Review of Economic Research on Copyright Issues, 2006, vol. 3(1), pp. 1-17

LICENSING AND ROYALTY CONTRACTS FOR COPYRIGHT

RICHARD WATT

ABSTRACT. This paper reviews briefly how the owner of the copyright to a creation can best market access to that right to licensees under a variety of assumptions concerning the market. After an introductory section, the paper considers a situation of full certainty, in which the value of the final product that is sold by licensees is fully deterministic. In that setting, we consider a very simple model in which the copyright holder himself may or may not compete with the licensee in the final product market. Above all, it is shown that a linear form for the royalty contract always suffices in equilibrium. After that, a model with uncertainty as to the market value of the final product is develloped. In this model, we consider Pareto efficient sharing contracts, and it is shown that now a linear form is unlikely to suffice. Throughout (i.e. in both sections), we shall be interested in exactly when a linear royalty contract is efficient, since these types of contract are so prevalent in the real world.

Finally, as an introduction to the papers contained in the symposium, I devote a few words to each of them in turn.

1. INTRODUCTION

The way in which access to copyrighted creations is remunerated is a particularly interesting aspect of the economics of copyright that can be studied using very simple tools of economic analysis. Copyright royalty contracts are the means under which the economic incentive for creating filters back down to the creator, and as such they are clearly a fundamental issue when the decision to create is taken. However, in as much as royalty contracts mandate how much money goes to the original creator, they also mandate how much goes to any other party involved in the process of creation-delivery-consumption of copyrighted products. And so they are also fundamental to the incentive for publishers, distributors, and even final consumers.[1] In a nut-shell, the royalty contract stipulates how the market surplus is divided between all parties involved.

The economics of copyright royalty contracts has been studied in the literature, but nevertheless, this topic has not received the attention that it really deserves. Economists have spent more time dealing with the effects of copyright infringement, and alternative protection mechanisms, under very rudimentary assumptions on how the creator is financially rewarded. Almost always in such models, there is no "general equilibrium" type effects between how the market works and how the original creators are rewarded. There are many fundamental issues to resolve – exactly how should the contract that provides for the creator's financial reward be structured? Are the typical contracts that are used in the real-world appropriate?

[1]Consumers are affected as the royalty contract will, in the end, have an effect on the price that is charged.

How does the royalty contract affect the level of infringement, and vice-versa? It is the purpose of the current symposium to consider some of these issues.

First and foremost, of course, it is relevant to ask exactly how important are copyright royalty contracts for authors? Recent writings would indicate that they are not important at all as a source of income for authors. Take, for example, the evidence given in Gantz and Rochester (2005, pp. 86-87) concerning rock-star Courtney Love. Apparently, Love attests that even under a 20% royalty rate, the way contracts are structured can still easily leave a rock band with zero income from a hit record. Also, in the same book (p. 81), the authors quote Joni Mitchell as stating "I never really had a good deal in the business". More rigourous economic analyses of the royalty earnings of artists are found in Filer (1986), Towse (1999), and Towse (2001a).[2] While the "starving artist" theory has some weight behind it, it seems clear that artists do earn income from copyright, and that they are motivated by the prospect of such financial reward.

In this introduction to the symposium, some very simply economic theory concerning copyright royalty contracts is discussed. The objective is to "set the scene" for the proper research papers that follow.

2. Royalty Contract Structure – What Do We Know?

Possibly the first economic consideration of the structure of copyright royalty contracts was Baumol and Heim (1967). However, that paper is purely informative rather then analytical. Baumol and Heim set out the general structure of royalty contracts at the time of writing, at least for contracts for written works. It is interesting to see that the status quo was clearly the use of "linear" contracts, indexed to sales revenue. In a simplified world, if sales revenue for the work in question is s, then the royalty contract stipulates that the creator receives a royalty payment of $r(s) = a + bs$. In this simply formulation, a is an "up-front" payment, perhaps designed to either subsidise the author's creation costs, and b is the "royalty rate", that is, the fraction of sales revenue that goes to the creator. Clearly, even this simplified structure allows for a contract with no up-front payment ($a = 0$), and ones that contain only an up-front payment, that is, the author sells outright ($b = 0$), or any intermediate solution with both features present. The linear structure of royalty contracts is something that is still very prevalent today, for all types of copyright works.[3]

It is important to note that the royalty contract structure is not always a choice variable for the contracting parties. Issues of compulsory licensing can imply that the structure of the contract is given exogenously (Gallagher 2002), or at least that the structure of the contract is severely limited (perhaps to adhere to a condition of "reasonability" in the royalty rate). Secondly, a huge issue in royalty contracts is that surrounding the way in which copyright collectives for performance rights in musical compositions act. Collectives market a blanket license to a large repertory of compositions, and as such are directly akin to a monopoly. It is typically assumed that, due to the huge transactions costs in music performance licensing, collectives are a "natural monopoly", and thus there is an efficiency argument for them to exist.

[2] See also Towse (2001b) for a discussion of the financial incentives of artists versus other motives for creation.

[3] See Towse (2001a), for example, for a wealth of information regarding modern day royalty contracts.

But they are also heavilly regulated, as far as the prices they charge to users of the blanket license, and as far as how the revenue is distributed among members (for excellent, and very informative, treatments of the economics of copyright collectives, see Hollander 1984, Besen and Kirby 1989b, Besen, Kirby and Salop 1992, and more recently, Towse 1999 and Einhorn 2002). Thirdly, there is the issue of "droit de suite" (see Perloff 1998 for a general discussion), a case of a generally accepted inalienable right of authors.[4] Droit de suite entitles authors to a percentage of revenue from downstream sales of their works when they are sold by current owners. Solow (1998) presents a particularly good model of optimal royalty structures when early and later works are either substitutes or complements. However "waivable" this right may be, it is generally stipulated by law the exact parameter under which downstream sales revenue should be split (for example, under California law, the author is entitled to 5% of revenue, so long as that revenue exceeds both $1,000 and the price paid by the seller; see Solow 1998).[5] It is, of course, interesting that when legal restrictions impose particular structures upon the distribution of copyright revenue, it typically does so in a linear fashion.

