Collective Management as a Business Strategy for Creators:

An Introduction to the Economics of Collective Management of Copyright and Related Rights

by Dr. Richard Watt

WIPO
WORLD
INTELLECTUAL PROPERTY
ORGANIZATION

TABLE OF CONTENTS

PREFACE

The present module contains an introduction to the economic theory surrounding collective management of copyright, for inclusion in educational material on collective management made available by the World Intellectual Property Organization (WIPO). The module is written with a non-academic audience in mind, but nevertheless it does set out the basic first principles of the economic theory involved. Above all, the module is designed to inform business managers as to the underlying economic rationales that support collective management of copyright as a beneficial business strategy, and to outline some of the implications of the theoretical background for the business decisions taken within collective management organizations in practice.

The outline of the module is as follows:

1. Introduction to the general economics of copyright

    1.1. Property rights theories

    1.2. The consequentialist theory, economic efficiency and incentives

    1.3. Economic welfare

2. Introduction to the general theory of collective management

    2.1. General theory and efficiency

    2.2. Transaction costs rationale

    2.3. Risk-sharing rationale

    2.4. Weighing the costs and benefits of collective management

3. Decision-making in copyright collectives

    3.1. A note on regulation, and the concept of natural monopoly

    3.2. Blanket licensing

    3.3. Optimal collective size

4. Conclusions

UNIT 1
INTRODUCTION TO THE GENERAL ECONOMICS OF COPYRIGHT

Economics is concerned above all with the efficient functioning of markets and the relationship between decision-making and incentives. In most analyses, at the forefront is some concept of the welfare of a set of economic agents,[1] and the objective is to consider how this welfare is affected by different decision-making environments. Therefore, the primary focus is upon the actual choices made by economic agents (consumers, producers, regulators, etc.) and the restrictions around the choices that are actually feasible to be made.

Quite generally, economics can be thought of as the study of how incentives can curb behavior, and given that one of the greatest incentive mechanisms that is present in all modern societies is the legal system – which gives us the general rules of acceptable human conduct and the possible repercussions if that conduct is violated – it is hardly surprising that economists have a strong tradition of studying the law. In particular, economists query what the final consequences of particular laws are, whether these consequences are those that were intended, and how the law can be altered so as to better address the intended consequences.[2] This general research agenda has come to be known as the subject area of 'law and economics', and within that subject area, the study of copyright law has held a prominent position.[3]

Economists have long understood that copyright is more than just a property right and can also be understood as an incentive mechanism. As such, it falls directly within their standard domain of analysis. Economists view copyright as something that should encourage the right amount of creative activities by the right people, and then allow the resulting works to flow through a market in an efficient way and end up being enjoyed by the consumers who most value them. The end result of such a system should be that the welfare of individual creators and consumers increases, and this should increase the general aggregate welfare in the economy ('social welfare'). It should be clearly noted, however, that copyright also involves certain costs that need to be balanced against the positive aspects that we expect to flow from the system. There are the costs of running the system in the first place (e.g. costs of policing for infringing activities and the costs of managing and administering the legal framework). And there is also the somewhat more subtle concept of

---

[1]   With possibly as few as one single agent in the set, but also possibly with many agents.

[2]   Economics has generally taken what is known as a 'consequentialist' approach to the law, in which copyright is justified (or criticised) in terms of desirable outcomes. The underlying assumption within this approach is that the law should be designed so as to maximise some logical metric of social well-being. The consequentialist approach conflicts with the 'deontological' approach, which often forms the foundation of legal scholarship, and which seeks to justify copyright, in particular, in terms of moral validity.

[3]   I will not be providing a survey here, as several excellent surveys already exist. The interested reader can consult, for example, Towse et al. (2008).

whether or not the evident monopoly power that the grant of copyright provides
to suppliers actually has the effect of decreasing consumer welfare more than it
increases supplier welfare (more on this below). Furthermore, copyright is often
argued to place a barrier in process of incremental creativity (where a new work is
based upon an already existing one), and thereby it may end up reducing rather than
increasing the overall amount of creative works that are *potentially* produced. The
economic gains in terms of additional welfare created need to be considered together
with any economic costs involved. The primary thesis that is staunchly defended by
most economists is that the benefits of having a copyright law outweigh the costs,
so long as that law is appropriately designed, and so the *net* social benefit from an
optimal copyright law is strictly positive.

Aside from some early articles (e.g. Plant, 1934; Hurt and Schuchman, 1966;
Breyer, 1970), the economics literature dedicated to the concept of copyright began
in earnest in the early to mid-1980s with a small series of papers in the *Journal
of Political Economy*, and since then it has grown to include hundreds of relevant
papers.[4]

1.1   Property rights theories

The development of the concept of copyright has moved in two parallel directions.
First, there is the very logical concept that whatever a person manages to produce,
using his or her own ingenuity and innate abilities, should naturally belong to that
same person. Thus, just like any other item that a person may construct (e.g. a house
or a piece of furniture), the outputs of an intellectual (rather than purely physical)
creative process (e.g. a story, a song, a poem, a drawing, etc.) should still be thought
of as being an item of property. The only real difference between the two types of
creative processes is that one brings into existence an item of tangible property,
while the output of the other is an item of intangible property. But both are items of
property, and just as there are laws that protect tangible property, so should there
be laws that protect intangible property from theft and general misuse by non-
owners. The relevant property right is, of course, known as copyright.[5] The concept of
copyright as a natural property right was originally championed by John Locke, and so
it often goes by the name of 'Lockean property right theory'.

---

[4]   Several surveys exist of specific areas within what has become known as the economics
of copyright. See, for example, Gordon and Bone (2000), Liebowitz and Watt (2007) and
Towse et al. (2008).

[5]   Of course, there is a legal difference between the acts of authorship, which generate
creations that contain valuable expression of some sort (protected by copyright) and
acts of innovation, which generate valuable ideas of some sort (which are protected by
patent), both of which are the fruit of intellectual efforts. Here, we associate with acts of
'creativity' only the authorship of copyrightable expression.

The second theoretical rationale for a property right in the outputs of intellectual creativity is that which is argued by economists (among others), and which is based upon the standard economic theory of incentives, and of costs and benefits. Interestingly, this perspective is captured in the concept of copyright in the Constitution of the United States of America, and it was also certainly in the minds of the original enactors of the early copyright statutes in Europe. In contrast to the Lockean natural rights theory, the emphasis of this perspective shifts to the economic value of the creative good, and how it can be best progressed along the value chain, from its production until it is in the end able to be consumed, thereby generating new welfare.

