Copyright Collectives and Contracts: An Economic Theory Perspective

by Richard Watt (University of Canterbury and SERCI)

Abstract

The existing economic theory of copyright collectives, or copyright management organizations (CMOs) is strongly focused on the benefits of sharing of transaction costs. Here, we appeal to the contractual environment of CMOs to offer a different perspective. Copyright collectives form contracts at two principle points along the supply chain. First, there are the contracts between the collective's members themselves (the copyright holders), for distribution of the collective's income. And second there are the licensing contracts that the collective signs with users of the repertory. Using standard economic theory, the paper argues that there are significant efficiency benefits from having copyrights managed as an aggregate repertory, rather than individually, based on risk-pooling and risk-sharing through the contracts between the members themselves. Similarly, there are also aggregation benefits (at least in terms of the profit of the CMO) of licensing only the entire repertory, rather than smaller sub-sets. Interestingly, there is a link between these two theories of the efficiency of collective, rather than individual, management, and it lies at the heart of the theory of syndicates, and the characteristics that imply that the group (or syndicate as a whole) can be considered as a valid "representative", sharing the same preferences as each individual syndicate member.

1. Introduction

Copyright collectives are formed when groups of copyright holders join together into a single unit for the purposes of exploiting the economic rights in their different copyrights. Reduced to their most basic structures, copyright collectives carry out three main tasks on behalf of the members (see Hollander, 1984):

a. They license to users access to the copyrights of their members and they collect royalty payments from users;
b. They distribute the royalty income among the collective members;
c. They monitor the use of the copyrights of their members, they enforce the legal copyright parameters, and they bringing action against copyright infringements on behalf of their members.

Under the standard economic theory of copyright collectives (see, for example, Besen, Kirby and Salop 1992), the foundational aspect upon which a copyright collective forms is the existence of transaction costs that can be efficiently shared when copyrights are exploited together. In short, for a contract to be written between a given copyright holder and a user interested in accessing the relevant copyright, there exist many transaction costs, including (i) initial search costs so that the user and the copyright holder can locate each other, (ii) bargaining costs to settle on an agreeable royalty, (iii) costs of monitoring use and collecting the relevant royalties, and (iv) the costs of ensuring that the contract is respected (both by the user, and by other non-contracting users). For some uses, but not all, these transaction costs can be greater than the benefit that contracting has for the parties concerned, and in such cases (absent collective management) the contract will not take place. Furthermore, since

many users want to contract with a similar set of many copyright holders (and vice-versa), if the contracts are carried out individually the aggregate transaction costs multiply unnecessarily, with many contractual actions that generate costs simply replicating actions already carried out for a different contract.

On the other hand, if the copyright holders join together into a unified group, and if all that is offered to users is a blanket license for access to the copyrights of all of the works together, then the transactions costs are hugely reduced, and the implied savings can be shared on both sides of the ensuing contract. This is a theory of natural monopoly based on the sharing of transaction costs, in that when transaction costs are factored into the business model, the costs of running a collective are sub-additive (average cost diminishes with the size of the collective). In such an environment, it is efficient that licenses are granted collectively rather than individually.

However, of course there are a few caveats to this efficiency argument. Most importantly, collective management reduces enormously the choices of licenses that users can negotiate. Under individual licensing, users can restrict themselves to licensing only those works that they are interested in, while under blanket licensing, which is all that CMOs typically offer, they are forced to license all works, those they actually want and also those that they do not want. It is often argued that this feature is unfair to users, and is a source of inefficiency. However, that has yet to be proven to be the case in general, and indeed it is unlikely to be able to be proved. One must look to the alternatives that are actually feasible, and for licensing, the costs of establishing differential licenses for different users, according to their individual preferences and desires, may well make the users worse off than under a blanket license, as the prices would have to reflect the transaction costs implied. It is not in vain that blanket licensing as a response to transaction costs is a very prevalent, acceptable and non-controversial feature of many economic transactions; bus tickets that allow a variety of travel distances for the same price, road user charges for private motor vehicles that give drivers the right to use roads that they have absolutely no intention whatsoever of using, gymnasium memberships that allow members to turn up as much as they like, pay TV channel subscriptions that cost the same whether you watch a lot of TV or a little, Microsoft Office which includes programs that many of us never use, and newspapers (and academic journals) which contain many articles that are not actually read by all readers, even though they could do so if they wanted to.

In this paper I will not consider any further the transaction costs rationale for collective management of copyright. That topic has been extensively covered in prior literature. The only aspect of the transaction cost theory that is of any real current interest is whether or not digitization and the digital environment, which undoubtedly reduces transactions costs on many dimensions, is sufficient to destroy the natural monopoly aspect of copyright management (see Katz, 2006).[1] If so, one would expect to see individual contracting becoming the norm, and the "demise" of copyright collectives. While it may be true that individual contracting is more prevalent now than a decade or so ago, it is only for very specific and determined types of uses, and I also think that it is undeniable that collective copyright management is showing little sign of demise.[2] Indeed, collective management can also take good advantage of digitization to streamline their business, to the benefit of both copyright holders and users.