Perhaps the most comprehensive and complete analysis of the contractual means under which copyright revenue is distributed among along the creation-distribution-consumption chain is the book by Richard Caves (Caves 2000), which applies the now traditional principal-agent model and contract theory to the particular case of copyright royalty contracts.[6] Caves' book touches on just so many of the issues relating to royalty and licensing contracts that it becomes difficult to put each of the insights into their relevant subsections below, and so I elect to mention this important book in the introduction of the current paper. Caves presents a wealth of theoretical and empirical information regarding the intricasies of the contractual relationships between authors and distributors, widely defined, in the real world. In the interests of readership, I suppose, he abstains completely from offering any formal models from traditional contract theory, and instead concentrates on a discussion of how certain of the often noted features of typical real-world royalty contracts answer some of the theoretical issues brought up in the literature of contract theory. In particular, perhaps, Caves mentions the issue of informational asymmetries at the time at which the contract is signed. Of course, the principal-agent model *is* relevant for a full analysis of copyright royalty contracts, which *are* fraught with problems of asymmetric information, both of the adverse selection and the moral hazard type. However, the fact that real-world contracts do not tend to conform to the non-linear structures predicted by principal-agent theory under either adverse selection or moral hazard remains a puzzle.

2.1. **"Optimal" royalty contracts under certainty.** Most of the papers on royalty contracts consider the simplified setting of certainty. In a certainty environment, the demand for the creation in question is deterministic, and so the amount

[4] According to Hansmann and Santilli (1997), some juristictions do "grant artists a right of integrity that is nontransferable yet at the same time explicitly or implicitly permit artists to waive that right."

[5] See also Karp and Perloff (1993).

[6] Notwithstanding the importance of Caves' book, it appears that the first time that the principal-agent framework was suggested as an analytical tool for copyright royalty contracts is Towse (1999).

of revenue that will be available for distribution via the royalty contract is also deterministic, and will only depend upon the pricing decision. While clearly not the most relevant assumption for real-world settings, nevertheless these types of models do certainly give good account of many of the underlying features of optimal royalty contracts.

Besen and Kirby (1989a) attempt what is probably the first serious analysis of what an "optimal" royalty could look like. Using a general linear contract structre, they consider the royalty rate that would maximise the profits of the copyright holder. In their model, Besen and Kirby consider the payment of a royalty to be a license to copy, and so the copyright holder can sell both with and without such licenses, under the contract that only in the second case will copying take place with the appropriate transfer of royalty payment. Besen and Kirby's focus is on the issue of indirect appropriation (the ability of the copyright holder to appropriate copying revenue via price discrimination), and so they concentrate exclusively on the licensing royalty fee both when indirect appropriation is possible and when it isn't, and they set out the comparative statics of the royalty rate with respect to such things as the price of originals, the degree of substitutability between originals and copies, the demand for originals and the costs of copying. However, they do not address the more interesting issue of the general structure that the contract itself should take.

Chapter 3 of Watt (2000) is also dedicated to an analysis of copyright royalty contracts. There, an analysis of the optimal linear royalty contract is attempted (with no fixed-fee component) under conditions of deterministic demand, and it is shown how the existence of the threat of piracy may alter the optimal royalty contract (assuming it takes a linear structure). Clearly, since the royalty rate implies a marginal cost increase to the producer-distributor of originals, the choice of royalty rate affects the balance of competitiveness between the producer of originals and any producer of pirate copies. The model in Watt is based on two periods; in the first only the producer of originals operates (this is a "lead-time" period), and in the second a pirate may enter to compete. Intuition suggests that the greater is the royalty rate parameter, the lower will be the producer of original's output in the stage at which she holds a monopoly in the market (and hence the greater will be the second period market in which a pirate may compete), and the less favourable will be the terms at which the producer of originals competes with the pirate in the second period. Thus, the choice of a larger royalty parameter may be detrimental to the total revenue that is finally earned.

The basic idea in Watt (2000) is significantly improved by two of the papers in the current symposium. Firstly, Norbert Michel's paper also takes good account of how the existence of piracy affects the royalty contract, but he does so in a solid bargaining setting, whereas in Watt the model is simplified so that the copyright holder has total bargaining power when the royalty rate is set. Secondly, in Alan Woodfield's paper, it is noted that a constant royalty parameter over both periods, as is assumed in Watt, is certainly sub-optimal. Allowing for a more general (and thus non-linear) structure in which the royalty rate changes over time, is welfare improving for the team of the copyright holder and the producer of originals.

Scotchmer (2005, chapter 6) offers a series of small, but appealing, models relating to licensing of intellectual property in general, under a variety of assumptions on the situation at hand. Of particular relevance to the current discussion, Scotchmer discusses a simple Cournot type setting of royalty contracts. The model in question derives from earlier work (Maurer and Scotchmer 2004), and can be directly adapted to account for the particular case of copyright royalties under certainty, especially for the consideration of the optimality of a linear contract form. Here are the basic principles of that model, duely adjusted to the particular case of copyright (rather than patent, as is the original focus of the idea in Scotchmer 2005).

Assume firstly that the copyright holder licenses access to his creation without competing himself in the market for the final good concerned. For example, this would be the case for the performance rights to songs that are licensed to be played at nightclubs, when the copyright holder is not himself the owner-operator of a nightclub. For simplicity, assume that there are 2 licensees operating in the market for the final good, and that they compete a la Cournot. In this model, the copyright holder retains all bargaining power, and thus establishes a "take it or leave it" offer to the prospective licensees. If the contract is refused, then the prospective licensees each get a payoff of 0, which is then their reservation level of profit. We also assume that the licensees have identical cost funcions (no fixed costs, and a constant common marginal cost), and that the final demand for their output is linear, and of course, totally deterministic. We concentrate on linear royalty contracts.

