Review of Economic Research on Copyright Issues, 2018, vol. 15(2), pp. 57-79

INDIRECT COPYRIGHT INFRINGEMENT LIABILITY FOR AN ISP: AN APPLICATION OF THE THEORY OF THE ECONOMICS OF CONTRACTS UNDER ASYMMETRIC INFORMATION

RICHARD WATT AND FRANK MUELLER-LANGER

Abstract. Under current copyright law in many countries, Internet Service Providers (ISPs) can be found liable for the traffic on the websites that they host. While the ISPs themselves are not undertaking acts that infringe copyright, indirect liability asserts that they either contribute to, or encourage in some way, infringing activities, and thus they are liable to claims of indirect involvement by the affected copyright holders. The present paper explores indirect liability in a standard principal-agent setting, where both moral hazard (the act of monitoring) and adverse selection (differential costs of monitoring over ISPs) are present. The model considers the kinds of contracts that could be signed between the copyright holders (acting through a collective) and the ISPs (acting individually). We specify the contracts that are self-selecting and incentive compatible for the set of feasible scenarios.

1. Introduction

Under the current copyright laws of many countries, under certain conditions Internet Service Providers (ISPs)[1] can be found liable for the actions of their clients (firms and individuals that own websites), and even for the actions of the clients of their clients (firms and individuals that use those websites).[2] When an ISP hosts a website that is then used for downloading/uploading copyrighted content illegally, the affected copyright holder may have recourse against the ISP, which is often a much less costly legal course of action than attempting to identify and then sue the individuals or businesses that actually carry out the

The authors have benefitted from the insightful comments of an anonymous reviewer at RERCI, which have greatly improved the paper.

[1]In this paper we use the term ISP rather loosely. It refers to any business that provides a service which allows customers to upload or download data to and from the internet. Thus, it could be a telecommunications company providing general access to the internet, including webpage hosting. Or it could refer to a particular website that allows others to upload and download content, like for example YouTube, or the now infamous megaupload.com.

[2]See Dinwoodie (2017) for a comprehensive overview of the different standards adopted in different countries for establishing liability of ISPs. For instance, he indicates that US and UK courts have typically adopted standards that establish liability based on contribution or inducement. In addition, he suggests that, in France, ISPs can be held liable under the general fault provision in the Civil Code, i.e., when they disregard their duty to act reasonably.

illegal activity on the website. It is often argued (see, for example, Landes and Lichtman 2003) that secondary liability of this type is an efficient solution, as the threat of liability performs the role of an incentive for the ISP to monitor and control for the traffic and content on the websites that it hosts. Assuming that the ISP itself is better placed than is any third party (e.g. the copyright holder, or the copyright collective) to carry out this monitoring, then indirect liability may be an efficient mechanism for avoiding misuse of copyright content.[3]

But this assumes, perhaps rather blithely, that an ISP does actually have a cost-efficient monitoring mechanism at its disposal. If it happens that effective monitoring is too expensive even for the ISP, then indirect liability may still not be enough to give the ISP any incentive to monitor. In such a situation, it becomes debateable that indirect liability should even exist, or in other words, showing that monitoring is overly costly for the ISP should be a valid defence against any claims of indirect liability (e.g. when the only relationship is at the point of sale), as could be, for example, the case of a safe harbour provision.

Indirect liability will generally rest on a claim from the plaintiff that an ISP did not take all of the actions that it could have to avoid the copyright infringements.[4] Notice that this is a two-pronged attack – (1) were there any feasible actions that the ISP could have taken, and (2) were all such feasible actions carried out? Thus, when a claim of indirect liability is brought against an ISP, the defence will typically be that, either monitoring did in fact take place to the best of the ISP's ability and so there is no indirect liability, or that no monitoring took place since monitoring was just overly expensive for this particular ISP, that

[3]It is important to note that safe harbor provisions, as stipulated for instance in Articles 12, 13 and 14 of the E-Commerce Directive in the EU or Section 230 of the Communications Decency Act in the US, immunize ISPs from liability. For instance, see Dinwoodie (2017), and Paynter and Foreman (1998). Clearly, then, indirect liability is the opposite side of the coin of safe harbour clauses, which are in place precisely to give greater legal certainty to service providers, in order to encourage them in their business. It is a matter for the law to decide where exactly to set the boundary that defines the division between safe harbour protection and indirect liability. The present paper does not enter into the debate of exactly where safe harbours should end and liability should begin, but rather simply concentrates on scenarios that would be susceptible for claims of liability, that is, that go beyond anything that could be deemed to be immune from liability.

[4]Of course, indirect liability takes more than one form – contributory liability, vicarious liability, and more recently, inducement liability (see, for example, Dixon (2009), and Landes and Lichtman (2003)). In the current paper we abstract from any differentiation between these types of indirect liability.

is, monitoring should not be expected of the ISP, and so again there is no room for indirect liability. In short, if it is found that there were in fact no actions that the ISP could have reasonably taken to avoid the copyright infringements, then the ISP should not be liable.

On top of this, it is also true that, whenever there exists a diverse set of agents to whom we want to give incentives, a single incentive mechanism is unlikely to give an efficient outcome in terms of incentives and risk bearing. That is, if there are different "types" of ISP, differentiated by their ability to monitor website traffic, then the incentives given to each will be different when a single incentive mechanism is used. For a legal system to be "first-best" efficient, one would feasibly need a different set of laws for each and every individual and firm, something which is obviously not a realistic possibility.[5] However, there may still exist ways in which we can avoid the inefficiencies of the current second-best solution, making use of the now well established economic theory of contracts under asymmetric information.

In effect, if a suit is brought against an ISP, then that ISP can look to defend itself by either showing that the ISP did in fact carry out monitoring to the best of its ability, or alternatively by showing that although no monitoring took place, the ISP did not in fact have any simple manner in which to monitor for copyright violations on the websites that it hosts. The very fact that such a defence strategy can be devised, and may be successful, clearly indicates that in reality there must be more than one type of ISP, differentiated by their ability to monitor website traffic. In economics, this is an example of the very well-known problem of adverse selection. It is equally well-known that the solution for an adverse selection problem does not involve waiting for the bad state of nature to occur and then requiring the agent to defend in a court of law which type she actually is, against the threat of punishment for an agent who claimed to be of one type but is found to be of another. Rather, it is efficient for the principal to offer a contract menu at the outset, and to allow each agent to choose from the contract menu, in a manner in which self-selection occurs.