The general argument is often based upon the counter-factual as follows. If there were no enforced property right in the work, then it would be all the more challenging to be able to take the work to a market since the creator would have a much more difficult job to capture the value of the work from those who would like to benefit from having access to it (the consumers). So acting as a creator becomes less financially rewarding, and as a consequence we should expect that some potentially creative individuals would defect to alternative (more lucrative) uses of their time, with the result being a reduction in the number (and possibly in the quality as well) of creative goods that are available. In such a scenario, everyone who enjoys these sorts of works would be worse off because of having fewer of them to consume. Society in general would suffer, compared to a situation in which creativity flourished. Thus, while this argument supports the existence of a property right, it does so for reasons that are fundamentally different from those behind the natural rights theory. In short, the underlying idea is that the incentives that are created with the grant of copyright should lead to greater social welfare as creators create more (and perhaps better) outputs, and consumers are provided with a means of obtaining access to these outputs.

This second rationale for copyright is often known as the 'consequentialist theory', since it is concerned with the actual consequences (for example upon the behavior of economic agents) of the property right. We now go on to discuss the consequentialist theory in greater detail.

1.2 The consequentialist theory and the shape of copyright protection

The consequentialist theory as developed by economists studies the very justifications and purpose of copyright law in terms of the outcomes that are achieved from having copyright protection. Specifically, it is argued that copyright law is justified as a way in which creative individuals are able to receive payment in exchange for their efforts in creating useful works, and that without this promise of payment there is a risk that society would be worse off by the loss of potential useful creative works that would ensue as creative individuals defect to more attractive

employment options. Notice that there is an important implicit assumption in place here – society is assumed to gain more from those otherwise displaced individuals when they are creators of intellectual products than when they are engaged in their other alternative employments. Specifically, it is the job of an appropriately designed and managed copyright system to ensure that indeed only enough protection is granted so that only the more socially valuable creators remain as creators, and any other individuals who are more valuable in other jobs do indeed decide to defect from creative activities.

Of course, the consequentialist theory also addresses the flip-side arguments under which copyright can actually have a negative effect on creativity and welfare, through the effects of a possible creation of monopoly power. In so doing, the consequentialist theory conceives of a different-looking copyright law than might be expected under a pure theory of natural rights.

In a similar vein to the Lockean perspective, the consequentialist rationale for copyright also begins with the idea that the output (actually, only 'most' of the output – see below) of an act of intellectual creativity should belong in the first instance to the creator. Economists view ownership as a bundle of rights that determine what someone may do with some well-defined item of property – property may be possessed, used, transformed, sold, rented, destroyed and bequeathed (among a great many other options).[6] Above all, the owner has the right to exclude others from these same acts upon the property in question, and also the right to include others at a price (for example, the right to rent transfers the right to possession and use, but not the right to sell or destroy). The law is designed to protect the free exercise of such property rights from interference by others. Without the concept of ownership, meaningful transactions (which are the vehicle for realizing economic value) would be largely impossible.

But at this point the consequentialist theory of copyright envisages an extremely important difference from the Lockean rationale for copyright, and indeed from practically any other form of legally recognized property right. The theory proposes that the concept of ownership of the work should disappear after a set period of time has passed, and the property in question should then fall into the 'public domain'. The owner of the property right at that moment of time suddenly loses all rights in the work, and it thereafter belongs equally to each and every current and future individual or firm. Such a concept of social ownership of course destroys any market for the work itself, and so everyone is able to legally access all of the economic rights in the work in full at a price of 0.

---

[6]   Generally speaking, we differentiate between 'economic' rights and 'moral' rights. The former are rights from which some pure economic benefit can arise (e.g. a sale, or rental agreement), whereas the latter involves the right for the author to be always recognised as such. The consequentialist theory is only concerned with the economic rights.

It is pure economic efficiency that mandates that copyright should be time-limited. The copyright is granted so that, via the price paid, final consumers are able to contribute to the rewards that the creator receives in exchange for his/her creative efforts, and that this financial reward is sufficient for the creator to go ahead and create the work in question. The copyright allows the creator to gather income over time to compensate for the costs of creation suffered at the outset when the work was created. However, as soon as the expected present value of this income stream is sufficient for the creator to decide to go ahead with the creative process, then any further payments to the creator would be irrelevant for the objective of generating the work. Indeed, continuing with the payments beyond the limit of the copyright term would be counter-productive, since as soon as the copyright expires the price to access the work goes to 0, and consumption increases to its theoretical maximum with the corresponding gains in consumer welfare. So retaining protection beyond the limit date has the effect of artificially decreasing consumer welfare with no offsetting gain in the creative work (since it is assumed to exist already). Therefore, once the limit point has been reached there is no further need for the copyright to continue to exist, and so it should not continue to exist.

As mentioned above, not only does the property right expire after some specified passage of time, there are also further restrictions on the rights that are conferred to the creator (and any subsequent right-holder) right from the start. These restrictions are known as the limitations and exceptions to copyright,[7] and they normally amount to some of the original rights that are born along with the work being assigned in the first instance and for evermore to the public domain. Again, a property right with this type of structure does not happen in Lockean theory (or indeed in any other property right that I know of), and again a justification for such a structure can be found in economic theory; specifically the efficient functioning of markets and simple cost-benefit analysis. Without any limitations, cumulative creation could be impeded. So could education, since copyright-protected works are important as inputs to the production of education. The transaction costs of attempting to identify, and then to charge for, small snippets of copying could be enormous. The social costs of these sorts of things have been understood by law-makers to exceed the benefits that would accrue to creators should they be allowed to license them at a price, and so these sorts of access fall into the public domain right from the start.

Before leaving this discussion of the consequentialist theory, it is relevant to mention the debate concerning the purported negative effect that the grant of copyright has in terms of market power. Essentially, copyright protection erects a barrier to entry to a market, and this in turn allows the supplier of the work to increase its price above the competitive level. There is then a correspondingly lower level of consumption of the work than if the work were to be be priced competitively – the work is said to be

---

7   In many countries they come under 'fair use', or 'fair dealing' clauses in the copyright law.

'under-consumed'. It has been argued by some that this is inefficient, since there are welfare losses in comparison to a hypothetical competitive market for the work in which consumption is maximized.

