Digital Copyright
&
The Alternatives

An Interdisciplinary Inquiry

Peter Daniel Eckersley

April 2011, revised January 2012

Submitted in total fulfillment of the requirements of the degree of Doctor of Philosophy

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

Abstract

This thesis compares the existing institution of digital copyright to possible alternatives. Four policy regimes are considered: (1) the status quo 'weak copyright' regime, in which exclusive rights and piracy exist in parallel and are both economically significant; (2) a 'strong copyright' regime, in which Digital Rights Management (DRM) technologies and other forms of enforcement prevent piracy from being economically significant; (3) an 'information anarchy' regime, in which copyright is not a meaningful restraint on non-commercial copying and sharing of works; and (4) a 'virtual market' regime, in which file sharing is legalised, and public funding is used to pay artists and authors, based on decentralised measures of the the popularity and value produced by each copyright work (download counts, usage measurements, or voting).

These regimes are studied for their technical feasibility, and compared for ethical and economic desirability. Different regimes may prevail at different times, and in different places and industry sectors, even if copyright law is the same. All of the regimes are feasible in at least some cases, but their practicality and relative merit turns out to vary surprisingly across copyright-based industries.

The normative criteria that are applied include the levels of artificial scarcity that the regimes impose on information goods; the financial incentives they offer to information producers; the kinds of transaction costs they involve; the price of their technological and non-technological infrastructure; and, in the case of the virtual market, the taxation overheads that it would imply.

Artificial scarcity costs are found to be very high, especially in the music industry: the social value of music would rise by 55–98% if copyright law were abolished, and 18–32% would be lost if present laws were fully enforced. But copyright does serve an important purpose in providing incentives and fair reward to authors, and this appears to *probably* outweigh the arguments for information anarchy. However it is shown that virtual markets could probably perform this function better than exclusive rights-based copyright. Such public funding proposals do have overhead costs associated with taxation, however, and those are quantified.

It is found that there is a strong case for experimenting with virtual markets in some copyright industries, particularly for music, books, and websites. It is also concluded that the international trend towards the 'harmonisation' of copyright laws has been a mistake: the best answer to disruptive and transformative technologies is to experiment with a broad range of regulatory responses, but international treaties on copyright have harmed this type of regulatory biodiversity.

Declaration

This thesis is entirely my own work, except of course where I quote, cite and/or review the work of other authors. A summary of original contributions may be found in Section 1.5. The thesis is 99,776 words long, excluding the Glossary and Bibliography.

Signed:

Acknowledgements

Looking back, I am indebted to a great many people for discussions, feedback and suggestions over the course of this lengthy research project. I will not compile a list because I fear it would simultaneously be too long and too incomplete. Particular thanks go to my supervisors, and to those of you who reviewed chapters as they were completed.

Contents

Table of Acronyms

ACS — Alternative Compensation System

ADSL — Asymmetric Digital Subscriber Line

API — Application Programming Interface

ARIA — Australian Recording Industry Association

CC — Copy Control or Creative Commons

CD — Compact Disc

CD-R — Compact Disc (Recordable)

CE — Consumer Electronics

CODEC/codec — COder/DECoder

CV — Contingent Valuation

DMCA — Digital Millennium Copyright Act

DOI — Digital Object Identifier

DRM — Digital Rights Management or Digital Restrictions Management

DVD — Digital Video Disc

DWL — Deadweight Loss

EFF — Electronic Frontier Foundation

EUCD — European Union Copyright Directive

EULA — End-User License Agreement

FTP — File Transfer Protocol

GNU — GNU's Not Unix

GPL — GNU General Public License

HTTP — Hypertext Transfer Protocol

IFPI — International Federation of the Phonographic Industry

IP — Intellectual property and Internet Protocol

ISP — Internet Service Provider

LAN — Local Area Network

NLP — Natural Language Programming

M4P — MPEG-4 Part 14 (with Fairplay DRM)

MCF — Marginal Cost of Funds

MEB — Marginal Excess Burden

MIH — Metadata-Independent Hash

MPAA — Motion Picture Association of America

MPEG — Motion Picture Experts Group

MP3 — MPEG-1 Audio Layer 3

OS — Operating System

PC — Personal Computer

P2P — Peer-to-Peer

RAID — Redundant Array of Inexpensive Disks

RIAA — Recording Industry Association of America

SEC — Securities and Exchange Commission

SPP — Street Performer Protocol

SSH — Secure Shell

TCP — Transmission Control Protocol

TC — "Trusted" or "Treacherous" computing

TPM — Technical Protection Measure and Trusted Platform Module

TRIPS — Agreement on Trade-Related aspects of Intellectual Property

TTP — Trusted Third Party

UDP — Universal Datagram Protocol

URL — Uniform Resource Locator

USB — Universal Serial Bus

USTR — [office of the] United States Trade Representative

WAN — Wide-Area Network

WMA — Windows Media Audio

WCT — WIPO Copyright Treaty

WTO — World Trade Organisation

WIPO — World Intellectual Property Organisation

WPPT — WIPO Performances and Phonograms Treacherous

VCR — Video Cassette Recorder

VFAT — Virtual File Allocation Table

Glossary

Normally, one would not include a glossary of technical terms in a PhD thesis. After all, such terms should be common knowledge amongst readers of research literature, and search engines, encyclopedias, and specialised dictionaries exist to fill any gaps. Since this thesis is inter-disciplinary, however, it may be helpful to collect a number of relevant definitions here. The main purpose is to alert readers from one disciplinary background when they encounter terminology from another, since on occasion such terms may be mistaken for ordinary or unusual English. Brief definitions are included for the reader's convenience.

anonymity — a property in a computer system or network in which the particular users or machines conducting a transaction cannot be determined precisely. The level of imprecision must be large enough, and is often measured through an "anonymity set size", the number of possible actors that could have been the author of the transaction.

assumptions — when a mathematical model is used to study the real world, the assumptions are the link between the world and the model; they simplify reality so that it can be described formally, but if they are poorly selected, they may also make the description misleading.

asymmetric cypher — a cryptographic system in which the resources required to produce and interpret the cryptographic information are different.

broadcast encryption — a cryptographic system in which a message is encoded in such a way that a large (but controlled) set of people will be able to decrypt it, but those outside the set will not.

bundling — the practice of selling several different goods together, even though they could have been sold separately.

common knowledge — in modal logic and game theory, knowledge which is not only known to all parties but all parties know the others know, that the others know they know, etc. Common knowledge can be achieved by making a statement in public.

compromise — to breach security measures on a computer system.

contingent valuation — a survey-based method for estimating the value of public goods by asking participants how much they would be willing to pay to obtain the good, or how much they would be willing to accept as compensation for its loss.

consequentialism — a family of ethical theories which hold that, when choosing from a set of options (such sets of actions, rules, or institutions), one should select the option which leads to the best overall consequences.

corpus — a standardised body of text or documents to be analysed by computer programs.

crowding out — a psychological phenomenon in which the creation of incentives for some purpose actually has the countervailing consequence of discouraging voluntary performance of the incentivised act or task.

deadweight loss — a property of an economic situation such that some people could be made better off without harming others; in copyright economics, the term is used particularly to refer to the monopoly deadweight loss which results from prices that are above the marginal cost of production (taxation economists use it to refer to distortionary costs, but I attempt to avoid this usage).

demand — for an individual, what they are willing to pay for a particular good or service; in aggregation, a *demand curve* expresses how many units of a good or service would be purchased as a function of the market price.

denial of service — an attack on a computer system or service which prevents it from performing its normal function (rather than changing or gaining control of it).

digital rights management — a very broad term for hardware and/or software which enforces copyright or copyright-like restrictions on what users can do with digital material.

distortion — the effect of a tax which induces consumers to change their behaviour to reduce the tax; measured relative to what it would have been if they had simply paid the amount of.

dominant strategy — in game theory, a strategy is *dominant* if it is the best thing a player can do, regardless of what other players are doing.

endogenous — in economic theory, something is said to be *endogenous* when it results from the internal operation of a model; the opposite is *exogenous*.

endogenous demand — demand for goods that is dependent on what a person has consumed in the past. Acquired tastes and addictions are examples.

excludability — a good is said to be *excludable* when it is feasible for one person to prevent others from enjoying access to it.

exclusive right — the right to exclude all others from particular actions. Exclusive rights are the core building block of both copyright and material property systems.

exogenous — in economic theory, a property is said to be *exogenous* when it is fixed in the specification of a model (or is outside the scope of the model). The opposite is *endogenous*.

exploit — the use of a vulnerability (a bug or a design flaw) in a computer program or system in order to make it do something it was not supposed to do. Also, n. the knowledge of or means to use such a vulnerability.

externality — a side effect of an action or transaction (such as the production or con-

sumption of a good) on second or third parties who are not involved in it.

fingerprint — a short piece of information obtained from a media file, which can be matched against other fingerprints in order to recognise if the underlying audio, video or image file is the same.

first-best — a state of affairs which is in the best conceivable with regard to the variables under discussion.

fixed cost — an initial cost of setup or entry to begin producing in a particular market; for example, the cost of building a factory.

free rider problem — an economic phenomenon in which it is observed that large groups of people cannot produce optimal or even close-to-optimal amounts of public goods through voluntary cooperation or market-based mechanisms, because individuals face strong incentives to enjoy these goods without contributing to them.

free software — software which users are free to use, copy, modify and redistribute. All but synonymous with "open source software".

hack — v. to solve a problem by unexpected means. A hacked solution may be extremely inelegant and pragmatic, or extremely clever, astounding and pragmatic. The word also came to mean the use of hacking exploits to gain control over others' computer systems.

harmonisation — a process by which different nations' laws and regulatory principles are made more similar (or identical).

hash function — in general, a function which maps arbitrary input values into outputs of finite size, in an "unpredictable" way. These functions are many-to-one, in that each possible output will be produced by many possible inputs. See also fingerprint, metadata-independent hash, secure hash.

hit — a result which matches a query to a database or search index.

honeypot — a computer system which simulates another system that suffers from various security vulnerabilities. If attackers attempt exploit those vulnerabilities, their success is simulated and their methods and subsequent behaviour is monitored.

ID3 tag — the metadata fields included in MP3 files.

incentive-compatible — said of an economic process or mechanism in which participants fare best if they comply with the stated rules of that mechanism.

information — computer science and information theory: the means to determine which of several possible states something (a system or the world) is in; in economics: knowledge about the possibility and/or value of various goods, services, actions or transactions.

information good — a thing which can be placed in a mathematical form consisting

solely of information. Includes texts but not manuscripts, sounds but not vinyl records. Most copyrightable things are information goods.

instrumentalism — in the philosophy of science, the position that ideas are justified by their usefulness in predicting the world; hence, an approach to the law which holds that laws should be judged by their usefulness for serving particular social ends, rather than their "rightness" in some other sense.

integrity measurement — in trusted systems, integrity measures are tests which the trusted platform module may execute to ensure that the content of the system has not been altered in some "unauthorised" way. This may be performed for and certified to a remote party, thereby giving them some control over the sysetem.

market power — in economics, firms are said to have market power when their pricing decisions are able to substantially alter the market-wide price of some good or services. Market power exists in a spectrum from oligopoly to monopoly.

marginal cost of funds — the social cost of raising an extra dollar of taxation, which is typically more than one due to distortions.

marginal excess burden — another measure of distortionary costs of taxation. MCF = MEB + 1.

mechanism — in game theory and economics, a type of game whose rules are chosen to try to achieve a particular social outcome, typically the *incentive-compatible* disclosure of valuations from the participants. These are studied in the field of "mechanism design". Common examples include auctions, and processes that governments may use to decide how much of a public good to produce using taxation revenue.

metadata — data about data; for example, tags saying what a file is, where it came from, the format it is in or sometimes more complicated information such as aspects of what it actually means.

metadata-independent hash — a secure hash function over a media file in which the media file is first stripped of its metadata, ensuring that the output will be constant if the underlying audio/video encoding is constant.

Nash equilibrium — in game theory, a state in which every player is following their best strategy, on the assumption that other players are also following their best strategies (also known as a Nash-Cournot equilibrium).

natural monopoly — a monopoly that results from the high fixed costs of entering a particular market. Supplying water pipes to everyone in a city is a classic example of natural monopoly.

negative — in law, negative rights or freedoms are those which involve freedom from prevention or interference by others — for example, the right to free speech or freedom from sexual harassment.

neighbouring rights — a term used to describe exclusive rights closely related to copyright, such as those that cover performance/sound recording or broadcasts.

network effect or *network externality* — applies to a good whose value to one person is a function of who else uses the good. The classic example of a positive network externality is a telephone system, but negative examples are also possible (such as the use of cars in the face of road congestion).

non-rivalry — a good is said to be non-rivalrous if one person's use of it does not decrease its availability to others.

normative — of or pertaining to what *should* be, as opposed to what actually is (the opposite of *positive*).

open source — see *free software*.

Pareto-optimal — a situation in which no person can be made better off without making someone else worse off.

path dependence — path dependence is phenomenon in which past choices or events have effectively altered the options in the present state of affairs, generally with the connotation that optimal decision-making cannot be conducted with reference to current circumstances.

pirate, piracy — in common usage, these terms are approximately synonymous with copyright infringement. I will generally use them to conote copyright infringement in combination with an individual or cultural attitude that these acts should be permitted or perhaps even regarded as virtuous.

positive (law) — a right or freedom which requires active supply or social support; for example, the right to clean water or the freedom to obtain an education (as opposed to negative freedoms or liberties)

positive (social sciences) — in the social sciences, descriptive (the opposite of *normative* or prescriptive).

private copying scheme — a system of remuneration for artists and/or copyright holders for the reproduction of their works by private citizens. Private copying schemes are usually funded by levies on consumer recording devices and media (ACC 2001).

private key — a piece of secret information held by one party in an asymmetric cryptosystem; other parties require the matching public key.

public lending right — a system in which authors and/or copyright holders are paid for the lending of their works from (usually public) libraries.

public good — a good which is naturally non-excludable and non-rivalrous.

public key — a piece of public information that allows parties to use an asymmetric cryptosystem with the holder of the corresponding private key.

progressive taxation — a tax which is payed disproportionately by those with greater income or wealth.

pseudonymity — a system is said to offer pseudonymity if individuals participate in it using one or more artificial or chosen identities, rather than a legal name or equivalent. The term does not in general connote whether any parties have the ability to link pseudonyms to their holders or not.

rational — in economics, an agent which acts to maximise their utility is said to be rational. Frequently, rationality implies self-interest, but the concept can apply to agents with 'altruistic' utility functions that depend on the conditions of others.

rootkit — a piece of software that gives its installer the ability to hide files and running programs from standard operating system functions, and thereby, from any programs that do not re-implement OS functions at a low enough level to be unaffected.

second-best — in economics or ethics, a situation which is amongst the best *feasible* outcomes, even though better outcomes could be imagined. Second-best optimality is sometimes modelled with respect to some realistic constraints, such as the limits of human reasoning ability or the inevitability of transaction costs.

secondary liability for copyright infringement — applies in situations where a party did not actually infringe copyright themselves, but is found to have had some sort of liability for another's infringement. In Anglo-Australian law, the relevant concept is "authorisation" of infringement. In U.S. law, secondary liability has been found under doctrines of "contributory" liability, "vicarious" liability, and "inducement" of infringement.

secure hash — A secure hash is a cryptographic function that calculates a small, uniquely identifying piece of data from a much larger input. The identifying output data is usually a few hundred bits in size and is referred to as the "hash" of the input.

In semi-formal terms, a secure hash is a function $f(x)$ which (1) maps from bit strings to a finite domain; (2) is easy to compute; (3) is infeasible to invert (*i.e.* given y such that $y = f(x)$, it is infeasible to calculate x); and (4) is collision-resistant (*i.e.* it is infeasible to find a pair of strings a, b such that $f(a) = f(b)$.

sniffing — the use of a computer system to intercept traffic between two other parties. Sniffing can be used to eavesdrop on communications, to steal passwords, or to collect data with which to attack weak cryptosystems.

social welfare — a formal measure of the good or well-being of a group. Social welfare functions can be based on the fulfilment of desires, on happiness, or on other variables.

stopword — a word which is discarded from a search index because it is so common that it occurs in most documents.

substitutability — in economic models, the extent to which obtaining one good A has

the effect of decreasing demand for another good B. Antonym: complementarity.

swarming — technique for downloading files in which different portions of the same file can be obtained from different sources or uploaders, often simultaneously. Consistency is preserved by matching a hash. In the presence of malicious uploaders, structured hashes such as TigerTree are most desirable, because they allow the detection of mismatching file fragments.

symmetric cypher — a cryptographic system in which all participants share certain secrets that allow them to participate.

tamper resistance — computer hardware that has a physical design that prevents.

technical protection measure — a legal term for digital rights management.

transaction cost — either the cost of creating, administering and enforcing economic property rights, *or* just the cost of transferring these rights. See Section 8.1 for further discussion of these definitions.

trusted platform module — a chip in a computer that allows the computer to function as a trusted system.

trusted system — a computer system containing special hardware to facilitate external constraints or control over the code the system runs, or to establish that the code is in a state that is approved by the remote party before data is transmitted to the system.

Turing-complete — a computational system is said to be Turing-complete if it is capable, with the right input, and with sufficient time and storage space, of emulating any other Turing-complete system. The set of Turing-complete systems includes actual computers as well as idealised programming languages, systems of logical rules, certain equations, and other abstract constructions. In the real world, unlimited supplies of time and storage are not forthcoming, but the concept remains meaningful for categorising the kinds of calculations that can be performed by various devices.

utilitarianism — utilitarian ethical theories are variants of consequentialism in which the good is dependent on human pleasure, happiness or fulfilment.

versioning — the deliberate production of several versions of a product (often including inferior versions) in order to facilitate price discrimination.

watermark — in computer science, a modification to a file or document in some (typically invisible or not readily perceptually distinguishable) way in order to make distributed copies of the file recognisable or traceable.

work — in Anglo-Australian copyright law, a subset of copyright subject matter that receives the strongest levels of regulation: literature, musical compositions, computer programs, dramatic and artistic works. In U.S. copyright law, anything that is covered by copyright.

zombie — a computer system which has been compromised (directly or with tools such as worms) by a hacker, who is using a large collection of these systems for nefarious purposes such as sending spam or launch distributed denial-of-service attacks.

Part I

Introduction

Chapter 1

Introduction

Contents

1.1 The Research Question

Copyright is a very broad, property-like system of exclusive rights in expressive information. Copyright is 'digital copyright' when the expression it covers takes the form of data in computers. This thesis aims to determine whether this property system is in fact the best kind of regulatory framework for the production and distribution of digital

information goods.

This is a line of inquiry suggested by the numerous controversies and widespread disobedience that attend modern copyright, and also by the inefficiencies, both visible and hidden, which are effectively required by digital copyright laws as they presently stand. That such inefficiencies exist is beyond reasonable dispute, but a large part of this thesis lies in reckoning their size and the extent to which they might be avoided.

The project has explicit policy goals. At the end, we should expect it to tell us how copyright ought to be enforced in cyberspace — or alternatively, how and why file sharing and other profound infringements of copyright should be legalised.

1.1.1 The means to an answer

Answering the question satisfactorily requires several steps, which I will sketch very briefly here. First, we should know what we are investigating — what "digital copyright" is, and which aspects of it are to be examined.

With a clear model of digital copyright in mind, deciding whether it is in some sense the "best" solution to a problem requires both a definite notion of what it ought to achieve, and then some way of measuring it against that standard. There is a semi-standard answer to this question: copyright is supposed to provide the right incentives for authorship. Economists typically formalise that notion as providing optimal incentives for the production of informational public goods, though of course some commentators differ on the answer or on how to formalise it.

A comparison with alternatives obviously requires that there be some alternatives. To that end, I will spend some time and ink considering a few of those in detail. One possibility is the adoption of/surrender to a deregulated state of "information anarchy" in which digital copyright, or at least a substantial portion of it, is absent. Another is to replace exclusive rights with some "alternative compensation system" (or "ACS") that is based, directly or indirectly, on government funding or another collective source of revenue. A third possibility, not so much an alternative as an upgrade, is to build the existing rules as thoroughly and completely as possible into society's technological infrastructure. We will learn if these alternatives are feasible, and what would be required to make them work.

Finally, the answer to my research question lies in comparing exclusive rights in their present and hypothetically stronger forms, to the posited alternatives of deregulation or public funding. This turns out to be a complex and multi-faceted problem,

and an investigation of it occupies the second half of the thesis. Economic theory and computer science provide important assistance, though no straightforward answers. In any case, such a question cannot be answered exactly — it is too intricate and too imprecise — but strong conclusions can be reached about many dimensions of it. The bottom line to which we will fall is that there is good reason to believe that publicly funded alternative compensation systems would work better than markets based on exclusive rights, for at least some kinds of copyrightable works in at least some contexts. "Reason to believe" is not the same as living proof, but it does constitute a justification for further research, and an argument for practical experimentation.

1.1.2   The public good problem

This thesis is largely — but not entirely — about the best methods for funding the production of informational public goods. That is because the legitimate purpose of copyright is largely — but perhaps not entirely — to produce those public goods.[1]

The reader will recall that public goods are by definition both *non-rivalrous* and *non-excludable*. Non-rivalrous goods are such that the benefit each person gets from the good is undiminished if others also gain access to it: the security of one's homeland, however that might really be achieved, is no less valuable because others also live there. A morsel of food, in contrast, can only be eaten by one person and is therefore rivalrous. Non-excludable goods are those which, once created, cannot be denied to anyone: a painting can easily be placed in a locked room so that only those admitted will enjoy it, but the discovery that the bicycle is a stable, ridable device cannot be properly enjoyed without the world at large also coming to share it.

Neither of these two properties is exactly black-and-white, but anything which possesses a substantial amount of both can reasonably be called a "public good" and studied with economic models thereof. Much traditional copyright subject matter (paper books, vinyl records, sculptures) cannot be described as public goods, because they are embodied in physical objects that are rivalrous and excludable. But the acts that first produce them (authorship, composition, recording of performances, perhaps even sculpting) most certainly are public goods. Digital copies of the works are also naturally public goods.

Economic models for the most part conclude that capitalist markets work well in

---

[1] Some competing positions — such as the claim that copyright should be about fairness, or the view that the free rider problem for authorship should not be "solved" at all — will be discussed in Chapter 7.

providing the services of the butcher, the baker, and the candlestick maker. The things they produce are scarce, and will be valuable for exchange according to society's need for them. The same cannot be said of public goods like authorship.

Left to their own devices, markets do not necessarily encourage the production of public goods, because those who manufacture them cannot sell them (they're not excludable!). Sometimes, there is an incentive to produce public goods by coincidence; perhaps, for instance, an advertiser would pay a film studio to produce a cinematic piece, so that it could be used to promote various products. But there is no guarantee that the sum offered by the advertiser would be sensibly connected to the social value of the film. If the studio is unable to make money from selling tickets, DVDs or downloads, the incentives are likely to be too low, or of the wrong sort.

It would be in every beneficiary's interest for all beneficiaries to contribute towards the cost of public goods, but most individual beneficiaries would prefer to sit back and allow the others to pay for the public goods that they enjoy. The incentives are therefore for the users not to pay for the good. This is of course called the *free rider problem*.

Copyright is what computer programmers might call "hacking" the marketplace to fix the free rider problem — or attempt to fix it. If capitalist markets cannot manufacture non-excludable goods well, then legislate to make them excludable. In the past, that meant the imposition of controls on the use of printing and record pressing machines so that their users might pay royalties to authors (not a trivial project, but an elegant hack perhaps). In the Internet age, copyright requires the imposition of artificial scarcity on files that contain digital books and films and music, so that they can be sold, and paid for. This is a very very complicated piece of hackery.

In either the traditional or the digital setting, the copyright solution is imperfect because artificial scarcity is always in some sense unnecessary scarcity. In the best imaginable world, nobody would be prevented from downloading non-rivalrous digital goods.[2] Physical books and records would be sold for the lowest possible prices that would cover the cost of manufacturing them.[3] Of course, we cannot live in the best imaginable world. The research question can be restated in unconventional terms: "how ugly a hack is digital copyright? Is there a different solution (maybe a hack, maybe something designed from scratch) that works better in practice?"

---

[2]The cost of that exclusion is usually called a *deadweight loss*.

[3]Note that since these are conditions related to the distribution and sale of goods, neither of them addresses how much authors and artists and programmers would be paid for their work.

1.1.3 Aside: An impractical question?

At the outset, there may be an immediate concern about a policy-oriented investigation into whether exclusive rights are the right regime for organising information production. The objection is that such a question must be either purely theoretical or pointlessly quixotic. After all, a system of digital copyright based in exclusive rights has already been globalised; much of it follows logically from pre-digital copyright law, while the rest has been propagated through international treaties and the widespread passage of national legislation. Rather than spending a great deal of effort in trying to determine whether the system is fundamentally the ideal one, the objection runs, we should focus instead on ways to apply, interpret or incrementally improve it.

This preemptive criticism can be answered in several ways. One is to emphasise that the state of the law cannot and should not override concerns that are grounded in economic or technical reality; it is not enough to have laws — they must be harmonised with society and technology. While there is still widespread infringement, it is incorrect to regard digital copyright as an established system that therefore deserves to be treated with conservative deference.

A second point is that copyright historically has had other less proprietarian aspects to it, at least in some countries. Public lending rights and private copying levies (these are essentially specialised forms of public funding for authors, artists and publishers) still exist, and perhaps those parts of copyright systems could be preserved, nurtured, and countenanced as alternatives to the primary mechanism of exclusive rights.

A third response is that the correct way to formulate copyright policy is to begin with objectives and to then look for rules and institutions that achieve them.[4] By and large, the incremental nature of both legislation and jurisprudence follow a very different pattern. While policy analysis does affect the evolution of the law, copyright long ago found a life and a logic of its own that has no guaranteed connection to any satisfactory design goal. Sometimes it even seems that the rules come first and the objectives second. The lopsided politics of intellectual property has exacerbated this problem.[5] In this environment, governments and other regulatory institutions are

---

[4]This is especially true of copyright policy because the nature and implications of copyright regulation change significantly with the development of technology.

[5]On the nature of this politics, see, for example, (Boyle 1997; Braithwaite and Drahos 2000; Litman 2001; Drahos and Braithwaite 2002; Sell 2002). The main plot line is that the industries that benefit from extremely strong copyright are highly concentrated and therefore far more organised than the very diffuse communities which can be harmed by it.

unlikely to revisit fundamental assumptions, and it is the duty of academia to ensure that sufficient thought is given to the premises of copyright.

All of these replies are valid, but it must also be conceded that the original objection has a genuine basis. Copyright law does possess a great deal of inertia, and the industries that are interwoven with it contribute further to that fixity. If we do conclude that digital copyright (or at least certain aspects of it) should be drastically reformed, then there is a separate question of whether and how that might be possible. This is a matter I will reprise at the end of the thesis.

1.2 Motivation

Digital copyright is a powerful and problematic institution, and there are plenty of motivations both for studying it and for judging it against alternatives. That project is important because copyright is enveloped in conflicts which show little sign of resolving themselves. It is important because copyright is a primary source of finance for many kinds of information workers. It is important because digital exclusive rights prohibit billions of useful copies of books, articles and essays, songs, films and computer programs. There are other, secondary, motivations too. It is important because the policy processes which created digital copyright law were insufficiently thoughtful and excessively tendentious for such a complicated task. It is of great interest because it is a test of capitalism under conditions where scarcity, its original cause, is no longer present. And finally, as a researcher, it is attractive because — despite increasing scholarly attention to alternative copyright policies — there are many questions that remain unanswered.

Some of these motivations merit closer examination.

1.2.1 Copyright in Crisis

It seems that wherever one finds copyright in contemporary society, conflict is not far away. Disputes are being played out in the courts, on the streets,[6] in legislatures and cabinets, across the pages of academic journals, through the creation and use of soft-

---

[6]See, for example, (Kirsner 2005; Vance 2003; Sullivan 2004); `http://www.taika.org/~lonewolf/pics/folder/misc/copyright_demonstration_4_10_2005`.

ware, in regulatory agencies,[7] international governmental organisations and standards-setting bodies, in the media and over the hearts and minds of the citizenry.[8]

On their surface, these conflicts look like a fractal of micro-issues. Should computer programmers or technology firms be held liable if their creations or products facilitate widespread copying by the public? If copyright law does place restriction on what software is and is not permitted to do, how will that regulation be enforced? How are the rights of users to be defined and balanced with the rights of cultural industries? How long will the copyright term be? Should ISPs be required to disclose their subscribers' details so that they can be sued for copyright infringement, and under what conditions? Are rights holders allowed to launch software attacks on peer-to-peer networks? Is free/open source software in some sense superior to proprietary software? Should governments be taking active steps to support it? Does the success of the free software imply that the basic principles of intellectual property law are economically counter-productive under real-world circumstances? Are the principles of copyright at odds with those of publicly funded academia, and should scholars be attempting to expel copyright (or at least publishing businesses that do well with it) from the academy?

This discordant staccato is the result of a few deeply set causes: the fact that exclusion from information goods is non-consensual and unstable; the fear amongst those in established copyright industries that online copying will destroy their businesses — and the hope that the Internet, if conquered, will massively boost their profits; incompatibility between legal rules and the nature of new technologies; the side-effects of intellectual property in constraining and in some cases monopolising information production; the slowly increasing presence of civil society opposition to the interest groups that have written intellectual property laws in the past.

We may wonder whether the crisis-bound state of copyright is sustainable. Grandiose predictions of collapsing industries and cultural poverty now seem a little improbable, but so do predictions of any elegant resolution, such as the deployment of DRM systems that are so robust, ubiquitous, and accepted by consumers that piracy can be marginalised. The conflicts around copyright — the War On Copying and the widespread civil disobedience of exclusive rights — are unlikely to abate without dramatic change of some sort.

---

[7]See, for example, the contest over the U.S. Broadcast flag (Crawford 2003); `http://www.eff.org/IP/broadcastflag/`.

[8]See, for example, `http://www.respectcopyrights.org/content.html`, `http://video.google.com/videoplay?docid=1879605023056768074`, `http://www.eff.org/share/`

This fraught state of affairs makes the search for viable regulatory alternatives particularly attractive. Not only is there a thorny knot of problems to attack, but the dysfunctionality of the system as it stands suggests a higher-than-usual chance that such alternatives to established policy may be tested at some point in the future.

1.2.2   What is at stake?

The stakes in copyright policy are high wherever information is produced for a living and wherever it is obtained (or not obtained) at a price. This abundant train of consequences is a second powerful motive for attempts to improve the state of knowledge about plausible institutions for financing cultural and informational industry.

Take literary copyright, for instance. In that domain, exclusive rights are responsible not only for providing much of the income and financial incentives for authors,[9] but at the same time for preventing universal accessibility of their books. Were it not for those rights, the world's libraries would — probably years ago — have digitised the bulk of humanity's published writings. Locating and reading a text, even a very rare text, would be a matter of finding an Internet connection and typing its author or title. Search engines would be able to report, with some authority, which pages of what books discuss particular subjects. We could carry libraries around with us on our laptops and e-books.[10] Copyright has not entirely precluded all of these forms of access; scanning projects by libraries, Google, Amazon, and Internet Archive are finally making some progress. But copyright has undoubtedly slowed them greatly, and it may still be some time before copyright allows those who wish to to build comprehensive libraries of non-public domain digital books.[11]

The relationship between copyright and the World Wide Web is essentially the opposite of what has occurred to date with digital books. The web has flourished, but copyright is not of much assistance in paying web authors. Exclusive rights are naturally inconsistent with the open accessibility that makes the web what it is.[12] While

---

[9]Other traditional financial incentives have also been provided by literary prizes and by positions at universities that allow tenured employees to write literature.

[10]The laptop on which I am writing this could hold well over 150,000 compressed texts of 150,000 words each on its 40Gb hard drive. This is about 1/200th of the size of the Library of Congress (`http://www.loc.gov/about/`; it is uncertain whether people will ever get to carry *that* in their pockets (Toigo 2000).

[11]These issues are discussed further in Chapter 8.

[12]See for instance (Benkler 2006, Chapter 3). The nature of that openness is complex, even for plain hypertext (and more so for the dynamic hypertext that gets called "web 2.0"); it includes basic building

some publishers make money by collecting fees for access to material served by HTTP servers, most of this material isn't really part of the World Wide Web because it cannot be read, linked to, discussed, indexed and searched in the same way that conventional web pages are. Copyright law fails the authors of the web, because it does not offer them a way to get paid that is compatible with their medium of choice. We should consider if there are alternatives which would work better.

Another troubling aspect of digital copyright is the effect it may have on the structure of software production. No case symbolises that so clearly as the Microsoft monopolies. Despite increased competition from Apple and Google, in 2009, for every dollar that Microsoft spent on developing its core products (the operating systems and "office" applications for personal computers), it reaped around three dollars in net profit.[13] For a single firm to preserve a position like this, free from any plausible competition, in such broadly-based global markets, for over a decade, is unprecedented. It should not be presumed that copyright is responsible for the scale of the Microsoft monopolies; there are, of course, monopolies in other industries where intellectual property is not so central. But the role of copyright, as the chief value appropriation mechanism for such an extreme example, must be carefully scrutinised.[14]

In total, copyright industries account for 5 percent of the GDP of developed countries such as Australia or the United States (IIPA 2004). Piracy now provides us with at least half of our music,[15] and probably much more. Depending on how one interprets the situation, societies may be losing billions of dollars a year, either from piracy or from a shortage of it. In fact, both of these claims may be true at the same time, in the sense that both investment in cultural production *and* the extent of distribution of existing works may be too low.

---

blocks in the form of equality of access and experience for different users, as well as more subtle features such as bi-directional linkability. The fact that these features can usually be relied upon is what makes the web such an efficient means of communication. Sometimes, of course, they are absent. But their absence comes at a cost which threatens the very nature of the medium: a friend sends a link to some photos, but the hoster demands subscription before they can be seen; two people discuss a URL but each sees different text there; a search engine reports promising results but the web server asks for $25 before you can see the indexed document. If these cases were the rule rather than the exception, the World Wide Web would no longer deserve its name.

[13]See Microsoft 2009 Annual report, `https://www.microsoft.com/investor/reports/ar09/`; profit margins for the "Client" and "Microsoft business division" units were 73% and 68% respectively. The proportions were down from the early 2000s, when those numbers were closer to 80%.

[14]This thesis will only say a little about software copyright, because it turns out to be such a sufficiently different subject with such unique problems that it will have to be dealt with on its own.

[15]See Section 6.3.1.

1.3   Requirements For A Satisfactory Theory Of Digital Copyright

Any serious attempt to understand the interactions between digital technology and the institutions of copyright must take place at the intersection of many disciplines. This section describes the role of those disciplines in this thesis. Inevitably, some useful tools and bodies of scholarship must be excluded or de-emphasised in order to limit a creeping scope of analysis; the intention at this point is to identify the most important fields and to acknowledge those which are relevant but arguably less so.

Copyright is first and foremost a structure of law. While laws can be fruitfully studied in the abstract, their internal structures probed for logical, moral and even meta-ethical consistency, they are inevitably surrounded by stories which belong at least in part to other disciplines. Around copyright, those stories are particularly rich.

Historically, copyright law has formed the core of a system of regulation for commerce in human communication. Today, it can be argued, copyright has expanded in new directions, so that it also plays a central role in the construction of *functional* technological infrastructures and in the overlapping spheres of private and non-commercial communication.

These complex roles necessarily implicate several fields of study for a satisfactory picture of what was once called "literary and artistic property". Some theory of economics must be employed to understand the consequences of regulating the exchange of such a diverse set of informational entities as "copyright subject matter". In order to reach any normative conclusions, assumptions — perhaps subtle, disjunctive or complex assumptions — must be made about the importance of broad categories of informational objects to human beings and human societies. This is a slippery landscape of inquiry.

Contemporary copyright is also profoundly political, at least as politicised as it has been since its genesis, if not more so. Some theorists go further, to believe that copyright law has fundamental implications for political systems.[16] These facts must be borne in mind, if not explicitly emphasised, when one sets out to study the future of copyright.

As a crowning interdisciplinary complication, copyright is currently in the process

---

[16]Netanel, for example, argues that copyright has a "structural" impact in democratic societies (Netanel 1996, Part IV.C.2) that involves both its facilitation of authorship independent of the patronage of the state and wealthy elites, and its possible role in media conglomeration and homogenisation.

of a collision (or a merger) with something entirely different — let me call it "digital technology" for the time being — which possesses its own disparate origins and intrinsic logical properties. Given the situation of copyright as a nexus of codes and narratives, the researcher must choose their tools with care.

There are three disciplines which I draw on primarily in this work: legal theory, computer science, and economics. They are inescapable: the study of the law because it claims copyright amongst its subjects and much of the research on digital copyright occurs there; computer science because it studies what computers can and cannot do: what processes they can be made to serve with what difficulty; economics because the purpose of copyright is largely economic, and it should, correspondingly, be judged by its economic consequences. Although these are the obvious fields of inquiry, one comes across important sub-problems that might be more appropriately answered in the literatures of political, psychological, philosophical, historical or information systems scholarship.

1.3.1   The Law of Copyright and Legal Scholarship

A great deal of time may be spent exploring the details of copyright law. It is an activity which inherits the complexities of addressing any body of modern law, but it carries more than the usual quota of curious philosophical issues. Amongst these are the nature and definition of authorship, the boundaries between different categories of informational objects, the distinctions between ideas and the expressions that convey them, the balance between regulating the acts of humans and the nature of machines, the paradoxes of copyright as purposeful yet arbitrary rules. Each of these puzzles evolves through precedent and is periodically perturbed by technology and legislation. The subject is further enriched by its deeply international scope — by the story of ongoing harmonisation flavoured with enduring subtle diversity in national copyright systems.

These intricate and genuinely important topics of positive and comparative law should have as little connection as possible with this thesis. The simple reason is that they distract from the main subjects, which are the systemic consequences of digital exclusive rights and of alternatives.

Different systems of exclusive rights are not equivalent, and the details of the law have economic and normative implications which would be relevant to the present inquiry. It might therefore be argued that skipping over details is necessarily inap-

propriate. Frequently, though, the magnitudes of those normative considerations are comparatively small, and while they may entail some losses or gains in efficiency, a similar collection of imperfections and optimisations would also accompany any alternative compensation system.

There are a few situations in which I do focus on specific aspects of copyright law and jurisprudence. Sometimes it is because they impose significantly and insistently on the question at hand — because a detail of copyright carries very substantial consequences, by preventing (or not preventing) large categories of behaviour that have extensive economic effects, for example.

Although study of the "black letter" of the law plays only a small part in this work, there is a great deal of relevant research by legal theorists who have engaged with the inter-disciplinary issues that surround the law, or with questions of how it might be reformed. This thesis owes much to those scholars and their literature, as the reader should find evident through many citations to and discussions of its claims.

1.3.2 Computer Science

Computer science is the theory of algorithms and their implementation in physical machines. At its heart, it provides results (or bounds) on what computers can and cannot do, problems they can solve, problems they can solve with a certain efficiency, and problems which they cannot solve at all. From this starting point, the field extends to many other engineering problems — problems both fuzzy and circumstance-dependent — about the ability of computers to serve particular human ends. Problems near the "core" of information science have a fairly rigorous mathematical character (How quickly, if at all, can a computer evaluate a particular function? Under what circumstances, and in what sense, is a particular communications protocol secure?), while problems at the applied end of the discipline encounter the uncertainty that characterises the social sciences (How can particular intentions of a human user be conveyed to a piece of software? How much is it likely to cost to break into a real-world network?) .

Computer science is relevant to an understanding of digital copyright in many ways, and in some situations it is indispensable. Its role is to tell us about the power of software: how effectively can peer to peer networks distribute a huge catalogue of files? How much effort and resource is required to jam them? What is required to make a computer enforce exclusive rights? How hard is it for a hacker to change that?

How costly is it to stop them?

Most of the computer science which is relevant to this thesis relates to the broad sub-discipline of computer security. In addition to providing the theory of many crucial building blocks for copyright-related systems (such as secure hashes, signatures, and encrypted channels), computer security encompasses at least four topics which are highly pertinent: the architecture and effectiveness of DRM technologies for preventing unauthorised copying, the security of electronic auditing systems that are an essential part of the "virtual markets" introduced in Part III, the resistance or susceptibility of file sharing networks to software-based attacks by copyright holders, and the architecture and effectiveness of anonymising file sharing networks. One building block which will be of recurring interest is trusted computing — a non-traditional kind of hardware which is relevant, at least in theory, to all four of those subjects.

There are a few sub-disciplines aside from computer security that are of at least passing relevance. These include information retrieval (especially distributed information retrieval, which includes the architecture of file sharing networks), compression, scanning and digitisation technologies such as OCR, and the development of tools for editing various media. By and large though, developments in these research areas are only tangentially connected to the present investigation: we will meet them in footnotes.

1.3.3   Economics

Paul Samuelson and William Nordhaus defined economics as "the study of how societies use scarce resources to produce valuable commodities and distribute them among different people".[17]   That depicts a discipline which is largely descriptive or 'positive', although it can be applied to normative problems through the choice of some ethical criteria, such as social welfare functions or notions of "efficiency". Normative economics, then, sets for itself the task of determining what should be produced, and to whom it should be distributed — or at least, the kinds of institutions that should determine those things.

Copyright is, at its heart, an economic instrument. For some, its proper role is *exclusively* economic: solve the free rider problem, and do no more. Any examination of its properties, either descriptive or prescriptive, therefore requires some theory of

---

[17]This formulation appears on page 4 of the 16th Edition of their text *Economics*, although copyright restrictions on book searches have prevented me from efficiently checking when they first used it.

economic events. In some copyright scholarship, these theories are implicit, but such work is the poorer for its unstated and possibly unexamined premises.

The academic discipline of economics provides a language, a conceptual toolbox, for analysing copyright, and for teasing out some of the elusive aspects of its operation. But the field of economics is invariably a contested one, both from within and from without. Its epistemological status is troublesome, since economic results tend to depend upon strong assumptions about the behaviour of human minds, which are themselves emergent and changeable physical phenomena — perhaps the most complicated phenomena we know of. A strong simplifying assumption about human beings may hold quite true in one context, on one day, in one culture — and then collapse under different circumstances.

What this means, from the point of view of an interdisciplinary researcher writing about copyright, is that economic theory cannot be produced and waved like a wand to answer all of the practical questions on the table.

Much of economic theory is constructed with formalised mathematical models. Some of these models provide quite subtle and powerful insight into the social phenomena they address. The difficulty, however, is that each insight cannot easily be unified with the others; plausible models of marketplaces, particularly marketplaces for sophisticated products (a class of which many copyrightable works are members), involve the non-linear superposition of effects such as non-rivalry, imperfect information, variable transaction costs, endogenous demand and multiple, distinct network externalities. This recurring complexity makes theoretical economics a tantalising tool; it promises concrete and powerful conclusions — such as a proofs of the Pareto optimality of particular institutions — but when one tries seriously to align models with the entire drama which can be observed unfolding in the music industry, for example, matters quickly become much messier.

As a result, economic theory cannot escape being narrative or rhetorical discipline.[18] Each model or result tells a particular story or highlights a particular effect. When studying copyright policy, a number of these effects may be relevant (in the coming chapters we will come across numerous models involving public goods, distortionary taxation, and asymmetrical information, for example), but they may well be superimposed in intractable ways, and even if they are not they must still be considered alongside a great many inelegant impositions by the real world.

---

[18]Boyle makes this case very thoroughly in his comment on the role of economic reasoning in policy conflicts over the expansion of intellectual property laws to facilitate price discrimination (Boyle 2000).

Once the basic themes and variations of copyright economics have been set out and explained, it is perhaps empirical economic research, rather than detailed microeconomic modelling, which is of most assistance. The recent surge of interest in digital copyright has brought with it a wave of econometric studies examining piracy, file sharing, and music sales.[19] While there are some limitations to the data upon which some of these studies are based, and debatable aspects to their conclusions, the chief drawback of this research is that there is not yet enough of it.

The methodology followed by this thesis does not closely resemble that of most economists, but many of the questions it addresses are economic in nature. Wherever possible, I have cited and discussed the relevant economic arguments, models and data, and at a few points I have developed rough but original estimates of the scale of the economic effects in question.

1.3.4   Absent and under-represented disciplines

I have chosen three disciplines — law, computer science and economics — to emphasise, and from which to draw most of my sources. They by no means cover all of the knowledge relevant to the research question; they are simply the three that are most important. The concerns of other fields cannot be escaped; they can only be treated less satisfactorily. A truly complete understanding of digital copyright would comprise not only the perspectives of law, economics and information science, but also several others.

The most important of these may well be politics. It can be fairly maintained that the politics of copyright are a perfect illustration of the weaknesses of policy formation in democratic societies and in processes of global regulation. Vested interests haggle over rents, while the diffuse but very substantial public interest is neglected. It is also arguable that in the expansion of intellectual property laws, one finds the "last frontier" of capitalism, the frontier on which the case for the simple formula of "property and markets" is at its weakest, and yet it has made progress anyway. A view of copyright law which does not deal with these perspectives may be in danger of naivety, and to the extent that political science research on digital copyright can apply informative scrutiny to these sorts of propositions, it is of great value.

---

[19]See, e.g., (Liebowitz 2003b; Oberholzer and Strumpf 2004; Rob and Waldfogel 2004; Boorstin 2004; Peitz and Waelbroeck 2004; Rainie and Madden 2004; Rochelandet and Guel 2005; Michel 2006; Zentner 2006; Rob and Waldfogel 2007; Waldfogel 2010).

Of course, in order to make any judgements about the politics (or economics) of copyright, one must be working from philosophical premises. Without any notion of what is better and what is worse, we can pass no kind of judgement on anything. To say that that human desires for culture should be satisfied according to the measure of willingness and ability to pay is not a neutral claim. We may regard it as the best theory available to describe what good may come of culture, art and information, but such claims are a little tenuous and they should not generate consensus.

If we attempted to study the ethics of cultural commerce in a scientific manner, we would rapidly end up practising psychology and sociology. How does music (or film, or the humble book, or the comic, or pornography) affect human happiness and human fulfilment? Do these things always make us better off, or is it perhaps that a song we've heard but cannot hear again will just haunt us a little, sticking in our heads and making us just a shade grumpier until we get our fix? If the answer is sometimes the latter, are we sure we should be trying to provide incentives for the creation of more music? Does it matter how sophisticated the music is, how we listen to it, how it changes us?

To be frank, this type of difficult question may well be as important as the the fluctuations in Disney's business that follow the latest developments in file sharing, litigation and licensed media downloading. Equally puzzling but totally different conundrums exist in the justification of intellectual property regulation of software.

Chapter 7 does dabble in the ethical issues that underwrite a normative analysis of copyright, but only to a sufficient extent to complete my analysis — and largely in conformance with the conventional presumptions in the literature. This approach may not have tested the more suspicious scaffolds on the premises of copyright with the thoroughness they deserve.

1.4 Methodology

Complicated research questions are best addressed with structured methodologies, so this thesis will follow one. It is briefly described here. The chief objects of inquiry are four policy regimes:

1. a weak DRM/exclusive rights regime (the status quo);

2. a strong DRM/exclusive rights regime ("information feudalism");

3. an information anarchy regime;

4. a public funding regime (the "virtual market" or an "alternative compensation system" (ACS)).

In introducing these regimes I will argue that they are plausible and that they usefully represent important paths that copyright might follow. Their preconditions, in terms of patterns of behaviour as well as legal and technical regulation, will be examined. Three of the four regimes (strong DRM, weak DRM, and information anarchy) will be defined and checked for feasibility in Part II. The last alternative — the "virtual market" public funding system — requires more explanation and justification than the others; I therefore devote all of Part III to it.

With the regimes defined, I will try to determine which of them is preferable and for what reasons. Part IV undertakes this systematic normative comparison. Each Chapter in that Part examines a possible reason to prefer one alternative to another. Taken together, these chapters produce a balance sheet by which digital copyright systems can be compared.

My list of regimes is short, and even a long list could never enumerate all of the forms thats digital copyright could take; in theory and practice we could find situations that are hybrids, or none of the above. All I can claim is that these regimes do a satisfactory job of capturing the types of digital copyright we have seen to date, and many of the alternatives discussed in the literature; and that a great deal can be learned from comparing them.

Before this project actually commences, several things about it deserve some further explanation.

1.4.1 Policy regimes

In order to compare different regimes — different kinds of digital copyright[20] — it is important to know what digital copyright actually *is* (and therefore, what each regime is actually made out of).

We would hopefully all agree that digital copyright is a kind of law. But what is the law, beyond words written in statute books? It is uncontroversial to call a broader human institution by that name: the work of the court that reads the law, that develops its precedents and jurisprudence, that gives it life by resolving disputes with it; the

---

[20]In this instance, and others, I refer to both exclusive copyright regimes and non-exclusive alternatives as kinds of "copyright". Information anarchy would be the 'null' kind of copyright. It should be clear from the context whether the regime in question is specifically an exclusive rights regime.

shared beliefs of the legal profession who in most cases resolve issues before litigation is even considered.

For most scholarly purposes, our policy interest is not in the law for its own sake but rather for its effect on society or commerce or knowledge or, ultimately, on individual human lives. It follows that we must pay as much attention to whether and what and how the law regulates, as to the details of the legal rules themselves.

Modalities of regulation

Lawrence Lessig has proposed a fairly sophisticated picture of the regulatory nature of cyberlaw in general (Lessig 1999). It goes a long way to answering the question "what is cyberlaw?", or in our case, "what is a digital copyright regime made out of"? He identifies four "modalities" to the regulation of people's use of technology: (1) the law, as enforced by courts or just by legal threats; (2) the power of social norms; (3) marketplace constraints; and (4) the architecture of technical devices or "code".[21] He also points out that there can be important interactions between each of these modalities (Lessig 1999, pp 511–13). So, for example, the effect of the law on the nature of technological artefacts (on "architecture", as Lessig would term it) may be greater than its direct effect. If the law motivates the producers of software to do things differently, then law is *indirectly* regulating the users of that code.

This model is an elegant and satisfactory (for some purposes, at least) framework for understanding the means by which copyright rearranges people's behaviour.[22] Digital copyright is not just the law, but also the puppetry of the law in directing regulation through norms, markets, and technological artefacts.

The digital copyright regimes I will examine in the following chapters are not directly equivalent to any of these modes of regulation. They are the *result* of regulation: the patterns of behaviour and experience induced amongst the regulated members of

---

[21]See (Lessig 1999, pp 506–10). Lessig makes this classification in support of an argument that cyberlaw is a distinct kind of law for precisely the reason that it affects the nature of technology (whereas the "law of the horse" does not affect the nature of horses), and that one reason to study cyberlaw as a distinct subject is because of what it tells us about the limits of governmental regulation in general. These arguments are a little different to the use I make of his classification.

[22]Lessig's theory has prompted quite an amount of scholarly discussion. Just on the inclusion of software as a kind of "architecture" for example, Tim Wu has tried to explain when and why this architecture is constructed for political ends (Wu 2003), while others have argued that it should have the status of a separate, fifth mode of regulation (Cohn and Grimmelmann 2003; Grimmelmann 2005). The various qualifications, amendments and elaborations on the theory are not especially pressing for the explanatory use I am making of it.

society. As we shall see, the same regime can sometimes result from different combinations of regulatory instruments. Of course, the nature of the regulated society also matters: the same decisions by lawmakers (in two countries, say) might result in a different regime because the populations subject to regulation were different to begin with.

1.4.2   Which kinds of copyright will be studied?

The use of a 'regime' model to represent the choice of copyright policies does hide certain matters from constant view. This must be borne in mind.

Perhaps the most important of these is the fact that different regimes can be prevailing simultaneously — in different countries, within different copyright-dependent industries, or even in different segments of the same market.

The regime-oriented methodology is a way of staying focused on the commonalities between the music industry (heavily disrupted by MP3s, file sharing and other digital innovations) and the book industry (only beginning to be disrupted) and other cultural copyright industries — despite the fact that their prevailing circumstances are quite distinct. The differences are de-emphasised, becoming parameters rather than the story that is told explicitly.

It is therefore important at the outset to be clear about which copyright industries, and which parts of those industries, will remain in the main flow of the analysis. It turns out that the simplest case for alternative copyright regimes can be made when the subject matter in question is for entertainment, artistic or educational purposes. Software, and copyright works intended to be used in the course of business, have economic properties that make them less amenable to production through simple alternative compensation systems. These economic particulars will be discussed in Sections 5.4 and 7.5.4.

The design and evaluation of alternative compensation systems that might apply to those classes of works will have to wait for future research.

1.4.3   Comparing Regimes

Part IV sets about determining which regime is preferable, and on what grounds. Each Chapter there works through one dimension of the problem. Which system provides the best incentives for artists and publishers? Which system gives people the most

access to valued cultural works? Which requires the most technical and enforcement infrastructure? And so forth.

Dollars (often in the form of percentages of the dollar values of particular markets) will be used as the common denominator for the results of these Chapters, even though they are not the ideal metric. It would be better to measure the smiles on people's faces, or the number of new and interesting thoughts that spring to life, than to squeeze the entire analysis into society's wallets. Unfortunately, the planet has many established mechanisms for counting dollars, and far less infrastructure for counting smiles or new and interesting thoughts (a problem to be addressed, but not here and now).

The tools employed to compare the regimes vary from chapter to chapter. In some, good economic data is available; it is possible to sift through it and dig out a nice conclusion about the dollar-value advantage of one copyright model or another. In other cases, the raw material is stories (or arguments from other disciplines) rather than hard numbers. In those instances, dollar figures can at best be obtained by educated guesswork. While not ideal, this is not always unsatisfactory either, because what really matters are the orders of magnitude of the figures: some of the economic differences between the regimes turn out to be vastly more consequential than others, and it is really only effects in the top couple of tiers that drive the conclusions.

The results of the normative analysis are compiled at the end of the thesis, in Chapter 11. They include estimates of the advantages and disadvantages of each copyright regime compared to the others. There is a very strong case for the public funding regime! The conclusions also offer a map of the factors that affected the results the most. Some of these are fairly well-understood and well accounted for. A few remain poorly understood and some of these are singled out as deserving further, specialised, research.

1.5   Summary of Original Contributions

For the convenience of reviewers, this Section will attempt to briefly summarise which portions of this thesis represent original contributions.

Chapters 2 and 3 largely review existing market phenomena and literature, though the particular choice for partitioning the regimes introduced there is, I believe, original. The argument that the web essentially requires information anarchism may be new. Chapter 4 is also a review chapter, though hopefully the reader will find its com-

prehensiveness and perspective informative.

Chapter 5 contains many significant original contributions. Although many authors became interested in public funding proposals in the wake of Napster, the proposal published in (Eckersley 2003; Eckersley 2004b), which forms the core of the chapter, contains a number of novel aspects. These include deeper considerations of how to perform the accounting for such a system and how to achieve that securely; the arguments for the use of voting in addition to other measures of popularity; a more thorough analysis of the available methods for setting and adjusting the total levels of funding for the schemes; identification of the circumstances that make particular copyright marketplaces (and not others) amenable to these types of proposals, and the particular point that the Web deserves inclusion.

Chapter 6 is largely novel. It builds on pre-existing data from several sources, but the synthesis is original and this is to my knowledge the first attempt to estimate the overall systematic cost of artificial scarcity.

Chapter 7 also contains many original contributions. The opening sections of the Chapter explain the connections between information economics, the structure of copyright marketplaces, and the option of public funding much more precisely and thoroughly than previous scholarship. Section 7.3 contains an original examination of the relevance of various types of non-copyright incentives to copyright policy, and original formal models for comparing anarchy to copyright systems. Section 7.4 is a review, though the interpretation of de Trenqualye's results is original. Section 7.5 is also largely original, though it builds on a number of other cited works.

The observations that make up Chapter 8 are mostly not original, though I believe the complete picture obtained by combining them may be. The point about the avoidable inhibition of literary and scientific "gestalts" by copyright transaction costs may have been original when I first discussed it (Eckersley 2004b, §III.C.3), but it has subsequently received much more attention with the slow and limited development of Amazon, Google and the Internet Archive's book scanning projects.

The contributions in Chapter 9 are largely original. The numerical estimates on the cost of secure DRM (Eckersley 2004b, §III.C.2) are, to my knowledge, the only published attempts to seriously estimate how much secure DRM infrastructures are likely to cost. There may have been some estimates undertaken at Microsoft, Sony and other companies that have developed complicated DRM platforms, but if those exist they are not public.

Chapter 10 is the most serious attempt in the literature to study the interactions between the economics of taxation and proposals for alternatives to copyright.

Hopefully the conclusions in Chapter 11 are original too.

1.5.1 Timeliness

The bulk of the research for this thesis was undertaken between 2001 and 2005 (and much was published in (Eckersley 2004b)), although some aspects and portions were completed later, between 2006 and 2011.

Given the fast-changing nature of the underlying subject, there is a potential for earlier aspects of this thesis to have been contradicted or obsoleted by subsequent events. Did this happen?

There were certainly many events during the decade that are candidate causes for this kind of obsolescence: the emergence of mobile phones as general-purpose computers and media consumption devices; the rise of iTunes Music Store in 2004–2006, and that service's phasing out of DRM during 2007–2009; the movement of much file sharing traffic from decentralized P2P networks to centralized websites; the emergence of free advertising-funded and hybrid-subscription media services like YouTube, Pandora, Spotify, GrooveShark, Vimeo, and company.

The "regime" based model appears to have remained valid across these transitions. Surprisingly the data discussed in Section 6.3.1 shows that in most cases, transitions have not even changed the state of the regimes very much. For instance, music stayed a weak DRM/copyright regime, before and after iTunes ceased using DRM, and before and after Spotify reintroduced widespread DRM for music.[23] The film industry has remained a strong copyright/DRM regime, with most viewings of films remaining licensed and new widely-deployed DRM platforms including BluRay and Netflix streaming; an amount of piracy on YouTube and the huge portion of Internet bandwidth that came to be used by BitTorrent did not change these numbers. There may be an imminent change of regimes in the book industry, but it has not happened yet.

---

[23]Although conceivably, Spotify and similar services could turn lead to a virtual market-like regime with suitable regulatory influence.

Part II

Information Feudalism and Information Anarchy

Chapter 2

Digital Rights Management: Locks and Chains in Cyberspace

Contents

This chapter introduces the regimes that I will use as models of exclusionary digital copyright. It deals with the question, "what are the important features of economic systems based on legal, DRM-based, or other methods of enforcing copyright in technological systems?" We need an answer, before we can decide if those systems are a good idea.

I should note at the outset that I am going to use the term 'DRM' somewhat in a very broad way that includes sales mechanisms like payment-based access controls, and enforcement technologies like the monitoring of file sharing networks for litigation

purposes, in addition to the technologies that try to prevent people from copying media files that they already possess. One slightly unintuitive consequence of this definition is that the music industry could still be deemed to be using DRM even though it no longer insists on encrypting all of the compressed audio tracks it sells in restrictive formats like WMA or M4P.

I believe this is the best approach: the question is all about finding a reasonable simplification of far too many variables. DRM by any definition is an umbrella term for a remarkably diverse menagerie of hardware and software; the motivations may be similar but the means are not.

Marketplaces that use these tools can manifest in as many ways as a motorised vehicle can be designed. Working with the analogy: forget each rivet, wire, tube and cog; where are the wheels, the propulsion system, and the seats? Who pays the registration, and who gets to drive?

Rather than trying to capture all of the possible DRM regimes within a single model, or getting lost in a huge number of details, I introduce two that can be clearly specified and delineated. These are 'strong digital copyright', 'strong DRM' or 'information feudalism' on the one hand, and 'weak digital copyright', or 'pragmatic DRM' modelled on the prevalent status quo. They could be, respectively, the gleaming red sports car and the station wagon (beaten up and a little rusty, but it might get you there).

The chapter begins with a brief discussion of five families of technology that have been designed with the goal of preventing or inhibiting people from accessing and copying files (Section 2.1). Any DRM regime is going to rely on some combination of these. Also integral to the system are the kinds of legal rules and other regulatory practices that have been developed to complement and reinforce the operation of DRM itself (Section 2.2). Section 2.3 suggests a way to understand consequences of mixing these technical and legal components in different ways; it introduces two categories for DRM regimes ("pragmatic" and "feudalist") that are representative of the important economic options, and which will be evaluated in Part IV of the thesis. A brief conclusion can be found in Section 2.4.

2.1 DRM Technologies

2.1.1 Access and Copy Controls

The simplest way to pursue the copyrightly goal of ensuring that only users who pay for a work get to enjoy a digital copy of it, is to write software that tries to limit reproduction. Generally speaking, such software falls into two categories — access controls, and copy controls.

Access controls are a category of architecture that is designed to prevent a user from getting a first copy of a work unless they have a license to do so. Copy controls are snippets of software that try to stop audience members from making a reproduction of work once they have obtained a copy.

Access controls, for what they are worth, are comparatively easy to implement. A website that requires customers to pay a fee before being offered a download is a perfect example. True access controls can also be implemented in physical media by using broadcast encryption schemes. No payment, no key, no access.

One philosophy on the matter is that there is little point in trying to make access controls robust if piratical users will just share the works they have legitimately purchased access to.[24] In order for perfect access controls to deliver much more power to copyright holders, they need to be accompanied by strong copy controls.

Unlike access controls, true copy controls are impossible to implement on general purpose computers. Ordinary computers are capable of replicating any information they have access to.

Instead, those who want to design copy controls without assistance from hardware are reduced to designing copy inhibitors instead. Copy-inhibition techniques resemble access controls but with some conceptual differences. The DVD CSS scheme, for instance, was predicated on the idea that user controlled software would *never* get access to the work. It would only be displayed by certain "authorised" restricted player software — so the user could see a film on their screen or TV but would never obtain a digital copy. No first copy, no later copy. Of course, the "authorised" restricted player software can always be edited into unrestricted player software.

---

[24]This is most especially true if there is some widely used distribution system — like a P2P network or a general purpose file storage system — that means that one single sharing user can give a copy to any other user who wants a pirated copy (Biddle, England, Peinado, and Willman 2002).

2.1.2 DRM in Hardware

So the serious flaw in any arrangement of software intended to allow some access to a copyright work but to prevent unauthorised reproduction, is that software is malleable by its very nature. If a program — a sequence of instructions to be stored and performed on a Turing complete computer — is ever capable of reading the digitised version of a work, then it can easily be modified to save a copy of the work at the same time. The diverted copy can be used for any purpose. Another way of putting it is that attempting to write software that plays a media file but refuses to allow copying, is necessarily just wrapping a limiting user interface over a more powerful system. If the user can change the interface, they can copy the file.

The only reliable way to prevent such copying is to somehow change the hardware so that it is no longer a true general-purpose computer, and cannot be reprogrammed arbitrarily.

The idea of harnessing hardware in this fashion is not new, dating at least to the 'copy protection' of software in 1980s. Various tricks were attempted in that era, such as producing floppy disks with holes in them,[25] or 'dongles' that had to be plugged into a PC before a program would run. These forms of DRM tended not to be effective, because hackers could identify the code that tested for the presence of the unusual hardware and disable it.

More effective versions of hardware support were built into video games consoles during the 1990s. Sony's PlayStation, released in 1994, combined something analogous to the old trick with holes in disks — a non-standard code on the end of each authorised CD, that could not be written onto a CD-R — but they backed it up with an immutable booting arrangement which would refuse to read CDs that lacked the code. The PlayStation copy protection system proved sufficiently secure that physical interventions were necessary to defeat it. Theoretically, dongle-type copy protection systems could also be made comparably secure.[26]

A trade known as 'mod-chipping' sprang up to provide hardware modifications to disable the 'refuse to play' features found in PlayStations, and the more sophisticated encryption-based variants that appeared in DVD players, and X-Boxes. Usually, a

---

[25]The idea being that attempts to read from the affected parts of the disk produce errors; the program would only run if these errors occurred.

[26]The way to achieve this would be to make the dongle a small computer that performs a non-trivial component of the software's actual role. Breaking the scheme would then require reverse engineering to the point that the code on the dongle could be extracted or re-implemented.

cheap IC is spliced onto a device's circuitry at a few key locations.[27]

In order to make mod-chipping impractical, devices must in one way or another possess the property of *tamper resistance*. There are basically two ways to attempt this. One is to make it hard for an assailant to figure out what they need to do to change the hardware's behaviour; the other is to make it physically difficult for an assailant to make those changes. The first approach has, to date, proved ineffective.[28] The second and approach is more promising though also more expensive; it is discussed further in Section 9.2.1.

2.1.3   "Trusted" Computing

A customised and tightly controlled platform like the X-box (and especially the X-box 360, with its non-standard physical media) is in many ways the ideal environment for DRM. Making DRM stick in any other context will be harder. Phones are definitely a more hostile environment for DRM, with code drawn from more sources, generally cheaper and more varied hardware and less centralised control. The unavoidable weakest link is the PC.

The aspiration to implement secure DRM on PCs has led to a major technical R&D effort by a consortium of firms known as the Trusted Computing Group. Notable members include Microsoft, Intel, IBM, HP, AMD, Sony, Nokia, Sun, Philips, Samsung, Motorola, Ericsson, National Semiconductor and TI. Although this group has attempted to play down the DRM-focused nature of its work,[29] the fact remains that the concept of trusted computing was originally inspired by the copy protection problem (Stefik 1997).

"Trusted computing" (or "treacherous computing", as Richard Stallman has argued it should be called) includes a combination of ingredients which make it especially

---

[27]See `http://www.infinitymod.com/cgi-bin/matrix/site.pl?page=chip_mxl2` for the specifications and installation instructions for a typical PlayStation 2 modchip.

[28]Attempts to encrypt all of the processor I/O on the X-box failed to produce results; mod chips were built for the console anyway, and then exploits were found that allowed "software only" compromise. Andrew Huang's book records these weaknesses in great detail (Huang 2003).

[29]When visited in 2005, the Trusted Computing Group's FAQ (`https://www.trustedcomputinggroup.org/faq/`) contained the following disingenuous entry: [question] "Was TCG formed to specify Digital Rights Management (DRM) technologies?" [answer] "TCG specifications do not provide all the necessary technical elements required for DRM. It is conceivable that developers could build their own DRM solutions that would operate on systems with Trusted Platform Modules, but TCG specifications alone are not DRM solutions." Those readers entertained by coincidences should know that as I was writing this paragraph, the following story appeared on Slashdot: `http://yro.slashdot.org/article.pl?sid=06/02/19/070202`.

potent for DRM applications. As Schoen (2003) explains, there are four essential components to DRM based on trusted computing:

- *Remote attestation* — the ability of the hardware to certify to a remote party that a machine is in a particular state, before that party initiates a cryptographic session during which important information (such as decryption keys for DRMed material) may be sent.

- *Secure key storage* — the ability of the hardware to store keys in a fashion such that only certain applications (those that have been authorised by a remote authority and are executing in an approved state) can make use of them to decrypt DRMed material. The secure storage of just a few keys in special hardware facilitates an arbitrary amount of encrypted storage on conventional disks that can only be accessed by processes authorised by the DRM implementor.

- *Memory cloaking* — the ability to have certain processes executing on a computer whose memory contents cannot be seen by any other software on the system, even if that software has "root" or "administrator" privileges.

- *Secure I/O* — the ability of processes to receive input from users, and send output back, in such a way that the process can be sure that no other software is eavesdropping or modifying the communication. Secure I/O can be used to prevent software like keystroke loggers from sniffing passwords; it can be used to prevent useful software tools like screen grabbers or sound recorders from working if there is DRMed material around.

At a hardware level, these can be provided if the trusted platform module has sufficient access to and control of various parts of the system (memory, CPU registers, etc) to allow the calculation of "integrity measures" ("IM"s),[30] remote attestation of those measures (so that remote parties can be sure of the precise software running on a computer before exchanging keys for the decryption of content), and key storage (so that those keys are inaccessible when unapproved software is running).

It should be noted that trusted computing facilitates an open-ended range of different applications. Amongst these are not only DRM and document control (Anderson

---

[30]Integrity measures can be any algorithm to confirm that the system is behaving as it "ought to"; *e.g.,*, that the operating system is the latest version available, that it has not been compromised, and is enforcing a memory cloaking convention.

2003), but other schemes of interest: reliable detection of rootkits,[31] secure Internet voting,[32], networked games that can exclude cheating clients, and — in an ironic twist — anonymising peer to peer networks which shield their users' habits from active surveillance (Schechter *et al.* 2003, Section 4). In this list, rootkit detection and secure voting do not require remote attestation,[33] while DRM, cheat-proof games, and anonymising P2P do.

The hardware for proper trusted computing environments has not yet been deployed. An important component, called a Trusted Platform Module, is present in a growing number of systems. A Trusted Platform Module on its own is insufficient to implement secure remote attestation and secure I/O. The Module must have the direct means to interrogate the state of almost any system component... the RAM, the binary executable files on disk, etc. All of these components must be capable of cryptographic handshaking with the Trusted Platform Module and must exhibit some degree of tamper resistance. Otherwise, an attacker can just replace that hardware component with a version that has been modified to notice and leak important cryptographic keys, or to break some sensitive invariant *after* the IMs have checked it.

2.1.4   Watermarks

There are several reasons why DRM architects might want to make each authorised copy of a work uniquely identifiable. One sort of motivation is to make devices refuse to play media other than those likely to be owned by a legal purchaser/licensee (along the lines of the PlayStation's access code). Another, grandiosely termed "traitor tracing", is to identify the sources of files that have been liberated from copy controls and have subsequently been circulated.

A *watermark* is identifying information embedded within the data (as opposed to the metadata) of a file. Watermarks can range from faint background images behind paper documents, or unchanging logos superimposed on video streams by television stations, to algorithmically sophisticated systems for hiding attack-resistant and user-specific data within media files of all sorts.

---

[31]See for example `http://www.it-observer.com/articles.php?id=977`.

[32]As anticipated for example in (Oppliger 2002, Section 6.4). Secure online voting is particularly dependent on local attestation and secure I/O facilities.

[33]In practice, rootkit detection only works locally if there is some secure interface between the user and the trusted computing hardware. Existing hardware for trusted computing does not provide these sorts of facilities.

Early attempts to deploy watermarks as effective security measures were unsuccessful. One group of techniques was famously proposed for "refuse to play" applications under the Secure Digital Music Initiative (SDMI). They were all promptly cracked (Craver *et al.* 2001). The folk wisdom is that, regardless of the methods used to embed imperceptible information in media files, adversaries will find ways to transform the watermark away or to overwrite it.

Even if this belief is true, there are some proposals in the computer security literature for attack resistant traitor-tracing watermarks, which are based on variations in the underlying work itself (rather than just imperceptible noise which can be erased by an appropriate filter). Such a watermark has been included in the "Advanced Access Content System" (AACS) DRM specification for HD-DVD and Blu-Ray video discs.[34] Rather than including a single copy of a film, each disc would contain many points at which the film could be subtly different. For example, there might be a twinkle in a character's eye at one moment — or not — or a twinkle in a different place. Each device capable of decrypting the disc would get one of these several possible versions of that moment in the film. A user who makes and circulates many copies of films from their player may be partly or entirely identified by statistical analysis of variations visible in the pirated version.[35]

Both the plausibility of watermarking schemes and their practical utility vary greatly with the kinds of works they are applied to and the business model in question.

---

[34]See Ed Felten, AACS: Sequence Keys and Tracing, `http://www.freedom-to-tinker.com/?p=1110`; (Lotspiech 2007).

[35]Under the fairly reasonable assumptions that these schemes are based upon (which the literature terms "the marking hypothesis"), the only avenue of attack against the watermark is to use several differently-marked master copies. If keys (or equivalently, alternate moments from the film) from many users are combined, assailants can reduce the amount of information available to the AACS authority about the identity of the attackers. Conversely, including more marks in the films increases the amount of information. Boneh and Shaw (1998) introduced several schemes for arranging this information to identify groups of attackers colluding against a watermark. They aim to identify one of c attackers among a total population of n users with a probability better than $1-\epsilon$. They construct a marking scheme involving $2n^2(n-1)\log(2n/\epsilon)$ variations in the work, which is in $O(n^3\log(n/\epsilon))$. This scheme clearly requires too many marks for any use where n is a number of consumers. They offer a second scheme which is only logarithmic in n, but it requires more than $8c^3(c-1/2)\log(2N/\epsilon)\log(1/\epsilon)$ variants. A vibrant literature ensued, offering improved variants and alternatives to this scheme (Schaathun 2003). But it remains that to trace a group of hackers who have cracked 10 devices, copyright holders would need to see hundreds of thousands of the variant moments in works released by the hackers. Clearly, no Hollywood studio is going to mix this many variations into a single film (in fact, AACS only supports 1024). But if hackers liberate huge numbers of DRMed works, it may in fact be feasible to trace their keys this way.

Traitor Tracing

The most promising applications of watermarking are in traitor tracing. If the source of a file leaked into public circulation can be identified, various forms of retaliation are possible. That user might find themselves on the pointy end of a civil lawsuit, or they might just find that all of their devices' keys have been disabled, so that they cannot play any media released in the future.[36] Their existing libraries might be disabled (although this may only be a nuisance if they have already decrypted them). Individuals who liberated enough material might be targeted for criminal prosecution.

There are a few situations in which traitor tracing is almost certain to be both effective and cost-effective as a source for litigation targets. A good example can be found in the film industry, which relies on an inter-temporal price discrimination system where the highest valuers of their products are steered to purchase cinema tickets, followed by DVDs for medium valuers, and finally advertising-funded television screenings are used to "sell" to the lowest valuers.[37] Piracy is of course a threat to all of these stages, but the first round (cinema sales) is relatively defensible: if nobody has a digital copy of the film on their computers until it has been released on DVD, then studios are free to collect cinema revenue without being interrupted by file sharers. The difficulty has been that digital copies of at least some films *have* been leaking. This prompted aggressive responses: an attempt to ban 'screeners' and review copies of films;[38] and the arrangement of legislation making it a criminal offence to use video cameras in cinemas.[39] Leaks of cinema projection copies are also a concern. Watermarks are a feasible solution to this problem, because the population of authorised recipients of pre-DVD copies of films is small and precisely identifiable. The film industry appears to be turning to this solution.[40] The smaller number of adversaries decreases the chance that the

---

[36]This is in fact the scheme adopted in the AACS DRM used in HD-DVD and BluRay discs. AACS is a very sophisticated scheme that includes three revocation mechanisms to prevent attacks from being permanent(AACS Licensing Authority 2006). One is a cryptographic scheme known as a subset-difference tree (Naor, Naor, and Lotspiech 2001) that allows discs to be pressed so that they can be decrypted by any device's keys, except the keys that are known to belong to particular hardware or software players that are known to have been compromised. Discs also contain lists of hardware and software versions that are known to be insecure; the insertion of a newly released disc will permanently stop a drive from communicating with PC software that is known to have been cracked (even with other discs). The disc will also tell the PC software which drives it should stop talking to.

[37]See (Fisher 2004, pp. 67–79) for a more detailed discussion of this particular tiered structure.