Concretely, if q_i is the output of producer i, then the costs of this producer (outside of what is charged for access to the copyright under the licencing agreement) is given by

$$C(q_i) = cq_i \quad i = 1, 2$$

On the other hand, the (inverse) demand function for the final product is given by

$$p = 1 - a(q_1 + q_2)$$

Now, assume that the license agreement stipulates that each licensee must pay the copyright holder according to the linear function

$$L = F + rq_i \quad i = 1, 2$$

so that we explicity allow for an up-front fixed fee (F) as well as a royalty payment that is (marginally) independent of output (r). Our objective is to find the optimal (from the point of view of the copyright holder) values of (F, r).

Under the licencing agreement, licensee 1 has a profit function of

$$\pi(q_1) = q_1(1 - a(q_1 + q_2) - c - r) - F$$

and so his first order condition for a profit maximum[7] is

$$\pi'(q_1^*) = 0 \quad \Longrightarrow \quad 1 - 2aq_1^* - aq_2 - c - r = 0$$

Note that this first order condition defines the optimal output of producer 1 as a function of (among other things) the output of producer 2. It can be expressed as the linear equation

$$q_2 = \frac{1 - 2aq_1^* - c - r}{a}$$

In exactly the same way, we can express the optimal output of producer 2 as a linear function of the output of producer 1

$$q_2^* = \frac{1 - aq_1 - c - r}{2a}$$

[7]It is simple to show that the second order condition is satisfied (i.e. that the profit function is indeed concave).

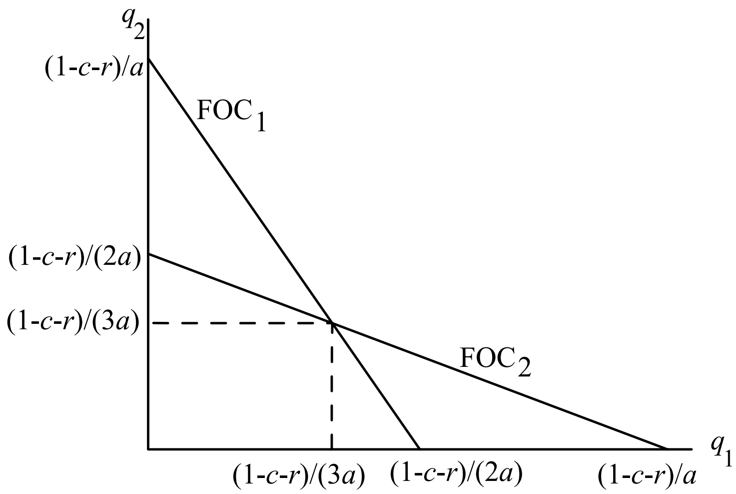

Figure 1

These two first order conditions are shown graphically in Figure 1, with the corresponding Cournot equilibrium at their intersection, $q_i^* = (1 - c - r)/(3a)$ for $i = 1, 2$. It should be immediately obvious that the only difference between figure 1 and a traditional Cournot model with 2 competitors is that here the marginal cost of each producer is $c + r$ rather than just c.

Now, we know that the maximal profit that can be earned in this market corresponds to what would be achieved with a monopoly producer who pays no royalty. Such a producer would produce an amount equal to $q^* = \frac{1-c}{2a}$, and would earn a proft of $\pi(q^*) = \frac{(1-c)^2}{4a}$. Thus, we can think of our copyright holder's problem in two steps – firstly he wants to guarantee that, together, the two licensees produce the monopoly output amount, and then he wants to transfer the entire profit over to himself in the form of licencing revenue.

The first of these two objectives is simple, and can be achieved by the relevant choice of royalty parameter r. Clearly, if each licensee is charged a royalty parameter such that his Cournot equilibrium production strategy is equal to one half of the monopoly output amount, then the total production will be the monopoly amount. Thus, r must be set such that

$$\frac{1 - c - r^*}{3a} = \frac{1}{2}\left(\frac{1 - c}{2a}\right)$$

This simply equation reduces directly down to

$$r^* = \frac{1 - c}{4}$$

Now, given this royalty parameter, we can calculate the optimal up-front fixed fee. Clearly, since each licensee is producing one half of the total monopoly output, the total profits generated are the same as monopoly profit, $\frac{(1-c)^2}{4a}$. This is being shared between the two licensees (who must be each earning the same amount of profit) and the license holder, who is earning the rest. The license holder's revenue from the royalty payment is

$$q^*r^* = \frac{(1-c)}{2a}\frac{(1-c)}{4} = \frac{(1-c)^2}{8a}$$

which is exactly half of the monopoly profit. So each licensee is earning one quarter of the monopoly profit, less any fixed fee that is charged. Quite simply, then, the copyright holder should charge a fixed fee to each licensee of exactly one quarter of the monopoly profit, $F^* = \frac{(1-c)^2}{16a}$, in order that the equilibrium profit of each licensee is set to 0, and the copyright holder gets to keep all of the monopoly profits.

In short, the optimal licencing fee function for the case of two licensees, when the copyright holder himself does not compete in the final product market is given by

$$L(q) = \frac{(1-c)^2}{16a} + \frac{(1-c)}{4}q$$

In short, *a linear licencing fee function is sufficient for the copyright holder to get all of the monopoly profits, which is a situation that cannot be bettered.* However, it is also clear that in order to do this, the fixed fee part of the licencing fee function is very important. Unless this is charged, the copyright holder only earns half of the maximum possible profits.

Now, consider what happens if the copyright holder himself also competes in the final product market. For example, this could be the case when a record label owns a night-club and also licenses access to its songs to other night clubs, or when a singer licenses his own songs to other singers to make cover versions. Again, if the licensee refuses the contract, we assume that his profits will be 0. Our objective once again is to see if a licencing fee function can be found that firstly retains full monopoly profits, and then allows the copyright holder to receive all of these profits as copyright licencing revenue. We now let producer 1 be the copyright holder, and producer 2 be the only licensee. Naturally, the copyright holder could set up an independent company to manage the rights to the copyright, charging both himself and the licensee the right to access. In this case we are straight back to the model just studied above. Here then, we will search for a more direct solution.