However, the comparison between a copyright-protected market, together with its monopoly power and under-consumption, and a perfectly competitive market with maximum consumption is erroneous for at least two reasons.

First, the grant of copyright is in the end a property right, and thus it gives no more monopoly power than the property right in any tangible item. In the same way that it is not considered to be an exercise of monopoly power that I am allowed to restrict use of my private car or house to only those whom I want to admit, neither should it be considered an exercise of monopoly power that a copyright holder can restrict use of his/her works to only those whom he or she chooses to admit.

Second, in a perfectly competitive market the supplier (i.e. the creator of the work) would earn a profit of 0, and since there are costs of creation, a perfectly competitive market cannot support an economically rational creator embarking on the creative process. The work would not be created, and there would be nothing to consume. Even with under-consumption (for the duration of the copyright term), a copyright-protected market still generates something to consume, and therefore some welfare for consumers, and this should be compared not to the welfare in a competitive market in which the work exists, but rather to 0, which is the welfare if the work were not to exist, which is the scenario that would likely occur under no copyright protection.

1.3   Economic welfare

The primary concern of most economic analyses, whether explicitly or implicitly, is to focus on decision-making that enhances welfare. The concept of welfare has already been used rather liberally above. It is then important to think about what is meant by welfare, and the welfare of whom?

When economists talk of welfare, they are referring to some measure of happiness or well-being that accrues to a given set of economic agents. When that set of agents contains a single individual, then we normally use the term 'utility' to refer to welfare. Utility is thought of as a measure that changes with the inclusion or exclusion of goods, or consumption generally, and therefore it can be usefully referred to as a 'utility function'. The more goods that are consumed, the greater is the utility of the individual concerned. Utility is an ordinal, not a cardinal, concept, and so the units of measurement are irrelevant. All that is important is that should one 'basket' of consumption generate greater utility for a given consumer than another 'basket', then we can conclude the first 'basket' is preferred to the second by that consumer.

When the objective set of economic agents contains more than one individual, we require a measure of their collective utility, or 'social welfare'. In that regard, there are many possible ways to combine the individual utility amounts in order to calculate social utility. All that is really important is that the measure of social utility again be ordinal, in the sense that if the utility of any one individual rises, and that of all others stays the same, then the measure of social utility should rise. However, at this point we run into the issue of inter-personal comparisons, since the measure of social utility needs to consider the trade-off should one individual's utility increase at the same time as that of another individual decreases (for example, if some consumption is somehow transferred between the two individuals). To what extent is it worthwhile in terms of social welfare to reduce one person's utility in order to increase that of another?

It is precisely this trade-off that is at the forefront of regulatory activities,[8] as for example the consideration of an economically optimal copyright system. Adding additional (or in economist's jargon, 'marginal') protection, over and above the level that would ensure that the work is actually produced, would logically increase the utility of creators at the expense of a decrease in the utility of consumers. The corresponding effect on social welfare would depend on exactly what assumptions are used when individual utilities are combined into the social welfare function.

In many instances, economists skirt this rather important issue by restricting themselves to considering 'efficient' solutions (or allocations), rather than those that necessarily maximize social welfare. An efficient allocation is one in which no individual utility can be increased without decreasing the utility of at least one other individual.[9] If an allocation is not efficient in this sense, then by definition we can adjust things such that at least one individual is made better off (his/her utility is increased) without causing any negative impact upon any other individual, and such a re-allocation is a socially good thing (it would necessarily increase *any* measure of social welfare that increases with the utility of each of the individuals concerned, without regard to considerations of intra-personal utility comparisons in social welfare). Essentially, efficient allocations remove any unnecessary wastage.

How can we think about copyright in terms of social welfare? We would want to consider a comparison of social welfare both with and without copyright, at least in terms of the efficiency of the resulting allocations. In order to simplify down to only what is strictly necessary, imagine that there are only two individuals – one is a potential creator of a work, and the other is a consumer of that work (should it exist). Each has an individual utility function, and their combined social utility is

---

8　　Here, and in all that follows, references to the 'regulator' are supposed to imply the bodies that establish and protect the law, be they parliamentary (i.e. the State), or the courts of law (for example in common law countries).

9　　More precisely, economists call such allocations 'Pareto efficient', in reference to the famous economist Wilfredo Pareto.

assumed to be increasing in each of the individual values of utility separately (with no other assumptions on social utility needed). If there is no copyright protection, the creator decides not to create the work, and undertakes some other activity, and the consumer spends her money on other goods. Now assume that if the work could be brought into existence somehow, then the consumer would gladly divert some of her budget from the other consumption goods to instead gain access to the work. That is, there must exist prices for access to the work at which the consumer would be made strictly better off by reducing consumption of other goods and instead accessing the work. If it also happens that that amount of money (or less) is sufficient to entice the creator away from his alternative employment and into the activity of creating the work, then the creator would also be made better off by accepting the consumer's payment and creating the work. In such a situation, *both* members of the society (the creator and the consumer) are made better off, and social welfare unambiguously increases in comparison to the scenario in which the work did not exist. It is the job of copyright law to give effect the relevant transfer of funds from the consumer to the creator, thereby leading to social welfare improvement.

There are several things to notice here.

First, it is not clear that there does actually exist a price that both the consumer is willing to pay to get the work and the creator is willing to accept to create the work. If no such price exists, then society values the potential creator more in his alternative employment than in creative activities. This is the case of a low quality creator, and it is efficient that he/her refrains from spending his/her time creating a work of low value.

Second, if there does exist a price that is acceptable to both the consumer and the creator, then there could be many such prices, and each such price gives rise to a different but equally efficient solution. Quite generally, economists have taken the view that copyright law should be structured such that the *lowest* such price is used – that price which makes the creator indifferent between creating the good and not doing so. In that way the work is created, and the consumer gets the benefit of it.

Third, assuming that an acceptable price can be found, then we can unambiguously conclude that it is socially efficient for copyright to be put into place, even though we cannot explicitly calculate the actual value of social welfare (since we have not made any assumptions at all on exactly how the individual utilities are combined into social utility).