[38]http://www.corante.com/importance/archives/2003/10/29/props_for_jack_valenti.php

[39]See 18 U.S.C 113 2319B (introduced by the Family Entertainment and Copyright Act of 2005).

[40]http://www.npr.org/templates/story/story.php?storyId=5151434

watermark will be stripped or changed, even if those tasks are relatively easy. We can therefore expect the film industry's deployment of watermarks to prevent these leaks to be highly effective.

Beyond these sorts of narrow application niches, traitor tracing is a theoretically tantalising technology. It promises to address the "break once run everywhere" problem by creating strong incentives not to be the person who breaks once. In practice, we do not yet know whether it can be made technically secure enough without producing inordinate numbers of variant moments in the underlying works; nor is it clear that the tracing schemes can be combined with efficient legal measures to realise these benefits.

2.1.5   Other Tracing Methods

All of the technical measures discussed so far have been located on the user's computer or in the user's files. None of these DRM systems has demonstrated much practical efficacy for entertainment goods.

There are other ways to use technology for copyright enforcement that do not suffer from this architectural weakness. For our present purposes, it makes most sense to categorise the surveillance that the entertainment industries have undertaken on P2P file sharing networks as a kind of DRM.[41] After all, the software they have built for this purpose serves the same objectives as "traitor tracing" watermarks that might be embedded in media files themselves.

Most P2P networks are vulnerable to this kind of tracing. Any network application built on top of TCP/IP has the property that each client can see the IP addresses of any machines with which it communicates directly. The most practical approach for building file sharing applications involves direct communication between the nodes that are sharing copies of files and the nodes that download them.

In these most practical P2P networks, copyright enforcers can always connect as a network client, search for apparently infringing files, download them, and see which IP addresses the data travels from. In the United States, "John Doe" lawsuits can in most cases be initiated based on the IP address, and subpoenas used to reveal the targets' identities. It doesn't matter whether the P2P protocols might be encrypted because,

---

[41]Entertainment companies have not for the most part carried out this surveillance themselves. Instead, smaller companies such as MediaDefender, BayTSP, and SafeNet have sprung up to provide these services to record labels and film studios.

when connecting as a client, the enforcer can see inside the encrypted channel.

Two means are known by which P2P network architectures can defend their members from this attack. One is to make the connection between the uploader and the downloader indirect and (at least partially) anonymous. Freenet has this property, as do ordinary P2P protocols used over Tor.[42] To date, anonymising networks have incurred very substantial performance penalties when compared to simpler non-anonymising P2P protocols. It is very likely that this tradeoff will remain burdensome.

The other defence that P2P networks can use against surveillance is to combine encryption with membership restrictions, so that only trusted individuals can participate. This strategy began with warez groups and ratio FTP sites, but may become more widespread as the war on copying continues. It works very well for dedicated file sharers, who have huge libraries of material and a commitment to exchanging it. The strategy is inherently less effective for casual teenage pirates, since many more of them need to band together to make searches fruitful, and because if most teenagers can figure out how to join a network, copyright enforcers will have little trouble doing so too.

Because of the great inconvenience to P2P users of adopting either of these defence mechanisms, surveillance technologies are among the most effective DRM technologies available. They have facilitated tens of thousands of lawsuits against ordinary P2P users, and are probably indirectly responsible for preventing hundreds of thousands or millions of people from file sharing.

2.2   Complementary Regulation

The various forms of digital rights/restrictions management systems grant rights holders differing degrees of control over what users do with copyrighted works. But there are limits to how far this control can go, if groups of users can band together to organise their escape. Files can be liberated — by analogue to digital conversion, if all else fails — watermarks erased, and the resulting bounty exchanged over file sharing networks. To achieve maximal economic impact, DRM must therefore be coupled with other measures that are intended to suppress these activities.

For many purposes, it would be clearest to treat the development of DRM technolo-

---

[42]See (Clarke, Sandberg, Wiley, and Hong 2001; Dingledine, Mathewson, and Syverson 2004). It should also be noted that the anonymity these networks provide is limited when they are subjected to traffic analysis.

gies as a topic separate from the development of digital copyright law. There are some connections, especially in the case of anti-circumvention rules, but most copyright scholars would for example regard TC DRM as a very different kind of regulatory instrument to the secondary liability regime applicable to P2P software developers. But from the viewpoint of systemic economic organisation, these interventions have a great deal in common.

Because the objective of this chapter is to identify the characteristics of markets based on DRM, it is necessary to include laws and other social modes of regulation alongside the technical protection measures themselves. This section briefly outlines a few of the most important examples.

2.2.1 Anti-circumvention Laws

The ease with which copy protection systems were broken motivated a complicated global legislative response: the development of anti-circumvention laws. With variations in detail, these laws prohibit the creation, use, and distribution of devices (or services) that extract works from the control of DRM systems.

The process of including those rules in the world's copyright statues began in earnest with the so-called Lehman "white paper" (Information Infrastructure Task Force 1995) and is ongoing. The most significant steps were agreement on the WIPO Copyright Treaty in 1996, the passage of the U.S. Digital Millennium Copyright Act in 1998, and the adoption of the EU Copyright Directive in 2001. Australia implemented anti-circumvention laws (and the other requirements of WCT) with the Digital Agenda Act in 2000, but was forced to adjust them with the US-Australia Free Trade Agreement in 2004. Similar harmonisation to the DMCA has been required of other nations entering bilateral trade agreements with the United States. As of this writing, a handful of developed countries are still navigating the turbulent politics of WCT implementation (the most prominent of these being Canada and Spain).

Anti-circumvention laws have proved extremely unpopular with the technical communities that are regulated by them. In logical terms, there is no clear distinction between the discussion of cryptography and the creation of tools to break it.[43] Technical minds have tended to interpret this fact as making it impossible or absurd to prohibit one without prohibiting the other. But while arguments from freedom of speech had

---

[43]David Touretzky and many assistants went to extravagant lengths to demonstrate this point during the conflict over tools to decrypt DVDs (Touretzky 2000).

earlier persuaded the U.S. courts that restrictions on the export of cryptographic software violated the first amendment,[44] they were ineffective in defending hackers who distributed tools for circumvention DVD copy protection.[45].

Anti-circumvention laws have also been criticised for their possible effect in relation to hopelessly insecure technical protection measures,[46] for threatening to eliminate user rights such as 'fair use' in the United States,[47] and for excluding the users of free/open source software from legal access to cultural works.

However fairly based these criticisms are, they do not directly interfere with the law's operation. With only a few hiccups, anti-circumvention rules appear to be achieving one of their proponents' highest priorities: the elimination of commercially purchasable means to opt out of whatever restrictions a copyright industry might decide to impose.[48] They have been much less effective at preventing the development and distribution of circumvention devices in general. Especially where DRM can be neutralised with software alone, that software has spread rapidly regardless of the law.[49] Preventing diffusion of that sort is far harder than preventing P2P file sharing, because the volumes of data involved are far smaller.

2.2.2   The Prohibition of File Sharing

A near-mandatory step in the construction of an economic system based on DRM is the suppression of file sharing networks. If people are free to build and use these tools without restraint, the task of selling data that is encumbered by DRM restrictions is a

---

[44]See *Bernstein v. U.S. Dept. of Justice*, case documents archived at `http://www.eff.org/Privacy/Crypto_export/Bernstein_case/`.

[45]See *Universal v. Reimeredes*, case documents archived at `http://www.eff.org/IP/Video/MPAA_DVD_cases/`.

[46]DRM manufacturers have encouraged this situation by litigating or threatening to litigate over such straightforward circumventions as removing ROT13 ciphers (Perens 2001) and holding down the "shift" key to prevent automatic execution of software on a CD (Dean 2003).

[47]17 U.S.C 1201, and especially the 'trafficking' causes of action in 17 U.S.C. 1201 a 2 A, effectively eliminate these rights, because users cannot obtain the tools necessary to enjoy them.

[48]This victory has been decisive in the United States, although there have been a few deviations elsewhere. Australia is one example — see *Stevens v. Kabushiki Kaisha Sony Computer Entertainment* [2005] HCA 58; the ruling is partly an artifact of the particular architecture of the access controls on the PlayStation (see paragraphs 130–144), but it has been sufficient to ensure that Australian consumers can purchase PlayStations and DVD players that are free from DRM and region-coding restrictions. A Finnish ruled that the CSS encryption used on DVDs did not constitute an "effective" technical protection measure; see Helsinki District Court case R 07/1004, 25 May 2007, `http://www.turre.com/css_helsinki_district_court.pdf`.

[49]This can be seen by the ease with which searches of many types yield up copies of libdvdcss, PlayFair, QTFairUse, FairUse4WM, and other widely illegal circumvention software.

daunting one. It only takes one decrypted copy to leak onto file sharing website or a P2P network, and then there are millions of copies.

P2P suppression has been a global effort comprising litigation, lobbying for stronger laws, the use of various kinds of denial-of-service attacks and of course the civil prosecution of individual network users. Although these efforts have not ended file sharing, they have certainly dented its progress.

The first stick which the music and film industries have wielded against file sharing has been litigation against network and site operators. These cases have rested on various legal theories of secondary liability for copyright infringement, according to which those who run a network are responsible for the conduct of its users.[50] In the United States, this assault on file sharing began well with the decisions in the *Napster* and *Aimster* cases, and the closure of other networks such as AudioGalaxy; it then faltered with the trial and appellate rulings in *Grokster*, before the Supreme Court intervened to create a doctrine of inducement to copyright infringement which created liability based in part on the intent, rather than just the architectural decisions of P2P developers.

Litigation in other jurisdictions has produced mixed results. KaZaa has been one of the most prominent defendants: after a hefty damages award at trial in the Netherlands, the original developers sold the network to an Australian company. An appeal in the Netherlands overturned the trial decision on technicalities, but a new case began in Australia.[51] The fly-by-night behaviour of KaZaa's Australian proprietors can only have contributed to findings of liability for authorisation of copyright infringement.[52] The conflict has been waged in the courts of many other states,[53] with most (but by no means all) of the rulings in the industry's favour.

The failure of lawsuits against P2P operators to eliminate the networks altogether soon led to massive campaigns of litigation against individual P2P users themselves. These campaigns have both technical and legal ingredients; technically they are only possible because of the feasibility of tracing users on P2P networks (as discussed in Section 2.1.5 above). Legally speaking, they require a regime that makes thousands of

---

[50]For a comprehensive survey, see (Strowel 2009).

[51]See *BUMA & STEMRA v. KaZaa*, Amsterdam Court of Appeal, unofficial translation at `https://w2.eff.org/IP/P2P/BUMA_v_Kazaa/20020328_kazaa_appeal_judgment.html`

[52]See *Universal Music Australia Pty Ltd v Sharman License Holdings Ltd* (with Corrigendum dated 22 September 2005) [2005] FCA 1242, `http://www.austlii.edu.au/au/cases/cth/federal_ct/2005/1242.html`; `http://www.usatoday.com/tech/news/2006-07-27-skype-kazaa-settlement_x.htm` (reporting details of a subsequent settlement)

[53]A pseudosurvey of sorts can be found in (IFPI 2006, p. 19).

lawsuits against individual file sharers a practical and profitable project. The statutory damages available in the United States[54] clearly meets these requirements, as does the German legal system,[55] under which many thousands of individuals have been sued. As I pointed out above, these suits are probably the most effective kind of DRM that copyright holders have found to date.

2.3   Categorising DRM Regimes

A phenomenon as heterogeneous as copyright enforcement based on technology cannot be treated as a yes-or-no policy proposition. Consideration must be given to the variety of forms that it may take.

One way to address that issue would be to quantify the extent of usage of each of the technologies discussed in Section 2.1. There could be a parameter for the strength of watermarking, a parameter for the strength of tamper resistance, a parameter for the amount of effort spent on debugging code (thereby determining the difficulty of finding exploitable weaknesses[56]), a parameter for the number of coders working to update P2P surveillance systems; non-technical factors, such as the vigour of copyright enforcement against individuals and file sharing firms, would also need to be included. This approach would be thoroughly comprehensive, and absurdly difficult. How would we *accurately* predict the outcome of each recipe that combines multiple technologies and legal rules, without being able to measure them?

Some authors have advanced models which abstract from the complex details of DRM in helpful ways. So, for example, Schechter *et al.*(2003) propose a two-parameter model that considers two costs: the "extraction" cost e, for obtaining a first unencrypted digital copy of a work, and a marginal "distribution cost" d that applies to its reproduction. Different interventions can then be understood as changing one or both of these numbers: DRM based on trusted computing or laws criminalising the use camcorders in cinemas[57] will increase e, while flooding file sharing networks with fake

---

[54]See 17 U.S.C. 504 (c).

[55]See `http://news.bbc.co.uk/2/hi/entertainment/5009224.stm` In Germany, constitutional privacy protections prevent private plaintiffs from obtaining the identities of ISP customers, but government agencies are able to initiate prosecutions and reveal identities so that private plaintiffs can continue the cases (see `http://www.p2pnet.net/story/10400`, translating German reporting).

[56]For an example of model treating exploits in this fashion, see (Anderson 2002).

[57]See 18 U.S.C 2319B (the criminal "camcorder" provisions introduced by the Family Entertainment and Copyright Act of 2005).

files, or filing suits against their developers, will tend to increase d. For a far more specific economic model, Duchêne and Waelbroeck (2003) also picked two parameters: a level of technical protection α and a level of copyright law enforcement ϕ. They then considered how rights holders' investment in DRM would vary with the strength of the law in different stylised copyright marketplaces. Similarly, a paper by Bergemann *et al.*(2005), which I discuss further in Section 2.3.3 below, uses a single parameter λ to control DRM's impact on both piracy and audience enjoyment of works, and a separate variable α to model the impact of other factors such as social norms, copyright enforcement and the availability of circumvention tools on piracy.

These modelling techniques are intuitively appealing because they promise some subtlety in representing DRM systems, without being trapped in interminable messy details. But parameterised cost models are hard to apply in practice. They are well suited to abstract optimisation, and economic theorists tend to arrange this as a centrepiece of their work, but it is not clear that this is helpful for inter-regime comparisons. Only with a a great deal of care and data collection could parameters like α and ϕ be linked to actual conditions in real markets.

When comparing DRM regimes to alternatives, it may be unhelpful to focus too closely on variables that cannot be adequately tied to real-world technologies and real world hackers. Less ambitious treatments of variations in DRM could in fact be more informative. By comparing the status quo to the strongest kinds of DRM imaginable, it is possible to see what policy consequences may arise from out of the soup of legal and technological details.

I have selected two such conceptual regimes for investigation in this thesis; they occupy distinct and interesting regions in DRM parameter space. The specific mechanisms that achieve these economic outcomes may vary — and in fact are greatly dependent on the market in question — but this is not problematic because the regimes are defined in other terms.

The first regime, which I term "status quo" or "pragmatic" or "weak" DRM, attempts to capture the kinds of effects observed from most actually existing copy protection efforts. That is to say, protection measures, surveillance and lawsuits are deployed, and have some degree of effectiveness at controlling what users can and cannot do with copyright works. At the same time, the widely opined belief that "any DRM can and will be cracked"[58] is assumed to hold true. Pragmatic DRM regimes are

---

[58]See, for example (Biddle, England, Peinado, and Willman 2002; Gray 2004; Bushing, Marcan, and

thus characterised by the practical option available to many consumers, to download cracks for their media players, install mod-chips, or simply rely on P2P media sources. The result is a system that is economically effective at moderating or staving off the strongest effects of piracy, reducing its impact without eliminating it.

The second regime, which I term "information feudalism" or "strong" DRM, is intended to model the philosophical endpoint and the ultimate success of DRM design. That is an environment where most users can access or copy works if and only if those activities are authorised by the system. Included in this classification are situations in which some circumvention option is available to users, but is so difficult or costly or risky that only the most dedicated hackers and pirates will take it.

2.3.1 Feudalism

The terms "lockdown" and "information feudalism" have been applied to a kind of expansive and uncompromising copyright law which also relies on a powerful super-position of DRM technologies to make itself essentially impermeable.[59] The use of the word "feudalism" to characterise modern intellectual property laws, though poetic, can be criticised on the grounds that there must be something deeply hyperbolic in any comparison involving the hierarchy of medieval lords, fiefdoms and serfs. But there are elements of sustainable metaphor. If the key elements of feudalism were "heavy cavalry, vassalage, enfeoffment, immunity, private castles and chivalry",[60] it is not beyond chutzpah to analogise the heavy cavalry as legal practitioners,[61] with chivalry as their training and professional ethos; recording contracts as vassalage, EULAs as serfdom and castles as the combined effect of DRM and the law itself. The Vikings, Magyars and other raiders, who motivated the rigidity of feudalism in the first place, translate quite nicely as "pirates," although the threat is now internal to the community requiring protection. Only *immunitas*, the conditional grant of governmental autonomy to local rulers, is lost in translation.

Whatever we are going to call it, the strong or locked-down or feudalist DRM

Sven 2010).

[59]A number of works have emphasised the "feudalistic" aspects of modern copyright (Drahos 1995; Drahos and Braithwaite 2002; Stallman 1997; Kretschmer 2001; Yen 2002). Many others have been similarly condemnatory with different words. Stallman famously argued that we are likely to see a "War on Copying" reminiscent of the "War on Drugs" (Stallman 1999). Such predictions have not been discouraged by subsequent events.

[60](Davies 1996, p. 311)

[61]*Cf.* (Barlow 1994, "From Swords to Writs to Bits").

regime is easy to recognise. It is the state in which something — probably a combination of extreme technical and legal measures — has actually succeeded in overcoming people's ability and inclination to pirate things. It is a state in which most people, most of the time, simply do not have the option of obtaining unauthorised copies of works.

Just how drastic a DRM regime would be necessary to reach a state of information feudalism? In the terms of the Schechter *et al.*model (2003), it appears that such outcomes could be reached by making either *d* or *e* very high, or by some combination of raising both.

The cost of extracting a first copy, *e*, could be made high by using tamper resistant trusted systems; but it could probably only be raised high enough in combination with watermarks that trace the source of analogue-to-digital leaks. This is at least a conceptual possibility, although the likelihood of such perfect deployments in the real world is much lower. Conceptually, this kind of DRM could hamper and delay the supply of newly published works into unencrypted forms suitable for file sharing.

It is harder to see how the distribution cost *d* could alone be raised far enough and uniformly enough to preclude most piracy. Open file sharing networks could perhaps be made irrelevant by a combination of legal prohibition of the software and vigorous prosecution of users. Prohibition might even be effective against anonymised networks. But it seems much harder to make privately organised copying using direct local connections, or general purpose encryption tools like `ssh` and VPNs, comparably risky.[62] These important classes of copies could only be inhibited by the mandatory installation of spyware that monitors and reports users' activities on their own computers. One hopes that such totalitarian measures cannot be seriously contemplated.

The most realistic way for rights holders to achieve feudalism would be to raise both *e* and *d* substantially. Perfect deployments may not be forthcoming, but trusted computing, tamper resistance, and collusion-resistant traitor tracing could raise *e* to the point where the decryption of new works became erratic. At the same time, ubiquitous surveillance and the closure of open file sharing networks could make the diffusion of those files that do leak much slower. The result could be that for most new publications, most users would have no option to pirate for a long time if at all.

---

[62]These kinds of copying cannot of course replace file sharing completely, because it will be much harder to obtain an obscure file from the other side of the planet. But in economic terms, the direct copying of entire media libraries (sometimes humorously termed "grand theft audio" or, when performed locally, the "sneakernet") can eliminate the need for many purchases and thereby transform a feudal DRM market to a pragmatic one.

2.3.2   Pragmatism

The status quo and other possible regimes which I classify as "information pragmatism" are more diverse than the feudalist ones, for the simple reason that they are economically much less demanding.

Pragmatic DRM environments are those that fall between the extremes of feudalism (discussed above) and anarchy (discussed in Chapter 3). They require a level of DRM and/or copyright law enforcement which is great enough to prevent dominant runaway piracy, but not so great as to eliminate piracy altogether.

Empirically, pragmatic DRM appears to be a very stable state. Despite a great number of apocalyptic prophecies (their own and their critics'), copyright industries are not about to disappear.[63] The only copying methods which could threaten to introduce true information anarchy are universal file sharing tools, and those do not appear to have grown (or shrunk) particular from the market share of 10–30% that they established in the early 2000s (Rainie and Madden 2005; Waldfogel 2010).

In terms of the model from Schechter *et al.*, the important variable in the stability of the status quo is the distribution cost d. This is chiefly set by the risk and inconvenience of finding and using file sharing tools. As discussed in Section 2.2.2, there are a number of regulatory processes acting to increase that inconvenience: the regular closure of networks (forcing users to find alternatives), the introduction of fake files for users to deal with, and the risk of being sued for sharing,

Both common sense and empirical observation would indicate that imperfect measures for raising e, such as making CDs hard to rip, or locking the files available from paid download sites, are not going to have much effect on P2P piracy, provided that *somebody* can obtain and share a plain digital copy of the work.[64] Since the commentary of Biddle *et al.*(2002), some term this the "darknet hypothesis". Copy control measures do of course interfere greatly with potentially piratical private copying outside of file sharing networks, but so long as the networks are available, the impact of technical protection measures can only be a small change in the volume of piracy.

Two things stand out about pragmatic DRM. One is that this regime is clearly the

---

[63]Lewis (2003), for example argues that rights holders do not need stronger DRM to preserve their businesses at present levels of health. This is consistent with at least some of the music industry's post-Napster sales; see `http://www.thelongtail.com/the_long_tail/2006/04/music_industry_.html`.

[64]Illustrative empirical observations here are that newly released works continue to make their way onto P2P networks at a steady pace, and that the total number of files on these networks is continually increasing (IFPI 2006, p. 21).

default option in most digital copyright markets. A corollary of the observations above is that, provided that enforcement and other anti-file sharing measures continued to keep d above zero, there is probably no need for any DRM at all in a "weak DRM" regime. In fact, determined regulatory changes would in most cases be necessary to make a transition to any of the three alternatives (anarchy, public funding, or strong DRM) discussed in this thesis. The other is that the normative and economic foundations of this status quo are deeply contradictory. Pragmatic DRM is made from strict rules that sometimes get ignored — about half the time, give or take. It is strange that it works at all.

2.3.3 The Issue of Permissive DRM

Another dimension along which theorists have tried to distinguish DRM implementations is flexibility. Does the system either allow a certain degree of circumvention, or actually allow a certain amount of copying within the bounds of its operation? Authors including Burk and Cohen (2001), Owens (2001) and Woodford (2004) have made arguments that flexibility in DRM (possibly required by law) should be used as a way to preserve "fair use" or similar rights in an environment ruled by technical protection measures.

Bergemann *et al.*(2005) model the strategic choice faced by an entertainment industry firm designing a DRM system which may allow consumers more or less flexibility in their use of the works covered by it. Flexibility is assumed to logarithmically increase consumer utility and linearly increase the supply of piratable copies. In their models, a profit-maximising copyright holder will increase DRM flexibility if consumers vary in their inherent propensity to pirate, and will decrease flexibility as external factors make piracy easier. Bergemann *et al.*also examine the relationship between flexibility and platform control. The assumptions are fairly restrictive (since there is no platform competition), but they do provide an intriguing story about entertainment companies regularly imposing inefficiently strict levels of DRM that harm consumer electronics manufacturers and social welfare.[65]

Many of the more commerically successful deployed copy control systems do fol-

---

[65]This model may not be entirely realistic, since it seems to predict that Sony, a company which controls both media platforms and a huge library of copyrighted works, should in fact use *less* DRM than anyone else wherever it can limit works to its own platforms. Sony, however, is infamous for championing platforms with the strongest DRM available at the time of development (including Mini Disc, MusicNet, XCP, MediaMax and BluRay).

low these principles of flexibility in allowing users to make certain kinds of copies. iTunes, for example allowed a licensed song to exist on up to three computers simultaneously, allows limited burning of audio CDs, and allows songs to be streamed to peers on the same LAN. Other services, such as Napster 2 and Yahoo!, use flexibility in conjunction with price discrimination: they offer more flexibility to customers who pay more money for their subscriptions. But there are important categories of uses that remain firmly prohibited by DRM. They include interoperability and format-shifting; playback on free/open source platforms; remixing, sampling and other transformative kinds of use; and the creation of backups that could survive in the long term as the platform itself becomes outdated.

These problems, in combination with some fairly injudicious technical decisions by Sony[66] have led governments to at least begin discussing whether they need to regulate DRM systems. A report by a British parliamentary committee calling for DRM-related labelling, warnings to entertainment companies about the criminal implications of DRM spyware, and requirements that DRM systems permit access to works by the visually disabled (U.K. House of Commons 2006). The French DAD-VSI legislation, implementing the EU Copyright Directive and the WIPO Copyright Treaty, was amended to require DRM manufacturers to facilitate interoperability by other platforms, although it is unclear how effective these amendments will prove in their final form.[67] Consumer protection authorities in Scandinavia, France and Germany have threatened to take action against Apple if they do not make the FairPlay DRM system interoperable with other DRM platforms.[68]

It should be noted that there is an important difference between the type of DRM that is permissive and DRM that is weak in a security sense. Although both allow a certain amount of private copying, security weaknesses are accompanied by a loss of central control and by a greater threat of unbounded P2P or Web piracy. The point of this observation is that according to the definition developed above, permissive DRM can easily be a part of a strong feudalist regime. It may even be a necessary part: if the powers of dominion are to be preserved, they must be exercised with some measure of restraint.

---

[66]See various settled class actions complaints against Sony BMG documented at `http://www.eff.org/IP/DRM/Sony-BMG/`.

[67]See `http://en.wikipedia.org/wiki/DADVSI#Interoperability_and_Apple_controversy`, visited Jan 2007.

[68]`http://biz.yahoo.com/ap/070122/europe_apple_itunes.html?.v=8`

2.4 Conclusion

This chapter has developed two representative, stylised patterns that DRM-based markets could follow. These patterns have been defined by their economic consequences rather than the precise combination of protection measures and legal rules employed.[69] Although they are just samples of the large and complicated space of possibilities, they allow for normative comparisons, not just between DRM and alternatives, but amongst different kinds of DRM. They capture the economically important ways in which technology could be harnessed to copyright law, and if conclusions can be reached about both of these regimes they can be fairly extrapolated to any serious economic application of copy protection technologies.

Feudalist DRM was defined as a state in which obtaining works by copyright infringement was either impossible or impractical for the great majority of users. It required a concerted (though not necessarily complete) combination of trusted computing, tamper resistance, traitor tracing and surveillance, suppression of file sharing and prohibition of circumvention tools. No important consumer-oriented digital copyright markets are presently displaying the characteristics of information feudalism.[70]

Pragmatic DRM was characterised by the fragmentation of a market between persistent pirates and copyright-obeying or DRM-constrained users. This has been the state of play for recorded music since the initial success of Napster; the subsequent succession of DRM deployments and enforcement efforts do not appear to have broken the basic equilibrium between authorised and unauthorised music. The film and television industries have moved in a similar direction, and books may eventually follow. At this point in history, information pragmatism is clearly the dominant regulatory pattern of consumer-oriented digital copyright.

Both the existing and hypothetically stronger DRM regimes will be considered

---

[69]In this sense, it is worth noting that a certain combination of legal enforcement and technological surveillance mechanisms might count as quite a strong DRM regime, even though it did not involve *anything* that would be recognised DRM under a narrow definition of the term as "copy and access controls".

[70]It could be argued that some markets serving specialist audiences, such as those for journal articles published by commercial academic publishers, are much closer to feudalism. The simple access control measures they use for electronic publication are overwhelmingly effective, because nobody seems to trade the pages of *Nature*, the *Harvard Law Review*, *Communications of the ACM* or the *American Economic Review* on P2P file sharing networks. This state of affairs can be readily explained with the observation that academics do not normally have to pay for their own access to journals; those expenses are borne by their institutional libraries, which cannot disregard the law so easily as individuals in the privacy of their own net connections.

more closely in Part IV, where they are compared to similar representative models of two alternatives: information anarchy and publicly funded 'virtual markets'.

Chapter 3

Information Anarchy?

Contents

This vessel, the accumulated canon of copyright and patent law, was developed to convey forms and methods of expression entirely different from the vaporous cargo it is now being asked to carry. It is leaking as much from within as from without.

Legal efforts to keep the old boat floating are taking three forms: a frenzy of deck chair rearrangement, stern warnings to the passengers that if she goes down, they will face harsh criminal penalties, and serene, glassy-eyed denial. (Barlow 1994)

John Perry Barlow has been the most famous prophet of the 'end of copyright', which — as of this writing — has not quite eventuated. What *has* happened is that there have been limited times and places in which copyright has lost its power to regulate people's behaviour. I will call that loss of control 'information anarchy'.

This chapter addresses some basic questions about information anarchy: how can

it come about? What are its important structural features? I will also briefly examine some arguments that have been made in favour of anarchy as a desirable or at least acceptable outcome.

Section 3.1 begins by discussing the origins of anarchy and how to recognise it. Section 3.1.1 proposes a definition. Information anarchy is actually a policy option, though usually an implicit one. It could be installed by removing certain acts from the ambit of copyright law, but it is far more likely to arise from a lack of determination in and expenditure on its enforcement.[71] Section 3.1.2 considers whether these two kinds of anarchy are in fact the same thing. Section 3.1.3 discusses how anarchy can arise from activities which are to a large degree private. This is one of the deepest flaws of digital copyright. Like laws governing private bedrooms, laws governing private living rooms are patchily effective.

Section 3.2 turns to the financial consequences of anarchy for those who produce copyright goods. Were anarchy to become universal, the impact on royalties and the profitability of cultural production would be large in most copyright industries, and very large in some.

The search for online artistic business models that do not depend on exclusive rights has produced a few interesting proposals and experiments. Sections 3.2.2 and Section 3.2.3 discuss some of the obvious (and not so obvious) possibilities. At this point, none of them are especially promising.

Section 3.3 discusses some of the arguments that have been made *in favour* of anarchy. These arguments vary in both their ambitiousness (*i.e.*, how much anarchy they are arguing for) and their persuasiveness. A complete evaluation of them depends on effects discussed in later chapters, and will therefore have to wait until the end of the thesis.

3.1   Characterising Anarchy

There are many degrees and stages of information anarchy. They can be seen as variations in the strength of Lessig's modalities of cyberlaw (which the reader will recall were the law itself, and its indirect effect on social norms, market structures and technical architecture).[72] The weaker the sum of these regulatory forces, the greater the

---

[71]Note that in some situations, very determined and expensive enforcement might still be insufficient.
[72]See Section 1.4.1 for a discussion of Lessig's theory of cyberlaw.

degree of anarchy. It follows that there would be tradeoffs between different weaknesses in copyright: one situation might be anarchistic because technical architectures make unlimited copying especially easy; another might make copying more difficult but be equally anarchistic because social norms against piracy were weaker.

As with DRM regimes, we are faced with the problem of how to reason about a category of situations that is so varied. Should we try to actually quantify how much copy protection systems are restricting people's behaviour? Or how little respect they have for copyright law? For the same reasons that I set out when discussing models of DRM in Section 2.3, I believe that trying to make such a precise analysis accurate would require inordinate labour if it is possible at all. One more practical alternative is to define a set of essential features for information anarchy — which could come about in the world in many different ways — and reason from those.

3.1.1   A definition

I adopt a definition of information anarchy as a situation in which obedience of the exclusive rights of copyright is unusual.[73] This does not require that copyright law be completely irrelevant. What it does imply is that copyright has little impact on the way that consumers get copies of works.

The kinds of 'situations' in which anarchy could conceivably occur range from a small market niche through to entire copyright industries or large geographical regions. As an economic phenomenon, anarchy matters most when it envelops most of the audience for some set of copyright works. Those are the situations in which it most completely demolishes the incentives provided by exclusive rights, and so it is those situations at which the definition is targeted.

The boundary between anarchy and 'pragmatic DRM' regimes of the sort characterised in Section 2.3.2 is inevitably blurry. A situation in which the median person in some audience was torn between licensed purchases and piracy would clearly count as a pragmatic DRM regime. A situation in which 90% of the population made very few purchases of digital works would clearly be anarchy. Somewhere in between, one

---

[73]The reader should note that this concept is unrelated to some positions that have been given similar names in the literature. For example, Goldsmith (1998) calls the view that cyberspace is resistant to regulation because of its international character, "cyberanarchy". But that usage is not very clear, and Post (2002) replied to Goldsmith's arguments with an article entitled "Against 'Against Cyberanarchy'", which did not use the word 'anarchy' once in its text, preferring instead to call the position "Exceptionalism".

institutional pattern turns into the other.

The presence of a grey area between regimes does not really present any conceptual problems for a normative analysis like the one I will perform in Part IV. It would, perhaps, lead to further questions if it was somehow determined that the best copyright system was a mixture of anarchy and pragmatic DRM. But, since this thesis does not reach such a conclusion, I shall not venture into the detailed geography of very weakly effective copyright systems.

3.1.2   Anarchy by choice and happenstance

One of the interesting features of information anarchy as a regulatory regime is that it can come into being by accident. Certainly, if policy makers wanted to achieve anarchy, they could do so by selectively removing parts of copyright law, or by weakening the institutions that are necessary to enforce it. But, to speak in generalisations, they show no inclination to do so.[74]

Information anarchy can also occur spontaneously. That is simply because digital copyright is a very ambitious project. Unless it is implemented in a particularly determined and competent (and therefore expensive) fashion, it may sometimes simply fail to work. If that happens, anarchy is a likely outcome. So, although no modern legislatures have enacted info-anarchist legislation, there have been places and times which can be held up as fairly good examples of information anarchy in action.

Certain college dormitories, after MP3s became widely playable, but before the music industry started its massive campaign of lawsuits, were a good example of anarchy writ small. China, for the most part, has been a good example of anarchy writ large.

The World Wide Web — or more particularly, the HTML pages that are published on it — are a somewhat different example of a medium which has developed under essentially anarchic conditions. The prevailing attitude was just to post, copy and repost things; copyright was not a concern unless you got a cease and desist letter.[75] Although there are now plenty of sites that charge for access, their shares of traffic are not large (see Figures 3.1, 3.2, and 3.3). Furthermore, it could be argued that those

---

[74]The Norwegian Liberal Party is the obscure exception that proves the rule: http://www.uv.no/politics/translated-items/culture-wants-to-be-free?set_language=en.

[75]It is beside the point whether the abundant reproduction that typified the web might be defended with legal theories such as implied license or fair use. The web succeeded because most of its users ignored such questions.

sites are not really part of the Web in the truest sense: people can't link to them in the same way, and the works and information on them cannot be propagated in the same ways that ordinary Web material can be, because they are no longer universally accessible and available for immediate and transparent citation.

So, it seems, the Web is predominantly anarchic not because exclusion is difficult, but because it was info-anarchist social norms that made it thrive.[76] This has the interesting implication that it would be important for us to understand the normative properties of information anarchy even if digital copyright law was working very well and piracy had been thoroughly marginalised.

3.1.3   In private or in public?

It is important to distinguish the kind of result one would obtain by abolishing copyright altogether, from the circumstances that prevail when it is very hard to enforce. That is because some kinds of enforcement are always going to be easier than others. At one extreme there is copyright infringement that is overt, commercial, and occurs on a large scale. It might involve setting up a factory to press Hollywood DVDs without a license. Or it might involve recording and selling copies of a musical work that is substantially but not entirely based upon another composition. Enforcement against those kinds of infringements tends to be cost effective. At the other extreme, there is unauthorised reproduction that happens in private, without commercial intent, and on a micro-scale — burning a mix CD for a friend, perhaps. In between these extremes, there are kinds of copying that inherit the economic significance of the former kind of piracy but the enforcement problems of the latter.

Because of the reproductive power of digital technology, the liberalisation of the more private and non-commercial kinds of infringement may be enough to achieve what I have defined as information anarchy.[77] That is to say, most people could be

---

[76]This is not to say that those norms are unambiguously successful; there are many things that society would benefit from having online, and for which people would be willing to pay, but which are not available because of the free rider problem.

[77]What exactly counts as "private" copying, and what exactly counts as "non-commercial" copying, can be debated. Is P2P file sharing private? Those who deny the proposition point out that it is sharing with "a million of your closest friends", but it is ultimately an interaction between the computers of two people sitting in the privacy of their own homes. The link between them may be encrypted. Is P2P file sharing non-commercial? Most uploaders and downloaders have no commercial intent — but their activities can have an impact on others' businesses. Some of the organisations involved in developing and operating P2P networks operate on a for-profit basis, while other clients and networks are developed by volunteers.

The predominant info-anarchy of the Web — sites with paid access controls are are read less than free alternatives, almost regardless of quality.
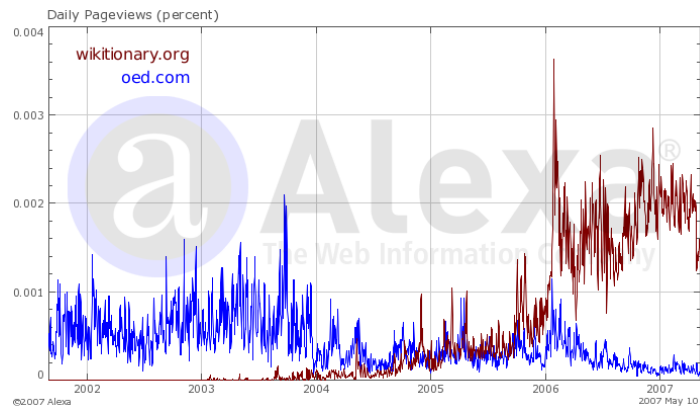
Figure 3.1: Wiktionary traffic outstrips OED traffic
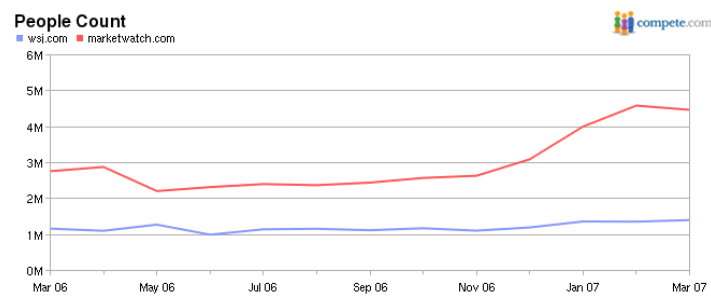
Figure 3.2: Marketwatch vs. the *Wall Street Journal*
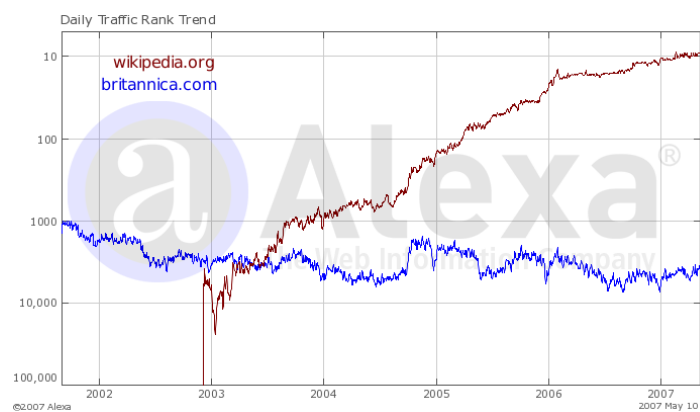
Figure 3.3: Wikipedia utterly dwarfs Encyclopædia Britannica

WSJ.com is one of a handful of successful "paywall" sites, and has subsequently overtaken Marketwatch, though largely on the strength of openly accessible content. As of 2011, it had approximately a million subscribers paying to read its website.

obtaining most of their music (etc) by piracy, while copyright law continues to be in force, and while the institutions of enforcement continue to be effective at preventing widespread and profitable piratical businesses. Indeed, a hypothetical info-anarchist government would only need to legislate an exception for copying for private or personal use, and file sharing networks would rapidly improve in quality until most people were using them.

This dependence of digital copyright upon the regulation of private behaviour is in fact one of the deepest causes of the crisis that has been unfolding for the past decade.[78] The need to put the law into people's living rooms brings with it many complications.[79]

So we know that digital copyright could be very thoroughly undermined by activities that were largely private. It follows that there are two importantly distinct kinds of anarchy: the kind that would result from unrestrained semi-private/non-commercial copying, and the kind that would result from the abolition of copyright in its entirety. Under the first, weaker kind of anarchy, the law would keep performing some important functions. In the music industry, for example, it would ensure that composers and music publishers continued to obtain a cut from musicians' live performance revenue; it would ensure that film-makers and advertisers needed to obtain permission before putting copyrighted music on their sound tracks; it would allow collecting societies to continue collecting and distributing public performance royalties. Any uses that were necessarily public or commercial would continue to be remunerated. Under the stronger kind of anarchy that would follow from a complete abolition of copyright, none of those revenue sources would persist.

The policy analysis in this thesis will work solely with the weaker kind of anarchy that extends to not-for-profit copying for private use. That variant is both far more probable (it can occur spontaneously, rather than requiring legislation) and is more

---

If greater precision was necessary, one could discuss copying that is "for private use" and "not for profit".

[78] Attempting to place a precise starting date on the digital copyright crisis would be a futile endeavour. Certainly, key events were the publication of the Lehman "white paper" (Information Infrastructure Task Force 1995); the appearance of `Winplay3`, the first real-time MP3 playing software (also in 1995; see `http://en.wikipedia.org/wiki/MP3`); and the launch of Napster in 1999.

[79] For discussions this state of affairs and its consequences, see, for example, (Stallman 1996); (Litman 2001); (Ginsburg 2002). The fact that mass copyright infringement is now performed as much by ordinary people as it is by professional "pirates" demanded significant expansion from the law of copyright itself: the difficulties with that expansion contributed greatly to the first two decisions in *MGM v. Grokster*, for example, and to the plaintiffs' difficulties in *RIAA v. Verizon* and *BMG v. Doe*. Looking further back, the legal necessity of keeping copyright law out of people's private lives was the reason that Germany first invented private copying schemes (Gaita and Christie 2004).

likely to be defended in a normative sense.[80]

3.2 Getting Paid by Anarchists

Anarchy may not be a better form of government, but it's better than no government at all. — Anon.

On its face, anarchy has the potential to deeply disrupt many of the businesses that operate in the entertainment sector. In theory, it could drastically reduce the royalty streams that flow from the sales of digital copyright works (including physical embodiments such as CDs, DVDs, or even books).

There are of course many sources of income which are not dependent on copyright. And it should be noted that the consequences of royalties disappearing would vary from industry to industry; some are far more dependent on exclusive rights than others. This is especially true under the weaker forms of information anarchy I described in Section 3.1.3, because copyright law would remain enforceable in regard to all sorts of for-profit activities. The music industry would keep performance revenues. The film industry would keep its share of box office takings. The book industry, which lacks any consistent and substantial income from performances, would be forced to adapt most dramatically.

I will investigate the importance of established non-copyright revenue sources in Chapter 7, and particularly Section 7.3. But there are also some other revenue sources which might play important roles in an anarchic information economy; some are largely unprecedented and others simply growing. These require closer examination.

3.2.1 Advertising

The most strongly established anarchy-compatible revenue source is advertising. Advertising can fund cultural production in many ways, ranging (in approximate order of intrusiveness) from text ads on websites, through online banner advertising, product placement in films and musical lyrics, commercial breaks on television stations all the way to copyright works of many sorts whose entire motive is the promotion of a particular product.

---

[80]Just about the only feature of strong anarchy that could be regarded as desirable is the very liberal environment for the creation of derivative works — and the creation of derivative works can be liberalised in other, far less drastic, ways.

Advertising is in a sense the floor through which anarchic funding of copyright works cannot fall. Even if there is no money from anywhere else, there will be advertising dollars on the table for authors and artists to take. But there may be several serious problems with taking it.

The biggest problem with advertising is that there is simply no guarantee that it solves more than a small fraction of the free rider problem. A great e-book may be worth ten dollars an hour for me to read, but advertisers might only be willing to pay a dollar an hour to place links in the corners of the pages.

Secondly, advertising may distort great work even while it funds it. Bloggers may replace long investigative pieces with short, attention deficit-inducing posts that generate more pageviews and ad revenue.[81] For instance, a narrative may be sliced up into simpler, less memory-demanding portions in order to fit television advertising in between. As one observer put it, "It is difficult to produce a television documentary that is both incisive and probing when every twelve minutes one is interrupted by twelve dancing rabbits singing about toilet paper"[82].

3.2.2   Tip jars and patronage

Copyright is obviously not needed for artists to solicit *voluntary* payments from their audience. Some observers have lauded tipping as an important 'post copyright' or 'parallel to copyright' source of artistic income.[83]

Tips are highly susceptible to free riding, so they cannot be expected to get close to the true financial demand for an artist's work. They may, however, be sufficient to cover some artists' living expenses, or handily complement other sources of income.

The solicitation of tips, in the form of donation buttons, has therefore become quite common amongst some artists whose work is aimed specifically at the web.[84] Others sell t-shirts. Because the medium is inherently so incompatible with exclusion (exclusion is possible but it cuts off the universal accessibility and inter-linkability that

---

[81]See, *e.g.,*, (Janelle 2007) although the extent of this phenomenon may be dependent on the subject matter; `http://www.successful-blog.com/1/the-short-post-vs-the-long-post-who-will-be-victor/`

[82]The Web attributes this statement to Rod Serling without citation.

[83]See, for example (Barlow 2000; Helman 2010).

[84]Some fairly prominent cultural projects work this way — *This American Life* and *Four Eyed Monsters* (`http://www.wired.com/entertainment/theweb/news/2007/06/youtubefest`) are excellent examples. Occasionally, they write up reports on their experiments: `http://www.kuro5hin.org/story/2003/4/27/195833/305`

make so many people want to use the web in the first place), these artists have little alternative.

The trick to successful tip-based fundraising may lie in the relationships that artists can form with their audiences. Readers give more, and more often, when they feel like they are helping someone they know, and building resources for a community of their peers.[85]

A dependence of artists upon tipping may also be philosophically problematic on grounds other than practicality. It requires artists to be especially competent at cajoling their audiences — as if it were not enough to ask that they are talented and dedicated to begin with. It makes their livelihood dependent on their ability to manage relationships, rather than just the merit of their work.

Some have argued that under anarchic conditions it should be online audiences, rather than just artists, who initiate the practice of tipping. Given the many millions of pirate file sharers around the globe, it is a reasonable presumption that a non-trivial number of them will want to send payments — to salve their consciences, if nothing else. This social institution has, sadly, not eventuated. Their most publicised attempt, by musiclink.com,[86] was to provide downloaders with a universal tip jar where they could specify which artist they wanted to pay; the site would aggregate the payments until there were enough of them to make tracking down the artist cost effective. Limited donations and difficulty in distributing them spelled the end of the project. Collecting society employees have also informed me that it would be difficult for their organisations to assist in such projects.

Patronage differs from tipping in that there are many fewer sources, making much larger payments; because of this the patron is usually in a position to negotiate over the work to be produced, exerting control if they wish. Patronage from wealthy individuals, corporations, governments, churches, and other organisations has often played an important part in cultural production. The importance (and proportional influence) of this process would certainly grow under information anarchy.[87]

---

[85]Rheingold called the phenomenon by which digital communities gather around members in need virtual "barn raising" (Rheingold 2008, Chapter 1). Successful fundraising for creative digital projects often has a similar character.

[86]MusicLink was originally called Fairtunes. Although their site is no longer archived, an account of their objectives can be found at `http://www.bizreport.com/news/66/`.

[87]Along with the concerns it raises; see, for example, (Netanel 1996).

3.2.3   "Street performer" protocols

Tipping and patronage are based on one-to-one relationships between artists and members of their audience. Other, more communal, institutions are also possible. Stephen Breyer, now Justice and resident copyright sceptic of the Supreme Court of the United States, described the possibility as follows:

> *If in the absence of copyright protection, an initial publisher of, say, a text with a long time horizon and large initial expenses fears that copying will destroy his profits, buyers — anxious to obtain the book — may find ways of assuring him enough revenue to operate profitably. Buyers, individually or in groups might contract to buy the book in advance of publication, before copying is possible. Or organised groups of buyers may place a large enough order for an initial printing (either before or after publication) to assure the publisher a profit.*
>
> *If many, or most, of the potential buyers for a certain particular type of book can be organised into a group that can negotiate with publishers, the danger that copying will make book production unprofitable diminishes. For if a publisher can convince the group that he needs a substantial "contribution" to cover the costs of production, the group will often prefer to pay the money than lose the book.* (Breyer 1970, pp 302–3)

Breyer's article goes on to develop a very rough sketch of how such processes would work for particular kinds of books. The idea was ahead of its time, both because copyright was working better in 1970, and because it is in most cases prohibitively expensive for audiences for particular works to organise themselves ahead of time — unless, that is, they can find a way to use computers to do the organising for them.

Some thirty years after Breyer's commentary, computer scientists, authors and entrepreneurs were rediscovering the idea. One of the clearest proposals is Kelsey and Schneier's (1999) "street performer protocol" (I will also use the abbreviation SPP).[88]

---

[88]Similar protocols have been used and discussed by others. An experiment by author Stephen King was the most widely discussed example (see discussion below). Patent applications were filed for at least two similar proposals, which do not cite Kelsey and Schneiers' work. One is a comparatively reasonable patent (Megiddo and Zhu 2000) assigned to IBM which anticipated Rasch's "Wall Street Performer Protocol" (Rasch 2001). The other is a late patent application by a startup company called Artistshare (Camelio 2003), for which Kelsey and Schneier's paper appears to count as invalidating prior art.

The name is a little confusing, because the protocol is an attempt to escape from the extreme free rider problem that might follow if artists simply 'hung out their hats' on the web.

To see how the street performer protocol works, imagine that a musician named Bethany has recorded an album of some merit, and is wondering how she can earn a reasonable amount for it without the assistance of copyright.

She begins by releasing a sample of the album to demonstrate its quality. That might be a couple of songs in full, and/or shorter clips from all of the tracks. At the same time, she announces the price she's asking, to cover risk and the costs of composing and producing the album — perhaps, in our example, $150,000.

Bethany's audience is invited to make pledges towards that sum. Rather than being simple payments, these pledges are contingent upon the release of the full album. If the audience bites, and the pledges received reach $150K, Bethany releases a free, high quality, version of the album online, and the pledges are called in.

In order to make the pledges easily enforceable, they could take the form of payments in escrow to a trusted third party (TTP). If the work is published, the TTP will transfer the cash to Bethany. If, after a pre-specified period, the work is not published, the money is returned to those who contributed it.

Two interesting variations on the street performer protocol should be mentioned. One version, described by Rasch (2001) under the name of the "Wall Street Performer Protocol", is designed to fund the development of free/open source software. Because consumers can specify *ex ante* the software they want to a degree that is simply not possible with cultural works, the order of pledging and production are reversed. Consumers offer amounts of money, in advance, for the addition of specified new features to existing free software. Other users can add their own pledges over the top, until a sufficient bounty accumulates to motivate software developers to write the code.

Trusted independent reviewers then examine the proposed code changes to see that they operate as claimed and are sound from a software engineering point of view. If so, the programmer collects the bounty. Because the WSPP deals with goods that can be specified in advance, it does have the useful advantage of reducing the risk that has to be borne by information producers — there is no need to do the work in advance of knowing that there is monetary demand for it. WSPP's applicability is in fact not limited to software, but to goods which have that property of *ex ante* specifiability.

Another variation of SPP was proposed by Paul Harrison (2002). The idea was to try to guarantee that each contributor could "gain a reward (in terms of production of the public work) in proportion to their contribution". In what Harrison called the "Rational Street Performer Protocol", participants get to make offers of the form: I pledge to pay x dollars for the creation of a public good, provided that each dollar is matched by contributions of y dollars from others. This appears to solve the free rider problem: each participant can limit the amount of free riding off their contribution.[89]

But, upon careful examination, the "rational" protocol is not really any more rational than the original SPP. Ordinary SPP is really a pledge of the same (x, y) form, but with y set to be the total price the artist has called for, divided by x, minus one. If there are incentive problems with the regular SPP (and I will describe them shortly) the same incentive-incompatibility also arises in RSPP. The one thing that RSPP does offer, in actuality, is a way to apply SPP to public goods of an open-ended nature, such as donations to charities or political parties. Rather than having a threshold of contingency, the pledges can be tied together at different total amounts.

The idea behind all of the variants of the street performer protocol, of course, is to introduce some of the contingency that is present in ordinary market sales, to a voluntary collective contracting process that does not depend on copyright. There is contingency in two directions: if people don't pledge, then they are less likely to get access to the work, and if they don't get access to the work, they don't have to pay for their pledges.

SPP does have some obvious drawbacks. It requires that people who are willing to pay for a work wait until the protocol succeeds, before they get a copy. It requires authors to invest in production of works in advance of payment, when sometimes those works will never be publishable.[90] It does not combine particularly well with traditional print/royalty publication.[91] But most seriously, it is not clear that the street performer protocol always solves the free rider problem it set out to address.

To see why, consider how the outcome for an individual audience member is determined by two variables: the sum of pledges offered by others, and the size of their

---

[89]The pledges in fact appear to form the conditions for a Lindahl equilibrium, which is known to be a Pareto-efficient way of funding public goods (Cornes and Sandler 1996, p. 204).

[90]If the protocol does not complete, it is important for authors not to publish their works anyway, because it reduces the credibility of their threat to not publish and thereby undermines incentives in future executions of the protocol.

[91]See for example (Zimmerman 2003, note 98 and accompanying text). Some innovative hybrids may be available, such having the online release after the paper release, but bringing the date forward as more pledges are made.

own pledge. If the sum is greatly below the cutoff, the work will never be released. If the sum is above it, the work will be published regardless of their offer. It is only when the sum of other pledges is just slightly below the cutoff that it becomes rational for a self-interested consumer to offer a pledge. As the number of participants in the protocol increases, the probability that the sum of pledges will be in that small critical zone shrinks.

The bottom line is that street performer protocols only work when audiences are altruistic, or when they are irrational in their understanding of the protocol itself. Either of these reasons can apply, but neither can be depended upon.

Real world experimentation

The most widely cited SPP-like experiment was conducted by the horror writer Stephen King. In 2000, King announced that he would publish a serialised version of his novel *The Plant*, with the condition that each chapter would be released when 3/4 of the people who had downloaded the previous chapter paid $1 for it. The results of the experiment were ambiguous. At first, a large proportion of downloaders paid, but the number dropped with each chapter. King relaxed his 3/4 threshold, but after a few more chapters, he decided to pursue other projects instead.

The experiment was in some ways successful. Stephen King did collect more than USD $700,000 for a few chapters of his book, and enjoyed very high profit margins (King 2001). It was widely reported that his experiment had failed (New York Times 2000), although King himself did not appear to take such a simple point of view (King 2000); he simply implied that he could make *more* money from traditional publishing projects.

It was probably unrealistic to expect three-quarters of a mass readership to pay what was effectively a paper book price for a digital novel. It makes more sense for an author to say how much they need to keep writing, and ask for that amount. It would never work for King — one of the most commercially successful writers of all time — because his income is so astoundingly large that his readers might rebel against it. Without the compelling personal story of a favoured author in need, the logic of free riding would inevitably come to dominate consumers' decisions.

Further real world experimentation with SPP variants has also yielded mixed results. A company called Artistshare, which aimed to be an SPP-based record la-

bel, did not appear to obtain many audience payments.[92] In the free/open source software world, a few high profile contingent fundraising drives succeeded: one by kuro5hin.org, a collaborative media site (Foster 2002); one for the open source release of Blender, a previously proprietary 3D modelling and animation package,[93] and one by Linux Weekly News.[94] However, these were fairly isolated success until the 2009 launch of the "Kickstarter" platform, which offers a compelling SPP implementation to projects in a many fields of creative endeavour.[95] Kickstarter has been successful at raising moderate amounts of money for a steady trickle of projects.

Attempts to construct general-purpose marketplaces for WSPP bounties have also come and gone.[96] A small project called the Software Bazaar operated with some success (many of the bounties were claimed by Russian programmers) until its operator, an academic mathematician, was distracted from running it. Source Exchange, a dot-com boom era startup, tried unsuccessfully to use bounties for larger commercial projects. The GNOME Desktop project offered a set of bounties for various feature implementations.[97] A few were claimed, but there is little evidence that they made any significant contribution to the GNOME development process; GNOME is a huge project and the size of the features involved seemed minuscule. For a time, a site called opensourceexperts.com collated information about bounties offered by different sources.[98]

It seems that contractual communal funding models have only been seriously effective for co-ordinating the activities of firms. Examples include the Apache Software Foundation (funded by numerous firms using Apache), the Mozilla Foundation (funded largely by AOL and Google), and OpenOffice.org (funded largely by Sun and Novell). One possibly significant difference is that firms participating in these alliances have far more information about each others' valuations of the proposed software, allowing them to insist that each participant pays their fair share. As yet, there have been few successes in collecting large amounts from large numbers of users.[99]

---

[92]http://www.alexa.com/data/details/traffic_details?q=&url=http://www.artistshare.com/home/default.aspx (click on "max" to see the entire dataset).

[93]See http://www.blender3d.org/Foundation/?sub=History (2003).

[94]http://www.lwn.net/Articles/5838/ (Jul. 26, 2003).

[95]http://www.kickstarter.com

[96]A number of these are recorded at http://www.ms.lt/en/workingopenly/markets.html.

[97]Unclaimed GNOME bounties can be seen at http://www.opensourcexperts.com/bountylist.html?bountytype=1&cat=49.

[98]http://www.opensourcexperts.com/bountylist.html?bountytype=1

[99]The Blender fundraiser appears to be an exception — but the asking price of 100,000 euro was

Even if ways were found to raise more money from broad-based WSPP processes, there might be unexpected and detrimental side effects. Paying some programmers can easily "crowd out" the voluntary contributions of others. There is clear evidence for crowding out in other kinds of volunteer organisations; in the case of free software projects it can take effect through psychology ("why should I do for free what others are paid for?") and through reductions in the openness of the development process.[100]

Conclusions on Street Performer Protocols

Street performer protocols are a briefly promising but ultimately limited source of income for producers of public goods. When the number of participants is large, they have no mechanism for requiring consumers to pay their fair shares of the bill.

Communal contracting may be of some assistance in two situations: (1) when the group of prospective participants is small and members have a lot of knowledge about each other; and (2) when cultural producers can establish strong psychological connections with members of their audience, thereby side-stepping the logic of free riding. In the latter case, it is not clear that SPP is much better than a tip jar anyway.

3.3 The Case for Anarchy

Anarchy is not without its defenders. Few authors go so far as to argue outright for the abolition of copyright,[101] but quite a number seem to adopt a gleeful tone when they suggest that its demise might be inevitable.[102]

The case for information anarchy, surreptitious or otherwise, can proceed on many grounds. It can be maintained (1) that adequate finance for copyright production is available from sources other than exclusive rights; (2) that financial incentives are much less important than is widely supposed; (3) that incentives may even work to undermine or corrupt the production of important art; (4) that production of particular information goods can in many ways be more efficient without the encumbrance of

---

really a liquidation price for such a large software system.

[100]On all of these points, see (Hill 2005). WSPP-type targeted incentives are probably not as bad in that regard as hiring some developers, because they will not partition the community around each program as drastically into paid 'insiders' and unpaid 'outsiders'.

[101]Though some do: see for example (Martin 1995).

[102]See for example (Moglen 1999; Kelsey and Schneier 1999; Barlow 2000; Love 2000; Clarke 2001). Other authors (Breyer 1970; Nadel 2003) come close to the same territory.

exclusive rights; and (5) that the burdens of copyright, in enforcement and loss of liberty, greatly weaken or even outweigh its incentive benefits.

The beliefs that underlie the first three arguments will be examined in greater depth in Chapter 7. I have already argued that SPP-type financing is not sufficient to sustain argument (1) in general, but there are many other funding sources which are. There is certainly evidence to support argument (2); not enough to dispense with the need for incentives, but enough to reduce art and culture's dependence on them. Argument (3) is greatly predicated on one's philosophical viewpoint, as I will explain in Section 7.2.

Argument (4) is very powerful. There is overwhelming evidence that certain kinds of information production thrive in environments that allow open access to, and adaptation of, large bodies of material. Without the various transaction costs required by exclusive rights, these kinds of open information ecosystems are the best ways to produce some copyrightable goods. Grand endeavours like the World Wide Web and the free/open source software diaspora, as well as smaller ones like Wikipedia and the mash-ups found on YouTube, illustrate that clearly. Benkler's work on the commons and "peer production" is the most thorough examination of the factors that determine when exclusive rights are detrimental to production efficiency.[103] Despite its great sharpness, this argument (4) for anarchy, is limited in its reach. It does not affect classes of copyright works which are not readily amenable to peer production. The distinction will be important in Section 5.4, where I argue that publicly funded remuneration systems are much easier to construct in those situations where peer production is not in play.

An evaluation of argument (5) requires an estimate of the direct and overhead costs of digital copyright. I provide one such account in Part IV, and particularly in Chapters 6, 8, and 9.

3.4   Conclusion

In this chapter, I defined a 'regulatory' regime called information anarchy. Its distinguishing feature is that under it, copyright is irrelevant to the way that most people obtain digital works. Anarchy can operate on many scales: from a market niche, to a medium, to an entire cultural economy.

Information anarchy could arise spontaneously, if the resources devoted to digital

---

[103]See (Benkler 2003; Benkler 2006). Many other authors have explored : Lessig, etc.

copyright enforcement were insufficient to make a significant proportion of people obey it, or if the law was altered to permit non-commercial copying. The abolition of copyright altogether would cause an even stronger kind of information anarchy, but even its strongest sceptics have hesitated to advocate such policies.

Under information anarchy, copyright industries would first of all fall back to existing income sources that are not related to the sale of digital works to consumers. The proportion of revenue that is 'safe' in that sense varies widely from industry to industry. Authors and book publishers are in the worst position.

An assortment of more exotic proposals have been made for ways that authors and artists could collect money from their audiences without exclusive rights. Some of these are basically gift or tipping models, which suffer from poor returns (since the free rider problem is still present), require competent cajoling from artists, and are also subject to more philosophical complaints about artists' sovereignty.

Of more interest are the street performer protocol and a related family of collective contracting processes. These business models try to reintroduce a degree of the explicit contingency that makes consumer purchasing behaviour different from tipping behaviour. Unfortunately, the protocol is not completely successful in this attempt. Section 3.2.3, explained why the free rider problem can persevere under these protocols.

It may not matter that the systems for financing cultural production are inadequate under information anarchy. Section 3.3 pointed out that there are other factors that could conspire to make anarchy attractive. The strength of these factors will be evaluated in Part IV.

To foreshadow the analysis in those chapters, anarchy as a 'regulatory regime' would have some large and concrete benefits. The *a priori* assumption would, however, be that they are outweighed by the harm that would be done in the loss of an incentive system that specifically ties payment for acts of authorship to the appreciation of audiences. Chapter 7 will show that this is surprisingly less certain than one might expect, at least in the case of the music industry. But, such surprises aside, information anarchism remains most attractive to those who adopt radical philosophical positions about the nature and role of valuable art, and are willing to condemn a great deal of what the existing entertainment industries do and do not do with their resources.

Part III

Virtual Markets for Virtual Goods

Mercatus Ex Machina

Digital copyright is at present suspended in an incoherent state between the extremes of 'feudalism' and 'anarchy'. If there is an elegant balance to be struck between the restrictive and the permissive in copyright, existing regulatory processes show few signs of reaching it. We should accede to information pragmatism and accept this messy compromise only if we discover that is has counter-intuitive virtues, or if there is no better choice.

It may be that, in order to find a more satisfactory middle ground between information anarchism and information feudalism, it is necessary to step away from the metaphor of property rights and into the territory of "compensation without control",[104] at least with regard to noncommercial copying. The contention of Part III of this thesis is that while digital technology is destabilising the economics of exclusive rights, it may be simultaneously opening dramatic and novel possibilities for alternatives of just that sort.

In the following chapters, I explore one such class of alternative compensation systems, which are publicly funded and not dependent on a strong central notion of 'property'. These systems would create what I term 'virtual markets' to provide incentives for information production, in much the same way that an actual marketplace provides incentives for the manufacture of physical goods.[105] At the same time, they would allow universal noncommercial access to information goods, avoiding the deadweight loss and high overheads of DRM exclusion systems.

With an appropriate design, these virtual markets could be decentralised, efficient and in some senses democratic[106]— the very qualities of successfully operating markets which lead so many observers to favour them. Within the proposal, there is a crucial role for government: the use of taxation to solve the underlying free-rider problem.

---

[104]The terminology is Lessig's (2001, p. 201), though he has no veto over the way I develop the idea here.

[105]The choice of the word 'virtual' is discussed further below — see *infra* § 5.1.

[106]Whether and how markets are actually democratic is a complicated question, and as much a matter of definition and context as a matter of fact. Clearly, there are differences in the way that marketplaces and electoral institutions aggregate preferences, as observed for example by Brennan and Hamlin (2000, pp. 83–4). The market may disenfranchise more participants than the ballot box, and give disproportionate power to wealthy individuals, but it can also provide a greater channel for the expression of the *intensity* of preferences.

At the same time, the architecture of the system ensures that the government has little or no control over the way that those funds are distributed; such power is devolved to the end-users and cyber-citizenry who pay for, and should benefit from, the scheme. This, I would argue, could amount to a kind of 'information democracy'.

The combination of decentralisation and the non-scarce nature of digital information goods allow publicly funded remuneration systems to survive the modern presumption that governments should not involve themselves in affairs that might otherwise be left to a more-or-less free market. This presumption appears regularly in the copyright literature when the possibility of public funding is discussed,[107] although it should be noted that in the case of digital copyright the position lacks important economic[108] and philosophical[109] justifications that it may enjoy in other situations.

So, placing the standards of copyright aside for the moment, imagine a tax-paying Internet user — I shall call her Alice — who wishes to use digital works in an unrestricted fashion. How might the society in which Alice lives reward the authors, musicians, composers, film crews, programmers and bloggers who produce things she values?

---

[107] This is a point which is sometimes made directly (Perlmutter 2001, p. 167; Davies 1994, pp. 149–55) and sometimes with regards to the comparison between market-based and judicially administered licensing systems (Gordon 1992, § 4; Merges 2004). The concern has both a direct political aspect related to freedom from censorship, and an economic dimension in the tradition of Hayek (1945).

[108] The economic case for free markets has been formalised in the form of the first and second fundamental theorems of welfare economics which state, respectively, that market equilibria are Pareto-optimal, and that any Pareto-optimal outcome can be achieved by first redistributing wealth using lump-sum transfers and then relying on market exchanges. There are, however, many causes for "market failure" which render these theorems inapplicable, and the public good nature of cultural works is prominent among them.

[109] A principal philosophical point in favour of property and markets is that they support a private ordering of affairs, allowing individuals to pursue their desires in life autonomously and without reliance on centralised, potentially flawed social decision making processes (Hodgson 1988). But it should be noted in reply that financing the production of cultural and informational goods, if it is to be done systematically and in the face of the free rider problem, is inherently a social problem. The exclusive rights of copyright, where they are able to fulfil that role, appear to be more 'private' than they are 'social': they vest in private persons and facilitate a privately organised cultural economy. But upon examination they are a social institution necessarily created and underwritten by the power of the state. This is in contrast to the rules of material property, which began as codifications of pre-existing practices and which remain closely connected to the way that humans interact with land and chattels. A weak or absent government does not normally lead to the end of proprietary possession — just greater expenditures in securing it.

An exception to this distinction is the exclusive right of first publication, which does correspond to a relatively natural property of the human world and was thus distinguished in *Donaldson v. Beckett* (but compare Deazley's observation that this may have been confused with a consequence of the ownership of manuscripts themselves (Deazley 2003, pp. 271–3), in which case the exception may not be so much of an exception)

If there was some way of measuring what Alice downloaded — or better yet, what she listened to and what she liked — it would be possible for a government to pay the producers of information goods according to the social value of their work.

This notion has a great many conditions attached to it, which I will examine in the following chapters. Chapter 4 begins by surveying the literature on the use of public funding for information production (which dates back at least as far as the 17th century); it also discusses some of the past and present "actually existing" systems which have performed this kind of role. Chapter 5 introduces the virtual market: a way of employing digital technology to measure the value that consumers place on digitised works, and to thereby achieve an efficient distribution of central funds to information workers. It also discusses in some detail the differences in feasibility of virtual markets for different classes of copyright works and users (Section 5.4), and the many factors that must be considered if they are to operate securely (Sections 5.2.2 and 5.2.3). Section 5.3 deals with the inevitable downside of public provision: the taxation that must be used to fund it.

If all of the technical and economic design objectives are taken seriously, the final picture is one of a complicated but nonetheless viable alternative to exclusive rights.

Chapter 4

The History and Literature of Public Funding for Information Production

Contents

This chapter briefly reviews historical experiments with and discussion of the use of public funding as an alternative to copyright and patent systems. The very idea might seem quaintly anachronistic to modern sensibility, but we may wish to dust it off for scrutiny if we conclude that there are persistent problems with the present system of exclusive rights enforced by DRM and law.

It is best to include patents in the discussion, at least briefly, because there is a great deal of basic similarity between the two cases; although important details do diverge, the premise is largely the same. Some authors have even examined the two subjects simultaneously: "public funding for intellectual property production", as it were. So I will begin with a discussion of governmental alternatives to patents (Section 4.1).

Section 4.2 then moves on to consider taxation funded remuneration systems for copyright subject matter. There are two systematic precedents for the use of taxation to pay the authors and publishers of copyrighted works: "private copying" levies and "public lending rights". The literature on more general systems of public funding to

replace copyright (of which this thesis is a part) has expanded greatly in recent years. This research is discussed in Section 4.2.3.

A related subject, which I will not address, is the character of publicly funded academia as an information production system. There is a small but certain overlap between the research outputs of universities and the fruits of copyright and patent incentives.[110] The institutions that guide those outputs financially are very different, of course. Academic research is generally driven by a strange alliance of intellectual curiosity and government prioritisation; exclusive rights drive the development of inventions and works with fairly immediate market values. When universities do systematically concern themselves with the "market" for their research, it is usually because they want to make money from copyright or patents! My point in making these observations is to suggest that the function and design imperatives of university systems are different enough from those of exclusive rights that they are a complement, rather than a replacement. This is especially true in the case of copyright, and for that reason, I will limit the discussion in this chapter to market-replacing government interventions.

In any of these cases, the efficacy of governmental reward systems depends greatly on the nature of the information produced and on the organisation of the reward mechanism.

4.1 Prizes and other Rewards for Invention

The concept of granting rewards as incentives for the production of information goods is not new; it has been discussed as an alternative to exclusive patent rights for centuries.[111] In 1660, a utopian essay (probably written by the prominent scientist and patent sceptic Robert Hooke[112]) described in some detail a fantastic society which benefited enormously from organising public rewards for valuable inventions.[113] The idea was not purely utopian; MacLeod gives examples of inventions for which retro-

---

[110]Confirmation of this claim is left as a diversion for the reader.

[111]A historical treatment of ideas about both patents and reward-based alternatives in England before 1800 may be found in (MacLeod 1988, Chapter 10).

[112]Robert Hooke (1635–1703) is remembered today for discovering the Newtonian mechanics of springs; for his observation that plants were made from microscopic structures, which he termed "cells"; and for significant inventive contributions to microscopy, telescopy, clockwork, and mechanical engineering.

[113]The work is described by Macleod, (1988, p 191). It was constructed as an extension of Francis Bacon's earlier utopia, *New Atlantis*, and credited to "R. H. Esquire".

spective rewards were provided,[114] while Wright points out two significant inventions which were prompted by "bounty" rewards, announced in advance.[115] Robert Macfie, a British MP and free trade advocate, agitated for an organised reward infrastructure during the mid-to-late 19th century,[116] but the schemes of Macfie and his fellow travellers were defeated by an inconstant zeitgeist.[117]

During the Second World War, Michael Polanvyi, who appears to have been aware of the free trade movement's antipathy toward patents, but not Macfie's work on developing alternatives, constructed a more extensive version of the argument for rewards.[118] Since then, economic analyses which address the question have periodically concluded that the case for the patent system is not clear, and that either publicly contracted research, or taxation-funded systems to reward inventors, might well be more efficient.[119]

The last decade has seen an explosion of interest in reward systems. Steven Shavell and Tanguy van Ypersele (1998, 2001) constructed a model to compare patents, rewards and a mixed regime. They concluded that the rewards would be superior under many circumstances.[120] Working with Shavell, Stephen Calandrillo (1998) set out a general argument for the superiority of rewards against both patents and copyright. Kremer (1998) proposed a more exotic scheme in which governments purchase patents for the public domain using an auction-based information revelation system. Abramowicz (2003) provided an excellent survey and synthesis of the papers by Kremer and Shavell & van Ypersele, as well as an earlier pharmaceutical-specific buyout proposal by Guell and Fischbaum (1995)[121] and an argument for a consumption subsidy system by Lichtman (1997).

Duffy (2004) provides a more critical reply. He argues that the debate over the deadweight losses of patents is in fact a resuscitation of an older controversy over marginal costs in markets with "natural monopolies" created by high fixed costs of

---

[114]*Id.*, at 191–3.

[115]See (Wright 1983), note 15 and accompanying text. See also (Abramowicz 2003, note 15), for further sources and commentary on the historical use of rewards.

[116]See, e.g., (Macfie 1869, 84–7).

[117]Machlup & Penrose provided an excellent account of the historical context for Macfie's arguments; see (Machlup and Penrose 1950).

[118](Polanvyi 1944).

[119]See (Arrow 1962) (discussing the desirability of governments employing contractors to perform research and development); (Wright 1983) (for a comparison of patents, rewards and contracts).

[120]Shavell and van Ypersele's model is discussed in greater depth in Section 7.4.1.

[121]Pharmaceutical patents are a comparatively promising domain for these proposals because the relationship between patents and products is fairly clear.

entry. Coase's (1946) influential argument against deploying public sector operations to replace those markets held that the deadweight losses were likely to be smaller than the inefficiencies caused by imperfect governmental information and the distortionary costs of tax-funded schemes. Duffy charges that the advocates of alternatives to patents have not adequately addressed these obstacles, or the relationship between the "IP" and "natural monopoly" problems. The former, he maintains, cannot be solved without solving the latter. These a claim is plausible on its face, but the details may or may not bear it out. In particular, variations in the availability of accurate *ex post* valuation data (see Section 7.4), and the presence of taxable complementary goods (see Chapter 10) mean that it is in fact hard to generalise about the achievable efficiency of government interventions from copyright to all natural monopolies in traditional industries.[122]

James Love and Tim Hubbard (Love 2003; Hubbard and Love 2004) developed an interesting proposal for an international regime that cross-references (on a country-by-country basis) government funding under reward regimes and in public sector medical research programmes against implicit funding of medical research through citizens' payments of high drug prices.[123] The idea was that existing treaty obligations to provide patents be relaxed for countries that contribute a comparatively large proportion of their GDP in direct funding for medical research and development. Proposals of this sort work from the view that we do not really know the best ways to fund medical research, and we will not really know until some of the more sophisticated options have been tested.

