THE ECONOMIC EVALUATION *of* ALTERNATIVES TO DIGITAL COPYRIGHT

Peter Eckersley

Department of Computer Science & Software Engineering
and Intellectual Property Research Institute of Australia,
The University of Melbourne

SERCIAC 2003
(Preliminary Version)

ABSTRACT

A certain air of controversy has arisen around copyright law, as a result of its interactions with digital technology. The body of literature claiming that existing copyright laws are economically sub-optimal is growing rapidly. Some authors are even claiming that it would be better to have no copyright system at all, while others argue that alternatives, such as various forms of public funding, would be preferable to exclusive rights in literary and artistic works.

This paper explores the economic differences between a system of consumer-regulating copyright based on "digital rights management", and alternatives based on public funding. I argue that the distinctions are sufficiently intricate that they elude any simple modelling technique.

Instead, this paper attempts a semi-analytic comparison which weighs a set of different factors which may favour either copyright or the alternatives. When these are summed, it is concluded that well-designed systems of public funding are almost certain to produce better social welfare outcomes than DRM-based copyright.

1 Introduction

The infrastructure for purely digital information economies is only beginning to crystallise[1]. As this occurs, it has become apparent that some of the central assumptions which have defined information industries in the past are in danger of unravelling. Perhaps the most prominent assumption of this kind is that states of artificial scarcity will be maintained by the exclusive rights of copyright law (in combination with "digital rights management").

It is too early to say yet whether technologically enforced copyright is a feasible long-term proposition. There are reasons to believe that there are some markets where it will be, and some markets where it will not.

But an essential, and often overlooked, point is that *feasibility* is not the correct test to apply to DRM-based copyright. Other means are available to ensure that the production of information goods is not hampered by free-rider problems. It is only after it has been shown that digital exclusive rights are likely to produce the best outcomes for the welfare of society, that they should be supported by policy.

Some authors argue that the free rider problem can be solved by decentralised, incentive-compatible contracting processes, in combination with "first mover" advantages and other coincidental incentives. These ideas are briefly addressed in Section 1.2.1.

The main intention of this article, however, is to evaluate proposals which call for public funding as full or partial replacements[2] for the exclusive rights which cover literary and artistic works. Questions about functional works which happen to fall under copyright law (such as academic texts, or software) are not addressed.

The comparison between DRM-copyright and publicly funded rewards is particularly intricate; it turns on a set of disparate technological and economic effects which are difficult to unify. Despite this complexity, I will argue that it is very likely that societies would be better off adopting such models in place of technologically enforced copyright systems.

[1] Some notes on semantics:

Infrastructure refers to both technological infrastructure (computers, networks, protocols, file formats, and application software), legal infrastructure, and socio-economic cultures of usage and production (the later may also, to some extent, require viable business models).

The term *information* has two important and distinctive meanings in discussions about copyright economics. The sense used here refers to broadly-defined entities which are composed of, or can be represented by, data — such as writing, music, films or software. The other sense, drawn from the economics literature, refers specifically to knowledge about prices, costs, preferences, or possibilities, which is held or transmitted by economic actors.

The word *digital* applies, first and foremost, to information which is constructed from discrete data. It also connotes the usage of such information in electronic computers that can reproduce it flawlessly and automatically.

[2] Most proposed alternatives to copyright are concerned with the kind of "private" or "non-commercial" copying one might find on peer-to-peer networks, but not with unambiguously commercial uses, such as advertising. Proposals for alternatives to copyright also differ in their approaches to the exclusive right of derivation.

1.1 Option 1: Technologically Enforced Copyright

The future of copyright, as advocated by the majority of significant publishing organisations, and presently sanctioned by law, rests in a world where exclusive rights are programmed into every computer (and every other other gadget used for manipulating and consuming information goods). This technology is often referred to as "DRM", which usually stands for "digital rights management", although its detractors argue that "digital restrictions management" is more accurate.

The core of technologically-enforced copyright is a device called a *trusted system*. The point of trusted systems is that, although they are owned by consumers, copyright holders retain a sufficient degree of control over them that the devices can be "trusted" to follow rights holders' rules.

Information can be moved into trusted systems through the use of secure cryptographic channels. But the devices must be resistant to physical tampering, so that the cryptographic keys (and the unencrypted content) they contain cannot be liberated by users. The cost of achieving practical tamper resistance may be very high, as I will explain in Section 3.1.

Another major problem with trusted systems, as Kelsey and Schneier (1999) have emphasised, is that they must always release their contents in human-accessible analogue form. At that point, an enterprising user may use recording devices to obtain an unencrypted, digital version[3]. This kind of copying is, however, likely to remain much more difficult than purely digital reproduction.

An auxiliary technology of some interest is algorithms for creating *watermarks* on information goods. The idea is that each user receives a version of the file which has been uniquely marked in some way — so that if they send copies to other people, a *traitor tracing* process can be used to identify (and hopefully prosecute) them for infringement.

If watermarks were secure, they could make DRM-based copyright significantly more feasible. But, to date, no watermarking techniques have been developed which cannot be automatically removed.

Ultimately, in order to be economically effective, technologically enforced copyright does not need to employ any of the techniques discussed here perfectly. Instead, it need only employ them to an extent that makes purchasing licensed copies of information goods the most attractive choice, for most consumers. The comparison at hand must, then, evaluate an imperfect-but-functional DRM system which achieves this goal.

A more detailed description and discussion of DRM can be found in the work of Stefik (1999), who also provides an economic argument for its adoption. The argument is essentially that

[3]This possibility lead the Motion Picture Association of America to propose that all analogue-to-digital converters (ADCs) be regulated, to require that they recognise and refuse to digitise material which carries a standardised watermark (MPAA 2002). To the knowledge of the author, no further studies have been undertaken to determine how practical or costly this would be. Since ADCs are a simple (and presently cheap) building block for many electrical devices, such regulation might prove quite expensive.

adopting copyright enforcement technologies must be a Pareto-improvement, since all relevant parties must consent to it. He argues that author/publishers will only distribute works in digital form if DRM is available, and that consumers will refuse to purchase digital versions if physical copies retain a more favourable combination of price and flexibility. This position implicitly assumes that there is no market power at work in the transition — and it says nothing about the relative desirability of DRM and any alternatives.