Finally, it is worthwhile to point out that the above discussion of how social welfare might improve with the creation and consumption of the work is really nothing more than the concept that purely voluntary exchange must be beneficial to society. If, by creating the work, both the creator and the consumer are better off, then it suffices for there to be a binding contract between them, voluntarily entered into, to ensure

that creation and consumption both occur. So why is copyright of any importance? The answer is that the real world is more complex than just one creator, one consumer and one contract. We only need to let there be more than one consumer to immediately recognize the importance of copyright. So say there are two consumers, and assume that each consumer still individually values the work at an amount that would be enough for the creator to create it. But of course if a consumer could access the work for free, he/she would be even better off than accessing it at a positive price. So the problem now is that if one consumer contracts to have the work created, there is a chance that the second consumer will be able to copy the work and consume it for free. Each consumer wants to be the second, and waits to see if the other consumer pays for the work. But in this waiting game neither consumer pays and the work is not created. This is a well-known coordination type of problem that goes by the name of the 'prisoner's dilemma'. It is copyright protection that comes to the rescue by forbidding the act of copying, and removing the option of consuming without paying. Under copyright protection, the creator is able to charge a price to each consumer who accesses the work, and thereby distribute the total cost of creation among all of the end users, to the mutual benefit of all concerned.

UNIT 2
INTRODUCTION TO THE GENERAL THEORY OF COLLECTIVE MANAGEMENT

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):

1.   They license to users access to the copyrights of their members and they collect royalty payments from users;

2.   They distribute the royalty income among the collective members;

3.   They monitor the use of the copyrights of their members, they enforce the legal copyright parameters, and they bring action against copyright infringements on behalf of their members.

Copyright collectives have not been without controversy. The aggregation of a great many individual copyrights into a single marketable repertory which is licensed as a single unit has been seen by some to be the creation of a monopoly, with unacceptable monopoly power. For exactly this reason, it has not been uncommon for the activities of copyright collectives to be subject to regulation.

2.1   General theory and efficiency

Whether or not a copyright collective is economically efficient depends upon a comparison of the outcomes that are achieved under collective management and under the next best alternative. To that end, let's assume that the next best alternative for each copyright holder is to attempt to license access to the work individually. In the end, the efficiency comparison is on two levels – first we have the idea from the previous section, that the expected present value of net licensing income that is received by the potential creator needs to be sufficient for the work to be created in the first place. And second, there is the actual level of costs involved in running the system. Let's circumvent the first problem and simply assume that, under both individual and collective licensing, enough licensing income is generated so that the work in question will be created and consumed. Then the only issue at stake is the cost that is suffered in total to get the work created, licensed and consumed. These costs may be actual financial outlays of money, and they may be non-financial,

like for example effort costs or the discomfort of maintaining a risky, rather than a certain, income stream.[10]

The costs involved in closing an economic transaction are known as 'transaction costs', and the greater they are, the lower is the resulting utility of the contracting party who suffers them. To be clear, transaction costs are not amounts of money paid by one party, for example the consumer, to the other, for example the creator (as is the case of licensing fees). Rather, they are amounts of money that are simply lost to the system, or paid to other economic agents outside of the environment of the exchange in question. The smaller we can make the transaction costs, the greater will be the resulting utility of one or both of the contracting parties, and that would directly yield a more efficient outcome.

The efficiency gains of collective copyright management as opposed to individual management are often largely (perhaps almost exclusively) centered upon transaction costs savings that can be achieved by collective management. If indeed it does hold that under collective management, the same (or more) works of the same (or greater) quality are created and consumed than would be the case under individual management, and that this is achieved with a lower level of costs, then collective management is more efficient than the alternative of individual management. We can state this in more concrete economic terms by noting that any transaction costs that must be paid are really just leakages of welfare from the system that consists of the economic actors involved in the production, distribution and consumption of the works in question. So when transaction costs are saved, then somewhere, someone is achieving a greater level of welfare (or individual utility), even though the same amount and quality of works is produced and consumed. Since it is not reasonable that a transaction cost saving by one economic agent within the system should lead to a reduction in welfare of any other economic agent within the system,[11] transaction cost savings are very clearly gains in welfare with no offsetting welfare reductions. At least one agent is better off, and none are worse off. And this is a 'Pareto improvement' (i.e., it is more efficient in aggregate welfare terms).

---

[10]　It is pervasive that economic agents are 'risk averse', in the sense that they would prefer a certain payoff (i.e. a given payoff amount received with certainty) to a risky one with the same expected value. The fact that individuals are willing to pay for insurance of the risks they face evidences these sorts of preferences, and shows that there is a measurable cost involved in keeping a risky, rather than certain, income stream.

[11]　Some transaction costs are non-financial losses – for example effort costs, or the costs of inefficiently high risk-bearing. These costs have absolutely no offsetting gain anywhere, within or without of the market for the copyright work. Other transaction costs may be financial, but paid to economic actors who are not really a part of the market in question – for example higher phone bill payments, or interest costs paid to a bank. If those costs are reduced then someone else in the wider economy is losing out, but we are only concerned here with the welfare of the participants in the market in question, and not the external economic agents.

We now go on to look in more detail at the transaction costs rationale for the efficiency gains of collective management of copyright over individual management. After that, we will consider the related topic of the superior risk-bearing opportunities that are afforded by collective rather than individual copyright management. In reality, the risk-bearing argument can also be thought of as a transaction cost argument, since risk is certainly costly, at least in terms of utility (or welfare). However, it is still reasonable to consider the risk-bearing rationale separately from the transaction cost rationale, simply because the standard transaction cost rationale in economics does not consider risk or risk aversion.

2.2   Transaction costs rationale

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 the following:

(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 transaction 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', which essentially means that the average cost diminishes with the size of the collective. In such an environment, it is *efficient* that licenses are granted

collectively rather than individually, where the efficiency derives from the same number of end users being supplied with what they want, but at a lower total cost.

However, 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 copyright collectives typically offer, they are forced to license all works, those they actually want and those that they do not want as well. 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.

One particular aspect of the transaction cost theory that is of current interest is whether or not digitization and the digital environment, which undoubtedly reduces transaction costs in many dimensions, is sufficient to destroy the natural monopoly aspect of copyright management (see Katz, 2006).[12] 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.[13] Indeed, collective management can also take good advantage of digitization to streamline their business, to the benefit of both copyright holders and users.