We shall assume that the copyright holder charges the licencing fee function only to producer 2, and accesses the copyright himself completely free of charge. Thus, the first order condition of producer 2 is exactly as in the previous section, while that of producer 1 is the same but with $r = 0$. The relevant graph is shown in Figure 2. It is easy to show that the intersection between the two first order conditions occurs in the positive quadrant only if the royalty parameter satisfies $r \leq \frac{1-c}{2}$. In particular, it should be noted that the first order condition of producer 1 lies everywhere above the locus of points such that $q_1 + q_2 = q^*$ (except, of course, for the point where $q_2 = 0$, in which case we get $q_1 + q_2 = q^*$ exactly). This locus is shown as a dotted line in Figure 2.

Since producer 1's first order condition is everywhere above the locus of points such that total production is equal to monopoly output, at any equilibrium (intersection between the two first order conditions) where both are producing strictly positive amounts, the aggregate production will exceed the monopoly quantity, and correspondingly the aggregate profits will be less than monopoly profits. Thus, the best strategy for the copyright holder is to charge a royalty of $r = \frac{1-c}{2}$, thereby ensuring that producer 2 closes down, leaving the entire market to producer 1 to enjoy. In this case, no licencing at all will take place.

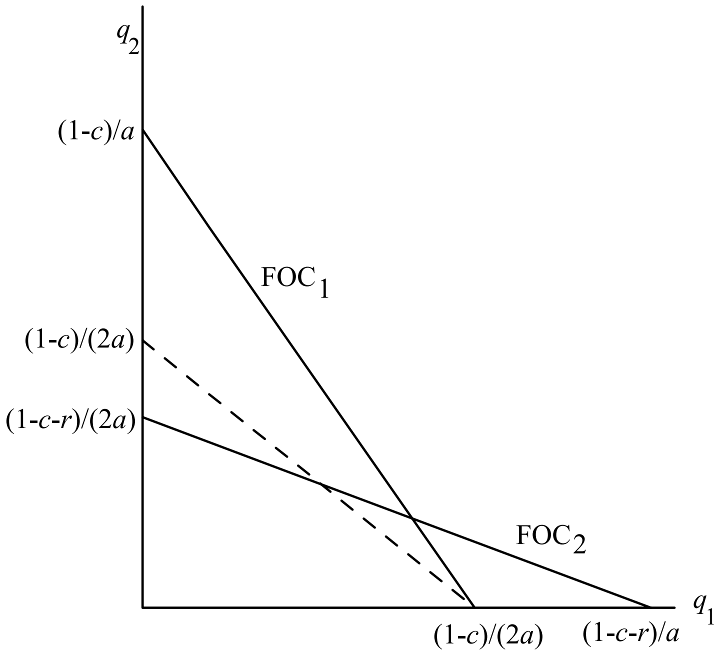

FIGURE 2

Most courts of law will interpret such a solution as being an abuse of monopoly power under anti-trust laws. Hence, the copyright holder may well be advised to search for another strategy under which some licencing takes place. This cannot be done using the licencing fee function alone, assuming that it is restricted to be linear. What is required is for the first order condition of the copyright holder (producer 1) to have at least one internal point at which it intersects the locus of points for which aggregate production is equal to monopoly output. This can be achieved quite simply by applying an output cap to producer 1 at some level $\overline{q}_1$ that satisfies $0 \leq \overline{q}_1 \leq q^*$. However, it will now be neccesary to set the royalty parameter accordingly at a particular level, $r(\overline{q}_1)$. The situation is represented graphically in Figure 3.

From the first order condition of producer 2, we have

$$1 - a\overline{q}_1 - 2aq_2 - r(\overline{q}_1) - c = 0$$

and since by definition $q_2 = q^* - \overline{q}_1$, reordering this becomes

$$\begin{aligned} r(\overline{q}_1) &= 1 - c - a\overline{q}_1 - 2a(q^* - \overline{q}_1) \\ &= 1 - c - 2aq^* + a\overline{q}_1 \end{aligned}$$

Finally, recalling that $q^* = \frac{1-c}{2a}$, it turns out that

$$\begin{aligned} r(\overline{q}_1) &= 1 - c - 2a\left(\frac{1-c}{2a}\right) + a\overline{q}_1 \\ &= 1 - c - (1-c) + a\overline{q}_1 \\ &= a\overline{q}_1 \end{aligned}$$

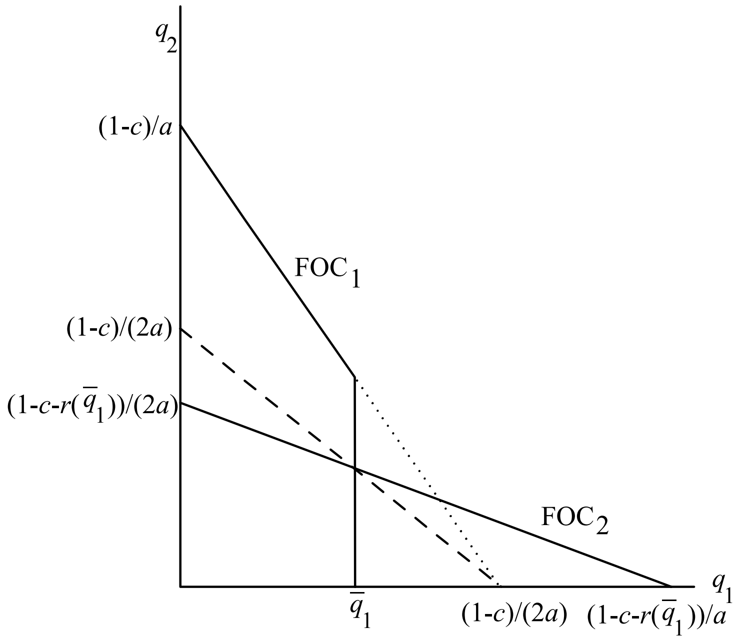

FIGURE 3

For example, using this optimal royalty parameter we get the following particular solutions. If the output cap is set at $\overline{q}_1 = q^* = \frac{1-c}{2a}$, then as was noted above, the royalty parameter is $r(q^*) = \frac{1-c}{2}$. Secondly, if the cap is set at $\overline{q}_1 = 0$, then the royalty is $r(0) = 0$, which is logical since when producer 1 does not produce, the only way to get producer 2 to produce exactly the monopoly output is not to give him any additional marginal costs. Finally, assume that the output cap is set at $\overline{q}_1 = \frac{q^*}{2}$. In this case, we get $r(\frac{q^*}{2}) = \frac{1-c}{4}$, which is exactly what was noted in the previous model with two licensees.