1.2 Alternatives to Artificial Scarcity

There are two important classes of alternatives to DRM-based copyright[4]. Both of them attempt to solve the free rider problem, without creating artificial scarcity in the distribution of information goods.

1.2.1 An Aside: Contracting Protocols

The first class of alternatives is not the focus of this paper, although I will discuss them briefly. They involve finding some way of creating a contract between a large group of consumers, each of whom is interested in overcoming the free-rider problem. These parties form a pact to make contributions to a public good under certain well-defined circumstances.

The idea of forming "buyers' groups" to solve the free-rider problem is an old one. Breyer, in his classic article on the economics of literary copyright, was optimistic that in many cases, this possibility would serve to minimise the damage were copyright to be abolished (Breyer 1970, pp 302–306).

The current crisis in copyright has seen a dramatic growth of interest in these approaches. Kelsey and Schneier (1999) have described what they call the *Street Performer Protocol* (SPP), in which artists produce a piece of work in advance, and calculate the return (call it X) they want from the public in exchange for it. They release a sample as a signal of quality, and accept "pledges" of the form "I, consumer c, will pay x_c dollars if you release the full work" from members of their audience. Once the pledges meet the artist's demand (ie, $\sum_c x_c > X$), the work is released and the pledges called in. In practice, it may help a great deal to have a "trusted third party" (TTP) who can accept the pledges as escrow payments.

In SPP, as proposed by Kelsey and Schneier, the risk of a transaction is borne primarily by the publishing artist. Factors such as reputation may also play a significant role in smoothing SPP's operation.

Rasch (2001) has discussed a "bounty market" variant of the protocol, which is particularly suited for tasks such as funding Free Software[5] development, where very low barriers to production

[4]More accurately, they are alternatives to exclusive-rights based copyright in general, but they are of particular interest in the face of the problems and controversy which surround digital copyright.

[5]Free (as in "freedom") Software is often referred to as *open source software*.

are important. Another variant ("RSPP") is described by Harrison (2002), which is tailored to pooling resources for the production of a continuous-quantity public good.

Decentralised contracting protocols may turn out to be very important for the digital copyright problem. Indeed, it is possible that they are superior to either of the other models examined in this paper.

Economic modelling does not provide a clear picture of SPP's effectiveness[6]. From some perspectives, however, *theoretical* examination of these protocols is less urgent than analysis of public funding options, because decentralised contracting processes can be implemented from the ground up in information markets, and do not necessarily require policy intervention[7]. At this stage, the most productive research which can be undertaken in this direction is probably experimentation — although there have been some successful, isolated deployments of these models[8], we do not yet know how difficult they are to achieve.

1.2.2 Option 2: Publicly Funded Rewards

There is a rapidly expanding literature arguing for various forms of direct government intervention to replace exclusive rights in a digital context.

There are two kinds of pre-existing, organised reward systems which serve as precedents. These are the "public lending right", and "private copying" levy schemes. They are both rewards, financed by taxation, which currently operate in various jurisdictions, addressing particular situations where exclusive rights are infeasible or unenforceable.

[6]The problem is that subtle changes in the information possessed by the agents, the model of rationality, or the role of transaction costs, can dramatically change the outcome. For example, Cornes & Sandler (1996, pp 215–217) demonstrate that a Lindahl voting game, which is quite similar to RSPP, will not produce honest preference disclosure (and hence only partially alleviates the free rider problem). It is not clear, however, that the same effect must occur with discrete public goods; if transaction costs are low enough, the protocol could be re-run until marginal public goods are produced. The outcome in this situation depends on agents' behaviour in a waiting game with imperfect information about the pool of potential contributors.

[7]It is conceivable that they are presently held back by the lack of certain kinds of infrastructure, such as cheap escrow payment systems, which governments could provide. It is also possible that these protocols would flourish, and produce excellent social welfare outcomes, in the absence of workable copyright systems — but that information producers are slightly better off under copyright, and so will not adopt them voluntarily.

[8]Perhaps the most famous experiment of this kind was undertaken by Stephen King, who collected over USD $700,000 in tips from a few chapters of his book *The Plant* (King 2001). It was widely reported that his experiment had failed (The New York Times 2000), although King himself did not appear to take this point of view (see `http://www.stephenking.com/sk_120400.html`). Other examples include the fundraising efforts of kuro5hin.org, a collaborative media site (see `http://www.kuro5hin.org/story/2002/6/21/10533/6651`); Blender, a 3D modelling and animation package (see `http://www.blender3d.com/`), and Linux Weekly News (see `http://www.lwn.net/Articles/5838/`). A recurring factor in all of these cases is that the donations have a strong element of contingency, and thus appear to have a great deal in common with the SPP. Also see Google's "Google Answers" (`http://answers.google.com`) service, which is currently the largest operating example of the Wall Street Performer Protocol.

Public lending rights are used by many countries to pay authors and publishers for the use of their works by public libraries. Their precise scope, method of operation and legal status varies, but in most cases they involve a pool of government funds which are divided according to the relative borrowing frequency of different works (von Lewinski 1992).

Private copying schemes are used in some states as compensation for uncontrollable private reproductions of musical or cinematographic works. They are usually funded by levies raised on reproduction hardware or blank media; the funds are allocated on the basis of numbers calculated by copyright collecting societies (Davies and Hung 1993).

Some of the proposals for alternatives to DRM-based copyright basically suggest expansions of one of these schemes to cover non-commercial activity on the Internet as a whole. For example, Foley (2001) suggests modelling the entire Internet as a library[9], with public lending rights for remuneration; Schulman (1999, pp 628–630), Fisher (2000), Lunney (2001, pp 911-918) and Ku (2002) consider the application of private copying schemes to peer-to-peer networks.

There are also some slightly more radical proposals in the literature. Calandrillo (1998) builds on the patent-oriented arguments of Shavell & van Ypersele (1998), in advocating a system of *ex post* rewards funded by income tax. In a previous working paper, I developed a proposal for a "virtual market" in which a combination of observable statistics and explicit votes by consumers, are used to allocate funds raised by composite taxation (Eckersley 2003). I will argue in Section 3.4 that the use of votes may result in significant improvements in efficiency[10].