Interest in drug- and vaccine- specific proposals has recently become very serious. Attempts to promote non-patent R&D financing systems within the World Health Organization appear to be making some headway.[124] Perhaps even more significant is the ambitious undertaking of the Bill and Melinda Gates Foundation, which announced in 2003[125] that it was looking for efficient ways to spend very large sums of money on

---

[122]It is possible, of course, that in some natural monopoly industries governments could have access to sufficiently good information about demand to be able to provide the good more efficiently than the private sector, or to design welfare-improving regulation of monopoly providers. This may even occur in regular practice, but the question is beyond the scope of this document.

[123]The level of R&D funding provided by a particular volume of royalty payments would be determined by multiplying by pharma's average level of re-investment in research. Investments could also be adjusted to give extra weight to work on neglected diseases, and to work which is scientifically open and therefore produces more positive externalities for other projects.

[124]See `http://www.cptech.org/ip/health/who/59wha/index.html` for information on the relevant developments at the WHO.

[125]`http://grandchallengesgh.org`.

developing treatments for neglected diseases.[126] Prizes are one of the mechanisms the foundation is using.

Unfortunately, the point at which patent rewards are considered as a concrete and detailed solution to problems in the economics of medical research is also the point at which the analogy to copyright-related schemes breaks down; the size, risk distribution, and market characteristics of the projects to be funded are completely different. There will therefore be little to learn about *copyright* from the recent debate about the best way to spend the Gates foundation's money (CGD 2005; Light 2005; Maurer 2005), or from the ongoing efforts to use these methods in drug and vaccine development.[127]

A rather different example of prizes being used to stimulate innovation was the system of "inventors' certificates" employed by the former Soviet Union[128]. The example is hardly inspirational. It does however seem that the problems of research and innovation in the U.S.S.R. stemmed from much larger difficulties with industrial organisation, rather than from the prevalence of rewards in place of patents.[129]

There is little doubt that the most significant difficulty with industrial rewards is the need to index them to the value of inventions. The extent to which an invention is adopted provides significant hints, but such information may be difficult to collect

---

[126]As an aside, it is amusing to note that Gates' endeavours in this area amount to a form of privatised government. The world pays Microsoft a Windows tax, a portion of which is spent on healthcare for the developing world. Rather than legislating the tax in the usual way, it has been enacted by a particularly monopoly-friendly combination of trade secret and copyright protection for software goods (Samuelson 1984; Gibson 2005). That hybrid intellectual property regime — now strengthened by the addition of patents — has been especially effective at limiting software interoperability and thereby capturing network externalities. The crispest illustration was drawn in the *Caldera v. Microsoft* litigation (Schulman 2000), which showed not only how competitors can be locked out by copyrighted trade secrets, but how anti-trust remedies work too slowly and assist the wrong parties to be of much use to the public; Caldera has since become infamous for filing lawsuits of questionable basis against members of the free and open source communities. But if the monopoly rents the public is paying are spent on neglected diseases, a normative critique of the status quo becomes less certain. If one can be sustained, it is that the Gates Foundation grants ($1.36 billion in 2005; see `http://www.gatesfoundation.org/MediaCenter/FactSheet/`) are only a small portion of Microsoft's $10 billion annual profits.

[127]See, for instance, the Advance Market Commitment for vaccines project, `http://www.vaccineamc.org`.

[128]The Soviet Union applied reward models to both inventions and writing. See (Baxter 1973, 14–21); (Boguslavsky 1979, 130–5); (Levitsky 1964); (Loeber 1984), for descriptions of Soviet "intellectual property" systems.

[129]On the problems of Soviet industrial R&D, see, for example, (Kabalina and Clarke 2000); (Graham 1993, 179–80). I have not been able to find any detailed information on the way that the Soviet reward system was operated (which is most likely to have existed in Russian). Personal communication with several investigators who have specialised in the area suggested that there was a persistent lack of high-quality introspective examination of the Soviet industrial R&D system.

and may lead to perverse incentives for reward recipients (it might be in an inventor's interest to actually pay people to use certain products in order to increase the size of their reward!) Even if complete demand curves for relevant items can somehow be measured, the magnitude of the utility provided by each use remains unknown, because saleable products usually comprise much more than a single invention.[130] As Abramowitz (2003) emphasises, many of the patent prizes could work under appropriate conditions, but they must be carefully selected for the particular economic environments in which they are to operate.

4.2   Rewards for Copyright Subject Matter

Suggestions that it might be desirable to replace patents with publicly funded alternatives have been relatively persistent, even in democratic capitalist societies. In contrast, copyright sceptics have historically been less enthusiastic in recommending that copyright be replaced wholesale with a system of public funding.[131]

The likely explanation for this distinction is that the perceived losses to society as a result of patent monopolies were much higher than those caused by copyright monopolies. Two interconnected reasons are the functional nature of patentable inventions,[132] and the relative strength of patent rights.[133] Furthermore, the visible example of rewards for authors in communist states was intimately linked with a highly ob-

---

[130]*Cf.* (Croskey 1993, 639–40) (criticising Polanvyi's proposal on account of the informational difficulties it raises). But *cf.* (Shavell and van Ypersele 2001, 541–2) (expressing optimism that *ex post* sales information in the hands of a government could be as effective as the *ex ante* information possessed by inventors).

[131]There are of course exceptions; scholars of copyright have occasionally mentioned the possibility of publicly funded provision, even with approval; see, for example, (Hurt and Schuchman 1966, at 424, 432). But it is unsurprising that the idea was not taken very seriously before the widespread adoption of the Internet, because the case for rewards was, at that point, much weaker; see, for example, (Breyer 1970, note 104).

[132]The impact of functionality on debates about monopoly incentives and reward systems for particular kinds of information goods can be seen clearly with the inclusion of software, which is predominantly functional, in the copyright system. Note that the Free Software Foundation (Stallman 1985) were among the first to take up modern political opposition to copyright; their identification of technological autonomy as their *casus belli*; and their simultaneous argument that a "software tax" could be used to fund code production.

[133]Following the terminology proposed by Drahos (1996, Chapter 6), patent protection is "exclusive", prohibiting the re-implementation of an idea. Copyright, in contrast, is "preventative", and covers only a particular expression of the idea. Artists can usually build on the ideas of their predecessors, even when they cannot literally sample and remix them. The grant of patents poses a more insistent threat to sequential innovation.

jectionable system of censorship, on one hand, and state patronage of ideologically acceptable writers on the other.[134]

Having said this, publicly funded reward systems were actually developed and deployed for copyright works in the second half of the twentieth century. Rather than replacing copyright, they were conceived and created as adjuncts to it and applied to situations in which exclusive rights were impractical. The two kinds of systems are "public lending rights", which remunerate authors for the use of their works in libraries; and "private copying" levies — taxes paid to copyright owners according to the (possibly dubious) theory that they were harmed by the uptake of technologies such as audio and video cassette recorders, and deserved to be compensated. I will discuss each of those systems in turn. They are important precedents, now that digital technology has upturned the furniture of exclusive copyright.

4.2.1   Public Lending Rights

Public lending rights (PLR) systems are perhaps the best existing example of a copyright reward system. The fact that they are called "rights" is a little confusing; in fact, public lending rights are government schemes to pay writers for their public good authorship, without interfering with users' ability to access large bodies of works through free public libraries. Public lending rights were first created in Scandinavia in the 1940s and 50s,[135] and have since spread to most of the developed world.[136]

The nature and legal status of PLRs varies widely with jurisdiction;[137] a minimal degree of harmonisation occurred with the European Union rental directive.[138] Although in any coherent account public lending rights would be regarded as part of the copyright system, most nations define them as separate in order to avoid the national treatment obligations of copyright treaties[139] (Germany is an exception). In most countries the payments are made exclusively to authors, although Australia and the Nether-

---

[134]See (Levitsky 1964); (Loeber 1984).

[135]The first PLR system was introduced in Norway in 1947; see `http://www.plrinternational.com/established/plradministrators/norway.htm`. Sweden followed in 1954.

[136]See `http://www.plrinternational.com/plradministrators/plradministrators.htm`.

[137]See (von Lewinski 1992; IPA 1997).

[138]See Council Directive 92/100/EEC of 19 November 1992 on rental right and lending right and on certain rights related to copyright in the field of intellectual property; `http://ec.europa.eu/internal_market/copyright/docs/docs/1992-100_en.pdf`.

[139]Article 5(1) of the Berne Convention on Literary and Artistic Works, and Article 3 of TRIPS, implement the *national treatment* principle, which requires that any rights granted domestically by a member state to its own authors, must also be granted for works originating in other signatory states.

lands split the payments 50/50 with publishers, and Germany has a more complicated division.[140] Other aspects of variation include the coverage of works other than books (which is unusual but does occur); the kinds of libraries which are included; the sampling techniques used to determine borrowing frequency, and the extent to which they are open to foreign authors (with approaches ranging from isolationism[141], through reciprocity, to national treatment[142])

In countries where these schemes are well-funded, they form a significant portion of authors' (and also, where applicable, publishers) incomes.

4.2.2   Private copying schemes

Statutory private copying regimes exist in many jurisdictions, and primarily address the reproduction, in private, of music or video.[143] They involve the collection of a pool of money from taxes or levies on goods such as blank CDs, cassette recorders or digital media players.[144] These funds are distributed by collecting societies to copyright and neighbouring rights holders.[145]

Private copying levies first appeared in the wake of successful litigation for secondary copyright infringement by the German collecting society GEMA against a manufacturer of audio recording equipment.[146] In that sense the foundational theory of private copying systems is that copyright owners are "harmed" when the public has access to technologies than enable people to make copies of works, and because it

---

[140]Germany's split is administered by the collecting society VG Wort; it includes a share ratio of 63:27:10 between authors, publishers and a social fund; see `http://web.archive.org/web/20040629090634/http://www.plrinternational.com/plradministrators/germany.htm` for details.

[141]The U.K. pays German authors, but not those from any other states.

[142]Foreign authors whose works are widely lent in Germany can obtain direct payments from the collecting society VG Wort.

[143]For a general treatment, see, for example (Davies and Hung 1993).

[144]The exact combination of levies varies from system to system; see, for example (ACC 2001, pp 9–10) for a survey of revenue sources employed in different jurisdictions; (Copyright Board of Canada 2003) adjusting Canada's levy rates on digital devices.

[145]The details of the distribution system vary widely. There is generally an arbitrary split of royalties; in the case of music, it is between record labels, performers, and music publishers (who have contractual obligations to pay composers, typically between "very little" and 50%). The split of remuneration between different works is determined by the collecting society; traditionally, the measures employed have comprised observing radio airplay and sales (see, for example, `http://neil.eton.ca/copylevy.shtml#show_me_the_money`). These are unsatisfactory.

[146]See *GEMA v. Grundig* Decision of May 18, 1955, I ZR 8/54, 17 BGHZ 266; 1955 GRUR 492. See (Gaita and Christie 2004, Part II.B) for an English language analysis of the case.

is impractical to prevent the harm,[147] rights holders deserve compensation by government intervention. This is a decidedly user-unfriendly philosophy. Its echo can even be heard in MPAA President Jack Valenti's infamous comment, while trying to repeat GEMA's success in 1980s America: "I say to you that the VCR is to the American film producer and the American public as the Boston strangler is to the woman home alone."[148] It followed that the American public should have been taxed for their use of the VCR, and were it not for the film industry's defeat in *Sony v. Universal City Studios* (the "Betamax" case) and related legislative contests, they would have been.

The information feudalist pedigree of the private copying levy does not mean that it cannot be co-opted for more subtle purposes.[149] In contemporary disputes over them, many groups have now changed sides: major copyright holders are firmly in favour of abolishing levies, and some consumer groups trying to preserve them![150] Why, the reader might wonder, has such a dramatic realignment occurred? The answer, in a single acronym, is DRM.

So long as a work is controlled by DRM, unauthorised private copying is — almost by definition — impossible. If consumers have to pay a higher price to convince the DRM system to let them burn a CD or copy a song to their portable media player, they might start to ask why they also have to pay levies on the blank CDs and the media players.[151] This "pay twice" possibility leads to a number of uncomfortable scenarios for DRM enthusiasts: consumers insisting that they be allowed to circumvent to do the things they have been billed for, or even insisting that works be free of DRM altogether.

The prospect that has most troubled major rights holders is that private copying levies could lead to the legalisation of P2P file sharing. The possibility is real. In 2003, the Copyright Board of Canada, which administers the Canadian private copying levy, decided that downloading from P2P networks, for personal use and onto media for which levies had been paid, would be legal at least temporarily.[152] In 2004, a Canadian court ruling on the privacy of file sharers issued a *dictum* holding that uploading to P2P

---

[147]At least in pre-DRM Germany, it may have been unconstitutional to intrude on people's privacy sufficiently to prevent copying (Gaita and Christie 2004, Part II.A).

[148]The comment was made before a Subcommittee of the U.S. House of Representatives. See `http://cryptome.org/hrcw-hear.htm`.

[149]For a discussion of ambiguities in the philosophical basis of levies, see (Christie 2004).

[150]Other interest groups have not switched sides: artists and collecting societies retain a generally positive view of levies, and electronics firms remain opposed.

[151]See, for example (Peukert 2005, Part IV.D.2).

[152]See (Copyright Board of Canada 2003, pp. 20–1). In practice, because the levies have not been charged on hard disks or MP3 players, the set of circumstances under which listeners can take advantage of the rule is rather limited.

networks was also legal, provided that the uploader themselves had a licensed copy.[153]

A real legalisation of file sharing is not quite at hand. These rulings were only possible because Canada had not yet implemented the exclusive rights of "making available to the public", required by the WIPO Copyright treaty, and the *dictum* in *BMG v. Doe* was roundly criticised on appeal.[154] In addition to any philosophical or emotional preference for exclusive rights, there are several reasons why Canada's major record labels may continue to oppose a levy: a fear that it would not raise enough money (as of 2011, the levy raised $30 million per year,[155] which is 1-1.5 orders of magnitude too low for funding the entire industry); the fact that arbitrary portions of the royalties go to performers and composers.

Other jurisdictions have continued to flirt with using private copying levies to legalise P2P. In December 2005, the French National Assembly passed amended WCT/EUCD implementation legislation to create a "global license" that would achieve exactly that,[156] but organised opposition from major copyright holders ensured that the final outcome was very different.[157] In June 2006, Swedish political parties started talking about revisiting their WCT implementation along similar lines.[158] None have actually crossed the line of running these experiments. In 2010, the Brazilian government introduced a proposal for a "sharing license" alternative compensation system (Paranaguá 2010), though as of this writing it is unclear if it will proceed.

Aside from these specific P2P proposals,[159] the European Union — home to most of the world's private copying schemes — is still trying to decide what to do about them. Some groups are keen to seem them phased out;[160] others want to see them transformed.

[153]BMG Canada Inc. v. John Doe, 2004 FC 488 (Ottawa, Mar. 31 2004) http://www.fct-cf.gc.ca/bulletins/whatsnew/T-292-04.pdf 14–15 (dicta stating that, until Canada implements the WIPO Copyright Treaty, making files available on a P2P network would not attract secondary "authorisation" liability).

[154]See BMG Canada Inc. v. John Doe, 2005 FCA 103 (Ottawa, May 19 2005) http://decisions.fca-caf.gc.ca/fca/2005/2005fca193.shtml, paragraphs 46–54.

[155]http://www.savethelevy.ca/en

[156]See http://www.assemblee-nationale.fr/12/cra/2005-2006/109.asp

[157]See http://en.wikipedia.org/wiki/DADVSI#The_.22global_license.22, last accessed 16th of July 2007.

[158]http://evalu8.org/staticpage?page=review&siteid=10254

[159]Without legislative adjustments, it was expected that they could not apply to the semi-public acts of copying found on file sharing networks (Hugenholtz *et al.* 2003, p. 41). There have been isolated instances of judicial dissent on this point; see, for example (Oates 2006).

[160]See (Hugenholtz, Guibault, and van Geffen 2003, Chapter 6). (an independent report commissioned by the Business Software Alliance; Chapter 6 addresses how the EU might implement its stated goal of phasing out private copying levies in the era of DRM).

4.2.3   Going Further

The near-consensus in favour of exclusive rights for authors has begun to weaken with the advent of digital networks. The prospect of universal, scarcity-free access to much of humanity's knowledge and culture has inspired a growing number of claims that government funding might become desirable.

The development of this idea can be seen in explicit policy arguments that compulsory or blanket licenses should be extended to widespread private (including peer-to-peer) digital copying. Chronologically, the proposals have included Stallman's argument that the U.S. AHRA (Audio Home Recording Act) should have been used to completely liberalise end user copying (Stallman 1992a); Schulman's argument for compulsory licenses to escape the transaction costs of digital licensing (Schulman 1999, 628–30); Fisher's argument that taxation-based blanket licenses would be preferable to the *satus quo* (Fisher 2000a, § IV. 2); similar conclusions by Lunney (2001, 911–8) and Ku (2002, §VI). Gervais argued that a private copying regime might be unavoidable now that the social norm of free copying has taken hold (Gervais 2004). Libraries, and public lending rights, have also been an inspiration for some proposals (Rothman 1992; Foley 2001).

Not all of the arguments have worked from precedents in existing copyright systems. Calandrillo (1998) and Shavell & van Ypersele (2001, at 541–2) advocate government rewards, funded by income taxation, as replacements for both copyright and patents; The Free Software Foundation (Stallman 1985) advocated taxes on computers, allocated by both users and government, to fund software production.

More recently, authors have begun developing very detailed proposals for taxation-funded blanket licenses. My own work for this thesis has been a part of that line of research.[161] Broadly speaking, the schemes discussed in this literature follow the same pattern that I will set out at length in the next Chapter (5). My emphasis on technical architecture, and the possibility of letting users vote, is original. Comparative references and discussions of the contributions of these authors are included at relevant points in this thesis.

---

[161]See (Love 2002); Eckersley (2003, 2004b); (Baker 2003); (Liebowitz 2003a); (Netanel 2002; Netanel 2003); (Fisher 2003; Fisher 2004); (Litman 2004); (Farchy 2004); (Peukert 2005) (focusing on the feasibility of alternative compensation systems under Berne and TRIPS) (Oksanen and Välimäki 2005) (attempting to relate abstract theory about ACSes to the evolution of levy systems); (McDaniel 2007) (a detailed review of several proposals in the literature); (Aigrain 2008) (a very detailed argument for alternative compensation systems in the French cultural, political and legal setting in particular).

The next chapter will discuss how alternative compensation systems based on taxation could work in practice.

Chapter 5

Decentralised Compensation Systems: Designing "Virtual Markets"

Contents

In this chapter, I will introduce the notion of a "virtual market" — an alternative compensation system for digital media based on decentralised measures of the value of digital art, culture and entertainment. We will closely examine the design requirements, technical and economic details of this class of alternative compensation systems. Section 5.1 discusses the basic question of how the value of copyrighted digital material can be measured. Section 5.2 sets out the requirements that a successful implementation of a virtual market would have to satisfy, and discusses how that could be

achieved. Section 5.3 considers how alternative compensations might be best funded. Section 5.4 asks which markets might be made virtual, or in other words, where can we observe prevailing economic conditions that would make digital ACSes a sensible proposal?

5.1 Mercatus ex machina

Let us return now to the thought experiment of Alice and the problem of paying for her usage of copyrightable goods (see p. 90). I suggested earlier that in order to adequately remunerate authors and publishers for their work, a government could tax Alice, and then distribute that revenue to the producers of the works that she likes best.

The simplest way to split royalties in an alternative compensation system would be to count how often Alice and others like her download different works. The government could track downloads from different sources,[162] and distribute rewards in proportion to the popularity of content. Although simple, such a system suffers from the drawback that it cannot determine how much people actually like the different works they have accessed. For example, if Alice downloads two songs, listens to the first, and then deletes it, but listens to the second song every day, the two artists would receive the same reward. The measurement is also limited to works which are downloaded directly, excluding those that are distributed between users using recordable media or direct file transfers.[163]

A significant improvement on download counting could be had by instead measuring Alice's actual usage of different works. Her media player and and portable electronic possessions could remember how many times she plays each track in her collection.[164] Her web browser could record her surfing habits. Her e-book reader could count the time that she spends poring over novels and essays. Her operating system could track the programs she uses. Employing this information would increase both the fairness and the economic efficiency of an alternative compensation system.

---

[162]The central authority would not necessarily need to serve the downloads themselves — making measurability a requirement for the legality of reproduction would be sufficient to obtain an audit trail from most peer to peer networks, for example.

[163]It would be possible to use other measures that look through people's collections to measure the total "reach" of each file, although the infrastructure required to do this reliably would require much of the more complicated security infrastructure discussed in Section 5.2.2, at which point usage measurement would be just as easy but more informative.

[164]When I first wrote about this, the idea of collecting this data on large numbers of people was hypothetical, but software like iTunes and AudioScrobbler/Last.Fm's plugins has made it a reality.

The most obvious way to collect usage data, at least for music and video works, would be to count the number of times the work is played. But it turns out that that measure is subject to arbitrage by creating many small works — for instance, if artists can split a track on an album into multiple tracks without greatly reducing the amount the album is listened to, their expected returns would be greater. Similarly, there would be perverse incentives to split television episodes into smaller segments, to serialise books, etc.

Appropriate solutions to that problem would be to normalise the weighting of audiovisual works by their length, or to measure or infer the number of hours of attention that people spend on them.[165] Employing this information would increase both the fairness and the economic efficiency of an alternative compensation system.

Essentially, counting normalised usage or attention is a proxy measure for the social value of the public goods produced by different authors. Compensation systems that attempt to measure the value of those public goods can be termed "virtual markets", since they attempt to simulate one of the most important functions of marketplaces: linking the incentives for producers of a class of goods to the benefit that the recipients of those goods receive.

A few observations can be made about the usage metering approach. It is not perfect: there may be short books that we read only once, but which change our lives; a piece of software which simplifies a complicated task may receive less usage than one which performs it inefficiently, and so on. Secondly, the usage metering metric, more than the download metric, is subject to conscious manipulation. There is nothing to stop someone playing a tune on repeat with the speakers turned off, and in fact that can easily happen by accident with some very widely used media player software. Thirdly, the information reported to a central authority is very personal; although it could be anonymised by technical means many users might remain highly unimpressed by it.[166]

All of these points suggest that perhaps it may be better to consider a system which involves elements of explicit voting and voluntary participation. By giving Alice a certain number of votes (say 100 per month, for the sake of example), she could express

---

[165]Measuring hours of attention is difficult, because computers can't usually tell what their users are doing. The reading of electronic books is an important exception, because the pattern of navigation that corresponds to reading is distinctive, if not unmistakable.

[166]Even a download measurement system would obtain highly personal information, but the issue here is as much one of perceptions as one of genuine compromise of privacy. This is not to say that the privacy consideration is unimportant, but rather that it may already be absent and that it can be designed in or out of almost any process.

her preferences in a more accurate and a more consensual fashion. If she has read a novel which is particularly disappointing, she might not reward it at all, or might give it only a symbolic vote or two. On the other hand, when a novel is extraordinary, she might give it all 100 votes, or an ongoing reward each time she re-reads it.[167] The disadvantage of regular voting is the effort it involves — continually rating numerous snippets of authorship and artistry could be a chore which many people were inclined to avoid.

In order to combine the benefits of voluntary voting with the efficiency of direct usage metering, a hybrid model could be employed. In that case, the same infrastructure would be employed to record Alice's usage habits, but rather than reporting this information directly, Alice would have control over it. She could use it to determine her vote, or not, as she wished.

To illustrate this idea, imagine that Alice is listening to some music, or perhaps intending to download a few new songs for her collection. Because she hasn't voted for the past month, her download client pops up with a notice mentioning that she should do so. Alice now has three choices. She could refuse to vote completely (in which case, her downloads alone would be counted).[168] She could spend the time to vote manually and from scratch, carefully considering which works had been of the most value to her recently. Finally, she could *allow her computer to suggest a vote*.

In this last case, the usage measurement infrastructure would kick in and provide its observations on her recent preferences. At that point, Alice would only have to click "yes" (or "use my iTunes records as a vote") to reward artists based on the time she's spent on them. Of course, she does not have to accept this suggested vote precisely; the rating for each work might be an adjustable quantity, allowing her to complete a customised ballot for remuneration. This combination of usage measurement and user control is the most desirable schema for a virtual market reward system.

I have termed this mechanism a "virtual market" not because it is a "market" which happens to operate on the Internet. Instead, it is *virtual* in stronger senses of the word — a sort of "market through the looking-glass".[169] Despite the involvement of public

---

[167]Votes might or might not be normalised, so that if Alice only cast 12 votes in a particular month, they would be re-weighted to be worth 8.5 votes each.

[168]Downloads might be assigned less weight than explicit votes by users, to reflect their lower degree of interest in information production, and to provide an incentive for voting. The economic effects of this policy are considered in Section 7.5.4.

[169]The *Shorter Oxford English Dictionary* provides the following definitions (amongst others) for the word 'virtual': "...so in essence or effect although not formally or actually", and the "apparent... image

funding, the rewards and incentives which flow from virtual markets are very similar to those which would result from the exchange of goods and currency in a marketplace, although these exchanges do not occur directly.[170] And while there is a centralised governmental authority which collects taxation and distributes royalties, the determination of which information goods should be produced remains a decentralised and emergent result of the privately held knowledge and preferences spread throughout society. At least, it works that way if the system runs according to plan.

5.2   Technical Design Considerations

Having identified an idealised design for the user-interface behaviour of a virtual market, the next question is whether the proposal is implementable, and what hardware, software, and human incentives that would require.

In terms of software componentry, a download-based virtual market would be the simplest architecture to deploy. Although it lacks some of the desirable economic properties of more sophisticated models,[171] slightly less sophisticated infrastructure is required to make it secure.

For either voting- or usage-based virtual markets, the system must at some point record the patterns of access to and usage of various files embodying copyright works. Under cooperative circumstances, where it is assumed that all actors play their role in the architecture faithfully, this would be a fairly straightforward task; recognising files reliably is the only non-trivial step, and it is solvable (see Section 5.2.1).

Under more realistic non-cooperative conditions, security must be one of, if not the, most important design objective for an alternative compensation system. A virtual market is a mechanism for distributing cash, and as such, there are firm incentives for people to find ways to 'game the system'. The music industry in particular already have their own traditions of audit-manipulation.[172] These practices need to be anticipated

---

resulting from the effect of reflection... upon rays of light".

[170]This type of proposal may therefore be able to evade the common objection (see note **??** above) to publicly funded authorship on the grounds that it is not "market based", and so hands control of publishing to the state.

[171]On this point see Chapter 7.

[172]Such as the role of "independent promoters" in the U.S. music industry (Boehlert 2001), and the conventions regarding private copying remuneration allocation in Europe. Anecdotally, a law professor who had also worked as a musician in Italy recounted how the accounting systems used by collecting societies there were rigged to benefit senior musicians. When the EU considered reforming its member states' patchwork quilt of levy systems, improved transparency was high on the priority list (European

and prevented.

There are essentially two aspects to design for security: ensuring that the human actors who have roles to play within the virtual market do not have substantial incentives to diverge from those roles, and ensuring that the various hardware and software systems that they use to communicate with one another operate accurately and reliably. These issues are discussed, respectively, in Sections 5.2.2 and 5.2.3

5.2.1   Recognising files

How can software on Alice's computer tell what a particular file is, in order for it to inform the virtual market that the work deserves compensation? If works in digital form were completely static files, this would be an easy problem — one could just rely on a standardised tag in the file or a canonical description like a URL.[173]

As it turns out, things are more difficult because metadata, and even data itself, can be and is changed by users on a semi-regular basis.[174] Some files (such as audio files ripped from CDs) may never have been tagged to begin with, or may have had incorrect metadata from the outset. The unreliability of metadata on its own is a slight nuisance, but it makes the variability of actual data, caused especially by the use of lossy compression algorithms for media such as music, film, and images, extremely inconvenient. The software must recognise many versions of each file, some of which may be mis-labelled.

In these hard cases where lossy compression is employed, the difficulty of the measurement and accounting task depends on how precise the results need to be. Some deployed systems for tracking musical tastes operate by simply trusting the metadata in files, adopting a "near enough is good enough" approach and collecting impressively detailed statistics on music listening habits.[175] But because there is money at stake, a virtual market should be designed to resist both honest and deliberate mislabelling.[176]

Commission 2006).

[173] Such as the Digital Object Identifier (Rosenblatt 1997), which was originally designed for electronic publishing but has been adapted for use with other media.

[174] This practice of metadata editing is a surprisingly complicated phenomenon and is in some ways reminiscent of arranging one's record, book or video collection. But the extent to which it occurs in any given domain is heavily dependent on technical standards. MP3s have unpredictable metadata because the CDs they are ripped from do not contain track labels. HTML files were rarely edited by anyone other than their original posters, until wikis started to become widespread. It is still fairly uncommon for users to edit PDF files.

[175] As mentioned above, both Last.Fm and Apple's iTunes system do this.

[176] For the purpose of illustration, three (probably honestly) mislabelled files the author has observed

The simplest solution to mutable data and unreliable metadata would be to insist on the use of "official" versions of files. The use of versions of files which lacked an accompanying certificate (or equivalently, a secure hash stored in a central database) from the virtual market administration would not lead to compensation. The serious drawback with that approach is that entire existing libraries of digital media would be invisible, introducing a potentially serious bias to the survey. Attackers could also mount a kind of financial denial of service by circulating unofficial versions of files, thereby reducing certain artists their fair share of remuneration.

We would like to do better. To state the problem precisely, there is a program on a PC or media device, which is assumed to behave honestly. When it plays a media file, its task is to recognise the file in some authoritative way (either with accurate metadata or with an abstract identification number which is unique to each remunerable work). In the easiest version of the problem, which I will term 'interactive examination', the software has a communications channel to a central database managed by the virtual market. Transmitting the entire media file to the database is not generally permitted, because the bandwidth requirements would be too large and because some media files are private objects. A harder version of the problem would allow 'offline examination', a process in which there are two stages: the media player examines the file in the first step, and records only a small amount of identifying information, while in the second stage the virtual market database receives that information and recognises the file from it.

There are two very useful computational tools which would be necessary for solving these problems: metadata independent hashes and fingerprints. Both of these are variants on the general concept of a hash function.

Metadata-independent hashes are a straightforward adaptation of the secure hash functions used in modification detection codes or MDCs.[177] They can be implemented for any media file format like MP3, Vorbis or Quicktime AVI; tags and any other metadata in the file itself are just discarded and then the hash is evaluated. The 'audio SHA1' hash is an example of this concept.[178] These metadata independent hashes

---

while using last.fm are confusing Jefferson Airplane's *Somebody to Love* for a Janis Joplin song, Cat Power's cover of *Wild is the Wind* for the Nina Simone original, and Screamin' Jay Hawkins *I Put a Spell On You* for a Credence Clearwater Revival song. An example of deliberate mislabelling was the passing off of Mandalay's album *Pearl* as a Portishead album, presumably to inflate its sales at record stores of ill repute.

[177]See, e.g.,, (Menezes, van Oorschot, and Vanstone 1996, § 9.2); see also the bitprint/TigerTree hashes that Gnutella combines with SHA-1 (`http://bitzi.com/developer/bitprint`).

[178]See Bitzi bitcollider source code, `lib/mp3.c`, CVS version 1.5 committed 31st July 2001, `http:`

(which I will term 'MIH's) allow trivial recognition of files with the same compression encoding but changed tags. A stronger family of MIH algorithms could operate by decoding files to their raw form, and calculating a hash of the decoded data. Those MIHs are slightly more accurate, recognising certain kinds of redundancy in media files, but they are also significantly more costly to evaluate because of the decoding overhead and much larger input to the hash function. The cost/benefit tradeoff is likely negative for ACS applications.

Fingerprints are a kind of 'fuzzy', imperfect and potentially insecure hash which should be robust over minor changes in the content of a media file, capturing essential subjective properties (Pye 2000; Cano, Batlle, Kalker, and Haitsma 2002; Gruhne 2003). They should survive changes in lossy compression (at least provided the data rate is not extremely low). Some fingerprints are designed to recognise content despite cropping, superposition, sampling, or transmission through distorting channels such as telephone lines, although they are of less relevance to the problem at hand.

The development of fingerprinting algorithms is an area of ongoing research. There are tradeoffs to be made between the size of the fingerprint, its accuracy in terms of both false-positive and false-negative matches, the computational resources required (especially at the database end), and robustness over re-compression. Databases have been deployed to perform song recognition using fingerprints, although at this point there are no systems which are both freely usable and reliably accurate.[179] There are, however, some algorithms in the literature which claim a high degree of robustness (Haitsma and Kalker 2002). For present purposes, it can be assumed that fingerprints will have a high recognition rate but conservatively it must also be assumed that there will be a small proportion of errors, so that fingerprints alone do not solve our problem.

I propose a pragmatic algorithm to solve the offline and interactive identification problems which augments fingerprints with MIHs and processes to detect and resolve inconsistent identifications[180]:

- Copyright owners initially register works for which they wish to be remunerated

---

//cvs.sf.net/viewcvs.py/bitcollider/bitcollider/lib/mp3.c?rev=1.5

[179]See, e.g.,www.musicbrainz.com, which uses the TRM fingerprinting algorithm (which isn't reliably accurate). Personal communication with researchers working in the area indicates that the barriers to better systems relate to business strategy, patent licensing and other social (as opposed to technical) obstacles.

[180]One system which is a little similar is the audio recognition process patented by Relatable (Ward and Richards 2004). That scheme is oriented towards combined media objects like CDs; its key weakness is that it fails to realise the importance of using both hash- and fingerprint-based recognition at the same time

with the virtual market. The ACS administration then builds a database of these works, their appropriate metadata, and their fingerprints. The process of file recognition builds a table mapping MIHs to these registered works.

- During an examination stage (which may be offline), clients record the metadata, a fingerprint and a MIH for each file seen.

- During the recognition stage, the virtual market first checks if the MIH is known. If so, it counts a usage or vote for that file.[181] If the MIH is unknown, the database is queried using the fingerprint; if it returns only a small number of hits with similar fingerprints, and one or all of these match the file's metadata, then that should be recorded as a tentative recognition.[182] The file's MIH is then (tentatively) associated with the matching copyright work.

- If a file's fingerprint does not match but the metadata does match a known work, that fingerprint and MIH are both marked as tentatively associated with the work.

- If a situation arises in which several different and competing copyright works wind up being tentatively associated with the same MIH, the virtual market will need to resolve the issue authoritatively. A larger and more informative fingerprint could be requested from clients that are performing interactive identification of that file. If that was insufficient, the virtual market could actually obtain a copy to compare with the master registrations; users could be asked to provide it voluntarily ("click 'yes' to upload a copy of this file so that it can be identified").

- Files which remain completely unrecognised are recorded, so that if at some point a copyright work is registered that matches them, remuneration can be assessed.

- A (high) limit might be placed on the number of unrecognised files allowed from each client, in order to prevent deliberate or bug-induced denial of service attacks on the central database.

An architecture of this sort would allow not only for a high recognition rate but also for a dynamic correction process in which commonly unrecognised files are prioritised

---

[181]If the file's metadata does not match the database's known identification, that event might be recorded if there is a chance the database is wrong.

[182]Several fingerprints may be returned and have very similar but not identical metadata because (for example) they might correspond to different recordings of the same song by the same artist.

for identification. It also foils the broad category of attacks on the file recognition system that involve attempts to circulate mislabelled files in order to redirect a portion of the income due to the creator of a particular work. Provided that a non-negligible number of users continued to label their files correctly, such attacks would lead to conflicting tentative identifications that could then be easily resolved.

5.2.2 Secure Measurement and Voting

Any virtual market infrastructure would need to not only foil the entrepreneurial musician who sets out to download fifty thousand copies of his own song — but also to be robust in the face of software-based attack on client systems. If this were not the case, then worms, viruses, trojan horses and direct computer security breaches that affect Alice's computer, could all grant the perpetrators control over hard financial resources. Although it might be challenging for attackers to collect the proceeds of this sort of ACS manipulation anonymously, the risk must be thoroughly mitigated if the system is to be dependable. This is particularly the case because, unlike many other forms of electronic financial fraud, the victims could have no idea that they are in fact victims. When an unauthorised transaction is charged to a credit card, the owner will usually notice — but when one artist's payments from a collecting society are too low because someone else has fraudulently masqueraded as a popular artist, the attack may be much harder to detect.

In some respects, the goal of designing an attack-resistant virtual market is quite similar to that of enabling secure remote Internet voting, which is a very challenging problem (Rubin 2002), although not necessarily an unsolvable one. To solve the remote Internet voting problem completely, one needs an algorithm that conducts an election in which participants are (1) authenticated; (2) permitted to vote at most once; (3) guaranteed that their votes are confidential; (4) able to confirm that their vote was counted correctly and (5) able to confirm that the tallied results are correct. In many remote voting situations, we must add (6) that even if attackers can run arbitrary, malicious code on the computers that the voters are using, the election tallies will still be correct, or the voter will be able to detect tampering.

Many technological solutions are applicable to both, although the risk mitigation considerations are subtly different. One factor that makes electoral systems slightly harder to defend is that typically a small fraction of the votes within them are disproportionately valuable. Another is that sampling methodologies may be politically

impractical in the case of electoral systems, where there is a compelling notion that every last vote must be counted.

If each voter controls a computer, there are cryptographic methods which can, when deployed appropriately, guarantee that their votes or usage measurements have been counted correctly while preserving the confidentiality of each vote, giving us properties (1)–(5).[183]. But the more daunting part of the problem is ensuring that people do in fact control their own computers, because such control can be subverted unless extensive steps are taken to make it dependable.

There is a recent branch of the cryptographic literature which attempts to verify the results of elections conducted with potentially malicious hardware (property 6) by using clever voter-verification algorithms.[184] However, the requirements that these schemes impose upon voters — such as checking mathematical properties of cryptographic receipt numbers — are daunting even when voting is an activity performed with special instructions, on special occasions, in special locations. They could not be implemented for regular, semi-automated voting on one's music or literary preferences. Instead, I will discuss other methods that might be used to work towards property (6), or at least to limit the errors that the attacker can introduce.

Voting on PCs and other general-purpose devices

Performing the usage measurement or voting tasks of a virtual market on personal computers, smartphones, and other general purpose, multi-use devices would be a very desirable outcome, because these are the predominant species in the planet's IT infrastructure. A large fraction of our consumption of digital copyright works occurs on or (in the case of media players such as iPods) is connected to general-purpose devices. The challenge of running a virtual market securely on PCs and phones is greater than it would be on simpler platforms because of their complexity and variability. Complexity contributes to insecurity because it ensures a more plentiful supply of flaws to exploit;[185] variability makes it harder to test for and fix those flaws. If a task can be performed securely in such a rabbit warren, it can be performed securely anywhere.

---

[183]See, for instance (Chaum 1981; Jakobsson, Juels, and Rivest 2002)

[184]See, for instance, (Chaum, Ryan, and Schneider 2005; Adida and Rivest 2006; Ryan and Schneider 2006; Benaloh, Moran, Naish, Ramchen, and Teague 2009).

[185]Apple's iOS devices are significantly more homogeneous and centrally-controlled than other general purpose computers, but even on those devices the wide variety of available apps has essentially the same consequences.

Unfortunately, it cannot. PCs in general are prone to so many security problems that significant amounts of professional labour is required to keep them from being overrun by worms and hackers.[186] If the assailants are resourceful and determined, only extremely careful operations are capable of defending against them. The PCs that most people have in their homes do not receive those levels of attention and are, as a result, woefully insecure. Offering attackers a way to make money from them would be very poor planning.

The security situation on phones is differently but comparably bad. Although smartphone operating systems have somewhat stricter security architectures than most desktop OSes, they suffer from other countervailing problems, including network carriers that block or delay the distribution of security updates, and the use of proprietary and poorly audited baseband processors.

General purpose devices, then, need to be augmented in some way in order to make them a suitable platform for virtual market usage metering or voting. I will now turn to a few candidates for such augmentation.

One of the most pragmatic augmentations for voting on PCs is to use different software. A bootable CD-ROM or other read-only media can supply a small, customised operating system just for voting.[187] There are difficulties: the need to ensure that machines cannot be subverted to run the voting operating system in a virtual machine while tampering with its operation,[188] and the inconvenience of establishing a network connection using a one-use operating system. Even assuming these could be solved, security-by-custom-OS would not be appropriate for an alternative compensation system since the elections are supposed to run frequently (prohibiting the inconvenience of a reboot for users). Also, the usage data the votes are supposed to be constructed from necessarily has to reside on the user's usual, insecure, operating system, because it is generated by day-to-day software usage. Other security mechanisms need to be found.

---

[186]Some reports have estimated that as many as 89% of PCs are infected with malware (Turner 2006). Whether or not such extreme numbers are credible, it is clear that the malware problem on PCs is severe (Jaatun, Jensen, Vegge, Halvorsen, and Nergard 2009; Caballero, Grier, Kriebich, and Paxson 2011), and applications must be engineered on the assumption that a large percentage of their client PCs could be compromised.

[187]The idea is patented (Babbitt, Roberts, and McClure 2006).

[188]One class of attacks, for instance, involves re-writing PCs' BIOSes so that they never truly boot from the voting CD, but pretends to do so while launching the virtual machine.

Separate voting hardware

There are at least two deep causes for the insecurity of PC software. One problem is the complexity of the system. The almost limitless array of applications, libraries, incautious users, and the interactions between them, ensure that there will all but always be ways for motivated attackers to remotely execute arbitrary code on PCs, and escalate privileges once they have done so. The other fundamental problem is the extreme difficulty of auditing a PC. There is in general no way of telling when a PC has been subverted by malicious code. The most straightforward tests (such as searching for incriminating code and data on the machine) are foiled by rootkits.[189] Rootkits may themselves be detectable, but this simply leaves attackers and defenders in an arms race. With the advent of widespread polymorphic code obfuscation, commercial antivirus vendors are struggling to win even a single round of that contest.[190]

One way to avoid these two impasses is to attach additional, separate, simple hardware to the machine. Because such systems could be much less complicated and much more uniform than PCs or phones, and their states correspondingly predictable, security auditing could be performed much more thoroughly. It might even be possible to employ automated verification for this purpose. By necessity, any such hardware has to be cheap.

A simple device[191] to perform secure voting might comprise:

- A microcontroller with an embedded private key,[192] to create digital signatures.

- A symmetric cipher implementation to provide a secure communications chan-

---

[189]A *rootkit* is a piece of software which gives its installer the ability to hide files and running programs from standard operating system functions (and thereby, from any programs that do not re-implement OS functions at a low enough level to be unaffected).

[190]See (Gutmann 2007) for a catalogue of the techniques that malware authors are using to stay ahead of antivirus software.

[191]One might be tempted to call it a "dongle"; it does have some properties in common with the gadgets commonly employed since the 1980s as software copy protection devices, although its purpose is, in a sense, the exact opposite, and it does not suffer from the same classes of categorical insecurity. Perhaps "anti-dongle" would be more accurate.

[192]*Private keys* form a part of *asymmetric cryptosystems*, which can provide both secure digital signatures and *public key encryption* (secure message "envelopes"). In this case, the private key is a unique secret stored in each device, while the public key is kept on record with the virtual market administration. Anyone possessing the public key can confirm that a message signature was produced by the (secret) private key (possession of the public key also allows the creation of messages which only the private key can decipher). Suitable algorithms might include a standard digital signature scheme or a special signature scheme with additional privacy-preserving properties (Chaum 1984; Stadler, Piveteau, and Camenisch 1995).

nel[193]

- A small LCD (liquid crystal display) to show the name of a creative work, and its author

- A "confirm" button

- A connection to a PC (this could be any standard data connection, such as a serial or USB (universal serial bus) port)

Whenever Alice wanted to download or vote for material in the virtual market, she would need to confirm that the transaction presented on the LCD was correct.[194] The device could then sign the details, creating an unforgeable receipt to be passed into the compensation system.[195].

When devices of this type are produced and distributed in volume, their cost can be fairly low — perhaps in the order of ten dollars. Some users might be willing to pay for extra features, such as wireless networking, or a more sophisticated user interface for adjusting their votes. These could be added to more expensive voting devices without compromising the security of the system. Though manageable, these cost overheads are substantial. ACS designs that avoid them may be preferred.

Costs aside, separate hardware would be an elegant solution for a virtual market based on download counts. It is less elegant as a means to report hybrid voting data, both because the user interface on a cheap device could be awkward for voting, and because securing the data streams into the device would be complicated. One might trust iPods and e-book readers to be fairly secure, and then equip them with cryptographic code so that voting gadgets could verify that usage data had come from them and not been tampered with by malicious code on a PC in between. The problem is

---

[193] A *symmetric cipher* allows two parties who share a common secret "key" to send messages which only the other can read. If any volume of data needs to be exchanged, symmetric ciphers have the advantage that they are much faster than public key methods (asymmetric cryptography is still used to exchange the symmetric key). A good choice might be AES/Rijandel (see `http://csrc.nist.gov/CryptoToolkit/aes/` (2003)), since it is efficient when used in minimalist hardware.

[194] The device thus defends against the class of attacks discussed in (Gobioff, Smith, Tygar, and Yee 1996).

[195] Astute readers may observe that, although secure authentication hardware can guarantee that Alice must have approved each important transaction, an attacker with control of Alice's PC might subtly alter the information provided for approval in the first place. Such mischief could be detected relatively quickly, either by vigilant users, or through the employment of networks of *honeypots* — computers which pretend to be participating in the virtual market, but which actually serve to identify assailants and analyse their behaviour; see Section 5.2.2 below.

that a protocol of that sort would either require fairly awkward key management in-frastructure,[196] or would face the risk of vote manipulation by replay attacks, where a legitimate report from one device was propagated by an attacker and fed into many users' voting software.[197] This picture is sufficiently tangled that I would not recom-mend the use of separate hardware beyond download-counting applications.

Voting with trusted systems

Another way to enable secure measurement and voting on a PC is to rely on built-in, rather than bolted-on, assistance from hardware.[198] As it happens, trusted computing architectures have been designed to provide precisely that kind of built-in facilitation for DRM.[199] The deepest difference between the two applications is that under DRM, a trusted system works against its owner's interests and may as a result need to be physically tamper resistant (see Chapter 9), while no such complication arises for vot-ing applications, which of course include virtual markets. It is somewhat ironic that systems conceived for the purposes of enforcing exclusion would in fact render easier assistance to the alternatives.

Trusted computing platforms could be used to support virtual market data collec-tion to different extents. In the simplest case, which is confirming that a download

---

[196]For example, requiring users to type an ID number from the back of their media players into their voting gadget.

[197]According to Syverson's taxonomy, this is a classic replay deflection attack (Syverson 1994).

[198]The potential application of trusted computing to electronic voting has been recognised, for exam-ple, by Rubin (2001) although interestingly it was omitted in a later version of the article (Rubin 2002) where a simpler line was taken: "The technology does not exist to enable remote electronic voting in public elections." Both claims are accurate, because presently deployed trusted computing hardware is insufficient for remote electronic voting.

[199]The Trusted Computing Platform Alliance has attempted to play down the significance of DRM as a motivation for trusted computing. For example, their 2005 FAQ contains the question and answer:

> "Was TCG formed to specify Digital Rights Management (DRM) technologies?" "TCG specifications do not provide all the necessary technical elements required for DRM. It is conceivable that developers could build their own DRM solutions that would operate on systems with Trusted Platform Modules, but TCG specifications alone are not DRM solutions."

this neglecting to mention that trusted systems make possible much stronger DRM than would other-wise be possible. Their technical documents also studiously avoid any direct mention of DRM, limiting explanations to oblique implication (such as the ability to ensure that users cannot run "improperly configured" software that does not report its status as a protocol requires (TCG 2005, p. 10)) This cir-cumspection is probably a public relations strategy to manage the very hostile responses which DRM provokes in parts of the technical community. The fact remains that the idea was conceived for just that purpose (Anderson 2003).

has been requested by a human being, the task of the trusted system is essentially to simulate the kind of hardware discussed in the previous section. The trusted subsystem must establish a confirmed link between an I/O device, such as a keyboard or mouse, and the file to be downloaded. The burden of that task is largely in making the user interface trustworthy — when a user clicks on 'yes', how can the TPM be certain that the question asked above the yes button is the right one? The difficulty of achieving this trust is a function of the number of software layers involved; doing it with a fully-blown kernel, windowing system, GUI toolkits and standard libraries is ambitious, because an adversary would only need to subvert one of these components. An easier alternative might be to allow a simpler and more easily verifiable program to step in and replace some or all of these layers for the step of vote confirmation.

A more complicated use of trusted computing for ACS usage measurement (rather than just download counting) would involve threading the chain of trust from the I/O hardware to applications that are used on a continuous basis, such as media players and web browsers. This connection would be necessary to allow the collection of reliable usage data, either for its own sake or for use as a 'default vote'. In such cases, it is no longer possible to use a simple user interface that takes a short cut through the normal mesh of libraries; browsers and media players need to exercise that intermediate infrastructure thoroughly. Security requires that all of these components be verified.

For this daunting task, a trusted system is only of partial assistance. It can provide hashes to identify all of the layered code at work in a particular application, and employ other 'integrity measurements' to ensure that it is configured to behave as expected.[200] But even if the code and its state are known, there is no guarantee that all of these sophisticated and rapidly evolving components can be relied upon.[201] If an exploit is discovered in one of the relevant libraries, trusted computing can ensure that no votes are accepted from systems on which the vulnerability has not been patched. But the ACS administration must find some other means to detect if some attacker, somewhere, has found and begun to exploit such a vulnerability. Fortunately, there are ways to perform such reconnaissance.

---

[200]Integrity measures are small snippets of code which can examine arbitrary aspects of a process or system's state or configuration in order to confirm that in this respect it is behaving as expected (TCG 2005).

[201]The set of integrity measurements applied to a particular program may in some instances be inadequate to detect all attacks but as they become more numerous and though-out, they certainly make an assailant's task more difficult.

Honeypots in architectures for collecting remuneration data

Honeypots are computer systems which feign some form of vulnerability in order to lure invasion by hackers and 'malware'.[202] When these assailants have occupied the system, their behaviour can be carefully monitored and studied.

In constructing an alternative compensation system, honeypots could be valuable both for identifying attacks 'in the wild' and to enable some limited statistical correction if remuneration payments had already been affected. The former use is a perfect counterpart to the trusted systems architecture discussed above; it can alert the community when the presently "trusted" versions of the media tracking code are in fact not trustworthy. The latter use provides a stop-gap in cases where download and usage data has been seriously compromised. In addition, honeypots could be useful for actually catching hackers who fail to cover their tracks completely.

The natural honeypot to employ would be a PC on which a known selection of copyright works was downloaded and played/watched/read.[203] The "remuneration votes" effectively cast by compromised honeypot systems could then be de-anonymised; if they vary from the known inputs, manipulation (or bugs) can be inferred. Examination of the memory and disks of those machines would then reveal a great deal about both the methods and consequences of that manipulation, and enable the development of effective countermeasures.

As one may come to expect with computer security, there are some subtle problems to be overcome. In this case there two key problems to solve:

1. ensure that assailants cannot distinguish honeypots from other systems; and

2. ensure that assailants cannot find ways to select systems for attack that are very unlikely to be honeypots.

The first problem is inconvenient, but as we shall see there may be some elegant solutions. One danger to note is the risk of over-reaching if the honeypot designer is ambitious and expects their traps to perform sophisticated spying functions. Inevitably, the features of honeypots that record and transmit surveillance information can also be

---

[202]For an overview of the concept and some implementations, see (Spitzner 2003).

[203]In many cases, honeypots *simulate* a virtual vulnerable machine; see for example Seifreid (2002), Jiang *et al.* (2004), or `http://project.honeynet.org/papers/honeynet/`. For the purposes of protecting and alternative compensation system this layer of simulation is optional, and it is in fact likely to increase the risk that the trap will be detected (Corey 2004, § 4).

used to detect them; attacks against published honeypot code have demonstrated this very clearly,[204] and it seems that such systems are doomed to be eaten in the game of cat-and-mouse unless they take the risk of relying on security through obscurity.

A comparatively light-handed approach to design is rather more promising. If the honeypot is barely different from an ordinary system, it is going to be difficult to spot. In a good design, the only thing that would distinguish the trap is a known selection of media files, and the input of that selection must lie outside the dimensions of the system (most especially the software) that a hacker can feasibly explore.[205] Extensive steps must be taken to ensure that the machines involved cannot be systematically traced to the organisation that runs them. Details of the observations made by the honeypots would be best left unpublished.[206] Even if these precautions are taken, there are detection issues to consider.

One interesting problem arises from the potentially blurred distinction between ACS-manipulating code and other malicious code. It is not possible to build a honeypot that works reliably for ACS attacks without also having it badly infested by other malware. For example, many of the PCs that are infected by worms are used as 'zombies' to send spam.[207] In most cases, one would not want honeypots to contribute to the spam problem by allowing spurious outgoing SMTP (email) transactions — but then a smart ACS manipulator might first test that a machine can send spam before changing its behaviour in the compensation system; a machine from which spam is blocked is more likely to be a honeypot. There could be many similar triggers. Can online banking passwords be sniffed from a machine? Will it launch denial-of-service attacks or exploit code against other machines? Against these kinds of tests, honeypots would need to behave as badly as the general population of unmaintained and insecure Internet-connected PCs. This kind of undercover work is far from desirable, but while the vulnerable PC population remains large, the honeypots' contribution to Internet delinquency would at least remain marginal.[208]

---

[204]See (Corey 2004; Holz and Raynal 2005; Dornseif *et al.* 2005).

[205]This could be achieved with a human who actually selects playlists and so forth by hand; or with a piece of hardware that replaces the honeypot's keyboard and mouse.

[206]See (Bethencourt *et al.* 2005; Shinoda and Itoh 2005) for attacks that reveal the locations of sensors based on the published reports of the sensor networks to which they belong.

[207]See, for example, `http://www.theregister.co.uk/2003/07/17/trojan_turns_victims_into_ddos/`.

[208]One difficulty here is that the administrators of the machines would be knowledgeable about the abuse of their systems, and might thereby face greater negligence liability for it than the average unwitting and unfirewalled broadband user. Of course, this would require the potential litigant to know that

The second requirement for honeypot effectiveness that I enumerated above is that there must be no way for hackers to pick on a class of computers that are very unlikely to be honeypots. This is harder than keeping particular honeypots secret. To see why it is so difficult, imagine a clever ACS manipulator who sets up a website which offers some alluring download to people who fill out a form;[209] those who do are then fed a trojan horse in response.[210] A passive honeypot could never catch this attack. Even a honeypot which was programmed to download and run random executables off the web (!) would not catch such a cautious adversary. For this reason, at least some of the honeypots would require human pilots.

To make honeypots most cheaply and thoroughly representative of PCs in general, it might be easiest to make some proportion of ordinary PCs and PC users into honeypots. That way, any attacks that affect some users (like the trojan website) will affect some honeypots. This could be done by adding just enough hardware to these PCs to ensure that the users' votes were independently measurable. A good example of this tactic for music is to install a sound card that records fingerprints of the files it plays. Such a device could be made safe from detection and would not be confused even if the speakers were off or the volume turned down. Any difference between the ACS data reported by the computer and the tracks identified by the sound card would indicate a compromise of the ACS system on that PC.

Sound card auditing would be effective for a usage-measurement virtual market, but not for one based on votes — there is no point in verifying the sources for the suggested votes if the votes themselves could be altered. Aaron Swartz (personal communication) suggested a way around this limitation, which is to place the auditing hardware on the monitor (and keyboard and mouse) cables rather than the sound card. The purpose of this monitor monitoring hardware is to ensure that the screen was displaying a correct voting screen before and while the user entered their vote. Different implementations are possible, although a concrete example would be to require the voting application to fill the entire screen, and to calculate a secure hash of every frame displayed during the voting process.[211] If a vote is recorded without a sequence of

---

the machine contributing to their inconvenience was in fact a honeypot.

[209]This would be a kind of malicious application of the "Completely Automated Public Turing test to tell Computers and Humans Apart" (CAPTCHA); see `http://en.wikipedia.org/wiki/CAPTCHA`.

[210]A *trojan horse* is some unexpected code inside another program. The attacks that could be used are not limited to trojans. Any active attack could be launched against targets that had been selected using the web page.

[211]If the computational cost of calculating secure hashes of high resolution screens is too high for

hashes that correlates with the user interface (keyboard and mouse) input that occurred while it was being entered, then the honeypot may have caught something and should be investigated further.

Several conclusions are at hand. Honeypots can certainly provide a valuable second-layer of defence for virtual market data collection systems that rely primarily on some other measures to keep costs-of-attack high. In order to make them dependable, they need to be included in real people's computers and adopt a minimally intrusive character. Music-auditing sound cards work for ACSes based on usage data. Vote-auditing keyboard, mouse and video cables work for ACSes based on voting.

Sampling vs. Complete Coverage

Another security option available to virtual market designers is to move from a complete coverage data collection model to a 'sampling' method. Instead of counting each and every Alice, a few are selected as representative. Fisher proposed data collection for an alternative compensation using a "Nielsen model" similar to the one used to measure television viewership (Fisher 2004, p. 226), and there is much to recommend this general approach, even though the particular types of devices used by Nielsen are not directly applicable to media consumed through a wide range of devices.

The chief benefit of sampling is that it allows greater energy and resources to be spent on ensuring that Alice's PC is secure and her vote will be counted correctly. The drawbacks are the introduction of sampling errors into the data (niche artists may either slip through the system undetected, or get lucky and be over-counted), and the loss of the desirable democratic sense that every taxpayer has control over the works their dollars will fund. As the sample size becomes larger, the first drawback decreases, but the second drawback remains troublesome.

Sampling was actually included in recent proposed legislation for extending France's private copying levies to cover P2P distribution.[212] It presents an attractive combination of accuracy and practicality for alternative compensation systems that a government might implement on a short timeframe. With a sampling system in place, it would be perfectly feasible to commence the longer-term project of building the security infrastructure necessary to allow universal virtual market suffrage, or at least suffrage for

---

realtime analysis, most of the screen could be filled with a flat colour, and a hash calculated only for the part that is not uniformly coloured.

[212]Jean-Baptiste Soufron, personal communication.

anyone who is using a trusted platform module to keep malware off their systems.[213]

Summary

Effective defence against technical vote-stealing attacks should be considered mandatory in the design of alternative compensation systems. By combining a strong system of defence — voting hardware or trusted computing systems — with the secondary protection of a web of honeypots, a virtual market might be rendered sufficiently resistant to technical manipulation. Although software flaws would potentially enable adversaries to change votes on occasion, manipulation on any serious scale could be detected and corrected rapidly.

The costs of this immune system will be discussed further in Chapter 9, but for the moment it is sufficient to say that they would be neither insignificant nor prohibitive. Trusted computing BIOSes (provided that they do not aspire to tamper resistance) are just a few extra transistors on a chip; the issue would be auditing its design and waiting for or encouraging them to be present in enough computers.[214] Extra security measures on client systems can be purchased more cheaply by collecting data from a limited sample population, rather than every single Internet user. The cost of a honeynet is a moderate fixed cost of developing those systems, plus the cost of deploying a population of them. Such ongoing costs are likely to be in the range of a few million dollars a year for a very substantial mesh of honeypot defences. All of these costs are manageable.

A final remark is in order about the computer security considerations that have been examined in this section. They have a distinctive character, which has much in common with the problem of remote Internet voting but few if any parallels in ordinary copyright marketplaces. We may wonder where it is that such security dynamics come from, and further, how their presence relates to any security issues in exclusive copyright systems. The answer, I believe, is that these obstacles are an unwanted shadow of one of the main benefits of a virtual market: the de-coupling of payments for works from the distribution of those works.

---

[213]Note that I am not suggesting that people be disallowed from participating unless they use a trusted platform module to remotely attest their software setup; rather, that people should have to declare that they have the means and inclination to keep their systems secure. Local attestation is just as useful as remote attestation for this purpose.

[214]Due to IBM's enthusiasm for the technology, one happens to be present in the laptop this thesis was typed on, but this is far from universal.

If works are not freely distributable, and every copy has its price, then the users who pay the bills have effectively been given responsibility for ensuring that remuneration is answerable to distribution. When affairs are so arranged that value and remuneration must be determined independently from the act of payment, the possibility arises for attackers to manipulate the processes that measure value.

As it turns out, this problem is not unique to alternative compensation systems. In fact, subscription based DRM models are equally susceptible. Recently launched services such as *Yahoo! Music Unlimited*,[215] Real Networks's *Rhapsody Unlimited*[216] and *Napster Membership*[217] allow users to download as many files as they want for a fixed monthly fee. Insofar as these services are paying royalties to artists that are based upon the popularity of their music,[218] they will also be the targets of hard-to-detect manipulation attacks that redirect those royalties. Indeed, the only thing mitigating the incentives for such attacks at present is the relatively small number of people using each of these services.

---

[215]See `http://music.yahoo.com/`

[216]See `http://www.listen.com/services`, archived copy on file with the author.

[217]See `http://www.napster.com/more_about_napster_ntg.html`

[218]Although I have not obtained architectural details of their royalty determination systems, they have stated that this is their approach. The Napster Privacy Policy (`http://www.napster.com/privacypolicy.html`) states,

> "In order to make sure that artists and copyright owners receive applicable royalties, this software identifies and counts the songs you have obtained and/or accessed including, if you are a Napster To Go member, tracks that you may have listened to offline on compatible portable devices. At the aggregate level (i.e., not tied to the personally identifying information of any user), we use this data to report and pay royalties, for internal analysis and we share this data with certain Partners for their own analysis."

The Rhapsody Privacy policy (`http://rhapreg.real.com/rhapsody_pages/policy.jsp?policy=privacy`) states:

> "RealNetworks uses your personal information to: ... Track content accesses, downloads and burns for the purpose of paying royalties and license fees to third party providers, such as record labels and other copyright holders and content distributors;"

And the Yahoo! Music privacy policy (`http://privacy.yahoo.com/privacy/us/music/details.html`) states

> "For licensing reasons, we send designated third parties anonymous track information about the music you listen to through LAUNCHcast radio or Yahoo! Music Unlimited, or about videos that you watch. We may also send aggregated, non-identifiable information about our users' declared zip codes and gender to third parties."

5.2.3   Human Security

In addition to the precondition that the ACS network be resistant to electronic attack, it is also particularly important to guarantee that there are no systematic incentives for human participants to trick the virtual market in some way. If votes can be exchanged for direct material assistance, if they are in practice fungible with cash, then they lose their particular social purpose — which is to facilitate the production of public goods. Incentives to "cash in" votes would, at the very least, reduce the quality of the information in the virtual market, and, at worst, render the whole system infeasible.

Straightforward examples of this are Alice voting for herself, or pre-arranged voting in small cliques.[219] Alice and her family might vote for her devious sister Delilah, who has created a fake artist's account. Some forms of clique voting may be automatically preventable without necessarily compromising pseudonymity.[220] The family plot to channel votes to Delilah could be foiled by a requirement that artists receive support from a significant number of users, before being eligible for remuneration. That way, the family can vote for Delilah only if many other people are also doing so, thereby indicating that she is actually an artist.

A related and perhaps more serious fraud risk is the deliberate transferal of identity. In this situation, Alice could "rent" her voting power to an "artist" (or a network of conspirators) in exchange for cash. This at first seems like a serious threat, because if it occurs with the consent of all parties, it will be almost impossible to detect. There may, however, be a simple and efficacious mechanism for preventing identity rental. The key is to make the agreements upon which trades are based very difficult to enforce — in this case, by making it costly to verify that Alice's vote has in fact been cast for a particular person.[221] If votes are independently unverifiable, and can be altered

---

[219]I use the word "clique" here in its ordinary English (as opposed to mathematical) sense. More precisely, the situation of concern is a smallish group of voters (less than a few hundred), whose votes flow exclusively, or almost exclusively, to other members of the clique.

[220]*Pseudonymity* here refers to the fact that although different votes made by the same person can be linked together, they cannot be linked to that individual. In this case, pseudonymity applies to all the ordinary users of the system, but not to performers who actually collect money from it (since these are the people who are actually able to vote for themselves or each other). This corresponds with Froomkin's definition of the term (see (Froomkin 1995)), although if this privacy cannot be compromised, a case can be made for describing it as "anonymity with persistent nyms"; see (Clarke 1999, § 3.6).

[221]Accounting information, telling artists about the number of votes they have received, forms a channel in the sense of Shannon; see (Shannon 1948). The problem of determining the verifiability of particular votes, as a function of the granularity of users' choices, the number of other votes being received by the same work or artist, and any quantisation and noise introduced by the virtual market, is deserving of a brief article in itself. Without conducting that analysis, it is probable that there exist

retrospectively, then one would have to actually keep the seller's secure voting hardware (in practice, probably the trusted computing BIOS chip from her motherboard), in order to reliably buy votes.[222]

So another risk to consider is that an unscrupulous Alice might be willing to sell her voting hardware. Firstly, let us consider the case in which that hardware is separate from her computer. If Alice is self-interested, she will sell if the price she is offered exceeds the costs of making the sale, which include lost opportunities, risks and side effects. To begin with, she might choose to sell because an artist could offer her half of the (present discounted) value of her future votes, which would easily outweigh the small risk of being caught. But the balance could be shifted, if the card has a dual purpose which gives it a direct value to the holder. It might function as a credit or debit card, a public transport card, or a link to some other valuable service;[223] if the costs associated with losing this gadget are greater than the financial benefits of selling it, Alice would be wise to keep it.

In addition to technical measures intended to raise the cost of buying virtual market votes, strong social incentives can also be deployed; it is usually practical to offer substantial rewards for evidence that leads to a conviction for attempted vote buying.

In virtual markets with universal suffrage, it is inevitable that some people will attempt to find ways to exploit the alternative compensation system. By making the barriers to entry high, it can be ensured that such attempts are not widespread. And by making attempted exploits both risky and unreliable, it can be ensured that business-minded criminals will find easier ways to make money. Like, perhaps turning to art.

The Role of Social Norms

Why should users participate in a virtual market? Surely it would be easier not to worry about voting — to download files by whatever means was easiest, and to save time by not rewarding the files' creators? In this section I will consider these issues, and some of the effects that mitigate them.

---

ways of making verifiability unprofitable, though there will be some tradeoffs against the ideal degree of specificity in the rewards which artists receive.

[222]If Alice was allowed to keep the hardware, she could take cash from the conspiracy, but simultaneously continue to reward her favourite artists.

[223]Although these proposals are inter-institutional, they may be beneficial to all parties, by providing secure, pseudonymous authentication, in an elegant and efficient manner. It might also be desirable to have the devices perform both pseudonymous and identifying authentications, for different applications, using separate keys.

Firstly, note that it is important to ensure that participation is easy — that downloads or usage can be tracked automatically, and that voting is not labour intensive. Previous sections addressed these design problems.

Provided that there are no barriers to participation, there is actually a direct and unambiguous incentive to allocate one's votes in accordance with one's preferences. If you read a book by an interesting writer, but fail to give them votes for their troubles, then it is less likely that the author will write another book; even if they do publish again, it is likely to take longer, because they will have to support themselves by other means. The difference between incentives to vote in an alternative compensation system and, say, the incentives to buy books in a marketplace is their strength, not their direction.

There would none the less be a role for education in encouraging public spirited participation in a virtual market. Enthusiastic following and support for authors and artists is widespread, even in the absence of direct incentives to encourage it, but this is a cultural phenomenon. One can easily imagine cultures in which rewards would be allocated "correctly", in the sense of voter honesty, and cultures in which they would not be.[224] This educative task is far less daunting than the idea of convincing teenagers to respect DRM, because a virtual market removes the need to "condition away" the free rider problem.[225] Whilst self-interest might dictate that many consumers should avoid technologies which prevent them from copying files between computers, there are no serious reasons for individuals to avoid participating in virtual markets.

Finally, it is interesting to observe that while social norms are highly relevant for mitigating the kinds of profiteering discussed in this section (clique voting, vote buying and voting hardware purchases), they are not particularly helpful for the computer security concerns discussed in Section 5.2.2. All of the 'human security' risks become rarer if they are stigmatised, because the threat they pose is proportional to the number of people willing to assist. In contrast, social norms are far less effective against

---

[224]Many similar problems have been faced in the field of *contingent valuation* (CV), where surveys are employed to value public (usually environmental) goods. Although there are a number of important differences (such as the role of passive use, and the fact that CV is used for centralised decision making, rather than decentralised allocations), the extensive research on the applicability of CV has shown that, if citizens are given appropriate information about the role of their contributions, they will provide high-quality information on the value of public goods; see (Carson, Flores, and Meade 2001).

[225]There have been regular claims that education will be critical to the success of technologically enforced copyright; see, for example, (Information Infrastructure Task Force 1995, Part III); (World E-Commerce & IP Report 2002); (Goodenough 2002) (a more exaggerated development of this notion, which goes so far as to suggest that neuroscience may have a role in conditioning people to obey copyright law).

network attacks because the hackers who are the exclusive perpetrators of such attacks might be engaged in private, highly profitable behaviour, or involved in small subcultures that do not respond to broader social values.

5.2.4 Conclusion: A Specific Proposal

Section 5.2 has discussed the design constraints that arise from trying to construct an alternative compensation system which allocates payments to the most valued cultural producers, and from anticipating and defending against technical and social modes of failure. The analysis leaves us with the skeleton of a specific design:

- Remuneration from the virtual market is determined by the result of a weighted ballot. Users who are voting have a vote suggested according to measurements of their digital reading, listening and viewing habits, but are able to override or adjust the results of those measurements.

- The 'election' should be conducted using a cryptographically verifiable protocol.

- Usage measurements are conducted using the family of algorithms proposed in Section 5.2.1.

- Suffrage in the election should initially be limited to a small representative sample population (in the order of tens of thousands of individuals).

- Trusted computing hardware should be employed for auditing the machines that collect usage data and votes from the sample population.

- Professional systems administrators should oversee the security of these systems.

- Some of the machines in the sample population should be honeypots, with sound cards modified to audit usage measurements, and monitor cables, mice and keyboards modified to audit the user interface for voting.

- Once the system is operating smoothly, anybody with appropriately secure hardware should be offered suffrage. Some of the new voters should also have honeypots installed on their systems.

- As suffrage becomes widespread, a minimum threshold should be applied to remuneration, so that it is only paid to artists who attract an audience larger than a family or a small circle of friends.

- Strong criminal penalties should apply to people who deliberately set out to purchase votes in the compensation system.

- Rewards should be offered for evidence that leads to convictions for vote-buying, so that vote sellers have an incentive to defect in vote-sale transaction.

5.3   Funding Virtual Markets

The reader may by now be willing to believe in Alice, but I would not wish to stretch her credulity to the existence of Wonderland — it is relatively easy to give money away, but first, one must obtain it from somewhere; where there is "public" funding, one usually finds a private taxpayer. With taxation, the questions "on whom, and how much?" must be answered, because they play a central role in determining the economic properties (and political fashionability) of any policy endeavour.

The specification for funding an alternative compensation system can usefully be partitioned into two problems: selecting a tax formula, or several tax formulae; and then deciding the rates at which each of them should be levied. Sections 5.3.1 and 5.3.2 will deal with these questions in turn.

5.3.1   Where should the burden fall?

There is a wide range of taxation models available for funding a virtualised copyright mechanism. The criteria for choosing among them include political feasibility, the degree to which the tax "fairly" charges those who benefit more or less from the virtual market, and the kinds of "distortionary" incentives the taxes might create for people to change their behaviour (for instance, a tax on blank storage media might discourage people from making backups).

These considerations are sufficiently complicated that I will put them aside until Chapter 10. For the time being, let me simply pluck a tax from a hat and predict that the best way to finance a virtual market would be to combine three different revenue sources: one-third coming from taxes on Internet connections; one-third coming from levies on hardware, such as iPods, stereos and hard disks; and one-third coming from

a progressive source such as income taxation, residential property tax, or general revenue.

5.3.2 How much to raise?

Having chosen the form that taxation will take, it is then necessary to determine the rate at which it is levied. Those rates are indirectly determinative of the total funding pool for the virtual market — and thus of the amount which each artist receives from the system. Needless to say, getting this right is both important and difficult.

Fisher has emphasised that extrapolation from previous sales in the music industry would be invaluable in ensuring a smooth transition from existing systems of cultural distribution.[226] The idea is that the amount of tax raised would be dynamically adjusted so as to replace the amount of revenue lost to the industry through file sharing, as the levels of sharing increase. While helpful in the short term, this approach gives little long-run guidance.[227]

There are several more sustainable strategies for setting the tax rates used to fund an alternative compensation system. The simplest is to rely on a governmental or administrative decision making process, in the hope that various interests groups will end up agreeing on a compromise which is fairly good. Artists and publishers want higher rates; ISPs and hardware manufacturers want lower rates. Taxpayers want lower rates, unless they are so low that they end up harming the supply of works. These groups could take turns shouting or deploying their charms to persuade regulators to modify the formula. While it is possible that the administrative compromise would be reasonable, there would be plenty of room for poor decisions and dissipation of resources through regulatory contests and rent seeking.

At the other extreme, one might opt for complete decentralisation in setting the total tax level; users of the system are informed about the amount that artists are currently earning, they know from experience the levels of subjectively relevant cultural output, and are then given a regular choice to vote "higher" or "lower" for funding levels. From the perspective of economic theory, this approach appears to be quite effective (see Section 7.4.2 and accompanying footnotes). But the task of ensuring that consumers are sufficiently well informed to vote rationally in such a system, and the improbable prospect of governments surrendering control over certain tax rates, makes

---

[226](Fisher 2003, at 14).
[227](Liebowitz 2003a, at 15).

this option more of a theoretical curiosity than a serious policy option.

There are intermediate options. One would be to try to establish a statistical picture of the "exchange rate" between willingness to pay and virtual market votes. This exchange rate could be determined by using contingent valuation (CV) surveys on relatively small groups of consumers.[228] They are asked hypothetical questions, such as "at what level would a reduction in taxes paid be sufficient compensation for the loss of this album?", or "if this particular album was *not* going to be made available online, how much would you be willing to pay to obtain a copy privately?".

Collecting data of this sort not only allows the audience's preferences to control the tax rate, but it provides an algorithmic test of the consistency of virtual market voting behaviour. If the willingness to pay values reported by consumers in CV surveys are *not* approximately proportional to the votes those consumers actually made,[229] then it is possible to infer that there are systematic flaws or inconsistencies in the voting process. The largest problem with contingent valuation methods is that they, like the historical sales data Fisher recommends, might eventually cease to be useful. Consumers who were used to paying for music or e-books or DVDs may have a reasonable notion of how much they are willing to pay for those things. But after, say, a decade of free unlimited access, valuations might be less informed by any actual experience of conditional exchange.