---

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

[13]   For example, the data on royalty collections through copyright collectives which is available on www.cisac.org shows world-wide year-on-year increases, in constant prices, of around 4 to 5 percent.

2.3   Risk-sharing rationale

The input suppliers for the general business model of a copyright collective are the
individual copyright holders, and they are thought of as the owners (like shareholders)
of the collective. In this way, a copyright collective is really a type of mutual firm, or
what is known in the economics literature as a 'syndicate'. The main focus of the
economic theory of syndicates has been on risk-sharing, and indeed the general
theory of risk-sharing is based upon the contractual relationships between the
members of a syndicate. Concretely, an 'optimal risk-sharing problem' in economics is
normally formulated as follows:

> Given an uncertain payoff X and m agents, divide X into m (possibly uncertain)
> shares, one share for each agent, such that the sum of the shares[14] is equal
> to X, and such that each agent's payoff is acceptable for him/her (i.e. it gives
> that agent a level of utility that is at least as great as what would be achieved
> otherwise).

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 copyright
management organizations (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 copyright collective is indeed a
syndicate, and therefore the economic theory of syndicates is entirely applicable. For
the case of a copyright collective, X is the total net royalty income gathered by the
collective, and the shares are the distribution of royalties to each individual member of
the collective.

There exists a sizeable economic literature on the contracting environment for the
members of a syndicate (i.e. the way that the aggregate surplus should be shared
among them, that is, the determination of the shares x_i), which we will not delve
into here.[15] Suffice it to say that a Pareto-efficient solution to such a problem is the

---

[14]   When dealing with random, or uncertain, variables by 'sum' we mean the sum of the
        shares in each different state of nature, i.e. for each different outcome of the aggregate
        value X.
[15]   A general outline of the problem of efficient risk-sharing within a syndicate is presented in
        Gollier (2001), chapter 21. For a simple model of efficient risk-sharing of copyright royalty
        income, see Alonso and Watt (2003).

requirement that the sharing rule, that is the set of all of the shares, should be such that it is impossible to alter it 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. A little more generally, we can state that there will be many Pareto-efficient sharing rules, and the syndicate members will likely undertake some sort of bargaining process to determine exactly which one ensues.

A copyright collective, or copyright management organization, is a collection of individual works that form an aggregate repertory, which is what is licensed. As such, aggregation is an important element in play, and it turns out that aggregation is also the major determinant in the level of risk that the collective must deal with. There are two effects that happen due to aggregation regarding the risk of the licensing income generated.

First, 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 is independently distributed), the more works that are collected and licensed together (i.e. the larger is the repertory), the smaller will be the variance (risk) of the average aggregate repertory licensing income. For a sufficiently large repertory, it becomes virtually certain that the average income-per-work is equal to the expected income-per-work.[16] 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 organization. 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. if the collective would like to reduce the risk suffered by each of the members as much as possible, then the collective should welcome new members as much as they can.

Second, a copyright collective can also benefit from *risk-spreading*, or the ability of each individual member to participate in the risks of all of the other members. For the case of a copyright collective this effect is confounded, to a certain degree, with the above law of large numbers effect, since adding a member also adds a work, along with its risk, to the repertory.[17] Nevertheless, since a larger repertory implies a smaller aggregate risk, and a greater membership base, there are both greater insurance opportunities among the members, and less total risk for these insurance opportunities to have to deal with.

---

16 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.

17 As opposed to, for example, the theory of public investment, where each 'work' might be a different investment project, independent of the members of the mutual that makes the investment.

Whether or not copyright collectives do actually take full advantage of the risk-pooling and risk-sharing opportunities that are available to them depends entirely upon how the sharing rule, under which total repertory licensing income is distributed, is determined.[18]

2.4   Weighing the costs and benefits of collective management

As we have seen, copyright collectives, that administer and manage a great many copyrights all together in a single repertory, have several beneficial features for general efficiency and welfare. But they also harbor certain possible downsides, which need to be weighed against the benefits. Here we will give a quick overview of the principal such trade-offs.

First, as already mentioned in regard to the very grant of copyright itself, a copyright collective is a legal monopoly, and as such it possesses monopoly power in its dealings with consumers who would like to access the works in the repertory. Monopoly power generally ends up in a higher than otherwise price, less consumption and correspondingly a lower level of welfare for consumers. However, this argument against collective management is largely the same straw man as the argument that the grant of copyright itself delivers monopoly power, since it may very well be that only by allowing the greater profit that can be earned under collective management will the supply side of the market (creators, publishers, distributors, etc.) find sufficient financial incentive to function. That is, it may well be that individual management is only financially viable for a tiny minority of works – those that are immensely popular, for example – and so without the option of the additional benefits that collective management offers (transaction cost savings and risk-sharing), a huge number of works would never come to market, and consumer welfare would be far less than if they were available but under a certain degree of monopoly control. The exact extent to which this is true depends upon exactly how significant the benefits are from collective management for each (potential) work in the supply side when compared to individual management. However, suffice it to say that it does seem to be largely agreed by economists that without collective management, the number of works that would be available for use would be very significantly reduced. This, of course, is evidenced by the fact that all over the world governments *have* decided to allow copyright collectives to form and to operate, even if it is just as often understood that their actions should be regulated.

Second, digitization might appear to have a significant impact upon the social value of aggregation of individual works into a single repertory, by reducing the transaction costs involved in individual management (or perhaps at least the formation of many collectives with much smaller, perhaps more specialized, repertories). It has been

---

[18]   On this point, see Snow and Watt (2005).

argued that this aspect could be sufficient to offset the benefits from collective management (see Katz, 2006). However, 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 organize. If that is so, then the greater is the digitization effect, the more members should be joined together 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 collective formation under a theory of risk-sharing syndicates to be largely immune to digitization effects.

In short, when one weighs the relative costs and benefits of collective management of copyright, essentially one is drawn to attempting to provide an answer to the question of which is the most economically efficient collective size. The smaller is the economically efficient collective size (with a limit at 1 work, i.e. individual management), then the less efficient is pure collective management, i.e. the costs of collective management would outweigh the benefits. But, on the other hand, if the economically efficient collective size is very large, perhaps even to the extent that there should only be a single collective comprising all works,[19] then we can conclude that the benefits of collective management outweighs the costs. Note that this is a query on the optimal size of a collective, something that will be considered further on in this module.