Just as the royalty parameter becomes a function of firm 1's output cap, so must the optimal fixed fee, $F(\overline{q}_1)$. For example, when producer 1 produces the monopoly output then no fixed fee at all is required, since producer 2 is closed down, so $F(q^*) = 0$. On the other hand, if producer 1 decides not to produce at all, as we have just seen the royalty parameter must be set at 0, so in this case the fixed fee will have to be the entire amount of the monopoly profit, $F(0) = \frac{(1-c)^2}{16a}$. Of course, intermediate solutions are also very easy to find, one for each possible choice of output cap between 0 and total monopoly output, and so while *in general a linear contract coupled with an output cap by the copyright holder is sufficient for the copyright holder to earn monopoly profits, in this case there is no single unique optimal contract.*

As a final comment, note that in this model, independently of whether or not the copyright holder actually competes in the market for the final product, if licencing is to take place (i.e. the copyright holder does not close the other producer down) then the optimal licencing fee function always involves both a royalty parameter as well as a non-zero fixed fee. In spite of this, in the real world, it is uncommon to see copyright licencing using a fixed fee.

2.2. **"Optimal" royalty contracts under risk.** As we have just seen, in a fully deterministic setting, it seems likely that a linear royalty contract will suffice. But the same is not true when we introduce risk into the picture. However, the consideration of royalty contracts under certainty is well developed, but the theory under uncertainty is still very much in its infancy.[8]

When risk is present, under symmetric information, the royalty contract can be thought of as a mechanism under which not only remuneration, but also risk sharing takes place. Basically, it is true that a copyrighted creation is a stochastic prospect as far as the income that it generates in the market is concerned. Before any marketing actually takes place (which is when the copyright royalty contract must be drawn up and signed), the final revenue that will be available to be distributed between the copyright holder and the producer-distributor is uncertain, or at best it is risky.[9] Given that, the royalty contract acts as a risk-sharing devise. This aspect of royalty contracts was first analysed by Watt (2000, chapter 3), and more fully by Alonso and Watt (2003).

It is quite well known from the theoretical literature on risk sharing that a linear contract does not share risk efficiently outside of the case in which both parties to the contract have constant relative risk averse utility functions, with the same degree of risk aversion. While an assumption of constant relative risk aversion for both parties is quite believable in the copyright setting, it is almost certainly not the case that copyright holders (individual song-writers and singers, for example), will have the same degree of risk aversion as producer-distributors (large corporate record labels, for example). Thus, assuming symmetric information, linear royalty contracts would generally be inefficient.

If we make the assumption that the producer-distributor is less risk averse than is the copyright holder,[10] while retaining the assumption of constant relative risk aversion, under an efficient royalty contract the former should insure the latter. This would imply that the copyright royalty rate (of revenue) to be paid to the copyright holder should be greater the smaller is the revenue that eventuates. Of course, whenever such non-linear structures actually appear in the real-world, they typically do exactly the opposite; a greater royalty rate is offered when sales revenue is higher. This apparent paradox can be (partially) explained if we remove the symmetric information setting, so that an incentive to create greater quality is given by a contract that discriminates against the contingency of poor market revenue. However, in reality, I have yet to find a sufficient explaination for the paradox for the (abundant) cases in which the royalty contract is actually signed after creation has taken place (and the quality of the creation is fully observable by the producer-distributor).

[8] In a twist on the literature of copyright collectives, which is traditionally studied under certainty, Snow and Watt (2005) analyse the risk sharing benefits of such collectives that currently go begging in the real world.

[9] Here I take the well-known Knightian definitions; uncertainty is a case in which the probabilities of each possible outcome are not known (we would say the probabilities are subjective), while risk is a case in which all parties share the same probability beliefs (probabilities are objective).

[10] Indeed, strong evidence of this is given by the fact that royalty contracts do often insure creators by offering an up-front payment that is taken as being a royalty advance. Alonso and Watt show that this exactly mimics an insurance contract with a deductible clause.

In what follows, I breifly outline the basic theory from Alonso and Watt (2003), in an attempt to provide a simplified treatment of how risk affects royalty contracts.

Assume that the total surplus that is generated in the market for the final copyrighted good is random. For simplicity, assume that there is only one licensee, and that the copyright holder himself does not compete in the final market. For example the copyright holder could be a singer (or an author, or a copyright collective), and the licensee is a record label (publisher, night-club). The copyright holder and the licensee act together to generate an uncertain surplus, which they then share according to a predetermined sharing rule (a royalty contract). We assume that the licensee receives the surplus, and then pays the copyright holder. In order to retain a simple graphical environment, we assume that there are only two states of nature. In state 1, the surplus to be shared is high (the copyrighted song is a hit, for example), and in state 2 it is low. Denote the surplus that is generated in the market by x, so that in state 1 we have x_1 and in state 2 we have x_2, where $x_1 > x_2$. Assume that the probability of state 2 is p, which is strictly between 0 and 1, and that both the copyright holder and the licensee fully agree on the value of p. Finally, we assume that both the copyright holder and the licensee are risk averse, that is, the utility function of player i is $u_i(z)$, which is strictly increasing and concave. We denote the copyright holder by $i = c$ and the licensee by $i = l$.