[1] Katz's argument that DRM would be able to substitute for collective management has not been without challenge, most recently for example by Towse (2012 and 2013).

[2] For example, the data on royalty collections through CMOs which is available on www.cisac.org shows world-wide year-on-year increases, in constant prices, of around 4 to5 percent.

Rather, I will concentrate exclusively upon the contracting environment of copyright collectives. There are two issues of maximum interest; (i) the contracts between the members of the collective themselves (or what is the same, the contract between a given member and the rest of the collective), above all for distribution of royalties generated from the blanket license, and (ii) the contracts between the collective as a whole and the users of the repertory of copyrighted works. Along the way, we will see that at both of these contracting points we find strong arguments for the efficiency of managing access to copyrighted works collectively rather than individually, that is, the transaction cost rationale for collective management of copyright is not the only reason why collectives are efficient.

2. Basic theories of the firm, which one(s) fit(s) for a collective?

Above we have already briefly outlined the theory of natural monopoly, based upon decreasing average costs, that it is often argued characterises copyright management. However there are several other possible theoretical scenarios of production that may also be relevant.

a. Standard theory of production

At its most basic perception, a copyright management organization (CMO) is a type of firm – it takes as inputs access to certain rights in a series of individual compositions, and sells as an output access to certain rights in an aggregate set of compositions. As an organizational unit, most definitions of a firm "share the idea that a firm should be able to produce (or sell) more efficiently than would its constituent parts acting separately" (Tirole 1988, pg. 15). The firm's ability to take advantage of economies of scale or scope determines the optimal size of the firm. Clearly, if that were not the case, then the firm would split up into smaller firms, or perhaps individual contractors. For the case of a CMO, the assumption that the firm is more efficient than its constituent parts directly argues in favour of collective, rather than individual, licensing.

Nevertheless, there are aspects of CMOs that might differentiate them from other more standard firms, and that might lead us to a different theory of the firm that is more specifically appropriate for CMOs. There are two principle special theoretical features that warrant mention.[3] First, the CMO's production activity is very restricted, all it does is re-package by aggregation, and does not actually "produce" anything new (aside, perhaps, from the services it offers its members which include negotiating license fees, collection of fees, and supply of monitoring and enforcement services). The efficiency gains are therefore not so much in terms of productive efficiency, but in terms of transactional cost savings and the aggregation economies of scale due to bundling. Second, the input suppliers (individual copyright holders) are thought of as the owners (like shareholders) of the CMO, and so a CMO is really a mutual firm, or what is known in the economics literature as a "syndicate". This has important implications for the objective function of the CMO,[4] and for the way in which input suppliers are remunerated, both of which may be significantly different from the standard theory of the firm.

[3] There are other concerns as well. For example, CMOs are severely constrained by copyright law, which operates within national boundaries. Whereas much of CMO business is cross-border. This creates a tension that is not present for many other types of firm.

[4] Rather than maximising total profit, a strong argument can be made for a CMO to maximise average profit, or profit per-member.

b. Theory of two-sided markets

It is tempting to think of a CMO as a two-sided platform, since its function is to bring users and copyright holders together. As such, a CMO provides service to both users and to copyright holders, and so each should be willing to pay a price to access the platform. However, I think that this point of view is a bit of a straw-man. Indeed, *any* firm that uses inputs to produce an output that has positive demand provides (to a certain degree) a function of bringing together and providing benefit for both demanders (of final output) and suppliers (of inputs). The only aspect of a CMO that makes it look somewhat more like a two-sided platform is the fact that it hardly produces anything new from its inputs. However, if the CMO were to be restricted to the standard functions of a two-sided platform, then its activity would only involve providing the meeting place for users and suppliers (at some cost to each), and then each user and each supplier would contract individually, rather than through the intermediary that is the CMO. Clearly, this is not what a CMO does, and thus the relevance of it being a two-sided platform is somewhat lost.

Furthermore, the standard theory of a two-sided platform only allows one to consider the optimal access price for users to the platform, and the optimal aggregate surplus that would be passed back to platform participants to share. Concretely, for the case of CMOs, the standard theory of two-sided platforms would not establish the manner in which copyright holders should each share in the aggregate surplus, nor what exactly is the final product that should be offered to the users – a single blanket license to the entire repertory, or a set of licenses to smaller parts of the repertory.

c. CMOs as libraries

I like to think of CMOs as a type of "library", like a normal library of books.

CMO	**Book library**
Does not sell copyrights to works, it only supplies access to certain rights in compositions for a limited period of time.	Does not sell books, it only provides access to them for a limited time.
Has a large set of users with similar needs.	Has a large set of users with similar needs.
Each user cannot possibly use every single individual work in the repertory, and so makes a choice out of what is made available.	Each user cannot possibly use every single individual book in the library, and so makes a choice out of what is made available.
Offers the benefit of a large choice.	Offers the benefit of a large choice.
Charges a single licensing fee for access to the whole repertory.	Charges a single membership fee for access to the whole book collection.[5]
Price discrimination is standard among different users (discotheques, radio stations, gymnasiums, etc.)	Price discrimination is standard among different readers (children, senior citizens, etc.)
Users have an alternative method of gaining access – individual (direct) licensing.	Users have an alternative method of gaining access – purchase a book.
Digitalization is an important consideration and may affect how business is carried out.	Digitalization is an important consideration and may affect how business is carried out.