---

[19] Of a similar type, of course – music, or literary, or visual arts, etc.

UNIT 3
DECISION-MAKING IN COPYRIGHT COLLECTIVES

3.1   A note on regulation, and the concept of natural monopoly

The efficiency of collective management, essentially, derives from the fact that
average costs (inclusive of transaction costs, and any other non-financial costs such
as the costs implied by risk) decrease as more and more works are added to the
collective's repertory. The average costs are the total costs required to license the
repertory divided by the number of works in the repertory. It is the cost-per-work in
the repertory. When this average cost is decreasing, it is more efficient that there be
a single licensing body offering all of the works together, rather than many small ones
each offering a small sub-set of works. In such a case, we have what economists
term a 'natural monopoly', which is a scenario in which the most efficient manner
(in terms of costs) in which consumers can be supplied is by a monopoly supplier. In
these cases there is a strong argument in favor of allowing collective management to
operate as a monopolist.

However, there is also the danger of allowing a monopoly to operate but then having
that monopoly exert excessive power over consumers in terms of how it operates
– above all its decisions on membership (i.e. repertory size), pricing and the specific
bundles of items that are offered to consumers. In order to get the best of both
worlds (efficient supply under a monopoly, but non-abusive activities), it is often
the case that copyright collectives are allowed to form, but are regulated in terms
of exactly how they run their business. It is then the job of the regulators to set
the limits of such elements as the pricing arrangements offered by the collective,
the membership rules that the collective uses, and perhaps even the sharing rules
used to distribute aggregate collective income among members (which in turn will
have an effect on membership), in such a way that consumer welfare is maximized
but subject to there being sufficient incentive for the supply side (from creators, to
publishers, to distributors) to offer and make available an appropriate bundle of works.
And then, once the limits to activities are set, the collective should be allowed to
make any decisions that fit within those regulatory limits.

3.2   Blanket licensing

Aside from the actual pricing of access to works, perhaps the most crucial element of
the day-to-day operations of a copyright collective is the decision to only offer a blanket
license to the entire repertory, rather than offering a myriad of different products, each
containing a sub-set of the repertory (perhaps right down to single works). However,
as has already been discussed above, it is somewhat strange that blanket licensing in a
copyright collective is controversial, when exactly the same practice is so commonplace
in other industries (from public transport, to gymnasium memberships, to newspaper

publishing). It is also the case that the only way that the natural monopoly efficiencies can be fully exploited is by using a blanket licensing arrangement. If smaller sub-sets of the repertory were to be offered, perhaps with each user determining exactly which works are to be included in his or her particular subscription, then the transaction cost savings that are implied by collective management are lost, since each individual subscription package needs to be negotiated separately and, above all, monitored for actual use. This would undoubtedly increase the price that would need to be charged to users, likely to an extent that they would gain access to only the sub-set of works that they require, but they would pay a price that is greater than had they purchased a license to the entire repertory. In short, the best decision that a collective can make as regards the menu of products offered to users is to continue with a blanket license to the entire repertory.

3.3   Optimal collective size

In terms of economic efficiency, as long as the natural monopoly element is present (average costs decreasing with repertory size), the optimal collective size is potentially unlimitedly large, although of course this depends on precisely what we mean by 'optimal' (more on this below). Under decreasing average costs, the more works are added to the repertory, the greater are the total efficiency gains that become possible when compared to supplying those same works using several smaller repertories. Similarly, with respect to the gains in risk-pooling and sharing, the more works that are added to a blanket licensed repertory, the less risky is the average of repertory income with no offsetting reduction in the expected value of the average income. That is, with more works, there is the clear opportunity to offer each copyright owner a better personal income stream from the point of view of risk-bearing. Both of these effects imply that the larger is the repertory, the greater can be the collective welfare gains, and these welfare gains can potentially be distributed across *both* the members of the collective (copyright holders) *and* the users of the collective's repertory in such a way that some (or all) are made better off, with no-one made worse off.

However, the above is only true under certain assumptions regarding the relationships between the works that make up the repertory, and (as mentioned above) the particular objective that the collective has, which defines what we should understand by 'optimal' collective size. For a start, regarding the transaction costs rationale, ideally each work would need to be more or less equal in terms of quality (i.e. they each offer the same utility to users), and they would all need to have essentially identical costs if licensed individually. As soon as there are some works that are more popular, or others that are for some reason more costly to move through the value-chain, it may not be true that adding works is always beneficial. It might be best to selectively leave certain works out of the repertory. Second, for the risk-bearing benefits to take place reliably, it is necessary that the demand for each work is independent of the demand for all other works. So if the demand for songs by, say, Michael Jackson increases,

this effect would have to have no implication for the demand for songs of, say, Bruce Springsteen. Where there are interrelationships between the values of different works by users, it is not necessarily true that total risk is reduced by aggregating works into a single repertory. Again, it might be best to leave some particular works out.

Nevertheless, in order for the benefits of aggregation of works into a single repertory to be completely offset, or even outweighed, by these factors, the asymmetries and interdependencies over works would need to be considerably more extreme than is evident in the real-world situation, at least for the case of musical compositions. The only meaning that we should attach to the real-world fact that not all songs are equal and that there are likely some complementarities and possibilities of substitution between different songs, is that the purely theoretical efficiency gains from aggregation are not maximized. The presence of asymmetric and interdependent works dilutes, but does not eliminate, the efficiency gains from aggregation. The net gains are still present, and they are still positive.

Perhaps the most important concept to understand when we discuss the optimal collective size is the actual objective that the collective sets itself. In terms of the economic objectives of the collective, it is important to always recall that a copyright collective is a group of individual copyright holders. Thus, it does not follow naturally that the collective is interested in maximizing its total profit, as would be the case for most businesses. Take the very simplified case in which a copyright collective is made up of a group of identical copyright holders (i.e. each member contributes a copyright of identical economic value to the repertory). In such a case, there is no logical reason why the sharing rule for distributing net royalty income should discriminate among the members, and so each member should receive a payment equal to the average profit of the collective. It then follows that what the collective should aim to maximize is average profit, and not total profit. In essence, if it came to a vote among members as to what the collective's objective should be, there would be a unanimous vote for maximizing average profit, which is what each member individually receives as payment. This objective does have important effects upon the optimal membership size of the collective.