[5]Current legal systems, copyright included, that offer a single set of laws to which all must abide, are a "second-best" efficient option, but where the enormous transactions costs of the first-best make it unfeasible.

However, for the case at hand, the situation is not quite so simple as straightforward adverse selection, as each ISP also has an effort choice to make, that is, whether or not to actually monitor. It is unlikely that the outcome of this choice (whether or not monitoring took place) is easily observable by any third party, so a hidden action is also present. This is an element of moral hazard that we must consider along with the adverse selection. Again, it is very well-known that an efficient solution to a moral hazard problem of this sort is not to wait for the bad state of nature to occur, and then to require the agent to defend that she did indeed carry out the action in question, but rather an incentive compatible contract should be offered at the outset.[6]

In short, we assume here that an ISP is characterized by a parameter (unobservable by anyone other than the particular ISP) that determines the ability to monitor, and then the actual act of monitoring or not is also unobserved by others. Our objective is to find contracts between the copyright holders (as represented by a collective) and the ISPs that are both self-selecting (i.e. each type of ISP signs a different contract) and incentive compatible (i.e. those ISPs that are contracted to carry out monitoring, and only those ISPs, actually do so). Notice then that the model here combines both adverse selection and moral hazard in a single contracting environment.[7]

Our analysis mimics the standard models of applied microeconomics known as the principal-agent setting under asymmetric information. The assumptions that we make are the same as those made in all such analyses, and our intention is to simply apply that setting to the case of copyright holders and ISPs. The outcome of the model, just as in all standard settings, is a set

[6]For a general discussion of both adverse selection and moral hazard in economics, see (for example) Macho Stadler and Perez Castrillo (2001), and Bolton and Dewatripont (2005).

[7]The economics literature abounds with models of adverse selection and moral hazard separately, but models that analyse the principal-agent setting with both problems present simultaneously are rare. This may well be due to the extreme complexity of a general analysis of such a situation. For early examples of models with both adverse selection and moral hazard, see Picard (1987) and Page (1991). Faynzilberg and Kumar (2000) give a general treatment of such a model, and chapter 6.3 of Bolton and Dewatripont (2005) gives a short survey and analysis of an application.

of equilibrium contracts that stipulate payments between the parties (here, licensing agreements) contingent upon the different outcomes (here, detection or not of indirect liability).[8] The model does not purport to describe any current real-life scenario of contracting between copyright holders and ISPs, but rather the objective is to show how such contracts might be used as an efficient solution to the problem of indirect liability for copyright infringement. In that sense, the model and analysis is not intended to be positive in nature, but rather it is a very normative model and it offers normative conclusions.

2. Analytical Framework

Here we present a simple model of the situation of ISPs and copyright holders, with the latter being represented by a single copyright collecitve. The model posits the collective as the "principal", and the ISP as the "agent". Assume that there are only two different ISP "types", and only one possible action choice (monitoring or not).[9] There are only two states of the world, defined by whether or not a copyright violation occurs and is detected on one of the hosted websites. Copyright violation detection is under the control of the copyright collective, or perhaps an external enforcement agency. We do not model the process under which violations are searched for or detected, but rather we treat detection as an exogenous random variable. Again, this is in order to place our model in exactly the same environment as the standard principal-agent setting. In state 1 no copyright infringements are discovered on the websites that the ISP hosts, and in state 2 copyright infringements are detected. The

[8]The authors are indebted to a referee who has pointed out the similarities between the underlying message of the model in this paper and copyright compensation systems (CCS). These are general surcharges paid by copyright users to a collective for posterior distribution to copyright holders, and they can be thought of as a form of licensing (see, for a very timely example, Handke et al. 2018). In exactly the same way, the model in the present paper can be thought of as a form of licensing, although it involves licensing contracts that may be contingent upon observed outcomes of a random variable (detection or not of indirect liability).

[9]Monitoring is defined as any act of the ISP under which it actively and purposefully analyses the activity on the sites it hosts in an effort to identify copyright infringements, which it would then remove from its services. Monitoring by an ISP will thereby reduce, but likely not eliminate, the probability of the copyright collective detecting infringements. Monitoring imposes costs upon an ISP, and those costs will depend upon a variety of factors that are different across ISPs, like for example the number of sites hosted and the amount and type of material on each site. While a very interesting concept, the endogenous issue of how monitoring choices by an ISP might affect their own cost of monitoring (as it might affect the number and type of customers using the service) is ignored in the present analysis.

probability of the two states of the world will be determined, at least in part, by the monitoring actions that the ISP can take to avoid such violations. Concretely, the probability of state 2 is reduced when the ISP carries out monitoring activities, but different ISPs have different abilities to carry out such activities. These differences may be determined, for example, by the type of material that is typically uploaded to websites hosted by the ISPs, the sheer number of websites that are hosted, and any number of other aspects. We assume that the only thing that differentiates ISPs is their ability to monitor, i.e. in regards everything else, they are identical.

Above all, it should be noted that this is very closely related to a scenario of pure risk bearing. The ISP faces an exogenous risk (that of indirect liability), and can choose to mitigate that risk by undertaking monitoring activities. The decision to monitor is identical to what is called "self-protection" in the insurance literature, in that it decreases the probability of the bad state of nature. The model described in this paper uses the similarity between the situation of the ISP and a situation of general risk bearing with the option of self-protection in much of the terminology used.

Appealing to the standard theory, the copyright collective should offer different contracts, with different payoffs for each state of nature, that the ISPs can choose from. Each different contract will be designed for a particular type of ISP, and self-selection occurs when each type of ISP freely chooses the correct contract for that type. Incentive compatibility is also an issue, and within each contract there is also a demand for monitoring – here, either monitoring is demanded or not. Thus, incentive compatibility happens if each ISP that chooses a contract that includes a demand of monitoring, actually prefers to monitor, and each ISP that chooses a contract with no monitoring prefers not to monitor.

