

# Information, Incentives, and the Economics of Security and Privacy

Rick Wash

September 15, 2005

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>The Economics of Security</b>	<b>2</b>
2.1	Digital Rights Management . . . . .	2
2.2	Vulnerabilities . . . . .	5
2.3	Spam . . . . .	7
2.4	Privacy . . . . .	8
2.5	Spending and Investment . . . . .	10
2.6	Security Breaches . . . . .	10
2.7	Global Insecurity . . . . .	11
2.8	Theory . . . . .	12
<b>3</b>	<b>Focusing on Content Provision and Digital Rights Management Technologies</b>	<b>14</b>
3.1	Model . . . . .	16
3.2	Results . . . . .	17
3.3	Various Effects of DRM . . . . .	19
3.4	Empirical Evaluation . . . . .	20
<b>4</b>	<b>Conclusions</b>	<b>20</b>
<b>A</b>	<b>Reading List</b>	<b>21</b>
A.1	Economics of Security . . . . .	21
A.2	Information Economics . . . . .	24
A.3	Other Referenced Works . . . . .	26

## 1 Introduction

Information Security is a major problem on the Internet today. The Computer Science community has spent a large amount of time and money identifying specific problems and devising numerous effective technical solutions to these problems. Unfortunately, security is still a large problem despite technologically effective solutions. This is because many of the security problems are both technical and social in nature.

Ross Anderson, in his important paper “Why Cryptosystems Fail” (Anderson, 1993) showed that in almost all of the cases when banks had security failures, effective technical solutions exist to prevent that failure. However, there was a largely-unnoticed social dimension to security that really caused the failure. This was either because the attacker was able to confuse and deceive

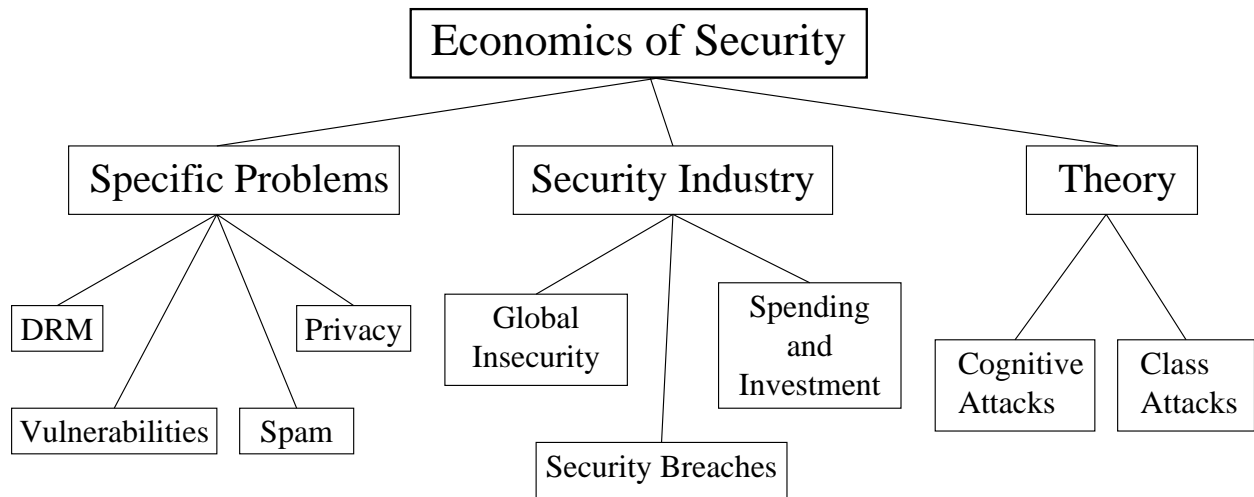


Figure 1: The Economics of Security

legitimate users, or (more commonly) the legitimate users chose not to use the technical security solutions.

Trying to understand and predict the behavior of users, which is necessary since all systems have users, is outside of the domain of traditional Computer Science. Fortunately, the field of Economics has developed a rich set of tools and techniques for doing this. Applying this toolbox to the problems of information security has yielded many interesting and useful results, but much work still needs to be done.

## 2 The Economics of Security

Concepts from Economics have proven very useful when applied to information security problems. People play a fundamental role in the use of all technologies, and understanding how and why people make decisions about information security technologies is essential to providing good security. Currently, the field of ‘Economics of Security’ is very young, with relatively few good papers. The primary workshop<sup>1</sup> in the field is now only four years old.

Most of the work in this area has been applying economic concepts to specific information security problems that have proven particularly difficult to solve with traditional computer science, such as Digital Rights Management, vulnerability management, spam email, and privacy. A second body of work in this field concerns the industry that has been built up around information security. This includes questions about how and where to spend money on information security, how to properly share information about security breaches, and the effects of insecurity on others in the industry. Finally, there is a body of work regarding general theories of information security and economics. My sub-categorization of this field is illustrated in Figure 1. This section deals with each of these areas in turn.

### 2.1 Digital Rights Management

Digital content, such as digital movies, music, and books, has been embraced by the public. Music can be downloaded from the Internet cheaply and easily shared with friends. However, this low-cost

<sup>1</sup>WEIS: Workshop on the Economics of Information Security

copying has proven itself a thorn in the side of the content industries as it allows people to easily copy their products from others without paying for them. Despite the fact that this is illegal under copyright law, it is impractical for the content industry to sue hundreds of thousands of teenagers to stop this practice.

Digital Rights Management (DRM) is a technological solution to this problem in copyright. However, experience has shown that effective (secure) DRM technologies are extremely difficult to achieve in practice. This has not prevented some companies (like Apple) from rolling out DRM solutions that they knew could be broken. These imperfect DRM systems increase the barriers to copyright infringement, hopefully just enough that most people would rather buy than infringe.

Politically, DRM is a hot topic. Many content industries claim that if Congress doesn't mandate some form of DRM, then they will not have enough profits to continue making new content. Consumer groups claim DRM is eroding their fair-use rights. DRM technology vendors are concerned with convincing both content producers and consumers to adopt their technologies. Overall, there are many interesting questions regarding these new technologies.

The overarching questions here are *How does the introduction of Digital Rights Management technologies change the landscape for digital content? How does the ability to circumvent (at a cost) DRM technologies factor into this?*

DRM technologies have proven difficult for computer scientists. Most computer scientists approached the problem by trying to provide 'security,' as if it were some binary state (its either secure or its not). Perfect security for DRM is extremely difficult, as it has competing goals: to provide access to digital content and to restrict access to that digital content. However, imperfect DRM is doable. Imperfect DRM provides access but makes certain types of access artificially more cumbersome than they would otherwise be.

The DRM industry has three classes of agents: Content Producers, Platform Providers, and Consumers. This is very similar to what economists are currently calling a "two sided market." (Rochet and Tirole, 2003) Content producers want to sell their content to consumers, but have the problem that consumers can then do many things with this content that the producers don't want them to, such as share it on the Internet with the consumer's closest million friends. Platform providers have the technology that enables consumers to use content, but only if producers have placed their content on the platform.

All digital content must exist on some type of platform, or combination of hardware and software that can be used to view/use/access the content. DRM platforms are content platforms that include some technology that attempts to artificially restrict the content's use beyond what would normally be implied by standard technology and legal controls. This technology usually involves encryption and asking permission of the content producer to use the content. DRM platform providers then have to convince both the content producers and the consumers simultaneously to use their platform to access content.

In this type of industry, there are many interesting questions that can be raised. What will the organization of the DRM platform industry look like? What strategies will businesses need to use to succeed in this industry? What changes will DRM technologies cause in the content industries, or for consumers? Is DRM technology socially desirable, and what kinds of policy levers for DRM are available that can affect this new technology?

### **2.1.1 Organization of the DRM Platform Industry**

One avenue of work regards competition in the DRM platform industry. Here, there are both strategic questions (What strategies can help businesses succeed?) and competition policy questions (What strategies should be regulated by anti-trust?) that have been asked.

First, I consider strategic questions for businesses in this industry. Acquisti (2004a) has a network effects model that looks at the incentives for content providers and consumers to adopt a given DRM platform, and gives some insight on strategies that DRM platform providers can use to gain leverage over open platforms. Sundararajan (2004) discusses general strategic pricing issues related to digital content in the face of piracy, and one line of concern includes pricing with DRM technologies.

Now, what questions does DRM raise for competition policy? Park and Scotchmer (2004) study various models of collusion among DRM providers and provide theoretical evidence that some methods of collusion might not be socially inefficient. However, they do not have any solid rules about what is inefficient and should be subject to anti-trust regulation, leaving the door open for further research in this area. Anderson (2003) attempts to understand some of the possible anti-trust effects that hardware security technologies developed for use in DRM may have. This technology can also strongly enable lock-in and many anti-competitive behaviors that must be carefully monitored. He mainly cites possible actions, but doesn't formally analyze what self-interested agents are likely to do, nor does he study possible interventions. These are open questions.

### **2.1.2 DRM and Content**

DRM technologies will cause changes in the digital content market. These changes can be in how the content is priced, what content is made available, or in what possible uses the content has. There are a few studies that attempt to address these changes and determine what implications these have for policy and strategy.

Park and Scotchmer (2004) and Acquisti (2004a) study the optimal pricing of digital content using DRM technologies, and have results that the price of content is likely to decrease relative to a perfect legal enforcement regime, because illegitimate copies of the content compete with legitimate content. It is still an open question if there are any strategic pricing concerns for DRM, and what effects these changes will have on the markets.