2 The Structure of a Comparison between Copyright and Publicly Funded Alternatives

2.1 Microeconomic modelling techniques

DRM-based copyright and the various proposals for public funding models are entirely different structures for cultural marketplaces. The task of making a normative comparison between the two is to some extent ambitious.

A standard economic approach to comparing these structures would be to develop a general, abstracted, microeconomic model, and thus evaluate the conditions under which each mechanism is superior.

There are several properties of the digital copyright question which make this standard approach particularly difficult. One is that explicitly technological factors play an important role in the

[9] An earlier, more "visionary" library-related proposal was made by Rothman (1992).

[10] The idea of using votes in copyright-replacing systems is also not entirely new; for example, Stallman argued that a vote-like scheme should be used to fund software development (1985). In practice, similar effects may be achieved when governments grant 100% tax credits for donations to organisations which produce public goods.

distinction between the options[11]. Another is that, with all but the simplest formulas for using taxation to produce public goods, the desirability of outcomes is closely tied to actual distributions of wealth and preferences[12]. A third, and particularly tricky, problem is that the operation of alternatives to copyright may depend on cultural or psychological factors, where assumptions of rationality are either incomputable or unrealistic.

Tractable microeconomic models of copyright (Landes and Posner 1989; Watt 2000) cannot be easily extended to incorporate the key factors required to illuminate the distinction between exclusive-rights based markets and the more sophisticated public funding models.

Another set of microeconomic models which primarily address patents (Wright 1983; Shavell and van Ypersele 1998) at first appear more promising for comparing copyright to reward systems. But these models are driven by information asymmetries which apply broadly to inventions, but not to artistic and literary goods. An inventor may have more knowledge about the value of their research than a government agency — even if that agency attempts to allocate *ex post* rewards. But it is not obvious that the same information asymmetries apply to "virtual markets" or other alternatives to copyright[13].

There are also some relevant models to be found in the extensive literature on "mechanism design" for the production of public goods, reviews of which can be found in Cornes & Sandler (1996, Chapter 7) or Campbell (1987). Whereas Wright and Shavell & van Ypersele make assumptions about the quality of the information possessed by the government, most of the mechanism design literature models explicit messages sent by "voters" to a "planner" who raises taxes and provides public goods.

This resource allocation literature has been largely concerned with general properties, such as Pareto-optimality, incentive-compatibility, or rational participation constraints, which are desirable for public good production mechanisms. Unsurprisingly, compromises must be made between these properties, especially where voters have sufficient information to engage in strategic long-term behaviour (Roberts 1979) or where there are several different public goods (Bucovetsky 1991). More optimistic results can be obtained through iterative voting if voters are only locally rational (de Trenquayle 1997). Unfortunately, these results are sufficiently abstract that they are of only passing usefulness in a policy examination of copyright.

[11]These technological factors are almost totally separate from the central microeconomic phenomena in cultural marketplaces. A comparison between them thus requires a common metric, such as utility or dollars. Although this form of analysis is possible, it requires further movement away from generalised modelling and towards partial simulation of real societies.

[12]Note that neither copyright nor the alternatives are close to constrained Pareto optimality under realistic assumptions.

[13]The question of what, if any, information asymmetries may exist in copyright-replacing reward systems, is addressed in Section 3.4

2.2 An *N*-Dimensional Approach

While it may be possible to construct unified simulations which integrate the different aspects of the comparison into a single measure of social welfare, a numerical project of this sort requires a great deal of care to be accurate.

The intention of this article is to reach an approximate conclusion — an indication of which copyright system is likely to be superior, and upon what that superiority depends.

The route I adopt for achieving this is to make comparisons along the different axes of distinction between DRM-copyright and reward-based alternatives. There are a broad range of functional differences which can be identified between the two choices — some of them obvious, others more subtle. By examining each of these aspects separately, and identifying the approximate magnitude of those which favour one side of the comparison, we can obtain a relative estimate of social welfare under each system.

For reasons of space, the emphasis of my analysis is placed upon those issues which have received less treatment in the literature on the economics of copyright *per se*.

3 Elements of the Comparison

3.1 The Cost of Security Technology

DRM-based copyright and publicly funded alternatives require very different sorts of technological and security infrastructure. The complexity, price and risk implications of these technologies is an important factor, which, for the most part, the literature does not adequately address.

As this section will explain, there are fundamental reasons why we should expect a technological copyright system which relies on exclusion for payment, to cost much more than one which separates the issues of providing access and raising funds. I will also provide two back-of-the-envelope methods for estimating how expensive the former approach might actually be.

Firstly, it is important to acknowledge that there are significant costs associated with the technological infrastructure for operating a publicly funded reward system for authors and artists. These costs include operating secure servers which catalogue works and track their usage volumes; authenticating users; and preventing users' identities from being hijacked for nefarious purposes.

In total, the cost of infrastructure for a reward system may be substantial. But in the comparison against exclusive rights, this cost does not count either way, because equivalent infrastructure is also required for a DRM-based system of copyright. These costs are associated with making the network operate the way almost all of the participants agree that it should, and preventing defectors from subverting the system.

DRM also requires another, more ambitious kind of infrastructure — one which ensures that the copyright system works, despite the fact that most users would be better off if it didn't.

3.1.1 Why DRM costs

The underlying theme in security is risk management. DRM is an exercise in computer security, and hence, it is also fundamentally about managing risks. Just as in other areas of security, the tools used to control risk can be technological or "social", and they can involve either reinforcing existing practices or re-writing the rulebook to achieve systemic advantages.

The historical practice of copyright has been to sell the (public) information good, *embodied* in an excludable physical object. This practice is, as Barlow (1994) emphasised, radically challenged when consumers have the power to extract the work from its physical "bottle" and manipulate the information inside. The DRM solution is to build a new, stronger bottle out of cryptography and tamper-resistant hardware. Leaks are chased with traitor-tracing algorithms.

But as risk management strategies go, it's far from optimal, because it places a weak point, a physical location at which the infrastructure can be compromised (and large amounts of information liberated), inside every single consumer's home. Each user has little or no incentive to respect the system of copyright enforcement. And, in the long run, if there are robust alternative channels of distribution — FastTrack, Gnutella, Freenet, or their successors — it only takes *one* user to crack the security layer, only one untraceable leak, for massive economic losses to follow.