Following the standard principal-agent theory results under asymmetric information, it would be expected that in a self-selecting equilibrium, the ISPs that have the lowest ability to effectively monitor the nature of the traffic on their websites would be offered a contract that

makes them immune to a claim of indirect liability. In other words, the copyright collective fully insures such an ISP against claims of indirect liability. It may also seem reasonable that an ISP that can show that it did indeed carry out the monitoring activities should also be immune to the claim of indirect liability. However, we shall see that this is not the case. In general, performing the monitoring activity will have the effect of reducing the amount of remedy that a plaintiff can claim under indirect liability, but without eliminating the liability outright. Thus, as is normal in adverse selection models, the better "type" (here the ISP with a greater ability to monitor) will typically need to bear some risk in the final outcome, but the risk borne will not be as great as the risk if no monitoring at all were carried out.

Note that although an ISP that chooses a full insurance contract would be immune to indirect liability, the cost for such a contract would be relatively high. On the other hand, an ISP that has chosen a contract that involves monitoring, and that consequentially has a low probability of being sued for indirect liability for copyright violation, would face a much smaller upfront cost, but would be liable to pay some compensation to copyright holders should copyright infringing material be detected on their sites. A proper choice by the copyright collective of the prices of each contract, and the level of the remedies that are charged to a partially insured ISP that has found to have hosted a copyright infringing website, will make the proposed system efficient.[10]

The payments that are made by ISPs who subscribe to contracts with the copyright collective should of course serve as compensation for the copyright holders.[11] The most obvious way to do this would be to understand that the initial payments for entering into such contracts are payments for blanket licenses for the use of copyright material that would be paid to the collective for distribution to the copyright holders. Those ISPs that contract to a full coverage

[10]Of course, we might also desire to have a null-contract, at zero cost, under which the ISP has no insurance at all against claims of indirect liability, and must defend his case in court in exactly the same way as is currently the case. However, in the model expounded here, such a null-contract would not be chosen by any ISP, and is thus redundant.

[11]We make no effort at all to consider the secondary issue of how the income that the collective earns from ISPs is distributed among its members. That decision is irrelevant to the issue of optimal monitoring and efficient contracting which is the focus of the present model.

contract will have paid the equivalent of a full blanket license, and so are immune to posterior claims of indirect liability for copyright violation. On the other hand, those that purchase partial insurance contracts (a discounted initial fee) will still be liable for further payments should copyright violations be detected on the websites that they host. Any payments that are awarded in the form of remedies from under-insured, or not insured, ISPs from law suits claiming indirect liability would be paid to the plaintiff, i.e. the copyright collective.

In the model, we shall assume that the copyright collective (the principal) is risk neutral. Thus, the collective is indifferent to a situation of no contracting at all and one in which contracting occurs but at the same expected level of contract and remedy payments. The ISPs (the agents) are assumed to be strictly risk averse, and so they strictly prefer a situation of the same expected wealth but less risk. Again, these assumptions are those that are standard and accepted in the principal-agent literature. For example, consider an ISP under the current system of no contracting with wealth of \$3 million if no copyright infringing material is detected on its websites, and of \$2 million if copyright infringing material is detected (that is, the indirect liability remedy paid to copyright holders is \$1 million). If each of these two states are equally likely, then the ISP's expected wealth is \$2.5 million, and the copyright holders' expected remedies are \$0.5 million. If, instead of waiting to see if copyright infringements are detected or not, the copyright holders and the ISP agree that the ISP should pay \$0.5 million to the copyright holders up front in exchange for total immunity from indirect liability, then the risk neutral copyright holders would be indifferent, and the risk averse ISP would be better off.

Notice that the system described above manages to give a different incentive to each ISP, depending on the contract chosen. But it would not be an unmanageably difficult system to put into practice. Indeed, from the point of view of the copyright collective and the copyright holders, there should be no difference in how they act currently. If the copyright holder detects infringements on a website, then he can look on a registry of the contracts with the ISPs to

see to what degree a remedy can be obtained from a claim of indirect liability. If the ISP in question, for example, was found to have fully insured, then the copyright holder has already been compensated for the indirect liability, although there may still be scope for a direct liability claim against the individuals or firms who upload copyright protected material.

3. Analysis

3.1. **Benchmark model (no adverse selection).** Assume that there are only two "types" of ISP, differentiated by their cost of monitoring the traffic on the websites that they host. There is a single act of "monitoring", independent of ISP type. A type i ISP suffers a (financial) cost of c_i in order to monitor, where $c_1 < c_2$, that is, a type 1 ISP has a lower cost of monitoring than does a type 2 ISP. Outside of the ability to monitor website traffic, the ISPs are otherwise identical. By monitoring the website traffic, an ISP is able to reduce the probability of copyright violations on the corresponding websites. Finally, we assume that the low cost ISP is willing to undertake the monitoring even in absence of any contract with the copyright collective, while the high cost ISP is not willing to do so. Whether or not an ISP monitors, copyright holders always have the ability to sue under indirect liability, but the remedies that are available are determined by the type of contract that the ISP has with the copyright collective.

As in many simple principal-agent models, we assume that the copyright collective acts as a risk neutral principal.[12] We also assume that the collective acts as if it were competitive, in the sense that it does not attempt to maximise the expected profit earned from the ISP from indirect liability.[13] Thus, the model works as if the collective were interested in improving the expected utility of the ISPs as much as is possible while respecting the participation and

[12]A model with a risk-averse collective would, of course, be entirely possible to set up and analyse. However, it would also be very cumbersome due to the added complexity of risk aversion on the side of the principal. Nothing of particular importance to the understanding of how asymmetric information is handled is gained from adding this complexity, and so in keeping with the general tradition of simple principal-agent models, we retain the assumption of risk-neutrality.