A simplified general theory of an average profit-maximizing collective (in absence of any risk-sharing considerations)[20] is given in Watt (2014), and is summarized here in Figure 1. In short, maximizing average profit leads to a strictly finite optimal collective size, n^*, that is also strictly smaller than the size that would be socially optimal, n^s.[21] It also holds that if the regulator is able to impose a rule that new members cannot be refused entry to the collective, then the collective will operate at a membership level,
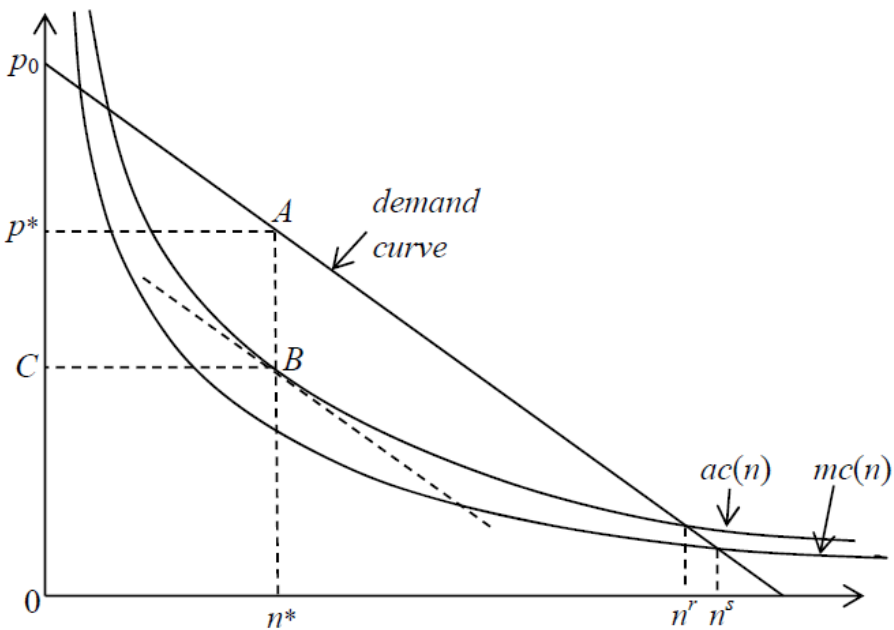
---

[20] Taking into account risk-sharing benefits as well, and risky individual song values, the outcome would be that the optimal collective size would increase over and above the size under certainty, but it would still be finite if each member has a finite level of risk aversion.

[21] Under quite reasonable assumptions, it is also strictly smaller than the repertory size that would maximize total, rather than average, profit.

n^r, that is somewhere between that which is optimal for the collective itself, and that which is socially optimal.

Figure 1 works as follows. For simplicity we assume a linear market 'demand' curve for a repertory of size n (that is, a repertory with n songs in it). This curve is based on the following. A repertory of size n can be licensed for a total payment from the market of $p{\times}n$, but where the effective price-per-unit of repertory, p, is assumed to decrease the larger is n.[22] The 'demand' curve then shows the graph of p as n increases. The height of the demand curve at any given level n indicates the price-per-song in the repertory that the *marginal*, or the lowest valued demander of those that consume, is willing to pay, and this is the price that is paid by all users, assuming that the collective cannot price-discriminate among users. Take for example a repertory of size n^*. The effective price-per-song that such a repertory can be licensed for is p^*, and the total licensing fees earned from users who access the repertory is given by the product $n^*{\times}p^*$, which is the area of the rectangle with corners p^*, A, n^* and 0. Also shown in the figure are the average cost curve for the collective, $ac(n)$, and the marginal cost curve, $mc(n)$. Both of these are downward sloping curves, with $ac(n)$ everywhere above $mc(n)$. These features capture the assumption of decreasing average costs, which is the natural monopoly element of the collective – the larger is the collective's repertory, the smaller are the average costs of licensing it.

Figure 1: The regulated optimum compared to the unregulated and the social optima

---

[22]   The price p reflects the average value of songs in the repertory. As the repertory grows, it is likely that less use is made of each one as there is a clear constraint on users' time that makes the songs compete with each other for the attention of users. Thus it is reasonable that the average value-per-song falls the greater is the number of songs included.

At any given repertory size n, the vertical distance between the average cost curve and the demand curve gives us the level of average profit at that repertory size, and this is what an unregulated collective is assumed to maximize. In Figure 1, this happens at the repertory size n^*, which is found where the slope of the average cost curve equals the slope of the demand curve. In this unregulated solution, the total licensing income (the area of the rectangle with corners p^*, A, n^* and 0) is split between the costs of operating the collective at that repertory level (the area of the rectangle with corners C, B, n^* and 0), and the total (net) profit that can be distributed among the collective members (the area of the rectangle with corners p^*, A, B and C). The users (the consumers of the repertory) receive a strictly positive amount of 'consumer surplus' equal to the area of the triangle with corners p_0, A, and p^*.[23] This consumer surplus measures the net benefit that is enjoyed by the users, who are (as a group) asked to pay less to access the repertory than what they are in fact willing to pay.

If the collective is unregulated, and under the assumption that each member of the collective (and we are assuming that each member is identified by exactly one song in the repertory) receives an equal share of the net profit, then the collective will stop admitting members once the repertory has reached the size n^*. Adding members beyond that repertory size would have the effect of decreasing the average profit, which means decreasing the amount of money that is distributed to each member. Thus the members would unanimously vote not to increase membership beyond n^*.

On the other hand, imagine that the objective were to maximize social utility, or the sum of profits and consumer surplus. The membership that achieves this is known as the 'social optimum', and it is given by n^s in Figure 1. Recall that the height of the demand curve tells us the effective willingness to pay per-song of the marginal consumer for a repertory of size n, and it happens that the height of the marginal cost curve, $mc(n)$, tells us the cost of adding a new song to the repertory. So long as the marginal consumer is willing to pay more to have a unit added to the repertory than what it costs to add that unit, it is socially valuable to increase the repertory by that one unit. However, as the size of the repertory gets larger, the gap between the marginal willingness to pay (the height of the demand curve) and the marginal cost (the height of the $mc(n)$ curve) gets smaller and smaller, and it goes to 0 at the repertory size n^s. At this socially optimal repertory level, notice that the average cost curve has gone above the demand curve, which indicates that average profit is now negative. Thus in order to achieve the socially optimal repertory level, it would be

---

23 Actually, consumer surplus is likely greater than the triangle indicated. The triangle indicates the consumer surplus that the marginal consumer would earn from a repertory of size n^* multiplied by the total number of users. If all consumers had equal values for repertory (so all were effectively marginal), then the triangle measures consumer surplus accurately. But if there are some consumers with higher values, then consumer surplus would exceed the triangle indicated.

necessary for there to be some cross-subsidization scheme from consumers (who at n^s earn a large amount of consumer surplus) to the collective members. This could, for example, be worked through a subsidy payment to the collective from taxation revenue.