The sharing agreement must stipulate two numbers, say (s_1, s_2), where s_i is the amount that the licensee pays the copyright holder in state i. Given this, the relevant graphical tool for analysing the situation is the Edgeworth box, in which total state surplus is measured on each axis. This is done in Figure 4, where the lower origin corresponds to the copyright holder, and the upper origin corresponds to the licensee. The equilibrium contract, s^*, is a point on the contract curve, that is determined by some relevant process (perhaps by a bargaining game).

In Figure 4, the indifference curves of each player are shown tangent to each other as they cut throught the contract curve. The lower indifference curve corresponds to the licensee, and the higher one to the copyright holder. The shapes of these two indifference curves show that each player is strictly risk averse. Secondly, the two parallel lines that extend from the two origins are the player's respective certainty lines (the loci of points such that each player would be indifferent to which state occurs). The fact that the two certainty lines do not coincide (i.e. that the box is a rectangle rather than a square) corresponds to our assumption of existence of aggregate risk, that is $x_1 > x_2$. Finally, although the contract curve has been drawn passing through the two origins, this is not necessarily the case. Indeed, all that is necessarily true is that the contract curve cannot touch either certainty line at an interior point.[11]

As in the previous subsection, here we shall be particularly interested in the possibility that the efficient solution, s^*, be a point such that the royalty function is linear, that is, if the efficient solution s^* can be expressed as $s_i^* = L(x_i) = F + rx_i$ $i = 1, 2$. Such a solution requires that we find two numbers F and r such that $F + rx_1 = s_1$ and $F + rx_2 = s_2$. It is trivial to solve these two equations to get

$$r = \frac{s_1 - s_2}{x_1 - x_2} \qquad F = \frac{s_2 x_1 - s_1 x_2}{x_1 - x_2}$$

[11]The formal proof of this statement, and of those that follow, can be found in Alonso and Watt (2003).

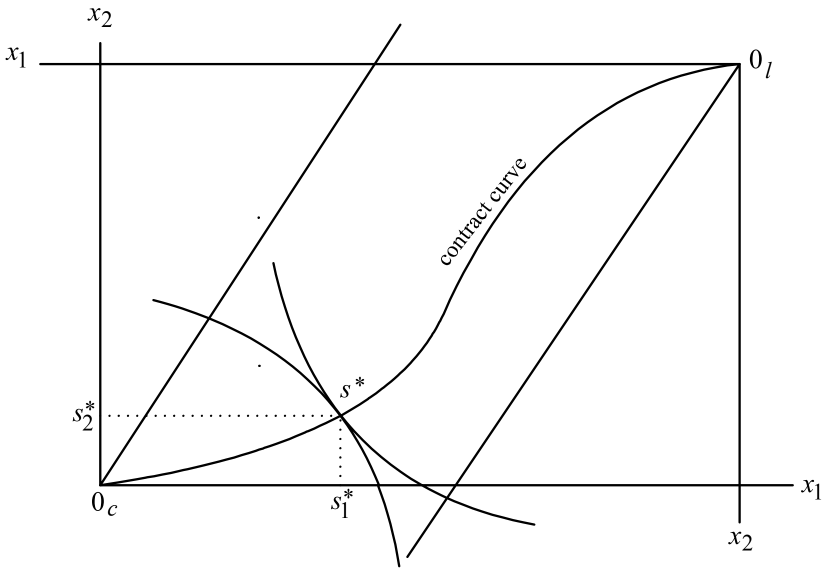

FIGURE 4

However, since both F and r depend on x_1 and x_2, any change in either of these parameters will alter the linear relationship. It is also clear that if we abandon the two dimensional nature of the situation, then we would have to compute 2 parameter values from more than 2 equations, and so the system would be overspecified.

In real world contracts, we often see a system used that is a contract with constant proportional sharing and no fixed fee. For example, author's rights in books are often set at approximately 10% of sales, independently of whether sales are high or low. In this two dimensional setting, this requires $s_i = rx_i$ for $i = 1, 2$. However, since this imples that $r = \frac{s_1}{x_1}$, we get $s_2 = \left(\frac{x_2}{x_1}\right) s_1$, that is, it is a point on the diagonal of the Edgeworth box. Given this, we now ask the question; can the general solution (a point on the contract curve) can also locate on the diagonal of the box? More concretely, can such a point be consistently the general solution to the problem?

The answer to both questions is yes, but while the first (that a point on the contract curve is also located on the diagonal) is a quite general property, the second is not (that such a point will consistently be the solution to the problem). Concretely, in all cases (i.e. independently of the exact utility functions of the two parties to the contract), the contract curve must intersect the diagonal of the Edgeworth box at least once (but it could do so at one of the origins), but in quite a host of cases the intersection between the contract curve and the diagonal is unique (that is, there is only one). For example, there is a unique intersection if both of the participants have constant or increasing absolute risk aversion.

What about the question of whether the solution will consistently be on the diagonal of the box? The answer is that it will if and only if both the copyright holder and the licensee have constant relative risk aversion, with the same risk aversion parameter. This is simply because under the stated condition, the contract curve coincides exactly with the diagonal of the Edgeworth box. And this is the only case under which this happens. The proof of the "only if" part of this satement is particularly difficult, and will not be attempted here. But an explanation of the "if" part is simple. Since both parties to the contract have constant relative risk averse utility functions, we have

$$u_i(z) = \frac{z^{1-R}}{1-R} + k_i$$

where the subscript i refers to the participant (copyright holder or licensee), and k_i are constants.[12] The parameter R is the value of relative risk aversion for each of the two participants. Using this utility function, marginal utility is simply $u_i'(z) = \frac{1}{z^R}$, and so since the contract curve is the set of points such that the two marginal rates of substitution (the ratios of marginal utility) are equal, recalling that the copyright holder earns s_i and the licensee earns $x_i - s_i$ in state i, we get the result that the contract curve is the set of points that satisfy

$$\frac{s_2^R}{s_1^R} = \frac{(x_2 - s_2)^R}{(x_1 - s_1)^R}$$

That is, the contract curve satisfies $(x_1 - s_1)s_2 = (x_2 - s_2)s_1$, which simplifies directly to $s_2 = \left(\frac{x_2}{x_1}\right) s_1$, which is the diagonal of the Edgeworth box.