[13]Copyright collective pricing is often regulated, in exchange for allowing the copyright monopoly to exist. Thus, holding the collective to zero expected profit may be thought of as the effect of regulation.

incentive compatibility constraints. Of course, the copyright collective will also not reduce the expected value of the indirect liability remedies that they would receive from the default of no contracting.[14]

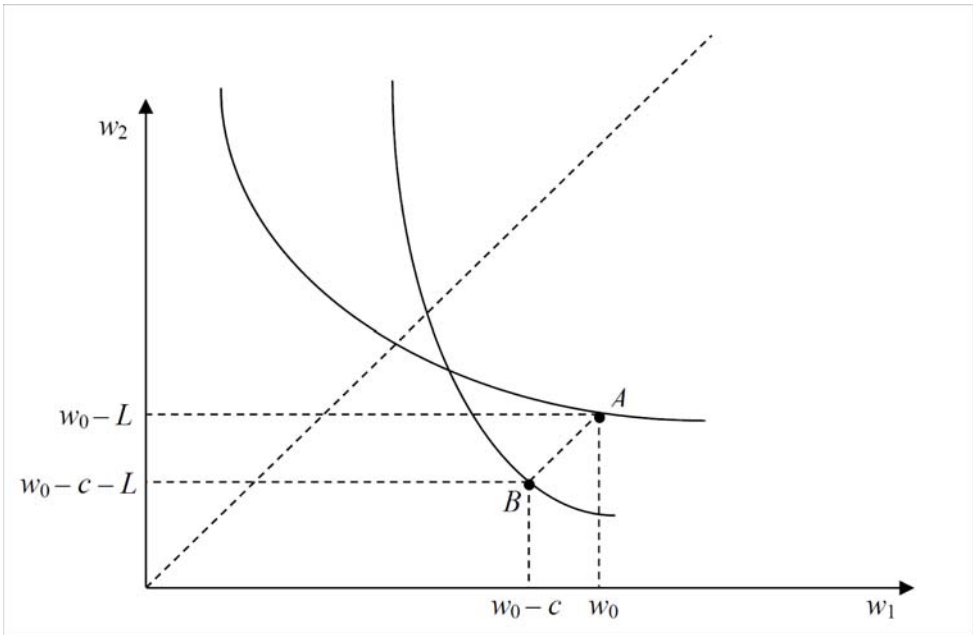

Figure 1

To help with our understanding, let's go through the model when there is only one type of ISP, and it is a low-cost ISP (i.e. even with no contract with the copyright collective, the ISP is willing to undertake monitoring). We use throughout a simple model of moral hazard due to Rees (1989). In Figure 1 we show the initial situation drawn in standard contingent claims space (the space of wealth of the ISP in each of the two states of nature). Here, state 1 wealth is denoted w_1, and state 2 wealth is denoted w_2. The ISP has an initial wealth of w_0, but faces an indirect liability risk valued at L. The probability that the ISP will have

[14]Throughout, we limit our analysis to a graphical exposition, rather than mathematical. Of course, a full mathematical environment lies beneath our graphs, but all of the understanding and logic behind all of the results can be much more easily gained using the graphs rather than the corresponding equations.

to pay the liability is initially p. Thus, the ISP is initially located at point A (with wealth of w_0 in state 1, and of $w_0 - L$ in state 2) with the indifference curve passing through that point. On the other hand, if the ISP were to undertake monitoring, it would lose c dollars in each state of the world, leading to point B, but in exchange the probability of suffering the indirect liability cost is reduced. The indifference curve becomes steeper to reflect the more favourable probabilities after having undertaken the process of monitoring. Notice that the two indifference curves intersect below the diagonal "certainty line" passing through the graph, indicating that the ISP prefers to monitor.[15]

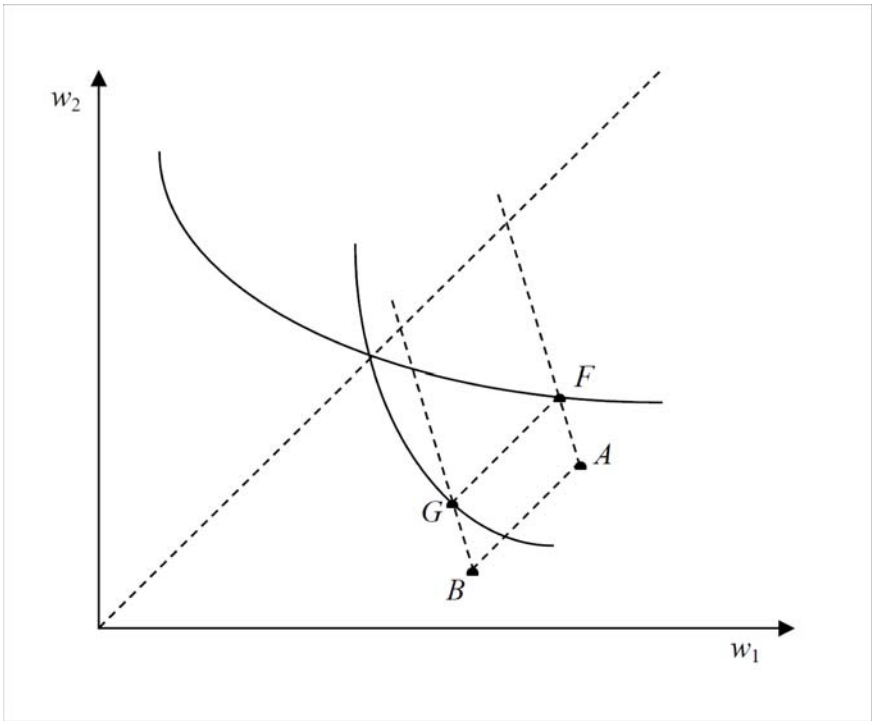

Figure 2

[15]The indifference curves are comparable where they go through the certainty line. The curve which intersects the certainty line at the highest point is associated with the more preferred level of utility.

Now, by contracting with the ISP, the copyright collective can move the ISP along the expected value line passing through point B (see Figure 2), thereby shifting the indifference curve upwards, implying an increase in the expected utility of the ISP. However, since the act of monitoring is not observable, it may be that the ISP decides to take the deal offered by the copyright collective but then not carry out the monitoring activity, which would imply an effective movement from point A upwards along a line that is parallel to the expected value line passing through point B. However, by accepting the deal and yet not carrying out the monitoring, of course the ISP suffers the same probabilities as at point A, i.e. a relatively flat indifference curve.