Acquisti was the first to my knowledge to comment on the types of content that are likely to exist on a DRM platform. Using a characterization from Mackie-Mason, Shenker, and Varian (1996), he suggests that high-value niche goods are more critical to the success of DRM platforms than low-value mass-market goods. I will attempt to expand on this insight and develop some policy questions in Section 3. A question along these same lines which hasn't been studied much yet is what effect DRM technologies will have on incentives for creation of new content aside from the obvious one of increasing the profitability of digital content.

Bergemann, Eisenbach, Feigenbaum, and Shenker (2005) study the options and possible changes that DRM enables in the uses of digital content. Since DRM can be used to limit certain uses but not others, content providers have a choice which uses to limit. This paper looks at flexibility very generically and studies what level of flexibility would content providers choose. It is an open and interesting question to determine which uses content providers will want to limit (what way will they operationalize the flexibility level they chose), and to see if these decisions are socially beneficial. Another interesting question is how the ability to restrict uses of digital content can change the incentives for innovation.

### **2.1.3 DRM Policy**

There are many interesting public policy issues related to Digital Rights Management. The first question concerns how desirable DRM technologies are in society, and society's need for these technologies. Elements of competition policy were address in Section 2.1.1. There are policy

questions around the changes in content from the previous subsection. Finally there are open questions as to mandating DRM technologies.

Oberholzer and Strumpf (2004) and Givon, Mahajan, and Muller (1995) both study the effects that the piracy of digital content has. Since DRM technologies were designed to prevent piracy, these studies directly apply to the social desirability of DRM technologies. Additionally, Samuelson and Scotchmer (2002) studies the social desirability of reverse engineering, which is another use that DRM technologies were designed to prevent. These study leave it unclear whether DRM technologies are needed or not.

The changes in content listed in the previous subsection leave an open question whether these changes are desirable for society. Will DRM technologies affect the development of new content adversely by reducing the pool of existing content to build from? Additionally, if certain types of DRM technologies are not socially desirable, what are some efficient methods of enforcement?

Finally, numerous people have called for mandating DRM technologies by law. One example is the recent attempts to pass a ‘broadcast flag’ law, which is a law that mandates support in broadcast television sets and recording equipment for a flag that can be set by the content providers instructing the devices to not record a given piece of content. Is such mandatory DRM technology socially beneficial? It is definitely more cost effective than forcing the content providers to develop effective anti-recording technology, but may not be socially desirable.

## 2.2 Vulnerabilities

Vulnerabilities are pieces of information that allow unauthorized hackers to access computers. More specifically, vulnerabilities are bugs in software that can be exploited by hackers to run arbitrary programs that the hackers write. Information about these vulnerabilities is valuable to attackers because it is necessary for hacking.

Information about these vulnerabilities is also valuable to defenders because it tells them about holes in their defenses that need fixing. This gives a number of good guys an incentive to publicize this information so “attackers aren’t the only ones who have it.” Information about vulnerabilities is also valuable in the market as it adds useful information about software quality to consumers purchase decisions, thereby putting pressure on software companies to improve their quality.

The overarching question for this area is *What value does information about vulnerabilities in software have?*

Understanding vulnerabilities has proven difficult for computer scientists. At their base, vulnerabilities are just software bugs, so much of the existing software engineering literature applies. However, they are special both because they are particularly bad bugs, and because they can provide specific extra value to other people, thereby incenting these people to discover and exploit them. Computer scientists are effective at discovering them, fixing them, and understanding just how potentially bad a problem vulnerabilities are (Beattie, Arnold, Cowan, Wagle, Wright, and Shostack, 2002). But, they are very ineffective at understanding the social effects of vulnerabilities, understanding the desirability of the good guys searching for them, changing the incentives for disclosure, or other such peripheral activities.

The research questions here center around three themes. The first concerns the public disclosure of vulnerability information. The second concerns incentives for discovering vulnerabilities. The last discusses the effects of vulnerability information on the software market.

### 2.2.1 Public Disclosure

Information about vulnerabilities has different value to different parties. Attackers love it when such information is made public, as it makes their jobs easier. Software companies hate it, as it exposes the flaws in their work. Security professionals are torn. On one hand, making such information public helps them to know what the problems in their systems are, and helps them to fix them. This information being made public also puts pressure on software companies to reduce security problems. On the other hand, this information also makes the hacker's job easier, leading to more security problems. This is particularly true when the information is made public before the software company (the vendor) can produce a fix.

Most of the research in this area addresses the question "Should someone who finds a vulnerability make it public, and if so how?" and the related question of "Should we actively look for such vulnerabilities?"

Ozment (2005) looks at the question of rediscovery. If the discoverer chooses not to make the vulnerability information public, what are the chances that it will be rediscovered and made public anyway? Ozment uses empirical data from an open source project to estimate the occurrence of rediscovery. This question could benefit from an economic model for better understanding.

Often, when people discover vulnerabilities, they notify the vendor and allow them to fix the problem before make the vulnerability information public. However, vendors have the incentive to slowly fix the problem so as to put off the public relations problem of having a known vulnerability. The community of people who look for vulnerabilities have responded to this by implementing 'disclosure policies' that specify maximum time lengths to wait for a vendor fix before public disclosure. Arora, Telang, and Xu (2004) asks what the optimal policy would look like considering the need for public disclosure and the interests of the vendors.

One unanswered question in this area concerns severity of vulnerabilities. Not all vulnerabilities were born equal – some have more severe security consequences than others, both because of how they are used and because of inherent properties. Should all vulnerabilities be treated the same, and if not, how should more severe vulnerabilities be treated differently? When is it worse to disclose the vulnerability information publicly than to keep it a secret? And if security professionals adopt a policy of keeping severe vulnerabilities secret, does that reduce the incentive for software companies to secure the highly-sensitive parts of their software?

### 2.2.2 Incentives for Discovery

Rescorla (2004) addresses the effects of disclosure on the software that contains the flaws, and discovers no evidence that this disclosure has lead to an increase in quality. This weakens the argument that discovery of vulnerabilities is good for social welfare. However, there are still many open questions about this. How can we accurately measure the impact of disclosure on the software that contained the flaw? Rescorla really is looking to see if we appear to be exhausting the possible vulnerabilities, and finds that we aren't. It would be interesting to see if public disclosure causes future versions to be more secure.

Anderson (2004) looks at differences in software (open source software versus proprietary software) and studies whether vulnerabilities are easier to find in one type or the other. He concludes that in a perfect world both types of software are equivalent, but under more realistic assumptions there is a difference. He also attempt to characterize what assumptions lead to which differences.

It is an open question whether easier-to-discover vulnerabilities benefit the attacker or defender more. Attackers may have less trouble finding vectors of attack with easy-to-find vulnerabilities, but defenders have less trouble finding and patching these bugs also.

### 2.2.3 Vulnerabilities and Markets

Schechter (2002) originally proposed having a market for vulnerability information. Since vulnerability information indirectly contains information about the quality of software, software companies can use this market to purchase vulnerability information and publicly prove statements about quality.

This opens up a host of new issues, though. The most important is ‘What kind of statements can software companies credibly make using a vulnerability market?’ Schechter suggests ‘The security vulnerabilities in my software are more difficult to find than the ones in my competitor’s software’, but this is not the only possible statement. Care must be taken with these statements, though. Many ‘snake-oil’ cryptography companies have tried to use ‘challenges’ to achieve the same end, but since they were never collected on they actually provided very little information.

Often, skilled hackers will not participate in these markets. How well can these markets work when there is only a small set of people looking for vulnerabilities? I suspect that absolute statements about quality cannot be made, but relative statements can be.

Ozment (2004) studies what the best method of implementing such markets would be, and suggests that a Vickrey auction would serve well.

## 2.3 Spam

Spam email is a problem that haunts everyone with an email address. It arises from an incentive problem where it is extremely cheap to send lots of untargeted messages to unwilling recipients, but these messages then cause an information overload problem for the recipients. Spam is a security problem since spammers and recipients are in an adversarial setting, and therefore spammers work hard to get around any anti-spam (security) measures that the recipient puts in place. Spam is also related to computer security because:

- Spam is often sent out by computers that have been hacked, much the same way viruses are sent out via email.
- Spam is sometimes used to gain vital security information through fraudulent messages and websites. This is called phishing for information. This information can then be used to access bank accounts and such.

The overarching question for this area is *How can we prevent unwanted spam messages from consuming our attention without harming third parties and their ability to communicate?* I focus first on technological solutions and then on economic solutions.

### 2.3.1 Technological Solutions

There have been a number of technological solutions that have been proposed and used. These solutions have varying degrees of success, but almost all of these solutions have unintended social consequences that can’t be measured in traditional measures like ‘number of messages.’ It is important to understand these consequences when evaluating these solutions.

One proposed solution is to require email senders to ‘pay’ by proving that they have done some amount of a time-consuming calculation. Laurie and Clayton (2004) gives evidence that this solution will not work because the amount of work required to noticeably reduce spam email will also noticeably reduce or eliminate legitimate email.

Another proposed solution is to block the IP address (Internet address) of the machines sending the spam, or more effectively to block the whole ISP of the machines sending spam. Serjantov and

Clayton (2005) attempt to calculate the effect that blocking an ISP will have on other consumers by looking at email communication patterns and determining the number of affected consumers. They conclude that email blocking is not a feasible strategy without affecting a large proportion of email communications.