A fourth option, which I believe is the most promising, combines several of these approaches. Following Fisher, begin with a tax rate which ensures that on the day of transition, the amount of revenue the music industry receives from the virtual market is the same as the amount it previously received from the sale of sound recordings (or more or less, if there is sound economic data to indicate that the current amount is non-optimal). Subsequently, adjust this number by using survey data to measure changes in the amount of time that people are spending listening to sound recordings. If the number goes up, adjust the tax rate to keep the funding level proportional, and vice versa. The question of how effective this method would be is investigated in detail in

---

[228]*Contingent valuation* is a method in which public goods are valued based on carefully constructed surveys administered to small focus groups. These surveys usually ask consumers how much they would be willing to pay for the particular good, or conversely, how much they would be willing to accept as minimal compensation for the loss of the good if it already exists. For examples of applications of CV for copyright-related purposes, see (Bohm 1972); (Brennan 2002). For an overview of the issues involved in ensuring that CV data is accurate, see (Carson, Flores, and Meade 2001).

[229]Note that if you ask a consumer about a single work, their CV report and votes are proportional because the constant of proportionality is unconstrained. But if they are asked about a number of different works in their collection, a meaningful test for correlation is possible.

Section 7.5.1.

5.3.3   One dollar, one vote?

An additional difference between funding mechanisms is that they can vary in whether and how they enable a link to the distribution of funds to artists, from the amount paid by each taxpayer — in other words, some funding models are inherently "one taxpayer/user, one vote"; while others allow a choice between "one dollar one vote" and "one taxpayer one vote" allocation formulae.

The simplest approach to administering the funding for a virtual market would be to leave the collection of revenue separate from the process of distribution. Everyone pays their taxes, and then Alice's votes (or usage, or downloads) determine a reward for each of the artists she likes, rewards whose size is independent of the amount of tax she herself paid. This is a "one user, one vote" system. But under some circumstances, it may be possible to use Alice's preference to allocate precisely the same amount of cash that she had contributed through taxation. This "one dollar, one vote" approach would only be feasible for certain kinds of taxes — it would be relatively easy for surcharges on ISP bills, harder for income taxes and impractically costly for hardware or blank media levies.

"One dollar, one vote" taxation has some advantages and some drawbacks. On some normative accounts, a democratic basis for culture would be particularly desirable, while others might hold that the tastes of well-educated (and hence wealthier) taxpayers are more likely to reach an underlying goal of aesthetic value. Utilitarian analysis would favour weighting by the psychological intensity of demand, a troublesome quantity which might nonetheless be inferred by willingness to part with dollars for complementary goods[230] or by indirect means.

The choice between one-dollar-one-vote and one-user-one-vote has some consequences for the nature of the incentives created by a virtual market. These implications are explored a little further in Section 7.5.4.

---

[230]Although purchases of complementary goods may still be less informative than purchases of the actual cultural items in question; see Section 7.5.3.

5.4   Which Information Markets Could Be Made 'Virtual'?

Both the complexity and the persuasiveness of propositions like the virtual market vary according to the economic details of each kind of copyright subject matter. So which industries should the proposal apply to? Part IV will find some variables that make the case for alternative compensation systems stronger or weaker in different industries. But in the meantime, there are also some fundamental preconditions which make the virtual market proposal more compelling in some situations and more sociologically complicated in others. I will address these here.

5.4.1   Derivative, 'Monolithic' and composite works

Some very interesting organisational questions arise when we consider the relationship any kind of compensation system and situations in which artists, authors and programmers collaborate, build upon, modify, adapt, derive from, sample, remix, "mash-up", wikify, embed, and otherwise incorporate others' work into their own, and their own work into that of others.

In the status quo copyright regime, these creative processes are often infringement unless licenses can be obtained from all of the chronologically preceding copyright holders. How should a virtual market treat them?

Principles for permitting remix culture

One fairly unsatisfying option would simply be to regulate modified works exactly as copyright law currently does, preserving an exclusive right of derivation[231] and legal norms about joint authorship,[232] leaving the objective of liberating copying separate from the objective of liberating remixing and transformation.

---

[231]Or its equivalent, in copyright systems with no distinct right of modification. In UK and Australian copyright, for instance, there is no single derivation right; there is an exclusive right of adaption that applies to substantial transformations such as translation from one language to another or adaptation from one medium to another. In other cases, acts such as sampling, remixing and transposition are controlled by the exclusive right of reproduction.

[232]See (Chen 2011, § III), surveying U.S. caselaw on disputes about joint authorship, and observing a tremendous reluctance to find that creative contributions lead to co-authorship, except in situations where the dominant author intended to bring a second author into an equal partnership, and understood the legal consequences of doing so.

There are better suggestions in the literature. Fisher proposed that artists who include elements of others' works within their own should be required to declare that fact, thereby passing a fraction of the revenue from the modified work back to the creator of the original (Fisher 2004, 234–236); the particular fraction would depend on the declared portions of the work that were sampled, and the initial onus would be placed on those producing derivative works to credit and apportion to their sources (Fisher 2004, 206), though administrative and judicial dispute resolution mechanisms would be available. Netanel proposed that noncommercial derivations be permitted without any requirement for permission, but that licenses should be required if the secondary authors wish to earn money from their work.[233] Netanel's proposal is the safer bet, in the sense that it is a guaranteed improvement on the status quo. Fisher's proposal is more ambitious and possibly better (it would, for instance, save professional documentary film makers from spending huge amounts of time negotiating licenses) but its practical and psychological dimensions are complicated and seem hard to predict.

It turns out that Hollywood runs a living laboratory for how rule-based credit division along Fisher's lines might play out. Union contracts for scriptwriters working on major motion pictures employ similar principles, although they are heavily biased to favour the first author to start working on the script: the first author retains a title credit provided their work remains 10% of the final product; subsequent authors may be included if they contributed more than 50%. Friend (2003) provides a fascinating account of this system at work. Such rules are functional, in the sense that they divide credit for scripts, but in that high-stakes setting, they also generate plenty of acrimony.

The problem of peer production

The Fisher or Netanel proposals are sufficient to marry financial incentives with "remix culture" in a lot of cases. But there may be limits to how complicated works of multiple authorship can get before institutions for remuneration start to impose consequences for industrial organisation. What if a collaborative project has thousands of participants who make contributions of varying shapes and sizes and at varying rates, over long periods of time? What if we wrote our encyclopedias and operating systems that way?

I propose that the strongest predictor of whether such royalty-splitting algorithms are sure to be workable is whether the copyright works are *monolithic*. That is, the

---

[233] See (Netanel 2003, 38–39); Netanel's proposal was echoed as a proposal for exclusive rights regime as the policy conclusion in (Lessig 2008, Section 9.4)

good is for the most part free standing, and created once, rather than undergoing a continuous process of maintenance and development. "Created once" is perhaps too strong: a song that is written and recorded once could be remixed into a second or a third monolithic work. But that would be different to putting an audio editor file online, and letting dozens of people move the notes around or transform the waveforms on a continuous basis.

In the Internet era, non-monolithic works have begun to often be produced by a process that Benkler (2003) has termed "peer production", a mode of industrial organisation based primarily on voluntary collaboration. Peer production is quite different to the organisation of industries into markets (where co-operation is organised and bounded by transactions and contrasts) and firms (where co-operation is organised by a management hierarchy and incentivised by wages and salaries) that one typically finds in capitalist industries.

How to divide credit or payments among the contributors to peer production processes is not a simple question: a fair apportionment is not only hard to measure but continually changing.

The German Wikipedia community has confronted this question directly. In 2007, the German collecting society VG Wort created a payment scheme called METIS to distribute a significant portion of the country's private copying levies to websites. In 2008 and 2009, there were extensive discussions between the METIS administrators and the German Wikimedia chapter about whether the encyclopedia's editors could receive a portion of this significant revenue stream. The German Wikipedian community voted against these proposals. One significant reason was that, while editorial/authorial contributions to Wikipedia[234] are recorded and conceivably measurable, it would be harder to fairly divide remuneration between editors and others who made contributions of other forms, including site administration[235] and software engineering for the project.[236]

Virtual market funding for free software projects encounters similar conceptual difficulties. Imagine Alice attempting to choose which software has been of great use

---

[234]Perhaps tellingly, the Wikipedia community seems to avoid the term "author"; contributors are called editors regardless of whether they focus on writing new material or polishing, maintaining and defending words written by others. As of late 2011, the term "wikipedia editors" occurs 32 times more often on the wikipedia domain than the term "wikipedia authors".

[235]For a summary of the roles performed by Wikipedia administrators, see `https://en.wikipedia.org/wiki/Wikipedia:Administrators`.

[236]See (Chen 2011, note 205) or, for original German sources, `https://de.wikipedia.org/wiki/Wikipedia:METIS`.

to her (so that she can vote for it). Perhaps she has found her brand-new 3-dimensional web browser to be particularly helpful, so she gives it lots of votes. But it may turn out that the 3D web interface was a relatively straightforward piece of code to write, because most of the work had already been done in developing libraries for sophisticated 3D graphics, providing network functionality, and for handling web-related protocols and file formats. Although it appears to Alice that the user interface level application is providing numerous useful features, the actual work has been done by many separate components. Furthermore, each of these software sub-structures may have been written by many different contributors. The virtual market lacks sufficient information to determine where the reward should go.

In a marketplace of proprietary software, a complex web of contracts, negotiated between software production firms with various levels of market power, acts to define the relative remuneration received by different contributing firms. Compensation for individual programmers is in turn defined by their employment contracts. This market structure can be criticised on the grounds that it interferes with the most efficient ways of writing code,[237] and because its answer to the credit allocation problem is heavily coloured by the effect of market power.[238] Still, at least it has an answer; free and open source software development models, which escape many of proprietary software's weaknesses, commonly skip the task of financially significant credit allocation altogether. In order for virtual markets to include software, they would have to find some way to mimic the proprietary web of contracts, without also importing the associated web of constraints.

Does this mean that virtual markets are inherently incompatible with free software, Wikipedia, or other massively collaborative peer production? Not necessarily. Benjamin Mako Hill (2005) has pointed out that, while the existence of financial incentives poses a threat of "crowding out" of voluntary contributions to a peer production endeavour, there are examples of very large and successful projects, including Ubuntu and Mozilla, which appear to combine the two modes of organisation fairly successfully.[239]

---

[237]For critiques of the proprietary software production model, which are widely cited amongst programmers, see (Stallman 1992b); (Raymond 2000). For theoretical work which attempts to explain when open or closed development is efficient, see (Benkler 2003).

[238]This may be a exacerbated by the poor architecture of intellectual property regimes for software. For serious normative treatments of those design issues, see (Samuelson, Davis, Kapor, and Reichman 1994; Gibson 2005).

[239]For an early account of Mozilla's hybrid organisation and the relationship between paid, full-time

The degree to which peer production can be sensibly and efficiently supported by a virtual market is an open question that I am not going to fully resolve here. But one could imagine that if one set out to use ACS funding to produce a free encyclopedia, it might be simpler to fund an authorship process that looked more like Encyclopedia Britannica's than like Wikipedia's. Similarly, using a virtual market to pay the submitters of apps to the iPhone App Store or the Android Market would be administratively simpler than trying to apportion credit and payments among the authors of the Linux kernel.

5.4.2   Commercial and Non-commercial use

Virtual markets could be constructed as a narrowly-tailored licensing solution to the 'problem of piracy in the music [or publishing, or film] industry', along the lines of a more ambitious private copying scheme. They could also be conceived as a complete replacement for an arm of the copyright system, with the objective of liberalising all uses of the relevant works.[240]

Casting the institution in those broader, exclusive-rights-replacing terms would incorporate several classes of activities within its ambit that are clearly beyond what might be termed 'private copying' or perhaps more helpfully 'non-commercial use'. Prominent among these would be the distribution of works designed for use by businesses (such as market intelligence reports, corporate accounting software, legal treatises, database services, and so forth); the use of cultural works in the service of commercial ends such as advertising; and the transformative use of material in producing derivative works or adaptations.

There are good reasons not to attempt to license these activities through a virtual market. In comparison to the narrowly cast "legalise file sharing" or "pay authors writing for the Web" model, an ACS that covered such things would involve more conceptually problematic and incentive-incompatible valuation signals, would risk interference with useful freedoms of contract, and would be standing in for much more complex and heterogeneous market processes.

It is safe to say that virtual markets would reduce the existing role of contracts in establishing conditions to accompany the exercise of exclusive rights by users. The

---

contributors and the wider development community, see (Mockus, Fielding, and Hebsleb 2002).

[240]Authors who appear to lean towards this position include (Calandrillo 1998; Shavell and van Ypersele 2001).

broader the virtual market, the broader the damping of contracts. In the case of uses of cultural works that are enjoyed by private individuals in their daily lives, it is difficult to see how subjecting those uses to contractual bargaining has any potential to create any sort of value in the first place.[241] But it is not so easy to discount the negotiations that might surround a complicated piece of software, an advertising deal or the for-profit cinematic adaptation of a novel. While there are no doubt cases in which haggling over commercial rights is inefficient, and while the interests of authors could be independently protected against the advertising industry with moral rights mechanisms, there are many situations in which licensing negotiations render essential assistance in co-ordinating and valuing businesses' activities.

Other difficulties arise from the nature of the measurement and voting procedure. When an individual reads a website or listens to a song, we have at least a statistically suggestive idea of the nature of their appreciation. If they actually vote for or against those items, the picture is even clearer. It is not so clear how a virtual market could collect similarly informative data from businesses. Usage measurements will certainly be possible in many cases, but because the price of digital goods aimed at businesses varies greatly, it is difficult to translate such usage into a valuation. Voting might be possible, but it would be fraught with incentive incompatibilities. With the provisos discussed in Section 5.2.3 above, it is safe to assume that most people will value cultural goods enough to be fairly honest in voting for them. But imagine that a corporation is given a stack of votes in exchange for the taxes it pays. What is to stop the relevant executive from giving 80% of them to a friend's company? No plausible regulation of the process could do much to change such possibilities.

On these grounds, virtual markets seem rather less promising as means to fund the commercial use of copyright works than they do as a way of avoiding the trauma of enforcing copyright against private citizens.

5.4.3   Conclusions on Scope

Virtual markets would be easier to construct for some categories of copyright subject matter than others. Funding the production of of digital books, music, web pages with identifiable authors, photographs and moving pictures, for non-commercial audiences,

---

[241]License agreements when they appear in consumer copyright markets tend to either be contracts of adhesion or versioning mechanisms for price discrimination. Such contracts are unlikely to promote utility when the free rider problem is separated from the equation.

are the simpler cases.

Software is a more complicated case. Virtual markets for software could be implemented, but they might alter existing modes of industrial organisation, particularly around free and open source software projects. Even under, say, Netanel's rule, payment streams that required agreement about credit and royalty division could be seen as a taxpayer-funded subsidy that favoured management hierarchies over community projects, cathedrals over bazaars.[242]

These issues are probably resolvable. Perhaps more free software projects would adopt the kind of semi-centralised organisational models that Mozilla and Ubuntu use; or perhaps bazaar-style projects would adopt their own conventions about how to divide remuneration for their work.

Like software, commercial uses of copyrighted works are also a more complicated case. But it is less problematic to exclude this category of works and uses. The digital copyright crisis is mostly a crisis around non-commercial use.

I am going to focus on analysing the four regimes as they apply to non-commercial use of literary, artistic, and entertainment works. This is the traditional core of copyright, and the domain in which contemporary political crises are most likely to lead to reform.

It is hard to imagine a proposal like the virtual market being seriously considered for software unless it had already been demonstrated to work for music or books. In part, this is because the critics of software copyright have already established a strong separatist position, using free and open source licenses to build themselves an independent corner of the world in which the usual logic of exclusive rights does not apply; this alters the priorities of software copyright skeptics and lessens the political imperative to "liberate" all of the proprietary software. In contrast, the Creative Commons movement will never build a CC-licensed replacement for the Beatles.

But in parallel, the political improbability of virtual markets for software is also increased by the fact that proprietary software firms already have business models that are quite virtual-market like in the sense that they involve collecting taxation,[243] and

---

[242](Raymond 2000)

[243]In the PC sector, Microsoft negotiates contracts with laptop and other computer manufacturers that strongly incentivise these manufacturers to include the cost of a copy of Microsoft Windows with every machine they sell. These contracts are often called the "Microsoft tax" or the "Windows tax". Customers who wish to buy the hardware but run free/open source operating systems on it struggle to obtain refunds, if it is possible at all (see `https://en.wikipedia.org/wiki/Windows_refund`). In this sense, the prevailing circumstances in the PC operating system market are actually a hybrid of the

are therefore not threatened by piracy.

So in summary, for the purposes of Parts III and IV of this thesis, I am going to analyse the regimes as they might apply to the non-commercial use of writing, music and film, including websites. Some complicated cases (such as large collaborative wikis) remain encompassed. Under a principle like Netanel's for works of multiple authorship, it would be up to the communities that produce the works to experiment and decide whether they wished to collect and distribute royalties to contributors, or not.

This conclusion about scope does not mean that alternative compensation systems are uniformly *desirable* for all of these media; there are technical and economic reasons why one might want to place some of those within an alternative compensation regime and leave others under copyright regimes.

5.5   Conclusion

This chapter has considered the possibility of a "virtual market" through which authors and artists could be remunerated from the public purse. There are difficult sides to the design and implementation of such a system: ensuring that high-resolution information about the relative and total value of works is available; preventing the manipulation of that information by technical or social means (or at the very least, ensuring that such manipulation is not profitable); preempting rent-seeking by various interest groups affected by the system; choosing tax or levy formulae that balance fairness and efficiency well. All of those problems appear to be solvable. A selection of practical security measures, including votes verified with trusted computing, honeypots to detect attacks, and sampling methods, are available to ensure that virtual markets could operate reliably.

The question to be resolved in Part IV is whether we should expect these "virtual market" alternative compensation systems — imperfections and all — to work better than other means of remuneration.

---

feudalist and virtual market models. The tax-based revenue streams are virtual market-like, although the receipt of those taxes by a single firm is decidedly feudalistic. Apple's vertical integration and bundling of hardware an operating systems creates the same result, of course.

In the mobile phone market, things are remarkably similar. Although Microsoft has been unsuccessful at bundling its software with many phones, it has been very successful at collecting royalties for handsets running Android (Brodkin 2011).

Part IV

Economic and Normative Analysis

In the preceeding two Parts, I have characterised and examined four possible digital copyright regimes. I now turn to the question of which of these regimes a society would do best to institute.

This question is more easily posed than answered. The differences between the regimes are multiferous. Some are stark and easily delineated, but some of the consequences of changing a cultural remuneration system are exceptionally subtle. Any difference, obvious or otherwise, might provide grounds for preferring one copyright system over another.

Reaching a satisfactory normative conclusion about the approach that society *should* take to digital copyright is rather like putting together a jigsaw puzzle. A few sections of the puzzle, such as the fact that more copies of works will be made if there are no restrictions on copying, or the fact that using taxation to fund remuneration will have some negative side-effects, are relatively shapely and therefore easy to assemble, weigh and measure. Other parts of the problem — such as working out how differences in the remuneration system may result in differences in the behaviour of artists, publishers and the financiers of cultural enterprise — are harder to slot into place. This Part attempts to assemble a coherent picture that includes deals accurately with these questions, too.

Each "region" of the puzzle has a chapter allocated to it. The way that I have chosen to make this decomposition is neither perfect nor unique. It does, however, serve to break the problem up into parts that can at least be comprehended in an adequately independent and systematic fashion. The components I have selected are as follows: the cost of creating artificial scarcity through exclusive rights; the effect of remuneration systems on the incentives for cultural production; the cost of infrastructure necessary to make each system function; the effect of each regime on the transaction costs of using the copyright system; and the cost of using taxation within public funding models. I believe this list is sufficiently comprehensive to include all of the important policy reasons for preferring one regime to another.

Chapter 6

The Cost of Artificial Scarcity

Contents

> *'It appears to me, Miss Leete,' I said, 'that if we could have devised an arrangement for providing everybody with music in their homes, perfect in quality, unlimited in quantity, suited to every mood, and beginning and ceasing at will, we should have considered the limit of human felicity already attained, and cease to strive for further improvements.'* (Bellamy 1888)

The quality of digital goods that sets them apart from so much else in our lives is that they are so thoroughly lacking in scarcity. Their non-rivalrous nature ensures that more or less everyone with a computer (or a even just a phone) can have them. When copyright prevents this access, it must be counted as harmful.

The possibility of abundant distribution does not, it should be remembered, ensure abundant supply. Until computers are capable not only of reproduction but of authorship, we will have to rely on creative labour — a finite resource — to compose the digital world. Many disagreements about digital copyright can be explained as one side minimising one of these points while the other side minimises the other. Of

course, a realistic perspective must take both seriously.

This chapter weighs the normative implications of the first issue: how much distribution occurs under each regime? Until Chapter 7, it will be assumed that changes in patterns of distribution do not change the supply of works.[244] Section 6.1 briefly reminds the reader of the economic concept of deadweight loss from monopoly and discusses its application to copyright markets. Section 6.2 considers the extent to which deadweight losses can be mitigated by charging different prices for different users. Section 6.3 provides some estimates of the actual dollar value and percentage costs of scarcity. The inefficiencies are large, particularly in the music industry — larger in fact than the total revenue of that industry, and in the order of billions of dollars per year for an economy like Australia's, or tens of billions in an economy like the United States. Section 6.4 then presents an argument for why we should in fact be even more concerned about artificial scarcity than dollar-value measures of its cost might recommend.

6.1   Deadweight Loss, Competition and Monopoly

The concept of *deadweight loss* in markets for public goods was first mentioned back in Section 1.1.2. In simple economic models, it is often represented as a portion of a graph that looks something like Figure 6.1.

Whenever the marginal cost of providing a good to more people is lower than the price that is set for the good, there is the possibility of "deadweight" inefficiencies. Some of those people will not pay the higher price, but would have paid the marginal cost. This is strictly Pareto inefficient, because those people could be made better off without making anyone else worse off. In the case of reproducing digital information goods, marginal costs are very close to zero. Unless the price is also close to zero, inefficiencies are guaranteed.

These kinds of deadweight losses are normally studied using economic models of monopoly pricing.[245] This has been a source of occasional confusion; some authors argue that because competition occurs within copyright industries, exclusive rights cannot really be monopolies (Kitch 2000), and cannot really cause deadweight losses.

---

[244]Chapter 7 removes that assumption, and Section 7.3.4 in particular models the multiplicative interaction between the number of works that are incentivised by copyright and the artificial scarcity cost of those same exclusive rights.

[245]See, for instance (Tirole 1988, Chapter 10); (Shavell and van Ypersele 2001, § II.B)
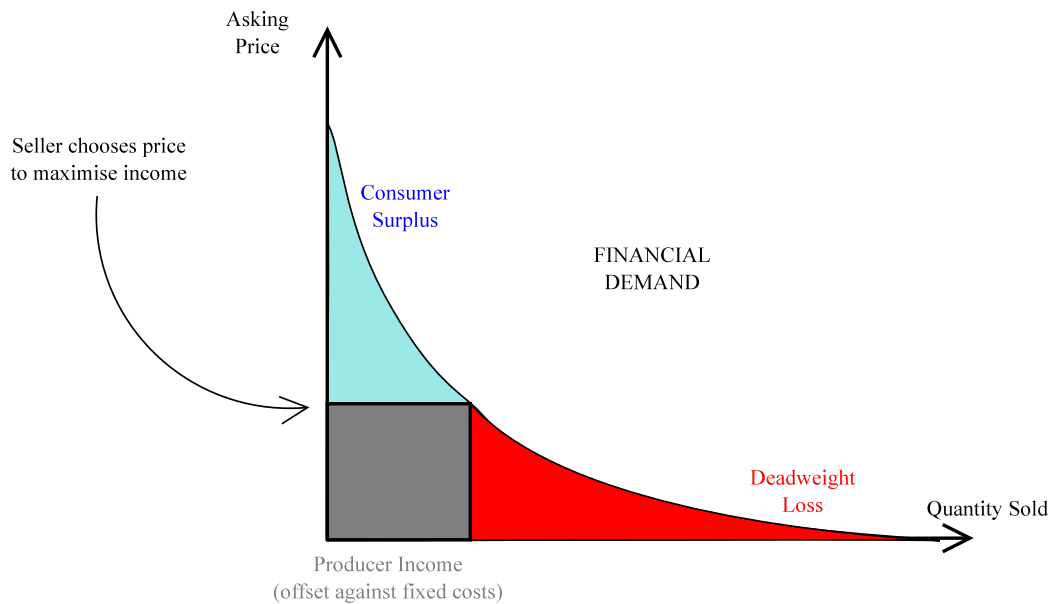
Figure 6.1: Deadweight loss under monopoly pricing with zero marginal cost

These are arguments about definitions, and are not especially enlightening.[246] It is clearer to note that what varies from copyright to copyright and patent to patent is the scope of the monopoly. This ambit may be narrow or broad depending on the degree of substitutability between different copyright works, on the amount of introspective knowledge that consumers have about their own tastes, and even on the amount of information that sellers have about consumers' preferences for other works.[247]

The most basic kind of monopoly, the one that is always present when copyright is enforced, is a monopoly on the particular work in question. Even though there may be many other similar and substitutable works, the rights holder always gets exclusive benefits from their particular computer program or book or song. Those monopolies allow producers to charge their average costs, which are higher than marginal costs, and thus, in equilibrium, to recoup their initial investments and stay afloat.[248] Although

---

[246]Kitch is arguing that a monopoly a situation in which people in practice *must* buy something from a particular vendor, rather than a situation in which there is only one vendor for something.

[247]For a powerful illustration of this last point, see (Detering 2001, § 1). Detering's model gives results which favour consumers heavily despite perfect price discrimination, by allowing sellers of imperfect substitutes to offer aggressive discounts on a consumer-by-consumer basis. But the kind of information-rich market which he models will be extremely rare in practice, if it occurs at all.

[248]This is a fairly stylised picture of affairs, though it is the standard one to be encountered in the economic literature. Note that investments include a risk-adjusted cost of funds, so "average costs" actually allows for standard return-on-investment profits.

these very narrow privileges need not have much in common with the history of Stuart monarchy, trusts and robber-barons that originally made the word "monopoly" so fearful, the fact remains that they are accompanied by inefficiency. Normatively, the extra copies that would sell for between marginal and average cost should still be made.

There are also situations in which copyright is accompanied by serious "market power" that allows producers to choose profit-maximising prices that appropriate substantial portions of consumer surplus. That happens if the copyright monopoly turns out to affect not only the market for the precise work it covers, but the larger market for works somewhat like it. Several dynamics may contribute to this outcome.

One is the inherently limited substitutability of many copyright works. Many artists' work is distinctive, and this leads to a degree of market power. In some cases, their talents are unique, in which case they receive a 'monopoly by endowment' which is potentially expansive. Brilliant writers, virtuoso performers, singers with extraordinary voices, and publishers and record labels with the foresight to sign them are all potential beneficiaries. In other instances, an artist's work can be emulated, but the reputation that they have captured by popularising it first is unique. A style of written expression, a musical innovation, a new mode of delivery or a cultural transplantation may be appropriated in this way. There is no better example than Elvis Presley.

It is clear that limited substitutability, endogenous demand[249] and network externalities[250] play an important role in copyright marketplaces, but the quantitative extent of that role has not been established.[251] And, while some economic models have attempted to capture substitutability, they are, of necessity, heavily stylised.[252] In any case, a reliable conclusion is that the uniqueness of artists increases the scope of copyright-based market power in cultural industries to some positive, and possibly large, degree.

The presence of network externalities in demand for copyright works is another source of market power and concomitant deadweight losses.[253] The staggeringly impressive profitability that Microsoft and later Apple have extracted from their platforms illustrates this well. But network externalities are not confined to software markets, and there are compelling indications that network externalities also apply to cultural

---

[249] *Endogenous demand* is the economic term for acquired tastes (and addiction).

[250] The role of network externalities in entertainment and cultural marketplaces is discussed further in Section 7.5.5.

[251] Compare, for example (Cohen 1998, 520–522).

[252] See, e.g., (Rosen 1981, 847–848).

[253] Compare, for example (Farrell 1995).

goods.[254] Some further deleterious effects may flow from firms' efforts to capture these externalities. It should also be noted that, although markets dominated by network effects are subject to competition, such competition must occur *ex ante*; by the time the network has formed, the results have been decided. This kind of competition is better than nothing but less effective than markets that allow entry at any point. For all of these reasons, network externalities increase the extent to which copyright is a strong grant of monopoly pricing power.

Finally, there are other examples of monopolies within copyright-based industries that are due to consolidation, large production costs, or barriers to entry from high fixed costs in marketing and distribution.[255] Those phenomena can arise in any sector of a capitalist economy and are not particularly associated with copyright. They are also likely to become less common as digital technology lowers the cost of distribution infrastructures and makes niche marketing more efficient.

In conclusion, the default situation for copyrighted works is that rights holders will sell them at a price which is at least their average cost and is higher than marginal cost. In many instances, where demand is endogenous and substitutability limited, or where network externalities are at work, the absence of close competition will allow producers to appropriate consumer surplus freely. Both of these phenomena (investment recovery and surplus appropriation) decrease distribution and therefore cause deadweight losses.

To be clear, this does not mean that the kinds of market power discussed here cannot serve useful incentive functions, or that the revenue they imply for producers is necessarily undeserved. The point to be emphasised is that, whatever the answers to those questions may be, copyright is a grant of monopoly that is intimately accompanied by deadweight inefficiencies that are, by definition, avoidable.

6.2   Mitigation By Price Discrimination

Many advocates of strong copyright rules argue that deadweight losses are overcome by price discrimination.[256] The idea is that publishers have an incentive to minimise

---

[254]See Section 7.5.5 below.

[255]See, for example, In Re: Compact Disc Minimum Advertised Price Antitrust Litigation, MDL Docket No.   1361 (D. Maine), 2003 `http://www.musiccdsettlement.com/english/finaljudgmentorder.pdf`.

[256]Duffy, for example, applies this specific argument against alternative compensation systems (Duffy 2004, Part I.A.4). An implausibly extreme variant of this reasoning can be found in (Easterbrook 1999,

deadweight loss by producing cheap versions of their works to appropriate the extra demand at the bottom end of the market. The price discrimination theory is thoroughly critiqued in the literature (Cohen 1998; Boyle 2000), but I provide a brief treatment of some of the issues here. The main objective is to convince the reader that artificial scarcity remains problematic, and it must still be weighed and measured, even when price discrimination occurs.
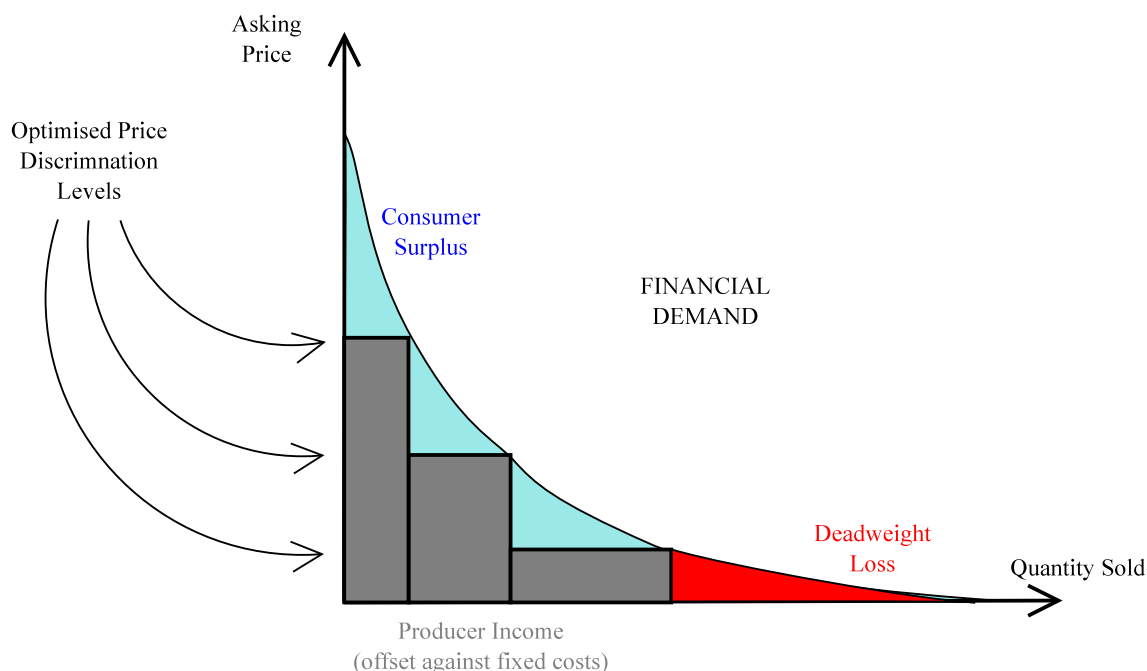
Figure 6.2: The effect of price discrimination on deadweight loss and surpluses

The effect of a typical price discrimination strategy is illustrated in the contrast between Figure 6.1 and Figure 6.2. It shows the demand curve for a copyrighted work, with a seller who has chosen three prices for three versions of the work, and has been able to perfectly sort consumers so that each consumer Alice purchases the most expensive version she is willing to pay for. Compared to a monopolistic vendor who only chooses one price, several changes are evident: producer income is higher; deadweight loss is lower; and the change in consumer surplus is dependent on the

p. 112). *Price discrimination* is said to occur when the producer of the good creates several versions, and sells those at different prices to consumers with different levels of demand. Alternatively, *intertemporal price discrimination* can occur when the producer lowers the price over time, selling to high-valuing consumers first and to lower-valuing consumers at later dates. A more comprehensive account of price discrimination for digital goods can be found in (Shapiro and Varian 1999, Chapters 2–3).

shape of the demand curve, though generally it will decrease as price discrimination becomes more complete.

The first point is that price discrimination is always limited. Methods for achieving it can often be defeated by arbitrage or by consumer learning. For example, the price of a good may start out being high, in the hope of selling it to high-valuing customers, and then drop over time, so that those with lower valuations also purchase a copy. But if consumers learn this pattern, they may decide to wait for a lower price, even though they would be willing to pay the higher one.[257] A more audacious model might involve producers closely monitoring the behaviour of each consumer, and then offering them a media file at a price determined from their wealth, their past purchases, and their actual watching/reading/listening habits (as reported by a DRM system). But even this scenario — which some people would find inherently objectionable on grounds of privacy or fairness[258] — could not be expected to make correct predictions all of the time.

Even if the laws of economics and psychology had allowed it, perfect price discrimination might not be a desirable outcome. One *potential* disadvantage of market structures involving price discrimination is that those with high demand, who might have had the good more cheaply, are forced to pay more for it, thereby reducing consumer surplus.[259] Another is that some kinds of versioning will involve deliberate degradation of the good in order to negatively differentiate the cheaper version.[260]

In summary, price discrimination has the potential to reduce but not eliminate deadweight losses. In many instances, it has negative side-effects. Although a degree of price discrimination will more often than not be desirable in utilitarian terms, the net effect is not even guaranteed to be positive, let alone grounds for dismissing the problem of artificial scarcity.

---

[257]*Cf.* (Takeyama 1997) (modelling a similar effect with piracy as a form of price discrimination)

[258]It is interesting to note, with respect to these kind of responses to price discrimination, that Apple was willing to enter a substantial dispute with record labels by refusing to allow them to price discriminate through iTunes (Borland 2005).

[259]This is an old-fashioned example of the dependence of social welfare outcomes upon the preexisting endowment of wealth. One might believe that dollars in the hands of consumers will produce more benefit than dollars in the hands of record industry executives and shareholders; alternatively, one might believe that dollars in the hands of artists will do more good than dollars in the hands of the public at large. Where the dollars are needed most, and who actually ends up with them, is a question well beyond the scope of this thesis.

[260]There are many examples of this: "special editions" of software that have some features removed; soft-cover books (which are not much cheaper to make than hardbacks); cassette tapes (that actually cost *more* to manufacture than CDs, but sell for less); self-destructing DVDs (Taub 2003); inter-temporal discounts that reduce the value of the good by making it available later.

6.3 Estimating deadweight loss under Copyright

Deadweight losses, and the fact that copyright necessarily creates them, are clearly the largest reason to investigate alternatives to copyright. But how large is large? The answer depends to a significant extent on the kind of copyright system. The greatest losses are to be expected in the strong DRM-based regime, wherein piracy is only a marginal activity. Under the status quo, one must account for the fact that not all copying is actually prevented by regulation. Works that are pirated do not count towards deadweight loss. This section provides some approximate but illustrative calculations of the scope of distributional inefficiency under both regimes.

These calculations are difficult because they involve the measurement of hypothetical circumstances. The optimal amount of copying, against which we must measure, occurs in environments where copying is completely unrestricted by the law or by its indirect moral, technical, or market-rearranging modes of regulation. But no such environments are available for study and comparison.[261] It is impossible to know how many socially valuable copies are not made because of the law. Instead, we must be content with beginning our estimates with lower bounds: the volume of works that are pirated in observable situations. Beyond this, we must add informed guesses at the amount of additional copying that would follow from complete deregulation.

6.3.1 Deadweight loss under strong DRM compared to the status quo

Under the status quo, artificial scarcity costs are mitigated by piracy. In other words, the law theoretically prohibits many copies of songs which in actuality are made and enjoyed. The strong DRM regime was defined by the assumption that the great majority of unauthorised reproduction is prevented (with any residues being economically insignificant). Following the transition to such a regime, most of the copies of works which are presently made illegally would no longer be made at all. A smaller fraction would be replaced by licensed sales.

In order to evaluate the scale of these changes, I will consider two estimates of the value of pirated musical works. The first is based upon detailed surveys of American

---

[261]Even though Canada has had moments where users' private copying rights appeared to extend to large scale file sharing, these situations have been transient, unclear, and disputed. And even though a handful of states, such as Vanuatu, have no copyright regulation whatsoever, it is difficult to extrapolate from them to developed countries or even better-resourced developing states.

college students undertaken by Rob and Waldfogel (2004, 2006). The other is an original analysis based upon data variously published and commissioned by ARIA.

Music industry estimates from Rob and Waldfogel's data

Rob and Waldfogel's research involved administering surveys to around 500 students at a small range of U.S. colleges. These students were asked to identify popular albums[262] that they had purchased or downloaded during the preceding year. Most of these surveys then collected a post-purchase willingness to pay figure for each album.[263] A small number (N=92) were asked both *ex post* and *ex ante* valuations for each album. This last dataset is of the greatest interest; it gives a clear illustration of the effects of imperfect information in music markets and it allows for an informed prediction of what downloading consumers would do in the absence of the option to pirate.

From this data, Rob and Waldfogel estimate that the unauthorised reproduction[264] of popular albums during the period 1999–2003 has increased consumer surplus in their sample by USD $70 per capita. Of this, $45 is avoided deadweight loss, while $25 is wealth transferred from the music industry (per capita expenditure on popular albums is estimated to have fallen from $126 to $101 due to piracy).

A few points must be borne in mind when interpreting Rob and Waldfogel's results as a measure of the deadweight loss that would occur when shifting from the status quo to information feudalism. Firstly, the albums included in this data are only those that sold 2 million copies during the period in question, and so the real total transfers and deadweight losses occurring within their sampled population would be higher if the dataset included all recorded music. It is difficult to say how much higher, but we can observe that, by volume, around three-quarters of RIAA-reported albums sales are certified — that is, are sales of albums that had reported U.S. sales of 500,000 or more copies (Rob and Waldfogel 2004, p. 16). Many other albums were sold by smaller

---

[262]The criterion for "popularity" was that the album had sold more than 2 million copies in the United States between 1999 and 2003

[263]A few of the students were asked how much they would be willing to accept as compensation for losing the album and foregoing any opportunity to hear the music again. These numbers were very high, with median values of USD $100. These extreme values are likely to be, at least in part, artefacts of the strange and impractical proposition being offered to the subjects.

[264]Although it is not clear from the working paper, the survey would have led respondents to include burnt CDs (questions for (Rob and Waldfogel 2004), obtained by personal communication and on file with the author). Participants were ask to chose between "bought", "gift" and "download without pay, or sharing".

independent labels or even by bands who produced their own CDs. So total, real 1999–2003 deadweight losses per capita for this sample are at least USD \$60,[265] and more probably over the \$90 mark by the time that independent albums, and albums selling between 500,000 and 2 million copies are also included. Correspondingly, transfers from the music industry are at least USD \$33,[266] and probably more than \$50.

We should, however, expect Rob and Waldfogel's college student samples to overstate these numbers for the population in general; for the US in general that deadweight loss number is probably in the \$40-\$55 range,[267] while lost sales are probably in the \$20–30 range.[268]

An Australian music industry estimate from ARIA data

The second method for estimating DWL is one I constructed in 2004 from several observations about the Australian music market (Eckersley 2004a). It is perhaps less rigorous than Rob and Waldfogel's approach, but is derived using distinct datasets and methodology and therefore useful as a comparison. The input observations are the retail price of CDs and MP3 downloads, and Australian Record Industry Association (ARIA) data which allows the derivation of ratios of pirated music files to licensed music sales. The price numbers place an upper bound on financial deadweight loss per person per track/album, and using an assumption that demand curves are linear below the available sale price for music,[269] allow an estimate of average financial deadweight loss per person per album. The estimate is as follows:

$$D_{\maltese} - D_{\copyright} \; = \; P_{\copyright} \cdot \bar{v} \cdot R_D \tag{6.1}$$

Where $D_{\maltese}$ is the deadweight loss in the feudalist regime; $D_{\copyright}$ is the deadweight loss under the status quo; $P_{\copyright}$ is the number of pirated copies in the status quo regime;

---

[265]\$45 $\times \frac{4}{3}$ = \$60, assuming that the uncertified albums have the same proportion of deadweight loss as the "popular" ones.

[266]\$25 $\times \frac{4}{3}$ = \$33.3

[267]PEW Internet and American Life surveys in the United States report that music downloading in general is 1.6 times higher in the 18–29 age bracket than it is within the sampled population as a whole (Rainie and Madden 2005, p. 7). The Australian ARIA survey, as a contrast, reports that the incidence of file sharing is almost 2.4 times higher amongst those aged 10–25 than it is across the population in general (Quantum Market Research 2003, p. 4). These estimates create a range: 90 ÷ 1.6 = 56.25; 90 ÷ 2.36 = 38

[268]50 ÷ 1.6 = 31.25; 50 ÷ 2.36 = 21.19.

[269]A less technical way of stating this assumption is that if an album sells for \$15, there will be an equal number of people who value it at \$11, \$4.50, \$13.89, \$1 and any other price below \$15.

$\bar{v}$ is the mean value of these copies to the people who receive them; and R_D is the proportion of the currently unauthorised copies that would not be replaced by sales in the feudalist regime.

Unfortunately, the ARIA data alone does not allow a complete quantification of the cost of artificial scarcity, because it does not allow an independent derivation of R_D.[270] However we can use an estimate of that quantity from Rob and Waldfogel's data or other studies. Even though this creates some cross-dependence on that work, the ARIA estimate nonetheless allows us to check the other parts of Rob and Waldfogel's numbers.

Let us begin with $\bar{v}$. CD and download prices place an upper bound on deadweight loss per album because in a simple account, financial deadweight loss cannot be higher than album prices — otherwise, consumers would choose to purchase the album. For most CDs this bound was in the AUD \$10–30 range.[271] Licensed MP3 downloads appeared shortly after that period, with an upper bound on deadweight loss of around AUD \$1.50 per track.

These upper bounds do not in and of themselves tell us the value of each copy that would be lost in the transition from the status quo to feudalism. But an assumption of linear demand functions provides a reasonable estimate. Under this assumption, the average value of these copies is 50% of the upper bound, giving us a value of $\bar{v} \simeq 0.75$ AUD per track.[272]

ARIA's detailed surveys on piracy allow us to obtain a quantity for $P_{\copyright}$. Survey data indicated that in 2003, approximately 770 million songs were reproduced and distributed without authorisation from copyright holders; 430 million over file sharing networks, and 340 million with CD burners.[273] This gives us $P_{\copyright} = 770,000,000,$

---

[270]Along similar lines, Watt emphasises the lack of data on demand curves as the primary hurdle for the development of economic understandings of copyright (Watt 2004, § 3.1).

[271]Anecdotal observation of the Australian market during the 2003 era would suggest that a mass-market album sold in the range of AUD \$20–\$30; discounted older or fringe CDs sell for AUD \$10–\$20, while a bootlegged or "unauthorised" CD, or a CD single, can sell for as little as \$5. The ratios are very similar in the United States, with the marginal cost of a CD falling in the range \$3-\$5, the maximum price at \$18, \$10–\$13 being typical of "cheap" discounted prices, and occasional heavy discounts to the \$6-\$10 range.

[272]The estimate is similar for downloaded tracks, which are priced at around \$1.50, and a CD album of 13 tracks priced at \$20

[273]These figures come from a study commissioned by ARIA; see (Quantum Market Research 2003). 11% usage amongst 16.6 million Australians aged 10+: 1.83 million downloaders × 19.6 average files per month × 12 months = 430 million songs downloaded. 40% of 16.6 million people received copied CDs: 6.64 million CD recipients × 3.9 average CDs per year × 13 songs per CD = 337 million songs on copied CDs.

although because this was a telephone survey of piracy, these numbers may be under-reported. For comparison, in 2003 — a record-breaking year — Australians purchased about 677 million licensed songs;[274] so in that market in that year, piracy accounted for just over 50% of music distribution.[275]

As noted above, this piratical distribution would not have occurred under the strong DRM regime. Most of it would have not happened at all, but a smaller proportion would have been replaced by sales. The ARIA data does not allow an estimate of those proportions, so I will borrow this single quantity from Rob and Waldfogel's study. They observe that by volume, 20% of pirated copies replaced a sale, while 80% replaced deadweight loss.[276] In other words, $R_D = 0.8$

So now, we are ready to substitute into Equation 6.1:

$$D_{\text{☠}} - D_{\text{©}} = P_{\text{©}} \cdot \bar{v} \cdot R_D \tag{6.1}$$

$$= 770,000,000 \times \$0.75 \times 0.8 \tag{6.2}$$

$$= \$460,000,000 \tag{6.3}$$

This indicates that, if copyright had been more-or-less completely enforced in Australia in 2003, social welfare would have suffered an artificial scarcity cost of close to half a billion dollars that year. Per GDP, the equivalent cost in the United States would be a little over 7 billion US dollars per year.[277] Per capita, these numbers are around

---

[274]From ARIA's 2003 sales figures (ARIA 2004): 50.6 million albums × 13 songs per CD + 9.5 million singles × 2 songs per CD = 677 million songs sold on CD. There were no licensed music download services operating in Australia at the time. The Quantum report itself states that the unauthorised copies it measured accounted for only 10.7% of the music "acquired" by Australians (Quantum Market Research 2003, p. 4); this figure would be consistent with ARIA sales data only if sources such as listening to radio broadcasts and watching video clips were included.

[275]53%, according to these approximate figures.

[276]Rob and Waldfogel are not the only authors to estimate this quantity, although their methodology is by far the most persuasive. For comparison, see (Oberholzer and Strumpf 2004) (concluding that piracy may not cause *any* lost sales, and may even increase them); (Boorstin 2004); but see (Liebowitz 2004) (surveying and critiquing empirical studies on the effect of file sharing on music sales, including those by Oberholzer & Strumpf and Boorstin, and concluding that at this stage, we should still expect that file sharing is at least partly responsible for the recent decline in U.S. record sales); (Peitz and Waelbroeck 2004) (estimating that file sharing had been responsible for a 20% decline in music sales between 1998 and 2002; because the number of pirated works and the number of sales is similar, this is essentially agrees with Rob and Waldfogel's figure); (Michel 2006) (estimating the switch to be only 4%); (Waldfogel 2010) (an updated version of the earlier survey by Rob and Waldfogel, finding $R_D \in [0.7, 0.85]$).

[277]According to the CIA World Factbook, the U.S. GDP is USD $10.98 trillion, while Australia's is

USD $25 per year, which is substantially lower than the range obtained from Rob and Waldfogel's estimate ($40–55). One possible explanation for this discrepancy is the fact that the ARIA telephone survey might have underreported levels of piracy compared to Rob and Waldfogel's album-by-album examination of their subjects' music collections. Aside from the fact that some people might deny a potentially illegal act when asked about it in general terms by a stranger administering a telephone survey, the ARIA survey questions only addressed CD burning and P2P file sharing, and not the direct copying of entire music collections between subjects' computers. For this reason I believe the higher estimates are more likely to be correct, but will conservatively proceed using a value of the deadweight loss under strong copyright in the music industry of $D_{\text{☹}} - D_{\text{©}} \in [\$25, \$55]$ per person per year, or USD 7–15.5 billion per year (within the United States). So stopping music piracy would increase artificial scarcity costs by between 60% and 130% of the total value of sales in the U.S. music industry,[278] which would be an 18–32% reduction in the current social benefit of music.[279]

So at least in the music industry, the social surplus that is currently derived from piracy is of approximately the same size as the financial benefit that accrues to all of the participants in that industry. A transition to a fully-enforced copyright regime would largely eliminate the social surplus from piracy.

The film and book industries

For other types of copyright works, the numbers may be different. In a second paper, Rob and Waldfogel (2007) collected survey data on the film industry, which I will use in conjunction with MPAA and other video sales data for a third estimate of artificial scarcity costs. That study found that a much larger proportion of film piracy replaced sales, with $R_D = 0.38$.[280] We can use this as an input to an estimate similar to that

---

USD $570 billion (CIA 2004).

[278] See (IFPI 2003), reporting world music sales of $32 billion and a U.S. share of 37%.

[279] Converting from a fraction of sales revenue to a fraction of social value requires an estimate of consumer surplus. The relationship between sales and consumer surplus in the music industry is estimated in Waldfogel's work; he finds (Waldfogel 2010, Table 6) that total consumer value from *purchased* songs is 179% of sales revenue ($\frac{11.59}{6.48}$). The value of pirated works, expressed as a proportion x of sales is therefore a fraction $\frac{x}{x+1+1.79}$ of social value; $\frac{0.6}{0.6+1+1.79} = 0.18$; $\frac{1.3}{1.3+1+1.79} = 0.33$. 60–130% of total sales is therefore 18–32% of the total present social value of the music industry.

[280] The results in (Rob and Waldfogel 2007) contain a very significant distinction between copies that are made for first and second viewings. They conclude that approximately all first-viewing piracy replaces a sale, while only around 0.20 sales are lost if a film is pirated for the purpose of watching it a

in Equation 6.1, though it needs to be split between subsectors (cinema tickets, DVD sales, and various forms of rentals):

$$D_{\text{☃}} - D_{\text{©}} = \sum_{\text{subsector}} P_{\text{©}} \cdot \bar{v} \cdot R_D \tag{6.4}$$

$P_{\text{©}}$ is observed to be 5.2% of the number of film-viewings reported in the survey (much smaller than the 50%+ observed for music). I will assume that this figure is constant across those who might have watched a film in a cinema, those who might have watched it on DVD, and those who might have rented it. Using MPAA per-capita U.S. cinema ticket sales data, we can estimate that in the cinematic subsector, $P_{\text{©}} = 68$ million and $\bar{v} = \$3.75$.[281] In the DVD sales subsector, $P_{\text{©}} = 48$ million and $\bar{v} = \$8.63$.[282]

Similar calculations can be performed in the video rentals subsectors though public data is only available for 2004, not 2005.[283] For DVD rentals $P_{\text{©}} = 91$ million and $\bar{v} = \$1.6$.[284] For VHS rentals $P_{\text{©}} = 43$ million and $\bar{v} = \$1.6$.[285] We can substitute these values into Equation 6.4

$$\begin{aligned} D_{\text{☃}} - D_{\text{©}} &= \sum P_{\text{©}} \cdot \bar{v} \cdot R_D \\ &= 68,000,000 \times \$3.75 \times 0.38 \\ &\quad + 48,000,000 \times \$8.63 \times 0.38 \\ &\quad + 91,000,000 \times \$1.64 \times 0.38 \\ &\quad + 43,000,000 \times \$1.39 \times 0.38 \\ &= \$330,000,000 \end{aligned} \tag{6.5}$$

---

second time. Across all scenarios, Rob and Waldfogel observe 5.2% of film-viewings being of pirated copies and 3.2% lost sales, allowing us to deduce $R_D = 1 - \frac{3.5}{5.2} = 0.38$.

[281] See (Motion Picture Association of America 2009) reports 4.4 tickets per person per year at an average of USD 7.5. So we can estimate $P_{\text{©}} = 0.052 \times 4.4 \times 296,000,000 = 68,000,000$ and $\bar{v} = 7.5 \div 2 = 3.75$.

[282] According to (Belson 2006), DVD sales were approximately 16 billion, at a mean price of $17.26 (Belson 2006). VHS sales were a negligible proportion of this. For DVDs, $P_{\text{©}} = 0.052 \times 16 \div 17.26 = 48,000,000$, and $\bar{v} = 8.63$.

[283] 1.75 billion DVD rentals and 842 million VHS rentals; US $8.1 billion total rental revenue, 71% of which was from DVDs; see `http://www.entertainmentscene.com/video_industry_facts.htm`.

[284] $P_{\text{©}} = 0.052 \times 1,750,000,000 = 91,000,000$ and $\bar{v} = \frac{8,100,000,000 \times 0.71}{1,750,000,000} \div 2 = 1.64$.

[285] $P_{\text{©}} = 0.052 \times 842,000,000 = 43,000,000$ and $\bar{v} = \frac{8,100,000,000 \times (1-0.71)}{842,000,000} \div 2 = 1.39$.

So during the 2003–2005 period, the deadweight loss difference between the status quo and strong DRM was *much* lower (around 20–30 times smaller!) for films than it was for music. This reflects both the greater time and inconvenience required to download a film, and the larger time commitment required for watching it — if Alice is willing to spend two hours watching a film, perhaps the asking price is less likely to be a prohibitive factor for her.[286]

It also illustrates something else quite interesting. Film piracy was not, during the 2003–2005 period, of particularly large economic significance! In fact, it would be fair to say that the 'status quo' for that industry was in fact *strong* copyright, not weak copyright.

In the book industry, during the comparable period, artificial scarcity costs were not yet reduced by piracy to any great extent either. As of 2010, the maturation of homebrew book scanning equipment and the growing adoption of e-book readers and other devices suitable for extended reading means this may have begun to change. Unfortunately there is not yet sufficient data on the phenomenon to allow a quantitative analysis.

6.3.2   Status quo deadweight loss

The previous section discussed two measures of the *extra* deadweight loss that would follow a change from the prevailing weak DRM regime to a strong DRM regime in the music industry. In other words, they are measures of the extra loss that would occur if piracy suddenly stopped, but not measurements of the amount of deadweight loss which is occurring at present. That number — deadweight loss under the status quo compared to information anarchy or an alternative compensation system — may be much larger, but can only be measured under counterfactual circumstances. That is to say, it could be observed if we could see into a world where copying technologies did not suffer from the same inhibited development; where legal and/or moral sanctions were not applied to copying; where iPods had pairwise copying features[287] and P2P

---

[286]Compare (Rob and Waldfogel 2007, 393–394).

[287]As of 2005, a scan of a site like `http://www.ipodhacks.com` indicated that it was possible to find ways to copy songs between iPods by using a PC as an intermediary, although Apple do their best to make this difficult. Short of ludicrous hacks such as "podshaking" — re-recording the audio signal from another iPod — (`http://www.themodgods.com/2005/03/podshanking-physical-pod-on-pod.htm`) or installing Linux (`http://ipodlinux.org`) on these gadgets, there was no commonly known way to get them to share files through a direct USB or Firewire connection.

networks were not deliberately flooded with fake files;[288] where those networks were not closed by litigation or driven underground, to be tended by spyware purveyors; where recommendation engines and indexes like Pandora, Last.fm, IMDB, Library-Thing, and Netflix's streaming service could supply people with media directly and without restraint, and where social networks could allow people to share works seamlessly with their friends.

Although this counterfactual measurement of deadweight loss under the status quo is nearly impossible, there is some information which I will use to estimate its scale in the case of music. The resulting number is the most uncertain of any in this chapter, because it is specifically the extent to which copying would increase if the technical and social development of file sharing systems had been allowed to proceed without regulatory prohibitions of any sort.[289]

To get some loose grip on all these effects, we could start with the peak and more recently reported rates of music file sharing; in the U.S. these are somewhere between 24%, and 28%, respectively, for uploading (Rainie and Madden 2005, q. WEB1) and about 11% and 32% (Rainie and Madden 2005, qs. WEB1 & DLVD5) for downloading; in Australia, they are somewhere between 11% and 22% for P2P sharing and 22% for CD burning.[290] To find the ratio between current and first-best quantities of distribution, take those figures and extrapolate them to the entire satisfactorily-equipped and musically-inclined population (a factor of 3–4 or more) and decrease them because music enthusiasts are likely to have self-selected as file sharers; increase them because music enthusiasts are also more likely to be abstaining from piracy on moral grounds; increase them because, if file sharing was easier, many of those who presently share files would share *more* files. A factor of three seems conservatively realistic.

That would leave status quo deadweight losses from artificial scarcity in the range of 1.5–3 billion dollars per year, for the music industry, in Australia. In the U.S., the number is around 20–40 billion dollars per year. These figures amount to around USD $60–120 per person per year, or 180–390% of the cash value of sales in the music industry! As a proportion of the total value of the music industry (sales plus consumer

---

[288]And even if hackers attempted to do so, it would be easy to build hash catalogues to foil their efforts.

[289]By way of contrast, it is interesting to note that Rob and Waldfogel are able to ignore this relatively troublesome quantity, because they consider only two states: the status quo downloading regime, and one in which there is no downloading (Rob and Waldfogel 2004, p.14). They do not consider the first-best regime in which consumers obtain all of the music valued above the minimum possible (expectation) costs of finding and downloading.

[290]See (Quantum Market Research 2003, p. 4). The peak number is probably less than 22%, because that figure is the percentage who reported *ever* having used a P2P network.

surplus from sales plus consumer surplus from current piracy), this deadweight loss is 55–98%.[291]

I believe that levels of film file sharing remain too low for a meaningful estimate of the parallel percentage in the film industry. Deadweight loss is likely to be a lower proportion of the value of the industry than is the case for music, but not the 20–30 times lower that was observed for $D_{♚} - D_{©}$.

In summary, with respect to artificial scarcity there is an ordering of regimes with anarchy and alternative compensation systems at the top; a significant drop to the status quo, and a smaller drop to the strong DRM regime. In the music industry these two drops are very large (180–390% and 60–130% of sales, respectively). In the film industry the drops are much smaller.

6.4   The Compounding Effect of Inequality

The estimates in Section 6.3 of the deadweight loss prescribed by copyright are quite substantial. But there is a sense in which deadweight loss as calculated by the usual financial measure — dollars that people were willing to pay for copyrighted works — understates the utilitarian scope of the problem. The fewer dollars a person has, the less money they will be willing to pay for copyright works, regardless of the amount of joy or wisdom that those works would actually convey. Suppose that a teenager with an annual income of $500 is willing to pay up to 50¢ for a particular song, and a professionally employed adult is also willing to pay 50¢ for it. Suppose further that the minimum asking price is 89¢. We should place greater normative significance on the exclusion of the teenager, because the same amount is a far larger proportion of her income. Under a utilitarian social welfare function, the teenager has a larger marginal utility of dollars, and a 50¢ deadweight loss is therefore of greater normative importance.

In societies with significant inequalities of wealth, we should expect a disproportionate number of the cases of exclusion to be exclusion of poorer individuals (and correspondingly, a disproportionate amount of consumer surplus at the other end of the demand curve to be accumulated by wealthier individuals). Call this sate of affairs X.

In a different, more equal, society O in which the marginal utility of dollars was

---

[291]Thrice the proportions identified in footnote 279, in fact.

the same for all individuals, the utilitarian and financial measures of deadweight loss would be equivalent to each other. But in X, the marginal utility of dollars is not constant across individuals, and the two measures will disagree, with the financial measure of deadweight loss underestimating the true utilitarian inefficiency.

One way to make this underestimation mathematically precise is to pick a fair dollar-based unit with which to measure utility, such as 'the marginal utility of a dollar to the median income earner'. In real (unequal) societies, utilitarian deadweight loss in dollars-for-median-income-earners is larger than financial deadweight loss in dollars.
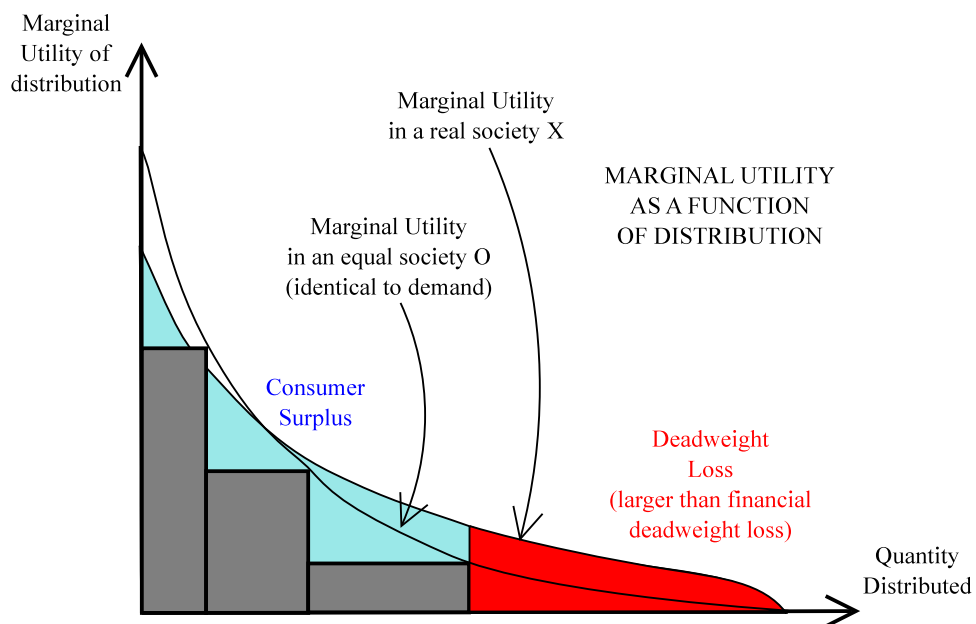
Figure 6.3: Utilitarian Deadweight Loss

The situation is illustrated in the contrast between Figure 6.2, showing financial deadweight loss in some market, and Figure 6.3, which illustrates utilitarian deadweight loss in the same market.

Exemplary instances are abundant: there is little direct incentive for a textbook publisher to distribute cost-price textbooks to those who cannot otherwise afford them, little incentive for record companies to give teenagers with minimal pocket money free or cheap access to all of the music they would like, and (embarrassment aside) little incentive for pharmaceutical firms to allow cost-price distribution of drugs to impoverished nations.

Ideally, if a welfare system was acting to redistribute wealth to those "devalued" by

the demand curve, then price discrimination could be expected to reasonably constrain deadweight loss. But given that the reach of the welfare state is limited, we are faced with a situation where the use of exclusive rights to finance information production turns out to exacerbate preexisting inequalities.[292]

What this means is that the numerical estimates of the distributional advantages of anarchy or virtual markets over the status quo, and in turn of the status quo over strong copyright, are too low. They would need to be adjusted upwards if they were going to be measured in utilitarian units such as the "value of the dollar to a typical middle class person".

6.5   Summary

This chapter estimated the artificial scarcity costs associated with digital copyright. In the music industry, these inefficiencies are presently in the range of 180–390% of the total sales in that industry, or 55–99% of the social surplus value of the industry. Moreover, if the music industry found a way to prevent piracy, artificial scarcity costs would rise by a further 60–130% of sales, or 18–32% of total surplus. Because of inequality, these proportions would be even higher if expressed in utilitarian rather than dollar-value terms.

These results can be placed in the following slogan form: *If piracy in the music industry stopped today, the enjoyment produced by a typical piece of music would fall by about a third. If music sharing were completely legalised, the value and enjoyment of a typical piece of music would double.*

It was also shown that artificial scarcity costs vary tremendously between copyright industries. While the cost of moving from the status quo to a strong DRM regime for music was estimated to be enormous ($7–15 billion per year in the U.S.), that was not the case for the film industry (just $330 million per year). This may reflect the fact that prevailing circumstances in the film industry are much closer to a feudalistic regime, but it may also be a result of the differences between the way that people consume these two kinds of works. We also cannot yet estimate how large status quo deadweight losses are in the film industry, because we simply cannot guess how

---

[292]For example, when the United States employed the threat of trade sanctions to persuade developing nations to enact strong copyright and patent laws, and ultimately to sign TRIPS (Braithwaite and Drahos 2000, Chapter 7), it did not consider combining this with aid packages to correct the billions of dollars of inequitable wealth redistribution that TRIPS might cause.

many movies people would watch if all of classic cinema were free and conveniently available. For the time being, there aren't enough film pirates whose behaviour can be observed. Book piracy is also presently too rare an activity to be used for measuring the characteristics of demand curves — though this may change as e-book readers mature.

The possibility of eliminating deadweight loss is the most obvious argument for considering alternatives to digital copyright, and the most compelling. An important question to be addressed in the following chapters is whether exclusive rights have powerful arguments of their own that might counterbalance the burden of artificial scarcity.

Chapter 7

Paying the Piper: Information and the Incentives for Cultural Production

Contents

Most people would agree that the legitimate purpose of a copyright system is to either incentivise or reward cultural producers for their contributions to society. I shall assume this to be the case.[293] Either way, a believable normative comparison of different copyright regimes requires a comparison of the rewards they offer to information producers. This chapter provides one.

What is society getting from copyright law? In utilitarian accounts, the incentive function of each copyright regime should be judged by the quality and quantity of works produced as a result of it. It is possible that the utilitarian perspective on incentives should be adjusted or combined with some other ethical view of copyright's purpose, but such endeavours will have to wait for some other thesis.[294] Making judgements about the incentive benefit of copyright is already exceptionally difficult — more difficult than any of the analyses in the preceding or following chapters. To see why the task is so hard, consider the following:

- First, it is not inherently obvious how we should measure the quality of cultural works — but some measure is necessary, because rewards that increase with quality are desirable. The standard economic reply — "The value of works is what people are willing to pay for them" — is just one answer among many, and it is unsatisfactory. This matter of ethics and aesthetics is discussed in Section 7.2.

- Second, the fact is that many authors and artists are motivated by considerations other than financial reward. Existing methodologies cannot tell us whether any particular work would have been produced without copyright, or even what frac-

---

[293]It is not obvious that continental European moral rights theories of copyright, cast in terms of "personality rights" or "labour theories", could fall into this category. I would simply observe that those theories are hard to defend (Fisher 2000b, Parts B & C) unless they emulate the essential tradeoffs present in more instrumentalist or consequentialist theories. Compare (Gaita and Christie 2004; Christie 2004), discussing how civil law regimes invented private copying levies to protect users' rights to privacy, despite the fact that these schemes seem to constrain the rights of authors and would therefore superficially make more sense under a utilitarian theory of copyright.

[294]Timothy Brennan, for instance, has argued (Brennan 2004) that the purpose of copyright should be to promote "fairness" rather than the common good. This is an intriguing suggestion, although Brennan does not work with a compelling theory of what fairness actually is. There is a complicated literature on that point, and an analysis following from Rawls' theory of "justice as fairness" (Rawls 1971) might reach quite different conclusions about the role of copyright to one grounded in humans' everyday intuitions on the subject (Konow 1996).

tion of works would have been produced.[295] Nor can existing methodologies tell us how many authors and artists are changing the nature of the works they create in order to increase their expected incomes.

These limitations to our knowledge make it hard to interpret results about the accuracy of incentives. Suppose we could show that market incentives were ten percent more 'accurate' than a levy system in matching audience appreciation. We still could not say how that compared, for example, to a system which provided no incentives whatsoever. This difficult problem is discussed and modelled in Section 7.3.

Only once the preliminaries listed above have been navigated is it possible to confront the core problem in comparing the incentives produced by the four copyright regimes: how well do they offer incentives that are proportional to the social value of works produced? This is itself a subtle and intricate question, which I will address at length in Sections 7.4 and 7.5. Section 7.4 reviews some high-level models in the literature that answer the question very abstractly. Section 7.5 considers a list of low-level effects that modify the high-level picture and may in practice make an important difference in the incentives created by the different regimes. Section 7.6 ties all of this together.

7.1   Information, Incentives, and Calculation

It is easy to provide an incentive to do something.[296] One can announce an offer of payment contingent upon some act, or promise to bake scones for whoever does it, or apply whatever other means are at one's disposal. But at a systemic level, it is impossible to bake scones for the right people at the right times unless one knows what incentives *ought* to be offered. This problem of information collection was first emphasised during mid-twentieth century debates about the feasibility of centrally planned socialist economies. As Hayek put it,

> *Which of these systems* [a market or central planning] *is likely to be more efficient depends mainly on the question under which of them we can ex-*

---

[295]This is a very old point; see for example (Plant 1934, 167–168,191).

[296]For the purposes of this chapter, I will define an *incentive* simply to be a thing that encourages people to do something. This is different, for example, to Zimmerman's narrower definition of "a reward without which the desired level of an activity would not occur" (Zimmerman 2003, note 57) — which could less confusingly be called a 'necessary incentive'.

pect that fuller use will be made of the existing knowledge. And this, in turn, depends on whether we are more likely to succeed in putting at the disposal of a single central authority all the knowledge which ought to be used but which is initially dispersed among many different individuals, or in conveying to the individuals such additional knowledge as they need in order to enable them to fit their plans with those of others. (Hayek 1945, p. 521)

The remarkable function of markets, noticed by Adam Smith and later captured formally by Arrow and Debreu,[297] is that they create incentives for the efficient production and provision of goods and services, in accordance with the information revealed about people's wants and needs by the price system and the accompanying negotiations of trade. Hayek's point was to emphasise the flows of information in that picture. A central planning authority cannot search for a good economy-wide production strategy unless it knows what each worker and each firm is capable of doing in the future. The reason central planning is likely to fail, and the reason Hayek is usually right, is incentives: without the right incentives, workers and firms will not even discover what tasks they are capable of, let alone disclose specifications of those tasks to a planning authority.

Later theorists studied these mechanisms much more closely, and concluded that it was only when producers and consumers in efficient marketplaces share a great deal of basic common knowledge that the relationship between price information and incentives is completely straightforward.[298]

It is interesting to note that the virtual market model of public funding for copyright works is in some ways reminiscent of the "market socialism" proposals which Hayek criticised.[299] In fact, one of the leading socialists in the calculation debate later realised the tremendous relevance of computers to their side of the argument (Lange 1967), and economists inspired by Hayek responded in kind.[300] It is clear that the marketists won the argument: there are no credible suggestions in the modern literature that societies could replace all of their marketplaces with central planning institutions

---

[297]See (Mas-Colell, Winston, and Green 1995, Chapter 16).

[298]See for example (Stiglitz 1994) for an overview of how information economics has built on top of Hayek's work, and modelled many of the situations where common knowledge is insufficient.

[299]This famous historical dispute in economics is now called the "socialist calculation debate". For a review of the debate and Hayek's contribution to it, see (Caldwell 1997).

[300]See, for example, Lavoie (1990). The reader is cautioned that the standards of the computer science-dependent arguments made by both sides of this debate were uniformly low.

without suffering dramatic losses of productivity.[301] Regardless of whether computers could crunch the numbers satisfactorily, persuading people to put all of the pertinent information about their productive capabilities into the system, and giving them the means to do so, seems fantastic.[302]

The similarity between the virtual market/exclusive rights comparison and the socialist calculation debate is deceptive. The reason is that virtual markets are proposed as a replacement for a small set of markets with highly distinctive characteristics, including the presence of public goods, and a simple dichotomy between producers and consumers.[303] The importance of the second point is that it makes it feasible for a central agency to obtain all of the information necessary to offer efficient market-like incentives for production. There are no complicated factors of production to consider: all that needs to be measured is the value of each good to consumers. There is, in short, very little to calculate. In fact, as we shall see in Sections 7.4.1 and 7.5, there are informational reasons why virtual markets could provide *better* incentives than a copyright-based marketplace.

7.2   What is 'Valuable' Cultural Production?

In order to count the value of the works produced as a result of copyright incentives, we will need a definition of that value. This is not an entirely straightforward problem.

---

[301] See, for example Temkin (1996). But see Stiglitz (1994) for a discussion of numerous informational problems in actually existing marketplaces, which provide potential grounds for efficiency-improving governmental interventions, though these remain far less drastic than central planning for an entire economy.

[302] It is clear that contemporary computers could manage the data structures necessary to represent an entire economy: after all, banks' systems already handle a large proportion of society's economic transactions; personal electronic devices could trivially handle those transactions we currently execute with cash. *If* those structures included all of the available production options (*i.e.* immediately available mappings from factors of production to new commodities), it would probably be possible to use some nifty pragmatic optimisation algorithm to provide good economic instructions. But the 'if' is fearsome; too much of the knowledge about what each person might be able to do with their own labour is completely private and continually changing. In the absence of the necessary data structures for optimisation, it should be presumed that it cannot be performed well. This is not to say that science fiction authors have not tried to imagine ways of doing it; I can point to one picturesque attempt (Sterling 1999), but it would probably require full-blown artificial intelligence. Less ambitious (and conceivably remotely feasible) conceptions might appear in the future.

[303] Of course, there are some copyright markets where this distinction does not apply; see Section 5.4. In those cases, the development of satisfactory policy regimes is made all the more difficult.

Partly it is about dealing with the discrepancies between prices and use-values[304] but it is also about deciding on the normative or social welfare significance of aesthetic qualities.

Most of the theoretical literature on the economics of copyright employs the simplest possible answer: the value of a work to society is the sum total that consumers are willing to pay for it (in cases where distribution is limited, it is only those consumers who actually receive the good whose valuations count).[305] This measure is convenient for calculation but it also suffers from certain limitations.

One stark problem with dollar value social welfare is that money is defined to be of equal significance regardless of whether it is held by wealthy or poor individuals. Most people would disagree. In the previous chapter, I discussed some of the evaluative changes that follow from correcting for non-uniform marginal utilities of cash (see Section 6.4).

Within this chapter, I will discuss several further adjustments that result from the fact that the works people desire are not necessarily the best ones. This may sound paternalistic, but that is not *always* the case. In Sections 7.5.2 and 7.5.5, I discuss two situations (arising from imperfect information) in which many consumers would agree that their willingness to pay was not proportional to the worth of a copyright work.[306]

There are also some genuinely 'paternalistic' or 'elitist' objections which can be brought against preference-satisfaction theories of the value of cultural works. Suppose, as a thought experiment, that people start off with a preference for trashy rock music, and that listening to nothing but jazz, operatic, and classical music bores them for a year or two, until they have acquired a taste for it. Suppose further that once they

---

[304]The best way to define a use-value for a copyright work would be to take the integral of the difference in utility for consumers over time as a result of its creation. Therefore, if there are a thousand 'Westerns' available at a particular enjoyability level, but none any better, and someone makes one (call it "El Dorado") that is 10% better. El Dorado would have a use-value that was much more than 10% higher than one of the others. The creation of future higher-quality westerns might slow the benefits to society stemming from El Dorado's existence (although in practice being the first film to introduce new ideas or cinematic techniques tends to result in some enduring value)

[305]For a defence of the use of price as an index of the true aesthetic value of art, see (Grampp 1989). For a characterisation of the implicit and unsatisfyingly restrictive assumptions that accompany social welfare functions formed by adding up satisfied willingness-to-pay, see (Mas-Colell, Winston, and Green 1995, Chapter 4). Another criterion that is used, more often in the literature on pure public goods than in the literature on copyright, is Pareto optimality. Under strict assumptions, states that maximise dollar value social welfare are necessarily Pareto optimal and also more strictly Kaldor-Hicks efficient. This section is largely a criticism of those assumptions.

