L_\alpha \geq B_{\beta\beta}^I - L_\beta$, and select β otherwise.

Now consider Time 1. The owner of the β technology can guarantee itself $\phi_{III} B_{\beta\alpha}^{III}$ even if is not chosen as the standard. If it offers a royalty that induces the producer to accept it as a standard, however, it can increase its share from ϕ_{III} to $(1 - \phi_{II})$. The lowest royalty it is willing to offer to be the standard is:

$$\tilde{L}_\beta = \frac{\phi_{III}}{1 - \phi_{II}} B_{\beta\alpha}^{III}. \tag{3}$$

If the owner of the β technology offers \tilde{L}_β , then to become the standard the owner of technology α would have to offer:

$$L_\alpha = \tilde{L}_\beta + B_{\alpha\alpha}^I - B_{\beta\beta}^I = B_{\alpha\alpha}^I - (B_{\beta\beta}^I - \tilde{L}_\beta) \tag{4}$$

as long as:

$$[B_{\alpha\alpha}^I - (B_{\beta\beta}^I - \tilde{L}_\beta)] (1 - \phi_{III}) > \phi_{II} B_{\alpha\beta}^I. \tag{5}$$

If the owner of technology α offered any royalty above that given in Eq. (4), the owner of technology β would undercut it in order to induce producers to accept it as a standard. Thus, given (5), no royalty for α can exceed the value given by Eq. (4). Since we have already established that the owner of the β technology will not offer a royalty below that given by Eq. (3) and that the royalty given by (4) is the best response by the owner of the α technology to the royalty given by Eq. (3), we have established:

Proposition 2. At Time 1 in the Standards Competition Game, the owner of technology α offers the royalty given by Eq. (4) and the owner of technology β offers the royalty given by Eq. (3). The Type I manufacturers choose the α technology as the standard, and the owner of the β technology sets a royalty of $B_{\beta\alpha}^{III}$ at Time 3.

Since we have assumed away the possibility of inventing around the standard, the pricing at Time 1 in the Standards Competition Game corresponds to value pricing in the No Opportunism Game, with $B_{\alpha\alpha}^I$ corresponding to B and to $B_{\beta\beta}^I - \tilde{L}_\beta$ corresponding to S . Note that the opportunity cost of not adopting the alternative standard is the value net of the royalty (rather than the entire value). Introducing invent-around investment would yield results conceptually analogous to those presented above. The intuition can be either that the α technology is easy to invent around for the current standard, or that it could be replaced in the next generation of the standard. Since we formulated the game with Bertrand competition at Time 1, the royalty that the owner of β is assumed

to command requires that β have some value even if it is not the standard. If each patent has no value when it is not the standard, then under the Standards Competition Game, patent holders would be willing to offer a royalty-free license at Time 1. More generally, if the patent owners compete against each other to be included in multiple standards, licensing competition in practice might be “softer” than what we have modeled. In practice, competition may be over technologies with some outside value and hence patent owners are unlikely to set their rates at royalty free.

6. RAND-committed and non-RAND committed patents: The some opportunism game

Having formalized the meaning of a RAND commitment on a single patent, we can now analyze tying a RAND-committed patent to a non-RAND committed patent. To keep matters simple, we assume a sunk cost of complementary inputs (as in Section 4) without specifying whether they are out-of-pocket costs or the opportunity cost of not adopting an alternative technology.

6.1. Analysis

Assume that the patent owner owns patents over two technologies needed to capture value B with investment S into complementary inputs (as in GK). Let R_1 and R_2 be the R&D expenses needed to invent around each patent and p_1 and p_2 be the respective probabilities of success. Suppose the patent-owner makes a RAND commitment on patent 1 but not patent 2.

Table 3 lays out the timing of the “Some Opportunism Game.” At Time 1, the patent owner sets the license fee for patent 1. At Time 2, the manufacturer decides whether to invest in R&D with respect to each patent and whether to invest in complementary inputs. It can try to invent around both patents, patent 1, patent 2, or neither. At Time 3, the results of the R&D become known (by everyone). At Time 4, the patent owner then gets to choose a license fee for patent 2²¹. At Time 5, the manufacturer decides whether to produce.