Many commentators have come to the conclusion that Digital Rights/Restrictions Management is ultimately impossible[14]. It's not yet completely clear that this is the case, most particularly if copyright owners are just trying to use it to retain a slice of their monopoly rent[15]. But in the long run, even for this purpose, DRM is going to need to be very good to stave off the inevitable security compromises that can liberate huge slices of their catalogue.

So, there are good reasons to believe that functional DRM will be expensive. But how much does "expensive" cost? I have identified two methods for "back of the envelope" calculations of the overall cost of DRM. Interestingly, despite very different methodologies, they produce remarkably similar answers.

3.1.2 Extrapolation from existing hardware

The platform for effectively secured DRM (or most of it) exists today. It has been developed by IBM for use in the financial cryptography market. It's called the 4758 cryptographic co-

[14] For a clear, explicit argument of this explicit argument, see (Kelsey and Schneier 1999)

[15] The work of Anderson & Kuhn (1996, 1997) could be interpreted as suggesting that even if rights holders resign themselves to accepting the release of content due to analogue-to-digital conversion, it may still be impossible to maintain a secure network of "trusted systems". But as I will argue in Section 3.1.2, it may be possible to mass-produce devices which are sufficiently tamper-resistant to ward off all low-cost attacks.

processor (Smith and Weingart 1999), and it is *very* difficult to break into. It has protection against physical intrusion[16], as well as changes in temperature, power supply and timing signals. It can detect attackers trying to interfere with it using ionising radiation.

If the device detects an intrusion attempt, it can erase all of the private keys stored inside it immediately[17]. The mini operating system it uses has been verified using automated theorem provers[18].

It is rather unlikely that any info-anarchists have anything in their back sheds that will get the cryptographic keys, or original music, out of a gadget like this, provided it's been developed and deployed properly.

How much does it cost? The currently available version, which is a PCI card that fits inside a desktop computer, sells for around USD $4,000. This is clearly too much for DRM — but the question is, how cheaply could it be made in volume?

Some unofficial estimates undertaken at IBM's Thomas J. Watson Research Center[19] indicate that prices for the existing hardware could be brought down to $500. With some further R&D, it could be miniaturised and made available for $100–$150. If it's possible to develop a single-chip device which has the same kinds of defences (this would be a major research project), then the price might get as low as $20–$30.

I would contend that, if the "trusted systems" model of DRM is going to work, it will be necessary to include something like this in every device that handles copyright subject matter. Each sound card, stereo, and portable music device will need to contain a device which is roughly equivalent to the 4758.

Even if the wholesale price tag for a "trusted system" is $20, the resulting rises in the prices of consumer devices are very significant. This is a substantial reason to prefer publicly funded rewards to DRM-based copyright.

3.1.3 Inference from other computer security costs

Another approach to estimating the price tag for effective DRM, is to examine the costs of computer security in broad terms, and look for realistic inferences about the cost of enforcing copyright. Although this method is subject to greater error than looking at a particular technology like the 4758, it may capture essential points which are otherwise missing. Remarkably, it also produces very similar numbers.

[16] Tampering by intrusion is prevented by a shell of four layers of different, interwoven sensors; each of these is resistant to, or can detect, different mechanical or chemical disassembly techniques.

[17] It also stores the keys in such a way that they do not leave residual memory imprints which are subject to forensic analysis.

[18] A highly publicised attack on this device depended on the fact that not all of the libraries deployed with it had been similarly verified (Clayton and Bond 2002). Once this problem was rectified, the 4758 became significantly more secure.

[19] These figures were obtained by personal communication.

The raw data is a recent survey of 503 organisations' experiences in dealing with computer crime (Power 2002). These organisations reported that a total of USD $375.6 million was lost annually in incidents of a kind which might be applicable to a DRM network[20]. These measurable losses were spread over 44% of the population surveyed, amounting to $1.7 million per organisation affected. The importance of these loss figures is that they provide some indicative *lower bound* for the price of achieving security; if effective protection is cheaper than the expected losses due to security breaches, then most organisations will quickly deploy it.

Adopting this conservative lower bound, fully effective security for a controlled corporate network costs between $750,000 (average measured losses per organisation) and $1.7 million (average loss for organisations which measure their losses). The cost for securing a "trusted" device in a consumer's home might in some respects be higher (since these locations are not controlled by rights holders) and in other respects may be much lower (because some security risks scale with the number of computers/users on the network). The weakest assumption is that costs per device are the same as costs per employee. The average number of employees for the organisations in the survey was about 5000[21]. Hence, if we divide the minimum organisational cost of close-to-bulletproof security ($750,000), by the number of employees per organisation (5000), we obtain a ballpark conservative prediction of effective security costs for an embedded consumer device: USD $150.

Both estimates suggest that, if secure DRM is generally achievable, we should expect it to cost rather sizeable amounts (USD $20–$150) for every device which is integrated into the DRM network. This is an important argument in favour of the alternatives.

3.2 Deadweight Loss

The problem of monopoly-induced deadweight loss is recognised by almost all authors commenting on the economics of copyright[22]. Throughout the literature, it forms the core objection to granting exclusive rights as incentives for the production of public goods.

It is clear that deadweight loss is a major economic weakness in copyright. One important question is whether price discrimination may effectively overcome much of the deadweight loss. There is no doubt that price discrimination can reduce exclusion costs; but the combined effects of arbitrage and a lack of sufficient incentives for low-price versioning (Eckersley 2003, pp. 28–30) prevent price discrimination from solving the problem.

[20] See Power (2002, pp 10–11). This excludes losses from "insider abuse of Net access", laptop theft and denial-of-service attacks

[21] Following (Power 2002, p. 3), and assuming median numbers for each interval, and 15,000 employees for organisations in the 10,000+ category.

[22] There are some exceptions: see, for example (Easterbrook 1999, p. 112). There is also some debate about the semantics of the word "monopoly" (Kitch 2000, Part I) — but regardless of which words are used, the fact remains that in almost all markets constructed by "intellectual property", many consumers who are willing to pay the marginal cost for goods will be denied them. The same cannot be said for other kinds of markets.