[306]Where the wealth-related effects discussed in the previous chapter adjust the average value of copyright works in general, the effects discussed here tend to adjust the relative value of different works.

have acquired the taste, their lives are incomparably richer. In such endogenous prefer-
ence situations, the true value of hard rock music could conceivably be negative.[307] It
might be best to appoint a board of music connoisseurs to decide whether each record-
ing should receive a reward, or should instead be subject to a harm-minimisation tax.

In that somewhat humorous example, the connoisseurs were to be employed be-
cause they could *predict* the way that listeners' preferences would evolve if properly
nurtured. An even more elitist position might hold that there is such a thing as objec-
tive beauty, which is somehow linked to the fundamental truths of the universe. They
might hold that all other things are distractions from the quest for this beauty, but that
only a tiny number of people are wise enough to see that fact. According to such theo-
ries connoisseurs might be employed, not for their predictive power, but because their
taste is simply 'better'.

I have presented both of the above positions in exaggerated and parodic terms,
but they are in fact widely adhered to and have resulted in the creation of numerous
institutions around the world to support elitist art and culture. Governmental support
for the arts is, by and large, motivated by paternalistic theories of just this sort: the art
is either regarded as intrinsically worthy beyond or aside from its market valuation, or
its promotion is seen as a enriching peoples' tastes.

There are other positions which can perhaps reconcile some of these contradictory
views of the value of cultural goods. One might, for example, argue that 'art' and
'entertainment' are two distinct things (though some works may embody both). En-
tertainment might then be valued because it satisfies people, while art might be valued
because it challenges them, makes them think, or inspires them to develop as spiritual
beings.[308]

This not being a thesis on aesthetics, I will not travel any further with these sub-
jects. Pragmatically, I am going to avoid any satisfactory discussion of 'art' and deal
only with the economics of entertainment. The rest of this chapter works with the
assumption that satisfaction of the preferences of audiences is what copyright policy
should seek to maximise (with appropriate consideration of imperfect information and

---

[307]In fact, many people believe in all seriousness that heavy metal *is* bad for its audience; see for
example (Longhurst 1995, 195–202), discussing the "hypodermic syringe" and subsequent theories of
media impact. There is empirical evidence to support the proposition that taste for music is endogenous
(Holbrook and Schindler 1989), although it does not go so far as to show that some tastes are "better"
than others.

[308]Klamer (1996), for example, develops an argument along these lines against the use of prices as a
measure of the value of artistic works.

path dependence/endogenous demand). It should be noted that regardless of whether *entertainment* is funded by copyright or by alternative compensation systems, there may continue to be good arguments for other publicly (or charitably) funded institutions that support *art* for its own sake.

7.3   How Much Do Copyright Incentives Matter?

We hope that the presence of copyright (or an alternative) motivates authors to write, musicians to compose and record songs, and so forth. When cultural or informational projects require large amounts of funding for equipment or non-creative labour inputs, we hope that the prospect of future licensing revenue motivates investors to write cheques to pay those bills. This section discusses the degree to which those motivations are indispensable to society's stock of books, films, and sound recordings.

7.3.1   Intrinsic and non-monetary incentives

Question: *What do you call a musician without a girlfriend?*
Answer: *Homeless.*[309]

Monetary incentives are by no means the only reason why authors take up their pens. There are many benefits from engaging in creative labour, and beyond that, some people may write or sing or make art because it is in their nature to do so. In forming informed policy, it would be very helpful to know how much difference the addition of financial rewards actually makes. Unfortunately, there are many holes in the available evidence on this subject.[310]

There are plenty of possible sources of non-monetary motivation for cultural production. They include the inherent enjoyability of creativity, the prospect of fame and increased social status, increased self-confidence, and the fact that artistry can make people more romantically attractive. These effects may consciously affect people's decisions about authorship, or they may have evolved their way into human nature.

---

[309]My thanks to Emily Hudson for passing on this joke.

[310]Towse, for example, attempts to work from labour market surveys of artists to shed some light on the tradeoff between financial and non-financial motivations (Towse 2001, p. 159 & Chapter 3). She finds strong evidence that non-financial motivations make people professional singers, but the data is not sufficient to form quantitative conclusions about the economic importance of monetary incentives.

The strength of non-pecuniary incentives is well demonstrated by the massive over-supply of would-be entrants into the market for artists, authors, scriptwriters, film actors, and other glamorous creative roles.[311] This widely observed oversupply does not necessarily mean that more explicit incentives are unimportant, because the group of interest is not the large body of would-be artists, but the much smaller group who have the talent and commitment to produce particularly valuable works. Monetary incentives may still be important for bringing some members of this group into creative vocations, or for keeping up the effort levels of those who would be artists either way, or for raising capital for projects that require it.

Turning to the specific forms of intrinsic motivation, the connection between artistry and sexual attractiveness stands out as clear, amusing, and relevant, but it has not to my knowledge been discussed in the copyright literature at all. Charles Darwin first posited that music might have evolved as a means of advertising to prospective mates:

> *...many cock birds do not so much pursue the hen, as display their plumage, perform strange antics, and pour forth their song in her presence.* (Darwin 1882, Chapter 8)

> *...it appears probable that the progenitors of man, either the males or females or both sexes, before acquiring the power of expressing their mutual love in articulate languages, endeavoured to charm each other with musical notes and rhythm* (Darwin 1882, Chapter 19)

A great deal of evidence has accumulated to show that it plays precisely that role in musically inclined animals such as songbirds.[312] Until very recently, Darwin's theory about humans has received much less attention and research than it should have (Miller 2000). Anecdotally, many popular musicians report sex as having been a major motivation for their choice of career.[313] And, as Sean Connery's character quips in *Finding Forrester*[314] authors only go to book readings "because they want to get laid."

---

[311]See (Towse 2001, pp 58–61, 160) (observing that there are 500 trained classical singers for every one who works regularly); (Menger 1999).

[312]See Slater (2002) for an entry to that literature.

[313]Gene Simmons, lead singer of *Kiss* and one of the most infamous examples, wrote "We all picked up guitars because we all wanted to get laid. Plain and simple." (Simmons 2001, p. 39). Google quickly provides a supply of performing artists who acknowledge sex as a primary or significant motivation for their work: Billy Joel, Sinéad O'Connor, Ray Manzarek of *The Doors*. Luciano Pavarotti never confessed.

[314]`http://www.script-o-rama.com/movie_scripts/f/finding-forrester-script-transcript.html`

How important are these animalistic motivations? Some limited attempts have been made to measure whether artistic and musical performance are in fact correlated with the search for a partner. At this point, the literature seems to regard the theory as a plausible evolutionary explanation of the development of human music, but initial results also suggest that it is unlikely to be *the* dominant cause for performance in contemporary society.[315]

However important specifically romantic motivations may be, the operation of non-financial incentives in general is clear when one examines data on artists' labour markets. Typical results are that 18% of serious artists in a particular city had received no payments whatsoever for their work in the preceding three years; or that 20% of full-time male artists and 40% of full-time female artists depend on their partners for financial support.[316] And even in the great majority of those cases where artists do make a living from their vocation, they would be earning more if they were doing something else.

The analogous problem of understanding and measuring non-financial motivation also exists in the case of free/open source software production.[317] The first theories proposed by academics were that most free software authors work according to their own intrinsic motivations[318] or that they were doing it to make themselves more employable.[319] Survey data has since shown that the situation is more complicated. There appear to be four fairly distinct clusters of motivations amongst contributors to free software projects;[320] the free software movement has succeeded by combining them.

So there is clear evidence of non-financial motivations amongst cultural producers, and evidence that roughly half of the workers in part of another copyright marketplace (open source software) are not doing so for personal gain.[321] In this light, I believe the

---

[315](Miller 2000; Benzon 2001)

[316]See Towse (2001, pp 64–65), surveying the results of artists' labour market surveys.

[317]Although it might be added that fewer people are likely to be writing free software in order to make themselves more attractive to members of the opposite sex.

[318]As Eric Raymond put it, "every good work of software starts by scratching a developer's personal itch" (Raymond 1997).

[319]This position was taken by Lerner and Tirole (2000) and modelled more elaborately by others (Lee, Moisa, and Weiss 2003). See also Johnson (2002) and Mustonen (2003).

[320]See (Lakhani and Wolf 2003, p. 14). The four clusters they identify are (1) 29% who enjoy the intellectual stimulation of coding and wish to improve their skills; (2) 27% who want software for their own personal, non-work related purposes; (3) 25% who use the code for their work and are mostly paid to write it; (4) 19% who have community-oriented, reciprocal-obligation or ideological motivations for writing free software.

[321]See *Id.* I would count intellectual stimulation and community participation (groups 1 and 4 in Lakhani and Wolf's study) as not being materially gainful and therefore being comparable to volunteer

default expectation should be that at least several tens of percent of the output value of cultural copyright industries would exist independently of financial incentives.

These observations about non-financial motivation are important because, to the extent that they are causally significant, they weaken any rationale for copyright that involves the free rider problem. We will see how much, in a mathematical sense, in Sections 7.3.4 and 7.3.5.

Certainly, some information anarchists base their denials of the necessity of exclusive rights on the fact that cultural production is something that people *just do*.[322] Even if we only accept these arguments in part, they may still be highly relevant to determinations of the appropriate scope of copyright monopolies (or levels of remuneration provided by alternatives).

7.3.2   Co-incidental financial incentives

Aside from intrinsic, personal motivations for artistic labour, there are some incentives for production that are financial but not a result of consumer-facing copyright law. Clear examples include the income that musicians can earn from live performances and merchandising, the income that free and open source software authors can earn from consulting and support services based on the code the have written, revenue from advertisements included with free material, the range of voluntary payments ranging from tip jars to schemes like the street performer protocol, and cross-subsidies for authorship from other parts of organisations (Clarke 2007). Cinemas are perhaps a similar kind of non-copyright-dependent income for the film industry.[323] Note that these incentives, which I will call co-incidental, are equivalent to the "first mover advantages" discussed in the literature on the economics of patents.

Co-incidental incentives should be analysed because, like intrinsic motivations, they mitigate the free rider problem that makes copyright necessary. At the completion of this thesis, I have not encountered satisfactory data to quantify the proportions

---

cultural production. Writing open source software for work is clearly a for-profit activity; writing tools for one's own use also has concrete benefits. These last two cases (groups 2 and 3 in the study) are more closely analogous to the coincidental incentives discussed in Section 7.3.2.

[322]See, for example, (Moglen 1999, "Because It's There: Faraday's magnet and Human Creativity").

[323]A state of complete copyright anarchy might have some negative impact on cinemas because pirated copies would leak to pirate theatres. But this would only be true of anarchy that extended beyond the consumer sphere; and in more likely info-anarchic scenarios it would likely just increase the fixed costs to film studios of making their distribution and projection networks secure (see also Chapter 2 note 38 and accompanying text).

of financial incentives that are provided by these revenue sources, though anecdotes abound in claiming that they are becoming more and more important, especially for the music industry.[324] This issue is therefore flagged as an important aspect of further work.

It should in any case be noted that one serious problem with an anarchist reliance on intrinsic and co-incidental motivations for cultural production is that the absence of formal market-oriented incentives will sometimes allow other, secondary and less helpful financial incentives to dominate. For example, cultural works designed around advertising objectives will grow in influence.[325] Those kinds of works already exist, but a decline in the production of expensive, glitzy cultural material for the market-place would give advertisers a new incentive to fill the void. Organisations with political objectives might find that they could benefit greatly from paying talented authors for their work (when otherwise, those authors might have been comparatively free to pursue the projects of their choice, and works funded for political purposes would have more competition from copyright-funded material). Libraries of works produced for advertising and political purposes are not necessarily the ones that we would best off with.[326]

7.3.3 An industry-dependent answer?

While individuals are often motivated by considerations other than profit, capital tends to be different. Whenever the investment required for particular cultural projects becomes large — more than the kinds of sums that people spend on their hobbies — and is required for expensive equipment or non-creative labour, spontaneous production driven by intrinsic motivation is less common. Huge numbers of skilled and talented individuals form bands to write, rehearse, record and perform music without much prospect of being paid; it is rarer to see the stage, lighting, makeup, camera crew and other support staff required to make most high quality feature films working as a hobby. It is more often the case that these workers need to be motivated and organised

---

[324]One sign of this has been the appearance and predominance of so-called "360 degree" recording contracts, which assign all revenue streams (including touring and merchandising revenues) from artists to record labels.

[325]This presently appears to be the case with popular music in China (Crampton 2003).

[326]This is not to say that political patronage is necessarily destructive of great art; such a library would after all include Dante Alighieri's *Divine Comedy* and Shakespeare's plays. In truth, institutions of political patronage are too complex and varied for me to treat them satisfactorily here; see for example (Westfall 2002, 39–41).

by wages, and that means that many films need to attract substantial investments.

Some limited degree of patronage from rich individuals or foundations might still be available to fulfil this investment role in the absence of copyright — but it is improbable that existing big-budget films or TV series would often be funded out of public-spirited charity. Co-incidental incentives are sometimes sufficient to organise large investments.[327] An ample supply of entertainment goods that are very expensive to produce may always depend on some system of financial incentives — unless copyright or public funding is available, they will simply not be produced in adequate quantities.

The need for large-scale investment is not limited to the film and television industries. The popular music industry has a recent history of spending large fractions of its investment on video clips and marketing campaigns. To some extent, these expenses are actually inefficient (see Section 7.5.5 below), but, if video clips are goods that people actually value, the failure to provide adequate incentives for their production would be regrettable.

In the end, it is clear that incentives from sources other than copyright and government intervention make a large contribution in the music, publishing and software industries, but have much less of an effect in the film and high-quality television markets, which require large and co-ordinated expenditures of capital. The fact that copyright law is so un-systematic in the way it treats these industries differently may be an illustration of how difficult it is to implement (or even design) economically efficient systems of exclusive rights.

For our present purposes, the most important corollary of the variable necessity of incentives in different industries is that the normative status of information anarchy varies correspondingly. Anarchy would have a calamitous impact on the production of Hollywood action films, but far less impact on the production of novels, comics or folk ballads.

---

[327]TV programs are of course partially funded by co-incidental advertising revenue that is independent of copyright. But the standards of television programming have risen as copyright-based sources of funds have become available (Gore Vidal reportedly offered an alternative explanation for the rise of high quality TV programming: "television is now so desperately hungry for material that it is scraping the top of the barrel." `http://en.wikiquote.org/wiki/Gore_Vidal`) Revenue from DVD sales has without doubt increased the investment in and quality of top-of-the-market TV production significantly.

7.3.4 Models to compare anarchy vs. the other regimes

Although we do not know how much of cultural production is motivated by copyright incentives, it turns out that we can say how much difference copyright would need to be making in order to be better, on utilitarian terms, than information anarchy. The answer depends on the size of artificial scarcity losses, but Chapter 6 provided estimates those losses in the case of the music industry.

The models introduced here are not especially complicated from a mathematical perspective, but they take a different approach from the rest of the literature by treating the effectiveness of copyright incentives as an independent input variable, and examining the tradeoff between that variable and artificial scarcity. They allow us to determine the minimum effectiveness copyright would need to have in order to outweigh deadweight losses of the size discussed in Chapter 6.

Anarchy vs Weak Copyright

The universe of possible creative works can be partitioned into four categories: A works which only exist when copyright-like incentives are present; B works which exist under either anarchy or copyright, but whose consumption will be lowered by copyright and DRM when those are available; C works in the "commons" which exist under anarchy or copyright and which are never affected by copyright and DRM restrictions; and D works which exist under anarchy but not under copyright. Thus there would be $A + B + C$ works produced under copyright, and $B + C + D$ works produced under anarchy.

Suppose that each of the A, B, C and D works has (on average) a potential surplus value of α, β, γ and δ respectively.[328]

Next, let us set out two simplifying assumptions which conservatively favour the copyright regime. Firstly, assume that the group of anarchy-only works is empty, *i.e.* that $D = 0$. Secondly, note that we will assume that the social value contributed by the B and C works is independent of whether the A works are also present.[329] In reality, it is likely that there is some diminishing marginal utility with increased cultural

---

[328]Potential surplus values are the social surplus from the work if it is distributed to everyone who benefits from it. Note that a work would have negative surplus value if the costs of producing it (including opportunity costs) was larger than the benefit that could ensue from it.

[329]Technically, relaxing this assumption would require either that Equations 7.1 and 7.2 use some form of mathematical composition other than addition to combine the welfare contributed by the different works, or equivalently that α, β, γ and δ not be constants.

consumption, making the B and C works less valuable under the copyright regime, and strengthening the relative case for anarchy.

Finally, assume that artificial scarcity costs are a proportion $x \in [0, 1]$ of the value of a given work (and that this proportion is the same for the A and B works). Then society's total welfare $W_{\copyright}$ under a copyright regime is given by:

$$W_{\copyright} = (1 - x)(\alpha A + \beta B) + \gamma C \tag{7.1}$$

Under anarchy, the total welfare $W_{\oplus}$ is:

$$W_{\oplus} = \beta B + \gamma C \tag{7.2}$$

The degree to which copyright is preferable over anarchy is given by:

$$W_{\copyright} - W_{\oplus} = (1 - x)\,\alpha A - x\beta B \tag{7.3}$$

When this quantity is positive, a copyright regime is preferable to anarchy. That is true if and only if:

$$(1 - x)\,\alpha A > x\beta B \tag{7.4}$$

$$\frac{\alpha}{\beta}\frac{A}{B} > \frac{x}{1 - x} \tag{7.5}$$

The term on the left, $\frac{\alpha}{\beta}\frac{A}{B}$, is the ratio of the social value of works which require copyright or similar incentives for their existence, over the social value of works which do not require the incentives but which nonetheless suffer from the same artificial scarcity when copyright is present. The condition is illustrated in Figure 7.1.[330]

Any given marketplace, in which copyright and DRM effect a particular increase in cultural production, and in which there is a particular level of artificial scarcity, is a single point on this two-dimensional chart. If we knew the values of x and $\frac{\alpha}{\beta}\frac{A}{B}$ in, say, the marketplace for jazz recordings, then we could locate it within the Figure 7.1 graph and thereby determine whether anarchy would be better than copyright for those recordings.

---

[330]Note that this model just compares artificial scarcity costs and incentive benefits — the effects discussed in Chapters 8 and 9 will of course perturb the graph (largely in favour of anarchy). Relaxing the assumptions above ($D = 0$ and the form of the welfare functions as discussed in note 329) would favour anarchy.
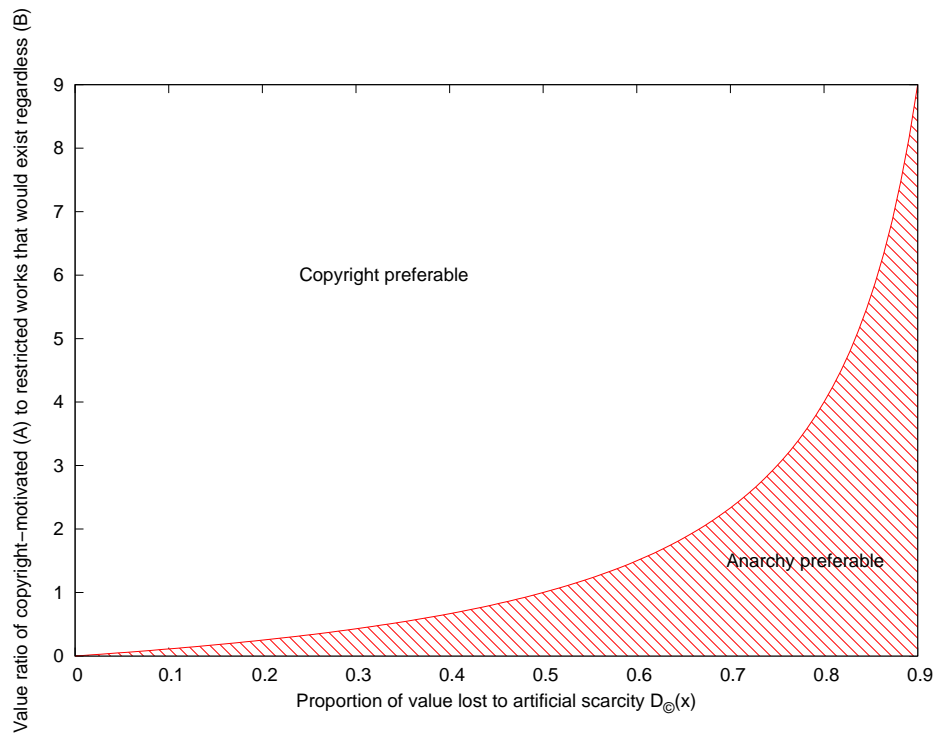
Figure 7.1: A simple utilitarian model comparing Anarchy and Weak Copyright

We can in fact do this for the music industry, because Section 6.3.2 obtained estimates of total deadweight loss as a percentage of the value of that industry (55–98%), from which we can calculate x:

$$\frac{55}{100 + 55} < x < \frac{98}{100 + 98} \tag{7.6}$$

$$0.35 < x < 0.49 \tag{7.7}$$

The implications of these bounds for x in the music industry are shown in Figure 7.2. As illustrated there, $0.35 < x < 0.49$ implies that copyright incentives need to produce a value ratio of more than some threshold between 0.5 and 1, in order for copyright to be preferable to anarchy. Equivalently, we need to have between about 150% and 200% more music (weighted by quality) in order for the artificial scarcity costs to be worth paying. It follows that the utilitarian case for anarchy is a very slightly
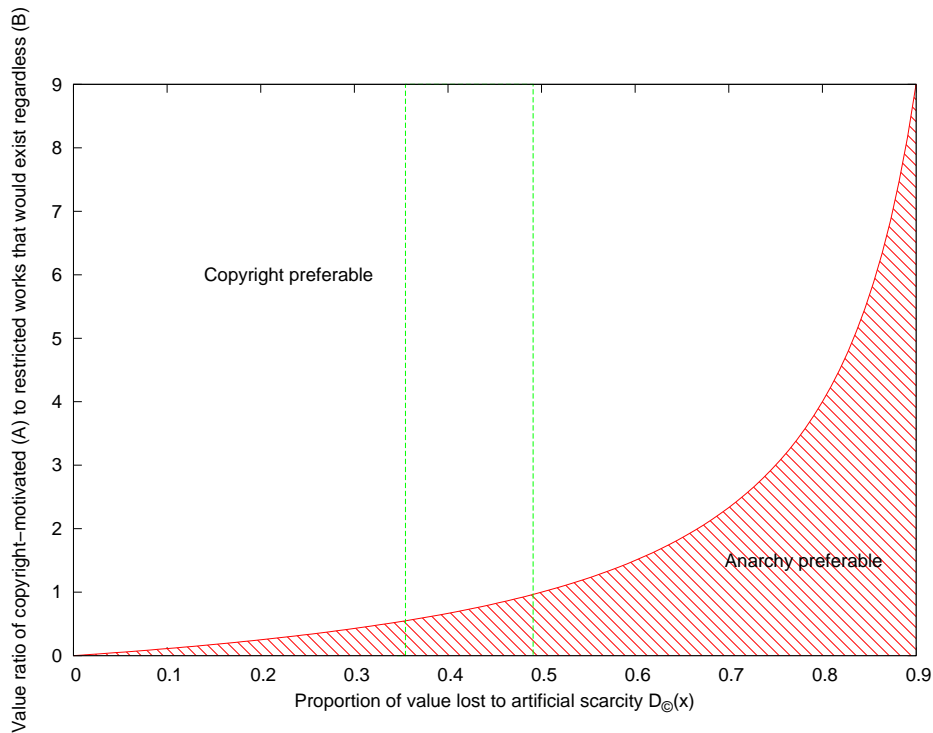
Figure 7.2: Anarchy vs Weak Copyright: the music industry lies between the green lines

stronger than one might have presupposed, at least in the music industry; copyright may well increase the amount of music produced by more than 50–100%, though that is not certain. By contrast, this is quite certain not to be the case in the film industry both because x is lower and $\frac{\alpha}{\beta}\frac{A}{B}$ is likely to be higher.

Anarchy vs Virtual Markets

Now let us turn briefly to the comparison between anarchy and virtual markets. Under a virtual market, a different number of works A' will be motivated and have an average surplus value of α'. Assume the excess taxation and infrastructure cost of the virtual market is T. Then welfare is given by $W_{\circledcirc}$:

$$W_{\circledcirc} = \alpha' A' + \beta B + \gamma C - T \tag{7.8}$$

And the advantage of a virtual market over anarchy is simply

$$W_{\circledcirc} - W_{\oplus} = \alpha' A' - T \tag{7.9}$$

Or as a constraint,

$$\alpha' A' > T \tag{7.10}$$

This shows that the virtual market/anarchy tradeoff is not driven by deadweight loss x (as in the previous comparison), but simply by the question of whether the incentives provided by the alternative compensation system are large enough to outweigh the tax overhead of creating it. This is the standard cost-benefit test found in public economics: if we start in a state of anarchy, is there a socially constructive project of collecting some tax and paying it to some artists? This tradeoff and the distortionary cost of taxation will be discussed further in Chapter 10.

7.3.5 Virtual vs copyright markets

The previous section established the conditions under which anarchy works best. But what about comparisons amongst the other three regimes?

The welfare advantage of a virtual market regime over a copyright regime is given by:

$$W_{\circledcirc} - W_{\copyright} = \alpha' A' - (1 - x)\alpha A + x\beta B - T', \tag{7.11}$$

where T' is the net cost of the ACS.[331] This advantage for virtual markets is positive if and only if

$$\alpha' A' + x\beta B > T + (1 - x)\alpha A \tag{7.12}$$

Rearranging the terms, we have:

$$\left(\frac{\alpha'}{\alpha}\frac{A'}{A} - 1 + x\right)\alpha A + x\beta B > T \tag{7.13}$$

The left hand side of this inequality has two terms. The second, $x\beta B$, is just the impact of copyright on the surplus from works that would have been created without copyright. The first term is more complicated and deserves some further examination.

---

[331]T' includes all the taxation, infrastructure, transaction costs and all the other costs discussed in Part IV. It is essentially certain that $T > 0$ and $T' < T$, and likely but perhaps not certain that $T' > 0$.

It is the total possible surplus from copyright incentives αA, multiplied by a combination of other variables:

$$\left(\frac{\alpha'}{\alpha} \frac{A'}{A} - 1 + x \right) \tag{7.14}$$

This expression is a measure of how much better (or if it is negative, worse) virtual markets function at delivering the potential surplus of the copyright regime. Note that the subexpression $\frac{\alpha'}{\alpha} \frac{A'}{A}$ represents the ratio of the surplus from virtual market incentives to the potential surplus from copyright incentives. If $\frac{\alpha'}{\alpha} \frac{A'}{A} > 1$, virtual markets gain advantage both from offering better incentives than markets, *and* for attaining higher distributional efficiency; if $\frac{\alpha'}{\alpha} \frac{A'}{A} < 1$, there is a tradeoff between exclusive rights offering better incentives but the alternative offering better distribution.

Some further substitution allows us to neaten up this equation. Following the observation above, it makes sense to rewrite $\frac{\alpha'}{\alpha} \frac{A'}{A} = 1 + \gamma$, where $\gamma < 0$ in cases where copyright provides better incentives; $\gamma > 0$ when virtual markets provide better incentives, and $\gamma = 0$ if they are equivalent. Suppose that copyright-like incentives are responsible for some proportion k times more production value than coincidental and intrinsic incentives; that is:

$$\alpha A = k \beta B \tag{7.15}$$

Substituting into constraint 7.13 and rearranging:

$$\left(\frac{\alpha'}{\alpha} \frac{A'}{A} - 1 + x \right) \alpha A + \frac{x}{k} \alpha A > T \tag{7.16}$$

$$\gamma + x + \frac{x}{k} > \frac{T}{\alpha A} \tag{7.17}$$

There are still many variables present, so it may be helpful to illustrate this constraint by substituting in some plausible values. Suppose that $x = 0.42$ (around the middle of the range for music deadweight losses identified in Chapter 6, and non including inequality effects that would effectively make x larger), and that $k = 3$ (which would correspond to copyright incentives being responsible for three quarters of the value of the industry in question). Then we would have:

$$\gamma + 0.55 > \frac{T}{\alpha A} \tag{7.18}$$

The left hand side measures the distribution/production advantage of the virtual market. It is positive unless $\gamma < -0.55$, that is unless the incentives under the copyright system are 55% more effective. The right hand side measures the implementation cost advantages of the copyright system, as a fraction of the social value of the works produced.

It will be a little easier to understand the implications of this constraint once we have an idea of which system of incentives is more accurate (*i.e.*, is γ positive or negative?), and once we have some measurements of the costs of taxation, discussed in Chapter 10.

7.3.6   Weak vs Strong Copyright

The modelling in the previous sections compared weak copyright to anarchy and virtual markets. What about the strong DRM/copyright regime? Until more and better results are obtained from various natural experiments, the most obvious assumption to make in addressing the question is that the feudalist copyright regime would produce more works of value in direct proportion to the extra revenue raised.

In the case of the music industry, Rob and Waldfogel's survey data again gives us some insight into the matter (Rob and Waldfogel 2006). Their data predicts that the status quo provides 20% less revenue for music publishers than strong musical copyright would. Given the 18–32% decrease in distributional benefit that would come from strong musical copyright (see Chapter 6), it is quite unlikely that feudalism would be preferable, all other things being equal. And, as we will see in Chapter 9, all else is not equal. At least for music, weak copyright is superior to strong copyright.

7.3.7   A pragmatic conclusion on the necessity of copyright-like incentives

Section 7.3 has examined the question of how production induced by copyright regimes relates to production that would occur regardless of that induction. It began by discussing some of the reasons why some production occurs independently of incentives that are based on payment for the product. It then examined how the relative levels of copyright-motivated and copyright-independent production interact with the phenomenon of artificial scarcity to make the different regimes more or less desirable.

Accurate data on the importance of financial incentives (from copyright or copyright-

replacing institutions) is elusive. We know that strong financial incentives are necessary for some instances of cultural production and probably not very necessary in others (*Titanic* vs *Kiss*). We can take informed guesses[332] about the kinds and relative numbers of instances in each category, but in the end we have little idea about how much difference these incentives are actually making to society's supply of entertainment. Some difference, certainly; possibly a lot — but we do not know. As a result, the mathematical treatment of the problem contains a parameter such as $k = \frac{\alpha}{\beta} \frac{A}{B}$ to represent this uncertainty.

Simple models *were* able to say, with a reasonable degree of precision, how large k needs to be — how large a difference exclusive rights need to make (as a function of the degree of artificial scarcity they impose) in order to be preferable to states of information anarchy. The answer, illustrated in Figure 7.2, is that, if copyright incentives are raising music production by a factor of less than some threshold between 2 and 3.2, the benefits do not justify the expense, and we would be better off if copyright law did not prohibit the non-commercial reproduction of music. The case for anarchy is probably weaker in other copyright industries, though there is less satisfactory data on them at present.

The comparison between copyright and virtual markets is slightly more complicated than the comparison of either of them against anarchy. But it was possible to derive some parameterised constraints that relate the levels of deadweight loss, the relative effectiveness of the incentives produced in each regime, and levels of taxation inefficiency under which the alternative compensation is preferable to an exclusive rights regime.

7.4   Some Relevant Theoretical Results

The previous Sections have attempted to set out a sufficiently nuanced foundation for understanding when incentives may give grounds for choosing one copyright system over another. With that foundation in place, I will now shift to a specific economic question: how well do copyright markets and alternatives translate consumer preferences into incentives for producers? This Section begins that task by reviewing some of the previous results in the economics of public goods that answer formalised versions of the question, explaining why these models reach different conclusions, and

---

[332]They are, at least, informed by anecdotes or surveys of information producers about their motives.

looking for those which most convincingly apply to digital copyright.

There are many theoretical economic models in the literature that directly address the effectiveness of different institutions in generating incentives for the production of public goods. Most of these models work with the assumption that public goods are pure: they posit that no exclusion is possible and investigate the efficiency of various governmental solutions. The literature is extensive and technically complex; satisfactory exposition demands a fair amount of time and mathematical background from the reader. For that reason, I will just discuss results and implications from mechanism design models without explaining how they work in detail. Since this thesis is not attempting to add to the ranks of those results, this approach should not be problematic.

Section 7.4.2 makes some general observations about that literature and briefly summarises results that are relevant to the present inquiry. The wide variation in those results indicates that the efficiency of governmental provision of public goods is closely entwined with the nature of the public good in question and especially with subtle properties of the information possessed by, and strategies pursued by, various parties. This means that there is significant uncertainty about whether public funding institutions can theoretically provide efficient incentives, although there are grounds for cautious optimism.

A couple of models (Wright 1983; Shavell and van Ypersele 2001) deviate from the pure-public-goods approach by aiming specifically at "intellectual property" and comparing publicly funded reward systems to marketplaces based on legislated exclusive rights; one of those is examined in detail in Section 7.4.1. It is also optimistic about the prospect of alternative compensation systems.

In the end though, there are many complications in the real-world application of proposed publicly funded alternative compensations. Although some of those issues have been modelled in isolation, there is as yet no unified theoretical framework for a complete comparison. This Section essentially deals with the available theoretical baselines; Section 7.5 will deal with the messier complications.

7.4.1 A microeconomic model of prizes under asymmetrical information

One model of particular interest is Shavell and van Yperseles' (2001) comparison of publicly funded rewards to a patent system. Although it is heavily stylised, it specif-

ically addresses the same kind of policy problem that I am examining in this Part of the thesis.[333] They consider four regimes: patents, government rewards based on the *ex ante* anticipated value of inventions, a combined patent/prize regime in which inventors get to chose which form of remuneration they prefer,[334] and a regime where governments base *ex post* rewards on observed sales measurements.

The normatively relevant effects in their models are monopoly deadweight losses and non-optimalities in the profit-maximising level of investment chosen by the inventor (because R+D returns[335] are assumed to be positive and diminishing with investment, there is an optimal level against which actual investment can be compared). This means that they are addressing the economic effects discussed in this chapter and the preceding one, in a single stylised picture.

They start with a probability distribution for a parameter t which represents the usefulness of and economic demand for a possible invention. In their model, inventors know the exact value for t, while the government knows only the probability distribution of t.[336] Given t, the demand function is common knowledge. Perhaps problematically, the probability of a research project's success as a function of investment is also assumed to be common knowledge (Shavell and van Ypersele 2001, p. 532).

In the resulting model, the patent regime suffers from two well-known inefficiencies: deadweight loss, and the fact that non-existent or imperfect price discrimination leads to sub-optimal levels of investment by inventors (if a larger proportion of the surplus could be appropriated, inventors would spend more — distributional effects are not considered).

The efficiency of the *ex ante* reward regime depends on how widely the social value of inventions varies from its mean.[337] The efficiency of the *ex post* regime is argued, by mathematical hand waving, to be higher than that of the *ex ante* rewards

---

[333]Their model does not specifically address copyright goods; they are discussed in passing but the model is limited to patentable inventions "for concreteness" (Shavell and van Ypersele 2001, p. 529). In their models, an invention is a yes-or-no good with a probability of being discovered that depends on investment. This can be transposed into a copyright setting by replacing the 'probability of discovery' with a 'quality of the work' variable that is similarly dependent on investment. Social welfare is computed similarly in each case.

[334] Compare to Litman's (2004) subsequent suggestion of a similar scheme for copyright.

[335]R+D returns are measured in units of probability of success.

[336]This assumption is not entirely realistic in that it causes some of the informational challenges of constructing alternative compensation systems to disappear, such as the issue of market size which I discuss in Section 7.5.1

[337]See Shavell & van Ypersele, Equation 10. This aspect of the paper is perhaps the most unrealistic, because the distribution of t should not be exogenous. If governments offered fixed rewards, firms would go looking for cheap, low quality inventions.

because sales data helps the government constrain the possible range of t values for each invention.[338] The distortionary costs of taxation are not modelled.[339]

Shavell and van Ypersele's paper provides several interesting theoretical results. One is that governments can conceivably do a better job of offering incentives than markets based on copyright or patent-style exclusive rights — even though they have less information than producers in those markets. It is shown that the comparative efficiency of the rewards offered by different regimes is intimately linked to the distribution of demand functions induced by the distribution of t.[340] The more predictable market demand is, or the more predictable characteristics it has, the better governments will be able to emulate it[341] (this is precisely the reason that alternative compensation systems are more promising for non-commercial access to cultural and entertainment works than they are, say, for patentable inventions or business software). Finally, once deadweight losses are included, their results show an interesting triangular order of preferability between *ex ante* rewards, patents and hybrid institutions. In their models, hybrids are strictly preferable to patents; rewards are superior to hybrids when the government's information is sufficiently good; but patents (and hybrids) are preferable to rewards if the government's estimations of the value of inventions are too error-prone.

7.4.2 Demand revelation/resource allocation mechanisms

Background

Models like Shavell and van Ypersele's, discussed in the previous section, or that of Wright (1983), make assumptions at the outset regarding the information available to governments and inventors about the demand for proposed public goods. With these assumptions in place, they proceed to examine the efficiency of prize based institutions based on those information endowments.

A much larger body of literature on public goods takes a structurally different approach to the problem, studying processes by which governments could obtain that

---

[338] See Shavell & van Ypersele (2001, Part II.H). This strategy on behalf of the government would also at least mitigate the problematic incentives for firms to prefer low-cost, low-value inventions.

[339] Shavell & van Ypersele concede this weakens their conclusions on rewards.(Shavell and van Ypersele 2001, p. 544). I discuss the size of distortions in Chapter 10.

[340] *Id.* Equations 10, 11, 10′.

[341] Shavell and van Ypersele observe that demand must be predominantly distributed close to its expectation value for rewards to be efficient. Section 7.5.1 observes that for *ex post* rewards, the more important variable is demand per consumer accessing the work, a quantity which some copyright *marketplaces* seem to be very poor at measuring.

kind of demand information in the first place. The models in that literature, which are often based on game theory, include explicit messages or votes that consumers can pass to governments, and which are the only source of information for valuing public goods.

This literatures is a part of the sub-discipline of economics and game theory known as "mechanism design". Authors discuss what are variously known as "resource allocation mechanisms", "demand revelation mechanisms", or "planning procedures" for public goods production. These mechanisms define a way that consumers can send messages to a planner/government,[342] and then specify how much taxation will be raised from each citizen, and how much of the public good will be produced, as a function of these messages. Extensive surveys of the relevant literature can be found in Campbell (1987) and Cornes & Sandler (1996, Chapter 7), although it should be noted that new contributions continue to appear.

Unlike the work of Wright or Shavell and van Ypersele, these models tend not to be designed to compare markets and public provision of semi-excludable "intellectual property" goods; the results about government provision usually apply to pure public goods for which markets are assumed to be impossible or extremely inefficient due to free riding.

The inclusion of explicit messages allows the study of the incentives for consumers to reveal their preferences truthfully: is it in their interest to honestly state their willingness to pay for a public good, or would they be better off either exaggerating or understating it? This is an especially important question in mechanisms where the amount of taxation each consumer pays is highly dependent on the messages they send, or where the final output of public goods is strongly dependent on individual messages.

Ideally, we would be able to read out a few results from the mechanism design literature on public goods provision, demonstrate that one of them corresponds closely to the virtual market regime, and conclude that virtual markets therefore constituted an efficient resource allocation mechanism (or not). Unfortunately, the translation is not so simple. Firstly, it is not clear which of the many available sets of assumptions best describes a particular alternative compensation system such as a virtual market.

---

[342]For instance, the message might be a dollar amount that the consumer is willing to pay; a yes or no vote saying whether they would like a public good produced under a given funding structure; a message saying whether they would like *more* or *less* of some public good to be produced; or sequences of any of these over time.

Secondly, when it is shown that a mechanism lacks desirable properties, it is hard to know how serious the problem is: a result showing Pareto non-optimality does not say how far from optimal the outcome is; a result showing a lack of incentive compatibility does not say how far participants' messages will deviate from those the mechanism asked for.

The variation of assumptions in the literature arises because public goods planning problems can be specified in a great many ways. The game can be a one-off project (how much will the population spend on refurbishing their town square?), an iterated process (what will this quarter's budget for tree planting be?) or a continuous-time approximation of rapid iterations (which is not entirely unrealistic for some online activities). Consumers' strategies can be local (maximise their pay-off at this iteration) or global (maximise their expected future pay-off). The bulk of the literature examines situations where there is only a single public good, although there are a few results about planning for the provision of two or more. Which of these variants is more realistic depends of course on the phenomenon to be modelled.

As for results, most of the resource allocation mechanism literature has focused on looking for mechanisms that have certain ideally desirable properties. The three most commonly examined ones are Pareto optimality,[343] incentive compatibility (the best thing each consumer can do is to disclose their preferences honestly) and rational participation (whether everyone is better off under the mechanism than they would have been, had it not existed). These properties are frequently contradictory.[344] Furthermore, they may be present at equilibria, but whether they hold under other conditions, and whether resource allocation games converge to those equilibria is a different matter.

An additional measure of confusion arises because the meaning of Pareto optimality is not the same in every model. In some cases Pareto optimality implies efficient taxation, while in others the taxation system is externally specified and therefore cannot give rise to inefficiency[345] (consumers' preferences for public goods are assumed to have been formed with the knowledge of the tax formula that will pay for them).

---

[343] Pareto optimality indicates that no individual can be made better off without causing someone else to be worse off. It is the predominant measure of efficiency used in the economics literature, although it does not address important concerns about the equity of wealth distribution; see, for example, (Cornes and Sandler 1996, at 220).

[344]For example, Roberts (1979, Part 4) built on earlier work by Hurwicz (1972) to show that if participants in an iterated demand revelation game are well-informed and employ long term strategies, then Pareto optimal resource allocation mechanisms will not in general produce honest preference disclosure.

[345]Compare, for instance, Clark-Groves with de Trenqualye.

Results of particular interest

Some of the early results about planning procedures were due to Bowen (1943), who discovered a set of necessary conditions for Pareto optimal production of public goods.[346] He also showed that if majority voting[347] is used to identify a satisfactory level of tax-funded provision of a single public good, then, under strict but not completely implausible assumptions, the result will be Pareto optimal.[348]

One of the conditions necessary for Bowen's model to reach Pareto optimality is that consumers' ideal amount of the public good is symmetrically distributed about the median,[349] and that each consumer's preferences for the public good have a "single peaked" maximum. A later proposal by Clark and Groves found processes that would produce efficient quantities of a public good with more general populations, provided that each person's preferences fitted certain mathematical constraints.[350] Their mechanism ensured that it was *always* in consumers' interest to reveal their preferences truthfully. But in order to work, it forced governments to run inefficient budget surpluses! Another mechanism developed by Groves and Ledyard fixed that problem, producing balanced budgets and efficient production regardless of preferences. But in their scheme, consumers only have an incentive to be truthful if the mechanism is in Nash equilibrium. Outside of it, participants have incentives to lie.

Perhaps the biggest problem with applying any of these results (regarding Clarke-Groves mechanisms, Groves-Ledyard mechanisms or Bowen's model) to a virtual market is the "single public good" assumption. Although it is tempting to regard existing public sector institutions as fixed, and to propose an alternative compensation system funding pool as a single public good which should be regulated by a single stand-alone

---

[346]These conditions later became known as the "Samuelson conditions" in attribution to Paul Samuelson (1954); history is gradually correcting to the term "Bowen-Samuelson conditions".

[347]*Majority voting* refers to finding an outcome such that it will win a two-choice election against any other alternative. It is easy to show that if there is a single parameter to be chosen, and each voter's preferences are single-peaked, then the median vote will command a majority; see (Bowen 1943). Note that a majority vote is a criterion, not a voting system.

[348]In Bowen's model, taste for public and private goods are independent, each individual pays an equal share of the cost of producing the public good, and preferences are assumed to be normally distributed. This result is generalised in (Cornes and Sandler 1996, at 205–10), where it is shown that majority voting can be Pareto optimal whenever *individual variation* in preferences for the public good are independent of consumption of private commodities, using a more general Lindahl taxation formula which combines lump-sum and fixed-rate income taxation. Also, the requirement for bell-curve preferences is relaxed to include any symmetrical preference distribution.

[349]Or equivalently, that the mean and median are the same.

[350]See (Cornes and Sandler 1996, pp. 226–229).

plebiscite, it is not correct to do so. Glitzy pop music, politicised hip-hop and electronic jazz are not substitutes for one another, and none of them are substitutes for existentialist fiction. Each of these must be regarded as separate public goods if the notion of efficient production is to be meaningful.

Once several public goods are present, the design of optimal resource allocation mechanisms becomes more difficult. Indeed, Bucovetsky (1991) has shown that a majority-voting equilibrium exists only if the many-dimensional space of public good preferences can be reduced to two 'taste variables'.

Bucovetsky's result seems to suggest that a virtual market could not be optimal, because it will inevitably face a population with highly heterogeneous and complicated tastes. There is, however, some cause for optimism if we move away from models with once-off majority voting and replace them with an incremental approach where the levels of different public goods are adjusted gradually.

De Trenqualye has shown that if cost sharing (ie, tax) rules are fixed in advance,[351] then an incremental voting system, which adjusts the budget allocations to many different public goods in a continuous fashion, will reach a Pareto-optimal equilibrium.[352] I will show that virtual markets can be interpreted as a kind of de Trenqualye voting process, to which these optimality results apply.

The applicability of de Trenqualye's efficiency result to a virtual market

Constructing this interpretation is not entirely trivial, because de Trenqualye's model is dynamic and gradual, adjusting the previous budget slightly at each infinitesimal increment. The virtual market is on its face quite different to this type of dynamic mechanism — instead of gradual budget adjustments, it is what the literature would term a repeated static rule, under which the budget allocations for different public goods can change drastically at each iteration, and today's budget is not in any direct sense an input into the calculation of tomorrow's budget. In order to show that de Trenqualye's optimality results apply to a virtual market, it must be shown that the repeated static mechanism is equivalent to a dynamic one.

---

[351] That is, each person i knows that if it is decided that amounts x, y and z of different public goods will be produced, then they will be paying a total tax of $f_i(x, y, z)$ (where f_i is some function they can calculate).

[352] See (de Trenquayle 1997). De Trenqualye's model assumes voters make decisions based on their local preferences in the space of possible public goods; Pareto optimality at equilibrium and local incentive compatibility depend on convex preferences; guaranteed existence and inevitability of equilibrium depends on Euclidean preferences.

This equivalence can be achieved by transforming the definition of the set of public goods being provided and voted upon. A virtual market which provides one set of public goods approximates a de Trenqualye process with respect to a different but related set of public goods. The public goods provided by each iteration of the virtual market are remuneration to the authors and copyright holders of specific works; the public goods provided in the dynamic de Trenqualye sense are the medium-term incentives observed by those who are considering producing new works in the future. Conveniently, it is this second set of public incentive goods that we are most concerned about providing as a matter of public policy. Payments to individual rights holders might theoretically fluctuate in a drastic fashion from month to month, but this does not cause corresponding drastic fluctuations in the supply of new copyrighted works that consumers can avail themselves of. That is because those who engage or invest in cultural production with a conscious aim at profit most commonly make their decisions by observing the medium-term history of remuneration for previous works, and then estimating the likely rewards for their own work based on that history. Turning the virtual market into a de Trenqualye process requires an integration or summation operation, but the persistent nature of cultural goods happens to perform this operation automatically.

Figure 7.3 shows in more detail how de Trenqualye's model can be used as an approximate interpretation of the operation of a virtual market. There are two steps being illustrated: one is the translation of a virtual market vote of the type proposed in Section 5.1 into one of the votes in de Trenqualye's model. The second step is the existence of a point S, which represents the medium term incentives observed by producers, and which gradually moves in response to the sum of the votes in precisely the way that de Trenqualye's model requires.

Observe that the ratio of each user's votes (for different works) defines a line in the N-dimensional positive real space $\mathbb{R}^N+$, where N is the number of different information goods which can be chosen.[353] An example is shown by the ray V from the origin. If users are given a way of choosing how much revenue should be allocated to the virtual market,[354] their choice uniquely determines a point P where that ray intersects the

---

[353]Note that if some of these goods are substitutes for one another, the taste space will be compressed during the transformation step. The diagram shows "glitzy pop", an example axis for the incentive goods, where Alice might actually have sent her money to Britney Spears or the Spice Girls.

[354]As per the 'complete decentralisation' voting option discussed in Section 5.3.2. The accuracy of this data is further discussed in section 7.5.1 below.

hyper-sphere[355] whose radius is equal to the total budget B which Alice would prefer. The normalised vector i, towards P from the currently prevailing state of incentives S, is analogous to one of the votes used in de Trenqualye's model.

The point S moves gradually as the system operates, in a direction determined by summing all of the votes i from consumers. Under the conditions summarised in footnote 352, S settles at a Pareto optimal destination. De Trenqualye's assumption of non-strategic local voting is particularly justified in the virtual market transposition because of the extreme difficulty of collecting useful information about the global preference landscape, and because preferences are likely to be dynamic anyway.[356]

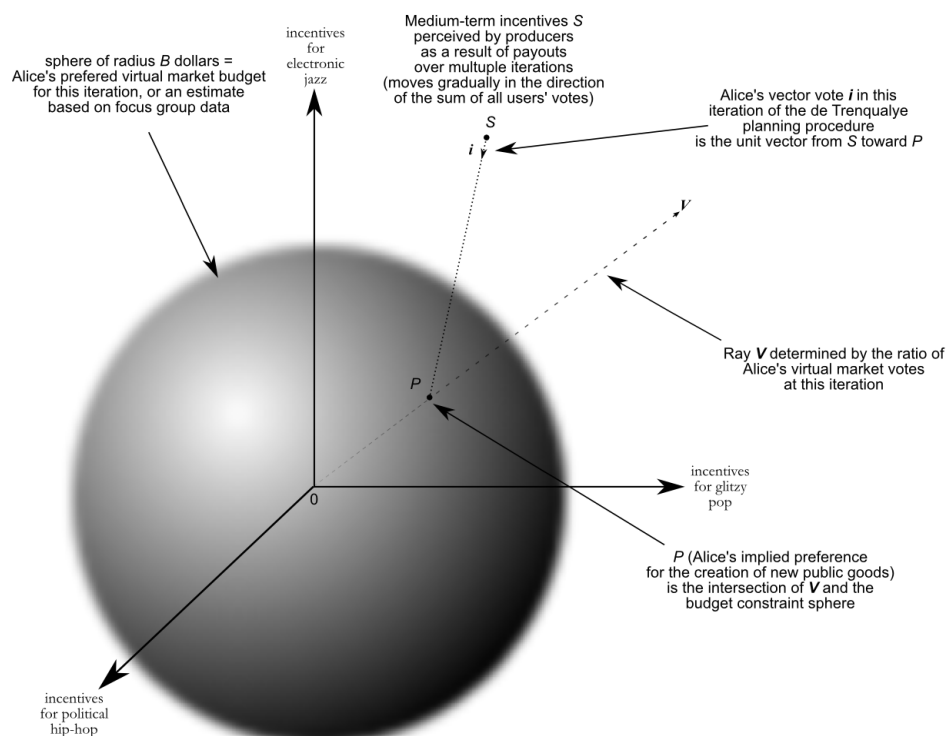Figure 7.3: Interpretation of a virtual market as a de Trenqualye planning process
The three axes indicate a 3-dimensional cross-section of the many-dimensional taste space

One limitation in this interpretation is the fact that consumers' preferences in cul-

---

[355]A sphere is a circle extended to three dimensions; a *hyper-sphere* is a circle extended to four or more dimensions.

[356]Compare (de Trenquayle 1997, Section 4.2).

tural marketplaces evolve over time; de Trenqualye assumed that preferences were fixed. If the rate at which preferences change is faster than the rate at which incentives converge towards equilibrium, then the fact that equilibria are Pareto optimal is no longer relevant. For comparative purposes, it should however be noted that marketplaces face exactly the same prospect of perpetual disequilibrium.

The applicability of de Trenqualye's result, showing Pareto efficiency for voting mechanisms that set budgets for incentivising public goods production, gives significant grounds for optimism about the baseline efficiency of virtual markets or similar compensation systems in setting budgets for many different copyright public goods.

The theoretical economic results discussed in this section were essentially optimistic about the quality of incentives under government-funded compensation systems: Shavell and van Ypersele's result showing that a well-informed government administered reward is more effective than intellectual property rules for any single public good; and the applicability of de Trenqualye's results, which show that virtual markets are potentially a very efficient mechanism for informing governments about which copyright goods it should reward.

These are essentially simplified, high-level mathematical stories. There are many caveats, complications, and details that become apparent when carefully considering how marketplaces and alternatives provide incentives. These may turn out to be important, and the next section turns to address them.

7.5   Incentive-generating information in real and virtual markets

In this section, I will examine how the incentives provided by actual copyright marketplaces, and actual plausible alternative compensation systems, deviate from the idealised models of Shavell and van Ypersele, Bowen or de Trenqualye — or for that matter, from the idealised model of Arrow and Debreu which underpins the belief that markets are Pareto efficient. There are surprisingly diverse reasons why the information that determines remuneration is not as good as it could be. I will discuss five of them: the information that the regimes use in determining the total level of funding available in cultural marketplaces (Section 7.5.1); differences in the way that the various regimes interact with the transparency of information goods (Section 7.5.2); differences between the way regimes encourage consumers to report their preferences

honestly (Section 7.5.3); differences that arise when consumers' influence on the cultural economy is essentially egalitarian, rather than proportional to their willingness to spend money (Section 7.5.4); and the degree to which different regimes' information collection systems respond to and encourage excessive advertising (Section 7.5.5).

Many of the effects discussed in this section (particularly 7.5.2–7.5.5) add a degree of uncertainty or probabilistic variation to the amounts that consumers pay, or are willing to pay, for particular copyrighted goods, when compared to the benefit they actually get from the good. The effect of such variation is that some producers are paid too much for their work, and others too little. It should be noted that in a simple analysis, paying some people too much and others too little does not necessarily change producers' behaviour, particularly if the variation is unpredictable. If investors are just as likely to benefit from the variation as they are to suffer from it, then their *expected* remuneration from a project is the same as it would be if the variation were removed. Uncertainties only matter if producers are risk-averse[357] or if an effect makes it *systematically* likely that certain works will be paid too much and others too little.[358] It is only when skewed remuneration patterns become entrenched that social welfare suffers greatly.

7.5.1 The most basic information: how much are individual works (or entire markets) worth?

Institutions for funding the production of cultural public goods have one core informational task, and that is to determine how much to pay for the production of each good. Alternatively, we can conceive of this task as being to determine how large the total funding in a given marketplace should be, and then how that total should be divided.

According to the mathematical results discussed in Section 7.4.2 above, virtual markets should be able to determine the answer to both of these questions in optimal or close-to-optimal fashion. In order for de Trenqualye's result to apply, however, two pieces of input information were required from Alice: her ratio of preferences for different works, and her preference for the total virtual market budget.

The first of these, the ratio of preferences for different public goods, is easy to obtain with any of the data collection methods proposed in Section 5.1. The second

---

[357]Which they are in practice; this is one of the possible reasons that cultural copyright industries have come to be dominated by a small number of very large firms.

[358]Abramowicz reaches the same conclusion about patent prizes; see (Abramowicz 2003, at 123).

category of input data, which can be conceived of as the total level of funding that is to be split up, or the tax rate to be applied to various commodities that are used to fund the virtual market, is more troublesome. During an initial transition period to an ACS, it would be possible to determine this number by selecting a tax rate that causes total funding for recorded music to be the same as it was under the previous copyright system. But over the longer term, this extrapolation would become increasingly arbitrary, and updated information would have to be obtained by one of the methods discussed in Section 5.3.2: contingent valuation surveys, elections amongst the audience as to whether levels of funding should be increased or decreased at any moment, or some kind of purely bureaucratic process.

Some authors have criticised alternative compensation systems on the grounds that they would presumably follow the third of those three methods, and that bureaucratic means for determining total levels of funding are inherently unable to adjust funding levels well and continuously in response to changing circumstances.

This section investigates the merit of this claim on its own terms: how much better are copyright based markets than government bureaucracies for determining total funding levels for entertainment industries?

Liebowitz's Argument Against Alternative Compensation Systems

Several authors have expressed concerns about alternative compensation systems on the grounds that they would not be able to reproduce the efficient patterns of production expected from markets (Liebowitz 2003a; Merges 2004; Liebowitz and Watt 2006). Stan Liebowitz stakes out this position most clearly when he argues that the most significant shortcoming in collective licensing schemes is their inability to match the market's ability to set efficient levels of total funding.

Liebowitz sets up the issue by stating that the primary goal in setting funding levels in collective licensing schemes is "often thought to be to mimic markets since markets are thought to provide efficient solutions that maximize economic values."[359] Though he is attributing this position to unspecified third parties, one gets the impression that Liebowitz himself agrees with it. In his account, a marketplace arranges for optimal levels of funding, and a bureaucracy for determining ACS tax levels should be compared to those optimal levels.

Liebowitz emphasises his argument by showing a graph of the revenue earned from

---

[359](Liebowitz 2003a) p 15.

album sales in the United States over a number of decades. The graph is reproduced here as Figure 7.4. He points out that the graph exhibits unpredictable features, with music sales rising and falling in response to economic circumstances and the appearance of new media formats (especially cassette players and CD players).
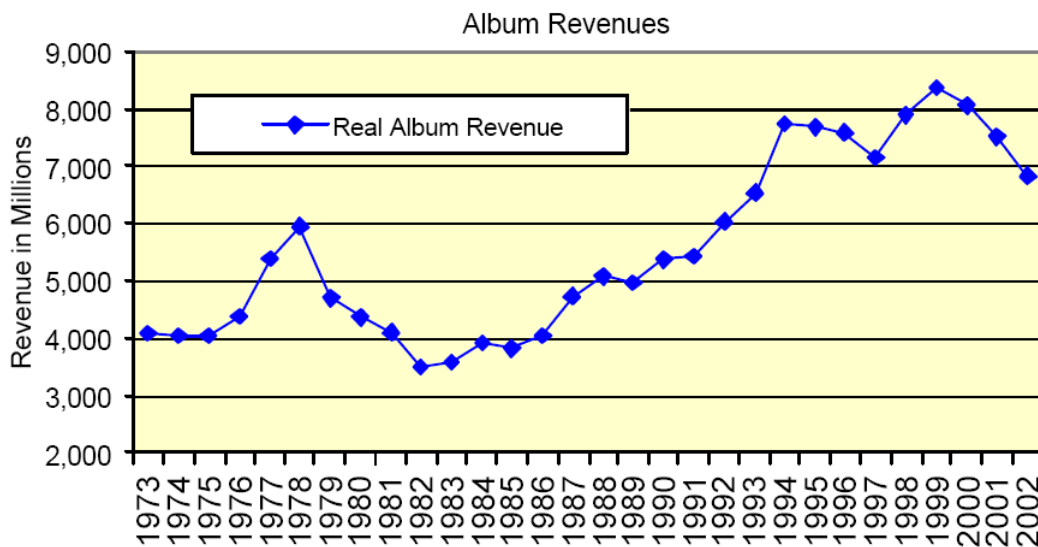
Figure 7.4: Real revenue from U.S. album sales over time
Copied from (Liebowitz 2003a, p. 16)

Liebowitz's argument may lead us to conclusions which are somewhat too pessimistic about alternative compensation systems, for two reasons. One is that the technical literature has shown that markets for excluded public goods (such as copyright goods) generally provide non-optimal levels of funding for those goods. Typically such markets under-fund production, although there are hypothetical scenarios where it is over-funded. In other words, the numbers in Figure 7.4 are not optimal, and are probably (but not certainly) lower than society might have preferred. It is easier for a bureaucracy to match the efficiency of public goods markets, which suffer from these kinds of informational problems, than it would be for a bureaucracy to second-guess the size of the market for private goods like nails, or hammers, or doll houses.

The second reason why Liebowitz overstates the inefficiency of bureaucracies is that such bureaucracies do not need to "fly blind" in determining funding levels. There are informative proxy variables available which are strongly correlated with expenditures on copyright goods. By following those variables, bureaucracies could do a

better job than one might expect at determining market sizes.

Let us examine each of these considerations in turn.

Production is not optimal in copyright marketplaces

The reasons why copyright markets are likely under-fund production is most simply depicted in the case of single good which may or may not be produced by a single producer with perfect foresight. The producer sets one or a small number of prices, so as to maximise profit, but is limited in their ability to price discriminate. It follows that there is some deadweight loss amongst consumers who are willing to pay less than the minimum chosen price, and some consumer surplus among those who are willing to pay more than the price they are offered. If the social surplus from the production of the good is positive, but smaller than the sum of deadweight loss and consumer surplus, the producer's profit will be negative and the good will not be produced. This form of under-funding previously appeared, for instance, in Shavell and van Ypersele's model of patent prizes (Section 7.4.1).

Models in the public goods literature generalise that phenomenon beyond the case of a market with a single producer. Oakland (1987, p. 515–517), for instance, summarises the literature's conclusions with a model of a market with N consumers and many producers, in which the price for each excludable public good (each song, say) is constant and does not involve price discrimination. All songs cost the same amount c to make, and are interchangeable, except that there is no benefit from having two copies of the same song. Consumers vary in their demand for songs. In equilibrium, songs are produced and sold at a range of prices from c/N — a song that is purchased by all or almost all of the consumers — to c — a very expensive song purchased only by the individual with the highest demand. In this equilibrium, songs are under-produced, though less severely than in the case of a single producer (Oakland 1987, p. 518–519).

Thompson (1968) shows that a similar market, in which the producers have perfect information about each consumer's demand and the ability to engage in perfect price discrimination, leads to the opposite inefficiency: over-production of the public good. In such a market producers would charge each consumer Alice for each song an amount equal to her *average* valuation for all songs, resulting in the consumers purchasing all of the songs but obtaining zero consumer surplus. The entry of more and more firms into this lucrative business eventually drives producer profits to zero, too. The equilibrium corresponds to significant over-production of the public goods

in question. Cornes and Sandler (1996, at 243–8) review Thompson's model and cite numerous authors who disagree with his premises as unrealistic. Interestingly, the extreme assumptions about what sellers know about demand underlying Thompson's analysis are probably more plausible for cultural works covered by omniscient DRM than for any other kind of public good. It is unclear whether any such markets would realistically reach states of supra-optimal production.

Although the assumptions that underpin these models vary, they all agree that production levels in markets based on the sale of excluded public goods are unlikely to be optimal. This consensus is inconsistent with the assertions of authors such as Liebowitz or Merges (2004), who believe that copyright markets are efficient on the production side, when in fact microeconomic theory does not appear to support that belief.

In practice, the situation is further complicated by the fact that there are many aspects of copyright law which affect production levels and are the subject of political decision making — including the duration of rights, thresholds of originality, or the scope of fair use and other exceptions. Most of these decisions ultimately affect the size of the market. One need look no further than the recent wave of retrospective copyright term extensions[360] to see that governments can pull these regulatory levers in a fashion that is far from optimal.[361]

How inefficient are production levels in copyright markets? The precise answer is unknown because we do not know which set of microeconomic assumptions is most applicable. Of the three models mentioned above, the simple monopoly model is probably the most realistic. In that model, where each work that may be produced is a unique opportunity for a single producer,[362] under-production applies to those works whose expected social return is positive, but smaller than the sum of deadweight loss and consumer surplus. Given that deadweight loss as estimated in Chapter 6 is already very large, and consumer surplus is another large total on top of it, we should expect

---

[360]See EU Council Directive 93/98/EEC of 29 October 1993 harmonizing the term of protection of copyright and certain related rights; Sonny Bono Copyright Term Extension Act (1998). Similar retrospective extensions were subsequently required by numerous bilateral trade agreements between the United States and other nations.

[361]Retrospective copyright term extensions decrease incentives for contemporary artists by increasing the rights clearance costs associated with sampling or adapting old works. Romer makes a similar point about the analogy between imperfect funding levels for reward systems and legislative failures resulting in burdensome state-supported DRM regimes (Romer 2002).

[362]This is essentially the same model used in (Shavell and van Ypersele 2001) and discussed in Section 7.4.

very problematic under-investment in works with moderately positive social surpluses. We don't know what proportion of opportunities this represents, but we should expect incentives to be more than 50% too low[363] in the case of marginally useful creative opportunities.[364]

There is an important question about the extent to which governments administering a virtual market could actually address this problem of socially desirable but not profitable works without paying a great deal more money to all publishers and authors.

If virtual market administrators were in a position to know the true production cost of a project, they could increase remuneration for those works which suffer from the problem, but not for all works. Unfortunately, this would create perverse incentives for producers not to control their costs.

Another approach would be to pay authors and publishers an extra amount corresponding to a portion of the deadweight loss avoided under the virtual market. This would occur naturally if the size of the virtual market was indexed to a measure like the number of hours spent watching, listening to or reading copyrighted works, because the avoidance of deadweight loss would increase those numbers. Such an approach would leave all of the consumer surplus available under copyright regimes to the audience, and a similar proportion of the newly available surplus, but at the same time ensure that more works of moderate net social valuable works become profitable to produce.

A third approach, which might be advisable in practice, would be to couple a small increase in payments of the sort just mentioned with a scheme of limited progressive redistribution of revenue to lower-income artists. This assumes that there is a correlation between low artist incomes and works which are of positive net social value, but are not profitable. Such a correlation is not guaranteed to exist — the question requires empirical study — but some redistribution of royalty revenue might be desirable

---

[363]In the case of a linear demand curve for a good sold at a single price, the appropriable surplus is precisely 50% of the social surplus. To see why, take a general linear demand curve $y = -ax + b$, $a > 0$, $b > 0$. Then revenue π is equal to units sold times price: $\pi = xy = -ax^2 + bx$, which is maximised when $\frac{dp}{dx} = -2ax + b = 0$, that is when $x = b/2a$. At that point profit $\pi = -b^2/4a + b^2/2a = b^2/4a$. The total social surplus of the product S is the area under the entire triangle, $S = \frac{1}{2}(b/a)b = \frac{1}{2}b^2/a$. Appropriable surplus is $\frac{\pi}{S} = \frac{b^2/4a}{b^2/2a} = \frac{1}{2}$, precisely 50% of total surplus. For convex demand curves, which are considered to be the normal case, appropriability is even lower.

[364]Note that this does not necessarily mean that the copyright system should be made 50% "stronger", or by matching every dollar of copyright revenue with 50 cents of government subsidies; the consideration only applies to works that are of moderate-but-positive social benefit, or prospective projects which investors believe are likely to be in that range.

regardless on other utilitarian grounds.

Are bureaucracies necessarily very poor at determining market sizes?

The literature tells us that markets for entertainment based on the artificial exclusion of public goods will often be of a non-optimal size. But are these non-optimalities large enough to make a bureaucracy look efficient by comparison? Possibly. Liebowitz used Figure 7.4 to support the claim that governments could not reproduce the year-on-year changes in market size that would be necessary for efficient cultural production. But the theoretical results reviewed in the previous section show that the market is likely to be wrong by a large amount in each year shown in that graph. It is conceivable that a virtual market could do a better job of securing the production of cultural works that have moderate social surpluses in an early year of its operation. But could a government change its funding levels dynamically, to respond to year-on-year fluctuations in the same way that the market does?

Liebowitz may be right that it could not, but the answer is greatly affected by the way the question is framed. If a bureaucracy tried to determine funding levels from scratch, given only the number from the previous year, there would be ample reason for pessimism. But there are indicative variables that could be used as proxies, and we should be less pessimistic about bureaucracies if they used those variables to determine funding levels.

One particularly promising indicative variable is the total number of hours that measured participants spend watching, listening to, and reading works that are funded by the virtual market. Those numbers should be expected to track, at least to a significant degree, any changes in the aggregate social value of those entertainment industries. Another possible variable is total expenditure on consumer electronic (CE) devices that are complementary to the entertainment goods in question. This second variable has the convenient virtue of actually having been recorded during some of the years that Liebowitz points to, and can be compared to the market size history which he argues would have been hard to predict. Since I argued above that the market size history is itself unlikely to be optimal, measures of how close a bureaucracy could get to it do not show that the bureaucracy is optimal. What these measures really evaluate is the ability of a bureaucracy to make year-on-year changes to the size of a virtual market, in a manner that is responsive to the same decentralised causes which influence the sizes of markets.
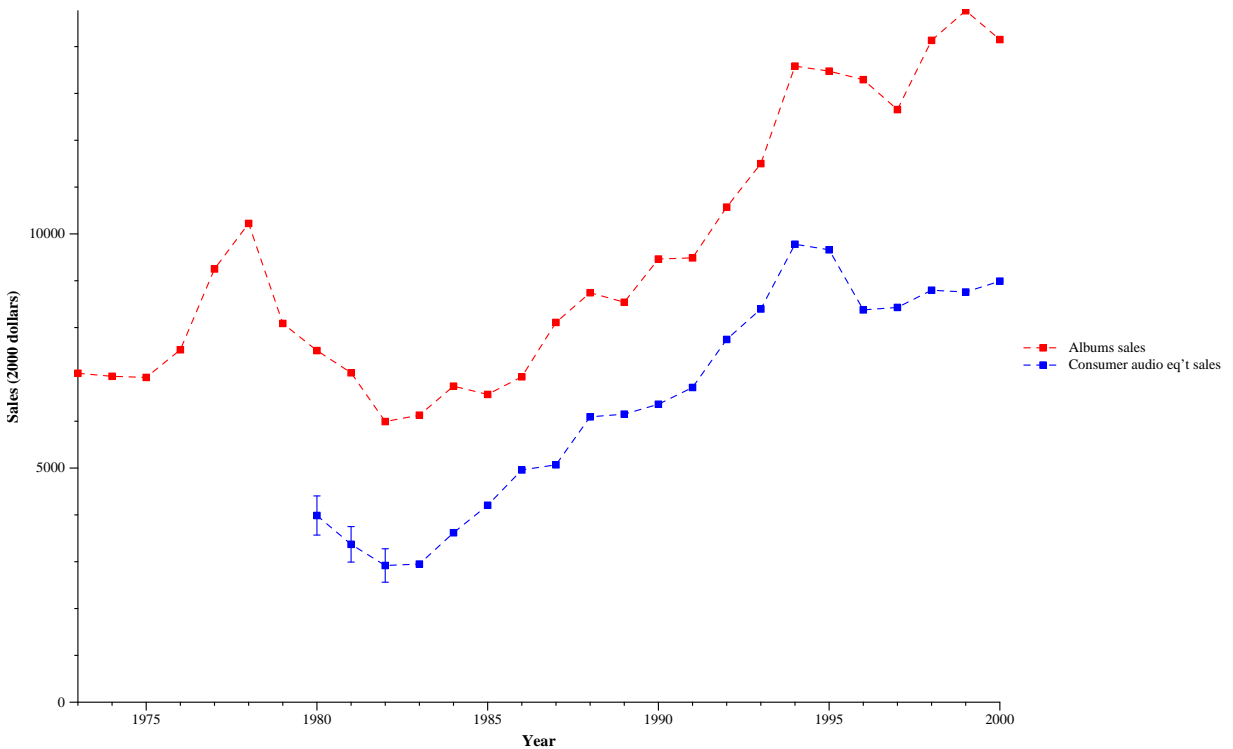
Figure 7.5: US album revenue and wholesale consumer audio revenue.
Album data courtesy of Stan Liebowitz. Consumer audio revenue data are wholesale U.S.
figures courtesy of CEA Market Research.

The actual comparison can be found in Figures 7.5 and 7.6.[365] These graphs show that between 1985 and 2000 the ratio of consumer spending on albums and audio devices was virtually constant, despite numerous technological changes during that period (such as the introduction of the "walkmen", CD players and "discmen"). They allow us to estimate what might have happened if copyright had been replaced by taxes on a set of commodities, such that during the year of transition the total level of funding for entertainment goods was constant, and such that in subsequent years, the tax rate was left constant, with resultant fluctuations in entertainment revenue.

---

[365]Several remarks are in order about these datasets. The CEA has records of wholesale sales of radios back to the 1950s, but only began keeping data on other kinds of devices in 1980. The graphed data includes all the categories which are strongly complementary to album sales (but I did not include sales of radios or MP3 players). Data on the sales of "boom boxes", an important category of devices, only goes back to 1983, which was mid-way through the "boom box craze" of the early 80s and by which year CEA data shows that 35% of US homes already had one of these devices. Rather than excluding the years 1980–1982 from the graph, I have added estimates of $900 million plus or minus $200 million for boom box sales during those years, allowing the first three data points to be included.
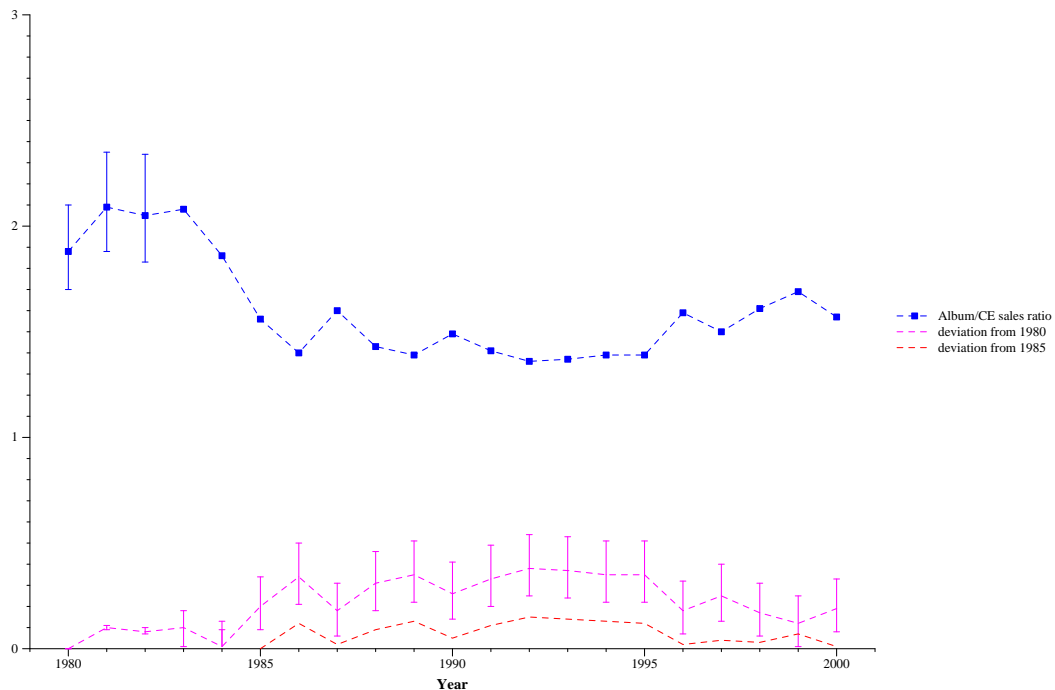
Figure 7.6: Ratios of album revenue to wholesale consumer audio revenue
Accompanied by the proportion by which the 1980 and 1985 ratios deviate from that year's
ratio.
Error bars are due to uncertainty in 1980 boom box sales.