A very popular spam solution is to use technological filters. It is an open question to what extent these filters cause problems in email. Richardson, Resnick, Hansen, Derry, and Rideout (2002) conducted a study of web filters to determine how much legitimate web traffic they blocked. A similar study of email would be useful. If technological email filters disproportionately block certain types of communications, then it is likely that they will be ineffective for people who regularly partake in such communications.

One promising solution involves using ‘secure email addresses.’ (Gabber, Jakobsson, Matias, and Mayer, 1998) This paper proposes a largely technical solution that uses separate email addresses for each possible use, and then allows users to block certain incoming addresses if spammers use them. This paper includes an economic adversarial model of email and analyzes the solution in this model.

### 2.3.2 Economic Solutions

Since these purely technical solutions have not been effective enough, there have been attempts to use economic mechanism design to provide a solution. Loder, van Alstyne, and Wash (2004) propose a payment mechanism for email, and give an economic model for studying this mechanism. They also compare this mechanism with technological filtering and a government tax. This model only considers individual emails, and does not account for non-email considerations. It is an open question if people’s use of email will change with such a mechanism in place.

There have been a number of proposals for regulatory solutions to spam, including forced labeling, taxes, and increased penalties for spammers. It is an open question as to the effectiveness of these solutions in an adversarial setting, for example when spammers can leave the jurisdiction fairly easily to avoid these laws.

## 2.4 Privacy

Privacy is often a hotly-debated topic. Identity theft is a common worry on the Internet, and many companies exist to fight this problem. Consumers say they are very reluctant to disclose their personal information for fear that it can be used against them. Inadequate security often leads to unintended disclosure of personal information. Yet despite all of this, consumers have repeatedly shown themselves willing to trade their privacy for very small amounts of money, a candy bar, or a couple cents discount at the grocery store.

Firms have the incentive to collect personal information. It can allow them to better distinguish between customers, which enables price discrimination. Additionally, firms have little incentive to actually protect this private data once they have collected it compared to the people the data is about who stand to suffer if the data is in the wrong hands.

The overarching question in this area is *How valuable is the privacy of personal information and what is the best way to ensure this privacy?*

In this section, I focus on three main issues related to the economics of privacy. First, what are the incentives for firms to collect personal information? Second, I look at the underlying causes of the privacy paradox and for possible solutions. Finally, I look at the practical problem of protecting private information from abuse and insuring its value is properly taken into account.



### 2.4.1 Incentives to Collect Personal Information

Odlyzko (2003) and Acquisti and Varian (2005) both study how the ability to price discriminate, or charge different types of people different prices, motivates firms to collect personal information. Acquisti and Varian (2005) actually find that online, firms normally do not have the incentive to collect personal information for price discrimination, but if they can customize their service for individuals then they do have such an incentive. An open question is, in practice, how strong are these incentives to collect personal information? What reasons other than price discrimination and customization would firms collect personal information?

Most firms that collect personal information have a posted ‘privacy policy’ that states what the firm is allowed and not allowed to do with the information. However, there is little legal enforcement of these policies. An interesting question would be to see how violating such a policy affects the firm’s profits, if at all. Also, how does putting stronger penalties for violation influence their incentives to collect the information? One strategy to prevent violating such policies is to change the policy when needed rather than violate it. How can this loophole be closed?

### 2.4.2 The Privacy Paradox

There is a well-known paradox related to privacy. When asked a hypothetical question about their personal privacy, people are adamant about how it should be protected. They tend to value their privacy highly. However, when put to the task and forced to make a decision about revealing their private information for a small reward, they frequently choose the reward over keeping their information. An important question in this area asks why people behave like this. Which answer is really accurate, the high valuation in a hypothetical, or the low valuation in a choice scenario? A second question is how to correct for this, or to induce people to act appropriately.

Acquisti (2004b) looks at the first question and proposes a theory of immediate gratification. This theory states that people disproportionately value having something now, and do not rationally discount future periods. He describes experiments to validate this, and some of the consequences of such a decision rule. This work assumes that the high hypothetical value of private data is accurate, and the choice that people make is not in their self interest. Is there a similar theory based on the opposite assumption?

### 2.4.3 Protecting Privacy

In the end, what people really care about is protecting their privacy. However, since private data is extremely useful, people don’t just want to keep it a secret at all costs. Revealing this information to companies can be beneficial to all parties involved. However, consumers want to know that companies will properly protect their personal information.

Varian (1996) proposes giving consumers property rights over personal information about them. This is the current solution in use in Europe. Varian analyzes this idea and studies various policy levers that are available. He concludes with a warning that too rigid of property rights could stifle business. It is an open question what type of property rights would be most appropriate for personal information. It is also not clear if such a protection scheme would provide the appropriate incentives for companies to protect personal data when you take into account transaction costs of enumerating permissions.

Samuelson (2000) looks at two different methods of protection: property rights, and a market-based scheme. The market-based scheme is based on default licenses and explicit rules for violating these licenses. These would provide legal remedies for consumers against companies who violate their posted privacy policies. She advocates the market-based scheme since it adapts better.

For both of the above schemes, it is important to take into account the privacy paradox. How does people's irrationality affect the efficiency of these protection schemes? Do either of these schemes get around the paradox problem, and if not can we devise a scheme that does?

## 2.5 Spending and Investment

Spending on information security technologies is a risky but necessary business. First of all, IT (and security in particular) is not a profit center. Also, security, unlike IT, is not an enabling technology, so its benefits cannot be measured by increases. Security is a risk management technology, meaning that it attempts to prevent or control the losses associated with security breaches, and its impact is a reduction in the cost of doing business. This distinction is even more important for home users, who don't have the time or skills to do complex risk management.

Firms are trying to figure out how much they should spend on protecting their information resources. The security industry says that businesses aren't spending enough (obviously). Stockholders are saying that they are spending too much. Until, that is, the firm has a security breach, at which point the stockholders are upset that more wasn't spent. Overall, the question that firms have is *How much should we spend on information security, and how should we spend it?*

To answer the first part of the question, Adkins (2004) develop a model based on insurance and financial derivatives to determine how much the security would be worth (as a risk management strategy) to a business. To answer the second part, Gordon and Loeb (2002) develop a return-on-investment (ROI) system for determining whether a specific technology is worth investing in. They conclude that often it is not worth spending to protect the most vulnerable systems, and that spending should be significantly less than the value being protected to maximize the expected value of the investment.

These papers do not fully answer the question, though. They are more like starting points. Camp and Wolfram (2000) shows that information security has strong externalities, and individually rational spending decisions are unlikely to allocate the proper amount to the problem. I am also interested in how complementarities between targets affects vulnerability. If two systems have information that is complementary, then a hackers incentive to break into both is greater than their incentives to break into them separately.

Another important question concerns the feasibility of 'security insurance.' Businesses will want to insure against security breaches, but security breaches have a number of issues that complicate insurance. Specifically: there is often correlated risk, there is little data about their frequency of occurrence, it is difficult to value losses due to stolen information, and losses are due to the malicious behavior of other people, not due to exogenous random forces. An open question is 'What is the best way to get around these problems to provide affordable insurance?'

## 2.6 Security Breaches

Security breaches are a fact of life. Most large companies have to deal with them occasionally (and sometimes more than occasionally). After a breach, there is a lot of obvious cleanup that needs to be done, such as reinstalling the compromised computers, figuring out what the attacker has done, and repairing any damage caused. But there are also a number of non-obvious decisions that need to be made by the compromised firms.

One beginning question is how much damage do these security breaches cause? This question is important in risk calculations, for buying insurance, and for planning. Campbell, Gordon, Loeb, and Zhou (2003) used stock market data to estimate the damage that various security breaches caused. They found that certain types of breaches, specifically those involving the compromise of

sensitive data, are more damaging than others. This is not the end of the story, however. It would be useful to break these costs down into its components: How much was spent on cleanup? How much was lost due to damage? How much did people's opinions of the firm decrease due to the security problem?

Another important question in this area regards sharing information about these security breaches with others. This information can be very useful in preventing future breaches. However, this information also generally has negative public relations value, and therefore firms are naturally reluctant to share it. Gal-Or and Ghose (2003) uses a game theory model to study the incentives that firms have to share such information, specifically in the context of industry organizations created for this purpose. They find results that security investment and security breach information are 'strategic complements', sharing alliances are more beneficial in more competitive industries, and benefits increase with the size of the firm.

Open questions still remain. Can the government use any public policy mechanisms to increase the amount of sharing of security breach information? And is such sharing beneficial? California has a mandatory disclosure law that states all firms in California must notify all California citizens when their personal information is stolen via a computer security breach. While this law is mainly a privacy law, it also serves to inform other firms about security breaches in California.

## 2.7 Global Insecurity

Camp and Wolfram (2000) point out that information security has a strong externality component to its value. Since the Internet breaks down geographic borders, hackers in Asia can break into 10,000 vulnerable machines in Europe and then use those machines to conduct a Denial-of-Service attack on an American business. In fact, in early 2000 hackers used even more machines than that in a distributed denial of service attack on eBay, Yahoo!, CNN, and more. The machines used were distributed across the globe.