However, this result is rather delicate, since any divergence between the two degrees of relative risk aversion (both still being constant), leads to the result that *a fixed proportion sharing rule is never efficient*. There is a simple intuition for the result. If, say, the copyright holder is more risk averse than the licensee, then the efficient contract will always require that the former insures that latter (in the sence that the contract will asign more risk to the licensee than to the copyright holder). This means that the contract curve will locate everywhere above the diagonal of the box, bowed towards the copyright holder's certainty line (see Figure 5). On the other hand, if the licensee is more risk averse, then for the same reason the contract curve will be bowed towards his contract curve. In both situations, the contract curve only touches the diagonal of the box at the origins (i.e. situations in which one party retains all of the sales revenue, which will never be an optimal contract since it will not satisfy participation of the party that is not paid).

Clearly, this analysis leads to two interesting empirical anomalies. While it is believable that both participants have constant relative risk averse utility functions, the two anomalies eventuate from the comparison of their relative values of risk aversion. Firstly, as was noted above, it is often seen that the contract used is located on the diagonal of the contract curve. However, this can only be efficient (under the assumptions used here, concretely symmetric information) if the copyright holder and the licensee are equally risk averse. This is hard to believe in those cases in which the licensee is a large corporation like a record label, while the copyright holder is an individual songwriter or singer.

Secondly, and even more strangely, we also often see contracts under which the copyright holder recieves a greater royalty percentage the higher are sales. In order for this to be efficient, the contract curve must be located below the diagonal of the box, that is, it would require that the copyright holder is insuring

[12] It is well known that this is the only utility function that is consistent with constant relative risk aversion.

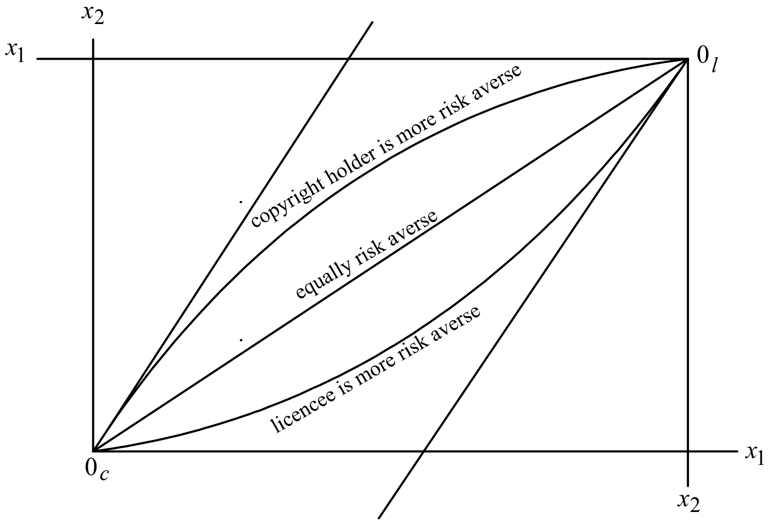

FIGURE 5

the licensee (and therefore that the copyright holder is the least risk averse of the two).[13] This anomally is accentuated by situations in which the licensee pays an advance payment against future royalties, something that can only be consistent with the licensee being less risk averse than the copyright holder, since it involves the former insuring the latter.

3. The Symposium Papers

The economic theory of copyright royalty contracts is still very much being developed. There are many aspects of general contract theory that still could be incorporated into a satisfactory theory of royalty contracts (asymmetric information is, at least in my opinion, right at the top of the list). The present symposium attempts to point to some of the most salient features of contract theory in general, of royalty contracts in particular, and how they relate to the more general economic theory of copyright.

In the next paper, an expert in the economic theory of bargaining, Abhinay Muthoo, sets out the principal aspects of the royalty contract environment that allow the equilibrium contract to be analysed as a bargaining problem. Certainly this is the most appropriate way in which the general problem of the structure of the royalty contract should be analysed. The simplistic setting in which one party to the contract has all bargaining power that is so often used in the existing literature may have had its day. The elementary model has proven useful for establishing some basic understanding and insights, but it is not the way forward. Copyright holders *do* bargain with the firms that produce and distribute their intellectual

[13]This may be the response to a problem set in asymmetric information. However, it is well known that the high effort (or good type, in an adverse selection model) contract must locate below the agent's certainty line, but not necessarily below the diagonal of the relevant Edgeworth box. Thus it is difficult to see why the diagonal should be a reference point for an incentive compatibility problem.

property, either individually or collectively, and it is natual that economists now go ahead and attempt to include such bargaining into the theory of royalty contracts.

Following on from Muthoo's paper, Norbert Michel does a wonderful job of explaining how the existence of a threat of piracy affects the royalty contract, when that contract is indeed the result of a bargaining process. Although quite formal in nature, a careful reading of Michel's paper is truely enriching. This is the first attempt of which I am aware to develop a complete theory of the economics of copyright, from the particular aspects of the demand for copyrighted products, through to the way in which authors are remunerated, with all of the interrelations between different aspects of the problem in place.

The next paper in the symposium, by Alan Woodfield, points out another very interesting direction for research. Going back to a simple setting in which the royalty contract parameters are given, Woodfield confirms that relaxing the condition that the contract takes the same royalty rate in each period is certainly welfare enhancing for the supply side of the market (copyright holder and producer-distributor). This insight appears because of the marginal cost disadvantage that a producer of originals faces as compared to a pirate producer when the former pays a copyright royalty and the latter does not. Indeed, the supply side can do better if the copyright holder subsidises the producer of originals by accepting a negative royalty rate when a pirate is present in the market, in exchange for a larger than normal royalty rate before the pirate has established himself. This is indeed an interesting avenue for future research, as Woodfield's paper turns up a host of numerical simulations that indeed point to a general result, yet to be established theoretically.