The manufacturer’s choice at Time 5 limits what the Patent Owner can charge at Time 4, but it does not prevent hold-up. Suppose that at Time 4 the manufacturer needs a license for patent 2, but has invented around patent 1. In that case, the patent owner can set the license fee for patent 2 at B . If, on the other hand, the manufacturer needs a license to both patents, the patent holder can set the license fee for patent 2 at $B - L_1$. Either way, the patent holder earns B and the manufacturer loses all its sunk costs (in complementary inputs and any R&D it performed).

At Time 2, the manufacturer in principle has eight possible choices. Without successful invention around patent 2, however, the patent owner can expropriate any costs that the manufacturer incurs at Time 2; we can therefore immediately rule out any choice that does not entail investing in R&D to invent around patent 2. As a result, the manufacturer has only three realistic options to consider at Time 2. It can invest in complementary assets and R&D to

²¹ Assuming that the manufacturer can wait to learn the outcome of its R&D before deciding whether to invest in complementary inputs would not alter the results qualitatively. Here, we assume that the manufacture decides whether to invest in the complementary inputs at the same time that it makes its R&D decisions to accommodate the interpretation of S as the opportunity cost of not having selected an alternative possible standard.

Table 3
Some opportunism game.

Time	1	2	3	4	5
Decision maker	Patent Owner	Manufacturer	Random	Patent Owner	Manufacturer
Decision	Set royalty L_1	R&D for each patent, complementary investment	R&D succeeds or fails	L_2	Produce or shut down

invent around both patents, it can invest in complementary assets and R&D to invent around patent 2 only, or it can not invest at all. With the first option, there are four possibilities for the resolution of uncertainty. With the second option, there are 2.

Suppose the manufacturer invests only in R&D to circumvent patent 2, which implies that it will need to pay L_1 if it chooses to produce. If the R&D is successful and it produces, the manufacturer's payoff is $B - S - L_1 - R_2$. The patent owner's pay-off is L_1 . If the manufacturer's R&D is not successful, the patent owner sets $L_2 = B - L_1$. In this case, the payoff to the manufacturer is $-S - R_2$ and the payoff to the patent holder is B . The manufacturer's expected payoff if it invests in complementary assets and just patent 2 R&D is $p_2(B - L_1) - S - R_2$.

Now suppose the manufacturer invests in R&D to get around both patents. If both R&D efforts succeed, the manufacturer does not need either license. Its payoff is $B - S - R_1 - R_2$ and the patent owner gets 0. If it successfully invests around patent 2 only, the manufacturer needs a license to patent 1. In this case, its payoff is $B - S - R_1 - R_2 - L_1$ and the patent owner gets L_1 . If the manufacturer's patent 2 R&D fails then, regardless of whether its patent 1 R&D succeeds, the patent owner chooses an opportunistic license fee for patent 2. The patent owner gets total license fees of B and the payoff to the manufacturer is $-S - R_1 - R_2$. The manufacturer's expected payoff if it invests in complementary assets and both patent 1 and patent 2 R&D is $p_1 p_2 B + (1 - p_1) p_2 (B - L_1) - (S + R_1 + R_2)$.

Let $L_1^v = B - (S + R_2)/p_2$. With $L_1 = L_1^v$, which is value-based licensing for patent 1, the manufacturer's expected return from investing in both complementary assets and patent 2 R&D equals its expected return from not investing at all. Let $L_1^c = R_1/p_1 p_2$. If $L_1^c < L_1^v$, then the manufacturer is indifferent between investing and not investing in patent 1 R&D when $L_1 = L_1^c$, which is cost-based licensing.

As in the No Opportunism Game, the patent owner cannot charge more than the value-based license fee for the RAND-encumbered patent. Thus, the cost-based license fee is only a consideration when it is less than the value-based license fee. When it is, the condition for the patent holder to get equal expected payoffs from value-based and cost-based licensing is $[B - (S + R_2)/p_2](1 - p_1) = R_1/p_1 p_2$, which has roots:

$$p_1^* = \frac{1 \pm \sqrt{1 - 4(R_1/(p_2 B - (S + R_2)))}}{2} \tag{6}$$

Thus, we can state:

Proposition 3. At Time 1 in the Some Opportunism Game, the patent owner charges $L_1^v = B - (S + R_2)/p_2$ for patent 1 when $B - (S + R_2)/p \geq R_1/p_1 p_2$. It also charges $L_1^v = B - (S + R_2)/p_2$ for patent 1 when $B - (S + R_2)/p > R_1/p_1 p_2$ and p_1 lies within the range defined by Eq. (9). Otherwise, it charges $L_1^c = R_1/p_1 p_2$ for patent 1. At Time 4, the patent owner charges B for patent 2 if the manufacturer has successfully invented around patent 1 and $B - L_1$ if it has not.