This fact is even more apparent when you consider computer viruses and worms. There have been a number of major worm outbreaks. The Code-Red worm was extremely damaging: "359,104 hosts were compromised in approximately 13 hours."<sup>2</sup> Tracking and stopping these worms is becoming more and more difficult as the adversaries get better at writing these worms. The Sapphire worm exemplified this problem. It was the fastest-spreading worm observed yet: "It infected more than 90 percent of vulnerable hosts within 10 minutes."<sup>3</sup> Fortunately, techniques for after-the-fact analysis are improving (Kumar, Paxson, and Weaver, 2005). Weaver and Paxson (2004) studied the question of how much economic damage could possibly be done by a worm, and the result is rather large (on the order of \$50 billion).

The causes of this global insecurity are both technical and economic in nature. Technically, it is very difficult to secure that many machines. But economically, it is even more difficult because all of these machines are subject to distributed control, and their owners do not feel most of the damage their insecurity can cause. Geer, Bace, Gutmann, Metzger, Pfleeger, Quartermann, and Schneier (2003) describes how diversity is beneficial to prevent these types of problems, but individuals do not have any incentive toward diversity in their software purchasing decisions.

It is an open question what the best way is to provide stronger incentives to individuals to account for this externality. There have been proposals for liability whenever your computer is used to attack another. This solution has problems (the transactions costs of suing 100,000 people is prohibitive), but a full analysis has yet to be done. Another proposal concerns mandatory

---

<sup>2</sup>[http://www.caida.org/analysis/security/code-red/coderedv2\\_analysis.xml](http://www.caida.org/analysis/security/code-red/coderedv2_analysis.xml)

<sup>3</sup><http://www.caida.org/outreach/papers/2003/sapphire/sapphire.html>

information security insurance to attempt to centralize the risk. More work along these lines would be very valuable.

## 2.8 Theory

Information security has some interesting theoretical properties that make it quite unique and interesting to study. First of all, there is an adversarial relationship between the parties involved. This does not usually mean that one person's utility function appears in another's (such as in negative altruism), but that they are both correlated with some external effect, and in opposite directions. For example, a hacker's utility is positively correlated with their success in hacking, and the defender's utility is negatively correlated. However, these agents make rational choices, and as such do not always act in strictly opposed ways.

Another property that makes it interesting is that security is almost always characterized by imperfect and asymmetric information. Without the ability to observe the actions of attackers, it is difficult to tell the difference between the hacker choosing not to attack (the security system working), and the security system failing silently. This causes difficulty in evaluating security systems, and consequently difficulty in practically reasoning about them. Also, there is imperfect information over time, and systems once thought secure are no longer after a vulnerability is discovered. The information about the security of systems is also asymmetric, with hackers often knowing different things than defenders.

Information security is also characterized by the low marginal costs of information goods. Protecting information from compromise is difficult because it is so costless to transfer it anywhere on the Internet. A related problem is that it is possible to separate 'skill' and 'ability'. Skill is the knowledge of how to be a hacker. Ability is the means to carry out the attack. Ability can be written into a computer program and automated, but only by a skilled attacker. This means that it is possible to turn otherwise benign computers into 'zombie hackers' that attack other computers at the skilled hacker's will. It also means that there exist 'script kiddies,' which are mostly unskilled hackers that use programs written by more skilled hackers to break into systems. This separation of 'skill' and 'ability' also enables class attacks, when a hacker attacks a whole class of vulnerable systems simultaneously. Since the ability to attack can be automated, these attacks can be carried out rapidly, one after the other, with only microseconds between them. Class attacks are similar to correlated risk.

Globalization also poses problems for information security. The Internet enables hackers from anywhere on the globe to attack systems. This larger scale of interest is daunting for people and businesses who otherwise would only care about local issues. A local pizza delivery business who wants to accept orders on the Internet now has to suddenly worry about teenage pranksters in Russia. It also causes problems for the law, as hackers can 'jurisdiction shop,' or move to the country that has the most favorable hacking laws.

Finally, information security has many externalities that complicate market solutions. There are many network effects in the adoption and use of security products and solutions. For example, DRM technologies will never be used unless there is a widespread belief that they will because the content producers would not choose to limit their audience by placing their content on an unpopular platform. A negative version of this is the monopoly argument in Geer et al. (2003), where a 'monoculture' of software creates increased problems. Another example, discussed in more detail above, is the externality in vulnerability from Camp and Wolfram (2000), where I am more insecure the more insecure you are.

Anderson does a good job of describing the difficulties that information security face in his two main papers on the subject (Anderson, 1993, 2001). He discusses what types of problems cause

security systems to fail, and concludes that they are mostly problems with people and incentives, not with the technology. He also discusses how the market for security products is similar to a lemons market (Akerlof, 1970) because of the difficulty in evaluating security products.

Some of the theoretical work on this topic came from Varian (2002). He was concerned with the amount of effort the agents involved chose to exert to defend something. He also considers the adversarial setting where the effect of effort goes in opposite directions for different agents. In his overly simple world, he concludes that either attackers will give up and defenders will do the minimum amount of security, or attackers will work hard to succeed, and defenders will work hard to make sure the attackers don't have it easy. Which of these two actions will be chosen depends fully on the relative cost-benefit ratios.

Dekel and Scotchmer (1999) study risk behavior in the context of game theory. They find that winner-take-all games encourage risky behavior, and in an evolutionary context most people who play winner-take-all games are risk-loving. This is relevant to security since security is mostly a winner-take-all game. Therefore, most of the types of people that will become hackers are risk-loving.

### 2.8.1 Class Attacks

One type of attack that is more prevalent in information security than traditional security is the class attack. A class attack is a set of related attacks by a hacker that all exploit the same vulnerability, or a small set of vulnerabilities. Class attacks are particularly dangerous in information security because they can be automated and carried out at extremely fast speeds.

Some of the early work on this was done by Schechter and Smith (2003). He develops a model of a firm evaluating their attractiveness to an attacker that includes the choices of other firms. If everyone else is using the same software package, then choosing that package increases the attractiveness by decreasing the average cost of breaking into all of those firms. They also study the incentives of defenders to collude and share information.

Chen, Kataria, and Krishnan (2005) analyze this problem within a firm. They consider the case where a firm has many systems and a queue for service and repair. They conclude that having a diversity of systems rather than a monoculture greatly protects against class attacks and decreases the repair time for large attacks.

Geer et al. (2003) take this idea and apply it in a larger setting, that of global viruses and worms. They look at the current situation of having a monopoly vendor for operating systems (Microsoft), which are the most common target of current attacks on the Internet. They conclude that the US is vulnerable to attacks on the infrastructure of our information systems because of this. They recommend diversifying the market.

This brings up an interesting question for innovation policy. Many current thinkers in innovation policy believe that for these types of information goods, 'competition for the market' can replace 'competition in the market.' (Scotchmer, 2004) Even though it results in a monopoly it should not be subject to anti-trust because new innovations can overtake the incumbent technology. It is an interesting question to consider how class attacks on such a 'natural monopoly' change this reasoning. Is it still OK to allow one firm to dominate a market such as the one for operating systems if it leaves the US information infrastructure open to attack?

Class attacks are a form of correlated risk. It is interesting to see how they change analysis. For example, in the DRM work below, I discuss a possible class attack that might change the incentives for DRM platform providers. Diversity is the most common defense against correlated risk, but software firms are often not interested in that solution since it means fewer sales for them. Are there other workable solutions to class attacks and common vulnerabilities that don't require users

reducing their purchases?

## 2.8.2 Cognitive Attacks

Thompson (2002) studies a special type of information security attack – one where the attacker spreads misinformation for his own gain. The example he uses is a hacker that forged a press release about a company being investigated by the SEC, and sold that company short. He develops a general model of this type of attack that uses a horse-race metaphor. He also briefly discusses countermeasures.

Many attacks fit into this category. One example is email phishing schemes, which are fake emails that attempt to get users to enter private information on a fake website. Famed hacker Kevin Mitnick calls these attacks ‘social engineering.’ These attacks take advantage of people’s trusting nature and/or ignorance. Con artists are the real-world equivalent of hackers who use such attacks. They fit in the category of ‘information security’ when they involve the use of information technology, which has a total different set of social assumptions than the real world does. The globalization of the Internet also means that what is considered unsavory in one culture but not in another can cross the border very easily.

These attacks are particularly dangerous because they are very difficult to protect against with technology. There is no technology that can make sure that someone who legitimately has access to a system is not using it because they were duped by hacker. However, there may be incentive mechanisms that might be useful against these attacks. It is an open question whether mechanism design can be used to protect against such attacks.

## 3 Focusing on Content Provision and Digital Rights Management Technologies

Piracy of digital goods has become a large problem on the Internet due to its ability to distribute large quantities of content to large quantities of people at very low cost. As a result, many digital content producers have turned to technology to insure their revenue stream. These technologies, collectively known as *Digital Rights Management (DRM)* technologies, use encryption and other security technologies to artificially restrict the uses of a digital good.

An example of DRM is Apple’s iTunes FairPlay technology<sup>4</sup> for its iTunes Music Store. All songs downloaded from the store are encrypted by Apple. Apple’s iTunes Music Player, when so authorized by Apple, will decrypt and play these songs. Apple will only authorize computers after users have paid for the music. Apple has a set of policies listing what they permit users to do with songs, including (as of September 6, 2005):

- Authorize up to 5 computers to play the songs
- Transfer any songs to an iPod portable music device
- Burn songs to CD an unlimited number of times
- Burn any single playlist of music to CD up to 7 times
- Re-download the encrypted songs from iTunes Music Store

---

<sup>4</sup><http://www.apple.com/iTunes/>

Note that in this example, Apple does not permit their users to copy the music to friends very easily. Normal music files do not have this restriction, and the only reason this restriction exists is that it was artificially put there by Apple's DRM technology. Also notice that the iTunes Music Player (which runs on the local user's computer) has the decryption keys stored on the hard drive so that it can play the music. A sufficiently sophisticated user can always find these keys, decrypt the music files themselves, and save un-encrypted and therefore unrestricted copies. However, this act of 'circumvention' is costly and time-consuming, and many users either don't have the requisite skills or decide that they would rather just pay for the music and live with Apple's restrictions.