A bureaucracy that had taken that approach would have obtained very similar results to the marketplace during the period sampled, with the average deviation in the ratio of consumer audio equipment to album sales being only 8%![366] The level of accuracy would, however, depend on the year of the transition. A transition during or after 1985 would have produced better results than a transition during the earlier 1980–1984 period, when consumers appeared to be spending more money on albums for each dollar spent on audio equipment (ratios of around 2:1, compared to around 3:2 from 1985 onwards). There are several possible explanations for this,[367] and the

---

[366]Of course, if most of the revenue was raised by actually taxing consumer audio devices, the tax would have to be carefully designed so as not to distort the market that was used as a source of input data. But as discussed in Chapter 10, there are many other possible sources of revenue for ACSes.

[367]One explanation might be changes in the real per-album prices faced by consumers; another might be a wave of early adoption of cassette players (prior to 1980) which was followed by a period of more cassette purchases and fewer consumer electronics purchases; it is also possible that there were changes in methodology or completeness for the CEA data during the earliest years of its collection. These hypotheses cannot be tested without further data; if such data was available, it would represent very valuable future work.

implications depend on which explanation is true. In the worst case (from the point of ACS efficiency), if a bureaucracy had simply mimicked the 1980 ratio, the average annual deviation from market production levels would have been between 13% and 34%.[368]

It is often thought that government bureaucracies should not be placed in a position to determine important economic variables because those institutions will do a poor job of aggregating all of the decentralised knowledge held by private parties that is pertinent to a good setting for the variable in question. This may be true in general, but if we restrict our attention to the narrow case of the annual production levels of entertainment goods, it appears that accurate indicative proxy variables — such as hours spent paying attention to those entertainment goods, or dollars spent on complementary technological devices — exist and would allow bureaucracies to respond to essentially the same constellation of individually hard-to-observe events that the market does.

In conclusion

The indicative variables discussed above are simply economic observations that could be made about entertainment marketplaces, and which greatly simplify the task of predicting their size. There are other strategies that might be employed in trying to determine the size of a virtual market, which make more direct reference to users' financial willingness to pay for the goods funded by an alternative compensation system. These were discussed in Section 5.3.2, and include contingent valuation surveys as well as mechanisms — which appear more frequently in the theoretical economic literature than in policy proposals — that explicitly call on citizens to vote on how much money they want to spend on particular public goods.

Given the surprising predictive accuracy of a proxy variable like audio CE sales, it does not seem that a virtual market would need to turn to these more direct measures of preferences in order to act responsively in determining total levels of funding. It might be sufficient to determine the total tax rate using a proxy variable, and then employ contingent valuation in an attempt to detect moments in history (such as the one that appeared to occur between 1983 and 1985) when the relationship between the proxy variable and real demand for entertainment goods appeared to have changed slightly.

---

[368]The averages vary widely depending on which estimates for boombox sales one uses in the 1980 year; see Figure 7.6.

It is somewhat surprising that the problem of determining total levels of production in alternative compensation systems does not offer strong grounds for preferring copyright regimes. But, when the details are examined, the superficially reasonable presumption that markets are the socially preferable institution for determining the total volume of resources dedicated to cultural production no longer seems correct.

7.5.2 Transparency

One peculiar property of copyright marketplaces as information processing mechanisms derives from the inherently non-transparent nature of many copyrighted goods: Alice can't really tell what they are, until she has "consumed" them.[369] This fact about copyright subject matter leads to a certain amount of inefficiency in copyright markets. If consumers have to pay for a piece of writing before they read it, for example, then they will be signalling their anticipated valuation, rather than an actual valuation. The two quantities are not necessarily the same.

To some extent, lack-of-transparency problems are attenuated by the provision of samples: chapters of books, trailers for films, or an imperfect way to listen to songs before buying them (such as the 30 second clips available on iTunes, the headphones provided by record stores, or perhaps even radio broadcasting). The effectiveness of these strategies varies from medium to medium, but none of them are ideal. The situation in the music marketplace is probably the best and most transparent. It is still common for people to regret their CD purchases, but new sampling and recommendation services are reducing that problem. The situation for books and films is less fortuitous.

Rob and Waldfogel (2004) have collected survey data that beautifully illustrates how — even in the recorded music market, perhaps the most transparent copyright market — consumers' purchasing decisions are highly inefficient despite the availability of samples. They have collected data comparing the *ex ante* anticipated valuations and the *ex post* informed valuations for a set of albums purchased and pirated by U.S.

---

[369]See (DeLong and Froomkin 2000) (emphasising the importance of transparency in modern information economies); (Arrow 1962, at 615) (highlighting similar problems in markets for industrial information); (Takeyama 2002) (arguing that these effects should be counted against strict copyright systems in economic analyses). A similar concept is captured by Nelson's distinction between "experience" and "search" goods; see (Nelson 1970). Kretschmer *et al.*(1999) go further, arguing that copyright works are "credence" goods which cannot be evaluated reliably even once they have been consumed. There is some truth to this claim, although it is difficult to be precise about the normative implications which might ensue.
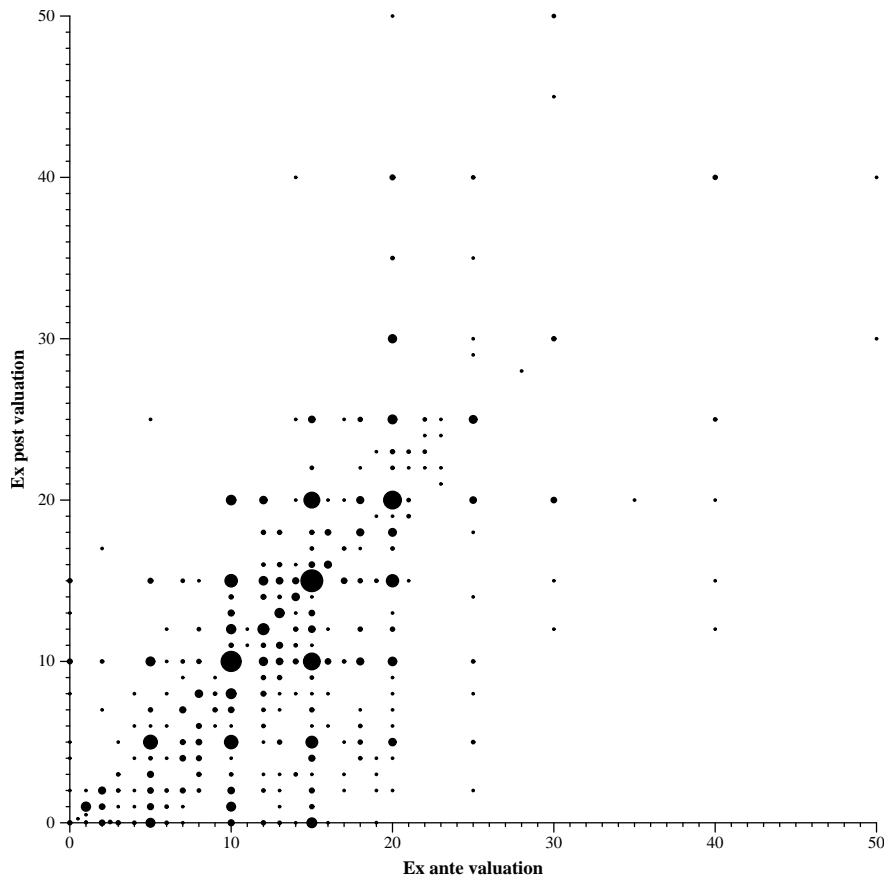
Figure 7.7: *Ex ante* vs *ex post* valuation for albums.
Created from the same data as (Rob and Waldfogel 2004, Figure 4).
In this depiction, the area of each data point illustrates its frequency.

college students. Their data establishes several facts. Firstly, music consumers do not make their choices well. The correlation between *ex post* and *ex ante* valuations was only 0.63.[370] Secondly, the value of the average album depreciates over time — in many cases the audience is 'consuming' an album by listening to it until they begin to tire of it.[371] Thirdly, they confirm that consumers are more likely to pirate albums when their anticipated valuation is low, and more likely to buy when they expect a highly enjoyable product.[372] Fourthly and most importantly for the present argument, consumers err *systematically* in their anticipation of value: there are some artists whose albums are consistently over-valued, and others whose are consistently under-valued.

---

[370] See (Rob and Waldfogel 2004, p. 24); the actual data is shown in Figure 7.7.
[371] See *id* p. 25.
[372] See *id*, Table 3.

artist	adjusted depreciation	Pleasantly surprised	Grew on me	Familiar before got it	Guessed right	Disappointed from start	Got tired of it	N
RED HOT CHILI PEPPERS	31.3%	30.0%	43.3%	26.7%	10.0%	0.0%	13.3%	30
BEATLES, THE	26.2%	25.9%	22.2%	33.3%	0.0%	0.0%	22.2%	27
JONES, NORAH	17.4%	35.3%	58.8%	11.8%	5.9%	11.8%	11.8%	17
U2	9.8%	18.8%	43.8%	18.8%	6.3%	6.3%	18.8%	16
LINKIN PARK	9.8%	37.0%	14.8%	29.6%	7.4%	7.4%	14.8%	27
DION, CELINE	9.7%	18.8%	0.0%	25.0%	12.5%	12.5%	31.3%	16
COLDPLAY	6.3%	25.0%	27.8%	27.8%	22.2%	0.0%	8.3%	36
2 PAC	1.9%	10.7%	25.0%	21.4%	7.1%	0.0%	42.9%	28
EMINEM	0.7%	22.9%	21.4%	21.4%	8.6%	0.0%	31.4%	70
SOUNDTRACK	0.0%	28.3%	23.9%	19.6%	6.5%	6.5%	23.9%	45
DOORS, THE	-1.2%	12.5%	12.5%	25.0%	12.5%	6.3%	37.5%	16
MATTHEWS, DAVE BAND	-5.2%	7.1%	19.0%	26.2%	14.3%	19.0%	26.2%	42
MOBY	-6.2%	5.9%	5.9%	47.1%	5.9%	17.6%	17.6%	17
MAYER, JOHN	-8.7%	31.3%	25.0%	18.8%	6.3%	6.3%	25.0%	16
NELLY	-17.2%	15.8%	21.1%	15.8%	5.3%	5.3%	31.6%	19
DESTINY'S CHILD	-20.2%	0.0%	6.7%	20.0%	20.0%	6.7%	46.7%	15
50 CENT	-20.5%	17.4%	21.7%	17.4%	8.7%	0.0%	39.1%	23
'N SYNC	-20.7%	4.8%	0.0%	4.8%	4.8%	14.3%	71.4%	21
AGUILERA, CHRISTINA	-21.6%	10.3%	0.0%	10.3%	13.8%	10.3%	55.2%	29
BLINK 182	-23.0%	6.7%	13.3%	6.7%	6.7%	26.7%	46.7%	15
CAREY, MARIAH	-23.5%	0.0%	10.5%	26.3%	5.3%	31.6%	26.3%	19
BACKSTREET BOYS	-24.5%	0.0%	0.0%	17.6%	5.9%	11.8%	64.7%	17
DMX	-24.8%	10.5%	5.3%	21.1%	10.5%	10.5%	52.6%	19
SPEARS, BRITNEY	-28.3%	0.0%	3.4%	6.9%	3.4%	3.4%	82.8%	29
JA RULE	-48.4%	6.7%	0.0%	26.7%	0.0%	26.7%	40.0%	15
JAY-Z	-52.4%	17.4%	0.0%	8.7%	17.4%	26.1%	34.8%	23
VARIOUS	-86.9%	10.0%	0.0%	5.0%	5.0%	10.0%	45.0%	15

618

Figure 7.8: The difference between *ex ante* and *ex post* valuation for music varies systematically by artist. Copied from (Rob and Waldfogel 2004, Table 12).

Note that "adjusted depreciation" is not ordinary percentage change but instead $\log_e \frac{ex\ post}{ex\ ante}$ valuation. This means that a 100% appreciation is actually an increase by a factor of *e* (2.7183...) while a -100% appreciation is multiplication by $1/e$.

Systematic errors in consumers' *ex ante* judgement have serious normative implications, because they mean that market incentives are inefficiently favouring 'superficially appealing' works over those with enduring aesthetic merit. A table illustrating this effect is excerpted in Figure 7.8. Albums by artists in the middle of the table, with adjusted depreciations close to 0, receive a payoff that is proportional to the social value of the work. The artists at the top of the table are being paid proportionally less, while those towards the bottom are selling a lot of CDs that the purchasers regret having bought. In both cases, the deviations mean that publishers and artists have incentives to develop poorer-quality works than audiences would ultimately prefer.[373]

A measure of the 'accuracy' of the consumer anticipations observed in Rob and Waldfogel's survey can be found by taking the average of the absolute values of the depreciations in Figure 7.8. That mean deviation is 18.5%,[374] which means that, if a random college student is asked about a random (popular) album in their collection, we would expect an 18.5% difference between *ex ante* and *ex post* valuations. Even though a portion of this variation should be accounted for as "consumption" of the work (Alice enjoys it until she's sick of it), these systematic deviations are grounds to favour public funding systems that base rewards on ongoing measurements of value, over marketplaces that force consumers to decide on their willingness to pay in advance, or over public funding systems that only use *ex ante* measures such as download counts.

Lack of transparency creates a degree of inefficiency in exclusionary copyright systems of all sorts, and sampling only partially alleviates the problem. It should however be noted that strong DRM makes it easier to offer samples (such as difficult-to-copy streaming media, time-limited copies of songs or self-destructing DVDs[375] and could even theoretically use these to overcome the *ex ante* vs. ongoing valuation problem.[376]

---

[373]Because of the criterion for inclusion in Figure 7.8 (15 album owners in Rob & Waldfogel's survey), we can be fairly sure that all of the projects in the table were profitable investments for their record labels. Even the performers are likely to have been well-paid. But, for each of those artists, many less successful acts were signed/created. The relevant incentives are for that entire portfolio of investments.

[374]Or 17.0% as a log ratio deviation. Because the deviations in Figure 7.8 are log ratios, the usual arithmetic mean can be meaningfully employed. The percentage has been weighted by the sample frequency for each artist.

[375]See `http://www.theregister.co.uk/2004/06/01/self_destructing_dvd/` ; `http://www.engadget.com/2008/06/02/staples-to-stock-flexplay-self-destructing-dvds/`.

[376]This would involve renting songs to consumers before selling them. The initial price would be low, but not zero (because the seller would want to charge for the passing enjoyment that comes with a new song) and then would rise over time, to select out those consumers whose ongoing valuation turns out to be higher than they had expected *ex ante* (the ones that are above the $y = x$ line in Figure 7.7). Unfortunately, sellers only have an incentive to improve transparency in cases that benefit them; they would probably not adopt the strategy because they are benefitting from many of the points below the

Most of the business models that could operate in a deregulated, anarchistic informa-tion economy also suffer from inefficiencies due to low transparency.[377]

Because a virtual market based on usage or voting allows users to decide if they like things after they have experienced them, it is capable of providing incentives which are more closely attuned to consumer's preferences than those under the other three regimes. As Figure 7.8 shows, these incentives would sometimes be noticeably greater, and in other cases much lower, than the incentives provided by a copyright market. The average difference would be somewhere around 18.5%. This advantage of public funding is a little greater over pragmatic DRM and anarchy than it is over strong DRM (because strong DRM allows for more extensive sampling).

7.5.3 Do consumers have incentives to create the right incentives for producers?

As de Trenqualye showed, virtual markets and similar systems can be made incentive compatible for consumers[378] without complicating their basic design. However, the difference in payoffs from truth-telling and dishonest voting is very small. Suppose that Alice really enjoys Metallica's music, but it also happens that she'd like some other band to receive more money for reasons unrelated to her relative preferences for the music they record. She might prefer another band on political grounds, for instance.

What does this entail? In a virtual marketplace, she could download Metallica's music while giving all of her votes to the other band, yet suffer negligible personal consequences; although there would be a minuscule reduction in incentives for the artists she likes (Metallica), it would only hurt her in the unlikely event that the loss of her dollars led to the cancellation of Metallica's recording contract.

In contrast, in a strong DRM marketplace, Alice cannot redirect money from Metal-

---

$y = x$ line in Figure 7.7. The relevant information revelation games are however beyond the scope of this thesis.

[377]One possible exception is the Wall Street Performer Protocol (see Section 3.2.3; Rasch 2001), where the specifiable and functional nature of software or other non-cultural information goods allows a degree of *ex ante* transparency.

[378]See Section 7.4.2 above. Truthful revelation is only optimal at Nash equilibrium and is not guaran-teed to be a globally dominant strategy in virtual markets, but it is in truth extremely improbable that any consumers or groups would ever have enough information to engage in strategic voting with the intention of changing *others'* votes.

lica to the other band without losing her copies of their albums.[379] If she refuses to honestly reveal her preferences, then she doesn't get to enjoy the works in question. Her incentives follow the same pattern as in the virtual market, but they are much stronger. If we regard each of her actions as sending a message about her preferences, then formally, her preference ordering over messages is the same, but there is a much greater difference between the payoffs or utilities for different messages under DRM regimes.

The dollars Alice pays to access works in copyright marketplaces appear to act as a limited sort of 'guarantee of sincerity' for her messages. In this respect, we can be relatively confident that an exclusive rights–based marketplace is reporting the public's preferences accurately.[380] The guarantee appears to be weaker in the virtual market case.

One should also acknowledge a subtle interaction between the artist and the audience which affects preference disclosure. In much the same way that successful street performers cajole their audience into making payments, or some recording musicians discourage their audiences from free-riding with pirated music downloads, many artists working in a virtual market would no doubt encourage their audiences to vote "early and often". It could be argued that this creates an incentive distortion which disadvantages artists who are either unable or unwilling to guide their audiences in this manner. Although this argument is certainly valid, it applies almost as extensively to copyright-based markets as it does to their virtual alternatives.

The weak penalty for dishonest votes in alternative compensation systems has led some observers to caution that alternative compensation systems could allow various organised groups to "game the system" and collect votes from their members for unintended purposes, such as political organisation.[381] This could be done under all sorts of covers: a band might announce that they were donating most of their income to Sinn Féin, and Sinn Féin might then make an effort to encourage their supporters to report listening to, or vote for, the band's music as much as possible (regardless of whether they actually enjoy and listen to it).

---

[379]Note that in weak DRM marketplaces, the situation is somewhere between these two extremes. In that regime, Alice might decide to pirate Metallica's albums while buying some other band's albums.

[380]This is not a rigorous guarantee, especially under copyright of the more crisis-bound contemporary kind; if piracy was more prevalent amongst certain demographics — teenagers being an extreme example — their preferences might receive a discounted response from the market. There have been some claims that this already occurring in music markets (Davis 2002). And conversely, no doubt, one could find examples of people paying for things which they do not really want.

[381]Private mailing list postings, on file with the author.

The most overt instances of vote diversion of this sort could probably be defeated by a prohibition on the donation of royalties to NGOs or political parties in cases where consumers would have had any knowledge that that was where they would end up. But there would always be a handful of borderline cases that blur the line between creative labour and other civil society activities.

The risk of vote diversion would be further reduced by the use of sample populations in place of a universal voting or measurement system for a virtual market. Alice faces similar incentive structures whether she is a member of the sample population or the public at large, but there is some prospect of screening to ensure that participants in a sample group are not deliberately trying to achieve unanticipated ends through their listening/voting habits.

An interesting observation to make on the subject is that it may be rational for Alice to send royalties to Sinn Féin if she supports them, but only up to a point. The political activities that Sinn Féin promises to spend their royalties on are a public good, much like the cinema that Alice's favourite film-makers are thinking of producing. She would make a tradeoff between them. So it turns out that if virtual markets solve the free rider problem for copyright goods, they could accidentally end up solving part of the collective action problem at the same time.[382] It would only be the employment of careful regulations of the sort mentioned above that would keep these two solutions separate.

While it does not appear that there are grounds for regarding the role of non-payment signals as seriously problematic in a properly implemented alternative compensation system, it remains the case that DRM systems — absent piracy, at any rate — perform this particular task optimally and without so much regulation. This constitutes a reason, albeit a rather marginal one, for preferring the latter system.

7.5.4   What signals can a person send in each regime?

In each of the possible regimes, Alice gets to choose from a 'menu' of informational signals that she can send to the economy with respect to any given work. Previous sections have already discussed the question of incentive compatibility — whether Alice will *want* to send the right signal to the economy. There may also be a question

---

[382]The *collective action problem* is essentially the analogue of the free rider problem in politics. Ordinary people have too little incentive to contribute to political endeavours because they receive only a minuscule fraction of the benefits of their actions. See (Olson 1971).

of whether it is *possible* for Alice to send the right signal, because perhaps the signal she should send is not on the menu of options at all.

The chief piece of information which the economy wants from Alice is her valuation for each work — how much she is willing to pay for it. But it is only under the anarchic regime, with voluntary donations as a primary method of payment, that Alice can truly send any signal she wants for this quantity.

In the copyright regimes, Alice sends her signals through her purchases, and she will often only have a binary choice: whether to buy the work or not. Sometimes, she may have a choice to purchase the work at a small number of discrete prices (if there is price discrimination in operation). If the work is out of print, she cannot send any signal at all.

In a virtual market, Alice can choose to send any signal from zero (not using or voting for a work at all) through to some maximum (listening to a song 24/7 or giving all of her votes to it). This is an approximately continuous spectrum, which is clearly more desirable than the small number of discrete choices in a market. On the downside, however, virtual market signals are subject to a budget constraint, which is that Alice as an individual can never send a total signal larger than her portion of the alternative compensation system's supply of royalties. People buying copyrighted goods are subject to budget constraints too, of course, but they are often much looser in the sense that many people could afford to spend a lot more money on copyright goods than they currently do.[383]

On an individual level, it should be clear that neither of these menus of prices is all that close to optimal. With respect to copyright, it is only if a rights holder happened to choose a price which was close to but below Alice's maximum willingness to pay that her signal could be accurate. With respect to the virtual market, Alice's signal will generally be accurate only if her preferred budget for works was close to the one the ACS actually used.

Having said this, the relevant consideration is not how accurate Alice's signal is but what what happens when her signal is combined with those from the rest of her society.

In the virtual market, if it is the case that people who would have spent more money on copyrighted works have approximately the same distribution of preferences as the people who would have spent less money, the total signal will be the right one, though

---

[383]The inequality effects discussed in Section 6.4 could be characterised as a utilitarian inefficiency in copyright marketplaces which essentially derives from budget constraints.

Alice is above or below her ideal signal because of her budget constraint. But if there is a profound connection between tastes and total willingness to pay for entertainment, and especially if the differences in spending are a result of different levels of interest in or appreciation for entertainment goods (that is, *not* resulting from differences in disposable income), the virtual market could over-incentivise the production of goods that low-valuers want, and under-incentive the production of the goods that high-valuers prefer. This is, in a sense, a central problem of replacing democracies of dollars with democracies of individuals.

The marketplace also suffers from inefficiencies as a result of the menu of signals. The phenomenon, discussed in Section 7.5.1, of markets under-producing marginally valuable copyrighted goods because of the inappropriability of deadweight loss and consumer surplus, could be considered one consequence of the limited price menu.

Another problem that the marketplace suffers is the need for producers to estimate which prices to offer to place on the menu in the first place. In theory, they might select prices to conform to the demographics and culture of the audience for each work they release. But at least some groups of copyright holders seem unable or unwilling to do this,[384] probably because they lack the knowledge to choose those prices well on case-by-case basis.

The fact that sellers must as a matter of practicality use limited pricing strategies reduces the quality of the signals that Alice can send about which works she likes, which in turn reduces producers' revenue and profits, which in turn reduces incentives to invest in future creative projects, which in turn reduces Alice's benefits from those future projects.

Without further data on the issues discussed here — correlations between spending on entertainment goods and particular taste profiles; measurements of typical differences between willingness to pay and prices — it is not possible to say whether signal menu limitations favour one regime or another. Data collection of that sort would constitute valuable further work.

---

[384]Even after iTunes began offering tracks at price points of 69, 99, and 129 U.S. cents, very few rights holders chose to use the 69 cent price to any significant degree.

7.5.5   Network effects and contests over the definition of 'cultural space'

Let me now turn to an interesting and rather subtle argument that the role of advertising in cultural industries may create problematic incentives for publishers. Copyrighted works play a crucial role in defining and manipulating something that I will term 'cultural space' — the *lingua franca* of those interactions between individuals and amongst social groups which are mediated by culture. Mark Nadel has pointed out that that perfectly natural role may, indirectly, lead to a wasteful dissipation of resources through marketing contests (Nadel 2004, Part II.B). If this is correct, it would be helpful to know whether alternative compensation systems could avoid such problems.

The essence of this argument is that many cultural goods create a system of "network externalities" or "network effects" amongst their audience.[385] These network externalities are more subtle than those that attend tools like telephones and email clients (whose usefulness is in some sense proportional to the number of people who have them), but they are equally real. Human social interactions are filled with references to both popular and niche culture, and our perspectives on the word are unavoidably coloured by the art and entertainment we consume. Many adults read Harry Potter novels, rather than other works of fantasy, so that they can follow dinner party conversations, and not because they expect to enjoy those particular books more. It follows naturally that building and exploiting networks around their products is an valuable strategy for copyright owners.[386] So, as Nadel points out, the existence of cultural externalities results in marketing contests over the definition of cultural space; there are many tunes which are capable of capturing the human psyche, but only a few of them will top the charts.[387]

To be clear, the problem at hand is not that some works become enormously successful. Under any non-censorial policy regime, there will always be certain high-

---

[385]See, for example (Kretschmer, Klimis, and Choi 1999, S63).

[386]On the nature of these strategies, see (Kretschmer, Klimis, and Choi 1999, pp. S67–S69).

[387](Nadel 2004); compare (Adler 1985) (modelling a similar same phenomenon, but using imperfect information, rather than preference-altering externalities *per se*) ; (Salganik, Dodds, and Watts 2006) (confirming empirically that factors other than inherent quality are causally determinative in the success of cultural works). Adler only identified "luck" as the tie breaker in these contests, while I find Nadel's nomination of advertising more persuasive with regard to real-world marketplaces. Compare also Towse (2001, p. 88), who briefly suggests some difficult empirical tests to distinguish Adler's theory from several other explanations of stardom proposed in the literature (Rosen 1981; MacDonald 1988; Towse 1993). A more rigorous examination of Towse's criteria, and the addition of an extra one to test Nadel's hypothesis, would be a useful contribution to the literature.

quality works which are very widely watched/read/listened to. Instead the problem is that the enormous rewards available for producing a work that becomes culturally significant will lead to enormous mutually-cancelling expenditures on advertising. That kind of advertising does not make society any better off, because the quality of the most successful works would have been the same without it.

These expensive advertising contests are in some respects similar to the "race to invent" or "common pool problem" discussed in the economic literature on patent systems.[388] There is no general solution to this problem unless the institutions which provide financial incentives can somehow obtain and respond to comparative information about all of the players competing for each niche in cultural space.[389]

On one level, exclusionary copyright and publicly funded rewards appear to suffer equally from races to define cultural space. Because virtual markets mimic real 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 — a level playing field of more works, the best of which would compete and evolve to define 'cultural space'. In mathematical terms, optimal cultural production can occur only when the returns on information goods are not just 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. It is not at first obvious how institutions for financing creativity could achieve this, and I am not aware of any proposals to address this problem in a rigorous and formulaic manner.

There is however a less-than-obvious property of those alternative compensation systems that allow users to vote independently of consumption, which may mitigate inefficient contests over cultural space. 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. This is a behaviour that some people would be inclined to engage in naturally ("I'll vote for independent bands as a matter of principle; even though I enjoy Britney's music, she's doing fine as it is"), but it could also be explicitly encouraged.

The extent to which voter behaviour might succeed in short-circuiting marketing contests and the extent to which encouragement would be useful depend, of course,

---

[388]See, e.g., (Wright 1983) at 691, 693 (describing this effect, and surveying relevant comments in the literature).

[389]Cf. (Wright 1983, at 694) (identifying conditions for a solution to the race-to-invent problem). Wright's institutional solutions are not directly applicable to the present problem, because the roles of uncertainty and quality are different in the two kinds of market.

on the way that individuals make these tradeoffs. A calculation of the size of these effects would require good data on the distribution of advertising expenses for a range of works, and a model (theoretical or empirical) of how voting audiences would treat those works. Such work is beyond the scope of the present inquiry, but a methodology that provided numbers would be particularly valuable further work. In any case, it is apparent that the option given to audiences by virtual markets, to allocate remuneration that is not in strict proportion to consumption, could well reduce advertising inefficiencies in copyright industries.

7.5.6   Conclusions on incentive-generating information

Section 7.5 has studied a set of artifacts and imperfections in the way that markets, or virtual market alternatives, form a link from the audience's preferences to incentives for entertainment producers. Many of these phenomena have been discussed or even modelled in the literature, but not together. This Section has examined each of them in turn, and attempted where possible to derive numerical estimates for the size of the errors they introduce between socially optimal incentives and those actually produced by a regime. Table 7.1 summarises these observations.

Section	Sub-issues	Regime favoured	Approximate error (mean deviation) in incentives	Other notes
Market size (§7.5.1)	Limited appropriability of social surplus	Virtual markets (slightly)	Up to 50% for marginal works; less for others	ACSes do not get this error down to 0% unless they pay more than © markets
	Can ACS bureaucracies replicate market trends?	Copyright regimes	8 % (1985 baseline) 13–34% (1980 baseline)	-
Transparency (§7.5.2)	-	Virtual markets with *ex ante* data	18.5%	-
Incentives for sending correct signals (§7.5.3)	-	Copyright markets	Unknown (requires experimentation)	More serious for sector-specific ACSes than broader ones
Ability to send correct signals (§7.5.4)	-	Unknown	Unknown	Deserving of further research
Contests over cultural space (§7.5.5	-	Virtual markets with voting	Unknown	Deserving of further research

Table 7.1: Information → incentive interactions differentiating copyright from virtual markets

An attempt to tally up the table can be made as follows. The third and fifth rows, corresponding to Sections 7.5.3 and 7.5.5 respectively, favour copyright and ACS regimes in turn, but to an unknown degree. Further work would be required to properly account for these effects, but for the time being they must be regarded as a "wildcard" effect in favour of each regime.

Section 7.5.4 discusses a phenomenon which may favour either regime, and cannot be counted in favour of either of them without further research.

The second and third rows are more or less commensurate: the 18.5% expected error caused by *ex ante* payments under copyright are more or less equivalent to the 8–34% expected error in the size of a virtual market based on extrapolation using indicative proxy variables.[390]

The "up to 50%" number in the first row is not directly comparable to the percentages in the two rows below, because it is neither uniform across all copyright works nor entirely avoidable under any of the regimes, although the virtual market would mitigate some of this problem if some of the surplus from the removal of artificial scarcity went to producers. Practical attempts to reduce this inefficiency further — such as redistributing some royalties to low-income artists — might reduce the number further, but the subject requires research.

Either way, when the rows of the table are summed, this last information efficiency effect appears to give the virtual market regime a slight expected edge over copyright in the quality of the incentives it produces. This result can be tied back to the simple model developed in Section 7.3.5.

Let us return to Inequality 7.18. Establishing a small expected advantage for the incentives created by the virtual market implies that γ is positive in that equation:

$$\frac{T}{\alpha A} < \gamma + 0.55 \tag{7.18}$$

In other words, if the costs of implementing a virtual market in terms of taxation,

---

[390]For most of the years during the 1980–2000 observation period, the introduction of an alternative compensation system of a size derived from indicative proxy variables would have produced an error in the order of 8%. But had the system been introduced between 1980 and 1983, the error would have been higher. In principle, data from other sources, such comparative analysis of contingent valuation surveys and virtual market votes, might be able to detect shifts like the one that occurred between 1983 and 1985. In practice, it would be safer to assume that alternative compensation systems would fail to adjust for such shifts. Even so, the overall expected error in market size across all years remains comparable to the 18.5% caused by the lack of transparency in copyright markets.

infrastructure and transaction costs[391] is not extremely high — at least, more than 55% of the social value of the music industry — then virtual markets are preferable to the status quo.

7.6   Conclusions

This Chapter has attempted to determine which of the digital copyright regimes would furnish us with a better supply of cultural and entertainment goods. The analysis required to answer that question is more complicated than those in other chapters of this thesis. Firstly, it is necessary to say what makes one supply of works "better" than another. Then it is necessary to grapple with the fact that legislated incentives are only partially determinative of what will be produced. Finally, it is necessary to determine how the incentive-creating structures actually vary across the regimes.

Section 7.2 examined several competing notions of the value of copyrightable things, and concluded that the idea that incentives should be generated by reference to consumer preferences makes far more sense for 'entertainment' than it does for 'art' that isn't also entertainment. Questions about what institutions should exist to promote art in that stronger sense are important, but not the subject of the present inquiry. Willingness-to-pay measurements in dollars are a workable first approximation for the value of entertainment goods, although corrections for identifiable effects such as social inequality, lack of transparency, and cultural externalities are necessary.

Section 7.3 scrutinised the causal connection between payment and cultural production, which is usually taken for granted in the economic analysis of copyright. It was concluded that, within entertainment industries, the incentives created by copyright legislation are less important than often supposed. Two other kinds of motivation — "intrinsic" non-financial motivation, and "coincidental" financial motivation — were discussed and shown to be significant in many entertainment sectors. Simple mathematical models were introduced which allow some quantification of how the incentive benefit of copyright trades off against copyright's artificial scarcity cost, if the strength of intrinsic and coincidental incentives is used as an input variable. Surprisingly, the anarchic regime might almost be competitive with copyright regimes on narrowly utilitarian terms in industry sectors where those non-copyright motivations

---

[391]Taxation is definitely a net cost for virtual markets relative to the status quo, but infrastructure is more ambiguous and transaction costs are probably a net benefit; see Chapter 8.

are strong. The model for the virtual market regime is more complicated, but it also shows that incentives in ACSes would have to be significantly worse than in copyright regimes before the copyright regime was preferable.

The subsequent sections then turned to the question of whether ACS incentives would, in fact, be worse than those provided by copyright marketplaces.

Section 7.4 reviewed the theoretical economic literature on the ability of governments to provide efficient incentives for the production of public goods. Two models in particular stand out: Shavell and van Ypersele's (2001) comparison of the efficiency of distribution and incentives under patent and reward systems; and De Trenqualye's (1997) results on the efficiency of iterative voting for choosing a budget for many public goods. De Trenqualye's result does not superficially apply to ACSes like virtual markets, but an interpretation was proposed under which his efficiency results do in fact apply. Both papers give grounds for believing that public funding can provide incentives that are to first order at least as good as those created by markets.

Further consideration indicates that there are many informational effects at work in both copyright marketplaces and the alternatives, which have not yet been unified in simple, tractable mathematical models — many of the effects have been modelled, but not together. A number of these phenomena were discussed in Section 7.5. Some of them (transparency, contests over cultural space, and the problem that works of small-but-positive social value are not profitable to make) clearly favoured alternative compensation systems, provided sufficiently good sources of data were used. Some other phenomena clearly favoured marketplaces (the need in an ACS to involve bureaucracies in setting market sizes, and the question of whether individual audience members have incentives to send the right signals). When summed, it appears that virtual markets offer incentives of slightly better quality than real marketplaces!

That conclusion made it possible to refine the mathematical model introduced in Section 7.3. It shows the tradeoff between the scarcity and incentive benefits of virtual markets, on one hand, and the taxation and other overhead costs of introducing it, on the other. It remains necessary to provide an estimate for those taxation and other costs, and that will be the subject of the next few Chapters.

How should this Chapter's conclusions about incentives be understood in Hayekian terms? Hayek was correct in emphasising that markets should be seen as information processing systems, passing data about people's needs along to those engaged in production, and offering producers incentives to meet those needs. Entertainment in-

dustries possess a number of unusual characteristics — especially with regard to the measurability of the value of the goods produced by those industries — which make it conceivable that the incentives could be provided as well or better by other forms of institution. On its own, this fact would be insufficient motivation to rush about proposing that entertainment markets be replaced by alternative compensation systems. But, when coupled with the fact that alternative compensation systems make it possible to avoid the artificial scarcity of digital public goods, a more pointed policy argument begins to emerge.

Chapter 8

Transaction Costs and the Anticommons

Contents

To endorse Google's library initiative is to say "it's OK to break into my house because you're going to clean my kitchen," said Sally Morris, chief executive of the U.K.-based Association of Learned and Professional Society Publishers. "Just because you do something that's not harmful or (is) beneficial doesn't make it legal." (AP 2005)

While its competitors went through the "painstaking" and "costly" process of obtaining permissions before scanning copyrighted books, "Google by comparison took a shortcut by copying anything and everything regardless of copyright status." [citation to Microsoft] *As one objector put it: "Google pursued its copyright project in calculated disregard of authors' rights."* — Authors Guild et al. v. Google Inc.[392]

The central idea of copyright is that people cannot do certain things to works — copy them, publish them, change them — without the permission of rights holders, be they authors or publishers. Commerce has its ways and means, and the power of exclusion is often used to obtain payments for the usage of works, rather than preventing them altogether. One inconvenience in this notionally elegant ordering of affairs is the difficulty of organising "permission" — a task that includes the cost of potential users and rights holders finding each other, agreeing on license terms, and making payments.

These overheads are called "transaction costs." They are of interest to us, because they are likely to be systematically different under different digital copyright regimes. Strong DRM might somehow shrink them by automation. Virtual markets might do away with them altogether. Anarchy might remove them in some cases, and increase them greatly in others. This chapter investigates the comparative burden of transaction costs across each of these policy options.

8.1   A Definition

Before examining the role of transaction costs in copyright systems, it may be helpful to consider how they are defined and what they encompass. Given how much has been

---

[392]See `http://thepublicindex.org/docs/amended_settlement/opinion.pdf`. In this decision Judge Chin rejected the class action settlement between Google and authors' and publishers groups, which would have allowed Google to proceed with the scanning of books whose copyright holders could not be located, and created a collective licensing system that would (among many other things) have permitted Google to offer subscription access to the full text of all of the books it had scanned. Even once Sally Morris's colleagues had seen the wisdom of doing things that way, the copyright system did not permit it.

written about them, they are a notoriously slippery conceptual device.[393]

From the perspective of clarity, things got off to a bad start when Coase, who is usually credited with the introduction of transaction costs into economic theory, used the potentially enigmatic "cost of using the price mechanism"[394] to explain why some tasks are aggregated within particular firms, while others are organised across marketplaces. Later, the literature bifurcated into two camps. There was agreement that transaction costs were of some importance, but not on how to define and understand them.

Allen associates the two schools of thought with what he terms the "property rights" and "neoclassical" definitions of transaction costs. Both definitions relate to property (note however that property rights are here defined not as specifically legal entities but rather in terms of the economic practicality of access and control[395]). Each of them has its analytic and explanatory advantages, and a selection should largely depend on the question one is attempting to answer.

In the former, broader, "property rights" definition, a transaction cost is any cost associated with "establishing and maintaining property rights."[396] Most of the law and economics literature employs this definition. A number of factors are often seen as contributing to these costs — locating goods and their owners, negotiating contracts, monitoring performance in those contracts, enforcement, and the inefficiencies caused by incomplete contracts. The property rights definition of a transaction cost is very useful for explaining institutional structures and for studying the efficiency consequences of different allocations of property rights. Unfortunately, it is particularly unhelpful for studying copyright; it aggregates so many different issues together that (in our case) it threatens to swallow whole the comparison between digital copyright systems and alternatives.

The latter, narrower, "neoclassical" definition of a transaction cost is limited to the expenses of "*transferring* property rights".[397] Because the property in question is eco-

---

[393]The relevant chapter of the *Encyclopedia of Law and Economics* (Allen 1999) provides both an introduction to the very extensive research on transaction costs, and a detailed exposition of why their scope is so ambiguous.

[394](Coase 1937, p. 38)

[395]In the economic definition, a property right exists when and only when it is practically enforceable. A squatter who is charging for tickets to see an unlicensed film he projects in his living room is exercising economic property rights in the house and the film performance, but probably not legal property rights.

[396](Allen 1999, p. 898–9)

[397](Allen 1999, p. 901) (emphasis added).

nomic, not legal, granting a non-exclusive license is a kind of transfer. This definition is more helpful for our present purposes, because it selects out a set of licensing-related costs which are more (though not completely) separate from the issues discussed in the other chapters of Part IV.

Transaction costs and artificial scarcity

Even having adopted the narrower, neoclassical definition of transaction costs, there are some definitional obstacles to employing the concept in a normative analysis of digital copyright. In particular, there is a partial overlap with the artificial scarcity effects discussed in Chapter 6.

A large portion of scarcity deadweight losses occur because the transactional and informational costs of beneficial licensing exchanges are too high. Those inefficiencies, it could reasonably be argued, should be counted as transaction cost effects. But other aspects of artificial scarcity — especially the inequality effects discussed in Section 6.4 — are not transaction costs at all.

For the purposes of this chapter, I will simply avoid counting any costs twice. The basic losses from prohibitive transaction costs were dealt with in Chapter 6; this chapter considers costs that are paid rather than prohibitive, as well as some prohibitive costs that are not associated with individual works.

Surprising Transaction Costs

Scholars of copyright would not necessarily anticipate some of the strange transaction costs imposed by the present regime.

One very interesting category is the expenditure of time, effort and resources by various individuals and organisations in worrying about copyright law itself. These sorts of expenses have always existed, but they have become much more widespread now that copyright regulation has entered people's homes, and businesses to which it has no obvious connection. Sometimes, these are just overheads of operation: most museums, libraries and cultural institutions have one or two staff working full-time on navigating copyright issues.[398] Sometimes, the widespread deliberation about copyright law leads to absurd results.

A telling example of this inefficiency is the decision, by Wal-Mart and other chains

---

[398]Emily Hudson, personal communication.

of U.S. photo shops, to refuse to print digital photographs if they display a high degree of photographic talent,[399] or to scan antique photographs without permission from the studio that took them.[400] The theory behind the first of these policies is that their customers might be downloading photographs from the Internet and having them printed, or making an end-run around wedding photographer's printing fees. In the vast majority of the cases prohibited by these chains' policies, the customer's request would be entirely legal, but the stores lack the resources to understand and correctly apply the law in each case. The end result is that many constructive transactions are blocked.

8.2   Mechanisms for Reducing Transaction Costs

If there were no processes or institutions working to reduce them, copyright transaction costs would make licensing harder than swimming through glue. Of course, there are societal responses to such strong imperatives. In this section I discuss them, and consider whether they can solve the transaction cost problem — or what they leave out.

8.2.1   Reduction by Market

Copyright's first defence against the charge that it brings with it burdensome transaction costs is the possibility that the marketplace provides appropriate and sufficient incentives for the construction of institutions that in turn minimise these costs. Such minimisation might occur in different ways. If a particular kind of usage is especially common, then the copyright holder may be motivated to unilaterally make that kind of license easily purchasable. If the transactions in questions are more diffuse, and more complicated — and especially if they require some kind of standardisation amongst many rights holders, then the institutional solution may involve concentrating some of the rights and bargaining power over many works in a single place. Concentration has occurred both through the conglomeration of very large catalogues of rights in the hands of a few firms and through the creation of collecting societies (which I discuss

---

[399]See "Snap Judgements" `http://web.archive.org/web/20050609032250/http://www.signonsandiego.com/news/computing/personaltech/20050530-9999-mz1b30snap.html` (discussing chain stores' policy of refusing to print copies of photographs that look too "professional")

[400]See also "Wal-Mart: you can't scan century-old photos of your ancestors because copyright lasts forever" `http://boingboing.net/2008/08/12/walmart-you-cant-sca.html` (discussing refusals to scan 100-year-old family photos because the copyright might be held by a studio).

below)

As a result of sellers' most basic incentives, when it comes to distributing works, the need to formally clear all of the rights required for the use of a work is not a great problem in many situations.[401] A majority of users will fit neatly into one of a few categories; their terms of access can easily be codified, and transaction costs are limited to the cost of purchasing a standard commodity. The marketplace provides strong incentives for rights holders to organise this reasonably well.

So in cases where rights holders can easily anticipate the kinds of licenses that prospective users will be willing to pay for, more-or-less unilateral action by sellers does a reasonable job of overhead minimisation. Most CDs come with most of the economic rights that most users desire; even while DRMed, iTunes sales came with a bundle of legal and economic rights that many users appear to find satisfactory.[402]

There are still some problematic limitations to these market solutions. The worst instances arise when rights holders' desire to get paid is incompatible with the best uses that could be made of a work. Refusal to license motivated by a fear of piracy is one way that this can occur.[403] The problem is even more serious when it comes to indexing, search and analysis tools.

During the first fifteen years of the web, many rights holders have taken the approach of licensing their works exclusively to proprietary search engines (whether those run by conglomerates like Thompson and Reed Elsevier or niche firms like Xrefer) but not through the web to open search engines. There are negative consequences: the usefulness of web search engines is lower, and the creation of barriers to entry in the search business that reduce competition to index works as well as possible,

The marketplace works reasonably well for common denominators (although some transaction costs are still present). It cannot be relied upon to short-circuit transaction costs in more complicated cases.

---

[401]Distribution-side rights clearance is of course very different to production-side rights clearance. Identifying exactly who owns the various rights applicable to a particular work, in order for them to be paid for a performance or derivative work (by a collecting society, a firm running a DRM system, or a virtual market administration) may sometimes be difficult, but these costs apply equally to exclusive rights based and alternative compensation systems.

[402]The "FairPlay" DRM used by the iTunes Music Store limited users to copying tracks to five computers, burning no more than seven copies of a playlist to CD, and not listening to songs on devices other than iPods or computers running iTunes; see `http://en.wikipedia.org/w/index.php?title=ITunes_Store&oldid=303411491#Digital_rights_management` .

[403]For example, until 2007 the music industry refused to give consumers MP3 copies of songs (because they are inherently DRM-free), regardless of whether there were good reasons for consumers to want them. The rationality of these fears varies.

8.2.2 Reduction by Technology

Another reply to concerns about the scope of the transaction costs imposed by copyright is to emphasise that the problem must be minimal, because the power of digital technology has been lowering the transaction costs that affect copyright works. In an argument against compulsory licensing, Robert Merges provides a particularly unambivalent example of this rhetoric:

> ...*Compulsory licensing supposedly addresses the "market failure" of high transaction costs.*
>
> *But markets for digitised works do not suffer from market failures. Furthermore, the Internet has reduced the transaction costs that once served as a key rationale for compulsory licensing. Recent developments suggest that fears of excessive control of digital content are overblown. Without enhancing compulsory licensing, the digital landscape is diverse, as the case of music demonstrates. There is free music, temporarily free music, and low-cost music online. Offline, music companies are lowering the prices of CDs.* [404]

Stepping around the audacious and eminently citable claims about market failure,[405] one finds a more interesting fallacy in Merges' position. There is an explicit claim that transaction costs are falling, and an implicit one that a "diverse" digital landscape is necessarily a sign that the sum of transaction costs must now be small, because many different kinds of transactions are occurring.

Although it may be true that digital technology is lowering the cost of particular transactions,[406] it does not follow that the burden of licensing overheads must be insignificant. The fact that different commercial strategies can be observed in the marketplace provides no proof of this point either: the successful execution of some transactions does not imply that all worthwhile transactions are executed, or that those that are observed have not come with high overheads.

---

[404](Merges 2004, p. 1)

[405]Market failure is defined as any situation in which a market does not function efficiently. Merges' claims here seem rather close to a denial of deadweight losses from artificial scarcity — though in other articles he has set out a more nuanced view of the matter (Merges 1997, Part I.A (predicting that some kinds of transaction costs will be reduced or eliminated online while others will persist); Part III.A (defending fair use as a solution to enduring transaction cost problems).

[406]Though certainly not all of them if the author is using the "property rights" definition (note 396 and accompanying text)! Even under the neo-classical formulation, the need to be mindful of piratical leaks may increase the reluctance of rights holders to make certain sorts of trades.

Counter-intuitively, the total weight of transaction costs may even have risen with the success of the Internet. Digital technology gives with one hand, by reducing the cost of finding a work or licensing it. But it beckons with other, by greatly increasing the number of transactions which are possible in the first place. A smaller transaction cost can in theory even be *more* problematic in normative terms when it is necessarily multiplied across a larger number of feasible transactions.

How does digital technology increase the number of possible transactions? One fairly straightforward way it does this is by allowing consumers to know about and therefore consider purchasing many more works. Another is that it makes the production of derivative works easier. But perhaps the most transactionally complicated opportunities opened up by digital technology are those that involve computer programs which are trying that operate over very large numbers of works — indexing them, making them searchable, and analysing them on a large scale. These tasks produce a kind of gestalt — a whole that is greater than the sum of its parts[407] — which potentially crosses the exclusive rights of tens or hundreds of thousands of rights holders. Or more.

8.2.3   Reduction by Collecting Society

The spontaneous formation of collecting societies in the marketplace provides a distinctive example of how a completely novel institution can spring up to enable a particular flow of transactions. In this case, by explicitly or implicitly moving exclusive rights to a central point — the collecting society — a potential user needs (in theory) to deal with only a single negotiating partner to secure the rights they need for their activities.

Consider a market for an information good in which there are N potential users of works controlled by M different rights holders. If every license must be negotiated directly, the transaction cost burden is proportional to NM. With an established collecting society, the cost is proportional to $N + \alpha M$ (where α is the ratio of the average transaction cost of the collecting society obtaining a set of rights to the average cost of a user licensing them). This can be a very large saving, provided that N is much bigger than α, but there are difficulties to be overcome.

---

[407]This is not a translation of the German word; the meaning was only acquired as a side effect of the word's use within a particular school of thought in the field of psychology. But it useful to have an English word meaning "a whole that is greater than the sum of its parts", so I will use it, possibly even as an adjective.

One is that the cost of establishing the collecting society may be very large (in other words, α is large). This is especially the case when the rights in question need to be obtained retrospectively. Because a new digital right of communication to the public was created by the WIPO Copyright Treaty,[408] existing collecting societies in many jurisdictions do not hold all of the rights necessary for important Internet "gestalt" applications.[409]

Another complication is the transnational nature of modern digitisation, indexing and search facilities. If search engines (for example) had to clear rights to the works they indexed, they would have to conduct successful negotiations not with one collecting society, but with dozens of them. Variations in applicable law, and a daunting array of contracts between collecting societies, authors and publishers, suggest that even if these costs were not prohibitive, they would be wildly burdensome. Fortunately, web search engines appear to have mostly avoided needing to clear rights for essentially the same tenuous legal reasons that allow the web to function in the first place.[410]

8.2.4   Reduction by Fair Use

Wendy Gordon has famously argued that the very purpose of fair use — or at least a large part of it — is to remove exclusive rights from situations where they are likely to cause prohibitively large transaction costs.[411]

Limitations and exceptions to copyright (a very large family, of which fair use is just the most-discussed member) should in theory be a good solution to large scale transaction cost problems. In the United States, fair use has certainly been of use for shielding search engines,[412] though the shield is far from certain or completely reliable,[413] and therefore cannot be regarded as the ideal institutional arrangement.

---

[408]See (WCT 1996, Article 8). Some nations, including the United States, have not implemented the WCT requirement as a distinct exclusive right, but this does not necessarily mean that older copyright assignments in those jurisdictions are inclusive of digital publication rights; see *Random House Inc v. Rosetta Books LLC* 150 F. Supp. 2d 613 (SDNY 2001).

[409]Google Books' creative solution to the general problem of rights aggregation in the United States is discussed in Section 8.3.1 below.

[410]See Section 8.2.5. But see note 413.

[411]Gordon originally set out this argument in (Gordon 1982), but later set out a more nuanced theory about the appropriate scope of market (Gordon 2003), distinguishing amongst other things the important distribution functions served by fair use.

[412]See, e.g., Kelly v. Arriba Soft 280 F.3d 934 (9th Cir. 2002) (holding that making cached "thumbnail" copies of images, in order to provide them as search results, constituted fair use).

[413]See, e.g., Perfect 10, Inc. v. Amazon.com, Inc., 487 F.3d 701 (9th Cir 2007) (overturning a district court's holding that Google's automated creation and hosting of thumbnails failed the four-factor fair

The difficulties with the "fair use" solution to transaction costs are exacerbated by that fact that copyright limitations and exceptions have been eroded by the introduction of anti-circumvention laws.[414] Once important corpora of works have been locked away with DRM, it no longer really matters whether indexing them is fair use, because activist rights holders have overwhelmingly succeeded in obtaining anti-circumvention rights that trump fair use.[415]

It is possible that in the future, courts will expand fair use and other doctrines of limitations and exceptions to copyright, so that uses of copyright works in gestalts is clearly allowed.[416] It is also possible that legislatures will intervene to that same end. But such possibilities are entirely speculative, and the present reality is that limitations and exceptions are insufficient to stave off all of digital copyright's anticommons effects.

8.2.5   Reduction by Commons

Another solution to the transaction cost problem is to change the default rules. Under exclusive rights, most acts require permission. In an information commons, many or all kinds of acts are presumed to be permitted.

Copyright commons can arise in different ways, and display different mixtures of the legal and economic anti-property. The Web is — for the most part — a *de facto* commons. Almost anything on a website can be linked to, cut and pasted, deeplinked, framed, scraped, translated, aggregated, trackbacked or mashed up. The legal framework that governs this commons is a little sketchy: it is clear that implied licenses do much of the work, augmented by fair use in U.S. jurisdiction, but it is not clear how

---

use test in instances when the indexed sites were themselves infringing, and Google benefitted because those sites participated in Google's advertising)

[414]See (Nimmer 2000), Universal City Studios, Inc. v. Corley, 273 F.3d 429 (2d Cir. 2001)

[415]Superficially, the DMCA appears to allow some exceptions to anti-circumvention rules; see 17 U.S.C. §1201 (a)(1)(C,D), (d, f, g, h, j). But the structure of the act renders them all but meaningless; see 17 U.S.C. §1201 (a)(1)(E) (ensuring that exceptions for the act of circumvention do not apply to manufacturing, importing, offering to the public or otherwise trafficking the devices necessary to perform those acts, which are separately actionable under §1201 (a)(2)(A) and §1201 (b)(1)). The exceptions in §1201(d, f, g, h, j) are similarly limited. The same Kafkaesque rules have since been exported by U.S. trade negotiators. When an Australian House of Representatives committee realised that they had executed a Free Trade Agreement with the United States that had the same logic about circumvention and circumvention devices, they described the requirement as a "lamentable and inexcusable flaw" in the Agreement (Australian House of Representatives 2006, paragraph 3.118).

[416]A sensible outcome would be the addition of a fifth factor to fair use, such as "the existence of a public interest in allowing the use" to be weighed against the others that U.S. courts have adopted to date.

far these doctrines extend.[417]

P2P file sharing networks are largely illegal commons in which reproduction and some forms of modification are barely inhibited. Unauthorised file sharing downloads are no doubt helping many people to save on their transactional budgets! But inevitably, mass prosecutions against users and findings of secondary liability against the operators of these networks have limited how far those savings can travel in the status quo regime.

Let me elaborate a little further on the results of these institutions. It is only when primary materials are available in open information ecosystems that competition can freely drive the development of the most usable and sophisticated "value added" search services.[418] To varying extents, commons are able to deliver these benefits.

Due to the operation of what is effectively "implied" or "opt-out" licensing, Web-based search and analysis tools have been able to avoid many transaction cost problems. Noteworthy examples include Google™,[419] CiteSeer/Research Index[420] and the Internet Archive.[421] Part of the beauty of these systems is that they are not "deposit" collections that depend on the participation of the holder of the work to achieve inclusion. Instead, they actively harvest everything that has been published in a particular form, achieving a degree of reliability and utility that might not otherwise be reachable. A poetic illustration of that point is that often, the best way to find computer science research is to use google to find articles on Citeseer (even though, in order to be on citeseer, they had to be on the web in the first place). Web site operators are given the option of excluding their sites from these commons by the use of `robots.txt` files.[422] Crucially, the rule is opt-out with `robots.txt` rather than opt-in with a copyright license.[423]

Unfortunately a huge portion of humanity's art, culture and wisdom is not on the open web, and does not fall within the scope of this peculiar solution. Publishers have in fact taken strenuous exception to any incursion of the opt-out norm into their

---

[417]See, for example, (Kiritsov 2000; Waladan 2005; Norgaard and Garcia 2004; Lin 2005).

[418]See, for example, (Elkin-Korin 2001).

[419]`http://google.com`; see also (Brin and Page 1998).

[420]CiteSeer was originally developed in NEC's research laboratories, though it is now operated by Pennsylvania State University; see `http://citeseer.ist.psu.edu`; see also (Giles, Bollacker, and Lawrence 1998).

[421]See `http://archive.org`.

[422]See `http://en.wikipedia.org/wiki/Robots_exclusion_standard`.

[423]But, in response to this issue, Bracha (2007) has argued that there is no reason why property should necessarily be synonymous with an opt-in model.

domain: it "knocks the notion of copyright on its head", as Patricia Schroeder, head of the Association of American Publishers, put it (Martens 2005).

The Creative Commons (CC) is an effort to draft, and persuade people to use, liberal licenses for cultural copyright works.[424] It took its inspiration from the success of free and open source software licenses in providing legal cover for the communal authorship of code. It is not so much a substitute for web-based publication as a way of removing concomitant legal uncertainty — an attempt in retrospect to make the web legal. The CC project reduces transaction costs for transformative users who obey the law. It may reduce other transaction costs by securing the online publication of some new material, if publishers decide to build businesses around dual licensing with CC terms for non-commercial use only.

I will not talk at length about the Creative Commons here, but I will make one observation. If CC licenses reduce transaction costs (and they surely do for the makers of derivative works), they are a long way from eliminating them. The overheads of understanding, selecting and applying a CC license alone is enough to guarantee that. One anecdote may serve to illustrate this point: a friend of mine, who also happened to be an extremely productive and talented scholar, once emailed me to ask about his choice of creative commons licenses. He had taken a number of good photographs, and wanted to put them online. He didn't know how best to trade off the effects of non-commercial use clauses (which would prevent inclusion in Wikipedia) against their benefits, should someone wish to print a book containing his photos. He had spent a week puzzling over the problem, and was planning to spend longer. That was quite a transaction cost!

The establishment of commons (especially the web and P2P networks) has probably been the most effective solution to the transaction cost problem to date. But existing commons do not include everything. It is the works and use-cases not covered by them that accrue transaction expenses most quickly.

8.2.6   The bottom line

Each of the mitigatory methods discussed in this section has some effect on transaction costs.[425] It is consistently apparent that one class of gestalt uses (indexing, search

---

[424]See `http://creativecommons.org`; (Carroll 2007).

[425]In truth, they overlap a little: markets organise technology, and fair use contributes to the web commons.

and analysis) is insufficiently assisted. The creation of commons helps a great deal, but commons do not just spring out of thin air. They must be planted, nurtured and defended. A new factor for fair use, or other broad limitations and exceptions to copyright, would be one way to ensure that copyright gestalts can exist in the public domain. Alternative copyright regimes are another.

8.3 A Digital Anticommons?

If transaction costs are not sufficiently mitigated, the worst case outcome is not just unnecessary expenses and services that are less sharp due to a lack of competition and open information ecosystems. It is that important kinds of services are greatly delayed or never eventuate because the transactional burden, accrued over many different copyright works, is prohibitive.

That kind of outcome — one in which beneficial projects are prevented by a thicket of interlocking property rights — has been called the "tragedy of the anticommons.[426]

As I will explain in this section, humanity is close to allowing anticommons effects to interrupt one of the most incredible projects in history: digitising and indexing books, and making them searchable (all of them). It is very likely that the creation of a universal, searchable digital library has been much delayed by the need to clear rights.[427] It is possible that copyright will continue to prevent this project from being completed, or if it is completed, from being made available to the citizens of many countries. The digitisation of books is the grandest endeavour that is caught in the copyright anticommons — but it is not the only one.

8.3.1 Books

Since the printing press transformed the process of book production in the 15th century, publishers have set somewhere around 30 million books in type. Today, the text of all of those books could be stored on a server costing less than $1,000.[428] All of

---

[426]Heller first noticed this phenomenon when studying privatisation in post-communist states (Heller 1997). Anticommons effects are also recognised to occur in other intellectual property regimes; see for example (Heller and Eiseneberg 1998); (Reichman 2000) (advocating liability-rule regimes as an alternative to creeping exclusive rights which are likely to give rise to anticommons effects).

[427]In the absence of copyright issues, it is a fair bet that some firm would have digitised the Library of Congress by the end of the dot-com boom.

[428]Assuming that the average book has 80,000 words and can be compressed into a 100 kilobyte file. Over a total of 30 million texts, this amounts to a total of 3 Terabytes. An off-the-shelf PC server,

the novels, reference works, scholarly journals, textbooks, cookbooks, biographies, memoirs and illustrated primers would be instantly accessible. We could so easily combine the convenience of Google with the completeness of the word in print and the authoritativeness of its most reliable authors.

The problem of out-of-print and hard to find books is also now just an artefact of copyright law. The cost of printing and binding a single copy of a book on demand is now only $3-$5. The equipment for doing so costs only between $5,000 and $50,000.[429] Every library could be equipped with one of these machines.

How, it must be asked, is humanity faring in realising these opportunities?

Amazon's Search Inside the Book

Amazon's "Search Inside The Book" program aimed to take advantage of some strategic advantages in their marketplace position to make book searching a practical (and profitable) endeavour. The two cards that they hold are a strong set of relationships with publishing firms, and the immediate possibility of turning search hits into sales of books from their site.

Launched in 2003, the project is an opt-in scheme for publishers; those who chose to participate send copies of their titles which are scanned with specialised equipment and indexed. Users searching through Amazon's catalogue receive hits which are matched not only from the titles and descriptions of books, but also references to and snippets of text from pages of books which match the search terms.

Users can then access up to 10% of the full text from each title, in the form of page images covered by weak browser DRM. The 10% limitation is enforced through the use of credit cards for authentication (once again, the fact that this is unified with Amazon's payment system makes this practical for them to deploy).

Search Inside The Book is already a useful tool — especially so for academics

---

with two USD $150 1.5TB disks, would do the trick. Including more redundancy, co-located hosting and professional administration would raise the cost by an order of magnitude. The illustrations would require more space, considerably more so for some books, but the overall storage requirements would be only an order of magnitude larger. The cost of placing a digital copy of the world's stock of literature on the net would appear to be in the order of $100,000 Australian dollars for setup. Over time a figure like that would probably be recurrent every three years or so.

[429]A device like the Espresso Book Machine, which produces books of the same quality as traditional printing processes, costs around US $50,000 `http://palimpsest.stanford.edu/byform/mailing-lists/bookarts/2008/05/msg00164.html`; an automated printing and binding device can be had for a tenth of that price: `http://palimpsest.stanford.edu/byform/mailing-lists/bookarts/2008/05/msg00188.html`.

working on highly inter-disciplinary topics (the reader may wish to guess at which of the references in this thesis were only discovered because of its assistance). The important question is how it compares to the services that would be available under other copyright regimes. Printing, for example, was rapidly disabled after the service was launched as a result of what could only be described as copyright paranoia.[430] The coverage of the service, though extensive, is still a minuscule fraction of published literature (though a more respectable portion of new titles).

Google Books

Not to be upstaged by Amazon, Google also spent 2003 initiating a book search program, entitled Google Print. The initial model appeared similar to that used by Amazon, though with some minor differences — they required publishers to make at least 20% of each text available; they used browser cookies for authentication;[431] offered publishers a cut of advertising revenue as an incentive, and cross-promoted sales through others' online stores (including both publishers' pages and sites like Amazon and Barnes & Noble).

In December 2004, Google announced a different strategy for taking their book search project further. In a pilot program with libraries at Harvard, Stanford, Oxford and the University of Michigan, and with the New York Public library, Google is scanning a large number of texts using non-destructive scanning equipment. By relying on these libraries, with their very extensive collections, Google obtain the obvious benefit of escaping the need for material co-operation from publishers, some of whom would be reluctant to participate, some of whom would never get round to sending texts, and many of whom no longer exist.

Google and the authors' and publishers' groups did not wait for a ruling on whether the scanning and indexing was fair use, but instead proceeded to settle the class action. An important property of U.S. class action settlements is that they are binding on all members of the class who do not opt-out during the settlement process. The task of negotiating licenses with the copyright holders of all printed books is essentially impossible — not only because of the difficulty of reaching agreements, but because many works are "orphans" whose copyright is held by publishers which no longer

---

[430]This was a result of an investigation by the Publishers Guild, observing that if one used ten separate accounts (requiring ten credit cards) it would be possible to obtain a complete printout of any book in the system (Olsen 2003).

[431]See (Duffy 2005).

exist or parties (often the heirs of rights holders) who are unaware of that fact or otherwise uncontactable. By settling with a class, Google obtains licenses from this entire universe of copyright holders in a single stroke.

The Google Book Search settlement attempted to create a novel collecting society — the Books Rights Registry — and to give it permission to grant certain kinds of licenses to Google. The settlement was controversial on many grounds,[432] though probably the largest was that the structure of class action law prevent the same terms from being offered to Google's competitors. That no doubt was a significant factor in the court's decision to reject the settlement. But, unfortunately, it prevents the experimentation with collective licensing models that the settlement would have afforded. And prevented or delayed the appearance of a comprehensive, full-text, digital library.

8.3.2 Research Databases

Transaction costs have slowed the development and deployment of information services that are built on top of scholarly research. The features of research tools such as CiteSeer and ArXiv (which index freely available papers, primarily in computer science and physics respectively), and to a lesser extent Google Scholar, have often gone well beyond those of proprietary databases that index copyrighted research material,[433] but they cannot extend those benefits to the entire corpus of scholarly literature, because it is largely owned by the proprietary competitors.[434] The first-best regime would clearly allow users to access whichever writings through whichever service they chose.

Research databases are not strictly within the scope of inquiry that I set out for the thesis in Section 1.4.2. Much of their function lies within the academy, and is therefore not personal and non-commercial. There is a danger of overstating the case

---

[432]See *e.g.,*, `http://thepublicindex.org/documents/responses`.

[433]Such as Thomson ISI's *Web of Knowledge*®, Thomson West's *Westlaw*®, or Reed Elsevier's *LexisNexis*™. Because proprietary databases are the only viable option in many disciplinary areas, researchers in those fields may not have compared their performance to services like CiteSeer. Admittedly, the comparison between research services is also complicated by the role of patented algorithms (see, for example, U.S. patents #5,265,065; #5,794,236; #6,285,999 and #6,289,342), and the significant amount of human labour involved in some of these databases (the *Web of Knowledge* includes manually indexed citations, while *Westlaw*'s *KeyCite*® depends upon manual categorisation of legal material).

[434]Recently, some proprietary publishing companies have started including their papers in search engines' indexes using a "bait and switch" approach. If a URL from their server is requested by a major search engine's spider (presumably identified by its source IP address), it is answered with the full article text. Google then starts offering that article in its search results, but when a user clicks on the link, they are greeted with a page offering them the article for a fee (usually in the US $20 range). Sometimes, the abstract is available for free. One wonders what fraction of would-be readers pays the fee.

against exclusive copyrights if the output of academia is included, because the problem there is as much a failure of universities to collectively assert their rights against the academic publishing industry, as it is a failure of the copyright system. University communities hold all of the rights, and as they begin acting in the public interest, and their own collective interest, all of the problems of accessibility, indexability and searchability of scholarly research will gradually disappear. This has already begun to happen in the U.S. with the passage of legislation requiring future research funded by the National Institutes of Health to be openly accessible after a year.[435]

Having said that, the issue is not entirely out of scope. A fraction of the benefit of these outcomes would not arise through more efficient research, but through better public access to the fruits of academia for personal interest and educative purposes. These kinds of access *are* within the scope of inquiry (non-commercial use of literary copyright works) and would in fact be rapidly achieved by the virtual market and anarchic regimes. For this reason, a part of the scholarly copyright problem must also be considered when reckoning the transaction costs of digital copyright in people's private lives.

8.3.3 Other domains

There are many other gestalt applications which transaction costs may delay, prevent, or simply encumber.

There are other indexing and analysis applications which, while not as grand as the book search efforts, are still feasible with existing technologies. These include, particularly, universal databases of components for the creation of more complicated media: video snippets,[436] audio samples, or models for 3D animation.[437]

Looking to the future, there are several fields of computer science research chipping away at problems that, if well solved, lead to very useful analysis tools over the corpus of human writing. These benefits will also run into thickets of exclusive rights.

---

[435]See `http://publicaccess.nih.gov/policy.htm` for details of the NIH open access legislation and its implementation.

[436]Publicly accessible video search tools have made some progress with the success of *YouTube* and *Google Video*, although they are limited in scope to videos uploaded by their users. The author has seen three private demonstrations of independently developed video search engines based on DVD subtitles which are much better than anything publicly available; none of them has been launched because of copyright concerns.

[437]There are some existing proprietary marketplaces for 3D models, such as 3d02, but they are very limited.

"Natural language processing" (NLP) also known as "computational linguistics", is one such field. It deals with computer programs that pragmatically parse and manipulate human language, without tackling the much more difficult issue of semantics. Some of the relevant problems that NLP researchers are working on include question-answering search[438] and tools for extracting summarised factual data from text.[439]

"Machine learning" techniques are another field from which useful corpus analysis tools might be expected to appear. Machine learning is the study of algorithms which can recognise and categorise individual objects by learning from larger collections of them. In its application to data mining and search, machine learning can help us answer queries such as "find an object of type *x* that has property *y*".

Of course, the future development of nifty algorithms for and implementations of search and analysis systems would not preclude Google, or Amazon, from taking them and applying them to the writings they have scanned and/or licensed.[440] Deals could also be struck with academic publishing conglomerates — although they strike very hard bargains with researchers.[441] Rather, there are two other concerns which present themselves. One is that these companies may have been receiving, from copyright, a role as gatekeepers that allows them to monopolise the effective deployment of these future technologies.

Secondly and perhaps more worryingly, the encumbrances of licensing may work to slow the development of these technologies in the first place. When the default terms of access are liberal, or people just presume that they are (as with the World Wide Web), a plethora of commercial and research-oriented efforts can flourish, striving to analyse the corpus more effectively. Systems like Google are in fact the fruits of that openness.[442] It may be that unless the world's books are as openly accessible as the world's web pages, that the next PageRank-like breakthrough will be missed.

---

[438] See Radev *et al.*(2005, §2.1) for a survey of techniques based on specialised indexes (other approaches, such as those discussed in section 2.2 of that paper, would not necessarily be affected by transaction costs)

[439] This field is known as "Information extraction"; see (Cardie 1997) for an early survey.

[440] In fact, the first proposed Google Book Search class action settlement offered this kind of access for academic research purposes, but not for commercial start-ups that want to turn discoveries in these fields into products. See Google Book Search Original Settlement Agreement, *available at* `http://www.googlebooksettlement.com/intl/en/Settlement-Agreement.zip`, Articles 1.90, 1.121, 1.129-130, 7.2 (b) vi.

[441] Personal communication with bibliometrics researchers who use Thompson Corporation data.

[442] As Larry Page put it, the road to building Google started when he "got this crazy idea that I was going to download the entire Web onto my computer." (Vise and Malseed 2005, pp. 11,36–37) Fewer PhD students try to download the entire Sony music catalogue, or the entire Lexis Nexis corpus, or even the still-in-copyright books that Amazon and Google have scanned, onto their computers.

This result of transaction costs is only speculative. We cannot be certain that the researchers and entrepreneurs with the capability to achieve this kind of progress will not find ways to start with smaller corpora, strike deals with Google or Amazon, and keep the ball in flight continuously. But the real possibility that they will not, suggests that transaction costs must count as a real downside to exclusive rights based digital copyright.

8.4   Comparison of Regimes

In the other chapters in Part IV, I have been able to present some estimates, however approximate, of how much society would gain or lose by adopting one digital copyright regime or another. I know of no methodology for obtaining a similar measure for transaction costs. How can we infer the sum total of expenses like my friend's deliberation over Creative Commons licenses, Wal-Mart's absurd policies on digital photography, the delays in the digitisation of books, and the onerous terms that academic publishers impose on bibliometrics researchers? In the stead of a more satisfactory quantitative comparison, this section offers a qualitative one.

The Status Quo

At present, digital copyright is erecting transactional barriers in numerous places. Mitigatory institutions exist, but they have not removed the large-scale overheads that result from the basic requirement to obtain permission at every turn. The result is that people are spending a lot more time and effort on licensing, and worrying about licensing, than they should be. Anticommons effects are also in play. Book search, research databases, and future gestalts have not been prevented by copyright, but their timeliness, coverage and quality has undoubtedly suffered.

The situation could be greatly improved through the creation of copyright exceptions that are both broader and more clearly delineated than existing mechanisms like fair dealing or four factor fair use. But such developments are politically unlikely,[443] and it would be highly optimistic to include their benefits when estimating the normative desirability of the pragmatic/*status quo* regime.

---

[443]For instance, despite widespread agreement that "orphaned" copyright works present a major transaction cost problem (Peters 2008), the U.S. Congress has been unable to pass legislation addressing the problem; see `http://www.techdirt.com/articles/20080930/1946332421.shtml`.

Transaction costs in Feudalism

Information feudalism does not alleviate the basic necessity for permission that makes transaction costs so troublesome at present. The creation of strong, unified DRM systems might nonetheless help to reduce the cost of making some of those transactions, in the mode discussed in Sections 8.2.1 and 8.2.2. This is a quantitative, rather than a qualitative improvement.

It is important to realise that while technology may greatly reduce the costs of establishing a negotiation of some sort with a rights holder (or their electronic agent), it cannot magically perform the negotiation itself. The fact that producers and consumers are haggling over how to divide a surplus, and the fact that producers have imperfect information about consumers' valuations, guarantee that non-standard licensing must be an inefficient process: one that sometimes succeeds, sometimes fails, and along the way is costly.[444]

Strong DRM also has some countervailing and counterproductive effects. In these regimes, copyright holders are especially concerned to prevent leaks that flow outside their secure distribution channels. But proper indexing and search tools actually facilitate those leaks. Any search facility which provides information about the context of results (as Google does with search terms, and CiteSeer does with citations) is also providing parts of the actual text of the document. It is a relatively straightforward process to write software which combines these snippets of text, recreating the entire original document.[445] Preventing these forms of revelation (for example, by attempting to track the users of search facilities) may be possible, but is likely to be costly.