6.2. Interpretation

Proposition 3 summarizes our results about the meaning of a RAND commitment on one patent when the licensee needs access to a second technology to reap commercial value from the first

patent. As in Proposition 1, the RAND royalty is sometimes cost-based and sometimes value-based. However, the need for access to the second technology affects both license fees.

The value-based RAND royalty is $B - (S + R_2)/p_2$ rather than $B - S$. In effect, the technology 2 R&D becomes another sunk cost needed to commercialize patent 1. The expected cost of successful technology 2 R&D is R_2/p_2 and not just R_2 because of the risk that the R&D fails. Moreover, given our assumption that the manufacturer must invest in complementary assets and patent 2 R&D at the same time, the risk that the technology 2 R&D will fail also creates a risk that the manufacturer will not be able to realize any value from its sunk investment in complementary assets. As a result, the value-based license requires subtracting out S/p_2 rather than S from B^{22} .

The need to invest in risky technology 2 R&D also affects the cost-based license for patent 1. In Proposition 3, the cost-based royalty is $R_1/p_1 p_2$ whereas it is R/p in Proposition 1. If it does not make a RAND commitment on patent 2, the patent owner can charge an opportunistic royalty of B for patent 2. As a result, the manufacturer's investment in patent 1 R&D only saves it the cost of a patent 1 license fee when it successfully invests around both patents.

As in Proposition 1, the manufacturer cannot charge more than the value-based fee because the manufacturer would never accept it. However, for some parameter values when the cost-based royalty is less than the value-based royalty, the patent owner might still opt for the higher value-based royalty with probability $(1 - p_1)$ than the lower cost-based royalty with certainty.

6.3. Comparison with RAND royalty on bundle of RAND-committed patents

If the patent owner makes RAND commitments on both patents 1 and 2 and licenses them together as a bundle, then we can interpret Proposition 1 to indicate what the RAND royalty would be for a license to a bundle of the two patents. We can interpret R in Proposition 1 as $R_1 + R_2$ and p in Proposition 1 as $p_1 p_2$. In that case, the value-based royalty would still be $B - S$ and the cost-based royalty would be $(R_1 + R_2)/p_1 p_2^{23}$.

The value-based license fee for patent 1 in the absence of a RAND commitment on patent 2 is less than the value-based license fee for the bundle of patents 1 and 2 given RAND commitments on both. As a result, if a patent owner does offer licenses to its RAND-committed patents only in bundles that include licenses to its non-RAND-committed patents, it cannot justify its proposed license fee on the grounds that it would have been RAND had it made RAND commitments on both.

At first, this result might seem paradoxical because it suggests that a particular royalty for a license to a bundle of two patents can be reasonable in some circumstances but not in others. But the result is not paradoxical at all. Analyzing RAND royalties entails in effect going back in time to imagine what terms the patent owner and manufacturer would have reached if they had contracted at an earlier date. Answering that question requires an understanding

²² If the manufacturer can delay its investment in S until after it learns whether its technology 2 R&D is successful, the value-based royalty would be $B - S - R_2/p_2$.

²³ GK in effect impose the restriction $(R_1 + R_2)/p_1 p_2 > B - S$, which implies that bundle licenses in their model are value-based. We do not impose this restriction and believe that the distinction between value-based and cost-based RAND royalties is an important one.

of what the expectations of both parties would have been at that earlier time. The presence or absence of a RAND commitment on patent 2 affects expectations and therefore the meaning of a RAND commitment on patent 1. Absent a RAND commitment on patent 2, the manufacturer would expect to need to invent around patent 2 in order to commercialize patent 1; and that expectation would limit what the manufacturer would be willing to pay for a license for patent 1.

One might wonder why the patent owner would not make a RAND commitment on patent 2 if doing so would increase the royalty it could command from patent 1. Assuming the patent owner has a choice of whether or not to commit to RAND for patent 2, the logic for doing so is identical to the possible rationale for choosing value-based license fees when the cost-based fee would be lower. In this case, the choice for the patent owner entails a trade-off between a lower-fee with certainty and a higher fee with some probability below 1 but above 0. In the context of a cooperative standard, however, the patent owner may have no choice: it may need to commit to RAND if the patent is disclosed for use in the standard under development. Alternatively, patent 2 may be non-essential for standard compliance in a technical sense, but still complementary and highly useful from a commercial sense, in which case the patent owner need not make a RAND commitment.