In Figure 2 we draw the equilibrium movement. The equilibrium contract is at point G, which is a risk reducing movement from point B. The point F is the point that would be obtained by an ISP that takes the deal but does not monitor. Notice that the equilibrium is found when the indifference curve corresponding to actually carrying out the monitoring and the indifference curve corresponding to not monitoring (but taking the deal) intersect exactly at the certainty line. This intersection implies that the ISP is exactly indifferent between taking the deal and monitoring or taking the deal and not monitoring, and both of these are better than not taking the deal and not monitoring. As is normal in economics, we resolve the indifference in favour of the ISP taking the deal and then undertaking the monitoring.

What sort of deal has been offered to the ISP? It is essentially an insurance arrangement, under which the risk-neutral copyright collective insures the risk-averse ISP against the risk of discovery of copyright infringements. However, the insurance arrangement is partial in nature, that is, if a copyright infringement is detected, the liability payment that the ISP should make is smaller than without the insurance, but it is not reduced to 0 (as would be the case under full insurance coverage).

In Figure 3 we show all of the information. The distance between H and J is the initial liability cost (the amount of money that the ISP would have to pay in indirect liability remedies

if copyright infringements were discovered on its websites in absence of any contracting, L). The movement from point A to B corresponds to the ISP actually paying the cost c and undertaking the monitoring. Then, the movement from B to G is the "deal" offered by the copyright collective. The deal consists of the ISP paying a further cost equal to the distance between points J and K to enter the deal in exchange for a reduction in the liability remedies of an amount equal to the distance between G and J. That is, the liability is now reduced to the distance from H to G. In essence, the copyright collective is offering an insurance arrangement to the ISP (priced at the fair premium), at a premium payment of distance JK, and with an indemnity of distance JG.

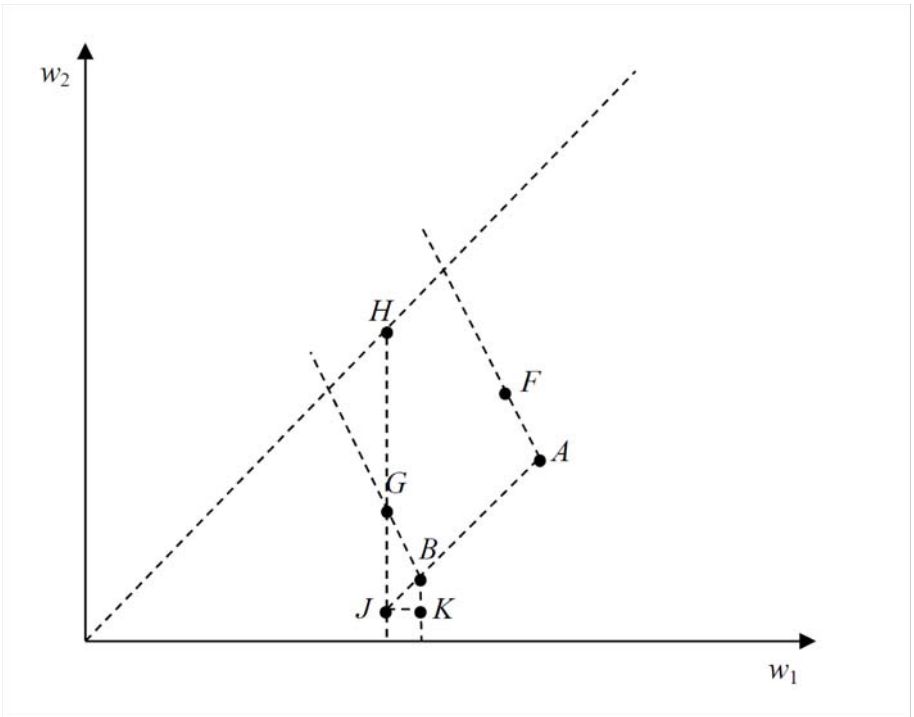

Figure 3

Now, let's look at the situation of a single type of ISP, but where this time the ISP has a high cost of monitoring, that is, the ISP is unwilling to monitor. In this case, the indifference curve through point A intersects the indifference curve through point B to the left of the certainty line. Now, attempting to move the ISP indifference curve upwards from point B is fraught with failure, since by accepting the deal and yet not monitoring, the ISP will always achieve a greater level of utility. That is, there is now a clear issue of incentive compatibility that the collective needs to bear in mind when designing the contract to offer. Effectively, there is no incentive compatible deal involving partial insurance (i.e. requiring monitoring) that the copyright collective can offer this ISP in return for monitoring to be carried out.

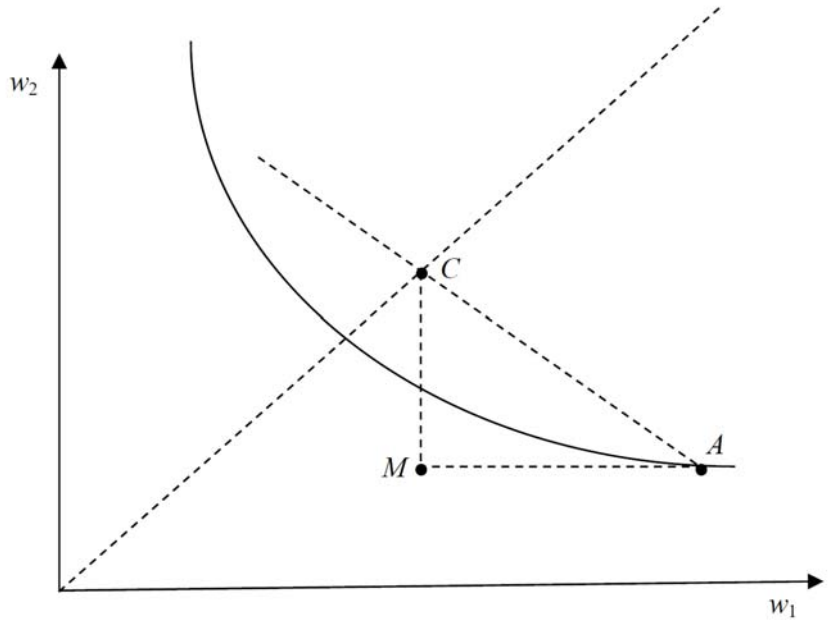

Figure 4

This, however, does not mean that the copyright collective cannot offer any deal at all. Just as in a regular adverse selection model, the collective can move the ISP along the expected

value line passing through point A to the point of intersection with the certainty line. The deal located at point C in Figure 4 keeps the expected value of the ISP's wealth constant at the level it is at initially, but removes all risk of indirect liability. At point C, the ISP is fully insured, but at a cost (equal to the distance from A to M), which is the statistically fair premium for such insurance coverage. Never-the-less, since the ISP is assumed to be risk averse, and this deal removes all risk but retains the expected value of the ISP's initial situation, this is a most beneficial deal for the ISP, and it would be gladly accepted.