DRM technologies have many effects that must be understood. They reduce the utility to consumers of digital content by restricting its use. By preventing piracy, they can increase the content producer's profit. A side effect of increased profit is an increased incentive for innovation, leading to more digital content being available. (Scotchmer, 2004) DRM technologies also have strategic considerations for pricing (Park and Scotchmer, 2004) and flexibility of use (Bergemann et al., 2005).

Not all digital piracy is bad for the content producers. Oberholzer and Strumpf (2004) do an empirical study of the music industry and find that despite widespread piracy, the industry is quite healthy. Givon et al. (1995) use an information diffusion model to estimate that the UK software industry sees increased sales (on the order of 70%) due to word-of-mouth recommendations of pirated users. And Sundararajan (2004) develops an economic model of piracy that includes these beneficial effects of piracy and attempts to develop strategies for profit maximization that include anti-piracy efforts. This paper, however, does not deal with this aspect of piracy. I assume that there are no positive producer benefits from piracy (obviously the pirate user benefits).

I am interested in an idea that was first observed by Acquisti (2004a). In that paper, Acquisti develops a model of platform adoption in the presence of network effects. The DRM technology makes transferring between two networks (open and DRM networks) possible, but one direction is costless (open  $\rightarrow$  DRM) and the other costly (DRM  $\rightarrow$  open). In addition to other results, Acquisti derives a result that popular content is more likely than niche content to be transferred to the open platform if it is initially only available on the DRM platform. He therefore concludes that the success of a DRM platform depends on 'user-generated' content (as opposed to 'widely-popular vendor-generated' content).

I use this idea as a basis for studying incentives for innovation in a content industry that uses DRM technologies. I study the situation where there is a negative network effect in the costs of 'breaking' the DRM, or transferring the content from the DRM platform to the open platform. As content becomes more popular, the average cost of breaking decreases since the content will more likely fall into the hands of people capable of breaking it, hackers will be more interested in breaking it because it can be shared farther, or there exist returns to scale to breaking DRM technologies. In this situation, I hypothesize that content producers will prefer high-value 'niche' content over low-value 'mass-market' content when using DRM. Here I study the implications of such a preference.

I first develop an theoretical model that builds upon the model of DRM technologies initially developed by Park and Scotchmer (2004). I modify this model to include network effects in the cost of breaking the DRM protection. This is a supply-side network effect, which has not been studied much in the literature. In order to study different types of content, I use a model of content originally developed by Mackie-Mason et al. (1996). I expand on their definitions slightly to allow greater flexibility in decisions of content providers.

After discussing this theoretical work, I move into the real world and discuss some possible avenues of empirical / econometric research ideas.

### 3.1 Model

The DRM industry has three classes of agents: Content Producers, Platform Providers, and Consumers. This is very similar to what economists are currently calling a “two sided market.” (Rochet and Tirole, 2003) Content producers want to sell their content to consumers, but have the problem that consumers can then do many things with this content that the producers don’t want them to, such as share it on the Internet with the consumer’s closest million friends. Platform providers have the technology that enables consumers to use content, but only if producers have placed their content on the platform.

All digital content must exist on some type of platform, or combination of hardware and software that can be used to view/use/access the content. DRM platforms are content platforms that include some technology that attempts to artificially restrict the content’s use beyond what would normally be implied by standard technology and legal controls. DRM platform providers then have to convince both the content producers and the consumers simultaneously to use their platform to access content.

I model the decision that faces a content producer. Following Park and Scotchmer (2004), my model of DRM technology has the producer choosing some level of protection,  $e$ . This protection has cost  $K(e)$ , with  $K'(e) > 0$  and  $K''(e) > 0$ . Consumers can purchase this content at price  $p(q)$  from the producer, or they can break the DRM at cost  $b(e)$ , with  $b'(e) > 0$ . (Note that Park and Scotchmer have  $b(e) = e$ ). Quantity  $q$  will be purchased at price  $p(q)$ . Therefore, the producer’s decision problem is

$$\max_{p,e} p(q) \cdot q - K(e)$$

such that

$$p(q) \leq b(e)$$

I now modify this model to include network effects in the cost of breaking. Formally, I redefine the function  $b(\cdot)$  to take a second parameter,  $\hat{q}$ , which is the expected quantity of content that consumers possess. Therefore, the producer maximization problem is that I consider here is:

$$\max_{p,e} p(q) \cdot q - K(e) \tag{1}$$

such that

$$p(q) \leq b(e, \hat{q})$$

Finally, in equilibrium, I set  $\hat{q} = q$  to become a fulfilled-expectations equilibrium. (Katz and Shaprio, 1985)

I assume that  $\frac{\partial b(e, \hat{q})}{\partial \hat{q}} < 0$ , which means that the cost of breaking decreases as the content becomes more popular.

Next I model the possible content types that a producer can choose to produce. For simplicity, I assume that content has a linear demand parameterized by  $\bar{p}$  and  $\bar{q}$ . The demand function is therefore

$$p_{\bar{p}, \bar{q}}(q) = \bar{p} \left( 1 - \frac{q}{\bar{q}} \right) \tag{2}$$

Following Mackie-Mason et al. (1996), I vary content along two axes: content can either be high-value ( $\bar{p} = p^H$ ) or low-value ( $\bar{p} = p^L$ ), and content can either be mass-market ( $\bar{q} = q^M$ ) or niche ( $\bar{q} = q^N$ ). These four types of content are illustrated in Figure 2.






Figure 2: Four types of Content

### 3.2 Results

To solve this model, I proceed using a method similar to that which Economides used for network effects. (Economides, 1996) I first solve for the optimal level of protection  $e$  in terms of both the actual quantity sold, and the quantity expected to be sold. I then proceed with the maximization of producer utility. Finally, I impose the fulfilled expectations condition.

Now, I define a function  $f(q, \hat{q})$ :

$$\begin{aligned} f(q, \hat{q}) &= \{e | b(e, \hat{q}) = p(q)\} \\ &= b^{-1}(p(q), \hat{q}) \end{aligned}$$

$f$  is a function that gives the level of protection  $e$  needed to achieve sales of  $q$  items given consumer expectations at level  $\hat{q}$ .

The producer decision now can be written as

$$\max_q p(q) \cdot q - K(f(q, \hat{q}))$$

After solving this for  $q$ , I can impose the fulfilled expectations condition to arrive at the final profit.

For comparison, a pure monopolist's decision (absent illegitimate copying) is

$$\max_q p(q) \cdot q$$

I can now show some interesting comparative statics about my model. But first I will describe a useful theorem from the paper 'Monotone Comparative Statics' by Milgrom and Shannon. A function  $f(x, t)$  has *increasing differences* in  $(x, t)$  if  $\forall x' \geq x'', t' \geq t'' f(x', t') - f(x', t'') \geq f(x'', t') - f(x'', t'')$ , or equivalently  $f(x, t') - f(x, t'')$  is increasing in  $x$ . Now the useful theorem:

**Theorem 1** *Let  $f : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ . If  $f(x, t)$  has increasing differences in  $(x, t)$ , then  $\operatorname{argmax}_x f(x, t)$  is weakly increasing in  $t$ .*

PROOF This theorem is shown in more general form as Theorem 5 in Milgrom and Shannon (1994).■

Now I can compare the producer's decisions about quantity and price under monopoly and under DRM:

**Lemma 1** *In a world where a producer has to use DRM to prevent copying, he sells more units at a lower price than he would as a monopolist, as long as  $\frac{\partial f}{\partial q} + \frac{\partial f}{\partial \hat{q}} < 0$ . If this condition is reversed, then he sells fewer goods at a higher price.*

PROOF Define  $h(q, t) = p(q) \cdot q - t \cdot K(f(q, \hat{q}))$ . Notice here that  $t = 0$  corresponds to the monopolist's decision problem, and  $t = 1$  corresponds to a DRM producer's decision problem. I want to show that  $h(q, t)$  has *increasing differences* in  $(q, t)$ . If  $h(q, t)$  has increasing differences, then by the Theorem 1 above,  $\text{argmax}_q h(q, t)$  is increasing in  $t$ . This would mean that the producer's optimal  $q$  is greater in a DRM world than as a monopolist. A direct consequence of this is that the price charged is less (assuming downward-sloping demand curves).

To show increasing differences of  $h(q, t)$ , I need to show that  $h(q, t') - h(q, t'')$  is increasing in  $q$  for  $t' > t''$ :

$$\begin{aligned} h(q, t') - h(q, t'') &= p(q) \cdot q - t' \cdot K(f(q, \hat{q})) - \\ &\quad (p(q) \cdot q - t'' \cdot K(f(q, \hat{q}))) \\ &= (t'' - t') \cdot K(f(q, \hat{q})) \end{aligned}$$

Since  $t'' - t' < 0$ , what remains is to show that  $K'(f(q, \hat{q})) \cdot f'(q, \hat{q}) < 0$ . Now  $K' > 0$  by assumption. Now, I need  $f'(q, \hat{q}) = f_1(q, \hat{q}) + f_2(q, \hat{q}) < 0$ . Remember that  $f(q, \hat{q}) = b^{-1}(p(q), \hat{q})$ . Now  $b_e > 0$ , so  $b_p^{-1} > 0$ , and  $b_q^{-1} = b_p^{-1} \cdot p_q < 0$ . Therefore,  $f_1 < 0$ .