Finally, if we assume that cross-subsidization is not possible (perhaps for political reasons), what is the optimal repertory size from the point of view of a social regulator (one who is interested in maximizing social welfare)? The answer is that the regulator would want to set the membership at n^r, which is exactly where average profit goes to 0.[24] This is the closest repertory size to the socially optimal repertory size, but that does not imply negative profits for the collective.

In passing, we note two important things. First, the actual placements of the three repertory sizes n^*, n^r and n^s in Figure 1, in terms of their relative sizes, is totally arbitrary. Nothing at all is implied in the graph as to the relative values of these three repertory sizes. All we know is that n^* is smaller than n^r which in turn is smaller than n^s. But we really do not know how close n^r is to either of the two extremes. Nor do we know how large these numbers are in comparison to 0. Thus it is entirely feasible that moving from, say, n^* to n^r could imply a small (say, 1 or 2%) increase in repertory size. Second, notice that at very small repertory sizes, the average cost curve and the marginal cost curve both go above the demand curve. Thus, at these repertory sizes again the collective would earn negative profits, and indeed it would not even be socially optimal to supply the users. With this in mind, imagine that collective management were for some reason not allowed, and that the only option for each author is individual licensing. This is exactly the same as saying that songs are licensed in a series of 'collectives' with membership equal to 1 single repertory element. If it happens that such repertories lie in the socially ruinous area of the graph, this would indicate that banning collective management of copyright, and thereby obliging any licensing to be individual only, would end up in no licensing at all being optimal – that is, the market for licensing music would disappear entirely. Thereby, we can see that collective management is indeed a socially optimal institution.[25]

---

24   Average profit of 0 is only a reference point. What would be implied is that the value to the marginal collective member (average profit once he/she is in the collective) is exactly equal to the value of not being in the collective (what economists call the 'opportunity cost'). This implies that the net benefit of being in the collective is 0, or that the marginal member is indifferent between joining or not.

25   This can also been seen somewhat more directly. The socially optimal repertory size is very large. In fact, right up to *ns* social welfare is increasing in repertory size. Thus, until ns is reached, the more songs that are licensed together as a single repertory the better. In other words, collective management is certainly socially optimal.

UNIT 4
CONCLUDING COMMENTS

In this module, we have discussed the general economics of the grant of copyright, and the economic rationale for collective management of copyright-protected works. The underlying reasons why both of these institutions are economically efficient are essentially the same – they allow a mechanism to operate under which users of copyright-protected works can contribute financially to the costs of creation and distribution of these works, thereby allowing them to exist in efficient quantities (and qualities). Copyright protection is required to allow market transactions to take place, without fear of free-riding happening to an extent that creators could no longer expect to capture a fair reward for their efforts. And collective management is an economically efficient organization of the supply side of the market, since it allows a given set of works to be licensed to users at the lowest possible cost (including the cost of risk). Both institutions (copyright protection, and collective management) are, and should be, regulated in order that the outcomes achieved are indeed in line with some consideration of social welfare.

With particular reference to collective management, the traditional rationale for the efficiency of blanket licensing by a single supplier is based on transaction cost savings, and a resulting theory of natural monopoly. However we have also seen that blanket licensing has additional efficiency benefits that derive from allowing superior risk allocations to take place. Since risk-aversion is certainly present in copyright holders (and creators themselves), and since copyright-protected works are inherently risky, it is economically efficient that the supply side of the market be organized in such a way that risk-pooling and risk-sharing opportunities, under which the risk that each copyright holder suffers is reduced (without need to reduce their expected payoff), can be fully taken advantage of. Collective management is a very efficient way in which such opportunities can be offered.

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.

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

Breyer, S. (1970), 'The Uneasy Case for Copyright: A Study of Copyright in Books, Photocopies, and Computer Programs', *Harvard Law Review*, 84; 281-351.

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

Gordon, W. and R. Bone (2000), 'Copyright', in Bouckaert, B. and G. De Geest (eds.), *Encyclopedia of Law & Economics, Vol. 2: Civil Law and Economics*, Cheltenham, UK: Edward Elgar; pp. 189-215.

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

Hurt, R. and R. Schuchman (1966), 'The Economic Rationale of Copyright', *American Economic Review*, May; 421-32.

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.

Liebowitz, S. and R. Watt (2007), 'How to Best Ensure Remuneration for Creators in the Market for Music? Copyright and its Alternatives', *Journal of Economic Surveys*, 20(4); 513-45.

Plant, A. (1934), 'The Economic Aspects of Copyright in Books', *Economica*, 1; 167-95.

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.

Towse, R., C. Handke and P. Stepan (2008), 'The Economics of Copyright Law: A Stocktake of the Literature', *Review of Economic Research on Copyright Issues*, 5(1); 1-22.

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.

Watt, R. (2014), 'Copyright Collectives: Some Basic Economic Theory', Chapter 10 (pp. 167-78) in Watt, R. (ed.), *Handbook on the Economics of Copyright: A Guide for Students and Teachers*, Cheltenham UK and Northampton Mass.: Edward Elgar Publishing Ltd.

About the author

Richard Watt is Associate Professor of Economics, and Head of Department, of the Department of Economics and Finance at the University of Canterbury in New Zealand. His research interests span the general field of applied microeconomic theory, with particular emphasis on the economics of copyright and the economics of risk and insurance. In 2001 he founded the Society for Economic Research on Copyright Issues (SERCI) an international group of academics interested in the economic fundamentals of copyright. SERCI holds an annual congress which shifts across the globe each year, and it publishes a peer-reviewed journal on the topic of the economics of copyright.[26] Dr. Watt is currently the General Secretary of SERCI.

---

[26]   *The Review of Economic Research on Copyright Issues* (RERCI). See www.serci.org/rerci.htm