We now know how the moral hazard game would be solved if there were only one type of ISP. On the one hand, if the ISP has a high cost of monitoring, then full coverage should be offered at the actuarially fair premium, and no monitoring should be demanded. Correspondingly, no monitoring will be carried out. On the other hand, if the ISP has a low cost of monitoring, then partial coverage against indirect liability should be offered, again at the actuarially fair premium, and a certain level of monitoring will be demanded, and will indeed be carried out.

Of much more interest is the case in which we have both moral hazard and adverse selection together, that is, the collective cannot identify, ex ante, the type of ISP that it is contracting with – low or high cost of monitoring – although the ISP itself does know its type, and at the same time, the collective cannot observe whether or not monitoring is carried out (and of course the ISP does know whether or not it has monitored). We now turn to an analysis of exactly this sort of situation.

3.2. **Simultaneous adverse selection and moral hazard.** When both types of ISP exist simultaneously, the copyright collective needs to consider the self-selection constraints that determine that the high cost ISP does not prefer the contract designed for the low cost one, and vice versa.

To start with, recall from what we have done above that the point of full coverage at the initial expected value is a contract that cannot be improved upon for the high cost ISP. Let's say then, that this point (point C in Figure 4) will always be offered by the copyright

collective with the high cost ISP in mind. The question then is, what contract can be offered for the low cost ISP? Call point D the certainty equivalent point for a low cost ISP that undertakes monitoring but without any contract with the copyright collective. Point D is the point at which the indifference curve passing through point B, with slope reflecting the fact that monitoring has taken place, cuts the certainty line (see Figure 5). The assumptions that we have made on the model are consistent with point D being either above or below point C (or of course, equal to point C). We need to analyse each case separately.

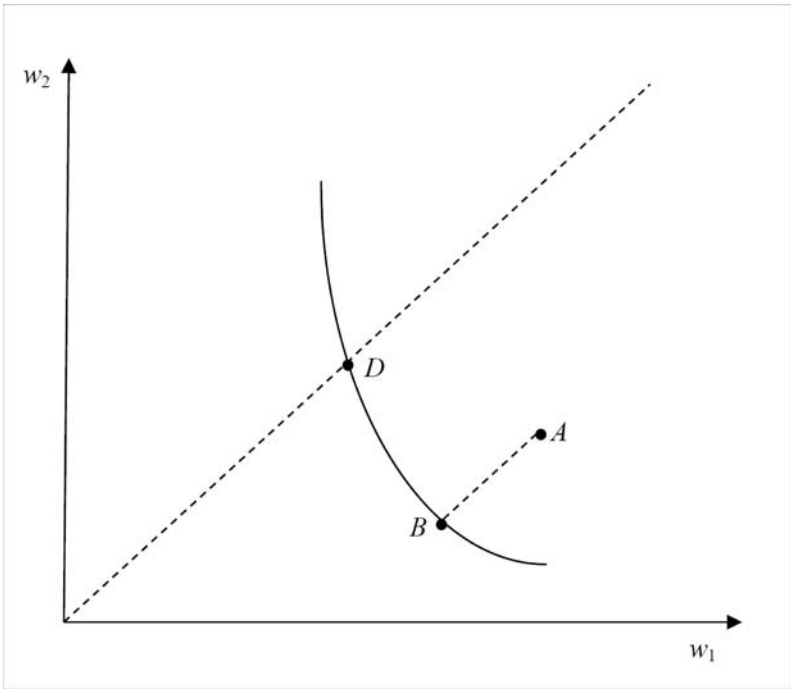

Figure 5

3.2.1. *D is above C (see Figure 6).* In this case, even if we were not to offer any contract other than point C, the low cost ISP would still carry out the monitoring (locating at point B), and would of course reject contract C. However, it is certainly possible in this case to improve the welfare of the low cost ISP over and above that obtained at the non-contracting point B. The relevant analysis is shown in Figure 6.

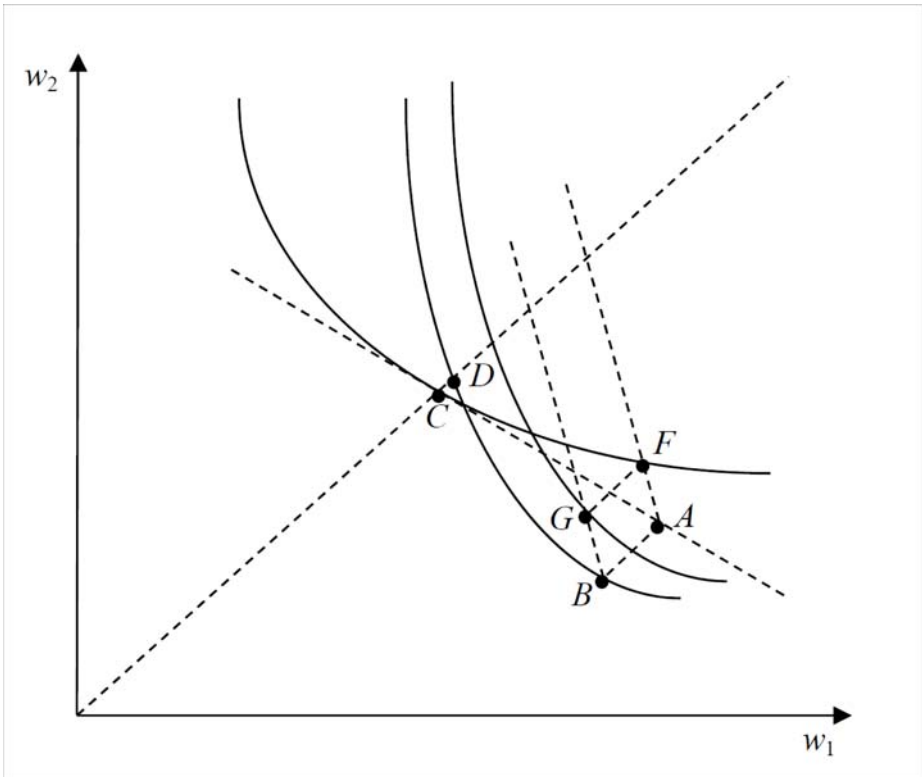

Figure 6

Recall that by moving the low cost ISP upward along the new expected value line (that corresponding to having carried out monitoring), the low cost ISP can either take the contract and carry out the monitoring (in which case it moves upward along the negatively sloped line through B in Figure 6) or take the contract and not monitor (thereby moving upward along the steepest negatively sloped line through A). In the first case the indifference curve is relatively steep (since monitoring was carried out) while in the second case the indifference curve is relatively flat (since no monitoring was carried out).