[5] While a CMO charges directly to users, many libraries are publicly funded, so users don't pay (or pay a heavily discounted fee).

There are, of course, significant differences as well. CMOs don't ever "own" the copyrights that they license, while libraries of course typically do own the copies of the books that they lend. CMOs typically don't price discriminate in terms of which works are used while libraries often charge additional fees for high-demand items. CMOs can grant licenses to many users consecutively, while libraries can only grant sequential licenses (this is changing, however, with e-books). Nevertheless, I do think that the similarities are sufficiently close for lessons to be learned from an economic theory of libraries. However, as far as I am aware, there has only been one attempt to write an economic theory of libraries; Varian (2000),[6] thus there is a very scant literature on this topic.

d. Risk sharing syndicates

An "optimal risk sharing problem" in economics is normally formulated as follows:

> Given an uncertain payoff X and m agents, divide X into (uncertain) shares x_i i=1,…,m, such that $\sum x_i = X$, with x_i being acceptable for agent i, i=1,…,m.

In most risk sharing problems, the individuals who will receive the shares of the aggregate payoff have each contributed something in the form of a membership fee for the right to share in X. The membership fees may be monetary payments, or contributions of risky assets or lotteries, or indeed any asset or service of value to the group. Of course, in most applications the aggregate payoff, X, is the result (or outcome) of the collection of the initial contributions of the m agents (which may, or may not, themselves be uncertain). One can think of the cases of insurance firms (where each individual contributes a loss lottery and a premium), mutual investment funds (where each individual contributes a deposit), or of course CMOs (where each individual contributes the rights to a risky copyrighted composition). The group of risk sharers is normally called a "syndicate".

It is quite clear from the definition given above that a CMO is indeed a syndicate, and therefore the economic theory of syndicates is entirely applicable. This is the general theory that will be explored in all that follows.

3. Contracts between collective (syndicate) members.

Having decided to pursue syndication and risk sharing as a model for CMOs, there are several important issues that are to be discussed. Firstly, we have the question of "efficient" risk sharing, which in effect determines the contract between the syndicate members. Second, there is the issue of the optimal size of the syndicate, for which we will appeal to the well-known Arrow-Lind theorem. Third, we should consider the effect (or not) of digitization upon the question of optimal risk sharing. And finally, we look at the syndicate as a representative agent for decision making.

a. Efficient "risk" sharing between syndicate members

A general outline of the problem of efficient risk sharing within a syndicate is presented in Gollier (2001), chapter 21. A Pareto efficient solution to such a problem is the requirement that the sharing rule, $x = (x_1,\ldots,x_m)$ should be such that it is impossible to alter x in such a way that at least one individual is made better off without making at least one other individual worse off. This is, of course, a most reasonable restriction to impose upon any sharing rule.

[6] See also Bakos et al. (1999).

Notice that the sharing rule x is precisely equivalent to a "contract" between the m agents, since it stipulates exactly how much each member of the syndicate will receive, given every feasible outcome of the aggregate variable X. The critical point, then, in the theory of syndicates is the determination of x, i.e. exactly how the social, or group, surplus should be shared between the members.

A fundamental result regarding the determination of the Pareto efficient sharing rule when each of the syndicate members is risk averse is the "mutuality principle", which states that, in order for a sharing rule to be efficient, in each state of nature each member of the mutual should receive an individual payoff that depends **only** on the level of the aggregate surplus to be shared in that state of nature. So if there are two states of nature in which the aggregate surplus is equal, then in one of those two states each member must get the same payoff as what he would get in the other of the two states. In essence, the mutuality principle says that the only thing that can cause mutual members to have risky final wealth levels is the existence of systemic aggregate risk. It also makes the clear point that if there is no aggregate risk, then it is efficient (and entirely possible) to have a sharing rule such that no individual suffers risk. On this point, Snow and Watt (2005) point out that the standard practice of CMOs (at least for musical works) of monitoring individual usage of works that are licenced as a part of a blanket license to the entire repertory and using that data to calculate the payoffs of the individual members is a clear violation of the mutuality principle. Simply put, basing member payoffs on individual usage of works is not Pareto efficient.