If I define  $h(q, t) = p(q) \cdot q + t \cdot K(f(q, \hat{q}))$ , then  $t = 0$  is still the monopoly situation, and  $t = -1$  is the DRM problem. It can be shown that  $h(q, t)$  has increasing differences iff  $f_1 + f_2 > 0$ . By the theorem above, this  $h(q, t)$  having I.D. means that the producer's chosen  $q$  is higher under monopoly than DRM. ■

This result makes sense. In the DRM world, a content producer has to compete with illegitimate copies of its own content. It can set the price of the illegitimate copies, but only at a cost. (This is similar to the literature on raising rivals costs (Salop and Scheffman, 1983) in antitrust economics.) This competition forces it to lower prices and increase demand. A similar result was in Park and Scotchmer (2004).

The condition also makes sense. It basically says a producer will only lower quantity to fight off piracy when lowering quantity has more benefit in the battle than raising the strength of DRM does.

Now I use a numerical example to prove my main proposition:

**Proposition 1** *There exists a situation where a monopoly content provider would be indifferent between niche and mass-market goods, but a content provider using DRM would prefer a niche good to a mass-market good.*

PROOF Assume the following:

$$\begin{aligned} p(q) &= \bar{p} \left(1 - \frac{q}{\hat{q}}\right) & b(q, \hat{q}) &= e - \frac{1}{4}\hat{q} \\ K(e) &= \alpha \cdot e^2 & b^{-1}(p, \hat{q}) &= p + \frac{1}{4}\hat{q} \\ b^{-1}(q, \hat{q}) &= f(q, \hat{q}) = \bar{p} \left(1 - \frac{q}{\hat{q}}\right) + \frac{1}{4}\hat{q} \end{aligned}$$

Solving the profit maximization problem, I find that

$$q = \frac{\bar{p} + 2\alpha \frac{\bar{p}}{\hat{q}} (\bar{p} + \hat{q})}{2\frac{\bar{p}}{\hat{q}} + 2\alpha \left(\frac{\bar{p}}{\hat{q}}\right)^2}$$

Here I let  $\alpha = 1$ . Now, I have two goods that were defined above:

$$\begin{aligned} \text{Niche: } & \bar{p} = 2, \bar{q} = 1 \\ \text{Mass Market: } & \bar{p} = 1, \bar{q} = 2 \end{aligned}$$

With the linear demand curve I use, a monopolist will always set  $p^* = \frac{1}{2}\bar{p}$ ,  $q^* = \frac{1}{2}\bar{q}$ , and  $\pi = \frac{1}{4}\bar{p}\bar{q}$ . As such, a monopolist would be indifferent between these two goods.

By plugging in these numbers, it is easy to see that profit for the niche good is greater than the profit for the mass market good under DRM. ■

### 3.3 Various Effects of DRM

In this research, I take a consumer-centric approach, studying the effects of this insight on consumer welfare. Since the use of DRM technology is voluntary on the part of the content producer, it will only undertake this technology if it improves its welfare. I am primarily concerned with public policy questions: Is this change in content provision troubling? What policy levers exist for lawmakers to correct for this content provision?

To understand the effects on consumer welfare, I compare the DRM situation with the no-DRM situation. The no-DRM situation is not straightforward to model either. I begin by looking at two simplified extremes as strawmen: that in which illegal copying doesn't happen (this corresponds to the 'perfect legal enforcement' world of Park and Scotchmer (2004)), and that of fully-rational consumers (who, in the absence of DRM, copy everything and never pay). These are both extremes, and neither of these exist in the real world. However, it is illustrative to study them as the real world is somewhere between them.

I model the no-copying world as a normal monopoly situation. The producer attempts to maximize  $p \cdot q$ . Since my demands are linear, he would always choose  $\frac{1}{2}\bar{p}$  and consequently  $\frac{1}{2}\bar{q}$ , leaving profits  $\frac{1}{4}\bar{p}\bar{q}$ . He will choose content in the 'normal' ordering.

In the fully-rational world, consumers will choose to copy content if at all possible. To incorporate this into my model, force  $e = 0$ . Consumers have the choice of paying price  $p$  or copying for  $b(0, \hat{q})$ . For simplicity, assume that when  $b(0, \hat{q}) = b_0$  for all  $\hat{q}$ . This means that the base cost of copying  $b_0$  is constant and independent of both quantity sold and DRM technology. In this world, a producer will be forced to choose  $p = b_0$  if they want any sales at all. Therefore, his surplus will be  $b_0 \cdot p^{-1}(b_0) = b_0 \cdot \bar{q} \left(1 - \frac{b_0}{\bar{p}}\right)$ . He will choose content ordered primarily by  $\bar{q}$ , preferring mass-market content over niche content.

Another consideration is an intermediate world. For example, consider the world where there are  $\alpha$  consumers who refuse to copy, and  $1 - \alpha$  consumers who are fully rational.

I also consider heterogeneity of consumers breaking cost. Consider a cumulative distribution function  $F(b; e, \hat{q})$  where  $F$  is the proportion of the population who's breaking cost is less than  $b$ . If I assume  $F(\cdot; e', \hat{q}) >_{\text{FSD}} F(\cdot; e, \hat{q})$  when  $e' > e$ , then I have a similar model to that given above. This heterogeneity can be because of different skills, or it can be because of differences in consumers distaste for illegal copying.

I am concerned with the choice of competing DRM technologies. Consider multiple DRM platform vendors, each of which offer different levels of protection at different costs. Are their strategic considerations that competition in DRM platforms bring about that effect consumer welfare? There are many issues here including adoption decisions including network effects.

Another line of inquiry concerns the concept of a 'class break.' A class break is what happens when a hacker breaks the underlying security technology, thereby also breaking all current and future instances of that technology. (Chen et al., 2005) A class break would break the encryption

protecting the content, leaving all content on the DRM platform available for copying. If the cost or likelihood of a class break depends on the total number of copies of content sold (or on the average quantity per piece of content, or on the quantity of the most popular piece of content), then a DRM platform provider may have an incentive to restrict what content is available on his platform. They may prefer high-value niche content to try to keep the incentives for a class break low. How does this situation compare to an ‘aware’ network from Mackie-Mason et al. (1996)? Does this lead to a ‘natural oligopoly’ where there is a diversity of DRM platforms each of which have a moderate amount of content?

This research also raises a question of content provision under different innovation incentive schemes. For example, are there likely to be different types of content created when there is a reward scheme rather than an intellectual property scheme, as in Shavell and van Ypersele (2001)?.

### 3.4 Empirical Evaluation

Another method of determining if this predisposition to using DRM on niche content really exists is to do an econometric study of a content industry that uses DRM, such as the online music industry, or the DVD sales market. Apple’s iTunes Music Store is a good instance of such an online market.

Apple’s iTunes Music Store (iTMS) is an online interface to purchasing and downloading music in electronic form. Apple uses its FairPlay DRM technology to restrict the uses of the songs that are downloaded. iTMS uses a one-size-fits-all DRM technology, and all of the music has the same restrictions. Music producers and distributors can choose which songs in their catalog to make available via iTMS. They also have the option of making certain songs only available as a bundle – you must buy the whole album to get the song.

I think it would be interesting to validate this idea by looking at the data generated by iTMS to determine if there is an effect on content provision. It is possible to get price data directly from iTMS. Quantity data, however, is not as directly observable. However, there are many music industry sources that publish some amount of quantity data on overall industry sales of specific songs. Or at least, they have rank indicators like the Billboard Top 100. I believe that this whole-industry quantity data is more relevant to my study than just the iTMS quantity data would be.

One complication is how to measure a change in content provision. I initially considered looking at time-series data to see how the distribution of quantities in the music industry changed as they began introducing music on the Internet. However, it would be very difficult to disentangle the effect of DRM on content provision from the effect that piracy has on content provision. The model above may give some predictions that can be measured in such a way, for example if piracy and DRM have opposite effects on content provision.

A better approach that can be done now is to look at the choices of the individual content producers in placing their content on iTMS. Some producers have chosen to only list certain songs on iTMS, and keep other songs unavailable through these electronic means. I would look to see if there is a correlation between these choices and the popularity of the songs. If more popular songs are excluded from being listed on iTMS, then this offers evidence toward validating my model.

## 4 Conclusions

The Economics of Security is a recent area of study that has many practical implications and many open questions. I have identified 8 major sub-areas of study in this field, looked at major results and open question in these sub-areas. Then I focused on one open question in the area of Digital Rights Management technologies. I developed a theoretical model that can be used to answer a number

of real-world questions in this area, and identified and motivated some of these questions. Finally, I sketched an empirical study that would begin to validate some of the ideas in this theoretical model.

## A Reading List

### A.1 Economics of Security

#### A.1.1 Trusted computing and DRM

Alessandro Acquisti. Economic incentives for platform providers. In *Workshop on Information Systems and Economics*, 2004. URL <http://opim-sun.wharton.upenn.edu/wise2004/sat611.pdf>.