Now, consider point F in Figure 6. This point mirrors point G, with the only difference between them being that at F no monitoring is carried out, while at G it is. The intersection

between the indifference curves passing through F and G must lie below the certainty line, since by assumption point D is above point C. Thus, given the option of signing the partial coverage contract defined by point G, the low cost ISP would take that contract and would prefer to carry out the monitoring as demanded. This option is also preferred by the low cost ISP to the contract located at point C. Note also that the high cost ISP would be indifferent between taking this contract and then not carrying out the monitoring, and taking contract C.

Therefore, the solution offered by points G and C achieves self-selection (that is, the high cost ISP chooses point C, the low cost ISP chooses point G), and incentive compatibility (the high cost ISP chooses not to monitor, the low cost ISP chooses to monitor). Any further movement upward above point G will violate incentive compatibility, as it will have a mirror point above point F, which would be accepted by the high cost ISP without monitoring. Thus the equilibrium (in the case in which point D is above point C) is made up of points G and C.

Notice that again, in this equilibrium, the high cost ISP is fully insured and does not undertake any monitoring. On the other hand, the low cost ISP is only partially insured, and does monitor. The low cost ISP must suffer risk, but not as much risk as if there were no contract at all. The reader will readily recognise this sort of situation, often referred to as signalling, from regular models of pure adverse selection, where the more attractive agent type signals her type by accepting to be only partially insured, and the less attractive agent type is fully insured.

3.2.2. *D is below C (Figure 7).*

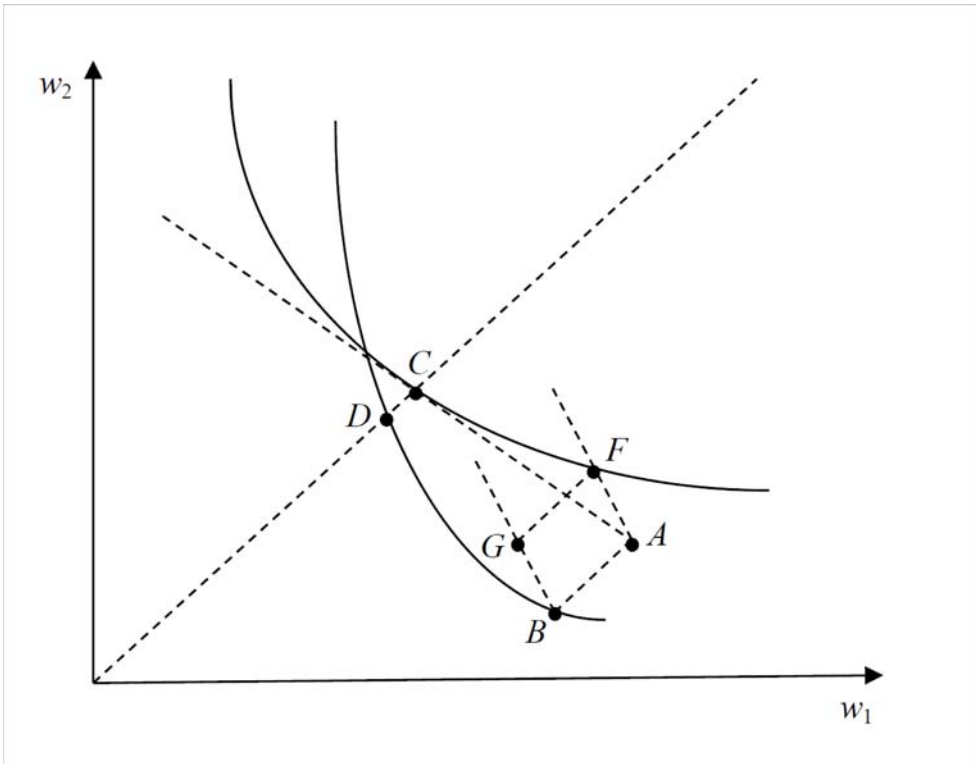

Figure 7

When point D lies below point C, the outcome without any contract offerings besides C would be that both types of ISP would choose point C, and neither of them would carry out any monitoring. This may be considered to be not a good outcome, since no monitoring at all will take place, however it is simply a reflection on the assumption that now even the low cost ISP has a relatively expensive monitoring technology. Never-the-less, again it may be possible to offer a second contract that will inspire monitoring by the low cost ISP.

To start with, note that again we can clearly continue to offer point G as well as point C (Figure 7). Point G, which is mirrored by point F, will be rejected by the high cost ISP in favour of point C, so there is no concern that G will be accepted but with no monitoring

undertaken by the firm accepting it. Point G is also strictly preferred by the low cost ISP to point B, but it may or may not be preferred to the contract at point C. This preference would be evaluted by drawing an indifference curve for a low cost ISP that monitors through G and seeing whether that curve intersects the diagonal of the graph above or below C. We have not drawn this curve, simply because to be coherent we would have to try to draw two – one intersecting above C and the other intersecting below C, which could be somewhat confusing on a single graph.[16] Thus, we now have two sub-cases; either (a) the utility of the low cost ISP at point G (with monitoring) is greater than the utility of the low cost ISP at the fully insured point (point C), or (b) the utility of the low cost ISP at point G (with monitoring) is lower than the utility of the low cost ISP at the fully insured point (point C). In the first case (case (a)), the equilibrium is again to offer points G and C. In the second case (case (b)), the equilibrium is to only offer point C, and all ISPs are pooled at the full-insurance contract with no monitoring.