Assuming that a Pareto efficient sharing rule (or contract) is indeed sought, then as mentioned above the syndicate members can only have risky final wealth levels if there exists aggregate risk, that is, if X itself is risky. The question of how the syndicate should share this aggregate risk among its members was addressed in a very influential paper by Wilson (1968), who shows that the sensitivity of an individual's final wealth to the aggregate wealth should be inversely proportional to that individual's Arrow-Pratt degree of absolute risk aversion.[7]

Notice that under the mutuality principle (and the Wilson result) the efficient sharing rule does not tie down a single unique payoff for each member. There are as many Pareto efficient sharing rules as one would like. Therefore, the actual contracted sharing rule will still depend upon some sort of negotiation between the members. For example, if we simply add to the requirement of efficiency the very logical condition that the sharing rule satisfies a participation constraint for each member (each member gets expected utility that is at least as great by joining the syndicate under the given sharing rule than by not doing so), then the set of feasible sharing rules is reduced considerably. Considerations such as these might also invoke some guise of co-operative bargaining theory to determine the final sharing rule (for example, see Harsanyi 1963). Therefore, the final sharing rule that is contracted to among members of the CMO will presumably depend upon the outcome of a co-operative bargaining game.

While on the point of participation constraints for syndicate members, an important consideration is whether or not the sharing rule is coalition proof, that is, whether the sharing rule is stable in the sense that under the sharing rule that is in place there is no incentive for individuals to separate out into alternative CMOs (including the default CMO which is individual acting alone). Essentially, for each sharing rule we can expect that there is a resulting stable coalition size and composition, and this therefore is a fundamental determinant of CMO membership. Given that, a natural question to ask is

[7] The result is often expressed in terms of risk tolerance, which is the inverse of risk aversion. The sensitivity of an individual syndicate member's final wealth to the aggregate wealth should be proportional to that individual's risk tolerance.

what is the optimal membership of a CMO?[8] Once that question is answered, then the question of the sharing rule and the participation conditions that jointly achieve the optimal membership can be tackled.

A very obvious focal point for the sharing rule is a linear one, in which the income of syndicate member i in state j is a linear function of the value of X in state j, regardless of the value of aggregate income to be shared. That is; $x_{i,j} = c_i + a_i X_j$. Notice that in this formulation a_i is independent X_j, so the percentage share of aggregate income going to individual i is the same regardless of whether aggregate income is large or small. This, for example, is the sort of sharing rule that is common in partnerships, where profit between equal partners may be split (say) 50-50, regardless of how large profit actually is, rather than (say) 30-70 when profit is small, and 70-30 when it is large.

As it happens, the efficient sharing rule is linear whenever the syndicate members all have "equi-cautious" utility functions of the HARA (hyperbolic absolute risk aversion) class.[9] These are utility functions that display linear risk tolerance with the same slope parameter, and this class of functions encompasses as special cases most of the standard utility forms including constant relative risk aversion (power utility and logarithmic utility) and constant absolute risk aversion (negative exponential utility), among others. While this sounds quite technical, in fact it may not be so far-fetched. For example, if all syndicate members were to have constant relative risk aversion (a reasonably common finding in empirical studies), with equal coefficients of risk aversion (which may be the case if syndicate members are all individuals of similar characteristics), then the sharing rule should be linear.[10]

It is very important to clearly note that neither the mutuality principle nor the results mentioned regarding HARA class utility and linear contracts depend in any way at all upon an assumption of the individual items of repertory being independently and/or identically distributed. Thus, regardless of any possible interrelationships, spill-overs and externalities between the different songs in the repertory, it still holds that for the sharing rule to be Pareto optimal it must satisfy the mutuality principle. Likewise, regardless of inter-song externalities, if all members' utility functions are in the same equi-cautious HARA class, then the sharing rule should be linear.

b. The Arrow-Lind theorem and the optimal size of the syndicate

The Arrow-Lind theorem states that under certain assumptions, the social cost of a risk moves to zero as the population involved tends to infinity, so that projects can be evaluated on the basis of expected net benefit alone. More specifically, as the number of participants in a risky venture grows, the risk premium of a risk-averse participant (corresponding to a share of the syndicate income) decreases, and, in the limit, goes to zero. Moreover, the social risk premium also goes to zero.

Essentially, in as far as a CMO is concerned, the Arrow-Lind theorem works as follows. A CMO is a collection of individual works that form an aggregate repertory, which is what is licenced. There are two effects that happen due to aggregation in regards the risk of the licencing income generated. First,

[8] And, of course, "optimal" for whom? The existing literature points clearly to the result that what is optimal for the CMO members may not be optimal for wider society. If that is the case, there is a clear scope for governmental regulation of the sharing rule such that the outcome is a socially optimal CMO.

[9] This result was alluded to in Wilson (1968), but is more formally treated by (for example) Amershi and Stoeckenius (1983), Pratt (2000) and Gollier (2001).

[10] A closely related contractual issue is the way copyright income is shared between a copyright holder (or author) and the distributor (publisher). This problem is analysed by Alonso and Watt (2005), who show that the sorts of contracts that are common in reality in such relationships are unlikely to be efficient.

there is *risk pooling*, or in effect the ability for a bad outcome in one work to be compensated by a good outcome in another work. From the law of large numbers (and assuming that the income from the individual works are independently distributed), the more works that are collected and licensed together, the smaller will be the variance (risk) of the average aggregate repertory licencing income. For a sufficiently large repertory, it becomes virtually certain that the average income per work is equal to the expected value of each work separately.[11] The implication is that the more works are collected into the repertory, the lower is the per-work risk (and the cost of financing that risk) faced by the CMO. Since risk is costly and risk averse participants welcome reductions in risk (with no reduction in expected value), this effect then directly points to it being efficient for the repertory to be as large as possible, i.e. the CMO should welcome new members as much as they can.