Another problem is that strong DRM regimes completely neutralise many important existing limitations and exceptions to copyright. Fair use is replaced by iTunes' 3 copy rule, or its descendents. The flexibility lost in this transition is a tremendous inconvenience for those who happen to make unusual uses (often transformative ones) that are not permitted by the DRM rules.

Even with the possibility of quantitative reductions from standardised licensing, it is not at all clear that a strong DRM regime is any better than the status quo.

---

[444]In fact, the Myerson-Satterthwaite theorem (Myerson and Satterthwaite 1983) in mechanism design shows that such negotiations over non-standardised contracts are inherently inefficient.

[445]The disabling of Search Inside The Book's print feature (Olsen 2003) is a real-life example of the consequences of this concern.

Transaction Costs Under Information Anarchy

On the surface, information anarchy would seem to eliminate transaction costs. There would have been no restrictions on individuals or organisations digitising books and putting them on the web, for searching engines to collate. The best writing (provided it was still being written) would have been uploaded rapidly by enthusiastic readers. There would likely have been large projects, funded if necessary by charitable foundations or governments, to scan works in a more systematic fashion. But publishers who were still selling paper books would have had little incentive to aid these efforts in any way. Copyright owners in films, audio samples, 3D models and many other kinds of usefully databaseable would also withhold their cooperation.

As further weakness of information anarchy is that it shifts some transaction costs to the process of production. So long as authors and publishers are trying to get paid (using mechanisms like the SPP, as discussed in Section 3.2.3), their efforts would carry transaction costs both for themselves and their audiences. When those costs involve such contortions as having to wait for months between paying for a book and receiving it, they are extremely inconvenient.

Transaction costs in a virtual market

Alternative compensation systems are not completely free of transaction costs, but they are close to second-best optimality. No permission needs to be obtained in those use-cases that are covered by the virtual market license (including many or all transformative uses, as discussed in Section 5.4.1). The full benefits of market forces and of open information ecosystems are therefore available at least to minimise the costs involved in identifying, obtaining and enjoying copyrighted works.

The only transaction costs unique to a virtual market are the time and effort expended by those users who chose to vote explicitly.[446] It cannot be guaranteed that the total time spent by users on voting and deciding how to vote will be negligible (after all, some people do spend a lot of time organising their record collections), but from a normative point of view, the fact that the activity is optional is of great reassurance.

---

[446]There are some substantial costs involved in setting up the virtual market system, but since these are fixed rather than marginal costs, I discuss them in Chapter 9.

8.5 Conclusion

Were it not for the transaction costs required for copyright licensing, our ability to find the writing that answers our questions or pertains to the specific subjects that interest us would be much greater than it is today. We would likely be able to sit down at our computers, and type into our choice of search engines queries to be answered from the reservoir of all of human writing — not just the web, but the entire universe of extant books and scholarly articles. After those search engines had returned their answers, we would be able to click through to the full text of the matching documents, either at a reasonable price — or, under a non-exclusive system, for free. The additional cost of obtaining a bound and printed copy of a relevant book — no matter how obscure — would only be a few dollars. The questions of curious teenagers would be answered, not just by the vicissitudes of website authorship, but by the best popular science writing, the best textbooks, and the best knowledge from research papers and monographs. As things stand today, finding those materials in online databases tends to require extensive research skills, where it is possible at all. That is the anticommons cost which must be counted against digital copyright. There are many other smaller-scale transaction costs, too.

Under a virtual market, these issues would have been resolved years ago. The incentives for publishers to get their materials indexed by search engines and into the browsers of readers, and the absence of legal complications for organisations wanting to scan older collections, would — probably by the time the dot-com bubble burst in 2001 — have dissolved the boundaries between the web and the printed book. Instead, we are still waiting for a Google Book Search settlement, and hoping that if it goes ahead, it does not produce a monopoly.

Stronger DRM regimes change the balance of transaction costs a little, but it is not clear that they are actually an improvement on the status quo. Information anarchy reduces licensing expenses, but it may not provide the right incentives for publishing firms that hold master copies of works to aid in digitisation and third-party indexing. It also imposes a whole new and problematic set of costs on the creation of transformative derivative works. Again, it is not clear that the alternative would be any improvement on the status quo.

Virtual markets are clearly the preferred regime from a transaction cost perspective. But what we do not know is how much there is to be gained. No methodologies are available to say how much time is wasted by people worrying about copyright and

licensing, or how to place a value upon the gestalts that could be created without the need to clear exclusive rights.

The idea that any book (or audiovisual sample, etc) could be instantly accessed through the best existing and future search systems is a powerful one. But there is a risk that researchers (including the author and readers of this thesis) will intuitively overstate the value of these tools to humanity as a whole. The learning of books is only useful if one is ready to read them, and is easy to romanticise the benefits of giving everybody access to all of them.

These are grounds for caution. In the absence of persuasive data, I would count transaction costs as a noteworthy but not overwhelmingly large disadvantage of exclusive digital copyright regimes.

Chapter 9

The Price of the Rules: Infrastructure, Enforcement and Security

Contents

Mr Johnson, that fuzzy region on your X-ray indicates one of two things. Either you have multi-drug resistant tuberculosis, or the copy-protection system on our computer thinks that part of your left lung looks like Mickey Mouse.

Cartoon by Philip Dorrell, `http://www.1729.com/blog/LookingForAWinWin.html`

Each of the copyright regimes we are considering has its own characteristic structures of information distribution, its own roles for participants to play, and its own flow of payments. These patterns do not just spring into existence; infrastructure is necessary to make an information economy dance in time. Important components may include hardware, software, effort spent on maintaining the security of a range of computer systems, education and legal advice, policing, and civil law enforcement. Some such infrastructures may be relatively standardised and centralised;[447] others may be *ad hoc* arrangements that spring up in the marketplace.[448]

---

[447]Such as private copying levies, virtual markets, the Apple iTunes Music Store or Microsoft Windows DRM.

[448]Such as those operated by bands who sell CDs at their gigs. There may be a natural gravitation

The choice of policy regimes clearly has a profound effect on the infrastructures that manage the fruits of cultural industry, regardless of whether this is or is not explicitly specified by the law.[449] And those structures are not all alike. They may have radically different price tags, in terms of the technology and institutions required to set them up, and the resources that must be devoted to keeping them operating correctly.

Providing an exhaustive ledger of all of the infrastructural costs of copyright and alternatives is beyond the scope of this thesis. The task may not even be possible within my analytic framework of a few distinct 'regimes' because each of those regimes groups together diverse possibilities amongst which those costs would differ widely (in other words, the structural blueprints and price tags for feudalism, public funding, pragmatism and anarchy, are not fixed). What this chapter does provide is some observations of major and systematic tendencies to differ that matter most for policy.

The chapter is divided into five sections. 9.1 briefly discusses the major infrastructural components of each of the four regimes. Section 9.2 considers how hardware and software security costs many be different for different kinds of copyright. Section 9.3 turns to what has come to be known as the 'war on copying' — the large and many-fronted campaign to enforce copyright in cyberspace. Section 9.4 takes a slightly more holistic approach, and discusses some fundamental reasons why infrastructure needs to be more or less expensive in different digital copyright models. Section 9.5 summarises and concludes.

9.1   The infrastructure of different copyright regimes

This section briefly identifies the kinds of facilities, technologies and institutions that are distinctively necessary for each copyright regime.

Strong DRM

A strong copyright regime of the sort introduced in Section 2.3.1 (in which obedience towards exclusive rights is a practical necessity) requires a combination of robust DRM

---

towards systematic solutions, although this is usually just a result of businesses looking for efficient ways of operating. For example, many bands that want to sell their own CDs end up using CDBaby (`http://www.cdbaby.com`)

[449] A statutory public lending right is a good example of strong specification, while anti-circumvention rules do not appear to specify much but do in fact guide the market in a particular direction. Other instruments, such as statutory mechanical licenses, fall somewhere in between — there is no legal *requirement* to use them in all cases but commercial practise has crystallised around them anyway.

systems and determined policing. Enforcement is required for plugging any leaks, pursuing those who promote them, and for scaring people away from alternative channels for content exchange or shutting them down altogether.

DRM controls must be constructed using trusted computing systems or other kinds of specialised hardware, both to prevent people from using alternative software to make unauthorised reproductions of works, and to ensure that if loopholes make such software possible, their closure can be rapidly forced upon uncooperative users. Those trusted systems must also be tamper-resistant, so that they cannot be "mod chipped" or otherwise physically subverted.

Cutting off or severely curtailing the supply of unrestricted digital copies would be an important first step, but it would not alone be sufficient to make piracy commercially irrelevant. The vast back-catalogue of unencrypted works, the possibility of analogue-to-digital reproductions, and the possibility of "insider" leaks from publishers together necessitate the constriction of channels for peer-to-peer copying, either by closing them or by preventing people from using them.

During the past decade, a proportion of P2P networks have been closed by litigation and judgements that firms operating and developing the networks are secondarily liable for their users' copying (Strowel 2009). But secondary liability doctrines have not driven the development and distribution of P2P software entirely underground.[450]

Some individuals have been scared away from P2P networks by the threat of litigation against them personally, although a significant expansion to existing litigation campaigns, or huge successes for proposed "three strikes" rules, could conceivably shrink open P2P networks to the point that they were commercially insignificant.

Even then, various kinds of private and privacy-preserving networks would continue to offer individuals a low-risk way to share files. Closing those networks would require either that existing doctrines of secondary liability be greatly expanded, or that merely running certain kinds of privacy-protecting software be deemed illegal, regardless of whether one was using them to infringe copyright.

These components, or at least a significant portion of them, are the infrastructure

---

[450]For instance, a large number of free/open source P2P clients produced by volunteer programmers have thus far escaped litigation. Furthermore, the organisations and services that facilitate free software development and distribution, such as Sourceforge and Debian, have no record of promoting infringements in the manner that made Grokster and Streamcast secondarily liable (Supreme Court of the United States 2005, Parts II.C, III.A), so under existing doctrines there is little to halt the distribution of existing free P2P clients, and some development might continue if programmers could demonstrate non-infringing intentions to suspicious courts.

required for an information feudalist regime that is free of economically significant piracy.

Status Quo Copyright

Under the status quo, copyright depends on a few different kinds of infrastructure. In additional to the bricks-and-mortar mechanisms that make possible the sale of books, CDs, DVDs and cinema tickets, there are structures focused on the digital world: on-line music services such as iTunes, Rhapsody, Amazon and E-Tunes; DRM code in operating systems and media players that some of those services use; institutions filing lawsuits against file sharing networks and file sharers; "education" campaigns to convince people that copying is wrong or to simply scare them out of it.

Trusted computing does not yet play an extensive role in deployed DRM systems. The closest live technologies are the built-in hardware authentication and decryption systems found in games consoles and dedicated media devices such as the X-Box 360, PlayStation 3, and BluRay drives and players. To date, strong, hardware-enforced DRM has remained too difficult to implement inside general-purpose PCs. But to-date unsuccessful efforts have consumed a great deal of R+D effort at Microsoft, Intel, Sony, Phillips, Apple, AMD, and many other firms.

Virtual Markets

Virtual markets require some kind of central administration authority, which would combine some of the characteristics of existing copyright collecting societies and electoral authorities.[451] This organisation would need to maintain the databases of MIHs and fingerprints necessary to perform media file identification. An independent, survey-based process would be required to periodically review the tax or levy rates that fund the system. A network of honey pot systems would be highly advisable for fraud detection.[452] The virtual market could include its own centralised file sharing servers, but this might be unnecessary. Community and marketplace organisations have demonstrated that they have all the expertise necessary to set up P2P networks

---

[451]This comparison may instill fear into U.S. readers, who are accustomed to a patchwork of electoral authorities administered by elected, partisan officials. They should perhaps take heart from the fact that other developed countries manage to provide these services in a centralised and comparatively trustworthy fashion.

[452]See Section 5.2.2.

that perform the distribution role efficiently, without ever needing an expensive central server farm from which to download files.

Virtual markets also require some support from users' computers. Their computers would, ideally, either include trusted computing modules or separate hardware to authenticate transactions.[453] Importantly, however, a virtual market that collected data only from a sample population would be able to employ these security measures at comparatively low cost.

Finally, alternative compensation systems would require anti-fraud staff to investigate data from honeypots and other incident reports; develop patches to trusted computing integrity measures that, at the least, detect these attacks and allow votes from compromised computers to be ignored until they are patched.

Anarchy

It is tempting to think that a state of anarchy in cultural production would, almost by definition, have no infrastructure at all. This is not in fact the case. The well-demonstrated human commitment to art and culture would lead to the formation of a constellation of institutions to bring about exchanges, financial and otherwise, between artists and their audiences.

Some of those institutions are familiar or at least clearly anticipated: rejuvenated patronage; more sophisticated tip jars; trustworthy intermediaries that facilitate the street performer protocol or its variants; expanded markets for product placement in films and artist placement in advertisements. Others would no doubt be inspired by necessity.

While we can be sure that such institutions would spring up, their relative importance in an anarchistic information economy cannot be so easily predicted. The different bottom-up remuneration processes would often exist in competition with one another, and it is unlikely that the outcomes of such contests could be determined in advance.

9.2 Systemic computer security costs

Each of our four possible copyright regimes has distinctive requirements for the hardware that must be distributed throughout userland. It is in this analysis that the first

---

[453] See Section 5.2.2.

signs of highly asymmetrical infrastructure costs become apparent. Pragmatic DRM and virtual markets both require moderate amounts of hardware to be installed. Strong DRM needs gadgetry that is more extensive (and far more expensive) on consumers' computers. Anarchy requires nothing at all.

9.2.1   Tamper resistance and security for strong DRM

The burdensome costs of a strong DRM infrastructure relate to the extreme difficulty of designing, building and maintaining these systems so as to be secure.

A single point of failure can cause the collapse of a strong DRM network. Once files are leaking out in digital form, even if it is only from a tiny number of holes, they can be distributed indefinitely through file sharing networks and other channels. To prevent this, millions of consumer devices need to not only be built around trusted systems,[454] but also be virtually tamper-proof.[455] This degree of security is a costly proposition, to say the least; but just *how* costly is difficult to say without having achieved it.[456] Informally, many security professionals have claimed that it is sufficiently difficult to be considered impossible. There has, however, been some technical progress on tamper-resistance, and a number of approaches can be used to predict a possible price tag. I will discuss two of these.

The most direct route is to make an extrapolation from the price of existing hardware. This estimation is based upon state-of-the-art tamper resistant devices for financial cryptography. The IBM 4758 co-processor is an example of an extremely sophisticated "trustworthy" cryptographic hardware platform.[457] The security features of the 4758 include multiple layers of physical shielding that are capable of detecting physical interference with drills, cutting instruments and solvents; sensors to detect unusual power supply and clock signals; fast-response thermometers to detect incineration or immersion in liquid nitrogen; Geiger counters to detect bit-flipping radiation. If any of these sensors detect an attack, the device instantly erases all of its private

---

[454]On the necessary role of trusted systems in making DRM an economic possibility, see, e.g., (Stefik 1997); (Schechter, Greenstadt, and Smith 2003, Parts 1–2); (Roemer 2003, Part II).

[455]Designing consumer devices which are literally tamper proof is almost certainly impractical, if not impossible; see (Anderson and Kuhn 1996). Instead, the degree of tamper resistance must be so high that only attackers with substantial resources can succeed — and, as Anderson & Kuhn went on to demonstrate, even this will be very challenging. See (Anderson and Kuhn 1997). In addition, large amounts of information must remain enclosed when security compromises do occur, and every significant instance of tampering must be traceable before it inflicts massive economic damage.

[456]See (Schneier 2000, Chapter 14).

[457]See (Smith and Weingart 1999).

keys. Those keys are stored in memory with a rotating scheme so that no "burnt in" forensic evidence is available to recover them. This device would stand a reasonable chance of preventing serious economic damage due to digital content leakage — although in deployment it is not necessarily immune to software flaws[458] In 2003, the 4758 was selling for around USD $4000 into a relatively small market.

Unofficial estimates undertaken at IBM's Thomas J. Watson Research Center (obtained by personal communication), indicated that the 4758 itself, even in very large volumes, would still cost over USD $500. A miniaturised version would probably cost $100–$150. It might be conceivable to get as low as $20 or $30 if a single-chip equivalent device could be designed (although this would require significant research and development). Even with the most optimistic figures, this kind of hardware would represent a very sizeable tax if it was required in all consumer media devices. In any case, we can reasonably conclude that the cost of employing miniaturised versions of existing tamper resistant technologies for DRM applications would fall between U.S. $20 and $150 per consumer device.

That first estimate focused solely on the cost of hardware. A more ambitious back-of-the-envelope calculation can attempt to account for a wider range of possible weaknesses by inferring costs from other areas of IT security.

A recent survey of 503 organisations' experience in dealing with computer crime indicated that a total of USD $375.6 million was lost annually in incidents of a kind which might be applicable to a DRM network.[459] These measurable losses were spread over 44% of the population surveyed, amounting to $1.7 million per organisation in that group; the other 56% of organisations were unable to quantify or did not disclose the size of their losses. 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. Strong DRM in particular needs to prevent all such breaches that are pertinent to the material it controls, because almost any lapse of security can be very serious.

Conservatively adopting the lower bound, effective security for a controlled corporate network costs between $750,000 (average measured losses per organisation)

---

[458]See (Clayton and Bond 2002).

[459]See (Power 2002, at 10–1). The losses included as potentially relevant were "theft of proprietary information", "sabotage of data networks", "telecom eavesdropping", "system penetration by outsider", spoofing, viruses, and unauthorised insider access; but I have excludes losses reported from "insider abuse of Net access", financial fraud, laptop theft and denial-of-service attacks.

and \$1.7 million (average loss for organisations that measured and disclosed their losses).[460] The cost for a 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.[461] 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 of the estimates made for the cost of strong DRM are a little audacious; that is the price of futurology. What is striking about them is that they give such similar answers. In the conclusions below (Section 9.5), I discuss the implications of these figures a little further.

9.2.2   Security for pragmatic DRM

The systemic computer security costs of pragmatic DRM are, by definition, the sum total that is currently being spent on the research and application of copyright protection systems — chiefly by private sector firms but also within other organisation such as universities.[462] These expenses accumulate in a range of contexts: through extra components that must be included in devices to make copy protection work; through extra hardware and software engineering effort that must go into systems whose behaviour is to be restricted by DRM; through negotiations and competition to set DRM standards and even through the transaction costs and deadweight loss of the patents that results from DRM R&D. Such disparate expenses are exceptionally hard to tally. Even the firms involved would have great difficulty in knowing how much of the effort expended by their engineers on relevant projects was ultimately attributable to the complexities

---

[460]One of these numbers assumes that the 56% of organisations who do not measure their losses in dollar terms suffer similar losses to those that do; the other number assumes that organisations which do not measure their losses suffer zero losses. The \$750,000 number is unrealistically conservative, but again, functions as a strong lower bound. The \$1.7 million number is not a strict upper bound, because it is conceivable that organisations which do not measure and disclose losses are those with poorer computer security practices and therefore suffer more than the reporting 44%.

[461]From the table in (Power 2002, at 3); assuming median numbers for each interval, and 15,000 employees for organisations in the 10,000+ category.

[462]Publicly funded R+D on DRM should be counted, but only to the extent that such research would not occur under alternative regimes.

introduced by copy protection and access control. Collating this non-existent data is more difficult again.

One cost measure which is deserving of discussion, but does not ultimately deliver a useful number for the cost of developing DRM, is the observation of market values for patents and similar indispensable components of DRM technologies. There are not very many of these data points available, and it turns out that they aren't strongly indicative of the costs of DRM.

The largest transactional datapoint was Microsoft's payment of US $440 million to settle a patent infringement suit brought by Intertrust and obtain licenses to Intertrust's DRM patents.[463] This was one expense among many others for Microsoft's DRM efforts, and should be considered alongside the fact that Microsoft barely charges compatible device manufacturers anything for using its DRM standards.[464] So for Microsoft we have to interpret this as a part of the cost of keeping content industries publishing on Microsoft's platforms when those industries refused to publish without DRM.

A smaller data point is the $45 million dollar price that Macrovision paid to purchase the BD+ BluRay DRM business from Cryptography Research (Cryptography Research 2007). BluRay discs have two DRM standards: mandatory AACS (see note 36) and optional BD+. BD+ is only used on a small portion of BluRay discs,[465] so total royalties on it are likely to be smaller than those from the mandatory AACS. The $45 million number should be seen as the expected future value of royalties paid to Macrovision on the BD+ patents and trade secrets, and not the entirety of the expenses that consumer electronics firms have to outlay for implementing it.

A third data point, of a slightly different nature, is the fact that Sony was willing to delay the launch of its Playstation 3 console for 6 months in order to wait for the

---

[463] See *Intertrust Technologies Corp. v. Microsoft Corp.*, Nos. C 01-1640 SBA, 2003 U.S. Dist. LEXIS 22736 (N.D. Cal. Nov. 26, 2003).

[464] See eg `http://www.microsoft.com/windows/windowsmedia/licensing/final.aspx` (royalties of 10–20 cents per compliant device). Even if every media device in the world became compliant with Windows Media 10 "Plays For Sure" DRM, Microsoft would only be earning a few tens of millions of dollars per year, which is not concomitant with paying $440 million as part of their entry fee. The strange economics of Microsoft's DRM licensing was further illustrated when one of their executives admitted that a key goal of the fee structure was to prevent small-scale implementations of their DRM (Brown 2006).

[465] As of 2009, an engineer who worked on BD+ reported that despite the renewable nature of the design, a number of studios had concluded that each generation of BD+ was being cracked too quickly, and that it wasn't worth the effort and cost of using it on their discs (personal communication).

AACS and BD+ DRM standards to be finalised.[466]

As stated above, data points do not give us a reliable window into the true engineering costs for adding complex DRM to already-complex systems, and even the firms that are paying those costs are unlikely to be able to measure them. But they do indicate by their orders of magnitude that total society-wide expenses for DRM engineering would be counted in units of billions of dollars.

9.2.3   Security for virtual markets

Virtual markets require non-trivial expenditures on security (both in hardware and software) to provide barriers against certain kinds of royalty-stealing attacks. There are two main differences between DRM security and virtual market security: (1) virtual markets only need to defend against remote attackers, while DRM needs to defend against device owners, who have physical access to their own property; (2) virtual markets can collect their data from sample groups that are much smaller than the total population, thereby allowing much lower total expenses even if much more is spent on each device to be secured.

Although a quantitative costing for virtual market implementation is beyond the scope of the present inquiry, the similarity to the pragmatic DRM case is striking. Both ultimately require some carefully written software and benefit to some degree from trusted computing hardware support. It is likely that the systemic costs of virtual market security would be lower than current expenditures on DRM, but it is conservatively safe to conclude that they are of the same order of magnitude.

9.2.4   Systemic security costs for information anarchy

By definition, information anarchy does not place any regulatory requirements or restrictions on users' computers and other electronic devices. A small amount might be spent by authors and publishers on tip jars or by the trusted third parties operating street performer protocols (see Section 3.2.3). But these expenditures would be non-systemic and truly minuscule compared to the security expenses required under the other three regimes.

---

[466]See `http://www.playstation3-ps3.com/2006/02/17/playstation-3-launch-date-delayed/`.

9.3 Policing, Enforcement and the "War on Copying"

As matters stand, many private citizens have the means and inclination to make large numbers of unauthorised and often illegal copies of copyrighted works. Such wide divergence between social norms and the law is reminiscent of nothing so much as the prohibition of popular recreational drugs, especially alcohol during the early 20th century and marijuana since then.

The way that this inconsistency is resolved is the key variable that distinguishes between moderate information anarchy (allow social norms to prevail), information pragmatism (limit piracy using imperfect DRM and vigorous private enforcement of the law) and information feudalism (use strong DRM to make social norms as irrelevant as possible).

9.3.1 The cost of the pragmatic war on copying

Under the status quo, the sustainability of copyright depends in part on the behaviour-changing power of legal threats and lawsuits, on the creation of anti-copying social norms, and in some cases, on the use of public or private police forces. Because pragmatic DRM technology is insufficient on its own to prevent copying, and because only a small fraction of the population abstain from copying for moral reasons, these modes of enforcement must act to change the behaviour of enough people to ensure that cultural industries remain profitable. In the process, a burden is placed upon those whose behaviour is changed.

In the United States, during the mid-2000s, over 700 ordinary people were sued each month for their use of P2P file sharing networks. Globally, the IFPI boasted of 19,400 "actions" against file sharers in 2005 (IFPI 2006, p. 21). Such suits were no doubt effective at changing most targets' behaviour, and scaring others away from copying. It also cannot be doubted that they are inflicting a great deal of anguish and psychological harm on the targets for what is normal (and to many, ethical[467]) human behaviour.[468] Typical settlement amounts in those cases were only a few thousand U.S. dollars and the loss of a media library,[469] but any realistic account of the utilitarian costs

---

[467]Only 30% of the U.S. public believe that users of P2P networks should face legal sanctions for that behaviour (Rainie and Madden 2005, Q42).

[468]Even Hillary Rosen, former chief executive of the RIAA, came out against the lawsuits (Rosen 2006).

[469]The IFPI (2006, p. 21) reports an average settlement of USD $3,000.

from the stress of being sued by a large and deliberately intimidating organisation must be much higher than the immediate material losses.[470] This is more especially the case because some of the defendants are falsely accused.[471]

The campaign of lawsuits against consumers occurred outside of the United States too, although did grow to quite the same scale. In Australia, there have only been cases against defendants who have been a little more proactive in infringement than simply running a P2P client with a well-stocked directory of shared files.[472] The absence of extreme statutory damages was, no doubt, a consideration.

The RIAA gave up its campaign of direct lawsuits at the end of 2008. The cost of legal process made this method of enforcement unprofitable, even considering the substantial sizes of the settlements.[473] Other businesses, known as "copyright trolls", have tried harder to make this business model work. These organisations do not necessarily start out as copyright holders, but instead *purchase* rights to works for the purpose of litigation. They have employed numerous creative legal strategies to minimise their costs and attain profitability.[474]

Meanwhile, the major rights holders have focused their enforcement objectives on a model known as "three strikes". Three strikes laws are minimal-process mechanisms

---

[470]It is typical for the legal threats made by the RIAA or other plaintiffs to cite the truly incredible statutory damages which they could claim against file sharers if they actually went to court (EFF 2003). Under 17 U.S.C. §504 (c), rates set under the No Electronic Theft Act of 1997, the minimum damages awardable to a copyright holder are $200 per work, in the improbable case that the defendant is able to show that they "had no reason to believe that his or her acts constituted an infringement of copyright", and $750 per work otherwise (the sample letter published by the EFF mentions only the latter figure). The maximum is a great deal higher ($30,000 or $150,000 if the plaintiff can prove willful infringement). One might have expected a court to stick to the lower bound when dealing with a casual P2P user, but in the one case where statutory damages were assessed against a file sharer, the jury awarded damages close to $10,000 per song shared, and the litigation has subsequently been complex (see `https://secure.wikimedia.org/wikipedia/en/wiki/Capitol_v._Thomas` for a summary, accessed March 2011).

In practice, if copyright industries regularly used these powers to this full tornadic effect, the law would probably change, perhaps through courts accepting arguments that the statutory damages for copyright infringement are so severe as to be unconstitutional. But it should be remembered that disproportionate asset seizures are common in the American war on drugs — and that the possibility of statutory or judicial law reform in no way diminishes the shock experienced when some poor soul finds the threat in their letterbox, and when their lawyer informs them that it could be sustained in court.

[471]The EFF, for example, cites news reports of a Massachusetts grandmother who was accused and threatened over demonstrably false allegations, and of the RIAA's poetic response: "when you go fishing with a driftnet, sometimes you catch a dolphin" (EFF 2005, p. 4).

[472]It is difficult to count these cases because they tend to appear, unreported, in magistrates' courts all over the country.

[473]`http://www.techdirt.com/articles/20100713/17400810200.shtml`.

[474]Several active U.S. copyright trolls are discussed at `https://www.eff.org/issues/copyright-trolls`. These businesses also exist in other jurisdictions.

for copyright holders to have Internet users 'disconnected' (preventing from having an ISP account) if rights holders accuse them of infringement three times. Statutes of this sort have passed in a number of European jurisdictions, though it is too early to know how many will survive constitutional challenges, or to know how much collateral harm they will inflict.

The less forceful side of a vigorous enforcement and "intellectual property education" effort is the strengthening of proprietarian anti-copying social norms. More people at dinner parties, more teachers with industry-sponsored class materials, will opine that "copying is theft". Institutions and their computer systems administrators will go above and beyond the requirements of the law to prevent copying by their users. More parents will discipline their children for downloading songs.[475]

In many instances these attitudinal changes may be harmless *per se*. In others, we might want to scrutinise the black and white moral thinking that often accompanies such attitudes towards piracy. Of the greatest normative concern, perhaps, is the conflict inflicted upon children whose understanding of cyber-ethics is more nuanced and better-informed than that of their parents.

There is also extra-judicial enforcement by institutions that are outside of the copyright system but motivated either by fear of it or by its normative authority. The scale and severity is difficult to judge. A relevant though completely anecdotal observation is that within the University of Melbourne, I know of one department in which IT staff were fired for file sharing; another in which several students were fined thousands of dollars each and lost their computer access for downloading films; and another in which music, legal or otherwise, is banned from all computers (including laptops owned by staff and used both at home and at work). It is not unreasonable to extrapolate a similar pattern to other departments within the university (with which I have less contact) and across many other institutions in the public, semi-public, and private sectors. To be fair, some organisations might have made productivity-motivated decisions to prevent the circulation of cultural works irrespective of copyright; if that were the case, any harmful side effects would be a result of the search for organisational efficiency, rather than a result of copyright policy. But having said that, it is unusual to ban people from listening to CDs at work — or to fire them for illegal and far more objectionable habits, such as speeding while commuting to the office. One could at-

---

[475]A common piece of advice appears to be to take away children's access to computers if they are too keen on sharing! `http://askville.amazon.com/stop-kids-illegally-downloading-music/AnswerViewer.do?requestId=19028393`

tempt to measure this cost of the war on copying in dollar terms, but these kinds of impositions may be easier to understand as incremental losses of freedom.

This Section has identified a range of social costs associated with enforcing copyright against consumers. It would be best to have some idea of just *how* serious they are. Unfortunately, there is very little data on which to base estimates of the cost of subjecting people to this grand social engineering project. Even with data, psychological theories of the harms that accompany social regulation and the loss of freedom do not easily lend themselves to such tabulations. We can really only take informed guesses about most parts of this calculation.

One aspect of these social costs is *relatively* easy to gauge — the direct social costs of lawsuits against P2P users in America. One can wave one's hands, and say: there were around 500 cases a month,[476] multiplied by somewhere between ten and twenty thousand dollars per case.[477] The total direct social cost of the U.S. lawsuits is likely to have been somewhere in the range of a hundred million dollars per year. The troll lawsuits are comparable.

Unofficial copyright enforcement by third parties (like that observed within universities) will typically be less harmful but probably more widespread. It is not clear that we can easily keep a count of disciplinary proceedings against students, sackings of employees, or domestic disputes. But a measure of their prevalence would be the first and crucial step in placing a value upon avoiding them.[478] The development of such tools is well beyond the scope of this thesis. In their absence, all that can be reasonably concluded is that the "war on copying" is a cost of unknown magnitude — significantly less than the war on drugs, but also a fair deal more than just the cost of all the lawsuits and their direct consequences.

---

[476]There were 20,000 file sharing lawsuits in the United States between mid-2003 and mid-2006 (Sag 2006), but there was considerable acceleration of the litigation programme after the early part of that period.

[477]This includes a few thousand dollars' worth of settlement payment, a few thousand dollars' worth of files that must be deleted and therefore revert to deadweight loss, a few thousand dollars worth of legal fees, and a few thousand dollars' worth of stress and distraction from being threatened with destitution. Including the settlement involves the assumption, which may be a little bold, that the marginal utility of dollars paid to record companies in these settlements is close to zero. But in defence of this accountancy, I would point out that in all probability these payments are insufficient to defray the costs of filing the suits in the first place.

[478]In fact, the problem is not uniquely tied to copyright; any credible methodology for weighing and measuring what is lost when society uses substantially non-consensual rules to change behaviour would be of great interest for a wide range of policy problems.

9.3.2   The costs of enforcement under other regimes

In pragmatic regimes, enforcement has the task of deterring enough consumers from piracy to ensure that markets for digitised works remain economically significant. This is not a straightforward task, and we should therefore expect the costs of enforcement to be greatest under sustained versions of the status quo.

In contrast, under the strongest imaginable feudalist regimes, enforcement would no longer be necessary to modify the behaviour of the public at large. Infringement would simply be impossible for most people.[479] Given such circumstances, policing and legal action could be reserved for talented hackers, subversives, and professional pirates. Such extremes of feudalism are, however, unlikely; it is more probable that enforcement against the public would need to continue on some scale to prevent the sharing of works that had leaked out of DRM, one way or another, or which had previously been published in unencrypted forms. It follows that strong DRM regimes would probably have social enforcement costs which are similar to those under the status quo.

Virtual markets have policing and enforcement costs that are accrued when staff must be paid to set up honeypot computers to detect fraud, pore over the logs of those systems looking for malware, cancel the effect of that manipulation on the remuneration system (either with vote cancellation or trusted computing patches), and try to track and prosecute those who are responsible. In virtual markets based on sampling, significant expenses might imposed on systems administration to keep the computers belonging to the sample group secure. Although these direct costs of enforcement are significant, there is not the same potential for social collateral damage from the endeavour. Most people's behaviour does not need to be changed; in the same way that electoral vote buying is not a widespread social phenomenon, we could expect attempts to manipulate a public funding system to be restricted to a small number of more serious perpetrators.

Whether or not there would be any policing and enforcement costs at all under anarchy is a matter of definition. Certainly, there could be some expenses associated with enforcing various contracts or contingent-payment pledges, but it could be argued that these are more properly categorised as transaction costs. Either way, these costs do not promise to be large by comparison to those under other regimes.

---

[479] As discussed in Section 2.3.1, these variants of information feudalism would require not only secure DRM, but the eradication of all forms of P2P software, the discovery and deployment of secure watermarks and traitor tracing systems.

9.4   Why do security infrastructure costs vary widely?

This chapter has identified some fairly large differences in the costs of technological infrastructure, especially between the the strong DRM regime and alternatives. In this section I will point out a couple of fundamental reasons why this must be so.

9.4.1   Fragility

The rough estimates above for the price of effective leak prevention for DRM[480] are by their nature imprecise. But they accurately reflect the kinds of expenses necessary to avoid "competing with free" (that is, having to make licensed purchases more attractive than available pirated copies).[481] Because of what some commentators have termed the "break once, run anywhere" or "break once, run everywhere" principle, a DRM network is only as strong as its weakest link. It is an elaborate and fragile structure for providing files to millions of people.

In order to reduce the losses from break-in at a few unknown but crucial points, extreme precautions must be taken everywhere. Under any DRM regime, this involves reducing the impact of break-ins by vigorous law enforcement; Under strong DRM, it also involves making break-ins so hard that they are genuinely rare. This requirement for security everywhere is a fundamental problem of exclusive digital copyright, which I have attempted to illustrated in Figure 9.1.

It is useful to compare this problem to the task of securing an alternative compensation system. Both DRM and virtual market networks contain security-critical points of failure. In a DRM regime, these are all located in the devices in users' homes (illustrated in Figure 9.1). In a virtual market, the security-critical systems are the handful of government-run computers which allocate the rewards (as shown in Figure 9.2). The actions of small groups behaving contrary to the specifications of the system cannot threaten the stability of the network as a whole.

A structural consideration which further exacerbates this gap between DRM and public funding regimes is the absence of transaction reversibility. Because copyright regimes act to make copies of works valuable in and of themselves, once pirates have started distributing a particular file widely, they have already made off with its value.

---

[480]See chiefly Section 9.2 for strong DRM, and Section 9.3.1 for pragmatism

[481]Note of course that 'free' in this expression only means zero dollars as an up-front payment; it is not free in terms of the effort required to find a piratable copy, or in terms of the risk of being sued for copyright infringement.
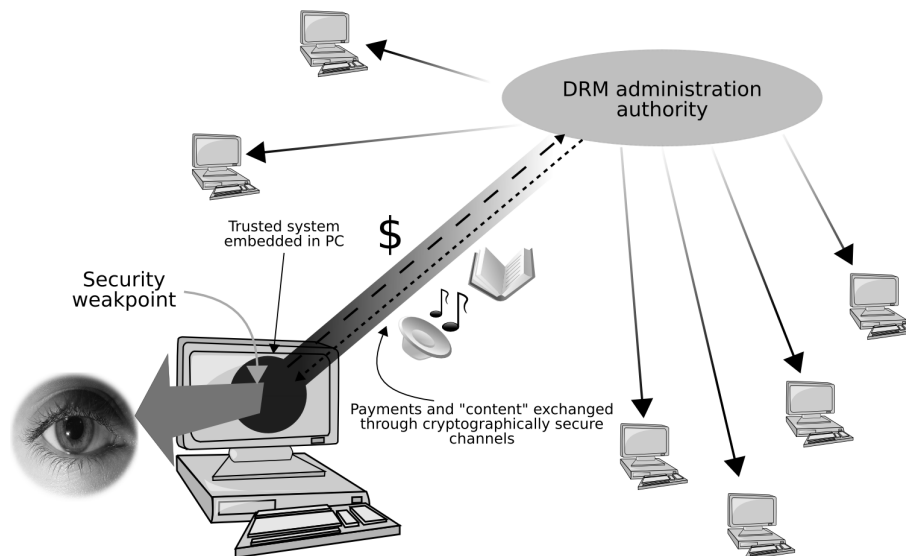
Figure 9.1: Information flows in a DRM system: one weak point per client

In almost all cases, there is no way that their piracy can be "rolled back" even if the means by which they obtained the file are discovered.[482]

In contrast, if it is at any point discovered that computer security breaches have led to inappropriate allocations of virtual market funding, it is a comparatively straightforward matter to either reclaim the cash, or at least adjust future remuneration to compensate.

9.4.2   Incentives and the Stability of Network Protocols

There may be another fundamental reason why the costs of security in secure DRM are much greater than in alternative compensation systems. Every stable, successful digital network ever built, has operated on an unstated principle — *the vast majority of participants want the network to function*. The Internet is, of course, the most striking example of this phenomenon. Participants attach computers to the Internet, and those computers execute code which is generally compliant with a set of agreed standards for communication. Even a small proportion of defecting nodes that subvert

---

[482]An exception to this point would be authorisation hacks that trick DRM systems into licensing works to people for free, without actually decrypting them out of their cryptographic/TC cages. If such files were discovered, the licenses could be revoked.
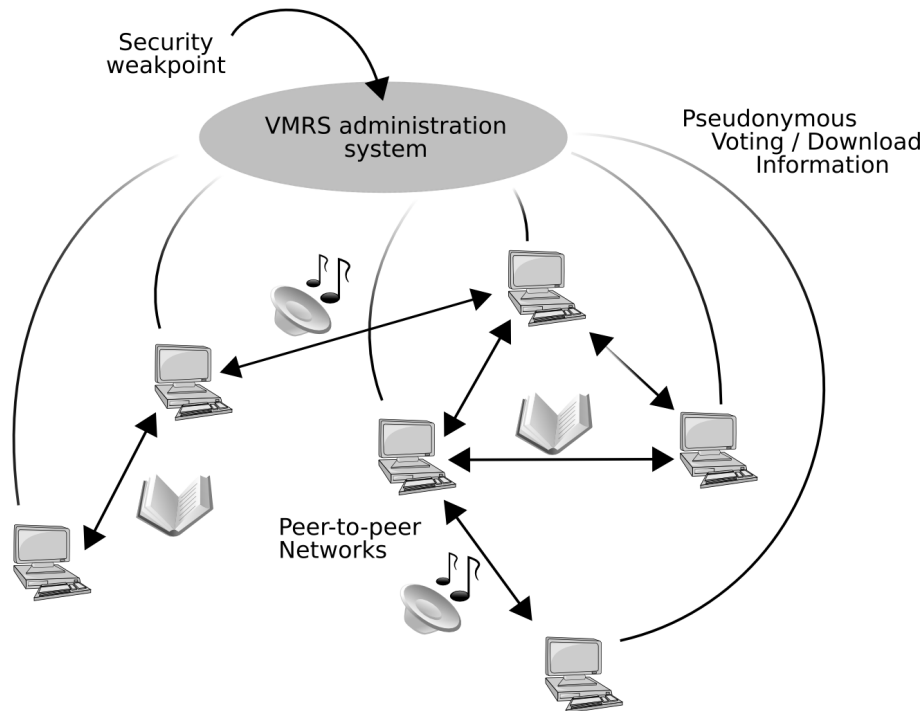
Figure 9.2: Information flows in a virtual market.

this arrangement can cause serious network problems.[483] One example of this is the effect that a tiny minority of spammers have had on SMTP-based email infrastructures (Geer 2004).

A DRM network cannot expect cooperative behaviour from its participants. As peer-to-peer file sharing has demonstrated, many users are eager to exchange copyrighted information, with disregard to publishers' and authors' legal privileges. Whilst most of these actors will not have the skill to write software which attempts to redefine the network, they would be more than willing to download and run it. Designing a network which functions whilst many of its participants work to cause its downfall is a problematic and expensive proposition.

The incentive viewpoint also explains why anarchy has the lowest overall computer security expenditure. Information anarchy is extremely cheap because it doesn't regulate people's behaviour at all. This suggests another way of thinking about the design goals of publicly funded compensation systems: governments should aim for a

---

[483]The work of Nisan and Ronen has sparked a growing literature on incentive-compatible network protocols, but this literature does not consider the messier question of enforcing participation in an intrinsically unpopular protocol. See (Nisan and Ronen 1999).

solution to the free rider problem that requires the smallest changes to consumer's private incentives and behaviour. The downside, in the case of virtual markets and private copying levies, is requiring large changes to some firms' behaviour (such as making ISPs and electronics manufacturers pay levies), but those changes are much cheaper to effect.

9.5   Conclusion

This chapter has investigated the comparative infrastructural overheads that are associated with organising the finance of digital culture in different ways. The differences between the four regimes under consideration are large.

Strong or "feudalist" DRM regimes fare most poorly in this analysis. Two different back-of-the-envelope estimates were made of the cost of preventing pirates from ever obtaining DRM-free perfect digital copies of works. One figure was US $20–$150 per consumer device, for tamper-resistant trusted systems. The other was a more 'holistic' analogy from economic losses reported due to computer security breaches; it suggested that such breaches could not be prevented for less than US $150 per computer.

If the numbers at the higher end of this range are accurate, the only reasonable conclusion is that strong DRM is impossible. Expenses at the lower end — starting at US $20 or a little over AUD $25 per device — are not so inconceivable. Copyright industries would be very keen to add costs at that level to consumer electronics if it meant the end of widespread piracy, and perhaps heavily DRMed platforms like the XBox 360 and Playstation 3 already include DRM expenditures of that order. In Australia, the consequent overheads would be in the order of hundreds of millions dollars per year.[484] Costs of this magnitude leave strong DRM at a decided normative disadvantage.

The weaker cousin, pragmatic DRM, presents a rather different balance sheet. The hardware costs for pragmatic DRM need not be nearly so exorbitant. With research and development costs amortised globally, the overheads within Australia might remain in the tens of millions of dollars per year.

---

[484]In the 2004–5 financial year, Australians purchased 3.6 million relevant digital media devices: 1.7 million DVD players, 0.8 million consoles and 1.1 million digital audio devices (Maskill 2005). These figures do not include computer sales, which are also relevant. Australian PC shipments stood at 3.29 million in 2004; these should not be included wholesale because it is likely that even if tamper-resistant trusted systems became mandatory for many media downloads, a fraction of PCs would continue to be built without them.

The most worrying aspect of a pragmatic digital copyright regime is its need to persecute behaviour that falls well within ordinary social norms. This is an invitation to a "war on copying" to sit alongside the "war on drugs". Indeed, that war is well underway in America. Although the harmful side effects of copyright law enforcement are unlikely to ever reach the amazing levels associated with drug policy, the injurious potential of coercive intellectual property rules should not be underestimated. Satisfactory quantitative measures are beyond the scope of this thesis, but losses in the U.S. are presently at least in the order of hundreds of millions of dollars per year.

With respect to computer security expenditures, virtual markets are quite comparable to the status quo. Public funding does require extensive infrastructure, including trusted systems or other means to collect reliable usage and voting data from participating users or sample populations. Like pragmatic DRM, but unlike strong DRM, this architecture need not meet the radically challenging design goals of incentive-incompatible security in everyone's living rooms and economic reliability without the ability to reverse fraudulent or piratical transactions.

The costs of infrastructure associated with information anarchy are the hardest to gauge as we cannot be certain of the design principles they would follow. We can be certain that the costs of policing and law enforcement would be close to non-existent. But we cannot tell how much would be spent on escrow publishing, sophisticated tip jars, advertising agencies, and numerous other gap-stopping institutions. It is nonetheless difficult to see how, if subject to competition, they could swallow much of the smaller pie of cultural expenditure that would be available in the absence of effective exclusive or other remuneration rights in digital works.[485]

When all of these observations are combined, there appears to be a fairly clear ranking of regimes by their infrastructural price tags.

In Australia, strong DRM would cost hundreds of millions of dollars per year, mostly spent on designing and building millions of devices to be "trustworthy" and resistant to physical interference, but also through enforcement against hackers and people who share leaked or otherwise unencrypted files. Pragmatic preservation of the status quo would cost somewhere in the high tens of millions of dollars per year, mostly due to the direct and indirect costs of a 'war on copying' but also through DRM expenses. Virtual markets would require similar expenditures (perhaps a few tens of

---

[485]An important exception to this point would be if either network externalities or very high fixed costs applied to the intermediate businesses supporting cultural production, allowing these support services to be monopolised.

millions per year) to those currently made on DRM R+D and deployment, but they do not require major re-education and law enforcement campaigns to recast consumer behaviour. Although numbers are hard to obtain, anarchy can be expected to have the lowest infrastructural costs of all.

Despite the aspects of uncertainty, there appears to be a clear order of preferability with respect to the systemic costs of making each regime function. In order of decreasing desirability: anarchy, public funding, the status quo, and finally, strong DRM.

Chapter 10

The Cost of Taxation

Contents

The tax man's taken all my dough,
And left me in this stately home,
Lazing on a sunny afternoon.
And I can't sail my yacht,
He's taken everything I got,
All I've got's this sunny afternoon.

— The Kinks

One of our four regimes — the virtual market — involves the collection and use of taxation to pay artists and authors. Tax systems are complicated social institutions of their own, and their use to solve economic problems comes at a price. One might wonder, how large is that price?

This is a thesis about copyright systems, and the objective is not to essay on the economics of taxation. Unfortunately, the answer to the research question depends, in a not entirely simple way, on the economics of taxation and at least a brief excursion in that direction is necessary.

The literature on alternative compensation systems has not confronted this subject in any serious manner. Some authors mention the cost of taxation, but not in a way that attempts to actually weigh and measure it against the upsides of a liberalised network. Others make assertions about these costs that are poorly supported[486] or that rest on controversies in the literature.[487]

Following a survey of the types of costs of taxation (Section 10.1), this chapter makes two contributions: the first is an attempt to identify what kind of taxes would work best for ACSes (Section 10.2). This is a compromise that requires simultaneously trying to be as fair, non-distortionary, equitable, and politically feasible as possible.

The second objective, in Section 10.3, is to provide at least a basic estimate of how much it would cost society to use taxation to pay for a virtual market. Unintuitively, this is not just a number that can be looked up in a textbook on public finance. If

---

[486]For instance, Duffy (2004, Parts I.A.1, III.D) argues that a failure to measure distortionary costs of taxation is a major weakness in the literature advocating ACSes. This is a reasonable, but he also tries to argue, without any substantial evidence, that the limited nature of IP monopolies means that deadweight losses will be less than distortionary costs of taxation.

[487]Shavell & van Ypersele recommend income taxation as the most efficient way of funding a reward system to replace patents; see (Shavell and van Ypersele 2001), note 45 and accompanying text, citing the optimistic results of Kaplow on the efficiency of funding public good production through income taxation; see (Kaplow 1996). It should be noted that Kaplow's claims are controversial (Browning and Liu 1998; Kaplow 1998), and in particular that his proposal requires each person's valuation of cultural information goods to be a function solely of their income. Benefit offsetting income taxation is not feasible when funding public goods that are complements to leisure.

economic theory is to be believed, the answer depends not only on (a) the particular kind of taxation that is used, but also on (b) what the taxation pays for, and on the interactions between (a) and (b)! Informed by this complexity, I will nonetheless just use some appropriate ranges from a textbook on public finance.

10.1   Types of Taxation Overheads

10.1.1   Distortionary Costs

Much of the complexity of the cost of taxation is due to a phenomenon that is variously called the distortionary cost of taxation, excess burden, the cost of funds, or the deadweight loss of taxation.[488] Taxation is said to be distortionary when it causes shifts between the production and consumption of one kind of good and another, causing divergence from the "natural state" of a free market. Distortions are usually expected to decrease social welfare, unless they act to correct externalities (such as the pollution caused by driving a car) or redistribute wealth in a way that decreases inequality. But the nature of these costs is somewhat subtle, and perhaps best introduced with some examples.

If apples are taxed while oranges are not, say, then people are likely to eat more oranges. It isn't initially obvious or intuitive that this is a bad thing, but there are some fairly general assumptions under which distortions are, in fact, undesirable.

A fairly compelling illustration of this undesirability is called the "welfare trap". Suppose that a government raises its revenue via an income tax, and that one of the services it provides to its citizens is a welfare payment to the unemployed. Suppose, as is typical, that these payments are only made to people who earn less than a certain amount of money each year.[489] Consider then the decision faced by Amy, an unemployed singer who currently earns, $15,000 per year in welfare payments. Suppose that she is offered a 30-hour-a-week job that would pay $25,000 per year, and that if she took the job she would cease to receive unemployment payments and would, instead, pay an effective tax rate of 20%, leaving her with $20K to spend. Because of the loss of welfare payments and the onset of income tax, Amy's effective hourly

---

[488]I will try to avoid this last term, since distortionary costs of taxation are not at all the same thing as the deadweight losses of monopoly that are commonly discussed in literature on the economics of copyright.

[489]In many cases, such welfare phase out with a sliding scale as the the recipient's income rises, but this isn't important for the example.

wage would only be $3.33 per hour, and she might quite rationally decide that it wasn't worth spending 30 hours of her week that way! If she decided to remain unemployed we may have identified a situation in which everyone would be better off if the tax system were different — for instance, if as an inducement to work Amy's tax rate was set to 0%, and she continued to receive $5,000 of her previous unemployment payments. In that different world, she took the job, the government would save $10,000 and Amy (by presumption, since she took the job) would be better off for working 30 hours for the marginal rate of $10 per hour.

The example of the welfare trap illustrates how taxes can "distort" taxpayers' incentives in counterproductive ways, and also how the distortion is a product of the interrelationship between taxation and the services that governments use taxation to provide.

10.1.2 Compliance costs

Compliance costs are the effort and resources expended by taxpayers in the act of paying their taxes. These costs might not increase significantly if virtual markets were funded with existing types of taxation, but there would be new compliance costs if new types of taxes — most probably levies — were used. If the levies were incorporated into the prices of various goods and services (Internet connections, blank media, electronic devices), the direct compliance costs for consumers would be negligible. But manufacturers and service providers would nonetheless bear them, and pass them along indirectly.

Representative compliance cost estimates for taxation have been 1.7–3.3% in Canada (Vaillancourt, Clemens, and Palacios 2008); 5-7% in the United States (Slemrod and Sorum 1984), and 7.3–9.9% in Australia (Pope 1993, p. 74). I will use an estimate of 2–10%.

10.1.3 Administrative costs

Most of the relevant administrative costs for a virtual market were discussed and compared across regimes in Chapter 9. There would be a few extra costs involved in taking receipt of taxation from levy-paying firms, but those processes are far less complicated than actually distributing funds to the beneficiaries of an alternative compensation system. I will therefore not at attempt to model or estimate them, although the typical

figures in the literature are in the order of 1–3%.

10.1.4   Potential redistributive benefits

Quite apart from their distortionary consequences, taxes usually have the effect of redistributing wealth. That is both because some people pay more tax than others, and because some receive more benefit than others from the goods that governments provide.

If the net redistribution is away from wealthier individuals and towards poorer individuals, we may have grounds to regard it as inherently beneficial. The nature and conditions on which such a preference rests depend on whether we take a utilitarian or a justice-based ethical position. In the utilitarian view, the redistribution is good if the beneficiaries have greater marginal utilities of wealth than the net-payers. On average, poor people benefit more from dollars, so redistributing dollars to them is good.[490] In a justice-based view the causes of poverty matter a little more, but wherever they are beyond the control of a poor person,[491] a movement of wealth to that person is likely to be good.

Some advocates of free markets acknowledge the need for wealth redistribution, but argue that it should be kept completely separate from the provision of services by governments. The idea is to have society decide democratically upon how much redistribution it wants, and then to use markets to provide private goods such as healthcare and education according to how much people are willing to pay from them. This argument effectively claims that redistributive effects should not be counted as a benefit from a virtual market, because they could be provided on their own.

There is however an important counter-argument to that position, because democracies can be very bad at setting a satisfactory level of wealth redistribution. The most dramatic example is perhaps the United States, where voters demonstrate a staggering lack of understanding of the scale and nature of inequality within their own society.[492] In the face of such failures of democratic common sense, it may be desirable to estab-

---

[490]Although some causes of poverty, such as mental illness, may also lead to low marginal utilities of cash.

[491]One example of such a cause is being born to financially poor parents. Another is being born to less savvy parents than others. Another is attending a school that offers less support and encouragement, or fewer opportunities. Etc.

[492]Brooks (2003), for example, reports a survey finding that 19% of U.S. citizens believe they are within the wealthiest 1% of the population, while a further 20% believe that they will be, at some point during their lives.

lish linkages between wealth redistribution and the provision of basic services, such as education, health, and public transport. Making progressive wealth redistribution a side-effect of funding such services may ameliorate democracy's failure to provide sufficient redistribution as a stand-alone project.

For this reason, it is legitimate to count redistributive effects as a potential benefit of virtual markets and other publicly funded compensation systems.[493] Whether the benefits are actually there depends, of course, on the particular combination of revenue sources that are used to fund the system. It is only when the taxes are raised on income, or wealth, or on goods and services whose consumption is disproportionately higher amongst wealthier people, that these benefits will actually be realised.

10.2   Which Tax Model is Best?

There is a wide range of taxation models available for funding alternative compensation systems. The criteria for choosing among them include political feasibility, the degree to which the tax "fairly" charges those who benefit more or less from the things funded by the virtual market, and the size and nature of distortionary costs.

The first criterion is important both in an inherent and a political sense. It is that consumers should pay tax that is closely in proportion to the amount they would have spent on relevant information goods, had a non tax-based system been employed.[494] The intuitive "fairness" of the system, to a large extent, depends on this proportionality.[495] But there are other considerations which should also be given weight. Charging more from those with greater ability to pay, and adopting a tax formula which is easy to enforce and difficult to sidestep, will improve both normative fairness and utilitarian efficiency. Tradeoffs must be made between these criteria.

The most straightforward solution is to raise levies on goods and services which are directly complementary with the consumption of digital culture.[496] This is the

---

[493]For a theoretical parameterisation of distortionary costs as a function of the degree to which inequality is considered undesirable, see (Dahlby 2008, 38–42).

[494]Ginsburg, for example, criticises the use of digital private copying levies on the basis that not all consumers want to pay for "all you can eat" downloads (Ginsburg 2001).

[495]Some theoretical treatments of fairness — along Rawlsian lines, for example — would lead to conclusions that were strongly dependent on the circumstances prevailing before the system's introduction. The simpler, intuitive kind of fairness is of political importance regardless of whether there are good arguments to say that it is misconceived.

[496]In general, taxes used to fund public goods production will be less distortionary if they are raised on goods whose benefits are *complementary* to the public good (Boadway and Keen 1993). That is,

approach used in existing private copying systems,[497] and has been advocated for use in alternative compensation schemes.[498] The obvious candidates for these levies include Internet connectivity, bandwidth, blank storage media (recordable CDs and DVDs, and perhaps even hard disks), Hi-Fi equipment, specialised devices for watching, listening to or reading digital culture; or even computers in general.

The chief drawback with these sources of revenue is that they are only imperfect proxies for the underlying consumption (or otherwise) of the information goods in question. Taxing a whole class of activities or gadgets, some members of which did not implicate copyright in the first place, amounts to a cross-subsidy to those that do. Beyond a certain point, such cross-subsidisation may be seen as unfair, and it may, even at small rates, also cause distortionary economic side effects. To illustrate: taxing blank storage devices and media discourages backups.[499] Taxing bandwidth usage could encourage people to switch to less data-intensive activities: listening to music rather than watching a film, or from downloading songs to downloading books. Taxing MP3 players and e-books too heavily could encourage the use of general-purpose computers in their place; taxing computers avoids that problem but may lead to others if it means that some people delay upgrading their machines, suffering with an old, slow computer on account of the tax.[500]

Alternatively, virtual markets could be funded solely by general revenue sources, such as progressive income taxation.[501] Although income taxation causes distortions in the way that people choose to work, these may be balanced by potential redistributive

---

consumers who benefit from the public good are more likely to wish to consume the private goods that are being taxed to produce it. Suppose, for example, that the public good was a set of upgrades and improvements to street lighting infrastructure. Funding that public good with a tax on cars would be less distortionary than funding it with a tax on televisions, because the street lighting is complementary to the cars.

[497]See, e.g., (Copyright Board of Canada 2003) (setting levy rates on cassettes, recordable CDs, and memory in dedicated MP3 players); (ACC 2001, 9–10) (tabulating the kinds of levies used in national private copying schemes).

[498]See, e.g., (Netanel 2003, at 43–4); French "global license" proposal (see footnotes 156–157 and accompanying text)

[499]Blank media taxes might also, theoretically, affect those who use such media for their own artistic creation. But because the cost of digital media is a small component of the cost of creative activity, this effect is unlikely to be substantial.

[500]This would be a serious problem if those affected included business users of computers, because the particular information goods provided by the virtual market would not be of much use to businesses, assuming the scope discussed in Section 5.4.3.

[501]Taxation is said to be *progressive* when the rate of incidence is higher on wealthier individuals. Fisher, for example, advocates this kind of revenue source; see (Fisher 2003, at 24–5) (settling on income taxation as the most desirable, if not the most politically feasible, way to raise revenue for an alternative compensation system).

welfare improvements.[502]

An argument from fairness against the use of general revenue sources is that there is a redistribution of resources, from taxpayers who do not use the Internet for cultural consumption, to support the creation of digital culture. This will occur, although it is not necessarily problematic. Incentives to produce digital writing and music will (almost always) result in more and cheaper works in tangible forms; and if a progressive form of taxation is employed, the demographics which pay disproportionate taxes are precisely those in which Internet usage is most pervasive. But the point may be moot, if it is possible to ensure that the progressive tax is only paid by those with Internet access.

An income tax used to fund artistic production could simultaneously be constructed as a surcharge on Internet connectivity. This could be done directly, by having the tax agency calculate a surcharge on each Internet connection, or indirectly, by offering a tax credit or reduction to citizens who do not have Internet connections. Theoretically, that approach could combine the merits of both progressive and levy-based funding models, though its novelty makes it less politically conceivable and would increase compliance costs.

Another way to achieve progressive levies on Internet access would be through a surcharge on residential connectivity that was proportional to the valuation of the connected residential property. This structure has the benefit of being less distortionary than an income-dependent levy,[503] although in the United States, it would be impractical for historical legal reasons.[504] The other hurdle for these particular levies is the

---

[502]Shavell & van Ypersele recommend income taxation as the most efficient way of funding a reward system to replace patents; see (Shavell and van Ypersele 2001), note 45 and accompanying text, citing the optimistic results of Kaplow on the efficiency of funding public good production through income taxation; see (Kaplow 1996). It should be noted that Kaplow's claims are controversial (Browning and Liu 1998; Kaplow 1998), and in particular that his proposal requires each person's valuation of cultural information goods to be a function solely of their income. Benefit offsetting income taxation is not feasible when funding public goods that are complements to leisure.

[503]See Chapter 10. I would like to thank Jamie Love for suggesting this solution to the problem of distortions in labour supply.

[504]See U.S. Const. art. 1, §§2, cl. 3 & 9, cl. 4. (requiring that direct federal taxes be apportioned among the states according to their populations); (Kinsler 2003, § I.B) (explaining the origins of the apportionment clause in Southern fears of a disproportionate burden from taxes on land and slaves, and the enduring constraints on the nature of federal property taxes). Constitutional complications to ACS taxes exist in Australia, too. In the 1990s, the courts held that legislation placing private copying levies upon blank audio cassettes was unconstitutional (High Court of Australia 1993) because these levies were neither a quid-pro-quo royalty collected in exchange for licenses (since home copying for private and domestic purposes was non-infringing) nor a constitutional tax (since the levies were not placed into a consolidated revenue fund). Presumably alternative compensation systems for file sharing could

difficulty of associating wireless Internet connections with physical residences; were it not for the possibility that wireless networks will become a primary component of Internet access infrastructure, I would advocate this mechanism as the best revenue source for virtual markets. As it stands, property value levies might be practical only once Internet connectivity is effectively universal — since at that point, the levy could be charged on all occupied residential properties, rather than just the ones with network connections.

An important property of tax distortions is that under many conditions, they often grow non-linearly with the tax rate.[505] It follows that broad based lower taxes will often be less distortionary than more narrowly based, higher ones.

Given that there are a range of possible ACS funding sources (Internet connections, blank media, media devices, certain kinds of computer equipment, general income or residential property), distortion may be practically minimised by drawing portions of revenue from several of these sources.

The use of combined tax sources would have ambiguous political consequences. On one hand, a broader range of interest groups would be affected by parts of the tax, making them into potential veto constituencies against the ACS. On the other hand, the actual tax rate that any group — such as ISPs or equipment manufacturers — would be asked to accept on their products would be significantly lower. This might make the ratio of benefits to costs more persuasive for them.

When these political and economic considerations are taken together, a hybrid such as a one-third-each split between consumer electronics/blank media, Internet connections, and general revenue seems like the best way to fund alternative compensation systems.

Conclusion

The taxation options to support virtual markets present a spectrum; different choices will hold different implications in terms of wealth distribution, subsidy effects for various industries, distortions, and constraints imposed on the total level of virtual market funding. This section should have convinced the reader that there is a wide range of choices available, that these choices are capable of serving a range of different normative goals, and that particular social and political contexts will play a major role in

---

be constructed in a way that met either or both of these constitutional requirements.

[505]See for instance (Creedy 2003), offering a simple model to illustrate this effect.

deciding between them.[506] A completely satisfactory investigation of these issues is beyond the scope of this thesis, but the provisional conclusion that ACSes should be funded from hybrid revenue sources seems promising.

10.3   Accounting for Distortionary costs

The theoretical literature on funding public goods with distortionary taxation presents a complicated picture. The conclusions depend not only on the precise taxes being used, but on the precise public goods being funded, the mathematical structure of consumer preferences,[507] and on the set of taxes and public sector services that are already in existence.[508]

These difficulties with mathematical microeconomic approaches, combined with the inevitable discrepancies between rational-actor models and actual human behaviour,[509] make it doubtful that we could use purely theoretical approaches to get accurate figures for the distortionary overheads of tax-based ACSes. Empirical methods are the only reasonable way around this problem, although good standards for those methods do not yet exist.[510] Further complicating matters is the persistent use of non-comparable unit systems in analysing distortionary costs.[511]

Despite these difficulties, the literature does provide at least a range of estimates for distortionary costs of different kinds of taxation. I am going to work with estimates of excess burdens from two different kinds of taxes: levies on complementary goods like broadband connections, hi-fi equipment, digital audio players and e-book readers; and progressive income taxes.

---

[506]This situation is in fact typical of the state of normative knowledge about taxation generally; see, for example Hettich and Winer (2005).

[507]Suppose that there is a private good a and a public good b, a few obvious possibilities are that consumer welfare is proportional to $a + b$; $\sqrt{a^2 + b^2}$; ab, or $\min(a, b)$. These cases led to radically different theoretical conclusions about distortions from funding the public good with particular kinds of taxation, and of course these are just a few of the simplest possibilities.

[508]See (Browning 1994)

[509]For a discussion of some of the likely implications of behavioural economics for distortionary tax theory, see (Congdon, Kling, and Mullainathan 2009).

[510]See, *e.g.,* (Dahlby 2008, Sections 3.8, 5.6, 7.4), disclaiming the lack of consistent methods in empirical studies of tax distortions generally; (Baylor 2005), making the following observations about the dynamic macroeconomic studies in particular: "A comprehensive assessment would require fitting the myriad of existing frameworks to a particular economy at a particular point in time, simulating equivalent tax experiments and reporting the results in a standardised fashion. Since such a task is precluded by its enormity, one must veer towards cruder methods of comparison."

[511]See (Triest 1990; Browning, Gronberg, and Liu 2000, Dahlby 2008, pp. 43–47)

10.3.1   Distortionary costs of bandwidth or hardware levies

In general, empirical studies of the distortionary costs of general consumption taxes obtain numbers of between 5% and 25% for the marginal excess burden (or equivalently, 1.05–1.25 for the marginal cost of funds) of these taxes (Dahlby 2008, Table 3.1).

However, it must be remembered that virtual markets and alternative compensation systems are complementary to private goods like Internet connectivity or consumer audio equipment. It follows that distortions from taxing these private goods would be lower than they would be from taxing consumption in general (Boadway and Keen 1993). Put simply, although the tax would make Internet connections or iPods more expensive, it would also make them more useful. The net effect on consumer behaviour would be smaller than that of the same levy used for some unconnected purpose.

I will therefore use a range at the lower end of the results in the empirical literature (5-15%) for estimating the marginal excess burden of such levies.

10.3.2   Distortionary costs of income taxes

The range of values obtained in empirical studies of the distortionary costs of income taxes are generally somewhat higher than for consumption taxes. Most of the figures in the literature range from 10–50%, with some outlying numbers below or significantly above that (Dahlby 2008, Table 5.3).

As discussed in Section 10.1.4 above, the use of income taxes to produce public goods may have significant benefits in countries with higher levels of inequality. There is a three-variable correlation between lower income taxes, higher inequality, and a lower marginal costs of funds for income taxes.[512] Clearly, there would be a stronger case for using income taxes for virtual markets in countries with more inequality and a lower marginal cost of raising taxes through income tax. Unfortunately, with the exception of one Japanese study (Bessho and Hayashi 2005), most empirical papers on the excess burdens of income taxation do not model this tradeoff.

I will proceed with an estimate of income tax distortions as being between 20-40% of the revenue raised, which is the middle band for the figures in the literature. On

---

[512]See (Kleven and Kreiner 2006), estimating the marginal cost of income taxation in five EU countries in 1998 as being: 1.26 for the UK, 1.52 for Italy, 1.72 for France, 1.85 for Germany, and 2.20 for Denmark. It would make a lot more sense to use income taxes to fund virtual markets in the UK (or the US) than it would in Denmark.

one hand, this figure may be pessimistic because it does not account for the redistributive benefits of the taxation; on the other hand, because entertainment goods are complementary to leisure time, using income taxes to fund entertainment may be more distortionary than it would be for funding other public goods. The distortion might be higher than this, especially in societies with high tax rates and low levels of inequality. In those cases, income taxes would clearly be the wrong choice of revenue source for a virtual market.

10.3.3   The net benefit of an ACS after tax distortions

Recall that Equation 7.17 identified a condition for the preferability of virtual markets over the status quo:

$$\gamma + x + \frac{x}{k} > \frac{T}{\alpha A} \tag{7.17}$$

T was the total cost of the tax components of the virtual market (adjusted downwards to account for the transaction cost benefits that were identified in Chapter 8, but assume for now that those are zero). αA was the social value of all the works produced as a result of copyright incentives, and $\gamma > 0$ was a parameter capturing the net incentive benefit of the virtual market over copyright; x was the proportion of the social value of works that is lost to artificial scarcity, and k was the value ratio between works that require copyright incentives to exist and those that do not, but are affected by copyright anyway.

What is the relationship between T and the marginal excess burden estimates (5–15%, or 20–40%) for levy or income tax funding?

$$\text{MEB} = \frac{\text{distortion}}{\text{amount raised}} = \frac{T}{I \cdot (\alpha A + \beta B)} \tag{10.1}$$

I is the fraction of the social value of works that is paid to artists; this factor is present because it is only that portion of the social value that needs to be collected in taxes. An estimate of this quantity for music downloads can be inferred from (Waldfogel 2010, Table 6), which is $I = 0.34$.[513]

---

[513]The largest distortion would occur in Waldfogel's "file sharing not possible" regime (with sales displacement at 0.28). That scenario gives $I = \frac{7.34}{13.13+8.5} = 0.34$.

$$\text{MEB} = \frac{T}{0.34 \cdot (\alpha A + \beta B)} = \frac{2.94}{(1 + \frac{1}{k})} \cdot \frac{T}{\alpha A} \qquad (10.2)$$

Recall also that the working assumption from Equation 7.18 was that three-quarters of the social value of entertainment was due to copyright incentives ($k = 3$). Reinstating this assumption and rearranging gives us:

$$\frac{T}{\alpha A} = 0.45 \cdot \text{MEB} \qquad (10.3)$$

Substituting back into Equation 7.17 and

$$\gamma + x + \frac{x}{3} > 0.45 \cdot \text{MEB} \qquad (10.4)$$

Equation 7.7 found bounds for x: $0.35 \leq x \leq 0.49$

$$\text{MEB} < 1.32\gamma + [1.03, 1.45] \qquad (10.5)$$

So, given these figures and reasonable assumptions in the music industry, and taking $\gamma = 0$, virtual markets would certainly be preferable if the marginal excess burden of taxation was less than some fraction between 103% and 145% of the amount raised. With the estimates of MEBs of the available revenue sources (2–10% for compliance costs plus either 5–15% for levies, or 20–40% for income taxes), this is very likely to be the case.

10.4   Conclusion

This chapter aimed to identify some reasonable kinds of taxation that could be used to fund alternative compensation systems, and to identify some reasonable estimates for the overhead costs that those taxes would carry.

Such overheads are a significant reason for disfavouring public funding as a solution to the problem of digital copyright. Compliance costs are a non-trivial concern, but their magnitude remains 10% or less of the virtual market size. Distortionary costs are a more complicated phenomenon whose size is highly dependent on the particular set of taxes used to collect revenue, the public goods produces, and also on the particular society where they are employed.

Although detailed measurements of the distortions that might accompany ACS

levies or taxes are beyond the scope of this thesis, one can work with the ranges reported for comparable taxes in the extensive, if dissonant, literature. Those were 5–15% for levy-type taxes on complementary consumption, and 20–40% for income taxes. Adding 2–10% compliance costs, we get 7–27% overheads for levies or 22-50% overheads for income tax. Expressed as fractions of the total value of the music industry, these ranges would be 3–12% and 10–23%, respectively.

Those are substantial costs, but at least in the case of music production, they are not large enough to outweigh or even nearly outweigh the factors that discussed in previous chapters that favour alternative compensation systems.

Part V

Conclusions

Chapter 11

Conclusions

Contents

This thesis set out to determine how we should pay our artists and authors now that the Internet has become the primary means of distributing their work. What kinds of arrangements are possible, and how well would they function?

The method employed to answer the question involved two steps. The first, undertaken in Parts II and III, was to propose four "regimes". Those categories are not precise, and they may not be exhaustive, but they do capture the essential features of important kinds of digital copyright and alternatives, while abstracting away the rest of a very complicated space of possible practices, laws, technologies, and social norms. Chapters 2 through 5 motivated and introduced these regimes.

The second step, undertaken in Part IV, was a systematic normative comparison of the regimes. Each chapter there discussed a different category of effects and phenomena that might offer grounds for preferring one regime over another. These were mostly economic considerations, although technical, ethical, and other issues turned out to be important in many places. Part IV collected the pieces required to say which regime

we might prefer, and on what grounds. Now it is time to put those pieces together.

11.1   Normative Results: Which Regime is Best?

Table 11.1 summarises the findings from Part IV of the thesis with respect to the music industry in particular. Note that these measures cannot all be expressed in precisely the same units: some are measured as a proportion of total social value and others are measured in mean deviation from optimal incentives. These would only be equivalent if an extra dollar of incentives for some type of material led directly to an extra dollar of social value from that material.

CHAPTER (and sections)	STATUS QUO	FEUDALISM	VIRTUAL MARKETS	ANARCHY
Artificial scarcity (Chapter 6)	0	- 18–32%	+ 55–98%	+ 55–98%
Incentives (Chapter 7)				
Transparency effects (§ 7.5.2)	0	+ slight improvement	+ 18.5% more accurate remuneration for the average artist	0
Market size (§§ 7.3.4–7.3.6, 7.5.1)	0	+ 20 %	- 8–34 % (all works)[514] up to +50% (marginal works)	- 30–90%[515]
Contests over cultural space (§ 7.5.5)	0	0	+? (only with voting)	+?
Non-payment signalling (§ 7.5.3)	0	0	- marginal	0
Signal menu effects (§ 7.5.4)	0	?	?	+?
Transaction costs (Chapter 8)	0	?	+?	+?
Infrastructure, enforcement + security (Chapter 9)	0	-$20 or more per device	?	+?
Taxation overheads (Chapter 10)	0	0	-3–23%	0

Table 11.1: Summary of the normative conclusions of Part IV relative to the *status quo* in the music industry

Note also that percentage adjustments cannot be added linearly; they need to be multiplied out. For instance, a 98% increase followed by a 90% reduction is a net 81% reduction.

Some of the considerations in Part IV turned out to be much weightier than others. The largest differences were observed in Chapter 6 (against the two DRM regimes and in favour of the two alternatives, and in favour of weak DRM over strong DRM); in Chapter 7 (against anarchy and in favour of the other regimes, and in favour of strong DRM over weak DRM); and in Chapter 9 (against feudalism and in favour of the other regimes). In dollar value terms, those differences in an economy like Australia's are in the billions per year. Their sizes are comparable to the total revenues of copyright-based entertainment industries, and are therefore the principal determinants of copyright regime efficiency.

A second tier of comparative costs was observed, whose sized are an order of magnitude below the first set. Within that band, the excess burden of taxation (Chapter 10) is the largest; anarchy finds some computer security savings over weak DRM or public funding. Also, the enforcement costs of the pragmatic "war on copying" (Section 9.3).

Some other effects were more difficult to quantify. For instance, the transaction cost advantages of the non-DRM regimes discussed in Chapter 8 might well be in the same range (hundreds of millions of dollars per year in Australia and tens of billions globally), but this remains speculative.

11.2   The Case for Alternative Compensation Systems

When the factors in favour and against the various regimes are tallied up, virtual market alternative compensation systems come out significantly ahead. Compared to the status quo, virtual markets eliminate artificial scarcity, slightly improve the quality of incentives, and reduce the burden of transaction costs. The downside is the excess burden of taxation, which is smaller than the artificial scarcity benefits alone. The advantage over the other two regimes is similar or larger.