Ross Anderson. Cryptology and competition policy-issues with ‘trusted computing’. In *Workshop on the Economics of Information Security*, 2003. URL [http://www.cpppe.umd.edu/rhsmith3/papers/Final\\_session1\\_anderson.pdf](http://www.cpppe.umd.edu/rhsmith3/papers/Final_session1_anderson.pdf).

Dirk Bergemann, Thomas Eisenbach, Joan Feigenbaum, and Scott Shenker. Flexibility as an instrument in digital rights management. In *Workshop on the Economics of Information Security*, 2005. URL <http://infoecon.net/workshop/pdf/50.pdf>.

Moshe Givon, Vijay Mahajan, and Eitan Muller. Software piracy: Estimation of lost sales and the impact on software diffusion. *Journal of Marketing*, 59:29–37, 1995. URL <http://links.jstor.org/sici?sici=0022-2429%28199501%2959%3A1%3C29%3ASPEOLS%3E2.0.CO%3B2-V>.

Felix Oberholzer and Koleman Strumpf. The effect of file sharing on record sales – an empirical analysis. In *Second Workshop on the Economics of Peer-to-Peer Systems*, 2004. URL [http://www.unc.edu/~cigar/papers/FileSharing\\_March2004.pdf](http://www.unc.edu/~cigar/papers/FileSharing_March2004.pdf).

Yooki Park and Suzanne Scotchmer. Digital rights management and the pricing of digital products. Technical Report 04-09, NET Institute, October 2004. URL <http://ssrn.com/abstract=618466>.

Pam Samuelson and Suzanne Scotchmer. The law and economics of reverse engineering. *Yale Law Journal*, 111(7):1575, 2002. URL <http://socrates.berkeley.edu/~scotch/re.pdf>.

Arun Sundararajan. Managing digital piracy: Pricing and protection. *Information Systems Research*, 15(3):287–308, 2004. URL <http://oz.stern.nyu.edu/papers/mdpfinal.pdf>.

Also see papers on Network Effects, Service Architecture, and Innovation.

#### A.1.2 Vulnerabilities

Ross Anderson. Open and closed systems are equivalent (that is, in an ideal world). In *Proceedings of Open Source Software: Economics, Law and Policy*, 2004. URL <http://www.cl.cam.ac.uk/ftp/users/rja14/toulousebook.pdf>.

Ashish Arora, Rahul Telang, and Hao Xu. Optimal policy for software vulnerability disclosure. In *Workshop on the Economics of Information Security*, 2004. URL <http://www.dtc.umn.edu/weis2004/xu.pdf>.

Ashish Arora, Ramayya Krishnan, Rahul Telang, and Yubao Yang. An empirical analysis of vendor response to disclosure policy. In *Workshop on the Economics of Information Security*, 2005. URL <http://infosecon.net/workshop/pdf/41.pdf>.

Steve Beattie, Seth Arnold, Crispin Cowan, Perry Wagle, Chris Wright, and Adam Shostack. Timing the application of security patches for optimal uptime. In *Proceedings of LISA 02: Sixteenth Systems Administration Conference*, pages 233–242, 2002. URL <http://www.nxnw.org/~steve/papers/lisa2002-time-to-patch.pdf>.

Dmitri Nizovtsev and Marie Thursby. Economic analysis of incentives to disclose software vulnerabilities. In *Workshop on the Economics of Information Security*, 2005. URL <http://infosecon.net/workshop/pdf/20.pdf>.

Andy Ozment. The likelihood of vulnerability rediscovery and the social utility of vulnerability hunting. In *Workshop on the Economics of Information Security*, 2005. URL <http://infosecon.net/workshop/pdf/10.pdf>.

Eric Rescorla. Is finding security holes a good idea? In *Workshop on the Economics of Information Security*, 2004. URL <http://www.dtc.umn.edu/weis2004/weis-rescorla.pdf>.

Stuart Schechter. How to buy better testing: using competition to get the most security and robustness for your dollar. In *Infrastructure Security Conference*, 2002. URL <http://www.eecs.harvard.edu/~stuart/papers/isc2002.pdf>.

Also see papers on Bundling, Mechanism Design, and Network Effects.

### **A.1.3 Spam**

E. Gabber, M. Jakobsson, Y. Matias, and A Mayer. Curbing junk e-mail via secure classification. In *Proceedings of the 2nd International Conference on Financial Cryptography*, pages 198–213, 1998. URL <http://citeseer.nj.nec.com/27735.html>.

Ben Laurie and Richard Clayton. ‘proof-of-work proves not to work. In *Workshop on the Economics of Information Security*, 2004. URL <http://www.dtc.umn.edu/weis2004/clayton.pdf>.

Thede Loder, Marshall van Alstyne, and Rick Wash. Information asymmetry and thwarting spam. In *Proceedings of ACM E-Commerce*, 2004. URL <http://portal.acm.org/citation.cfm?id=988780>.

Andrei Serjantov and Richard Clayton. Modeling incentives for email blocking strategies. In *Workshop on the Economics of Information Security*, 2005. URL <http://infosecon.net/workshop/pdf/emailblocking.pdf>.

Also see papers on Information Asymmetry and Mechanism Design.

### **A.1.4 Privacy**

Alessandro Acquisti. Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of ACM Electronic Commerce Conference (EC 04)*, pages 21–29, New York, 2004. ACM Press. URL <http://www.heinz.cmu.edu/~acquisti/papers/privacy-gratification.pdf>.

Alessandro Acquisti and Hal R Varian. Conditioning prices on purchase history. *Marketing Science*, 2005. URL <http://www.heinz.cmu.edu/~acquisti/papers/privacy.pdf>.

Andrew M Odlyzko. Privacy, economics, and price discrimination on the internet. In N. Sadeh, editor, *ICEC2003: Fifth International Conference on Electronic Commerce*, pages 255–366, 2003. URL <http://www.dtc.umn.edu/~odlyzko/doc/privacy.economics.pdf>.

Pamela Samuelson. Privacy as intellectual property? *Stanford Law Review*, page 1125, 2000. URL [http://www.sims.berkeley.edu/~pam/papers/privasip\\_draft.pdf](http://www.sims.berkeley.edu/~pam/papers/privasip_draft.pdf).

Hal R Varian. Economic aspects of personal privacy. In *Privacy and Self-Regulation in the Information Age*. U.S. Dept. of Commerce, 1996. URL <http://www.sims.berkeley.edu/~hal/Papers/privacy/>.

Also see papers on Network Effects.

### A.1.5 Investment and Spending

Roger Adkins. An insurance style model for determining the appropriate investment level against maximum loss arising from an information security breach. In *Workshop on the Economics of Information Security*, 2004. URL <http://www.dtc.umn.edu/weis2004/adkins.pdf>.

L Jean Camp and Catherine Wolfram. Pricing security. In *Proceedings of the CERT Information Survivability Workshop*, pages 31–39, 2000. URL <http://www.ljean.org/files/isw.pdf>.

Lawrence A. Gordon and Martin P. Loeb. The economics of information security investment. *ACM Transactions on Information and System Security (TISSEC)*, 5(4):438–457, November 2002. URL <http://portal.acm.org/citation.cfm?id=581274>.

### A.1.6 Security Breaches

Katherine Campbell, Lawrence A. Gordon, Martin P. Loeb, and Lei Zhou. The economic cost of publicly announced information security breaches: empirical evidence from the stock market. *Journal of Computer Security*, 11(3):431–448, 2003. URL <http://portal.acm.org/citation.cfm?id=876669>.

Esther Gal-Or and Anindya Ghose. The economic consequences of sharing security information. In *Workshop on the Economics of Information Security*, 2003. URL <http://www.andrew.cmu.edu/user/aghose/Infosec.pdf>.

### A.1.7 Global Insecurity

L Jean Camp and Catherine Wolfram. Pricing security. In *Proceedings of the CERT Information Survivability Workshop*, pages 31–39, 2000. URL <http://www.ljean.org/files/isw.pdf>.

Dan Geer, Rebecca Bace, Peter Gutmann, Perry Metzger, Charles P. Pfleeger, John S. Quartermann, and Bruce Schneier. Cyberinsecurity: The cost of monopoly. Technical report, Computer & Communications Industry Association, 2003. URL <http://www.ccianet.org/papers/cyberinsecurity.pdf>.

Nicholas Weaver and Vern Paxson. A worst-case worm. In *Workshop on the Economics of Information Security*, 2004. URL <http://www.dtc.umn.edu/weis2004/weaver.pdf>.

### A.1.8 Theory

Ross Anderson. Why cryptosystems fail. In *ACM First conference on Computers and Communications Security*, 1993. URL <http://www.cl.cam.ac.uk/users/rja14/wcf.html>.

Ross Anderson. Why information security is hard: An economics perspective. In *Proceedings of the 17th Annual Computer Security Applications Conference*, 2001. URL <http://www.cl.cam.ac.uk/ftp/users/rja14/econ.pdf>.

Pei-yu Chen, Gaurav Kataria, and Ramayya Krishnan. Software diversity for information security. In *Workshop on the Economics of Information Security*, 2005. URL <http://infoecon.net/workshop/pdf/47.pdf>.

Eddie Dekel and Suzanne Scotchmer. On the evolution of attitudes toward risk in winner-take-all games. *Journal of Economic Theory*, 87:125–143, 1999. URL <http://socrates.berkeley.edu/~scotch/wta.pdf>.