Second, the CMO can also benefit from *risk spreading*, or the ability of each individual CMO member to participate in the risks of all of the other members. For the case of a CMO this effect is merged, to a certain degree, with the above law of large numbers effect since adding a member also adds a work to the repertory.[12] The risk premium of any given risk averse CMO member for taking on a zero mean risk with positive variance is proportional to that variance,[13] and under risk pooling the variance in question decreases with the number of members that are admitted to the syndicate. So not only do more members imply less risk for each, it also implies greater insurance opportunities and lower risk premiums.

To take the simplest possible example, assume that there are n members of a CMO, and each has a copyright whose market value is identically and independently distributed. Denote the random variable that is the value of each copyright by ω_i, where i=1,2,…,n. Assume that the expected value of each of the copyrights is μ, and each of their variances is σ^2. If the copyright holder were to market his song alone, without the CMO, his expected utility would be $Eu(\omega_i)$, where E is the expectations operator, and u is the utility function. In that case, the risk premium of this copyright holder acting alone is given by the r that solves $Eu(\omega_i)=u(\mu - r)$. On the other hand, assume that all n members join into a syndicate (a CMO), and market all of the copyrights as a single repertory, and then share the resulting aggregate income (which is still a random variable) equally. In that case, each member earns an amount of money equal to the average value of the sum of the individual random variables, which is $(1/n)\Sigma\omega_i$. Denote this average aggregate income by ϖ. From elementary statistics, the average of the sum of independent random variables is a random variable with expected value equal to the average of the individual expected values. That is, under the CMO, and given our assumption that all of the copyrights have the same expected value, the expected value of ϖ is just μ. On the other hand, the variance of ϖ is equal to $(1/n^2)n\sigma^2 = (1/n)\sigma^2$. So the risk premium for a copyright holder in such a CMO is given by the q that satisfies $Eu(\varpi)=u(\mu - q)$. Since random variables ω_i and ϖ both have the same expected value (μ), but ϖ has a lower variance (since n is strictly larger than 1), we know that q must be smaller than r.[14]

---

[11] In terms of what we have above, the law of large numbers implies that as the number of members becomes very large, the level of risk inherent in X is arbitrarily close to 0, even though each component of X has strictly positive risk.

[12] As opposed to, for example, the theory of public investment, where each "work" might be a different investment project, which are independent of the members of the mutual that makes the investment.

[13] Formally, the Arrow-Pratt approximation to the risk premium for a zero mean positive variance risk is equal to one half of the variance times the measure of absolute risk aversion at the income in absence of the risk.

[14] The risk premium decreases when expected value is held constant and variance decreases. Consider the Arrow-Pratt approximation to the risk premium, $(\sigma^2/2)A(\mu)$, where A is the absolute risk aversion function. For given μ this is smaller the smaller is σ^2.

More generally, think about what happens as members are added to the CMO. The expected value of the income each receives (the average of the sum of the values of their individual copyrights) is unchanged at μ, but the variance of their income decreases, since $(1/n)\sigma^2$ decreases with n. In fact, in the limit as n goes infinitely large, $(1/n)\sigma^2$ goes to 0. That is, by risk pooling and risk sharing, as the number of CMO members gets very large, they each get an expected income of μ with almost certainty.

Now, consider the value of the group's risk premium, $n \times r(n)$, as n grows. As n goes tends to infinity, $r(n)$ tends to 0. Thus, in principle, $n \times r(n)$ looks to be indeterminate. However, what the Arrow-Lind theorem does is to prove that in fact $r(n)$ tends to 0 at a faster rate than n tends to infinity. And so the product of the two tends to 0. That is, not only do the individual risk premiums of the CMO members all tend to 0 as the number of members rises, but so does the group's risk premium. As such, the Arrow-Lind theorem refers to the risk preference of the entire syndicate (the CMO). The implication of the theorem is that when a CMO has large enough membership it becomes virtually risk neutral, and so there is no need for the CMO to adjust its behaviour for risk. A necessary condition for this to hold is that for each member, the covariance of the repertory income with their own marginal utility absent a CMO is zero. This clearly links back to the mutuality principle, which states that for the CMO sharing rule to be efficient, each member's CMO income should be independent of their own contribution, which is essentially what they would have absent a CMO. So if the sharing rule respects the mutuality principle then indeed the covariance of the repertory income with the marginal utility of each member is 0 as required by the Arrow-Lind theorem.