Although the problem of deadweight loss exists in the absence of the larger problem of financial inequity — few people purchase all of the cultural goods they would appreciate — the two interact in a particularly troublesome manner. The cost of artificial scarcity is amplified when there are significant numbers of consumers (unemployed, teenagers, those outside the first world) who would benefit from access to information goods, but who can express little financial demand for them.

While its precise extent varies, deadweight loss must certainly be counted as a significant factor in favour of alternatives to DRM.

3.3 The Distortionary Cost of Taxation

A major issue which must be considered in comparing exclusion-based models (such as exclusive copyright) with publicly funded alternatives is that raising taxation may have deleterious side-effects.

Any tax formula used to fund public goods can be decomposed into a component which is "benefit offsetting" (citizens are left indifferent after the introduction of the tax and the creation of the public good), and a component which is both *distortionary*[23] and redistributive (Kaplow 1996). An ideal tax system must both fund public goods which pass a cost-benefit test, and perform redistribution which is "socially desirable", although this objective is inevitably subjective.

The nature of the actual distortions and redistributions caused by various digital copyright systems are quite subtle.

DRM-based systems are likely to have some distortionary implications as a result of their employing versioning for price discrimination purposes. Distortions may also be caused whenever exclusive rights lead to market power.

Public funding alternatives have distortionary consequences which depend greatly on the incidence of the taxation involved. Simple levies will be sub-optimal to the extent that taxable purchases of hardware or bandwidth, do not predict users' valuations of information goods. For example, a per-megabyte tax on internet traffic creates incentives to switch from high-bandwidth to low-bandwidth media. A charge on hardware may discourage the use of that hardware for purposes which do not implicate copyright law.

If income taxation is used to fund literary and artistic production, it is likely to affect labour supply. Although Kaplow (1996) has shown that for certain utility functions, public goods can be financed by "benefit offsetting" income taxes, this result is not entirely applicable to copyright

[23] Taxation is said to be *distortionary* when it causes shifts between the production & consumption of one kind of good, and another, causing divergence from the natural state of the market. Distortions are usually expected to decrease social welfare, unless they act to correct *externalities* (side effects of actions, such as the pollution caused by driving a car) or redistribute wealth in a way which increases social welfare.

reward systems. Firstly, preferences are certain to vary within each income bracket[24]. Secondly, the particular utility functions associated with cultural goods may be particularly difficult to offset, because these goods are often complements to leisure.

The best approach may be a combination of the levy and income tax models — such as a progressive tax on Internet connections — which would be benefit-offsetting in a wider variety of situations. Nonetheless, they will inevitably carry distortionary and redistributive impacts.

Do the distortionary side effects of taxation count in favour of the exclusive rights based DRM approach? Under certain assumptions, the answer is certainly 'yes'. If a society with a previously optimised taxation system[25] is faced with the choice between technologically enforced copyright and an alternative, then distortions should be weighed in favour of DRM.

But in real societies, the answer may be quite different. There is strong evidence to suggest that ordinary democratic processes produce systematically inadequate degrees of wealth redistribution (Brooks 2003). Under these conditions, Kaplow's argument — that negative distortions are accompanied by positive redistributive effects — is likely to hold. Publicly funded reward systems thus address two issues: the free rider for information production, and the need for stable institutions which produce an equitable redistribution of wealth. The side-effects of taxation produce a change in social welfare, but this change is not necessarily negative.

3.4 Information Revelation

In order for any resource allocation mechanism (including a market), to function properly, information about what people value, must somehow be used in making decisions about what is to be produced. Neither copyright, nor the alternatives, do this perfectly. But if one model were to employ more complete information for resource allocation, this would be grounds for preferring that model.

3.4.1 Transparency

Many information goods display a lack of what DeLong & Froomkin (2000) have called "transparency" (Arrow 1962; Takeyama 2002). In artistic goods, this is associated with higher search costs, and *ex ante* signalling, which degrades the quality of information which the market uses to reward artists and publishers.

Many people purchase CDs without having first heard all of the music on them. In many cases, if they had listened to that CD (and others) first, their purchase decisions would have been better.

[24] Those within each salary range who place greater value upon cultural goods are thus relatively favoured by a public funding system. The consequences of discounting this fact are difficult to determine, without knowing to what extent demand for cultural goods is endogenous.

[25] That is, a taxation system which trades of the redistributive benefits off income tax with the goal of minimising harmful distortions, to maximise social welfare.

Similarly, many people purchase best-selling books on the strength of recommendations; but they might afterwards wish that their dollars could be given to other artists.

Amongst alternatives to copyright, the simpler public funding models, which (for example) allocate fixed rewards for each download, are also subject to these problems. However, as I have argued (Eckersley 2003, Section 2.2), it is possible to construct "virtual markets" in which consumers make *ex post* valuation signals. In these cases, lack-of-transparency effects should be counted against copyright.

3.4.2 Reliance on non-payment signals

While copyright may suffer from *ex ante facto* signalling, the alternatives are likely to suffer from signals which are unclear. In a copyright market, price setting strategies allow for some variation in the message sent by each user; all other things being equal, consumers will be willing to pay more for works which they enjoy more.

If rewards are determined solely by counting downloads, the information returned is binary: each person either downloaded it, or they did not[26]. This may encourage the same kind of "lowest common denominator" production which DeLong & Froomkin (2000) identify as following from the advertising-based business models common in the television industry.

The situation again improves when alternatives to copyright allow consumers to "vote" on values, rather than simply measuring the number of copies distributed. In this model, consumers have a great deal of control over the signals they send, and are able to correctly report the benefit they obtain from digital cultural goods[27].

If there are imperfections in this process, then they are likely to result from skewed or misrepresentative voting behaviour. Users not voting is only a problem if there are certain sociological demographics who systematically participate more than others, relative to their demand for cultural goods. This effect, if it occurred, would be unambiguously cultural; sociological methods would be required to understand it. Consumers have a clear, but not strong, incentive to vote for the artists they appreciate. And, to the extent that voters are "self selecting", they are signalling higher valuations for cultural goods.

Although a well designed reward system which employs non-payment signalling may be effective, and may even have some advantages over copyright/sales mechanisms[28], it is safest to assume that DRM has some informational advantages in this area.

[26] Note that if people are somehow allowed to download things several times, and each download is counted, the system is now a voting mechanism — although possibly one in which voting early and often is easy.