4. Conclusions

This paper explores a model of contracts between copyright holders and ISPs as copyright users under asymmetric information on ISP monitoring activity and costs. ISPs are firms that host the websites of other firms, and they can be found liable for the copyright infringing activities that may take place on the websites they host. An ISP may mitigate the risk of infringing activities on the hosted wbsites by engaging in monitoring activities that are costly to the ISP. Undertaking monitoring is a clear example of "self-protection" against the risk of indirect liability. The self-protection analogy is used in the current paper to then study the types of contracts that might be written that can insure the ISP against the risk it faces.

[16]Essentially, just shift the curve through B and D upwards until it passes through G, and then see if it intersects the diagonal above or below C.

Monitoring activities are non-observable for copyright holders, so there is moral hazard present. On top of that, there are more than one "type" of ISP, differentiated by their ability to monitor. This is an aspect of adverse selection.

This paper explores the possible contracts that may be signed between an ISP and the copyright holders (acting via a copyright collective) such that the contracts solve, simultaneously, both the adverse selection and the moral hazard problems, that is, contracts that are simultaneously self-selecting and incentive compatible. We find that such contracts do exist. As a general rule, in the equilibrium, low cost ISPs will have a contract that demands monitoring. The level of indirect liability in the contract will be less than the level of indirect liability in absence of contracts. On the other hand, high cost ISPs will have a contract that insulates them to a much larger degree from indirect liability, but at a higher initial cost. In essence, the contracts in question are a mechanism under which the copyright holders can insure, either partially or fully, the ISPs against the threat of indirect liability.

The copyright collective will consider firstly offering a full coverage contract at a high premium in order to attract high cost ISPs. Such a contract would involve no monitoring, and would make an ISP that accepts it immune to indirect liability. Then there are basically two types of scenario. Firstly, if the low cost ISP would prefer to reject the full coverage contract, then the low cost ISP can obtain a higher utility by monitoring and not contracting at all with the copyright holders than by accepting the full immunity contract. In this case, it will be possible for the copyright holders to offer a partial coverage contract at a low premium and that demands monitoring that would be accepted by the low cost ISPs, and rejected by the high cost ISPs. Thus, in this case, we get a fully self-selecting equilibrium in which the high cost ISPs choose the full coverage contract and do not monitor, while the low cost ISPs choose the partial coverage contract and do undertake monitoring.

The second case occurs when the low cost ISPs would prefer to accept the full coverage contract than to monitor and not contract with the copyright holders. In this case, again

we can always construct a partial coverage contract for the low cost ISPs such that the high cost ISPs would always be willing to accept the full coverage contract with no monitoring instead of the partial coverage contract with monitoring. But it is not clear that the low cost ISP would prefer the partial coverage contract. If it does, then again we have a self-selecting and incentive compatible equilibrium. However, if the low cost ISP does prefer the full coverage contract to the partial coverage one, then we can have a pooling equilibrium at the full coverage contract, such that all ISPs accept this contract and none of them undertakes monitoring.

One general result that appears in this paper, and that may have some implication for the way the current laws are interpreted, is that fact that although the low cost ISP can be given the correct incentive to monitor, the act of monitoring (to the best of the ability of that ISP) will not imply full immunity from indirect liability, only a reduced level of liability. Full immunity only occurs for high cost ISPs, and in the cases in which the full coverage contract is implemented. In contrast to this, currently (that is, without any insurance type contracting between ISPs and copyright holders) an ISP may be able to successfully defend against a claim of indirect liability if it can successfully prove that it did in fact undertake all of the monitoring activities that it was able to undertake (that is, actually carrying out the monitoring is not a sufficient defence, but rather monitoring to the best of one's ability would need to be proved, which may be difficult since monitoring is by assumption here an act that is not observed by third parties). If, on the other hand, the ISP contracts with the copyright holders under a partial coverage contract then there would be no need to prove monitoring at all, but the liability remedy that can be sought will be reduced.

References

Bolton, P. and M. Dewatripont (2005), *Contract Theory*, Cambridge, MA: MIT Press.

Dinwoodie, G.B. (2017), "A Comparative Analysis of the Secondary Liability of Online Service Providers", in Dinwoodie G.B. (Ed.), *Secondary Liability of Internet Service Providers*, Ius Comparatum – Global Studies in Comparative Law, Vol 25, Springer, Cham (Switzerland).

Dixon, A. (2009), "Liability of Users and Third Parties for Copyright Infringements on the Internet: Overview of international developments", in Strowel, A. (Ed.), *Peer-to-Peer File Sharing and Secondary Liability in Copyright Law*, Edward Elgar Publishers Ltd.; pp. 12-42.

Faynzilberg, P.S. and P. Kumar (2000), "On the Generalized Principal-Agent Problem: Decomposition and existence results", *Review of Economic Design*, 5; 23-58.

Handke, C., J.P. Quintais and B. Balazs (2018), "Truce in the Copyright War? The Pros and Cons of Copyright Compensation Systems for Digital Use", *Review of Economic Research on Copyright Issues*, 15(2); 23-56.

Landes, W. and D. Lichtman (2003), "Indirect Liability for Copyright Infringement: Napster and beyond", *Journal of Economic Perspectives*, 17(2); 113-24.

Macho Stadler, I. and D. Perez Castrillo (2001), *An Introduction to the Economics of Information*, 2nd Edition, Oxford. UK: Oxford University Press.

Page, F.H. (1991), "Optimal Contract Mechanisms for Princial-Agent Problems with Moral Hazard and Adverse Selection", *Economic Theory*, 1; 323-38.

Paynter, H. and R. Foreman (1998), "Liability of Internet Service Providers for Copyright Infringement", *University of New South Wales Law Journal*, 21; 578-92.

Picard, P. (1987), "On the Design of Incentive Schemes Under Moral Hazard and Adverse Selection", *Journal of Public Economics*, 33; 305-31.

Rees, R. (1989), "Uncertainty, Information and Insurance", in Hey, J. (Ed.), *Current Issues in Microeconomics*, London, MacMillan.

Watt; University of Canterbury, richard.watt@canterbury.ac.nz. Mueller-Langer; Max Planck Institute for Innovation and Competition.