Another important assumption for the Arrow-Lind theorem to hold true for sure is that the individual elements (songs) in the aggregate repertory should be independently distributed. That is, there should not be any interrelationships, spill-overs, or externalities of any kind among the different songs in the repertory. As soon as this does not hold, the validity of the result that the optimal membership should increase as much as possible is placed under doubt, and may not hold true. For example, if a new member is added with a song that decreases the market value of other songs (in the sense of a stochastic worsening of the associated probability density), then the additional risk-pooling and risk-spreading benefits of adding the song must be weighed against the negative effect upon the values of other songs. In such a case it can easily turn out that the optimal repertory size is strictly finite.[15]

c. The (non)effect of digitization

Under the transaction cost rationale for formation of a CMO, digitization might appear to have a significant impact upon the social value of aggregation of individual works into a single repertory under a CMO. This may indeed be so (see Katz, 2006), but it is hard to imagine that digitization also reduces the risks of each work individually. Indeed it is more likely to increase the riskiness of a given work, since under digitization, piracy becomes so much cheaper to organise. If that is so, then the greater is the digitization effect, the more members should be joined together in the CMO in order to achieve the same risk savings as under a less digitized environment. In that way, there is a strong argument for the rationale for CMO formation under a theory of risk-sharing syndicates to be largely immune to digitization effects.

[15] Of course the opposite might also hold, that is, a song that has positive externalities for the value of songs already in the repertory. In such a case the addition of the new member is doubly valuable. On the other hand, if there are any costs involved in incorporating and servicing new members, then again this might lead to a finite optimal CMO size.

Of course, digitization may also have other cost saving effects for the CMO, in terms of managing the huge amount of information required for their day-to-day business. However there is another particular effect that digitization cannot help for in an efficiently organized CMO, and that is the distribution of royalty income. As has been mentioned above, most (perhaps all) CMOs monitor the final use of each individual work in the repertory, and then use that information to calculate the royalty payments for each member. This is not efficient, as it violates the mutuality principle. So if, rather than bothering with monitoring, each member is paid a negotiated royalty amount that is independent of the actual use that is made of his/her particular works, there is absolutely no need to monitor individual title use, and so the new digital capabilities for doing so are largely redundant (at least for the purpose of royalty distribution).

In short, I find that the effects of digitization upon a CMO are diluted, so long as the CMO is organized efficiently. In particular, there is no reason to think that digitization might cause the optimal CMO membership to drop, and much less to drop to the extent that the CMO disappears. On the contrary, if digitization causes the risk inherent in managing copyrights to increase, then it should lead to larger, not smaller CMOs.

d. Adding a decision rule

In interesting situations (of which the CMO is a case), the uncertain aggregate payoff X depends also upon a decision, a, that is made by the group. For example, Wilson (1968) defines a syndicate to be "… a group of individual decision makers who must make a common decision under uncertainty, and who, as a result, will receive jointly a payoff to be shared among them." This appears to be a reasonable description of a CMO. In this case, one not only wants to determine the optimal sharing rule, but also the optimal decision to be taken. For example, the CMO must decide upon such things as the price at which a blanket license is offered, the efforts to be taken in monitoring and enforcement, and even the set of licenses that will be offered (i.e. a single blanket license to the entire repertory, or blanket licenses to several mutually exclusive sub-sets of the repertory). These decisions may affect both the expected value of X and its variance (and, of course its other higher statistical moments), which have a clear bearing upon the risk sharing problem. For example, depending on exactly what exactly is the sharing rule between members, a group decision that increases the expected value of X but that also increases its variance, may be a preferred course of action for some members and not for others.

Given that a collective decision will be made, it is important to consider how the syndicate represents the preferences of all of the syndicate members, that is, can the group as a whole be considered to be a "representative individual" for the members? This of course is an age-old problem, going back at least to the path-breaking work of Ken Arrow (1951). As it happens, the syndicate is indeed a representative individual if all of the syndicate members' utility functions are in the equi-cautious HARA class, which is the case in which the optimal sharing rule is a linear function of the aggregate surplus. Thus, in that case, whatever decision is made by the syndicate as a whole is also optimal for each of its members. Recall that this result does not depend in any way upon the individual elements of the repertory being independently distributed.

We will come back to the issue of decision making at an aggregate level in the next section, when we discuss the licences that a CMO typically offers to users (bundling into a single blanket license).

4. Contracts between the CMO syndicate and users

On the other side of the CMO market is the relationship between the CMO and the set of repertory users. There are at least two elements of the contract between the CMO and its users that are of interest. First, the price at which access to the CMO's repertory is obtained is an important aspect of the contract with users. Above all, it is important to consider if the CMO exerts monopoly power when setting its price, and if so, how this would impact upon social welfare. Normally the pricing aspect of CMO operations is subject to regulation, and so it is often not a perfectly freely contracted factor, but rather it is contracted under particular regulatory conditions.

Nevertheless, it is interesting to note that CMOs do very often engage in price discrimination according to observable features of different users. For example, music playing bars and gymnasiums might be charged according to square meter size which will typically be different for different establishments. This leads to different users paying different prices to access the same repertory in the same conditions.[16] Similarly, radio stations are often charged according to the amount of advertising revenue they earn. In those cases the royalty paid to the CMO amounts to a tax on revenue, which again typically leads to different radio stations paying different amounts to access the same repertory in the same general conditions. Also, in some countries radio stations are taxed at a different rate according to the fraction of their air-time that is actually dedicated to music (as opposed to other content, like talk). It is of course well known that price discrimination is a profit maximizing strategy, so it is not surprising that CMOs engage in this practice whenever possible.