The theory of copyright royalties is intimately linked to the theory of collective administration of copyright. When a collective is in force, the bargaining power of the supply side of the market is enhanced, transactions costs are saved, and we are faced with the problems of pricing the blanket license and of exactly how the revenue from a blanket license is to be distributed to the collective members as royalty income. The paper by Michael Einhorn, a well known researcher and consultant on issues of collective management of copyright, retakes the economic regulation issue of the possible anticompetitive use of blanket licensing for generating royalty income.

The final paper in the symposium is by Martin Kretschmer, who presents a new set of data concerning the actual earnings of artists from copyright. Kretschmer calculates the Gini coefficient for artists' copyright earnings, and is thus able to quantify the disparity of earnings over artists. Not surprisingly, it is found that copyright royalty income is an extremely risky source of earnings (if Gini coefficients, being closely related to variance, can be taken as a proxy for riskiness), with some huge winners and plenty of cases of non-earners. This in itself suggests an interesting line of research yet to be properly initiated; if copyright income is so risky for authors, is it also risky for the producer-distributors? How can royalty contracts be modified to appropriately account for the spreading of the income risk from copyrighted creations?

Some of the papers in this symposium are somewhat demanding, especially for non-technical readers. This is purely a reflection on the nature of the topic itself. Royalty contract theory is, after all, contract theory, which has traditionally been presented in formal guise. However, where possible, all of our more technical authors provide plenty of understandable intuition, and the insights that can be gained from an attentive reading of these papers are indeed important. I sincerely hope that this symposium proves to be a mere seedling in a future forest of research on this stimulating topic for economic theoreticians.

References

Alonso, J. and R. Watt (2003), "Efficient Distribution of Copyright Income," in Gordon, W. and R. Watt (eds.), *The Economics of Copyright: Developments in Research and Analysis*, Cheltenham UK and Northampton MA: Edward Elgar Publishing Ltd.

Baumol, W. and P. Heim (1967), "On Contracting with Publishers: Or What Every Author Should Know", *American Association of University Professors Bulletin*, 53; 30-46. Reprinted as chapter 22 in Towse, R. (ed.), *Baumol's Cost Disease: The Arts and Other Victims*, Cheltenham UK and Northampton MA: Edward Elgar Publishing Ltd.

Besen, S., S.N. Kirby and S. Salop (1992), "An Economic Analysis of Copyright Collectives", *Virginia Law Review*, 78; 383-411.

Besen, S. and S.N. Kirby (1989a), "Private Copying, Appropriability, and Optimal Copying Royalties", *Journal of Law and Economics*, XXXII; 255-280.

Besen, S. and S.N. Kirby (1989b), "Compensating Creators of Intellectual Property: Collectives that Collect", *RAND ReportN^{o}R-3751-MF*, Santa Monica: RAND.

Caves, R. (2000), *Creative Industries: Contracts Between Art and Commerce*, Cambridge MA and London UK: Harvard University Press.

Einhorn, M. (2002), "Music Licensing in the Digital Age", in in Towse, R. (ed.), *Copyright in the Cultural Industries*, Cheltenham UK and Northampton MA: Edward Elgar Publishing Ltd.; pp. 165-77.

Filer, R.K. (1986), "The 'Starving Artist' – Myth or Reality? Earnings of Artists in the United States", *Journal of Political Economy*, 94; 56-75.

Gallagher, T. (2002), "Copyright Compulsory Licensing and Incentives", in Towse, R. (ed.), *Copyright in the Cultural Industries*, Cheltenham UK and Northampton MA: Edward Elgar Publishing Ltd.; pp. 85-98.

Gantz, J. and J.B. Rochester (2005), *Pirates of the Digital Millennium*, Upper Saddle River NJ: Prentice Hall.

Hollander, A. (1984), "Market Structure and Performance in Intellectual Property: The Case of Copyright Collectives", *International Journal of Industrial Organization*, 2; 199-216.

Karp, L. and J. Perloff (1993), "Legal Requirements that Artists Receive Resale Royalties", *International Review of Law and Economics*, 13; 163-77.

Maurer, S. and S. Scotchmer (2004), "Profit Neutrality in Licensing: The Boundary Between Antitrust Law and Patent Law", Working paper CPC04-43. Berkeley, CA, Department of Economics, University of California.

Perloff, J.M. (1998), "Droit de Suite", in Peter Newman (ed.), *The New Palgrave Dictionary of Economics and the Law*, Volume 1, London: Macmillan; pp. 645-8.

Scotchmer, S. (2005), *Innovation and Incentives*, Cambridge MA: MIT Press.

Solow, J.L. (1998), "An Economic Analysis of the Droit de Suite", *Journal of Cultural Economics*, 22(4); 209-26.

Snow, A. and R. Watt (2005), "Risk Sharing and the Distribution of Copyright Collective Income", in Takeyama, L. et al (eds.), *Developments in the Economics of Copyright: Research and Analysis*, Cheltenham UK and Northampton MA: Edward Elgar Publishing Ltd.; pp. 23-36.

Towse, R. (2001a), *Creativity, Incentive and Reward. An Economic Analysis of Copyright and Culture in the Information Age*, Cheltenham UK and Northampton MA: Edward Elgar Publishing Ltd.

Towse, R. (2001b), "Partly for the Money: Rewards and Incentives to Artists", *Kyklos*, 54(2-3); 473-90.

Towse, R. (1999), "Copyright and Economic Incentives: An Application to Performers' Rights in the Music Industry", *Kyklos*, 52(3); 369-90.

Watt, R. (2000), *Copyright and Economic Theory: Friends or Foes?*, Cheltenham UK and Northampton MA: Edward Elgar Publishing Ltd..