[27] This of course depends on the fact that the signals are easy to send correctly. This issue is addressed in (Eckersley 2003).

[28] One advantage is that piracy does not affect the quality of signals; another is discussed in Section 3.7.

3.4.3 One-user-one-vote effects

Finally, there is a set of information revelation related effects which apply to virtual markets and some other reward systems, that result from their "egalitarian" structure. The ordinary copyright marketplace is a "one dollar, one vote" environment. Almost all of the models for publicly funded rewards replace that with a "one user, one vote" arrangement[29].

What are the consequences of this difference?

One result is a form of cultural wealth redistribution whose consequences are difficult to predict, because it depends on the relationship between wealth and preferences.

Another, more problematic, result is that some goods — those which have a very high value, for a small number of consumers — may not be sufficiently remunerated, because the reward system constrains the contribution of each of those consumers. In a world with many information goods of this sort, a "one user, one vote" mechanism would be seriously flawed.

There is, however, a remarkable empirical law which applies to cultural goods, and which suggests that we do not live in such a world:

The Law of Near-Constant Prices

The variation in prices for non-rivalrous cultural goods is generally much lower than variation in the production costs of these goods. [30]

This can easily be seen, for example, in the film industry. Consider two highly successful films: *The Blair Witch Project* cost US $22,000 to produce, while *Titanic*, finished two years earlier, cost $200 million[31]. This is a difference of four orders of magnitude, but the retail prices of cinematographic works differ by less than a single order of magnitude[32].

Similar disparities can be observed in the market for musical compositions and recordings. The situation is less clear for writing, where retail prices vary more widely, and production costs are less extreme; but to a large extent, these distinctions are caused by variations in the cost of the physical object itself (books are diverse, rivalrous objects), and the fact that the publishing industry really supplies many separate kinds of goods (books of photography have little in common with novellas).

[29] Of course, people do not actually vote in a marketplace, or in most of the alternatives. Nonetheless, purchase or consumption decisions for cultural goods are functionally similar to votes about what kinds of cultural production should be rewarded.

[30] In contrast, the prices of rivalrous cultural goods, such as paintings, vary much more than their production costs.

[31] The figure for *The Blair Witch Project* comes from the film's DVD liner notes. The cost for *Titanic* comes from the online encyclopedia *Wikipedia* (`http://www.wikipedia.org`).

[32] Some readers may object on the grounds that the cost of *The Blair Witch Project* should include the $1.5 million spent on marketing (Wired News 1999) — in contrast, the marketing budget for *Titanic* was about $60 million (see `http://www.petesmoviepage.com/Titanic/news2.shtml`, mirrored at `archive.org`). Even if this is taken into account, there is still over two orders of magnitude difference between variation in production cost and variation in retail price.

The empirical phenomenon of near-constant prices suggests quite strongly that the kinds of culture which would be threatened by egalitarian reward systems, are not the same kinds of culture which can be constructed from purely digital information and "napsterized". Only in cases where retail prices vary widely, does technologically enforced copyright have an advantage due to "one user, one vote" effects.

3.4.4 Information revelation as a whole

The information revelation properties of the two systems vary in rather subtle ways. Dependence on non-payment signals, and the "one user, one vote" nature of reward systems, should be counted in favour of exclusive rights. On the other hand, the ability to overcome transparency is a significant point in favour of *ex post* reward systems. Conservatively, it is safest to conclude that DRM-based copyright probably has some small advantage in informational terms.

3.5 Transaction Costs

There are several different kinds of transaction costs which occur in cultural information economies.

There are classes of transactions which are necessary to identify and pay artists and authors. These are likely to be very similar under DRM or any alternative. Related to these is the process of negotiations for splitting credit amongst various contributors to the process of cultural production. Again, there is unlikely to be any systematic differences in these costs under different copyright systems.

Of more interest for the present comparison, are the transaction costs associated with user-side rights clearance. By definition, these costs only exist when users require permission to exercise various exclusive rights which are attached to works, and may thus count against DRM.

In most situations, these costs remain unproblematic, because rights holders address them in their networks of distribution and sale. It is only when third-party users are attempting to exercise rights in unusual ways, that these transaction costs arise.

Perhaps the most immediate example of this is the operation of search engines, caches and other indexing tools. While some of these activities are covered by limitations and exceptions to copyright, there are major problems with access. The owners of collections managed using technical protection measures will inevitably be reluctant to allow third parties complete access for the purposes of indexing and analysis — particularly because the best indexing and analysis services make large fractions of the "full text" directly available to users.

Absent transaction costs, it would be possible for the providers of indexing and analysis tools to negotiate arrangements with both rights holders and users, to provide the best possible search and meta-search facilities.

Although this problem is a probably not as significant as the primary costs of exclusion, it is certainly a factor which must be counted against DRM-copyright, and in favour of the alternatives.

3.6 Management Issues

Both copyright, and publicly funded alternatives, require important parameters to be set through governmental management.

In copyright systems, a balance must be struck between the incentives produced by exclusive rights, and the benefits which can be obtained by maximising distribution (and the creation of derivative works). The political processes required to strike this balance correctly are seriously hampered by collective action problems (Boyle 1997).

In publicly funded reward systems, one important parameter — the total level of funding — is likely to be set through policy. Sustained, inadequate, levels of funding would certainly have a negative impact on social welfare. But, in contrast to the exclusive rights case, the interests of the key lobby groups and the public as a whole coincide, reducing the risk of management failure.

I therefore argue that management errors produce a relatively constant negative effect under DRM, and a stronger (but less likely) negative effect under a system of public funding. It is not clear that either system should be preferred on these grounds.

3.7 Contests over the Definition of "Cultural Space"

There is evidence that demand for many copyright works is endogenous. Nadel (2003) has argued that contests in the definition of cultural space may result in a wasteful dissipation of resources. Are these sources of market failure, and do any of the alternatives genuinely have the potential to reduce them?

It seems that many cultural goods create a subtle system of network externalities. Ordinary human interactions are regularly filled with references to popular (or niche) culture, and our perspectives on the word are unavoidably coloured by the art and entertainment we consume. Naturally, exploiting networks around their products is an important strategy for copyright owners.