The second element of the contract between the CMO and its users is the fact that often the only product that is offered to users is access to the aggregate repertory under a blanket license, rather than access to individual works, or smaller subsets of works. Just like the standard argument that a CMO should exist because, due to the existence of transaction costs, it is essentially a natural monopoly, the defence of the all-or-nothing offer of a blanket license is normally based upon the transaction costs savings that such a strategy implies. However, as we will note here, there are other reasons why a bundled repertory might be more efficient than an unbundled one.

The main question is whether the choice to offer only blanket licensing is efficient, or whether it would be better (either for the syndicate itself, or for the users, or both) to offer a set of licenses to sub-sets of the total repertory. The first point to recall is that, under the condition that the utility functions of each of the syndicate members is in the HARA class,[17] then the group utility represents the preferences of each member (i.e. the syndicate as a whole can be thought of as a representative individual). In such a case, the syndicate's decision to only offer a blanket license to the entire repertory is also the best decision for all members. Whether or not it is also optimal for users is quite another issue.

The literature on the issue of bundling is convincing (see, for example, Bakos and Brynjolfsson 1999) – bundling of information goods is both economically efficient and profit enhancing for copyright holders, because it reduces the heterogeneity in user willingness to pay for individual titles (a bundle allows a sale at the average, rather than the minimum, willingness to pay).[18] But, there is still scope for separate "specialist" bundles, subsets of the entire repertory, to be efficient so long as

[16] The idea seems to be that the intended unit of taxation is an end consumer – a member of a gym, or a patron in a bar. The size of a local is a proxy to measure the number of customers. However the fact remains that what is licensed is blanket access to a repertory (and not blanket access for each final customer), and so different users (bars, gyms, etc.) do pay different prices for what amounts to the same thing.
[17] In which case, as we discussed above, efficient sharing involves linear contracts.
[18] In parallel work, Bakos and Brynjolfsson (2000a) show how bundling can also have strategic value that can lead to greater profits.

differentiated bundles can be clearly identified and that there is sufficient demand for them from users. This would be the case, for example, if users have systematic differences in their valuations for different sub-sets of the repertory, in which case the CMO can do better by offering a menu of different sub-bundles of the entire repertory.[19] Notice that the theory of bundling relies upon the economies of aggregation, and so is not closely tied to the theory of transaction costs. Therefore, so long as the conditions for optimal bundling are present, it might be optimal to bundle even if transaction costs were negligible.

Interestingly, the underlying reason why bundling can be optimal for CMO repertories is remarkably similar to why CMO membership should optimally be left to grow very large (as we looked at above). Both effects rely upon the law of large numbers. For the issue of membership, if the individual surplus available from each work is drawn from the same distribution, then adding new members' works to the repertory implies that the variance of the aggregate surplus from the total repertory will reduce without reducing the average expected value. For the issue of bundling, the law of large numbers implies that the average valuation of users for a bundle of goods with valuations drawn from the same distribution will be increasingly concentrated near the mean valuation as more goods are added to the bundle. Thus, bundling allows for pricing at the average, rather than at the minimum, willingness to pay. We should also point out that, in the same way as for our argument for the optimality of increasing the membership of the CMO, the optimality of bundling relies upon there not being too many negative externalities between the different elements of the repertory. While additional songs aid in terms of the aggregation (price discrimination) effects, if they serve to reduce the value of existing songs, then it may well turn out that bundling them together can be less profitable.

Furthermore, as Bakos and Brynjolfsson (2000b) show, it happens that while bundling is more efficient the lower are the marginal costs of production, it is less efficient the lower are the marginal costs of delivery. So the move to digitization, which reduces both the marginal costs of production and of delivery has, in principle, an ambiguous effect upon the optimal bundle size.

5. Conclusion

This paper has considered the two most prominent contracting environments faced by copyright management organizations – contracting among the members of the CMO themselves to share in the surplus generated by the licensing of the CMO repertory, and contracting between the CMO and the users of its repertory. In both instances, we find that aggregation, and the economies of aggregation, are important for establishing that CMOs themselves should exist, and that they should offer a blanket license rather than disaggregated licenses. However these results are diluted, and possibly even invalidated, when there are significant negative externalities among the individual copyrights involved. It is interesting to note that the economies of aggregation are only partially dependent upon the sharing out of transaction costs, and so an interesting point of difference is found with the traditional literature on CMOs in that we find a rationale for existence of collective management and blanket licensing that is not particularly dependent upon the existence of transaction costs.