7. Conclusions and policy implications

The question we set out to answer is whether a patent owner honors a RAND commitment if it offers its RAND-committed patent in a bundle with non-RAND-committed patents. The question is, however, inherently incomplete.

First of all, offering a bundle of RAND-committed patents and non-RAND-committed patents does not violate a RAND commitment if the patent owner also offers the RAND-committed patent separately on RAND terms. Given the separate offering of RAND-committed patents, the RAND commitment places no additional restriction on the license fee the patent owner can seek for bundles that include the RAND-committed patents²⁴.

Even if we restrict attention to the pure bundling of RAND-committed and non-RAND-committed patents, however, the question we set out to answer remains incomplete. For example, even though RAND commitments are not commitments to offer royalty-free licenses, a patent owner would honor its RAND commitment by offering a royalty-free license to a bundle of patents that includes both RAND-committed and non-RAND-committed patents. Moreover, as we argued in Section 2, bundling (and especially pure bundling) can reduce transaction and litigation costs; and some patent users may demand more inclusive licenses. Thus, the issue to address is not pure bundling in and of itself but, rather, the terms on which a patent owner can offer a pure bundle of RAND-committed and non-RAND committed patents. In particular, can the patent owner take a credit for the value of its non-RAND-committed patents to determine an implicit license fee for its RAND-committed patents and then claim to honor its RAND

commitment if this implicit license fee is RAND? For three reasons, our answer to this question is, “No.”

First, the reason a RAND commitment not to impose opportunistic license terms is important is precisely because of the anticipation that the patent owner will have an incentive to impose opportunistic license terms after the licensee sinks investment. The opportunity to bundle other patents with RAND-committed patents and claim credit for their value in determining RAND terms for the bundle creates a clear risk of using bundling to circumvent the commitment²⁵.

Second, the credit given for the non-committed RAND patents would require an assessment of their value. Absent a RAND commitment on those patents, however, the patent owner can claim an opportunistic value for such patents as their market value. The patent owner is within its rights to claim such a stand-alone value for a non-RAND-committed patent, but forcing the manufacturer to pay such a fee in order to access the RAND-committed patent would amount to renegeing on the commitment not to impose an opportunistic license fee.

Third, imposing a reasonableness requirement on the credit given for the non-RAND-encumbered patent does not solve the problem. As we show, the RAND license rate on patent 1 when there is no RAND commitment on patent 2 is less than what the RAND rate would be on the bundle if both patents had RAND commitments. Had the patent owner made a RAND commitment on patent 2, it would have reassured the manufacturer that it did not need to develop the technology covered by patent 2 on its own and could thus save the expense and risk of making investments to do so. Not having made a RAND assurance, it is not reasonable for the patent owner to extract the same value as if it had. Put differently, if the patent holder is willing to “throw in” other patents—either non-RAND committed patents or patents that are committed for separate and distinct standards—“for free,” such that the licensing terms for the bundle would be RAND with or without the inclusion of the additional patents, the mere inclusion of additional licenses does not necessarily harm consumers. This is not a mere theoretical possibility: given the transaction costs involved in enforcing license agreements, as well as the complementarities across patents in a given technology portfolio, licensors may well find it is in their best interests to add in related non-RAND patents “for free.” Symmetrically, licensees cannot demand a discount for RAND-only licenses, as compared to RAND-plus licenses, without first establishing that the broader portfolio license rates and terms are higher than the RAND-only rates and terms. As a practical matter, however, the complexities of establishing that a RAND-plus license has RAND terms and conditions may push licensors more toward mixed bundling, but the door to RAND-committed patent tying should nonetheless remain open.

We view our results as clarifying and extending the GK arguments for tolerating patent bundling rather than either contradicting or even qualifying them. The GK argument is that allowing patent bundling encourages long-term patent licensing. To the extent that RAND commitments are a form of long-term licensing (or are analogous to such licenses), allowing pure bundling of RAND-committed patents encourages patent owners to make RAND commitments. Restricting the ability to increase licensing fees on RAND-committed patents by claiming credit for the value of non-RAND-committed patents that they bundle with their RAND-committed patents forces patent holders to honor those commitments. Lastly, our model clarifies the arguments that licensees can and cannot make in pressing for a RAND license offer.