As Nadel points out, the existence of cultural externalities results in "marketing contests" over the definition of cultural space. There are innumerable tunes which are a capable of capturing the human psyche, but only a few of them will top the charts.

These expensive advertising contests are very similar to the "race to invent" or "common pool problem"[33] found in patent systems. There is no general solution to this problem which is not a function of the conditions in the particular marketplace in question.

[33] See Wright (1983) for a survey and discussion of possible solutions.

On one level, exclusionary copyright and publicly funded rewards are likely to suffer equally from races to define cultural space. Because the public funding schemes mimic market returns, there will be an excess of investment in marketing a few costly cultural products, while society would be better off with more diverse investment in cultural creation — combined with an open market for works, the best of which will evolve to define "cultural space".

Optimal cultural production can occur only when the rewards for information goods are not only an increasing function of demand or value for the good, but also a decreasing function of the resources dissipated in embedding them in networks of cultural externalities. This is not the case in an ordinary marketplace.

It could conceivably occur in a "virtual market" system, but only if a culture of voting evolves to achieve this end. If users reward works according to both their subjective quality *and* the fact that they are not supported by strong marketing, then resource dissipation through races will decrease.

I would hesitate to count contests over cultural space as a clear point in favour of publicly funded reward systems. At the same time, I would emphasise that there are potential benefits from encouraging the explicitly democratic construction of culture, which should be considered together with the drawbacks discussed in Section 3.4.2.

3.8 Disclaimer: Some Media are More Equal than Others

Copyright law affords very broad groups of information goods the same legal privileges. Writing, musical composition, visual art and computer programs receive essentially equivalent breadth and depth of exclusive "protection". Other kinds of subject matter, such as film, photographs and recordings receive slightly different (and generally weaker) copyright monopolies.

When attempting to compare DRM-based copyright to alternatives, however, the differences between media become increasingly important. It is quite possible that watermarking algorithms, for example, will be much more feasible for cinematographic works than for writing. Similarly, analogue-to-digital recording of music is much easier than the re-recording of film.

In addition, the strength of "first mover" advantages varies greatly from medium to medium. Musicians are able to sell tickets on the strength of their recordings; cinemas will continue to be important even if there is widespread piracy of DVDs. In contrast, if display technology improves to the point where digital books are as readable as paper ones, authors will find that their principal source of revenue — royalties — is seriously threatened.

When these long-term effects are combined, the digital copyright problem appears most serious for writers — and the advantages of reward systems are strongest in that medium. I have attempted to make my approach here as general as possible, but the info-economic peculiarities of the different media are certainly deserving of further consideration.

4 Conclusions

This paper has enumerated and examined various axes of economic comparison between technologically-enforced exclusive rights and alternatives based on public funding.

Of those dimensions, I identified three (expected management efficiency, distortionary taxation and contests over cultural space) where neither system could claim a clear advantage. There were two factors where one model had a small efficiency lead (information revelation, for DRM, and transaction costs, for the alternatives). Finally, there were two areas where public funding models were clearly greatly superior — deadweight loss, and technological infrastructure costs.

After cancelling terms, there are two major economic factors which favour reward systems, without significant countervailing effects. Even if we assume that there are flaws in the analysis, the strength of this result would suggest that it is unlikely to be reversed; the outcome gives a high degree of certainty that a utilitarian choice between the two systems must favour replacing technologically enforced consumer copyright with one of the alternatives.

What practical corollaries follow from this result? In the short term, I would suggest cautious optimism towards proposals for levy-based remuneration for private copying, provided that these schemes come with licenses or exemptions excusing the covered conduct. Attempts should be made to ensure that the taxation used to fund these schemes is progressive as well as levy-based, and to grant consumers a direct role in deteriming the way rewards are allocated.

As technological developments continue to challenge the enforceability of existing copyright laws[34], the best policy response is not the enactment of new legislation to mandate the use of Digital Rights/Restrictions Management technologies. Instead, legislatures would be well advised to begin considering how they might implement alternatives.

5 Acknowledgements

I would like to thank Andrew Clausen and Ryan Lampe, for their helpful comments on this text — and the many other people who helped the ideas contained in it to evolve.

[34] See, for example, the summary judgement on *MGM v. Grokster* in the US District Court for the Central District of California (`http://www.eff.org/IP/P2P/MGM_v_Grokster/030425_order_on_motions.php`), and the Amsterdam Court of Appeal's appellate judgement in *BUMA & STEMRA v. Kazaa* in the Netherlands (`http://www.eff.org/IP/P2P/BUMA_v_Kazaa/20020328_kazaa_appeal_judgment.html`).

References

Anderson, R. and M. Kuhn (1996). Tamper resistance - a cautionary note. In *Proc. Second USENIX Workshop on Electronic Commerce*. USENIX Association.
http://www.cl.cam.ac.uk/users/rja14/tamper.html.

Anderson, R. and M. Kuhn (1997). Low-cost attacks on tamper resistant devices. In *Proc. 5th International Workshop in Security Protocols*, LNCS, pp. 125–136. Springer Verlag.
http://www.cl.cam.ac.uk/∼mgk25/tamper2.pdf.

Arrow, K. J. (1962). Economic welfare and the allocation of resources for invention. In *The Rate and Direction of Inventive Activity*. NBER / Princeton Univ. Press.

Barlow, J. P. (1994, March). The economy of ideas : a framework for rethinking patents and copyrights in the digital age (everything you know about intellectual property is wrong). *Wired Magazine 2.03*.
http://www.wired.com/wired/archive/2.03/economy.ideas_pr.html.

Boyle, J. (1997). A politics of intellectual property: Environmentalism for the net? *Duke Law Journal 47*, 87–116.
http://www.wcl.american.edu/pub/faculty/boyle/intprop.htm.

Breyer, S. (1970, December). The uneasy case for copyright: A study of copyright in books, photocopies and computer programs. *Harvard Law Review 84*(2), 281–351.

Brooks, D. (2003, 12th of January,). The Triumph of Hope Over Self-Interest. *New York Times*.

Bucovetsky, S. (1991). Choosing tax rates and expenditure levels using majority rule. *Journal of Public Economics 46*, 113–131.

Calandrillo, S. P. (1998). An economic analysis of intellectual property rights: justifications and problems of exclusive rights, incentives to generate information, and the alternative of a government-run reward system. *Fordham Intellectual Property, Media and Entertainment Law Journal 9*(1), 301–360.