Perhaps the most significant issue debated in the paper is the issue of the contracts that are written between members of a CMO for the purpose of distributing the surplus from licensing the repertory

[19] Some of the results in Bakos and Brynjolfsson (1999) have been shown (by example) to be incorrect, and were indeed corrected, by Geng et al. (2005), who show that pure bundling is only optimal if the variance per unit of mean of the valuations of users is sufficiently small. Furthermore, if consumers' valuations for successive additional goods in the bundle do not decrease too quickly, then pure bundling is "approximately optimal", meaning that the profit difference between a pure discriminating monopolist and a pure bundler becomes very small.

among themselves. The paper has argued that a CMO is essentially a standard syndicate, and as such risk sharing among the members is primordial to understanding the contracts that should be in place for sharing CMO surplus. The main result to point out is that, in order for the sharing rule (or contracts) to be Pareto efficient, it must be the case that the payoff going to each member should be independent of the outcome of that member's particular copyrights. Thus monitoring licensing use, and paying members according to the use made of each copyright, as is done in the real-world, is Pareto inefficient. Second, linear sharing rules should be in place if (but not only if) all CMO members have utility functions that are of the same equi-cautious HARA class. While this may seem to be restrictive, since most CMO members may actually be similar in other characteristics (for example they are mainly authors and creators, and most are not particularly wealthy), it may well be the case that equi-cautious HARA utility functions is not such a bad assumption. Finally, under the same condition of equi-cautious HARA utility, it happens that the CMO as a unit can be thought of as a representative individual, in the sense that it has the same risk preferences as all of the members. In that case any decision that is optimal for the CMO, is also optimal for all of the members individually. One such decision is in regards pricing and bundling of the repertory.

References

Alonso, J. and R. Watt (2003), "Efficient Distribution of Copyright Income", chapter 5 (pp: 81-103) in Gordon, W. and R. Watt (Eds.), *The Economics of Copyright: Developments in Research and Analysis*, Cheltenham UK and Northampton Mass., Edward Elgar Publishers Ltd. (related problem of sharing of copyright income between creator and distributor)

Amershi, A.H. and J.H.W. Stoeckenius (1983), "The Theory of Syndicates and Linear Sharing Rules", *Econometrica*, 51(5); 1407-16.

Arrow, K.J. (1951), *Social Choice and Individual Values*, New York; John Wiley & Sons.

Bakos, Y. and E. Brynjolfsson (1999a), "Bundling Information Goods: Pricing, Profits and Efficiency", *Management Science*, 45(12); 1613-30.

Bakos, Y. and E. Brynjolfsson (2000a), "Bundling and Competition on the Internet", *Management Science*, 19(1); 63-82.

Bakos, J.Y. and Brynjolfsson, E. (2000b), "Aggregation and Disaggregation of Information Goods: Implications for Bundling, Site Licensing and Micropayment Systems", chapter 6 (pp: 114-37) in B. Kahin and H. Varian (Eds.), *Internet Publishing and Beyond: The Economics of Digital Information and Intellectual Property*, Cambridge MA; MIT Press.

Bakos, Y., E. Brynjolfsson and D. Lichtman (1999), "Shared Information Goods", *Journal of Law and Economics*, 42(1); 117-56.

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

Geng, X., M.B. Stinchcombe and A.B. Whinston (2005), "Bundling Information Goods of Decreasing Value", *Management Science*, 51(4); 662-7.

Gollier, C. (2001), *The Economics of Risk and Time*, Cambridge, MA; MIT Press.

Harsanyi, J.C. (1963), "A Simplified Bargaining Model for the n-Person Cooperative Game", *International Economic Review*, 4(2); 194-220.

Hollander, A. (1984), "Market Structure and Performance in Intellectual Property: The Case of Copyright Collectives", *International Journal of Industrial Organization*, 2(3); 199-216.

Katz, A. (2006), "The Potential Demise of Another Natural Monopoly: Rethinking the Collective Administration of Performing Rights", *Journal of Competition Law and Economics*, 2(2); 245-84.

Pratt, J. (2000), "Efficient Risk Sharing: The Last Frontier", *Management Science*, 46(12); 1545-53.

Snow, A. and R. Watt (2005), "Risk Sharing and the Distribution of Copyright Collective Income", chapter 2 (pp: 23-36) in Takeyama, L., W. Gordon and R. Towse (Eds.), *Developments in the Economics of Copyright: Research and Analysis*, Cheltenham UK and Northampton Mass., Edward Elgar Publishers Ltd.

Tirole, J. (1988), *The Theory of Industrial Organization*, Cambridge, MA; MIT Press.

Towse, R. (2012), "Economics of Copyright Collecting Societies and Digital Rights: Is There a Case for a Centralised Digital Copyright Exchange?", *Review of Economic Research on Copyright Issues*, 9(2); 3-30.

Towse, R. (2013), "The Economic Effects of Digitization on the Administration of Musical Copyrights", *Review of Economic Research on Copyright Issues*, 10(2); 55-67.

Varian, H. (2000), "Buying, Sharing and Renting Information Goods", *The Journal of Industrial Economics*, 48(4); 473-88.

Wilson, R.B. (1968), "The Theory of Syndicates", *Econometrica*, 36(1); 119-32.