²⁴ Strictly speaking, “mixed bundling” entails offering all the components of a bundle separately in addition to the bundle. In the GK model and our simplified extension of it, mixed bundling would entail offering not only a license for the bundle and a stand-alone license for patent 1, but also a stand-alone license for patent 2. Moving away from these simple models, however, a practical application of mixed bundling would entail offering a license to the portfolio of RAND-patents alone, offering a second license to the non-RAND patents (either as a portfolio or in rational subsets), and offering a third license with a bundle of the RAND and non-RAND patents together, rather than having to offer individual licenses of each RAND-patent and each non-RAND patent, for the product offering cost reasons we explain above.

²⁵ Farrell and Weiser (2003) refer to a similar exception to the single rent principle as the “Baxter exception.”

Acknowledgments

We thank Jorge Contreras, Bernhard Ganglmair, Rich Gilbert, Georg von Graevenitz, John Harkrider, Elizabeth Wang and two anonymous referees for helpful suggestions. We thank Jonathan Kowarski for research assistance. We are also grateful to Google for financial support for early versions of this research. The opinions represent the views of the authors and not those of any organization.

References

- Bowman, Ward S., 1957. *Tying arrangements and the leverage problem*. *Yale Law J.* 67, 19–36.
- Choi, Jay Pil, Stefanadis, Christodoulos, 2001. *Tying, investment, and the dynamic leverage theory*. *Rand J. Econ.* 32, 52–71.
- Contreras, Jorge, 2015. *A market reliance theory for FRAND commitments and other patent pledges*. *Utah Law Rev.* 2015, 479–558.
- Elhauge, Einer, 2009. *Tying, bundled discounts and the death of the single monopoly profit theory*. *Harvard Law Rev.* 123, 397–481.
- Elhauge, E., 2015. *Treating RAND commitments neutrally*. *J. Compet. Law Econ.* 11 (1), 1–22.
- Evans, David S., Salinger, Michael A., 2005. *Why do firms bundle and tie? Evidence from competitive markets and implications for tying law*. *Yale J. Regul.* 22, 37–89.
- Evans, David S., Salinger, Michael A., 2008. *The role of cost in determining when firms offer bundles and ties*. *J. Ind. Econ.* 56, 143–168.
- Farrell, Joseph, Weiser, Philip J., 2003. *Modularity, vertical integration, and open access policies: towards a convergence of antitrust and regulation in the internet age*. *Harvard J. Law Technol.* 17, 85–134.
- Gilbert, Richard J., Katz, Michael L., 2006. *Should good patents come in small packages? A welfare analysis of intellectual property bundling*. *Int. J. Ind. Org.* 24 (5), 931–952.
- Harkrider, J., 2013. *REPs not SEPs: a reasonable and non-discriminatory approach to licensing commitments*. *Antitrust Chron.* (10), <https://www.competitionpolicyinternational.com/reps-not-seps-a-reasonable-and-non-discriminatory-approach-to-licensing-commitments/>.
- Layne-Farrar, Anne, 2014. *Moving past the SEP RAND obsession: some thoughts on the economic implications of unilateral commitments and the complexities of patent licensing*. *George Mason Law Rev.* 21 (4), 1093–1110.
- Lerner, Josh, Tirole, Jean, 2004. *Efficient patent pools*. *Am. Econ. Rev.* 94, 691–711.
- Krattenmaker, Thomas C., Salop, Steven C., 1986. *Anticompetitive exclusion: raising rivals' costs to achieve power over price*. *Yale Law J.* 96, 209–293.
- Nalebuff, Barry, 2004. *Bundling as a barrier to entry*. *Q. J. Econ.* 119, 159–187.
- Quint, Daniel, 2014. *Pooling with essential and non-essential patents*. *Am. Econ. J.: Microecon.* 6 (1), 23–57.
- Shapiro, Carl, 2001. *Navigating the patent thicket: cross licenses, patent pools, and standard setting*. In: Adam, Jaffe, Lerner, Josh, Stern, Scott (Eds.), *Innovation Policy and the Economy*, 1. MIT Press, Cambridge, MA.
- Swanson, Daniel G., Baumol, William J., 2005. *Reasonable and nondiscriminatory (RAND) royalties, standards selection, and control of market power*. *Antitrust Law J.* 73, 1–58.
- Whinston, Michael, 1990. *Tying, foreclosure, and exclusion*. *Am. Econ. Rev.* 80, 837–859.