Campbell, D. E. (1987). *Resource Allocation Mechanisms*. Cambridge Univ. Press.

Clayton, R. and M. Bond (2002). Experience Using a Low-Cost FPGA Design to Crack DES Keys. In *Proc. of the Workshop on Cryptographic Hardware and Embedded Systems (CHES 2002)*.
http://www.cl.cam.ac.uk/∼rnc1/descrack/DEScracker.html.

Cornes, R. and T. Sandler (1996). *The theory of externalities, public goods, and club goods* (second ed.). Cambridge Univ. Press.

Davies, G. and M. E. Hung (1993). *Music and video private copying: An international survey of the problem and the law*. Sweet & Maxwell.

de Trenquayle, P. (1997). An extension of Bowen's dynamic voting rule to many dimensions. *Social Choice and Welfare 15*(1), 141–159.

DeLong, J. B. and A. M. Froomkin (2000, February). Speculative microeconomics for tomorrow's economy. *First Monday 5*(2).
http://www.firstmonday.dk/issues/issue5_2/delong/.

Easterbrook, F. H. (1999). Cyberspace versus property law? *Texas Review of Law and Politics 103*(4).

Eckersley, P. (2003, February). Virtual Markets for Virtual Goods: An Alternative Conception

of Digital Copyright. IPRIA Working Paper 02-03.
http://www.cs.mu.oz.au/∼pde/writing/virtualmarkets.pdf.

Fisher, W. W. (2000). Digital music: Problems and possibilities. In *Proc. of the "A Free Information Ecology in the Digital Environment" Conference*. Information Law Institute, New York University.
http://www.law.harvard.edu/Academic_Affairs/coursepages/ (...)
tfisher/Music.html.

Foley, J. H. (2001). Enter the library: Creating a digital lending right. *Connecticut Journal of International Law 16*, 369.

Harrison, P. (2002). The rational street performer protocol.
http://yoyo.cc.monash.edu.au/∼pfh/rational.html.

Kaplow, L. (1996). The optimal supply of public goods and the distortionary cost of taxation. *National Tax Journal 49*(4), 513–533.

Kelsey, J. and B. Schneier (1999, June). The street performer protocol and digital copyright. *First Monday 4*(6).
http://www.firstmonday.dk/issues/issue4_6/kelsey/.

King, S. (2001). The Plant Income/Expense Report Through 12/31/00.
http://www.stephenking.com/PlantNumbers_010101.html.

Kitch, E. W. (2000). Elementary and persistent errors in the economic analysis of intellectual property. *Vanderbilt Law Review 53*(6), 1727–1741.
http://law.vanderbilt.edu/lawreview/vol536/kitch.pdf.

Ku, R. S. R. (2002). The creative destruction of copyright: Napster and the new economics of digital technology. *University of Chicago Law Review 69*, 269–324.

Landes, W. M. and R. A. Posner (1989). An economic analysis of copyright law. *Journal of Legal Studies 18*(2), 325–363.

Lunney, G. S. (2001). The Death of Copyright: Digital Technology, Private Copying, and the Digital Millenium Copyright Act. *Virginia Law Review 87*, 813–920.

MPAA (2002). Motion Picture Association of America, Content Protection Status Report. Filed with the US Senate Judiciary Committee.
http://judiciary.senate.gov/special/content_protection.pdf.

Nadel, M. S. (2003, January). Questioning the Economic Justification for (and thus the Constitutionality of) Copyright Law's Prohibition Against Unauthorized Copying: §106. AEI-Brookings Join Center, Related Publication 03-1.
http://www.aei.brookings.org/publications/abstract.php? (...)
pid=302&aei_brookings=e0e42d86e9133e582c86da044e8233d2.

New York Times (2000, 1st of December,). King Closure. Editorial.
http://www.nytimes.com/2000/12/01/opinion/01FRI3.html.

Power, R. (2002). CSI/FBI Computer Crime and Security Survey. *Computer Security Issues & Trends 8*(1).
http://www.gocsi.com/press/20020407.html.

Rasch, C. (2001, June). The wall street performer protocol: Using software completion bonds to fund open source software development. *First Monday 6*(6).
http://www.firstmonday.dk/issues/issue6_6/rasch/.

Roberts, D. J. (1979). Incentives in planning procedures for the provision of public goods. *Review of Economic Studies 46*, 283–292.

Rothman, D. H. (1992, 6th of July). Information access for all. *Computerworld*. http://www.teleread.org/computerworld.htm.

Schulman, B. M. (1999). The Song Heard 'Round the World: The Copyright Implications of MP3s and the Future of Digital Music. *Harvard Journal of Law and Technology 12*(3), 590–643.

Shavell, S. and T. van Ypersele (1998). Rewards versus intellectual property rights. Discussion Paper 12/98, John M. Olin Center for Law, Economics and Business; later published in *44 J. Law Econ.* 525 (2001). http://www.law.harvard.edu/programs/olin_center/papers/pdf/246.pdf.

Smith, S. W. and S. Weingart (1999). Building a high-performance, programmable secure coprocessor. *Computer Networks 31*, 831–860. http://www.research.ibm.com/secure_systems/papers/arch.pdf.

Stallman, R. M. (1985). The GNU Manifesto. Free Software Foundation. http://www.gnu.org/gnu/manifesto.html.

Stefik, M. (1999). *The Internet Edge: Social, Technical and Legal Challenges for a Networked World*. MIT Press.

Takeyama, L. N. (2002, June). Piracy, assymetrical information, and product quality revelation. Society for Economic Research on Copyright Issues. http://www.serci.org/2002/takeyama.pdf.

von Lewinski, S. (1992). Public lending right: A general and comparative survey of the existing systems in law and practise. *Revue Internationale du Droit D'Auteur 154*(3), 3–81.

Watt, R. (2000). *Copyright and Economic Theory*. Edward Elgar.

Wired News (1999, 2nd of August,). *Blair Witch* casts a strong spell. http://www.wired.com/news/culture/0,1284,21045,00.html.

Wright, B. D. (1983, September). The economics of invention incentives: Patents, prizes and research contracts. *American Economic Review 73*(4), 691–707.