How should this result be read? A study — even a fairly lengthy and exhaustive one — does not provide certainty that this is the best way forward. A fairer way to

---

[505]The 8–34% range is not uniform. As discussed in footnote 390, the expectation value is towards the lower end of this range.

[515]This is not an estimate based on data, but instead a statement of the broadest conceivable bounds. It corresponds to a value of k between 0.5 and 9 in the model in Section 7.3.5.

view the result is as a conditional claim: *if* the way the question was structured in this thesis, and the assumptions that were made in answering it, are correct *then* we have strong reason to believe that alternative compensation systems would be a much better way to pay our authors and artists.

This structure tells us not simply to believe that public funding could work better, but what that belief should depend upon. For example, if the tax revenue pool used to fund the virtual market was completely arbitrary, Liebowitz's arguments (see Section 7.5.1) would be much stronger, and the case for alternative compensation systems would be weaker.

Similarly, it might seem that the case for virtual markets was also fragile with respect to the deployment of a satisfactory computer security architecture, although it turns out that same security obstacles exist in DRM copyright regimes that allow unlimited "subscription" downloads or any other form of all-you-can-download pricing (Section 5.2.2).

11.2.1   A minimum standard for alternative compensation systems

An understanding of the factors that make virtual markets efficient remuneration systems also enables us to vet real-world ACS policies. Because alternative compensation systems are debated in various countries, and may perhaps be implemented one day, it is off interest to extract a checklist of preconditions that such proposals should satisfy:

- the scheme is accompanied by a blanket license which allows, at least, the reproduction of literary and musical copyright works in digital form, for non-commercial purposes;[516]

- works restricted by DRM or "technical protection measures" are ineligible for funding from the scheme;[517]

- access to the scheme is non-discriminatory, so that authors publishing on the web, as well as independent musicians and film makers, have a fair opportunity to compete for funds against larger commercial entertainment firms;[518]

---

[516]See Section 1.4.2.

[517]It is desirable to discourage the kinds of DRM-related expenses discussed in Chapter 9. Forcing rights holders to choose between ACS remuneration and DRM creates the kind of regime discussed by Shavell and van Ypersele and Litman (see footnote 334).

[518]As observed in Chapter 3, the web currently operates as an info-anarchic environment and is funded

- the revenue collected under the scheme is distributed to authors, artists and copy-right owners on the basis of decentralised metrics such as time-normalised usage metering, usage sampling or user votes (download counts should be avoided);[519]

- the mechanisms for measuring these quantities (whether by sampling or universal measurement) are designed and implemented with thorough and transparent peer review from the computer security community, to ensure that they preserve users' privacy and are resistant to disproportionate manipulation by particular individuals or groups;[520]

- the total amount of revenue available under the system is subject to regular adjustment based on rigorous economic measures of demand for works funded by the system.[521]

There are other important variables to consider in the construction of publicly funded alternative compensation systems. Probably the most important of them is the impact of the scheme on the inequitable distribution of income within entertainment industries. Fixing the deep problems of digital copyright does not necessarily require dealing with those issues, but it might provide an opportunity to do so, with additional criteria like the following:

- the schemes give a fair minimum portion of their royalties to artists, authors, performers and composers, even if these individuals no longer hold the relevant rights to the works they have produced;

- the schemes allocate remuneration on a progressive basis, so that artists with relatively low incomes receive a higher return from increasing popularity than superstars do.

---

by anarchy-compatible revenue sources such as advertising, voluntary production, and donations. However, as discussed in Section 5.4.1, the nature of much web publishing would make it well-suited to alternative compensation systems.

[519]A decentralised metric is necessary for the efficiency results discussed in Section 7.4.2. The use of *ex ante* metrics like usage measurement and voting lead to the benefits discussed in Section 7.5.2. Voting has further benefits as discussed in Section 7.5.5

[520]The necessity of these properties is fundamental. Thorough and transparent review processes are necessary for achieving public confidence in the system's security — especially because the program's administrators would suffer fewer negative consequences from a breach than artists would. The greatest confidence can be achieved with open code, even though the practical implications of openness for security are complicated (Anderson 2002).

[521]See Section 7.5.1.

11.3   Strong vs weak DRM

Quite apart from any conclusions about alternative compensation systems, the thesis included a comparison between two distinct kinds of exclusive digital copyright, the strong and weak regimes.

Each of these modes of organisation has advantages over the other. Feudalism benefits from the possibility of lower enforcement costs against end users (Section 9.3), slight reductions in transaction costs (Section 8.4) and improvements in consumer informedness (Section 7.5.2).

Most significantly, feudalism could be expected to produce a slight increase in investment — in the order of 20%.[522] The social surplus from this increased cultural production (the benefits minus the opportunity costs of that investment being made on other things) is likely to be somewhat lower.

The pragmatic status quo can counter these effects with its own advantages. First and foremost is the saving on strong DRM's elaborate security apparatus. That saving alone is gigantic (see Section 9.2). Secondly the status quo is made more desirable by the fact that piracy is presently serving a valuable purpose in alleviating artificial scarcity (see Section 6.3.1). Again, the effect is very large, in the range of 18–32%.

In the end, the weak copyright regime is clearly preferable to the strong one.

What this result indicates is that the tradeoff between the benefits of increased DRM (higher investment in cultural production) and the costs (less distribution of works, higher expenditure on the DRM itself) is optimised at a point which is much closer to the status quo than to feudalism.

What the result does not answer is where the "optimal" point of regime strength actually lies: if we are going to have digital copyright, should it be as it currently stands, or weaker, or stronger?

My methodology was not constructed with the intent of answer questions of that type (see Sections 2.3 & 3.1.1). The search for an accurate continuous parameterisation of DRM and enforcement strength is a matter for future research. It might build on work like that of Schechter *et al.*(2003) or Duchêne & Waldbroek (2003) but it must at least take into account the variables summarised here.

---

[522]See section 7.3.6.

11.4   How does Anarchy fare?

Of the four policy regimes I have examined, anarchy is the most distinct from the others, because it lacks a strong connection between consumers' taste for digital works and producers' incentives to provide them. In most regards, anarchy is at least as good as, and often superior to, the other regimes. After all, regulation is costly!

Where anarchy fails spectacularly of course is in the incentives it offers to artists and publishers, and the lower levels of production that result. Despite the presence of non-copyright motivation — monetary and otherwise — the loss of production makes anarchy unacceptable in most markets on conventional utilitarian accounts of the worth of cultural production.[523] There was only a distant possibility that anarchy approaches the status quo in desirability in some music markets, and basically no possibility in the market for films. Non-utilitarian accounts may be even more critical of information anarchy because is viewed as unfair even if it did not greatly reduce the supply of cultural works.

The Street Performer Protocol was a serious but mostly unsuccessful attempt to devise a way around that barrier. A definite negative conclusion on its ability to deliver much more from mass audiences than a good old fashioned fundraising drive should redirect the attention of those interested in info-anarchistic business models.

11.5   Harmonisation, Regulatory Biodiversity and the Future of Copyright

"Having been involved in that 'conspiracy' [TRIPS], *I would argue that it was not a conspiracy. It was something a little different."* — Jacques J. Gorlin (2002)

The clear story of copyright — and indeed all of the "intellectual property" laws — since the 1980s has been one of worldwide harmonisation around specific and strict regulation, enacted through treaties like TRIPS and WCT, and through bilateral trade agreements. That trend was the fruit of long-term planning by the industries benefiting from these laws (Braithwaite and Drahos 2000, 2002), and of their regulatory entrepreneurship supporting the entire *idea* of intellectual property. This corporate

---

[523]See Chapter 7 and especially Section 7.3.4.

activism gradually imbued the idea with tremendous momentum of its own. The IP movement was motivated by different factors: the problem of preventing international free riding on cultural and technological investment, the desire by corporate actors to strengthen established business positions, and a visionary sense by a handful of intellectual property lawyers that they could globalise their area of law.

The strength of the case for alternative compensation systems, as a potentially much more efficient and sustainable way to fund copyright industries, raises an interesting possibility about the entire project of global harmonisation. Is it possible that entertainment and publishing executives — in their pursuit of strong copyright rules, strong treaties to globalise those rules, and strong enforcement — might have made a mistake of a counter-intuitive sort. They might have prevented experimentation with, and stunted the development of, alternative regulatory and business models, at the precise moment in history when that was most necessary for the health of their own industries.

Understandably, the consensus view among copyright holders in the 1980s and 1990s was that the future prosperity of their businesses depended on maximising the reach of copyright into cyberspace and into countries outside of the developed world. But what we have subsequently learned is that the vision of perfectly and globally enforced copyrights is a very difficult and expensive one, if it is achievable at all. In the long run, more probably, major copyright holders will have to fight to sustain weak copyright regimes, and outbreaks of information anarchy will be an ongoing threat to their businesses.

In this era of persistent uncertainty about the sustainability of copyright businesses, and persistent incompatibility between traditional regulatory models and disruptive technology, the most prudent thing these industries could possibly have done would have been to run as many regulatory experiments as possible. Try strict enforcement in some countries, take a more liberal approach in others, and try different kinds of collective licensing and alternative compensation systems in at least a few places. Using different countries as laboratories, the content industries could have been much more informed about the practicality, financial viability, and other consequences of the different options. But to a community that had spent the preceding two decades fighting for the expansion and enforcement of a particular kind of copyright law, that path, and the possibility that they might not have known what was best for their own industries, was inconceivable.

At this juncture, eleven years after the appearance of Napster, it remains the case that *nobody* knows for sure what would be best for each of the copyright industries. We are stuck with a global regulatory monoculture that keeps us ignorant about the best available policy options. The problem is not just that the three-step test in TRIPS Article 13 discourages countries from implementing alternative compensation systems,[524] but that the culture of copyright law and policy discourages experiments of all sorts. The failure of the Google Books settlement is a particularly poignant example.

Had we had regulatory biodiversity to learn from, there would be more data about what works, and much more hope of evolving more sustainable ways of paying artists and publishers, and minimising the wasteful dimensions of copyright law. We may still need that data, but we will not be able to begin collecting it until ten or twenty years later than necessary.

11.6   Final Words

What to do about digital copyright can, without exaggeration, be characterised as the most important policy question arising from the advent of networked digital computers. Lawmakers have set out to strengthen copyright while nature conspires to do away with it, and the shape and size of several fields of human endeavour hang in the balance.

Our predicament highlights a subtle but profound limitation in the way that law and policy respond to technological development. Especially in English-speaking countries, none of the alternatives to the feudalist model have received any serious political consideration. Whether or not the "virtual market" is the right way forward, whether or not we can live with the status quo, and whether or not anarchy should even be considered as a policy option, it seems that wise societies should be searching for ways to ensure that technology grants their citizens the greatest possible access to art, knowledge and learning, greater opportunities for creativity and collaboration, and greater sovereignty over the world around them. This, notion of an informed choice is, perhaps, an ideal of "information democracy" that lies between the extremes of

---

[524]Although the Berne three-step test in TRIPS might complicate the introduction of virtual markets or other ACS, it could certainly be worked around, particularly if enough rights holder groups were supportive. One avenue would be to reach a sufficiently flexible interpretation of the test (Eckersley 2004b, IV.A.2). Another would be to fund an alternative compensation system, and pay the entire pool to whichever rights holders opted to participate — it is likely that many of them would quickly opt to do so.

"information anarchism" and "information feudalism".

The deepest challenge is that of inertia. The metaphor of physical property, in its application to books, sheet music, vinyl records and video tapes, may have done past societies more good than harm. It would be an irresponsibility of the highest order if that increasingly stretched and ill-fitting metaphor were allowed to preempt proper experimentation to determine how societies should govern the relationships between authors, artists, and publishers, and the rest of their citizens.

Readers with the professional habit of reasoning in the logic of copyright law will no doubt perceive some of the notions proposed in this thesis as radical departures from established institutions. This perspective, however, is inherently susceptible to underestimation of the degree to which "digital rights management" or vigorous enforcement of copyright statues on ordinary people is in turn a radical imposition on the natural logic of the computer and the natural uses that humans would like to make of them. In this matter, it seems, we are all radicals. Given the weight and novelty of the issues at work, we should choose our stripes carefully.

Bibliography

AACS Licensing Authority (2006, Feb. 17). Advanced Access Content System (AACS) Introduction and Common Cryptographic Elements. Revision 0.91, `http://www.aacsla.com/specifications/specs091/AACS_Spec_Common_0.91.pdf`.

Abramowicz, M. (2003). Perfecting Patent Prizes. *Vanderbilt Law Review 56*, 115–236.

ACC (Australian Copyright Council) (2001). Remuneration for private copying in Australia: A Discussion Paper. `http://www.copyright.org.au/PDF/Articles/PrivCopDiscPprAV.pdf`.

Adida, B. and R. L. Rivest (2006). Scratch & Vote: Self-Contained Paper-based Cryptographic Voting. In *WPES*, pp. 29–40. `http://ben.adida.net/research/AdidaRivest-scratch-and-vote.pdf`.

Adler, M. (1985). Stardom and Talent. *American Economic Review 75*(1), 208–212.

Aigrain, P. (2008). *Internet & Création*. InLibroVeritas. `http://paigrain.debatpublic.net/?page_id=171`.

Allen, D. W. (1999). Transaction Costs. In B. Bouckaert and G. D. Geest (Eds.), *Encyclopedia of Law and Economics*, Chapter 0740, pp. 893–926. Edward Elgar & The University of Ghent. `http://encyclo.findlaw.com/0740book.pdf`.

Anderson, R. (2002). Security in Open versus Closed Systems – The Dance of Boltzmann, Coase and Moore. In *Proc. Open Source Software: Economics, Law and Policy*. `http://www.ftp.cl.cam.ac.uk/ftp/users/rja14/toulouse.pdf`.

Anderson, R. (2003). 'Trusted Computing' Frequently Asked Questions. Version 1.1, `http://www.cl.cam.ac.uk/users/rja14/tcpa-faq.html`.

Anderson, R. and M. Kuhn (1996). Tamper Resistance - a Cautionary Note. In *Proc. 2nd 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`.

AP (Associated Press) (2005). Google Takes On Copyright Laws. `http://www.wired.com/news/business/0,1367,68901,00.html`.

ARIA (2004). Australian Record Sales – 2003 Full Year Figures. Australian Record Industry Association. `http://www.aria.com.au/news/stats2000.htm`.

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.

Australian House of Representatives Standing Committee on Legal and Constitutional Affairs (2006). Review of technological protection measures exceptions. `http://www.aph.gov.au/house/committee/laca/protection/report/fullreport.pdf`.

Babbitt, V. L., S. H. G. Roberts, and N. McClure (2006, Dec. 9). Secure internet voting system with bootable disk. United States Patent Number #7152156.

Baker, D. (2003, Nov.). The Artistic Freedom Voucher: Internet Age Alternative to Copyrights. Center for Economic and Policy Research. `http://www.cepr.net/publications/ip_2003_11.htm`.

Barlow, J. P. (1994, March). The Economy of Ideas : A framework for rethinking patents and copyrights in the digital age. *Wired Magazine 2.03*. `http://www.wired.com/wired/archive/2.03/economy.ideas_pr.html`.

Barlow, J. P. (2000, Oct.). The Next Economy of Ideas. *Wired Magazine 8.10*. `http://www.wired.com/wired/archive/8.10/download.html`.

Baxter, J. W. (1973). *World Patent Law and Practice* (2nd ed.). Sweet & Maxwell.

Baylor, M. (2005). Ranking Tax Distortions in Dynamic General Equilibrium Models: A Survey. `http://www.fin.gc.ca/wp/2005-06-eng.asp`.

Bellamy, E. (1888). *Looking Backward 2000–1887*. Ticknor & Co.

Belson, K. (2006, Jun. 13). As DVD Sales Slow, Hollywood Hunts for a New Cash Cow. *NewYork Times*. `https://www.nytimes.com/2006/06/13/technology/13disc.html`.

Benaloh, J., T. Moran, L. Naish, K. Ramchen, and V. Teague (2009, December). Shuffle-Sum: Coercion-Resistant Verifiable Tallying for STV Voting. *IEEE Transactions on Information Forensics and Security 4*(4).

Benkler, Y. (2003). Coase's Penguin, or Linux and the Nature of the Firm. *Yale Law Journal 112*(3), 369–446. `http://www.benkler.org/CoasesPenguin.PDF`.

Benkler, Y. (2006). *The Wealth of Networks*. Yale Press. `http://www.benkler.org/Benkler_Wealth_Of_Networks.pdf`.

Benzon, W. L. (2001, Aug.). Review of Geoffrey Miller's "Evolution of Human Music through Sexual Selection". Artwithbraininmind mailing list posting. `http://web.archive.org/web/20070806234642/http://kh.bu.edu/artwithbraininmind-l/1477.html`.

Bergemann, D., T. E. J. Feigenbaum, and S. Shenker (2005). Flexibility as an Instrument in Digital Rights Management. In *Proc. Workshop on Economics and Information Security*.

Bessho, S. and M. Hayashi (2005, Aug. 17). The Social Cost of Public Funds: The Case of Japanese Progressive Income Taxation. Technical report, Policy Research Department, Ministry of Finance, Japan. `http://www.mof.go.jp/jouhou/soken/kenkyu/ron123.pdf`.

Bethencourt, J., J. Franklin, and M. Vernon (2005). Mapping Internet Sensors With Probe Response Attacks. In *Proc. of the 14th USENIX Security Symposium*, pp. 192–208. `http://www.usenix.org/evens/sec05/tech/full_papers/shinoda/shinoda.pdf`.

Biddle, P., P. England, M. Peinado, and B. Willman (2002). The Darknet and The Future of Content Distribution. In *Proc. ACM Conference on DRM*. `http://crypto.stanford.edu/DRM2002/darknet5.doc`.

Boadway, R. and M. Keen (1993). Public Goods, Self-Selection and Optimal Income Taxation. *International Economic Review 34*(3), 463–478.

Boehlert, E. (2001, Mar.). Pay for Play. *Salon*. `http://archive.salon.com/ent/feature/2001/03/14/payola/print.html`.

Boguslavsky, M. M. (1979). *Copyright in International Relations: International Protection of Literary and Scientific Works*. Australian Copyright Council Ltd. Translated from the 1973 Russian edition by N. Poulet.

Bohm, P. (1972). Estimating Demand for Public Goods: An Experiment. *European Economic Review 3*, 111–130.

Boneh, D. and J. Shaw (1998). Collusion-Secure Fingerprinting for Digital Data. *IEEE Transactions on Information Theory 44*(5), 1897–1905.

Boorstin, E. (2004). Music Sales in the Age of File Sharing. Senior Thesis, Princeton Univ. `http://www.princeton.edu/~eboorsti/thesis/Music Sales in the Age of File Sharing.pdf`.

Borland, J. (2005, Apr. 15). Music moguls trumped by Steve Jobs? CNet News. `http://news.cnet.com/Music-moguls-trumped-by-Steve-Jobs/2100-1027_3-5671705.html`.

Bowen, H. (1943). The Interpretation of Voting in the Allocation of Economic Resources. *Quarterly Journal of Economics 58*, 27–48.

Boyle, J. (1997). A Politics of Intellectual Property: Environmentalism for the Net? *Duke Law Journal 47*, 87–116. `http://www.law.duke.edu/journals/dlj/articles/dlj47p87.htm`.

Boyle, J. (2000). Cruel, Mean or Lavish? Economic Analysis, Price Discrimination and Digital Intellectual Property. *Vanderbilt Law Review 53*(6), 2007–2039. `http://www.vanderbilt.edu/Law/lawreview/vol536/boyle.pdf`.

Bracha, O. (2007). Standing Copyright Law on its Head? The Googlization of Everything and the Many Faces of Property. *Texas Law Review 85*, 1799. `http://papers.ssrn.com/sol3/papers.cfm?abstract_id=931426`.

Braithwaite, J. and P. Drahos (2000). *Global Business Regulation*. Cambridge Univ. Press.

Brennan, D. J. (2002). Fair price and public goods: a theory of value applied to retransmission. *International Review of Law and Economics 22*, 347–375.

Brennan, G. and A. Hamlin (2000). *Democratic Devices and Desires*. Theories of Institutional Design. Cambridge Univ. Press.

Brennan, T. J. (2004). "Fair use" as policy instrument. Society for Economic Research on Copyright Issues. `http://www.serci.org/2003/brennan.pdf`.

Breyer, S. (1970). The Uneasy Case for Copyright: A Study of Copyright in Books, Photocopies and Computer Programs. *Harvard Law Review 84*(2), 281–351.

Brin, S. and L. Page (1998). The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems 30*(1–7), 107–117. `http://www-db.stanford.edu/pub/papers/google.pdf`.

Brodkin, J. (2011, Oct. 23). Microsoft collects license fees on 50% of Android devices, tells Google to "wake up". `http://arstechnica.com/microsoft/news/2011/10/microsoft-collects-license-fees-on-50-of-android-devices-tells-google-to-wake-up.ars`.

Brooks, D. (2003, Jan. 12). The Triumph of Hope Over Self-Interest. *NewYork Times*. `http://www.nytimes.com/2003/01/12/opinion/the-triumph-of-hope-over-self-interest.html`.

Brown, I. (2006, Jan. 30). DRM-a-go-go. Blogzilla blog post. `http://dooooooom.blogspot.com/2006/01/drm-go-go.html`.

Browning, E. K. (1994). The non-tax wedge. *Journal of Public Economics 53*, 419–433.

Browning, E. K., T. Gronberg, and L. Liu (2000). Alternative Measures of the Marginal Cost of Funds. *Economic Inquiry 38*(4), 591–599.

Browning, E. K. and L. Liu (1998). The optimal supply of public goods and the distortionary cost of taxation: comment. *National Tax Journal 51*(1), 103–116.

Bucovetsky, S. (1991). Choosing Tax Rates and Expenditure Levels Using Majority Rule. *Journal of Public Economics 46*, 113–131.

Burk, D. L. and J. Cohen (2001). Fair Use Infrastructure for Copyright Management Systems. *Harvard Journal of Law and Technology 15*, 41–83.

Bushing, Marcan, and Sven (2010). Console Hacking 2010: PS3 Epic Fail. Presentation at 27th Chaos Communications Congress. `http://www.youtube.com/watch?v=LuIlbmn-4A4`.

Caballero, J., C. Grier, C. Kriebich, and V. Paxson (2011, August). Measuring Pay-per-Install: The Commoditization of Malware Distribution. In *Proc. of the 20th USENIX Security Symposium*. `http://www.usenix.org/event/sec11/tech/full_papers/Caballero.pdf`.

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.

Caldwell, B. (1997). Hayek and Socialism. *Journal of Economic Literature 35*, 1856–1890.

Camelio, B. (2003, Mar. 31). Methods and apparatuses for financing and marketing a creative work. United States Patent Number #7885887.

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

Cano, P., E. Batlle, T. Kalker, and J. Haitsma (2002). A review of algorithms for audio fingerprinting. In *International Workshop on Multimedia Signal Processing*. `http://www.iua.upf.es/mtg/publications/MMSP-2002-pcano.pdf`.

Cardie, C. (1997). Empirical Methods in Information Extraction. *AI Magazine 18*(4), 65–80. `http://www.aaai.org/Library/Magazine/Vol18/18-04/Papers/AIMag18-04-008.pdf`.

Carroll, M. (2007). Creative Commons as Conversational Copyright. In P. K. Yu (Ed.), *Intellectual Property and Information Wealth: Issues and Practices in the Digital Age*, Volume 1, pp. 445–461. Praeger. `http://papers.ssrn.com/sol3/papers.cfm?abstract_id=978813`.

Carson, R. T., N. E. Flores, and N. F. Meade (2001). Contingent Valuation: Controversies and Evidence. *Environmental and Resource Economics 19*(2), 173–210. `http://econ.ucsd.edu/~rcarson/cvconfinal.pdf`.

CGD (Center for Global Development) (2005). Making markets for vaccines: Ideas into action. `http://www.cgdev.org/section/initiatives/_archive/vaccinedevelopment`.

Chaum, D. (1981, February). Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM 4*(2).

Chaum, D. (1984). Blind Signature System. In D. Chaum (Ed.), *Proc. CRYPTO 83*, New York, pp. 153–153. Plenum Press.

Chaum, D., P. Ryan, and S. Schneider (2005). Untraceable electronic mail, return addresses, and digital pseudonyms. In *Proc. 10th European Symposium on Research in Computer Science*. `http://epubs.surrey.ac.uk/1905/1/fulltext.pdf`.

Chen, S.-L. (2011). Collaborative Authorship: From Folklore to the Wikiborg. *University of Illinois Journal of Law Technology and Policy 2011*, 131–168.

Christie, A. F. (2004, July). Private Copying Licence and Levy Schemes: Resolving the Paradox of Civilian and Common Law Approaches. IPRIA Occasional Paper 02-04. `http://www.ipria.org/publications/workingpapers/Occasional paper 2.04.pdf`.

CIA (2004). *U.S. Central Intelligence Agency World Factbook.* `http://www.cia.gov/cia/publications/factbook/`.

Clarke, I. (2001). Fairshare - Rewarding artists without Copyright. `http://freenetproject.org/index.php?page=fairshare`.

Clarke, I., O. Sandberg, B. Wiley, and T. Hong (2001). Freenet: A Distributed Anonymous Information Storage and Retrieval System. In H. Federrath (Ed.), *Proc. ICSI Workshop on Design Issues in Anonymity and Unobservability.* International Computer Science Institute (Springer). `http://freenetproject.org/index.php?page=icsi-revised`.

Clarke, R. (1999). Identified, Anonymous and Pseudonymous Transactions: The Spectrum of Choice. In *Proc. User Identification & Privacy Protection Conference.* `http://www.anu.edu.au/people/Roger.Clarke/DV/UIPP99.html`.

Clarke, R. (2007). Business Models to Support Content Commons. *SCRIPTed 4*(1), 59. `http://www.law.ed.ac.uk/ahrc/script-ed/vol4-1/clarke.asp`.

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

Coase, R. H. (1937). The Nature of the Firm. *Economica 4*, 386–405.

Coase, R. H. (1946). The Marginal Cost Controversy. *Economica 13*(51), 169–182.

Cohen, J. E. (1998). Lochner in Cyberspace: The New Economic Orthodoxy of "Rights Management". *Michigan Law Review 97*(2), 462–563. `http://www.law.georgetown.edu/faculty/jec/Lochner.pdf`.

Cohn, C. A. and J. M. Grimmelmann (2003). Seven Ways in Which Code Equals Law (And One in Which It Does Not). In *Code: The Language of Our Time*, pp. 20. `http://90.146.8.18/en/archiv_files/20031/FE_2003_Cohn_en.pdf`.

Congdon, W. J., J. R. Kling, and S. Mullainathan (2009). Behavioral Economics and Tax Policy. Technical Report 15328, National Bureau of Economic Research. `http://www.brookings.edu/~/media/Files/rc/papers/2009/0616_behavioral_economics_kling/0616_behavioral_economics_kling.pdf`.

Copyright Board of Canada (2003, Dec. 13). Tariff of Levies to Be Collected by CPCC in 2003 and 2004 on the Sale, in Canada, of Blank Audio Recording Media. Canada Gazette Part I (Supplement). `http://www.cb-cda.gc.ca/tariffs/certified/c13122003-b.pdf`.

Corey, J. (2004). Advanced Honey Pot Identification and Exploitation. *Phrack 11*(63), 79–108. http://www.phrack.org/unoffical/p63/p63-0x09.txt.

Cornes, R. and T. Sandler (1996). *The Theory of Externalities, Public Goods, and Club Goods* (2nd ed.). Cambridge Univ. Press.

Crampton, T. (2003, Feb. 21). Pop Stars Learn To Live With Pirates. *International Herald Tribune*.

Craver, S. A., M. Wu, B. Liu, A. Stubblefield, B. Swartzlander, D. S. Wallach, D. Dean, and E. W. Felten (2001). Reading Between the Lines: Lessons from the SDMI Challenge. In *Proc. of the 10th USENIX Security Symposium*.

Crawford, S. P. (2003). The Biology of the Broadcast Flag. *Hastings Communications and Entertainment Law Journal 25*, 603–652.

Creedy, J. (2003, Dec.). The Excess Burden of Taxation and Why it (Approximately) Quadruples When the Tax Rate Doubles. New Zealand Treasury Working Paper 03/29. http://www.treasury.govt.nz/publications/research-policy/wp/2003/03-29/twp03-29.pdf.

Croskey, P. (1993). Institutional Utilitarianism and Intellectual Property. *Chicago-Kent Law Review 68*, 631–657.

Cryptography Research (2007, Nov. 19). Macrovision to Acquire Blu-ray Disc Security Technology from Cryptography Research, Inc. Press Release. http://www.cryptography.com/newsevents/press_releases/2007/11/19/20071119-macrovision.html.

Dahlby, B. (2008). *The Marginal Cost of Public Funds: Theory and Applications*. MIT Press.

Darwin, C. (1882). *The Descent of Man, And Selection in Relation to Sex* (2nd ed.). London: John Murray. http://pages.britishlibrary.net/charles.darwin/texts/descent/descent_front.html.

Davies, G. (1994). *Copyright and the Public Interest*. Max Planck Institute for Foreign and International Patent, Copyright and Competition Law & VCH.

Davies, G. and M. E. Hung (1993). *Music and Video Private Copying: An International Survey of the Problem and the Law*. Sweet & Maxwell.

Davies, N. (1996). *Europe: A History*. Oxford University Press.

Davis, D. (2002, Jun. 23). Moby Attributes Sales Slide To The 'Pearl Jam Effect'. Yahoo! Launch News. http://launch.yahoo.com/read/news.asp?contentID=209387.

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

Dean, K. (2003, Oct. 11). Shift-Key Case Rouses DMCA Foes. *Wired*. http://www.wired.com/entertainment/music/news/2003/10/60780.

Deazley, R. (2003). Re-reading Donaldson (1774) in the Twenty-First Century and Why it Matters. *European Intellectual Property Review 26*(6), 279–279.

DeLong, J. B. and A. M. Froomkin (2000, Feb.). Speculative Microeconomics for Tomorrow's Economy. *First Monday 5*(2). `http://www.firstmonday.dk/issues/issue5_2/delong/`.

Detering, D. (2001). New Method of Lindahl-Pricing for Intellectual Property Goods. `http://medieninhalte.eventme.com/article.htm`.

Dingledine, R., N. Mathewson, and P. Syverson (2004). Tor: The Second Generation Onion Router. In *Proc. of the 13th USENIX Security Symposium*, pp. 303–320. `https://svn.torproject.org/svn/projects/design-paper/tor-design.pdf`.

Dornseif, M., T. Holz, and C. N. Klein (2005). NoSEBrEaK – Attacking Honeynets. In *Proc. 2005 IEEE Workshop on Information Assurance and Security*. `http://md.hudora.de/publications/2004-NoSEBrEaK.pdf`.

Drahos, P. (1995). Information Feudalism in the Information Society. *The Information Society*. `http://slash.autonomedia.org/print.pl?sid=02/07/30/2024235`.

Drahos, P. (1996). *A Philosophy of Intellectual Property*. Dartmouth Publishing Co.

Drahos, P. and J. Braithwaite (2002). *Information Feudalism: Who owns the knowledge economy?* Earthscan Publications.

Duchêne, A. and P. Waelbroeck (2003). Does Peer-to-Peer harm Copyright Owners? Protecting and Distributing Digital Products. Society for Economic Research on Copyright Issues. `http://www.serci.org/2003/duchenewaelbroeck.pdf`.

Duffy, G. (2005, Mar. 8). Hacking Google Print. `http://www.kuro5hin.org/story/2005/3/7/95844/59875`.

Duffy, J. F. (2004). The Marginal Cost Controversy in Intellectual Property. *University of Chicago Law Review 71*, 37–56.

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

Eckersley, P. (2003, Feb.). 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`.

Eckersley, P. (2004a). Calling a Spade a Spade: Winners and Losers under a Retrospective Copyright Term Extension. Submission to the Australian Senate Select Committee on the USFTA. `http://www.aph.gov.au/senate_freetrade/submissions/sub487.pdf`.

Eckersley, P. (2004b). Virtual Markets for Virtual Goods: The Mirror Image of Digital Copyright? *Harvard Journal of Law and Technology 18*(1), 85–166. `http://jolt.law.harvard.edu/articles/pdf/v18/18HarvJLTech085.pdf`.

EFF (2003). EFF Sample Notification of Copyright Infringement Claims. Electronic Frontier Foundation. `http://www.eff.org/IP/P2P/sample_riaa_letter.pdf`.

EFF (2005). RIAA v. The People: Two Years Later. Electronic Frontier Foundation. `http://www.eff.org/IP/P2P/RIAAatTWO_FINAL.pdf`.

Elkin-Korin, N. (2001). Let the Crawlers Crawl: On Virtual Gatekeepers and the Right to Exclude Indexing. *University of Dayton Law Review 26*, 179–209.

European Commission (2006). Fair Compensation for private copying: copyright levies reform. EC Proposal 2006/MARKT/008. `http://ec.europa.eu/internal_market/copyright/docs/levy_reform/roadmap_82_83_en.pdf`.

Farchy, J. (2004). Seeking alternative economic solutions for combating piracy. Society for Economic Research on Copyright Issues. `http://www.serci.org/2004/farchy.pdf`.

Farrell, J. (1995). Arguments for Weaker Intellectual Property Protection in Network Industries. *StandardView 3*(2), 46–49.

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

Fisher, W. W. (2000b). Theories of Intellectual Property. In S. Munzer (Ed.), *New Essays in the Legal and Political Theory of Property*. Cambridge Univ. Press. `http://www.law.harvard.edu/Academic_Affairs/coursepages/tfisher/iptheory.html`.

Fisher, W. W. (2003). An Alternative Remuneration System. Chapter 6 of PROMISES TO KEEP: TECHNOLOGY, LAW AND THE FUTURE OF ENTERTAINMENT (forthcoming). `http://tfisher.org/PTK.htm`.

Fisher, W. W. (2004). Promises to Keep: Technology, Law and the Future of Entertainment.

Foley, J. H. (2001). Enter the Library: Creating a Digital Lending Right. *Connecticut Journal of International Law 16*, 369.

Foster, R. (2002, Jun. 21). The fundraiser ends, and the next stage begins. KURO5HIN. `http://www.kuro5hin.org/story/2002/6/21/10533/6651`.

Friend, T. (2003, Oct. 20). Credit Grab: How many writers does it take to make a movie? *New Yorker*. `http://reworld.org/mirrors/friend03credit.pdf`.

Froomkin, A. M. (1995). Anonymity and Its Enmities. *Journal of Online Law*. `http://www.wm.edu/law/publications/jol/95_96/froomkin.html`.

Gaita, K. and A. F. Christie (2004). Principle or Compromise? Understanding the original thinking behind statutory licence and levy schemes for private copying. *Intellectual Property Quarterly 4*, 422–447. `http://www.law.unimelb.edu.au/ipria/publications/workingpapers/IPRIA WP 04.04.pdf`.

Geer, D. (2004, Feb.). Will New Standards Help Curb Spam? *Computer*, 14–16. `http://www.computer.org/computer/homepage/0204/TechNews/r2014.pdf`.

Gervais, D. J. (2004). The Price of Social Norms: Towards a Liability Regime for File-Sharing. *Journal of Intellectual Property Law 12*, 39–73.

Gibson, J. (2005). Once and Future Copyright. *Notre Dame Law Review 81*(1), 167–244.

Giles, C. L., K. Bollacker, and S. Lawrence (1998). CiteSeer: An Automatic Citation Indexing System. In I. Witten, R. Akscyn, and F. M. Shipman III (Eds.), *3rd ACM Conference on Digital Libraries*, pp. 89–98. ACM Press. `http://citeseer.nj.nec.com/giles98citeseer.html`.

Ginsburg, J. C. (2001). Copyright and Control over New Technologies of Dissemination. *Columbia Law Review 101*, 1613–1647.

Ginsburg, J. C. (2002). Can Copyright Become User-Friendly? Review of Jessica Litman, Digital Copyright. *Columbia-VLA Journal of Law and the Arts 25*(1), 71–89. `http://papers.ssrn.com/paper.taf?abstract_id=288240`.

Gobioff, H., S. Smith, J. D. Tygar, and B. Yee (1996). Smart Cards in Hostile Environments. In *Proc. 2nd USENIX Workshop on Electronic Commerce*, pp. 1–11. USENIX Association. `http://www-2.cs.cmu.edu/~hgobioff/papers/usenix_ecommerce_96.ps`.

Goldsmith, J. L. (1998). Against Cyberanarchy. *University of Chicago Law Review 65*, 1199–1250.

Goodenough, O. R. (2002). The Futre of Intellectual Property: Broadening the Sense of "Ought". *European Intellectual Property Review 24*(6), 291–293.

Gordon, W. J. (1982). Fair Use as Market Failure: A Structural and Economic Analysis of the Betamax Case and Its Predecessors. *Columbia Law Review 82*, 1600–1657.

Gordon, W. J. (1992). Asymmetric Market Failure and Prisoner's Dilemma in Intellectual Property. *University of Dayton Law Review 17*, 853–869.

Gordon, W. J. (2003). Excuse and Justification in the Law of Fair Use: Transaction Costs Have Always Been Only Part of the Story. *Journal of the Copyright Society of the U.S.A. 50*, 149–198.

Gorlin, J. J. (2002). Symposium speech: U.S. Industries, Trade Associations, and Intellectual Property Lawmaking. *Cardozo Journal of International and Comparative Law 10*, 5–11.

Graham, L. R. (1993). *Science in Russia and the Soviet Union: A Short History*. Cambridge Univ. Press.

Grampp, W. D. (1989). *Pricing the Priceless: Art, Artists, and Economics*. New York: Basic Books.

Gray, P. (2004, May 5). DRM 'will be cracked' says iTunes hacker. The Register. `http://www.theregister.co.uk/2004/05/05/drm_will_be_cracked/`.

Grimmelmann, J. (2005). Regulation by Software. *Yale Law Journal 114*, 1719–1758. `http://www.yalelawjournal.org/pdf/114-7/Grimmelmann.pdf`.

Gruhne, M. (2003). Robust Audio Identification for Commercial Applications. In *Proc. Virtual Goods 2003*. `http://VirtualGoods.tu-ilmenau.de/2003/Robust_Audio_Identification_for_Commercial_Applications.pdf`.

Guell, R. C. and M. Fischbaum (1995). Toward Allocative Efficiency in the Prescription Drug Industry. *Milbank Quarterly 73*, 213.

Gutmann, P. (2007). The Commercial Malware Industry. In *Proc. DEFCON 2007*. `https://www.cs.auckland.ac.nz/~pgut001/pubs/malware_biz.pdf`.

Haitsma, J. and T. Kalker (2002). A Highly Robust Audio Fingerprinting System. In *Proc. 3rd International Conference on Music Information Retrieval*. Institut de Recherche et Coordination Acoustique/Musique. `http://ismir2002.ismir.net/proceedings/02-FP04-2.pdf`.

Harrison, P. (2002). The Rational Street Performer Protocol. `http://yoyo.cc.monash.edu.au/~pfh/rational.html`.

Hayek, F. A. (1945). The Use of Knowledge in Society. *American Economic Review 35*(4), 519–530.

Heller, M. A. (1997). The Tragedy of the Anticommons: Property in Transition from Marx to Markets. *Harvard Law Review 111*, 621.

Heller, M. A. and R. S. Eiseneberg (1998). Can Patents Deter Innovation? The Anticommons in Biomedical Research. *Science 280*, 698–701. `http://www.sciencemag.org/content/280/5364/698.full`.

Helman, L. (2010). Pass the Hat: Voluntary Payment as a Complementary Model for Music Copyright. Technical report, Columbia Law School. `http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1664034`.

Hettich, W. and S. L. Winer (2005). Rules, Politics and the Normative Analysis of Taxation. In *Handbook of Public Finance*, pp. 109–137. Springer. preprint available at `http://ideas.repec.org/p/car/carecp/00-12.html`.

High Court of Australia (1993). Australian Tape Manufacturers Association Ltd and others v. the Commonwealth of Australia. 176 CLR 480 FC 93/004, `http://www.austlii.edu.au/au/cases/cth/high_ct/176clr480.html`.

Hill, B. M. (2005, Jun. 10). Problems and Strategies in Financing Voluntary Free Sofrware Projects. revision 0.2, `http://mako.cc/writing/funding_volunteers/`.

Hodgson, G. M. (1988). *Economics and Institutions*. Cambridge: Polity Press.

Holbrook, M. B. and R. M. Schindler (1989). Some Exploratory Findings on the Development of Music Tastes. *Journal of Consumer Research 16*(1), 119–124.

Holz, T. and F. Raynal (2005). Detecting Honeypots and other suspicious environments. In *Proc. 2005 IEEE Workshop on Information Assurance and Security*. `http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.82.262&rep=rep1&type=pdf`.

Huang, A. (2003). *Hacking the Xbox: An Introduction to Reverse Engineering* (unlimited ed.). No Starch Press.

Hubbard, T. and J. Love (2004). A New Trade Framework for Global Healthcare R&D. *PLOS Biology 2*(2). `http://biology.plosjournals.org/perlserv?request=get-document&doi=10.1371/journal.pbio.0020052`.

Hugenholtz, P. B., L. Guibault, and S. van Geffen (2003). *The Future of Levies in a Digital Environment: Final Report*. University of Amsterdam Institute for Information Law. `http://www.ivir.nl/publications/other/DRM&levies-report.pdf`.

Hurt, R. M. and R. M. Schuchman (1966). The Economic Rationale of Copyright. *American Economic Review 51*(1), 421–432. `http://www.compilerpress.atfreeweb.com/Anno Hurt & Schuchman Econ Rationale Copyright.htm`.

Hurwicz, L. (1972). On Informationally Decentralized Systems. In C. B. McGuire and R. Radner (Eds.), *Decision and Organization: A volume in Honour of Jacob Marschak*. North Holland.

IFPI (2003). Recording Industry World Sales. International Federation of the Phonographic Industry. `http://www.ifpi.org/content/library/worldsales2003.pdf`.

IFPI (2006). Recording Industry World Sales. International Federation of the Phonographic Industry. `http://www.ifpi.org/content/library/digital-music-report-2006.pdf`.

IIPA (International Intellectual Property Alliance) (2004). Copyright Industries in the U.S. Economy: The 2004 Report. `http://www.iipa.com/pdf/2004_SIWEK_FULL.pdf`.

IITF (Information Infrastructure Task Force) (1995). Intellectual Property and the National Information Infrastructure: The Report of the Working Group on Intellectual Property Rights. `http://www.uspto.gov/web/offices/com/doc/ipnii/ipnii.pdf`.

IPA (International Publishers Association) (1997). A comparative study on Public Lending Rights in the European Union. `http://hdl.handle.net/10.1014/4`.

Jaatun, M., J. Jensen, H. Vegge, F. M. Halvorsen, and R. Nergard (2009). Fools Download Where Angels Fear To Tread. *IEEE Security and Privacy 7*(2), 83–86.

Jakobsson, M., A. Juels, and R. L. Rivest (2002, August). Making mix nets robust for electronic voting by randomized partial checking. In *Proc. of the 11th USENIX Security Symposium*.

Janelle, R. (2007). A Writers Guide to Freelance Blogging. `http://freelanceswitch.com/start/a-writers-guide-to-freelance-blogging/`.

Johnson, J. P. (2002). Open Source Software: Private Provision of a Public Good. *Journal of Economics and Management Strategy 11*(4), 637–662.

Kabalina, V. and S. Clarke (2000). Innovation in Post-Soviet Industrial Enterprises. In *Proc. 6th World Conference of the ICCEES*. `http://www.warwick.ac.uk/fac/soc/complabstuds/russia/Innopap.doc`.

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

Kaplow, L. (1998). A note on the optimal supply of public goods and the distortionary cost of taxation. *National Tax Journal 51*(1), 117–125.

Kelsey, J. and B. Schneier (1999, Jun.). The Street Performer Protocol and Digital Copyright. *First Monday 4*(6). `http://www.firstmonday.dk/issues/issue4_6/kelsey/`.

King, S. (2000). The Plant: Getting a Little Goofy. `http://www.stephenking.com/sk_120400_2.html`.

King, S. (2001). The Plant Income/Expense Report Through 12/31/00. `http://www.stephenking.com/sk_120400_2.html`.

Kinsler, J. S. (2003). Circuit-Specific Application of the Internal Revenue Code: An Unconstitutional Tax. *Denver University Law Review 81*, 113–217.

Kiritsov, S. (2000). Can Millions of Internet Users Be Breaking the Law Every Day?: An Intellectual Property Analysis of Linking and Framing and the Need for Licensing. *Stanford Technology Law Review 1*, 1–9.

Kirsner, S. (2005, Jan. 10). Downhill Battle wages uphill fight vs. music industry. *Boston Globe*. `http://www.boston.com/business/globe/articles/2005/01/10/downhill_battle_wages_uphill_fight_vs_music_industry/?mode=PF`.

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`.

Klamer, A. (1996). The Value of Culture. In A. Klamer (Ed.), *The Value of Culture: On the relationship between Economics and the Arts*, Chapter 1. New York: Basic Books.

Kleven, H. J. and C. T. Kreiner (2006). The marginal cost of public funds: Hours of work versus labor force participation. *Journal of Public Economics 90*, 1955–1973.

Konow, J. (1996). A Positive Theory of Economic Fairness. *Journal of Economic Behavior and Organization 31*, 13–35.

Kremer, M. (1998). Patent buyouts: a mechanism for encouraging innovation. *Quarterly Journal of Economics 113*(4), 1137–1167.

Kretschmer, M. (2001). Feudalism, Retreat or Revolution? In *Proc. Symposium on A New Feudalism of Ideas*. Bournemouth University. `http://www.cippm.org.uk/pdfs/martin.pdf`.

Kretschmer, M., G. M. Klimis, and C. J. Choi (1999). Increasing Returns and Social Contagion in Cultural Industries. *British Journal of Management 10*(supplement 1), S61–S72. `http://www.cippm.org.uk/pdfs/BritishJournalManagement.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.

Lakhani, K. R. and R. G. Wolf (2003, September). Why Hackers Do What They Do: Undertanding Motivation Effort in Free/Open Source Software Projects. MIT Sloan School of Management Working Paper 4425-03. `http://freesoftware.mit.edu/papers/lakhaniwolf.pdf`.

Lange, O. (1967). The Computer and The Market. In Feinstein (Ed.), *Socialism, Capitalism and Economic Growth*. Cambridge University Press.

Lavoie, D. (1990, Jan.). Computation, Incentives and Discovery: The Cognitive Function of Markets in Market Socialism. *Annals of the American Academy of Political and Social Science 507*, 72–79.

Lee, S., N. Moisa, and M. Weiss (2003, 21 Mar.). Open Source as a Signalling Device: An Economic Analysis. Goethe University Frankfurt am Main Department of Finance Working Paper Series: Finance and Accounting number 102. `http://www.finance.uni-frankfurt.de/wp/275.pdf`.

Lerner, J. and J. Tirole (2000). The Simple Economics of Open Source. Technical Report 7600, National Bureau of Economic Research. `http://www.nber.org/papers/w7600`.

Lessig, L. (1999). The Law of the Horse: What Cyberlaw Might Teach. *Harvard Law Review 113*, 501–546.

Lessig, L. (2001). *The Future of Ideas*. Random House.

Lessig, L. (2008). *Remix: Making Art and Commerce Thive in the Hybrid Economy*. Bloomsbury. `http://ebooksgo.org/commerce/Remix-o.pdf`.

Levitsky, S. L. (1964). *Introduction to Soviet Copyright Law*. Sythoff-Leyden.

Lewis, S. R. (2003). How Much is Stronger DRM Worth? In *Proc. 2nd Annual Workshop on Economics and Information Security*. `http://www.cl.cam.ac.uk/users/srl32/eis1.pdf`.

Lichtman, D. G. (1997). Pricing Prozac: Why the Government Should Subsidize the Purchase of Patented Pharmaceuticals. *Harvard Journal of Law and Technology 11*, 123.

Liebowitz, S. and R. Watt (2006). How Best to Ensure Remuneration for Creators in the Market for Music? Copyright and Its Alternatives. *Journal of Economic Surveys 20*(4).

Liebowitz, S. J. (2003a). Alternative Copyright Systems: The Problems with a Compulsory License. `http://www.utdallas.edu/~liebowit/intprop/complpff.pdf`.

Liebowitz, S. J. (2003b). Will MP3 Downloads Eliminate the Record Industry? The Evidence So Far. In G. Libecap (Ed.), *Advances in the Study of Entrepreneurship, Innovation and Growth*.

Liebowitz, S. J. (2004). Pitfalls in Measuring the Impact of File Sharing. `http://www.utdallas.edu/~liebowit/intprop/pitfalls.pdf`.

Light, D. W. (2005). Making Practical Markets for Vaccines. *PLOS Medicine 2*(10). `http://medicine.plosjournals.org/perlserv?request=get-document&doi=10.1371/journal.pmed.0020271`.

Lin, A. R. (2005). Who Owns the Cow When We Give Away the Milk for Free? Fair Use and the Protection of Web-Posted Materials. *Buffalo Intellectual Property Law Journal 3*, 46–82.

Litman, J. (2001). *Digital Copyright*. Prometheus Books.

Litman, J. (2004). Sharing and Stealing:. *Hastings Communications & Entertainment Law Journal 27*, 1–50.

Loeber, D. A. (1984). "Socialist" Features of Soviet Copyright Law. *Columbia Journal of Transnational Law 23*, 297–313.

Longhurst, B. (1995). *Popular Music and Society*. Polity Press.

Lotspiech, J. (2007). Advanced Access Content System Tutorial. Version 0.85, requestable from, `http://www.lotspiech.com/AACS/`.

Love, C. (2000, Jun.). Courtney Love does the math. *Salon*. `http://dir.salon.com/tech/feature/2000/06/14/love/index.html`.

Love, J. (2002). Artists Want to be Paid: The Blur/Banff Proposal. `http://www.nsu.newschool.edu/blur/blur02/user_love.html`.

Love, J. (2003). From TRIPS to RIPS: A better Trade framework to support Innovation in Medical Technologies. In *Proc. Workshop on Economic issues related to access to HIV/AIDS care in developing countries*. `http://www.cptech.org/ip/health/rndtf/trips2rips.pdf`.

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

MacDonald, G. M. (1988). The Economics of Rising Stars. *American Economic Review 78*(1), 155–166.

Macfie, R. (1869). *Recent Discussions on The Abolition of Patents For Inventions in the United Kingdom, France, Germany and the Netherlands*. Longmans, Green, Reader and Dyer. With Suggestions As To International Arrangements Regarding Inventions and Copyright.

Machlup, F. and E. Penrose (1950). The Patent Controversy in the Nineteenth Century. *Journal of Economic History 10*(1), 1–29.

MacLeod, C. (1988). *Inventing the Industrial Revolution: The English Patent System, 1660-1800*. Cambridge Univ. Press.

Martens, C. (2005, Aug. 12). Google provides opt-out for publishers. *Computerworld*. `http://www.computerworld.com/s/article/print/103895/Google_provides_opt_out_for_publishers`.

Martin, B. (1995). Against Intellectual Property. *Philosophy and Social Action 21*(3), 7–22. `http://www.uow.edu.au/arts/sts/bmartin/pubs/95psa.html`.

Mas-Colell, A., M. D. Winston, and J. R. Green (1995). *Microeconomic theory*. Oxford Univ. Press.

Maskill, A. (2005, August). GfK Marketing Canon Digital Lifestyle Index. `http://www.canon.com.au/pdf/editorial/canon digital lifestyle index 1h2005.pdf`.

Maurer, S. M. (2005). Choosing the right incentive strategy for research and development in neglected diseases. *Bulletin of the World Health Organization 84*(5). `http://www.who.int/bulletin/volumes/84/5/376.pdf`.

McDaniel, K. (2007). Accounting for Taste: An Analysis of Tax-and-Reward Alternative Compensation Schemes. *Tulane Journal of Technology and Intellectual Property 9*, 235–312.

Megiddo, N. and X. Zhu (2000, Jun. 21). System, method and program product for software development. United States Patent Number #6658642.

Menezes, A. J., P. C. van Oorschot, and S. A. Vanstone (1996). *Handbook of Applied Cryptography*. CRC Press. `http://www.cacr.math.uwaterloo.ca/hac`.

Menger, P.-M. (1999). Artists' Labor Markets and Careers. *Annual Reviews of Sociology 25*, 541–574. `http://www.law.berkeley.edu/journals/btlj/articles/vol12/Merges/html/reader.html`.

Merges, R. P. (1997). The End of Friction? Property Rights and Contract in the "Newtonian" World of On-line Commerce. *Berkeley Technology Law*

Journal 12, 115. `http://www.law.berkeley.edu/journals/btlj/articles/vol12/Merges/html/reader.html`.

Merges, R. P. (2004, Jan. 15). Compulsory Licensing vs. the Three "Golden Oldies": Property Rights, Contracts and Markets. *Policy Analysis*.

Michel, N. J. (2006). The Impact of Digital File Sharing on the Music Industry. *Topics in Economic Analysis and Policy 6*(1), Art. 18. `http://www.bepress.com/bejeap/topics/vol6/iss1/art18/`.

Miller, G. F. (2000). Evolution of Human Music through Sexual Selection. In N. L. Wallin, B. Merker, and S. Brown (Eds.), *The Origins of Music*. Cambridge, Massachusetts: The MIT Press.

Mockus, A., R. T. Fielding, and J. D. Hebsleb (2002). Two case studies of open source software development: Apache and Mozilla. *ACM Transactions on Software Engineering and Methodology 11*(3). `http://mockus.org/papers/mozilla.pdf`.

Moglen, E. (1999, Aug.). Anarchism Triumphant: Free Software and the Death of Copyright. *First Monday 4*(8). `http://www.firstmonday.dk/issues/issue4_8/moglen/`.

MPAA (2009). Theatrical Market Statistics 2005–2009. `http://www.mpaa.org/Resources/091af5d6-faf7-4f58-9a8e-405466c1c5e5.pdf`.

Mustonen, M. (2003). Copyleft — the Economics of Linux and other open source software. *Information Economics and Policy 15*, 99–121.

Myerson, R. and M. Satterthwaite (1983). Efficient Mechanisms for Bilateral Trading. *Journal of Economic Theory 29*, 265–281.

Nadel, M. S. (2003, Jan.). 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/admin/pdffiles/Nadel.pdf`.

Nadel, M. S. (2004). How Current Copyright Law Discourages Creative Output: The Overlooked Impact of Marketing. *Berkeley Technology Law Journal 19*(2). `http://btlj.boalt.org/data/articles/19-2_spring-2004_4-nadel.pdf`.

Naor, D., M. Naor, and J. Lotspiech (2001). Revocation and Tracing Schemes for Stateless Receivers. In *Proc. CRYPTO 01*, LNCS. Springer Verlag. `http://www.wisdom.weizmann.ac.il/~naor/PAPERS/2nl.pdf`.

Nelson, P. (1970). Information and Consumer Behaviour. *Journal of Political Economy 78*(2), 311–329.

Netanel, N. W. (1996). Copyright and a Democratic Civil Society. *Yale Law Journal 106*, 283–387. `http://www.history.ox.ac.uk/ecohist/readings/ip/netanel.htm`.

Netanel, N. W. (2002). Impose a Noncommercial Use Levy to Allow Free Peer-to-Peer File Sharing. University of Texas Public Law Research Paper #44.

Netanel, N. W. (2003). Impose a Noncommercial Use Levy to Allow Free Peer-to-Peer File Sharing. *Harvard Journal of Law and Technology 17*, 1. `http://www.utexas.edu/law/faculty/nnetanel/null.pdf`.

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

Nimmer, D. (2000). A Riff on Fair Use in the Digital Millenium Copyright Act. *University of Pennsylvania Law Review 148*, 673–742.

Nisan, N. and A. Ronen (1999). Algorithmic Mechanism Design. In *Proc. 31st ACM Symposium on the Theory of Computing*, pp. 129–140. `citeseer.nj.nec.com/nisan99algorithmic.html`.

Norgaard, C. and S. J. Garcia (2004). The Ninth Circuit's Decisions in *Foad v. Musil Govan Azzalino* and *Gardner v. Nike, Inc.*: The Creation, Interpretation and Assignment of Copyright Licenses under State and Federal Law. *Southwestern University Law Review 33*, 347–378.

Oakland, W. H. (1987). Theory of Public Goods. In *Handbook of Public Economics*, Volume 2, pp. 485–535. North-Holland.

Oates, J. (2006, Feb. 8). France rules in favour of P2P. The Register. `http://www.theregister.co.uk/2006/02/08/france_legalises_p2p/`.

Oberholzer, F. and K. Strumpf (2004, Mar.). The Effect of File Sharing on Record Sales: An Empirical Analysis. Draft working paper. `http://www.unc.edu/~cigar/papers/FileSharing_March2004.pdf`.

Oksanen, V. and M. Välimäki (2005). Copyright Levies as an Alternative Compensation Method for Recording Artists and Technological Development. *Review of Economic Research On Copyright Issues 2*(2), 25–39.

Olsen, S. (2003, Nov. 5). Amazon jams print feature on book search. CNet News. `http://news.com.com/2100-1032_3-5102917.html`.

Olson, M. (1971). *The Logic of Collective Action: Public Goods and the Theory of Groups* (revised ed.). Harvard Univ. Press.

Oppliger, R. (2002). How to Address the Secure Platform Problem for Remote Internet Voting. In *Proc. 5th Conf. on Security in Information Systems (SIS 2002)*, pp. 153–73. `http://www.ifi.unizh.ch/~oppliger/Docs/sis_2002.pdf`.

Owens, R. (2001). Digital Rights Management (DRM): A Look Ahead. *International Intellectual Property Law and Policy 7*.

Paranaguá, P. (2010, Sep. 2). Brazils proposal for P2P file-sharing. `http://pedroparanagua.net/2010/09/02/brazils-proposal-on-monetizing-p2p/`.

Peitz, M. and P. Waelbroeck (2004). The Effect of Internet Piracy on Music Sales: Cross-Section Evidence. *Review of Economic Research On Copyright Issues 1*(2), 71–79.

Perens, B. (2001, Aug. 2). Dimitry Skylarov: Enemy or friend? *ZDNet.* `http://www.zdnet.com/news/dimitry-sklyarov-enemy-or-friend/116424`.

Perlmutter, S. (2001). Convergence and the Future of Copyright. *Columbia-VLA Journal of Law and the Arts 24*(2).

Peters, M. (2008, Sep.). The Importance of Orphan Works Legislation. `http://www.copyright.gov/orphan/`.

Peukert, A. (2005). A Bipolar Copyright System for the Networked Digital Environment. *Hastings Communications and Entertainment Law Journal 28*, 1–80.

Plant, A. (1934). The Economic Aspects of Copyright in Books. *Economica 1*(2), 167–195. `http://www.compilerpress.atfreeweb.com/Anno Plant Copyright.htm`.

Polanvyi, M. (1944). Patent Reform. *Review of Economic Studies 11*, 61–76.

Pope, J. (1993). The Compliance Costs of Taxation in Australia and Tax Simplification: The Issues. *Australian Journal of Management 18*(1), 69–90.

Post, D. G. (2002). Against "Against Cyberanarchy". *Berkeley Technology Law Journal 17*(1), 1. `http://www.law.berkeley.edu/journals/btlj/articles/vol17/Post.stripped.pdf`.

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

Pye, D. (2000). Content-based methods for the management of digital music. In *Proc. of IEEE International Conference on, Acoustics, Speech, and Signal Processing*, Volume 4, pp. 2437–2440. `http://www.icsi.berkeley.edu/~dpwe/research/etc/icassp2000/pdf/4001_611.PDF`.

Quantum Market Research (2003). Understanding CD Burning and Internet File Sharing and its Impact on the Australian Music Industry. Summary of Research Commissioned by ARIA. `http://www.aria.com.au/documents/AriaIllegalMusicResearchReport_Summary.pdf`.

Radev, D. R., W. Fan, H. Qi, H. Wu, and A. Grewal (2005). Probabilistic Question Answering on the Web. *Journal of the American Society for Information Science and Technology 56*(3). `http://tangra.si.umich.edu/~radev/papers/radev-jasist.pdf`.

Rainie, L. and M. Madden (2004, Apr.). PEW Internet Project and Comscore Media Metrix Data Memo. Pew Internet and American Life Project. `http://www.pewinternet.org/pdfs/PIP_Filesharing_April_04.pdf`.

Rainie, L. and M. Madden (2005, Mar.). PEW Internet Project and Comscore Media Metrix Data Memo. Pew Internet and American Life Project.

Rasch, C. (2001, Jun.). 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/`.

Rawls, J. (1971). *A Theory of Justice* (revised (1999) ed.). Belknap.

Raymond, E. (1997). The Cathedral and the Bazaar. *First Monday 3*(3). `http://www.firstmonday.org/issues/issue3_3/raymond/`.

Raymond, E. S. (2000). On Management and the Maginot Line. In *The Cathedral and the Bazaar*. O'Reilley & Associates. `http://www.tuxedo.org/~esr/writings/cathedral-bazaar/cathedral-bazaar/x340.html`.

Reichman, J. H. (2000). Of Green Tulips and Legal Kudzuu: Repackaging Rights in Subpatentable Innovation. *Vanderbilt Law Review 53*(6), 1742–98.

Rheingold, H. (2008). *The Virtual Community*. MIT Press. `http://www.rheingold.com/vc/book/`.

Rob, R. and J. Waldfogel (2004). Piracy on the High C's: Music Downloading, Sales Displacement and Social Welfare in a Sample of College Students. NBER Working Paper 10874. `http://www.nber.org/papers/w10874`.

Rob, R. and J. Waldfogel (2006, April). Piracy on the High C's: Music Downloading, Sales Displacement and Social Welfare in a Sample of College Students. *Journal of Law and Economics 49*, 29–62. `http://bpp.wharton.upenn.edu/waldfogj/jle_piracy.pdf`.

Rob, R. and J. Waldfogel (2007). Piracy on the Silver Screen. *Journal of Industrial Economics 55*(3), 379–395.

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

Rochelandet, F. and F. L. Guel (2005). P2P Music Sharing Networks: Why The Legal Fight Against Copiers May Be Inefficient. *Review of Economic Research On Copyright Issues 2*(2), 69–82.

Roemer, R. (2003). Locking Down Loose Bits: Trusted Computing, Digital Rights Management and the Fight for Copyright Control on Your Computer. *UCLA Journal of Law and Technology 2003*, 8. `http://www.lawtechjournal.com/articles/2003/08_040223_roemer.php`.

Romer, P. (2002). When Should We Use Intellectual Property Rights? *American Economic Review 92*(2), 213–216.

Rosen, H. (2006, Jun. 4). For the Record, for What It's Worth. *Huffingon Post*. `http://www.cs.mu.oz.au/~pde/mirrors/rosen06for-the-record.html`.

Rosen, S. (1981). The Economics of Superstars. *American Economic Review 71*(5), 845–858.

Rosenblatt, B. (1997). The Digital Object Identifier: solving the dilemma of copyright protection online. *Journal of Electronic Publishing 3*(2). `http://www.press.umich.edu/jep/03-02/doi.html`.

Rothman, D. H. (1992, Jul. 6). Information Access for All. *Computerworld*. `http://www.teleread.org/computerworld.htm`.

Rubin, A. D. (2001). Security considerations for remote electronic voting. preprint, `http://avirubin.com/e-voting.security.html`.

Rubin, A. D. (2002, Feb.). Security considerations for remote electronic voting. *Communications of the ACM 45*(12).

Ryan, P. and S. Schneider (2006). Prêt à Voter with re-encryption mixes. In *Proc. European Symposium on Research in Computer Security*. `http://www.pretavoter.com/publications/esorics06.pdf`.

Sag, M. (2006). Twelve Year-Olds, Grandmothers, And Other Good Targets For The Recording Industry's File Sharing Litigation. *Northwestern Journal of Technology and Intellectual Property 4*(2). `http://www.law.northwestern.edu/journals/njtip/v4/n2/1/Sag.pdf`.

Salganik, M. J., P. S. Dodds, and D. J. Watts (2006). Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market. *Science 311*, 854–856. `http://www.sciencemag.org/cgi/content/full/sci;311/5762/854/DC1`.

Samuelson, P. (1954). The Pure Theory of Public Expenditure. *The Review of Economics and Statistics 36*(4), 387–389.

Samuelson, P. (1984). CONTU revisited: The case against copyright protection for computer programs in machine-readable form. *Duke Law Journal 1984*, 663–769.

Samuelson, P., R. Davis, M. D. Kapor, and J. H. Reichman (1994). A Manifesto Concerning the Legal Protection of Computer Programs. *Columbia Law Review 94*, 2308.

Schaathun, H. G. (2003). The Boneh-Shaw Fingerprinting Algorithm is Better Than We Thought. *IEEE Transactions on Information Forensics and Security 1*(2).

Schechter, S. E., R. A. Greenstadt, and M. D. Smith (2003). Trusted Computing, Peer-To-Peer Distribution, and the Economics of Pirated Entertainment. In *Proc. 2nd Annual Workshop on Economics and Information Security*. `http://www.eecs.harvard.edu/~stuart/papers/eis03.pdf`.

Schneier, B. (2000). *Secrets & Lies: Digital Security in a Networked World*. John Wiley & Sons.

Schoen, S. (2003). Trusted Computing: Promise and Risk. Electronic Frontier Foundation. `http://www.eff.org/Infrastructure/trusted_computing/20031001_tc.php`.

Schulman, A. (2000). The Caldera v. Microsoft Dossier. O'Reilly Network. `http://www.oreillynet.com/pub/a/network/2000/02/07/schulman.html`.

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.

Seifreid, K. (2002). Honeypotting with VMware - basics. `http://www.seifried.org/security/ids/20020107-honeypot-vmware-basics.html`.

Sell, S. K. (2002). Industry Strategies for intellectual property and trade: the quest for TRIPS, and post-TRIPS strategies. *Cardozo Journal of International and Comparative Law 10*, 79–108.

Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal 27*, 379–423 and 623–656. `http://cm.bell-labs.com/cm/ms/what/shannonday/shannon1948.pdf`.

Shapiro, C. and H. Varian (1999). *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business School Press.

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.

Shavell, S. and T. van Ypersele (2001). Rewards versus Intellectual Property Rights. *Journal of Law and Economics 44*, 525.

Shinoda, Y. and M. Itoh (2005). Vulnerabilities of Passive Internet Threat Monitors. In *Proc. of the 14th USENIX Security Symposium*, pp. 209–224. `http://www.usenix.org/evens/sec05/tech/full_papers/shinoda/shinoda.pdf`.

Simmons, G. (2001). *Kiss and Make-Up*. Crown Publishers.

Slater, P. J. B. (2002). Fifty years of song bird research: a case study in animal behaviour. *Animal Behaviour 65*, 633–639.

Slemrod, J. and N. Sorum (1984). The Compliance Costs of the U.S. Individual Income Tax System. *National Tax Journal 37*(4), 461–474.

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`.

Spitzner, L. (2003). *Honeypots: Tracking Hackers*. Addison-Wesley.

Stadler, M., J.-M. Piveteau, and J. Camenisch (1995). Fair Blind Signatures. In *Proc. EUROCRYPT '95*, LNCS, pp. 209–219. Springer Verlag. `http://www.ubilab.org/publications/print_versions/pdf/sta95.pdf`.

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

Stallman, R. M. (1992a, Jul./Aug.). The Right Way to Tax DAT. *Wired Magazine 1.03*. `http://www.gnu.org/philosophy/dat.html`.

Stallman, R. M. (1992b). Why Software Should Be Free. Free Software Foundation. `http://www.gnu.org/philosophy/shouldbefree.html`.

Stallman, R. M. (1996). Reevaluating Copyright: The Public Must Prevail. *Oregon Law Review 75*(1). `http://www.gnu.org/philosophy/reevaluating-copyright.html`.

Stallman, R. M. (1997, Feb.). The Right to Read. *Communications of the ACM 40*(2), 85–87. `http://www.gnu.org/philosophy/right-to-read.html`.

Stallman, R. M. (1999). Free Software and Beyond. Presented at Wizards of OS: Operating Systems and Social Systems, Berlin. `http://mikro.org/Events/OS/ref-texte/stallman.html`.

Stefik, M. (1997). Shifting The Possible: How Trusted Systems and Digital Property Rights Challenge Us To Rethink Digital Publishing. *Berkeley Technology Law Journal 12*, 137. `http://www.law.berkeley.edu/journals/btlj/articles/vol12/Stefik/html/reader.html`.

Sterling, B. (1999). Maneki Neko. In *A Good Old-Fashioned Future*. New York: Bantam Spectra.

Stiglitz, J. E. (1994). *Whither Socialism*. MIT Press.

Strowel, A. (2009). *Peer-to-Peer File Sharing and Secondary Liability in Copyright Law*. Edward Elgar Pub. Co.

Sullivan, B. (2004, Jan. 9). Music Industry Puts Troops in the Streets. *L.A. Weekly*. `http://www.laweekly.com/2004-01-08/news/music-industry-puts-troops-in-the-streets/`.

Supreme Court of the United States (2005). MGM Studios Inc v. Grokster, Ltd. 545 U.S. 913, `http://w2.eff.org/IP/P2P/MGM_v_Grokster/04-480.pdf`.

Syverson, P. (1994). A taxonomy of replay attacks. In *Computer Security Foundations Workshop VII*. IEEE Computer Society Press. `citeseer.ist.psu.edu/syverson94taxonomy.html`.

Takeyama, L. N. (1997, October). The Intertemporal Consequences of Unauthorized Reproduction of Intellectual Property. *Journal of Law and Economics 40*(2), 511–22.

Takeyama, L. N. (2002, June). Piracy, Asymmetrical Information, and Product Quality Revelation. Society for Economic Research on Copyright Issues. `http://www.serci.org/2002/takeyama.pdf`, `http://www.serci.org/2002/takeyama.pdf`.

Taub, E. (2003, Jul. 20). DVD's Meant for Buying but Not for Keeping. *NewYork Times*. `http://query.nytimes.com/gst/fullpage.html?res=9F0CEFDC1F3CF932A15754C0A9659C8B63`.

TCG (Trusted Computing Group) (2005, May 3). TCG Trusted Network Connect: TNC Architecture for Interoperability. Specification version 1.0 revision 4, `https://www.trustedcomputinggroup.org/groups/network/TNC_Architecture_v1_0_r4.pdf`.

Temkin, G. (1996). Information and Motivation: Reflections on the Failure of the Socialist Economic System. *Communist and Post-Communist Studies 29*(1), 25–41.

Thompson, E. A. (1968). The Perfectly Competitive Production of Public Goods. *Review of Economics and Statistics 50*(1), 1–12.

Tirole, J. (1988). *The Theory of Industrial Organisatoin*.

Toigo, J. W. (2000, May). Avoiding a Data Crunch. *Scientific American 282*(5), 58–74.

Touretzky, D. (2000). Gallery of DeCSS Descramblers. `http://www.cs.cmu.edu/~dst/DeCSS/Gallery/`.

Towse, R. (1993). *Singers in the Marketplace: The Economics of the Singing Profession*. Oxford University Press.

Towse, R. (2001). *Creativity, Incentive and Reward: An Economic Analysis of Copyright and Culture in the Information Age*. Edward Elgar.

Triest, R. K. (1990). The relationship between the marginal cost of public funds and marginal excess burden. *American Economic Review 80*, 557–566.

Turner, S. (2006, Aug. 15). The State of Spyware According to Webroot. *ZDNet*. `http://www.zdnet.com/blog/spyware/the-state-of-spyware-according-to-webroot/848`.

U.K. House of Commons (2006, June). "Digital Rights Management" Report of an Inquiry by the All Party Internet Group. `http://www.apcomms.org.uk/apig/current-activities/apig-inquiry-into-digital-rights-management/DRMreport.pdf`.

Vaillancourt, F., J. Clemens, and M. Palacios (2008). Compliance and Administrative Costs of Taxation in Canada. In *The Impact and Cost of Taxation in Canada: The Case for a Flat Tax Reform*. The Fraser Institute. `http://aysps.gsu.edu/ComplianceAdministrativeCosts.pdf`.

Vance, A. (2003, Oct. 12). RIAA hires guns, alcohol and smokes expert to fight piracy. The Register. `http://www.theregister.co.uk/2003/12/10/riaa_hires_guns_alcohol/`.

Vise, D. A. and M. Malseed (2005). *The Google Story*. Delacorte Press.

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.

Waladan, S. (2005). Copyright Reform Won't Cover Innovation. Australian Digital Alliance. `http://www.digital.org.au/issue/documents/InternetBulletinArticle.doc`.

Waldfogel, J. (2010). Music file sharing and sales displacement in the iTunes era. *Information Economics and Policy 22*, 306–314.

Ward, S. and I. Richards (2004, Mar.). System and method for fingerprint based media recognition. United States Patent Application #20040249859.

Watt, R. (2004). The Past and Future of the Economics of Copyright. *Review of Economic Research On Copyright Issues 1*, 151–171.

Westfall, S. R. (2002). "The Useless Dearness of the Diamond": Theories of Patronage Theatre. In P. W. White and S. R. Westfall (Eds.), *Shakespeare and Theatrical Patronage in Early Modern England*, pp. 13–41. Cambridge University Press.

Woodford, C. (2004). Trusted Computing or Big Brother? Putting the Rights back in Digital Rights Management. *University of Colorado Law Review 75*, 254–300.

World E-Commerce & IP Report (2002). Education Programs Needed to Combat Consumer Piracy, Experts Say. *World E-Commerce & IP Report 2*(5), 25–26.

World Intellectual Property Organisation (WIPO) (1996). Copyright Treaty. `http://www.wipo.int/treaties/ip/copyright/copyright.html`.

Wright, B. D. (1983, September). The Economics of Invention Incentives: Patents, Prizes and Research Contracts. *American Economic Review 73*(4), 691–707.

Wu, T. (2003). When Code Isn't Law. *Virginia Law Review 89*, 103–170.

Xuxian Jiang, D. X. (2004). Collapsar: A VM-Based Architecture For Network Attack Detention Center. Technical report, Purdue University. CERIAS Technical Report 2004-16, `http://www.cs.purdue.edu/homes/dxu/pubs/Security04.pdf`.

Yen, A. C. (2002). Western Frontier or Feudal Society?: Metaphors and Perceptions in Cyberspace. *Berkeley Technology Law Journal 17*, 1208–1263. `http://www.law.berkeley.edu/journals/btlj/articles/vol17/Yen.stripped.pdf`.

Zentner, A. (2006). Measuring the Effect of Online Music Piracy on Music Sales. *Journal of Law and Economics 49*(1), 63–90.

Zimmerman, D. L. (2003). Authorship Without Ownership: Reconsidering Incentives in a Digital Age. *DePaul Law Review 52*, 1121–1169.