Geoffrey Heal and Howard Kunreuther. You only die once: Managing discrete interdependent risks. Technical Report 9885, National Bureau of Economic Research, 2003. URL <http://ssrn.com/abstract=419240>.

Stuart E. Schechter and Michael D. Smith. How much security is enough to stop a thief? In *The Seventh International Financial Cryptography Conference*, January 2003. URL <http://www.eecs.harvard.edu/~stuart/papers/fc03.pdf>.

Paul Thompson. Cognitive hacking and the value of information. In *Workshop on the Economics of Information Security*, 2002. URL <http://www.sims.berkeley.edu/resources/affiliates/workshops/econsecurity/econws/15.doc>.

Hal R Varian. System reliability and free riding. In *Workshop on the Economics of Information Security*, 2002. URL <http://www.sims.berkeley.edu/resources/affiliates/workshops/econsecurity/econws/49.pdf>.

## A.2 Information Economics

### A.2.1 Information Asymmetry

George Akerlof. The market for lemons: quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3):488–500, 1970. URL <http://links.jstor.org/sici?sici=0033-5533%28197008%2984%3A3%3C488%3ATMF%22QU%3E2.0.CO%3B2-6>.

A M Spence. Job market signalling. *Quarterly Journal of Economics*, 83:355–377, 1973. URL <http://links.jstor.org/sici?sici=0033-5533%28197308%2987%3A3%3C355%3AJMS%3E2.0.CO%3B2-3>.

Joseph Stiglitz. The theory of screening, education, and the distribution of income. *American Economic Review*, 65(3):283–300, June 1975. URL <http://links.jstor.org/sici?sici=0002-8282%28197506%2965%3A3%3C283%3ATTO%22EA%3E2.0.CO%3B2-0>.



### A.2.2 Mechanism Design

- E H Clarke. Multipart pricing of public goods. *Public Choice*, 2:19–33, 1971.
- Drew Fudenberg. *Game Theory*, pages 274–291. MIT Press, 1991. Bayesian Implementations.
- T Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973. URL <http://links.jstor.org/sici?sici=0012-9682%28197307%2941%3A4%3C617%3AIIIT%3E2.0.CO%3B2-E>.
- Andreu Mas-Colell, Michael D Whinston, and Jerry R Green. *Microeconomic Theory*, chapter 23. Oxford University Press, New York, 1995. Incentives and Mechanism Design.
- R Myerson and M Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29:265–281, 1983. URL [http://www.sciencedirect.com.proxy.lib.umich.edu/science?\\_ob=MImg&\\_imagekey=B6WJ3-4CYGD4J-13J-1&\\_cdi=6867&\\_user=99318&\\_orig=browse&\\_coverDate=04%2F30%2F1983&\\_sk=999709997&view=c&wchp=dGLbVzb-zSkzW&md5=2f6a1cecc403c7d6b6159a50cdf9685f&ie=/sdarticle.pdf](http://www.sciencedirect.com.proxy.lib.umich.edu/science?_ob=MImg&_imagekey=B6WJ3-4CYGD4J-13J-1&_cdi=6867&_user=99318&_orig=browse&_coverDate=04%2F30%2F1983&_sk=999709997&view=c&wchp=dGLbVzb-zSkzW&md5=2f6a1cecc403c7d6b6159a50cdf9685f&ie=/sdarticle.pdf).
- Noam Nisan and Amir Rosen. Algorithmic mechanism design. *Games and Economic Behavior*, 35:166–196, 2001. URL [http://www.sciencedirect.com.proxy.lib.umich.edu/science?\\_ob=MImg&\\_imagekey=B6WFW-458NM77-2P-1&\\_cdi=6805&\\_user=99318&\\_orig=browse&\\_coverDate=04%2F30%2F2001&\\_sk=999649998&view=c&wchp=dGLbVlz-zSkzS&md5=5562485d41459b9f47ec0d445f0a53a6&ie=/sdarticle.pdf](http://www.sciencedirect.com.proxy.lib.umich.edu/science?_ob=MImg&_imagekey=B6WFW-458NM77-2P-1&_cdi=6805&_user=99318&_orig=browse&_coverDate=04%2F30%2F2001&_sk=999649998&view=c&wchp=dGLbVlz-zSkzS&md5=5562485d41459b9f47ec0d445f0a53a6&ie=/sdarticle.pdf).

### A.2.3 Network Effects

- Joseph Farrell and Garth Saloner. Installed base and compatibility: Innovation, product preannouncements, and predation. *American Economic Review*, 76:940–955, 1986. URL <http://links.jstor.org/sici?sici=0002-8282%28198612%2976%3A5%3C940%3AIBACIP%3E2.0.CO%3B2-0>.
- Joseph Farrell and Garth Saloner. Standardization, compatibility, and innovation. *RAND Journal of Economics*, 16:70–83, 1985. URL <http://links.jstor.org/sici?sici=0741-6261%28198521%2916%3A1%3C70%3ASCAI%3E2.0.CO%3B2-Q>.
- Michael L Katz and Carl Shaprio. Technology adoption in the presence of network externalities. *Journal of Political Economy*, 94(4):822–841, 1986. URL <http://links.jstor.org/sici?sici=0022-3808%28198608%2994%3A4%3C822%3ATAITP0%3E2.0.CO%3B2-D>.
- Michael L Katz and Carl Shaprio. Network externalities, competition, and compatibility. *American Economic Review*, 75(3):424–440, 1985. URL <http://links.jstor.org/sici?sici=0002-8282%28198506%2975%3A3%3C424%3ANEACAC%3E2.0.CO%3B2-M>.
- Stan J Liebowitz and Stephen E Margolis. Are network externalities a new source of market failure? *Research in Law and Economics*, 17:1–22, 1995. URL <http://www.utdallas.edu/~liebowit/netwextn.html>.
- Geoffrey Parker and Marshall van Alstyne. Two-sided network effects: A theory of information product design. *Management Science*, 2005. (Forthcoming).
- J-C Rochet and Jean Tirole. Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4), 2003. URL <http://www.univ-tlse1.fr/idei/Commun/WorkingPapers/2002/Platform.pdf>.

#### A.2.4 Sharing and Bundling

- Yannis Bakos and Erik Brynjolfsson. Aggregation and disaggregation of information goods: Implications for bundling, site licensing, and micropayment systems. *Management Science*, 1999. URL <http://www.stern.nyu.edu/~bakos/aig.pdf>.
- Yannis Bakos, Erik Brynjolfsson, and D Lichtman. Shared information goods. *Journal of Law and Economics*, 42, 1999. URL <http://ssrn.com/abstract=130904>.
- J C-I Chuang and M A Sirbu. Optimal bundling strategy for digital information goods: Network delivery of articles and subscriptions. *Information Economics and Policy*, 11(2):147–176, 1999. URL <http://www.sims.berkeley.edu/~chuang/pubs/ediip2.pdf>.
- Hal R Varian. Buying, sharing and renting information goods. *Journal of Industrial Economics*, 48(4):473–488, 2000. URL <http://www.sims.berkeley.edu/~hal/Papers/sharing.pdf>.

#### A.2.5 Service Architecture Effects

- Jeffrey K Mackie-Mason, Scott J Shenker, and Hal R Varian. Network architecture and content provision: An economic analysis. In G Brock and G Rosston, editors, *The Internet and Telecommunications Policy*. Lawrence Erlbaum, 1996. URL <http://www-personal.umich.edu/~jmm/papers/tprc.pdf>.

#### A.2.6 Innovation

- S M Besen and S N Kirby. Private copying, appropriability, and optimal copying royalties. *Journal of Law and Economics*, 32:255–280, October 1989. URL <http://links.jstor.org/sici?sici=0022-2186%28198910%2932%3A2%3C255%3APCAA0C%3E2.O.CO%3B2-K>.
- J Bessen and E Maskin. Sequential innovation, patents and imitation. URL <http://www.researchoninnovation.org/patrev.pdf>. Draft, 2002.
- Suzanne Scotchmer. *Innovation and Incentives*. MIT Press, 2004. Chapters 2, 4–6, 10.
- Steven Shavell and Tanquy van Ypersele. Rewards versus intellectual property rights. *Journal of Law and Economics*, 44(2):525–547, October 2001. URL <http://www.journals.uchicago.edu/JLE/journal/issues/v44n2/010101/brief/010101.abstract.html>.

#### A.3 Other Referenced Works

- N Economides. The economics of networks. *International Journal of Industrial Organization*, October 1996. URL <http://raven.stern.nyu.edu/networks/top.html>.
- Abhishek Kumar, Vern Paxson, and Nicholas Weaver. Exploiting underlying structure for detailed reconstruction of an internet scale event. Technical report, Georgia Tech University, 2005. URL <http://www.cc.gatech.edu/~akumar/witty.html>.
- Paul Milgrom and Chris Shannon. Monotone comparative statics. *Econometrica*, 62(1):157–180, January 1994.

Andy Ozment. Bug auctions: Vulnerability markets reconsidered. In *Workshop on the Economics of Information Security*, 2004.

Caroline Richardson, Paul Resnick, Derek Hansen, Holly Derry, and Victoria Rideout. Does pornography-blocking software block access to health information on the internet? *Journal of the American Medical Association*, 288(22):2887–2894, 2002. URL <http://jama.ama-assn.org/cgi/content/abstract/288/22/2887>.

Steven Salop and David Scheffman. Raising rivals costs. *American Economic Review*, 73(2):267–271, 1983.